Nuvalence

ARCHITECTURE GUIDANCE SERIES

Built for the Long Haul: Architecting an Intelligent Fleet Management Platform

WHITEPAPER

Purpose & Audience

Fleet managers are poised to lead technology adoption in the mobility and transportation space. With telematics, geographic information systems (GIS), and mobile apps already standard features of any commercial car, advances in augmented reality, artificial intelligence, and machine learning are showing potential to revolutionize fleet optimization, driver safety, and cost reduction.

Today's economic pressures, market competitiveness, and customer expectations are all too high for organizations with big fleets to not invest in the future. At the same time, the road to modernization can seem more like an endless journey requiring constant, wide-scale retraining. Where to begin?

What's required first, is a shift in perception. In order to realize fast results from its technological investments while also gaining the adaptability to position itself for long-term agility and integration of emerging tech, it's important for today's fleet-reliant organization to view its systems as components of a reimagined platform. With AI not a mere add-on but playing a vital role at the heart of its architecture, this modern fleet management platform won't age like a legacy system, but evolve.

The purpose of this whitepaper is to propose a practical framework for the strategic architecture, design, and execution of an intelligent fleet management platform. With AI at its center, this platform will help the organization transform mobility with efficiency, elevating customer experiences and streamlining the fleet manager's user experience. The intended audience for this whitepaper is a technical leader within the automotive or general mobility sector. Individuals in this audience usually have Chief Technology Officer, Head of Digital Platforms, Chief Software Architect titles, or similar backgrounds.

Nuvalence specializes in architecting and developing AI platforms and products designed to create strategic value for fast impact while enabling continuous innovation for long-term results. The details in this paper are based on real-world experience in mobility and structured thinking on the general topic of building fleet management platforms.

While it's helpful to have familiarity with platform architecture and data-driven architecture, any technical leader will derive valuable insights from this paper. To learn more about these concepts, we invite you to explore companion pieces in our Architecture Guidance Series, especially The Anatomy of a Digital Platform[1], which defines the technology model for this approach. For an industry-specific complement to this paper, Architecting a Transportation Cloud[2] focuses on the type of modern, backend platform necessary for an organization to take advantage of connected vehicle data.

How Intelligent Platforms Support the Fleet Management Lifecycle

Fleet managers oversee an organization's vehicles, vehicle operators, hardware, additional staff, and other relevant data. Organizations with typically large fleets and comprehensive management needs include:

- **Freight companies** and delivery services whose fleets transport goods to consumers.
- **Government agencies and municipalities** whose fleets include police vehicles, plow trucks, parking enforcement vehicles, and other city or county-owned vehicles.
- **Service providers** such as plumbers, electricians, or maintenance specialists that provide employees with vehicles equipped with tools and devices required to perform their duties.

1 "Generalizing the Architecture of a Digital Platform," Nuvalence, 2022: https://nuvalence.io/white-papers/generalizing-the-architecture-of-a-digital-platform/

2 "Architecting a Transportation Cloud," Nuvalence, 2022: https://nuvalence.io/white-papers/architecting-a-transportation-cloud/

Fleet managers in any of these areas lose money if their vehicles are not on the road, routes are not optimized to offset rising fuel costs, or they fail to meet the demands of their respective customers and constituencies. To meet their business goals, they must keep several key performance indicators (KPIs) in mind through cyclical periods of planning and execution. While organizations may value and prioritize each area differently, the fleet manager will watch metrics such as **cost**, **time spent**, **compliance**, and **utilization**.

The fleet management lifecycle models the manager's business processes. This may differ from company to company, but based on our experience and knowledge of fleet manager's business processes, this lifecycle will consist of seven stages:

Exhibit 1: *The fleet management lifecycle*

Planning: forecasting and strategizing for the coming phases of operations.

Growth: acquiring the team and resources needed to execute the plan.

Dispatching: assigning suitable operators and vehicles to perform jobs.

Monitoring: ensuring the fleet is performing as expected.

Maintenance: scheduling routine maintenance, upkeep, and compliance; responding to unplanned breakdowns and outages.

Reporting: producing retrospective statements on KPIs and overall fleet performance.

Optimization is the pivotal lifecycle stage in which the fleet manager identifies areas for improvement from the outputs of prior stages before entering the next iteration of planning. The primary goal of a fleet management platform is to drive this optimization in accordance with the KPIs that guide management – offering features such as vehicle maintenance tracking, real-time location monitoring, fuel usage monitoring, and expense and finance tracking.

AI offers a prime opportunity to optimize the fleet manager's user experience across this lifecycle, which in this case also represents our user journey. Focusing specifically on the service provider persona mentioned above, let's examine the fleet manager's goals at each lifecycle stage, and how AI can help accomplish the goal. It will be helpful to organize this exercise within a framework. In our experiences with generative AI and product management, Nuvalence has identified six taxonomic categories[3] that most organizations can use to classify AI use cases and structure such analysis.

This taxonomy can be visualized as concentric rings depicting a progression of the involvement of AI in an activity. Advisory, the innermost ring, describes interactions where generative AI acts as an expert sharing information with users, building to Cooperative, where generative AI works alongside the user as their copilot. One of the powerful aspects of this taxonomy is that it highlights the potential for generative AI and traditional machine learning (ML) to cooperate and augment each other's capabilities, creating the best leverage from fleet data and AI advancements.

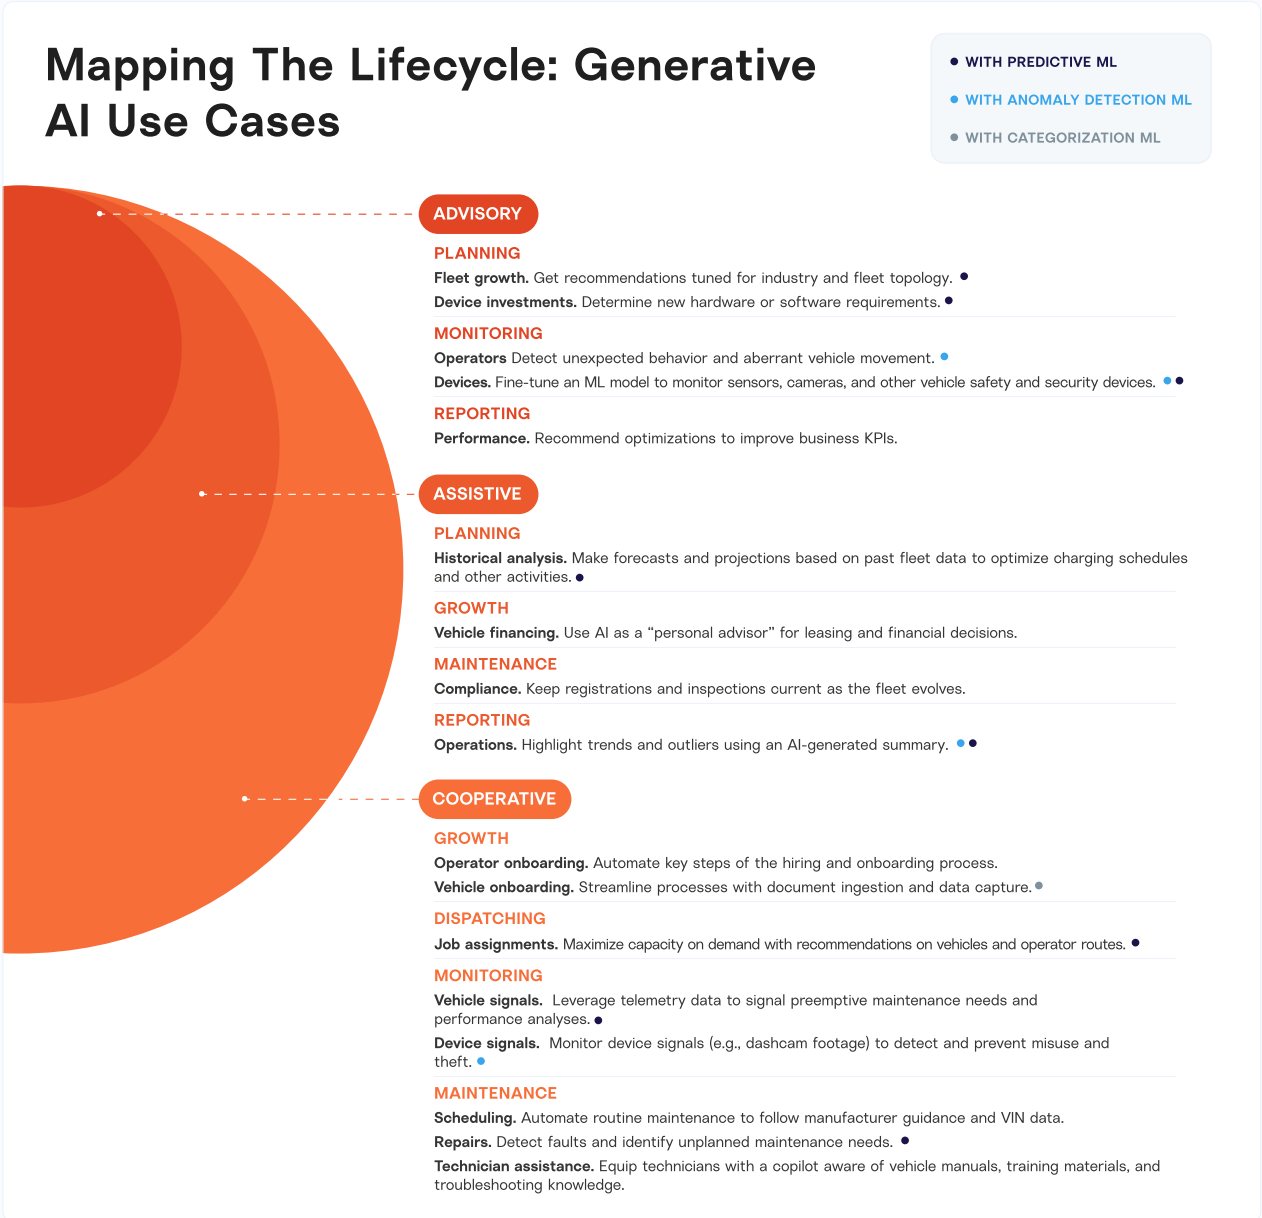

Exhibit 2: *Using a taxonomy to map lifecycle stages to opportunities for intelligent improvements*

[3] "A 6-Category Taxonomy for Generative AI Use Cases," Nuvalence, 2023: https://nuvalence.io/insights/a-6-category-taxonomy-for-generative-ai-use-cases/

Architecting the Intelligent Platform

To enable all of these AI-centric use cases for fleet managers, the platform must provide a foundation for ingesting relevant data, managing common workflows, and providing insights.

Key Considerations Driving Architectural Components

The operating context, as defined by the personas and interactions in the fleet management lifecycle described above, will strongly influence architectural decisions for the platform. This context includes extensibility – a primary feature of the intelligent platform being its ability to adapt to other data streams, such as those from service providers, regulatory agencies, and other internal and external partners that impact fleet operations. This section highlights considerations that are core to the fleet management experience, as well as a supporting architectural component for each.

Consideration	Component
Fleets generate data – from vehicles, devices, and operators. Ingesting, normalizing, and aggregating this data is a key first step toward building any fleet management platform. The data is essential for fleet managers in the lifecycle's monitoring and reporting phases.	**Data Ingestion** enables partners to quickly and efficiently provide fleet asset data to the platform.
Intelligence requires data. To train or fine-tune any AI/ML models, the platform must first be able to cleanse and process data into the format needed to derive value.	**Data Services** enable products and the platform to design and build data pipelines while ensuring that both raw data and processed data are available and cataloged for future use cases.
Fleet managers are always looking for opportunities to optimize. As key decision makers, fleet managers need actionable insights on high costs that can be reduced, manual processes that can be streamlined or automated, and above all, ways to reduce vehicle downtime.	**Artificial Intelligence and Machine Learning (AI/ML) Services** enable the platform (and products built on the platform) to leverage data, analyze trends, and suggest product, vehicle, and resource optimizations.
Fleet management is often a supporting operation of a business. Fleet managers need to quickly and efficiently manage their vehicles and operations so they have time to support broader business needs.	**Reporting Services** enable partners and users to build dashboards and report on key performance indicators for their business.
Fleet management isn't just about vehicles. While vehicles are a key component, fleet managers are also tracking operators of the vehicles, as well as devices that are key to business operations. Understanding different fleet topologies will also be necessary to categorize and refine AI recommendations.	**Topology Services** enable partners to catalog metadata about different assets to provide context for data ingestion.
Fleet management processes can be physical and have tangible ramifications, often requiring a human-in-the-loop (HITL). Different applications of AI/ML can identify when a vehicle requires maintenance or when a vehicle's registration is overdue, but the fleet managers and operators must be engaged in the overall process to validate the decisions and carry out the required physical actions.	**Workflow Services** enable partners to build cooperative capabilities that accelerate standard processes for fleet managers.

Reference Architecture and Architectural Component Deep Dive

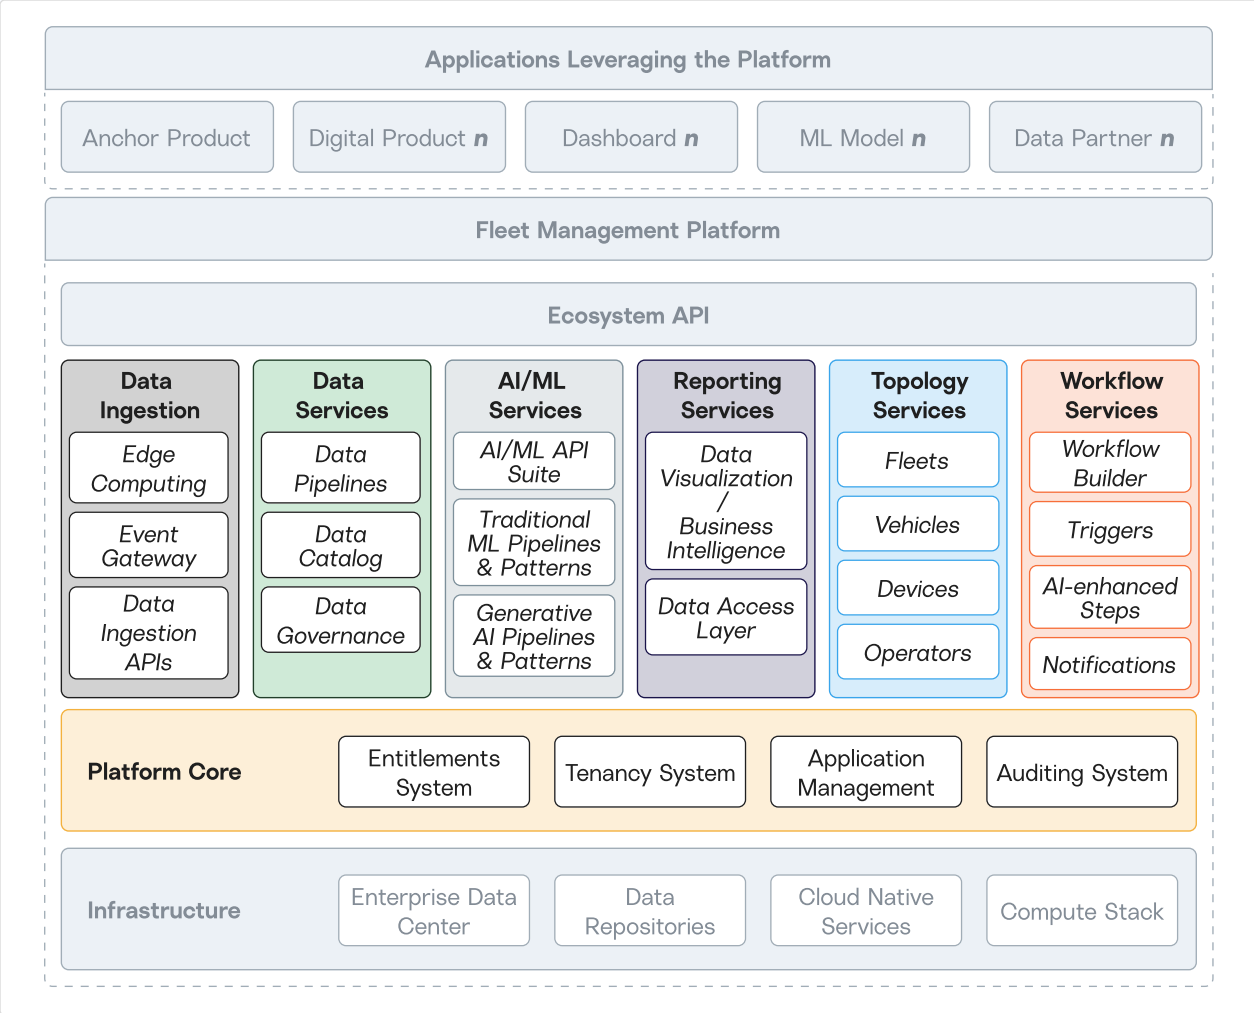

Exhibit 3: *The fleet management platform's generalized architecture and components*

PLATFORM CORE

In any platform implementation, several use-case-agnostic core systems exist to ensure the scalability and security of all of the use-case-specific capabilities to come. A comprehensive list of such capabilities can be found in [The Anatomy of a Digital Platform](#).[4] The core components of any digital platform often include:

1. **Entitlements System:** responsible for tracking access control logic and making authorization decisions.
2. **Tenancy System:** responsible for tenant-level tracking and isolation – partitioning partner application and customer data to ensure data security for the platform.
3. **Application Management:** responsible for cataloging and managing the applications that exist on-platform – both the applications that comprise the platform (described here) and products built on the platform.
4. **Auditing System:** responsible for keeping an audit record of important activities on the platform. In the fleet management domain, this may include changes to the fleet topology (new vehicle added, vehicle removed), compliance records, job status events for service providers, etc.

[4] "Generalizing the Architecture of a Digital Platform," Nuvalence, 2022: https://nuvalence.io/white-papers/generalizing-the-architecture-of-a-digital-platform

On top of this core, the platform's generalized architecture can be broken into several subsystems, each supporting one of the considerations described in the previous section.

DATA INGESTION

For the platform to be effective, it must be able to ingest a large amount of data. The Data Ingestion services are responsible for making it easy to consume new types of data as we see an increase in the number of platform integrations.

1. **Edge Computing:** can greatly reduce the volume of data and any latency of data processing through data filtering and computations performed on vehicles and other devices, before transmitting that data to the cloud.
2. **Event Gateway:** responsible for ingesting raw data from connected devices and directing the data to the appropriate data services in order to normalize, transform, and persist it.
3. **Data Ingestion APIs:** could be real-time APIs, batch APIs, etc. Most data ingested by the fleet management platform will likely be time-series data; such as vehicle events, device signals, workflow events, or product usage data.

DATA SERVICES

Any artificial intelligence requires training, so Data Services are a critical pillar of an intelligent platform. This subsystem includes:

1. **Data Pipelines:** enrich data that the platform has ingested. Use cases include providing additional context about the asset as defined by some static data set, calculating real-time KPIs, or processing data with an AI/ML API to produce new or enhanced data. Each new class of data introduced to the platform likely requires a data pipeline, so it's essential to have common patterns defined and documented to accelerate building new pipelines.
2. **Data Catalog:** responsible for cataloging all data assets on the platform. For any developer to effectively build data pipelines, ML models, or APIs on the data, they need a data catalog to tell them what assets exist.
3. **Data Governance:** includes policies and a framework for ensuring data security, compliance, quality, and data handling protocols for the platform. For example, a data governance policy may require that all data is anonymized before it is used for training.

AI/ML SERVICES

Artificial Intelligence and Machine Learning Services reduce the barrier to entry for building a product enhanced with artificial intelligence.

1. **Traditional ML Pipelines & Patterns:** processes for gathering and preparing data (beyond what was done by data pipelines) for machine learning, extracting features, training models, and serving them into production. Cleansed, anonymized, and categorized data may be worth storing and cataloging in the data catalog to enable partners building on the platform to use it in training. As with data pipelines, identifying and documenting common patterns is likely advantageous to accelerate building new ML pipelines. A few patterns already appearing in discussing the fleet management lifecycle include:
   a. *Predictive ML* – patterns for identifying trends and forecasts in a data stream, and optionally correlating that with other data. For example, identifying growth in jobs and vehicle utilization may result in purchasing decisions to grow the fleet.

b. *Anomaly Detection ML* – patterns for detecting when a data point is abnormal in a stream of similar data points. For example, an on-vehicle sensor begins reporting an abnormal value and should be inspected or preemptively replaced.

c. *Categorization ML* – patterns for attributing one or many categories to a piece of data from a finite set of categories. In theft prevention, this might look like categorizing into low, medium, and high thresholds for suspicious behavior to determine when an operator should review.

2. **Generative AI Pipelines & Patterns:** these patterns will build upon the traditional ML Pipelines & Patterns and may frequently change, as generative AI is relatively new and fast-moving. This practice might include:

   a. *Prompt Engineering Patterns* – proven resources and templates for prompts to create the best outcome for the end user. Many best practices are shared across verticals, but some specific patterns may be developed that use a broadly trained large language model (LLM), and focus it on fleet management personas and topics.

   b. *Fine-tuning* – pipelines to customize an LLM focus on a particular task or use case by tuning the LLM with expected prompt/response pairs.

   c. *Retrieval Augmented Generation* – design patterns for improving the quality of LLM responses by decorating a prompt with additional data and reference materials. This content could range from device manuals and training materials to database schemas and API responses, supporting a range of use cases requiring different data types.

   d. *LLM Orchestration* – a platform service that simplifies the process of orchestrating a set of actions that result in a high quality response from an LLM. This can include gathering data to augment the prompt; parsing generated responses to inform other system actions; or combining prompts, knowledge, and actions in new and compelling ways using other mechanisms.

3. **Tenancy in ML Models:** when training models for a platform, it's important to understand and consider how tenancy comes into play. Just as the platform may segment data for each tenant, there are different patterns for designing models to ensure tenant data privacy:

   a. *Isolated* – the model is trained with data for a specific tenant and used only for that tenant. This approach requires tenants to have amassed significant data on the platform before they can leverage AI/ML insights.

   b. *Shared* – the model is trained with all tenants' data and used for all tenants.

   c. *Tuned* – a generic model is trained with tenant-specific data to tune for that specific tenant. The generic model may be a public or shared model built with platform data.

4. **AI/ML API Suite**: ensures that any developer can leverage trained models via API, even without first understanding how to build and use them. APIs could include AI encapsulation of pre-trained models (such as document parsing, chat, etc.) or models trained on the platform.

REPORTING SERVICES

The fleet management platform is rich in data, but fleet managers seek timely and targeted insights. The Reporting Services subsystem encapsulates the systems required to generate these views of data for end users. This is covered in greater detail in our whitepaper on data-driven platform architecture[5].

[5] "Why Data Needs A Platform," Nuvalence, 2022: https://nuvalence.io/white-papers/why-data-needs-a-platform/

1. **Data Visualization & Business Intelligence (BI):** the presentation layer for this platform subsystem, the BI services gather data from relevant sources and produce dashboards and reporting relevant to the business. In an intelligent platform, the AI can make data correlations that humans cannot or would, at a minimum, take a long time to accomplish. It can also expand capabilities of the BI suite to auto-discover data sources from infrastructure, or suggest new reports and insights based on the available schemas.
2. **Data Access Layer:** APIs for managing and accessing analytical data, such as data for reporting and dashboards. This service could be leveraged by consumers who need to combine data insights across domains in interesting ways, but would not be suitable for transactional use cases that would make up the majority of on-platform traffic.

TOPOLOGY SERVICES

These domain-aligned services maintain an intrinsic hierarchy of assets, used for contextual data, authorization decisions, and data privacy partitioning.

1. **Fleets** are the groups by which fleet managers organize their vehicles. A large business may have multiple fleets representing different vehicle classes or geographies.
2. **Vehicles** may be grouped into fleets by fleet managers.
3. **Operators** are the people who operate vehicles; fleet managers may associate operators with one or many vehicles or fleets.
4. **Devices** (such as dash cams, charging stations, or sensors) may be associated with one or many fleets, vehicles, or operators.

In addition to the topology of how these assets relate, the assets will have data needed to inform decisions made in various workflows and AI use cases on the platform. A domain-driven design (DDD) approach can be taken to align the platform to the assets modeled in Topology Services, as well as the attributes and events that define those assets that will be available to integrated applications.

WORKFLOW SERVICES

These services provide the building blocks for developers to create workflows for manual processes performed by fleet managers and operators.

1. **Workflow Builder:** enables developers to define new workflows for the platform. The workflow builder provides common steps and triggers, giving developers an easy plug-and-play experience for defining processes as workflows. This service is beneficial because it can inject logic to feed workflow data back to the platform for reporting or training ML models.
2. **Triggers:** encapsulate logic for triggering workflows in a reusable way. Including, but not limited to:
    a. Triggers on recurring intervals (e.g., weekly).
    b. Triggers at a specific date and time (e.g., January 2 at noon EST).
    c. Triggers in response to a piece of data ingested or emitted by data pipelines (e.g., on a detected anomaly or AI-suggested action).
    d. Manual triggers
3. **AI-Enhanced Steps**: leverage the AI/ML services platform to automate what might otherwise be a manual or time-consuming step. For example:
    a. Leverage a document parsing API to gather data from an image rather than requiring users to input data manually.

b. Leverage past approval data to predict the outcome of manual approval steps, then enable automatic approval/denial at certain confidence thresholds.
c. Leverage anomaly detection to flag suspicious or abnormal activity in workflows and predict when a part will fail and require service.
d. Leverage an LLM to give fleet managers and operators insights, recommendations, and referenceable supporting content.

2. **Notification Service:** communicate with users at critical points within workflows. For example, when manual activity needs to occur (review, approval, vehicle maintenance, etc.).

What the Intelligent Architecture Delivers

Accelerated Application Development

Using this architecture, an anchor application could be a dispatching application that enables fleet managers to assign operators and vehicles to perform jobs. Building such an application on the platform enables developers to focus on the dispatching functionality and leverage the platform for workflow management, data, and insights.

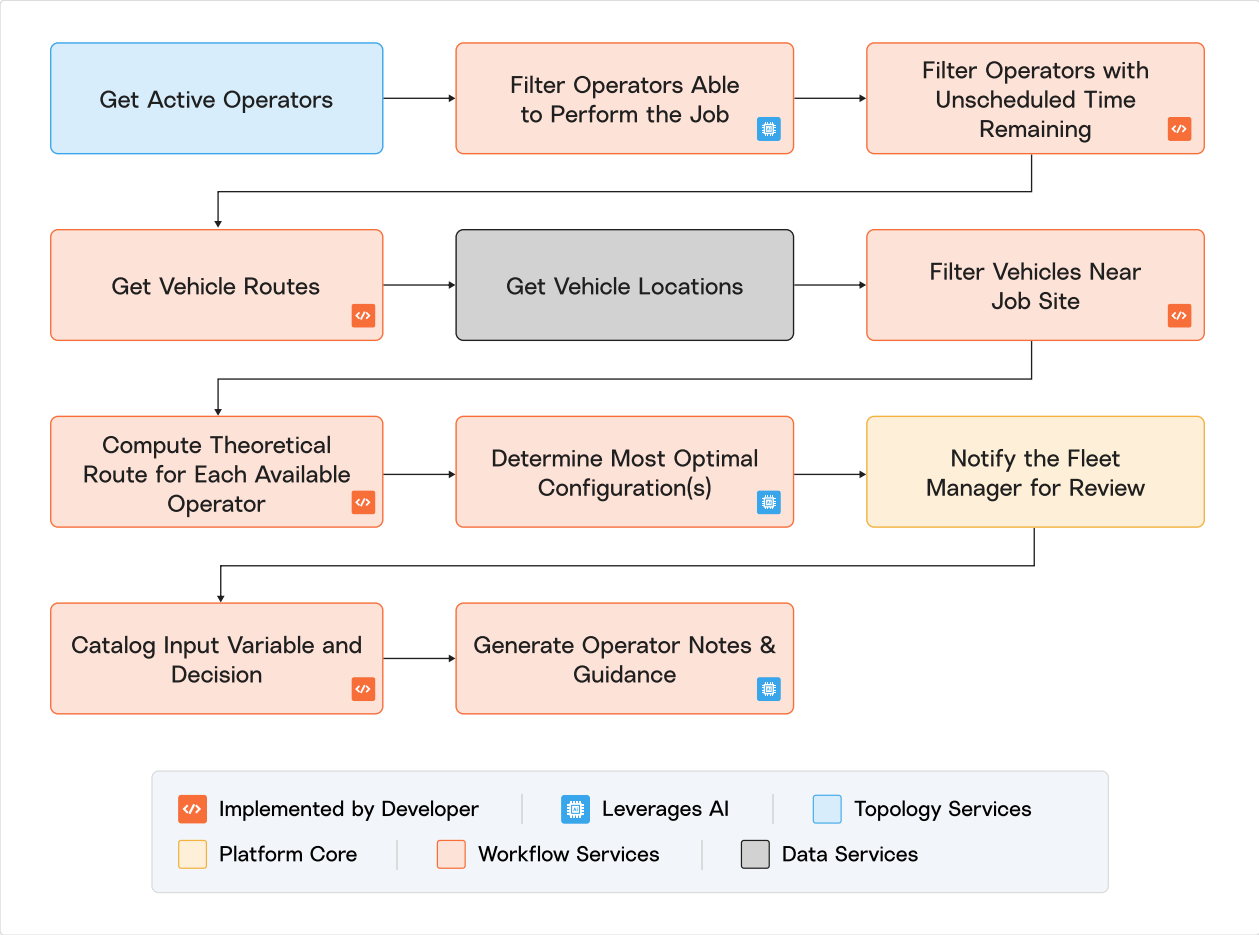

Exhibit 4: *A proposed workflow supporting intelligent on-demand fleet dispatching*

In a large fleet, this workflow may be executed several times per day or even per hour. Performing this entire workflow manually is incredibly time-consuming and error-prone. The steps range in

complexity from simple filtering to those steps requiring computing and scheduling routes and updating routes, factoring in multiple data points. However, the platform handles the complexity of tracking vehicles, operators, and their data; notifying the fleet manager when necessary; and orchestration of the workflow.

An application that meets the needs of fleet managers in this lifecycle stage would have components supporting each of these steps, as well as the user interface with which fleet managers, operators, and their customers interact.

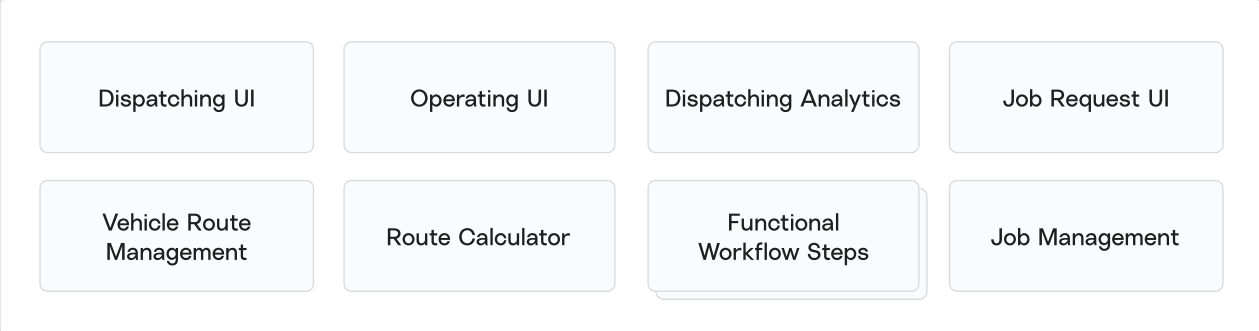

Exhibit 5: *Dispatching application architecture*

Each application component has a specific responsibility:

1. **Dispatching UI:** a view of dispatching history and analysis to inform fleet management decisions.
2. **Operator UI:** a vehicle operator view of their jobs, hours worked, and notifications.
3. **Dispatching Analytics:** a dispatching-focused dashboard built on the platform's Reporting Services.
4. **Job Request UI:** a view for customers to make requests.
5. **Vehicle Route Management:** the source of truth for calculating efficiencies on active or planned vehicle routes.
6. **Route Calculator:** computes routes for vehicles and relevant business metrics (to be used in choosing optimal routes).
7. **Functional Workflow Steps:** a set of functional components that interface with the platform workflow services.
8. **Job Management:** the source of truth for job status, from request to fulfillment.

So far, the dispatching application leverages nearly every platform subsystem:

1. **Workflow Services** for orchestrating dispatching processes.
2. **Topology Services** for getting information about vehicles and operators.
3. **Reporting Services** for reporting on dispatching KPIs.
4. **Data Services** for retrieving vehicle signals, as well as for reporting workflow and dispatching metrics needed to generate reports.

Creating Leverage with AI/ML Services

Some workflow steps could be implemented by leveraging **AI/ML Services** to reduce development efforts, improve the system's adaptability, and ultimately save time for fleet managers performing dispatching tasks. Examples include:

- **Filter Operators Able to Perform the Job:** could use generative AI to perform filtering rather than traditional programming methods. Fleet managers and operators could enter notes on the operator profiles or jobs to filter based on use cases and criteria the platform may not yet support.
- **Determine Most Optimal Configuration(s):** could enable the fleet manager to quickly review and implement an informed dispatching decision following guidance refined by the fleet manager, configuration options, and historical data.
- **Generate Operator Notes & Guidance:** could use generative AI to ensure the operator has the information they need to perform the job in question, including relevant details from the job history or customer interaction history, thus improving customer experience and operator confidence.

Since the application is already integrated with the platform and publishing data, the value of the **AI/ML Services** subsystem will only increase. In future product iterations, the dispatching workflow can leverage AI-enhanced steps to enable optimal dispatching decisions without recomputing every route and delta. Dispatching analytics will be able to leverage job management historical data to train a model that predicts scheduling needs. AI can also help address urgent scheduling requests by identifying which vehicle operator will have a time gap in their schedule to do the job. Moreover, historical routes, paired with fuel usage and expense reports, can be used to train a model to detect anomalies in expenses from operators and flag them to the fleet manager.

In this example, the developers of the dispatching application could save time in initial development by leveraging several of the subsystems of the platform. However, the true value of building an intelligent fleet management platform is realized over time. The platform collects data from the dispatching applications and unlocks additional features and functionality through AI insights with minimal additional effort for developers.

Benefits for the Business that Owns the Platform

Delivering a platform that provides such value to application developers and their products will likely attract more and more contributors to the ecosystem[6]. In turn, each new application built on the platform provides a few key benefits back to the platform.

1. Adds **features for end users** with little to no effort from the platform owner.
2. Creates further **opportunities for integrations**. As the ecosystem of applications on the platform grows, it becomes more obvious that new applications should be built on the platform to leverage the ecosystem of seamless integrations.
3. **Generates data** for the platform about product usage and fleet management operations.

The tremendous potential for this virtuous cycle of platform value-add can be unlocked by

[6] "A Practical Guide to Platform Assessment," Nuvalence, 2023: https://nuvalence.io/white-papers/a-practical-guide-to-platform-assessment/

tapping into newly generated data using **Reporting Services** and **AI/ML Services** in order to derive product insights, build targeted marketing campaigns, and generate recommendations. Leveraging the data generated on the platform allows the business to pinpoint the most impactful efforts. At the same time, it helps foster trust with users that the platform is focused on improving their experience and efficiency.

Business Processes Optimization Insights for Fleet Managers

The benefits of data and AI insights for the platform and application developers are amplified and passed on to all fleet managers. Revisiting the goals from the service provider perspective, the platform provides value to fleet management end users throughout the journey with different combinations of platform services. The platform composes data and services to create AI insights that save time and improve decision-making for fleet managers in each of these scenarios, as follows:

Planning

- Leveraging a model trained on historical utilization data, the AI/ML API suite will offer **predictions on an upcoming period's business** volume.
- The AI/ML API suite leverages a model trained on fleet topology and product usage data to **recommend hardware, software, and new vehicle investment** needs. Combined with generative AI, the platform will explain the recommendation and how it provides value to the fleet manager's business.

Growth

- Workflow Services can enable a workflow modeling the operator onboarding process. This workflow could ingest and process **license and certification** document uploads, leverage AI-enhanced steps to make high-confidence **background check** decisions instead of manual steps, and trigger actions such as notifications.
- Leveraging traditional ML and generative AI models, the AI/ML API suite can enable fleet managers to determine their eligibility for **leasing and financing** and understand, compare, and contrast the various options available to their business.
- The AI/ML API suite, leveraging traditional ML and generative AI models, can create a cooperative **vehicle onboarding** experience, where the platform can meet fleet managers halfway, ingesting the information they have readily available about their fleet and transforming it into the data the platform requires.

Dispatching

- Workflow Services, Topology Services, and AI/ML Services may be used in combination with one another to accomplish dispatching goals, as described earlier in this paper. This includes the **assignment of scheduled jobs** to **maximize vehicle capacity** and operator availability, **optimal route** determinations, as well as incorporation of **on-demand job assignments** throughout the day.

Monitoring

- Traditional ML Pipelines & Patterns can be used to **monitor vehicle signals**, potentially combined with some reactive functionality described in the Maintenance & Upkeep phase below to interpret data abnormalities and **predict unplanned maintenance**.

- Traditional ML Pipelines & Patterns can be used to **monitor operator activities and on-vehicle anti-theft devices**, potentially combined with some reactive functionality modeled in Workflow Services.

Maintenance

- Workflow Services, leveraging a variety of topology data, AI steps, and manual steps, can be leveraged to **model routine maintenance** schedules and **ensure registration compliance** actions are efficiently performed.
- **Unplanned maintenance functionality** would be very similar to the scheduled maintenance workflow. This might be enhanced with AI steps to process events based on vehicle history, signaling the type of maintenance the vehicle requires (whereas for scheduled maintenance, that is strictly specified).

Reporting

- Reporting Services are primarily responsible for aggregating data and generating reports and visualizations. This function can leverage the AI/ML API suite and underlying models to cater the reports to the business by **highlighting and making recommendations on important trends**, anomalies, and other details.
- The AI/ML API suite can be leveraged in combination with historical data from data services, generated reports, and a generative AI model to provide guidance to fleet managers in this phase of the journey on how to best **adjust and refine their fleet operations to attain KPI targets**.

The **Optimization** phase then provides a feedback loop, reflecting on data from monitoring, maintenance needs; and reporting outputs that can help inform the next round of planning, purchasing, and dispatching. Considering that the primary goal of the fleet management platform is to improve efficiency, optimization is the immediate area where AI insights will come to the aid of the platform and its users. Therefore, each stage in this lifecycle provides data or takes learnings from reflection on reporting and optimization.

Evaluating Positive Outcomes

To help evaluate return on investment, an organization starting down this path needs to define KPIs to measure success. These should align with organizational goals, but may include the following:

- **Percentage of new features leveraging platform capabilities** - This metric should indicate whether the right platform capabilities are being prioritized; the number should increase as the platform grows. If the number is low or stagnant, the platform teams' roadmaps should likely be reviewed against organizational goals to ensure better alignment.
- **Engineering Lead Time** - This is the interval between request and delivery. As the percentage of new features leveraging platform capabilities increases, the engineering lead time should decrease.
- **Net Promoter Score (NPS)** - Measured for integrators (developers building platform integrations) and end users, an AI-centric platform should offer intelligence that improves overall user experience, increasing NPS over time.

- **Monthly Active Users and Daily Active Users** - The number of users regularly returning to use the products should increase as the customer base grows. It should also grow as new and relevant features are identified that help users more effectively perform in their roles.

An intelligent fleet management platform leverages the data it already has from consumers using its products to provide value to the platform owner, integrated products, and end users. No matter what KPIs or criteria are used to measure its effectiveness, the efficiency of the intelligent platform provides a considerable competitive advantage over a fleet management platform that does not leverage the capabilities of AI/ML.

Conclusion

We have explored many benefits to building an intelligent fleet management platform in this paper. While specific outcomes may vary, they should accrue to two overarching advantages that have the transformative effect of making fleet management operations a competitive differentiator, and not merely a cost center. The intelligent platform enables the organization that owns it to:

Enable adaptability to many use cases – in minimal time.

Machine learning, and generative AI in particular, has unlocked tremendous potential for platforms to rapidly build new features supporting many use cases. Functionality that, in the past, may have taken weeks to design, build, and bring to production – or may not have even been possible before – may now be delivered in record time by leveraging appropriate data, insights, and models on a platform built with AI at its core.

This benefit is not limited to the platform itself, as the use of AI in product features passes on exactly the same benefits to end users – giving users more versatility and saving them time.

Identify opportunities to grow the business.

A key differentiator of artificial intelligence and machine learning is that it can identify trends and opportunities from massive and disparate data sets – opportunities that may not be as obvious to humans. For a fleet management platform, this understanding can be leveraged for growth in two ways: to grow the platform, as well as to help fleet managers build new features that enable them to grow their own business.

In the AI use cases described along the phases of the fleet manager's journey, the platform provides several opportunities to augment fleet managers' skills in supporting this growth.

Nuvalence

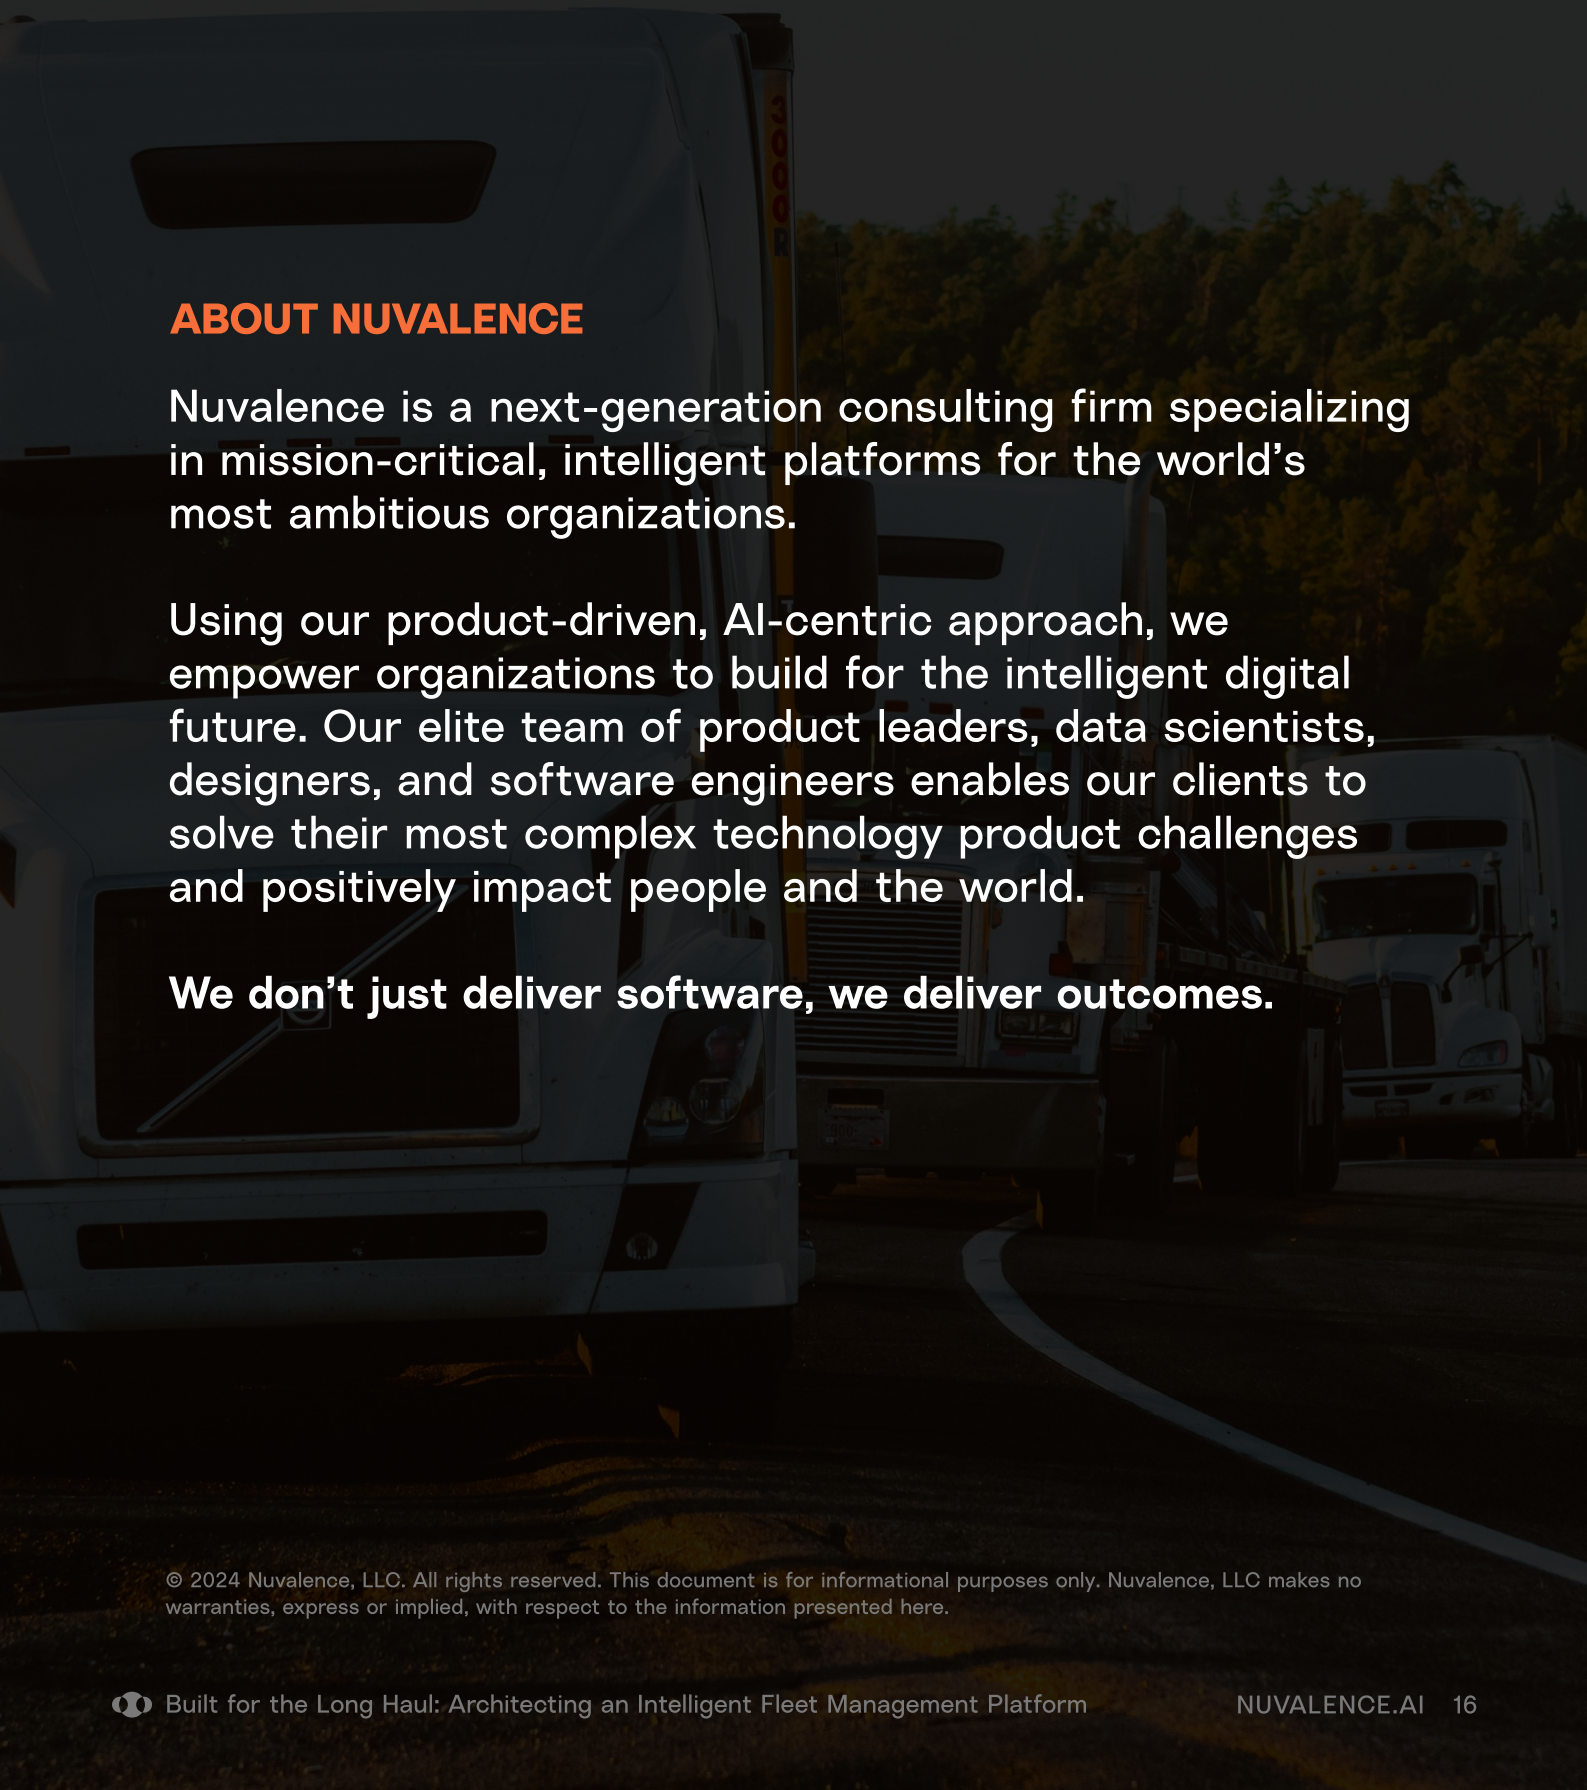

ABOUT NUVALENCE

Nuvalence is a next-generation consulting firm specializing in mission-critical, intelligent platforms for the world's most ambitious organizations.

Using our product-driven, AI-centric approach, we empower organizations to build for the intelligent digital future. Our elite team of product leaders, data scientists, designers, and software engineers enables our clients to solve their most complex technology product challenges and positively impact people and the world.

We don't just deliver software, we deliver outcomes.