



User Guide

Amazon Connect Decisions



Amazon Connect Decisions: User Guide

Copyright © 2026 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

What is Amazon Connect Decisions	1
Benefits	1
Capabilities	2
Key Concepts and Terminology	3
Decisioning	3
Planning	4
Understanding Your Homepage	6
Key metrics	6
Top topics for today	7
Natural Language Interface (NLI) / Decisions Teammate	8
Navigation bar	9
Onboarding	11
Onboarding for Amazon Connect Decisions	11
Prerequisites	12
Decisioning	13
What are Insights?	13
How Decisioning Works	13
What You'll Find in This Section	14
Configuring Insights	15
The four configuration areas	15
How configuration connects to Insights output	15
Best practices	15
How are Insights generated?	16
Knowledge Sources	19
Detection (Metrics and Rules)	20
Guidelines	24
Prioritization and Severity	26
Managing Insights	27
What Managing Insights Covers	28
In This Section	28
Filtering and Sorting Insights	28
Insights Details	36
Acting on Recommendations	43
Providing Feedback on Insights	47

Plans	49
Planning with Amazon Connect Decisions	49
Demand Planning	49
Prerequisites	49
What is Demand Intelligence?	50
Data Entities for Demand Planning	50
Creating and Configuring a Demand Plan	60
Reviewing and Analyzing Plans	61
Editing Forecasts	62
Monitoring Forecast Performance	63
Publishing Demand Plans	64
Rerun Demand Plans for the Next Planning Cycle	64
Understanding Your Forecasts	65
Plan Improvement Actions	65
Other Capabilities from Supply Chain Decisions Teammate	71
Supply Planning	73
What You'll Find in This Section	73
Data Entities for Supply Planning	74
Creating a New Plan	100
Configuring Plan Parameters	101
Reviewing and Editing Plans	101
Publishing Supply Plans	103
Common Issues and Solutions	104
Common Data Issues	104
Demand history has mixed date formats	104
Negative quantities in order history	104
Record counts don't match your source system	104
Orders appear in the system that don't exist in ERP (or vice versa)	104
Plan input files include products from other plants or business units	105
Common Exception and Recommendation Issues	105
Same product+site appears multiple times in exception list	105
Recommendation doesn't match what I see in the chart	105
Impact date or "Act By" date seems wrong	106
Recommendations reference orders I can't find in ERP	106
Common Accuracy Issues	106
Forecast is significantly worse than a simple moving average	106

Over-forecasting bias (forecast systematically higher than actuals)	107
Under-forecasting bias (forecast systematically lower than actuals)	107
Forecast accuracy degrades significantly at longer horizons	108
External forecast (EDI/customer forecast) not improving accuracy when used in consensus rules	108
Planning rules not taking effect	109
Consensus plan output is identical to the baseline forecast	109
Consensus rules apply to wrong products or sites	110
Forecast values are fractional (e.g., 2,500.37 units)	110
Product coverage drops significantly after adding preprocessing rules	111

What is Amazon Connect Decisions

Amazon Connect Decisions is a supply chain planning and intelligence solution with AI teammates that work alongside your team harmonizing demand signals into consensus forecasts, generating constraint-aware supply plans, and actively monitoring your operations to prevent problems, resolve issues, and identify root causes for continuous improvement.

AI teammates adapt to how you work: driven by your operating procedures and rules, integrating with your existing systems and learning from your practitioners' decisions. It leverages Amazon supply chain expertise to deliver effective plans and recommendations from day one.

You lead a team of AI agents that handle coordination and execute actions, freeing you to focus on strategic decisions that drive service levels and working capital efficiency.

Benefits

Adapt with your business

AI teammates adapt to your business, systems, and decisions, quickly delivering transformation within your existing workflows. Driven by your operating procedures and business rules, teammates work with your existing planning and ERP systems. They learn from planners' decisions, aligning with your priorities and institutionalizing knowledge so your operations get better over time, and every team member is more effective.

Stay ahead of what's coming

Lead a team of AI agents that handle coordination work 24/7—harmonizing demand signals, generating constraint-aware plans, and executing approved actions. Teammates also detect patterns invisible to manual analysis, like cascading supplier disruptions or inventory misalignment across thousands of SKUs, and surface what matters before it becomes a problem. Planners shift from reactive firefighting to proactive planning, running more effective operations and improving service levels through strategic decisions that drive business outcomes.

Build trust from day one

AI teammates are backed by science refined over 30 years of managing 400M+ Amazon SKUs and equipped with 25+ specialized supply chain tools. Your team benefits from domain expertise honed by one of the world's most complex supply chain networks, delivering

actionable insights and recommendations out-of-the-box with clear, explainable reasoning you can trust and verify.

Capabilities

Adaptative Intelligence

AI teammates learn from every plan and decision through a continuous loop: observing patterns, detecting what matters, recommending actions, and executing approved decisions. This institutionalizes knowledge—when planners leave, expertise stays; new members are productive from day one, improving outcomes without manual retraining.

Demand Intelligence

AI teammates dynamically orchestrate 18+ forecasting tools and foundation models trained on Amazon data, autonomously optimizing at the SKU level. Generate accurate forecasts from day one, even with limited history. Harmonize statistical forecasts, customer commitments, and cross-functional inputs through consensus planning while continuously monitoring accuracy.

Supply Intelligence

AI teammates generate forward-looking supply plans (preview) and intelligently cluster exceptions into strategic decisions ranked by impact. Run optimization models for constraint-aware planning across your network. Monitor operations 24/7, surface what matters, then execute approved decisions directly in existing systems—reducing cycles from weeks to hours.

Agentic Onboarding

AI-powered onboarding agents guide you through setup via natural language conversation. Connect fragmented data through JDBC or Amazon S3, prepare it for Connect Decisions, and automatically flag issues before they affect forecasts. Configure business rules and policies in plain language with real-time previews—getting operational in weeks, not months.

Key Concepts and Terminology

Decisioning

- **Insights:** The core resource for proactive supply chain monitoring, and the primary entity you work with in Amazon Connect Decisions. An insight brings together detection, root cause analysis, and recommended actions for a supply chain issue, helping you identify and resolve problems before they impact operations. An insight is generated when a rule detects that a metric has breached its defined conditions (for example, "Forecast accuracy for Product A dropped to 72% at Site B"). Related detections that share the same product, site, and recommendation type are consolidated into a single insight, giving you one place to investigate and act.
- **Metrics:** Quantifiable measures that evaluate supply chain performance. Metrics define what to monitor (e.g., forecast accuracy, inventory levels, days of cover) and how to calculate these measures across your data.
- **Rules:** Business logic that defines when exceptions (insights) should be generated. Rules specify the thresholds, conditions, and criteria that trigger insights based on your metrics (e.g., "Generate an exception when forecast accuracy falls below 85% for A-class products"). Metrics and rules work together to generate insights. Metrics define what to measure while rules evaluate those metrics against thresholds to determine when to generate insights. For example, a forecast accuracy metric calculates performance, and a rule triggers an insight when that metric falls below 85% for A-class products.
- **Root cause:** The AI-powered investigation that explains why an exception occurred. Root cause analysis examines contributing factors, data patterns, and contextual information to identify the underlying reasons for performance deviations.
- **Recommendations:** AI-generated suggested actions to resolve exceptions. Recommendations provide specific, actionable steps to address root causes and restore normal operations.
- **Agent guidelines:** Instructions that guide the AI agent's behavior when analyzing exceptions and generating recommendations. These guidelines incorporate your business policies, operational constraints, and decision-making preferences.
- **Severity Factors:** Configuration settings that determine how financial impact influences exception prioritization. Severity factors help the system surface the most business-critical issues first.

Planning


The following is the common terminology used in Demand Planning:

- **Baseline forecast** – It refers to the use of the historical data by the system to generate a forecast. It provides initial demand prediction before you apply any overrides.
- **Dataset** – A collection of data used for generating forecasts, such as historical sales orders or product information.
- **Demand planning cycle** – The time taken to create and finalize demand plans, which include forecast generation, and collaborating with stakeholders to adjust and publish demand plans.
- **Demand plan** – A plan is created for a defined time horizon and refreshed periodically through rolling time windows (planning cycles). Each demand plan can contain multiple versions within a planning cycle to support refinement of plan for any incremental data received within a planning cycle. The demand plan can display statuses including:
 - **Pending** – The plan configuration is created but not submitted for plan creation.
 - **In progress** – The plan configuration is submitted for plan creation and the plan is in progress for forecast generation.
 - **Failed** – The plan forecast has failed. Look for email and in-app notification to resolve issue.
 - **In-review** – The planning cycle is open and you can edit your forecast.
 - **Final** – The planning cycle is closed, and you cannot edit your forecast. However, you can view the demand plan.
- **Forecast configuration** – The set of planning conditions that govern the plan for forecast generation. This includes the planning cycle configuration, time horizon granularity, and that hierarchy configuration that influences how Demand Planning will generate the forecast.
- **Forecast granularity** – Defines how you want to create and manage the forecast. You can use a combination of product, location, customer, and channel dimensions. You can also choose the time interval for the forecast data to be aggregated by day, week, month, or year for each product in the dataset. For example, if your forecast granularity is set as Daily, you will see the forecast daily for each product in the dataset.
-  **Note**
Demand Planning uses the Gregorian calendar for planning. The default start day of the week is Monday.
- **Override** – A modification that you make to the system generated forecast.

- **Planning Horizon** – The total length of time into the future for which forecasts are generated, measured from the forecast start date. The planning horizon is determined by combining the time bucket (Daily, Weekly, or Monthly) with the plan horizon length. For example, a weekly plan with a 26-week plan horizon creates forecasts covering the next 26 weeks from the forecast start date.
- **Product lifecycle** – The product lifecycle refers to the various stages of a product from introduction to End of Life (EoL).
- **Published demand plan** – The final output of the plan. You can choose to publish the finalized demand plan to downstream inventory and supply planning systems for implementation.

Understanding Your Homepage

Welcome to Amazon Connect Decisions. The homepage is your central hub for supply chain decision intelligence. This guide walks you through each section of the homepage so you can quickly orient yourself and start taking action. The homepage is comprised of four key sections: Section 1: Metrics, Section 2: Top topics for today, Section 3: Natural Language Interface (NLI) / Decisions Teammate, and Section 4: Navigation bar.



Key metrics

The top of the homepage displays a maximum of five metric cards that give you a real-time snapshot of your supply chain health. Depending on the rules defined by the manager persona, these metrics can vary for your account compared to other accounts.

Metric	Description
Total Open Insights	The total number of active insights across all monitors. Includes a daily count of new insights created today.
Low Safety Stock (sample)	Tracks SKUs where inventory has fallen below the safety stock threshold.

Metric	Description
Projected Out of Stock (sample)	Flags items at risk of stocking out based on current demand and supply signals.
Excess Safety Stock (sample)	Identifies items carrying more inventory than needed, tying up working capital.
Days of Cover (sample)	Monitors how many days of supply remain for tracked items.

Each card shows:

- The **current total** of open insights in that category
- A **daily increment** (for example, +308 created today) indicating how many new insights were generated since midnight

 Tip

Click on any metric card to drill into the full list of insights for that category.

 Note

The Admin persona also sees metrics related to data validation errors and the total number of users in the account.

Top topics for today

Below the metrics, the "Top Topics for Today" section surfaces the most important insight groups for your attention. Each topic card displays the number of insights and their associated financial impact for the planner persona. These topic cards vary if you are an Admin.

Topic card	Number of insights (sample values)	Financial impact (sample values)	What it means
Highest Financial Impact	73	\$10,530,538	Insights with the greatest dollar value at stake. Start here for maximum ROI.
Most Urgent to Act On	100	\$991	Time-sensitive insights requiring immediate action, regardless of financial size.
In Progress	13	\$1,694,240	Insights your team is actively working on.
Overdue	100	\$0	Insights that have passed their action deadline and need immediate follow-up.

Each card includes a **Review** button. Clicking it takes you to the full filtered list of insights for that topic.

Note

Financial impact values are system-calculated estimates based on inventory exposure and demand signals. This can only be shown if your admin has provided financial information to Amazon Connect Decisions, otherwise it will be zero.

Natural Language Interface (NLI) / Decisions Teammate

The Decisions Teammate panel is located on the right side of the homepage. It is an AI-powered teammate that helps you navigate the application, interpret insights, and answer supply chain questions in plain language. It is always available across the application.

How to use it:

1. Type a question in the **"Ask a question"** text field.

2. The teammate responds with contextual answers drawn from your data and insights.
3. You can ask questions like: *"What are my top stockout risks this week?"*, *"Show me insights for a specific product or region"*, *"What does days of cover mean for my portfolio?"*

Best ways to use NLI:

- When you need a quick summary of your current supply chain insights and plans.
- When you want to explore insights without manually navigating filters.
- When you need help understanding a specific insight, root cause, or recommendation.
- When you want to know more about the demand plan or supply plan.
- When you want to share feedback about system-generated insights or plans.

As an agentic application, leverage the NLI whenever you want to learn more, take actions, and share feedback. The capabilities of NLI will continue to be enhanced over time.

Navigation bar

Amazon Connect Decisions uses two navigation elements to help you move through the platform.

Top Navigation Bar

- Amazon Connect Decisions logo – Click to return to the homepage from anywhere.
- Home – Returns you to this homepage.
- Insights – Opens the full Insights view with filtering, sorting, and detail capabilities.
- Plans – Navigates to your Demand Plans and/or Supply Plans.
- Notification bell – Alerts you to system updates.
- User profile – Access your account settings, preferences, and sign-out options.

Left Sidebar Menu

- Home – Homepage shortcut.
- Insights – Expands to show sub-navigation including Configuration, where admins can manage insight monitors and thresholds.
- Plans – Access demand and supply planning workflows.

- Collapse/Expand button – Toggle the sidebar to maximize your workspace.

 Tip

The left sidebar can be collapsed to give you more screen space.

Onboarding

Onboarding for Amazon Connect Decisions

Amazon Connect Decisions's onboarding agent serves as your AI setup teammates, abstracting away the technical complexity of configuring insights monitoring. Instead of manually defining SQL queries, mapping data schemas, or translating business requirements into technical specifications, you work conversationally with the agent to set up your monitoring system.

What the onboarding agent can help you with:

1. **Understanding your business context:** Share your standard operating procedures, business policies, and operational guidelines in natural language or by uploading existing documents. Agents analyze your materials and extract relevant configuration parameters automatically.
2. **Translating requirements into configurations:** Describe what you want to monitor in plain language; for example, "I need to track forecast accuracy for high-value products" or "Alert me when inventory drops below safety stock levels." The agent translates these business requirements into the appropriate metric definitions, rules, and thresholds without requiring SQL knowledge.
3. **Iterating on configurations:** Refine your monitoring setup through conversation. Ask questions like "Can we make this rule more sensitive?" or "Should we add filtering criteria for specific product categories?" The agent helps you explore options and understand trade-offs before committing changes.
4. **Validating changes before production:** Use the preview function to see how your configurations would perform against real data, enabling you to review insights that would be generated with your new settings before you save changes to production.
5. **Troubleshooting and optimization:** When configurations don't produce expected results, describe what you're seeing and the agent helps diagnose issues and recommend improvements.

This conversational approach means you can configure sophisticated supply chain monitoring without deep technical expertise. AI teammates handle the complexity while you focus on defining your business requirements and operational priorities.

Prerequisites

Prerequisites

Before configuring insights, ensure:

- Amazon Connect Decisions instance is set up with data flows configured and running
- Manager role permissions to access configuration settings
- Required data entities uploaded to your Supply Chain Data Lake, including:
 - Historical demand or sales data (minimum 12 months recommended)
 - Product master data with classifications and attributes
 - Site and location information
 - Current inventory positions
 - Supplier and lead time data (if configuring supply monitoring)
- Standard operating procedures or business policies prepared for upload (optional but recommended to accelerate configuration)

Decisioning

Amazon Connect Decisions continuously monitors your supply chain to detect issues that require attention and generate actionable recommendations to resolve them. The Decisioning section of this guide covers everything you need to get the system working effectively — from teaching it what to watch for, to reviewing and acting on the insights it surfaces.

What are Insights?

Insights are the primary output of Amazon Connect Decisions. Each insight represents a detected supply chain issue, paired with root cause analysis and recommended actions to resolve it. Insights are generated when your configured monitoring rules detect that a metric has crossed a defined threshold — for example, when projected inventory falls below safety stock minimums, or when days of supply drop below an acceptable level.

Each insight includes:

- A description of the detected issue and its business context
- Root cause analysis explaining why the issue occurred
- Recommended actions with specific parameters such as quantities, locations, and timelines
- Financial impact to help you prioritize your response (if data provided by the admin)
- Related insights that may share the same underlying cause

Insights cover both demand and supply monitoring. They are ranked by severity based on financial impact, so your team can focus on the issues that matter most to your operations.

How Decisioning Works

Decisioning operates in two phases that work together: configuration and management.

Configuration is how you teach Amazon Connect Decisions what to monitor, how to interpret what it finds, and how to prioritize what surfaces to your team. You define the metrics and rules that trigger insights, provide guidelines that shape root cause analysis and recommendations, and assign financial impact factors that determine severity ranking. Configuration is completed once during setup and refined over time as your operational needs evolve.

Managing Insights is the day-to-day experience of reviewing, filtering, and acting on the insights the system generates. Once configuration is in place, your team uses the Insights page to find relevant issues, understand their root causes, and execute or dismiss the recommended actions.

What You'll Find in This Section

Configuring Insights

- **How are Insights Generated:** Understand the four-stage process that transforms your supply chain data and configuration into actionable insights
- **Knowledge Sources:** Share SOPs and documented business best practices so Amazon Connect Decisions understands your operational context before generating metrics, rules, and guidelines
- **Detection:** Define the metrics and rules that determine when an insight is triggered for demand or supply monitoring
- **Guidelines for Root Causes and Recommendations:** Shape how root cause analysis and recommendations are generated so they align with your business practices and operational constraints
- **Prioritization and Severity:** Assign financial impact to insight types so Amazon Connect Decisions can rank insights by business priority

Managing Insights

- **Filtering and Sorting Insights:** Navigate the Insights page, apply filters by product hierarchy, site, severity, and custom segments, and sort by impact or urgency to focus on what matters most
- **Understanding Insights Details:** Review root cause analysis, key metrics, recommendations, and related insights for a specific insight
- **Taking Action on Recommendations:** Accept, discard, or mark recommendations as complete, and track the progress of resolution activities
- **Providing Feedback on Insights:** Rate the quality of insights and recommendations to help the system improve over time

Configuring Insights

Insights configuration is how you teach Amazon Connect Decisions what to watch for, how to interpret what it finds, and how to prioritize what surfaces to your team. Configuration applies to both demand and supply monitoring. There are four areas to work through in sequence: Knowledge Sources, Detection (Metrics and Rules), Root Causes and Recommendations (Guidelines), and Prioritization (Severity). A preview is available at each step so you can validate before activating any configurations.

To open configuration: From the home page, navigate to **Insights** in the left-hand navigation, then select **Configuration**.

The four configuration areas

- **Knowledge Sources:** Share SOPs and documented business best practices so Amazon Connect Decisions understands your operational context before generating metrics, rules, and guidelines.
- **Detection:** Define the metrics and rules that determine when an insight is triggered for demand or supply monitoring.
- **Guidelines:** Shape how root cause analysis and recommendations are generated so they align with your business practices.
- **Severity:** Assign financial impact to insight types so Amazon Connect Decisions can rank Insights by business priority.

How configuration connects to Insights output

Each area feeds the Insights experience directly. Detection determines what surfaces; guidelines shape the analysis and recommended actions; severity determines the order in which Insights appear. Changes to any area take effect as soon as the relevant items are set to Ready status.

Best practices

- **Start with clear monitoring objectives:** Before creating metrics and rules, define what supply chain performance issues matter most to your operations. This helps you configure focused, actionable monitoring rather than generating excessive alerts.
- **Leverage your onboarding teammate:** Use natural language to describe monitoring needs rather than attempting to write technical specifications. Ask questions like “Can we make this

more sensitive to A-class products?" or "What does this field do?" Connect Decisions agents translate your business requirements into appropriate configurations.

- **Preview extensively:** Always preview metric and rule changes before activating them. Use the preview function to:
 - Validate that metrics calculate correctly against your data
 - Confirm that rules trigger insights at appropriate thresholds
 - Filter preview results by products and sites to understand impact across your supply chain
 - Iterate on configurations until preview results match your operational needs
- **Use explainability instead of SQL:** When reviewing how metrics and rules operate, use the Explainability tab rather than examining SQL queries directly. This helps you understand the logic in plain language without technical complexity. When you see something that doesn't match your expectation work with the agent to correct that via explainability.
- **Iterate with draft status:** Take advantage of the Draft status to configure and test metrics and rules without impacting production insights. Only change status to Ready when you're confident in the configuration.
- **Balance sensitivity:** Set thresholds that catch meaningful insights without generating excessive noise. Use preview to find the right balance between alert coverage and operational manageability.
- **Start simple, then expand:** Begin with a few key metrics and rules for your highest-priority monitoring needs. Once you're comfortable with the configuration workflow, expand to additional monitoring areas.
- **Review and refine regularly:** After activating metrics and rules, monitor the quality and relevance of generated insights on your Planning Intelligence home page. Return to Monitoring Configuration to refine settings based on operational experience.

How are Insights generated?

Amazon Connect Decisions uses a systematic process to monitor your supply chain data, detect issues, and generate insights with actionable recommendations. Understanding this process helps you configure effective rules and interpret the insights you receive.

The Insight Generation Process

Insight generation follows a four-stage process that transforms your supply chain data into actionable intelligence:

1. Metric Calculation

The system continuously calculates metrics based on your supply chain data. These metrics are quantifiable measurements that assess performance across your operations, such as:

- Projected inventory levels
- Days of supply
- Inventory turns
- Lead time variability
- Forecast accuracy

Metrics are calculated at the granularity you define, such as by product, site, or product-site combination. The system updates these calculations based on the frequency you configure (daily, weekly, or as new data arrives).

2. Rule Evaluation

Once metrics are calculated, Amazon Connect Decisions evaluates them against your configured metric-based rules. Metric-based rules define the specific conditions under which you want to be alerted to potential issues.

Each metric-based rule includes three essential components:

Metrics: The quantifiable measurements being monitored

Thresholds: The boundary values that trigger an insight when crossed

Scope: The products, sites, or other dimensions the rule applies to

For example, a rule might state: "Alert when projected inventory falls below safety stock minimum AND days until stockout is 14 or fewer AND customer impact risk exceeds \$25,000."

When a rule's conditions are met, the system initiates the insight generation process for the affected items.

3. Root Cause Analysis

When a rule is triggered, Amazon Connect Decisions automatically performs root cause analysis to understand why the issue occurred. The system:

- Examines relevant supply chain data across multiple dimensions
- Reviews historical patterns and recent changes
- Analyzes relationships between different factors (demand, supply, inventory, orders)
- Applies your policy-based rules to provide business context

Policy-based rules guide this analysis by providing qualitative guidelines on how the system should consider and analyze problems. For example, a policy-based rule might state: "For inventory shortage insights, always analyze the following root causes: demand forecast error, supplier lead time issues, production capacity constraints."

The root cause analysis identifies the primary drivers behind the issue and provides detailed explanation of contributing factors.

4. Insight Creation and Recommendation Generation

After completing the root cause analysis, the system creates the insight with:

- A clear description of the issue
- The root cause explanation
- Relevant metrics and data visualizations
- Priority classification based on your configured prioritization factors
- Recommended actions to resolve the issue
- Alternative actions for consideration

Recommendations are generated based on your business rules, operational constraints, and the specific context of the issue. The system considers factors such as available inventory at other locations, supplier lead times, production capacity, and financial impact when formulating recommendations.

Timing and Frequency

Insights are generated based on the frequency you configure in your metric-based rules (typically daily or weekly). The system processes new data according to your data refresh schedule, recalculates metrics, evaluates rules, and generates insights for any new issues detected.

Existing insights are automatically updated or marked as complete when new data shows that the issue no longer meets the configured thresholds.

Knowledge Sources

Knowledge Sources are documents you share with Amazon Connect Decisions that give it context about your operations and monitoring requirements. Amazon Connect Decisions analyzes them and uses that understanding to inform how metrics, rules, and guidelines are configured in subsequent steps. Typical sources include SOPs, business policies, and operational guidelines.

What to share: Effective Knowledge Sources include documents that describe how your business operates and how you manage supply chain issues. Examples include:

- Inventory management SOPs (e.g., how you respond to stockouts, when to expedite replenishment, safety stock policies)
- Demand planning policies (e.g., how you handle seasonal demand, promotional lift, or new product introductions)
- Supplier management guidelines (e.g., lead time assumptions, preferred carrier rules, escalation procedures)
- Operational constraints (e.g., warehouse capacity limits, order minimums, regional fulfillment rules)

The more specific your documents are to your actual processes, the more accurately Amazon Connect Decisions can generate relevant metrics, rules, and guidelines on your behalf.

- **Supported formats:** PDF, Word, and plain text files are supported.
- **Template:** A Knowledge Source template is available for download to help you structure your inputs. Use it as a starting point if you do not have existing documentation.

To share a Knowledge Source:

1. From the Insights Configuration page, select the **Knowledge Sources** tab.
2. Download and complete the template if needed, or gather your existing documents.
3. Click **Upload** and select your file. Amazon Connect Decisions will process the document and confirm when it is ready to generate further configurations.

Tips

- Share multiple sources to give Amazon Connect Decisions comprehensive context; a single SOP is a good starting point, but adding related policies and guidelines improves output quality.
- You can share additional Knowledge Sources at any point during configuration, not only at the start. If you update a source document, share the new version to keep Amazon Connect Decisions current.
- Knowledge Sources provide reference context; they do not directly create metrics, rules, or guidelines on their own. Use Detection and Guidelines configuration to act on what Amazon Connect Decisions learns.

Detection (Metrics and Rules)

Detection configuration defines what Amazon Connect Decisions measures and when those measurements become actionable Insights for both demand and supply monitoring. It has two components that work together:

- **Metrics:** Quantifiable calculations Amazon Connect Decisions runs against your supply chain data; for example, forecast accuracy over the last four weeks for A-class products, or inventory days of supply by site. Metrics can function as standalone KPIs or serve as inputs into rules.
- **Rules:** Conditions that trigger an Insight by referencing one or more metrics and firing when thresholds are met; for example, a rule that triggers when forecast accuracy drops below 85% or when days of supply falls below the reorder threshold.

Metrics provide the measurement; rules determine when that measurement becomes something your team should act on.

Configuring Metrics

The Insights configuration tab displays your configured metrics and rules alongside a preview of the insights they generate under the Detection tab. This unified view allows you to see both your monitoring configurations and their operational impact in one place.

From this page, you can:

- View all configured metrics and rules for both demand and supply monitoring
- Edit individual metrics or rules

- Review the overall insights that would be generated with your current configurations

Creating a New Metric

To create a new metric, you can either:

- **Use Natural Language:** Ask the onboarding agent to create a metric for you. For example: "Create a metric to track forecast accuracy for A-class products over the last 4 weeks." The agent translates your requirements into a metric configuration.
- **Use the Create Metric Button:** Click "Add metrics" to open the metric creation interface and define parameters directly.

When you create or edit a metric, the system saves it with a "Draft" status. This allows you to configure and test the metric without impacting production insights.

Reviewing an Existing Metric

When you click on a metric, you arrive at the metric review page where you can:

- **Review the configuration:** See the current setup of your metric, including its descriptors, dimension grain, time span, and other parameters.
- **Ask questions with natural language:** Work with the onboarding teammate to understand and refine your metric. Ask questions like:
 - "How does this metric work?"
 - "Can we confine this to only A-class products?"
 - The agent helps you iterate on the metric definition conversationally, translating your business requirements into the appropriate configuration.
- **Review explainability:** Instead of reviewing SQL queries directly, view the Explainability section that describes how the metric operates in plain language. To do so,
 - Un-collapse "SQL Query" section
 - Click on "Explainability" tab
- **Preview individual metric impact:** Use the metric preview to see how the metric would be generated at runtime with your current configuration. Filter the preview by specific products and sites to understand how the metric will impact different parts of your supply chain. To do so:
 - Expand the "Metric Preview" element at the bottom of the screen.
 - Click "Preview metric" to generate runtime metrics.

- Refine your preview with search criteria across products and sites.
- **Iterate with preview:** As you work with the onboarding agent to refine the metric, use preview to validate changes. Continue iterating until the preview results match your operational needs, then save your configuration.

Edit fields directly (optional): If you prefer to configure the metric yourself, click "Edit" to change fields directly without agent support. Even when editing manually, you can still ask the onboarding teammate questions like:

- "What does this field do?"
- "Can you help me edit my forecast accuracy SQL?"
- "How should I set the temporal grain for weekly reporting?"

Your onboarding teammate provides interim configuration support even when you're working in the direct editing interface.

Activating Your Metric

Once you're satisfied with your metric configuration and preview results:

1. Change the metric status from Draft to Ready.
2. Click "Save changes" to activate the metric.

The metric will now recalculate for your production insights based on your configured thresholds.

Configuring Rules

The "Insights" configuration tab displays your configured metrics and rules alongside a preview of the insights they generate under the "Insights" tab. This unified view allows you to see both your monitoring configurations and their operational impact in one place.

Creating a New Rule

To create a new rule, you can either:

- **Use natural language:** Ask your onboarding teammate to create a rule for you. For example: "Create a rule to alert when forecast accuracy drops below 85% for A-class products." The agent translates your requirements into a rule configuration.

- **Use the create rule button:** Click *"Add rule"* to open the rule creation interface and define parameters directly.

When you create or edit a rule, the system saves it with a "Draft" status. This allows you to configure and test the rule without impacting production insights.

Reviewing an Existing Rule

When you click on a rule, you arrive at the rule review page where you can:

- **Review the configuration:** See the current setup of your rule, including metric connections, thresholds, conditions, and granularity settings.
- **Ask questions with natural language:** Work with the onboarding teammate to understand and refine your rule. Ask questions like:
 - "How does this rule determine when to trigger insights?"
 - "Can we add filtering for specific product categories?"
 - The agent helps you iterate on the rule logic conversationally, translating your monitoring requirements into the appropriate configuration.
- **Preview individual rule impact:** Use the rule preview to see how insights would be generated at runtime with your current configuration. Filter the preview by specific products and sites to understand how the rule will impact different parts of your supply chain.
 - Expand the "Preview" element at the bottom of the screen.
 - Click "Preview insight" to generate runtime insights.
 - Refine your preview with search criteria across products and sites.
- **Iterate with Preview:** As you work with the onboarding agent to refine the rule, use preview to validate changes. Continue iterating until the preview results match your operational needs, then save your configuration.

Edit Fields Directly (Optional): If you prefer to configure the rule yourself, click "Edit" to change fields directly without agent support. Even when editing manually, you can still ask the onboarding teammate questions like:

- "What does this threshold field control?"
- "Can you help me limit this to Class A products?"

Your onboarding teammate provides interim configuration support even when you're working in the direct editing interface.

Activating Your Rule

Once you're satisfied with your rule configuration and preview results:

1. Change the rule status from Draft to Ready.
2. Click "Save changes" to activate the rule.

The rule will now generate insights in production based on your configured thresholds and conditions.

Reviewing Overall Insights

After configuring and activating your metrics and rules, review the combined impact of all your configurations:

1. Navigate to the **Detection** tab on the Monitoring Configuration page and expand the preview panel.
2. This global view preview displays the complete set of insights that would be generated with all your current configurations, showing you the combined effect of your metrics and rules.
3. Use this view to assess whether your monitoring setup produces the right volume and insights for your operations.
4. If you need to adjust the sensitivity or coverage of your monitoring, work directly with the onboarding agent to tweak them or navigate to individual metrics or rules to refine them further.

Guidelines

Guidelines for Root Causes and Recommendations

Guidelines help shape how Amazon Connect Decisions generates root cause analysis and recommendations for insights. These guidelines act as policies that ensure AI-generated explanations and suggested actions align with your business practices and operational constraints.

Accessing Guidelines: From the Insights Configuration page, select the **Root Causes & Recommendations** tab.

What Guidelines Do: Guidelines provide context and constraints that influence how the system:

- Analyzes root causes of insights and alerts
- Generates actionable recommendations for resolution
- Prioritizes remediation approaches based on your business rules
- Adapts guidance to your operational environment and supply chain constraints

Working with your onboarding teammate

Use your onboarding teammate to create and refine guidelines conversationally. Ask questions like:

- "Create a guideline to prioritize replenishment for high-velocity SKUs during promotional periods"
- "Add a guideline to flag slow-moving inventory exceeding 45 days of cover for markdown review"
- "Can we create a guideline that ensures minimum order quantities align with vendor lead times?"

Share your standard operating procedures or business policies to help the teammate understand your requirements. The teammate analyzes your inputs and translates them into guidelines and guardrails that shape AI teammate's behavior.

Creating Guidelines

To create a new guideline, you can either:

- **Use natural language:** Ask the onboarding teammate to create guidelines for you. For example: "Create a guideline that prioritizes inventory reduction strategies over expedited shipping when addressing excess inventory."
- **Use the Create Guideline button:** Click **Add guidelines** to define parameters directly.

Example Guidelines

- "Prioritize replenishment for high-velocity SKUs during promotional periods"
- "Ensure minimum order quantities align with vendor lead times"
- "Monitor for demand spikes driven by seasonal or promotional activity"
- "Escalate stockout risks for A-class products to planning managers within 24 hours"

Reviewing and Editing Guidelines

Select any guideline to view its configuration. Work with the onboarding teammate to understand how the guideline influences insights. Ask questions like:

- "How does this guideline affect recommendations for excess inventory insights?"
- "Can we make this guideline more specific to A-class products?"
- "Show me examples of recommendations before and after applying this guideline"

Activating Guidelines

Guidelines operate with a Draft/Ready status model:

1. Create and refine guidelines in Draft status to test their impact
2. Iterate with the onboarding teammate until the guideline produces desired behavior
3. Change status to **Ready** to apply the guideline to all future root cause analysis and recommendations
4. Click **Save changes** to activate

Active guidelines immediately influence how the system analyzes insights and generates recommendations.

Prioritization and Severity

Severity factors enable prioritization based on financial impact, helping you focus on insights that affect your highest-value operations.

Accessing Severity Configuration

From the Insights configuration page, navigate to the **Prioritization** tab.

What Severity Factors Do

Severity factors translate insights into business impact by calculating:

- **Revenue impact:** Lost or at-risk revenue from supply chain disruptions (e.g., stockouts, delayed shipments, unfulfilled demand)

- **Cost impact:** Additional costs incurred from supply chain inefficiencies (e.g., excess inventory carrying costs, expedited shipping, obsolescence)

Working with Your Onboarding Teammate

Use your onboarding teammate to configure severity factors conversationally. Ask questions like:

- "How should we calculate revenue impact for different product categories?"
- "What cost factors should we include for excess inventory situations?"
- "Can we weight revenue impact differently for A-class versus C-class products?"

Share your product financial data, cost assumptions, or historical impact analysis. The teammate analyzes your inputs and recommends severity factor configurations that align with your business priorities.

Using Financial Impact in Prioritization

Once configured, severity factors enhance insight filtering and sorting on your Insights pages. The system displays calculated financial impact for each insight based on the type:

- Stockout insight show revenue impact (estimated lost sales)
- Excess inventory insight show cost impact (carrying costs and obsolescence risk)
- Supply disruption insight show revenue impact (delayed fulfillment)
- Overstock insight show cost impact (additional warehousing and capital costs)

You can filter and sort by financial impact to create a remediation roadmap that maximizes ROI, focusing first on insights with the highest financial implications.

Example: A stockout insight might show \$50,000 in daily revenue impact based on lost sales. An excess inventory insight might show \$2,000 in daily cost impact from carrying costs plus potential markdown exposure.

Managing Insights

Once Amazon Connect Decisions generates insights based on your configured rules, your day-to-day work centers on reviewing, prioritizing, and acting on them. The pages in this section walk you through that process — from finding the right insights to taking action and closing the loop.

What Managing Insights Covers

Working with insights follows a natural progression: you start by locating the insights most relevant to your role, dig into the details of a specific issue, act on the recommendations the system provides, and share feedback to help the system improve over time. Each page in this section covers one part of that workflow.

In This Section

Filtering and Sorting Insights The Insights page can surface hundreds of issues across your supply chain at any given time. This page explains how to narrow that list using filters, search, and sorting — so you can focus on what matters most to your role and responsibilities.

Understanding Insight Details Once you've identified an insight to investigate, this page explains what you'll find on the Insight Details page — including the root cause analysis, key metrics, supporting data, related insights, and the activity log. Use this page to understand the full context of an issue before deciding how to respond.

Acting on Recommendations Amazon Connect Decisions generates AI-powered recommendations alongside each insight. This page covers how to evaluate those recommendations and execute the suggested actions

Providing Feedback on Insights Your feedback helps the system improve over time. This page explains how to share input on insight quality and recommendation accuracy directly from the Insights Details page.

Filtering and Sorting Insights

Amazon Connect Decisions provides comprehensive filtering and sorting capabilities that help you quickly focus on the insights most relevant to your supply chain responsibilities. You can navigate to the Insights page from multiple entry points and apply filters based on product hierarchies, site hierarchies, insight properties, and custom business segments.

Accessing the Insights Page

You can access the Insights page through several pathways:

- Select **Insights** from the left navigation menu or from top navigation bar
- The Insights page displays with your default view settings

From the Homepage

The homepage provides quick access to insights through status summary cards displaying insight counts by status or severity. Select any cards to navigate to the Insights page with that filter pre-applied.

When you navigate from a homepage card, the Insights page opens with relevant filters already applied.

Amazon Connect Decisions Home | Insights | Plans 🔔 | Shirley Temple ▼

Menu <

- Home
- ▼ **Insights**
 - Configuration
 - Plans
 - Data management
- ▼ Settings
 - Users
 - Roles
 - Application

Insights

Not started: 9,049 | In progress: 11 | Pending resolution: 2 | Completed: 0 | Dismissed: 1 | Archived: 88 | Processing: 0

Save filter sets Choose an option Filter by created on date range Q Product = X

Insight ID	Description	Status	Severity	Assigned to	Impact	Impact date	Created on	Actions
838f9ea6-c839-4b95-8448-1c158f7be5e6	HOKA-CHALLENGER-WB inventory critically below safety stock at DC-SAC-001 due to supply gap and demand surge	In progress	Critical	Unassigned	USD 302,520	April 13, 2026	April 12, 2026 at 01:29 AM GMT+5:30	⋮
619f73c3-ba23-4a1a-a76c-2c7e01c7a2e4	HOKA-SPEEDGOAT-M11 at DC-BOI-001 Near Stockout Due to Chronic Vendor Underdelivery	In progress	Critical	Unassigned	USD 298,418	April 13, 2026	April 12, 2026 at 01:22 AM GMT+5:30	⋮
89d56b7f-a557-49fe-ab33-c199152bee72	Critical Stock-Out Risk: SAL-SENSE-11 at ST-SAC-001 Due to Severe Vendor Underdelivery	In progress	Critical	Shirley Temple	USD 284,333	April 13, 2026	April 12, 2026 at 01:22 AM GMT+5:30	⋮
35480800-8bcf-43a2-a0f1-2a160a0d760	HOKA-TECTON-W9 at ST-SAN-001 Below Safety Stock with 41-Day Replenishment Gap Due to Vendor Underdelivery	In progress	Critical	Unassigned	USD 275,857	April 13, 2026	April 12, 2026 at 01:23 AM GMT+5:30	⋮
79265549-d647-4c8b-wc21-c27b3acdda95	Critical Stock Shortage: HOKA-CHALLENGER-WB at MFC-LA-001 with 89-Day PO Lead Time Delay	Not started	Critical	Unassigned	USD 275,419	April 13, 2026	April 12, 2026 at 01:22 AM GMT+5:30	⋮
01c7209e-d46e-4a84-a76d-4983d5773c9	HOKA-SPEEDGOAT-M10 at ST-BOI-001 Inventory below safety stock due to multi-vendor delivery failures and demand spike	Pending resolution	Critical	Unassigned	USD 249,094	April 13, 2026	April 12, 2026 at 02:26 AM GMT+5:30	⋮
45e43a1e-37b0-4a5d-a172-097660362662	SAL-SENSE-11 at MFC-LA-001 Inventory 32% below safety stock with 68-day supply gap until next delivery	In progress	Critical	Unassigned	USD 245,466	April 13, 2026	April 12, 2026 at 01:22 AM GMT+5:30	⋮
c56f45ce-ad50-482c-b952-29f0a59baa47	HOKA-TECTON-W9 stockout at DC-PDX-001 due to massive demand spike and vendor fill rate collapse	Not started	Critical	Unassigned	USD 228,669	April 13, 2026	April 12, 2026 at 01:22 AM GMT+5:30	⋮
88aec99d-45a5-46c5-9c78-79ac7c0364270	Stockout Risk: HOKA-MAFATE-M10 at ST-BFL-001 Due to NATIONAL/TRADE Underdelivery & Demand Spike	Not started	Critical	Unassigned	USD 224,380	April 13, 2026	April 12, 2026 at 01:22 AM GMT+5:30	⋮
8a84541a-0a85-47b4-933a-21a2fc3676d0	SAL-SENSE-11 Below Safety Stock at ST-LA-001 Due to Vendor Fill Rate Failures & Delays	Not started	Critical	Unassigned	USD 223,117	April 13, 2026	April 12, 2026 at 01:22 AM GMT+5:30	⋮
6a224473-4585-4644-9b85-3a5b398b8f1cc	SAL-SENSE-10 at DC-SAN-001 Stock below safety level due to ALLIED-VENDORS delivery failure on PO_00000227	Not started	Critical	Unassigned	USD 222,581	April 13, 2026	April 12, 2026 at 01:27 AM GMT+5:30	⋮
a7402e3-9ff5-433f-87ec-38c88f5b7144	Critical stock shortage of SAL-SPEEDCROSS-10 at ST-BFL-001 due to vendor ATLANTIC-TRADING delivery failure	Not started	Critical	Unassigned	USD 217,077	April 13, 2026	April 12, 2026 at 01:22 AM GMT+5:30	⋮
67b7b494-9e45-41b7-b59d-5dfe8874e9be	SAL-XA-PRO-W9 at DC-LA-001 Stockout Risk from 93-Day Replenishment Gap and Vendor Shortfall	Not started	Critical	Unassigned	USD 217,024	April 13, 2026	April 12, 2026 at 01:23 AM GMT+5:30	⋮
b59f4453-4664-41ca-9641-d689f7765db2	SAL-SPEEDCROSS-11 at ST-BOI-001 Inventory below minimum due to vendor under-delivery and demand surge	Not started	Critical	Unassigned	USD 215,208	April 13, 2026	April 12, 2026 at 01:22 AM GMT+5:30	⋮

Portion of insights not visible due to pagination

Show me all available insights

Decisions teammate Response events ▼

Available Insights

Here are the first 10 insights retrieved. More insights exist beyond this sample – see the status breakdown below.

Status Breakdown

Status	Count
Not started	9,049
In progress	11
Pending resolution	2
Dismissed	1
Archived	88

First 10 Insights

Title

HOKA-CHALLENGER-WB inventory critically below safety stock at DC-SAC-001

HOKA-SPEEDGOAT-M11 at DC-BOI-001 Near Stockout Due to Chronic Vendor Underdelivery

Critical Stock-Out Risk: SAL-SENSE-11 at ST-SAC-001 Due to Severe Vendor Underdelivery

HOKA-TECTON-W9 at ST-SAN-001 Below Safety Stock with 41-Day Replenishment Gap Due to Vendor Underdelivery

Critical Stock Shortage: HOKA-CHALLENGER-WB at MFC-LA-001 with 89-Day PO Lead Time Delay

HOKA-SPEEDGOAT-M10 at ST-BOI-001 Inventory below safety stock due to multi-vendor delivery failures and demand spike

SAL-SENSE-11 at MFC-LA-001 Inventory 32% below safety stock with 68-day supply gap until next delivery

HOKA-TECTON-W9 stockout at DC-PDX-001 due to massive demand spike and vendor fill rate collapse

Stockout Risk: HOKA-MAFATE-M10 at ST-BFL-001 Due to NATIONAL/TRADE Underdelivery & Demand Spike

SAL-SENSE-11 Below Safety Stock at ST-LA-001 Due to Vendor Fill Rate Failures & Delays

SAL-SENSE-10 at DC-SAN-001 Stock below safety level due to ALLIED-VENDORS delivery failure on PO_00000227

Critical stock shortage of SAL-SPEEDCROSS-10 at ST-BFL-001 due to vendor ATLANTIC-TRADING delivery failure

SAL-XA-PRO-W9 at DC-LA-001 Stockout Risk from 93-Day Replenishment Gap and Vendor Shortfall

SAL-SPEEDCROSS-11 at ST-BOI-001 Inventory below minimum due to vendor under-delivery and demand surge

I examined the first 10 insights. There are more insights available – would you like me to load the next page, filter by a specific product, status, or severity?

Load more insights Filter by severity Filter by status

🔔 📄 📁 Sources (1) ▼

Ask a question

Understanding the Insights Page Layout

The Insights page is organized into four main sections:

Status Summary Bar

At the top of the page, status cards display insight counts by status category:

- **Not started:** Newly created insights with no user interaction
- **Processing:** Insights currently being analyzed or acted upon by the system
- **In progress:** Insights where you have initiated action on a recommendation
- **Pending resolution:** Insights that have been actioned but are not yet fully resolved

- **Completed:** Insights that are resolved and no longer meet the rule thresholds
- **Dismissed:** Insights you have explicitly chosen to discard
- **Archived:** Insights that have been moved to archive storage

Each card shows the status name and count in blue. Select any status card to filter the insights table to that status.

Filtering and Search Area

Below the status bar, you'll find filtering controls:

Saved filter sets: A dropdown on the left allows you to apply pre-configured filter combinations

Date range filter: In the center, "Filter by created on date range" displays the currently selected date range (for example, "2025-01-01 — 2025-12-31"). Select to open a date picker where you can choose start and end dates.

Search filter properties: On the right, a search bar allows you to search for specific insights by Task ID, product name, site name, or other properties

Active Filters Display

Below the filtering controls, applied filters appear as removable chips showing:

- The column name and filter criteria (for example, "Severity = High")
- An X button to remove individual filters
- A results count showing matching insights (for example, "279 Matches")
- A "Clear filters" button with dropdown for filter management options

Insights Table

The main table displays all insights with the following columns:

- **Insight ID:** Unique Task ID as a clickable link to view insight details
- **Description:** Brief summary of the issue detected
- **Status:** Current workflow state with icon indicator
- **Severity:** Priority level (Critical, High, Medium, Low) with color coding
- **Urgency:** Time-sensitive priority with sort capability

- **Assigned to:** User or team responsible for the insight
- **Impact:** Financial or operational impact when calculable
- **Impact date:** When the issue is expected to occur
- **Created on:** When the system detected the insight (with date and time)
- **Product:** Product identifier affected by the insight
- **Actions:** Three-dot menu with quick actions for each insight

Filtering Insights

Amazon Connect Decisions supports multiple filtering approaches to help you find relevant insights quickly.

Using the Search Bar

The "Search filter properties" bar provides quick access to filtering:

1. Select the search bar
2. Begin typing your search term (Task ID, product name, site name, keywords)
3. The system provides real-time search results as you type
4. Applied filters appear as chips below the search area

Hierarchical Filtering (Product and Site)

When filtering by product or site hierarchies, the system searches across all five hierarchy levels simultaneously.

To apply a hierarchical filter:

1. Use the search bar to enter a product or site name
2. The system displays matching results grouped by hierarchy level (Level 1 through Level 5)
3. Select the hierarchy value you want to filter by
4. Click on "Apply"
5. A filter chip appears showing the complete hierarchy path

The system uses prefix-based matching with multi-word support. For example, typing "temp" matches "Temperature Sensors" because it matches the prefix of "Temperature."

When you select a hierarchy level, the filter automatically includes all child items within that hierarchy. For example, filtering by "Electronics" at Product Level 1 shows all insights for any products within the Electronics category, regardless of their specific subcategory or SKU.

Filtering by Insight Properties

You can filter by specific columns to find insights based on exact field values.

To apply a property filter:

1. Use the search bar or column-specific filters
2. The filter input adapts based on field type:
 - **Text fields** (Status, Severity): Provides dropdown with predefined options
 - **Numeric fields** (Impact): Provides numeric input with comparison operators
 - **Date fields** (Created on, Impact date): Provides date picker with range selection
3. Select or enter your filter value
4. Click on "Apply"
5. A filter chip appears showing the property name and value (for example, "Status: Not started")

Date Range Filtering

The date range filter provides a dedicated interface for filtering by insight creation dates.

To apply a date range filter:

1. Select "Filter by created on date range"
2. The date picker displays showing the current date range
3. Select the start date and end date from the calendar
4. The filter applies automatically, showing the date range
5. The results count updates to show matching insights

Segmentation-Based Filtering

If you have configured custom business segments, you can use them to filter insights based on meaningful business categories such as customer tiers, product lines, or geographic regions.

To apply a segment filter:

1. Use the search functionality to find segment types
2. Choose the segment type you want to use (for example, "Customer Tier")
3. Select the segment value (for example, "Tier 1")
4. Click on "Apply"
5. A filter chip appears showing "Segment Type: Segment Value"

Setting Up Segmentation

Segmentation allows you to group products, sites, customers, and channels based on business criteria that matter to your operations. To use segmentation filtering, you must first upload a Segmentation Rules table that defines your business groupings.

Segmentation table logic:

- **Within a single row (AND logic):** When you fill in multiple fields in one row, ALL of those fields must match for a record to belong to that segment
- **Across multiple rows (OR logic):** When you create multiple rows with the same segment_type and segment_value, a record qualifies if it matches ANY of those rows
- **NULL values (wildcards):** Leaving a field blank means "match any value"

The segmentation table supports up to five hierarchy levels for both products and sites, along with additional fields like city, state, country, trading partner, company, and channel.

Important considerations:

- Segment updates apply only to future insights, not historical data
- segment_type and segment_value must each be 30 characters or fewer
- A single insight can belong to multiple segments
- Use hierarchy levels when possible, rather than listing hundreds of individual IDs

Managing Applied Filters

All applied filters appear as removable chips below the search area.

To remove a single filter:

1. Select the X button on the filter chip

2. The system removes that filter and updates results

To remove all filters:

1. Select **Clear filters**
2. The system removes all filter chips and displays the unfiltered view

Note

The Access Control filter toggle (if enabled in your profile settings) operates independently and is not removed by the "Clear filters" action.

Using Saved Filter Sets

Saved filter sets allow you to store frequently used filter combinations and apply them quickly in future sessions.

To apply a saved filter set:

1. Select the **Saved filter sets** dropdown
2. Choose the filter set you want to apply
3. Click on "Apply"
4. The system applies all filters in that set at once
5. Filter chips appear for each filter in the set

Using the Access Control Filter Toggle

The Access Control filter toggle restricts visible data to products and sites matching your assigned permissions. This toggle operates independently from the filter chip system.

To configure the Access Control toggle:

1. Select your name in the top right corner
2. Choose **Profile** from the dropdown menu
3. Navigate to the **Assigned Scope** tab

4. Toggle **Filter views by my assigned scope** on or off
5. Select **Save**

When the toggle is enabled:

- A visual indicator appears on the Insights page showing that Access Control filtering is active
- You see only insights for products and sites matching your assigned scope
- This filtering applies in addition to any filter chips you manually apply

When the toggle is disabled:

- You see all insights across the organization (subject to your role permissions)
- Manual filter chips still apply normally

If your administrator has configured the system to enforce Access Control filtering for all users, you cannot disable this toggle.

Sorting Insights

All columns in the Insights table support sorting, helping you organize insights by priority, impact, or timing.

To sort the Insights list:

1. Select any column header in the Insights table
2. The system sorts insights by that column in ascending order
3. An arrow appears on the column header showing the sort direction
4. Select the same column header again to toggle to descending order

The default sort order is by Severity (Critical to Low). Your sort selection persists during your current session and works with applied filters.

Understanding Filter Results

The results count appears below the search area, showing how many insights match your current filters.

Results count format: "X Matches" where X is the count of insights meeting all applied filter criteria

When filters return no results:

The system displays: "No insights found for selected filters. Try removing some filters or adjusting your criteria."

The empty state includes:

- A **Clear filters** button
- A list of currently applied filters
- If the Access Control toggle is ON and you see no results, the message indicates: "No insights available for your assigned products/sites. Contact your administrator if you believe this is incorrect."

Insights Details

Understanding Insights Details

The Insights Details page provides comprehensive information about a specific insight, including root cause analysis, key metrics, recommendations, and related insights. This page helps you understand the issue, assess its impact, and determine the appropriate actions to resolve it.

Accessing Insight Details

To view detailed information about an insight:

1. Navigate to the Insights page
2. Locate the insight you want to review in the Insights table
3. Select the **Insight ID** (displayed as a blue hyperlink)
4. The Insights Details page opens

Page Header and Navigation

At the top of the page:

Breadcrumb navigation: Shows your current location (Home > Insights > Filtered Insights) and allows you to return to the Insights listing page

Last updated timestamp: Shows when the insight was last modified (for example, "Last updated April 16, 2025 03:07 UTC")

[illegible]

Insight Overview

The main header displays key information about the insight:

Insight Title: A descriptive summary of the issue (for example, "ACDN-M15-RPA-0070 at DC-SAN-001: No sourcing rules configured, zero inbound supply against massive demand causing stockout")

Key Metrics (displayed on the right):

- **Impact Date:** When the issue is expected to occur
- **Impact:** Financial impact amount (for example, "USD 3,561,955")
- **Insight ID:** Unique identifier with copy functionality
- **Created:** Date when the insight was generated

Summary: A brief paragraph explaining the core issue, business context, and why the insight was generated

Root Cause Analysis Section

Below the overview, the Root Cause Analysis section provides detailed investigation of what caused the issue.

Section Structure: The root cause analysis appears as expandable sections with descriptive titles

To view root cause details:

1. Locate the Root Cause Analysis section
2. Select the arrow icon next to any section to expand it
3. The expanded section displays:
 - Detailed explanation of contributing factors
 - Quantitative evidence and data points
 - Timeline of events or conditions
 - Related business context

Common root cause categories include:

- Forecasted Inventory Decline
- Identified Root Cause (for example, Missing Sourcing Rules)
- Evaluated Available Inventory
- Calculated Order Quantity Requirements
- Calculated Delivery Dates

Relevant Data Section

The Relevant Data section provides links to supporting information:

- **Data visualizations:** Links to charts and graphs showing inventory trends, demand patterns, or other relevant metrics
- **Analytical resources:** References to vendor lead time analysis, inventory policy constraints, or other supporting data
- **Related metrics:** Connections to performance indicators that informed the insight

Select any link to drill into the underlying data that supports the root cause analysis.

Recommendations Section

The Recommendations section presents system-generated suggestions for resolving the insight.

Recommendation Cards: Each recommendation appears as a card containing:

Recommendation title: A clear description of the suggested action (for example, "Create purchase order for 3,205 units of ACDN-M15-RPA-0070")

Recommendation details: Specific parameters including:

- Quantities and units
- Locations (sites, warehouses)
- Dates and timelines
- Product or supplier identifiers

Timestamp: When the recommendation was generated

Details link: Select "Details" to view complete recommendation information including rationale and expected outcomes

Action status: Current state of the recommendation (for example, "Actions available")

To expand a recommendation:

1. Select the arrow icon next to the recommendation title
2. The card expands to show:
 - Complete action parameters
 - Supporting rationale and analysis
 - Expected outcomes and benefits
 - Related data and calculations

Activity Log

The Activity Log section tracks all actions and updates related to the insight.

Log Entry Structure: Each entry includes:

- **Action description:** What occurred (for example, "Stock Out Risk item for ACDN-M15-RPA-0070 and DC-SAN-001")
- **Insight ID reference:** Links to related insights
- **System information:** API calls or system processes involved
- **Timestamp:** When the action occurred
- **Details link:** Select to view additional information
- **Actions available button:** When applicable, shows available next steps

To view activity details:

1. Locate the Activity Log section at the bottom of the page
2. Select "Details" on any log entry to expand it
3. Review the detailed information about that activity

The Activity Log displays entries in reverse chronological order, with the most recent actions at the top.

Related Insights Section

At the bottom of the page, the Related Insights section displays other insights connected to the current issue.

Section Header: Shows the count of related insights (for example, "Related Insights (6)")

Related Insights Table displays:

Columns:

- **Insight ID:** Task identifier as a clickable link
- **Description:** Brief summary of the related insight
- **Status:** Current workflow state with icon (for example, "Dismissed")
- **Product:** Product identifier
- **Site:** Site identifier

To view a related insight:

1. Select any **Insight ID** in the Related Insights table
2. The system navigates to that insight's Details page

Related insights help you understand whether the current issue is part of a broader pattern or connected to other supply chain challenges.

Understanding Visual Indicators

Throughout the Insights Details page, visual elements help you quickly interpret information:

Status badges: Color-coded indicators show the current insight state

- **Dismissed:** Gray indicator for discarded insights
- **In progress:** Blue indicator for active work
- **Not started:** Default state for new insights

Severity indicators: Color and icon communicate priority

- **Critical:** Red indicator for highest priority

- **High:** Orange indicator for significant issues
- **Medium:** Yellow indicator for moderate concerns
- **Low:** Gray indicator for minor issues

Expandable sections: Arrow icons indicate collapsible sections with additional details

Clickable elements: Blue hyperlinks indicate interactive elements like Insight IDs and details links

Action buttons: Gray "Actions available" buttons indicate where you can take action on recommendations

Navigating the Insights Details Page

From the Insights Details page, you can:

Return to the Insights listing: Select "Insights" or "Filtered Insights" in the breadcrumb navigation

View related insights: Select any Insight ID in the Related Insights section to investigate connected issues

Access supporting data: Select links in the Relevant Data section to open related dashboards and analysis views

Take action on recommendations: Use the action buttons to accept or mark recommendations as complete

Page Updates

The Insights Details page updates automatically when:

- New data becomes available
- Recommendations are generated or updated
- Actions are taken on the insight
- Related insights change status

The "Last updated" timestamp at the top of the page shows when the most recent update occurred.

Acting on Recommendations

The Insights Details page allows you to review and act on system-generated recommendations designed to resolve supply chain issues. This page explains how to evaluate recommendations and execute the suggested actions. You can configure Amazon Connect Decisions to interact with an external system to perform certain operations on your behalf; when enabled, Amazon Connect Decisions will offer to perform the actions recommended through the Insights (e.g.: creating, updating, or canceling Purchase Orders).

Note

When you use the Actions feature, you authorize Amazon Connect Decisions to use Boomi (boomibot.com) as our integration middleware to transmit and process the data needed to take actions on your behalf, such as creating a purchase order after approval in our platform; no data is stored within Boomi.

Understanding Recommendations

Recommendations are AI-generated suggestions for addressing insights. Each recommendation includes:

- **Specific action:** What to do (for example, create purchase order, transfer inventory)
- **Parameters:** Quantities, locations, dates, and other details
- **Rationale:** Why this action is recommended
- **Expected outcome:** How this action will resolve the insight

Recommendations are prioritized based on impact, feasibility, and urgency to help you focus on the most effective solutions.

Reviewing Recommendations

Before taking action, review the recommendation details:

- **Read the recommendation title and description:** Understand what action is being suggested
- **Expand the recommendation card:** Select the arrow icon to view complete details
- **Review the rationale:** Understand why this action is recommended

- **Check the parameters:** Verify quantities, locations, dates, and other specifics
- **Assess the expected outcome:** Evaluate whether this action aligns with your business needs

Available Actions

For each recommendation, you have several action options:

Accept

Select **Accept** to approve the recommendation and proceed with the suggested action.

When you accept a recommendation:

- The insight status changes to "In Progress"
- The recommendation is marked as accepted
- Related insights may be updated if applicable
- The system tracks the action for future reference

To accept a recommendation:

1. Review the recommendation details
2. Select the **Accept** button
3. Confirm the action if prompted
4. The system updates the insight status and begins tracking implementation

Action Taken

Select **Action Taken** when you have addressed the issue through other means outside the system.

When you mark action taken:

- The insight status changes to "Pending Resolution"
- The system tracks that action was taken manually
- The insight will be marked as complete once data confirms resolution

To mark action taken:

1. Complete the necessary action in your external systems
2. Return to the Insights Details page
3. Select **Action Taken**
4. The system updates the insight status

Recommendation Lifecycle

Recommendations follow a lifecycle based on your actions:

Not Started

The recommendation has been generated but not yet acted upon

In Progress

You have accepted the recommendation or initiated action

Pending Resolution

Action has been taken but the issue is not yet fully resolved in the data

Completed

The system has confirmed through data that the action successfully resolved the insight

Multiple Recommendations

Some insights may have multiple recommendations:

Primary Recommendation: The system's top suggestion based on analysis

Alternative Recommendations: Additional options if the primary recommendation is not feasible

To choose between recommendations:

1. Review all available recommendations
2. Compare the expected outcomes and feasibility
3. Select the recommendation that best fits your business needs and constraints
4. Take action on your chosen recommendation

You can only accept one recommendation per insight. Once you accept a recommendation, other recommendations for that insight become unavailable.

Grouped Recommendations

When multiple insights share the same root cause, the system may group recommendations:

Grouped recommendation benefits:

- Resolve multiple insights with a single action
- More efficient insight management
- Clearer view of systemic issues

When viewing a grouped recommendation:

- The recommendation shows all affected insights
- Taking action on the grouped recommendation affects all related insights
- All related insights move to "In Progress" when you accept the grouped recommendation

Tracking Action Progress

After taking action on a recommendation:

- **Monitor the Activity Log:** Check the Activity Log section for updates on your action
- **Review related insights:** See if related insights are also being resolved
- **Check the insight status:** The status indicator shows current progress
- **Wait for data confirmation:** The system monitors data for changes that indicate resolution

Once the system confirms through data that the action was successful, the insight status automatically updates to "Completed."

Best Practices

Review thoroughly before acting: Ensure you understand the recommendation and its implications

Consider business constraints: Verify that the recommended action aligns with your operational capabilities and business rules

Act promptly on critical insights: Critical insights with near-term impact dates require immediate attention

Provide feedback: Use thumbs up/down icons to help the system improve future recommendations

Monitor outcomes: Track whether accepted recommendations successfully resolve insights

Providing Feedback on Insights

Amazon Connect Decisions enables you to provide feedback on any insight — including Root Cause Analysis (RCA) findings and Recommendations generated for insights — directly through the Natural Language Interface (NLI). Your feedback helps Amazon Connect Decisions agents learn and improve the accuracy of future analyses and recommendations. There are two ways to provide feedback: **Natural Language feedback in the NLI** and **Thumbs Up/Down ratings** to the agent's response in the NLI.

Every piece of feedback — whether a quick thumbs up/down or a detailed natural language correction — is evaluated by Amazon Connect Decisions agents as a potential learning. Qualifying feedback that is deterministic, non-conflicting, and generalizable is curated into actionable learnings stored in the Knowledge Store. Over time, your feedback improves RCA accuracy, recommendation relevance, and reduces the need for repeated corrections — making Amazon Connect Decisions progressively smarter with every interaction.

Natural Language Feedback

You can also provide detailed, context-specific feedback by simply typing your correction or observation in the Decisions Teammate / NLI. This is ideal when you want to explain *why* an output is incorrect or share operational knowledge that the agent should learn.

For example, you review a projected stockout alert / insight for electronic connectors at your Frankfurt site. The root cause agent attributes the root cause to a demand forecast error. You respond via NLI:

"The root cause isn't a forecast error — Supplier TP_00001 has a contractual max order quantity of 291 units per PO, and the system generated a single PO for 1,600 units. The supplier will reject this. Split the PO into multiple orders of 291 units or fewer."

Decisions Teammate responds: *"Thank you for the correction. I've updated the root cause to reflect the supplier's max order quantity constraint. The recommendation has been regenerated as 6 split*

POs (5 x 291 + 1 x 145 units). This supplier constraint will be applied to all future recommendations for TP_00001."

Thumbs Up / Thumbs Down Ratings

Every RCA finding and Recommendation displayed in the NLI includes a thumbs up and thumbs down option. Use these to quickly signal whether the agent's output was accurate and helpful. You can add a comment to provide additional context.

For example, you receive a projected stockout alert / insight for running shoes at your Northeast distribution center. The RCA Agent identifies the root cause: *"Supplier TP_00012 shipment delayed by 8 days due to port congestion at Long Beach, combined with a 22% demand spike driven by the upcoming Boston Marathon."* The Recommendation Agent suggests: *"Reallocate 1,200 units from 6 low-demand Southeast stores to 9 high-demand Northeast stores."*

You give a **thumbs up** and add: *"Spot on — the port delay and marathon demand spike are exactly what's driving this. Good recommendation."*

Decisions Teammate / NLI acknowledges: *"Thank you for confirming. Your feedback strengthens the RCA Agent's confidence in correlating port disruption signals with regional demand events for future stockout analyses."*

Plans

Planning with Amazon Connect Decisions

Amazon Connect Decisions transforms supply chain planning through AI teammates that work alongside your team in a continuous learning loop. These AI teammates observe patterns across your operations, detect what matters most, recommend actions, and execute approved decisions—institutionalizing expertise so knowledge stays even when planners leave and new team members become productive from day one.

At the core of planning in Amazon Connect Decisions are three interconnected intelligence capabilities.

Demand Intelligence dynamically orchestrates 18+ forecasting tools and foundation models trained on Amazon data, generating accurate forecasts from day one even with limited history while harmonizing statistical forecasts, customer commitments, and cross-functional inputs through consensus planning.

Supply Intelligence generates forward-looking supply plans, intelligently clusters exceptions into strategic decisions ranked by impact, and runs optimization models for constraint-aware planning across your network.

Adaptive Intelligence ties it all together—continuously learning from every plan and decision, monitoring operations 24/7, surfacing what matters, and executing approved decisions directly in existing systems to reduce planning cycles from weeks to hours.

Demand Planning

A plan is created for a defined time horizon and refreshed periodically through rolling time windows (planning cycles). Each demand plan can contain multiple versions within a planning cycle to support refinement of plan for any incremental data received within a planning cycle.

Prerequisites

Before creating your first demand plan in Amazon Connect Decisions, ensure that you have the following prerequisites in place:

- Your Amazon Connect Decisions instance must be set up and configured.
- Your user account must be assigned a Manager role.
- Product or product-site data entities must be prepared and uploaded (see Data Entities for details).

What is Demand Intelligence?

Amazon Connect Decisions' Demand Intelligence offers customers an automated, accurate forecast powered by AI/ML. Using Demand Intelligence, you can improve the accuracy of the baseline forecast, automate consensus planning, proactively detect and manage exceptions, improve planner productivity and most importantly drive down excess inventory and out of stock scenarios.

Data Entities for Demand Planning

The following table lists the data entities and columns used by Demand Planning.

How to read the table:

- **Required** – The columns in this data entity are mandatory to execute a demand forecast without any failures.
- **Conditionally required** – The columns in this data entity are required depending on the configurations set under demand plan settings.
- **Recommended for forecast quality** – The columns in this data entity are required for the quality for the forecast.
- **Optional** – The column name is optional. For enhanced feature output, it is recommended to add the column name with values.

outbound_order_line (required)

How is this data entity used? Demand Planning uses this data as the primary source of historical demand for forecast. Additionally, fields selected as granularity are sent for training and are available as filters to review the demand plan.

outbound_order_line columns

Column	Is the column required?	How is this column used in Forecasting?
id	Required	<i>id</i> , <i>cust_order_id</i> , and <i>product_id</i> are used to uniquely identify a record in the data entity and this combination should always be unique. Make sure the column values do not have invalid characters such as asterisk and double-quotes.
cust_order_id	Required	
product_id	Required	
order_date	Required	Required for forecast creation. Identifies the period for time-series forecasting.
final_quantity_requested	Required	Required for forecast creation. Identifies the quantity used for time-series forecasting. This column must not contain null values and must be <i>numerical</i> . Make sure there are no commas in the values. For example, 500000.00 is an accepted value in Demand Planning.
ship_from_site_id	Conditionally required	This column is conditionally required for forecast creation <i>if</i> the column is selected for forecast dimension (Site Hierarchy). This column must have a value and is used for filtering and analysis of data.
ship_to_site_id	Conditionally required	
channel_id	Conditionally required	This column is conditionally required for forecast creation <i>if</i> the column is selected for forecast dimension (Channel Hierarchy). This column must have a value and is used for filtering and analysis of data.
customer_tpartner_id	Conditionally required	This column is conditionally required for forecast creation <i>if</i> the column is selected for forecast dimension (Customer Hierarchy). This column must have a value and is used for filtering and analysis of data.
ship_to_site_address_city	Conditionally required	This column is conditionally required for forecast creation <i>if</i> the column is selected for forecast dimension

Column	Is the column required?	How is this column used in Forecasting?
ship_to_site_address_state	Conditionally required	(Site Hierarchy). This column must have a value and is used for filtering and analysis of data.
ship_to_site_address_country	Conditionally required	
status	Recommended for forecast quality	This column is recommended for forecast quality. Orders with <i>canceled</i> status are not considered as forecast input.

product (required)

How is this data entity used?

Demand Planning uses the product attributes to establish hierarchy filters for demand plan review and for model training.

product columns

Column	Is the column required?	How is this column used in Forecasting?
id	Required	Required for data ingestion into Supply Chain Data Lake (SCDL). Make sure the column values do not have duplicate IDs and special characters such as asterix and double-quotes.
description	Required	Required for data ingestion into Supply Chain Data Lake (SCDL). This column can contain special characters such as asterix, hyphen, quotes, and double-quotes.
parent_product_id	Conditionally required	This column is conditionally required for forecast creation <i>if</i> the column is selected for forecast dimensions (Product Hierarchy). Make sure the column has values

Column	Is the column required?	How is this column used in Forecasting?
product_group_id	Conditionally required	and is used for filtering and analysis of data and model training.
product_type	Conditionally required	
brand_name	Conditionally required	
color	Conditionally required	
display_desc	Conditionally required	
product_available_day	Recommended for forecast quality	Recommended. The value in this column improves forecast quality by allowing the forecasting model to consider the timing of new product introductions.
discontinue_day	Recommended for forecast quality	Recommended. The value in this column improves forecast quality by allowing the forecasting model to consider the timing for product retirements.
base_uom	Recommended for forecast quality	Unit of measure for product. Default is Eaches. Currently demand planning only supports Eaches.
is_deleted	Recommended for forecast quality	Recommended. Enter Y if the product ID should be excluded from forecasting.
pkg_height	Recommended for forecast quality	Recommended. The physical characteristics of the product that the forecasting models can understand.

Column	Is the column required?	How is this column used in Forecasting?
pkg_length	Recommended for forecast quality	
pkg_width	Recommended for forecast quality	
shipping_dimension	Recommended for forecast quality	
casepack_size	Recommended for forecast quality	

product_alternate (Recommended for forecast quality)

How is this data entity used?

Demand Planning uses the data of product's predecessor(s) or alternate(s) to create forecast for new products. When data is ingested into the *product_alternate* data entity, Product lineage support for forecast is enabled. You can skip ingesting data into the *product_alternate* data entity and the forecast can still be generated.

product_alternate columns

Column	Is the column required?	How is this column used in Forecasting?
alternative_product_id	Required	Required for data ingestion into Supply Chain Data Lake (SCDL). Unique record identifier.
product_id	Required	Required for data ingestion into Supply Chain Data Lake (SCDL). ID of the new product or new version of

Column	Is the column required?	How is this column used in Forecasting?
		the product. Make sure <i>product_id</i> is populated in the <i>product</i> data entity.
<i>product_alternate_id</i>	Required	Required for data ingestion into SCDL. Identifier for a similar product or previous version of the product. To consider multiple similar products as a single <i>product_id</i> , enter the products in separate rows. Make sure <i>product_alternate_id</i> is populated in the <i>product</i> data entity.
<i>alternate_type</i>	Required	Required for applying product supercession or lineage. Use the static value <i>similar_demand_product</i> in all the rows.
<i>alternate_product_qty</i>	Required	Required for applying product supercession or lineage. Enter the proportion of history of the <i>alternate_product_id</i> you want to use for forecasting <i>product_id</i> . For example, if it is 60%, enter 60. When you have multiple <i>alternative_product_id</i> for a single <i>product_id</i> , the <i>alternate_product_qty</i> does not have to add up to 100.
<i>alternate_product_qty_uom</i>	Required	Required for applying product supercession or lineage. Use the specific static value "percentage".
<i>eff_start_date</i>	Required	Required for data ingestion into SCDL. Enter the start timeframe to consider the history of a similar product. Make sure this date is on or before the <i>eff_end_date</i> or you can leave this field empty and Demand Planning will auto-fill the year with 1000.
<i>eff_end_date</i>	Required	Required for data ingestion into SCDL. Enter the end timeframe to consider in history of a similar product. Make sure this date is on or after the <i>eff_start_date</i> .

Column	Is the column required?	How is this column used in Forecasting?
status	Recommended for forecast quality	Recommended. Enter <i>Inactive</i> to ignore the product supersession or lineage mapping.

supplementary_time_series (recommended for forecast quality)

How is this data entity used? Demand Planning uses this data as the primary source for tagging casual factors such as promotional events, discounts, holidays, and so on.

supplementary_time_series columns

Column	Is the column required?	How is this column used in Forecasting?
id	Required	Required for data ingestion into Supply Chain Data Lake (SCDL). Unique record identifier.
order_date	Required	Required for data ingestion into Supply Chain Data Lake (SCDL). Timestamp when the timeseries was recorded.
time_series_name	Required	Required for data ingestion into Supply Chain Data Lake (SCDL). Name of the specific type of time series. The <i>time_series_name</i> column must start with a letter, be 2 to 56 characters long, and can contain letters, numbers, and underscores. No other special characters are allowed.
time_series_value	Required	Required for data ingestion into SCDL. Value corresponding to the specific time series. Demand Planning only supports numerical input and time-series with categorical value is not considered.
product_id	Optional	Recommended. Unique identifier for a specific product. Use this column if the demand driver is available at product level.

Column	Is the column required?	How is this column used in Forecasting?
site_id	Optional	Recommended. Unique identifier for a specific site or location. Use this column if the demand driver is available at site level. This column can represent either <i>ship_from_site_id</i> or <i>ship_to_site_id</i> based on the lowest level site hierarchy configuration.
channel_id	Optional	Recommended. Unique identifier for a specific channel. Use this column if the demand driver is available at channel level.
customer_tpartner_id	Optional	Recommended. Unique identifier for a specific customer. Use this column if the demand driver is available at customer level.

Historical vs. Future Supplementary Time Series: Understanding Covariates in Forecasting

Accurate demand forecasting requires understanding not just historical sales patterns, but the external factors that drive demand changes. Supplementary Time Series (STS) data—also called covariates—captures these demand drivers like promotions, pricing, holidays, and inventory levels, enabling forecasting models to distinguish explainable patterns from random noise and predict how future business actions will impact demand. However, a critical distinction exists between covariates that are only known historically (like past inventory levels or competitor actions) versus those known in advance (like planned promotions or scheduled holidays), and understanding this difference is essential for building accurate forecasts that support proactive planning decisions.

A critical distinction in demand forecasting is between **past covariates** and **known covariates** (also called future covariates). Understanding this difference is essential for building accurate forecasting models.

Past Covariates (Historical STS Data)

Past covariates are supplementary time series values that are only known for historical periods. These variables are observed alongside your historical demand but cannot be predicted or known in advance for future periods.

Examples of Past Covariates:

- **Historical Inventory Availability:** You know what inventory levels were in the past, but future availability depends on demand, replenishment, and other uncertain factors
- **Actual Competitor Pricing:** Historical competitor price data is observable, but future competitor actions are unknown
- **Weather Conditions:** Past weather is recorded, but future weather (beyond short term forecasts) is uncertain
- **Website Traffic:** Historical traffic patterns are known, but future traffic depends on many unpredictable factors

Use in Forecasting Models: Past covariates help the model learn historical relationships and patterns. For example, if high inventory availability historically correlated with higher sales (due to better product visibility or fulfillment speed), the model learns this relationship. However, since these values are unknown for future periods, the model must forecast without them or make assumptions about their future values.

Known Covariates (Future STS Data)

Known covariates are supplementary time series values that are known or can be determined in advance for future periods. These are the most valuable inputs for forecasting because they provide concrete information about future conditions.

Examples of Known Covariates:

- **Planned Promotional Discounts:** Your marketing team has already scheduled promotional campaigns with specific discount levels for future dates
- **Price Index Changes:** Planned price adjustments are determined in advance based on your pricing strategy
- **Holiday Indicators:** Calendar based events (holidays, shopping seasons, fiscal periods) are known years in advance
- **Planned Marketing Spend:** Budget allocations and campaign schedules are predetermined
- **Store Opening/Closing Events:** Expansion or consolidation plans are known ahead of time

Use in Forecasting Models: Known covariates dramatically improve forecast accuracy because the model can incorporate actual future conditions rather than assumptions. For example, if you know

a 25% discount promotion is planned for next month, the model can predict the expected demand lift based on historical discount response patterns.

Practical Implementation Strategy

For Historical Periods (Training Data): Include both past covariates and known covariates in your supplementary time series data. This allows the model to learn relationships from all available demand drivers. Your dataset should contain actual observed values for all time series types up to the present date.

For Future Periods (Forecasting Horizon): Only include known covariates in your supplementary time series data. These are the demand drivers you can confidently specify for future dates. For example:

```
id,order_date,time_series_name,time_series_value,product_id,site_id,channel_id,customer_id,partner_id
1001,2025-02-01,discount_percentage,20.0,PROD_001,SITE_NYC,CHANNEL_ONLINE,CUST_12345
1002,2025-02-14,discount_percentage,30.0,PROD_001,SITE_NYC,CHANNEL_ONLINE,CUST_12345
1003,2025-02-01,holiday_indicator,0,PROD_001,SITE_NYC,CHANNEL_ONLINE,CUST_12345
1004,2025-02-14,holiday_indicator,1,PROD_001,SITE_NYC,CHANNEL_ONLINE,CUST_12345
```

This future data tells the model that a 20% discount is planned for February 1st and a 30% Valentine's Day promotion is scheduled for February 14th.

Practical Applications

- **Promotional Planning:** Track discount percentages over time to understand how promotional intensity affects demand. This helps identify optimal discount levels and predict the demand lift from future promotions.
- **Price Elasticity Analysis:** Monitor price index movements to quantify how price changes influence customer purchasing behavior across different products, locations, and channels.
- **Inventory Constraint Modeling:** Capture inventory availability levels to identify when stockouts or low inventory constrained sales, ensuring forecasts account for supply limitations rather than true demand signals.

Benefits for Demand Planning

By incorporating supplementary time series data, your Demand Planning system can:

- **Improve Forecast Accuracy:** Account for known demand drivers rather than treating them as unexplained variance
- **Enable Scenario Planning:** Model "what if" scenarios by adjusting future values of demand drivers
- **Identify Causal Relationships:** Understand which factors most significantly impact demand for different products and markets
- **Support Strategic Decisions:** Provide data driven insights for pricing, promotional, and inventory strategies

Creating and Configuring a Demand Plan

To create a demand plan:

1. Navigate to the Plans section in Amazon Connect Decisions and click **Create Plan**. Once a plan is created, you have the flexibility to edit it at any time to reflect changing business needs or incorporate new information.
2. **Configure Time Horizon settings:**
 - **Time bucket:** Select Daily, Weekly, or Monthly based on your planning needs.
 - **Plan horizon:** Specify the forecast horizon:
 - Daily: 1 to 28 days
 - Weekly: 1 to 26 weeks
 - Monthly: 1 to 12 months
 - **Forecast start date:** Must be on Mondays for weekly plans or the 1st day of the month for monthly plans. If a non-compliant date is selected, the system automatically adjusts to the start of the same time bucket. For weekly plans, the date will shift to the Monday of the same week. For monthly plans, the date will shift to the 1st day of the following month.
3. **Select Planning Grain:**

Choose the level of detail for your forecasts:

 - **Product:** Forecast by product across all locations
 - **Site:** Forecast by location across all products
 - **Ship from site ID:** Forecast by specific shipping location
 - **None:** Configures the plan at the product level only
4. **Set Prediction Lead Time:**

This parameter determines the input to forecast accuracy calculations. The prediction lead time is the number of periods ahead for which forecasts are compared against actual demand.

5. Configure Consensus Plan Rules (optional):

Consensus rules combine multiple forecast inputs into a unified forecast. You can create rules that specify how different forecast types should be combined.

Example rules:

- "For product category = Automotive, use average of sales forecasts, marketing forecasts, customer forecasts, and baseline forecasts"
- "If customer forecast is zero, then the forecast for that item must be zero"
- "For product categories Professional Power Tools, always use customer forecasts"

Requirements for consensus rules:

- The referenced data must exist in your Amazon Connect Decisions Data Lake (Entity Name = Forecast) for all forecast types. For example, if you specify a rule as "Use maximum of sales forecast and baseline forecast", both the "sales forecast" and "baseline forecast" must be present in the SCDL Forecast entity with their respective `plan_type` values (`plan_type` = "Sales Forecast" and `plan_type` = "Baseline Forecast").
- Rules should be arranged in the order they need to be executed in.
- Forecast data format must match your plan configuration granularity.

Reviewing and Analyzing Plans

To review a demand plan:

1. Navigate to Plans from your AWS Connect Decisions homepage.
2. Select your active plan from the available options.
3. The system displays a high-level view of forecasts aggregated across products and sites.
4. Use the filtering options to focus on specific products, locations, or time periods.
5. Review different forecast types:
 - **Baseline forecast:** AI-generated forecast based on historical patterns
 - **Consensus plan:** Combined forecast incorporating multiple inputs (if configured)
 - **Forecast inputs:** Individual inputs from sales, customers, and marketing

Understanding Quantile Selection

The system uses statistical quantiles to represent forecast uncertainty:

- **Plans without consensus rules:** The system automatically selects the optimal quantile based on your planning grain.
- **Consensus plans:** The system defaults all products to the P50 quantile.

To see which quantile is used for the product or product-site, filter to the finest grain level view to identify the optimal quantile.

Editing Forecasts

You can adjust forecasts based on your business knowledge and market insights.

To edit a forecast:

1. Open the chat panel in Plans.
2. Specify your adjustment by clearly indicating:
 - Time periods to adjust
 - Products and sites affected
 - New forecast values (as absolute numbers or percentage changes)
 - Reason for your adjustment

Example: "Increase the forecast for Product A at Site B by 15% for January 2026 due to planned promotion."

Monitor the status banner:

- **Blue banner:** The system is processing your change. The publish button is disabled during this time.
- **Green banner:** Your change has been applied successfully. Updated forecasts appear highlighted in light blue.

Providing feedback:

You can also provide feedback to improve future forecasts:

- Suggestions on how the system should handle missing data
- Which accuracy metrics are most important for your business
- Recommendations for new consensus rules or adjustments to existing rules

Concurrent Plan Editing

Since multiple planners can work on the same plan simultaneously, coordination is essential to avoid conflicts:

- Establish team protocols where each planner is responsible for specific product segments.
- Watch for the blue "Plan is being updated" banner, which appears across all users when any planner is making changes.
- Wait for the green "Forecast is updated" banner before starting your own edits.
- Coordinate editing schedules with your team to avoid conflicts.
- Establish a team communication practice where planners announce significant changes in a shared channel.

Monitoring Forecast Performance

Amazon Connect Decisions tracks five standard statistical metrics to measure forecast accuracy for the user specified prediction lead time in the plan configuration:

- **MAPE (Mean Absolute Percentage Error):** $(1/n) \times \sum |Actual - Forecast| / |Actual| \times 100\%$
 - Measures the average percentage error across all forecasts.

Note

Forecast references the forecast that was generated at your specified prediction lead time before the actual demand period occurred. Based on your plan configuration, the system measures accuracy by comparing the forecast created a specific number of time buckets in advance to the actual demand—for example, if your time bucket is weeks and your prediction lead time is 3 weeks, the "Forecast" value in accuracy formulas (MAPE, WAPE, RMSE, Bias, and MAE) references the forecast that was created 3 weeks prior to each demand period.

- **WAPE (Weighted Absolute Percentage Error):** $\sum |Actual - Forecast| / \sum |Actual| \times 100\%$

- Gives more weight to high-volume products.
- **RMSE (Root Mean Square Error):** $\sqrt{[(1/n) \times \sum (\text{Actual} - \text{Forecast})^2]}$
 - Emphasizes larger forecast errors.
- **Bias (Forecast Bias):** $\sum (\text{Forecast} - \text{Actual}) / \sum |\text{Actual}| \times 100\%$
 - Indicates whether forecasts systematically over-predict or under-predict.
- **MAE (Mean Absolute Error):** $(1/n) \times \sum |\text{Actual} - \text{Forecast}|$
 - Provides the average error in the same units as your demand.

Publishing Demand Plans

After planners have reviewed and refined forecasts, you can publish the final plan to make it available for downstream processes.

To publish a demand plan:

1. Review the aggregate performance metrics to ensure forecast quality meets your standards.
2. Coordinate with your planning team to confirm all necessary adjustments have been completed.
3. Click **Publish Plan** to process your request.
4. The system automatically saves the published plan to your configured Amazon S3 folder.

Rerun Demand Plans for the Next Planning Cycle

After the plan is published, you can rerun the plan to produce a new forecast that uses the most up-to-date historical demand data as well as any additional input updates.

To rerun a demand plan:

1. Navigate to your active demand plan in the Plans section.
2. Choose **Rerun Plan**.
3. The system redirects you to the Plan Configuration page with the forecast start date field automatically cleared.
4. Review and modify configuration settings as needed:
 - Manually select a new forecast start date.
 - Adjust time horizon settings if required.

- Update planning grain or prediction lead time if necessary.
- Modify consensus plan rules (if applicable).

5. Choose **Confirm & Generate Plan** to trigger the forecast generation.

Understanding Your Forecasts

AWS Connect Decisions provides forecast explainability to help you understand why the system generated particular forecasts.

To understand why a forecast was generated:

1. Navigate to the chat interface (right panel) in Plans. This is an AI-powered Demand Planning Teammate that can explain how your baseline forecast was calculated and what factors influenced it. You can also use it to edit forecasts or provide feedback to improve future predictions.
2. Ask a specific question: "Why is this final forecast X?"
3. Include the following information in your question:
 - Product name
 - Site name (if applicable)
 - Time range

The system will analyze the forecast and explain key influencing factors:

- Historical trends and patterns
- Seasonality effects
- Pattern changes and anomalies
- Product lifecycle stages

Use this explanation to validate whether the forecast makes sense based on your business knowledge.

Plan Improvement Actions

Amazon Connect Decisions proactively identifies opportunities to improve your demand forecast accuracy by detecting patterns and anomalies that require your business context. Plan

improvement actions are AI-generated tasks that surface where your input can most meaningfully enhance forecast quality, enabling your AI teammates to learn from your expertise and produce more accurate plans over time.

The system uses ML-driven detection to analyze your demand data, identify products and sites where forecasts would benefit from your input, and presents actions grouped by type. You provide input through the Supply Chain Decisions Teammate, and your feedback is automatically integrated into future forecast runs, creating a continuous improvement loop between you and your AI teammates.

How Plan improvement actions work

Plan improvement actions follow a three-step workflow:

1. **Detect** - The system analyzes your demand data using ML-driven detection and clustering. It identifies products and product-sites that need your input based on pre-determined scenarios including new products with limited history, products with extremely low forecastability, intermittent demand gaps, anomalous demand patterns, known global events, and high-variability products.
2. **Surface** - Detected actions are converted into digestible, pattern-level insights that you can confirm, correct, or contextualize. Actions are presented in the Demand Plan, grouped by action type.
3. **Act** - You apply validated corrections by providing input through the Supply Chain Decisions Teammate. Your inputs are integrated into planning configuration updates, data treatment rules, related time series inputs, and forecast override rules. The system then uses your input to refine models and improve future forecast accuracy.

Note

Plan improvement actions are refreshed every time a plan is run, ensuring that as incremental data is ingested into the system (such as additional products/sites added to planning or recent demand history uploads), you continue to receive updated actions on products and patterns needing your input.

Viewing Plan improvement actions

When you log into Amazon Connect Decisions, you can access Plan improvement actions from the Plans listing page.

To view Plan improvement actions:

1. Navigate to the Plans page from the landing page (Login→Home).
2. In the Plans listing page, locate the demand plan you want to review. Plan prioritization is based on pending Plan improvement actions as one of the factors.
3. Select the plan to open the plan details view.
4. Choose the Plan Improvement Actions tab.

The tab displays a list of actions grouped by action type alongside your forecast outputs. Each action includes:

- **Input Topics** - The category of input needed (static)
- **Description** - Dynamic description of the specific pattern or issue detected, including product, site, and product-site details
- **Status** - Current progress (Open or Closed)
- **Actions** - Available operations (View, Complete)

Understanding action types

Plan improvement actions cover the most common scenarios where planner expertise improves forecast accuracy:

#	Action Type	Description	Example
1	Provide input for new products/sites	Products launched recently with zero or limited history and no product lineage data available.	25 products launched in the last 2 months have zero or limited history and no product lineage data available .
2	Provide input on products with	Products with average demand per period since	30 Slow Sellers/Long Tail Products/Sporadic sellers

#	Action Type	Description	Example
	extremely low forecastability	launch/first order in history < 0.5, such as slow sellers and long tail products.	with average demand per period since launch/first order in history < 0.5.
3	Provide input on products with continuous periods of zero sales	Active products/product-sites with zero orders in the last year that may need discontinuation or scope adjustment.	100 Active product/product-site with zero orders in last 1 year.
4	Provide reasoning for intermittent gaps in historical demand	Random/non-repeating demand spikes and demand drops in history with no correlated causals (related time series).	Unexplained, intermittent periods of no sales (e.g., random/non-repeating demand spikes or demand drop in history with no correlated causals).
5	Validate "known" global events	One-time anomalies that coincide with known global disruptions or major differences in trend levels during specific periods.	One-time anomalies that coincide with known global disruptions/Major Difference in trend levels for 100 products during 2020-2021 periods coinciding with COVID.
6	Provide causals/events for high-variability products	Products exhibiting high variability in historical demand and "flat" forecast outputs.	26 products exhibit High Variability in historical demand and "flat" forecast outputs.

Completing a Plan improvement action

You provide input through the Supply Chain Decisions Teammate, which allows you to share business context for multiple products through natural language conversation.

Using the Supply Chain Decisions Teammate

1. From the Plan Improvement Actions tab, choose View on the action you want to address.
2. The Supply Chain Decisions Teammate panel opens on the right side of the screen, contextual to the selected action topic.
3. The agent will ask clear and specific questions about the products in scope. For example:
 - "What is the reason for the one-time 4x spike in demand for Product X in Feb 2024?"
 - "Do you expect this event to recur?"
4. Provide your response in natural language. The agent will:
 - Clarify through follow-up questions if needed (e.g., if your previous input contradicts new input)
 - Intelligently manage and prevent redundant questions (enabled by smart memory)
 - Present hypotheses for you to validate (e.g., "Should I treat this as promo uplift or true spike?")
5. When you have finished providing input for all product-sites in scope, the action status changes to Closed.

Note

You can also initiate a conversation with the agent to provide input proactively, even without navigating from a specific action. For example, you can ask "Why do I see a zero forecast for the entire forecast horizon for product XX?"

Note

If 10 products have similar input, the system groups them together as Agent groups to help simplify and ease planner actions.

Tracking progress

Plan improvement actions have the following status lifecycle:

Status	Description
Open	The action has been detected and surfaced but no planner input has been provided yet.
Closed	The planner has provided input through the Supply Chain Decisions Teammate, covering all product-sites in scope.

Tracking overall progress:

- The system continuously tracks closed versus open actions after every plan run.
- The action list refreshes based on the latest data. As incremental data is ingested into the system during a new planning cycle (e.g., additional products/sites added to planning or recent demand history uploads), you continue to receive updated/refreshed actions.

How your input improves forecasts

Your input from Plan improvement actions is seamlessly integrated into three key areas:

1. **Refinement of demand planning rules** - For example, your input on the reason for long periods of missing demand results in the product being removed from planning scope (product discontinuation).
2. **Enriching forecast inputs with causal data** - Enriching forecast inputs with causal data/Related Time Series (RTS) constructed based on your provided business context in natural language (e.g., automatically constructing a related time series for a recurring promotion).
3. **Capturing business rules** - Helps capture business rules unique to a select pattern of products (e.g., exclude from forecasting products with less than 3 orders).

Your inputs can result in different types of plan configuration updates, including:

- **Upload Product Lineage or Site Lineage Data** - Ingest Product Alternate, Site Alternate, or Product-Site Alternate data
- **Update Plan Configuration - Planning Rules** - Update forecast override rules (e.g., use sales forecast or marketing forecast as consensus plan input for 3 months from launch date)

- **Update Plan Configuration - Plan Scope** - Exclude from forecasting, forecast using last n period average
- **Construct Causal Data as Related Time Series** - Construct RTS for anomalies, COVID periods, or causal data (e.g., Price, Promotion, or Marketing Data)

Best practices

- **Provide context, not just answers** - When using the Supply Chain Decisions Teammate, explain the business reasoning behind your input. This helps the agent learn generalizable patterns.
- **Review actions each planning cycle** - New actions are generated with each plan run as new data is ingested. Regular review ensures your forecasts continuously improve.

Other Capabilities from Supply Chain Decisions Teammate

The Supply Chain Decisions Teammate, located at the chat interface (right panel) in Plans, goes beyond basic forecast explanations and edits. It provides advanced analytical capabilities to help you understand your demand patterns, identify issues, and make data-driven decisions.

Retrieving Top Products by Sales

- You can ask the Supply Chain Decisions Teammate to identify your highest-performing products.
- Example query: "Show me the top 10 products by sales volume for Q4 2025"
- The teammate will analyze your historical sales data and present the products ranked by sales performance, helping you focus your planning efforts on the most critical items.

Analyzing Forecast Accuracy Across Important SKUs

- The teammate can perform detailed forecast accuracy analysis for your key products.
- Example query: "Analyze forecast accuracy for my top 20 SKUs over the past 3 months"
- The system will evaluate multiple accuracy metrics (MAPE, WAPE, RMSE, Bias, and MAE) and provide insights into which products have reliable forecasts and which require attention.

Recommended Next Steps for Improving Forecasts

When the teammate identifies forecast quality issues, it can recommend specific actions, such as adding contextual data.

- If forecasts show high error, the teammate may recommend:
 - Supplementary time series data: Add promotional calendars, pricing changes, or holiday indicators
 - Product lifecycle information: Include product launch dates or discontinuation dates
 - Demand drivers: Incorporate marketing spend, inventory availability, or competitor actions
- Example query: "The forecast for Product ABC has high error. What contextual data should I add to improve it?"

Pattern Analysis Across Multiple Products

- The Supply Chain Decisions Teammate can identify systemic issues affecting multiple products.
- Example query: "Are there any common patterns causing forecast errors across my winter apparel category?"
- The teammate will analyze forecast performance across product groups and identify:
 - Seasonal patterns that the model may be missing
 - Common demand drivers affecting multiple products
 - Systematic biases (over-forecasting or under-forecasting trends)
 - Data quality issues impacting specific product categories

Creating Visualizations for Forecast Scenarios

- The teammate can generate graphs and charts to help you visualize potential forecast scenarios:
 - Baseline vs. adjusted forecasts: Compare the AI-generated forecast with your proposed changes
 - Historical actuals vs. forecasts: Visualize past forecast accuracy to build confidence
 - What-if scenarios: Show how different assumptions (promotions, price changes, inventory levels) would impact future demand
- Example query: "Create a graph showing how the forecast for Product GHI would look with a 20% promotional discount in March 2026"

- The teammate will generate visualizations that help you:
 - Communicate forecast assumptions to stakeholders
 - Evaluate the impact of business decisions on demand
 - Compare multiple planning scenarios side-by-side
 - Identify trends and anomalies more easily

Continuous Improvement Through Feedback

- You can provide feedback to help the teammate improve future forecasts.
- Example feedback:
 - "The model should give more weight to recent trends for fast-fashion products"
 - "Holiday seasonality patterns have changed; please adjust accordingly"
 - "When inventory is low, historical sales understate true demand"
- The teammate uses this feedback to refine its recommendations and help you build more accurate forecasts over time.

Supply Planning

Amazon Connect Decisions generates constraint-based supply plans that translate your demand forecasts into feasible, actionable production and procurement schedules. Supply plans respect real-world operational constraints like material availability, lead times, production capacity, and warehouse space limits so you can move from planning to execution with confidence.

What You'll Find in This Section

The following pages guide you through the complete supply planning workflow in Amazon Connect Decisions:

- **Creating a New Plan:** Initiate a new supply plan, define its scope across products and sites, and kick off the planning run
- **Configuring Plan Parameters:** Set the parameters that shape how Amazon Connect Decisions generates your supply plan, including planning horizon, time buckets, etc.
- **Reviewing and Editing Plans:** Explore the generated plan output and make targeted adjustments where business context requires refinement

- **Publishing Supply Plans:** Finalize and publish your approved supply plan for downstream execution

Data Entities for Supply Planning

The following table lists the data entities and columns used by Supply Planning.

How to read the table:

- **Required** – The entity or columns in the data entity are mandatory for the supply planning engine to function correctly.
- **Conditionally required** – The entity or columns in the data entity are required depending on the specific major functional requirements, policy type, order type, or sourcing rule type selected.
- **Optional** – The entity or columns in the data entity is optional. For enhanced planning accuracy and feature output, it is recommended to add the column with values.

site (required)

How is this data entity used? Supply Planning uses the site data to define the physical locations (plants, warehouses, and distribution centers) within the supply chain network. Sites serve as the origin and destination nodes for transportation lanes, sourcing rules, and inventory policies.

Column	Is the column required?	How is this column used in Supply Planning?
id	Required	Unique identifier for the physical site (plant, warehouse , or distribution center) within the supply chain network. Make sure the column values do not have invalid characters such as asterisk and double-quotes.
description	Optional	Human-readable name or label for the site.
geo_id	Optional	Foreign key linking this site to its geographic location in the Geography table.

Column	Is the column required?	How is this column used in Supply Planning?
is_active	Optional	Boolean flag indicating whether the site is currently operational and available for planning. By default, the site is considered as Active.
open_date	Optional	Calendar date on which the site became or will become operationally active. Used by the planning engine to exclude future sites from current plans while incorporating them into forward-looking network scenarios.
end_date	Optional	Calendar date on which the site was or will be decommissioned and removed from active planning. After this date, the planning engine will no longer route demand to or hold inventory at this site.
site_level_1	Optional	The highest level in the site hierarchy, typically representing the broadest geographic or organizational grouping. This denormalized field enables efficient filtering and sorting across the site hierarchy.
site_level_2	Optional	The second level in the site hierarchy, representing a regional subdivision within the top-level site grouping. Leave blank if your site hierarchy has fewer than 2 levels.
site_level_3	Optional	The third level in the site hierarchy, representing a more granular geographic or organizational classification within level 2. Leave blank if your site hierarchy has fewer than 3 levels.
site_level_4	Optional	The fourth level in the site hierarchy, representing a specific site cluster or sub-region within level 3. Leave blank if your site hierarchy has fewer than 4 levels.

Column	Is the column required?	How is this column used in Supply Planning?
site_level_5	Optional	The fifth and most granular level in the site hierarchy . Leave blank if your site hierarchy has fewer than 5 levels.

product (required)

How is this data entity used? Supply Planning uses the product data to identify and classify all items being planned. Product attributes are used to establish hierarchy filters for supply plan review, link products to inventory policies and sourcing rules, and support multi-level product segmentation.

Column	Is the column required?	How is this column used in Supply Planning?
id	Required	Unique identifier for the product/material (SKU-level). Make sure the column values do not have duplicate IDs and special characters such as asterisk and double-quotes.
description	Required	Human-readable name or description of the product.
product_group_id	Optional	Foreign key linking the product to its category in the Product Hierarchy table.
unit_cost	Optional	Standard cost per unit used for inventory valuation and cost planning.
unit_price	Optional	Unit Price or standard price or MSRP of product.
base_uom	Optional	Unit of measure for product. Default to Eaches. Currently supply planning only supports Eaches.
volume_uom	Optional	Product volume. Currently volume_uom has to be the same across all products.

Column	Is the column required?	How is this column used in Supply Planning?
unit_volume	Optional	Volume of product per volum_uom defined at the warehouse level.
is_deleted	Optional	Boolean flag indicating whether the product has been discontinued or marked inactive. By default, the product is considered as Active or Not deleted.
product_level_1	Optional	The highest level in the product hierarchy, typically representing the broadest product category or division. This denormalized field enables efficient filtering and sorting across the product hierarchy.
product_level_2	Optional	The second level in the product hierarchy, representing a subcategory within the top-level product category. Leave blank if your product hierarchy has fewer than 2 levels.
product_level_3	Optional	The third level in the product hierarchy, representing a more granular product classification within level 2. Leave blank if your product hierarchy has fewer than 3 levels.
product_level_4	Optional	The fourth level in the product hierarchy, representing a specific product family or sub-classification within level 3. Leave blank if your product hierarchy has fewer than 4 levels.
product_level_5	Optional	The fifth and most granular level in the product hierarchy. Leave blank if your product hierarchy has fewer than 5 levels.

product_hierarchy (optional)

How is this data entity used? Supply Planning uses product hierarchy data to organize products into structured category trees. This enables multi-level filtering and aggregated analysis of supply

plans, and allows sourcing rules and inventory policies to be applied at a group level rather than individual SKU level.

Column	Is the column required?	How is this column used in Supply Planning?
id	Required	Unique identifier for the product group/category level.
description	Optional	Name or label of the product group/category.
parent_product_group_id	Optional	Foreign key to the parent category, forming the hierarchical tree structure for product classification.

geography (optional)

How is this data entity used? Supply Planning uses geography data to establish hierarchical location classifications for sites. Geographic entities enable regional rollup analysis and allow lead times and inventory policies to be defined at geographic region granularity.

Column	Is the column required?	How is this column used in Supply Planning?
id	Required	Unique identifier for the geographic entity (city, state, country, region).
description	Optional	Name of the geographic location (e.g., city name, region name).
parent_geo_id	Optional	Foreign key to the parent geography, enabling hierarchical rollups (e.g., city to state to country).

transportation_lane (conditional required)

How is this data entity used? Supply Planning uses transportation lane data to model the physical routes between origin and destination sites in the network. This data defines transit times and validity windows that the planning engine uses to schedule inbound and outbound shipments accurately.

Column	Is the column required?	How is this column used in Supply Planning?
id	Required	Unique identifier for the transportation lane, typically a composite of origin, destination, product group, and vendor.
from_site_id	Required	Origin site (supplier warehouse or manufacturing plant) from which goods are shipped.
to_site_id	Required	Destination site (receiving warehouse or plant) where goods are delivered.
product_group_id	Optional	Product category this lane applies to, enabling lane-level filtering by product group.
product_id	Optional	Specific product this lane applies to, for product-level lane definitions.
transit_time	Required	Number of days required to transport goods from origin to destination. Default is 0.
eff_start_date	Required	Date from which this transportation lane becomes active/valid.
eff_end_date	Required	Date after which this transportation lane is no longer valid.

trading_partner (conditional required)

How is this data entity used? Supply Planning uses trading partner data to identify vendors, suppliers, customers, and carriers in the supply chain. Trading partners are referenced by sourcing rules, inbound order lines, and vendor lead times to associate supply activities with external business entities.

Column	Is the column required?	How is this column used in Supply Planning?
id	Required	Unique identifier for the trading partner (vendor, customer, carrier, or other external entity).
description	Optional	Name of the trading partner organization.
eff_start_date	Optional	Date from which the trading partner relationship is valid. Default 1900-01-01 00:00:00.
eff_end_date	Optional	Date after which the trading partner relationship expires. Default 9999-12-31 23:59:59.

vendor_product (conditional required)

How is this data entity used? Supply Planning uses vendor product data to define which vendors are authorized to supply specific products. This data enables the planning engine to validate sourcing rules and generate purchase orders only to approved vendor-product combinations within their effective date ranges.

Column	Is the column required?	How is this column used in Supply Planning?
vendor_tp artner_id	Required	Vendor/supplier who can provide this product.
product_id	Required	Product that the vendor is authorized/capable of supplying.
eff_start_date	Required	Date from which the vendor is authorized to supply this product.
eff_end_date	Required	Date after which the vendor is no longer authorized for this product.

vendor_lead_time (conditional required)

How is this data entity used? Supply Planning uses vendor lead time data to determine the expected duration between placing a purchase order and receiving the goods at the destination site. Accurate lead times are critical for the planning engine to recommend orders with the correct timing to meet demand without creating unnecessary stockouts or excess inventory.

Column	Is the column required?	How is this column used in Supply Planning?
vendor_tp artner_id	Required	Vendor/supplier whose lead time is being defined.
product_id	Optional	Specific product the lead time applies to (if product-level granularity).
product_g roup_id	Optional	Product group the lead time applies to (if group-level granularity).
site_id	Required	Destination site for delivery, since lead times may vary by receiving location.
planned_l ead_time	Required	Expected number of days from order placement to delivery receipt.
eff_start_date	Optional	Date from which this lead time definition is valid. Default is 1900-01-01 00:00:00.
eff_end_date	Optional	Date after which this lead time definition expires. Default is 9999-12-31 23:59:59.
region_id	Optional	Geographic region for region-specific lead time definitions.

company (optional)

How is this data entity used? Supply Planning uses company data to support multi-entity planning environments where different legal entities or business units operate within a shared supply chain instance.

Column	Is the column required?	How is this column used in Supply Planning?
id	Required	Unique identifier for the company/legal entity.
description	Optional	Name or label of the company.
address_line_1	Optional	Primary street address line for the company.
address_line_2	Optional	Secondary address line (suite, building, floor).
city	Optional	City where the company is located.
state_prov	Optional	State or province where the company is located.
postal_code	Optional	Postal or ZIP code for the company address.
country	Optional	Country where the company is registered or operating.
phone	Optional	Primary contact phone number for the company.
email	Optional	Primary contact email address for the company.

sourcing_rules (required)

How is this data entity used? Supply Planning uses sourcing rules to determine how each product at each destination site is replenished. Rules specify whether supply comes from an external buy (purchase order), internal transfer (transfer order), or manufacturing (production order), and define the sourcing priority and allocation ratios when multiple supply options exist.

Column	Is the column required?	How is this column used in Supply Planning?
sourcing_rule_id	Required	Unique identifier for the sourcing rule.
product_id	Optional	Specific product governed by this sourcing rule.
product_group_id	Optional	Product group governed by this sourcing rule.

Column	Is the column required?	How is this column used in Supply Planning?
to_site_id	Required	Destination site that will receive supply under this rule.
sourcing_rule_type	Required	Nature of the sourcing: buy (external procurement), manufacture (internal production), or transfer (inter-plant).
from_site_id	Conditionally required	Source/origin site for the supply (internal plant or warehouse). Required when sourcing_rule_type is "transfer".
tpartner_id	Conditionally required	Vendor/supplier ID; required when sourcing_rule_type is "buy".
transportation_lane_id	Optional	Link to the transportation lane used for this sourcing path.
sourcing_priority	Optional	Priority ranking used when multiple sourcing options exist for the same product/site.
sourcing_ratio	Optional	Percentage allocation (0-100) defining what proportion of total demand this sourcing rule fulfills when multiple active sourcing rules exist for the same product-site-period combination.
min_qty	Optional	Minimum order quantity constraint for this sourcing rule.
qty_multiple	Optional	Order quantity must be a multiple of this value (lot sizing).
production_process_id	Optional	Manufacturing process/route used when sourcing_rule_type is "manufacture".
eff_start_date	Optional	Date from which this sourcing rule is active. Default is 1900-01-01 00:00:00.

Column	Is the column required?	How is this column used in Supply Planning?
eff_end_date	Optional	Date after which this sourcing rule expires. Default is 9999-12-31 23:59:59.

inventory_policy (required)

How is this data entity used? Supply Planning uses inventory policy data to determine the safety stock calculation method and replenishment thresholds for each product-site combination. The policy type (absolute level, days-of-cover demand, or days-of-cover forecast) governs how the planning engine computes minimum inventory buffers and generates replenishment recommendations.

Column	Is the column required?	How is this column used in Supply Planning?
id	Required	Unique identifier for the inventory policy record.
site_id	Optional	The site/warehouse to which this inventory policy applies.
dest_geo_id	Optional	Geographic region this policy targets, enabling region-specific stocking strategies.
product_id	Optional	Specific product this policy applies to (if product-level policy).
product_group_id	Optional	Product group this policy applies to (if group-level policy).
ss_policy	Required	Safety stock policy type governing how safety stock is calculated: Abs_level, DOC_dem, or DOC_fcst. Currently sl (service level) is not supported. 1. Abs_level: Uses units specified in min/max inventory qty. Ordering will be suggested whenever inventory falls below min. 2. DOC_dem: Uses days of cover computed from historical demand as target level of inventory. 3. DOC_fcst: Uses

Column	Is the column required?	How is this column used in Supply Planning?
		days of cover computed from forecast as target level of inventory.
min_safety_stock	Optional	Absolute minimum safety stock quantity floor that overrides any policy-calculated value. Expressed in the product's base unit of measure.
min_inventory_qty	Conditionally required	Either min_inventory_qty/max_inventory_qty or target_inventory_qty is required when ss_policy is abs_level. Minimum inventory level (reorder point) that triggers a replenishment signal when on-hand quantity drops below this threshold.
max_inventory_qty	Conditionally required	Either min_inventory_qty/max_inventory_qty or target_inventory_qty is required when ss_policy is abs_level. Maximum allowable inventory quantity (order-up-to level) defining the upper cap for replenishment orders. Prevents overstocking by limiting the total inventory a site should hold for this product.
target_inventory_qty	Conditionally required	Either min_inventory_qty/max_inventory_qty or target_inventory_qty is required when ss_policy is abs_level. This is the target inventory quantity that the system will maintain.
min_doc_limit	Conditionally required	Either (min_doc_limit and max_doc_limit pair) or target_doc_limit is required when ss_policy is doc_dem or doc_fcst. Minimum days of cover floor ensuring that inventory always covers at least this many days of forward demand.
max_doc_limit	Conditionally required	Either (min_doc_limit and max_doc_limit pair) or target_doc_limit is required when ss_policy is doc_dem or doc_fcst. Maximum days of cover limit representing the upper boundary on inventory coverage.

Column	Is the column required?	How is this column used in Supply Planning?
target_doc_limit	Conditionally required	Either (min_doc_limit and max_doc_limit pair) or target_doc_limit is required when ss_policy is doc_dem or doc_fcst. Target days of cover (historical demand or forecast based).
eff_start_date	Optional	Date from which the inventory policy is effective. Default is 1900-01-01 00:00:00.
eff_end_date	Optional	Date after which the inventory policy expires. Default is 9999-12-31 23:59:59.

inventory_level (required)

How is this data entity used? Supply Planning uses inventory level snapshots to establish the current on-hand stock position at each site for each product. This data is the starting point for replenishment calculations, enabling the planning engine to determine the gap between current inventory and the target level defined by the inventory policy.

Column	Is the column required?	How is this column used in Supply Planning?
snapshot_date	Required	Date/time when the inventory snapshot was taken.
site_id	Required	Site/warehouse where this inventory is physically located.
product_id	Required	Product/material whose inventory is being reported.
on_hand_inventory	Required	Quantity of unrestricted, available-for-use stock on hand.
inv_condition	Optional	Condition or quality classification of the inventory (e.g., available, blocked, in-quality-hold).

forecast (conditionally required)

How is this data entity used? Supply Planning uses forecast data as the primary demand signal for replenishment calculations. The planning engine uses the forecasted demand quantities by product and site to determine how much inventory needs to be replenished and when, driving purchase, transfer, and production order recommendations across the planning horizon. This data is not required if forecast is generated by Amazon Connect Decisions Demand Intelligence.

Column	Is the column required?	How is this column used in Supply Planning?
snapshot_date	Required	Date when this forecast was generated or last updated.
site_id	Required	Site for which demand is being forecasted.
product_id	Required	Product for which demand is being forecasted.
mean	Conditionally required	Either mean or p50 needs to be provided. Arithmetic mean of the demand forecast distribution for this product-site-period combination. Serves as the primary demand signal.
p50	Conditionally required	Either mean or p50 needs to be provided. Median (50th percentile) forecast quantity representing the most likely demand scenario.
p10/20/30/40/50/60/70/80/90	Optional	Probabilistic forecast values at 10th, 50th, and 90th percentiles for demand variability analysis.
forecast_start_dttm	Required	Start date/time of the forecast bucket period.
forecast_end_dttm	Required	End date/time of the forecast bucket period.

outbound_order_line (conditionally required)

How is this data entity used? Supply Planning uses outbound order line data to track customer orders and demand commitments. This data provides visibility into confirmed demand, allowing the planning engine to reserve inventory for committed customer orders and factor fulfillment obligations into replenishment recommendations.

Column	Is the column required?	How is this column used in Supply Planning?
id	Required	Unique identifier for the outbound order line (delivery line item).
product_id	Required	Product being shipped to the customer.
cust_order_id	Required	Customer's sales order number this line belongs to.
status	Optional	Current fulfillment status of the line (e.g., Open, Cancelled, Closed).
order_date	Optional	Date the order was placed.
init_quantity_requested	Optional	Original quantity the customer initially ordered.
final_quantity_requested	Required	Current/adjusted quantity still to be fulfilled (accounts for partial shipments and order changes).
requested_delivery_date	Required	Date the customer originally requested delivery.
promised_delivery_date	Optional	Date committed to the customer after availability check.
actual_delivery_date	Optional	Date the goods were actually received by the customer.
ship_from_site_id	Required	Warehouse or plant from which the product will be shipped.

Column	Is the column required?	How is this column used in Supply Planning?
quantity_promised	Optional	Quantity confirmed/committed to the customer.
quantity_delivered	Optional	Quantity that has been physically shipped/delivered.

inbound_order_line (conditionally required)

How is this data entity used? Supply Planning uses inbound order line data to track open purchase orders, transfer orders, and manufacturing orders. This data provides visibility into planned supply receipts, enabling the planning engine to account for in-transit and confirmed inventory when generating replenishment recommendations.

Column	Is the column required?	How is this column used in Supply Planning?
id	Required	Unique identifier for the order line item (typically a composite of order ID and line number).
order_id	Required	Inbound order header number.
tpartner_id	Conditionally required	Mandatory for PO (Purchase Orders). Links the order line to the vendor/supplier. Default SCN_RESERVED_NO_VALUE_PROVIDED.
product_id	Required	The product/material being ordered on this line.
order_type	Required	Type classification of this order line. Should be PO (Purchase Order), TO (Transfer Order) or MO (Manufacture Order).
status	Optional	Current status of this specific line item (e.g., Open, InTransit, Closed, Cancelled).

Column	Is the column required?	How is this column used in Supply Planning?
from_site_id	Conditionally required	Mandatory for TO (Transfer Orders). Origin site from which the product is being shipped.
to_site_id	Required	Destination site for this line item's delivery.
quantity_submitted	Conditionally required	Original quantity requested/ordered. Either quantity_submitted or quantity_confirmed should be provided.
quantity_confirmed	Conditionally required	Quantity confirmed by the supplier for delivery. Either quantity_submitted or quantity_confirmed should be provided.
quantity_received	Optional	Quantity physically received at the destination site.
submitted_date	Optional	Date this line item was created in the system.
earliest_delivery_date	Conditionally required	At least one of earliest_delivery_date, latest_delivery_date or expected_delivery_date should be provided.
expected_delivery_date	Conditionally required	Planned or promised delivery date for this line item. At least one of earliest_delivery_date, latest_delivery_date or expected_delivery_date should be provided.
latest_delivery_date	Conditionally required	At least one of earliest_delivery_date, latest_delivery_date or expected_delivery_date should be provided.
earliest_ship_date	Optional	Earliest date and time when vendor shipped products for this order-line. If populated, planning engine will use this date to determine order start date.
latest_ship_date	Optional	Latest date and time when vendor shipped products for this order-line. If populated, planning engine will use this date to determine order start date.

Column	Is the column required?	How is this column used in Supply Planning?
product_group_id	Optional	Product group context for this alternate mapping.
to_site_id	Optional	Destination site for the alternate product delivery.
from_site_id	Optional	Origin site for the alternate product sourcing.
transportation_lane_id	Optional	Transportation lane to use when sourcing this alternate product.
min_qty	Optional	Minimum order quantity for the alternate product.
qty_multiple	Optional	Order quantity must be a multiple of this value for the alternate product.

product_bom (conditionally required)

How is this data entity used? Supply Planning uses Product Bill of Materials (BOM) data to explode finished goods demand into component-level requirements. When the planning engine generates production orders for manufactured products, it uses the BOM to calculate how much of each raw material or sub-assembly is needed, creating dependent demand signals for component procurement and production planning.

Column	Is the column required?	How is this column used in Supply Planning?
id	Required	Unique identifier for the BOM line record, typically a surrogate key or composite of parent product, component product, site, and effective dates.
product_id	Required	Foreign key to the Product table identifying the finished good or parent assembly that is manufactured using this BOM.

Column	Is the column required?	How is this column used in Supply Planning?
site_id	Required	Foreign key to the Site table identifying the manufacturing location where this BOM is valid.
component_product_id	Required	Foreign key to the Product table identifying the raw material, sub-assembly, or intermediate component required to produce the parent product.
component_quantity_per	Required	Quantity of the component product required to produce one unit of the parent finished good, expressed in the component's base unit of measure.
eff_start_date	Required	Calendar date from which this BOM relationship is effective and should be used by the planning engine for demand explosion.
eff_end_date	Required	Calendar date after which this BOM relationship is no longer valid and should not be used for planning.
production_process_id	Optional	Foreign key to the Production Process table identifying which manufacturing route or recipe uses this specific BOM.

production_process (not supported in Supply Planning right now, to be added soon)

How is this data entity used? This data entity is not supported in Supply Planning right now. But it will be added soon. When supported, the planning engine will use setup time and operation time to calculate the manufacturing lead times and evaluate production capacity feasibility.

Column	Is the column required?	How is this column used in Supply Planning?
production_process_id	Required	Unique identifier for the manufacturing process or production route. Represents a specific sequence of

Column	Is the column required?	How is this column used in Supply Planning?
		operations that transforms raw materials/components into a finished product at a given site.
product_id	Required	Foreign key to the Product table identifying the finished good or output product manufactured by this production process.
site_id	Required	Foreign key to the Site table identifying the manufacturing facility where this production process is physically executed.
setup_time	Required	Fixed one-time duration (in hours) required to changeover, calibrate, and prepare the production equipment before a batch run can begin.
operation_time	Required	Variable processing time (in hours) consumed per unit of output produced on this manufacturing line. Total production duration = setup_time + (operation_time x planned quantity).

segmentation_rule (not supported in Supply Planning right now, to be added soon)

How is this data entity used? This data entity is not supported in Supply Planning right now. But it will be added soon. Segmentation Rule data will enable differentiated planning by classifying products and sites into segments (e.g., by demand variability, customer tier, or product lifecycle). When supported, the planning engine will apply segment-specific inventory policies, forecast models, and service-level targets.

Column	Is the column required?	How is this column used in Supply Planning?
segmentation_rule_id	Required	Unique identifier for a segmentation rule mapping a specific segment_type + segment_value combination to product/site/channel dimensions.
segment_type	Required	Name of the segmentation dimension being defined (e.g., DemandVariability, CustomerTier, ProductLifecycle, ServiceLevel).
segment_value	Required	Specific value within the segment_type classification (e.g., High, Medium, Low for DemandVariability; A, B, C for CustomerTier).
segmentation_rule_description	Optional	Human-readable explanation of the business logic or criteria used to assign entities to this segment.
product_id	Optional	Foreign key to the Product table. When populated, the rule applies to this specific product only.
product_level_1	Optional	Highest level in the product hierarchy. Enables broad segmentation across entire product families.
product_level_2	Optional	Second-level product hierarchy classification. Provides intermediate product grouping for segmentation.
product_level_3	Optional	Third-level product hierarchy classification. Enables mid-granularity product segmentation.
product_level_4	Optional	Fourth-level product hierarchy classification. Enables fine-grained product segmentation.
product_level_5	Optional	Most granular product hierarchy level. Leave blank if the rule targets a coarser level or a specific product_id.
site_id	Optional	Foreign key to the Site table. When populated, the rule applies to this specific site only.

Column	Is the column required?	How is this column used in Supply Planning?
site_level_1	Optional	Highest level in the site hierarchy. Enables broad geographic segmentation across the network.
site_level_2	Optional	Second-level site hierarchy. Provides intermediate geographic grouping for segmentation.
site_level_3	Optional	Third-level site hierarchy. Enables mid-granularity geographic segmentation.
site_level_4	Optional	Fourth-level site hierarchy. Enables fine-grained geographic segmentation.
site_level_5	Optional	Most granular site hierarchy level. Leave blank if the rule targets a coarser level or a specific site_id.
city	Optional	City name of the ship-from site location. Provides geographic context for location-based segmentation.
state_prov	Optional	State or province of the ship-from site. Enables regional segmentation and compliance with jurisdiction-specific business rules.
postal_code	Optional	Postal/ZIP code of the ship-from site. Enables precise geographic targeting for segmentation rules.
country	Optional	Country of the ship-from site. Enables country-level segmentation for international operations.
trading_partner_id	Optional	Foreign key to the Trading Partner table. Enables segmentation rules based on supplier, customer, or carrier relationships.
company_id	Optional	Foreign key to the Company table. Enables multi-entity segmentation where different legal entities maintain separate classification rules.

Column	Is the column required?	How is this column used in Supply Planning?
channel_id	Optional	Foreign key referencing the sales/distribution channel. Enables channel-specific segmentation for omnichannel operations.
creation_date	Optional	ISO 8601 timestamp when this segmentation rule was created. Supports audit trail and version tracking.
connection_id	Optional	System-assigned identifier for the data source connection. Auto-populated by Prism upon ingestion — do not supply manually.
updated_at	Optional	System-managed timestamp of the last modification to this record. Auto-populated by Prism for data lineage — do not supply manually.
updated_by	Optional	System-managed identifier of the user or process that last modified this record. Auto-populated by Prism for data lineage — do not supply manually.

Direct upload data entities

Note

capacity_constraints, capacity_product_mapping, warehouse_space_limit, and planning_time_fence_constraints don't require setting up the data ingestion flows. These data entities can be directly uploaded on the supply planning configuration page through Decision Teammates.

capacity_constraints (conditionally required)

How is this data entity used? Supply Planning uses capacity constraint data to define the minimum and maximum throughput limits for each capacity resource at each site within defined time buckets. The planning engine enforces these constraints when generating production and storage recommendations to ensure plans remain within physical or contractual feasibility boundaries.

Column	Is the column required?	How is this column used in Supply Planning?
id	Required	Unique identifier for the capacity constraint record. References a combination of capacity resource, site, time bucket, and effective date range.
capacity_id	Required	Foreign key to the Capacity table identifying the specific resource being constrained (e.g., a production line, storage zone, dock, or labor pool).
site_id	Required	Foreign key to the Site table identifying the physical location where this capacity constraint applies.
bucket_type	Required	Defines the time bucket granularity for measuring and enforcing capacity utilization. Accepted values: Daily, Weekly, Monthly.
eff_start_date	Required	Calendar date from which this capacity constraint becomes active and enforceable by the planning engine.
eff_end_date	Required	Calendar date after which this capacity constraint expires and is no longer enforced.
min_capac ity_usage	Optional	Minimum throughput or utilization volume required within the defined time bucket. Represents a contractual or operational floor.
max_capac ity_usage	Required	Maximum throughput or utilization volume allowed within the defined time bucket. The planning engine must not schedule production or allocate inventory beyond this limit.

capacity_product_mapping (conditionally required)

How is this data entity used? Supply Planning uses capacity product mapping data to link products to the capacity resources they consume at each site. This mapping enables the planning engine

to aggregate total capacity demand across all products sharing a common resource, identify bottlenecks, and prioritize production sequencing when capacity is limited.

Column	Is the column required?	How is this column used in Supply Planning?
id	Required	Unique identifier for the capacity-to-product mapping record linking a specific product to the capacity resource it consumes at a given site.
product_id	Required	Foreign key to the Product table identifying the specific SKU or material whose production or handling consumes capacity from the mapped resource.
site_id	Required	Foreign key to the Site table specifying the location where this product-to-capacity relationship is valid.
capacity_id	Required	Foreign key to the Capacity table identifying which capacity resource this product consumes during production, storage, or handling.

warehouse_space_limit (conditionally required)

How is this data entity used? Supply Planning uses warehouse space limit data to constrain replenishment recommendations so that planned inventory never exceeds the physical storage capacity of a warehouse. This prevents the planning engine from generating supply orders that would result in overcrowding or infeasible storage requirements at a site.

Column	Is the column required?	How is this column used in Supply Planning?
site_id	Required	Identifier of the warehouse whose capacity is being defined.
max_capacity	Required	Maximum storage capacity of the warehouse, expressed in the volume unit of measure defined in the Product

Column	Is the column required?	How is this column used in Supply Planning?
		entity. Used to constrain supply plans from exceeding physical space.

planning_time_fence_constraints (conditionally required)

How is this data entity used? Supply Planning uses planning time fence constraint data to define time boundaries within which the planning engine is restricted from making changes to existing supply orders. These constraints enable planners to protect near-term supply plans from automatic rescheduling, preserving confirmed orders for specific products and sites within the defined time fence window.

Column	Is the column required?	How is this column used in Supply Planning?
rule_id	Required	Unique identifier for the planning time fence constraint rule.
product_id	Optional	Foreign key to the Product table. When populated, the time fence applies to this specific product only. When left blank, the constraint applies to all products at the specified site.
site_id	Optional	Foreign key to the Site table. When populated, the time fence applies to this specific site only. When left blank, the constraint applies to the specified product across all sites.
planning_time_fence_days	Required	Number of calendar days from the planning run date within which the planning engine will not automatically reschedule or cancel existing supply orders. Orders with order start dates within this window are treated as frozen and protected from automated changes.

Creating a New Plan

To create a supply plan:

1. Navigate to the Plans section in AWS Connect Decisions and choose **Create Plan** in Supply Plan section. Once a supply plan is created, you have the flexibility to edit it at any time to reflect changing business needs or incorporate new information.

2. Configure Time Horizon settings:

- **Time bucket:** Select Daily or Weekly based on your planning needs.
- **Plan horizon:** Specify the forecast horizon:
 - Daily: 1 to 365 days
 - Weekly: 1 to 52 weeks
- **Plan start date:** For weekly plans, a day of the week can be selected as the plan start date. It also indicates the beginning of the week.

3. Plan Schedule

Choose between One time run and **Recurring Schedule**

- **For Recurring Schedule,** choose **Schedule Frequency**
 - **Daily:** Enter the UTC schedule time
 - **Weekly:** Select the day of the week and enter the UTC schedule time
 - **Monthly:** Select the day of the month and enter the UTC schedule time

4. Demand Netting:

Configure what demands to drive the supply plan generation. Choose Forecast, sales order, or both

- **Forecast:** Forecasted demand
- **Sales order:** Actual sales order demand

5. Demand Time Fence:

The period during which the supply plan will ignore the forecasted demand.

6. Forecast consumption window:

The period in which actual customer orders replace or consume forecasted demand.

- **Forecast consumption window - forward days**
- **Forecast consumption window - backward days**

7. Historical Period for Demand

Number of historical days to consider for average historical demand calculation

8. Past due supply days

The numbers of days the supply orders can be past due but still considered as supply.

9. Planning time fence

The period in which the supply plan is frozen. No new planned orders will be created with the planning time fence.

10 Configuration rules

Constraint rules that guide the system in generating plans aligned with your business objectives. Add rules using the production capacity constraints template or warehouse space limit template.

- **Production Capacity Constraints:** Download the Production Capacity Constraints template csv files (zipped). Populate the data and upload 3 csv files (not zipped) back using Decision teammate.
- **Warehouse Space Limit:** Download the Warehouse Space Limit template csv file. Populate the data and upload it back using Decision teammate.

11 Generate Supply Plan

Choose **Generate supply plan** button on the top to start the plan generation job or wait for the scheduled plan run.

Configuring Plan Parameters

Configure an existing Supply Plan

You can navigate to plan configuration page by selecting **Configuration** from the meatball menu on the top of the planning hub or action menu from plan list. You can either manually update plan configuration change or chat with Decision Teammate to make the change.

Reviewing and Editing Plans

To review a supply plan:

1. Navigate to Plans from your AWS Supply Chain homepage.
2. Click on the plan name or select **Review Plan** from the **Action List**.

3. The system displays a high-level view of supply plan aggregated across products and sites.
4. Use the filtering options to focus on specific products, locations, or time periods.

Review supply plan details

- Demand section
 - **Total Demand:** Total summed up demand of the current time bucket
 - **Original Forecast:** Original demand forecast without forecast consumption
 - **Net Forecast:** Remaining forecast after forecast consumption
 - **Sales Orders:** Open sales orders which need to be fulfilled
 - **Open Transfer Order Demand:** Demand from open transfer orders
 - **Open Production Order Demand:** Demand from open production orders
 - **Planned Transfer Order Demand:** Demand from planned transfer orders
 - **Planned Production Order Demand:** Demand from planned production orders
- Supply section
 - **Total Supply:** Total summed supply of the current time bucket
 - **On Hand:** On-hand inventory
 - **Open Purchase Orders:** Open purchase order supply
 - **Open Transfer Orders:** Existing transfer orders which are either Open or InTransit
 - **Open Production Orders:** Open production order supply
 - **Planned Purchase Orders:** Planned purchase orders generated by the system
 - **Planned Transfer Orders:** Planned transfer orders generated by the system
 - **Planned Production Orders:** Planned production orders generated by the system
- Projected inventory section
 - **Projected Inventory:** Projected ending inventory balance
 - **Min Inventory Level:** Minimum inventory level calculated based on the inventory policy
 - **Target Inventory Level:** Target inventory level calculated based on the inventory policy
 - **Max Inventory Level:** Maximum inventory level calculated based on the inventory policy
 - **Safety Stock:** Safety stock requirements specified in the inventory policy
 - **Projected Days (Weeks) of Cover:** Projected Days(weeks) of Supply depending on plan time bucket configuration

Review order details

The system provides an Order View to allow planners to drill down into the transactional data behind their plans. This view is automatically filtered based on the user's selection in the main tabular view. The user can navigate to the order detail view by clicking on order quantities in Supply Plan view and select the link. The Order Detail View will show the corresponding supply orders tied to that particular supply quantity. For example, if the supply quantity represents the total of all open purchase orders in a particular week, the order detail view can be brought up to show the corresponding open purchase order details.

Edit order details

Users can review the generated planned orders and make modifications such as adjusting quantities and updating the order start date and expected delivery date. By default, when a user makes an adjustment, the modified order is automatically marked as a firm order, indicating that it will be preserved during plan rerun. Once the orders are adjusted, the user needs to click on **Rerun** button to recalculate the supply plan such that the impact of these changes is reflected across the whole plan. During this rerun, the system respects these firmed orders and does not recalculate them, ensuring planner intent is maintained.

Publishing Supply Plans

After planners have reviewed and refined the supply plan, you can publish the final plan to make it available for downstream processes.

To publish a demand plan:

1. Review the aggregate performance metrics to ensure supply plan quality meets your standards.
2. Coordinate with your planning team to confirm all necessary adjustments have been completed.
3. Click **Publish Plan** to process your request.
4. The system automatically saves the published plan to your configured Amazon S3 folder.

Common Issues and Solutions

Common Data Issues

Demand history has mixed date formats

Source systems may export dates as DD/MM/YYYY, MM/DD/YYYY, or YYYY-MM-DD, sometimes within the same file. The system may parse these incorrectly, assigning orders to wrong months.

Fix: Standardize date formats in your export process. If you can't control the source, add date validation in your data flow SQL.

Negative quantities in order history

Credit notes, returns, or reversals can appear as negative quantities. These can distort demand averages and confuse the model.

Fix: Filter to only positive quantities, or filter by order status (e.g., only Paid/Invoiced orders).

Record counts don't match your source system

- Most commonly caused by composite key collisions — if two records share the same unique identifier, one overwrites the other.
- Can also occur if filter criteria in the data mapping exclude records you expect to see.

Provide a specific product+site example and your expected record count so the team can trace the discrepancy.

Orders appear in the system that don't exist in ERP (or vice versa)

- Orders fulfilled or removed between report runs will disappear from the next refresh but may still appear in exceptions generated from the previous day's data.
- Newly created orders won't appear until the next data refresh.

This is expected behavior — exceptions will update on the next evaluation cycle after fresh data loads.

Plan input files include products from other plants or business units

If your source system exports include products outside the scope of your forecasting project:

- **The system will filter to the product master automatically.** Only products present in your product master file will be included in the forecast. However, if a large percentage of your input file is out of scope (e.g., 50%+ of rows), this indicates the source export needs tightening.
- **Check your product coverage rate regularly.** After each data load, verify what percentage of products in your sales and forecast input files match the product master. If coverage drops below 80%, investigate whether the source export scope has changed or the product master needs updating.
- **Out-of-scope products in plan inputs can cause inflated totals.** If your EDI or SIOP files include products from other plants, the aggregate forecast signal will be higher than it should be. Ensure plan input files are filtered to the same product scope as your product master before loading.

Common Exception and Recommendation Issues

Same product+site appears multiple times in exception list

This can happen when the underlying rule generates a separate exception for each qualifying date in the projection horizon.

Contact your support team to adjust the rule to flag only the earliest breach date per product+site.

Recommendation doesn't match what I see in the chart

The recommendation is generated by an AI agent that analyzes the data available at the time the exception was created. If data has changed since, the recommendation may reference orders or quantities that are no longer current.

- Check the timestamp on the exception — if it's more than a day old, the recommendation may be stale.
- If the recommendation is clearly wrong (e.g., ignores a large order visible in the chart), provide feedback using thumbs down and report the specific exception to your support team.

Impact date or "Act By" date seems wrong

- The impact date shows when the inventory issue begins (e.g., when stockout starts or excess exceeds threshold).
- The "Act By" date should account for lead time so you have time to act before the issue materializes. If Act By equals the impact date, lead time may not be incorporated — report this to your support team.

Recommendations reference orders I can't find in ERP

ERP snapshots change daily. An order referenced in yesterday's recommendation may have been fulfilled, cancelled, or rescheduled in today's ERP run.

This is a known limitation of ERP-based data. Historical consumption data can be added to provide better context.

Common Accuracy Issues

Forecast is significantly worse than a simple moving average

If your ASC forecast is losing to a 6-month moving average on aggregate WAPE, check these common causes:

- **Too many low-volume/inactive products in scope.** Products with sparse, intermittent demand are hard for any model to beat a simple average on. Use a preprocessing rule to scope the forecast to products with meaningful demand history (e.g., at least 6 months of non-zero demand).
- **Training on stale or contaminated history.** If your order history goes back many years, old demand patterns may not reflect current reality. Consider a preprocessing rule to limit training history to the most recent 3–5 years, or to replace anomalous periods (e.g., COVID) with normalized values.
- **Demand spikes from one-time orders.** A single large bulk order can create a false upward trend in the training data. Use a preprocessing rule to cap anomalous monthly demand values at a multiple of the trailing average (e.g., 5x).
- **Consensus rules applied in wrong direction.** The LLM agent can misinterpret rule language. "Decrease by 27%" may be applied as an increase. Always validate consensus output against

baseline by comparing specific products and months. Use explicit multiplication language ("multiply by 0.725") rather than directional language ("decrease by 27.5%").

Over-forecasting bias (forecast systematically higher than actuals)

A positive bias means you're ordering more than needed across the catalog. Common causes:

- **The model is trained on a growth period.** If recent years showed growth that isn't continuing, the model extrapolates a trend that no longer exists.
- **Consensus rules are stacking upward adjustments.** Multiple rules that each increase the forecast (stockout bias, trend boost, seasonal uplift) can compound. Review which rules are active and check if they're all applying to the same products.
- **Deleted/discontinued products still in scope.** Products with trailing-off demand that are still being forecast will show systematic over-forecasting.

Under-forecasting bias (forecast systematically lower than actuals)

A negative bias means you're consistently forecasting less than actual demand, leading to potential stockouts and expediting costs. Common causes:

- **External forecast signals not being incorporated.** If you have plan inputs (e.g., EDI customer forecasts, SIOp production plans) loaded but your consensus rules aren't applying them, the forecast defaults to the statistical baseline — which may not capture demand signals your planners see. Verify that consensus rules are actually modifying the output by comparing the ConsensusForecast export against the Forecast (baseline) export. If they're identical, the rules aren't firing.
- **Sparse product×site combinations pulling the aggregate down.** If you forecast at product×site granularity but many combinations have zero or near-zero demand, the model produces small non-zero forecasts for inactive combinations. These don't add up to much individually but collectively drag the total forecast below actuals. Use a preprocessing rule to exclude combinations with insufficient demand history, or use conditional zero-filling in your plan inputs to explicitly signal "no demand expected" for inactive combinations.
- **The model hasn't captured a recent growth trend.** Statistical models weight historical data. If your business has grown significantly in recent months but the model has years of lower-volume history, it will lag the trend. This typically improves over time as the model accumulates

more recent data. In the interim, consider a consensus rule that uses a trailing average of recent actuals as a floor for outer forecast weeks.

- **Year-over-year seasonality mismatch.** If this year's demand pattern differs from prior years (e.g., earlier seasonal ramp, new product launches), the model may under-forecast during the divergent period. Check whether the under-bias is concentrated in specific weeks or months that differ from the prior year pattern.

Forecast accuracy degrades significantly at longer horizons

It's normal for accuracy to worsen as the forecast horizon increases — week 1 is always more accurate than week 8. However, if the degradation is steeper than expected:

- **External signals only help near-term.** If you have consensus rules that incorporate customer forecasts (EDI) for the first few weeks, accuracy will be noticeably better near-term and drop off when the rules stop applying. This is expected — consider extending the rules to cover more weeks with a blended approach (e.g., 50/50 blend of external signal and baseline for medium-term weeks).
- **The baseline reverts to a long-term average at longer horizons.** Statistical models become less confident at longer horizons and tend toward the historical mean. If recent demand is above the historical mean, the outer weeks will appear under-biased. This is a model behavior, not a configuration issue.
- **Demand volatility makes longer horizons inherently harder.** If your demand has high week-to-week variability (coefficient of variation > 0.5), even a perfect model will show high error at longer horizons. Focus accuracy evaluation on the first 3–4 weeks, which is the actionable planning window for most operations.

External forecast (EDI/customer forecast) not improving accuracy when used in consensus rules

If you've added consensus rules to incorporate external forecasts but accuracy hasn't improved:

- **The external signal may not cover enough products.** EDI or customer forecasts typically only cover a subset of your product catalog (often 30–50%). Products without external signal still use the baseline. Check your coverage rate — if it's below 50%, the impact on aggregate accuracy will be limited.

- **The external signal may not be accurate enough to help.** Measure the external forecast's accuracy independently before using it in rules. If its WAPE is worse than the baseline, incorporating it will hurt rather than help. Consider limiting the rule to specific sites or products where the external signal is demonstrably better (e.g., volume-weighted WAPE below 50%).
- **The external signal doesn't report zeros.** Many EDI systems only send records for products with active orders — they omit products with zero demand rather than explicitly reporting zero. If your consensus rule says "when EDI = 0, set forecast to 0," it will never fire because there are no zero records. You need to generate synthetic zero records in preprocessing for product×site combinations that have no external signal AND no recent sales history.
- **The external signal accuracy varies by horizon.** Customer forecasts are typically most accurate for the immediate next week (essentially confirmed orders) and degrade rapidly. A rule that uses the external signal directly for all weeks may hurt accuracy at longer horizons. Consider a tiered approach: direct replacement for weeks 1–3, blended for weeks 4–6, baseline only for weeks 7+.

Planning rules not taking effect

If a consensus rule doesn't appear to change the forecast:

- **The rule may have been overridden by a higher-priority rule.** Rules are applied in priority order. A later rule can undo an earlier one. Check the rule ordering.
- **The rule condition may not match any products.** If the rule references a product attribute (e.g., `product_group_id`) that isn't in the item metadata, it will silently match nothing.
- **The rule language was misinterpreted.** The LLM agent generates code from natural language. Ambiguous phrasing can produce unexpected results. Be as specific and literal as possible. Use exact field names, explicit multipliers, and clear conditions.

Consensus plan output is identical to the baseline forecast

If the ConsensusForecast export has the same values as the Forecast (baseline) export, the consensus rules did not execute. Common causes:

- **Dimension mismatch in the join.** The consensus engine joins plan inputs to the baseline on dimension columns (product ID, site ID, date). If the column names differ between the baseline and the plan inputs (e.g., baseline uses `item_id` while EDI uses `product_id`), the join produces no matches and all rules fall through to the baseline default. Verify that the dimension mapping in your data flow configuration correctly maps between the two schemas.

- **Date format mismatch.** The baseline may store dates as 2026-03-02 while plan inputs store them as 2026-03-02T00:00:00.000Z. If the join requires exact match, timezone-aware and timezone-naive dates won't match. Check that date columns are converted to the same format before joining.
- **Plan inputs not loaded.** Verify that your plan input files (EDI, SIOP, etc.) were successfully ingested. Check the record counts in the system — if they show zero rows for a plan input, the file may have failed to load.
- **The consensus forecast_id matches the baseline forecast_id.** If both exports share the same forecast_id, the consensus engine produced a direct copy of the baseline without processing. This indicates a system-level issue — contact your support team with the forecast_id and demand_plan_run_id.

Consensus rules apply to wrong products or sites

If a rule that should only apply to specific sites or product categories is affecting the entire catalog:

- **The site/product filter condition may reference the wrong column.** If your rule says "apply to sites in [list]" but the generated code checks a column that doesn't exist or has different values, the filter may silently pass all rows. Verify by spot-checking a few specific products that should NOT be affected by the rule.
- **Rule priority order may be inverted.** Rules are applied as a chain where later rules override earlier ones. If a broad rule (e.g., "use baseline for everything") is applied after a specific rule (e.g., "use EDI for these 50 sites"), the broad rule will undo the specific one. Ensure your rule descriptions clearly state the priority order.

Forecast values are fractional (e.g., 2,500.37 units)

Statistical models produce continuous values, not integers. If your business deals in whole units, case packs, or minimum order quantities:

- **Add a rounding rule as the final consensus step.** A simple "round to nearest integer" rule applied after all other consensus rules will clean up fractional values. Values below 0.5 will round to zero, which is appropriate for very low-demand combinations.
- **Consider rounding to operational quantities.** If your products ship in standard pack sizes (e.g., cases of 12, pallets of 48), rounding to the nearest valid pack size may improve both the usability

and accuracy of the forecast. This requires pack size data in your product master. Share your MOQ or pack size data with your support team to explore this option.

Product coverage drops significantly after adding preprocessing rules

Preprocessing rules that filter the training data (e.g., "only forecast products with at least 8 weeks of non-zero demand") can dramatically reduce the number of products in the forecast if your data is sparse at the product×site level:

- **Check the granularity.** A product may have 52 weeks of demand at the product level but only 3 weeks at any individual product×site combination. A minimum history threshold applied at product×site level will exclude most combinations. Consider applying the threshold at product level instead, or lowering the threshold significantly.
- **Test before deploying.** Before activating a preprocessing rule, count how many product×site combinations pass the filter vs. your current total. If more than 20% are excluded, the rule is likely too aggressive. Start with a lenient threshold and tighten gradually.