AI FOUNDATION OVERVIEW

viaduct

AI foundation built for manufacturing

Turning data into decisions you can trust

Choosing the right foundation

AI is talked about as something that can fix almost any problem. But not all AI is the same. Each is built for different needs, requirements, and outcomes.

That distinction matters most in manufacturing, where data is complex, messy, and fragmented. Turning it into insights requires an understanding of how the data can work together to guide the decisions that are made every day.

A strong AI foundation turns your data into early anomaly and defect detection, fewer warranty claims, higher productivity, and better asset performance. However, the opposite can be costly and a major setback for an organization.

With so many options, choosing the right AI foundation can feel daunting. That's one reason so many AI projects in manufacturing stall before they deliver value.

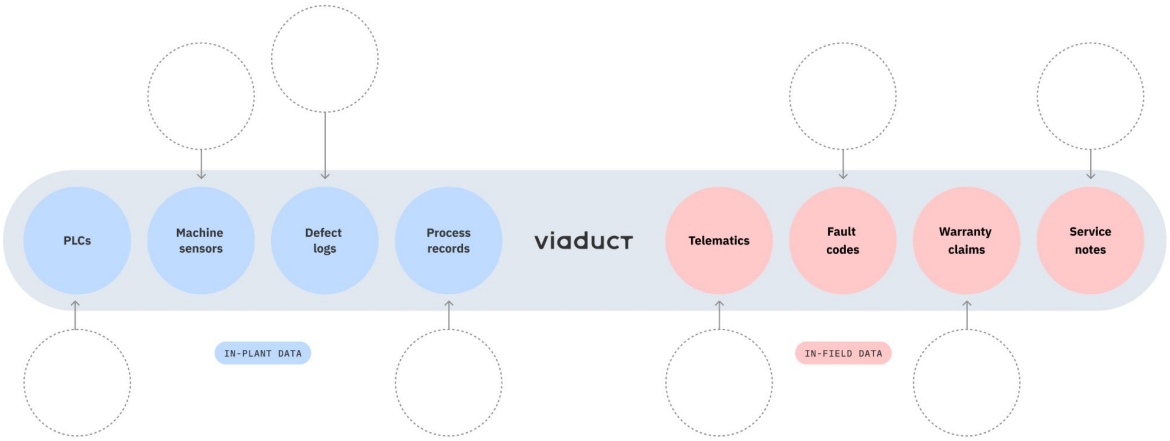
This document outlines what sets Viaduct's AI foundation apart and why it's built for the data and decisions manufacturers face:

- Single view of your manufacturing operations
- Multivariate time-series AI at the core
- AI that understands your language
- Predictive models tested, validated and proven
- GenAI that puts reliable answers in everyone's hands

Single view of your manufacturing operations

Bringing all of your data into one place is often where AI projects stall before they start.

Viaduct's flexible schema handles any data source, regardless of type or complexity, so teams can onboard new data sources themselves. Viaduct cleans and standardizes the data, eliminating manual data prep and speeding up the onboarding process. From there, teams can model the relationship among data sources, such as which work orders correspond to which assets or which fault codes map to which warranty claims.



The result is a single view that connects in-plant data (PLCs, machine sensors, defect logs, process records) with in-field data (telematics, fault codes, warranty claims, service notes). With everything in one place, teams can see why the same failure keeps happening, trace it to the root cause, and trust the insights are accurate and reliable.

Why it matters:
IT will spend less time building and maintaining custom integrations, and manufacturing teams can manage the data pipelines and reporting requests on their own.

Multivariate time-series AI at the core

In manufacturing, uncovering issues means understanding cause and effect across many signals at once. That takes AI that can analyze many variables together and track how they change and influence one another over time.

Viaduct's patented multivariate time-series models are trained natively on manufacturing and connected-asset data to identify trends, rates of change, and how signals interact. By combining these models with LLMs and statistical machine learning, Viaduct delivers analyses and predictions in minutes that would otherwise take weeks or months of custom data science work.



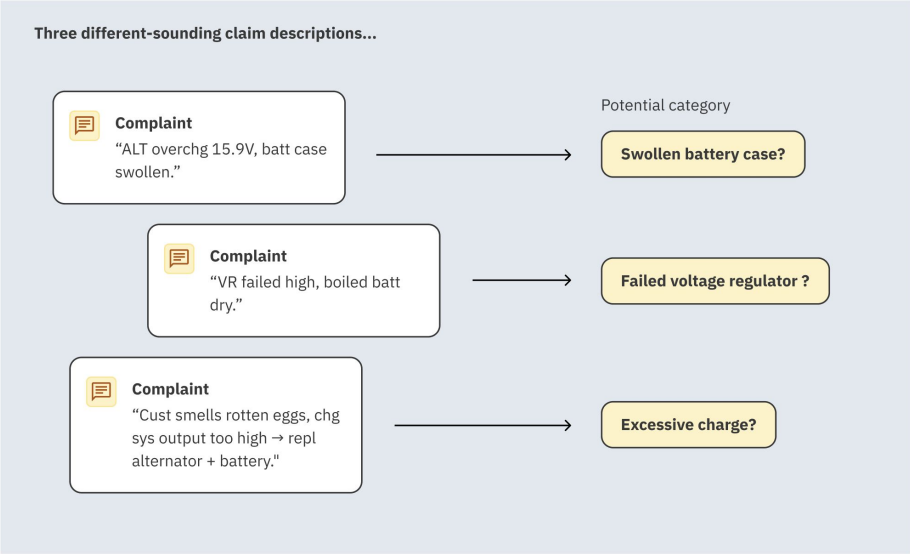
Why it matters:
By finding patterns across all signals, Viaduct highlights operational issues and anomalies that a generic model would miss.

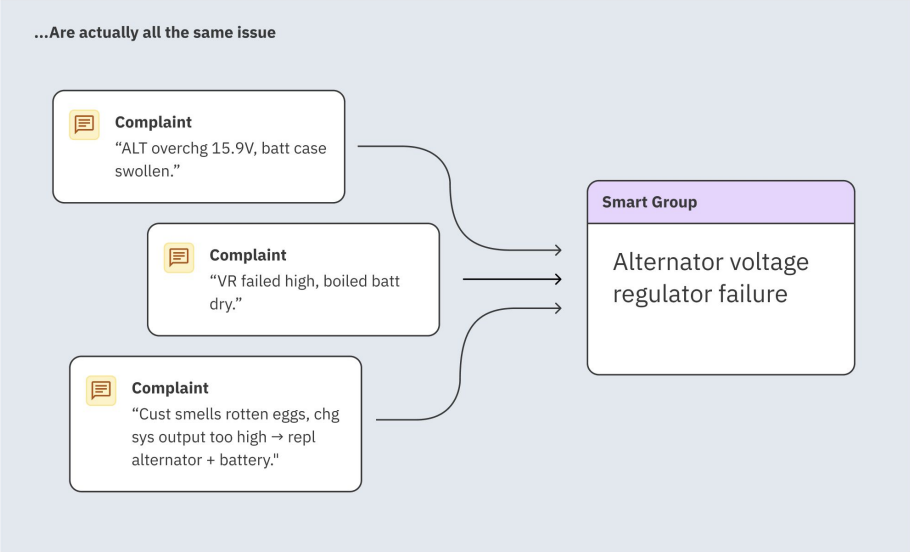
AI that understands your language

Every manufacturer has its own language, and even within the same company, people describe the same issue in different ways.

To most systems, those are three different problems. That makes it hard to see how widespread an issue is, which slows down response and resolution time.

Viaduct's Smart Groups automatically recognize these descriptions refer to the same issue, even with shorthand, typos or missing details, and group them together. These groups stay consistent so everyone is speaking the same language. On the other hand, generic, off-the-shelf AI can provide different answers every time, making results messy and hard to act on.





Why it matters:

Viaduct understands the context of your data so the insights and recommendations provided are both accurate and actionable. Additionally, Viaduct handles messy, inconsistent data as it comes in so IT doesn't have to spend time on data prep or data quality work.

Predictive models tested, validated, and proven

A predictive model is only as good as the accuracy of its predictions. Getting it right takes time, resources, and testing. It requires data science teams to define the failure, identify the data sources, test the models, and validate the results against what actually happened.

In-house teams often struggle to do this on their own. Data is spread across disparate systems, time-series modeling requires expertise, and maintaining them can be resource intensive and time consuming.

Viaduct's models are built and validated through a five-step process: scoping the failure, analyzing the data and engineering features, training and tuning, validation, and deployment. When benchmarked against in-house models, Viaduct's models are 44% more accurate and catch 46% more failures.

Viaduct continuously evaluates model performance by running manufacturing-specific tests on its workflows to ensure reliability and accuracy.

Why it matters:
Teams can trust that the predictions are accurate and repeatable, and can confidently act on the recommended next steps. IT doesn't need to build, validate, and maintain predictive models in-house or hire specialized data science talent to do it.

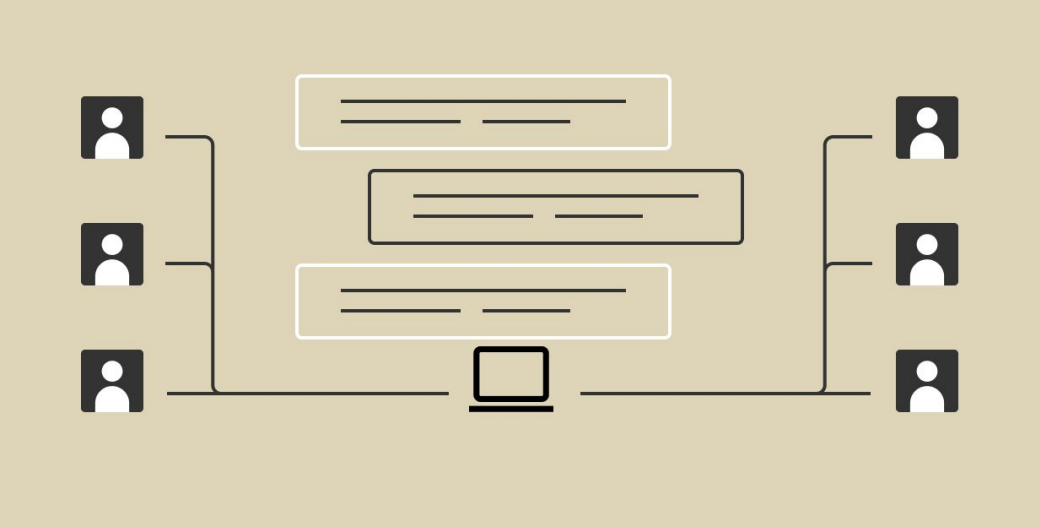
44%

Viaduct's models are 44% more accurate and catch 46% more failures.

GenAI that puts reliable answers in everyone's hands

For most manufacturers, getting answers about their data means waiting on IT or trying to figure it out on your own. Either option can take weeks and by then the problem has only gotten worse.

GenAI can provide answers, but they are not always reliable or consistent. ViaDuck, Viaduct's AI agent, addresses that by turning plain language questions into complex queries that run directly against your data. It gives LLM models the metadata they need to understand your data so the answers are based on fact rather than generated from scratch, and queries can be saved and rerun for repeatable results



Why it matters:
Users get answers to common questions in seconds in their preferred format (charts, dashboards, or raw data points), which can also be sent automatically. Meanwhile, BI and IT teams spend less time on routine requests and more time on high-value work.

How it works

Viaduct works with your MES, QMS, and ERP systems to get more value from them. If you have a data platform, Viaduct’s AI layer sits on top.

What sets Viaduct apart isn’t any single capability. It’s how Viaduct's AI foundation works to deliver manufacturing intelligence across your organization. The result is earlier detection, faster root cause analysis, fewer failures, and lower warranty costs.

About Viaduct

Viaduct provides a purpose-built AI platform for discrete manufacturers and fleet operators to improve productivity, quality, and asset performance. Founded in 2019 at the Stanford AI Laboratory, Viaduct is headquartered in Sunnyvale, California, with offices in Kobe, Japan, and Ljubljana, Slovenia.



The company's patented AI models are trained on large-scale time-series manufacturing and connected asset data to detect patterns and make predictions at scale, helping manufacturing, quality, and engineering teams detect emerging issues, identify root cause, and predict failures before they become costly disruptions.

Viaduct was acquired by Sumitomo Rubber Industries in 2025 and now operates as a wholly owned subsidiary. To learn more, visit www.viaduct.ai