



Platform Architecture

This white paper details the technical capabilities of the Alloy platform and how it connects heterogeneous retail and supply chain data to uniquely help customers make better decisions.

Table of Contents

Introduction	1	Modeling	13
Industry context		Product	
Design principles		Network	
Platform Overview	4	Time	
Integration		Aggregation and disaggregation	
Modeling		Prediction	17
Prediction		Industry-specific metrics	
Insight and action		Administration	
Identity and access management		Insight and Action	20
Best-in-class Software as a Service (SaaS)		Analysis	
Integration	9	Planning	
Pre-built connections		Conclusion	22
Historical data			
Failure recovery			
Data cleansing and validation			
Data export			

Introduction

Alloy is purpose-built to enable the demand-driven supply chain. The platform equips organizations to respond efficiently when changes in demand or supply inevitably arise, so you can carry less inventory and reduce costs, while improving customer service levels and meeting demand everywhere it exists.

"In the past, matching supply and demand has been extremely difficult given the long reaction time of supply chains and the inherent challenges that arise from communicating across the various IT platforms in a company extended chain.

A DDSC [Demand-Driven Supply Chain] offers real-time information on demand and inventory levels to all supply chain participants so that they can react quickly and effectively when unexpected changes arise."

- The Boston Consulting Group, *The Demand-Driven Supply Chain*

Many companies struggle to even start achieving the promise of the demand-driven supply chain due to the technical difficulty of processing real-time information on demand and inventory levels from across your supply chain network. The data comes from siloed sources, at different frequencies and granularities, in inconsistent formats that change sporadically, and with no connection to your existing data models. Turning it into useful information for business teams to inform their decision-making requires solving many diverse and ongoing engineering challenges.

Alloy develops and maintains a robust, intelligent platform architecture to address these needs and power demand-driven insights for Sales, Marketing and Supply Chain.

Industry context

In the consumer goods industry, two key trends have both made it harder to operate a demand-driven supply chain and elevated it as an imperative for success for today's businesses.

- **Fragmenting consumer demand.** Consumers have a multitude of options for purchasing a product. They can still go to a store and choose from a variety of retailers or a brand's own store, or go online and choose from an even wider range of e-commerce stores or a brand's own website. They can have it delivered, or buy online and pickup in store (BOPIS). And new ways of fulfilling demand are sure to continue emerging.

This choice increases the size and complexity of your supply chain network, adding to the work it takes just to collect, much less process, the data from this assortment of partners. At the same time, because consumers are so spoilt by choice, they have extremely high service expectations. If your brand is not available for purchase by their preferred method, it can easily lead to a lost sale as they will simply choose another available brand instead.

- **Rising retailer pressures.** Faced with these same high expectations from consumers, brick-and-mortar and online retailers have become more demanding of their brand suppliers. They have implemented on-time in-full (OTIF) service level requirements, with associated penalties if they're not met so they can keep products in-stock. Squeezed by competition, they are reducing inventory and placing smaller, more frequent orders, passing on the cost of carrying inventory and less efficient shipments (that may also need to be expedited).

Retailers give preferential treatment to the brands who are able to meet these demands and provide data-driven recommendations to improve replenishment and grow sales. However, they remain inconsistent in sharing their valuable data. They continually tweak how they do so, making it hard for brands to keep up with continual changes, and try to use their power to squeeze every dollar they can out of suppliers.

Design principles

The Alloy platform architecture was designed with five overarching principles in mind. These principles were informed by the co-founders' experience building analogous platforms for the stringent financial services industry and continue to guide our design choices and product development moving forward.

Integrated

The platform unifies multiple data sources to provide a "single source of truth" across an organization's different functions and teams, and even across partners, to eliminate common communication issues. As an end-to-end solution, it offers a single platform where teams can collaborate and complete their workflow, from insight to execution.

Flexible

Alloy recognizes the industry environment is continually evolving, and the platform must easily adapt to changing needs. It is flexible enough that existing functionality can be customized to a specific customer's requirements through a simple user interface, as well as for new product modules and not-yet imagined



applications to be added to the platform. Furthermore, while Alloy's current focus is integrating distribution and retail data for consumer goods manufacturers, the model is designed to work for upstream supplier and manufacturing relationships as well.

Scalable

As a modern platform, Alloy is built to maintain the same high levels of performance even as the volume of data processed, number of trading partners, variety of internal users and complexity of the computations multiplies. Alloy can scale up to meet the requirements of global enterprises, but scalability also means scaling to however big or small your particular business or team is, so you have a solution that is right-sized for your uses.

Open

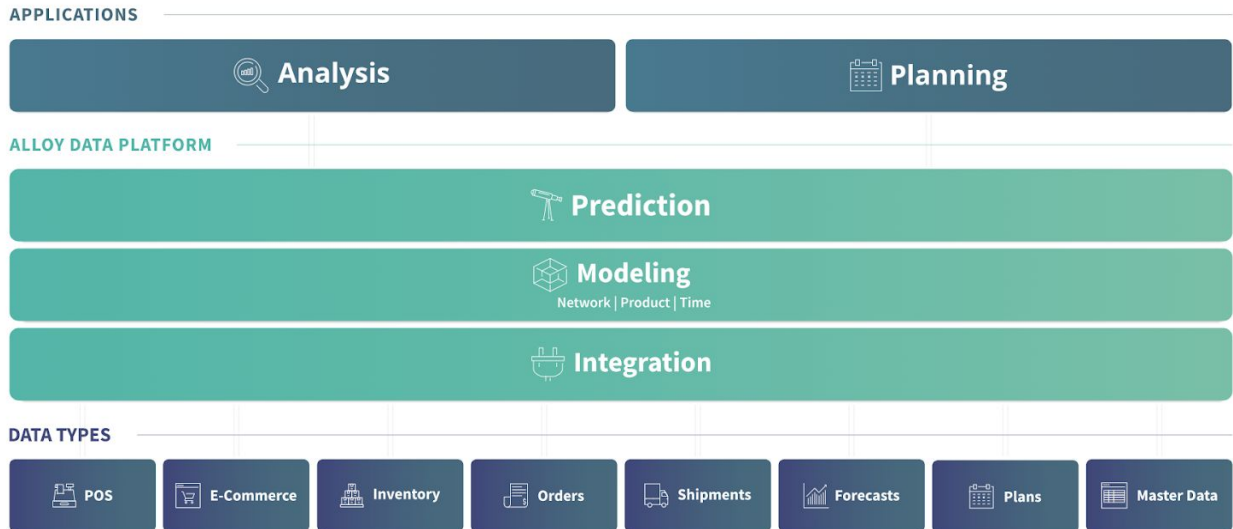
Alloy is an open platform that aims to democratize data. We offer applications that enable it to serve as an end-to-end solution, but also enable users and systems to export the data for use in third-party tools, such as planning software, or be made available in a central data lake. Within Alloy applications, such as forecasting, we leverage open source software and provide transparency into our approaches.

Actionable

Ultimately, Alloy aims to empower users to take action based on the data in the platform. That means it must be user-friendly and accessible to everyone, not just expert analysts, as well as predictive, so teams can influence what will happen in the future, not just retroactively report on what happened. It also means it requires minimal resources for setup and maintenance to deliver fast time-to-value and enable teams strapped for time and IT bandwidth too.

Platform Overview

The Alloy platform architecture consists of four key layers, which take data from its raw form, translate and map it to model your supply chain, use it to generate forecasts and business metrics and facilitate day-to-day workflows for demand-driven decisions.



Integration

The foundational layer performs data integration, automatically extracting relevant data to and from a breadth of sources already available to an organization, both external and internal. Alloy currently supports a wide variety of data types, such as sales, inventory and forecasts, to enable use cases for Supply Chain, Sales and Marketing at consumer goods manufacturers.

Modeling

The next layer is the “magic” behind the platform. Incorporating sophisticated data modeling and artificial intelligence, it processes the volumes of inconsistent data and transforms it into information that can actually be understood and used by business teams.

The data is modeled along three dimensions:

- **Product.** Translation across different product identifiers, units of measure and other products attributes used by different data sources.



- **Network.** Mapping of how products flow across your supply chain, including locations, shipment lanes and lead times.
- **Time.** Conversion between different fiscal calendars, time granularities and points in time.

Event modeling, the identification of special events impacting specific SKUs, locations and/or time periods, is currently available in beta mode.

Prediction

Prediction uses all the rich information that has been integrated and modeled to generate forecasts and simulate business metrics into the future. It combines proven forecasting models and business logic with machine learning algorithms to help users understand forward-looking states of the business.

Forecasts and metrics can be served at any level of granularity, according to a user's product, network and time preferences at the moment. This capability enables a key feature of Alloy that differentiates it from static reporting tools: real-time dashboards that can be customized on the fly.

Insight and action

The final piece of the platform architecture is the frontend where users can see and interact with Alloy, powered by the underlying Data Platform. It is made up of two applications today: Analysis and Planning. They are each designed with specific functionality and best practices to meet different purposes, but because they are all part of an integrated platform, provide a connected, seamless experience as you move between applications.

Identity and access management

The platform provides a password-based authentication mechanism that enforces password strength rules in accordance with current industry best practices. It also supports secure single sign-on (SAML SSO), allowing customers to delegate authentication to their own identity provider and requirements, including two-factor authentication, password strength rules and password expiration.

User roles

Each user is assigned a user role with a corresponding set of permissions.



- **Admin:** Can edit other users, assign admins and permissions. Has access to admin tools to define organization defaults for location attributes, product attributes, transaction event attributes and analysis settings.
- **Limited User:** Cannot edit other users and does not have access to admin tools. Can view existing dashboards and create new dashboards.

A user's role can be changed at any time by someone with the appropriate permissions.

Data permissions

A user can be given access to all data for the customer, or have access limited to specific data only, filtered based on any attribute or combination of attributes in Alloy. Data permissions are often used to restrict users to only the relevant data for their domain but not beyond, such as their particular:

- Retailer(s)
- Location type(s)
- Location tier(s)
- Geography(s)

Dashboard permissions

The applications in the **Insight and Action** layer of the platform architecture are made up of dashboards for users to interact with them. The creator of each dashboard can choose to make it public (accessible to everyone) or private (accessible to only themselves or other users they share it with).

Best-in-class Software as a Service (SaaS)

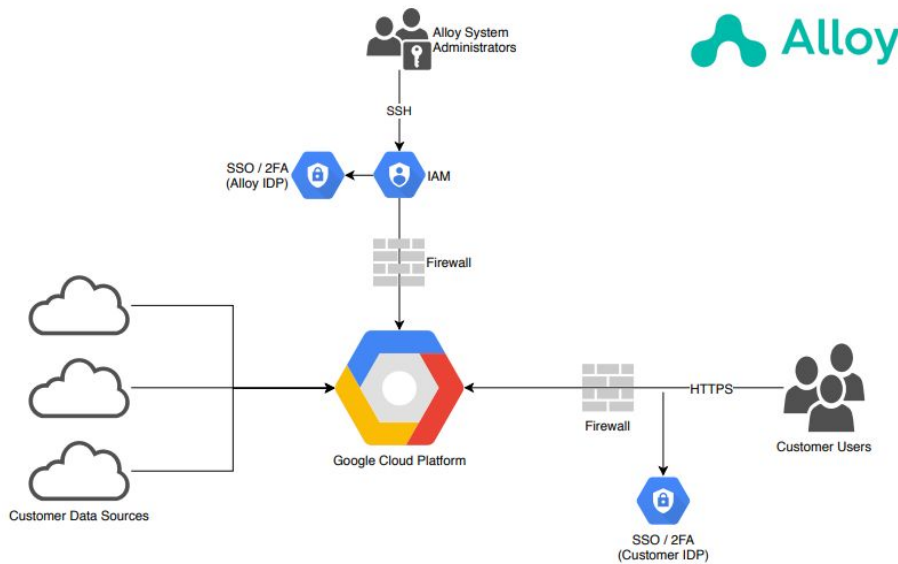
Alloy is hosted in the cloud and made available to customers on-demand via a desktop and mobile-friendly Internet browser application. The SaaS model enables us to offer a product that can be quickly implemented without IT configuration or resources, that scales and stays up-to-date across your organization and that delivers ongoing feature improvements, bug fixes and upgrades.



Google Cloud Platform

Alloy systems and services operate on Google Cloud Platform (GCP). We do not operate any on premise infrastructure. Most of our infrastructure is based in the us-central1 region located in Iowa, USA. We may utilize other GCP regions for high availability and disaster recovery purposes.

Google Cloud Platform is a major cloud provider with a strong reputation for security, scalability and reliability. Please refer to the [Google security whitepaper](#) for more information.



Data storage and transmission

Data is stored in Google BigQuery and [encrypted at rest by default](#). The multi-cloud data warehouse both seamlessly scales to the data volumes of large enterprises and enables real-time interactive applications that deliver value faster than traditional data warehouses.

Data in transit and requests to Alloy services are encrypted with TLS. It is not possible to use Alloy services over an insecure HTTP connection. HTTPS is enforced via the Strict-Transport-Security header and by directing all traffic to TLS-secured endpoints. We have an [A+ rating from SSL Labs](#).



Environment isolation

Alloy's servers are isolated from each other and the internet at the network level. A firewall is in place, and only those ports required to provide and operate the service are open. Within the network, individual services are only permitted to communicate with other services if there is a need to do so. For example, the service that processes incoming data is isolated at the network level from the service that hosts the frontend application.

Production, staging and development are hosted in completely separate Google Cloud projects, on separate networks.

Tenant isolation

Alloy is a cloud-based multi-tenant application. All customers share a single cloud environment. Logical separation between tenants is ensured by associating all data with a tenant ID, scoping all API access to a single tenant and restricting data queries to a single tenant at a low level.

SOC2 compliance

Alloy takes the protection and security of customer data very seriously. We are in the process of obtaining SOC 2 certification, which independently verifies our security, privacy and compliance controls adhere to global standards set by the [AICPA](#) (American Institute of Certified Public Accountants).

This rest of this white paper will discuss each layer of the platform architecture in greater detail.

 Integration

The integration layer automates what for many companies is a manual, time-consuming process of downloading data from different supply chain partners and adding it to a “master” spreadsheet on a weekly basis. Alloy not only frees up resources from this repetitive task, but also improves the recency, granularity, integrity, accessibility and security of the collected data.

Using Robotic Process Automation (RPA) and other automation techniques, Alloy extracts the data:

- As frequently as it is updated (up to hourly)
- At the most granular level available (down to the SKU-store-hour)
- For every user, regardless of functional silo (with permissioning controls)
- For storage and backup on highly secure servers (see [Data storage and transmission](#))
- With built-in failure recovery, data cleansing and data validation

Pre-built connections

Alloy can extract the original data from your primary sources, including retailers, e-commerce platforms, distributors, third-party logistics providers (3PLs) and your internal systems. We have established a network with hundreds of integrations into these partners for near real-time access to multiple data types, including sell-in, inventory and forecasts.

As a result, integration for a new customer can be “turned on” as soon as a customer enables access, with data securely flowing in as little as a few hours, depending on the trading partner. Similarly, existing customers can easily add new data sources as they grow (fee associated for additional data feeds).

Data access and compliance is based on a customer’s existing partner relationships and data sharing agreements. If customer credentials are required, such as for a partner portal, Alloy provides a secure link to input the credentials, which are encrypted and isolated at the network level so they can only be accessed by the integration service.

Each partner may provide data in one or multiple formats, all of which are supported by Alloy and considered distinct “data feeds.” The table below highlights the most common formats for different



partner types, with specific partner examples. For details on whether your partners are supported and implementation timeframes, please reach out to your Account Executive.

Table 1: Integration Formats

Data Source	Retailers	E-commerce	Distributors	3PLs	Internal
Partner portal	✓ e.g., Target VendorIQ Dashboard	✓ e.g., Amazon Retail Analytics	✓ e.g., Cin7	n/a	✓ e.g., data lake, ERP
Direct integration (API)	n/a	✓ e.g., Shopify API	n/a	n/a	✓
EDI	✓ e.g., EDI 852	n/a	✓ e.g., EDI 856, ASN	✓	n/a
Emailed report	✓ e.g., Sephora Best Seller Report	✓	✓	n/a	✓
CSV spreadsheet or TXT file	✓ (upon request)	✓ (upon request)	✓ (upon request)	✓ (upon request)	✓, e.g., ERP extract, sales targets

Historical data

In addition to extracting the most up-to-date data from your supply chain partners, Alloy extracts and stores historical data for future reference. It enables the automatic calculation of common metrics, such as year-over-year (YoY) changes in sales, that could otherwise require the maintenance of very large spreadsheets or finding and cross-referencing spreadsheets from multiple points in time.

During implementation, Alloy will conduct a one-time load of all historical data. It may come from a partner portal, if it provides historical data, or other sources such as spreadsheets, internal systems and email reports.

On a going forward basis, Alloy stores and retains all extracted data. We do not delete data unless our relationship with a customer ends, or a customer requests that we delete data. Alloy maintains daily

backups, which are encrypted at rest and retained for 30 days. Backups are scheduled to be obtained at low traffic times, to minimize impact on production services.

Failure recovery

The data pipeline is designed to automatically recover from failures by re-running until the data extraction process is complete. This process prevents any gaps in the data due to automation issues.

A common reason for failure is when a partner makes a change to the way they provide the data, such as the data format or metric name. Alloy immediately flags these changes and updates the integration tools to accommodate the new requirements. Any data gaps that may have occurred are filled in, and users can continue accessing their data without issue.

On occasion, fixing the failure requires input by a customer. In these cases, a member of the Alloy team will reach out to the right person and the data status will be updated on the [Data](#) application.

Data cleansing and validation

The integration ensures only clean, high quality data is brought into Alloy to deliver on the value of the applications in the Insight and Action layer.

Alloy cleans the data by detecting and excluding bad data from the platform. As data is extracted, it is placed into the **Alloy format**, and any data that does not fit within this strict model is automatically flagged, and if appropriate, removed. There are also checks for duplicate data, conflicting data and outliers, which are also flagged and resolved.

Furthermore, Alloy validates the following aspects of the incoming data:

- Headers
- As of date
- Time period and day of week
- Product identifier format
- Partner identifier



- Postal code
- Value features (e.g., positive, non-negative, Boolean), depending on the metric

As part of the process, Alloy may identify new attributes, such as a new product, and bring them to a customer's attention for validation. It ensures the new attribute will be integrated in the data model for automatic extraction in the future and prevents data gaps.

Data export

Data in Alloy can be exported in CSV or Excel format, at any time with no limits. Exports can occur via browser download, email attachment, or automated SFTP transfer to individuals or systems.

 Modeling

After integration extracts data in its raw form, the modeling layer of the Alloy platform transforms it to create a usable set of information. This process is a critical but oft overlooked step to ensure the data provides value and is actionable for business teams.

Each data point in Alloy is associated with a specific combination of product, network and time attributes, which are related to each other in a specific way. In some cases, these attributes and relationships are provided in the data, but in partner-specific terminology, such as a partner product ID. In other cases, the associated information is not explicitly provided, but can be intelligently inferred. Alloy can handle both of these types of data and associate the correct attributes and relationships to each data point based on **master data**.

When incoming data is unrecognized and can not be matched to the master data, such as a new product, Alloy identifies these exceptions and creates suggested values for the attributes based on known values for related products. The master data is updated, so Alloy can automatically model all future data. Users can also request updates to the master data to correct suggested values, reflect changes, or include planned additions.

Product

Alloy creates and maintains a unique Product Master for each customer, which typically includes the:

- Product SKU
- Product UPC
- Product name and description
- Product hierarchy, e.g., category / sub-category / family
- Product cost
- Product color, size, scent, flavor, style, season or any other defining attributes



Using this Product Master, Alloy can match partner data to internal SKUs based on the associated product attributes. Customers are not required to provide or maintain a master with their product SKU and each partner's unique product identifier, such as an Amazon ASIN.

Unit of Measure

In addition, Alloy maintains two types of Unit of Measure Masters to model product units. One is based on **standard conversion rates**, like:

- Grams to kilograms
- Kilograms to pounds
- Gallons to fluid ounces

These conversions are shared and consistent across all customers, and facilitates product planning based on sell-through.

The other Unit of Measure Master is unique to each customer and is based on **how a product is packaged** (or re-packaged) by the brand and its partners, such as:

- Eaches to packs / kits
- Cases to pallets

This type of unit of measure conversion is necessary to translate data coming from different tiers of the supply chain (e.g., brand warehouse, distributor, retailer) that each provide data in different units. It also ensures nuances like saleable inventory and order minimums are taken into account in subsequent calculations.

Network

Alloy creates and maintains a unique Network Map for each customer, which includes details on each location and the relationship between locations across plants, warehouses and distribution centers managed internally or by distributors, 3/4PLs or retailers. It typically identifies:

- Location type



- Location street address, city, state, zip code, Designated Market Area (DMA) and country
- Shipment lanes between locations, by product
- Lead times between locations, by product
- Location managers, territory, inventory targets or any other defining attributes

The Network Map is inferred and kept up-to-date based on actual orders and shipments from existing data sources. Customers are not required to provide the map or notify Alloy if network changes are made.

Retailer network

Alloy also maintains a shared Location Master for each retailer, which typically includes:

- Store ID
- Store name
- Store street address, city, state, zip code, Designated Market Area (DMA) and country
- Store region, regional manager, district and district manager
- Distribution Center that replenishes the store
- Store type, size, opening date, tier or any other defining attributes
- Distribution Center city, state, zip code and country

This Location Master is consistent across all customers and enables teams to easily conduct geographic analysis or visualize regional trends and inventory movements without additional data management.

Time

Alloy maintains different types of calendars so dates can be translated into weeks or months accordingly.

Available calendars include:

- Partner-specific fiscal calendars, e.g., Walmart fiscal calendar
- Retail calendars, e.g., 4-5-4 calendar



- Customer fiscal calendar
- Western calendar (Gregorian calendar)
- Global holiday calendar

The calendars facilitate consistent year-over-year comparisons and communication with partners.

Versioning

For forward-looking data, Alloy models the relationship between projections made at different points in time, including those from:

- Partner-provided forecasts
- Internal forecasts and plans

This versioning enables users to track how a provided number or calculated metric changes over time and assess accuracy with different time lags.

Aggregation and disaggregation

All data are stored at the most granular level available across these three dimensions—ideally, the SKU-location-day level. Modeling can aggregate this data up to a higher level, such as the product category, region, or annual level, or even across the customer’s entire business. Some types of data can simply be added together when aggregated, such as *dollar sales*, but many require more thoughtful aggregations, as it does not make sense to sum them, such as *price*.

In other cases, the data is not available at the most granular level. For example, some retailers only provide sales at the weekly level, or forecasts are only created at a product category level. Modeling can also disaggregate this data down to a more granular level, breaking it down based on historical trends and similar places where more granular data is available.



Prediction

Once data has been integrated and modeled, this layer of the Alloy platform synthesizes multiple data points to make predictions or inferences. It automates the complex computations necessary to produce different forecasts, forward-looking metrics and other business indicators.

Because it operates in real-time, Prediction immediately takes into account any new data that is available in Alloy when preparing metrics. It also provides the flexibility for users to “change their mind” and look at the information differently by simply adjusting a setting on their own.

Simulated metrics

Alloy simulates inventory levels for every item-location combination throughout your network and across your different partners. The *projected inventory on hand* is then combined with *forecasted sales* to calculate predictive measures and evaluate potential risks or opportunities, such as:

- *Projected Weeks of Supply (WoS)*, or how long *projected inventory on hand* is expected to last given *forecasted sales*
- *Lost sales \$*, or how much in retail sales could be lost unless *projected out-of-stocks* are fixed
- *Spoils \$*, or the value of goods that would spoil unless *projected overstocks* are avoided
- *Phantom inventory locations*, where *actual sales* and *forecasted sales* may vary because reported *inventory on hand* is inaccurate

Alloy is **flexible** enough to allow users to change the assumptions used in calculating a metric to align with partner preferences or business changes. For example, each of the inputs to calculate *Projected WoS*, can be customized by the user - to specify different forecast models, definitions of available inventory and/or shipment times to be used.

Forecasting

There are three different kinds of forecasts available in Alloy: partner-generated forecasts, your company's internally-generated forecasts and Alloy-generated forecasts. The first two are imported into Alloy via the Integration layer, and Alloy-generated forecasts are created in the Prediction layer.



Alloy generates nine different types of unit sales forecasts, which primarily forecast sell-through, but can be applied to sell-in data as well. Users have the option to choose from these different models to forecast different segments of their business:

- Alloy: Alloy's forecast of unit sales, based on optimal model selection
- ARIMA: Alloy's forecast of unit sales using the ARIMA forecasting method
- ETS: Alloy's forecast of unit sales using the exponential moving average forecasting method
- FFNN: Alloy's forecast of unit sales using a feed-forward neural network (FFNN)
- GAM: Alloy's forecast of unit sales using a Generalized Additive Model (GAM) including trend, seasonality and holiday additive components
- SMA: Alloy's forecast of unit sales using a Simple Moving Average (SMA)
- BSTS: Alloy's forecast of unit sales using a Bayesian Structural Time Series (BSTS) method including trend, seasonality and autoregressive components
- TBATS: The TBATS is an ETS method used for forecasting by combining (B) Box-Cox Transformation, (A) ARMA errors, (T) Trend and (S) Seasonality. The second T prefix is a modification for trigonometric seasonality, selecting Fourier components for more complex seasonal patterns
- Ensemble: The methodologies above are averaged to create a forecast that includes aspects from the component forecasts

These forecasts can be tracked against actual sales in Analysis or used for as a foundation for shipment planning in the Plan application.

Additionally, they are used to calculate a variety of simulated metrics.

Business indicators

Alloy automatically calculates dozens of other metrics that can be used as indicators predicting the direction of the business. The metrics are **specific to the consumer goods industry**, and the way brands and their supply chain partners evaluate performance and plan for the future. The complete list has



grown over time and continues to grow with the input of customers and industry experts, and as new data sources and types have unlocked new metrics.

As a result, the metrics reflect **best practices** from a broad cross-section of the industry on what metrics matter most and how each one should be calculated. For example, two different retailers may both provide a metric with the same name, such as *out-of-stock percentage (%)*, but calculate it differently. Alloy uses the provided data to calculate the metric in a consistent way for easy comparisons across retailers (while preserving the retailer provided metric to facilitate conversations with them) and to reflect learnings from across Alloy customers on the most accurate way to calculate it.

Some of the most commonly used business indicators include:

- *Year over year sales change*
- *Plan achievement*
- *Comp store sales, inferred based on actively selling stores*

Insight and Action

The Alloy Data Platform, made up of the integration, modeling and prediction layers, serves as a single source of truth powering business applications for insight and action. Alloy offers two connected applications today, focused on descriptive and predictive analytics and demand forecasting. As an open platform, Alloy can also power third party applications, such as planning tools.

Analysis

The Analysis application provides “dashboards” that users can customize with the data and metrics they want to see, in the way they want to see it, to identify business insights. These dashboards can be created by editing Alloy-provided templates for common types of analysis, editing existing dashboards, or entirely from scratch. There is no limit to the number of dashboards in Alloy, so users can create different dashboards for different purposes, such as:

- High-level analysis of the latest business performance and priorities to be addressed
- Root cause analysis of the source of identified issues and likely reason(s) for the problem
- Segmentation analysis of what actions and characteristics drive demand for different products and geographies
- ROI analysis of the performance of sales and marketing investments and experiments
- Supply chain analysis of inventory needs at each node of the network to meet downstream demand

Users can choose from many **visualization** “widget” options — tables, line charts, bar charts, pie charts, heat maps, location maps, timelines and calendars — and include any number of widgets on a dashboard. Furthermore, dashboards can be linked together to create “workflows,” so users can move seamlessly from one dashboard to the next and keep answering the next question.

Additional key features include the ability to drill down within a dashboard to show more granular data, dynamically filter entire dashboards or specific widgets based on any attribute, customize how charts are organized and designate a dashboard as a “favorite” for easy home page access.

Planning

The Planning application is a tool for users to customize the Alloy-generated POS forecasts and produce shipment forecasts to support demand planning. Users can customize both the forecast inputs themselves and make manual overrides, adjusting for:

- Inventory targets, such as safety stock levels
- Current conditions, such as lead times
- Future events, such as planned promotions or a competitive product launch

A user starts by choosing a forecast at the level of granularity of their choosing (by SKU, partner, or even location combination). Then,

- Alloy identifies the best forecasting method for that slice of the business based on historical accuracy
- Teams add any promotional events or manual adjustments to align forecasted sales to targets
- Alloy aggregates smaller forecasts into an overarching POS forecast and combines it with current inventory levels, lead times and inventory targets for each downstream partner to produce a Shipment Forecast.
- The Shipment Forecast can be disaggregated to recommend both timing and volumes for shipments to specific partners

Alloy provides complete **transparency** into how a forecast was selected and how forecasts were revised over time, so users have the necessary background when making changes.

Marketing events and promotions

Users can document past and upcoming marketing events, including the impacted products and/or locations, the measured or projected impact on demand and the cost, in addition to the inferred promotions based on historical data. All these events are taken into account when generating forecasts.



Conclusion

Alloy is an integrated, flexible, scalable, open and actionable platform that keeps up with modern requirements for performance, security and user value. We recognize that enabling the demand-driven supply chain starts with a robust data foundation and have developed deep expertise in extracting, translating and mapping raw data.

Our engineering works tirelessly to improve the platform by integrating more data sources, building additional applications while enhancing existing ones and leveraging the latest techniques and technologies to deliver more benefits to our customers. The Alloy platform is constantly updated, so please don't hesitate to reach out to us for the latest details as no document can stay up-to-date with our frequent releases.

About Alloy

Alloy equips consumer goods brands to capture demand and streamline supply by breaking down silos within and across supply chain, sales and marketing teams. Our cloud platform couples end-to-end inventory and shipment visibility with sell-in and sell-through analytics to surface relevant insights in real-time. Team members and leaders can proactively address risks and maximize opportunities as they monitor demand trends, evaluate performance against forecasts and drive collaboration internally and with partners to adjust execution and plans — down to the SKU/store/day level.

Companies ranging from the Fortune 500 to digital-natives use Alloy every day to grow customer relationships and improve service levels. While preventing lost revenue from out-of-stocks and phantom inventory, they also lower the costs of inventory holding, expedited freight, spoilage and markdown by efficiently responding to what and where consumers buy.

For more information, please visit www.alloy.ai or contact us at inquiries@alloy.ai.