



Grid Dynamics Whitepaper

The Essential Guide to Transforming IT Operations with AIOps

www.griddynamics.com



Grid Dynamics

September 2021

Contents

Introduction	3
What is AIOps?	4
Use Cases and Capabilities	5
How to Implement AIOps	6
Conclusion	7



Introduction

Modern IT operations have to deal with dynamic mixes of public cloud platforms and services, cloud-native and serverless applications, and on-premise deployments. These systems, services, and applications generate enormous amounts of data that are challenging to collect, analyze, and use for issue detection and remediation. In this white paper, we discuss how this challenge can be addressed using machine learning and artificial intelligence methods, what aspects of IT operations can be improved using such techniques, and how companies should plan their capability roadmaps in this area.

What is AIOps?

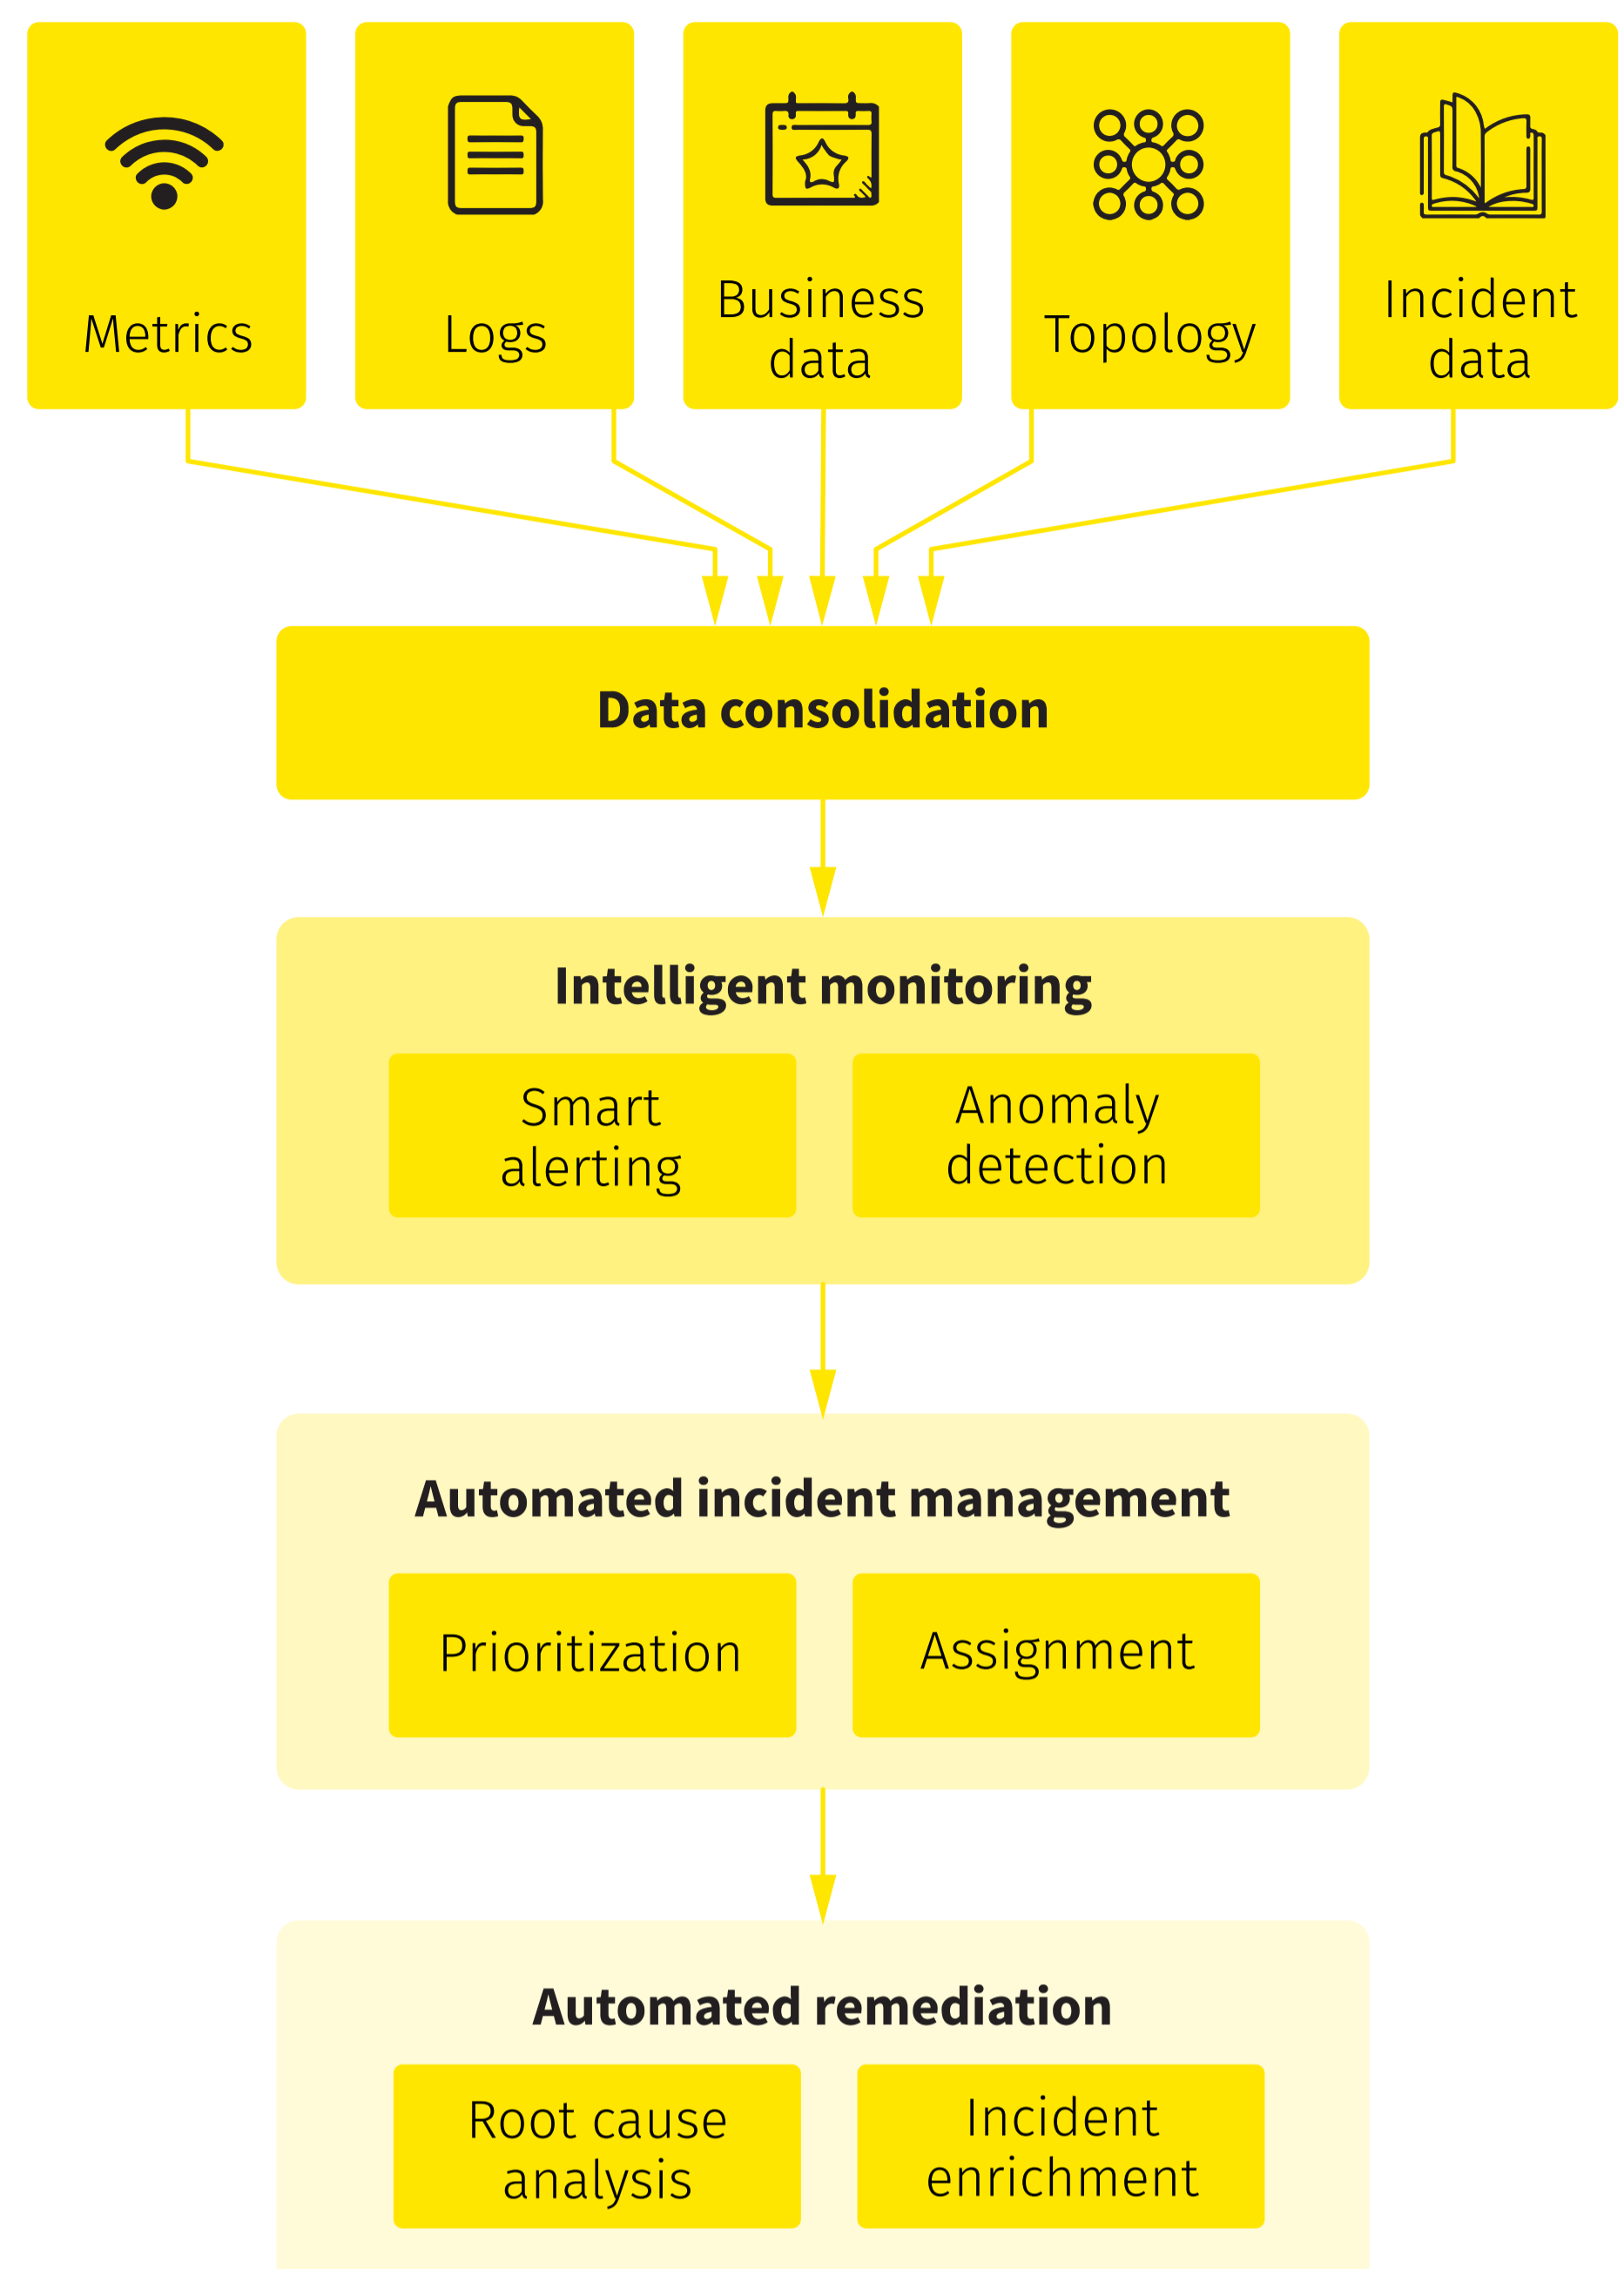
The term “AIOps” stands for “artificial intelligence for IT operations.” AIOps generally aims to combine big data, machine learning, and intelligent automation to improve the efficiency of IT operations and provide a better level of service. In this section, we discuss the areas of IT operations that can be enhanced using the AIOps technologies and the corresponding benefits of the AIOps transformation.

Data Consolidation at Scale

The prerequisite for using advanced analytics and data-driven methods in IT operations is a powerful data collection and consolidation infrastructure. AIOps assumes that traditional sources of IT data such as metrics and logs are consolidated and available for the analysis using machine learning methods and ongoing data streams can be processed in real-time. AIOps also underscores the importance of additional signals and data sources that can be used to explain the variance in the observed metrics. Examples of such signals and data elements include business events, release schedules, and environment topology.

Intelligent Monitoring

The ability to automatically perform deep analysis of thousands of interrelated metrics and signals is one of the key advantages of AIOps solutions. Machine learning algorithms can detect complex and implicit patterns that cannot be reliably identified using traditional rule-based methods and monitoring tools, and intelligently control the intensity of alerts to prevent an alert storm.



Automated Incident Management

The issues identified by anomaly detection algorithms need to be prioritized and assigned to relevant teams. AIOps solutions can substantially improve the efficiency of these steps using issue prioritization and issue classification models that use the contextual information about the incident (anomaly magnitude and duration, involved applications, correlated events, etc.) to make accurate decisions with regard to severity levels and remediation flow.

Automated Remediation

Finally, AIOps helps to automate the remediation workflows, decreasing the resolution time and the level of effort associated with the issue troubleshooting. This is usually achieved by providing advanced incident enrichment and root cause analysis tools. The incident enrichment algorithms and tools automatically identify the most relevant past cases, fragments of the runbooks, and segments of metrics and logs related to the incident. The root cause analysis tools help to reduce the resolution time by automatically tracing the dependencies between anomalies and detecting the source of the problem.

Use Cases and Capabilities

In the previous section, we outlined the main areas of IT operations, including monitoring, incident management, and remediation, that can be improved using machine learning methods. In this section, we go one level deeper and discuss specific use cases and operational capabilities that can be enabled by machine learning and advanced analytics.

Reliable Detection of Hidden Anomalies

The ability to alert on metrics that exceed manually configured thresholds is widely supported by conventional monitoring software. The threshold-based approach is a reasonable solution for detecting major failures, but it falls short when it comes to detecting more subtle anomalies such as unusual API call patterns or deviations from weekly cycles. These anomalies quite often indicate severe software defects and configuration errors that can have a major business impact. Detecting such anomalies using manually created rules and thresholds is usually not tangible, but machine learning methods can solve this use case very efficiently and detect hidden issues before they start to propagate.

Automatic Reconfiguration and Scaling

Automatic anomaly detection not only provides the ability to detect complex anomalous patterns and detect precursors of major failures, but also eliminates the configuration effort associated with manual threshold configuration and tuning. Modern anomaly detection algorithms support completely automated onboarding of new metrics and applications, so that the new data streams

can be added in a plug-and-play mode. Depending on the metric type, the system can either reuse models trained on historical data for similar entities, or train a new model on the fly.

Smart Alerting

Traditional alerting systems are prone to alert storms and missed issues. Machine learning offers several powerful capabilities for addressing these problems. The first one is advanced anomaly detection algorithms that are able to learn complex patterns including seasonality and cross-metric correlations, so that anomalies are intelligently scored, and only systematic deviations from the normal behavior is reported. The second important capability is learning from the feedback, so that the operations team can triage the received anomalies, and provide ground truth labels back to the model training process. The combination of these methods provides a strong foundation for building an accurate alerting system that outperforms the traditional solutions.

Proactive Monitoring

Machine learning helps not only to detect anomalies that have already started to unfold, but also detect trends that potentially can result in the actual anomalies and failures. One of the most common examples in this category is the analysis of performance trends with a goal to prevent overloading. The AIOps solutions are not only capable of analyzing the trends and predicting the future state of the system, but also providing specific recommendations on actions that need to be taken in order to prevent failures. Proactive capabilities can greatly improve the efficiency of IT operations helping to prevent losses, reduce outages, and minimize time to resolution (TTR).

Faster Root Cause Analysis

Root cause analysis is a challenging problem in complex environments with hundreds of applications and thousands of metrics because one failure or anomaly can trigger an alert storm. AIOps solutions help to mitigate this problem using topology-aware anomaly detection models and anomaly localization algorithms. Topology-aware models incorporate the information about application dependencies and automatically identify groups of entities involved in each anomaly. Anomaly localization algorithms identify representative time frames for each anomaly. These two techniques help to automatically compile a summary on each incident which includes the graph of involved entities, anomaly propagation path, and relevant segment of metrics.

How to Implement AIOps

The AIOps ecosystem is often implemented as a mix of proprietary products and custom components. In this section, we discuss a lightweight approach to building intelligent monitoring and automated remediation capabilities outlined in the beginning of this paper. We show how these capabilities can be established using generic cloud services and open-source components at a fraction of the costs associated with heavyweight proprietary systems.

1. Establish Data Consolidation

As the first step, metrics and logs collected from cloud and on-premise environments are processed in a streaming mode using cloud-native services such as AWS Kinesis and EMR, and consolidated in a storage such as ElasticSearch.

2. Establish Basic Anomaly Detection

Anomaly detection models are developed using components from open-source libraries, managed and retrained in a cloud ML platform such as AWS SageMaker using the ongoing data streams, and deployed in production for real-time metric scoring.

3. Implement Root Cause Analysis Tools

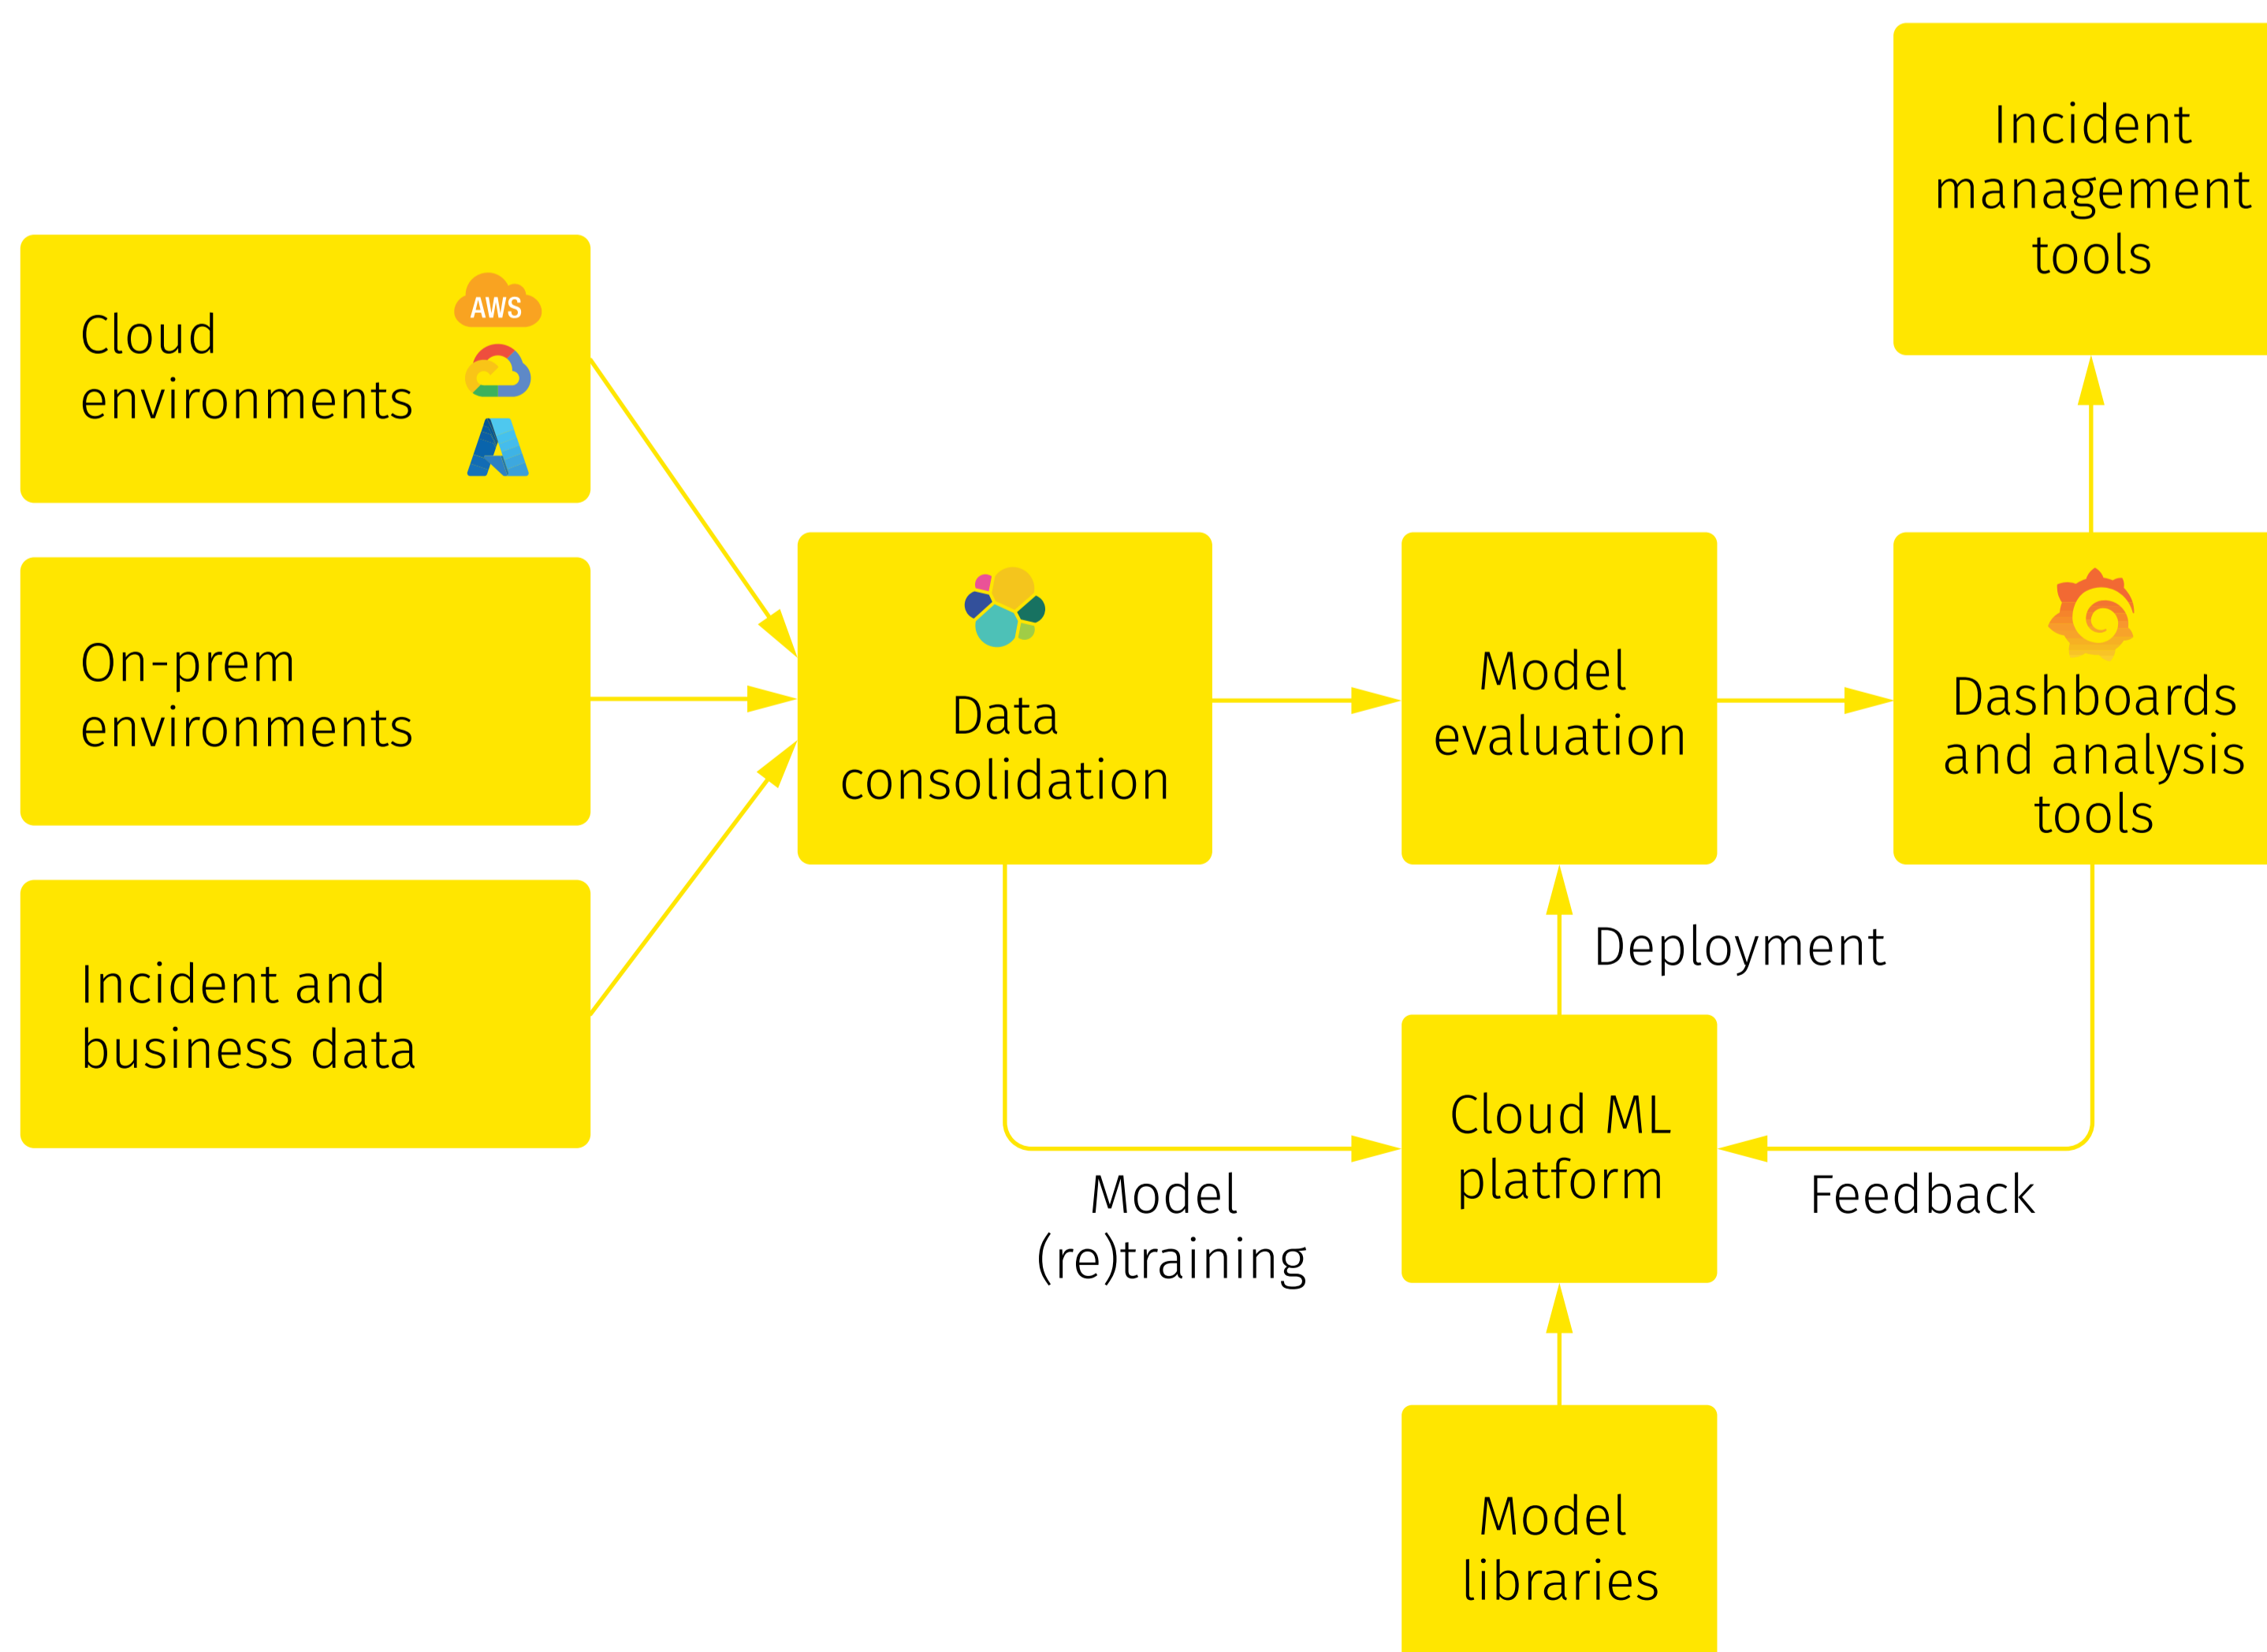
The metric storage and outputs of the models including alerts and anomaly summary data are integrated with visualization tools such as Grafana. Custom dashboards for root cause analysis and resolution feedback collection are created.


4. Improve Maintainability

The basic anomaly detection algorithms are extended with capabilities that improve scalability, efficiency and maintainability including plug-and-play onboarding of new metrics and entities and automatic fine-tuning of alerts based on the feedback received from the operations team.

5. Improve Incident Management

Finally, the outputs of the anomaly detection processes are integrated with the incident management tools. The incident prioritization and assignment capabilities provided by the incident management software can be extended with custom rule engines and models.





Conclusion

In this white paper, we discussed the typical layout of AIOps capabilities, specific use cases that can be improved using machine learning, and high-level implementation roadmap for a lightweight AIOps ecosystem. This analysis shows that existing IT solutions can be rapidly augmented with AIOps features using open-source components and basic cloud services.

About Grid Dynamics

Grid Dynamics (Nasdaq: GDYN) is a digital-native technology services provider that accelerates growth and bolsters competitive advantage for Fortune 1000 companies. It provides digital transformation consulting and implementation services in omnichannel customer experience, big data analytics, search, artificial intelligence, cloud migration, and application modernization. Grid Dynamics achieves high speed-to-market, quality, and efficiency by using technology accelerators, an agile delivery culture, and its pool of global engineering talent.

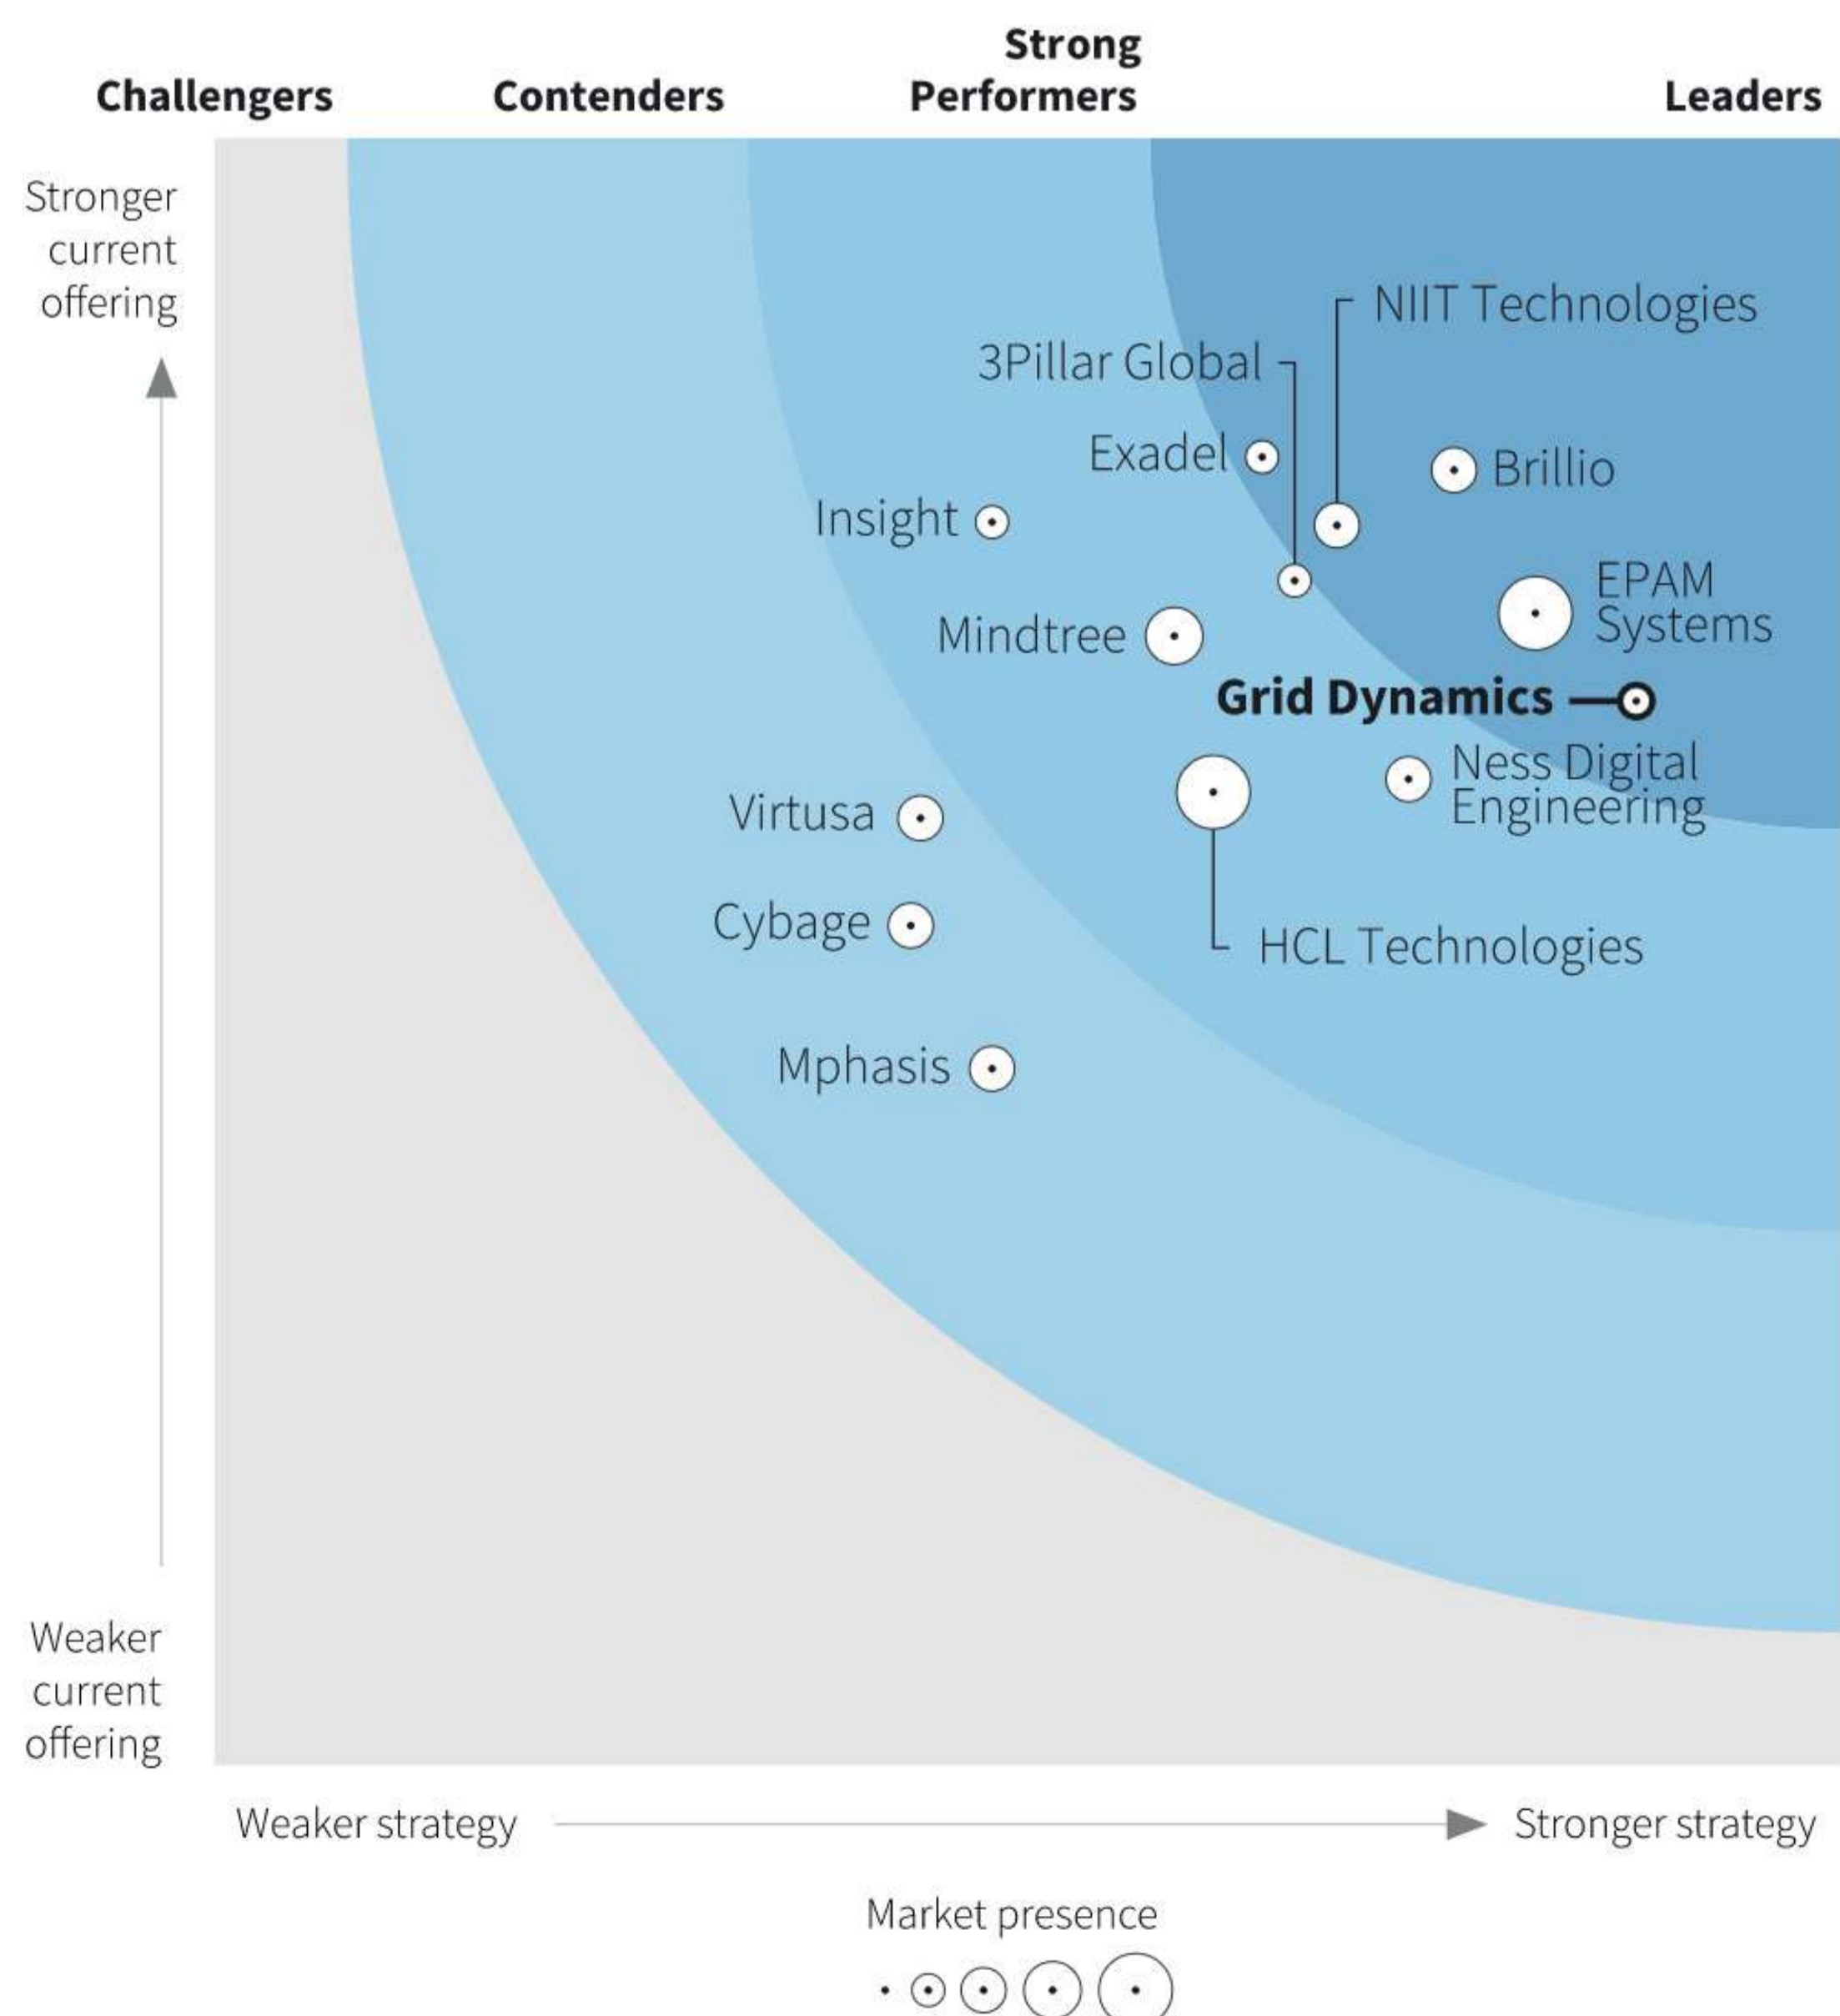
Grid Dynamics works in close collaboration with its clients on digital transformation initiatives that span strategy consulting, early prototypes, and the enterprise-scale delivery of new digital platforms. It helps organizations become more agile and create innovative digital products and experiences using its

extensive expertise in emerging technology, top global engineering talent, lean software development practices, and a high-performance product culture.

Headquartered in Silicon Valley, with delivery centers located throughout the United States, Western, Central, and Eastern Europe, Grid Dynamics is known for architecting and delivering some of the largest digital transformation programs in the retail, technology, and financial sectors to help clients win market share, shorten time to market, and reduce the costs of digital operations on a massive scale.

In 2019, Forrester named Grid Dynamics a leader among midsize agile development service providers. In 2020, Grid Dynamics went public and is trading on the NASDAQ under the GDYN ticker.

The Forrester wave™ Midsize Agile Development Service Providers Q2 2019



About Grid Dynamics











Key facts about us

- 9 countries across North America and Europe
- 2,510 employees in Q2 2021
- Forrester Leader Midsize Agile Software Development Service Provider Q2 2019

Our areas of expertise

- Data Science and Artificial Intelligence
- Analytical data and ML platforms
- Enterprise and site search
- Omnichannel experience
- Cloud enablement
- DevOps and QA automation

Our clients



Grid Dynamics

trusted engineering partner for digital transformation

Grid Dynamics LLC.

5000 Executive Parkway,
Suite 520 / San Ramon, CA
650-523-5000

info@griddynamics.com

www.griddynamics.com