

Case Study

Intel® Xeon® Scalable Processors
Intel® Software Tools



Anodot Boosts Machine Learning Anomaly Detection and Forecasting Performance

The Intel optimized solution ensures Anodot's AI platform scales with its business growth.

Highlights:

- Anodot used Intel® products to create a solution that accelerates ML algorithm training and inference runtime
- Training performance for ACF algorithm up to 127X faster with Intel® Integrated Performance Primitives (Intel® IPP) software
- Tests show inference time for XGBoost algorithm 4X faster with Intel® technology-driven solution
- Solution includes Intel® Xeon® Scalable Processors and Intel® software

About Anodot

Anodot's AI-driven, autonomous monitoring solution identifies revenue-critical business incidents, providing real-time alerts and forecasts. Organizations rely on Anodot for near real-time incident detection in their high volume and velocity data. They use Anodot to protect revenue, reduce costs, and improve daily operations. Customers include data intensive companies such as Pandora, Affirm, T-Mobile, Vodafone, UPS, Credit Karma, LivePerson, Payoneer, Vimeo, Puma, Atlassian and others. Two applications drive the company's platform: anomaly detection and forecasting.

The anomaly detection application looks at numerous patterns, identifying and correlating unusual activity and events. Anodot's forecasting application takes what the anomaly detection application has learned about the behavior of the data and then provides accurate forecasts for business growth and demand.

Challenge

Anodot's anomaly detection and forecasting applications consist of ensemble machine learning (ML) models that analyze time series data. These models analyze hundreds of millions of time series metrics every minute for their customers. Every new customer adds hundreds of thousands of new metrics that the models must analyze to learn and understand patterns in the data.

Anodot needed a solution that would give their ML platform unlimited scalability while allowing the company to manage its compute costs effectively as it grows. They also wanted to improve the speed, efficiency, and accuracy of model training



Faster Autocorrelation
Function Training

127X



Faster XGBoost
Forecast Inferencing

4X



Case Study | Anodot Boosts Machine Learning Anomaly Detection and Forecasting Performance

and inferencing — all crucial to detecting anomalies in real-time and predicting future business performance for customers.

Solution

While Anodot already runs their AI platform on Intel® CPUs, the team extended their collaboration with Intel to create performance tests to optimize the performance of their Autocorrelation Function (ACF) and XGBoost algorithm using Intel hardware and software.

Anodot's data science team identified the algorithms that utilized the most compute resource in each application: the ACF in anomaly detection and XGBoost in forecasting. The ACF algorithm accounts for 66% of the compute resource required to learn and understand patterns found in millions of metrics by analyzing time series data multiple times per week. For forecasting, Anodot uses XGBoost, a highly accurate algorithm used to perform roughly 50% of the forecasting tasks.

The team benchmarked ACF training performance improvements using their SDK implementation as the baseline and compared it to the Intel IPP optimized ACF, AutoCorrNorm, written in C++. The team tested the training runtime for varying lengths of time series data. Each series had anywhere from 1,000 to 38,000 data points — representing the most common lengths of time series data Anodot analyzes.

"Anomaly detection technology needs to understand what normal pattern is in a world of moving parts and constant change. Anodot has figured this out and patented this method in our AI algorithm. When done well, anomaly detection can also provide the foundation for other business applications such as forecasting."

— Ira Cohen, Chief Data Scientist, Anodot

Anodot also tested the XGBoost performance improvements using Intel® oneAPI Data Analytics Library (oneDAL) and the baseline XGBoost model. The team focused on inference as the Anodot's system performs this function constantly on

real-time data. They created a oneDAL-optimized version of their XGBoost model with a single line of code. The task involved forecasting a time series containing 24 hourly data measurements per day over a 30-day horizon that spanned four years of historical data.

The performance test results demonstrated a dramatic improvement in the ACF training performance and the XGBoost inference performance.

Results

"When choosing a machine learning platform, you need to think about scale as your business grows. So, model efficiencies and compute cost effectiveness become increasingly important. Our performance tests show the Intel software and Xeon platform provide us efficiency gains that will allow us to deliver an even higher quality of service at lower cost."

—Ira Cohen


The performance tests demonstrated that the specific combination of Intel's hardware and software delivered substantial performance improvements to the ACF and XGBoost machine learning implementations.

ACF Performance Results (Anomaly Detection Model):

- Up to 127X faster training performance with Intel IPP software.
- 66% reduction in the overall cost of running the training algorithm in a cloud environment — achieved by cutting the ACF runtime by almost 99%.
- Anodot's data scientists now have the freedom to test more complex algorithms without worrying about high runtime costs.

XGBoost Performance Results (Forecasting Model):

- 4X faster inferencing time with Intel software optimization.
- Forecast service can now analyze four times the amount of data at no additional cost for inference.



By cutting the ACF runtime by almost 99%, the runtime of the overall algorithms is reduced by almost 66%. Assumptions provided by Anodot. Instance cost/CPU hour (4 vCPU) \$0.050; # CPU hours per 1 M metric 500; # Training runs per month 10. Reference white paper for detailed information regarding tests conducted by Anodot.

Summary

The results of the performance test show that Intel® Xeon® platform and software frameworks improved the performance of Anodot's ACF and XGBoost algorithms, significantly reducing ML compute time and costs. The test also identified learning model efficiency improvements and solution scalability. This solution optimization ensures that Anodot's platform can deliver real-time anomaly detection and forecasting services at scale.

Solution Components:

- Intel® oneAPI Data Analytics Library (oneDAL)
- Intel® IPP optimized ACF, AutoCorrNorm, written in C++
- Intel® Xeon® Scalable Processors

*Anodot ran performance tests on AWS instances in November of 2021. More detailed information can be found in this [white paper](#).

Additional Information

[Intel White Paper: "Accelerate Real-Time Machine Learning Based Anomaly Detection and Forecasting at Scale."](#)

[Anodot Incident Detection](#), [Autonomous Forecast](#), [Payment Transaction Monitoring](#), [Cloud Cost Management](#), and [Digital Experience Monitoring](#).

[Intel® Developer Zone](#), [Xeon® Scalable Processors](#), [Integrated Performance Primitives](#) (Intel IPP), [oneAPI Data Analytics Library](#) (oneDAL), and optimized ACF, AutoCorrNorm from Intel.



Data from Anodot, November 2021. Reference white paper for detailed information regarding tests conducted by Anodot.

Notices and Disclaimers

Performance varies by use, configuration, and other factors. Learn more on the [Performance Index](#) site.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary. Intel technologies may require enabled hardware, software, or service activation.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.