DATAOPS FOUNDATION

Flexible, scalable data lake for centralizing, exploring, and reporting enterprise data

Transforming your organization to be more data-driven takes an evolution both in technology and culture. Data initiatives die without a guiding business objective. While every organization is unique, everyone wants to see and act on trustworthy data. That's where DataOps – a business' culture and process around access to clean, actionable, trustworthy data - comes into play.

2nd Watch's DataOps Foundation is a 4-8 week engagement that starts and ends with business objectives, hardened by years of industry best practices for cloud data lakes, yet flexible enough to meet your unique needs around performance, access, security, and governance. 2nd Watch helps enterprises transform their data analytics processes to:

Make better, data-driven decisions	Trust your dashboards and reports	Confidently bring in new sources for enhanced analysis	Create a culture of DataOps – broad access to clean, actionable data	Become AI-ready!

OUR SOLUTION

BUSINESS-DRIVEN

Our solution gets you moving towards business goals on day 1. We keep your business needs front and center, framing the conversation around the analytics, metrics, and reports you need to drive business goals.

- Consultants trained to keep your business goals front-and-center
- Flexible about underlying tools, but insist on surfacing actionable data
- Fresh insights across previously-siloed data sets
- Minimal maintenance – we connect your datasets so they automatically refresh
- Access control – we use LakeFormation to secure your data using sensible role definitions like Business Analyst and Data Engineer
- Works with your existing identity store – Okta, ActiveDirectory, etc.

SCALABLE, CLOUD-NATIVE, FLEXIBLE

Our solution uses a combination of AWS services like AWS Glue, S3, Key Management System, and Athena, and managed services like Redshift and RDS. These are all cloud-native and require no infrastructure maintenance, leaving you and your team free to create business value instead of managing servers.

- Not a single EC2 instance to manage
- Your stack is modern and stays modern
- Scales to hundreds or thousands of users and datasets
- Plugs in to any BI tools
- Prepares you for advanced Machine Learning
- Handles structured, semi-structured, and unstructured data of any volume

BEST PRACTICES

Our ready-to-deploy data lake templates capture years' of cloud experience and best practices, with integration from governance to data exploration and storage. We explain the reasons behind the decisions and make changes based on your requirements, while ingesting data from multiple sources and exploring it as soon as possible.

- Classic 3-zone data lake on S3
- Data is immutable by default – you'll never lose data
- Data is automatically compressed and archived to minimize storage costs
- Convenient search with always-up-to-date data catalog – easily discover all your enterprise datasets
- Data lineage – automated notification when data schemas "drift" / when data sources suddenly change format
- Only the most important "data marts" – approved datasets – get loaded into more costly Redshift or RDS, minimizing costs and complexity