AWS Prescriptive Guidance

Building an industrial Internet of Things (IIoT) digital transformation strategy

AWS Prescriptive Guidance: Building an industrial Internet of Things (IIoT) digital transformation strategy

Table of Contents

Building an industrial Internet of Things (IIoT) digital transformation strategy

Salih Bakir, Vladi Salomon, and Asim Kumar Sasmal, Amazon Web Services (AWS)

June 2022 (document history (p. 14))

The *industrial Internet of Things* (IIoT) refers to the use of internet-connected sensors and devices in the industrial sectors, such as manufacturing, energy, automotive, healthcare, life sciences, and agriculture. IIoT allows telemetry data collection from equipment, machines, and field devices in an operational environment. These environments are commonly subject to industrial regulations; compute, networking, and power constraints; and inclement conditions, and all of these challenges contribute to the complexity of designing an IIoT solution.

You can use IIoT data and other enterprise, IT, and operation technology (OT) data to create additional business value, such as optimizing operations, improving productivity, and increasing availability. According to The age of analytics: Competing in a data-driven world (McKinsey Global Institute study), manufacturers can use IIoT data to reduce product development costs by up to 50%, reduce operating costs by up to 25%, and increase gross margins by up to 33%. Therefore, many organizations in the industrial sectors are starting their digital transformation journey to solve business problems by using IIoT.

Targeted business outcomes

This guide helps you create a customized roadmap for your journey, from identifying your IIoT business objectives to realizing them. This guide is based on the experiences of the Amazon Web Services (AWS) Professional Services team partnering with customers to help with their IIoT digital transformation journey. Using the phased approach described in this document, you can:

- Identify business objectives, measurable key performance indicators (KPIs), and prioritized use cases.
- Evaluate your current systems and technologies, assess the skills of your team, and identify gaps.
- Implement a repeatable and reusable blueprint for rapid deployment at scale. A blueprint is an end-to-end IIoT system reference architecture you adopt on your digital transformation journey.
- Prepare your organization for continuous innovation.

Intended audience

This guide is for IT and business executives, program and project managers, architects, product owners, and decision makers in OT, such as heads of operations, plant managers, operations managers.

Whether you are at the beginning or in the middle of your IIoT digital transformation journey, you can use the phased approach described in this guide to either build a customized plan for your journey or identify any gaps in your current plan.

Phases of an IIoT digital transformation journey

AWS Professional Services uses a phased approach to build and realize a plan for an IIoT digital transformation journey:

- Phase 1: Identifying business objectives (p. 2) – Clearly identify and prioritize the business opportunities or problems to address. These are the main driver and the foundation for the overall process. Your business objectives should be ambitious yet achievable, with top-down shared goals for the organization. A typical IIoT digital transformation journey is more than a single project. We recommend a holistic approach for success by thinking big, starting small, and scaling fast.
- Phase 2: Evaluating your current state (p. 5) – Assess your current enterprise, IT, and OT systems, and assess the skill sets of your current team to identify any gaps. Invest in technology, training, or resources for the areas where you need long-term success. For success, ensure that you have strategic alignment between your IT and OT teams.
- Phase 3: Defining a blueprint (p. 8) – Define your blueprint, which is your target reference architecture. It should be repeatable and reusable so that you can rapidly deploy it at scale across industrial sites. This blueprint is the foundation of your journey, and it helps you realize your business objectives by using a think big, start small, and scale fast approach.
- Phase 4: Enabling continuous innovation (p. 11) – When your blueprint is operational, use the collected data to identify opportunities for continuous improvement and refinement. Continue to explore solutions to maximize insights from existing and new data.

Before you begin

It is important to have executive-level commitment to the long-term investment of an IIoT digital transformation journey. Executive sponsors must be aligned to a sustainable strategy and have patience for achieving the target outcomes. According to The age of analytics: Competing in a data-driven world (McKinsey Global Institute study), "Less than one-third of all respondents say their organizations have engaged a chief digital officer to support their transformations. But those that do are 1.6 times more likely than others to report a successful digital transformation." Therefore, prior to starting your journey, make sure the executive team understands and is aligned on the investment strategy, budget, and timeline. Confirm that all of the business stakeholders across IT and OT are committed.

Phase 1: Identifying business objectives

The business case for the digital investment report (Econsultancy website) states that "The lack of a clear longer-term business case and ROI, an absence of board-level understanding and sponsorship, and the perception of digital marcoms as tactical rather than strategic are all viewed as significant challenges to securing the right levels of investment for a digital strategy by at least three-quarters of responding companies." Therefore, it is important to identify and prioritize business opportunities based on measurable KPIs.

Note that business cases might differ for different industrial sectors, such as manufacturing, energy, automotive, healthcare, life sciences, agriculture. For an example of the business impacts to manufacturing, see the Achieved impact from *Successful Digital Industrial Solutions* section of the Hitachi Vantara Solution Brief.

In this phase, you do the following:

1. Identify business challenges (p. 3)
2. Identify measurable KPIs (p. 3)
3. Identify business objectives (p. 4)
4. Identify use cases (p. 4)

The outcome of this phase is that all stakeholders are aligned on the target objectives, understand expectations, and know how success will be measured.

Identify business challenges

The first step in defining the business objectives is making a list of the current business challenges that you want to solve and the new business challenges you might encounter by implementing an IIoT digital transformation solution in your environment.

The following are some common business challenges for manufacturing and industrial companies that are in the early phases of their IIoT digital transformation journey:

- Making legacy industrial machines and equipment smart
- Extracting trapped production data for new insights
- Reduced productivity and increased downtime due to chaotic operations and slow processes for root-cause analysis
- Asset management challenges due to data silos and lack of digital tracking of assets
- Lack of near real-time monitoring at different levels of operations, such as monitoring overall equipment effectiveness (OEE), throughput, and cycle time at the plant, line, and machine level

Identify measurable KPIs

Based on the identified business challenges, you can start asking *How would I measure a successful solution to this problem?* Answering this question helps you take data-driven approach to evaluating the success of the solution.

Determine the KPIs that you will use to measure the success of your journey, and make sure that they are measurable. The following are example KPIs for that apply to a variety of industrial sectors:

- % improvement for overall equipment effectiveness (OEE) or similar KPIs
- % reduction in operational cost
- % reduction in storage and compute cost for cloud, compared to on-premises
- % reduction in unplanned downtime due to proactive monitoring and maintenance
- % accuracy in demand forecasting and inventory management
- % reduction in latency observed by business users for business intelligence (BI) reporting
- % reduction in time to make the historical data available for advanced analytics, such as machine learning
- % reduction in time to scale compute and storage
- % increase in system uptime
- % increase in productivity
- % reduction in downtime

Identify business objectives

Now that you have identified the business challenges you want to solve and decided how to measure success, you can now define your business objectives. These objectives help you answer the questions *Why is this problem worth solving?* and *Who benefits from solving this problem?* You decide on a data-driven strategy for measuring success, such as comparing current state KPIs with the target state KPIs for the particular business objective.

For each metric or KPI you want to use, rephrase it as a business objective with a target measurable value. For example, if your business challenge is *product1 is frequently out of stock due to a manual detection process* and your metric is *% reduction in latency to detect the problem*, the business objective might be *95% reduction in latency to identify possible out-of-stock situations for product1*.

Prioritize your business objectives so that the team has a clear understanding of how to prioritize resource allocation.

Identify use cases

After you have defined your business objectives, you can now focus on the use cases. *Use cases* define the exact interactions end-users have with the system, and you use them to determine how to automatically create the expected business outcome. Use cases act as the main requirements when building your blueprint.

Each use case should consist of four key elements:

- One or more end-user personas who interact with the system
- Goals for each persona
- System actions that you want to implement, as experienced by the persona
- The expected result of the system actions, as experienced by the persona

Using the example business objective *95% reduction in latency to identify possible out-of-stock situations for product1*, the following is an example use case for this objective:

- **Persona** – Business analyst
- **Goal** – Use trend analysis to estimate out-of-stock situations for product1 within minutes
- **Action** – Use a BI reporting tool to generate a report that shows the trend
- **Result** – New solution should have 95% reduced latency to identify out-of-stock situations, compared to the previous manual process

After you have the list of use cases, evaluate them with the stakeholders based on the importance and feasibility of each use case. *Importance* is the value you expect to gain from the use case, such as its return on investment (ROI), and *feasibility* is the ease of implementation. Create a table like the following, and ask stakeholders to vote on the importance and feasibility of each use case. For example, in the following table, use case 1 received 4 votes for high importance and 3 votes for low importance. The majority vote indicates this use case has a high importance.

	Importance		Feasibility		Majority vote	
	High	Low	High	Low	Importance	Feasibility
Use case 1	4	3	5	2	High	High
Use case 2	5	2	1	6	High	Low
Use case 3	1	6	3	4	Low	Low

Next, you use the voting results to prioritize the use cases. Use cases with two high ratings are considered *quick wins*. Put use case with one high rating and one low rating in the *evaluate* category, and put use cases with two low ratings in the *consider* category. The following table shows a quadrant chart you can use to visualize this categorization.

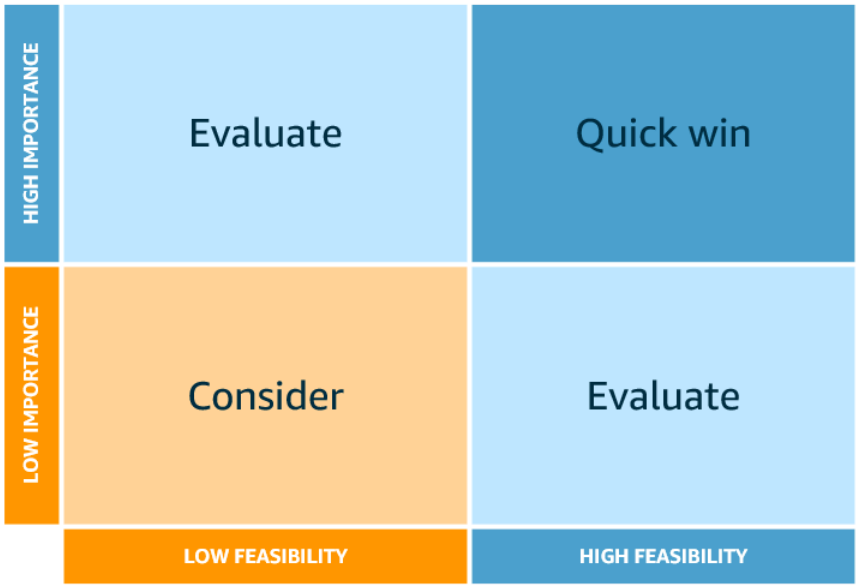

Prioritize use cases that are quick wins, and make sure that you consider dependencies. As you complete your journey, you start with the quick wins and as you progress, you can add use cases in the evaluate and consider categories, based on your budget and schedule.

Phase 2: Evaluating your current state

A complete IIoT digital transformation journey encompasses not only the IIoT-specific devices and strategy but also a holistic consideration of how those IIoT assets integrate with your IT and OT infrastructure and operations personas. Your infrastructure might be *on-premises* (local), or it might be *hybrid* of both on-premises and cloud infrastructure. Migrating your infrastructure to the cloud allows you to take full advantage of cloud-native features to improve agility, performance, and scalability.

In this phase, you do the following:

- Focus on people and culture (p. 5)
- Discover your current systems and technology stack (p. 7)
- Review other key considerations (p. 7)

AWS Professional Services uses a set of prescriptive offerings that are tried-and-tested with other customers. We can help you assess your current state and build a phased roadmap for your target state.

Focus on people and culture

One of the key challenges is not having the proper team skill sets to enable and sustain the IIoT digital transformation. You need to upskill your team, create new roles, and hire new talent to drive success. Unlocking success in digital transformations (McKinsey & Company survey report) states that "The digital transformation success is more than three times likelier when respondents say their organizations have invested the right amount in digital talent." We recommend you consider the following topics to assess the skill sets of your current team and take actions accordingly:

- Cloud technology stack expertise
- Primary technical skill sets such as:
  - IIoT
  - Machine learning (ML)
  - Data analytics tools and methods
  - Data lakes
  - Online analytical processing (OLAP) and online transaction processing (OLTP) systems, such as SQL/NoSQL databases and data warehouses
  - Business intelligence tools
  - Real-time monitoring tools
  - Web application development, including frontend and backend
  - Operating systems, such as Linux
  - Programming languages, such as Java, Python, JavaScript
- The resources to build software products and solutions, including:
  - Business analyst
  - Product owner
  - Project manager
  - UX/UI designer
  - Software architect
  - Data architect
  - IoT architect
  - Software developer
  - Software testing and automation engineer
  - Development Operations (DevOps) engineer
  - Data scientist
  - OT subject matter experts (SMEs), such as processing engineers, production engineers, plant managers, and line managers
- The team is sized and structured according to agile principle and practices
- Partners for long-term and short-term acceleration and training

Another important point is having an innovative culture to embrace digital transformation and drive it. Because even if you have the correct strategy, processes, and tools in place, if your organizational culture does not encourage innovation and adoption, the digital transformation is less likely to be successful. Consider some of the following strategies to encourage adoption of the digital transformation in your organization:

- Having a North Star vision, values, and principles (for more information, see North Star vision (p. 8))
- Having senior leadership support
- Having a roadmap that minimizes disruptions to operations
- Promoting an entrepreneurial mindset and accepting failures
- Having data-driven, customer-focused goals
- Adopting agile process and tools
- Recognizing individuals who advocate for the digital transformation, and providing them opportunities to lead or participate in the initiative

- Involving employees in the initiatives
- Providing more autonomy and flexibility for the teams
- Promoting teamwork, communication, and transparency
- Having strong and fast feedback mechanisms

Discover your current systems and technology stack

The technical capabilities of your existing systems define the scope of the future system architecture. Therefore, you need to discover your IT and OT infrastructure to understand its current technical capabilities.

Consider the following to assess the current edge infrastructure capabilities:

- Current edge architecture
- Existing IoT or IIoT systems or solutions and their capabilities
- Current data analytics and machine learning use cases, such as descriptive analytics, predictive analytics, anomaly detection, predictive and preventive maintenance, near real-time dashboard, and BI reporting
- Scale of existing solutions and future requirements
- Data sources and their capabilities for ingesting data, including:
  - Devices or tools, such as sensors, actuators, programmable logic controllers (PLCs), gateways, and OPC Unified Architecture (OPC UA) servers
  - Supported protocols for those device and tools, such as Modbus, BACnet, MQTT, and OPC UA
  - Data specifications, such as telemetry frequency, size of typical message, format, and volume
- Network infrastructure for clear isolation between OT and IT network
- Network connectivity, such as Ethernet, Wi-Fi, LoRaWAN, and 5G
- Existing historians and data storage systems
- Existing cloud connectivity options

Consider the following to assess the current cloud infrastructure capabilities:

- Current cloud architecture
- Data lakes
- Data analytics
- Data transformation
- Data service layer
- Data monitoring and BI
- Machine learning
- Web applications

Review other key considerations

In addition to the infrastructure considerations, there are also security, compliance, risk management, governance, and operational factors that you need to account for when assessing your current state. Assess the following topics in depth to address some of these considerations:

- Information security strategy that assesses and mitigates threats.

- High-availability requirements, such as recovery time objectives (RTOs) and recovery point objectives (RPOs) for the system.
- Data governance and access control.
- Identity and access management for the system.
- Data retention policies.
- Data classification and sensitivity.
- Data encryption, at-rest and in-transit.
- Compliance and regulatory requirements for processing and storing sensitive data are critical. This includes regulations such as General Data Protection Regulation (GDPR), personally identifiable information (PII), and Health Insurance Portability and Accountability Act (HIPAA).
- Service-level agreements (SLAs) for downstream data consumption and applications.
- Business risk management.
- Asset and device lifecycle management.

Phase 3: Defining a blueprint

Based on the evaluation of your current state in the previous phase, you can start building your blueprint. A blueprint is an end-to-end IIoT system reference architecture you adopt on your digital transformation journey. It serves as the foundation of your IIoT digitalization journey and helps you realize your business objectives. A blueprint:

- Is guided by your North Star vision (p. 8)
- Adheres to the core tenets of a successful solution framework (p. 8)
- Consists of repeatable and reusable building blocks (p. 9)

Sometimes, you might need a quick proof of concept to demonstrate value and feasibility for certain parts of the blueprint.

North Star vision

Your blueprint should be guided by your North Star vision, which is a clear, concise, and long-term goal that provides direction for making business decisions. If you don’t have a North Star vision, think big when creating one. This vision generally takes 3–5 years to realize. To achieve this vision, starting small and scaling fast are the keys to success.

Core tenets of a successful solution framework

To create a unified IT and OT data backbone in your blueprint, you need a functional architecture. Based on our experiences, we’ve identified the following three core tenets of the solution framework:

- Maximize insights
  - Democratizing access to data provides diverse insights and drives business value, such as SKU margin optimization.
  - Performing descriptive analytics on real-time or historical operational data helps you monitor KPIs, identify trends, identify potential areas of improvement, and take actions.
  - Performing diagnostic analytics on data helps you identify the root cause of operational events.
  - Performing predictive analytics on data helps you forecast future events in your business and operations.

  - Performing prescriptive analytics on data suggests multiple solutions for solving a given problem, based on descriptive and predictive analytics results.
- Minimize technical debt
  - Integrating seamlessly with the key existing IT/OT systems eliminates temporary solutions.
  - Automating the deployment pipeline removes manual process from your operations.
  - Standardizing tools prevents proliferation of tools and bespoke applications.
  - Using centralized management services to deploy standardized configurations across the environment, preventing the use of non-standard and potentially problematic configurations at the local site.
  - Creating patterns for updating and deploying infrastructure automatically or with minimal intervention for repeatable tasks. Examples include updating operating systems, periodically rotating device certificates, installing patches, or scaling data storage.
  - Designing and implementing repeatable and reusable patterns for rapid production deployment across sites at scale.
- Modular and future-proof blueprint
  - Designing for interoperability with existing IT/OT systems and infrastructures.
  - Designing for modularity, which helps you start small and scale fast, iteratively add new components, and select the best option for your use case.
  - Designing for flexibility with existing (*brownfield*) and new (*greenfield*) infrastructures.

Repeatable and reusable building blocks

The *building blocks* of an IIoT digital transformation journey are the various functional layers, considerations, and use cases that comprise the blueprint. The following image shows the high-level repeatable and reusable functional building blocks of a blueprint.

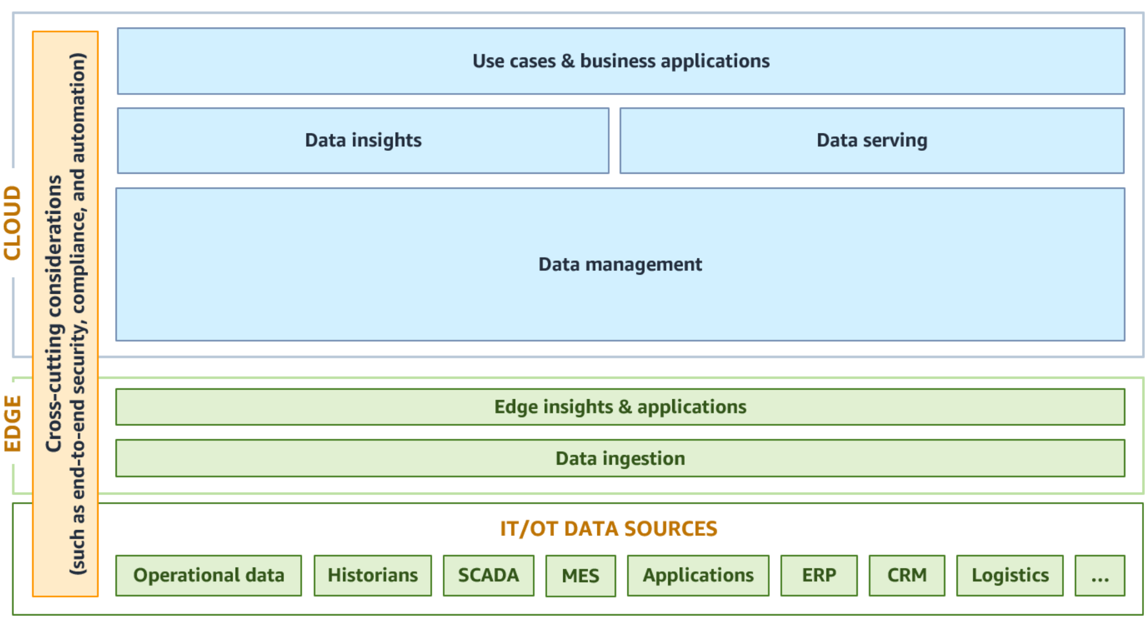

The following are the layers of a blueprint:

- **Data ingestion** – This edge layer collects data from various sources in your on-premises infrastructure or cloud environment. Typical IT/OT data sources might include telemetry data from supervisory control and data acquisition (SCADA) systems, distributed control systems (DCS), PLCs, secondary sensors, manufacturing execution systems (MES), software as a service (SaaS) and legacy applications, enterprise resource planning (ERP) systems, customer relationship management (CRM) systems, various supply chain systems, and data historians.

- **Edge insights and applications** – Depending on your use cases, you might want to deploy this edge layer. This layer is used to address any low latency and data residency requirements for your architecture, support production continuation when disconnected from the cloud, and enable innovation at the edge.
- **Data management** – This layer is responsible for various aspects of typical data management functions, such as:
  - Building and managing semantic data models (SDMs) for IT/OT resources for governance. Adding contexts to the machine data by using a semantic data model helps with downstream analytics for process and machine modeling.
  - Storing the data collected in the data ingestion layer. Use the data stored in this layer for processing and providing local insights, and for providing store-and-forward functionality when disconnected from the cloud.
  - Processing the data in the cloud to meet various consumption needs for end users, such as data integration, data normalization, data enrichment, data quality, data discovery, data catalog, and search.
  - Enabling a flexible data consumption service for external consumers to provide business insights.
- **Data insights** – This cloud layer is used for business insights that range from simple, such as near real-time KPI dashboards, to advanced, such as predictive maintenance, demand forecasting, and inventory management that uses the flexible data consumption service from the data management layer.
- **Data serving** – This cloud layer is used to democratize access to the data for various end users, such as various OT personas, data scientists, data engineers, and data analysts. This layer seamlessly serves data to other enterprise systems and third-party solutions to enable use cases and business applications.
- **Use cases and business applications** – This is the top layer of the architecture. This cloud layer contains the business applications and tools that address your business use cases. As needed, the applications and tools in this layer access the data and insights in the supporting layers.
- **Cross-cutting considerations** – This layer contains key non-functional requirements that apply to the data sources, edge, and cloud. This layer includes must-have elements, such as end-to-end security, configuration management, logging, compliance, and regulatory requirements. This layer helps you securely and efficiently operate your architecture, providing opportunities to enhance performance, reduce costs, or use automations that enable rapid deployment at scale across sites.

To create this unified data solution, we recommend using a unified functional architecture similar to the one presented. This holistic approach helps you think big, start small, and scale fast. Rather than taking on the entire digital transformation journey at once and making the journey impossibly difficult, you keep iterating on smaller deliverables that help you achieve your business outcomes. You might already have some of these building blocks in place today, and if so, you can reuse them.

AWS IDP solution offering

AWS Professional Services uses a tried-and-tested approach, AWS Industrial Data Platform (IDP), to discover, design, and implement a flexible and extensible unified data solution for Industry 4.0 (also known as smart manufacturing, smart factory, or smart industrial) success. The AWS IDP addresses a catalog of common use cases, such as:

- Operational and actionable KPIs for production and asset optimization, including overall equipment effectiveness (OEE), throughput, yield, and cycle time
- Automated quality and defect management solutions for predictive quality
- Predictive maintenance that reduces downtime and catastrophic equipment failures
- Energy optimization and carbon footprint reduction for sustainable manufacturing

- Supply chain optimization, including inventory management, demand forecasting, and track and trace

Your blueprint architecture might vary based on your use cases, your current state assessment, and the identified gaps. For more information about the relevant AWS services that you can use in your blueprint, see the AWS Industrial Data Platform (IDP) reference architecture.

Phase 4: Enabling continuous innovation

We recommend you consider your IIoT digital transformation initiative as a journey and not a single project. After defining business KPIs, enabling the organization, establishing the necessary skill sets, and building the blueprint; you increase the pace of innovation and explore more opportunities to build and scale throughout your business. When you have the blueprint running, you have a more transparent and data-driven business that you can monitor and track extensively.

You can use the IIoT data to identify challenges and opportunities. Furthermore, it is vital to always start with a business objective that is measurable to drive innovation. Then you can extend the blueprint to support it with a use case. It is also important to evaluate the current business model and evolve it to unlock new business opportunities so that you can be more competitive. For instance, you can shift from selling products to selling value-added services. Because IIoT capabilities provide extensive control of your products, you can extend or limit their features based on customers' requests with more competitive prices.

Finally, AWS Professional Services can help you build your own data flywheel by defining main elements to enable a virtuous cycle for innovation, just like the Amazon flywheel (YouTube video). The following image is an example of a flywheel.

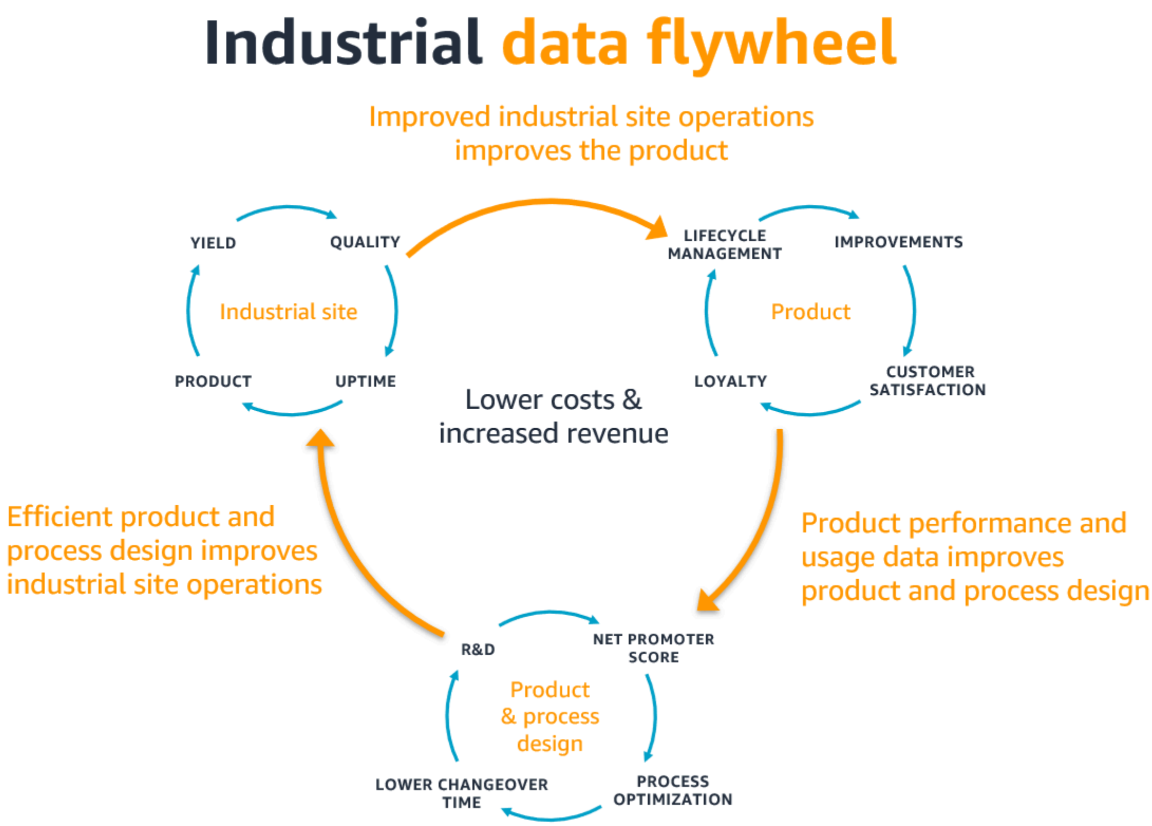

Conclusion and next steps

The latest industrial revolution, the IIoT digital transformation, is unlocking new possibilities for businesses of all sizes and in every industry. Using IIoT data helps solve business problems and create new opportunities. Businesses that embrace this IIoT journey are quickly realizing the benefits of improved insights and productivity.

However, the biggest challenges in any IIoT digital transformation is not knowing where to start and not defining solid business objectives. By defining your business objectives, the KPIs you use to measure success, and your North Star vision blueprint, you can adopt a systematic approach to your journey. You assess the current situation to understand any gaps, building the blueprint to close those gaps, and then focus on continuous innovation.

In this guide, we explained a phased IIoT digital transformation journey. In the experience of AWS Professional Services, this phased approach has been effective and successful. If you want help from AWS Professional Services to accelerate your journey, complete the contact form.

When required for your success, AWS Professional Services also works jointly with our broad AWS Partner Network. For more information, see AWS IoT Competency Partners and AWS Marketplace.

For more information about how AWS can help you achieve your operational objectives, see Resources (p. 13).

Resources

Customer case studies

- Yara and AWS to Digitalize Crop Nutrition Production System
- Volkswagen Group on AWS
- Coca-Cola İçecek Improves Operational Performance
- How Genie (a Terex brand) improved paint quality

AWS resources

- AWS Industrial Internet of Things
- AWS Industrial Data Platform (IDP)
- AWS for Industry
- AWS Migration Acceleration Program (MAP)
- AWS Cloud Adoption Framework (AWS CAF)
- Optimizing industrial operations with digital twins (AWS Online Tech Talks video)
- Ten security golden rules for Industrial IoT solutions (AWS blog post)

Whitepapers

- Manufacturing Transformation: Journey to the Cloud (AWS whitepaper)
- Enabling Manufacturing Innovation Through the Use of Cloud (IDC whitepaper)
- Mastering the Industrial Internet of Things (IIoT) (Roland Berger whitepaper)

Document history

The following table describes significant changes to this guide. If you want to be notified about future updates, you can subscribe to an RSS feed.

Change	Description	Date
Initial publication (p. 14)	—	June 20, 2022

AWS Prescriptive Guidance glossary

The following are commonly used terms in strategies, guides, and patterns provided by AWS Prescriptive Guidance. To suggest entries, please use the **Provide feedback** link at the end of the glossary.

AI and ML terms

artificial intelligence (AI)

> The field of computer science that is dedicated to using computing technologies to perform cognitive functions that are typically associated with humans, such as learning, solving problems, and recognizing patterns. For more information, see What is Artificial Intelligence?

binary classification

> A process that predicts a binary outcome (one of two possible classes). For example, your ML model might need to predict problems such as “Is this email spam or not spam?" or "Is this product a book or a car?"

classification

> A categorization process that helps generate predictions. ML models for classification problems predict a discrete value. Discrete values are always distinct from one another. For example, a model might need to evaluate whether or not there is a car in an image.

data preprocessing

> To transform raw data into a format that is easily parsed by your ML model. Preprocessing data can mean removing certain columns or rows and addressing missing, inconsistent, or duplicate values.

deep ensemble

> To combine multiple deep learning models for prediction. You can use deep ensembles to obtain a more accurate prediction or for estimating uncertainty in predictions.

deep learning

> An ML subfield that uses multiple layers of artificial neural networks to identify mapping between input data and target variables of interest.

exploratory data analysis (EDA)

> The process of analyzing a dataset to understand its main characteristics. You collect or aggregate data and then perform initial investigations to find patterns, detect anomalies, and check assumptions. EDA is performed by calculating summary statistics and creating data visualizations.

features

> The input data that you use to make a prediction. For example, in a manufacturing context, features could be images that are periodically captured from the manufacturing line.

feature importance

> How significant a feature is for a model’s predictions. This is usually expressed as a numerical score that can be calculated through various techniques, such as Shapley Additive Explanations (SHAP) and integrated gradients. For more information, see Machine learning model interpretability with AWS.

feature transformation

> To optimize data for the ML process, including enriching data with additional sources, scaling values, or extracting multiple sets of information from a single data field. This enables the ML model to

benefit from the data. For example, if you break down the “2021-05-27 00:15:37” date into “2021”, “May”, “Thu”, and “15”, you can help the learning algorithm learn nuanced patterns associated with different data components.

interpretability

A characteristic of a machine learning model that describes the degree to which a human can understand how the model’s predictions depend on its inputs. For more information, see Machine learning model interpretability with AWS.

machine learning (ML)

A type of artificial intelligence that uses algorithms and techniques for pattern recognition and learning. ML analyzes and learns from recorded data, such as Internet of Things (IoT) data, to generate a statistical model based on patterns. For more information, see Machine Learning.

multiclass classification

A process that helps generate predictions for multiple classes (predicting one of more than two outcomes). For example, an ML model might ask "Is this product a book, car, or phone?" or "Which product category is most interesting to this customer?"

regression

An ML technique that predicts a numeric value. For example, to solve the problem of "What price will this house sell for?" an ML model could use a linear regression model to predict a house's sale price based on known facts about the house (for example, the square footage).

training

To provide data for your ML model to learn from. The training data must contain the correct answer. The learning algorithm finds patterns in the training data that map the input data attributes to the target (the answer that you want to predict). It outputs an ML model that captures these patterns. You can then use the ML model to make predictions on new data for which you don’t know the target.

target variable

The value that you are trying to predict in supervised ML. This is also referred to as an *outcome variable*. For example, in a manufacturing setting the target variable could be a product defect.

tuning

To change aspects of your training process to improve the ML model's accuracy. For example, you can train the ML model by generating a labeling set, adding labels, and then repeating these steps several times under different settings to optimize the model.

uncertainty

A concept that refers to imprecise, incomplete, or unknown information that can undermine the reliability of predictive ML models. There are two types of uncertainty: *Epistemic uncertainty* is caused by limited, incomplete data, whereas *aleatoric uncertainty* is caused by the noise and randomness inherent in the data. For more information, see the Quantifying uncertainty in deep learning systems guide.

IoT terms

artificial intelligence (AI)

The field of computer science that is dedicated to using computing technologies to perform cognitive functions that are typically associated with humans, such as learning, solving problems, and recognizing patterns. For more information, see What is Artificial Intelligence?

brownfield strategy

The existing infrastructure in your environment. When adopting a brownfield strategy for a system architecture, you design the architecture around the constraints of the current systems and infrastructure. If you are expanding the existing infrastructure, you might blend brownfield and greenfield (p. 17) strategies.

cloud computing

The cloud technology that is typically used for remote data storage and IoT device management. Cloud computing is commonly connected to edge computing (p. 17) technology.

edge computing

The technology that increases the computing power for smart devices at the edges of an IoT network. When compared with cloud computing (p. 17), edge computing can reduce communication latency and improve response time.

greenfield strategy

The absence of existing infrastructure in a new environment. When adopting a greenfield strategy for a system architecture, you can select all new technologies without the restriction of compatibility with existing infrastructure, also known as brownfield (p. 17). If you are expanding the existing infrastructure, you might blend brownfield and greenfield strategies.

industrial Internet of Things (IIoT)

The use of internet-connected sensors and devices in the industrial sectors, such as manufacturing, energy, automotive, healthcare, life sciences, and agriculture. For more information, see Building an industrial Internet of Things (IIoT) digital transformation strategy.

Internet of Things (IoT)

The network of connected physical objects with embedded sensors or processors that communicate with other devices and systems through the internet or over a local communication network. For more information, see What is IoT?

machine learning (ML)

A type of artificial intelligence that uses algorithms and techniques for pattern recognition and learning. ML analyzes and learns from recorded data, such as IoT data, to generate a statistical model based on patterns. For more information, see Machine Learning.