Grid Dynamics

Analytics and ML Platform modernization:

A case study

www.griddynamics.com

powered by aws

Contents

Background

Our client, a Canada-based startup, built a global gaming loyalty platform that aims to help mobile gamers identify new games based on their gaming history. The platform offers built-in rewards to gamers for specific actions, which can be redeemed as Amazon, PayPal or Google Play gift cards.

Since the gaming industry has evolved and exploded over the last few years, the client's overarching objective is to make it easier for mobile gamers to discover new and relevant games in the sea of hundreds of thousands of new releases.

To personalize game recommendations, the client identifies customer preferences based on the types of games a customer has played and which applications are downloaded on their phones.

The loyalty platform helps customers identify new, unknown games that they are likely to enjoy, while also helping game studios promote new games to passionate audiences.

COMPANY

Gaming loyalty platform

INDUSTRY

Gaming

REGION

Global

VISION

Build a new mobile loyalty platform to help customers discover recommended game.

STRATEGY

Accelerate time-to-market by introducing a new AWS Analytics & ML platform, data governance processes, and automation.

The challenges

Much like many other startups, the client started their business by testing a hypothesis that naturally evolved into a stable, growing business over the years. Starting from a mobile application with a thin backend layer, it evolved into a sophisticated backend, with data and ML platforms to collect data and recommend games to customers. Starting from a single ML model, over the past decade, the number of models has significantly increased with plans to further increase models twice every six or twelve months.

Modern engineering principles require infrastructure-as-a-code management, continuous integration and continuous delivery, proper monitoring, and a standard approach to work with data assets; in other words: DataOps and MLOps. When it came to the client's data platform, however, the primary challenge was a siloed approach to working with data. Namely:

- No clear understanding of which data pipelines deliver data assets;
- No visibility of data issues that have occurred;
- No visibility of upstream/downstream pipeline dependencies.

Since there were no tools to identify and notify data consumers about data issues, consumers were reporting issues that weren't visible to the engineering team. The data platform didn't have a data catalog, so it was hard to discover data assets in the data platform. Due to these reporting issues, ML model development or analytical queries on top of the data required extra effort, leading to additional costs and impacting time-to-market.

The client's legacy ML platform also had a number of challenges:

- Lack of feature store and transparent process for tracking experiments and artifacts
- Difficulty in identifying who trained a model, as well as how and when it was trained, l what features were used for training, and exactly what artifacts were produced.

This was adding extra maintenance costs and complexities in scaling the number of models: manual handling of two to three ML models is manageable, whereas handling a decade's worth of ML models is significantly more challenging and further impacts time-to-market.

Observing the issues related to the manual development process for data management and ML models, the client came to a decision to build a brand new Analytics platform and ML platform in AWS, integrate them existing data sources and a backend platform serving mobile applications, and introduce reporting on top of the data.

The solution

Enhancements of the already existing platform were more costly than building a platform from scratch, so the client decided to build a brand new Analytics and ML platform in AWS. The client had a list of requirements for a new platform, and was open to new, promising technologies that haven’t become mainstream solutions.

In this case study, we refer to the Analytics platform and ML platform as separate entities, but they were delivered together as a whole solution with a specific focus on data engineering and ML model development.

Analytics platform

Businesses that have worked with legacy architecture for a long time understand what the requirements for a new platform should be. For example, engineers who work with data on a daily basis know which specific tools and processes are missing from legacy systems.

The requirements of a modern Analytics platform, therefore, should include:

- A data catalog with easy data discoverability, a glossary, and ownership tracking.
- Capabilities to support slowly changing dimensions, where needed, to be able to track the history of these changes.
- Reporting integration and the ability to export data to third-party data warehouses like Snowflake should be available out of the box.
- Self-service access to the data: once data is identified through a data catalog, it should be easily queried, with the ability to play around with query results.
- Personal Identifiable Information (PII) data management: Databricks offers column-level access control to limit access to PII data to only a specific group of users. Column-level access was designed as part of the whole security management process, including IAM and service accounts, and is implemented in infrastructure automation scripts.

Further, running data and ML pipelines on a legacy platform can result in issues with data inconsistency and difficulties identifying source data sets or data sources for target data tables that were used in reporting. Even though the client doesn’t have dozens of data sources, the number of sources is predicted to increase drastically over time, so data-source onboarding should be a standard operation requiring minimal effort.

With these requirements in mind, as well as existing workloads, and future plans, Grid Dynamics created the AWS-based architecture shown in Figure 1:

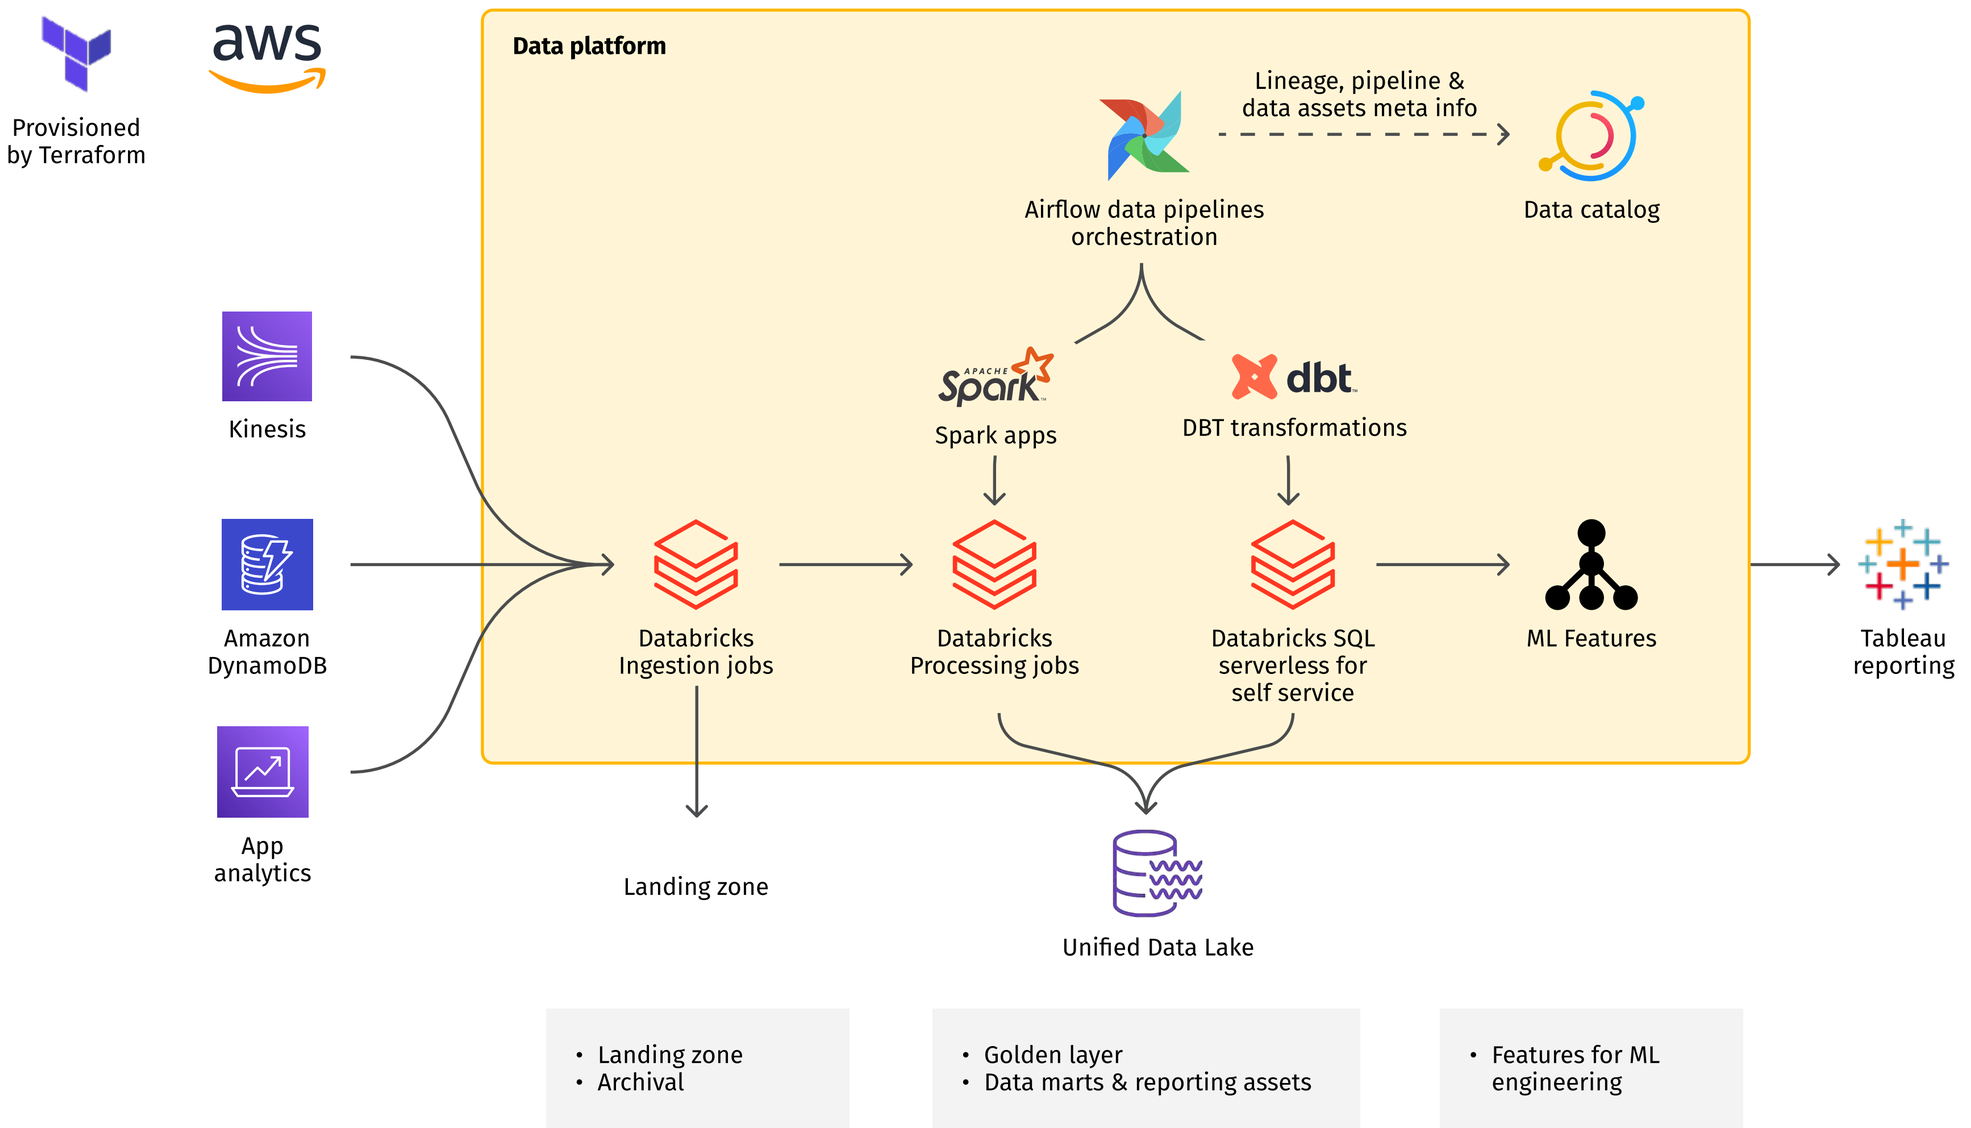

Figure 1: Analytics Platform architecture

The core part of the Analytics platform is Databricks. The compute and storage layers are implemented on top of Databricks, and orchestration is managed by Apache Airflow, which has become an industry standard in pipeline orchestration. Data transformations are implemented on top of DBT which makes it easy to rewrite legacy SQL scripts and introduce testing. DynamoDB is the primary database used to store mobile data, where Kinesis is the streaming data source that was introduced for streaming pipelines.

Today, the market provides many options for data catalogs. Our choice was Datahub: an open source, dynamically growing solution that easily integrates with Apache Airflow to fetch lineage information, with Databricks to track data assets, and Amazon SageMaker to track features created and data transformation jobs.

ML Platform

Compared to the client's fully-functioning legacy data platform, the ML platform was implemented partially with manual operations for artifact creation or releases.

So, according to the three levels of MLOps maturity, we needed to level up:

- MLOps level 0 (Manual process)
- MLOps level 1 (ML pipeline automation)
- MLOps level 2 (CI/CD pipeline automation)

Our goal, therefore, was to build an ML platform managed using infrastructure-as-a-code, with a fully automated CI/CD pipeline, the ability to track release artifacts and data science experiments, and share experiments across the organization. So, the general requirements were:

1. Enable fully automated infrastructure provisioning
2. Create a sandbox environment for data scientists
3. Enable the ability to track experiments, and the results of experiments
4. Enable the ability to track features
5. Enable the ability to track artifacts created, as well as features that were used, and other meaningful information
6. Enable batch and real-time model inference
7. Enable A/B testing support with a wide list of features that are not supported by AWS or Databricks out of the box

Having considered various leaders in the market who provide end-to-end MLOps capabilities, we made the decision to leverage Amazon SageMaker, which provides managed capabilities to run sandbox environments, various types of CPU and GPU instances, an A/B testing framework, and ML observability. The architecture of the implemented ML platform is shown in Figure 2:

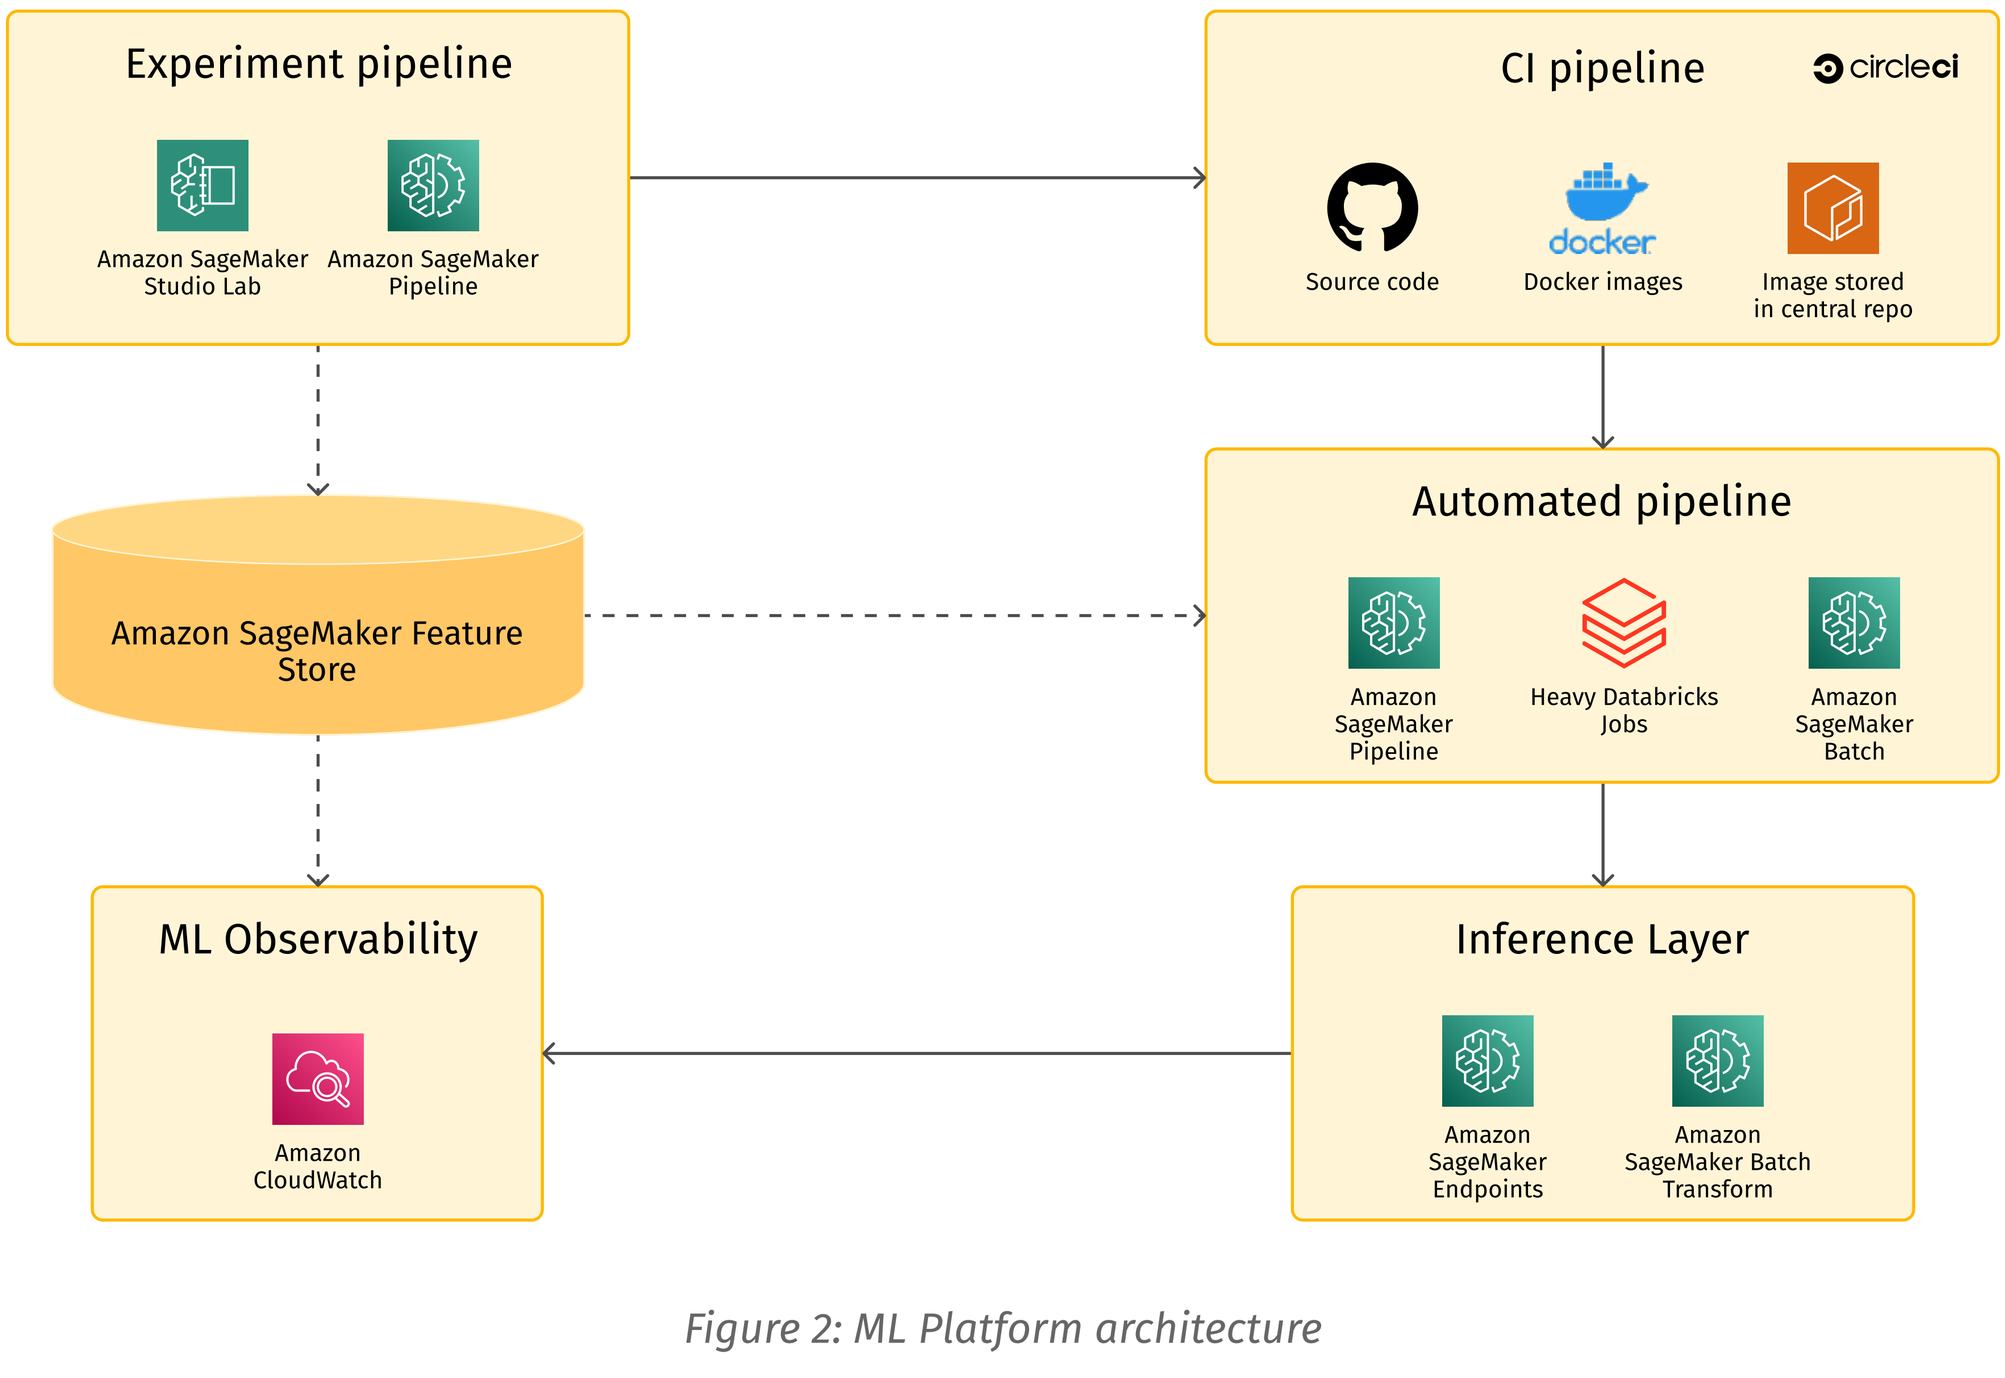

Figure 2: ML Platform architecture

Experiments are run in SageMaker Jupyter notebooks, and once the experiments are successfully completed, the code is committed to Github. After the commit, CircleCI starts the CI process which creates a model and all required dependencies, such as features, that need to be created.

Batch inference can be run in SageMaker or Databricks can be leveraged if needed. Real-time inference supports SageMaker and Amazon Elastic Kubernetes Service (Amazon EKS).

The client has various A/B testing scenarios that need to be supported as well. For example, sticky sessions should be served by specific ML models; i.e. a specific group of customers should be served only by specific models, so there is custom logic implemented during request routing to reach out to needed endpoints.

Infrastructure management

Modern industry practice requires managing infrastructure as-a-code: no manual changes of the infrastructure, configuration, security policies or service accounts used to run applications. Automated infrastructure management helps to avoid human errors and ensures that development environment creation and management is the same, or almost the same, as production environments. This way, during production release, there won't be issues with permissions that exist in the development environment that aren't in place in the production environment.

Our accumulated expertise from previous cloud migrations, enhancements of cloud platforms, and re-platforming programs has helped us build in-house knowledge and starter kits to automate the implementation of platforms for our clients. Automation is implemented for each cloud specifically, and modular elements make it easy to adopt and change according to business requirements.

For this project, we used Terraform automation to provision data and ML platform components. Automation helped to reduce time-to-market and CloudOps efforts required to build automation from scratch. The AWS Analytics Platform Starter Kit and ML Platform Starter Kit that were used for infrastructure provisioning, CI/CD adoption and components integration are available in AWS Marketplace, and can be implemented on top of AWS CloudFormation and Terraform for other hyperscalers.

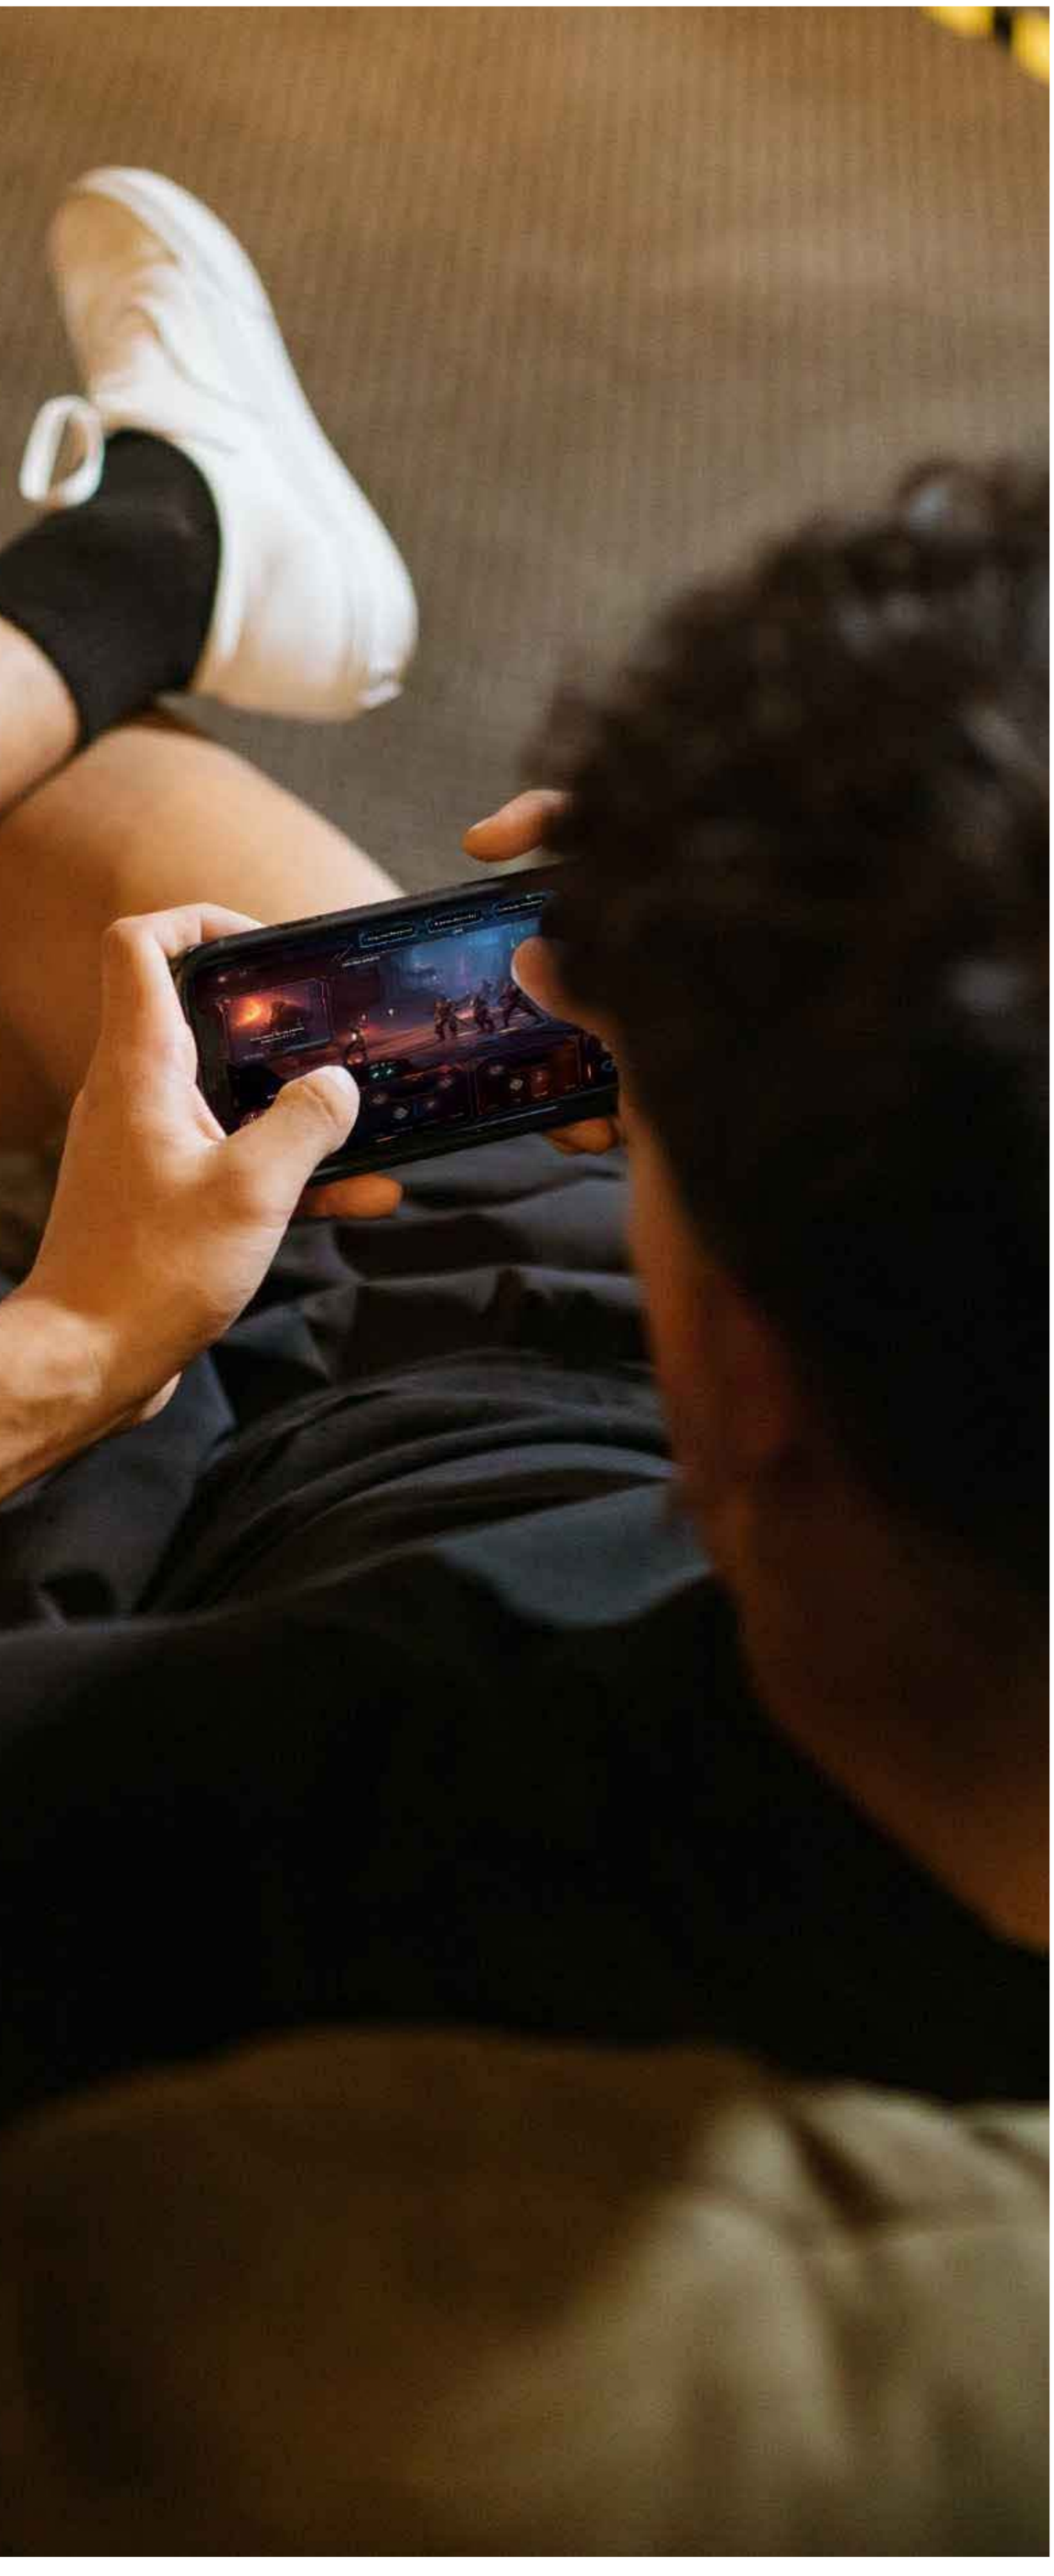

The results

Primary goals for startups include business growth, hypothesis testing, and rolling out new products and services to market as quickly as possible. Extensive growth inevitably leads to technical debt, over-engineering, and complexities that typically fall by the wayside. To prevent these issues from impacting business value and growth, startups generally need help from technical consultants to improve business KPIs such as reduced total cost of ownership and time-to-market.

In less than 6 months, Grid Dynamics helped the client build industry-standard Analytics and ML platforms from scratch, onboarded business users, migrated applications, and prepared for extensive customer base and business feature growth.

The brand new Analytics and ML platform are helping the client begin their expansion of new features, run experiments, and perform releases as they go.

Solution highlights

Use case:

Analytics and ML platform for rapidly growing North American startup

Solution features:

- Fully automated Analytics and ML platform in AWS
- Industry-standard best practices in data and ML engineering
- Out-of-the-box scalability
- User-friendly dashboards for non-engineering teams to work with the data
- Data governance and quality control

Outcomes:

- Reduced time-to-market and support required for new business features
- Technical capabilities to implement advanced features like enhanced customer recommendations or A/B testing
- Ability to integrate with 3rd party Customer360 solutions
- Secured with regulatory compliance

About Grid Dynamics

Grid Dynamics is a global digital engineering company that co-innovates with the most respected brands in the world to solve complex problems, optimize business operations, and better serve customers. Driven by business impact and agility, we create innovative, end-to-end solutions in digital commerce, AI, data, web UI and UX, and cloud to help clients grow.

Headquartered in Silicon Valley, with delivery centers located throughout the globe, Grid Dynamics is known for architecting revolutionary digital technology platforms for 7 of the 25 largest retailers in the US and 3 of the 10 largest consumer goods companies in the world, as well as leading brands in the digital commerce, manufacturing, finance, healthcare, and high tech sectors.

Our secret sauce? We hire the top 10% of global engineering talent and employ our extensive expertise in emerging technology, lean software development practices, a high-performance product and agile delivery culture, and strategic partnerships with leading technology service providers like Google, Amazon, and Microsoft.

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

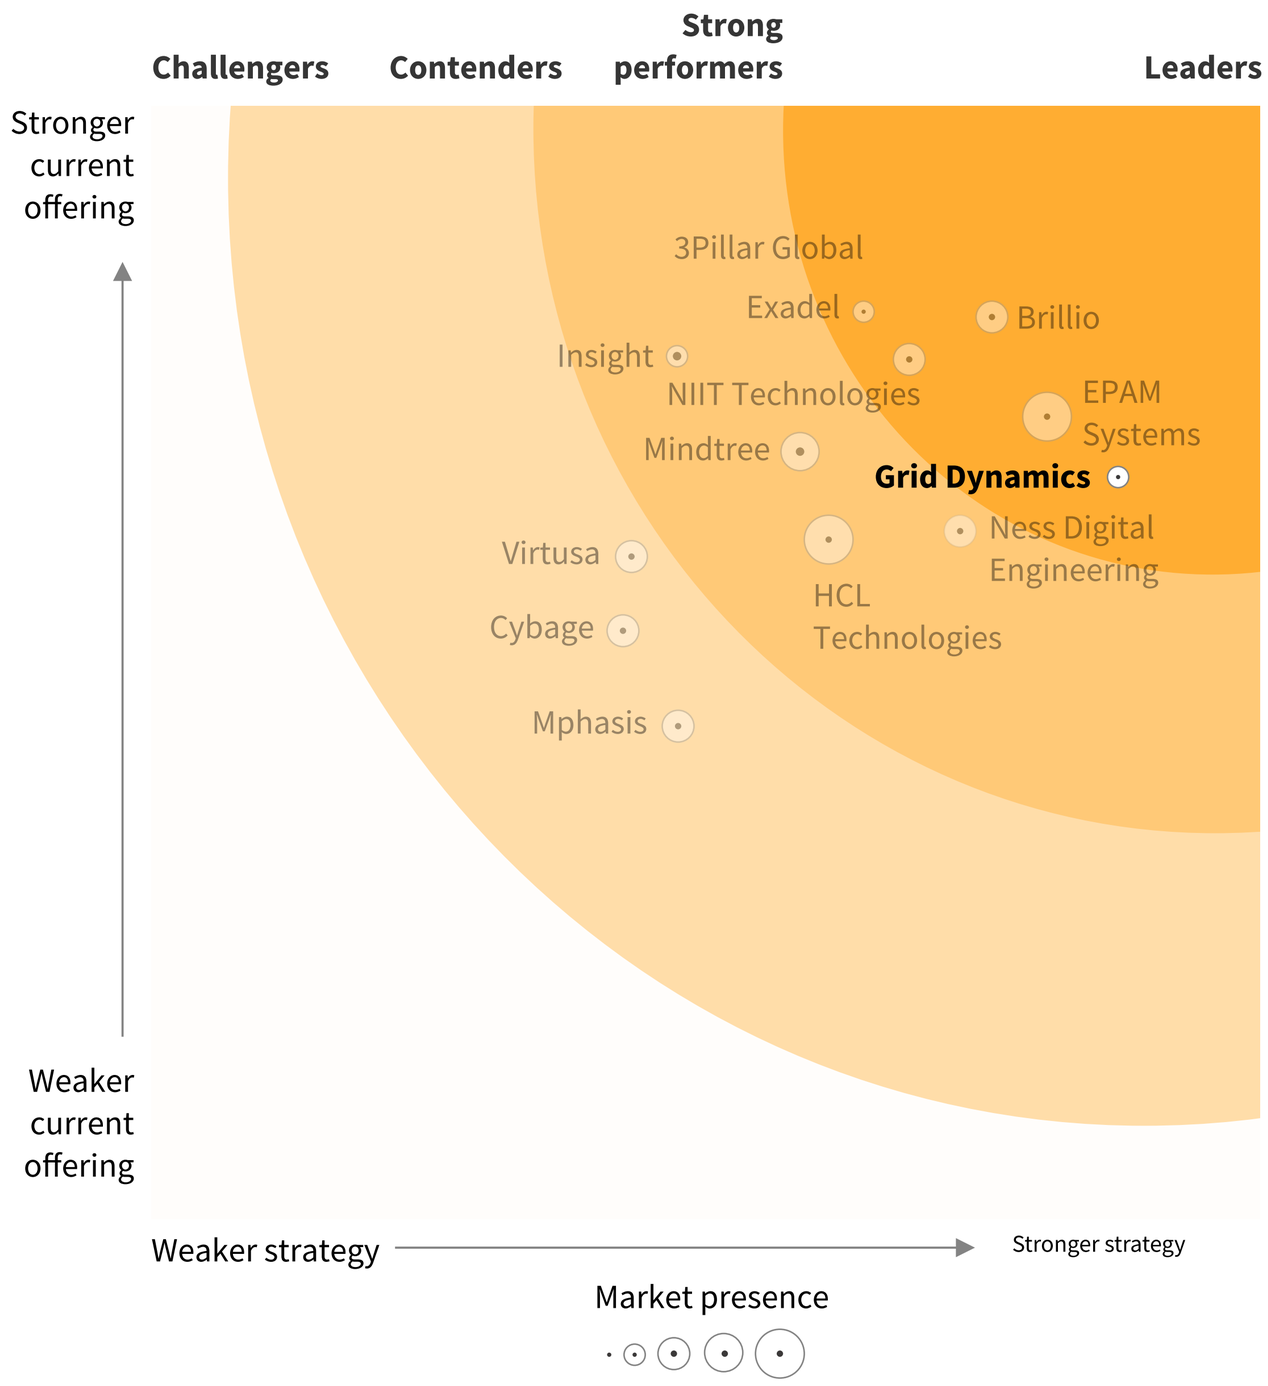

About Grid Dynamics

Key facts

- Offices across the US, Mexico, UK, Netherlands, Switzerland, India, and Central and Eastern Europe
- 4,000 employees in Q2 2023
- Forrester Leader Midsize Agile Software Development Service Provider Q2 2019
- Proprietary starter kits developed in collaboration with AWS, Google Cloud, Microsoft Azure, and others.

Areas of expertise

- **Experience engineering**
  Web UI | Mobile | UX | AR/VR
- **Data Science and AI**
  Search | Personalization | Supply chain | IoT
- **Platform engineering**
  Microservices | MACH | Composable
- **Data engineering**
  Big data | Streaming | MLOps
- **Cloud and DevOps**
  CICD | AIOps | SRE | QE

Clients

	Google	JABIL
align	Levi's	StanleyBlack&Decker
RAYMOND JAMES	fiserv.	AMERICAN EAGLE
Boston Scientific	INSPIRE Brands	HELLO FRESH

Grid Dynamics

powered by aws

Grid Dynamics Holdings, Inc.

5000 Executive Parkway,

Suite 520 / San Ramon, CA

650-523-5000

www.griddynamics.com