




1:1 Customer Engagement Optimization

STUDENT GUIDE



Copyright 2026
Pegasystems Inc., Cambridge, MA

All rights reserved.

This document describes products and services of Pegasystems Inc. It may contain trade secrets and proprietary information. The document and product are protected by copyright and distributed under licenses restricting their use, copying, distribution, or transmittal in any form without prior written authorization of Pegasystems Inc.

This document is current as of the date of publication only. Changes in the document may be made from time to time at the discretion of Pegasystems. This document remains the property of Pegasystems and must be returned to it upon request. This document does not imply any commitment to offer or deliver the products or services provided.

This document may include references to Pegasystems product features that have not been licensed by your company. If you have questions about whether a particular capability is included in your installation, please consult your Pegasystems service consultant.

PegaRULES, Process Commander, SmartBPM® and the Pegasystems logo are trademarks or registered trademarks of Pegasystems Inc. All other product names, logos and symbols may be registered trademarks of their respective owners.

Although Pegasystems Inc. strives for accuracy in its publications, any publication may contain inaccuracies or typographical errors. This document or Help System could contain technical inaccuracies or typographical errors. Changes are periodically added to the information herein. Pegasystems Inc. may make improvements and/or changes in the information described herein at any time.

This document is the property of:

Pegasystems Inc.

1 Rogers Street

Cambridge, MA 02142

Phone: (617) 374-9600

Fax: (617) 374-9620

www.pegasystems.com

Mission: 1:1 Customer Engagement Optimization

Product: Pega Customer Decision Hub™ '26

URL: [1:1 Customer Engagement Optimization | Pega Academy](#)

Date: 06 August 2026

Contents

Optimizing 1:1 customer engagement overview	4
Simulation testing	5
Monitoring action performance	7
Action Analysis	14
Analyzing customer distribution using Pega Value Finder	29
Pega Value Finder	30
Analyzing customer distribution using Pega Value Finder	31
Testing engagement policy conditions using audience simulation	40
Audience simulation	41
Running an audience simulation test.....	42
Analyzing the results of simulations using reports.....	56
Distribution test	57
Analyzing the results of simulation using reports.....	59
Detecting unwanted bias	62
Ethical bias.....	63
Detecting unwanted bias in engagement policies	65
Analyzing the effect of business changes using Pega Scenario Planner	71
Pega Scenario Planner	72
Analyzing the effect of business changes using Pega Scenario Planner	76

Optimizing 1:1 customer engagement overview

Description

Familiarize yourself with the various simulations that are used in one-to-one customer engagement to understand the effect of business change and analyze the areas of improvement to meet the business requirements. Also, learn how to monitor and understand the performance of live actions in Pega Customer Decision Hub™.

Learning objectives

- Explain the purpose of simulation testing
- Define the different types of simulations
- Monitor and understand action performance

Simulation testing

Simulation types


By running simulation tests in Pega Customer Decision Hub™, you can understand the effect of business changes on your next-best-action strategy framework. The Pega Customer Decision Hub portal offers a large variety of simulations. The simulation capability ranges from simulations that help identify under-served customers to simulation tests that allow you to investigate how an introduction of a new engagement policy might affect actions offered across a segment of customers.

The different types of simulation testing available in Pega Customer Decision Hub are: Value Finder, Audience Simulation, Distribution Test, Ethical Bias, and Scenario Planner.

Click on the hotspots to learn more about each of the simulation types.

Business-as-usual

Simulations are run in a Business Operations Environment (BOE) that is specifically designed to build, simulate and optimize changes. A sample dataset, which includes interaction history and adaptive models from the production environment, is created via a pipeline into the BOE. This dataset is used as the basis for the simulations.



Let's see how each of these simulations can be used in a business-as-usual use case.

U+, a retail bank, has recently implemented a project in which credit card offers are presented to qualified customers when they log in to the web self-service portal.

Now U+ would like to leverage the simulation testing capabilities of Pega Customer Decision Hub to:

- Market credit card offers to qualified customers
- Verify if their engagement policy conditions are presenting offers to customers as expected
- Check the offer distribution and prioritization ranges
- Check if business policies and regulations are being violated
- Show improvements in projected value

Each of these simulation types plays a specific role in a business operating model. This is U+ bank's operating model. An operating model supports businesses in the planning, development, testing, monitoring, and optimization of changes. The simulation capability can be utilized throughout the lifecycle but has greater significance when run at certain points.

Click the hotspots to see in which phase of the lifecycle U+ bank can use each of these simulation types.

Monitoring Action performance

With the Action Performance feature, you can monitor and understand the performance of live Actions in Pega Customer Decision Hub™. This tool provides an operational view of key performance indicators (KPIs) so that you can track customer engagement, identify trends, and ensure that new Actions receive adequate exposure. Action alerts notify you of significant performance changes, while advanced filters and detailed data help you to optimize marketing strategies and maintain transparency.


Action Performance overview

Transcript

This video shows you how to understand the performance of hundreds of live Actions in Pega Customer Decision Hub. Use Action Performance Action to monitor active Actions for changes in customer engagement activity. Identify important trends that require business attention, such as a decline in accept rate, and verify that new Actions receive enough exposure.

Action Performance includes multiple views that you can use to analyze your Actions from different viewpoints. You can review Action Performance by Always-on Actions, One-time Actions, Top Actions, Bottom Actions, Action Placement, and Action Analysis. Each view displays results based on the corresponding KPIs. You can also filter the results by scope, timeframe, Channel and audience. The Action alerts feature in Customer Decision Hub triggers notifications when an algorithm detects significant changes in Action performance, such as unusual impressions or click/accept rates, or low daily impressions for new Actions.

The algorithm monitors performance at the Channel and treatment levels. Notifications for established Actions include the date, reason, issue, group, Action, and affected KPI. For new Actions, the system trigger alerts when impressions are below a set number. Clicking an alert opens the Action Performance page with prepopulated data. A banner at the top highlights details of the detected performance change.



Analyzing the Action performance

This video helps you to analyze Action Performance of Always-on Actions. You can access the Action Performance landing page in Customer Decision Hub by clicking **Discovery > Action Performance** in the navigation pane. The default view shows the last two weeks of data in daily granularity, covering all active treatments, one-time Actions, Always-on Actions, groups, and issues, across all Channels.

Note: Action Performance Tracker is a component of the 1:1 Strategy Optimizer Bundle, and works in the Business Operations Environment only when a data migration pipeline is enabled.

Data can be presented for up to 15 months, depending on how you filter the data. You can filter the data by **Timeframe**, **Scope**, and **Channels**.


All filters are active for the duration of the Action Performance session. After removing or adding a filter, click **Apply** to refresh the page.

Action Performance


View: Always-on Actions | Action status: All | Timeframe: Daily (over two weeks) | End Date: 9/9/2024 | Scope: All issues | Channels: All channels | Apply

[Advanced filters](#)

The default view shows aggregate data for the following out-of-the-box KPIs, such as treatments that have been sent in outbound rounds, total number of impressions for inbound Channels, and the Click/Accept rate on offers for the selected scope:




When viewing a single item, such as an issue, hovering over a single bar in a graph provides data for that issue.



If you select multiple issues, the bars have gradations so that you can visually break down the data into different selections. The section at the bottom shows a detailed data sheet for each issue within the selected scope.


When comparing multiple items, click the arrow next to the date to see how the data compares across a selected period for each of the issues. Clicking on the aggregate number of sends or impressions opens a detail overlay, with a chart showing how the KPI is distributed across the selected scope. Hovering over a segment provides more detailed data. You can view up to 10 items at once.




Operational KPIs provide insights that give the Next-Best-Action operations team and stakeholders the confidence that Actions are performing as expected, and the possibility to identify trends and anomalies that require attention. Filtering and comparing current and past data ensures transparency.

In the Action Performance Tracker, you can apply the **Advanced filters** option to the data. By selecting the **Key Audience** filter, you can narrow down the data to a specific key audience segment that you created on the Segments landing page.

After selecting a key audience, the numbers and graphs that are displayed on the Action Performance Tracker landing page are specific to that key audience.



You can drill down into specific Actions to analyze how those particular groups or Actions engaged the selected key audience.



To automatically generate an offline Excel spreadsheet that contains the metrics for the selected timeframe, in the upper-right corner, click **Export Excel** file. The report retains Action Performance Tracker's filters, allowing you to inspect KPIs.

This demo has concluded. You have learned:

- How to understand and use the default view and KPIs.
- How to filter and compare data.
- How to use the **Advanced** filter and **Key Audience** selection.
- How to export data to an Excel file.

Action placements

This video helps you to understand how to use **Placements** view on the Action Performance landing page to monitor the performance of Actions on their placements.

Placement is a specific location where an Action is presented on an **Inbound** Channel. For example, the web Channel includes placements such as Hero, Tile, and Header.

The **Placements** view in Pega Customer Decision Hub helps you to analyze the performance of Actions across placements. By using the key performance indicators (KPIs) you can evaluate Action placement and refine it to drive higher customer engagement.

Select the **Placements** view Action and use the available filters to see placement performance.

Action Performance

View: Start Date: End Date: Placement types: Scope: Channels:

[Advanced filters](#)

Select a specific timeframe to see how Action placements occur during that period. Use the **Placement types** field to specify the placement type that you want to review. You can select all, single, or multiple placement types according to your requirement. In the **Scope** field, select **All issues** to see collective placement performance across all issues, **Issue** to see performance for specific issues, **Group** for specific groups, or **Action** for individual Actions. Use the **Channel** filter to narrow results by a specific Channel. Use the **Advanced filters** option to see how Action placements perform for specific key Audience.

The result displays a heat diagram that maps placement types across scope.



Three KPIs measure placement effectiveness. **Impressions** refer to the number of times Actions were shown in a placement. **Clicks/Accepts** refers to the number of positive customer responses to Actions in a placement. **Click/Accept rate** refers to the total number of clicks (for digital Channels, including outbound email and SMS messages and inbound self-service Channels such as web and mobile), and accepts (for assisted inbound Channels such as call center and retail), as a percentage of the total number of impressions in the respective Channels.

KPI

Impressions


Impressions

Clicks / Accepts

Click / Accept rate

You can see a table that represents the combined results of all the KPIs. Now, change the filters and see the Action placement performance for individual groups or Actions.

In this scenario, Actions in the credit cards group perform better in a carousel than in a sidebar placement.



Similarly, adjust the filters and analyze Action placements to align with your business priorities.

This demo has concluded. You have learned:

- About the Placement view in Action Performance.
- How to analyze Action Placements.

Action Analysis

Action Analysis provides insights to help you optimize your Next-Best-Action Strategy Framework and steer towards the most added value.

It helps you understand how Actions pass through the Action funnel, whether there are sufficient Actions for prioritization, and how propensity, weight, value, and levers influence prioritization.

You can run Action Analysis for a particular Issue, Group, Action, Treatment, Channel, Page Placement, or a Key Audience, to answer questions such as:

- Where am I losing opportunities to offer my Credit Card Actions?
- Why do high-value London-based customers have such a low acceptance rate of Credit Card Actions?
- Can I explore Engagement Policies to see if there are any opportunities to relax filtering for Credit Card Actions?

Transcript

This video demonstrates how to optimize engagement with Action Analysis in Pega Customer Decision Hub.


The Action analysis view in Customer Decision Hub helps you understand not just how your Actions perform, but enables you to discover insights into what has happened.

Action Analysis works only in the Business Operations Environment (BOE) and uses decision data from Production to give a near real-time view of Actions.

It shows how Actions move through the Action funnel, whether there are enough actions for prioritization, and how factors like propensity, weight, value, and levers influence outcomes.

With this insight, you can spot missed opportunities, uncover overly restrictive Engagement Policies, and refine your Next-Best-Action strategy to boost customer engagement and business results.

An action is analyzed based on metrics such as Sends, Impressions, Click accept rates etc. To analyze underperforming Actions and optimize your Next-Best-Action strategy, navigate to the **Customer Decision Hub**, from there go to **Discovery**, then select **Action Performance**. In the view list, click on **Action analysis**.



To narrow your focus and analyze specific data, apply filters in the Action Analysis view. Start with the **Start** and **End Date** filter — by default, it shows data from the last 30 days. Next, use the **Action Context** filter to choose the level of analysis, such as Customer or Account.

Then, refine your view using the **Scope** filter. You can select from All Issues, a specific Issue, Group, Action, or even a Treatment.

To narrow your analysis to a specific channel, apply the **Inbound Channels** filter. For this demo, we'll focus on the web channel.

Use the **Page Placement** filter to quickly target specific layouts, such as carousels or tiles, especially on Web or Mobile interfaces.

To focus on the most relevant segments, you can use the **Advanced Filters** and select **Key Audience**. However, this step is optional — for this demo, we'll proceed without applying any advanced filters.

Action Performance

View
Start Date
End Date

Action analysis
1/1/2025
6/30/2025

Advanced filters

Advanced filters


In the Action funnel, Key Audiences can only be reported on from the Engagement Policy stage onwards. Therefore, numbers in the 'Available Actions' and 'Channel availability and Product associations' stages include Actions for all customers. From Engagement Policies onwards, Actions are shown only for Key Audience members.

Key Audience

None

Submit

Click **Apply** to view the results. The Action Analysis view provides four key insights to help you refine your Next-Best-Action strategies.



First, Action Funnel Analysis — this shows you where Actions are being filtered out, helping you identify potential drop-off points in the decision flow.

Second, Actions per Decision Available for Prioritization — this insight helps you assess whether there are enough Actions available for prioritization. Typically, the more Actions that are available, the higher the likelihood of engagement.

Third, Engagement Policies — this highlights which Eligibility, Applicability, or Suitability rules are filtering out too many Actions. By adjusting these policies, you can improve Action availability and expand your reach.


And finally, Prioritization Influence — this insight reveals how factors like propensity, weight, value, and business levers impact Action ranking. Use this to fine-tune your prioritization strategy for better outcomes.

Exploring the Action funnel

Action funnel shows how many Actions remain after each stage group of the Next-Best-Action Strategy Framework, helping you identify where Actions are being filtered out.


Start by reviewing the **How do Actions pass through the Action funnel?** section. The insight box and chart reveal how Actions flow through each stage. For example, if most Actions are filtered out at the Engagement Policies stage, that's a signal to investigate further.

Click **Open Action funnel** to access a detailed breakdown. Use the upper section of the dashboard to toggle between different views.




Passing Actions show those that successfully move through each stage, Filtered Actions highlight those removed at specific stages, and Decisions without Actions indicate where no actions are available.

Let's switch to the Filtered Actions view to understand what's being excluded.



The lower section of the page displays the results in a table format, with columns for **Available**, **Passing**, and **Filtered Actions**, as well as **Decisions** for each stage group.

Clicking on any number in the table opens a pie chart that breaks down the data by business hierarchy, such as **Issue** or **Group**.




To dig deeper, click a stage group, such as **Engagement Policies**, to filter out the most Actions. The insight box will guide you in adjusting restrictive policies to improve Action availability.

Stage group	Available Actions	Passing Actions	Filtered Actions	Decisions
Available Actions ⓘ	73,745,116	73,745,116	0	4,406,400
Channel availability and Product associations ⓘ	73,745,116	69,354,084	4,391,032	4,406,400
Engagement Policies ⓘ	69,354,084	36,732,648	32,621,436	3,735,580
Journeys and Contact Policies ⓘ	36,732,648	22,039,684	14,692,964	3,371,712
Action Propensity ⓘ ⓘ	22,039,684	20,100,120	1,939,564	3,283,788
Available Treatments ⓘ	20,100,120	18,957,992	1,142,128	3,177,436
Treatment Eligibility ⓘ	18,957,992	13,580,960	5,377,032	2,431,680
Treatment Propensity ⓘ ⓘ	13,580,960	11,816,700	1,764,260	2,291,532

If a stage shows no data, it may be missing from the framework or not configured to store data, like outbound channel limits.


Finally, to refine your strategy, click **Open in Next-Best-Action Designer** in the upper-right corner and adjust policies as needed.



The Action funnel is a multi-stage filtering system that ensures only the most relevant and high-propensity Actions are presented to customers. It improves engagement and business outcomes by progressively narrowing down Actions through the various stage groups.


As you explore the Action funnel analysis, you can navigate through different stage groups using the left and right arrows in the Action funnel section to move seamlessly between the previous and next stage groups in the funnel.

The Action funnel begins with the **Available** Actions stage, which includes all active Actions generated by system decisions within the selected timeframe and filters. This is the starting point where every potential Action is considered.



Next, the funnel narrows through the **Channel Availability & Product Associations** stage. Here, Actions are filtered based on their compatibility with communication channels and product validity.

Once Actions pass these checks, they move into the **Engagement Policies** stage. Here, business rules are applied to determine whether each Action is eligible, applicable, and suitable for the customer across different contexts.



Finally, the **Journeys & Contact Policies** stage ensures that Actions are aligned with customer journeys and applies suppression rules to manage communication frequency. Actions not tied to a journey or those that exceed contact policy thresholds are excluded.

Exploring Action Availability for Prioritization

To determine whether there are enough Actions to prioritize in the Arbitration stage group of the Action funnel, you conduct an optionality analysis.

This analysis helps you assess whether the volume of available Actions is sufficient to support effective prioritization. Generally, the more Actions that are available, the higher the likelihood of achieving better click-through or acceptance rates.

In the Action analysis view, the **Are there sufficient Actions to** **prioritize?** section displays a chart showing how many Actions are available per Decision. For example, if the average is only two Actions per Decision, that's considered insufficient. Many Decisions may have none or very few Actions to be effective. The chart categorizes Decisions into three buckets. Decisions without Actions have zero associated Actions and appear in red. Decisions with Insufficient Actions include up to six Actions and are shown in yellow. Decisions with Sufficient Actions have seven or more Actions and are indicated in green.




Clicking **Show propensity** reveals how engagement likelihood increases with more available Actions.



If many Decisions lack Actions, this can indicate overly strict filters in earlier stages, such as **Engagement Policies** or **Journeys & Contact Policies**. Reviewing the full Action funnel can help you to identify and resolve these bottlenecks.

Exploring Engagement Policies

To understand how Engagement Policies filter Actions, click **Open Engagement Policies stage group** in the Action Analysis view.




This insight helps identify which stage – Eligibility, Applicability, or Suitability – is filtering out more Actions than expected.

Here, you'll see a chart and a table showing how many Actions are available, passing, or filtered at each stage.



You can also click on any number in the table to view a pie chart that breaks down Actions by business hierarchy, such as by Issue or Group. For instance, if the scope is set to All Issues, the chart will show how many Actions pass the Suitability stage from each Issue, such as Sales or Service.




If you find that too many Actions are being filtered out, consider loosening the restrictive policies. To do this, click **Open in Next-Best-Action Designer** in the upper-right corner of the Engagement Policies section.

Exploring Components that Influence Prioritization

Sensitivity analysis evaluates how four key factors—propensity, context weighting, business value, and business levers (PCVL model)—influences action prioritization.

The section in Action analysis view highlights which factor has the greatest impact. For example, if propensity dominates, it indicates a customer-centric approach where Actions most relevant to users are prioritized.

To explore how each component influences the prioritization of Actions, begin by clicking **Open Prioritization stage** in the Action analysis view.




This opens a detailed breakdown of how the system evaluates and ranks Actions during arbitration.

The system displays the Action funnel analysis at the level of the Prioritization stage, as shown in the figure.

Let's check the components that influence the Sales Issue.

Here, except for Business value, all the other components have a negative influence on prioritization.


The **Prioritization influence** tab helps you understand how much each PCVL component impacts the final ranking of Actions.




The system calculates this by temporarily removing one component at a time and measuring how many Decisions end up with a different top-ranked Action. For example, if removing propensity changes the winning Action in 60 out of 100 Decisions, its influence is 60%.

This analysis shows how heavily each factor contributes to prioritization. It helps you assess whether your model is balanced or if one component is dominating, which could affect the effectiveness and fairness of your Next-Best-Action strategy.


The **Loss report** highlights Actions that consistently rank second or lower and aren't returned to the channel. By analyzing these patterns, you can adjust prioritization components to improve their chances of success. Win/Loss reports are available only for a specific scope, such as an Issue, Group, Action, or Treatment. The Win/Loss reports aren't available for the All Issues view.



The report shows metrics like the **Median losing priority**, which reflects the typical priority score of losing Actions over time. To get a broader view, click **Prioritization comparison** in the **Median losing priority** section. This reveals how each prioritization component, such as propensity or business value, compares between your Actions and the Actions that they lose to.




In the lower section, a table displays how many Actions in your selected scope lose to competing Issues or Groups. You can see the loss rate and compare median priorities across categories. Clicking a number in the **Median priority** column lets you drill down into comparisons against specific competitors.



If you find that your Actions are consistently underperforming, consider adjusting the weaker components. You can do this by clicking **Open in Next-Best-Action Designer** in the upper-right corner of the page.

Use the **Win report** to see how often Actions win in the selected scope and how they compare to competitors.




Actions ranked first are considered wins and returned to the channel. The lower table shows win counts and median priorities across competing Issues or Groups.

For an overall view, click **Prioritization comparison** in the **Median winning priority** section.



This shows how arbitration components influence prioritization. Adjust components below competitors to improve outcomes.

Click a number in the **Median priority** column to compare against a specific Issue or Group.



You have reached the end of this video. What did it show you?

- How to optimize engagement with Action Analysis in Pega's Customer Decision Hub.
- How to explore the Action Analysis view, its purpose, and how you can use it to enhance your Next-Best-Action strategies.

Analyzing customer distribution using Pega Value Finder

Description

Learn how Pega Value Finder can be used to discover areas of improvement in the next-best-action strategy and adjust the engagement policies to create new actions and treatments based on the customer distribution analysis.

Learning objectives

- Explain what a value finder simulation is
- Analyze customer distribution using value finder

Pega Value Finder

What is a value finder simulation?

A Value Finder simulation allows you to engage more empathetically by identifying and profiling “under-served” customers, then suggesting actions for improvement – like adjusting engagement policies or creating new actions and treatments.

Using Value Finder, you can discover areas in which you can improve the next-best-action strategy by monitoring scenarios in which customers are presented with no actions or only low-propensity actions. This is particularly useful when planning new changes or optimizing existing parameters.

Value Finder identifies and profiles “under-served” customers. These are customers that either don’t receive actions, or only receive actions they have a low propensity to accept. It analyzes what happens at every stage of the next-best-action decision funnel, enabling you to:

- View offer distribution to well-engaged, under-engaged, or not-engaged customers
- Identify top opportunities for improvement
- Review details of under-served groups at each level of arbitration

Let’s consider an example in which a bank runs a value finder simulation and identifies groups of not-engaged and under-engaged customers. The details of the groups and how to address the findings are as follows:

Not-engaged customers: There are 7000 under-served customers, none of whom own a credit card. They all have credit scores over 650 but are blocked because the eligibility rules in place keep them from seeing specific offers they have a high propensity for. In this case, you might want to tweak the engagement policy to present them with more appropriate offers.

Under-engaged customers: There are 5000 customers that have good credit scores and own a credit card but have no propensity scores higher than 5%. You might need to create a new offer with different terms and test it to see if you can capture their attention.

Tip: To understand how the relevant action setting works, see [Defining the global default settings to determine relevant actions](#).

Analyzing customer distribution using Pega Value Finder

A Pega Value Finder simulation allows you to engage more empathetically with customers by identifying and profiling customers who receive irrelevant actions. A business can then analyze the Value Finder results to adjust engagement policies or create new actions and treatments.

Transcript

This demo shows you how to analyze customer distribution by using a Pega Value Finder simulation.

U+ Bank, a retail bank, has recently implemented a project in which qualified customers see credit card offers when they log in to the self-service web portal. The bank wants to check if any restrictive engagement policies result in customers with **irrelevant actions** in the web channel and find ways to serve them better.

This is the Pega Customer Decision Hub™ portal. In **Discovery**, you can create a new Value Finder simulation.

To create a Value Finder simulation run, select the issue and group and the channel in which you want to find opportunities. Then, select an audience on which to perform the simulation run. The audience is a list of potential target customers. You can easily modify the simulation name to identify the specific runs.

Create new Value Finder simulation

Note: It might take a moment to get your simulation running.

Scope & audience ?

Select the scope and audience for the Value Finder simulation.

Scope *

Grow / Credit cards

Channel

Web

Audience *

SampledCustomers_Inbound

Audience size

Not available for this audience type

Name


Value Finder

To see the simulation results, open the saved simulation. After the simulation run is complete, the Value Finder results display three areas: the scope, channels against which the simulation runs, and the number of customers in the simulation.

The right side of the screen shows the customer categories. Value Finder calculates these categories after each engagement policy stage. You can take a closer look at the categories of customers for whom engagement is low.


The pie chart on the left shows the results after Value Finder applies eligibility, applicability, suitability conditions, and arbitration.

In this example, 64.4 percent of customers are without actions. 14.2 percent of customers have only irrelevant actions. That means that the propensity to accept the action is low and is not good for customer engagement. Then, 19.2 percent of customers receive relevant actions and are more likely to accept the offers, which is good for engagement. Then, we have 2.2 percent of customers who receive only new actions. These are actions with an associated model that has not learned enough for you to rely on the propensity it produces.




To understand the results better, you can view the distribution of customers in the four categories for each engagement stage: eligibility, applicability, suitability, and arbitration. The bar chart illustrates the impact that each stage has on customer engagement and allows you to compare engagement across the stages.

In this example, you can see that most customers get at least one relevant action after eligibility and applicability. But after you apply suitability, a portion of the green and orange bars moves to red. As a result, these customers receive no actions or only irrelevant actions that they are less likely to accept.



Value Finder identifies customer categories after applying eligibility, applicability, suitability, and arbitration rules. These categories represent value potential; you can improve engagement by addressing the issues that Value Finder identified. The system broadly classifies the categories into customers with no actions and customers with irrelevant actions.

In this case, you must look at the suitability stage to understand what is happening. Notice that 355 customers receive no actions, and 26 have irrelevant actions because the suitability condition might be too strict.



Value Finder precalculates the minimum propensities that determine whether actions in specific channels are relevant. The system calculates that the minimum propensities take half the average model success rate for a given channel and direction, regardless of the action. Actions with higher propensities are considered relevant. The precalculated propensities from the global default settings are available here.

This example shows only the web channel per the simulation scope. When no data-based suggestion is available for a channel direction, the minimum propensity value defaults to 5 percent. You can modify these default settings for a new simulation. These modifications do not impact the simulations that are already created.

Value Finder settings

Note: Any modifications will cause your results to be reanalyzed and recalculated. It might take a moment to see a change in results

Relevant actions

Define the minimum propensity that considers an action to be relevant for a customer. A higher propensity results in more actions being classified as irrelevant.

Channel	Direction	Default propensity	Simulation propensity
Web	Inbound	35.00%	35.00%

Cancel
Submit

Extension simulations

You can dive deeper into the insights of Value Finder by running extension simulations for each customer group identified.

For example, you can run an extension simulation for the **Irrelevant actions after suitability** customer group. After running the extension simulation, you receive two sets of insights that help you to:


- Analyze which suitability policies filter out relevant action for these customers across groups and actions.
- Review descriptions of similar customers, of which a large part is only receiving irrelevant actions.
-

Now, view the analysis and update the engagement policies to understand the number of relevant actions that are filtered out after group engagement policies are applied.

This extension simulation shows that only 26 customers have irrelevant actions after applying suitability conditions. However, there were 38 relevant actions available for the customers before the suitability conditions were applied. 26 customers had about 38 relevant actions, which Value Finder filtered after applying the suitability conditions.

In the table, you can view a breakdown of the policies that prevent customers from receiving relevant actions at the group policies level, the **All issues** level, and the **All groups** level. The breakdown shows that the 38 actions filtered out were Rewards Plus card actions and Premier Rewards card actions. You can further analyze and view the suggested recommendation to determine at which level the suitability condition prevents customers from receiving an action. Use this information to identify the suitability conditions that prevent these groups of customers from receiving actions, whether at the

group or action levels. In this example, the Rewards Plus card action-level policy, *Is financially vulnerable is false* is where the system filtered out relevant actions.



To engage more customers, you can consider loosening a restrictive engagement policy by examining how many customers do not meet the policy conditions at the group or action level. You can add the findings to the value summary report to review them later. You can analyze the Premier Rewards card as well.

Customer groups

The **Analysis of customer** recommendation identifies customer groups that receive irrelevant actions. You can analyze these customer groups by using a distribution test to create new actions or treatments to serve these customers better. The system does not generate customer groups when there is insufficient evidence.

For example, if you create a new action or treatment for group A, there 75 customer engagements that you can improve address.

Details are provided for the **Only irrelevant actions after suitability** customer group to show how suitability conditions contribute to the issue. Value Finder supplies insights about each group.

The **Receiving only irrelevant actions** line gives the number of customers with irrelevant actions in a group. The percentage is the number of customers in the group divided by the total number of customers identified as receiving irrelevant actions. In this case, the system identifies 75 customers out of 154 who receive irrelevant actions in Customer group A.

The **Description** line gives the group characteristics, such as *CreditScore*.


Accuracy is the number of customers with irrelevant actions in the group divided by the total number of customers in the group. If the accuracy is 72.1 percent, then 72.1 percent of customers in the group receive irrelevant actions. The remaining $100 - 72.1 = 27.9$ percent have at least one action above the propensity threshold. If the accuracy, for example, is 100 percent, then all the customers in the group receive irrelevant actions.

 Customer group A		
Receiving only irrelevant actions	Accuracy	
48.7% (75 of 154)	72.1%	
Description		
CreditScore is less than or equal to 295.33		

You can also manage the group by removing a field from the description. Typically, the system lists all potential fields that are available for the next-best-action strategy. You can remove a field from the list if it is not required. If you remove a field, Value Finder recalculates the values for **Receiving only irrelevant actions** and **Accuracy**.

Manage Descriptions 	
Manage the fields that are used to identify and describe under-served customer groups.	
Used fields	
Age	
Annual income	
Credit score	
Relationship length in days	
Total liabilities	
Ignored fields	
No fields are ignored	

You can also save or export the customer groups with irrelevant actions as audiences.




You can then run distribution tests to get more insight into the current actions that these audiences receive. Use Value Finder recommendations and your distribution test results as feedback for business stakeholders. This information can inspire stakeholders to create new actions and treatments that are relevant to these customers.

This procedure applies to all Value Finder recommendations for customers with irrelevant actions and without actions, which you can view on the Value Finder landing page after you run a new simulation or display the results of a saved simulation.

Value summary

After you add the policies to the value summary, the value summary reflects the number of engagement policies in the summary.



Value summary shows the potential increase in customer engagement.

The value summary shows the insights you saved and the potential of addressing each issue. You can remove policies from the summary as needed.

You can download and share the value summary PDF with the Operations Manager team and then discuss whether to loosen this condition to let more irrelevant actions through. They can deploy this change and then monitor the effect.

This demo has concluded. What did it show you?

- How to configure and run a Value Finder simulation.
- What the top customer groups identified by Value Finder are.
- How to analyze the customer distribution across stages
- How to run extension simulations.
- How to interpret the analysis of customers.
- How to interpret the customer groups identified to receive only irrelevant actions.
- How to drill down to identify conditions that prevent customers from receiving irrelevant actions and no actions.
- How to add policies to value summary to examine how many customers do not pass the policy conditions at the group or action level.

Testing engagement policy conditions using audience simulation

Description

Learn how to improve the performance of your next-best-action strategy by testing the configuration of engagement policies against a set of customers.

Learning objectives

- Explain what an audience simulation is
- Run an audience simulation test

Audience simulation

You can improve the performance of your next-best-action strategy by testing the configuration of engagement policies against a set of customers. In this way, you can check how many potential actions are filtered out by each component of the policy and discover if a particular criterion is too broad or too narrow for your requirements.


Transcript

This video explains what an audience simulation is and how it improves the performance of your Next-Best-Action strategy.

Audience simulation tests the configuration of engagement policies against a set of customers.

Consider an example in which an audience comprises 1071 customers. Let's see how the audience will be filtered based on the engagement policies.

When the eligibility conditions are applied, only 350 out of 1071 pass through. These 350 customers become the input for the next stage, applicability conditions. Only 99 of the 350 pass through the second stage. In the final stage, suitability conditions, only 65 pass through. Of the original 350, only 65 customers qualify to receive at least one offer.




The filtering that happens in this simulation is similar to a funnel filtration for every engagement policy condition type. However, within an engagement policy condition type, each of the criterion is applied separately to the corresponding input audience.


Let's now look at the same example in detail to see how the funnel filtration happens within each engagement policy type.

The engagement policies can be defined for a specific group within an issue and/or for individual actions. Group-level conditions are applicable to all the actions within that group while the action-level conditions are applied to that specific action only.

In this case, the total number of the customers in the audience is 1071. Within the eligibility criteria, the first condition does not filter any customers, as all qualify. When the second condition is applied, 958 customers pass through.



Since both conditions need to be met, the intersection of these two conditions, 958, is the final number of customers who pass through at the group level. Now the system checks for individual action-level eligibility conditions that also need to be met.




The intersection of the group-level and action-level conditions form the final eligibility-level output population of 350.




The output population of the eligibility level is passed on as the input for the next level.

The input population that will be filtered by the applicability conditions is therefore 350. When the applicability condition is applied, 99 of the 350 customers qualify for at least one offer. In this case, let's assume there are no action-level conditions. Thus, the result from the individual group-level condition is the same as the overall applicability level, 99.




The output population of the applicability level, 99 customers, is passed on as the input to the suitability level.

When the suitability group-level condition is applied, 71 of the 99 customers qualify for the offers.



Now the system checks for individual action-level suitability conditions, and the intersection of both group-level and action-level conditions forms the final suitability-level output population, which is 65.



In summary, with the audience simulation test you can check how many potential actions are filtered out by each component of the engagement policy and discover if a particular criterion is too broad or too narrow for your requirements.


Running an audience simulation test

You can improve the performance of your next-best-action strategy by testing the engagement policy configuration against a set of customers. In this way, you can check how many potential actions are filtered out by each component of the policy and discover how a criterion might affect your requirements.

Audience simulation


This demo will show you how to run an audience simulation test. It will also explain how many potential actions are filtered out by each section of the engagement policy and reveal if a particular criterion is too broad or too narrow for your requirements.

U+, a retail bank, has recently implemented a project in which credit card offers are presented to qualified customers when they log in to the web self-service portal. The bank ran a value finder simulation to find under-engaged customers. In that simulation run, a set of customers is identified with no actions after engagement policy conditions are applied.



The bank would now like to run an audience simulation to investigate why these customers have no actions presented to them and correct the condition that causes the deviation. This is the Pega Customer Decision Hub™ portal. First, take a look at the set of eligibility criteria that has already been configured by U+.

In **Audience simulations**, you can create a new simulation run.



To create an audience simulation run, select an audience on which you would like to do the simulation. The audience is a list of potential target customers. You can modify the simulation name as required to easily identify the specific runs. Then, choose the scope of the run. That is, you can choose to simulate only the engagement policies to validate the eligibility, applicability and suitability conditions. Or, you can simulate on both engagement policies and arbitration to understand how the conditions work when arbitration across all actions is also considered.

Create simulation ✕

Audience *

Simulation ID prefix

SampledCustomers ▼

EngagementPolicyTest ✎

Population
1000

Next-Best-Action scope ?


☒ Engagement policies only

☐ Engagement policies and arbitration

Cancel

Run

Once the simulation run is complete, you can view the details of how the audience is filtered at the group level based on the configured engagement policy conditions.



For each component of the engagement policy, the simulation test shows a numerical or percentage value of the audience that will receive the action based on current criteria. For example, if the result of a criterion that checks if the action is active returns 100%, the component did not filter out any audience members.

The filtering process that happens in this simulation is similar to a funnel filtration for every engagement policy condition type. However, within an engagement policy condition type, each of the criterion is applied separately to the corresponding input audience.

In this case, the total number of customers in the audience is 1000. For this audience, the number of customers who received offers is 182. This simulation run also provides details of the audience filtration that happens with each engagement policy condition. When the eligibility condition is applied, 979 customers qualify for the offers. Within the eligibility criteria, the first condition does not filter any customers, as all qualify. When the second condition is applied, only 979 customers pass through. Thus, the intersection of these two conditions, 979, is the final number of customers who pass through to the eligibility level.


In this case, there are no eligibility conditions defined at the action level. If there are, the final result is the result of the eligibility conditions at the group level, combined with the results at the action level.

The output population of the eligibility level is passed on as the input to the next level.

The input population considered for the applicability condition is therefore 979. When the applicability condition is applied, 344 customers qualify for at least one offer out of the 979 customers. Note, since there are no specific action-level applicability conditions, the result of the individual group-level condition is the same as that of the overall applicability component level, 344.

The output population of the applicability level, 344 customers, is passed on as the input to the suitability level.

When the Suitability group-level condition is applied, 215 customers qualify for the offers out of the 344 customers. If you check the engagement policies at the action level, notice that there are additional suitability conditions. When all of these conditions apply, the total number of customers who qualify for at least one offer is 182.



At the bottom, you can also view the number of customers who qualify for the offers. When only engagement policy conditions are considered in the simulation run, the number you see here is the number of customers who qualify for at least one action. Notice that the Premier Rewards card offer is rarely presented. Let's try to find out why this is so.

Open the Premier Rewards card to view the audience simulation filtering at the action level.

Offers		
4 Offers (4 with specialized policies)		
Name	Specialized policies	
Rewards card	E	57
Rewards Plus card	S	168
Premier Rewards card	S	8
Standard card	E	57

You can choose the same audience simulation to view the action-level filtration details and investigate which engagement policy condition is causing the current outcome.

Offer: Premier Rewards card [Available]
Sales • CreditCards • PremierRewardsCard PegaCRM-Artifacts:01-01-01

Save as Delete Actions

Details **Engagement policy** Treatments Flow Test History

A Engagement policy

Open stored flow image
Refresh
Audience simulation...
View past simulations


Offer: Premier Rewards card [Available]
Sales • CreditCards • PremierRewardsCard PegaCRM-Artifacts:01-01-01

Details **Engagement policy** Treatments Flow Test History

Audience simulation
Simulation
Select... or Create simulation
Select...
EngagementPolicyTest-1
Eligibility ?

These numbers show how the engagement policies that are inherited from the CreditCards group are filtered. At the bottom, you can also view the final number of customers who qualify for the specific offers once the action-level engagement policy is applied.

Notice that at the suitability-level, only eight customers qualify for the Premier Rewards card offer.



This is because the configured engagement policy criterion is too narrow. The results of this simulation show us that from the 215 customers who passed the eligibility and applicability conditions, all qualify for the group-level condition, but only eight customers qualify for the action-level condition. As the action-level suitability condition is too narrow, a very small number of customers are presented with the offer. The bank then decides to reduce the CreditScore from 825 to 750 to ensure the Premier Rewards card offer is also presented to the customers.


Now, modify the credit score condition to reflect the correct value. Then, rerun the same audience simulation at the group level. Notice that now, 35 customers are presented with the Premier Rewards card offer.

The screenshot shows an 'Offers' interface with a header bar containing the title 'Offers' and a menu icon. Below the header, there is a section titled '4 Offers (4 with specialized policies)'. A table lists the offers and their counts:

Name	Specialized policies	Count
Rewards card	E	57
Rewards Plus card	S	168
Premier Rewards card	S	35
Standard card	E	57

Let's now run an audience simulation with engagement policies and arbitration in scope.

When you select engagement policies and arbitration as the scope, the **Include a second-pass simulation without business levers (for Scenario Planner)** check box is selected by default. When this option is selected, two simulations run in parallel. One simulation is based on your configuration in the Next Best Action Designer, with levers. The second simulation is based on your configuration in the Next Best Action Designer, without any levers. This simulation is used to determine the opportunity cost that is associated with applying levers. You can then use these two simulations in the Scenario Planner simulation later for analysis. It is a best practice to leave this option selected.



Create simulation

Audience *
SampledCustomers

Name
ArbitrationTest

Context
Customer

Population ?
Not available for this audience type

Next-Best-Action scope ?
☐ Engagement policies only
☒ Engagement policies and arbitration

☒ Include a second-pass simulation without business levers (for Scenario Planner)

Cancel Run


Note that the simulation results are different. This is because the simulation includes arbitration, adaptive analytics, treatment and channel processing, and constraints. That is, this simulation result shows the number of customers who will receive an action as their top action. Thus, the numbers at the group level and the action level tally. That is, the sum of all customers (12+118+22+30) is 182.

This demo has concluded. What did it show you?

- How to configure an audience simulation.
- How to view the simulation filtering details at the group and action levels.

Decision funnel explanation

This demo shows you how a decision funnel explanation simulation is auto-created for every audience simulation that is run at the Next-Best-Action Designer level. Every time an audience simulation is run, a decision funnel explanation simulation is created in the background. Now, navigate to the **Simulations** landing page to view the auto-created funnel simulation.



This is the auto-created funnel simulation. Alternatively, you can also create a simulation run from here and show the simulation results in the **Engagement policies > Audience Simulation**.

Name	ID
EngagementPolicyTest	S-33012

To view the funnel simulation, open the simulation.

Notice that the **NBA_Sales_CreditCards** strategy is used to run the simulation. This strategy is used is because you selected **Engagement only** in the audience simulation, without the arbitration.

Setup details	
Run simulation on revision version	
Strategy	
NBA_Sales_CreditCards	
Context	Results in
Data-Decision-Request-Customer	CRM-SR-Sales-CreditCards
Audience	Simulation ID
Data set	EngagementPolicyTest-1
Rule name	
SampledCustomers	
Purpose	
Explain	

When a funnel simulation is created, a set of reports is added to the run results. The simulation reports provide detail about the filtering that occurred.

Report	Report category	Output
Top level view	Simulations	ExplainDetails
Champion Challenger - Drill down view	Simulations	ExplainDetails
Proposition filter - Drill down view	Simulations	ExplainDetails
Prioritization - Drill down view	Simulations	ExplainDetails
Switch - Drill down view	Simulations	ExplainDetails

You can download these reports offline to analyze the outcome.

Top level view Edit report Actions									
Filtered by: (Report Sub Type = toplevel or (Report Sub Type = Overview and Rule Object Class = Rule-Decision-Strategy)) and Any Subject Type and Any Name									
Id	Group	Name	Strategy	Component	Type	#	#	P...	
SalesCreditCardsPremierRewardsCardNBA_AllIssues_AllGroups_CustomerResults					Output	10001,000	100.00		
SalesCreditCardsPremierRewardsCardNBA_Sales_CreditCards					Customer Actions	10001,000	100.00		
SalesCreditCardsPremierRewardsCardNBA_Sales_CreditCards					NBA_Sales_CreditCards_AllActions_PrimarySub Strategy	10001,000	100.00		
SalesCreditCardsPremierRewardsCardNBA_Sales_CreditCards_AllActions					NBA_Sales_CreditCards_A_All	10001,000	100.00		
SalesCreditCardsPremierRewardsCardNBA_Sales_CreditCards_AllActions					NBA_Sales_CreditCards_E_All	10001,000	100.00		

If you do not want to re-run the simulation in the future, you can wrap up the simulation. Any attached reports remain available for review.

Decision funnel (S-33012) PENDING-REVIEW

Actions

☐	Report	Report category	Output	
☐	Top level view	Simulations	ExplainDetails	⋮
☐	Champion Challenger - Drill down view	Simulations	ExplainDetails	⋮
☐	Proposition filter - Drill down view	Simulations	ExplainDetails	⋮
☐	Prioritization - Drill down view	Simulations	ExplainDetails	⋮
☐	Switch - Drill down view	Simulations	ExplainDetails	⋮

+ Add report

> Setup details

Wrap up

All the resolved simulations are accessible from the **Include resolved simulations** section.

Simulation Testing

Search

by Name, ID, Type, Strategy, Audience or Assigned to

☒ Include resolved simulations

Name	ID	Type	Strategy
EngagementPolicyTest	S-33012	Decision funnel	NBA_Sales_CreditCards

This demo has concluded. What did it show you?

- How to view the auto-generated decision funnel explanation simulation.

Analyzing the results of simulations using reports

Description

Learn how to unit test a strategy using distribution test simulation. Analyze the simulation reports to understand the impact of creating or updating an action and the implication of the same.

Learning objectives

- Explain what a distribution test simulation is
- Analyze results of simulations using reports

Distribution test


What is a distribution test simulation?

A distribution test simulation enables you to unit test a decision strategy. You can execute a simulation on an input audience, output the results to a database table and view the results using out-of-the-box simulation reports.


You can use these reports to inspect the number of actions offered in general, the number of actions across various channels, the average propensity of the actions, the ranks at which each action is offered, and so on. This is particularly useful when creating or changing an action and understanding the possible implications of that change.

Let's look at the Action distribution and Action priority reports as examples of the available reports.

Action distribution shows the volume per action. That is, how many times each action was output by the decision strategy. In the report you can also see the total volume of actions.



The Action priority report shows the average priority of each action. Here you can see that the RewardsPlus card offer has the highest average priority, meaning the expectation is that this will be the best performing action. If an action has a low average priority, the chance of this action being offered may also be lower.



Analyzing the results of simulation using reports


A distribution test enables you to run a simulation on a set of customers and generate various out-of-the-box reports to analyze the distribution and prioritization of actions.

Transcript


This demo will show you how to run a simulation and output the results to a database table and view the results using out-of-the-box simulation reports.

U+, a retail bank, recently updated one of their engagement policy conditions. They would like to unit test their changes by running a distribution test to check that the overall distribution is within the desired ranges.

This is the Pega Customer Decision Hub™ portal. Navigate to the **Simulations** landing page to create a new **Distribution test**.



To create a simulation, select the top-level strategy on which you would like to run the simulation. Use the top-level next-best-action strategy, as this simulation pulls in data from all relevant actions and treatments.



Then, select the input population on which you want to execute the simulation. Now, rename the simulation. This will help you identify specific simulation runs easily. Next, configure the output. Create a new output destination. Give the output destination a name.

You want to write the output to a database table, so select the **Database Table** option. Once you configure a new output destination, all the underlying artifacts are created for you (including the database table with the output columns). At this stage, you can select multiple output destinations, either another database table, or a Visual Business Director data source. For now, just continue with the **SimulationResults** database table.

Create new output

Name *

SimulationResults

Type *


Database Table

When one or more output destinations is a database table, you need to select which reports you want to have available for each destination. Select the output database table, then the report category, and finally one or more reports. Note that there are several out-of-the-box reports available. A decisioning architect can extend this list to create new reports specific to the business' needs. For now, just choose two of the standard reports: **Action distribution** and **Action priority report**.


☐	Report	Report category	Output	
☐	Action distribution	Simulations	SimulationResults	⋮
☐	Action priority report	Simulations	SimulationResults	⋮
+ Add report				

Run the simulation test.

When the execution is complete, you can inspect the results of the simulation by opening the assigned reports. You can also add reports to a simulation once the run is complete to view the output written to the database. This Action distribution report shows the volume of the individual actions. That is, how many times each action would be displayed to the customers. In the table you can also see the total volume at the group and business issue levels.



The Action priority report shows the average priority of each action. Here you can see that the RewardsPlusCard offer has the highest average priority, meaning the expectation is that this will be the best performing action. If an action has a low average priority, it may also be offered less frequently.



This demo has concluded. What did it show you?

- How to execute a distribution test simulation.
- How to output simulation results to a database table.
- How to assign out-of-the-box reports to the simulation output.

Detecting unwanted bias

Description

Learn how to detect unwanted bias in engagement policy conditions and avoid discrimination while keeping the business requirements and regulations in mind.

Learning objectives

- Explain what an ethical bias simulation is
- Detect unwanted bias in engagement policies

Ethical bias

An ethical bias simulation enables you to test your engagement policies for unwanted bias. That is, you can check if your conditions discriminate based on age, gender, ethnicity, or any other attributes specific to your business scenario. This is particularly useful during the testing stage, while developing and changing actions.

Transcript

This video explains what ethical bias is and how you can avoid unwanted bias in your engagement policies.

Ethical bias testing helps check your engagement policies for unwanted bias. That is, you can test if your conditions discriminate based on age, gender, ethnicity or any other attributes specific to your business scenario.

An ethical bias policy forms the base of an ethical bias simulation. This policy will include the bias fields and the thresholds for each field. You can include any property from your customer class. For instance, age and gender properties are typical properties you might want to include in bias testing.

As age is a numerical field, a Gini coefficient is used to calculate the bias. This is a method of measuring the statistical inequality of a value distribution. A Gini coefficient of 0 represents perfect distribution equality. You can select a warning threshold between 0 (warn if any bias is detected) and 0.7 (warn only if very high bias is detected). You can also choose not to check for bias within a particular business issue.



As Gender is a nominal value, a disparate impact is used to determine bias. A disparate impact is used to calculate bias for categorical fields by comparing the number of customers who were selected for an action to those not selected for an action, and then correlating the result to the selected bias field.

A disparate impact of 1 represents perfect distribution equality. You can select a warning threshold between 1 (warn if any bias is detected) and 0.50 - 2.00 (warn only if very high bias is detected). You can also choose to ignore this bias field for a particular issue in your business structure.

Even though no bias is ideal, it is not practical to set the threshold at that level as that would block any actions from reaching customers. Hence, in real business scenarios, a little bias is often allowed to prevent customers from receiving no actions.

Detecting unwanted bias in engagement policies

Ethical bias testing checks your engagement policies for unwanted bias. That is, you can test if your conditions discriminate based on age, gender, ethnicity, or any other attributes specific to your business scenario.

Transcript

This demo will show you how to run an ethical bias simulation and identify any unwanted bias in the engagement policy conditions.

U+, a retail bank, recently updated their engagement policy conditions to present credit card offers to qualified customers. Now they would like to run an ethical bias simulation to ensure there is no unwanted bias based on age or gender before they push the changes to the live production environment.

This is the Pega Customer Decision Hub™ portal.

To create a simulation, first configure the ethical bias policy. This policy will include the bias fields and threshold. You can select any property from your customer class. In this case, you will use the age and gender properties for the test, which are the most commonly used properties for bias testing. The age property value is a number. In the **Add bias** field window, add age as a bias field.

Next, add gender as a bias field.

Ethical Bias Policy

Bias fields

Bias threshold

Bias fields

Define fields for which bias will be checked.

Context

Data-Decision-Request-Customer

+ Add bias field

Field name	Field type	Bias measure
Customer.Age	Integer	Numeric
Customer.Gender	Text	Symbolic

Now, on the **Bias threshold** tab, review and configure the bias threshold settings for each property you selected. This configuration is done at the business issue level. This means, for every business issue, you can decide how much bias to allow. For instance, for the risk issue, you might want to disallow any bias, whereas for grow, you might allow some bias, if you designed a special offer for a specific age category or a specific gender. The bias threshold measurement depends on the type of field that you select. By default, the threshold is set for no detection for these fields.

In this case, the bank does not want to discriminate on age, so select the appropriate threshold value.

The bank also does not want to discriminate on gender, so select the appropriate threshold value. Save the changes.

Navigate to the **Simulations** landing page to create and run an **Ethical bias** simulation. On the Simulation Creation page, select the top-level strategy on which you would like to run the simulation. Then, select the input population on which you want to execute the simulation.

Now, rename the simulation. This will help you easily identify specific simulation runs.

Note that the simulation results are output to the **Insights** data set. This is an internal dataset where the results will be available for an hour by default. The ethical bias simulation has two bias reports automatically available as output. The results can be examined post simulation run.

Assign reports to outputs			
Reports		Configure	
Output	Report category	Report	
Insights	Bias	Bias report	
Insights	Bias	Detailed bias report	

Run the simulation. Once the run is complete, there will be an indication if any unwanted bias is identified. In this scenario, although business did not want any bias, it seems some bias has been detected.

Ethical bias check

 Bias was detected for one or more actions. Please check the bias report to learn more.

Open the generated reports and view the information in detail to understand where the bias was detected. You can sort the report in the bias detected column. Notice that there is bias detected on age. You can check the **Bias value** against the **Confidence interval** to check if the bias value is within the bias threshold range.

Navigate to Next-Best-Action Designer to view the engagement policy conditions. Note that there is an eligibility condition that uses age. Credit cards are valid only for customers with an age greater than 18. This is a hard eligibility rule that cannot be ignored. It is likely that bias has been detected due to this condition.

Engagement policy

E Eligibility ?

(isCustomer is true)

and (Age is greater than 18)

A Applicability ?

(Has Cards is equal to N)

S Suitability ?

No group criteria defined

Based on the bank's policies and regulations, you must review the bias threshold and decide how much bias the business will allow.

To modify the policy, open the Ethical Bias policy. Due to the eligibility condition, which cannot be ignored, business decided to allow a maximum bias of 0.1 Gini. So, increase the bias threshold, and re-run the same ethical bias simulation.

Ethical bias check

✓ All decisions are compliant with your bias policy.

Since the age bias threshold was increased, no bias should be detected.

Now, let's run an ethical bias simulation at the **Trigger_NBA_Sales_CreditCards** strategy level. This strategy includes engagement policy conditions, arbitration, adaptive analytics, constraints, and treatments and channels processing.


Note that bias is detected on gender. However, it is at the CreditCards group level, with only engagement policies in scope, thus, there was no bias.

Bias detected	Issue ¹ ↑	Group ² ↑	Action ³ ↑	Bias field ⁴ ↑	Category
false	Sales	CreditCards	PremierRewardsCard	.Customer.Age	--
true	Sales	CreditCards	PremierRewardsCard	.Customer.Gender	F
true	Sales	CreditCards	PremierRewardsCard	.Customer.Gender	M
false	Sales	CreditCards	RewardsCard	.Customer.Age	--
false	Sales	CreditCards	RewardsCard	.Customer.Gender	F
false	Sales	CreditCards	RewardsCard	.Customer.Gender	M
false	Sales	CreditCards	RewardsPlusCard	.Customer.Age	--
true	Sales	CreditCards	RewardsPlusCard	.Customer.Gender	F
true	Sales	CreditCards	RewardsPlusCard	.Customer.Gender	M
false	Sales	CreditCards	StandardCard	.Customer.Age	--
false	Sales	CreditCards	StandardCard	.Customer.Gender	F
false	Sales	CreditCards	StandardCard	.Customer.Gender	M

Navigate to **Prediction studio** to view the **Web Click Through Rate** adaptive model. Now, open the Premier rewards card model report to view the predictors.

Predictors	Score distribution	Trend			
Predictors	Status	Type	Performance (AUC)	Range/Symbols(#)	Bins(#)
Customer.AverageSpent	Active	Numeric	76.84	[1001.42; 1999.4]	12
Customer.Age	Active	Numeric	69.58	[19.0; 80.0]	9
Customer.Gender	Active	Symbolic	61.24	2.00	2

Notice that gender is one of the top predictors. As gender is a predictor, and the model is learning constantly, bias is detected even though there is no engagement policy condition on gender. Open the predictor to view the customer responses. Here you can see that there is a clear difference in the number of times the offer was presented to male customers versus female customers.



If you do not want the adaptive model to learn based on gender, ensure gender is not included as a predictor.

This demo has concluded. What did it show you?

- How to configure an ethical bias policy.
- How to run an ethical bias simulation.
- How to view the auto-generated ethical bias reports to understand the bias threshold deviation.

Analyzing the effect of business changes using Pega Scenario Planner

Description

Learn how Scenario Planner allows you to easily simulate “what-if” scenarios to accurately forecast results, optimize strategies to meet specific goals, and explore the potential trade-offs of each option.

Learning objectives

- Explain what a scenario planner simulation is
- Analyze the effects of business changes using scenario planner

Pega Scenario Planner

Pega Scenario Planner enables users to analyze simulation runs to determine how close the results are to meeting expected targets. Using Scenario Planner, you can browse the results of past next-best-action audience simulations on one page and compare the results to real outcomes in a simple, transparent interface. Also, you can view the results for all issues and groups or drill down to the details of a single action and channel. This view is useful when planning new changes or optimizing existing parameters.

Transcript


This video explains how Scenario Planner enables a business to meet its business priorities.

Scenario Planner allows you to easily simulate “what-if” scenarios to accurately forecast results, optimize strategies to hit specific goals, and explore the potential trade-offs of each option.

Scenario Planner also estimates the value that the subsequent next-best-action run creates using its current configuration. For example, you can observe whether prioritizing a specific action by increasing its business weight might result in an opportunity gain or an opportunity cost.



The Scenario Planner landing page shows three useful Key Performance Indicators: projected value capture, responses from Interaction History (including accepts, sends, and impressions), audience simulation counts (or projected reach), and projected accepts based on adaptive models.



Depending on the selected filters, the Scenario Planner view can display the results for a specific channel, action, group, or issue.

The projected value capture represents the estimated value that the upcoming next best actions (outbound run and inbound impressions) will bring with the currently selected configuration and the opportunity cost or gain of making changes to the Next-Best-Action Designer configuration.

The projected value capture is calculated using the action propensity, the action value of the top offer for each customer in the sample audience, and the projected accepts.



Here is an example of how the projected value capture is calculated.

Consider Customer1, who qualifies for three actions. The action propensity of each action is determined by the Artificial Intelligence (AI) engine. Also, the business has assigned an action value for each action.

With the action propensity (p) and the action value (v), the expected action value is computed using the formula $p \cdot v$. The expected action value is the anticipated average revenue that the business is likely to make from an action.

For example, if an action has a propensity of 0.09, and a business value of 51, the expected value of the action is $0.09 \cdot 51 = 4.59$. If the business value is not assigned for an action, the value is considered 1 by default.

The AI ranks the actions based on the expected action value. That is, the higher the expected action value, the better the rank.

Simulations are usually executed on a set of customers, also referred to as the audience. Similar to the first customer, the action details of all other customers are computed.

In this scenario, where there are no levers, the projected value capture is calculated by using the number of times an action is a top offer and the projected accepts. In the example, Action1 is the top offer for two customers. As a result, the projected accepts is the sum of the action propensities. That is, $0.09 + 0.10 = 0.19$.

This value means that the expected value capture for Action1 is the sum of the expected action value for Customer1 and Customer4. That is, $4.59 + 5.10 = 9.69$. The projected value capture is the overall sum of the actions' expected value capture.

Now, consider another situation where the business has assigned levers to the actions to boost them based on the business requirement. For example, the business has assigned a business lever of 0%, 30%, and 10% to Action1, Action3, and Action2, respectively.

Applying levers affects the expected value and the action ranking. The adjusted expected value of the action is calculated using the formula, $p \cdot v \cdot l$. That is, the propensity times the value times the lever. For example, if an action has a propensity of 0.08, an action value of 46, and a lever of + 30%, the new expected value of the action is 4.78.

Based on the adjusted expected value, the order of the actions changes. Notice that Action3 becomes the top 1 action for Customer1.

Similarly, the adjusted expected value is determined for all the other customers in the sample audience.

As in the first example, the expected value capture and the projected accepts are calculated using the number of times an action is the top offer, taking into account the new order. Per the example, Action1 is the top offer for Customer4 only. As a result, the projected accepts is 0.10.

This value means that the expected value capture for Action1 is the same as the expected action value for Customer4, 5.10. The projected value capture is the overall sum of all actions' expected value capture.

Note the projected value capture for each scenario. The projected value capture without levers is more than with levers. That is, if the levers are applied, there is an opportunity cost involved. If the projected value capture without levers is lower, there is an opportunity gain involved.

The Scenario Planner results help a business adjust Next-Best-Action Designer based on the metrics to meet its business priorities.

Pega Scenario Planner enables you to:

- Analyze simulation runs to understand how the business is performing
- Understand the opportunity cost or opportunity gain involved in adjusting the Next-Best-Action Designer configurations

Analyzing the effect of business changes using Pega Scenario Planner

Pega Scenario Planner allows you to easily simulate “what-if” scenarios to accurately forecast results, optimize strategies to hit specific goals, and explore the potential trade-offs of each option.

Transcript


This demo shows you how to use Scenario Planner to increase your understanding of the effect of business changes.

U+ Bank, a retail bank, has recently implemented a project in which credit card offers are presented to qualified customers when they log in to the self-service web portal. The bank wants to understand the opportunity cost or gain in using levers to adjust AI prioritization.

You can use Scenario Planner to analyze the opportunity cost or gain associated with applying levers. Scenario Planner also lets you understand the scope, to optimize for increased revenue, and trade-offs from adjustments due to levers.

This is the Pega Customer Decision Hub™ portal.

In the navigation pane click **Discovery** to create a Scenario Planner simulation with second-pass simulation enabled.




To analyze the simulations using Scenario Planner, select the scope of the run. The scope can be at the all issues level or for a particular issue and group. In this case, select the **Sales** issue and the **CreditCards** group. Now, select the scenario planner simulation run that you executed earlier.

You can also select a simulation without levers or another scenario planner simulation for a different next-best-action configuration and compare the results in general or for specific channels. For now, apply to view the results.


Scope	Simulation
Sales / CreditCards ▾	Sales_CreditcardsBaseline (... ▾)

A Scenario Planner simulation displays metrics that show the projected value capture based on the configurations within Next-Best-Action Designer, the actual responses from the interaction history, and the projected reach and responses calculated from the results of the simulation.



In this tile, you see the projected value that represents the estimated value that the upcoming next-best-actions (outbound run and inbound impressions) bring with the currently-selected configuration and the opportunity cost or gain of applying the levers. Prioritizing a specific action by increasing its business weight can result in increased value but often comes at the expense of the projected value capture for other actions, issues, and groups.


In this example, the pointer is under the bar, as there is an opportunity gain when levers are applied.



The projected value capture is \$43.71K without levers, represented by the blue bar until the pointer. In addition, there is a \$14.52K increase in value due to the implementation of business levers, amounting to a total of \$58.23K projected value capture.


If there is an opportunity cost involved, the bar in this section has a grey portion. The pointer in that area indicates a scope for value increase without levers.

This tile shows the number of accepted actions versus the aggregate number of outbound sends and inbound impressions over the last 30 days. That is, the actual or historical accepts and presentations versus the number of times an action was sent or presented to customers. In this case, there are 18.1K accepts out of 170K impressions.



You can change the timeframe for historical data from the last 30 days to the last 60, 90, 180, or 365 days. You can also aggregate data from the last full quarter up to a full year.

In this tile, you see the aggregate number of sends or impressions, with the projected accepts, that the next outbound run or presentation will accumulate with the currently selected configuration. The business can expect 429 accepts out of the 4.6K sends or presentations. This value is calculated based on the simulation run results.



The following table displays the results per the selected simulation, issue, group, action, and channel. In this case, you see the results for the actions in the CreditCards group.

The table provides the number of actual responses per card for the selected time frame, the projected reach of each card (that is, how many times a card is sent or presented to customers), and the projected responses for each card (the number of accepts that are expected for each card).


Action	Responses (available from last 30 days)	Projected reach	Projected responses
Premier Rewards card ⓘ	650 (7% of 10,000) 	1,080	70
Rewards card ⓘ	870 (9% of 10,000) 	992	86
Rewards Plus card ⓘ	5,350 (11% of 50,000) 	1,410	151
Standard card ⓘ	11,200 (11% of 100,000) 	1,092	122

For example, if the Premier Rewards card is sent or presented 1,080 times to customers, it is likely that 70 customers might accept the offer.

You can also analyze the opportunity cost or gain of each action.
Now, select the simulation without levers and compare the results.

Scope	Simulation	Comparison simulation
Sales / CreditCards ▾	Sales_CreditcardsBaseline (... ▾	(S-1) with no levers ▾

Note that in comparison mode, the Scenario Planner indicates the projected value and reach in green, as there is an increase when the lever is applied.
You can click the individual section menu to view the projected data comparison.



Similarly, in the table, the projected reach and responses are indicated per action with a green and red arrow to show which action is presented more and which action is presented less due to the business lever implementation. This information gives the business an idea of how well the actions might perform and how to decide on their trade-offs and prioritizations.

Now, navigate to the Results tab of the action without levers to view the effect on the business. Notice that there might be an opportunity cost if the levers are not applied. Also, the projected response is lower, leading to an increase in low-performing actions.

In this case, the business can see how implementing business levers might increase its revenue by presenting the right offers to the right customers.

Scenario Planner enables you to export the results to a CSV file for detailed analysis.

The CSV of the comparison simulation contains the deltas for all the values and metrics color-coded for each simulation, similar to the UI. Also, the CSV file contains the simulation data for each of the simulations.

The CSV file of the single simulation has a single view sheet that shows all the values and metrics, as did the Scenario Planner simulation result.

You have reached the end of this demo. What did it show you?

- How to analyze simulations using Scenario Planner
- How to compare simulations using Scenario Planner
- How to interpret the effect of business changes
- How to export and analyze the simulation results