AN INTRODUCTION TO BUSINESS MONITORING

Business monitoring focuses on the incidents and trends that impact your revenue. Why monitor only apps and IT—when you can monitor the KPIs that directly influence your business?

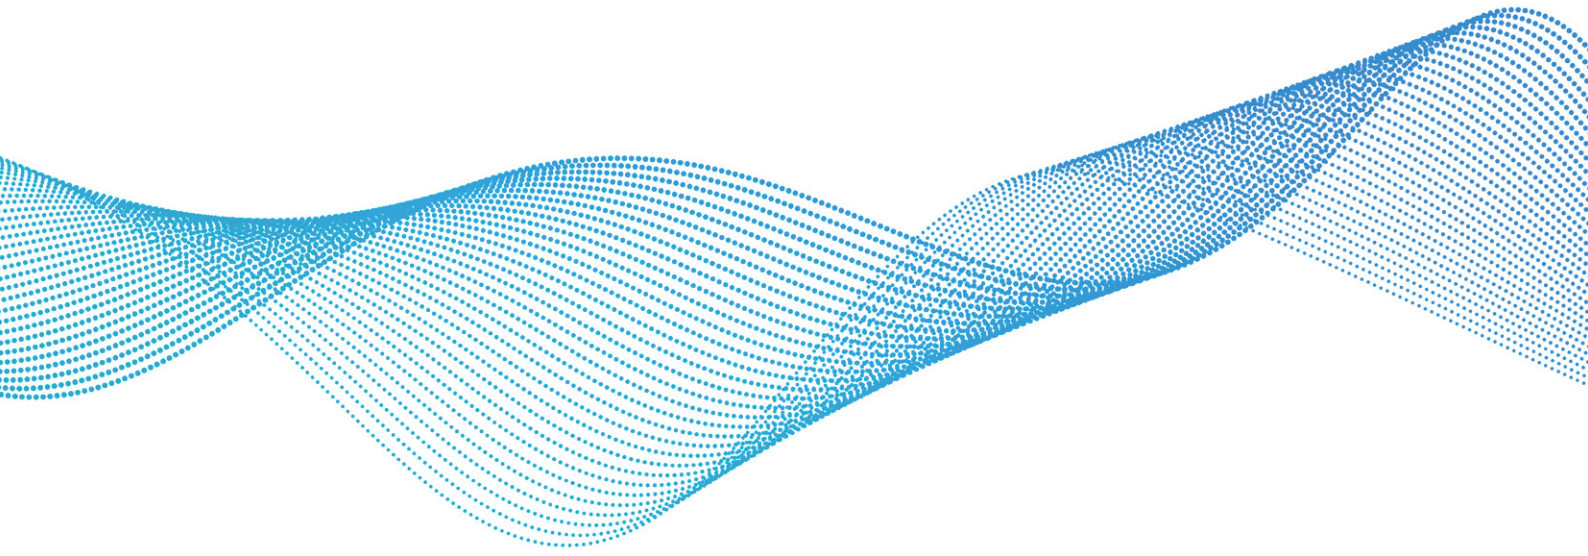

anodot

IT'S ALL ABOUT THE BUSINESS IMPACT

Whether you are in ecomm, telco, gaming, or adtech, every event or instance happening across your organization is eventually translated into its highest common denominator: the impact it carries for your business.

On a daily basis, every department deals with its own ups and downs: a faulty version release, an increase in conversions, server under-utilization, a decrease in customer support calls, unstable APIs. These events are monitored and measured according to their relevant attributes, and remediated - or doubled-down on - according to the specific use case. But at the end of the day, they all come together to define and determine the trajectory of the business. This is, usually, where things become blurry. Because while it is (relatively) feasible to monitor, for example, hourly costs of cloud resources, login success rates, or ad impressions—tying these metrics back to the overall business performance is altogether another kind of monitoring feat.

Many companies today try to feed business metrics - such as user activity, quality of service, or revenue - into APM or IT monitoring systems. Splunk, Datadog and others track your business in real time, based on log or application data – something that would seem to make sense. In practice, however, it fails to produce accurate and effective monitoring or reduce time to detection of revenue-impactful issues.

Why? Because monitoring machines and monitoring business KPIs are completely different tasks.

THE COMPLEX NATURE OF BUSINESS METRICS

Business metrics pose a unique monitoring challenge for three main reasons:

Context

Business metrics derive their significance from their unique context. They cannot be evaluated in absolute terms, but only in relation to a set of changing conditions. In that sense, they are subjected to interpretation. For example, their expected value is relative to the metric history. When we examine active users, we don't really care about the absolute number: what we want to monitor is how that number compares to the same time frame in the previous few weeks. In this scenario, an alert that informs us that the number has gone below 100 users is meaningless: we will always need to evaluate that number in context.

One of the reasons that business metrics are highly context-sensitive is that they're human-dependent. In most cases, business metrics are heavily influenced by human behavior, which is seasonal (day and night, weekday and weekend), and in itself influenced by a large number of parameters. For example, when we monitor data consumption on a telco network, human context teaches us that very low consumption during the night is fine. But the same rates during daytime might signal trouble. Unless it's a holiday, which might need to adjust metric thresholds to a different baseline altogether—etc.

Topology

As a monitoring framework, business topology is unknown. This concept can be easily demonstrated when compared to other types of data. When we monitor machine data we know the relationship between the different machines, and between the different operational metrics for each machine. For instance, we have a pre-figured framework for the relationship between server performance and utilization. That framework is lacking when monitoring business data. The relationships and correlations between the different metrics are too dynamic and volatile. We must rely on algorithms to learn and map them.

Volatility

Business sampling rate is irregular. When monitoring machines, we'll get a new data

point every x seconds as long as the machine is alive. This is not the case for business KPIs. We may experience full minutes and even hours without a purchase or a click. Irregular sampling rate poses a unique set of monitoring challenges: it requires significant adaptation in how data is stored, and how the algorithm works.

In short, business metrics necessitate a monitoring solution that can cope with these layers of complexity by applying:

1. Monitoring of 100% of granular data in real-time
2. Autonomous learning of metric behavior and seasonality
3. Full metric correlation and root cause analysis

MONITORING LAYERS ACROSS THE BUSINESS

Business Monitoring
- Revenue
- Cost
- Partners
- Audience

App Performance Monitoring
- Code
- Errors
- Latency
- Transactions

IT Monitoring
- Network
- Storage
- Logs
- CPU

DEEP 360 BUSINESS MONITORING

AI- and ML-driven platforms built for business monitoring apply these abilities in order to help businesses safeguard their revenues and costs, partners, and customer experience, including the audience journey and engagement.

Revenue and Cost Monitoring

Revenue and cost monitoring ensures a consistent revenue stream by identifying impactful issues as soon as they occur.

A business's revenue model is complex and fragmented into multiple streams consisting of micro-transactions, subscriptions, partners and affiliates. A glitch in any of these revenue sources can result in massive bleeds to the bottom line, either through major incidents or by small daily trickles that are overlooked for an extended period of time. Using AI to monitor these detailed revenue streams in real-time is the only way to ensure watertight revenue protection.

By monitoring revenue streams across segments, plans, products and payment providers, or marketing costs over various channels and campaigns, this type of business monitoring bulletproofs the business's revenue streams in the face of acute or chronic incidents. Typical use cases for revenue and cost monitoring include purchases monitoring, cloud cost monitoring, adwords monitoring, payment gateways monitoring—and many more.

Partners Monitoring

Every business is supported by dozens of connected systems. Marketing depends on publishers and email providers, customers rely on payment gateways and order management systems. A late, misguided or missing response to critical incidents has significant costs in revenue, time and brand reputation, often outlasting the incident itself.

That's why it's essential to monitor 3rd party tools and platforms that support and enable the business, including APIs, recommendations, js snippets, payments processing vendors, etc. Degradation in the level of service of a 3rd party can result in bad customer experience, loss of data and in many cases actual loss of revenues. The limited visibility towards 3rd party services, coupled with its potential impact on the business, makes this an area especially prone to critical incidents, and therefore extremely important to monitor efficiently.

Real-life examples of digital partner monitoring include partner networks monitoring, affiliate networks monitoring, ad performance monitoring—any use case where a 3rd party technology is tied into the business processes and performance.

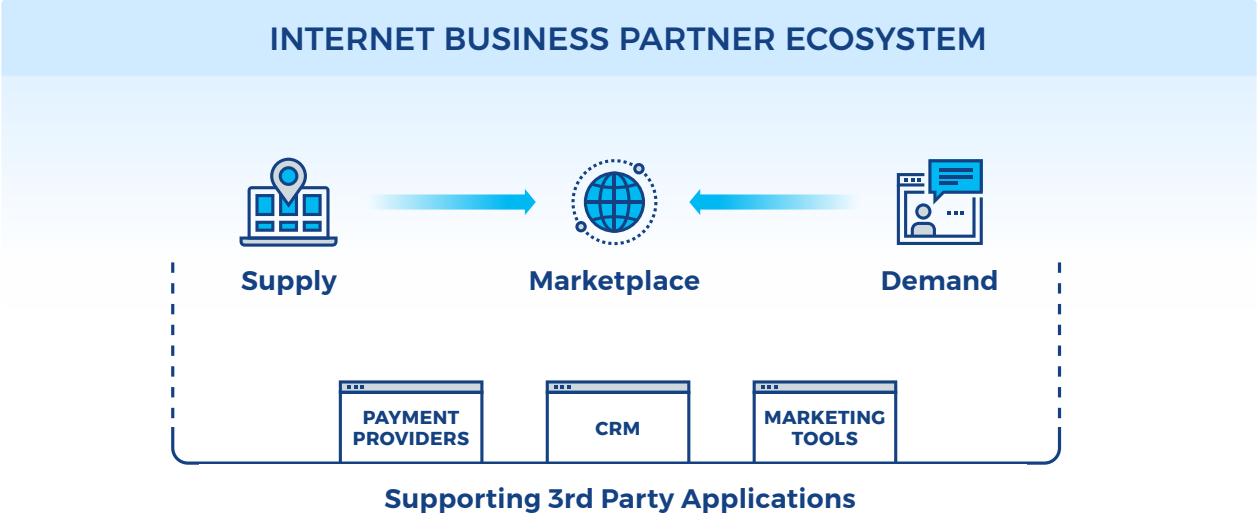

Customer Experience Monitoring

Marketing, product, and customer service are the crucial links in the chain that creates an outstanding customer experience. Customer experience lapses across the customer journey, throughout the product's versions and permutational complexities, will happen. A critical issue in any of the links breaks the chain, immediately

impacting customer experience, engagement, conversions, revenues, and reputation.

To ensure a seamless experience throughout the journey—from awareness to acquisition, retention, upsell/cross sell and advocacy—customer experience monitoring is non-negotiable. It's importance is emphasized when daily processes that have a potential to impact it are taken into consideration. These include events such as adding new features, monitoring A/B tests, releasing new versions, changing customer support processes, purchase funnel analysis, and more.

Customer experience monitoring can be leveraged for varied scenarios, including—but far from being limited to—active users monitoring, user activity monitoring, user retention monitoring, churn monitoring and usage monitoring.

BUSINESS MONITORING WITH ANODOT

As its name implies, Anodot's Autonomous Business Monitoring solution is built from the ground up to monitor business metrics. As opposed to other monitoring solutions (i.e. legacy solutions, APM, AIOps etc.), which are all about IT and application performance, Anodot takes a modern approach: focus on business monitoring first. Anodot leverages AI to constantly monitor and correlate business performance, providing mission-critical real-time alerts and forecasts. In order to provide industry-leading business monitoring, Anodot employs the following monitoring logic:

1. **Unsupervised machine learning built for business monitoring.** Learning every metric's normal behavior is a prerequisite to identifying anomalous behavior. However, business metrics run a wide gamut of signal types. That's why Anodot uses sequential adaptive learning algorithms on the fly that initialize a model of what is normal, and then compute the relation of each new data point going forward.

2. **Real-time monitoring of 100% of data.** Significant anomalies can occur in various metrics and "business depths". For example: revenues should be monitored per product, country, technology and also technology + country, product + technology and product + technology + country combined—across multiple layers of business data. In addition, disparate anomalies need to be constantly correlated to report on incidents in context, requiring complete data coverage. Anodot analyzes 100% of the business's metrics in real time and at scale by running machine learning algorithms on the live data stream itself, without reading and writing to a database.

3. **Detection of seasonality.** Anodot's patented Vivaldi method for seasonality detects every kind of seasonality, enabling the model to construct a true to life understanding of metric behavior, and freeing monitoring personnel from the reign of false positives, false negatives and alert storms.

4. **Anomaly Scoring.** Grading anomalies is critical for filtering alerts by significance. Alerts are scored according to deviation, duration, frequency and other related conditions. Anodot's patented anomaly scoring method runs probabilistic Bayesian models to evaluate anomalies both relative to normal based on their anomaly pattern, and relative to each other, to ensure that only—and all—significant incidents are detected and flagged.

The end result is a powerful business monitoring solution based on AI, which empowers

users to remedy urgent problems faster, and capitalize on opportunities sooner.

Anodot's Autonomous Business Monitoring solution is ideal for companies in various industries (e-commerce, online retail, software, adtech, digital entertainment, fintech and more). It automatically illuminates data blind spots so companies never miss another brand damaging incident or business opportunity. Its automated machine learning algorithms continuously analyze all business data and alert in real time whenever an incident occurs, even for questions that were never asked. Its built-in data science means that any user can easily gain actionable insights, even without any data science knowledge.

Over 40% of Anodot's customers are publicly traded companies, including Microsoft, Waze (a Google company), AppNexus, Comcast and many others.

www.anodot.com