

SAS® 360 Match

*Projecting Inventory*

*Updated June 23, 2020*



# Contents

<b>Forecasting Inventory .....</b>	<b>1</b>
<b>Overview .....</b>	<b>1</b>
<b>How Simulation Works.....</b>	<b>1</b>
<b>Setting Up and Optimizing Inventory .....</b>	<b>2</b>
Inventory Forecasting Options .....	2
Intra-Day or Real-Time Inventory.....	3
<b>Common Terms for Simulations .....</b>	<b>3</b>
<b>Checking for Available Inventory.....</b>	<b>4</b>
Inventory Queries.....	4
Contending Flights .....	5
Using the Availability Checking Tool .....	6
Estimating Tags.....	6
Estimating Frequency Caps.....	7
Tier Order .....	7
Out-of-Sync Contending Flights .....	8
Estimating Contending Flights .....	8
<b>Common Terms for Managing Inventory .....</b>	<b>8</b>
<b>Inventory Seeds.....</b>	<b>9</b>
<b>Use Cases .....</b>	<b>10</b>
Forecasting Optimization .....	10
Global Inventory Adjustments .....	11
<b>Constraints of Simulations .....</b>	<b>14</b>
Target Rules Without Queries or Estimates .....	15
Uneven Inventory Distribution and Pacing Rules .....	15
Excluding Other Business Rules.....	15
<b>Frequently Asked Questions .....</b>	<b>16</b>

# Forecasting Inventory

SAS® 360 Match enables you to forecast inventory for your customers. This document outlines the processes that are used to project inventory, their limitations, and how to fine-tune the process.

We start with a detailed introduction to the core simulation process, followed by an explanation of how the forecast is calculated and made available through the user interface. We also detail the options that are available to tweak your simulation and how solutions work in different settings.

## Overview

**Simulation** is the process that SAS® 360 Match uses to generate data to calculate capacity, availability, and overlap for a specified set of keys and values (also known as targeting or products). The simulation process provides you with the basis to forecast inventory.

SAS® 360 Match can simulate future ad traffic based on your historical request data. SAS® 360 Match creates an output table that includes all the ad-serving data for a defined time range. After the inventory table is generated, you can access and use the data for forecasting in ways that are described in the Use Case section.

Figure 1 shows the lookback window for a product and highlights some elements in the results.

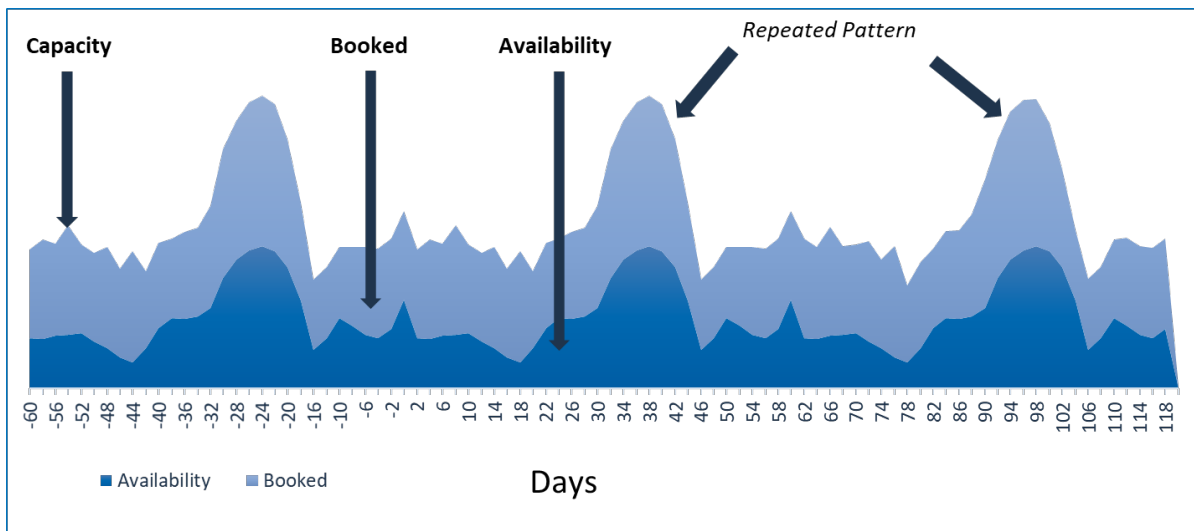


Figure 1. The simulation process for SAS 360 Match uses a product's lookback window to forecast inventory.

## How Simulation Works

A simulation process engages the business rules that affect ad delivery, including pacing flights to meet users' goals, imposing frequency caps, checking targeting, and enforcing tier options. Simulations are run overnight with the same process as serving live ads so that the results are as accurate as possible.

A simulation uses previous ad-serving logs to synthesize the inventory (daily impression volumes and tag values) for the forecast. This is the basis or lookback period.

The logs for the first basis day are used to simulate traffic for the first forecast day, the second basis day for the second forecast day, and so on. The forecast period can be longer than the basis period, in which case the basis is reused, as needed.

Each day of the basis is projected onto a corresponding day in the forecast. For example, if a basis day serves one million impressions, the corresponding forecast day that uses the basis day will also serve one million impressions.

The size of the basis period should be a multiple of 7 days, so that each basis day is projected onto the same day of the week in the forecast. Fluctuations in your inventory during the basis period can affect the accuracy of your forecast. Consider how your inventory changes when you select your base period. The longer the basis period the longer the simulation will take to run, but the more sensitive the forecast will be to seasonal variations.

The output of the simulation is similar to the data that is produced when serving live ads. A `future_bigsum` table summarizes the daily count of impressions for each creative and for each combination of tag values being logged. You can generate BI reports from this table or use it as the basis for the contending flights option when you have checked for inventory availability. This table is only updated once a day because simulation is run once a day. New flights and changes to flights made during the day do not affect the data in this table immediately, but only after the next simulation is run. SAS® 360 Match provides an intra-day and real-time update option that is explained in the section on Intra-Day or Real-Time Inventory.

Immediately after the simulation ends, a secondary process creates a smaller inventory table, derived from the data in the `future_bigsum` table. The inventory table contains the daily available and reserved impression counts for each of the tag combinations in `future_bigsum`. Impressions that were recorded for creatives in flights in preemptible tiers, including defaults, are regarded as available. All other impressions are regarded as reserved. The table does not contain creative or flight information.

The inventory table cannot be used directly for reports but serves as the basis for inventory forecasting queries. By default, the data in this table is not affected by new flights and edits to flights that are made during the day. Contact SAS Technical Support if you want any changes to be reflected immediately in the inventory table. However, given the nature of this table and the complexity of flight business rules, adjustments made to this table are nominal. You will receive more definitive and accurate results from the next simulation.

## Setting Up and Optimizing Inventory

Contact SAS Technical Support at [support@sas.com](mailto:support@sas.com) for assistance with migration and to configure your lookback window, which is also called **basis** in this document.

### Inventory Forecasting Options

Some configuration options can affect the simulation behavior, such as averaging impression volumes across days or performing linear regression to follow volume trends.

You can select a method for an inventory forecast using the BI Report feature. Contact SAS Technical Support to enable the specific configuration setting for your simulations.

The available simulation forecast methods are the following:

- **Default:** Each day of the forecast uses samples from the same basis day, where the first forecast day uses the number of impressions from the first basis day, the second forecast day uses the impressions from the second basis day, and so on.
- **Average:** Every day of the forecast uses the average daily number of impressions across the basis period.
- **Day-of-week average:** Each day of the forecast uses the same number of impressions as the basis average for the same day of the week. For example, the average of all basis Mondays is used for every forecast Monday, the basis Tuesday average is used for every forecast Tuesday, and so on.
- **Linear regression:** Each day of the forecast uses the projected number of impressions that is calculated using a linear regression of the daily impressions across the entire basis period.
- **Day-of-week linear regression:** Each day of the forecast uses the projected number of impressions that is calculated using a linear regression of the daily impressions for the same day of week of the basis. For example, a linear regression of all basis Mondays is used for every forecast Monday, a linear regression of all basis Tuesdays is used for every forecast Tuesday, and so on.

### Intra-Day or Real-Time Inventory

SAS® 360 Match also provides another backend process that can run in real time (during the day) to keep the available numbers up to date until the simulation can be run again the next night. This feature is optional and turned off by default. Contact SAS Technical Support to enable this feature.

For example, a simulation runs overnight and creates data that states 1 million of 2 million impressions for product X are available. That data is used the next morning to create a proposal and to check the inventory. If the proposal is created for 250,000 impressions and the probability set to 90 percent, the system will decrement the avails for product X appropriately. When another person creates a proposal for 800,000 impressions 45 minutes later, SAS® 360 Match will report that there are only 750,000 impressions available for that product.

The capacity fluctuates slightly each time the simulation is run because a slightly different set of historical dates is used each time. Specifically, capacity for product X might be reported as 2,000,345 today and may result in 1,999,896 the next day. In most cases, the differences should be small enough not to affect the general sense of what is available.

### Common Terms for Simulations

Here are some common terms that are used when discussing simulations.

**Availability:** The sum of all the impressions that are served during a simulation for a specified set of keys and values from tiers that are marked preemptible.

**Booked:** The sum of impressions that are served during the simulation from all tiers from currently live flights whose targeting matches a specified set of keys and values.

**Capacity:** The sum of all the impressions that are served during a simulation for a specified set of keys and values from all tiers.

**Delivery Patience:** Controls how much patience the system should have in pacing a flight to get on schedule. In this case, on schedule means that the flight is meeting its on-schedule priority in serving the flight and in the simulation. For more information, see [Add a Flight to a Campaign](#) in the SAS® 360 Match User's Guide.

**Forecasting:** The act of using the SAS® 360 Match inventory forecasting tools. This term is used interchangeably with inventory and simulation.

**Inventory:** The collection of tools and data that are used to forecast impression volume for specified sets of keys and values. Frequently used interchangeably with forecasting and simulation.

**Key Value Pairs:** An item of data that is identified by an arbitrary name. The key is the name and the value is the content. For example, the key is AGE, and the value is 25.

**Overlap:** The degree to which the targeting logic from one target or product allows impressions to serve to the same combination of keys and values that another target or product allows.

**Package:** A collection of products.

**Product:** A SAS® 360 Match system object that is a combination of a target and information about how much those impressions should cost an advertiser.

**Target:** A combination of key and value pairs and Boolean logic that defines a reusable rule that can be applied to SAS® 360 Match system objects such as products, flights, and creatives. This rule limits the places (combinations of keys and values) that the object can serve and count as ad impressions.

## Checking for Available Inventory

SAS® 360 Match provides a function to check inventory availability that estimates available and reserved impression inventory for some future period.

### Inventory Queries

The availability function is available from three places in the SAS® 360 Match interface:

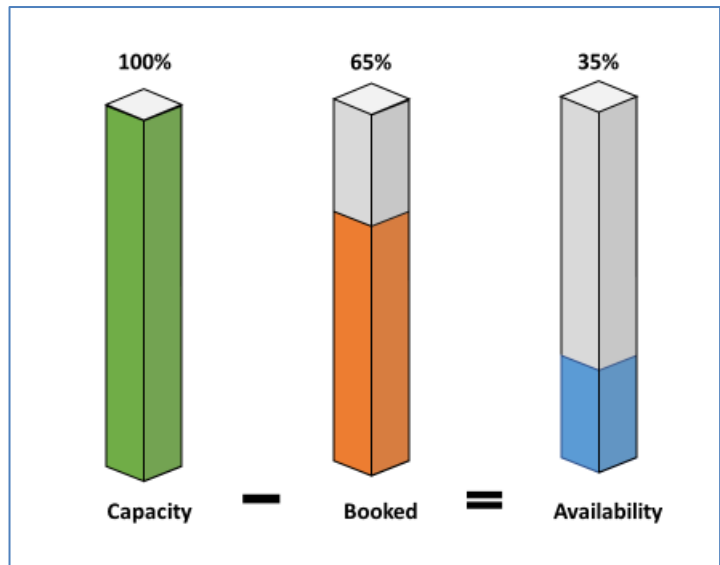


Figure 2. The relationship among availability, booked, and capacity.

Use **Sales > Check Avails** to check availability for an arbitrary date range and targeting. Available, reserved, and total inventory are shown. You can choose to break down the inventory view by day, week, or month. Use this option when you have a specific question, such as "What are the capacity and availability numbers for product X?" Create dummy line items and use the **Check Availability** button to receive a simple number to answer your question.

Use the **Available** button in the Inventory panel on the edit screens for flights and line items to determine if there are enough impressions available for the line item or flight you are currently creating. The check is based on your start and end dates, along with any targeting options that you specify.

- The **Available** button in the **Inventory** section on individual flights from **Traffic > Flights** checks availability for the flight's date range and targeting. In the results, you can view available, reserved, and total inventory for the date range. You can also see the projected inventory that the simulation anticipates this flight will consume. Served is how much inventory this flight has already served prior to the forecast period.
- The **Available** button in the **Inventory** section on individual line items from **Sales > Proposals** checks availability for the line item's date range and targeting. The results shown are like those on the Flight page.

## Contending Flights

You can check the **Contending Flights** option for all availability functions. When **Contending Flights** is checked, you receive a report of all flights in non-preemptible tiers that are projected to consume the same inventory. This information can be used to identify flights whose targeting could be changed to allow more of the queried inventory to be available.

The reported contending flights numbers do not necessarily represent the total inventory that will be consumed by those flights, but only the portion that overlaps with the inventory that is being checked.

For example, you need additional impressions for product pA. Some product pA is available but there is not enough to satisfy the advertiser's request. Product pC has significant overlap with product pA and therefore has contending flights.

Use the availability tool with the **Contending Flights** check box selected to determine if there are contending flights in pC that might be reallocated to free up inventory for pA.

To free up inventory for pA, you can make any of the following changes to the contending flights in pC, although each change may prevent the flights in pC from reaching their required impression goals:

- Move the flights to a lower tier. However, flights that are relocated to a lower tier will be evaluated less often.
- Add targeting rules to the flights to remove the overlap with pA. However, flights from pC can serve to fewer places.
- Apply frequency caps to the contending flights. However, flights can no longer serve as many times per unique user.

## Using the Availability Checking Tool

Here are some issues to consider when using the availability function.

### Estimating Tags

When an availability query is run, the date range and targeting are used to form a query of the table, as described above. When all the tags referenced in the targeting are in the simulation output tag set, a query is formed that directly produces results from data in the table.

But when the targeting includes a tag that is not in the simulation output tag set, the tag's impact on the availability results must be estimated. The following example illustrates the method to use and why it is used.

Consider targeting for `SITE = NEWS AND GENDER = MALE.`, where `SITE` is in the simulation output tag set, but `GENDER` is not. Since `SITE` is in the simulation tag set, the availability function directly queries the inventory table for the date range, yielding 10,000 available and 20,000 reserved impressions for a total of 30,000 impressions. Because the target also includes `GENDER = MALE`, the available inventory for that part of the target must be less than 10,000. `GENDER` is not in the simulation tag set, so its impact must be estimated by multiplying 10,000 by some fraction.

In this case, the availability function determines the fraction by considering what percentage of all impressions has historically been served while the `GENDER` tag was `MALE`. Such tag value frequencies are calculated once a day, based on the last seven days' worth of serving logs.

Assume that 48 percent of all impressions are served when `GENDER = MALE`. After multiplying the inventory numbers from the query with 0.48, the availability function shows that 4,800 impressions are available ( $10,000 * 0.48$ ), 9,600 are reserved ( $20,000 * 0.48$ ), and the total is 14,400 ( $30,000 * 0.48$ ).

Using `OR` targeting and other more complicated targeting can result in several frequency lookups that are logically combined, according to the target construction, to produce the fractions that are needed for the query. The fractions are applied to the availability numbers that can be directly queried, like in the above example, to yield the estimated available.

The estimation technique used by the availability function is accurate when the tags being estimated are independent of each other and independent of the tags that can be queried directly. However, the results are less accurate when the estimated tags are not independent. The following example illustrates this problem.

Using the example above, assume that the actual frequency of `GENDER = MALE` when `SITE = NEWS` is 47 percent, while the overall frequency of `GENDER = MALE` is 48 percent, as mentioned above. The 48 percent estimate is very close to the actual frequency of 47 percent. The numbers from the availability check are only about 2 percent higher, which is very close to the actual result if both `SITE` and `GENDER` were in the simulation tag set. In this case, the `GENDER` tag is highly independent of `SITE = NEWS`.

But consider a target that contains `SITE = SPORTS` and `GENDER = MALE`. Assume the `SPORTS` site skews `MALE` and that the actual frequency is 72 percent. The 48 percent estimate that the availability function uses yields figures that are probably 50 percent too low. Because



there is a higher correlation between `SITE = SPORTS` and `GENDER = MALE`, the `GENDER` tag is relatively dependent on `SITE = SPORTS`.

When possible, dependent tags should be favored for inclusion in the simulation tag set. Independent tags should be excluded, even if only to make way for more dependent tags to be included, because the number of such tags must be limited.

## Estimating Frequency Caps

The availability function applies a heuristic to estimate how a frequency cap affects availability. A frequency cap is a rule that limits the number of times a user can see a flight in a specific time period.

Generally, a more restrictive a target will be experienced by fewer visitors. Therefore, the function first calculates the fraction of all impressions that match the targeting. It multiplies that number with the estimated total number of unique visitors across the check period, yielding an estimate of unique visitors who will experience the matching target.

Finally, it multiplies that number with the number of impressions that are expressed in the frequency cap and the size of the period, yielding an estimate of the maximum inventory that could be consumed under the cap.

The availability function compares the max capped inventory estimate to the uncapped availability number. This is the availability number it would otherwise have returned if no cap was expressed, as described above. The function reports the lower number as available.

Consider an example checking availability of `SITE = NEWS` with a frequency cap of two per day, for a 7-day period:

`SITE = NEWS` matches about 40 percent of all impressions, and the total unique visitors for the period is estimated to be 100,000. The heuristic estimates that 40,000 unique visitors across the period ( $100,000 * 0.40$ ) will experience this target. The max capped inventory is then estimated to be 560,000 ( $40,000 * 2 * 7$ ).

If the uncapped available inventory is found to be 1,000,000, only 560,000 is reported as available because the frequency cap limits how much can be served. If the uncapped availability is 400,000, then 400,000 is reported.

This heuristic assumes a linear relationship between the targeting restriction and the number of visitors. That may be correct in some cases and not in others, so the resulting estimated availability may be accurate in some cases and not in others.

## Tier Order

The availability function regards all inventory consumed by flights in preemptible tiers as available and in non-preemptible tiers as reserved. When a preemptible tier is ranked ahead of a non-preemptible tier, availability checks for flights in the non-preemptible tier can produce numbers that are higher than what the tier may experience during the simulation or when serving live ads. This is because the inventory that was reported as available can end up being consumed by the preemptible tier, by virtue of its position ahead of the non-preemptible tier.

## Out-of-Sync Contending Flights

The future\_bigsum data used to produce contending flights is produced once a day by the simulation. Immediately after the inventory data is produced from future\_bigsum, these two tables are in sync and the numbers agree with each other.

During the day, as flights are created or modified, future\_bigsum cannot be updated. But if the real-time inventory option is enabled, the availability function adjusts the other inventory table's availability numbers. As the day progresses, this can cause the two tables to grow out of sync. The result can be that numbers from the contending flights reports don't make sense when compared with the reported inventory availability. These numbers converge again when the next day's simulation is run.

## Estimating Contending Flights

As with the availability numbers themselves, the availability function must apply the same tag estimation fractions to the queried flight serve counts. The same limitations of that estimation method apply to the contending flights numbers.

When viewed from the flight or line item screen, that flight or line item might also appear in the report with the same fractions applied. Because of the estimation fractions, the reported numbers may be lower than the actual projected number for the flight or line item. The projected number is the actual number that the simulation projects will serve.

A best practice is to ensure that all preemptible tiers are ranked below non-preemptible tiers. Business needs may require otherwise, but the availability function does not account for it.

## Common Terms for Managing Inventory

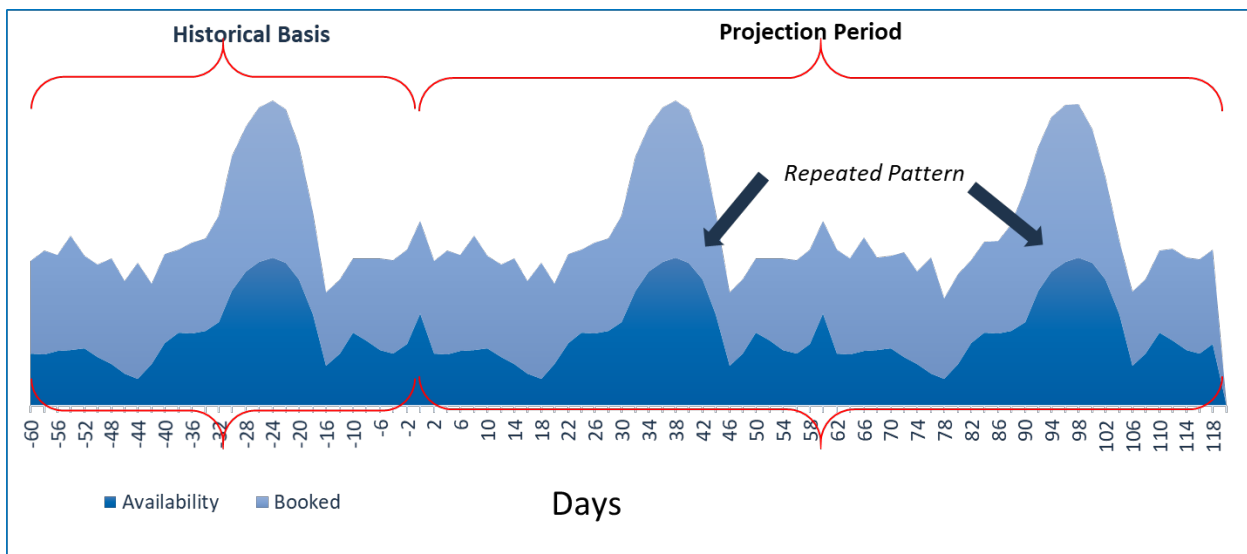


Figure 3: Inventory projection is calculated using the basis period.

Here are some common terms for settings and values to use when you configure a simulation.

## **Historical Basis**

This is the amount of time in the past, selected from historical request logs, that is used as future request data for the simulation. This data contains dates, cookie IDs, and TAG information. SAS® 360 Match generally selects a number between 28 and 84 days, which are multiples of 7, for this setting because the tags that are used in these requests are likely to be similar to the real requests that will come in tomorrow, next week, and next month. If the number goes back too far, the risk of the tag data being dissimilar and forcing the system to use estimations on some tags instead of actual data, increases.

## **Ignore Preemptible**

A setting that alerts the simulator to treat tiers set with the preemptible flag as normal and to only use system default ads served during simulation in calculating available impressions.

## **Preemptible Flag**

A tier setting that alerts the simulation process to treat ads served from this tier as available. If you have no tiers set as preemptible, only system default ads are considered available impressions. For example, you may have several tiers of ads that serve unpaid house banners or redirects to third-party monetization services. You generally do not want SAS® 360 Match to treat these ads as booked inventory because choosing these ads deprives your direct sales team access to these impressions when building proposals.

## **Projection Period/Forecast Length**

This is the amount of time in the future SAS® 360 Match attempts to predict using the request data from the sampling files (historical request data). For example, the Inventory\_Empty analyzer view in BI shows you data that looks like detailed delivery data for 30, 60, 180, or more days in the future. Keep in mind that the longer the simulation period, the longer it takes the process to run overnight.

## **Sampling Rate**

This is the number of requests from the sampling logs SAS® 360 Match skips for every request that is used. Using every request in the sampling file makes the simulation process take much longer and is not that much more accurate. The default setting of 53 requests can produce accurate forecasts. In testing, a sampling rate of 53 requests compared with 1 request showed no major difference in calculated availability numbers. Lowering the sampling rate is only recommended when you have many high-value, small-impression goal flights.

## **System Default Ads**

A 1x1 pixel ad is served when no other flights can serve to the request from any tier. For more information on inventory, refer to information in the SAS® 360 Match User's Guide on tiers and multivalue targeting.

## **Inventory Seeds**

Inventory seeds are a way for a publisher to define inventory that has not been sampled, such as a new customer, an acquired site, a new video program, a specific event, or offline inventory. A seed contains the following information:

- a path information string
- a date and an impression count for that date

The simulation generates count impressions for each date specified. Each impression has the tags provided by the path information string. The seed only specifies tag-related values. Therefore, visitor-related values are obtained from a copy of the previously sampled impression. This is useful so that you can keep your distribution of demographic data constant, such as the same `GEO_CITY` distribution. However, if you expect the distribution to change (for example, a new site that skews heavily female), then you must take this into consideration when allocating counts.

The feature is designed to evenly distribute the matching seeds throughout the given day. The day specified in the seed is for the forecast period. A count of 240 impressions would give 10 impressions to every hour. There is no way to specify the time of day.

The inventory seed feature uses the `seed_tags` and the `seed_times` tables.

A representative path information string is stored in the `tags` column of the `seed_tags` table. The `seed_times` table lists the `when` (date), how many (`count`), and which tags (`seed_id`) are used in the simulation. Each table has an ID column that uniquely identifies the row.

If you want the same seeds to apply to multiple dates, a row in `seed_times` must be present for each date.

Here are some issues to consider when using inventory seeds:

- Seeds are additive. You cannot subtract inventory with this feature.
- Seeds do not support supertags.
- Insufficient natural inventory may cause seeded impressions to appear on the wrong day or not at all. You cannot use seeds for the entirety of your inventory.
- Seeds copy the visitor ID from historical data, which may affect exposure policies.
- Inventory seeds are not part of the SAS® 360 Match UX. Instead, they are configured by uploading a CSV file.

## Use Cases

Here are some recommended ways you can use simulations to forecast your inventory.

### Forecasting Optimization

This use case shows you how to improve your SAS® 360 Match forecasting results. You can tune and optimize your simulation process towards the flights, tags, and products that make the most money and not go over your current capacity.

First, get source data from the BI cube `BIGSUM_EMPTY`. Use the top 20 revenue-generating flights from the past 30 days. You can export the cube to Excel if that is easier to work with. The results should look something like this:

Product Name	Flight	Floor Rate	Sold Rate	Goal Rate	List Rate	Impress Served	Total Rev
Main Homepage - CPM - 300x250	Honda - June - Home Page - 300x250 - 2015	\$9.00	\$9.15	\$12.00	\$15.00	1,576,414	\$14,424.19
Main Homepage - CPM - 300x600	Homepage - CPM - 300x600	\$15.00	\$12.50	\$20.00	\$25.00	1,144,301	\$14,303.76
Main Homepage - CPM - 300x600	Birthday Pop for Prizes - Homepage - 300x600 - 30 - 05 - 2015	\$15.00	\$17.50	\$20.00	\$25.00	363,597	\$6,362.95
Motors - CPM - 300x250	Honda - June - Motors Home Page - 300x250 - 2015	\$12.00	\$9.15	\$16.00	\$20.00	687,681	\$6,292.28

Next find all the tags, also known as keys, that are used in the targeting from those flights. Twenty standard tags and BT traits showed up in our samples of 20 flights. However, before changing the configuration, consider if the previous 30 days constitute a good sample of flights.

For example, determine if any BT traits that are frequently used were excluded during these 30 days. There may be special circumstances that might lead you to evaluate flights from 60, 90, or 120 days ago instead.

The current simulator configuration has the following tag details:

site, area, geo\_country, sc, a24, a8l, a18h, a9l, a49h, c, e, i, topic, ms, retargeting, pos, dev\_type, dev\_os, msg, body, a34, a35

Standard and BT tags from the top 20 flights:

Site, Area, E, Msg, ptype, Dev\_type, MS, P, C, Topic, Body, Model, abc, Vehicle, BT\_Autos, BT\_autos\_newcars, BT\_autos\_over40k, BT\_GI\_industry, BT\_make, BT\_phones

There is a discrepancy between the tags that the simulation uses and the tags based on the most profitable flights. Ideally, you want to include the tags that are commonly used in the flights and that provide the most value in the simulation. Other tags will use estimates and are not as accurate. Therefore, having an inappropriate set of tags in the simulation can reduce the accuracy of your most important forecasts.

## Global Inventory Adjustments

In this use case, modifiers are applied to the number of forecast impressions for specific days of the forecast. This data is taken from a flat file uploaded by the customer.

- You can provide input on specific global impression volume, on impression numbers by day, or by the percent increase or decrease for a specific day using the CSV format.
- The overall inventory is adjusted by the factors given and allows you to sell against it.
- Any unadjusted days are treated as is in the lookback window.

Each time the simulation runs, it checks for a new version of the file. If a new file is found, all existing modifiers are wiped out and replaced by the data in the new file. Therefore, any uploaded file must express the modifiers for all forecast days you would like to modify, regardless of whether a previous day's file expressed identical data for those days.

After processing, the file is moved into an archive subfolder.

If a file is not present when the simulation runs, but modifiers were uploaded on an earlier day, those modifiers are used in the simulation. To wipe out all modifiers, an empty file must be uploaded. The file must have the following elements:

- be named `invfactors.csv`.
- be placed in the `/public/inventory` folder.
- consist of plain text.
- NOT have a header row but consist only of data rows.
- have one row per day of the forecast to be modified, where each row has three comma-delimited values: `<date>`, `<absolute>`, `<relative>`. You can have leading and trailing spaces.
- `date` must be formatted as `YYYY-MM-DD` and represents the forecast (future) date whose number of forecast impressions is to be modified.
- `absolute` is an optional integer. If `absolute` is present and greater than 0, then it is the absolute number of forecast impressions to use for this day. The value of `absolute` overrides the number of impressions that would be derived from the basis data and averaging method. When `absolute` has a value, the value of the `relative` column is ignored and can be considered optional.
- `relative` is an optional decimal number (e.g., 1.14). It is ignored if `absolute` is present and greater than 0. If `relative` is present and greater than 0, it is the scale factor to apply to the number of forecast impressions derived from the basis and averaging method.

The file must include all the days of the forecast period that you wish to modify. When the file does not express data for a specific forecast day, the number of impressions forecast for that day follows the standard calculation made from the basis day and averaging method.

The file can express forecast days beyond the end of the forecast period. When the simulation is performed, the data for days that are beyond the end of the forecast period are ignored. In this way, you can preload data for future simulations, assuming there is no need to change the modifiers when those later simulations run.

This example uses a 7-day basis and 7-day forecast, where today is 2016-12-08. The basis period is 2016-12-01 through 2106-12-07, and the forecast period is 2016-12-08 through 2016-12-14.

Here are the impression counts for the basis period:

Date	Impression Count
2016-12-01	1,000
2016-12-02	1,100
2016-12-03	1,200
2016-12-04	1,050
2016-12-05	1,150
2016-12-06	1,000
2016-12-07	1,200

Assume that no averaging method is used during the simulation. This means that each basis day's number of impressions would be used as is for the corresponding forecast day.

Here is the data in a file named invfactors.csv:

2016-12-07,1200,1.0

2016-12-09,1050

2016-12-10,1150,1.2

2016-12-12,0,0.9

2015-12-13,,1.1

2015-12-15,1300,1.2

Here are the forecast period's impression counts for the simulation:

Date	Impression Count	Notes
2016-12-08	1,000	File has nothing for 12-08, so the basis count from 12-01 is used.
2016-12-09	1,050	File has an absolute value to use, so the basis count is ignored.
2016-12-10	1,150	File has an absolute value to use, so the basis count is ignored. The relative value of 1.2 is ignored because an absolute value is present.
2016-12-11	1,050	File has nothing for 12-11, so the basis count from 12-04 is used.
2016-12-12	1,035	File has a relative value that is used on the basis count (0.9 * 1,150).
2016-12-13	1,100	File has a relative value that is used on the basis count (1.1 * 1,000).
2016-12-14	1,200	File has nothing for 12-14, so the basis count from 12-07 is used.

The file's data for 2016-12-07 is ignored because it predates the forecast. Data for 2016-12-15 is also ignored because it is beyond this forecast period, but it will be used for tomorrow's simulation if no new data file is uploaded.

Now, assume that the simple averaging method is being used, which would ordinarily yield 1,100 (the average) impressions for every one of the forecast days. Here are the forecast period's impression counts that are used during the simulation:

Date	Impression Count	Notes
2016-12-08	1,100	File has nothing for 12-08, so the basis average is used.
2016-12-09	1,050	File has an absolute value to use, so the basis average is ignored.
2016-12-10	1,150	File has an absolute value to use, so the basis average is ignored. The relative value of 1.2 is ignored because an absolute value is present.
2016-12-11	1,100	File has nothing for 12-11, so the basis average is used.
2016-12-12	990	File has a relative value that is used on the basis average ( $0.9 * 1,100$ ).
2016-12-13	1,210	File has a relative value that is used on the basis average ( $1.1 * 1,100$ ).
2016-12-14	1,100	File has nothing for 12-14, so the basis average is used.

Regression methods used to extrapolate forecast impressions behave similarly. A relative modifier expressed for a day scales whatever number of impressions the regression would otherwise forecast for that day, and an absolute modifier overrides the regression's forecast.

## Constraints of Simulations

Certain data is not available in the logs that are used for the basis. This limits the effectiveness of certain types of targeting during the simulation process.

- USER AGENT, REFERER URL, and any ad hoc tags that are not explicitly defined in SAS® 360 Match are not available for simulation. When referenced in targets, they are treated as having empty values.
- No event data is available. When referenced in targets, all event counts are effectively zero.
- Tags or tag values that first came into use during the basis period may not be available during parts of the forecast period because they do not appear in some of the logs for the basis days.

In most situations, only some basis logs are used as input to the simulation to reduce the time that is needed to run a simulation. The simulation compensates by logging adjusted impression counts. This sampling process works well for general inventory. However, when specific tag combinations occur very infrequently but are used in targeting, the forecast results can be too low or too high. Rare inventory pockets that are direct sold sometimes need to have a higher resolution simulation than our 1:53 default.

The number and choice of tags that are logged in the future\_bigsum and inventory tables for your simulation are limited and configurable. Your selection affects the types of reporting that can be done after the simulation ends and how inventory forecasting queries are performed.

Tags that have multiple values might not be accurately reflected in the simulation. While a simulation uses multiple values for a tag for targeting, the simulation can only log a single value in by the output tables. This also occurs for the tables produced by serving live ads. However, you can exclude the tags from the simulation but review them using the statistically calculated option, which is used for all tags that are not in the simulation.

The same tag set is kept in the inventory table. Therefore, inventory forecasting queries can only assume that each chunk of inventory has a single value for each tag when, in fact, the



original (and future) ad requests might contain multiple values. For those reasons, it is best to exclude tags that frequently have multiple values from the simulation's logged tag set. Inventory forecasting queries have methods to estimate results for targeting that involves tags that do not appear in the inventory table, even tags that often have multiple values. For more information, see the section on Checking for Available Inventory.

## Target Rules Without Queries or Estimates

Some target rules do not involve logged tags and no estimation method is used. These are rules involving:

- DOMAIN
- KEYWORD
- TIME
- DAY OF WEEK
- RECURRENCE
- IP
- DISTANCE

When evaluating targeting, the availability function looks at the operator that is used on such rules. Rules with positive operators such as `Equal To`, `In`, `Like`, and `Greater than` are evaluated as false, on the assumption that such rules evaluate to false more often than true in live ad serving. Rules with negative operators such as `Not Equal To`, `Not In`, and `Not Like` evaluate to true.

## Uneven Inventory Distribution and Pacing Rules

Inventory typically is not uniformly distributed through the day or even across days. Therefore, a reported chunk of available inventory may not actually be usable to a flight that needs it during the simulation. This is also true during live ad serving.

During periods when the flight encounters enough inventory to maintain proper pacing, it may decline to serve any excess inventory. Later, when inventory availability drops, the flight may fall behind. This behavior is highly variable and influenced by the flight's pacing settings, whether top priority, delivery patience, or on-schedule priority.

The availability function reports aggregate inventory numbers for the date range queried and is not able to make adjustments that can account for these complex dynamics. As a result, a flight might not serve to its goal during the simulation or when serving live ads, despite the presence of sufficient inventory to satisfy it.

## Excluding Other Business Rules

The availability function does not take into consideration other business rules that affect the availability of inventory to a flight, such as the eliminate duplicate flights/advertiser options for the tier in which it resides, the companion flight options, or competitive exclusion.

## Frequently Asked Questions

### **What flights are conflicting for impressions in product X? Can I rearrange the flights?**

Refer to the section on inventory queries starting on page 7 for information on contending flights.

### **Why am I getting different numbers when I access inventory in different ways?**

Targeting might be slightly different in each case. Another possibility is that the requirement to estimate data for keys is not included in the simulation data set. Refer to the use cases and other sections of this document to find ways to minimize the use of tag estimation.

### **Why are campaigns under delivering despite there being available inventory?**

The engine wants to deliver a flight at with a specific number of impressions per second. If that rate is met, some impressions will go to other flights (maybe even preemptible or defaults, creating available inventory). SAS<sup>®</sup> 360 Match does not know if availability will decrease in the future. It serves it at the rate that is necessary to meet the calculated curve.

For example, if you have a goal of 1,000 and the flight runs for 10 days, ideally you should see about 100 impressions a day. If the availability of those 1,000 impressions occurs mainly on one end of the flight's run time, then you will see an under delivery. If all 1,000 impressions are available on the first day but no impressions were available on the subsequent days, you would see about 900 impressions available for that period even though the flight will under deliver. If the availability starts to decrease while the engine wants to increase the serve rate, you will see situations like this.

### **Why is a product that is oversold still giving me avails?**

Flights that are not set for top priority will allow some inventory to slip through when SAS<sup>®</sup> 360 Match (or the simulator) perceives that the flight is meeting whatever rate has been calculated to achieve its goal. For a flight to get as much inventory as it can, it needs to be marked as top priority. Short of that, a value for 0.0 for **Delivery Patience** might help. However, this will not eliminate inventory leaking to lower tiers, because available inventory does not present itself evenly. Traffic availability over time may affect the available inventory. For example, the forecast could show 1,000 impressions available on a Monday but not on subsequent days. If the flight has a goal of 1,000 impressions but is scheduled to be delivered from Monday through Sunday, all the 1,000 impressions are not available.

While there may be sufficient inventory when the flight is created, the actual inventory available is affected by subsequent changes, such as contending flights added, goals increased, and end dates shortened, during the duration of the simulation.

### **Why are availability and capacity numbers different from one day to the next?**

The numbers fluctuate slightly each time the simulation is run because a slightly different set of historical dates is used every time. Specifically, capacity for product X might be reported as 2,000,345 today, while running the same report tomorrow might result in a capacity of 1,999,896. In most cases, the differences should be small enough not to affect the general sense of what is available.



**THE POWER TO KNOW®**

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration. Other brand and product names are trademarks of their respective companies. Copyright © 2015, SAS Institute Inc. All rights reserved.

---