Decisioning Consultant

STUDENT GUIDE

PEGA® Academy

© Copyright 2021
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.pega.com

Mission: Decisioning Consultant
Product: Pega Customer Decision Hub™ 8.6
URL: https://academy.pega.com/mission/decisioning-consultant/v3
Date: 26 April 2021

Contents

One-to-one customer engagement....6
Next-Best-Action paradigm....7
One-to-one customer engagement paradigm....12
Next-Best-Action Designer....18
Next best action in an omnichannel environment....25
Next best action in a contact center....26
Next best action on digital....30
Next-Best-Action on digital....35
Business use case: Cross-sell on the web....36
Defining and managing customer actions....41
Action hierarchy....42
Managing business structure....52
Creating actions....54
Renaming actions....58
Presenting a single offer on the web....61
Real-time containers....62
Creating and configuring a real-time container....66
Presenting a single offer on the web....68
Defining customer engagement policies....73
Customer engagement policies....74
Defining eligibility, applicability, and suitability rules....79
Testing the next-best-action configuration with persona testing....84
Avoiding overexposure of actions....90
Contact policies....91
Defining action suppression rules....93
Arbitrating between actions....97
Action arbitration....98
Action prioritization with AI....104

Prioritizing actions with business levers....112
Essentials of always on outbound....114
Working of the always-on brain in outbound....115
Traditional marketing versus always-on outbound....120
Business use case: Always-on outbound customer engagement....125
Sending offer emails....128
Sending offer emails....129
Defining the starting population....138
Segmentation....139
Creating a segment....141
Defining an action for outbound....152
Creating an email treatment....153
Defining the action flow....170
Avoiding overexposure of actions on outbound....174
Avoid sending the same action for a time period....178
Limiting action volume on outbound....183
Volume constraints....184
Configuring a volume constraint....187
Sharing action details with third-party distributors....194
File and database templates....195
Writing action details to a file or database....200
Optimizing 1:1 customer engagement overview....210
Simulation testing....211
Analyzing customer distribution using Pega Value Finder....213
Pega Value Finder....214
Analyzing customer distribution using Pega Value Finder....215
Testing engagement policy conditions using audience simulation....220
Audience simulation....221
Running an audience simulation test....226
Analyzing the results of simulations using reports....236

Distribution test....237
Analyzing the results of simulation using reports....239
Detecting unwanted bias....243
Ethical bias....244
Detecting unwanted bias in engagement policies....246
Analyzing the effect of business changes using Pega Scenario Planner....255
Pega Scenario Planner....256
Analyzing the effect of business changes using Pega Scenario Planner....260

One-to-one customer engagement

Description

Familiarize yourself with the one-to-one customer engagement paradigm. Learn about the business problems Pega Customer Decision Hub™ can solve and discover how Pega's omni-channel AI delivers the right action during every customer interaction.

Learning objectives

- Explain the basics of the Next-Best-Action approach
- Describe the purpose of Next-Best-Action Designer and its user interface

Next-Best-Action paradigm

Introduction

The value of big data and analytics is fully realized when every customer conversation delivers exactly the right message, the right offer, or the right level of service to provide a great experience while maximizing the customer's value to the organization. With Pega Next-Best-Action, business experts develop decision strategies that combine predictive and adaptive analytics with traditional business rules to maximize this value.

Transcript

This is your customer. You want him to buy your products, use your services and have a great experience. And your competitors want the same thing. To compete, you have to take the right action at every customer touch, ensuring that each conversation delivers exactly the right message, offer and level of service. You want to provide a great experience, while maximizing the customer's value to your organization.

Artificial Intelligence, or AI, can help—if you can get past the hype. Pega has been using AI to create real business value for years, driving real-time decisions that deliver awesome

engagement on any channel and improving experiences for over 1.5 billion customers across the globe.

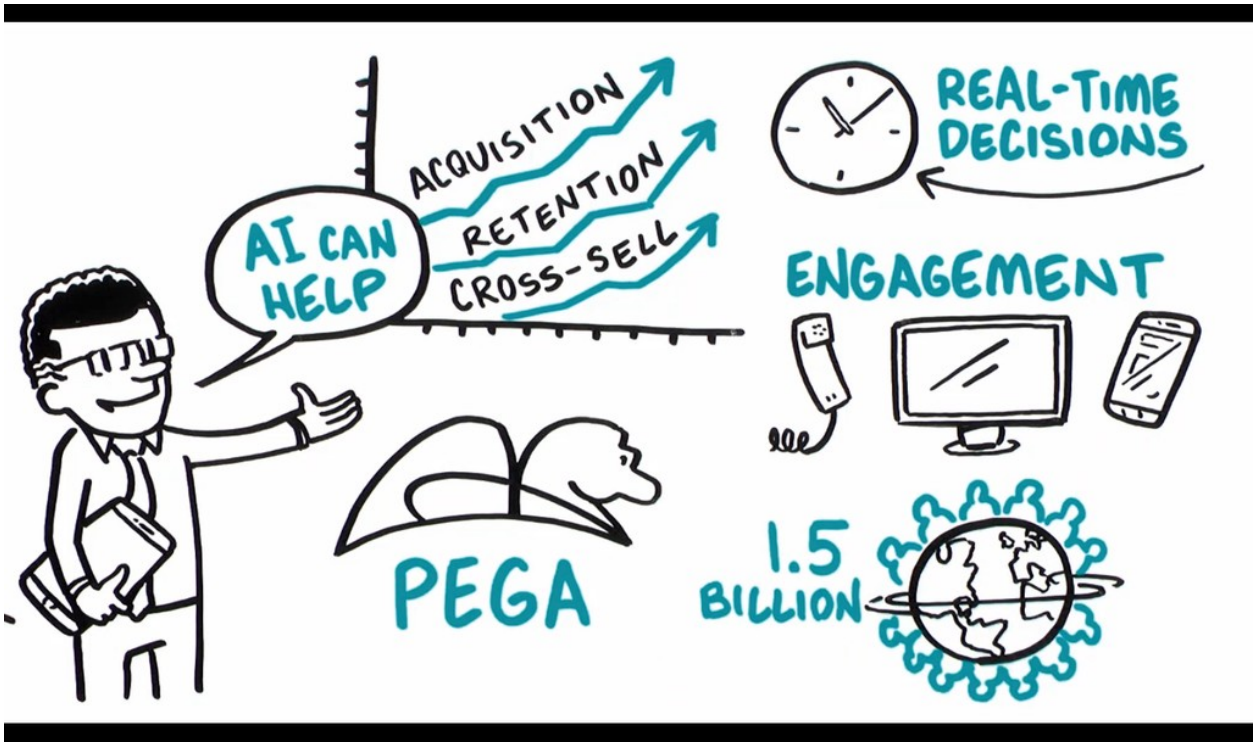

Pega's omni-channel AI delivers the right action at every customer touch by crunching millions of data points in real-time. Make an offer, initiate a retention plan, predict a problem before it happens. Every decision generates the next-best-action for your customer, and your business.

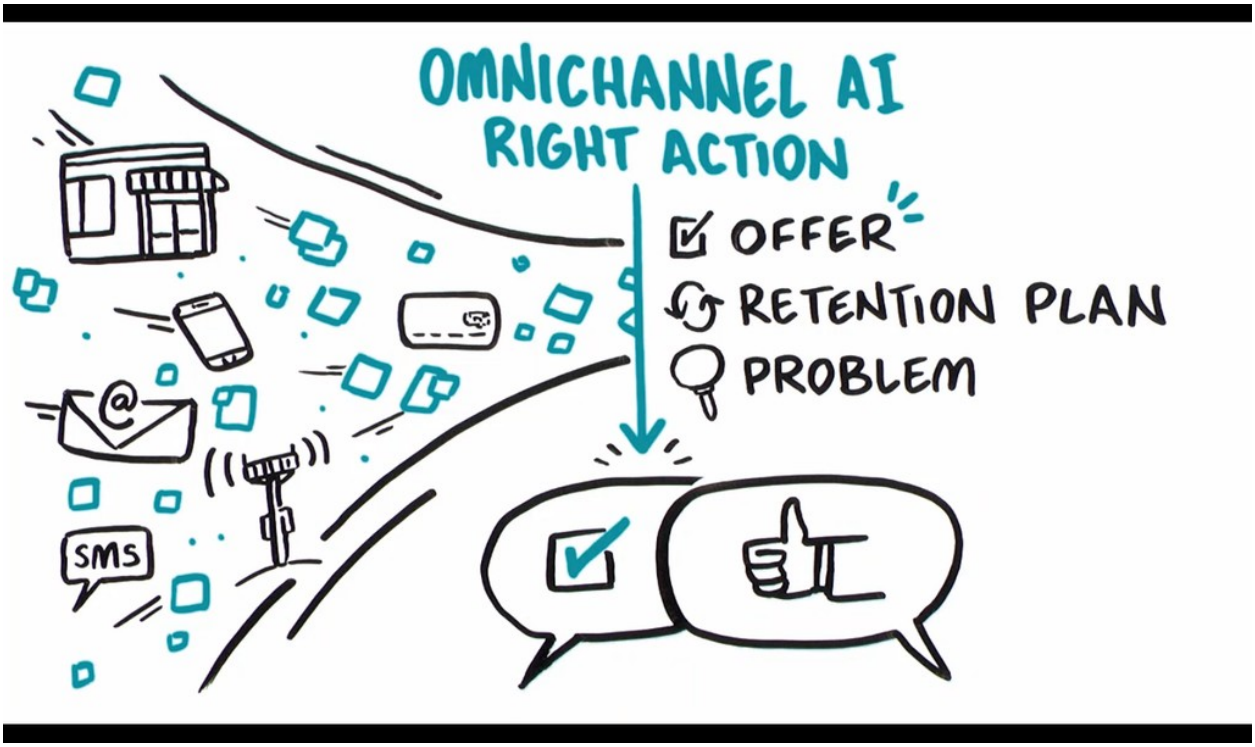

Pega's AI is built for business people, not scientists or developers. They design visual decision strategies that combine predictive analytics, algorithms developed through mining large sets of data, adaptive analytics, machine-learning algorithms that improve with each interaction, and traditional business rules that allow users to prioritize and arbitrate between decisions.

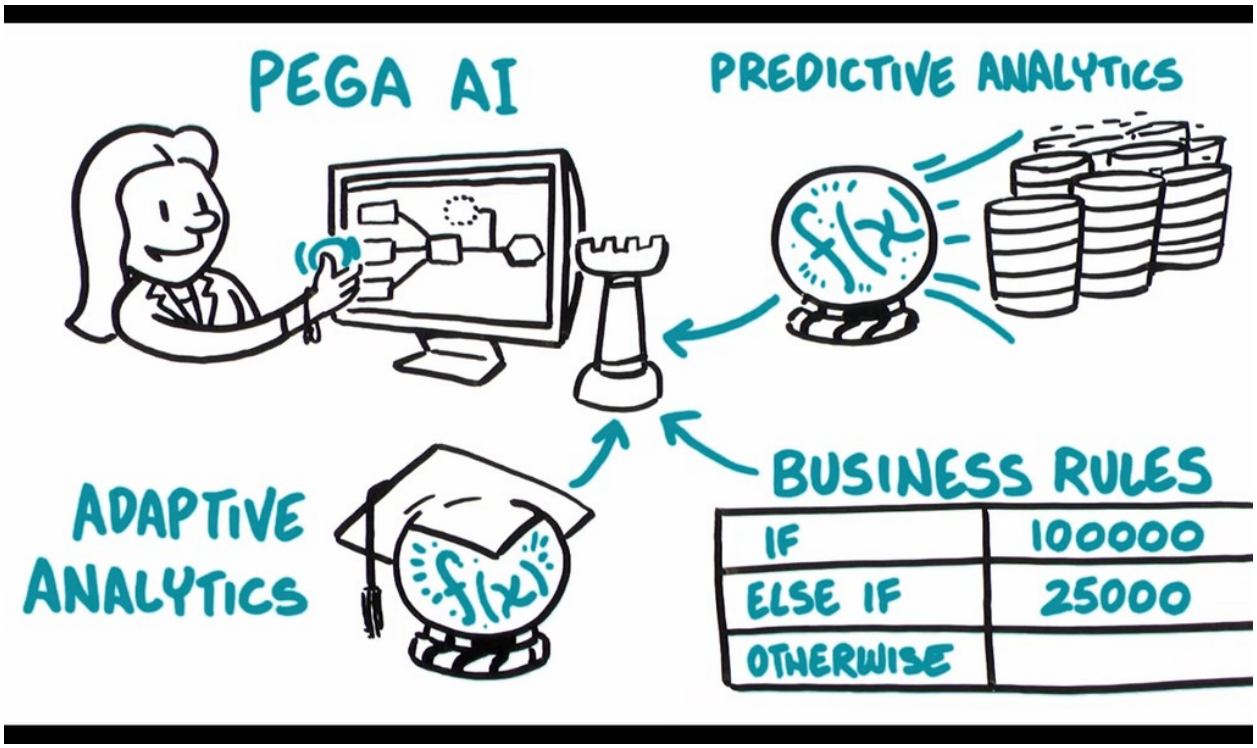

Pega uses the strategy to look across all the potential actions you may take with a customer, make an offer, initiate a retention plan, open a service case, place an ad, and ensure exactly the right action is taken at every interaction and it works across all channels to provide a consistent experience in a store, on the phone, on the web, mobile, with the chat bot, or just some crazy tech that hasn't even been invented yet.

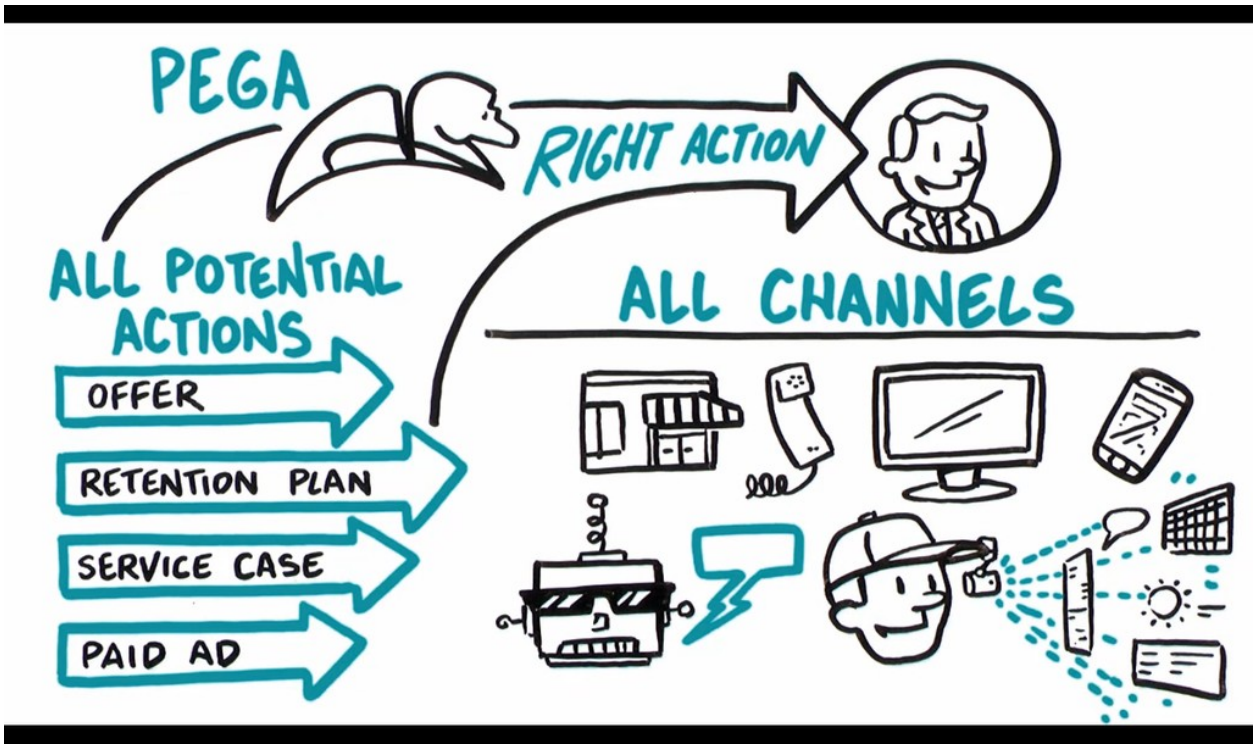

And Pega connects to streams like mobile locations or network events to detect patterns and drive the Next Best Action proactively. And strategies are completely contextual. Any change in the customer's context — a click, a reply, a location change, a Tweet — will trigger the Next Best Action. So, you can really listen to your customers and act accordingly.

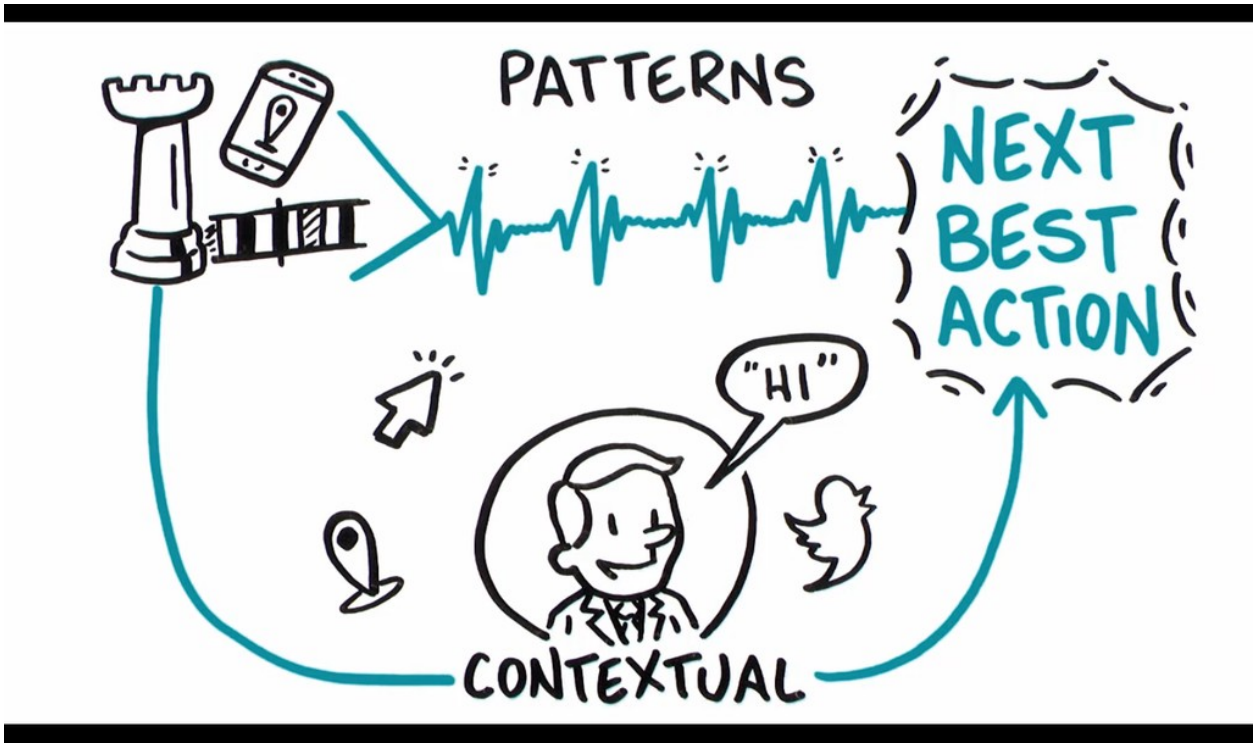

Pega's real-time, omni-channel AI puts the power in your hands, so you can optimize every customer interaction for experience, and value.

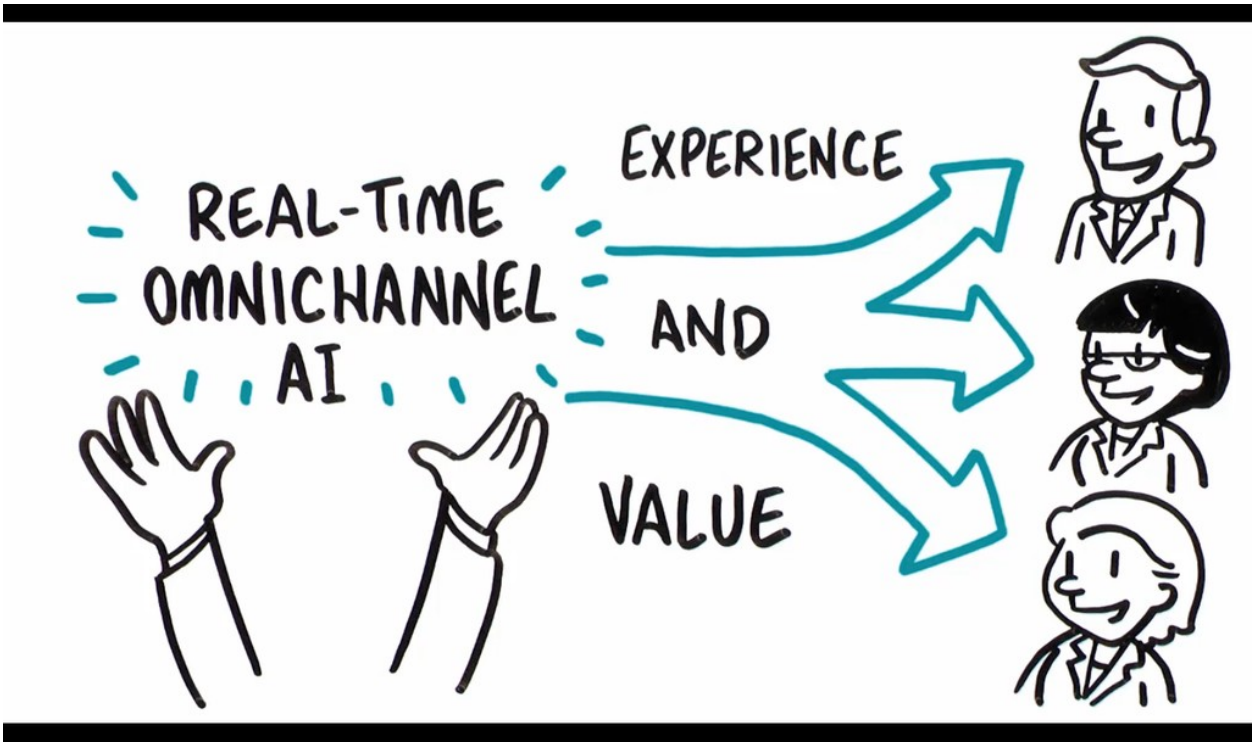

One-to-one customer engagement paradigm

Introduction

The optimal outcome of every customer interaction is to provide a great experience while maximizing the customer's value to the company. To achieve this, you have to be able to perform the right action in the right channel at the right moment for each customer. We call this capability, "1-to-1 Customer Engagement".

Transcript

In this video, learn about the 1-to-1 Customer Engagement paradigm and how the principles of Next-Best-Action are implemented using the Pega Customer Decision Hub™.

Customers are more empowered than ever before. As a result, they have very high expectations of the experiences they receive from their service providers. Their experiences must make sense within the context of their lives. This means they must be meaningful, consistent, and personalized across every channel they interact with.

In business, the optimal outcome of every customer interaction is to provide a great experience while maximizing the customer's value to the company. To achieve this, you have to be able to perform the right action in the right channel at the right moment for each customer.

We call this capability, "1-to-1 Customer Engagement".

1-to-1 Customer Engagement enables companies to transition their marketing away from a traditional one-to-many campaign-driven approach. A one-to-one approach allows companies to have consistent, contextual and relevant conversations with individual customers across any channel or touch point.

The key to achieving 1-to-1 Customer Engagement is an idea that's simple to conceive, but very difficult to execute: one centralized brain.

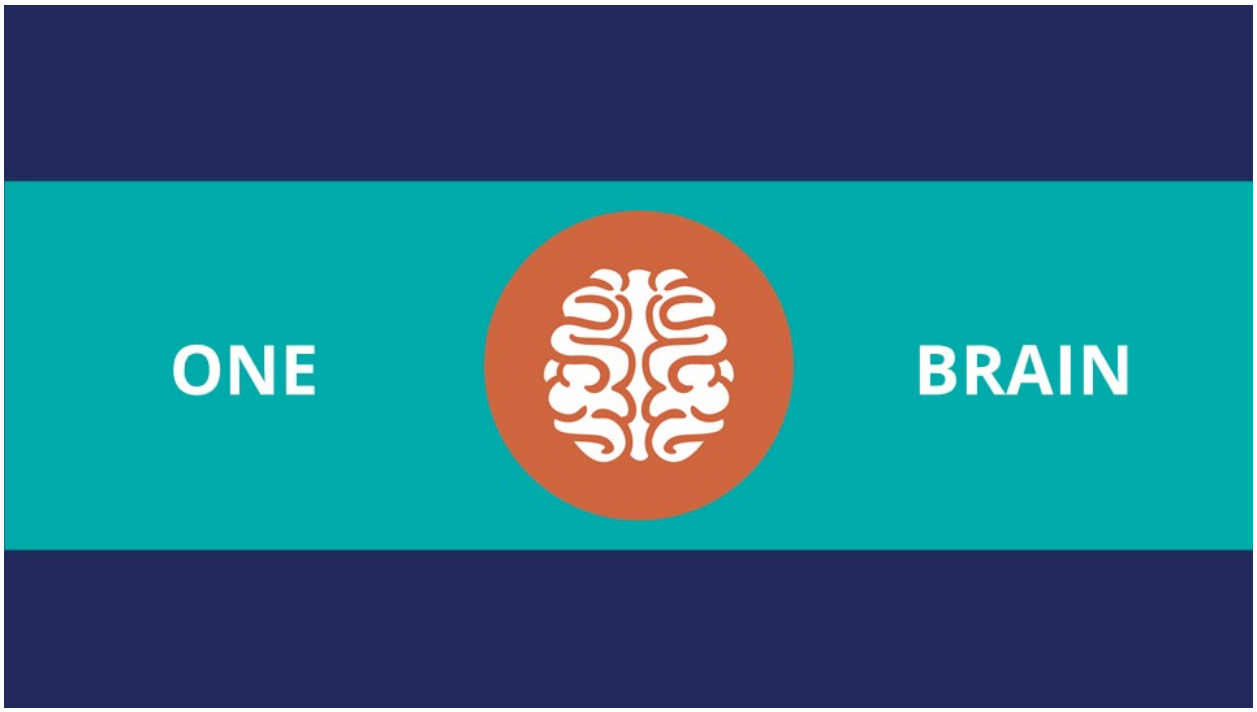

In other words, one piece of intelligence that acts as a single decision authority across your application ecosystem.

Each channel or system profits from this single source of customer intelligence and can leverage it to gain insights or perform relevant actions.

In Pega Marketing™, this centralized brain is called the Pega Customer Decision Hub, and it leverages AI to enable 1-to-1 Customer Engagement.

In Pega Infinity™, the Pega Customer Decision Hub forms the core of the customer engagement platform, which sits at the center of existing systems and channels in an enterprise.

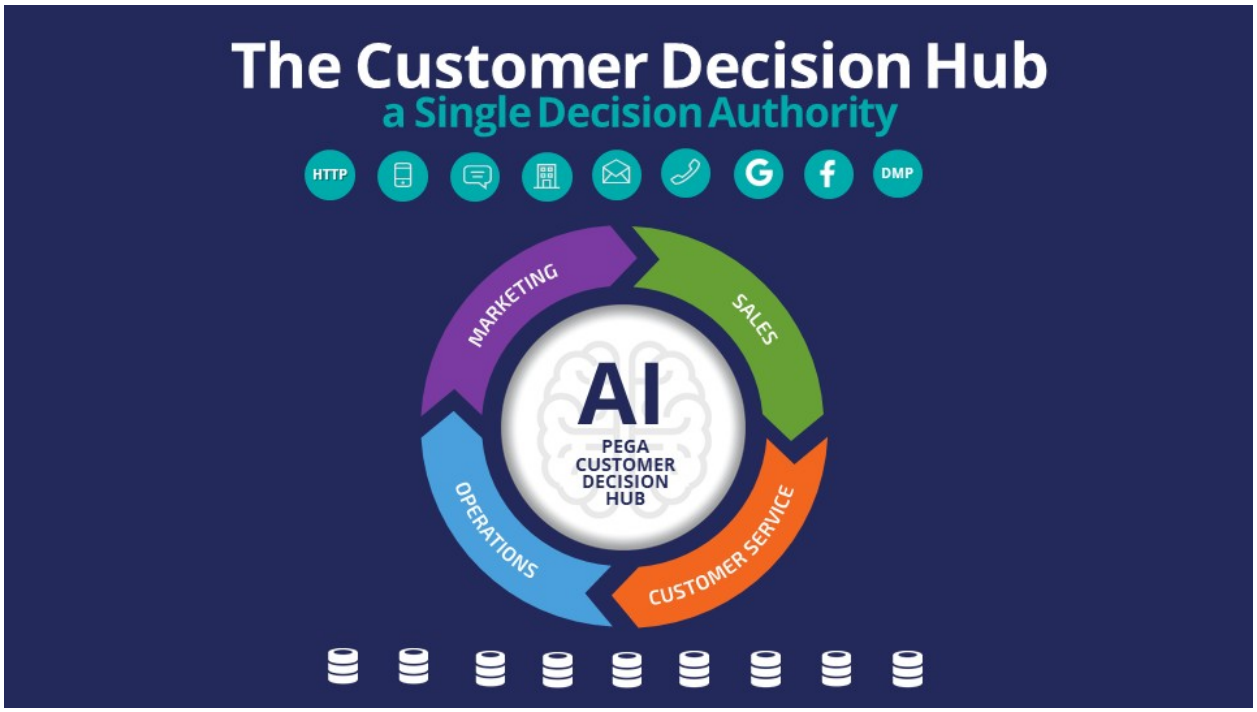

Data from every customer engagement across the enterprise is collected by the Brain and used to make predictions and decisions about every interaction in every channel.

Continuous learning and decision-making are the foundation of a 1-to-1 Customer Engagement solution.

The Customer Decision Hub combines analytics, business rules, customer data, and data collected during each customer interaction to create a set of actionable insights that it uses to make intelligent decisions. These decisions are known as the Next-Best-Action.

Every Next-Best-Action weighs customer needs against business objectives to optimize decisions based on priorities set by the business manager.

In the milliseconds before interacting with a customer, the Customer Decision Hub processes thousands of predictive and adaptive models to determine customer needs, considering the customer's immediate context to ensure the Next-Best-Action is relevant, timely, and contextual. These models can be propensity, risk, or churn models.

Next, the decision strategy considers business rules and matches those with the customer's context and higher-level business goals.

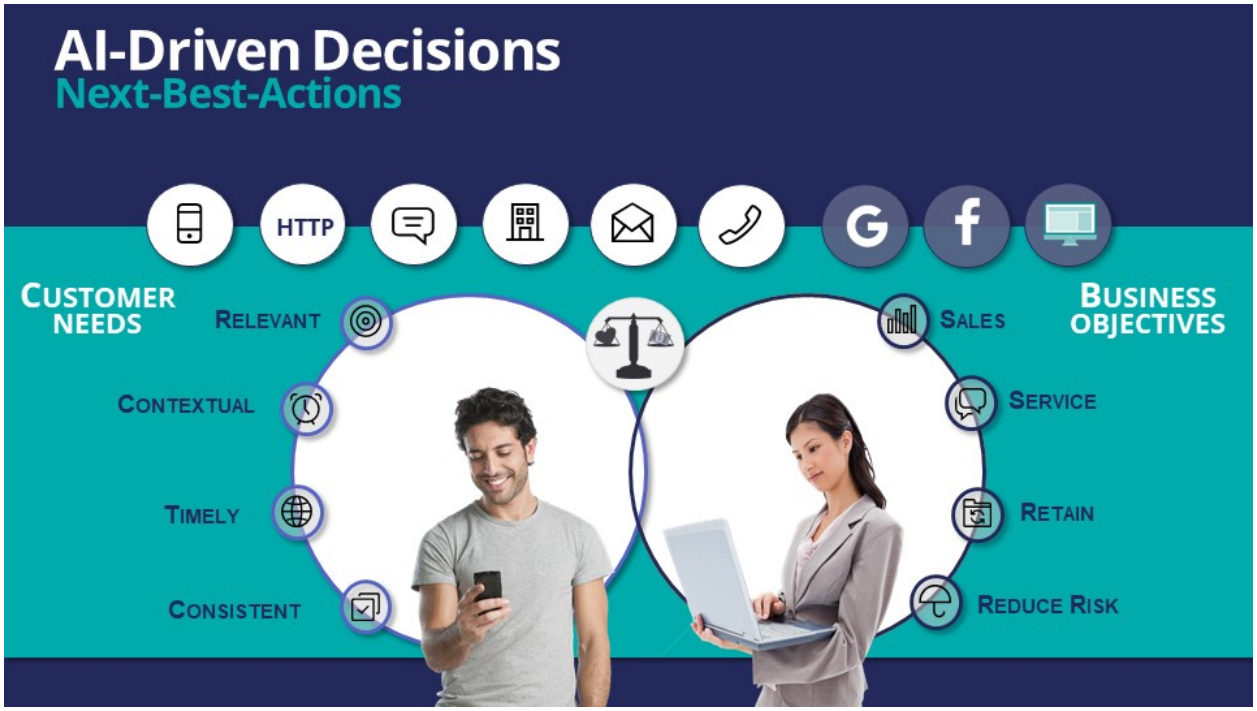

All of this information is used by the Next-Best-Action decision strategy to evaluate every potential action that could be taken with a particular customer in a given situation. The decision strategy then recommends the best way to interact with the customer to achieve the optimal result.

Using the Next-Best-Action approach, the Customer Decision Hub is able to identify the best moments for making a sale, providing a service, making a retention offer, or doing nothing at all (e.g. if nothing is relevant enough to warrant the customer's attention). Next-Best-Action is even able to select which offers are most likely to be accepted by the customer in a sales or retention situation. Next-Best-Action decisions are distributed, in real-time, to each of your real-time owned channels, such as web, mobile, and contact center. Through Pega Marketing, Next-Best-Actions can also be distributed to real-time paid channels such as Google, YouTube, Facebook, LinkedIn and Instagram. Pega Marketing also integrates with non-real time outbound channels such as data management platforms (DMPs) and email.

Once the Next-Best-Actions are distributed and customer responses have been received by the Brain, the whole process begins again, and new Next-Best-

Actions are distributed within milliseconds. Every outbound channel, including a data management platform, is dynamically updated with the Next-Best-Action to ensure consistency and an optimized customer experience no matter which channel the customer interacts with.

In summary, the Pega Customer Decision Hub is the Always-On Brain that acts as a single, centralized decision authority.

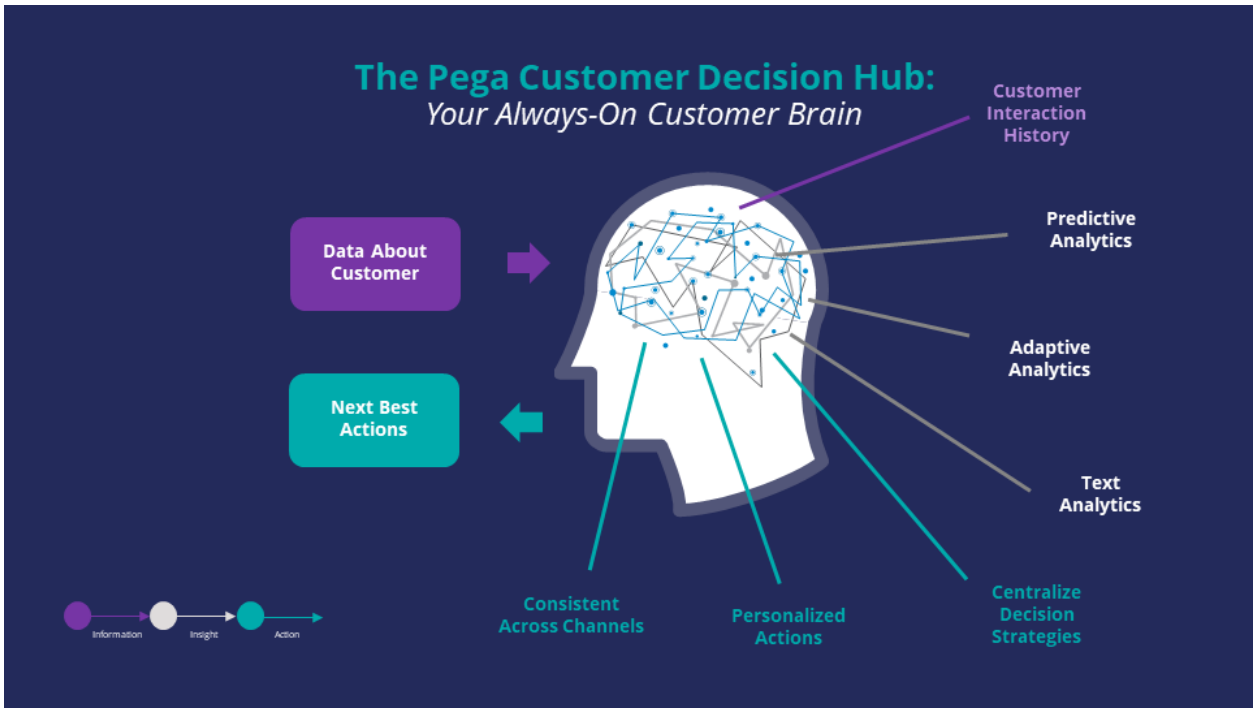

It uses data about the customer, including past interactions, as input.

It leverages advanced AI techniques to make predictions.

And it uses decision strategies (which combine traditional business rules with predictive, adaptive and text analytics), to deliver consistent and personalized Next-Best-Actions across all channels.

Next-Best-Action Designer

Introduction

Next-Best-Action Designer guides you through the creation of a Next-Best-Action strategy for your business. Its intuitive interface, proven best practices and sophisticated underlying decisioning technology enable you to automatically deliver personalized customer experiences across inbound, outbound and paid channels. Next-Best-Action Designer is organized according to the high-level sequence of steps needed to configure the Next-Best-Action strategy for your organization.

Transcript

Next-Best-Action Designer guides you through the creation of a Next-Best-Action strategy for your business. Its intuitive interface, proven best practices and sophisticated underlying decisioning technology enable you to automatically deliver personalized customer experiences across inbound, outbound and paid channels.

The Next-Best-Action Designer user interface allows you to easily define, manage and monitor Next-Best-Actions.

The tabs across the top of the user interface represent the steps that need to be completed to define Next-Best-Actions.

Use the **Taxonomy** component to define the business structure for your organization.

Use the **Constraints** component to implement channel limits and constraints.

Use the **Engagement policy** component to define the rules that control which actions are offered to which customers.

Use the **Arbitration** component to configure how actions are prioritized.

Use the **Channels** component to configure when and where Next-Best-Action is triggered.

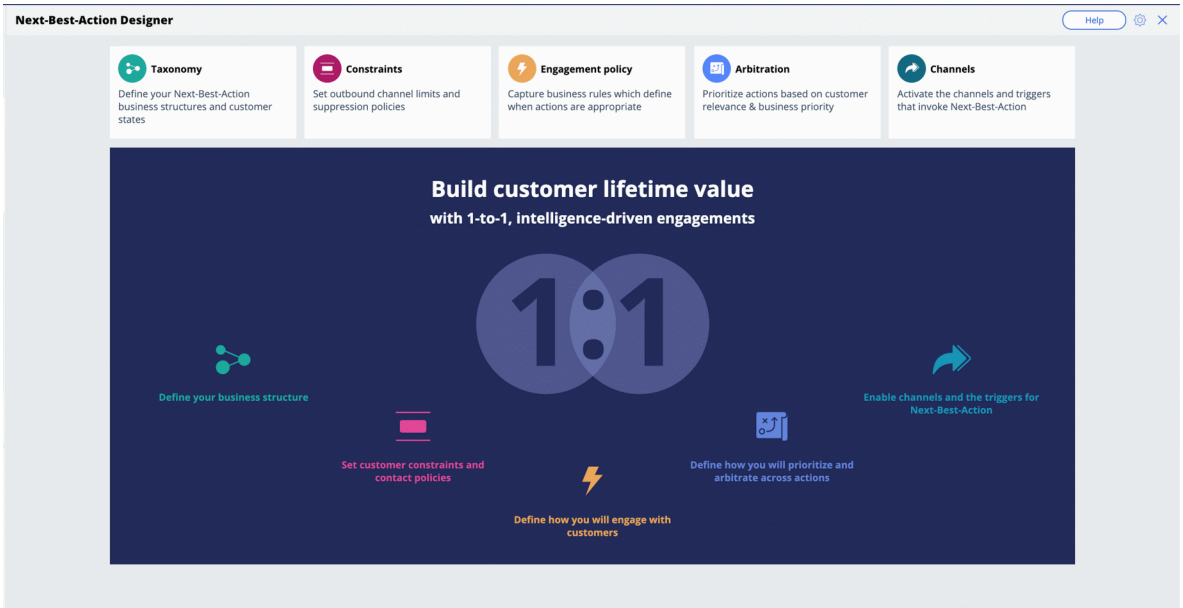

The system uses these definitions to create an underlying Next-Best-Action Strategy framework. This framework leverages best practices to generate Next-Best-Action decision strategies at the enterprise level. These decision strategies are a combination of the business rules and AI models that form the core of the Pega Centralized Decision Hub, which determines the personalized set of Next-Best-Actions for each customer.

Use the **Taxonomy** component to define the hierarchy of Business Issues and Groups to which an action belongs.

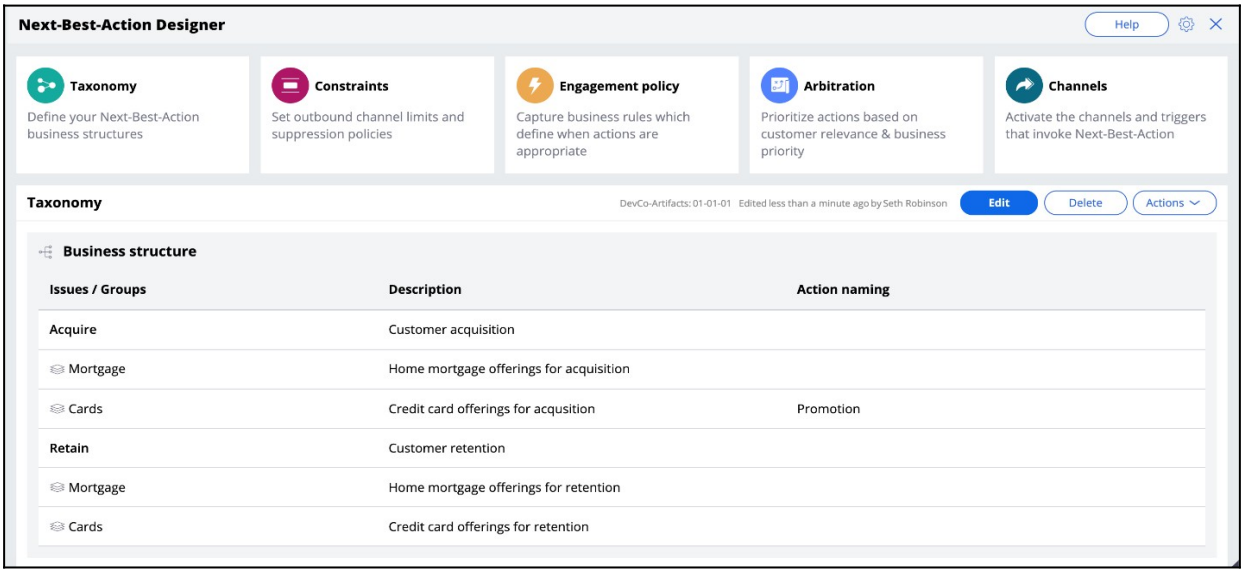

A Business Issue is the purpose behind the actions you offer to customers. For example, actions with the purpose of retaining existing customers should be grouped under the business Issue of Retention. Actions with the purpose of acquiring new customers belong to the business Issue of Acquisition.

Business Groups are used to organize customer actions into categories. For example, as part of the business Issue of Acquisition, you can create Groups for products like Credit Cards, Mortgages, or Personal Loans, with the intention of offering these to potential customers.

Use **Constraints** to specify outbound contact limits as well as to limit overexposure to a specific action or group of actions.

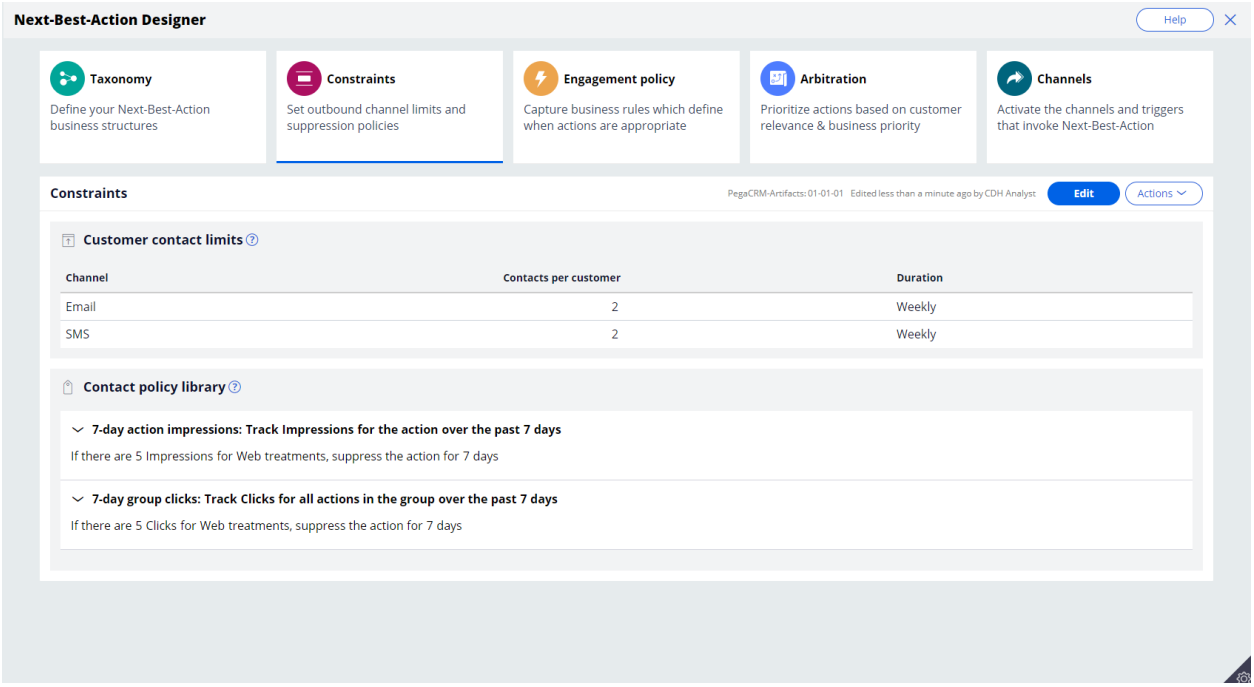

Customer contact limits allow you to limit the number of interactions that a customer can receive over a given period of time on a specific channel. For example, you can decide that you do not want your customers to receive more than two emails per week or two SMS messages per week.

On the Constraints tab of Next-Best-Action Designer, you can define more extensive suppression rules by creating Contact Policy rules in the library. Contact Policy rules are reusable across all Business Issues and Groups.

In the Contact Policy library, you define suppression rules that automatically put an action on hold after a specific number of outcomes are recorded for some or all channels. For example, an action can be suppressed for a customer for seven days after the customer has seen an ad for that action five times. Suppressing or pausing an action prevents over-exposure by limiting the number of times a customer is exposed to the same action.

Use **Engagement policies** to define when specific actions or groups of actions are appropriate for customers.

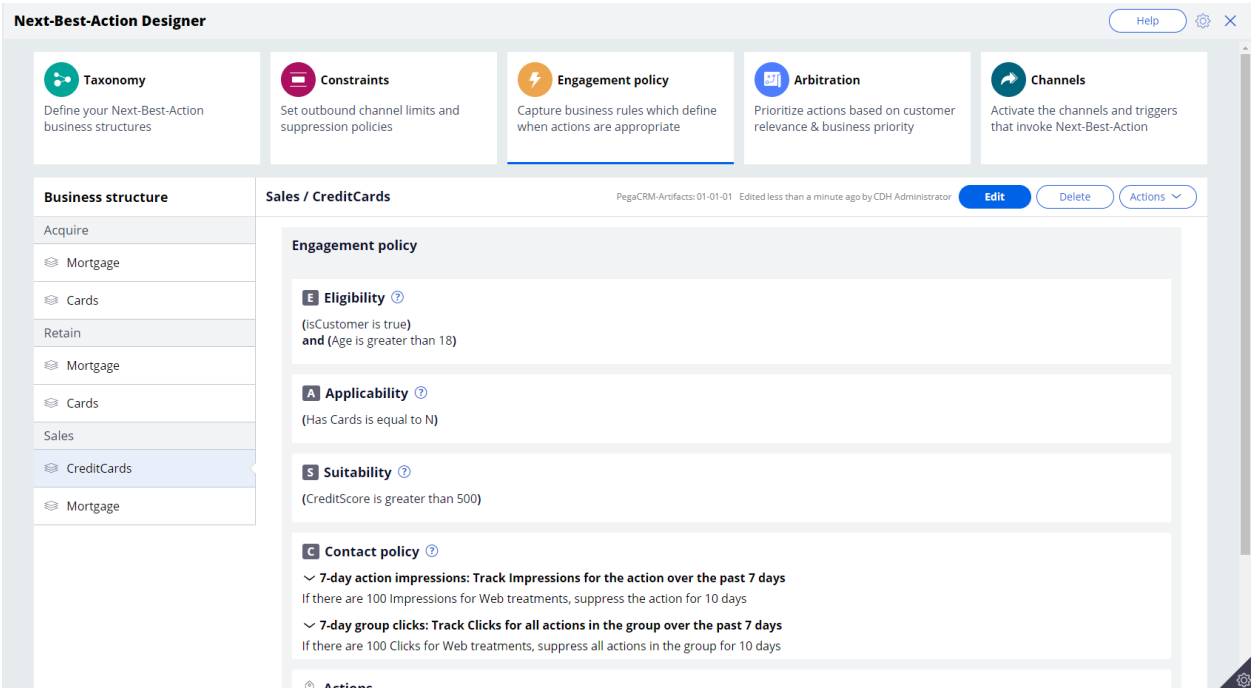

There are four types of engagement policies:

Eligibility determines whether or not a customer qualifies for an action or group of actions. For example, an action may only be available for customers over a specific age or living in a specific geographic location.

Applicability determines if an action or group of actions is relevant for a customer at a particular point in time. For example, a discount on a specific credit card may not be relevant for a customer who already owns a card.

Suitability determines if an action or group of actions is appropriate for a customer for ethical or empathetic reasons. For example, a new loan offer may not be appropriate for a customer whose credit score is low, even though it might be profitable for the bank.

Contact Policies determine when an action or group of actions should be suppressed and for how long. For example, you can suppress an action after a specific number of promotional messages has been sent to customers. To activate Contact Policy rules created in the library on the Constraints tab, add them to the Engagement Policy tab.

Arbitration determines how the Customer Decision Hub prioritizes the list of eligible and appropriate actions that come out of each group.

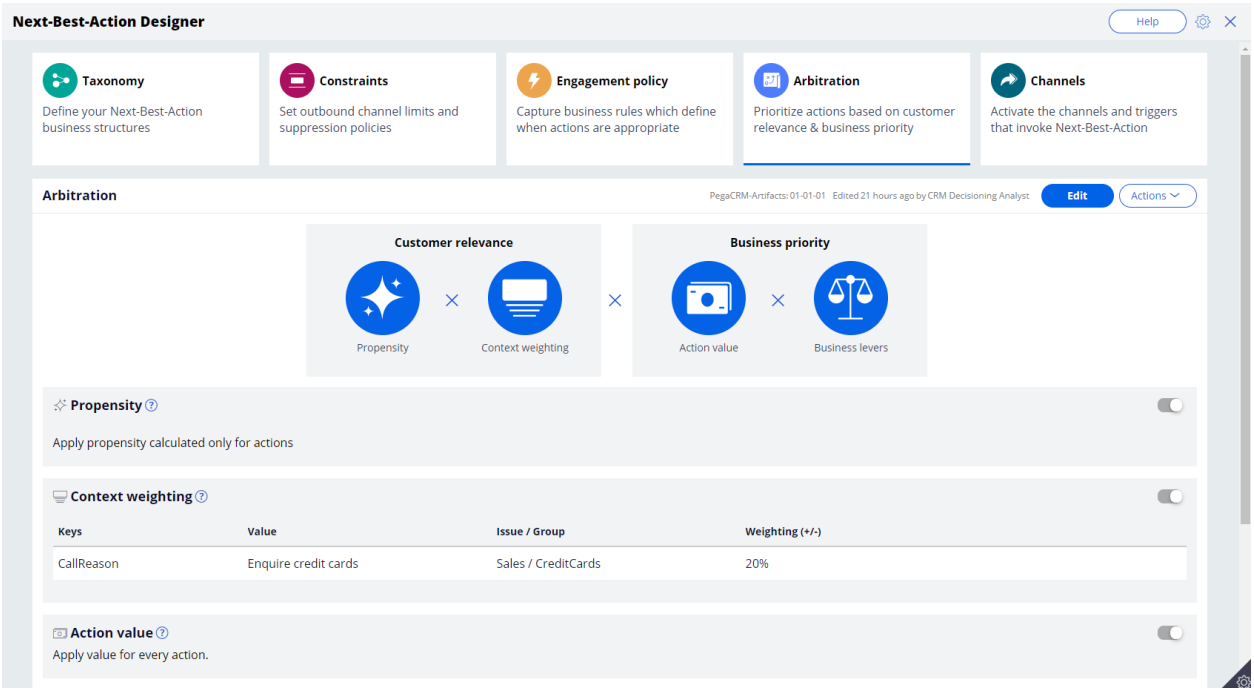

The factors weighed in arbitration are: Propensity, Context weighting, Action value, and Business levers, each represented by numerical values. A simple formula is used to arrive at a prioritization value, which is used to select the top actions.

Propensity is the likelihood of a customer responding positively to an action. This is calculated by Artificial Intelligence (AI). For example, a click on an offer banner or an accept of an offer in the contact center are considered positive behaviors.

Real-time contextual data is an important part of making highly relevant recommendations**.** **Context weighting** allows you to assign weighting to a specific context value for all actions within an Issue or Group. For example, if a customer contacts the bank to change their address, the weight of the Service context will increase, and the highest priority action will be to ensure that the relevant service is delivered to the customer.

Action value enables you to assign a financial value to an action and prioritize high-value actions over low-value ones. For example, promoting an unlimited data plan might be more

profitable for the company than a limited data plan. Action values are typically normalized across Issues and Groups.

Business levers enable you to accommodate ad hoc business priorities by specifying a weight for an action or Group of actions and/or its associated Business Issue.

Next-Best-Action Designer enables Next-Best-Actions to be delivered via inbound, outbound and paid channels.

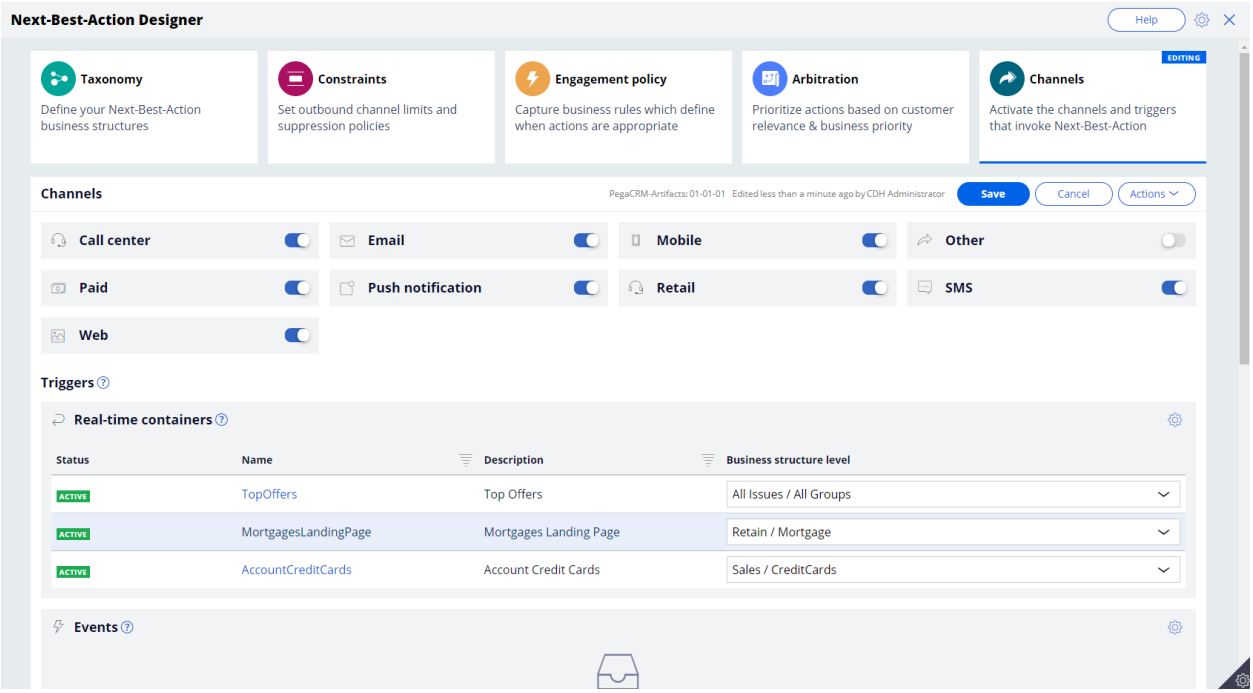

These channels can be toggled on or off. If a channel is toggled off, the Next-Best-Actions will not be delivered to that channel.

An external real-time channel is any channel that presents actions selected by the Customer Decision Hub to a customer. These channels can include a website, or a call-center or mobile application. A real-time container is a placeholder for content in an external real-time channel.

A trigger is a mechanism whereby an external channel invokes the execution of a Next-Best-Action decisioning process for specific Issues and Groups. The result will be delivered back to the invoking channel. For example, when a real-time container called “Mortgages Landing Page” is configured, the website invokes this real-time container before loading the mortgage page.

As you have seen in this video, Next-Best-Action Designer is organized according to the high-level sequence of steps needed to configure the Next-Best-Action strategy for your organization. These steps involve:

- Defining the business structure for your organization
- Implementing the channel limits and constraints
- Defining the rules that control which actions are offered to which customers
- Configuring how actions are prioritized
- Configuring when and where Next-Best-Action is triggered

Next best action in an omnichannel environment

Description

Pega's omni-channel AI delivers the right action at every customer touch by crunching millions of data points in real-time.

Learn how you can use Next-Best-Action to ensure a customer service representative takes relevant actions at every step during a customer interaction. Learn how easy it is to configure the brain by using the Next-Best-Action Designer to select the right offer for the right customer on a digital channel.

Learning Objectives

- Explain the benefits of Next-Best-Action in a contact-center
- Explain how adaptive models are used to predict customer behavior
- Describe the cross-sell on the web use case
- Explain the purpose of key configurations in the Next-Best-Action Designer

Next best action in a contact center

Introduction

The Customer Decision Hub is an “always on” centralized decisioning “brain” that calculates a 1-to-1 business case for every Next-Best-Action recommendation. To create the business case, the Customer Decision Hub combines customer profile with previous interaction results, the current call context, business rules, and then applies predictive analytics.

Transcript

Next-Best-Action is used to ensure a customer service representative takes relevant actions at every step during a customer interaction.

In this scenario, U+ is a retail bank that uses Pega Customer Service in its contact center.

A call comes in to the U+ service center from Sara Connor, a U+ customer. The call is immediately routed to a service representative.

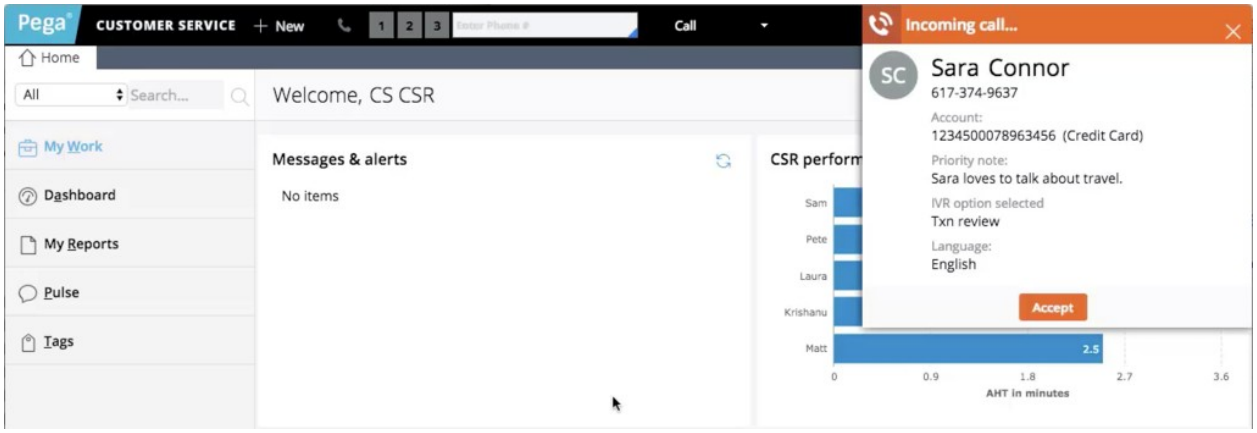

The call details from the IVR system indicate that Sara would like to discuss recent credit card transactions.

Once the service representative accepts the call, all relevant details about Sara are visible on screen.

Next-Best-Action then guides the representative to take the next step with Sara.

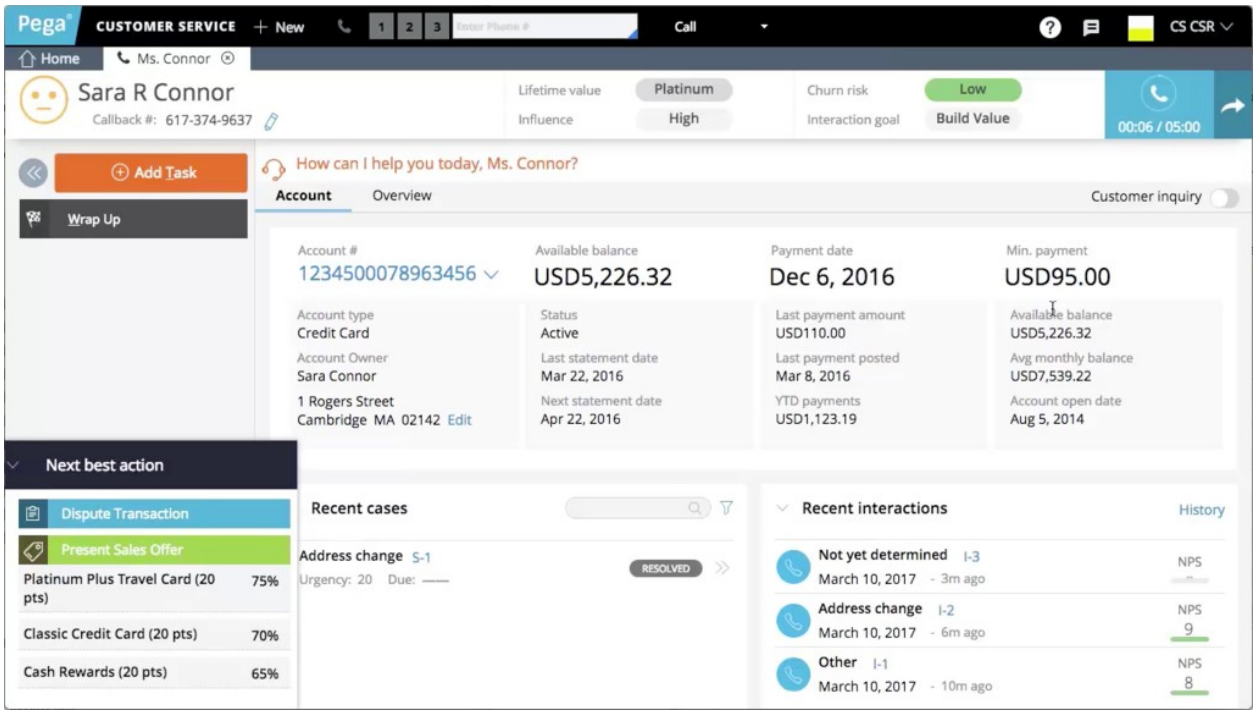

In the lower left corner of the screen you can see the Next-Best-Action that has been recommended to the agent by the Customer Decision Hub.

The Customer Decision Hub is an “always on” centralized decisioning “brain” that calculates a 1-to-1 business case for every Next-Best-Action recommendation. To create the business case, the Customer Decision Hub combines customer profile with previous interaction results, the current call context, business rules, and then applies predictive analytics.

The Customer Decision Hub re-evaluates the Next-Best-Action and delivers a new recommendation when any new information becomes available. For example, when the customer responds to the recommended action.

In this case, the recommended action is to start a service task to handle Sarah’s transaction dispute. So the service representative carries out the task. But the representative is always in control and can select other service actions as appropriate during the conversation.

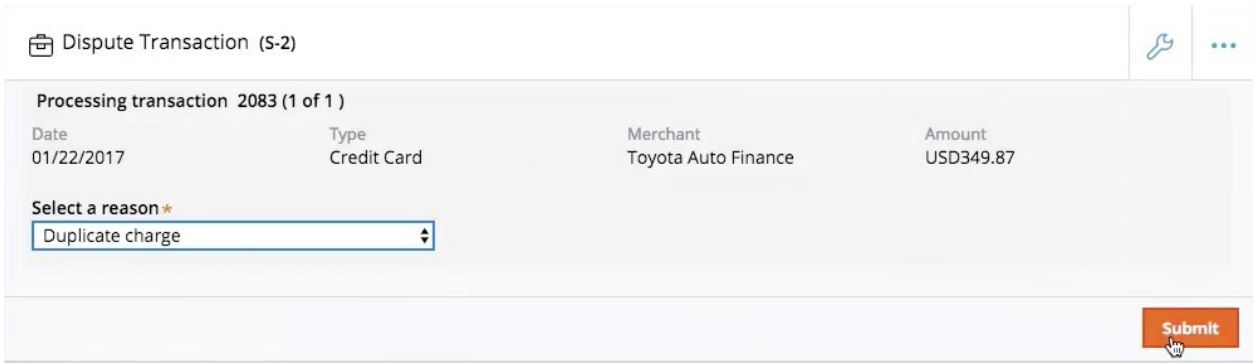

Once the service representative completes the task, the Next-Best-Action is refreshed to show the next recommended action, which is to present a sales offer.

The Customer Decision Hub has analyzed Sara's credit card usage patterns, which indicate that she is a frequent traveler.

As a frequent traveler she often pays a large amount in foreign currency transaction fees. So the highest recommendation is for a Platinum Plus Travel Card, which has no foreign currency transaction fees.

The percentages to the right are scores that are used to rank all relevant offers. The scores are calculated by balancing what the bank would like to promote with what Sara is likely to be interested in.

The representative can view more details about the recommended offer so he can further discuss its benefits with Sara.

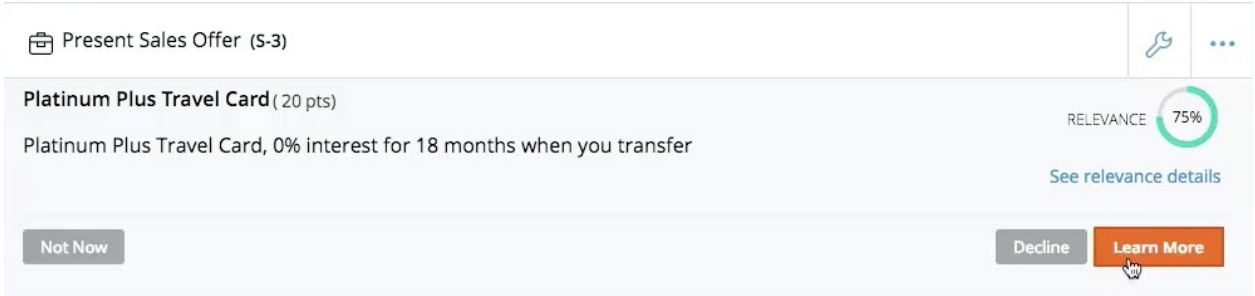

In this case, Sara is interested in knowing more.

After learning about the benefits of the offer, Sara is indeed convinced that it's a good offer for her, and she decides to accept it.

When the representative clicks Accept, the customer response is recorded in the Interaction History.

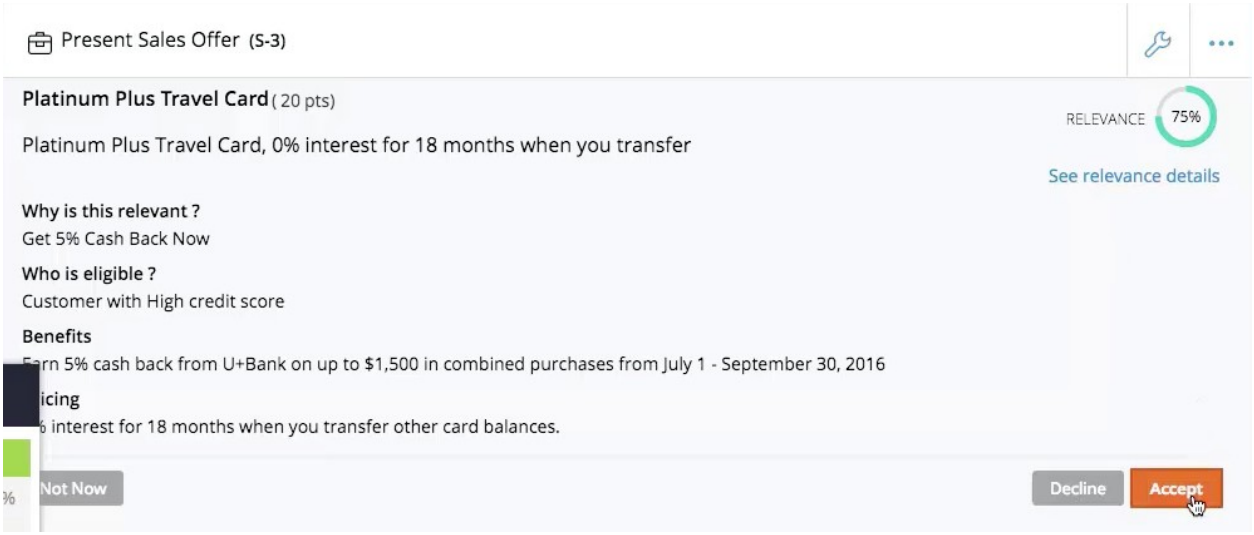

This new information is used by the Customer Decision Hub to make the next recommendation.

Notice that the Next-Best-Action is now refreshed.

The Platinum Plus Travel Card and all other credit card offers have been removed from the list.

To summarize, the always on, centralized Customer Decision Hub combines customer and contextual information with business rules and analytics to recommend the Next-Best-Action in real-time.

The Next-Best-Action is re-evaluated when new information becomes available.

Next best action on digital

Introduction

Leverage a website as a marketing channel to improve one-to-one customer engagement, drive sales, and deliver Next-Best-Actions in real-time. Use the Pega Customer Decision Hub™ to recommend more relevant banner ads to customers when they visit their personal portal.

Transcript

This video describes a typical cross-selling use case on the web channel.

U+ is a retail bank. The bank would like to leverage its website as a marketing channel to improve 1-to-1 customer engagement, drive sales, and deliver Next-Best-Actions in real-time.

The bank has decided to use the Pega Customer Decision Hub™ to recommend more relevant banner ads to its customers when they visit their personal portal.

Banner ads are shown on various pages throughout the website.

For example, on the home page, U+ can display a Hero banner at the top of the page, which is typically a larger image with bigger typeface.

Below that, there is space to display several Tile banners, which are typically smaller.

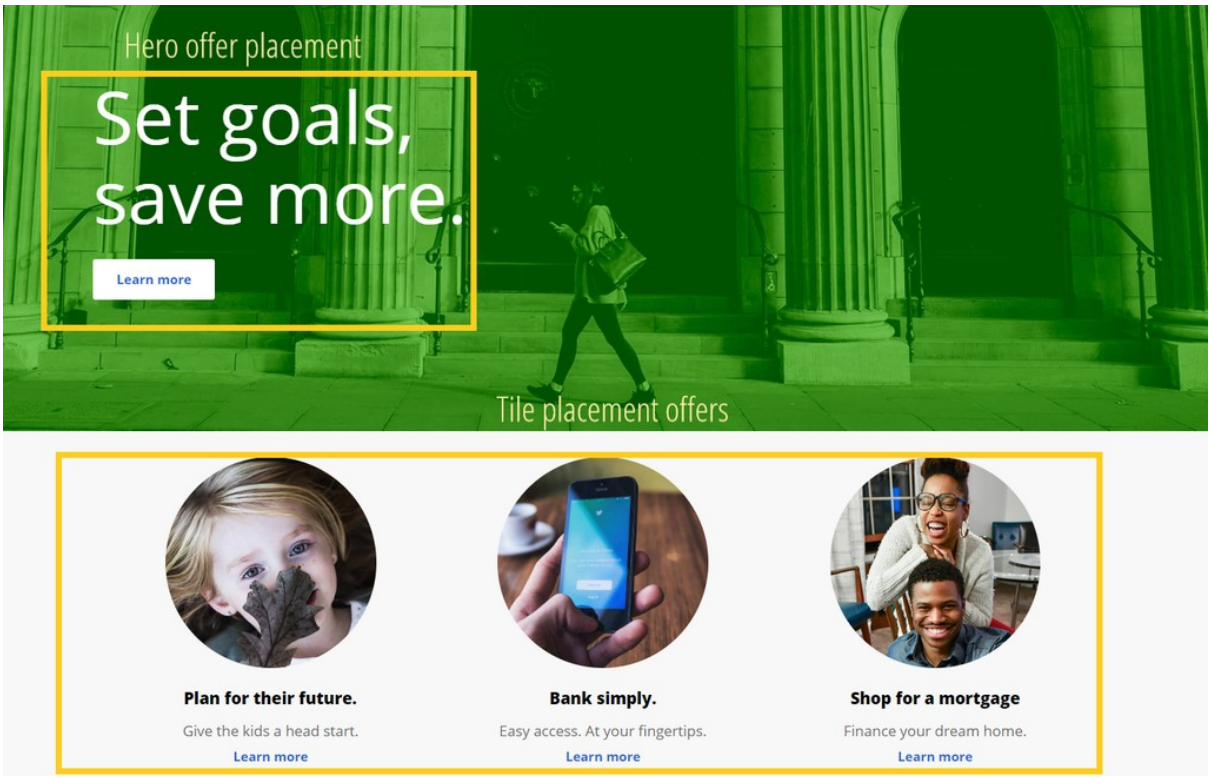

When customers log in to their personal portal, they also see a Tile banner on the **Account overview** page.

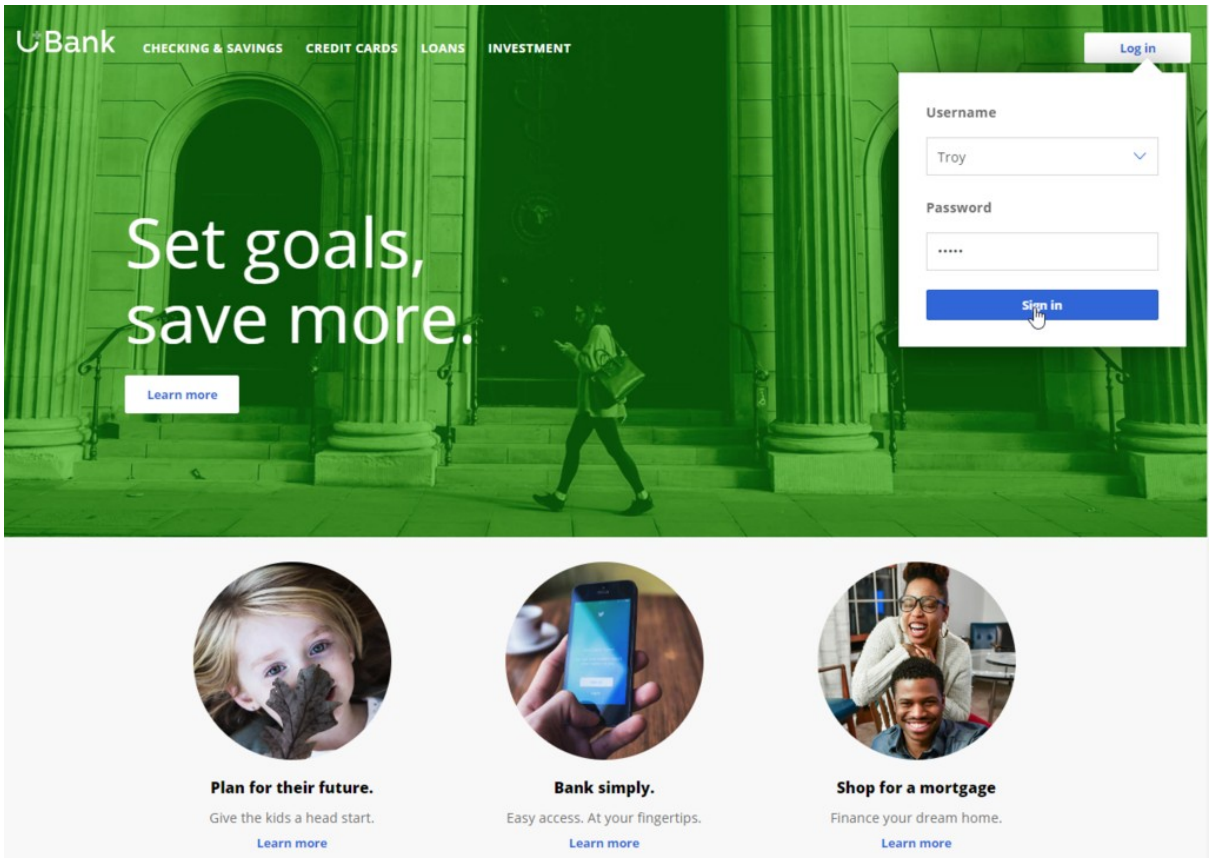
Bank
CHECKING & SAVINGS
CREDIT CARDS
LOANS
INVESTMENT
Log in
Username
Troy
Password
Sign in
Set goals, save more.
Learn more
Plan for their future.
Give the kids a head start.
Learn more
Bank simply.
Easy access. At your fingertips.
Learn more
Shop for a mortgage
Finance your dream home.
Learn more

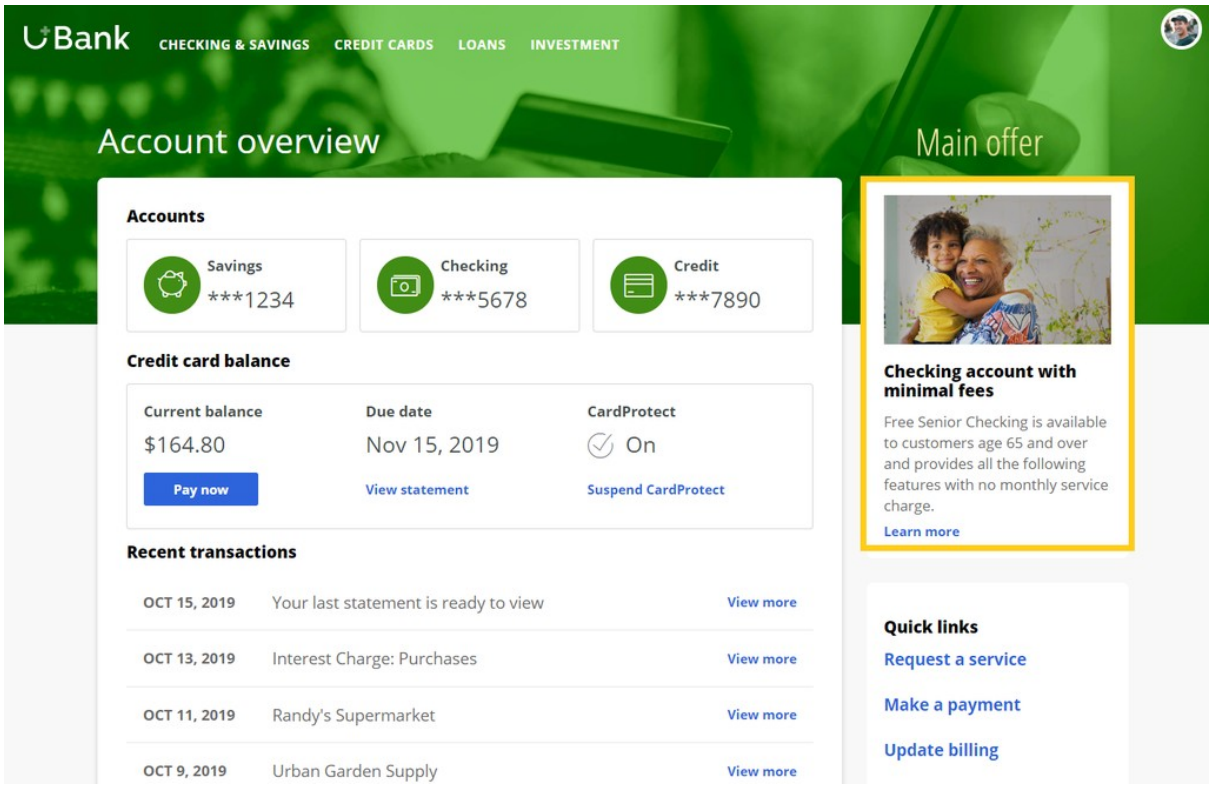
Bank
CHECKING & SAVINGS
CREDIT CARDS
LOANS
INVESTMENT
Account overview
Accounts
Savings
***1234
Checking
***5678
Credit
***7890
Credit card balance
Current balance
$164.80
Pay now
Due date
Nov 15, 2019
View statement
CardProtect
On
Suspend CardProtect
Recent transactions
OCT 15, 2019
Your last statement is ready to view
View more
OCT 13, 2019
Interest Charge: Purchases
View more
OCT 11, 2019
Randy's Supermarket
View more
OCT 9, 2019
Urban Garden Supply
View more
Main offer
Checking account with minimal fees
Free Senior Checking is available to customers age 65 and over and provides all the following features with no monthly service charge.
Learn more
Quick links
Request a service
Make a payment
Update billing

The main intent of U+ at this stage is to increase their web engagement. This can be measured by click-through rate. A Click-through is recorded when the customer clicks on the **Learn more** link.

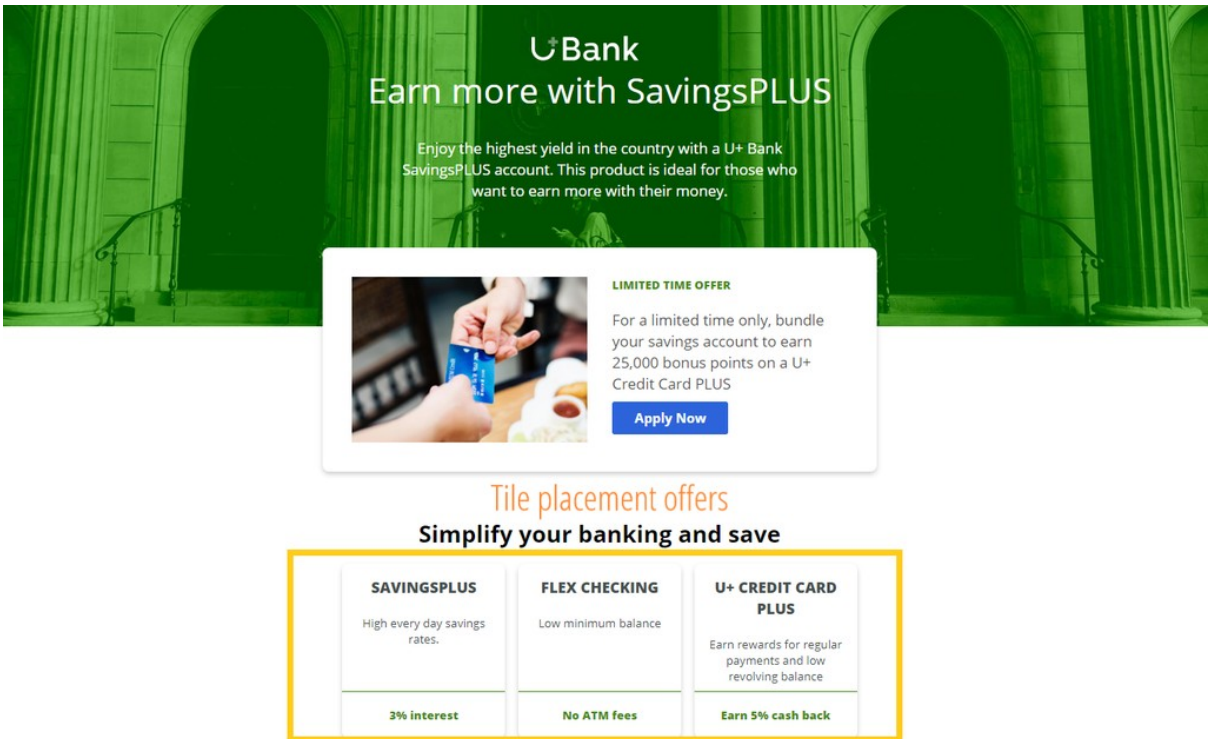

The bank would like to use these banners on the **Account overview** page, to display offers that are more relevant and likely to receive a positive response.

The offers will be selected by a combination of artificial intelligence (AI) and other business rules. The AI and business rules are defined in the Pega Customer Decision Hub.

The Pega Customer Decision Hub is the always-on customer brain that acts as a single, centralized decision authority. The always-on customer brain selects the right offer to be displayed to each customer who visits the bank's website.

Next-Best-Action Designer lets you configure how you want the always-on brain to select the best offer for a customer. The best offer is the result of a series of decisions that are executed in a hierarchical fashion by the brain.

The bank plans to implement the requirement in multiple phases.

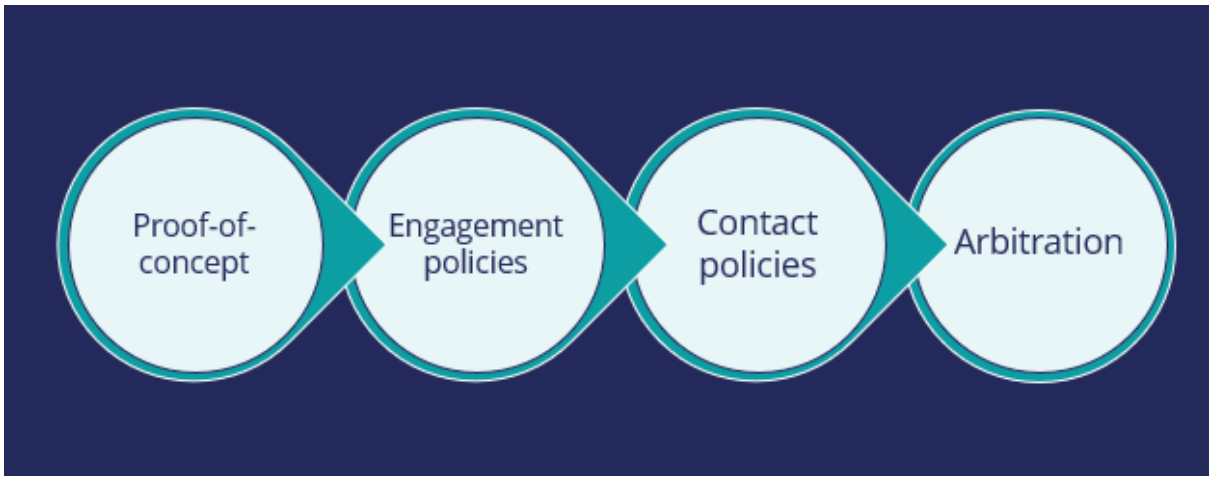

The first phase is a proof-of-concept phase. In this phase, the goal is to display a credit card offer on the U+ website. This requires getting the basic environment up and running, setting up the business structure, defining an Action and a Treatment, and enabling channels and triggers for Next-Best-Action.

As a result of this phase, a credit card offer will be displayed on the **Account overview** page to all customers who visit the U+ web site. For example, if customer Troy logs in to his account, the **Cash back** offer is displayed. If another user logs in, they will see the same offer. However, in practice, more offers should be displayed. Also, not all offers may be available to a customer for various reasons.

The next phase is to add customer engagement policies. Engagement policies are the set of conditions such as eligibility, applicability, and suitability that qualify an offer, or a group of offers for a customer. As a result of engagement rules, customers will see only those offers that the organization believes they should be exposed to. For example, Troy logs in, he sees the Rewards Card, but for Barbara this is not applicable, so it will never appear; instead, she sees the Rewards Plus Card.

Too many contact attempts over a short period of time can have a negative impact on a customer's attitude toward further offers by your company. Therefore, in the next phase, U+ implements some contact policies using suppression rules, which allow an offer to be put on hold after a specific number of outcomes. For example, if Troy ignores an ad a few times, then the ad will no longer be shown to him over a period of time. Instead, his **Account overview** page will show a different ad.

Basically, from a set of all available offers, the choice is narrowed down by engagement policies. Then the selection is further narrowed down by suppression rules.

After the engagement policies and suppression rules have “whittled down” the total possible offers to a few, Arbitration is used to choose the top offer based on what is relevant for the customer right NOW.

Arbitration is the last phase of cross-sell in the U+ web use case.

Arbitration aims at balancing customer relevance with business priorities. Specifically, Propensity, Context Weighting, Action Value, and Business Levers are given numerical values. A simple formula is then used to arrive at a prioritization value, which is used to select the top offer. For example, Troy qualifies for three credit card offers. When he logs in, he sees the top offer for him, the Standard Card. This offer is the Top 1 because the priority value is the highest among all other offers.

Next-Best-Action on digital

Description

Learn about a typical cross-selling use case on the web channel and its four phases of implementation to improve 1-to-1 customer engagement, drive sales, and deliver Next-Best-Actions in real time.

Learning objectives

- Describe a typical cross-selling use case on the web channel
- List the four phases of use case implementation

Business use case: Cross-sell on the web

Introduction

Pega Customer Decision Hub's Next-Best-Action Designer lets you configure how you want the always-on brain to select the best offer for a customer. The best offer is the result of a series of decisions that are executed in a hierarchical fashion by the brain. Cross-selling on the web channel will help you improve 1-to-1 customer engagement, drive sales, and deliver Next-Best-Actions in real-time.

Transcript

This video describes a typical cross-selling use case on the web channel.

U+ is a retail bank. The bank would like to leverage its website as a marketing channel to improve 1-to-1 customer engagement, drive sales, and deliver Next-Best-Actions in real-time.

The bank has decided to use the Pega Customer Decision Hub™ to recommend more relevant banner ads to its customers when they visit their personal portal.

Banner ads are shown on various pages throughout the website.

For example, on the home page**,** U+ can display a Hero banner at the top of the page, which is typically a larger image with bigger typeface.

Below that, there is space to display several Tile banners, which are typically smaller.

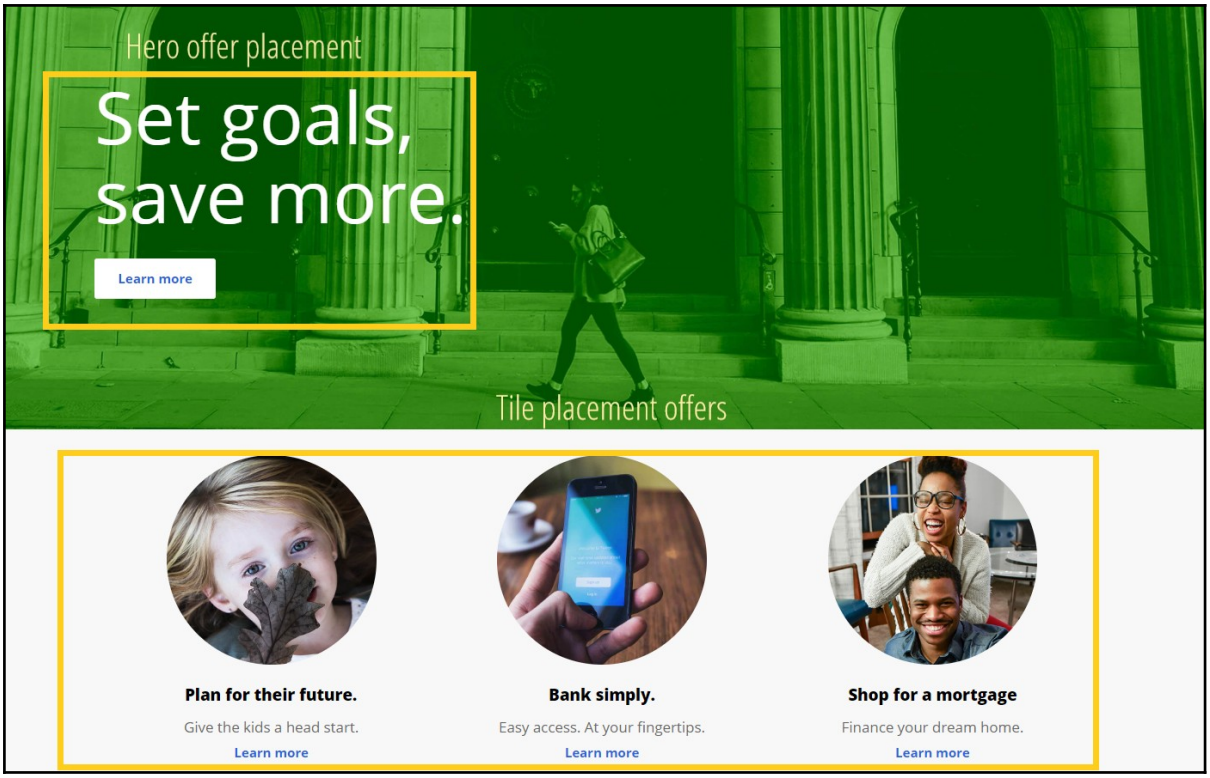

When customers log in to their personal portal, they also see a Tile banner on the ‘Account overview’ page**.**

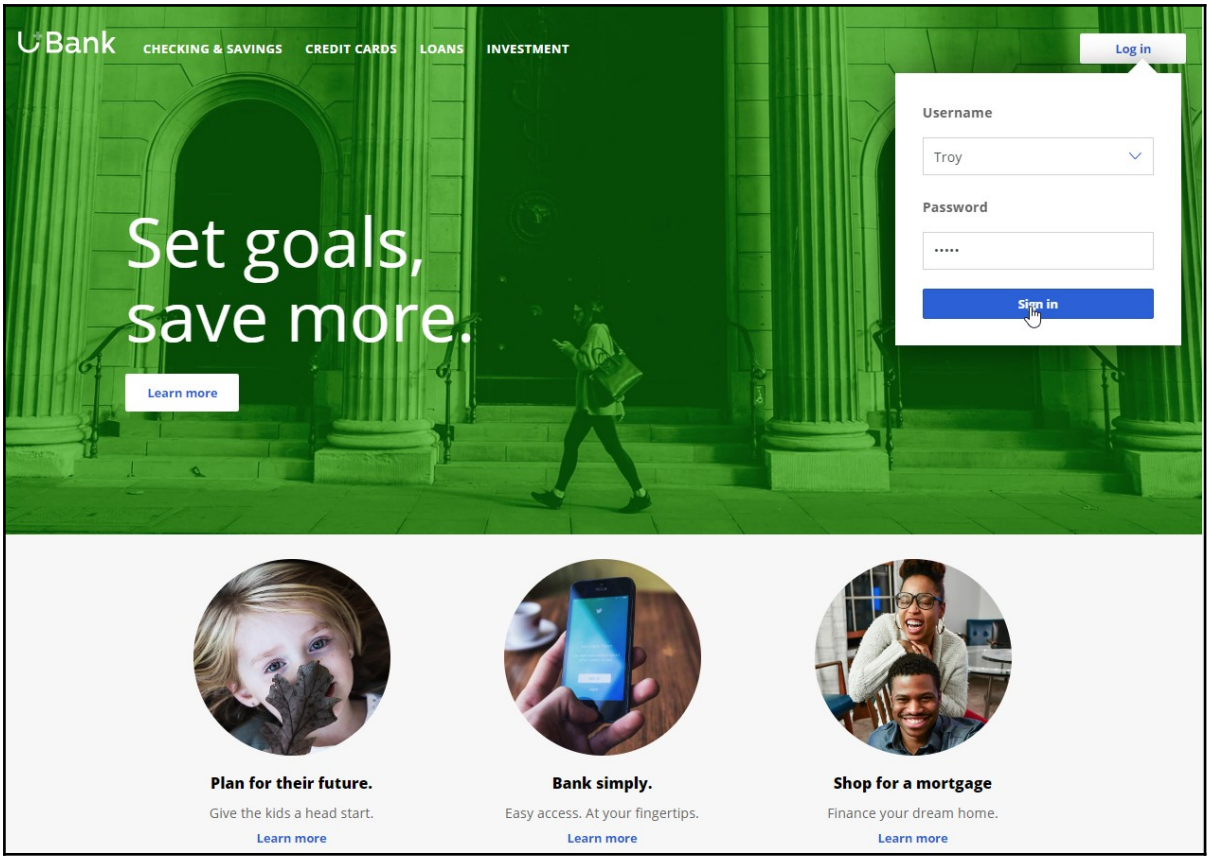
Bank
CHECKING & SAVINGS
CREDIT CARDS
LOANS
INVESTMENT
Log in
Username
Troy
Password
Sign in
Set goals, save more.
Learn more
Plan for their future.
Give the kids a head start.
Learn more
Bank simply.
Easy access. At your fingertips.
Learn more
Shop for a mortgage
Finance your dream home.
Learn more

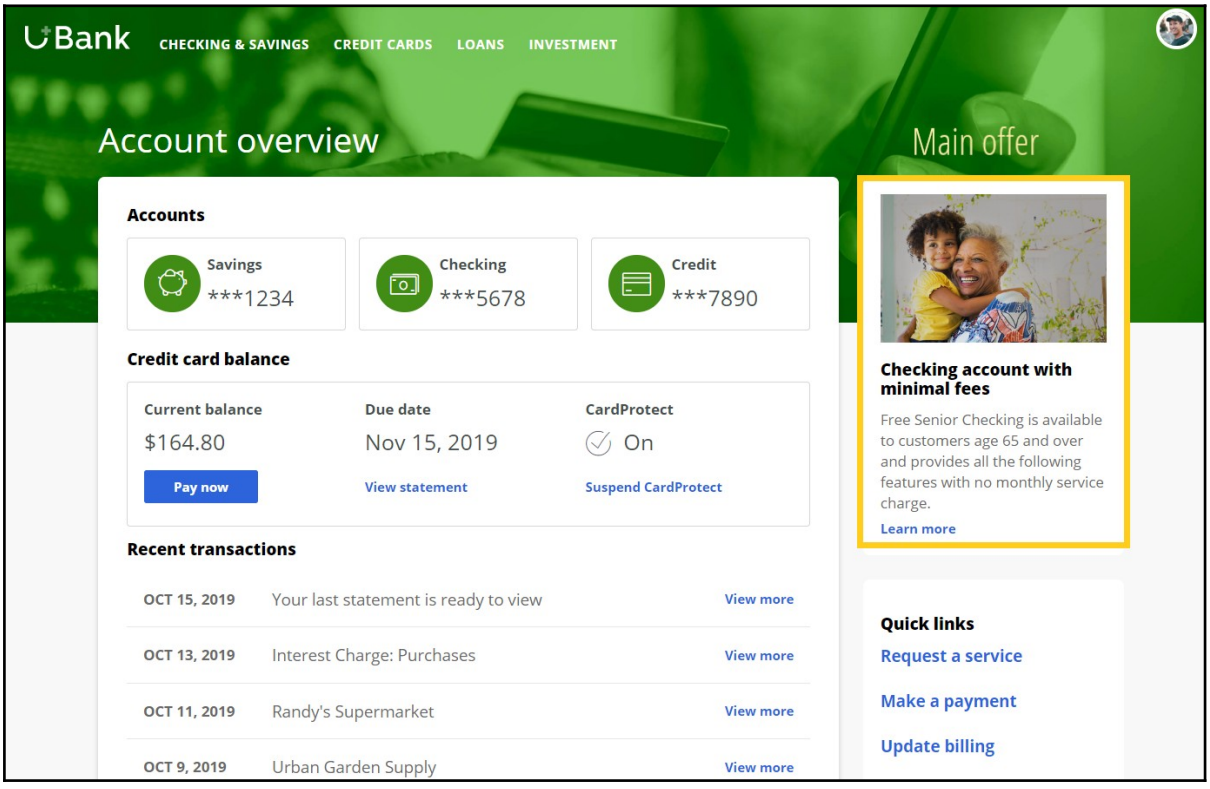
Bank
CHECKING & SAVINGS
CREDIT CARDS
LOANS
INVESTMENT
Account overview
Accounts
Savings
***1234
Checking
***5678
Credit
***7890
Credit card balance
Current balance
$164.80
Pay now
Due date
Nov 15, 2019
View statement
CardProtect
On
Suspend CardProtect
Recent transactions
OCT 15, 2019
Your last statement is ready to view
View more
OCT 13, 2019
Interest Charge: Purchases
View more
OCT 11, 2019
Randy's Supermarket
View more
OCT 9, 2019
Urban Garden Supply
View more
Main offer
Checking account with minimal fees
Free Senior Checking is available to customers age 65 and over and provides all the following features with no monthly service charge.
Learn more
Quick links
Request a service
Make a payment
Update billing

The main intent of U+ at this stage is to increase their web engagement. This can be measured by click-through rate. A Click-through is recorded when the customer clicks on the ‘Learn more’ link.

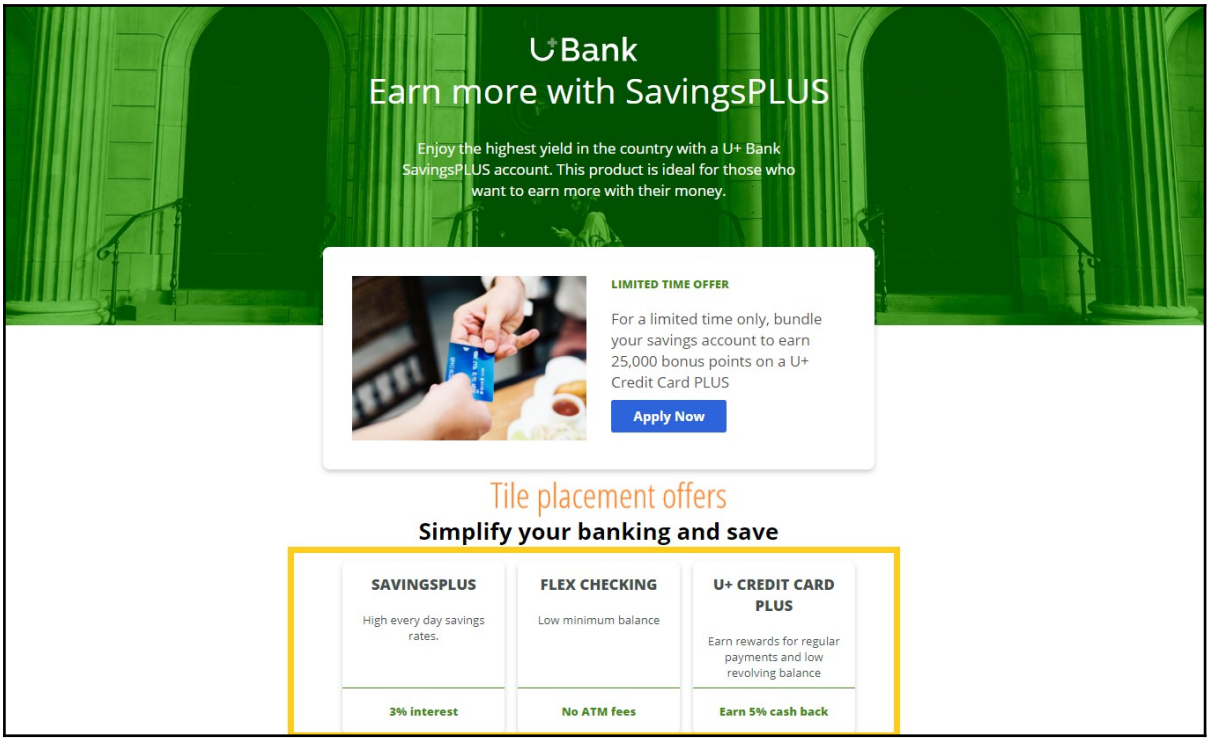

The bank would like to use these banners on the ‘Account overview’ page, to display offers that are more relevant and likely to receive a positive response.

The offers will be selected by a combination of artificial intelligence (AI) and other business rules. The AI and business rules are defined in the Pega Customer Decision Hub.

The Pega Customer Decision Hub is the always-on customer brain that acts as a single, centralized decision authority. The always-on customer brain selects the right offer to be displayed to each customer who visits the bank’s website.

Next-Best-Action Designer lets you configure how you want the always-on brain to select the best offer for a customer. The best offer is the result of a series of decisions that are executed in a hierarchical fashion by the brain.

The bank plans to implement the requirement in multiple phases.

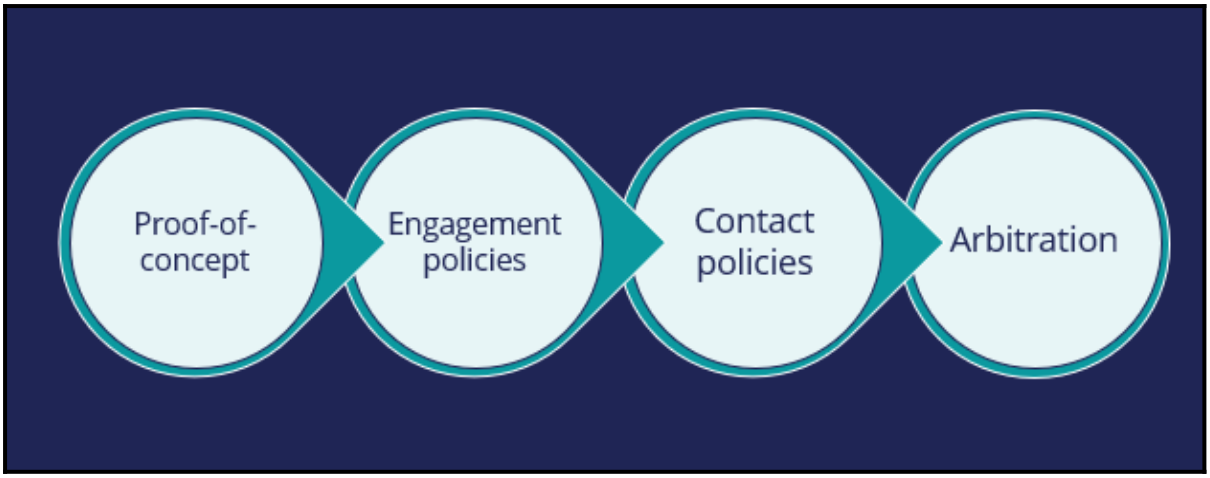

The first phase is a proof-of-concept phase. In this phase, the goal is to display a credit card offer on the U+ website. This requires getting the basic environment up and running, setting up the business structure, defining an Action and a Treatment, and enabling channels and triggers for Next-Best-Action.

As a result of this phase, a credit card offer will be displayed on the 'Account overview' page to all customers who visit the U+ web site. For example, if customer Troy logs in to his account, the 'Cash back' offer is displayed. If another user logs in, they will see the same offer. However, in practice, more offers should be displayed. Also, not all offers may be available to a customer for various reasons.

The next phase is to add customer engagement policies. Engagement policies are the set of conditions such as eligibility, applicability, and suitability that qualify an offer, or a group of offers for a customer. As a result of engagement rules, customers will see only those offers that the organization believes they should be exposed to. For example, Troy logs in, he sees the Rewards Card, but for Barbara this is not applicable, so it will never appear; instead, she sees the Rewards Plus Card.

Too many contact attempts over a short period of time can have a negative impact on a customer's attitude toward further offers by your company. Therefore, in the next phase, U+ implements some contact policies using suppression rules, which allow an offer to be put on hold after a specific number of outcomes. For example, if Troy ignores an ad a few times, then the ad will no longer be shown to him over a period of time. Instead, his 'Account overview' page will show a different ad.

Basically, from a set of all available offers, the choice is narrowed down by engagement policies. Then the selection is further narrowed down by suppression rules.

After the engagement policies and suppression rules have “whittled down” the total possible offers to a few, Arbitration is used to choose the top offer based on what is relevant for the customer right NOW.

Arbitration is the last phase of cross-sell in the U+ web use case.

Arbitration aims at balancing customer relevance with business priorities. Specifically, Propensity, Context Weighting, Action Value, and Business Levers are given numerical values. A simple formula is then used to arrive at a prioritization value, which is used to select the top offer. For example, Troy qualifies for three credit card offers. When he logs in, he sees the top offer for him, the Standard Card. This offer is the Top 1 because the priority value is the highest among all other offers.

This video has concluded. It showed you a typical cross-selling use case on the web channel and its four phases of implementation.

Defining and managing customer actions

Description

In Pega Customer Decision Hub™, learn how to use Next-Best-Action Designer to build the business issue/group hierarchies for your organization. Create an action under a specific business issue/group hierarchy and learn to rename an action to better reflect its business context.

Learning objectives

- Explain an actions hierarchy
- Configure a business structure with issues and groups
- Create actions
- Rename actions

Action hierarchy

Customer action introduction

In Pega Customer Decision Hub™, next-best-action customer recommendations can take many forms, such as a banner advertisement, a retention offer, or a service message. Learn how to recognize different types of actions in your organization and organize similar actions into action groups.

Actions in your organization

Click the **Play** icon to learn about customer actions and how to organize them.

Full-screen mode is available in the lower right of the interaction.

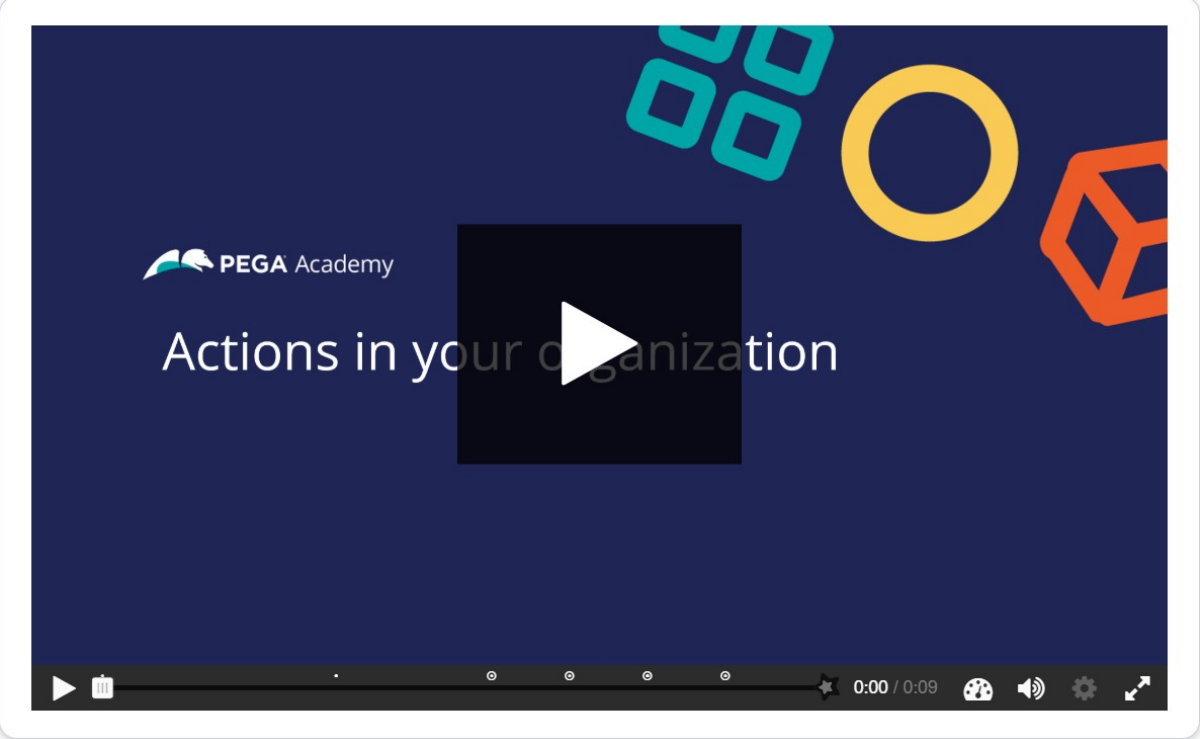

In Pega Customer Decision Hub, customer actions are offerings that list the details of your products. Next best action uses these properties to determine the priority of each offering for each customer and provide you with the next best action to take.

To learn more about the different type of customer actions, consider the following scenarios in various channels: contact center, mobile, and web.

Click **Continue** to see how actions are used in various scenarios.

Scenario 1: Contact center: up-sell

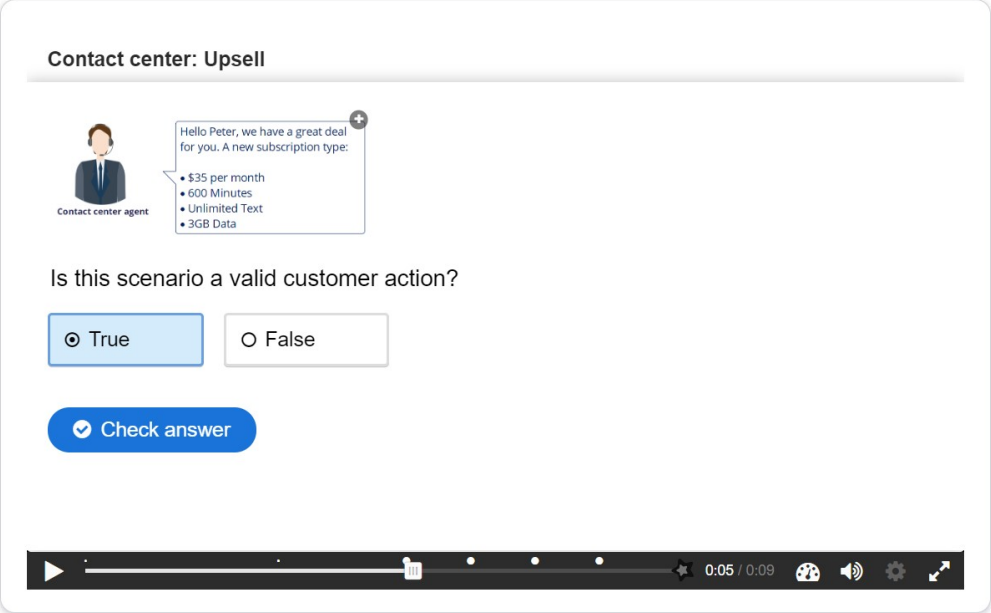

Is this scenario a valid customer action?

- True
- False

Correct feedback: **This scenario is an upsell offer made in a contact center and is a valid customer action.**

Incorrect feedback: **An upsell offer made in a contact center is a valid customer action.**

Scenario 2: Contact center: customer service

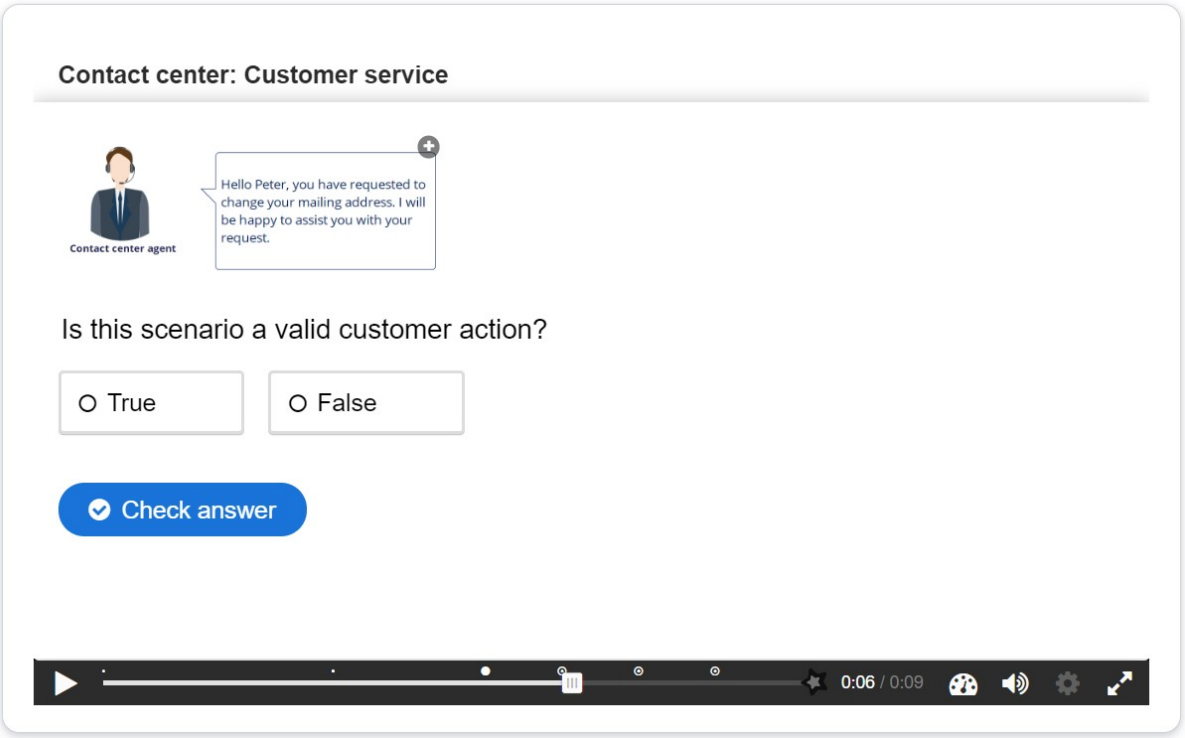

Is this scenario a valid customer action?

- True
- False

Correct feedback: **This scenario is a service message recommended by the Pega Customer Decision Hub and is a valid customer action.**

Incorrect feedback: **A service message recommended by the Pega Customer Decision Hub is a valid customer action.**

Scenario 3: SMS – retention

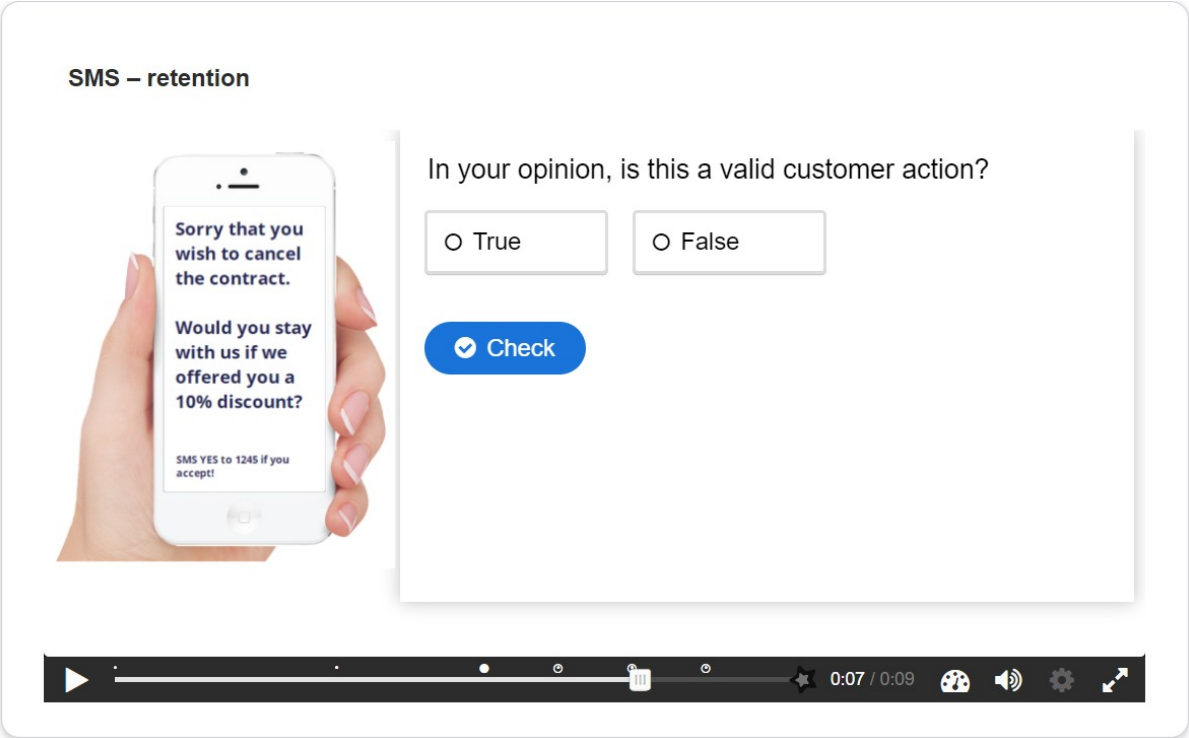

In your opinion, is this a valid customer action?

- True
- False

Correct feedback: **This scenario is a retention offer made on the mobile channel to retain a customer and is a valid customer action.**

Incorrect feedback: **A retention offer made on the mobile channel to retain a customer is a valid customer action.**

Scenario 4: Inbound-web: cross-sell

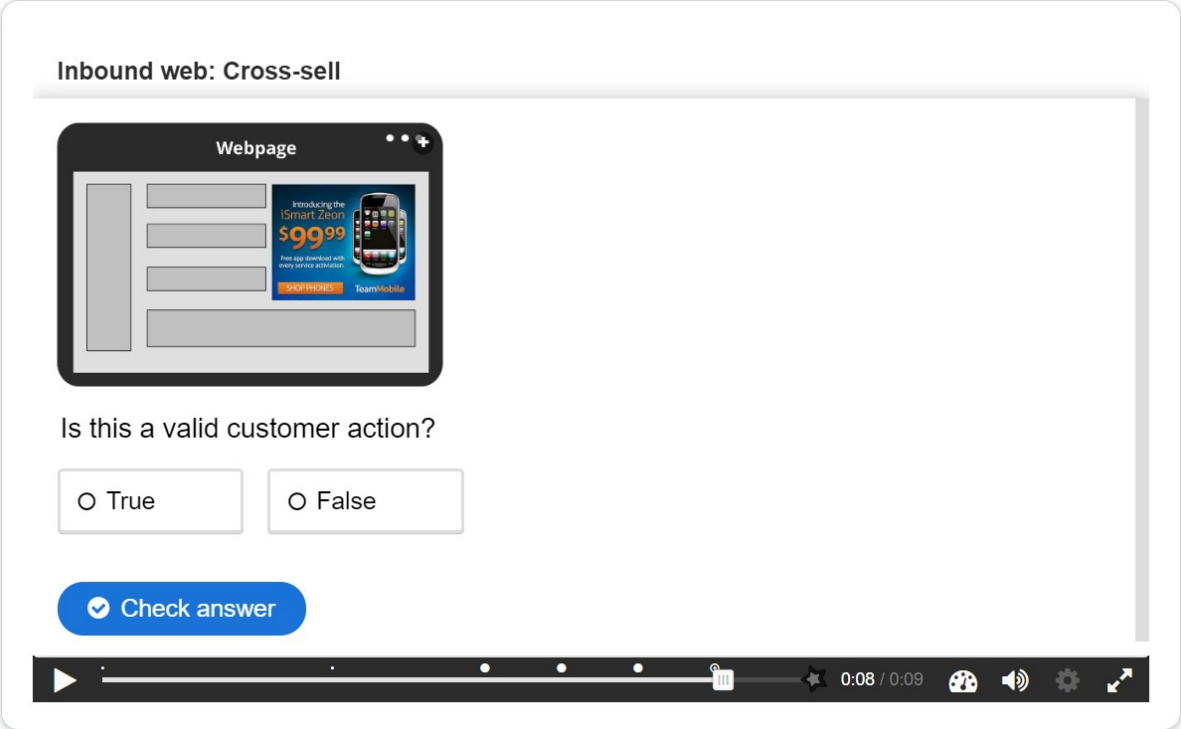

Is this a valid customer action?

- True
- False

Correct feedback: **This scenario is a banner advertisement on a website and is a valid customer action.**

Incorrect feedback: **A banner advertisement on a website is a valid customer action.**

Customer action properties

Every customer action has properties that define its characteristics.

In the following image, click the **+** icons to identify the valid properties of customer actions.

Image is a valid property and represents the image that is used for the banner.

Price is a numeric property that represents the price of the customer action.

Benefit is a valid property that represents the value proposition to the customer.

Short title is a valid property that represents a short description of the customer action.

Call to action is not a valid property. This button is a call to action for people who want to purchase the product.

Company name is not necessarily a valid property, as this text represents the company name.

Business hierarchy

In Pega Customer Decision Hub, customer actions apply to various business issues and are organized into a three-level hierarchy. The business issue is the purpose of the actions that you offer to your customers. Each action that you prepare for your customers is associated

with an issue and a group. Because of the association, you must always define the hierarchy

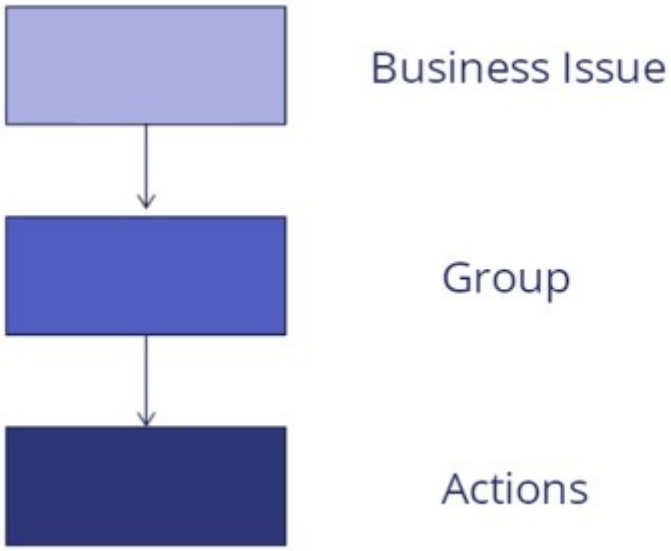

before creating actions.

In the following image, click the **+** icons to learn more about the three-level hierarchy. Then, check your knowledge with the following interaction.

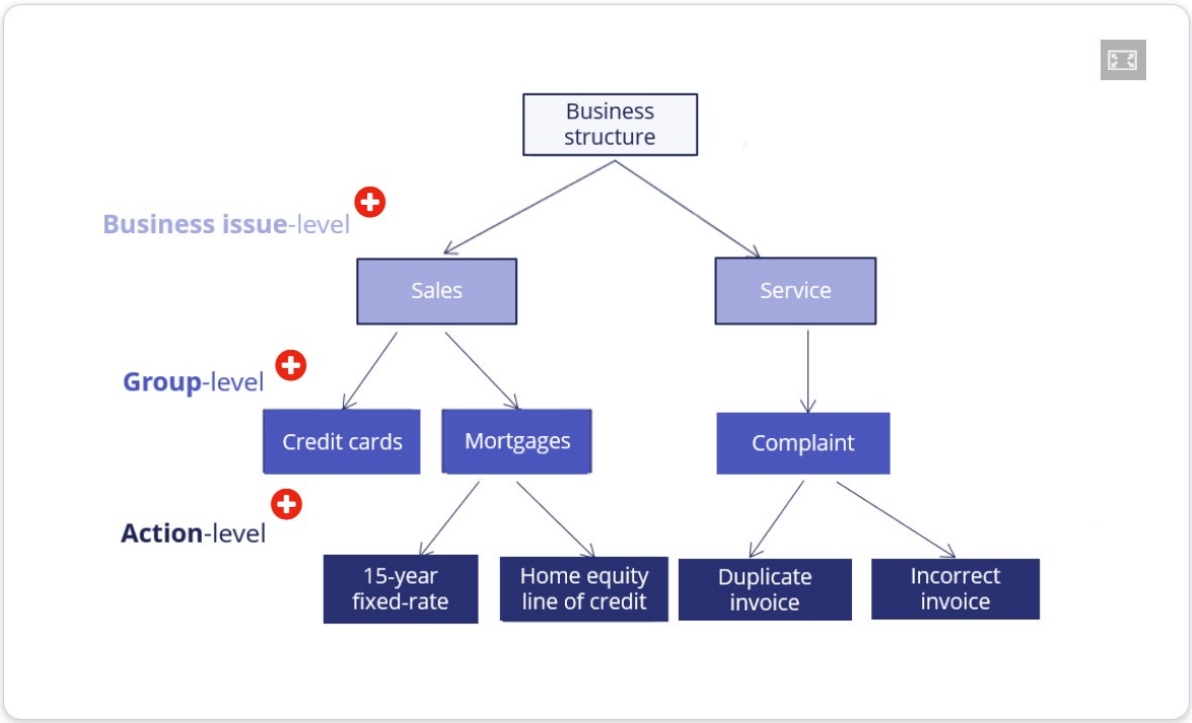

Business issue: In Pega Customer Decision Hub, a **business issue** represents the business area for which a customer action is applicable (for example, **Sales**, **Retention**, and **Service**).

Group: In Pega Customer Decision Hub, a **group** organizes actions into logical buckets (for example, **Credit Cards** or **Mortgages**).

Actions: In Pega Customer Decision Hub, an offering is referred to as an **action** (for example, **Reward card** or **30-year fixed rate)**. Each action that you prepare for your customers is associated with an issue and a group.

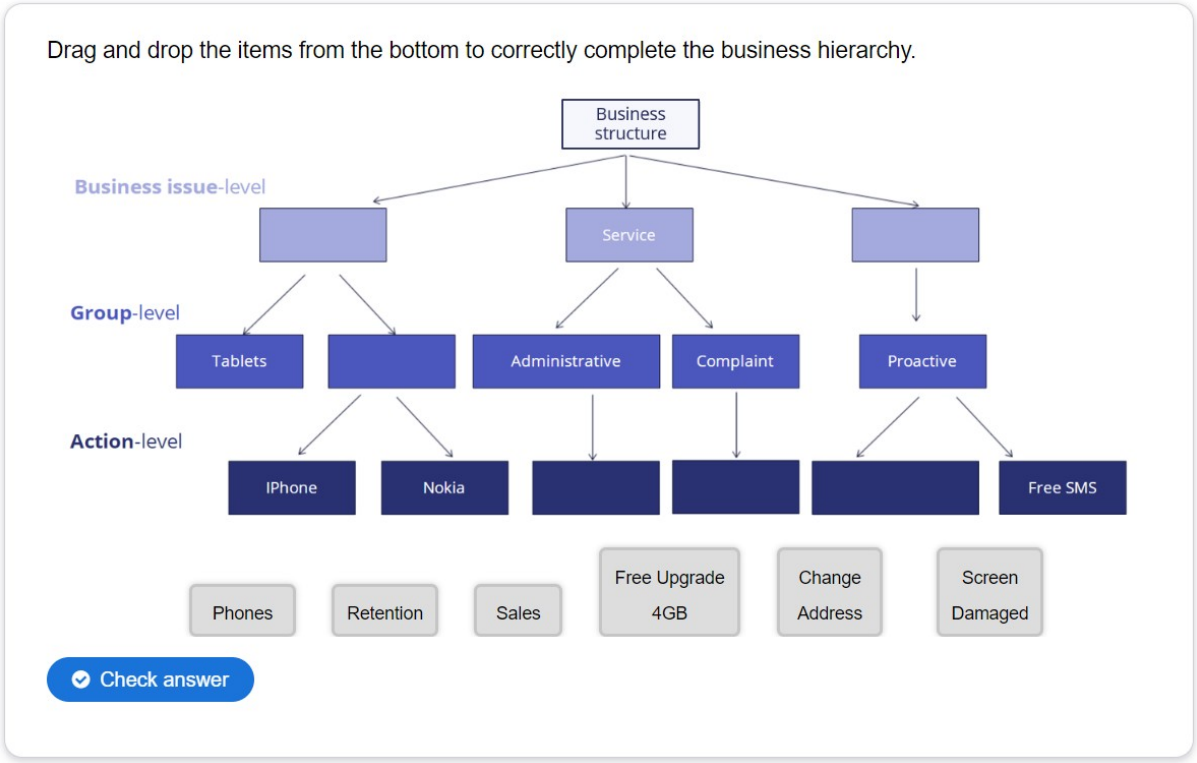

Drag and drop the items from the bottom to correctly complete the business hierarchy.

Correct feedback: **That's right! You selected the correct response.**

Incorrect feedback: **You did not select the correct response. Please try again.**

Customer actions and renaming

Pega Customer Decision Hub uses the default term actions to refer to messages delivered during customer interaction. If needed, you can change the term actions to a more appropriate term. For example, actions under a Loans group (Sales issue) can be changed to Loan offers whereas actions under a Loans group (Retention issue) can be changed to Loan loyalty program.

In the center of the following image, slide the vertical line to view a sample hierarchy with renamed actions. Then, check your knowledge with the following interaction.

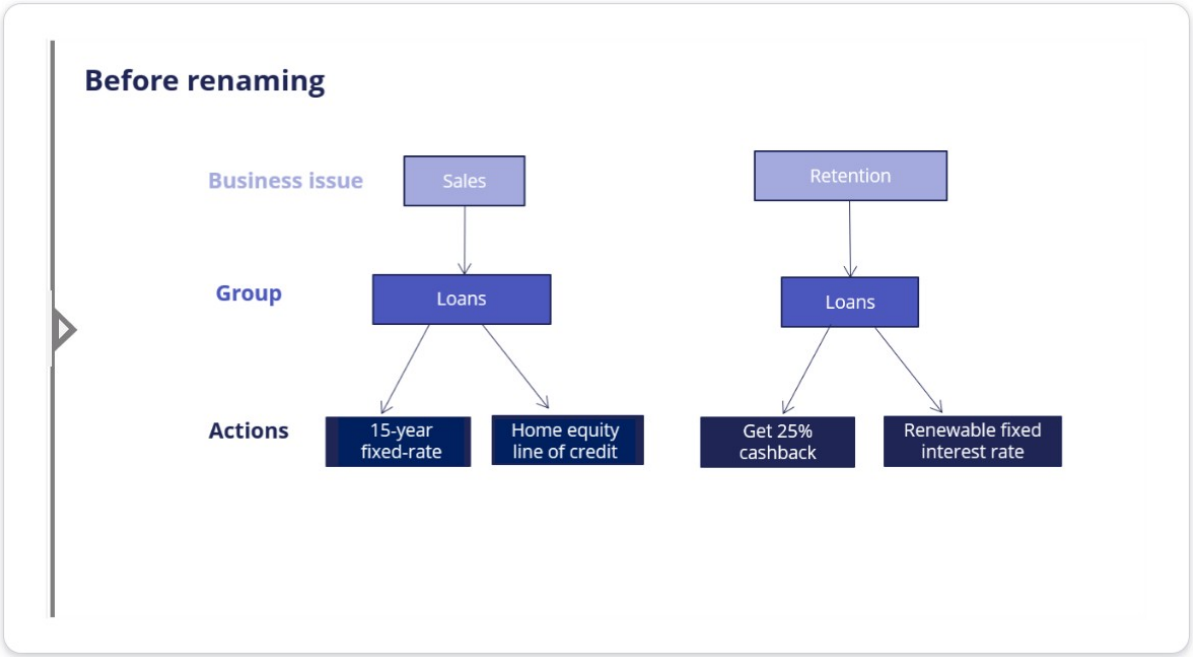

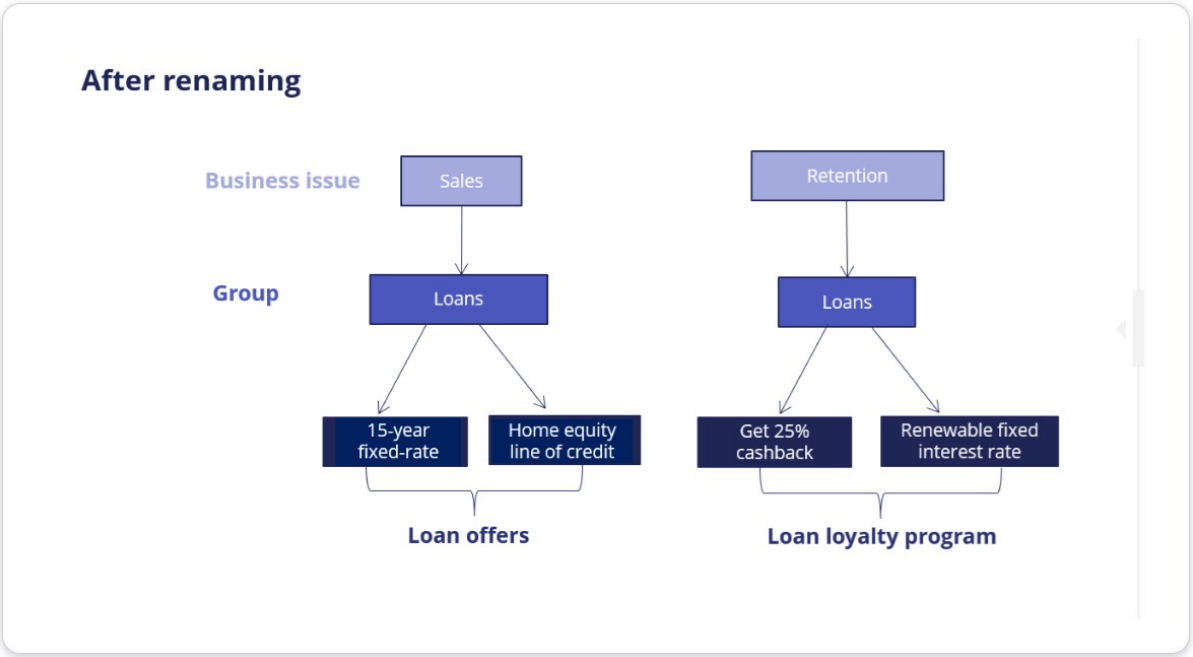

Set 1: A customer action is any banner ad, retention offer, or service message.

Customer actions are properties that define their characteristics.

Set 2: A customer action can be presented on any customer interaction channel such as Web, Mobile, Contact center, or Social.

A customer action can be presented only on Web.

Set 3: A customer action is organized into a hierarchical structure. The default hierarchy is: Business issue -> Group -> Action.

A customer action is organized into a hierarchical structure. The default hierarchy is: Business issue -> Action -> Group.

Choose the correct statement.

Correct feedback: **Well done!**

Incorrect feedback: **We recommend you go through the material once more.**

Managing business structure

Introduction

In Pega Customer Decision Hub™, Next-Best-Action Designer is used to build the business issue/group hierarchies for your organization. Together, business issues and groups form the organizing structure for your customer interactions. Each Next-Best-Action that is presented to a customer is associated with a business issue and group.

Transcript

This demo will show you how to set up your business structure in Next-Best-Action Designer.

This is the Pega Customer Decision Hub™ portal. To manage the business structure, navigate to Next-Best-Action Designer. Customer actions are organized using a hierarchical business structure called Issues and Groups. Under the Taxonomy tab, you define the Issues and Groups that will play a role in the Next-Best-Action decisioning process.

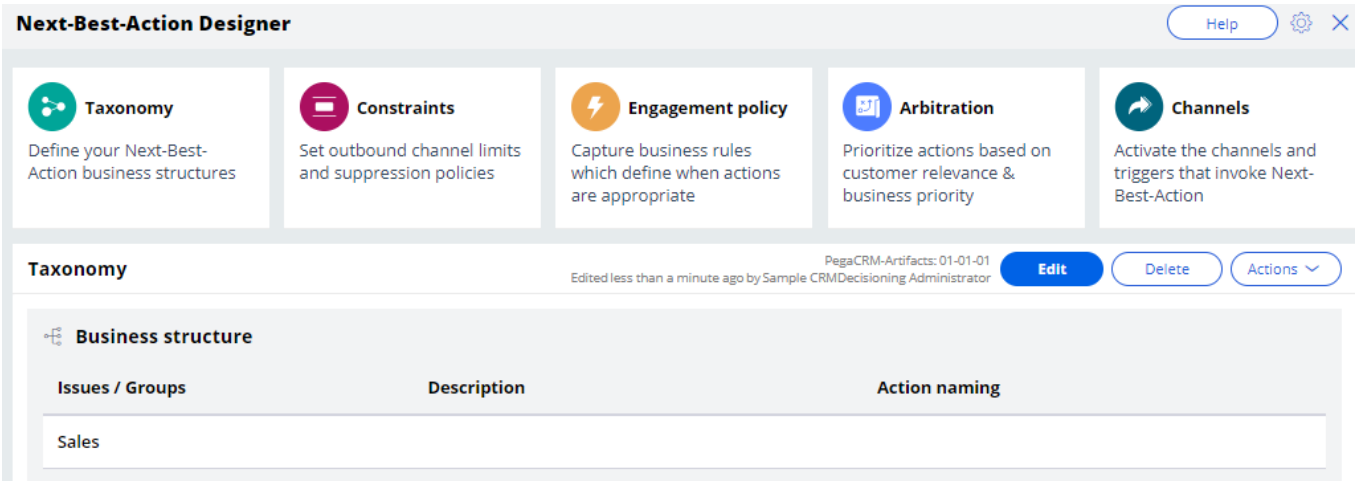

In this case, you will set up U+ bank's business structure based on their various areas of business focus. Currently, there is only one Issue in the business structure.

You can add, remove, or create new Issues here by editing and configuring the hierarchy. Note that, this can be edited only with an additional privilege. Here is a list of Issues that were created before. You can also create new Issues.

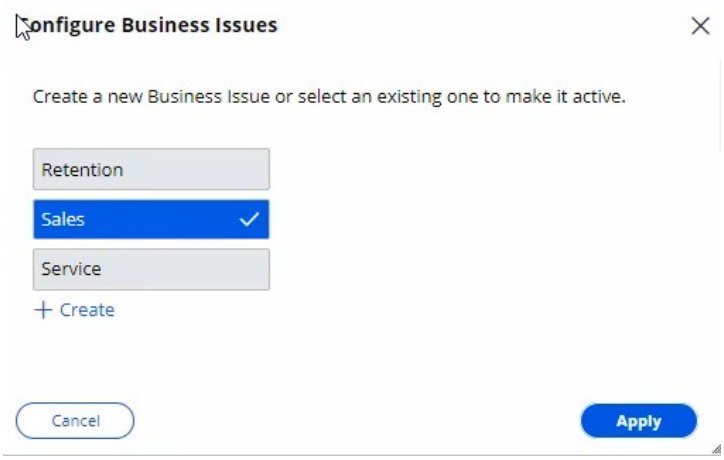

Now, add an existing Issue. Similarly, you can add, remove, or create business Groups. To add more groups, configure the groups.

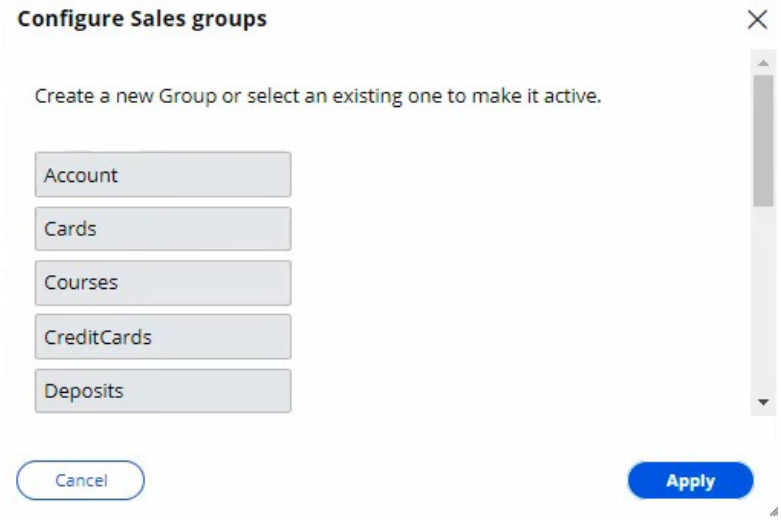

To complete the configuration, save the changes.

This demo has concluded. What did it show you?

- How to add an Issue to the hierarchy.
- How to add Groups to an Issue.

Creating actions

Introduction

An action holds various details about a particular offering, such as Start Date, End Date, and Expected Revenue. Each action is backed by a decisioning proposition. The system automatically manages this relationship by creating, deploying, and deleting the proposition instances as needed. Since an action is closely tied to a proposition, it must always be created in the context of an issue and group.

Transcript

This demo will show you how to create an Action. It will also explain how to create Actions in bulk.

U+, a retail bank, recently introduced several new credit card offers that they would like to display to their customers. As a decisioning consultant, you have been tasked with creating these offers in the **Sales** Issue and **CreditCards** Group as per the pre-defined business hierarchy.

This is the Pega Customer Decision Hub™ portal. Actions are created and managed on the Actions landing page.

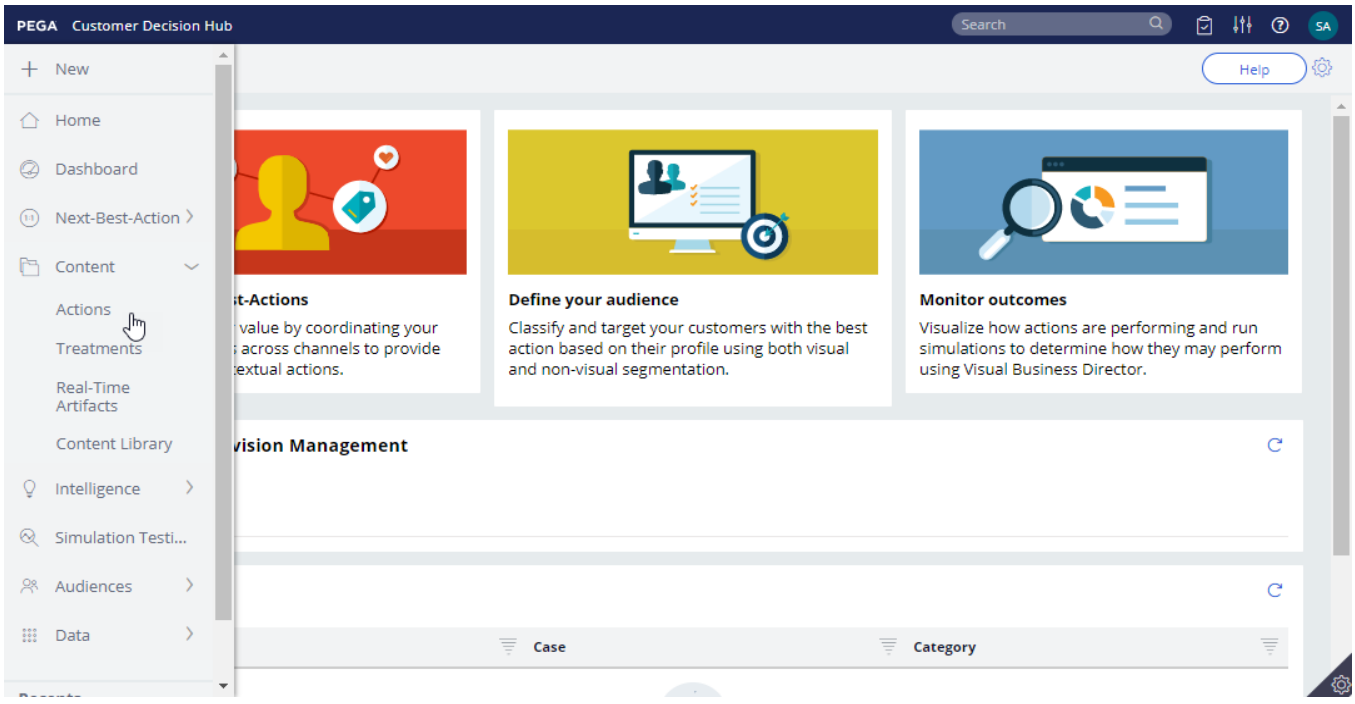

You can either create a single Action or Actions in bulk.

To create a single Action, enter a short description of the new Action. Select the appropriate **Issue** and **Group**. Now, open the Action.

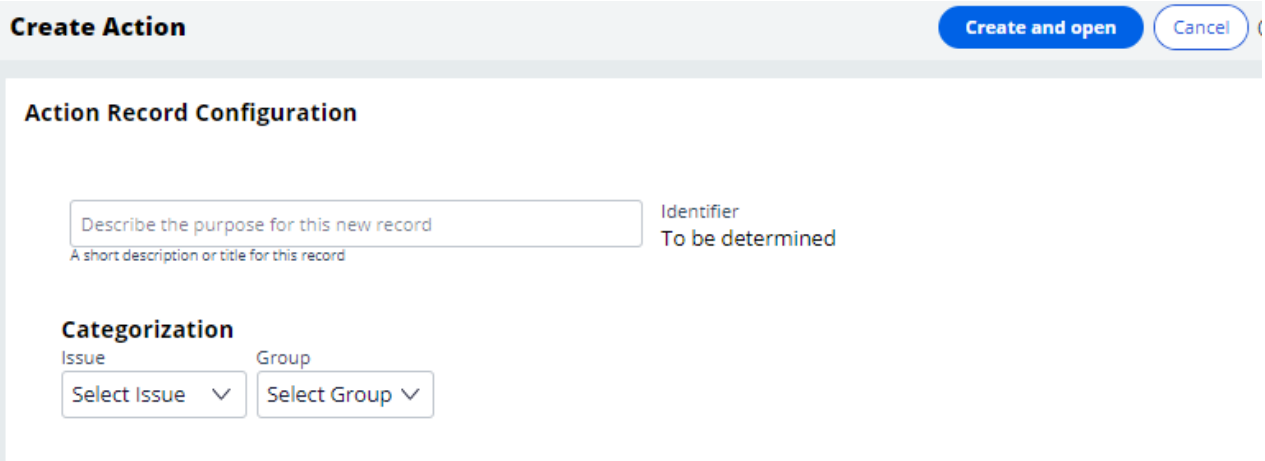

The attributes of an Action distinguish it from other Actions. Attributes are used by the Customer Decision Hub to select the right Action for a customer. Fill in the attributes relevant to this Action. For example, provide a **Description** that explains its purpose and **Benefits** that describe how this Action will benefit the customer.

Bear in mind that some of the values you enter may be customer-facing information. For example, U+ wants to display the content of the Benefits attribute when this Action is presented on its website. *To complete the configuration,* **save** *the Action.*

To view Actions within a certain Group, you can filter by a specific **Issue** and **Group**. This is the Action you just created. Thus far you have created a single Action. However, sometimes you may want to create multiple Actions at once. You can do that by uploading a list of Actions from a Comma Separated Value (CSV) file. To create such a CSV file you need to download a CSV template and fill it in.

Open the downloaded CSV template. You can use this template to add/update Actions. Add new Actions with corresponding attributes to the file.

Properties that start with py, such as pyName, pyIssue, pyGroup, pyLabel and pyIsPropositionActive are mandatory internal properties. It is important that you input these values in the correct format into the CSV file.

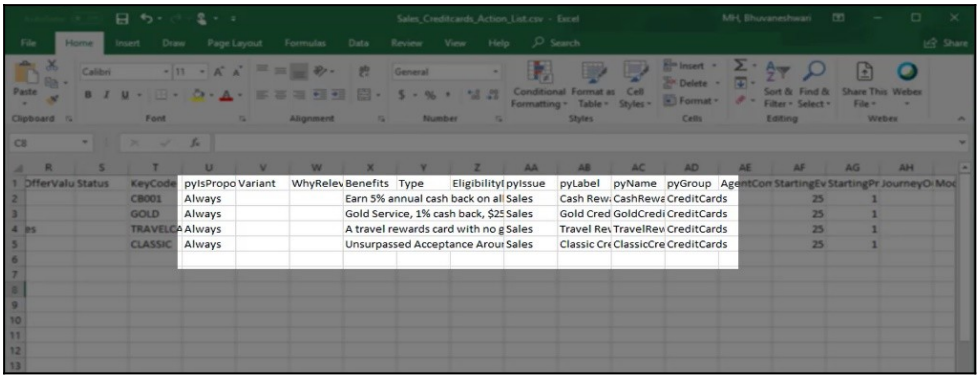

The property pyIsPropositionActive corresponds to the Availability property, which is visible on the Details tab of an Action. Possible values for this property are 'Always', 'Never' and 'Date'.

Choose 'Always' to ensure that the Action is selected during the Next-Best-Action decisioning process by the Customer Decision Hub.

Set it to 'Never' if you never want the Action to be selected, for example when you want to retire the Action

If you set the value to 'Date', you should enter a date range in the 'StartDate' and 'EndDate' fields. The Customer Decision Hub will select this Action only during that time period.

pyIssue corresponds to the Issue the Action was created for.

pyLabel corresponds to a short description of the Action.

pyName corresponds to the Action name.

pyGroup is the Group to which the Action belongs.

Save and close the CSV file.

Now, import Actions from the file**.** Select the CSV file that contains the list of Actions you want to import. If you want to delete existing Actions from this Issue/Group that have been created in the system but not in the CSV file, you can select this option.

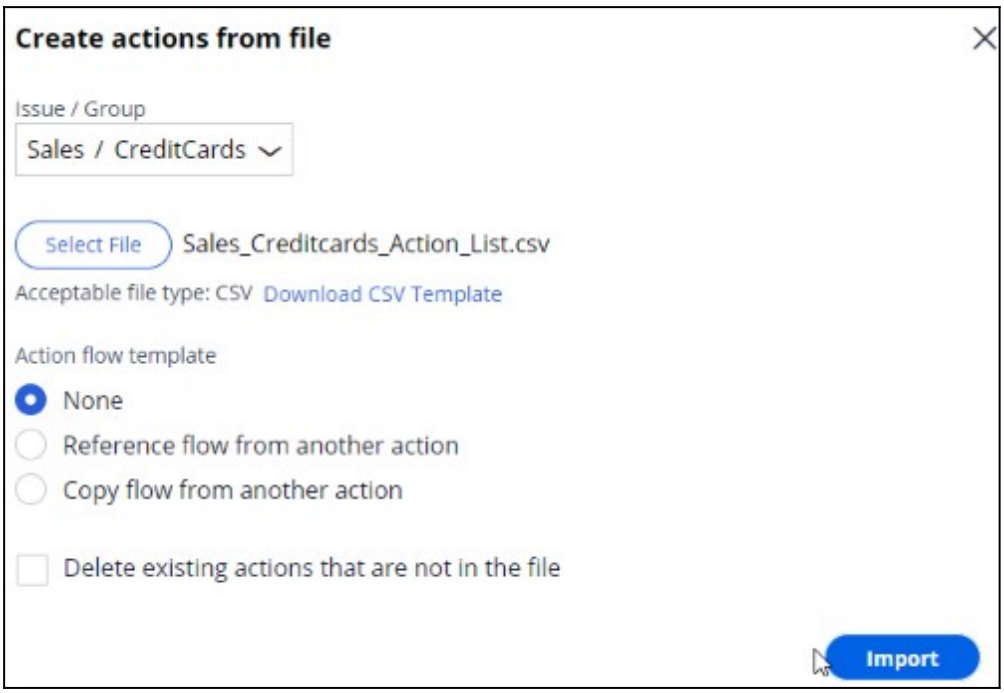

For now, you are only interested in adding new Actions.

Notice that the window displays a summary of the result of the import.

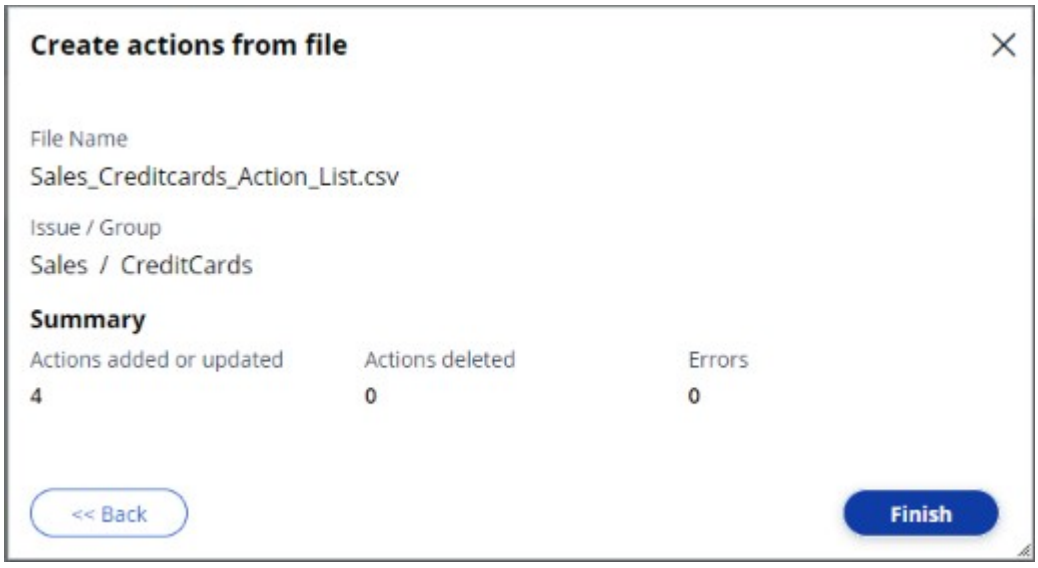

For example, when a duplicate Action is added, the Errors count is incremented, and you can download the error file to learn the exact issue. Complete the import. Notice that the new Actions are now listed on the **Actions** page.

This demo has concluded. What did it show you?

- How to create an Action.

- How to set Action attributes.
- The most important properties of Actions in a CSV file.
- How to create multiple Actions using a CSV file.

Renaming actions

Introduction

By default, Pega Customer Decision Hub uses the term ‘actions’ to refer to messages delivered during a customer interaction. If needed, you can change this to a more appropriate term, such as Promotions or Nudges. Actions can be renamed with different terms based on the business issue and group to which the product belongs.

Transcript

This demo will show you how to change the default naming of Actions to something that better reflects their business context.

This is the Pega Customer Decision Hub™ portal. On the **Actions** landing page, you can view all the Actions that have been created. To view Actions within a certain Group, you can filter by a specific **Issue** and **Group**. In this case, filter Actions in the **Sales** Issue and **CreditCards** Group. Open an Action to view the name. Notice that the Sales/CreditCards Group uses the default terminology ‘**Action**’.

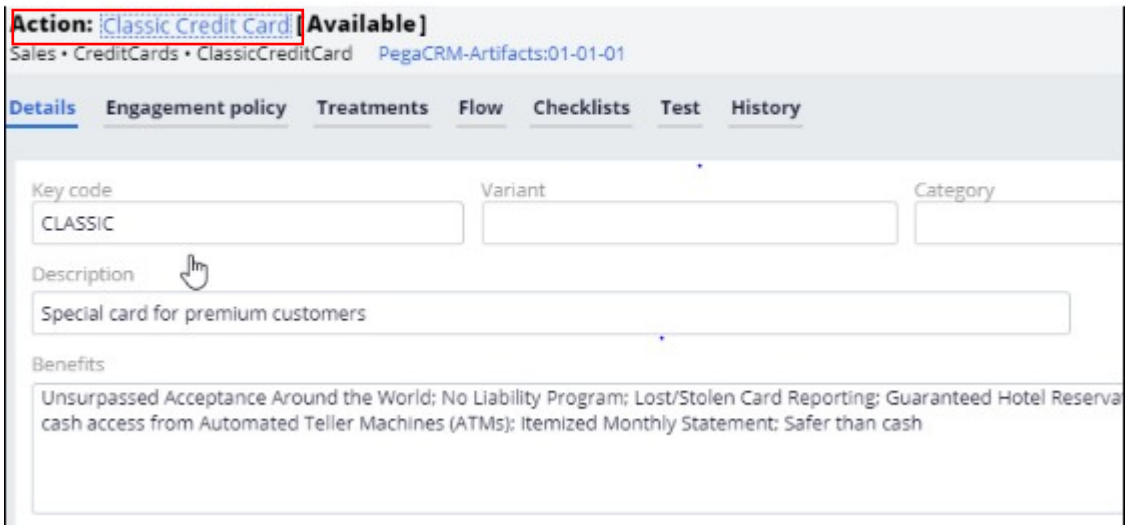

U+ bank wants to rename Actions under the **Sales** Issue and **CreditCards** Group to better suit their business purpose. The Actions under the **CreditCards** Group are always promotional credit card offers, therefore the bank wants to change the naming convention from ‘Action’ to ‘Offer’.

You can change the default naming convention of Actions in Next-Best-Action Designer.

To manage the business structure, navigate to the Taxonomy tab of Next-Best-Action Designer. The 'Action naming' convention is set in a business Group's configuration settings. This can be set for each group in the business hierarchy.

Edit the hierarchy to modify the Group attributes.

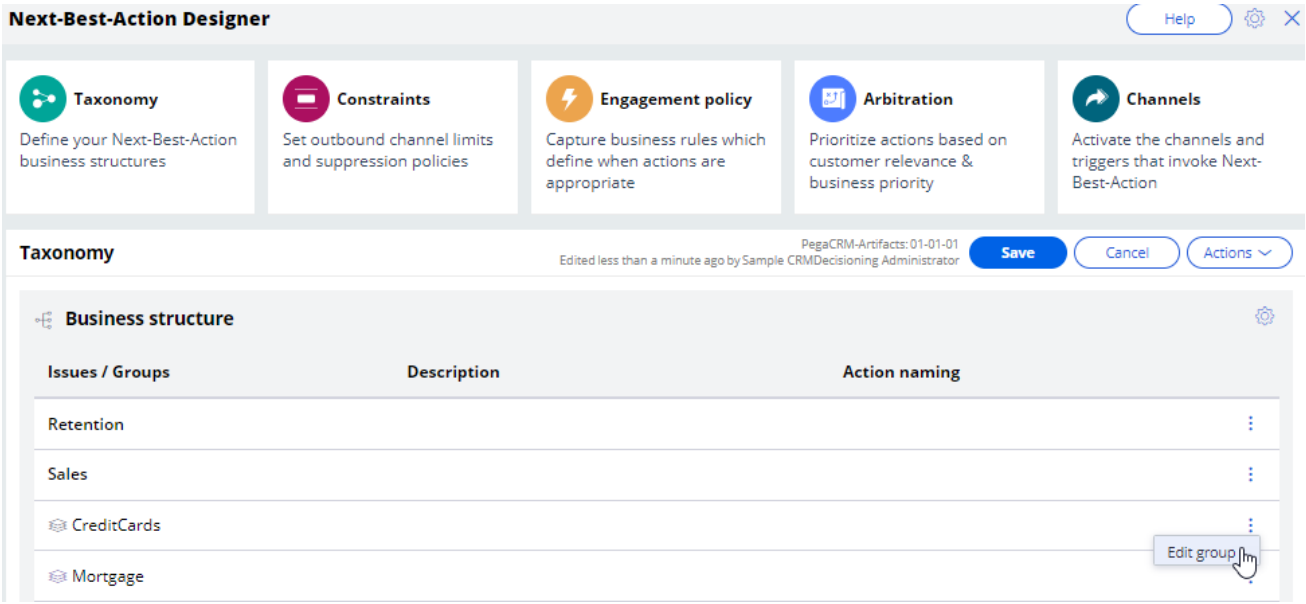

In this case, change the Action naming convention to Offer.

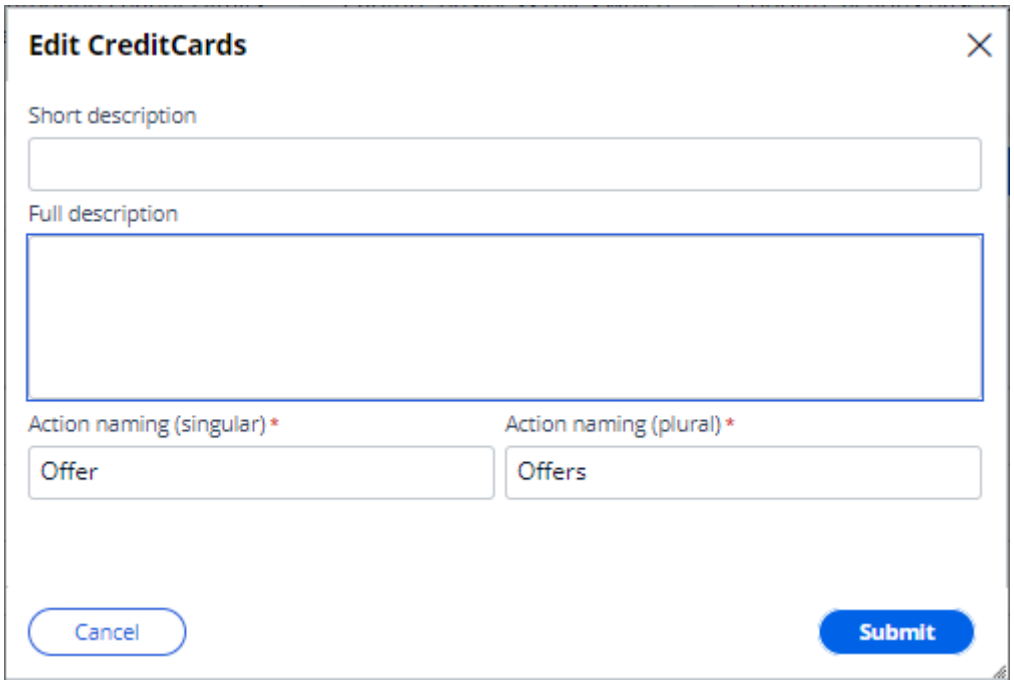

To complete the configuration, save the changes. Verify that the Action is renamed. Open any Action under the **Sales** Issue and **CreditCards** Group to view the effect of renaming.

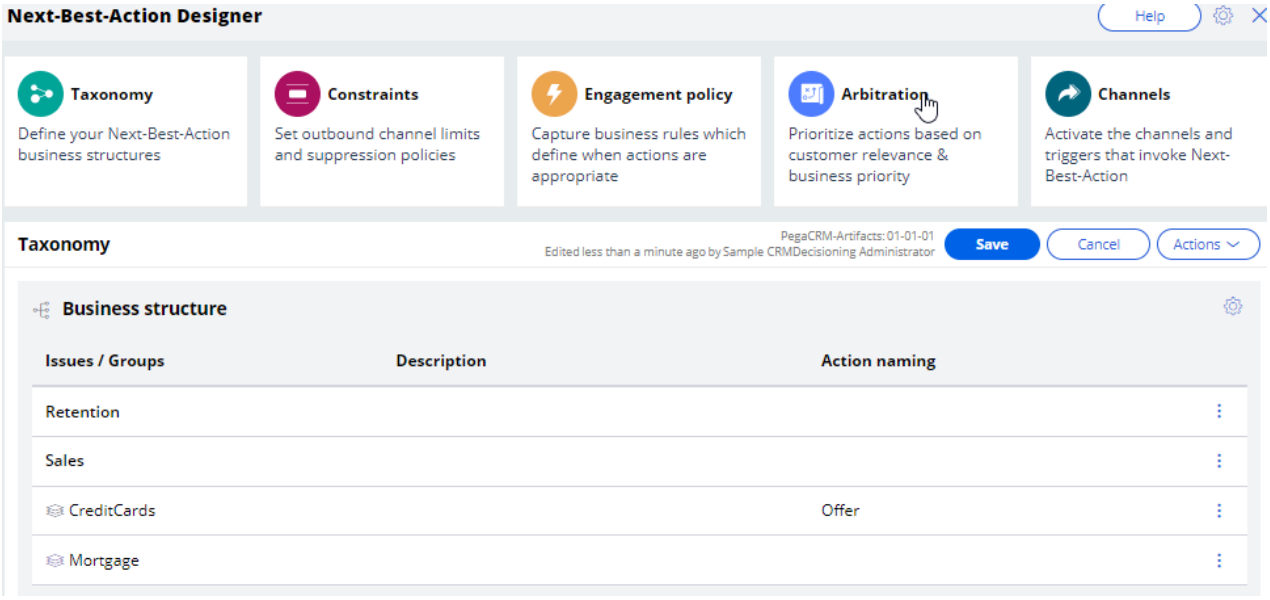

From now on, Actions under Sales/CreditCards will be referred to as Offers.

This demo has concluded. What did it show you?

- How to change the default naming convention of Actions, to a name that better reflects their business context.

Presenting a single offer on the web

Description

Learn about real-time containers and how they are used to display actions selected by Pega Customer Decision Hub in real-time channels. Gain experience creating and configuring real-time containers by learning how to configure Next-Best-Action Designer to select a single offer that will be displayed on a website.

Learning objectives

- Describe the concept of real-time containers
- Explain how to create and configure a real-time container
- Explain how to configure Next-Best-Action designer to display an offer on the web

Real-time containers

Introduction

A real-time container is a service that manages communication between Pega Customer Decision Hub and external channels. An external real-time channel is any channel that presents actions selected by the Customer Decision Hub to a user or customer. For example, a website, a call-center application, or a mobile application.

Transcript

This video explains the concept of real-time containers, which manage communication between the Pega Customer Decision Hub™ and external channels.

This is the website of a retail bank called U+ Bank.

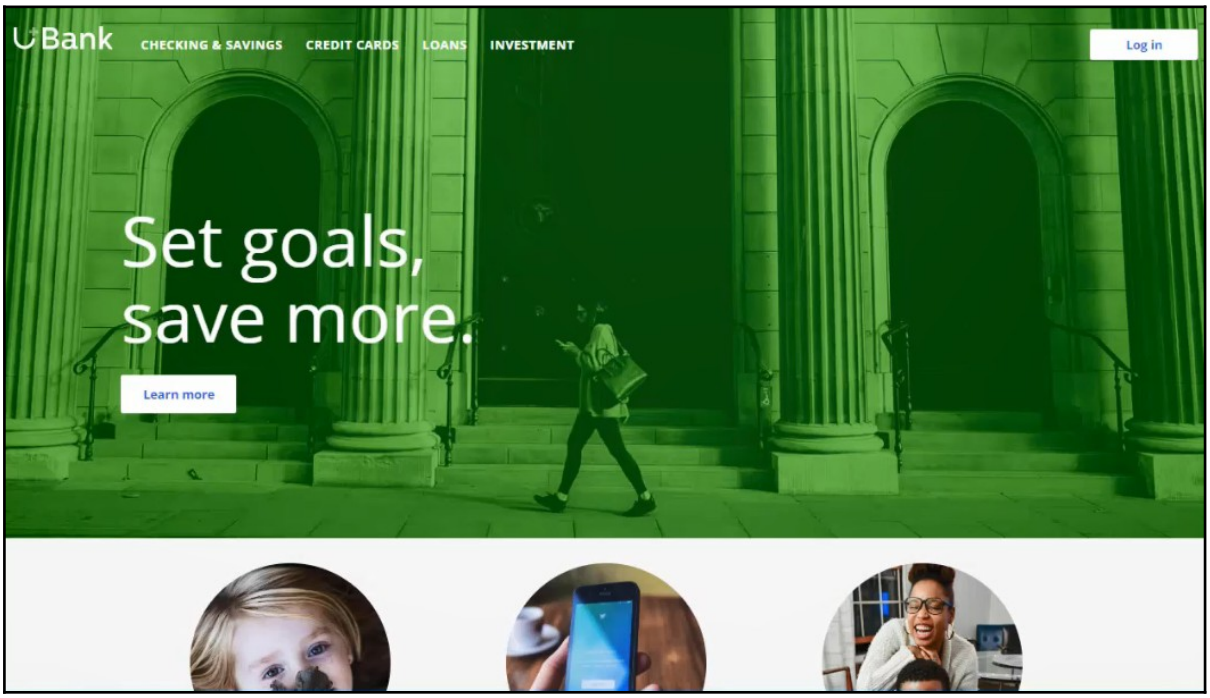

The bank plans to promote new offers on the account page, which is displayed when customers log in.

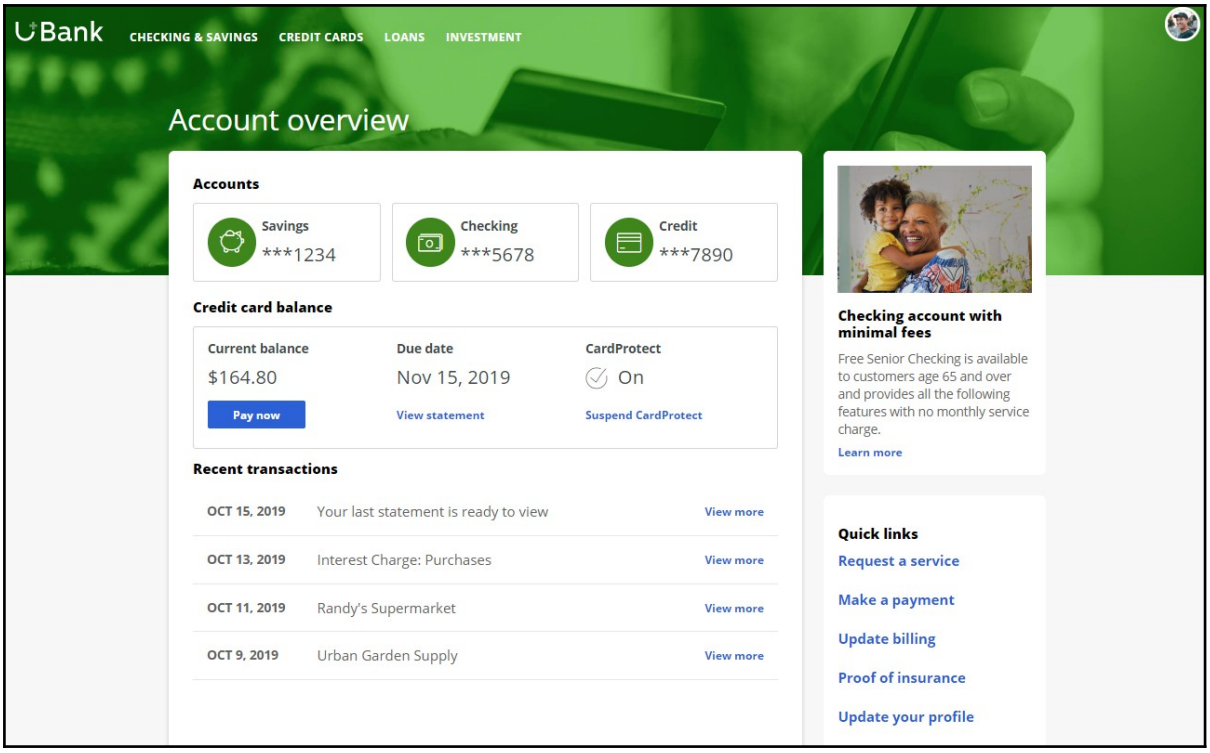

The marketing department wants to leverage the Customer Decision Hub's Next-Best-Action capability to display the right offer for each customer.

The Customer Decision Hub's real-time container functionality is used to implement this requirement.

A real-time container is a placeholder for content in an external real-time channel.

An external real-time channel is any channel that presents actions selected by the Customer Decision Hub to a user or customer. For example, a website, a call-center application, or a mobile application.

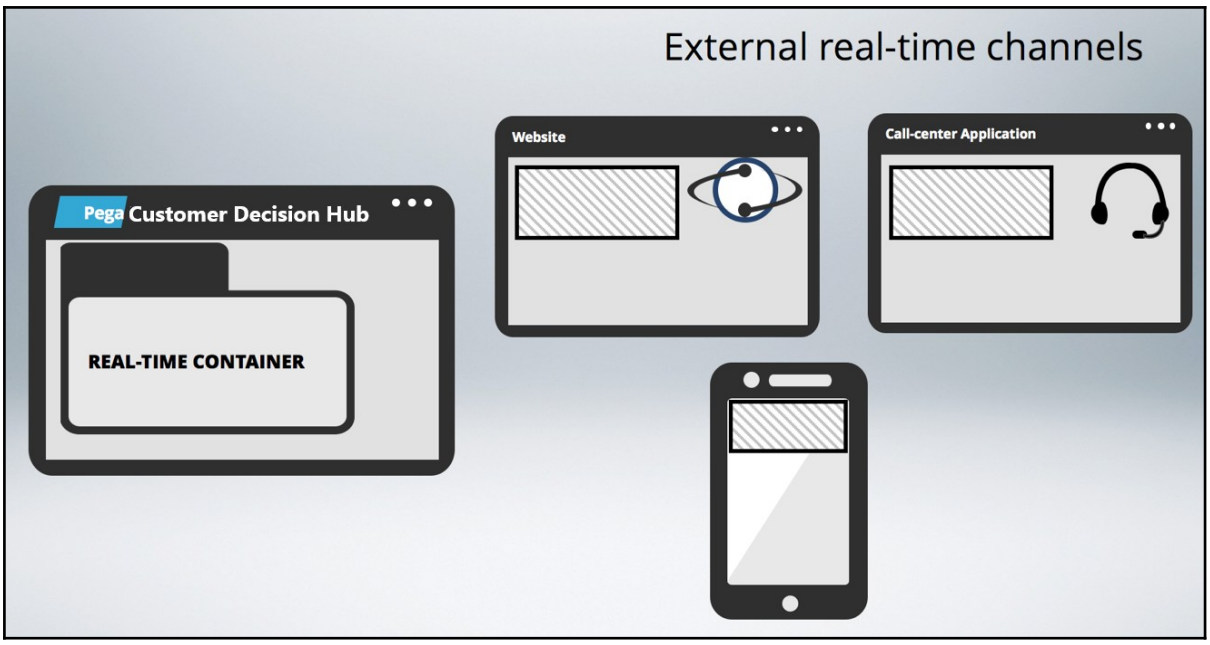

Here's how the website invokes the real-time container to present credit card offers on the account page.

In the Customer Decision Hub's Next-Best-Action Designer, a real-time container called "Account Page Container" is configured. The website invokes this real-time container before loading the account page.

The Customer Decision Hub then evaluates the actions from the associated Issue/Group, which in this case is Sales/CreditCards and returns the resulting offer details back to the website.

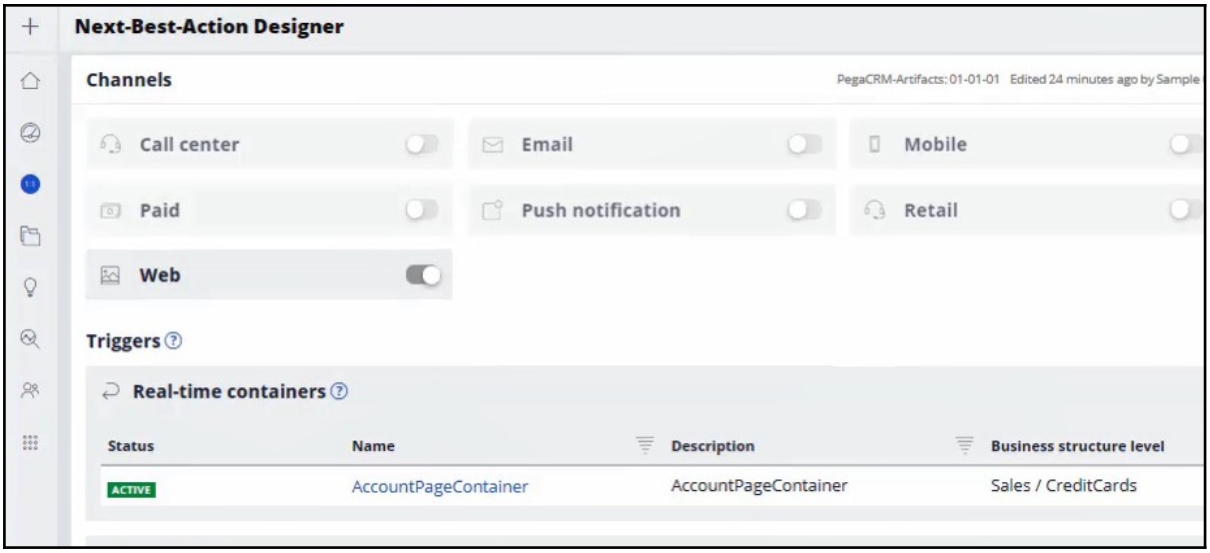

The website then loads the account page with the content returned by the Customer Decision Hub, such as the offer image, description and other relevant attributes.

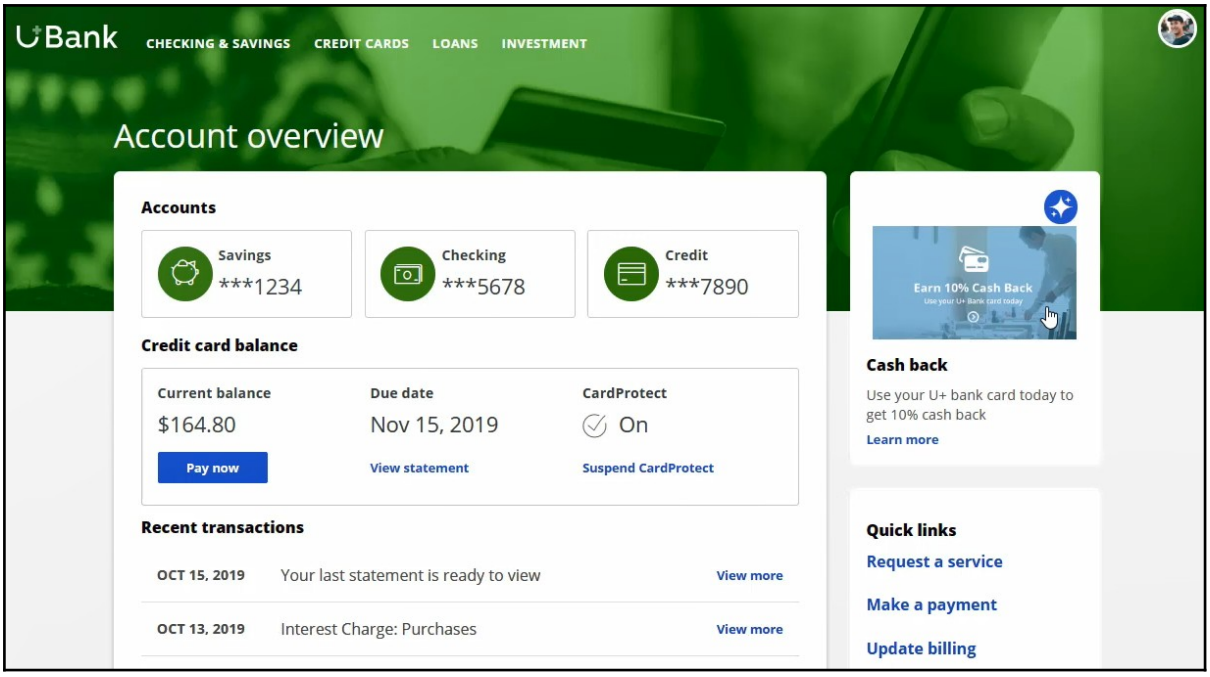

Meanwhile, the Customer Decision Hub records these customer interactions in the Interaction History. An Impression is recorded to indicate that the action was shown to the customer, a Click-through is recorded when the customer clicks on the action. Marketers use these metrics, i.e. Impressions and Click-throughs, to measure the level of customer engagement, and therefore, the success of the marketing effort.

Creating and configuring a real-time container

Introduction

Learn how to create and configure a real-time container which allows you to manage communication between Pega Customer Decision Hub and external channels.

Transcript

This demo will show you how to create and configure a real-time container.

To create real-time containers, navigate to the Real-Time Artifacts landing page.

Enter a short description of the new real-time container.

An external channel uses the real-time container's identifier when it makes a request to the Pega Customer Decision Hub.

You can change the identifier to a value that the external channel in your environment expects.

For example, the U+ Bank website is pre-configured to call a real-time container with the identifier Ubank_home_banner.

Now, open the real-time container.

Here you can turn the real-time container's ability to respond to requests from external channels on or off.

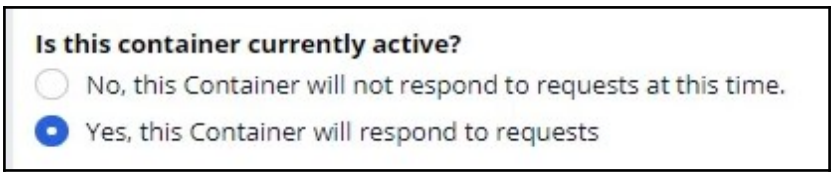

Keeping the default value will allow the Customer Decision Hub to serve requests from the website.

‘Impression capture’ allows the external channel to control how the Customer Decision Hub records impressions.

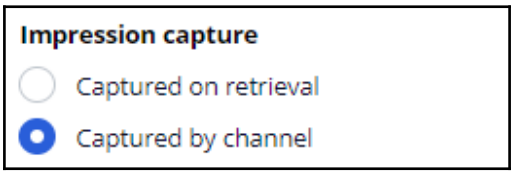

‘Captured on retrieval’ means that the Customer Decision Hub records an impression in the Interaction History immediately after sending the action details to the external channel.

‘Captured by channel’ means that the Customer Decision Hub does not immediately record an impression. Instead, the external channel can explicitly request that the Customer Decision Hub records an impression at a later point in time.

In this case, the bank wants the impressions not followed by a click to be captured, so keep the default value.

To record a click in the Interaction History without initiating an action flow for the customer, keep the default value.

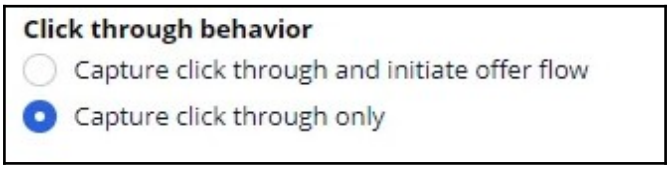

This option is suitable for actions that do not have a defined flow, for example, on the Web channel when a customer clicking on an ad does not trigger any follow-up steps.

Here you can view the list of Next-Best-Action strategies that you have associated with this real-time container.

Save the configuration.

Navigate to the Containers tab to see the newly created real-time container.

This demo has concluded. What did it show you?

- How to create a real-time container

- How to configure the real-time container settings

Presenting a single offer on the web

Introduction

Pega Customer Decision Hub is the always-on customer brain used to select the right offer to be presented to each customer in any real-time channel. Learn how to configure Next-Best-Action Designer to select a single offer that will be displayed on a website.

Transcript

This demo will show you how to configure Next-Best-Action Designer to select a single offer that will be displayed on a website.

U+, a retail bank, would like to use the Pega Customer Decision Hub™ to display a single offer on its website.

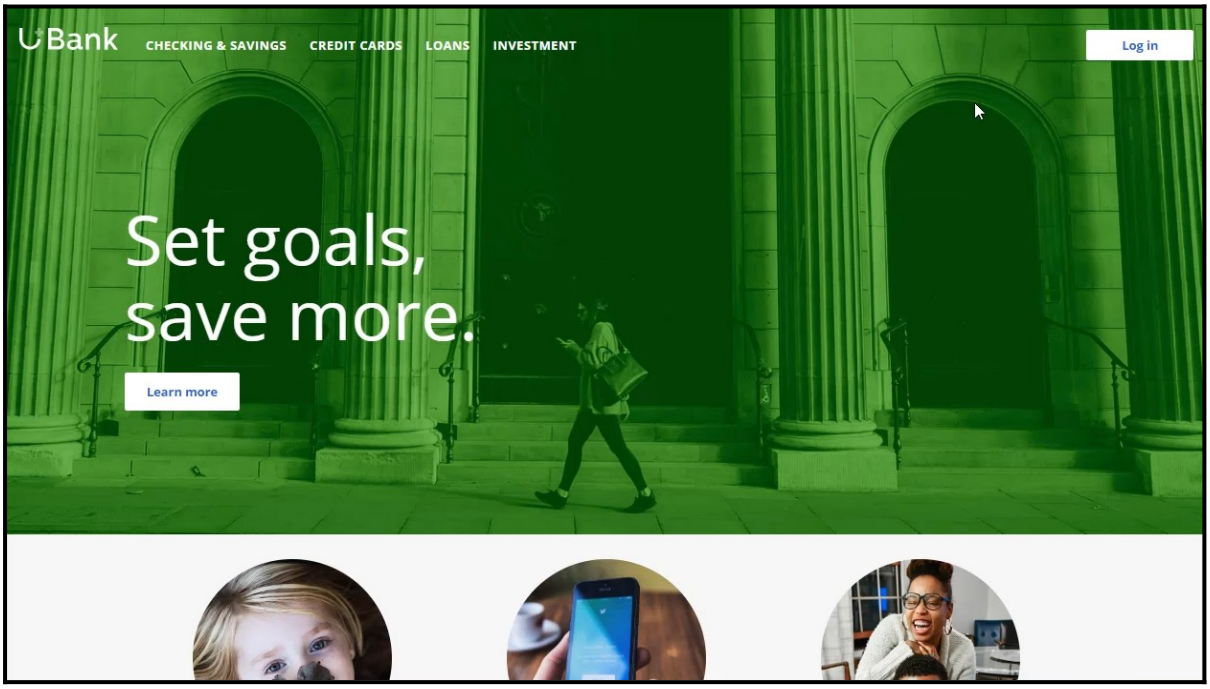

U+ bank wants to make offers related to credit cards and display the same ‘Cash back’ offer to every customer who logs in to the web site.

For example, if customer Troy logs in to his Accounts page, the ‘Cash back’ offer is displayed.

If Troy clicks on the ‘Learn more’ button, it takes him to the Offers page.

This page shows the offer details.

To implement this business requirement, navigate to Next-Best-Action designer.

Here you can configure business rules to define when specific actions or groups of actions are appropriate for customers.

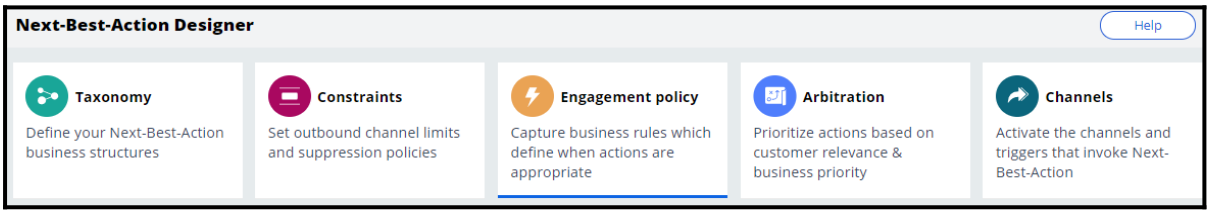

Notice that the NBA hierarchy currently has three Business issues with Groups under them. In this case, U+ wants to promote credit card offers. So, open the CreditCards Group.

You can configure the actions here.

A consultant has already created a few actions under Sales/CreditCards. However, in this scenario, the bank wants to show only the ‘Cash back’ card.

Now, open on the ‘Cash back’ offer.

To display an offer, you need to add a treatment that is specific to the channel.

In this scenario, U+ bank wants to present the offer on their website, so select the right treatment type.

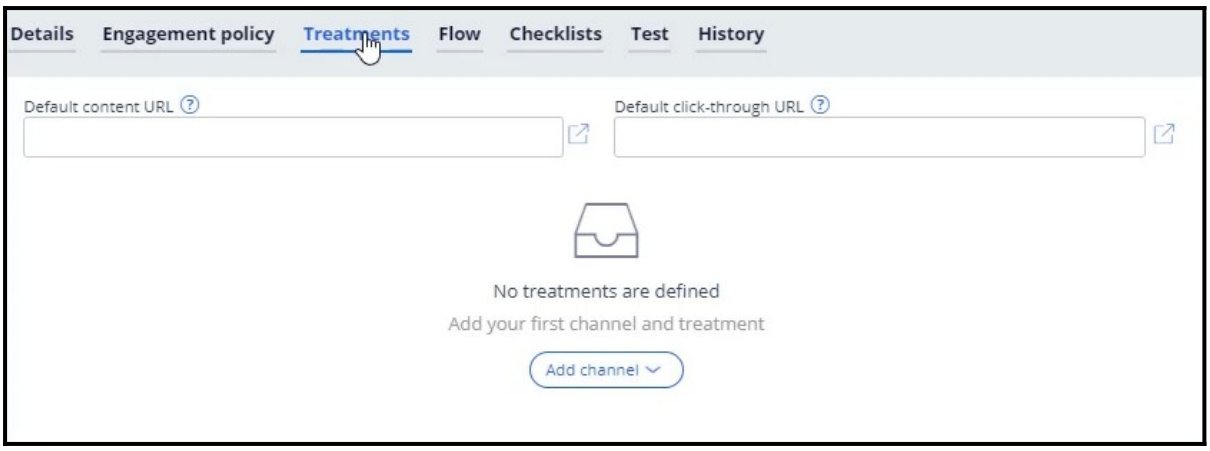

A web treatment represents the message that you want to communicate visually to your customer on the web channel. You can either use an existing treatment or create one here.

Now, fill in the required fields for the web treatment.

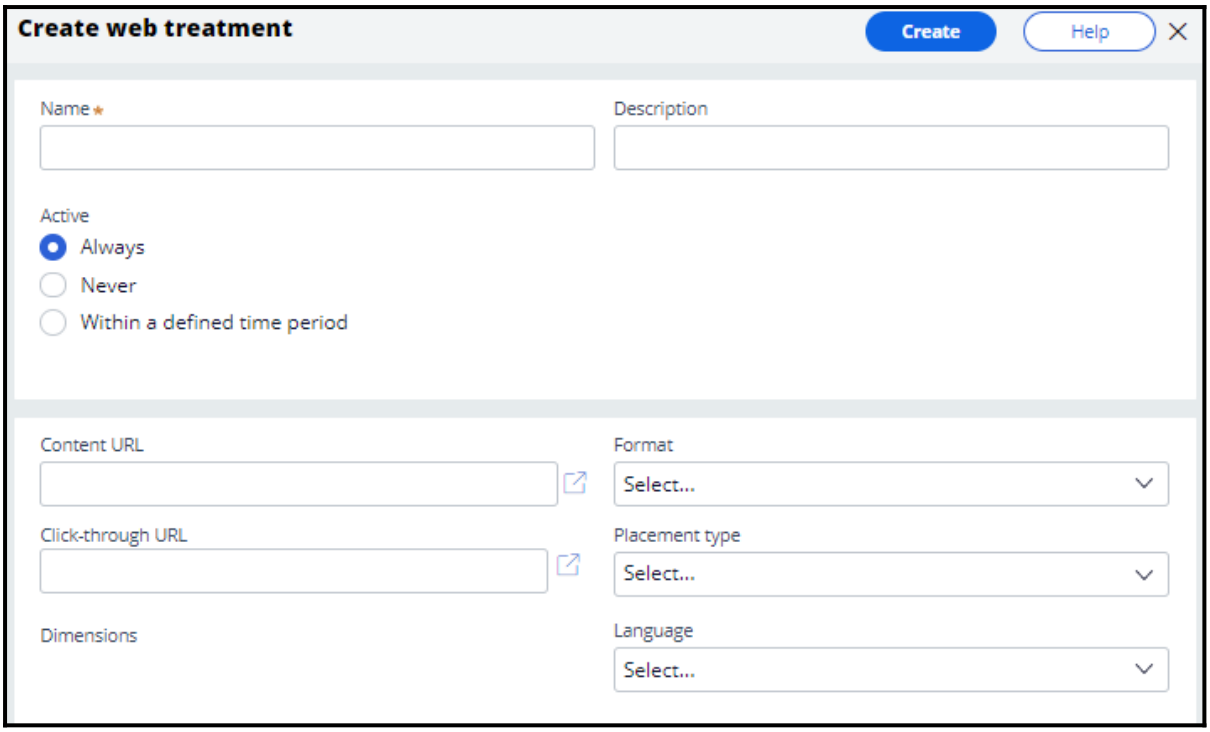

Here you provide a link to an image that will display the action for this treatment.

Once you provide the link, a preview of the content is displayed.

This is the URL that you want the customer to go to when they click on the action.

Now, select the intended location and style of the treatment design. For example, the treatment can appear as a large central banner, a rotating strip of images, or a footer on the website.

In this scenario, select Tile to display the ‘Cash back’ offer on the top right of the Account page.

You can specify the language of the treatment if required. Specifying the treatment language allows Pega Customer Decision Hub to consider the customer’s preferred language when selecting the treatment.

Add the web treatment you just created.

In this phase, U+ does not have any further eligibility or prioritization requirements for this action. Save the changes.

Here, you can enable the channels and triggers that will invoke Next-Best-Actions.

As U+ wants to display the offer on the web, enable the web channel.

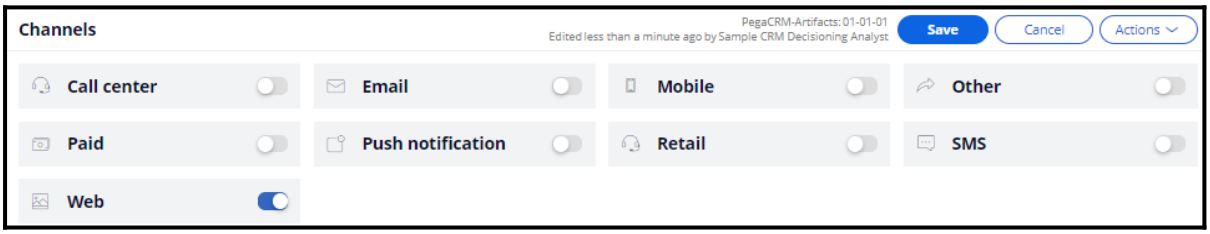

Now, configure the real-time container that the U+ website will use to display the offer banner on the account page.

The real-time container manages communication between the Pega Customer Decision Hub and external channels such as the web and call center. A decisioning consultant has already configured a real-time container for you, so select it to be added to your configuration.

Once the real-time container is added, configure it to select the results from an appropriate Business issue and Group. In this case, U+ bank wants to display the ‘Cash back’ offer, which is under Sales/CreditCards.

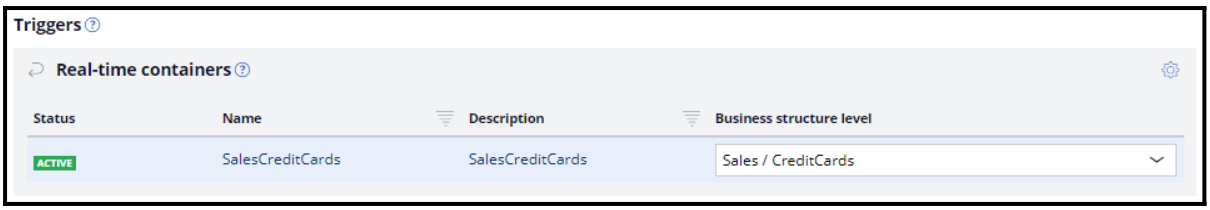

With that, all the necessary configuration for this scenario is complete. Save the changes for the configuration to take effect.

The account page on the U+ Bank website has been configured to use the real-time container with the name 'SalesCreditCards' and with Placement type 'Tile'.

The web treatment is shown as a tile on the top right of the page.

When Troy clicks on the ‘Learn more’ button, it takes him to the URL that was configured as the click-through URL for the treatment.

If another user logs in, they will see the same offer.

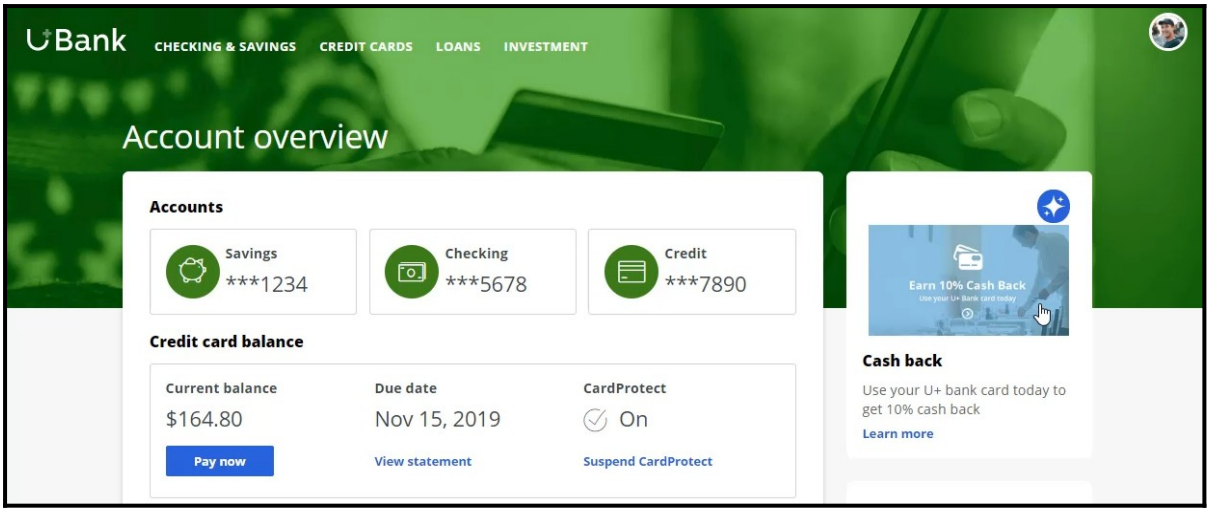

This demo has concluded. What did it show you?

- How to configure Next-Best-Action designer to display a web treatment on the U+ bank website.

Defining customer engagement policies

Description

Engagement policies are the set of conditions such as eligibility, applicability, and suitability that qualify an offer or a group of offers for a customer. As a result of engagement rules, customers will see only those offers that the organization believes they should be exposed to.

Learning objectives

- Define engagement policy conditions in Next-Best-Action Designer at the action group level
- Define engagement policy conditions at the action level
- Differentiate between eligibility, applicability and suitability conditions
- Conduct persona test for customer personas with specific characteristics to evaluate the next-best-action strategy results

Customer engagement policies

Introduction

Engagement policies are a set of business rules and practices used by the organization to determine which customers are eligible for which Next-Best-Actions. These policies allow you to specify the conditions under which an action or group of actions are eligible for a customer.

Transcript

This video explains the concept of customer engagement policies.

The Pega Customer Decision Hub™ combines analytics, business rules, customer data, and data collected during each customer interaction to create a set of actionable insights that it uses to make intelligent decisions. These decisions are known as the Next-Best-Action. Every Next-Best-Action weighs customer needs against business objectives to optimize decisions based on priorities set by the business manager.

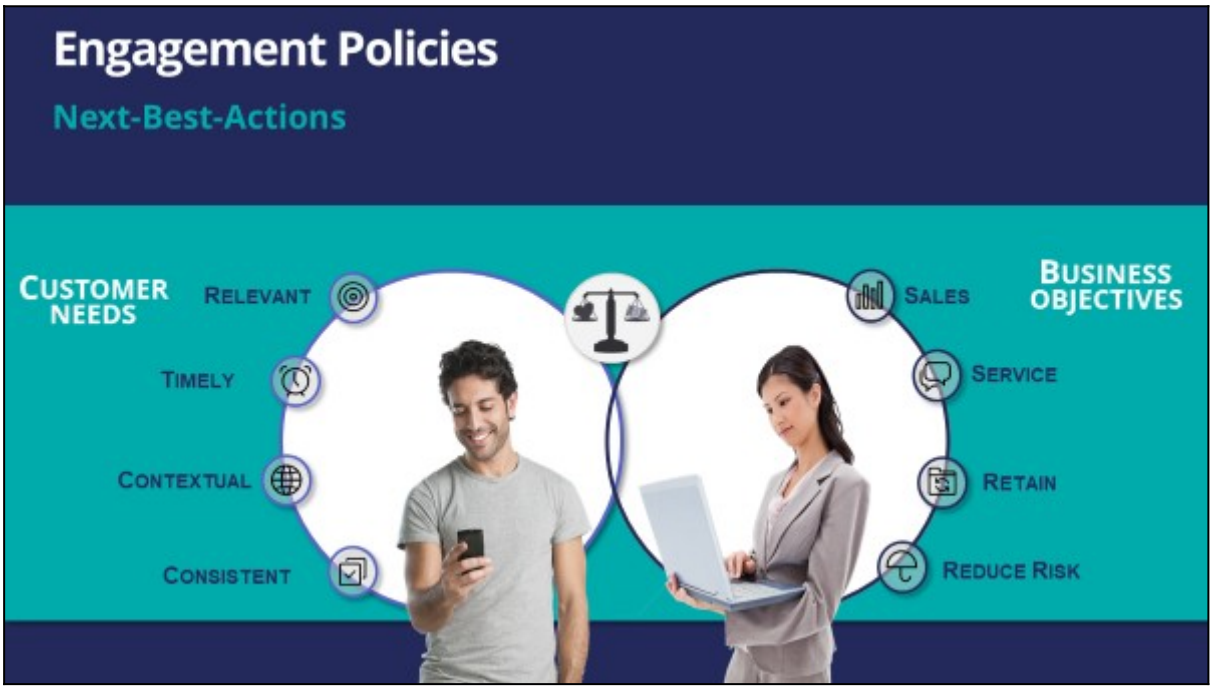

Typically, the business defines a set of rules that make certain actions available to certain customers. This set of rules is called an engagement policy.

As part of an engagement policy, three types of conditions are defined – Eligibility, Applicability, and Suitability.

Consider the following examples: a retail bank is promoting a Gold Credit Card; a telco is offering a new iPhone upgrade with an unlimited data plan; and a communications and media company is promoting a new bundle of HD channels.

Let's see how to define engagement policy conditions that will ensure the bank's Next-Best-Action decisions support these promotions.

In Eligibility, you define strict rules for what is legal, and even possible, to offer customers. For example, to be eligible for the Gold Card offer, customers must be 18 years or older.

Similarly, for the iPhone upgrade offer, customers are eligible for a new contract only if their old contract ends in less than three months.

For the TV channels offer, customers must already own a TV subscription. This offer is not available for customers who have only mobile or landline subscriptions.

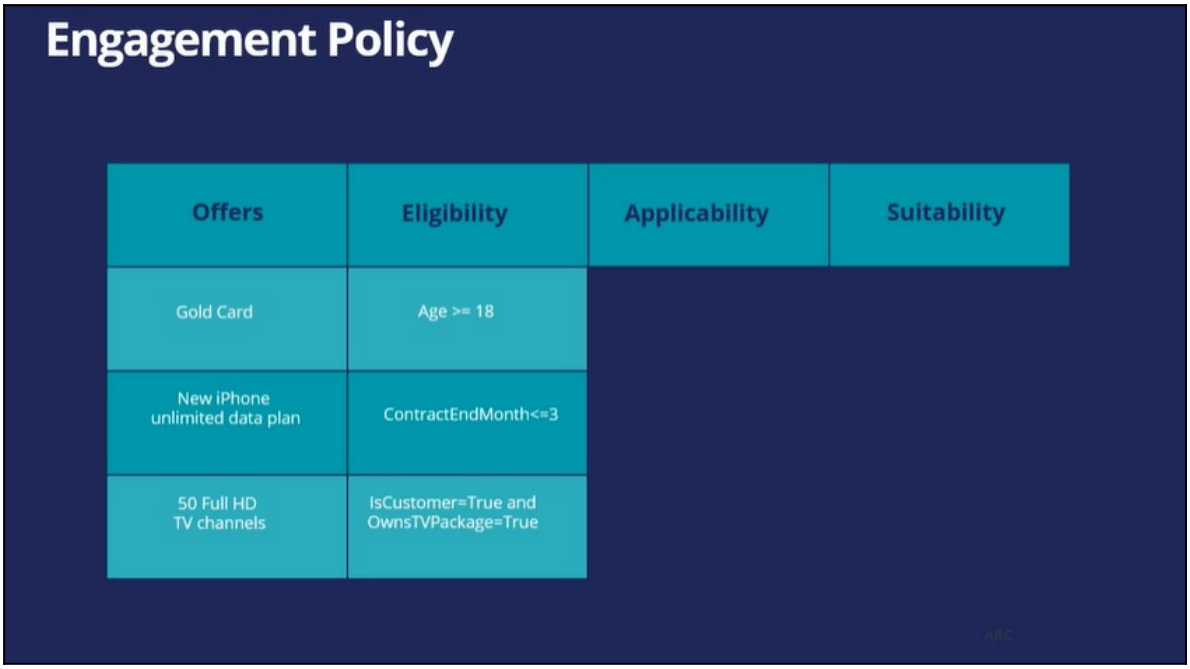

In Applicability, you specify rules for limiting what to offer based on a customer's current situation**,** which is often defined by the products they currently have. These rules are not as rigid as those for Eligibility.

For example, a Gold Card is not applicable if the customer already has a higher value card, such as a Platinum Credit Card. If a customer already has a Platinum Card, they might be eligible for the Gold Card, but the Gold Card is not applicable to them. If they ask for it, they may get it, but the business would prefer not to present them with the Gold Card offer.

Similarly, with the iPhone upgrade offer, if a customer explicitly expressed in the last survey that they weren't interested in an iPhone, this action is not applicable to them. For the TV channels offer, the business does not want to advertise HD channels to a customer who has recently bought a set top box that is not capable of HD.

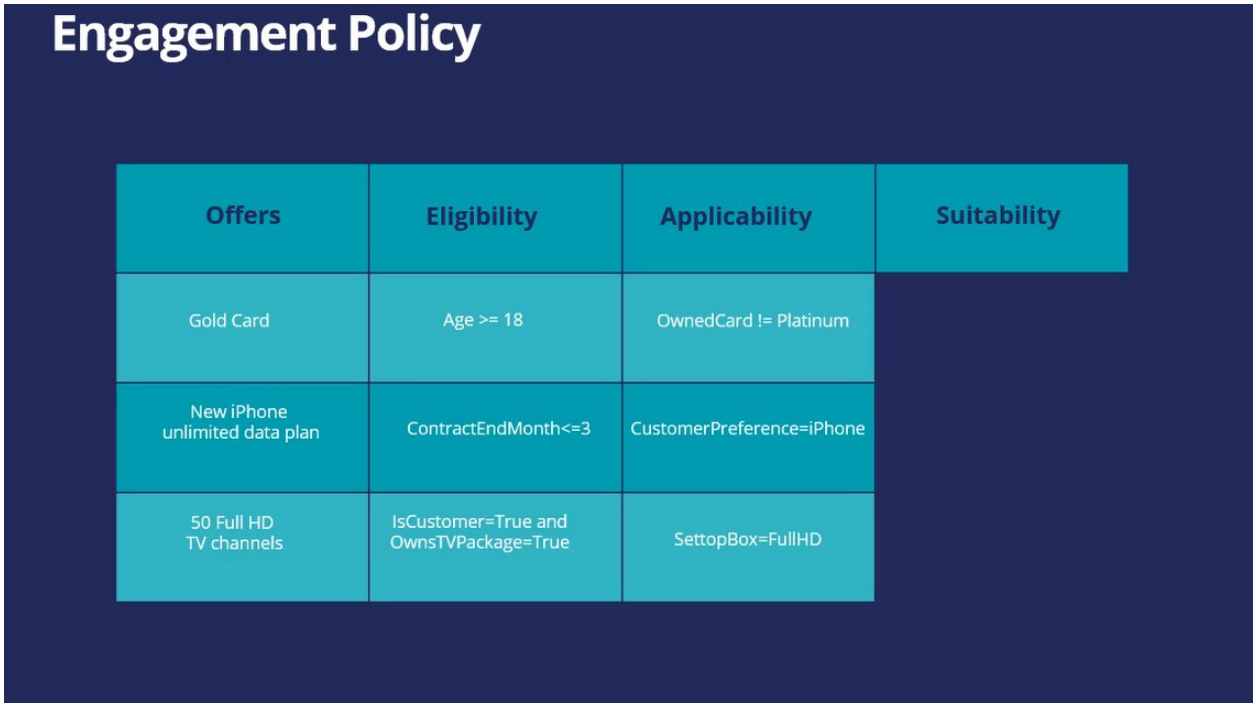

In Suitability, you specify conditions that define an offer as appropriate for a customer. Suitability rules are in place to promote the concept of empathy. That is, to help an enterprise be empathetic toward their customers and refrain from making offers that may not be a good fit.

For example, as the Gold Card is a high value card, it is only suitable for a customer whose debt-to-income ratio is below a certain threshold. Although a customer might be eligible for it, and the offer might be applicable to them, it would be inappropriate to market it to them, as there is a risk of default.

Similarly, an unlimited data plan is not suitable to be offered to a customer with low Internet usage. In the last example, if the customer's favorite TV shows are not available in HD, then it's not empathetic to offer them the new HD channels package.

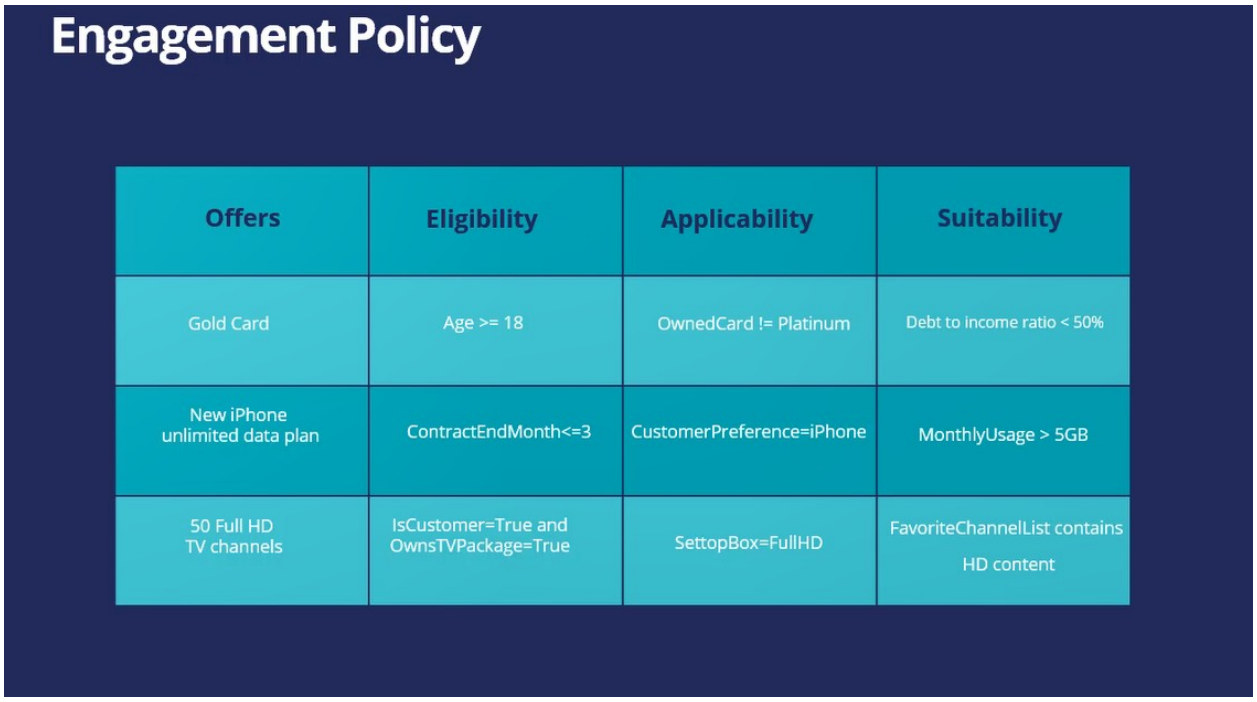

U+, a retail bank, has configured its engagement policy to suit its own business objectives as well as the needs of its customers.

In this scenario, a marketer for U+ has designed 200 actions that could be presented to customers. To select the Next-Best-Actions from these, the Pega Customer Decision Hub first checks eligibility conditions and filters the actions. Then, the applicability conditions are run to filter them further. Next, the suitability conditions are checked to derive the final set of available actions.

These actions will go through one final stage before being presented to customers: the arbitration stage. Arbitration is used to prioritize and choose the Next-Best-Actions based on what is relevant for each customer right now. This is decided by considering factors such as AI-calculated propensity, the action value, and various business levers.

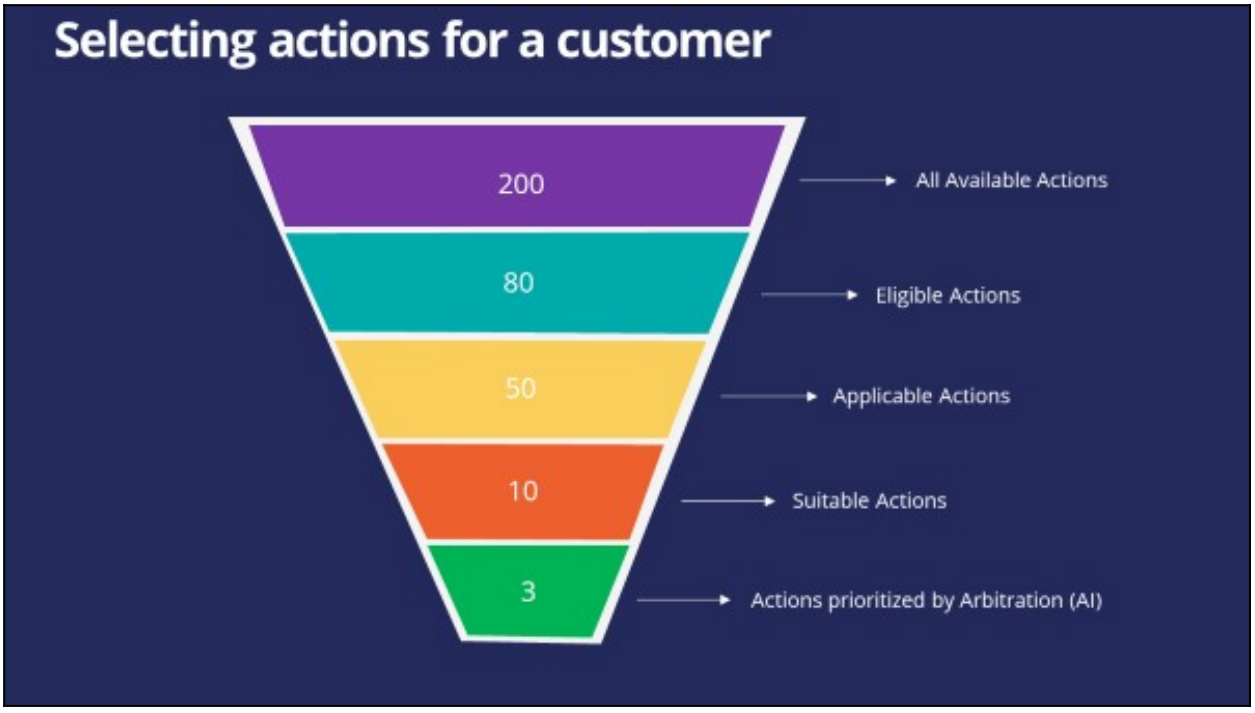
Selecting actions for a customer
200
All Available Actions
80
Eligible Actions
50
Applicable Actions
10
Suitable Actions
3
Actions prioritized by Arbitration (AI)

Defining eligibility, applicability, and suitability rules

Introduction

Engagement policies are defined in Next-Best-Action Designer. Engagement policies specify the conditions under which an action or group of actions is available for a customer. These policies should be defined in the following categories: eligibility, applicability and suitability, which represent the true nature of the associated conditions.

Transcript

This demo will show you how to define engagement policy conditions such as Eligibility, Applicability, and Suitability using Next-Best-Action Designer.

U+, a retail bank, has introduced two new credit cards and would like to offer them to customers based on certain criteria.

First, all credit cards are eligible only for existing U+ customers who are at least 18 years old.

In addition, the business wants to ensure that these two new credit cards will be available for new customers who do not yet have a credit card.

Last, the business understands that not all credit cards are suitable for everyone.

Due to the credit limits of each card, the business wants to offer the Rewards Card to customers with a credit score higher than 500 and the Rewards Plus Card to customers with a credit score higher than 750.

This is the Pega Customer Decision Hub™ portal. You define engagement policies in Next-Best-Action Designer. Engagement policies specify the conditions under which an action or group of actions is available for a customer.

Now, you will set up the engagement policies to reflect U+ bank's requirements. The engagement policies can be defined for a specific group within an issue and/or for individual actions.

Edit to configure the engagement policies. Notice that the two actions are listed under this group.

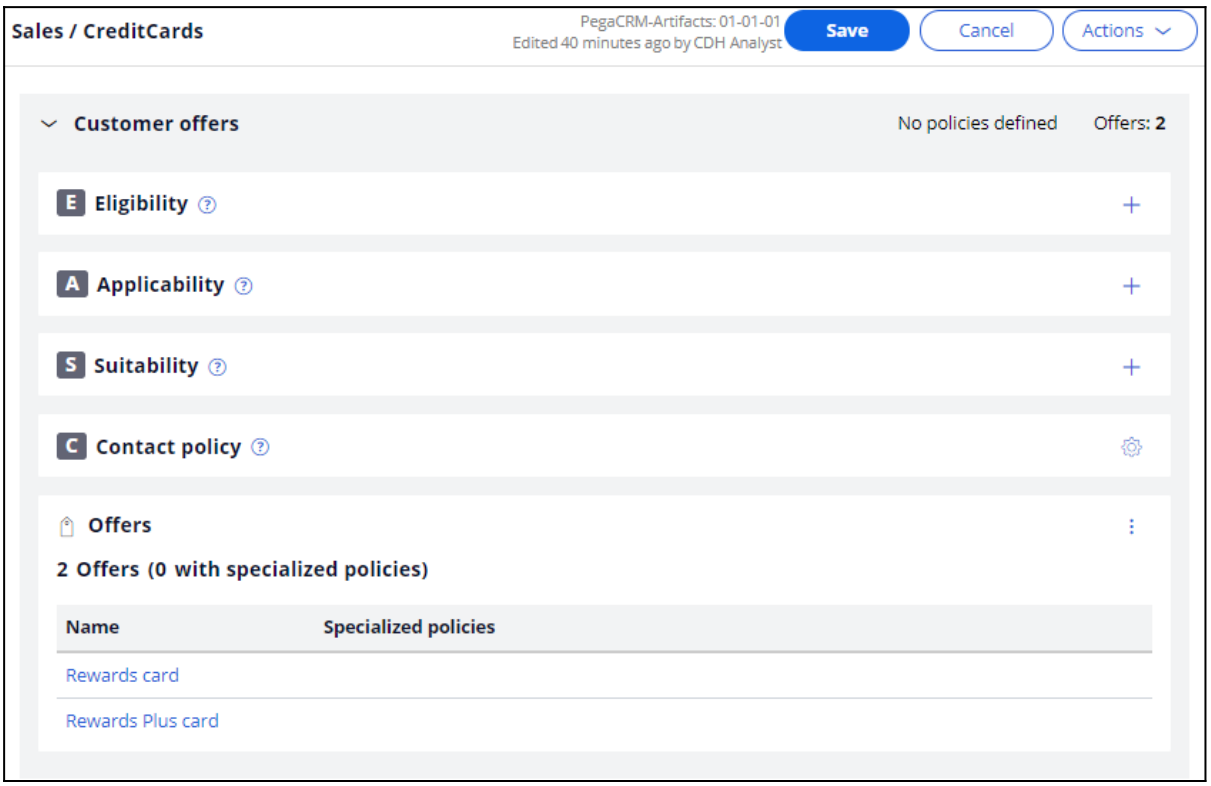

As a best practice, engagement policies should be defined in the following categories: Eligibility, Applicability and Suitability, which represent the true nature of the associated conditions.

First, define the Eligibility condition to ensure that only current customers are considered for the action, and that the customers are at least 18 years old.

Then, define the Applicability condition so that only new customers, who currently do not have a credit card, qualify for the actions. These criteria are also being defined at the group level.

The "new customer" check is done using the **LifeCyclePeriod** property. Add a condition to check if the customer already has a credit card.

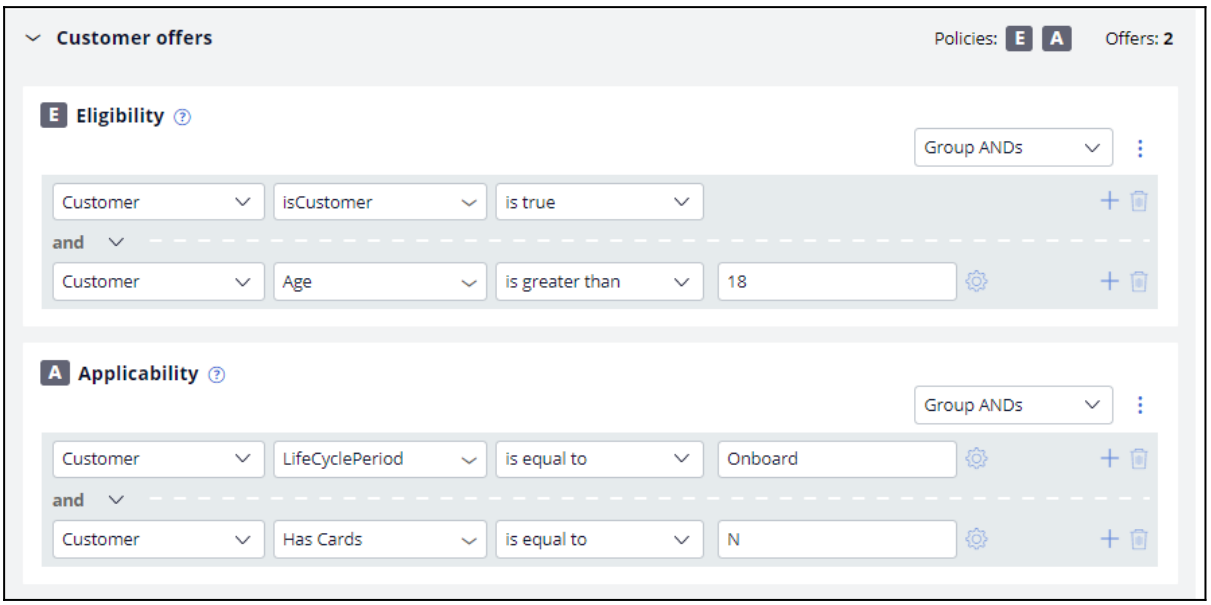

Lastly, business wants a different Suitability condition for individual cards depending on the card's limit. Since the conditions are specific to each credit card, the Suitability condition for each must be defined at the Action level.

Open the Rewards Card Action. This card comes with a certain credit limit. Therefore, U+ believes this card is only suitable for customers with a credit score higher than 500, even though customers may satisfy the Eligibility and Applicability conditions. Thus, define the Suitability condition accordingly.

Next, open the Rewards Plus Card action. This card has an even higher credit limit than the Rewards card. Thus, this card is suitable for customers with credit score higher than 750. To complete the configuration, save the changes.

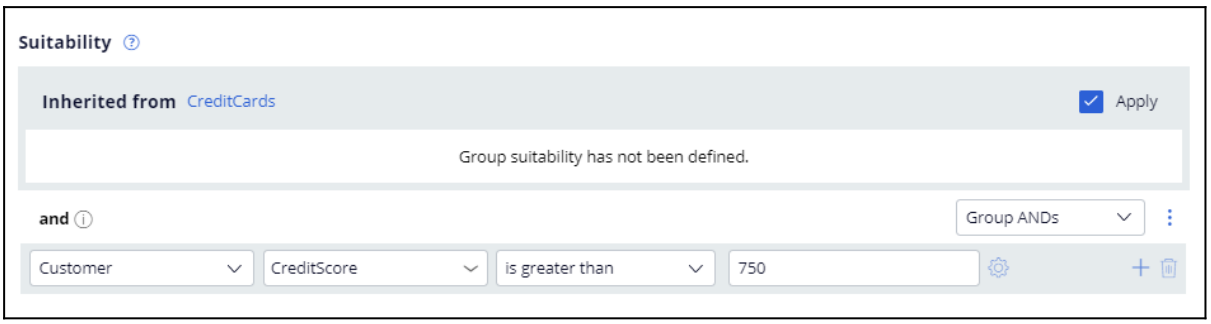

Log in as Troy, a 25-year-old Chef, who became a customer of U+ bank 3 months ago, and his credit score is 600. Notice that the Rewards Card offer is displayed, as Troy satisfies all Eligibility and Applicability conditions. Troy will only be offered this card because his credit score is higher than 500 but lower than 750.

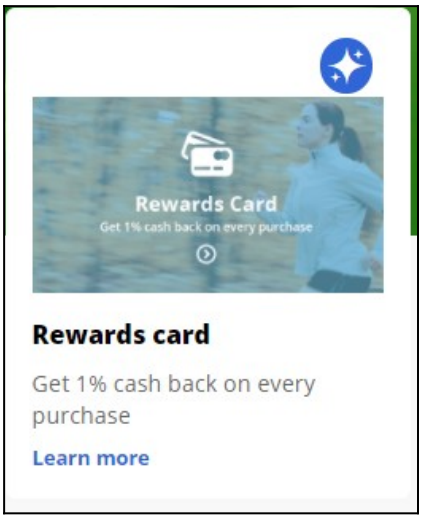

Now, login as Barbara, a 40-year-old Engineer, who became a customer of U+ bank a month ago, and her credit score is 800. Notice that the Rewards Plus Card offer is displayed, as Barbara also satisfies all Eligibility and Applicability conditions. Barbara on the other hand will be offered both cards, as her credit score is higher than 750.

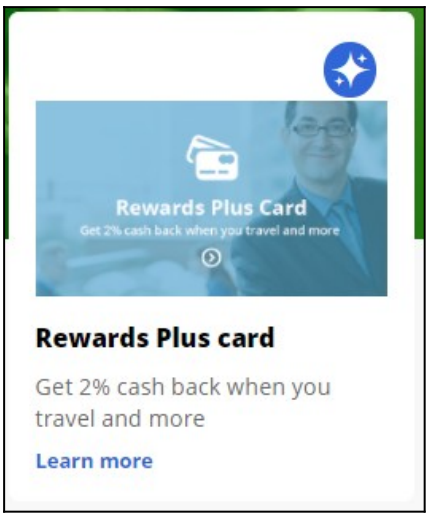

This demo has concluded. What did it show you?

- How to define engagement policy conditions in Next-Best-Action designer at the Action group level.
- How to define engagement policy conditions at the Action level.
- How to differentiate between Eligibility, Applicability and Suitability conditions.

Testing the next-best-action configuration with persona testing

Introduction

You can design persona-based tests to verify that the Next-Best-Action Strategy gives the expected results. You can design these personas according to your requirements and the strategy that you are testing to ensure you do not introduce any regressions as you define your engagement policies and arbitration.

Transcript

This demo will show you how a persona test is conducted on customer personas to evaluate the next-best-action strategy results.

U+, a retail bank, wants to verify that customers are receiving the correct offers. In other words, the bank wants to ensure that the strategy is sending offers to customers in line with their business requirements.

Personas are a representation of various customer profiles that you use to test the results of the next-best-action strategy framework.

Persona-based tests use customer personas with specific characteristics to evaluate next-best-action strategy results.

The bank decides to use Troy, Barbara, and John as personas to test the configurations.

Troy is a 25-year-old chef who became a customer of U+ bank 3 months ago, and his credit score is 600. Barbara is a 40-year-old engineer who became a customer of U+ bank a month ago, and her credit score is 800. John is a 45-year-old IT employee who became a customer of U+ bank 15 days ago, and his credit score is 600.

Persona-based tests are used to verify that the next-best-action strategy gives the expected results. On the Next-Best-Action Designer Engagement Policy tab, you can create test cases for any group.

You can create test cases for a specific group, or for all groups. Running test cases for all groups implies that you are running the next-best-action strategy for all issues and groups in your business structure.

In this case, as the bank wants to target the credit cards group, let's create a persona test case at the group level.

To configure a persona test case, select the action and treatment that you expect the test persona to receive according to your next-best-action strategy.

In this case, Troy is supposed to receive the Standard card and Rewards card actions.

In the Persona (Data transform) field, select the persona against which you want to test the strategy. Select Troy.

In the next-best-action scope section, select whether the test should only check engagement policy configuration, or include additional elements such as constraints and arbitration.

Selecting “Engagement policies only” will validate that your policy conditions are providing the desired results. This will ensure that eligibility, applicability and suitability are tested.

Selecting "Engagement policies and arbitration" will validate the effectiveness of your policies when arbitrating across all actions. The test will consider arbitration, adaptive analytics, treatment and channel processing, and constraints.

In this case, select "Engagement policies only".

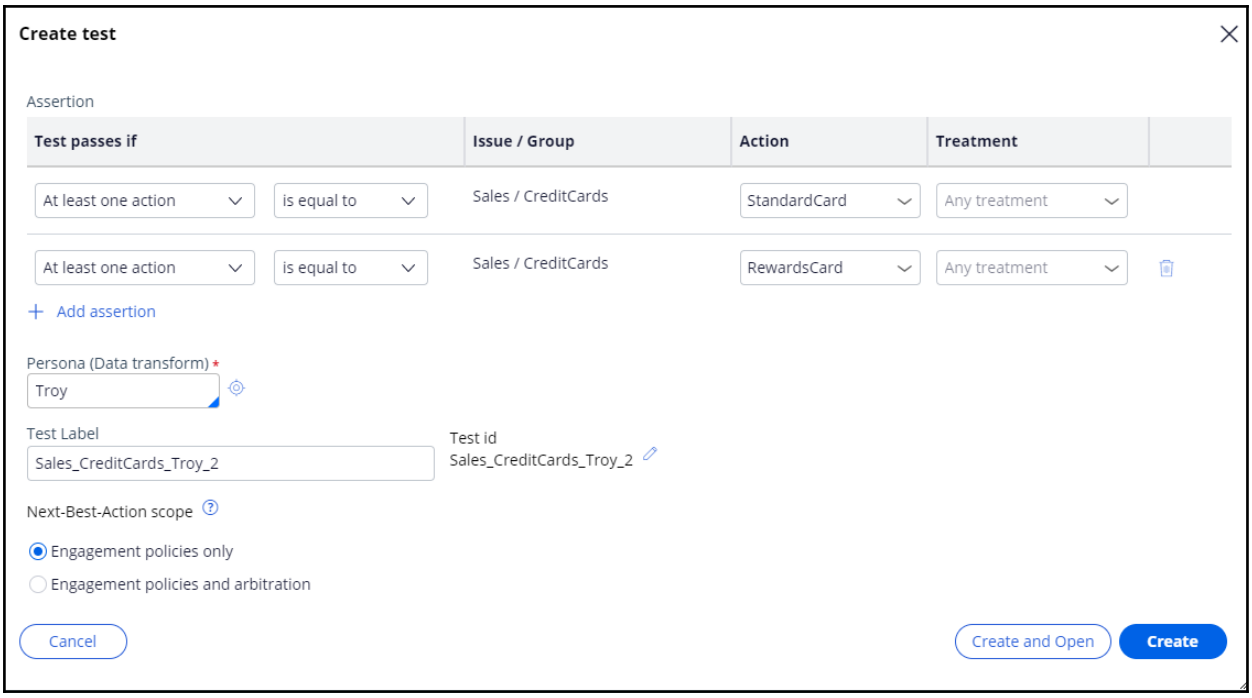

Once configured, select the test and click the **Run selected tests** icon.

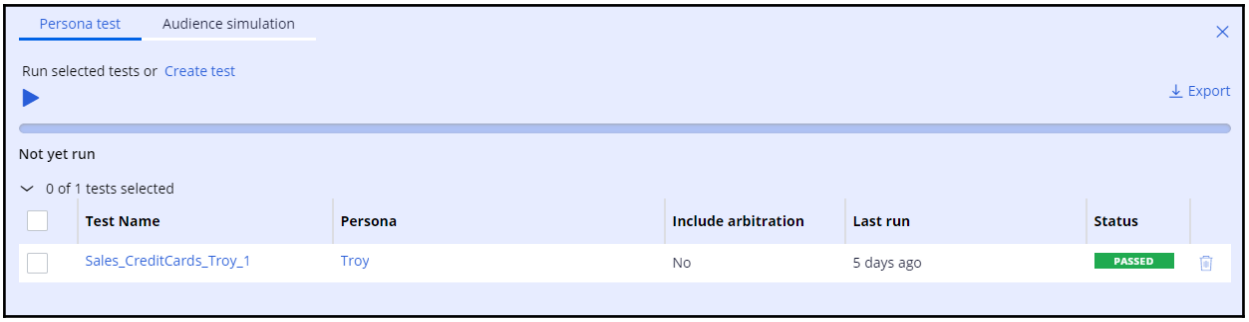

Barbara is a persona who is eligible for the Rewards plus card and the Premier rewards card. Let's create a test case for this persona.

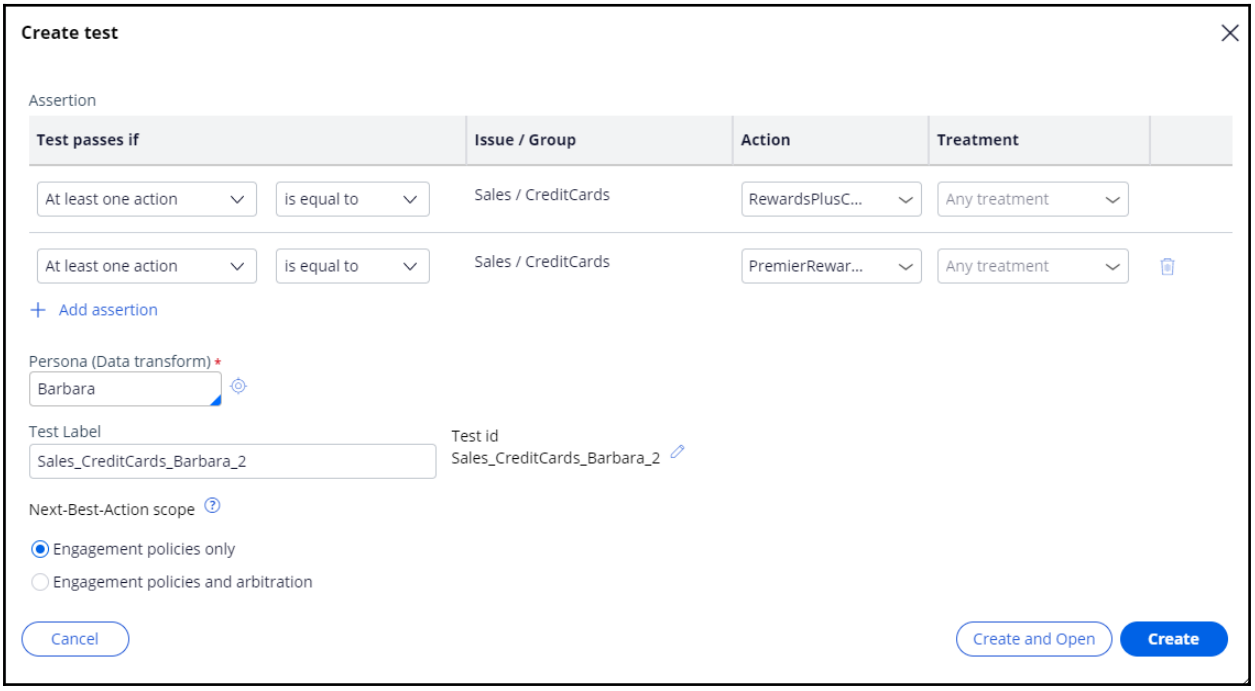

John is a persona who is eligible for the Standard card and the Rewards card. Let's create a test case for this persona.

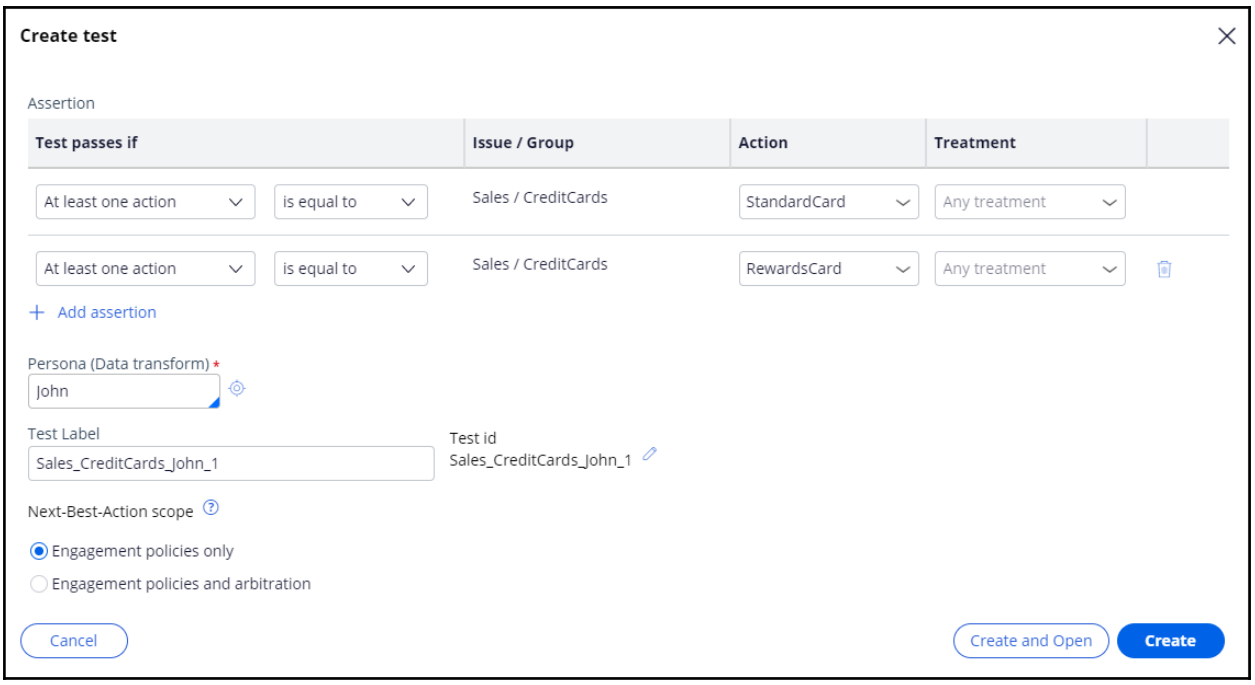

Run the test to confirm he passes the test.

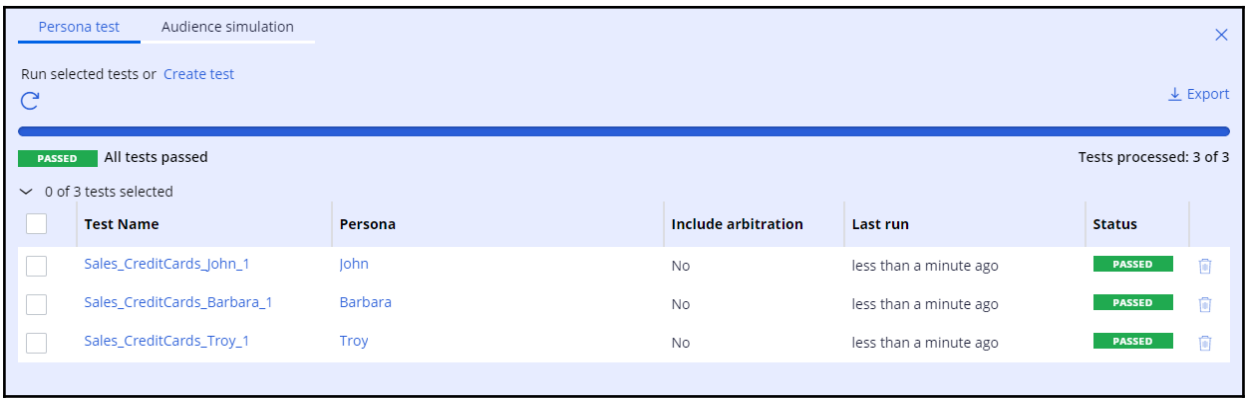

The bank has now decided to offer the Standard card to customers who have an average balance greater than $2000.

So, let's see how this engagement policy change affects these test results. To define an engagement policy, click Standard card.

Define an Applicability condition so that only customers who have a minimum average balance of $2000 qualify for the actions.

Run all the tests to see how this engagement policy configuration affects the test cases.

Observe that Troy fails the test. This means he does not satisfy all the existing engagement policies. Analyze the results to see why the test did not give the expected outcome.

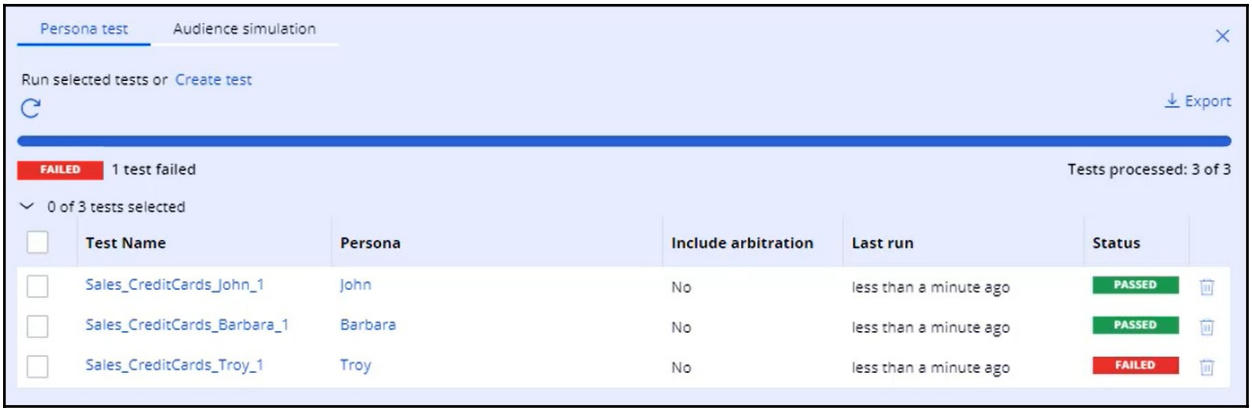

Per the report, he is not eligible for the Standard card and, thus, this test failed.

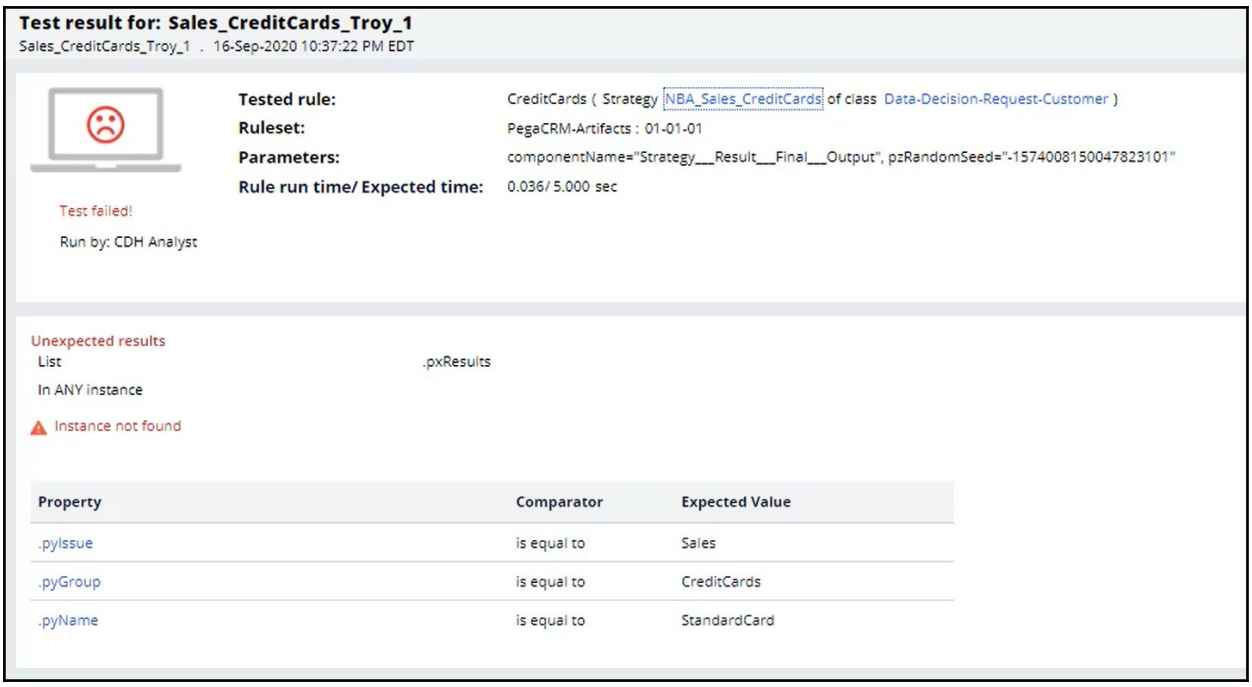

Troy's average balance is 1500, which is less than the defined applicability condition. Whereas, John's average balance is 3000 and Barbara's average balance is 3500, which are more than the defined applicability condition.

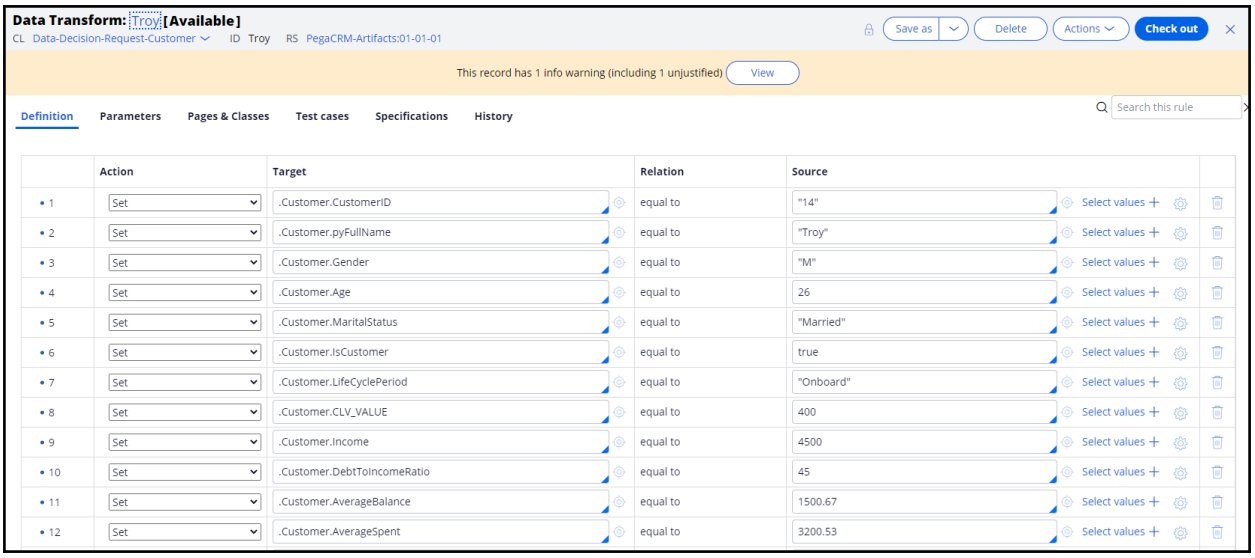

The Troy persona failed the test, but John and Barbara passed it.

Let's edit the test for Troy's persona to delete the Standard card assertion and run the test So, you can design all these personas according to the business requirements and test them every time an engagement policy has changed.

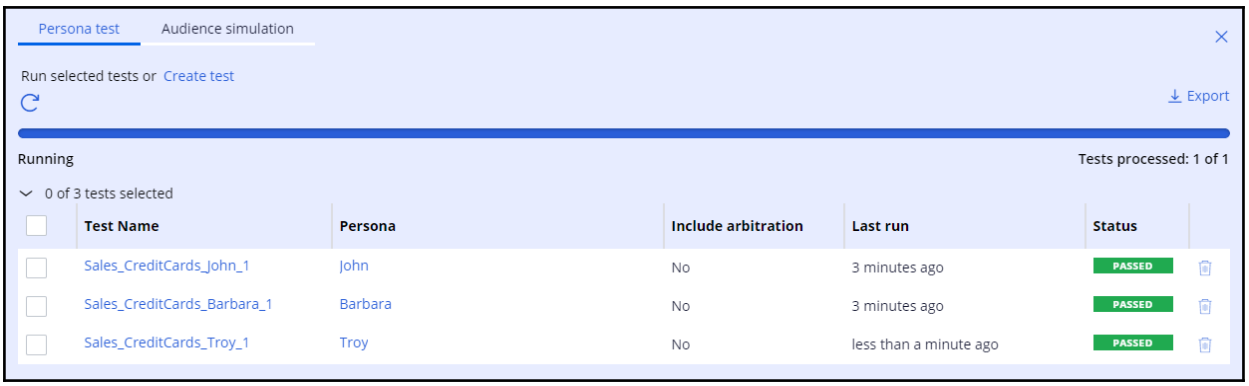

You can also export the results into an excel sheet for further analysis.

This demo has concluded. What did it show you?

- How to create persona test cases.
- How engagement policy changes affect the persona test results.
- How to edit persona test cases.

Avoiding overexposure of actions

Description

Learn about contact policies and how they are used in Pega Customer Decision Hub. Gain experience defining contact policy rules to suppress actions after a specific number of outcomes.

Learning objectives

- Describe contact policy requirements
- Explain how to define contact policy rules to suppress a single action or group of actions for a limited time period

Contact policies

Understanding contact policy requirements

Too many contact attempts over a short period of time can have a negative impact on a customer's attitude toward further actions by your company. To maximize the lifetime value of every customer relationship, organizations must prevent outreach fatigue by optimizing the number of actions taken.

In the Pega Customer Decision Hub, contact policies allow you to suppress actions after a specific number of outcomes.

Suppressing or pausing an action prevents oversaturation by limiting the number of times a customer is exposed to the same action.

Defining contact policies

Contact policies determine when and for how long an Action or group of Actions should be shown to a customer. Contact policies track responses to Actions over a specific period of time, allowing you to implement rules such as the following:

- Do not show an ad to a customer for two weeks if the customer ignores the ad five times in a one-week timeframe.

Note: If outcomes are tracked for an individual action, then the action is not shown once the suppression criteria are met.

- Do not show a group of ads for six months if a customer clicks on any ad in the group 3 times over a period of 30 days.

Note: If outcomes are tracked for all actions in the group, then all of these actions are not shown once the suppression criteria are met.

An Interaction History Summary rule is used to determine the number of impressions and clicks generated by a customer over a period of time. The default time periods are 7 and 30

days. There might be business requirements to track a customer's response to an offer over different time periods, for example, 14 days.

You can add more tracking periods by creating a new Interaction History Summary rule for the required time period and then updating the part of the Next-Best-Action strategy that references it.

Defining action suppression rules

Introduction

Suppression rules determine when and for how long an action or group of actions should not be shown to a customer. These suppression rules pust an action on hold after a specific number of outcomes are recorded for some or all channels.

Transcript

This demo will show you how to suppress a single action or group of actions for a limited time period.

U+, a retail bank, currently displays various credit card offers to each customer who logs in to the website.

For example, every time Troy logs in to his accounts page, a credit card offer is shown. Sometimes the same offer is shown multiple times.

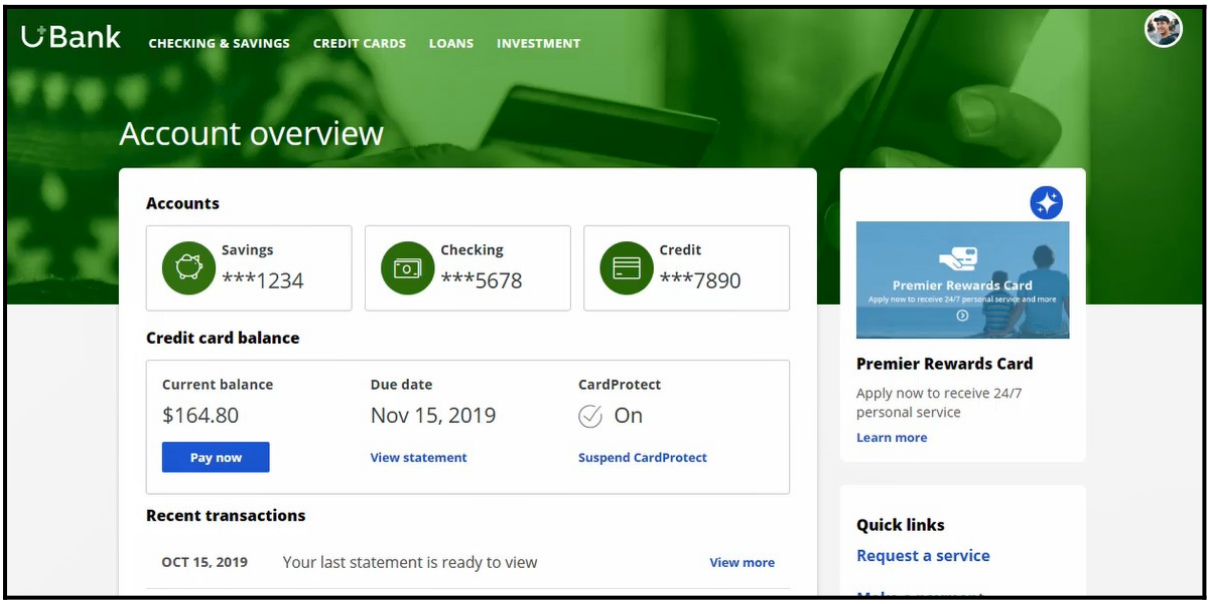

For a limited time period, the bank wants to automatically suppress offers that are shown or clicked too often.

In this scenario, the bank has two requirements. First, do not show a credit card for ten days if the card was shown three times in the last seven days.

Second, do not show any credit cards for ten days if a user has clicked on a credit card five times in the last seven days.

Contact policies are used to implement these business requirements. You create contact policies in Next-Best-Action Designer.

On the Constraints tab, you can define the suppression rules by creating contact policy rules.

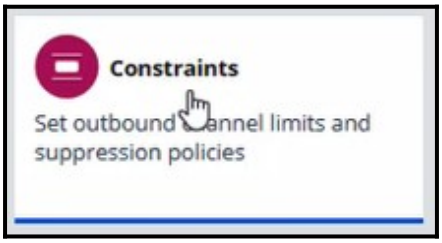

For the first requirement, showing a credit card a maximum of three times, configure a contact policy to track Impressions at the Action level.

Provide a name for the contact policy.

Then select the type of outcome that will be tracked by the contact policy, in this case Impressions.

You can specify whether the responses are tracked for one specific action, or for all actions in the group. Track the first requirement at the Action level, since you want to show one specific card a maximum of three times.

You can select the time period over which the responses should be tracked. In this case, responses should be tracked over a period of seven days.

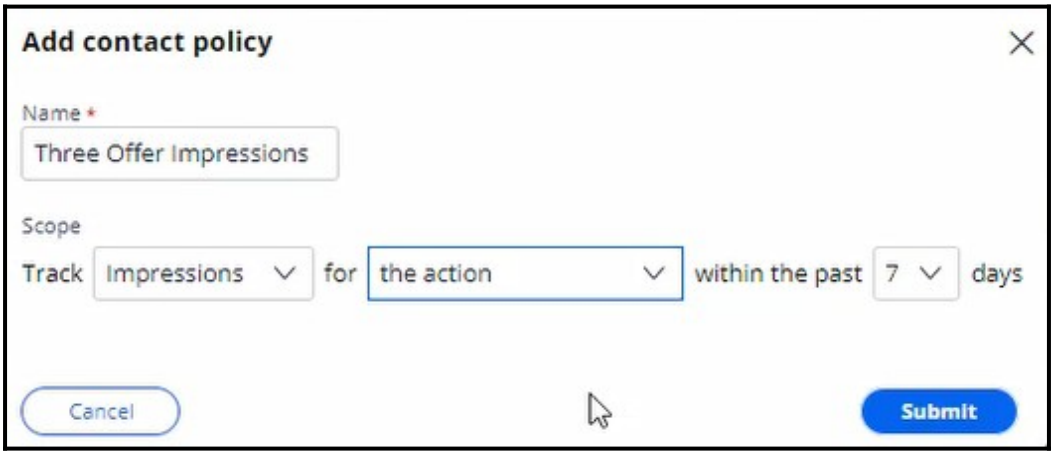

The newly configured contact policy has been added.

The first business requirement is to suppress the action for ten days if there are three impressions for the web treatment, so fill in the details accordingly.

Enter the number of responses required to fulfill the suppression criteria.

Select the channel for which the responses are tracked. Note that if you want to track impressions across multiple channels, you can select **Any**.

Enter the number of days for which an action should be paused after the suppression criteria are met.

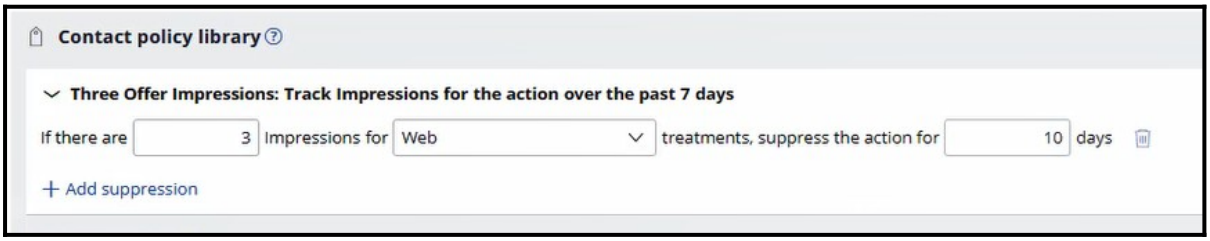

The next business requirement is to suppress the entire group of actions if there are five clicks for web treatments.

The first contact policy is configured to track Impressions, so add another contact policy.

For this requirement you will be tracking Clicks for all actions in the group. This because you want to hide all credit cards if there are five clicks on any one credit card.

Once the contact policy is created, fill in the suppression rule details. If there are five clicks on web treatments, suppress the action for ten days.

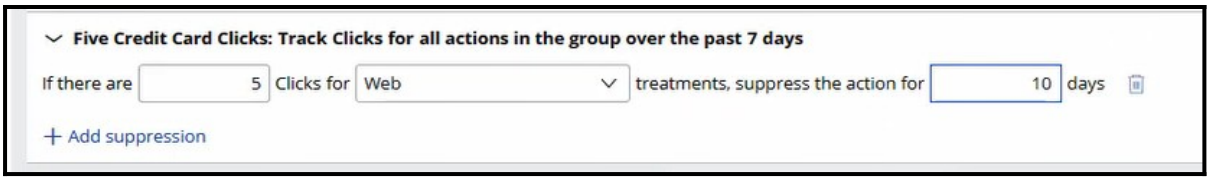

Save the changes.

The contact policy rules are reusable as policy rules across all business issues and groups.

As the bank wants to suppress credit card offers, open the CreditCards group.

Add the contact policy rules you just created.

With that, all the necessary configuration for this scenario is complete. Save the changes.

When customer Troy logs in to his accounts page, the Premier Rewards Card offer is displayed.

After showing him this offer three times, it is automatically suppressed, and a different credit card offer is shown.

When customer Barbara logs in to her accounts page, a credit card offer is displayed.

After clicking on any credit card offer five times, the credit card offers are not shown again.

This demo has concluded. What did it show you?

- How to define contact policy rules to suppress a single action or group of actions.

Arbitrating between actions

Description

After applying engagement policy rules, a customer may still qualify for more than one action. Learn how every Next-Best-Action weighs customer needs against business objectives to optimize decisions based on priorities set by the business.

Learning Objectives

- Describe what action arbitration is and how it works
- Explain how customer needs and business objectives are considered during arbitration
- Prioritize actions based on AI
- Prioritize actions based on AI & business levers

Action arbitration

Introduction

Pega Customer Decision Hub combines analytics, business rules, customer data, and data collected during each customer interaction to create a set of actionable insights that it uses to make intelligent decisions. Arbitration aims to balance customer relevance with business priorities by weighing numerical values for the following factors: propensity, context weighting, action value, and business levers. Learn to create a simple formula for arriving at a prioritization value, which is used to select the top actions.

Transcript

This video explains the concept of action arbitration.

Pega Customer Decision Hub™ combines analytics, business rules, customer data, and data collected during each customer interaction to create a set of actionable insights that it uses to make intelligent decisions. These decisions are known as Next-Best-Action.

Every Next-Best-Action weighs customer needs against business objectives to optimize decisions based on priorities set by the business manager.

U+ Bank, a retail bank, has several actions for its customers and has configured engagement policies to suit both customer needs and business objectives.

In this scenario, a marketer for U+ has designed 200 actions that can be presented to customers. To select the Next-Best-Actions from these, Pega Customer Decision Hub first checks the eligibility conditions and filters the actions. Then, the applicability conditions are run to filter it further. Next, Customer Decision Hub checks the suitability conditions to derive the final set of available actions.

These actions move through one final stage before being presented to customers: the arbitration stage. Arbitration is used to prioritize and choose the best actions based on what is relevant for the customer right now.

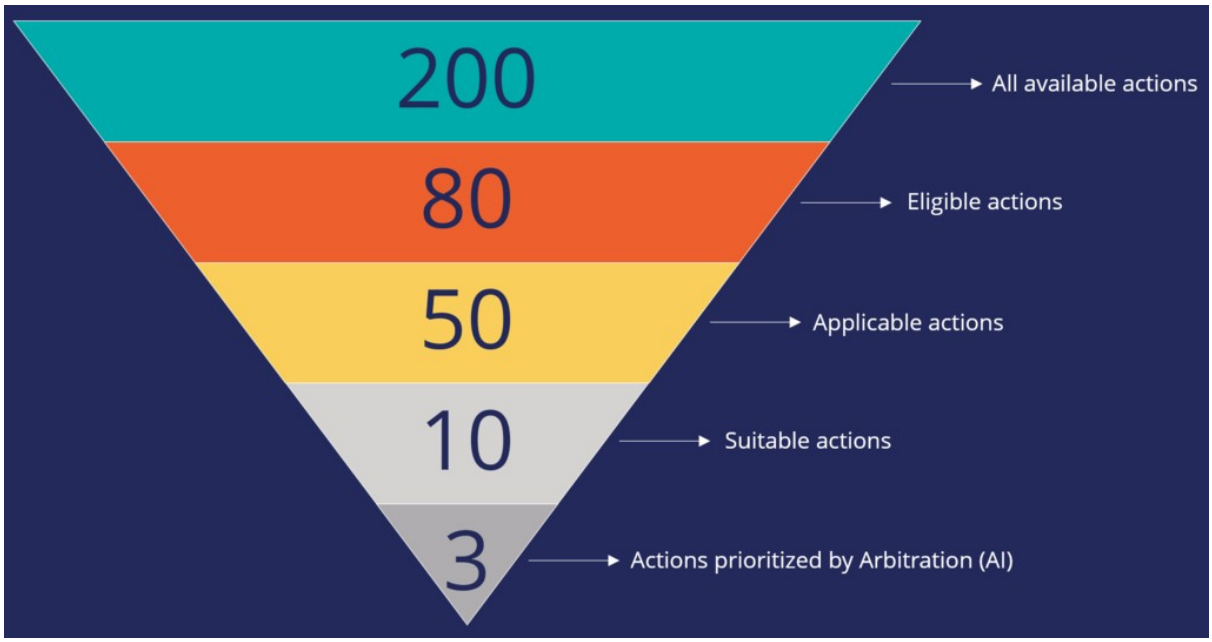

Arbitration aims to balance customer relevance with business priorities. The factors weighed are **Propensity**, **Context Weighting**, **Action Value**, and **Business Levers**, each represented by numerical values. A simple formula is used to arrive at a prioritization value, which is used to select the top actions. The number of top actions selected depends on the channel of interaction. For example, the top three actions, plus two tiles and one hero treatment, can be selected for display on a bank's website.

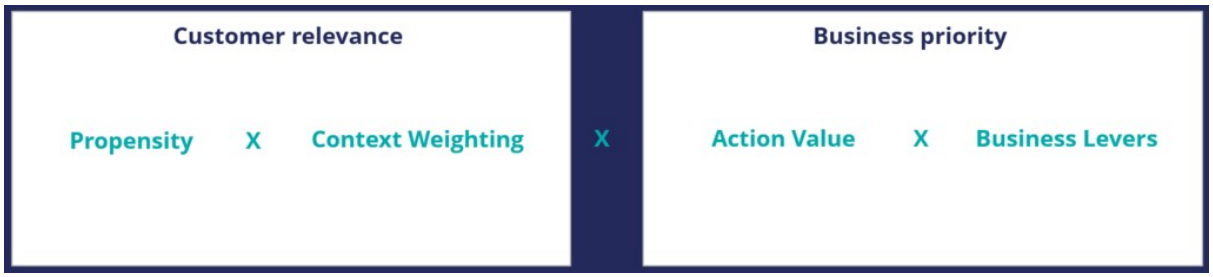

Propensity is the likelihood of a customer responding positively to an action; this is calculated by AI. For example, the higher the likelihood of a customer accepting an offer, the higher the Propensity value for that offer.

Context Weighting allows Pega Customer Decision Hub to consider the situational context for each action. For example, if a customer contacts the bank to close their account, the highest-priority action is to ensure that the customer is retained. The priority of an action is increased by a specified value when the context is detected.

Action Value enables you to assign a financial value to an action and prioritize high-value actions over low-value ones. This value is typically normalized across Issues and Groups. For example, an unlimited data plan is more profitable than a limited data plan. So, in a situation where a customer is eligible for both plans, the unlimited plan has higher priority.

Business Levers allow the business to assert some level of control over the prioritization of actions defined within the system. Levers are used to manually nudge Customer Decision Hub toward Next-Best-Actions based on external factors. For example, the recommended Next- Best-Action might be to offer a credit card to a customer when they visit the home page. But to meet a business goal, the Mortgage Line of Business favors a mortgage offer even if that offer is ranked a little lower on the list of possible actions.

Consider an example where three actions are selected for arbitration. At the moment, only the Propensity is used for prioritization.

Rank	Issues	Groups	Actions	Propensity	Context weighting	Action value	Business levers	Priority
	Sales	Credit cards	Gold Card	0.5	1	1	1	0.5
	Retention	Proactive	10% discount	0.55	1	1	1	0.55
	Service	Administrative	Address change	0.4	1	1	1	0.4

Action arbitration with propensity before prioritization

The result of the arbitration is that the top action is the one with the highest Propensity.

Rank	Issues	Groups	Actions	Propensity	Context weighting	Action value	Business levers	Priority
1	Retention	Proactive	10% discount	0.55	1	1	1	0.55
2	Sales	Credit cards	Gold Card	0.5	1	1	1	0.5
3	Service	Administrative	Address change	0.4	1	1	1	0.4

Action arbitration with propensity after prioritization

Examine what happens when Context Weighting together with Propensity are considered for arbitration. For example, if the intent of a customer calling customer service is to change their address, the Context Weight of a Service action increases.

Rank	Issues	Groups	Actions	Propensity	Context weighting	Action value	Business levers	Priority
1	Retention	Proactive	10% discount	0.55	1	1	1	
2	Sales	Credit cards	Gold Card	0.5	1	1	1	
3	Service	Administrative	Address change	0.4	2	1	1	

Action arbitration with context weight before prioritization

As a result, the Arbitration caters to the current need of the customer and presents a Service action as the top action for the customer. Thus, the Arbitration caters to the current need of the customer and presents a Service action as the top action for the customer.

Rank	Issues	Groups	Actions	Propensity	Context weighting	Action value	Business levers	Priority
1	Service	Administrative	Address change	0.4	2	1	1	0.8
2	Retention	Proactive	10% discount	0.55	1	1	1	0.55
3	Sales	Credit cards	Gold Card	0.5	1	1	1	0.5

Action arbitration with context weight after prioritization

Consider another scenario in which a customer is eligible for two credit cards and two other actions. Now, consider that the Action Value is also used in arbitration when prioritizing. In this case, the Platinum Card is assigned a higher value by the business than the Gold Card.

Rank	Issues	Groups	Actions	Propensity	Context weighting	Action value	Business levers	Priority
1	Sales	Credit cards	Gold Card	0.6	1	1	1	
2	Sales	Credit cards	Platinum Card	0.55	1	2	1	
3	Retention	Proactive	10% discount	0.2	1	1	1	
4	Service	Administrative	Address change	0.1	1	1	1	

Action arbitration with action value before prioritization

Thus, the arbitration selects the Platinum Card as the top action.

Rank	Issues	Groups	Actions	Propensity	Context weighting	Action value	Business levers	Priority
1	Sales	Credit cards	Platinum Card	0.55	1	2	1	1.1
2	Sales	Credit cards	Gold Card	0.6	1	1	1	0.6
3	Retention	Proactive	10% discount	0.2	1	1	1	0.2
4	Service	Administrative	Address change	0.1	1	1	1	0.1

Action arbitration with action value after prioritization

Finally, consider an example in which all four parameters are used for arbitration. In this case, U+ Bank wants to promote two new checking account offers under the Sales issue. The bank sets a higher Business Lever value for the Checking Accounts actions.

Rank	Issues	Groups	Actions	Propensity	Context weighting	Action value	Business levers	Priority
1	Sales	Credit cards	Gold Card	0.6	1	1	1	
2	Sales	Credit cards	Platinum Card	0.55	1	1	1	
3	Sales	Checking Accounts	Premium Checking	0.55	1	1	2	
4	Sales	Checking Accounts	Student Checking	0.5	1	1	2	
5	Retention	Proactive	10% discount	0.2	1	1	1	
6	Service	Administrative	Address change	0.1	1	1	1	

Action arbitration with business levers before prioritization

Although the Propensity of the Checking Accounts actions is low, they are selected as the top actions due to their high Lever values.

Rank	Issues	Groups	Actions	Propensity	Context weighting	Action value	Business levers	Priority
1	Sales	Checking Accounts	Premium Checking	0.55	1	1	2	1.1
2	Sales	Checking Accounts	Student Checking	0.5	1	1	2	1
3	Sales	Credit cards	Gold Card	0.6	1	1	1	0.6
4	Sales	Credit cards	Platinum Card	0.55	1	1	1	0.55
5	Retention	Proactive	10% discount	0.2	1	1	1	0.2
6	Service	Administrative	Address change	0.1	1	1	1	0.1

Action arbitration with business levers after prioritization

Action prioritization with AI

Introduction

Explore how AI-based arbitration works and how AI predicts customer behavior. Arbitration aims to balance customer relevance with business priorities. To select the top actions, a formula is used to arrive at a prioritization value. The formula uses the propensity value, which is calculated using AI. Propensity is the predicted likelihood of positive behavior, such as the likelihood of a customer accepting an offer.

Transcript

This video will explore how AI-based arbitration works and explain how AI predicts customer behavior.

U+, a retail bank, uses the Pega Customer Decision Hub™, to display marketing offers to customers on its website. The bank would like to display more relevant offers to customers based on their behavior.
Troy, a customer, qualifies for two credit card offers. When he logs into the bank's website, he sees the top offer for him, the **Standard Card.**

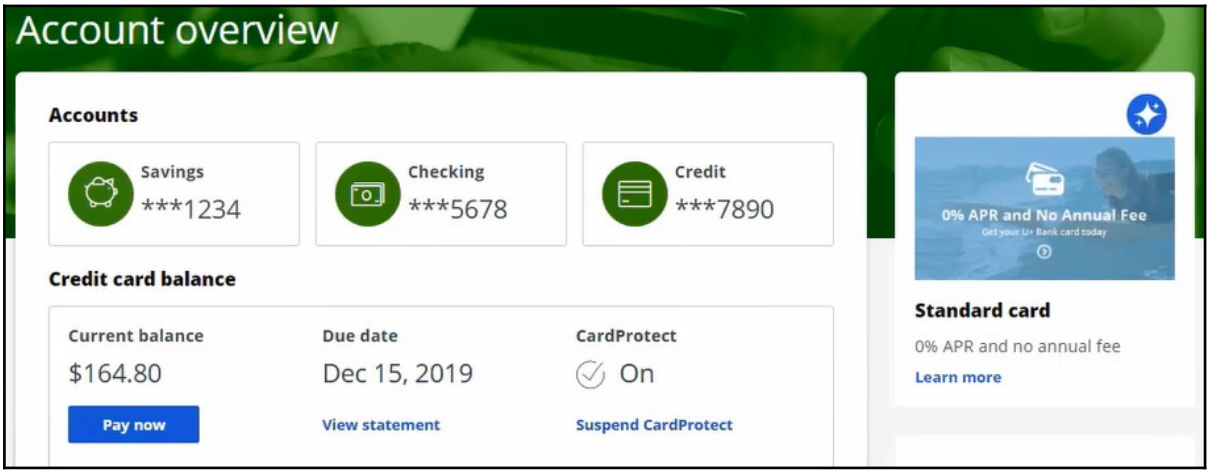

These are the Arbitration settings defined in Pega Customer Decision Hub's Next-Best-Action Designer. Arbitration aims to balance customer relevance with business priorities. To achieve this balance, Propensity (P), Context weighting (C), Action value (V), and Business levers (L) are represented by numerical values and plugged into a simple formula, P * C * V * L. This formula is used to arrive at a prioritization value, which is used to select the top actions.

Notice that only Propensity is currently enabled. Propensity is the predicted likelihood of positive behavior, such as the likelihood of a customer accepting an offer. The value of Propensity is calculated using AI.

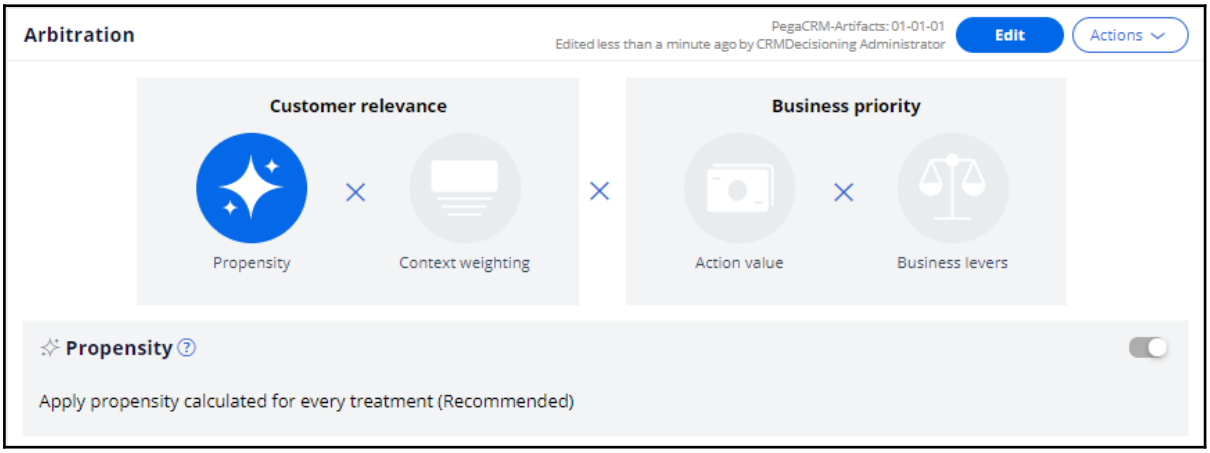

Note the Propensity and Priority values of the **Standard Card**. The Propensity for every action starts at 0.5 or 50%, the same as the flip of a coin. This is because in the beginning, the AI has no past customer behavior on which to base its predictions. Propensity is one of the factors used to arbitrate between relevant offers and select the top offer for a customer.

Though only the Propensity is enabled for arbitration, the value of Priority, which is currently based on Propensity only, does not match the Propensity value. This difference occurs because the Priority calculation does not use the raw model Propensity value directly. Instead, the calculation uses an adjusted Propensity. The adjusted Propensity value helps to equalize the sudden changes in Propensity values that the AI calculates during the initial phase of its learning when it has yet to gather enough customer behavior data to make accurate predictions.

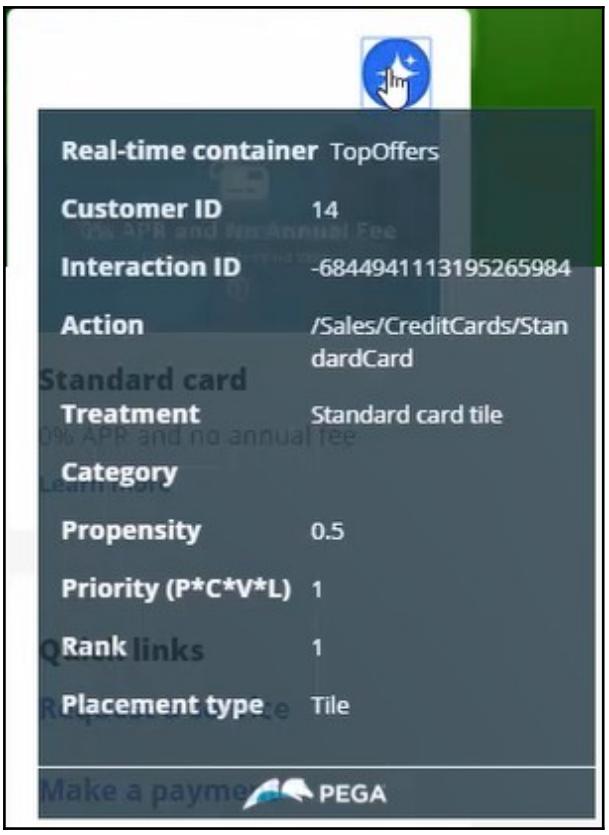

If Troy doesn't click on the current offer this time, a different offer will be shown the next time he visits the website. The next offer Troy is eligible for, the Rewards Card, is then selected for display. If Troy ignores this card as well, by not clicking on it, then the next time he logs in, the Standard Card offer will be displayed again. Why this behavior? First, Troy only qualifies for these two credit card offers. Second, the AI model behind these offers is configured to treat an impression not followed by a click as a negative behavior. In other words, when a customer is presented with an offer but doesn't click on it, the AI records this as a negative behavior. As a result, the Propensity, and therefore the Priority, of the not-clicked-on offer decreases. Notice that the Propensity value of the **Standard Card** offer dropped from 0.5 to 0.25.

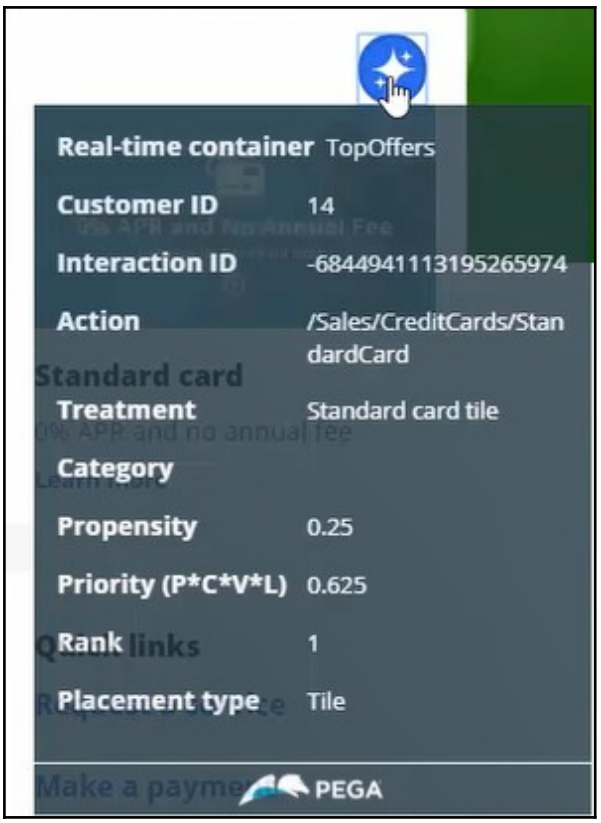

Now, if Troy clicks on the 'Learn more' link for the **Standard Card** offer, a positive response is recorded, and thus the Propensity value of the **Standard Card** increases.

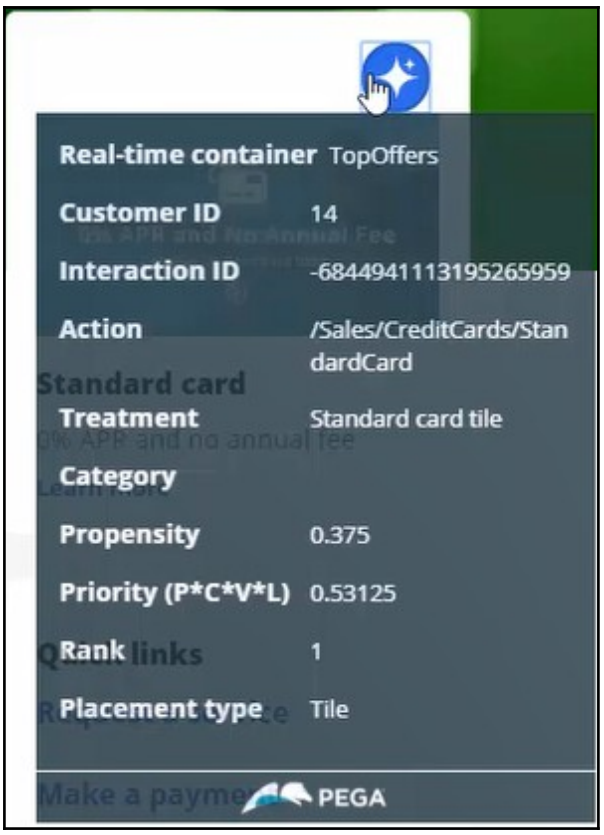

The Customer Decision Hub is configured to calculate the Propensity for each Treatment. To understand how this works, let's examine the AI behind a Treatment. This pop-up window provides a summary of the AI behind this Treatment. In the Pega Customer Decision Hub, the AI that determines the Propensity for positive behavior towards an action or Treatment is called an adaptive model. From here, you can navigate to the adaptive model itself.

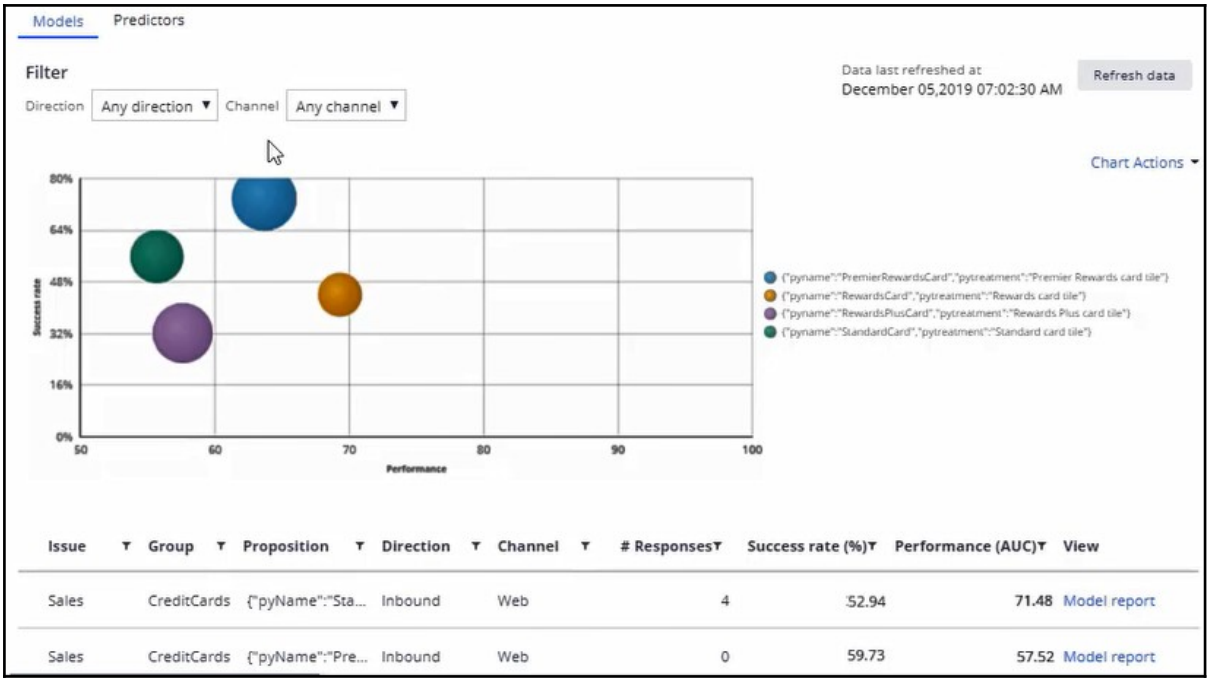

An adaptive model is a self-learning predictive model that uses machine learning to calculate Propensity scores. It automatically determines the factors that help in predicting customer behavior. These predictors can include a customer's demographic details, product and service usage, past interactions with the bank, and even contextual information such as the current channel of interaction.

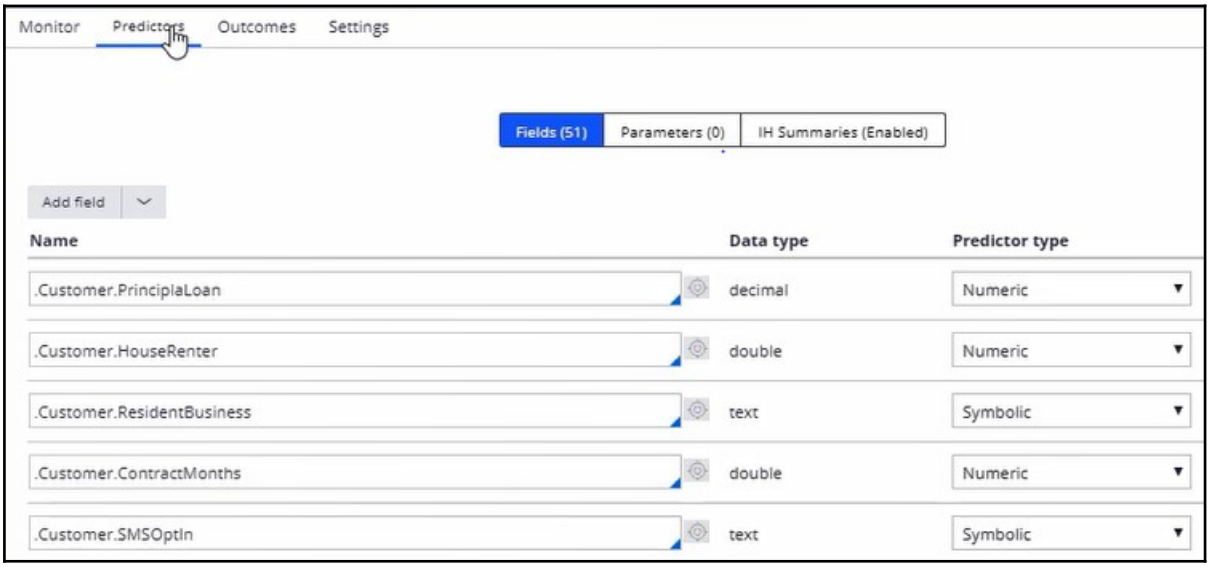

This adaptive model considers an impression not followed by a click, when a marketing offer is displayed on a website, a negative behavior. It considers a Click a positive behavior.

Therefore, when a customer sees an offer message but doesn't click on it, the model records a negative behavior.

The monitoring tab provides an overview of the model's performance. The business can use this information to assess the contribution of the model's predictions with respect to the success of the actions.

The model report provides more insight into the AI model itself. This AI model is automatically generated by the system, and it adapts its prediction algorithm in real-time, based on incoming customer responses. The report shows more information about the predictors, such as how they are grouped and details a data scientist can use to analyze the current health of the model and diagnose any potential problems.

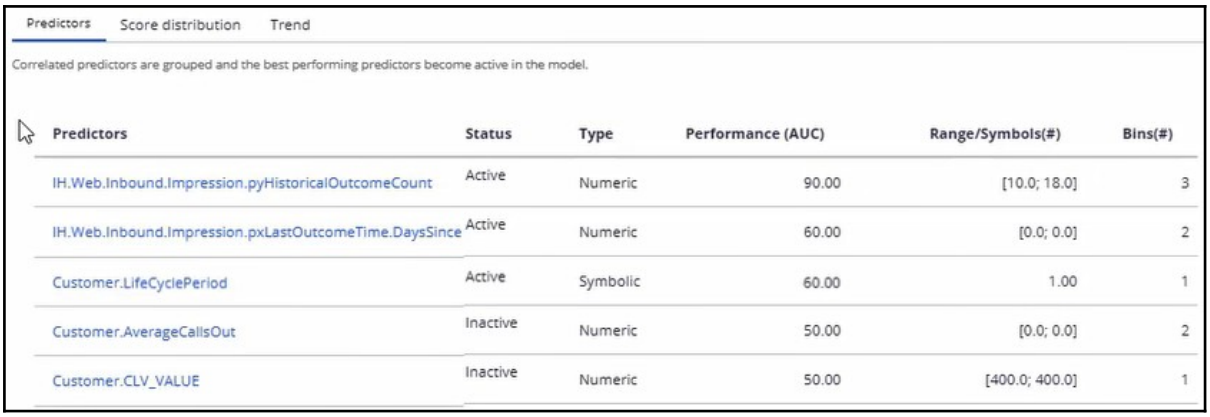

Predictors	Status	Type	Performance (AUC)	Range/Symbols(#)	Bins(#)
IH.Web.Inbound.Impression.pyHistoricalOutcomeCount	Active	Numeric	90.00	[10.0; 18.0]	3
IH.Web.Inbound.Impression.pxLastOutcomeTime.DaysSince	Active	Numeric	60.00	[0.0; 0.0]	2
Customer.LifeCyclePeriod	Active	Symbolic	60.00	1.00	1
Customer.AverageCallsOut	Inactive	Numeric	50.00	[0.0; 0.0]	2
Customer.CLV_VALUE	Inactive	Numeric	50.00	[400.0; 400.0]	1

In the Predictor report, you can examine the performance of individual predictors. Let's examine the LifeCyclePeriod predictor.

This a predictor of type Symbolic. The individual Predictor report shows that a customer whose lifecycle stage is RETAIN is most likely to accept the Standard Card action in the web channel.

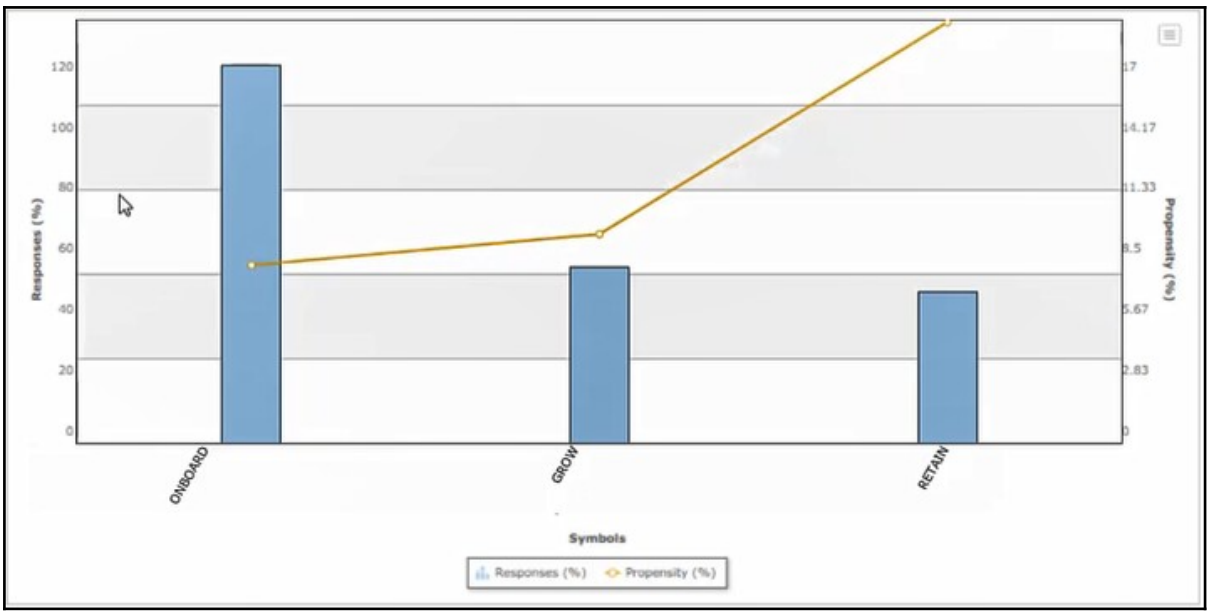

The behavior of one customer can influence the Propensity calculation for other customers with a similar profile. For example, when Robert, a customer with a profile similar to Troy, logs in, he is shown the same offer as Troy. The same AI model is behind the Treatments for both customers, so Robert's action will influence Troy's Propensity score.

This demo has concluded. What did it show you?

- How AI uses customer behavior to calculate Propensity.
- How Propensity smoothing is used to help jump-start the process of AI learning.
- How the Prioritization value is calculated using the (P*C*V*L) formula.
- What adaptive models are and how to monitor them.

Prioritizing actions with business levers

Introduction

Often, due to an internal ad-hoc priority, the business would like to boost the chance of certain actions being selected. To achieve this, they would like to present more relevant offers to customers based not only on their behavior but also on business priorities. Learn how to include business requirements in an action prioritization calculation to boost the chance of an action being selected.

Transcript

This video will show you how to include business requirements in an action prioritization calculation to boost the chances of an action being selected.

U+, a retail bank, noticed that one of the offers, the Rewards card offer, was not presented frequently enough due to its low propensity because customers ignored it during the initial launch.

For example, Troy, a customer, qualifies for two credit card offers – the Standard Card and the Rewards Card. When he logs in to the bank's website, he sees the top offer for him, Standard Card.

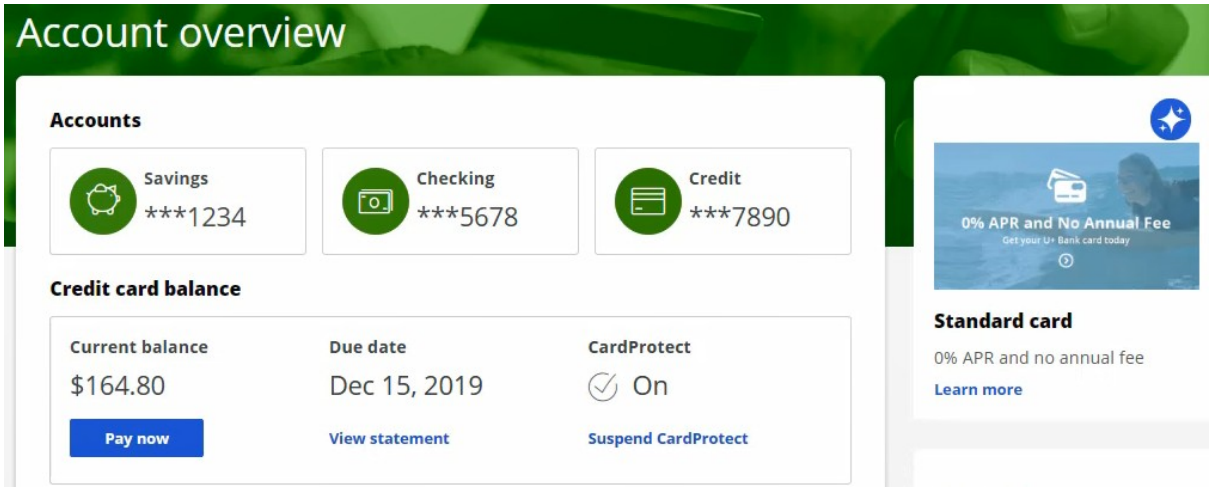

Now, due to an internal ad-hoc priority, the bank wants to boost the chances of the Rewards Card being selected as the top offer. That is, the bank would like to present more relevant offers to customers based on not only their behavior but also on business priorities.

To implement this requirement, you must first enable the Action Weighting, a Business Lever, in Next-Best-Action Designer's Arbitration equation. This ensures that an action's business weight is used in the priority value calculation.

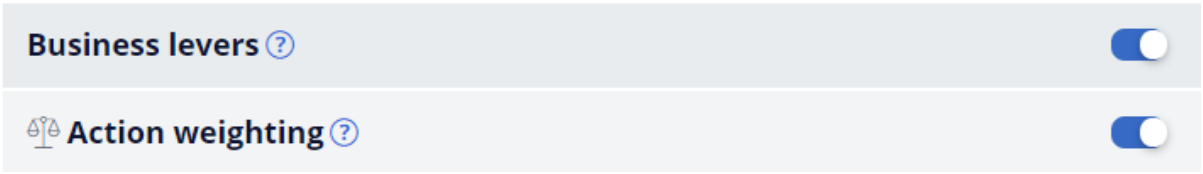

In this case, the bank wants to boost the Rewards Card. So, open the Rewards Card offer. Edit the offer to set a business weight, a value in percentage, that is required to boost the offer. In this case, U+ wants to increase the changes of this action selected by 10%.

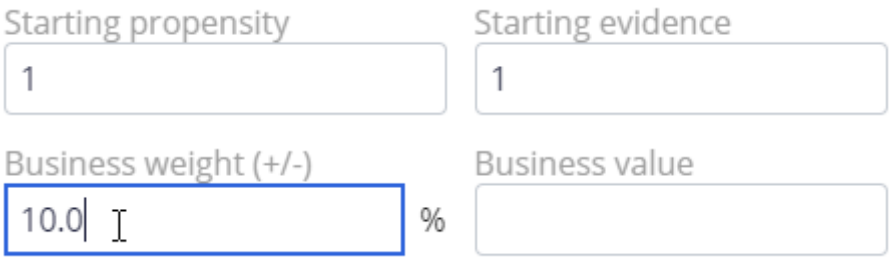

Save the offer for the changes to take effect.

Now, when Troy logs into the website, he will see that Rewards card is the top offer.

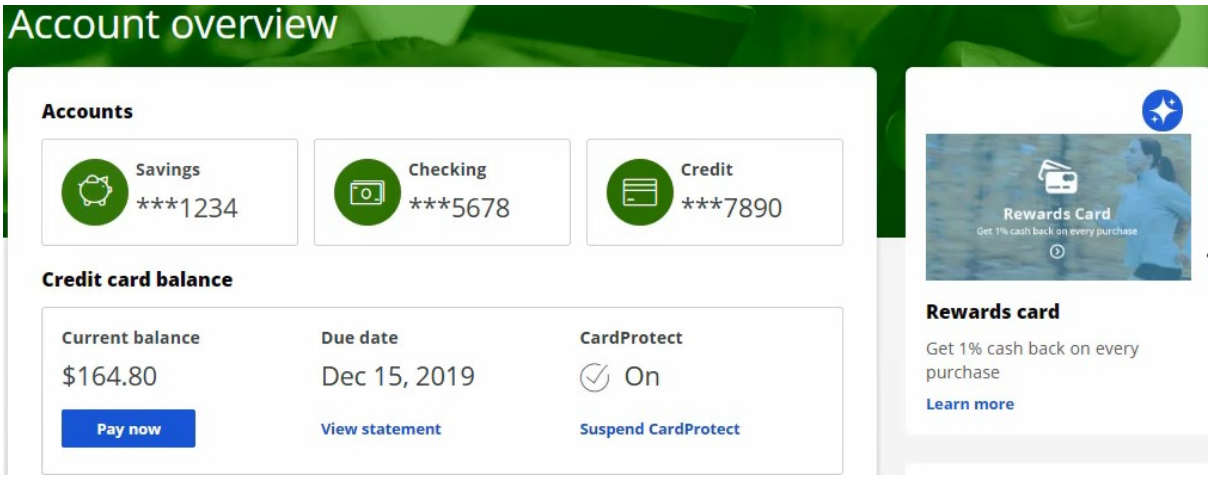

This demo has concluded. What did it show you?

- How to include business requirements in action prioritization calculation.
- How to boost an action using business weight.

Essentials of always on outbound

Description

Learn about the principles of always-on outbound customer engagement and its similarities and differences with traditional outbound marketing. Gain an overview of how the brain works in outbound customer engagement. Study a typical always-on outbound business scenario.

Learning objectives

- Describe the principals of always-on 1:1 customer engagement
- Compare and contrast traditional marketing concepts with always-on outbound
- Describe a typical always-on business use case in the outbound channel

Working of the always-on brain in outbound

What do customers expect from business?

Customers are more empowered than ever before. As a result, they have very high expectations of the experiences they receive from their service providers. These experiences must be meaningful, consistent, and personalized across every channel they interact with.

In alignment with customer expectations, Pega Customer Decision Hub™ is designed to deliver the right message, to the right customer, at the right time, on the right channel.

This approach not only provides a better experience for customers, it helps the business improve its customer relationships over the long term.

The always-on brain

Pega Customer Decision Hub is the “always-on brain” that acts as a single, centralized decision authority.

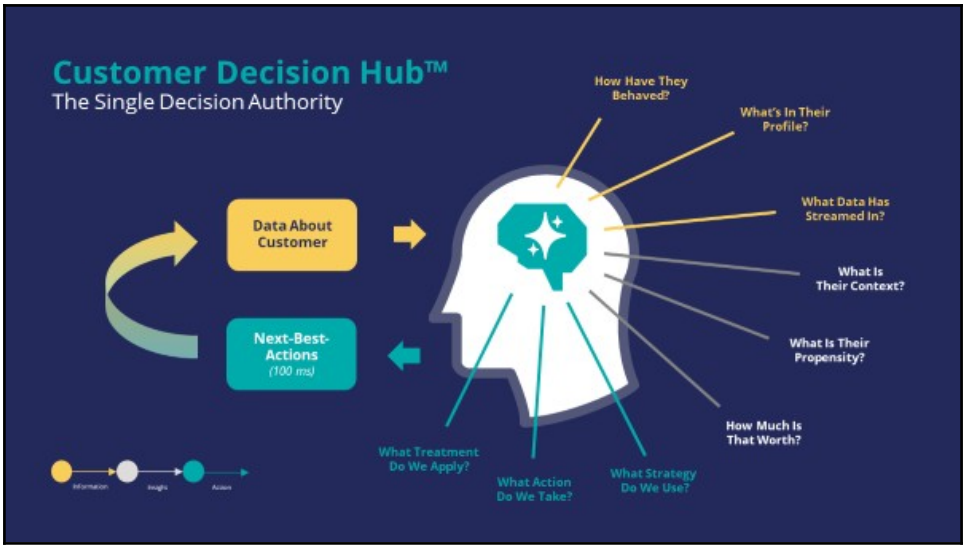

The always-on brain carries out 1:1 customer engagement across channels.

“Always-on” means that the brain constantly monitors each customer’s context across channels. When it detects a need, it selects and delivers the appropriate offer, retention message, or service recommendation in the customer’s preferred channel.

The brain uses data about the customer, including past behavior, profile information, and contextual data that may, for example, stream from their device, as input. It leverages advanced AI techniques to predict customer context, propensities, and relevance. It also uses stored decision strategies to trigger messages and actions at the moment the customer will be most receptive to them. As a centralized resource, the brain enables you to deliver consistent and personalized Next-Best-Actions across all channels.

Customer engagement channels

Customer interactions take place on various channels, which can be broadly classified as “owned” and “paid”.

Owned channels are the channels owned by the business, such as a self-service web portal, contact center, or mobile application. In these channels the business has more control over customer engagement activities.

Paid channels are owned by a third party, for example social media platforms such as Google and Facebook. Businesses have to pay to use these channels for customer engagement activities.

Owned channels can be further classified as inbound and outbound channels.

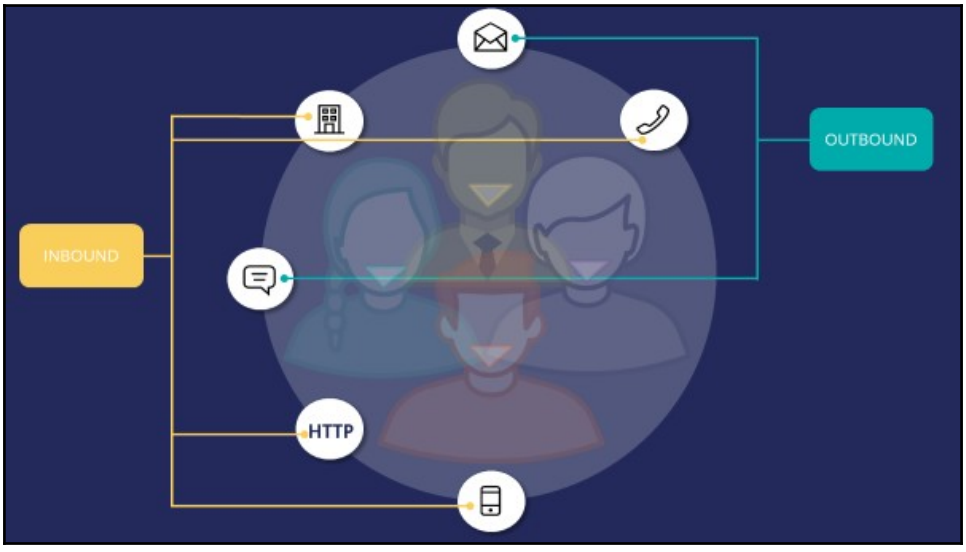

An inbound channel is used when a customer approaches the business, via, for example, a self-service web portal.

An outbound channel is used by the business to send a message to the customer, via, for example, email or SMS.

The flow of information and the mechanics of delivering Next-Best-Actions is different in inbound and outbound channels.

Inbound interactions

When a customer interaction takes place on an inbound channel, the channel identifies the customer, then the brain evaluates the Next-Best-Action for that customer and sends the result back to the same channel, in real-time. The customer's response to the action is recorded by the system, and the cycle repeats.

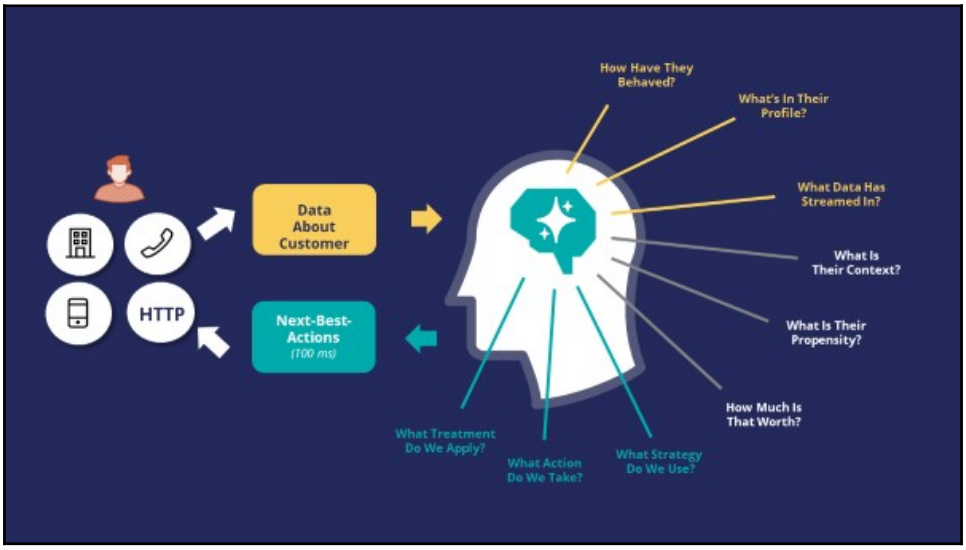

Outbound interactions

Always-on outbound means that the brain sends outbound messages, via email or SMS, to customers only when it's appropriate.

This can happen in three ways:

1) **Scheduled update**: The brain is configured to proactively send out messages on a scheduled basis when relevant. It evaluates the Next-Best-Action for each customer identified in a segment, which contains the potential list of customers to whom you want to send messages. In most cases, customer responses arrive after the message has been sent. When customers respond, the responses are recorded in the system and used in the subsequent Next-Best-Action evaluation.

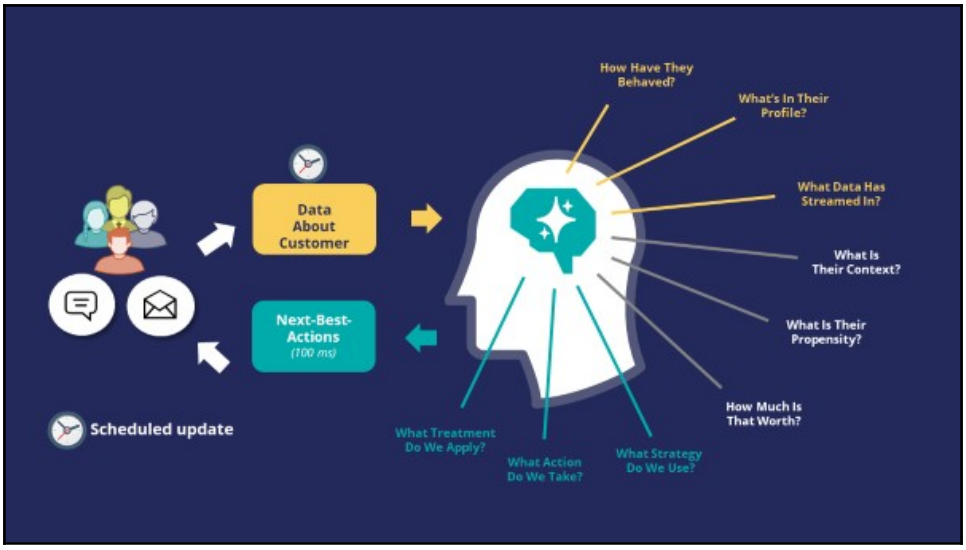

2) **Customer event**: When a customer event is detected that is of significance to the business, for example, a customer abruptly ends an online transaction, data is received via the event stream. The brain then evaluates the Next-Best-Action and sends the resulting action, such as a notification email with a link to continue the incomplete transaction, to the customer. In this case, the outbound message is triggered only for the customer for whom the event occurred.

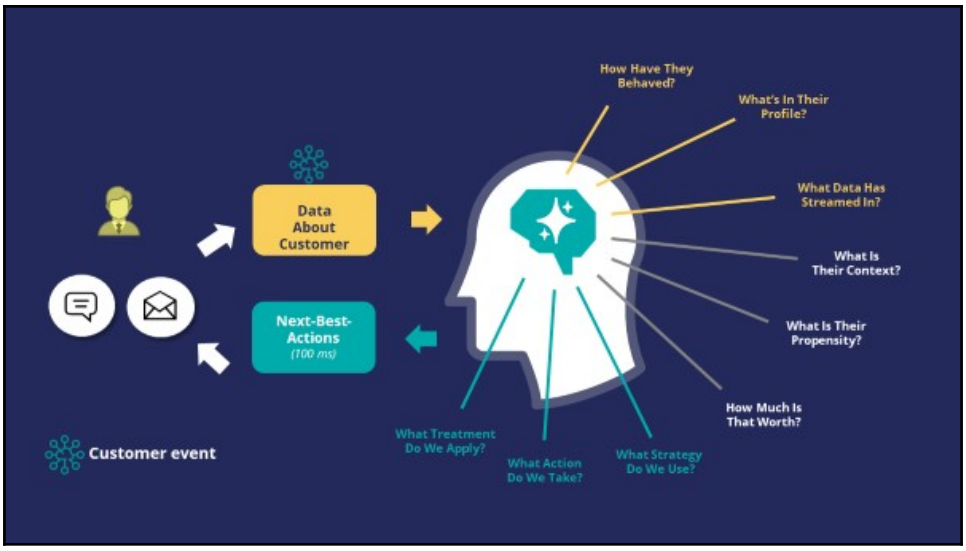

3) **Priority communication**: In extreme circumstances, you may want to contact all or a subset of your customers outside the regularly scheduled update. In such cases, you would configure the brain to send out a high-priority communication to those customers. This would be a one-off communication and won't be repeated automatically by the system.

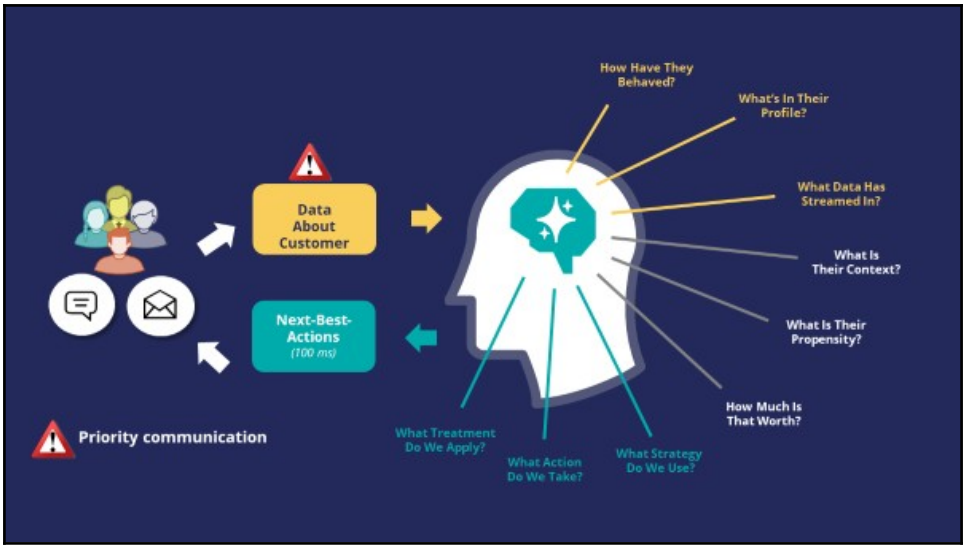
How Have They Behaved?
What's In Their Profile?
What Data Has Streamed In?
What Is Their Context?
What Is Their Propensity?
How Much Is That Worth?
What Strategy Do We Use?
What Action Do We Take?
What Treatment Do We Apply?
Data About Customer
Next-Best-Actions
(100 ms)
Priority communication

Traditional marketing versus always-on outbound

Overview

As more organizations switch from traditional marketing approaches to 1:1 customer engagement, there is a need to compare and contrast the terminologies used in the two approaches.

Here are some key terminologies and concepts used in traditional marketing, translated into concepts in the always-on, Next-Best-Action approach.

Marketing campaigns and programs

In traditional marketing, a marketing program is a set of activities designed to meet marketing objectives by grouping campaigns into logical business outcomes.

In the always-on approach, marketing programs translate into hierarchies of **Issues** and **Groups**. You can use your existing marketing objectives and desired outcomes to define the Issue/Group hierarchies that best meet your needs.

For example, a marketing program could be designed to address the business objective of decreasing dormancy in credit card usage. In this case, you would design a series of marketing campaigns that target infrequent credit card users, encouraging them to use their cards with various offers such as additional bonus points or discounts. In the always-on approach, this translates into defining the Issue as, in this case, *Dormancy*, and the Group as *Promotions*. Under this group, you would then create actions called *10% additional bonus points,* or *5% discount*, and so on.

As a marketing operations person, you receive campaign briefs, which are documents you can use to create **Actions**.

A campaign translates into an action or a set of actions delivered to qualified customers. In the always-on approach, you put an action into or out of play by changing its availability setting. The advantage in the always-on approach is that any adjustment to an action will propagate across all Issues and Groups, which means that you are optimizing across all programs and campaigns at once.

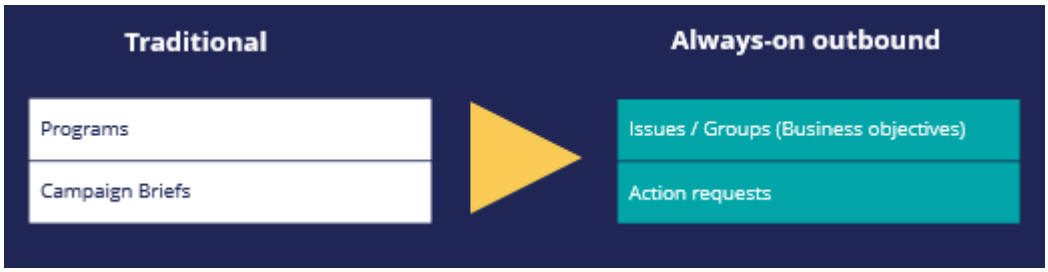

Segmentation

Traditionally, you would use the term segment, audience, or population interchangeably to refer to the target audience of a Campaign, and you would define granular criteria to identify this audience.

In the always-on approach, audience segmentation translates into engagement policies combined with Artificial Intelligence (AI).

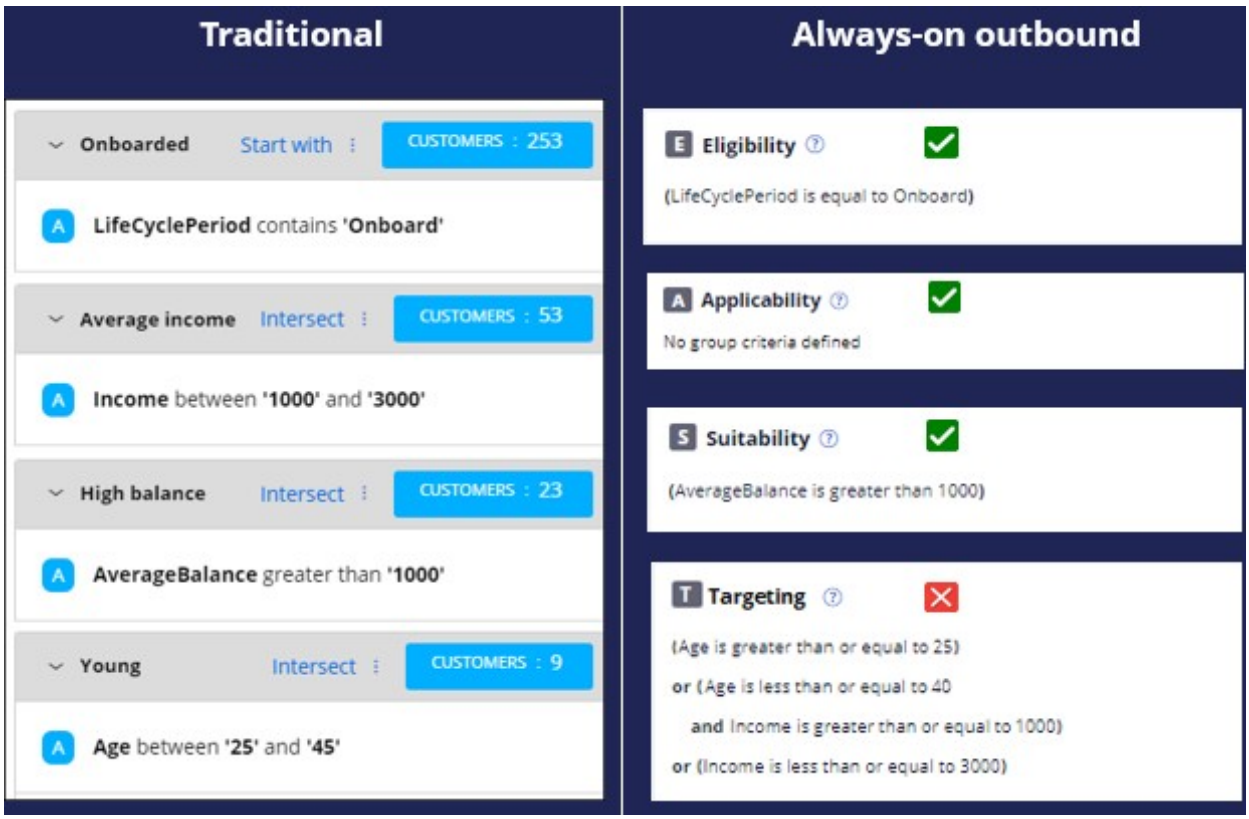

By understanding the purpose of traditional segments, you can translate them into eligibility, applicability, suitability and contact policy rules.

When selecting targets, any criteria that are not based on clear business rules should be avoided. For example, targeting a specific demographic such as an age range or income level based on market research or intuition is not necessary. In the always-on approach, it's

the job of the AI to choose the right action for each customer, once all disqualified actions are filtered out.

In the above example, the lifecycle period criterion is translated into an eligibility condition as it a definite requirement of the business. The Average balance criterion is translated into a suitability condition as the business thinks the group of actions will be appropriate for customers who have such a balance. However, targeting an age range or income-based demographic are not clear business requirements. These conditions in the traditional approach are coming more from intuition and some offline analysis. Therefore, in the new always-on outbound approach, these conditions are not part of engagement conditions but are left to the AI. The AI will take care of finding the right customer segments that are more likely to have a positive response to the configured offers.

Segments, in the new approach, are used only to identify the starting population for an always-on outbound schedule. For example, the starting population could be all customers who have opted in to receive promotional messages.

Segmentation reports

Traditionally, you might have used segments to generate reports on customers who match specific criteria.

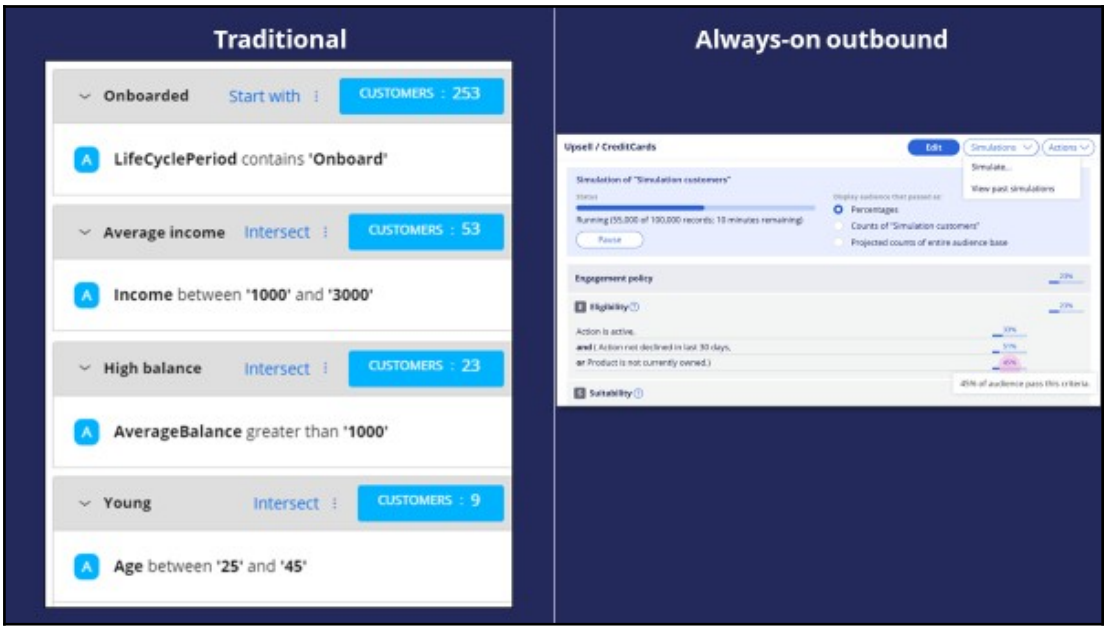

In the always-on approach, you use reports generated by the simulation facility in Next-Best-Action Designer to:

- Understand customer counts.
- Validate the effect of changes to engagement policies and AI controls.

- Visualize the action mix.

Volume caps

When dealing with any form of outbound communication, there is a need to limit the number of communications both on a per-customer basis and in terms of total messages generated by the system over a given time period.

In the always-on approach, you use volume constraints, contact limits and suppression rules to set caps on communications.

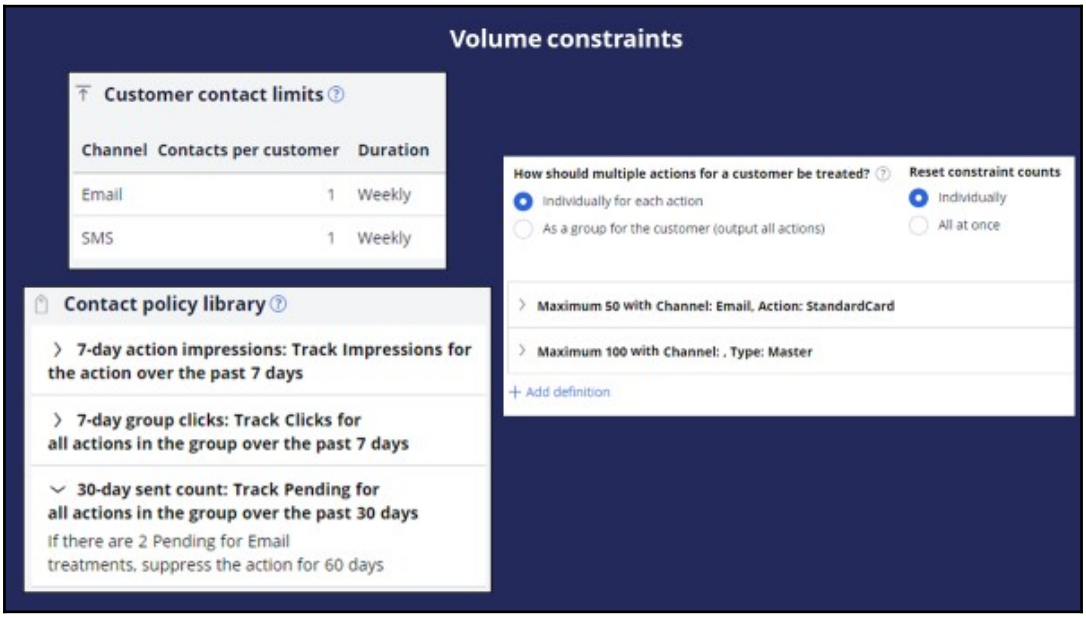

Scheduling

In a traditional marketing approach, there are several options to control the campaign schedule. For example, you can run a campaign as a one-off, or as recurring daily, weekly or monthly.

In the always-on approach:

- The schedule is always-on, it runs at a set frequency, for example on a daily basis.
- Actions are marked as in or out of play using the action 'availability' option.
- Emergency schedules can be set-up as secondary schedules.
- Contact policies help prevent communication fatigue in customers.

Journey guidance

Traditional campaign journeys can be long and complex. They are designed as deterministic journeys in which one campaign triggers the next, which triggers a third, and so on.

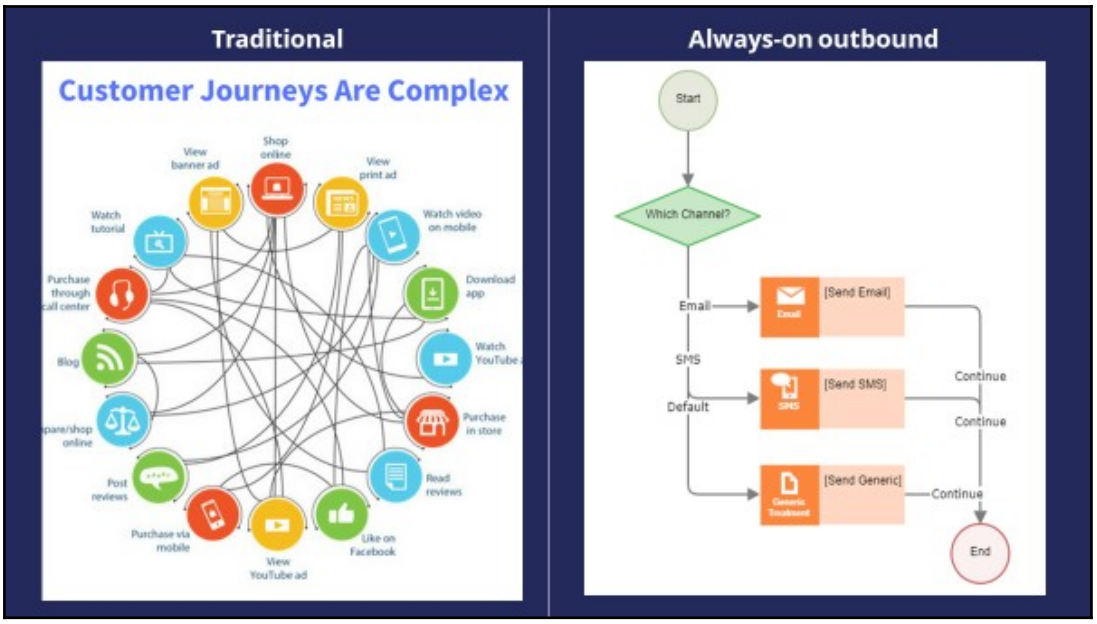

The always-on approach turns managing complex journeys into managing experiences:

- Logic in the channels should be minimized to focus on the overall experience. The number of steps involved in outbound communication, represented by action flow shapes, should be limited to 2 or 3 per channel.
- Next-Best-Action combines many actions to optimize experiences. Let the Pega brain decide the next step.
- Action flows shouldn't contain decision logic (that's in the brain) and should not try to adapt to customer or business changes.
- Use business hierarchies (Issues/Groups) to organize the business outcomes and actions that collectively make up experiences.

Business use case: Always-on outbound customer engagement

Introduction

U+ Bank has recently implemented the cross-sell on the web microjourney with the goal to increase the engagement with their customers on the web channel. The bank uses Pega Customer Decision Hub™ to recommend relevant banner ads to its customers when they visit their personal portal.

Customer Decision Hub combines customer information, information about their past communications with the bank, current contextual information, business rules and the predicted customer behavior, as calculated by Artificial Intelligence (AI) to select the Next-Best-Action, which in this case is the best sales offer for each customer.

Satisfied with the results of the first microjourney, the bank now wants to extend these Next-Best-Action recommendations to be delivered via one of their outbound communication channels, email.

In this microjourney, the main intent of U+ is to increase the customer engagement on email channel. This can be measured by the click-through rate. A click-through metric is recorded when the customer clicks on the ‘APPLY TODAY’ button in the email received.

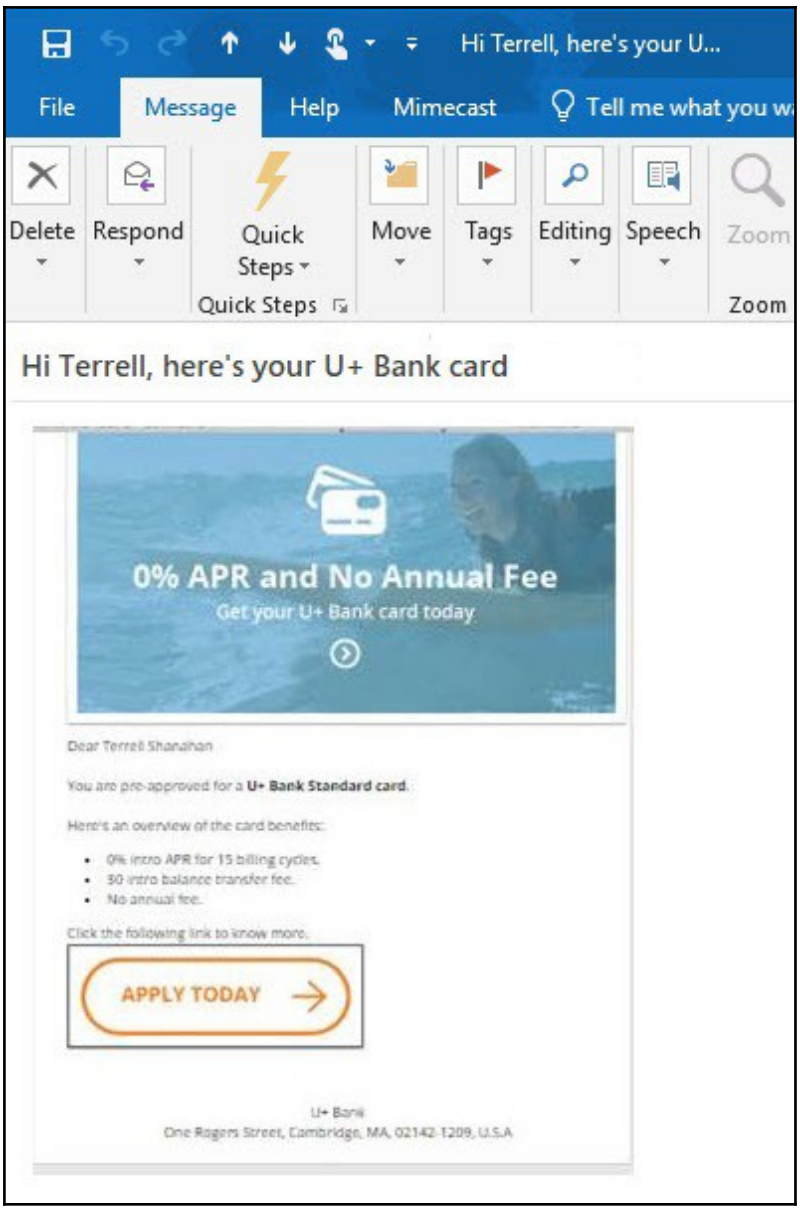

Implementation phases

The marketing operations team plans to implement the requirement in multiple phases.

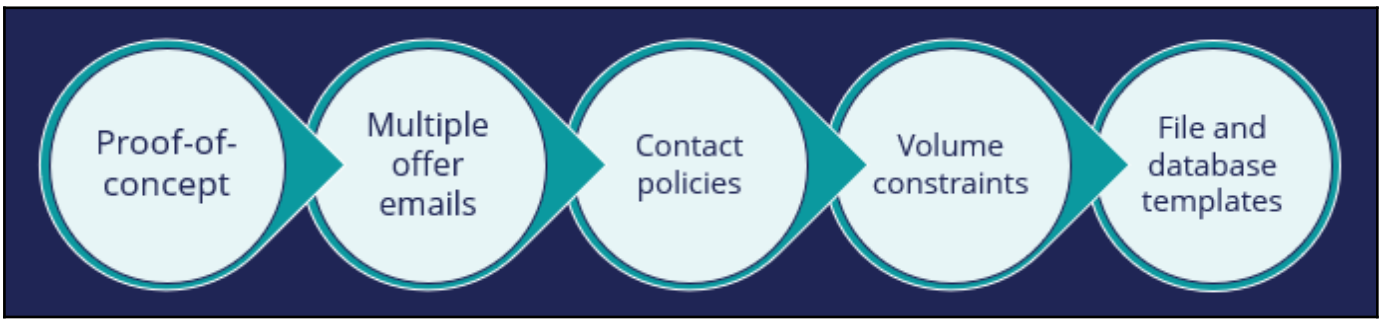

The first phase is a proof-of-concept phase. In this phase, the goal is to send a credit card offer via email to qualified customers. This requires adding an email treatment and configuring the flow to an existing action, enabling the email channel, and configuring the always-on outbound schedule.

As a result of this phase, the Standard card offer details will be sent to qualified customers in an email. For example, if a customer, Stephen John, qualifies for the Standard card offer, the always-on outbound will send him an email with the offer details. He can click on the 'APPLY TODAY' button, it takes him to the credit cards page on the bank's website where he can read more details about the offer and apply for it, if interested.

In the second phase, the bank plans to include more credit card offers to be promoted via email. For example, if a customer is eligible for more than one credit card offer, he will be presented with the top offer he qualifies for. The prioritization uses AI and other factors such as business weights and levers to select the top offer for a customer.

In the third phase, the bank would like to introduce contact policy rules to ensure that a customer does not receive the same offer message twice within a period of time. For example, a customer will not receive an offer email for 60 days if one was sent in the past 30 days.

In the fourth phase, the bank would like to apply some limitations on the outbound volume due to certain operational constraints. The IT department of the bank has imposed a restriction on the number of emails that can be sent in a day. The marketing team has been asked to apply some limitations on the outbound volumes of the email.

The fifth phase is to share action details with a third-party email distributor. To comply with the security and spamming regulations, the bank has decided to use a third-party email delivery service instead of using the inhouse email server. As a result, the action details need to be written to a database table. The action details will then be shared with a third-party email distributor.

In summary, you will learn more about how U+ bank will leverage the always-on outbound capability in this phase of the project.

Sending offer emails

Description

Learn how to send offer emails to qualified customers. Reuse existing actions and customer engagement policies configured in Next-Best-Action Designer and extend them for outbound customer engagement in the email channel.

Learning objectives

- Configure the always-on brain to send offer emails

Sending offer emails

To market your offers, you can send offer emails to qualified customers. Pega Customer Decision™ Hub enables you to re-use existing actions and customer engagement policies configured in Next-Best-Action Designer and extend them for outbound customer engagement in the email channel.

Transcript

This demo shows you how to send offer emails to qualified customers.

U+, a retail bank, wants to use Pega Customer Decision Hub to send emails related to a credit card offer to qualified customers.

The email contains the offer details and a call to action. When a customer clicks the **APPLY TODAY** call to action in the email, they are directed to the credit cards page on the bank's website, where they can learn more about the offer and apply if interested.

This is the Pega Customer Decision Hub portal. Before configuring an action to send offer emails to qualified customers, look at the Next-Best-Action Designer configured by U+.

The **Taxonomy** tab contains the business hierarchy of issues and groups. Every action presented to the customer is classified within an issue/group hierarchy. In this case, the bank wants to do cross-sell/upsell of credit cards. As a result, the bank created one Issue called **Sales** and a Group called **CreditCards**.

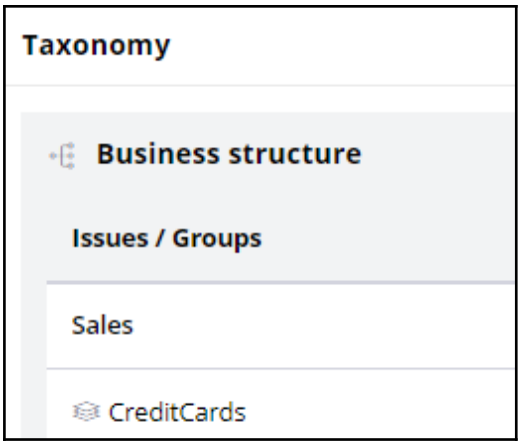

Now, navigate to the **Engagement policy** tab to view the business rules and actions defined for the CreditCards group.

These engagement policy conditions determine which customers qualify for the actions. For example, the eligibility condition **Age is greater than or equal to 18** implies that only customers who are equal to or over the age of 18 are eligible for the credit card offers.

Also, credit card offers are applicable to customers only if they do not currently own a credit card.

Notice that a consultant has already created a few actions under the Sales/CreditCards hierarchy.

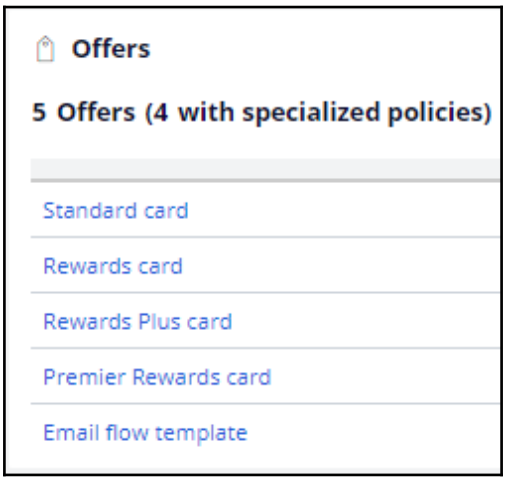

These actions correlate with the credit card offers that the bank wants to promote.

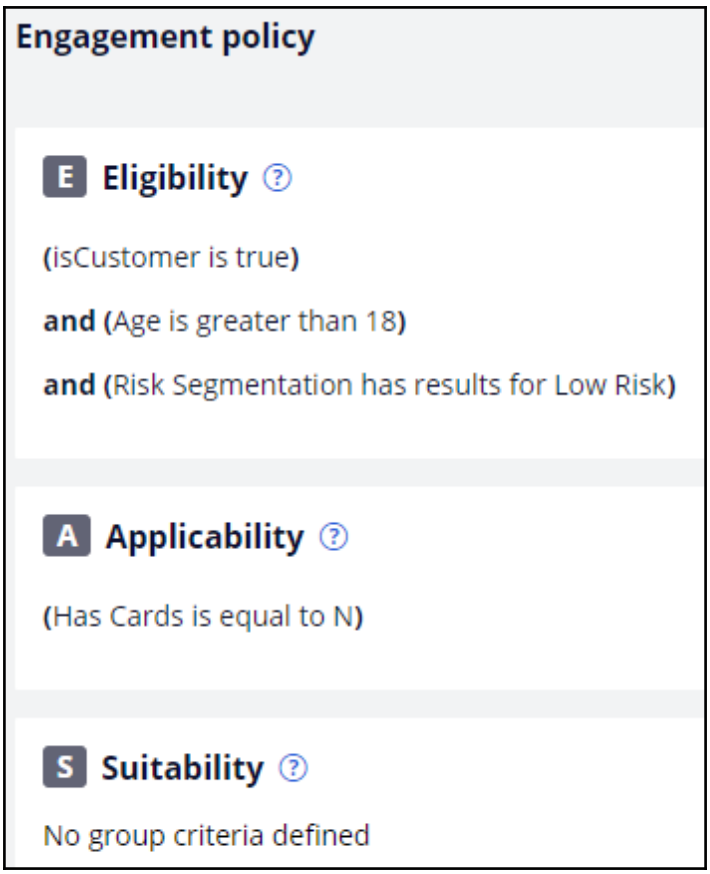

Currently, the bank promotes credit cards on the web channel only. The bank wants to extend this promotion to the email channel.

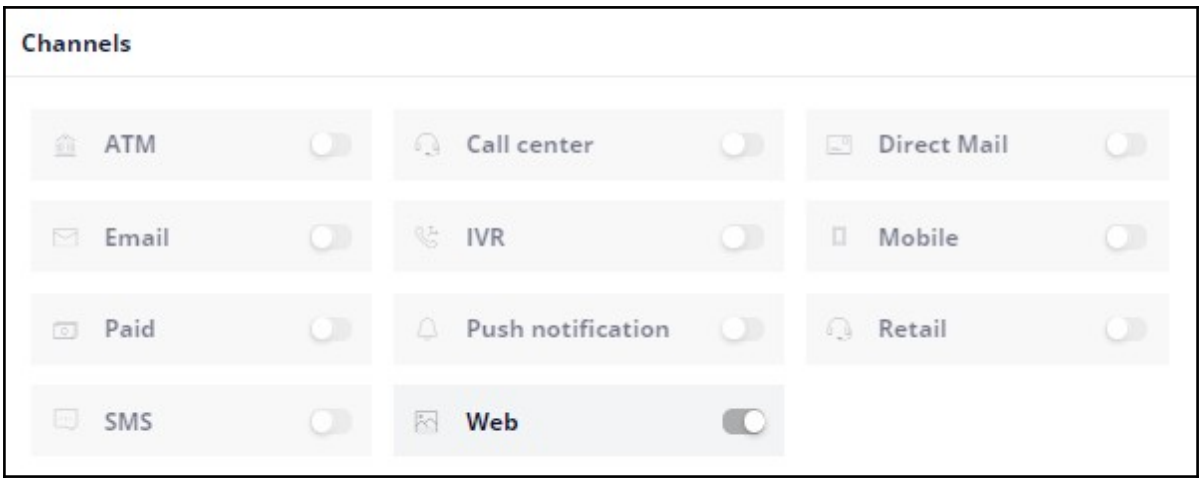

The bank wants to begin by sending the Standard card offer emails to qualified customers. To send emails in an always-on outbound environment, you need to add an email treatment to the action.

Open the **Standard card** action to add an email treatment.

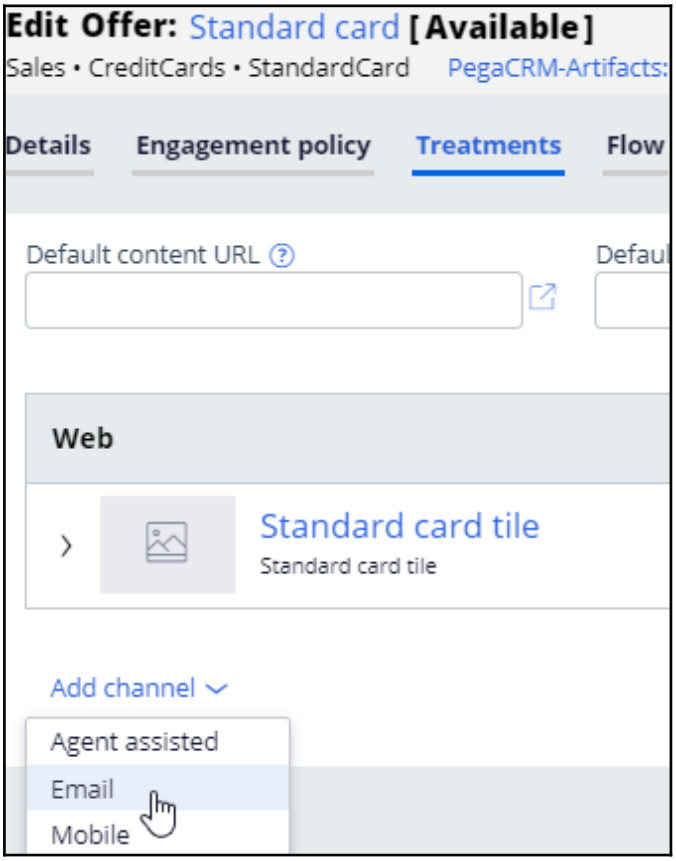

You can use an existing treatment or create a new one. In this case, you use an existing email treatment.

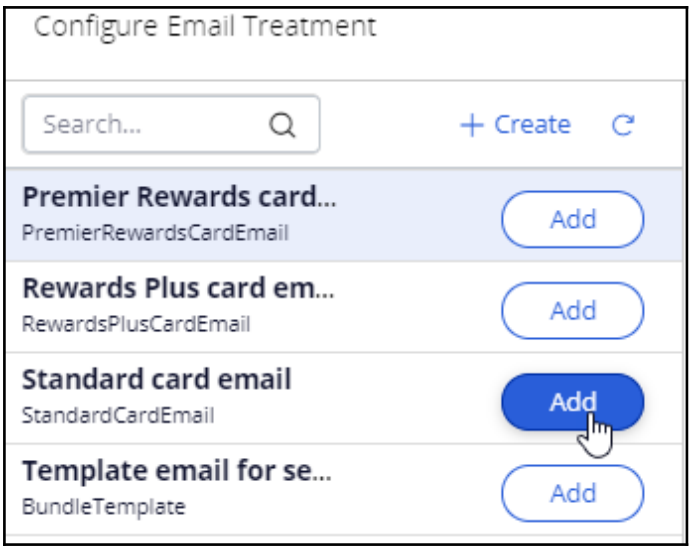

Next, you must configure the action flow. The action flow is a flow diagram that contains the sequence of steps carried out during the lifecycle of an action. You can create action flows once and reuse the flows across actions by using the **Dynamic template** option.

Now, click **Dynamic template** to add an action flow template.

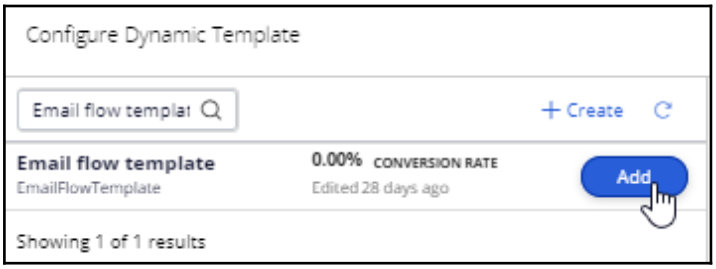

Notice that the flow has only one key shape: the **Send Email** shape. It represents the outbound email delivery.

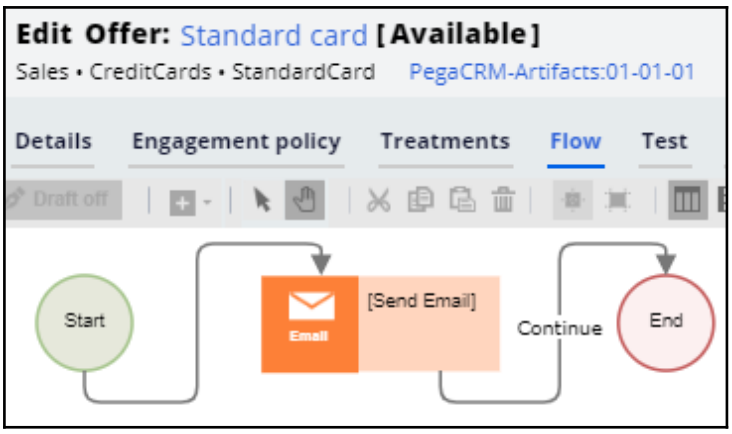

Now, on the **Channels** tab, you can enable the channels and triggers that invoke a Next-Best-Action evaluation.

In this case, as U+ wants to send out emails, enable the **Email** channel.

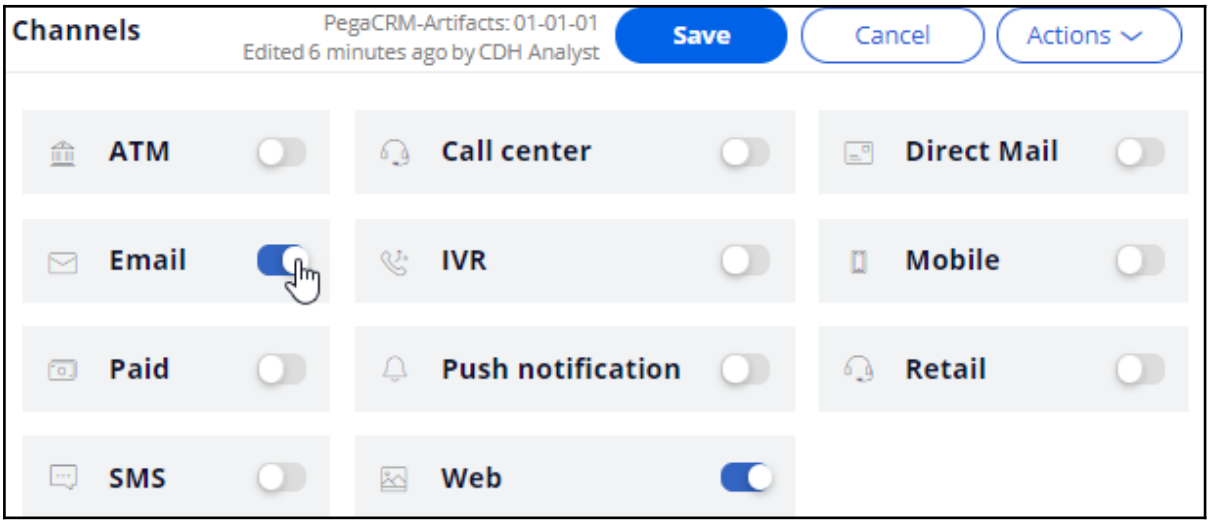

In the always-on outbound paradigm, the brain evaluates the next best action for each customer identified in a segment, which contains the potential list of customers to whom you want to send messages. The system can send these messages at a regular frequency (for example, daily). You configure the **Primary schedule** to set the frequency.

The schedule has two mandatory configurations: **Recurrence** and **Starting population**.

To configure an outbound schedule, first, set the required **Recurrence** frequency. In this case, retain the default frequency, *Daily*. Then, select the **Refresh the audience** check

box. Enabling the **Refresh the audience** option ensures that the latest customer information is always used when running the schedule.

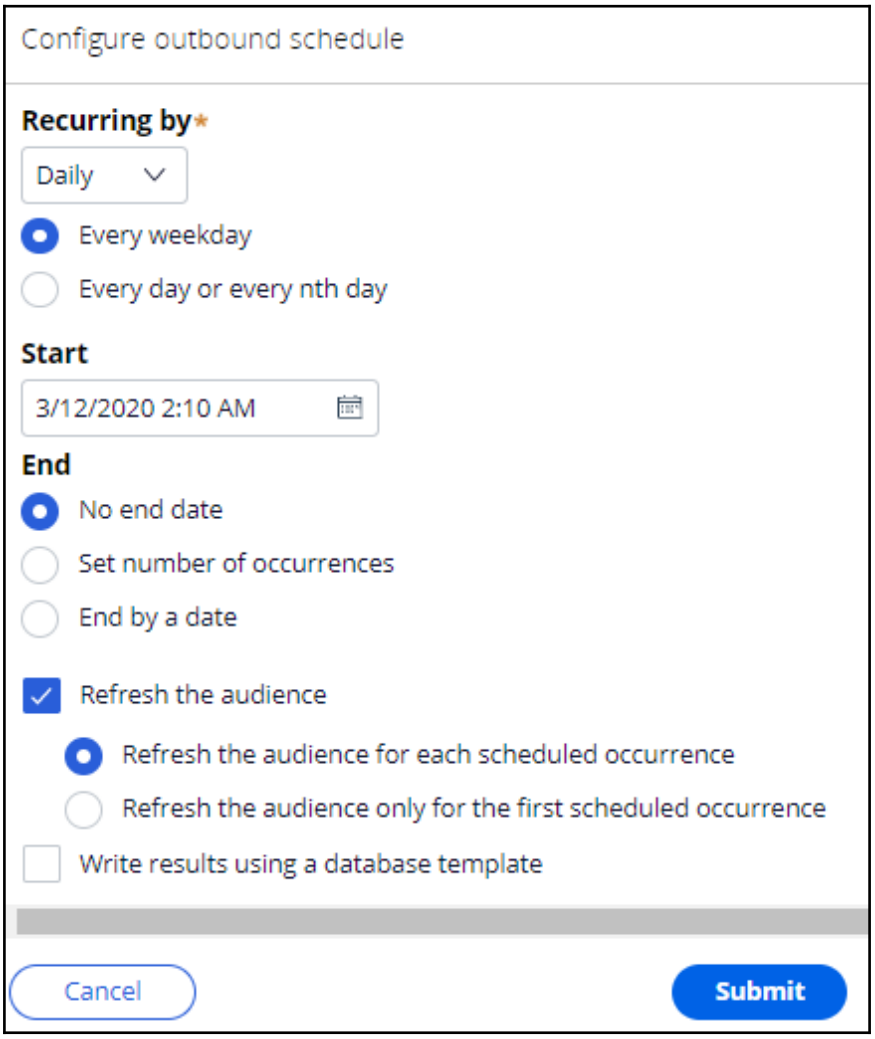

Next, select the segment that contains the list of customers to be considered by the brain for outbound messaging. Typically, this is a large population of all customers to whom you potentially want to reach out. In this case, select the **All mailable customers** segment.

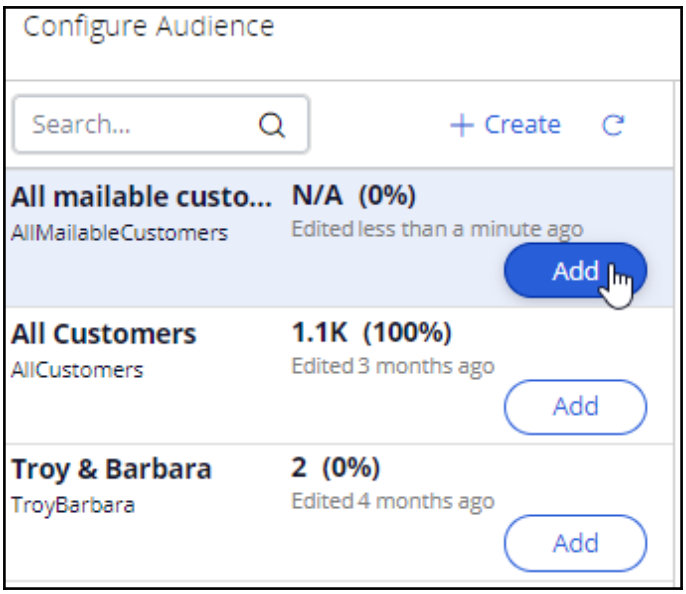

With that, all the necessary configurations for this scenario are complete.

Run the schedule. When the date and time for a scheduled run are reached, an entry is displayed with the status **Scheduled**, which indicates that the run is about to begin.

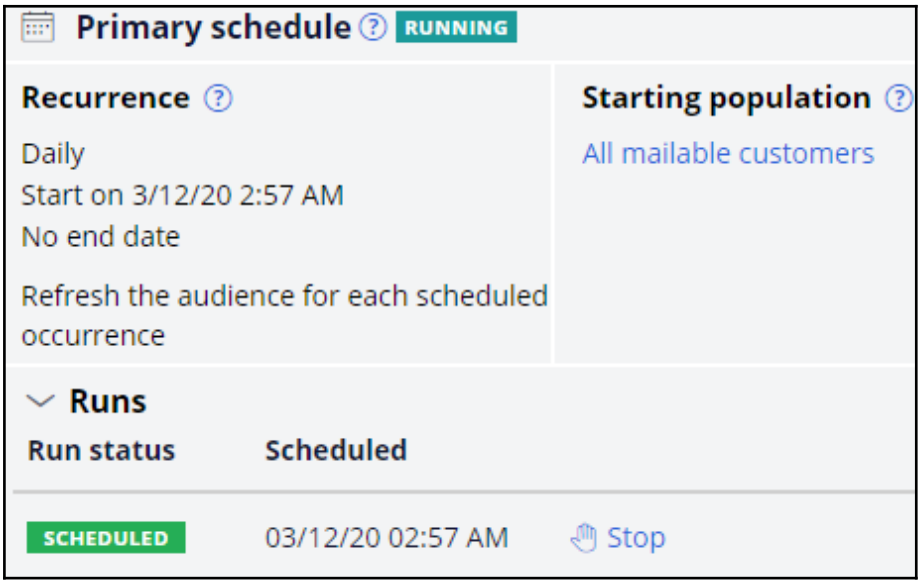

When you refresh, the status changes to **Running** once processing begins.

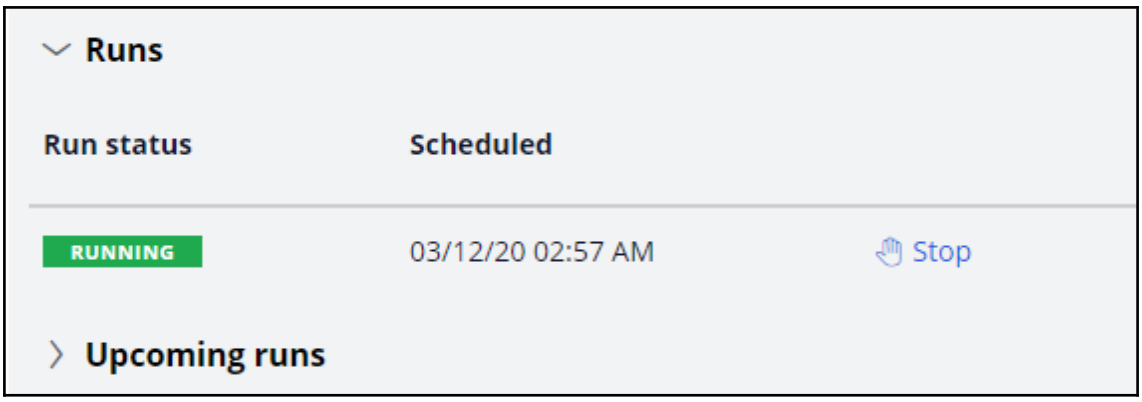

You can refresh to see the latest status. Once the run is complete, you can see the results of the run.

In the completed run results, you can see the number of customers selected in the **Audience** stage, which is the count of customers selected by the Segment.

You can also see the number of actions identified by the strategy, which is the count of customers selected by the brain.

The count in the **Initiated** section represents the number of actions initiated for outbound messaging. The count in the **Result** section is the number of actions that have completed all the steps defined in their action flow.

You can also see the actions in the current run, which the brain selected.

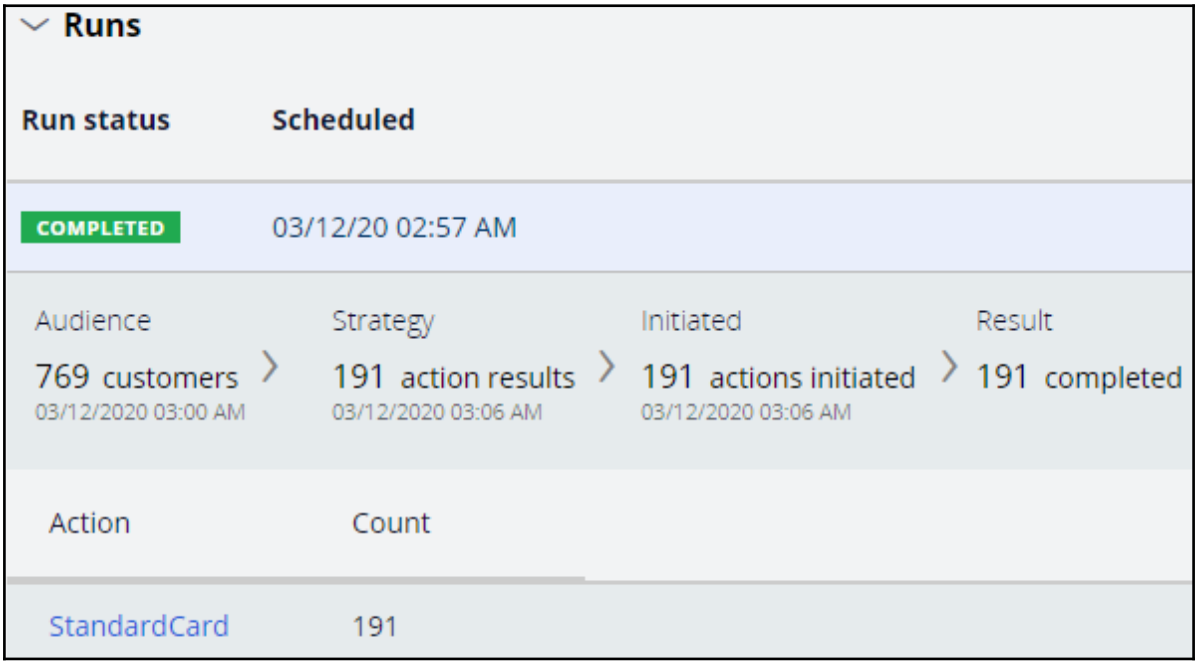

Qualified customers receive the email. When they click the **APPLY TODAY** call to action, they are taken to the preconfigured URL, which, in this scenario, is the Credit Cards landing page on the bank's website.

This demo has concluded. What did it show you?

- How to enable the email channel in Next-Best-Action Designer.
- How to add an email treatment to an action to have an email offer available for the outbound run.
- How to schedule and run an outbound run.

Defining the starting population

Description

The starting population is the potential list of customers you want the brain to consider for always-on outbound customer engagement. Learn how to create a segment and configure it with basic criteria to identify the starting population for Next-Best-Action outbound runs.

Learning objectives

- Describe the concept of segmentation
- Explain how to create a segment for the starting population of an outbound run

Segmentation

What are segments?

In Pega Customer Decision Hub™, the audience is a list of potential target customers, implemented using the segment component. A segment is a reference to a database table containing a list of customer IDs that match specific criteria.

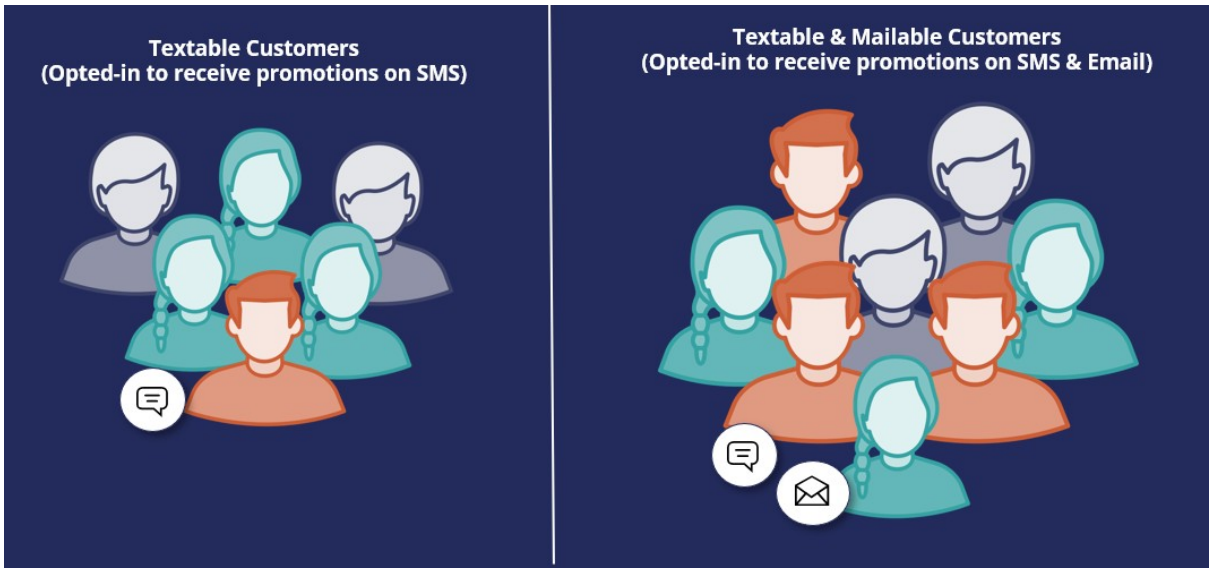

In the always-on outbound paradigm, the Pega brain evaluates the Next-Best-Action for each of the customers identified in a segment, which contains the list of potential customers to whom you want to send messages.

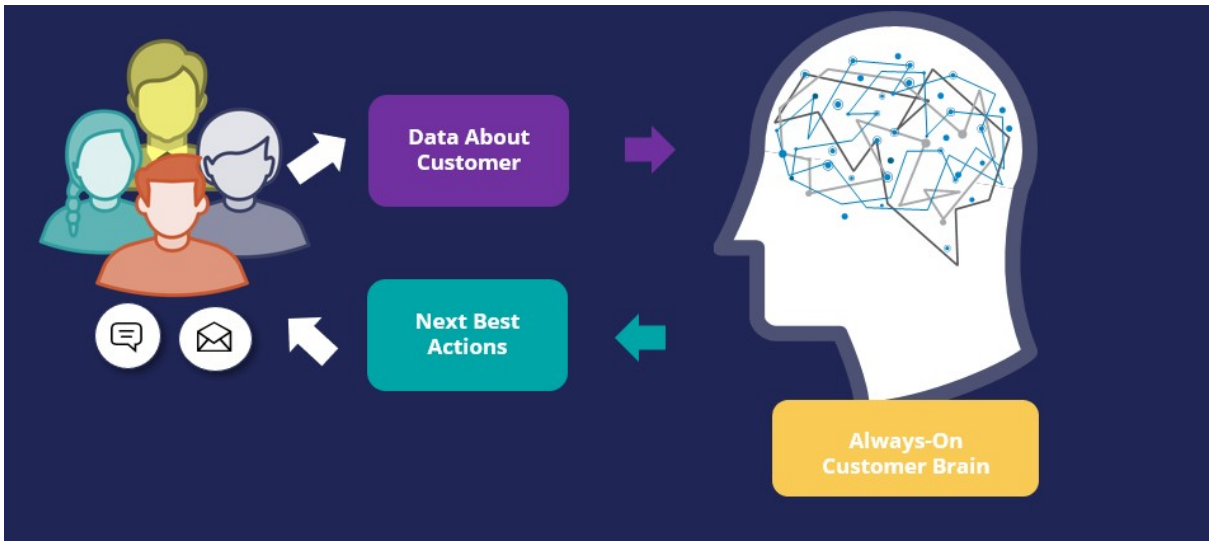

This set of customers is known as the starting population. The customers in this group will be further qualified through eligibility, applicability, and suitability criteria configured in Next-Best-Action Designer.

Criteria in a segment

To select a starting population, you apply a few high-level criteria to a segment. It's best practice not to get too granular at the segment level, rather use eligibility, applicability, and suitability rules configured in Next-Best-Action Designer to fine-tune your criteria once the starting population has been selected.

You can define one or more conditions that must be true for a customer to be included in a segment.

For example, to send emails related to a credit card offer to a target audience, you should first define a criterion that will filter for customers who have opted-in to receive promotional emails.

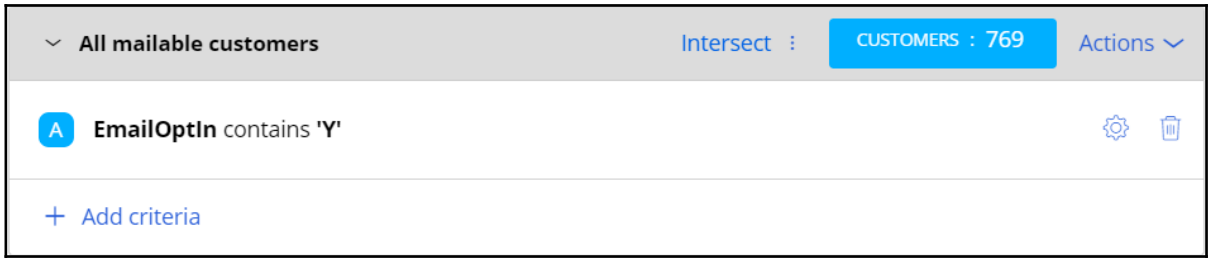

Segments hierarchy

A segment can reference a "parent" segment, enabling you to create a hierarchy of segments.

For example, you can create top-level segments that identify common lists of customers for exclusion or inclusion in various "child" Segments. Standard top-level segments allow global communication rules to be implemented by forming the basis of customer lists that automatically adhere to these rules. The rules can include, for example, standard exclusions, which are the core customers who should not be contacted, or standard inclusions, which are the core customers who may always be contacted. The resulting customer lists can be used as the starting point for creating the final segment.

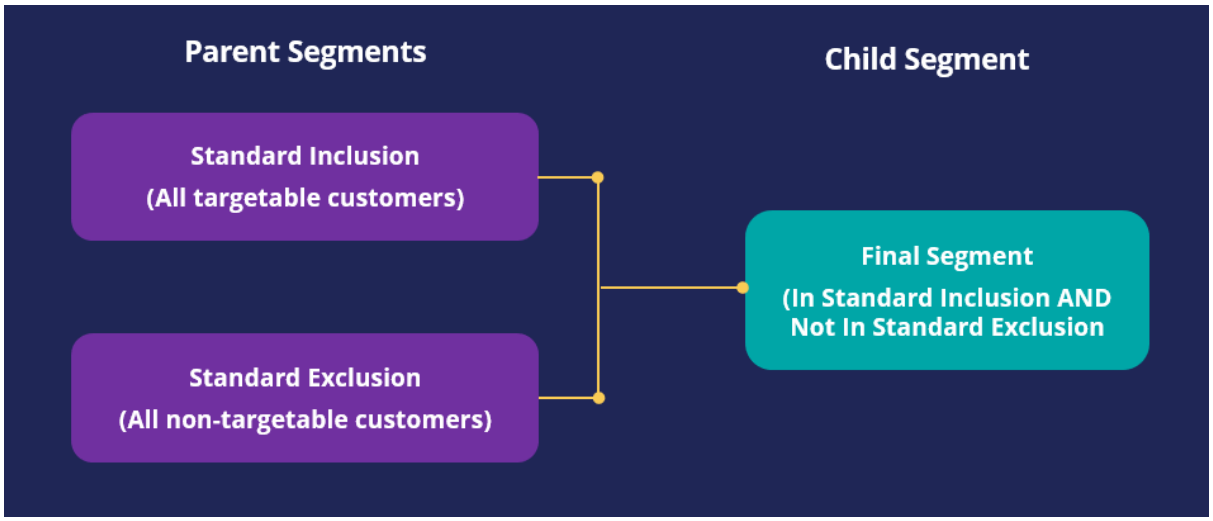

If changes are made to global rules, the definitions of the standard top-level segments can be modified accordingly, and as a result, every segment that references them will automatically be re-aligned to the new rules.

Creating a segment

Business scenario

U+, a retail bank, wants to use the always-on outbound approach to send promotional emails to qualified customers. As the starting population for the outbound schedule, the bank wants to include all customers who have opted in to receive promotional emails.

Creating a segment

To meet the business requirement, you need to create a segment with two criteria.

1. A criterion to reference the existing top-level segment, ***All Customers,*** which contains a list of all customers in the customer base.
2. A criterion to check if customers have opted in to receive promotional emails.

Segments are managed on the Segments landing page. The Segments landing page can be accessed from **Audiences** -> **Segments** in the navigation menu.

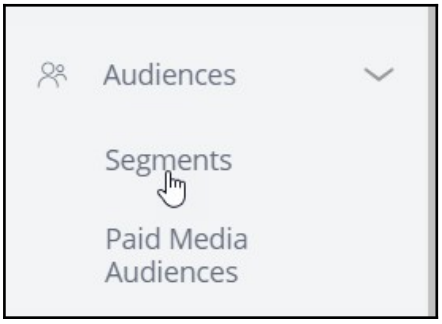

You can open existing segments or create a new segment from here.

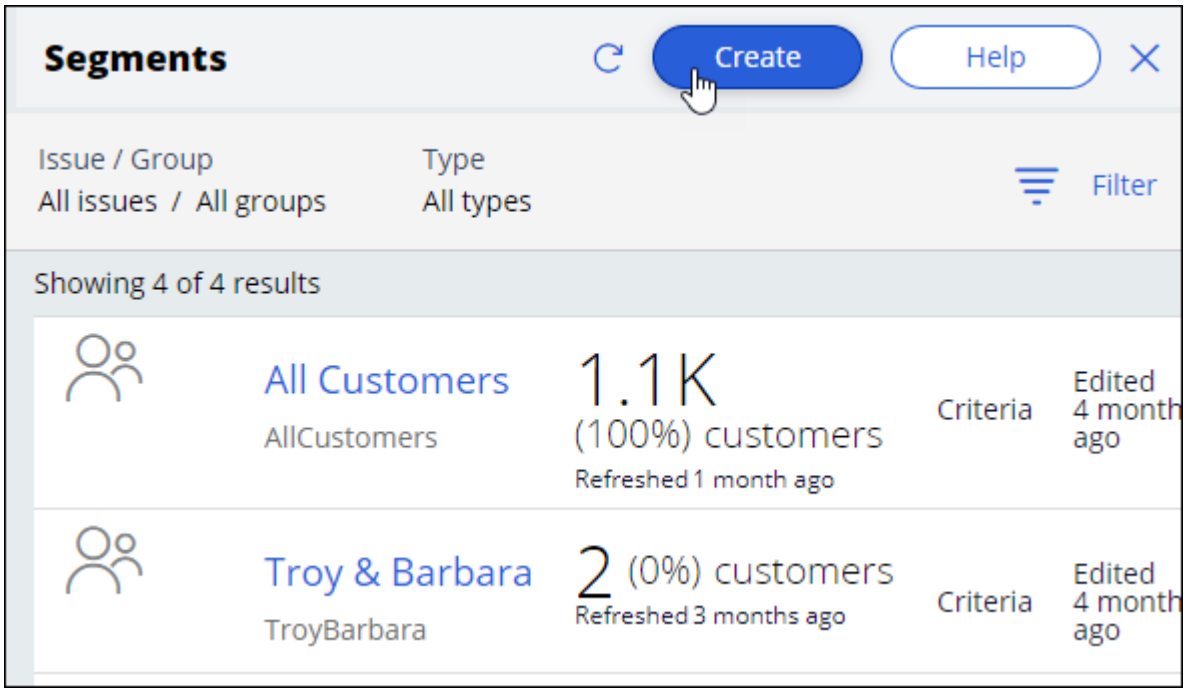

Provide a short description for the new segment, and the identifier is automatically populated when you tab out.

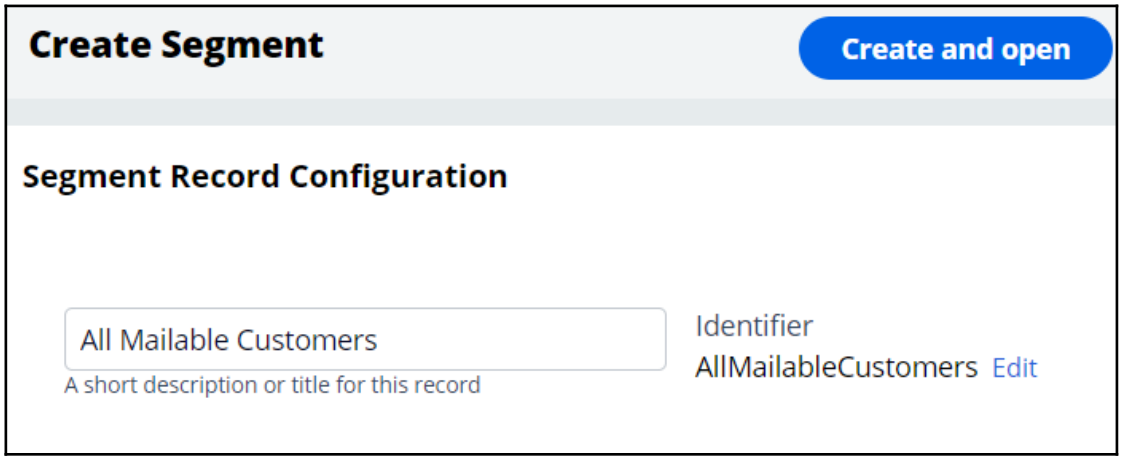

Entering the segment criteria

The **Design** tab is where you assemble the criteria required to identify customers of interest.

Segment criteria are grouped for ease of reading. So, before you can add the first criteria, add a criteria group.

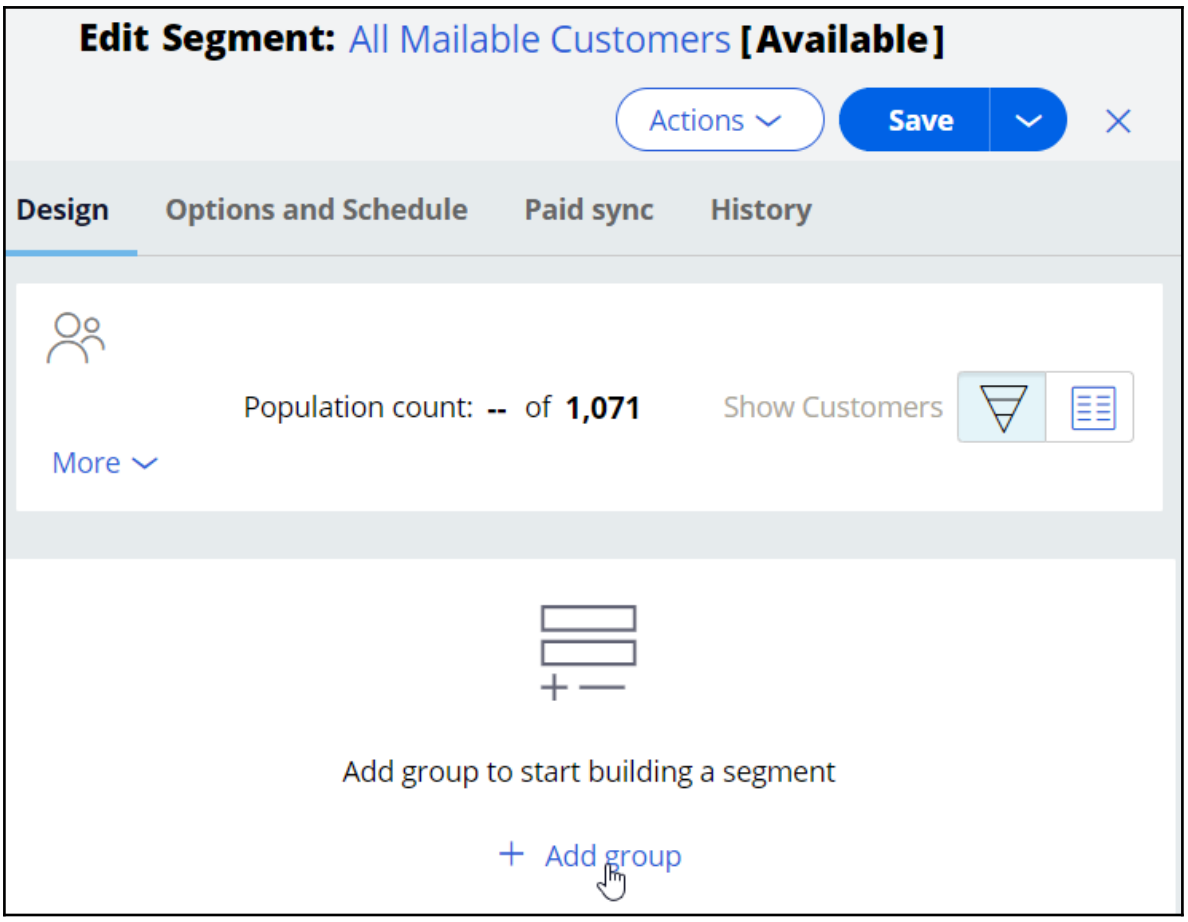

You can provide a more descriptive name for the criteria group.

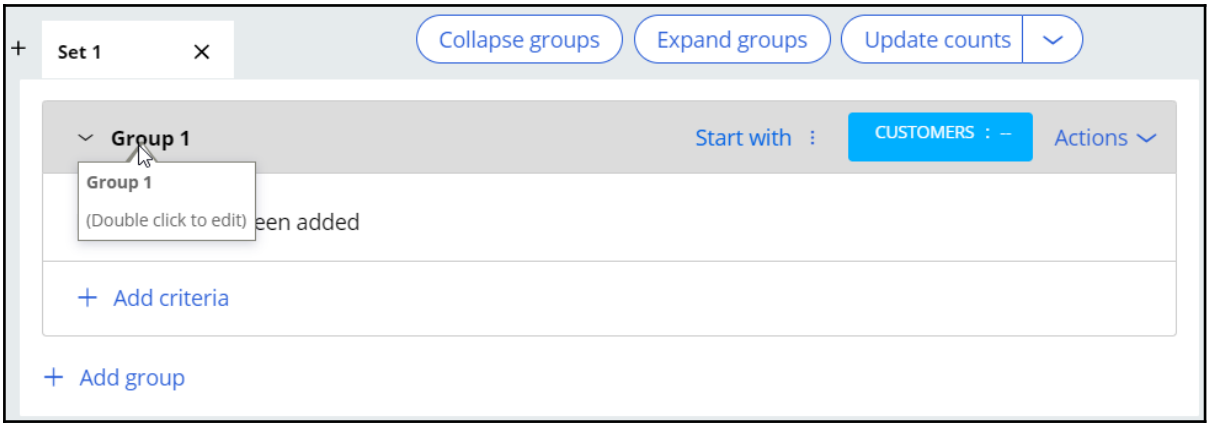

In this scenario, the bank wants to consider all customers in the customer base, then select the customers who have opted in to receive promotional emails. The bank has already created the ***All Customers*** *segment*, which contains the list of all customers. You can reference this segment in the first criterion.

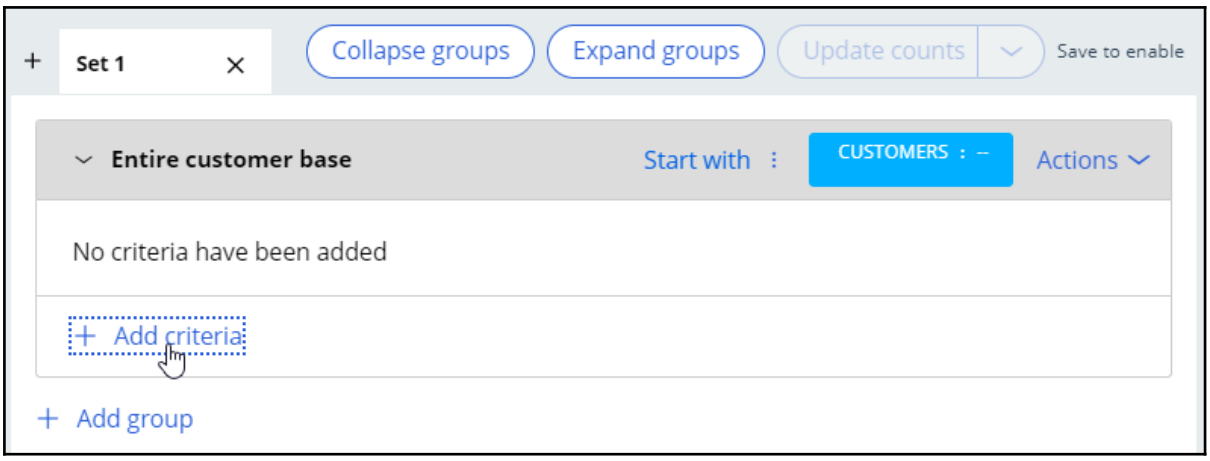

In the **Add Criteria** window, the **Segments** tab lists all available segments. Select the existing ***All Customers*** segment.

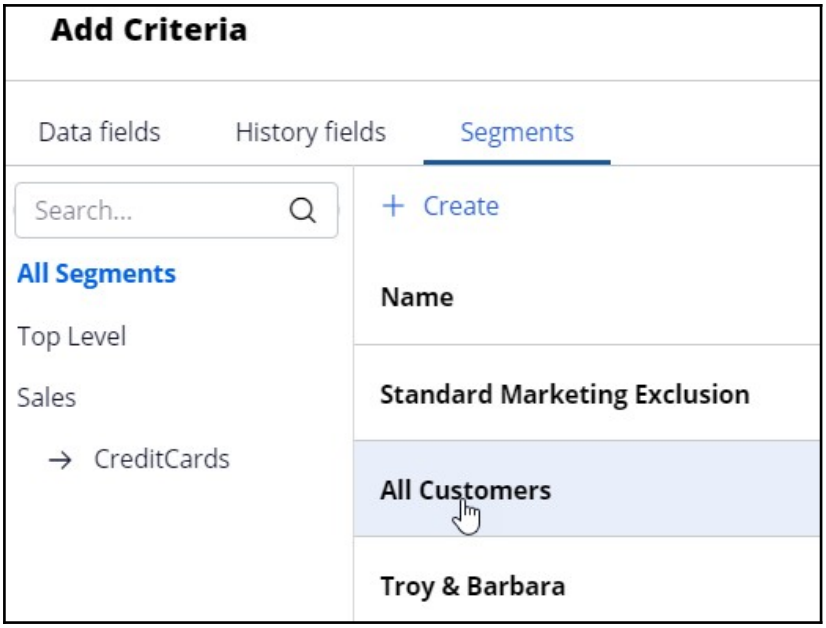

The **All Customers** segment is added as the first segment criterion. The **Exclude** option allows you to exclude the result of a criterion.

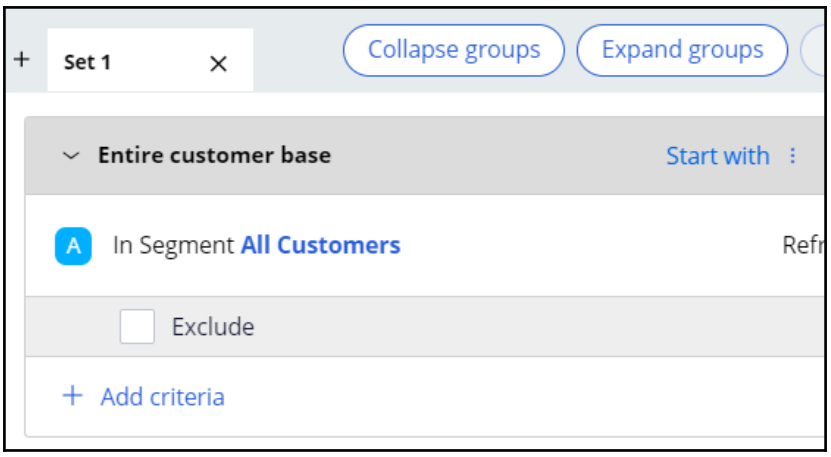

The next step is to add another criterion to identify customers who have opted in to receive promotional emails. You can add the new criterion in the existing criteria group or in a new group.

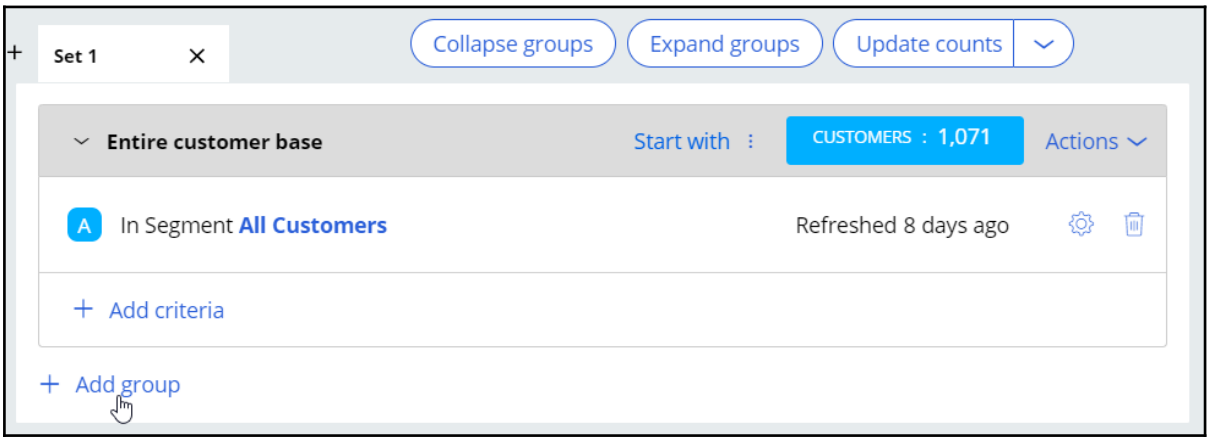

In the **Add Criteria** window, the **Data Fields** tab lists all available customer properties. You can either select the field directly or search for it by entering the partial name in the search field. Then, select the desired field from the search results. In this system, the field **EmailOptIn** contains the information about whether a customer has opted in to receive emails or not.

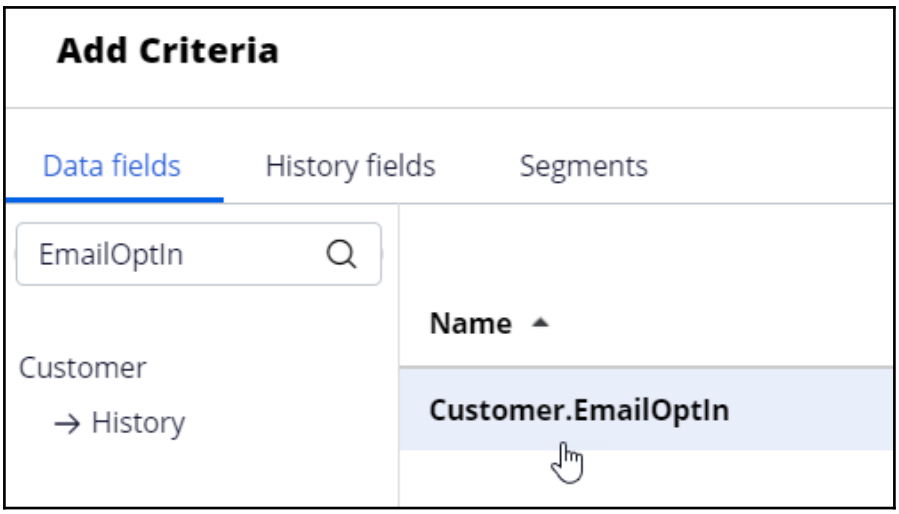

Once the field is selected, you can configure it to meet the business requirement. In this system, the value ***Y*** indicates that the customer has opted in, and ***N*** indicates that he/she has not opted in to receive promotional emails.

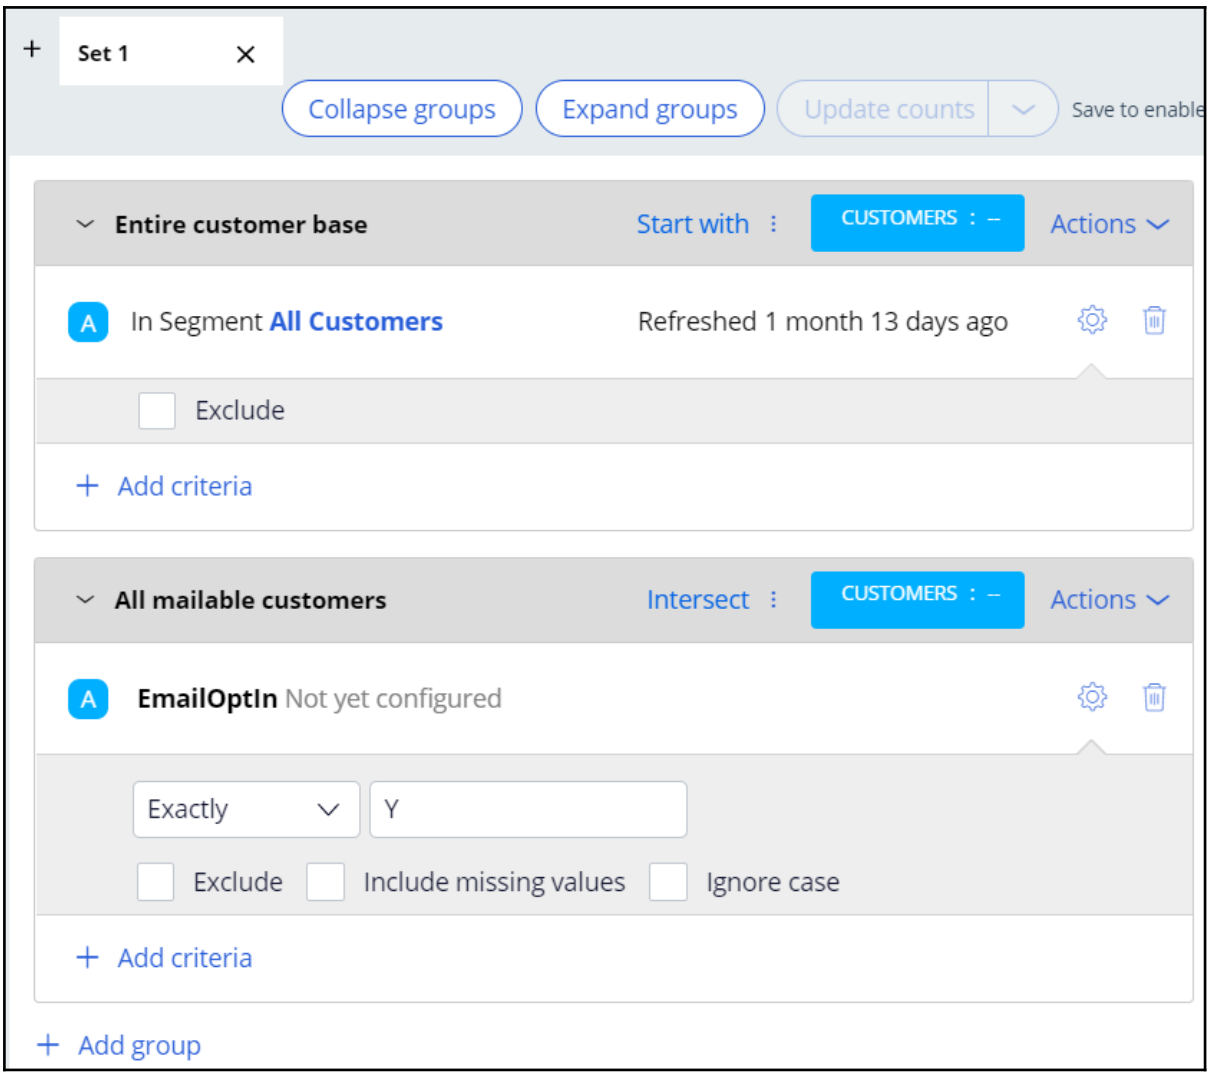

The **Include missing values** option is used to include results that do not have a value specified for the selected field.

The **Ignore case** option is used to perform a case-insensitive search. You can leave both options unselected for this scenario.

The **Update Counts** option is used to update the population count of the segment while you are building it.

> Note: The **Update Counts** option is enabled when you save the segment.

You can combine the result of a criteria group with the result from the criteria group above it using one of the three options:

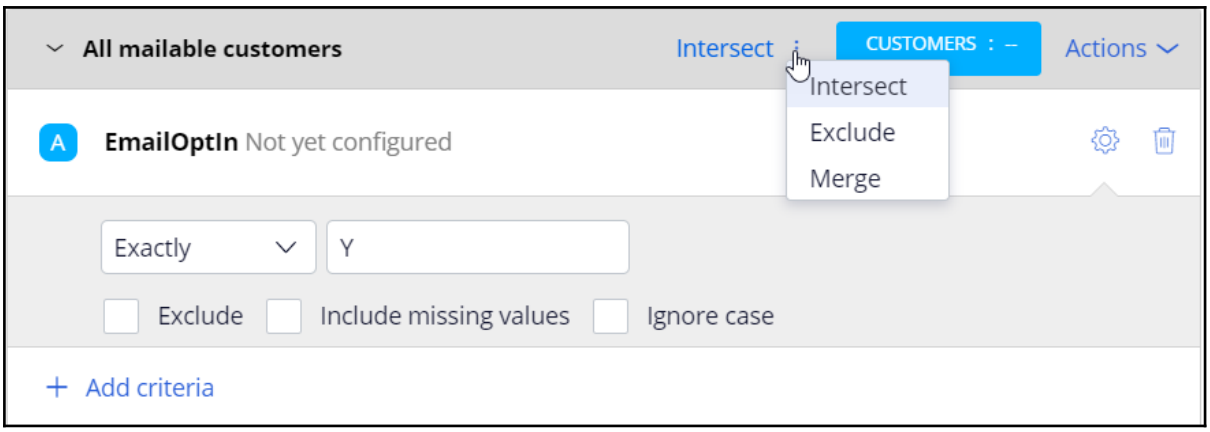

Intersect selects the customers the two results have in common.

Exclude excludes the result of the current criteria group from the result above it.

Merge combines the results of both criteria groups.

Save and run the segment

When you are done defining all the segment criteria, you can save and run the segment. Click **Run** to run the segment.

While the segment is running, you can stop the execution by clicking the **Stop** button.

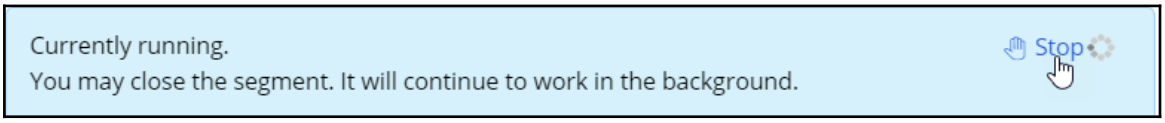

Once the run completes, the number of customers in the resulting segment is displayed as **Population count**.

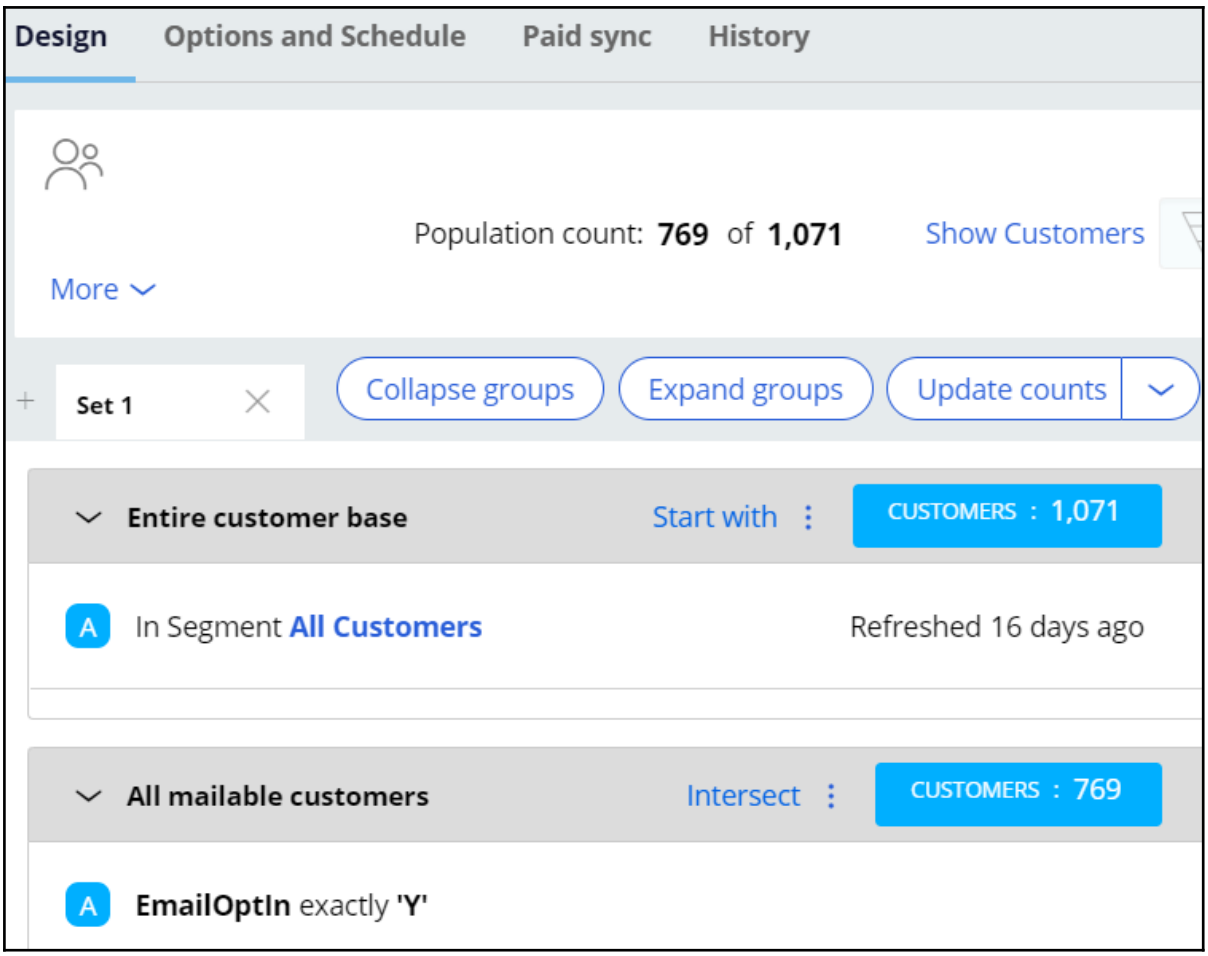

There are ***1071*** customers in the ***All Customers*** Segment, out of which ***769*** customers have opted in to receive promotional emails.

Click the **Show Customers** link to preview the list of customers selected for this segment.

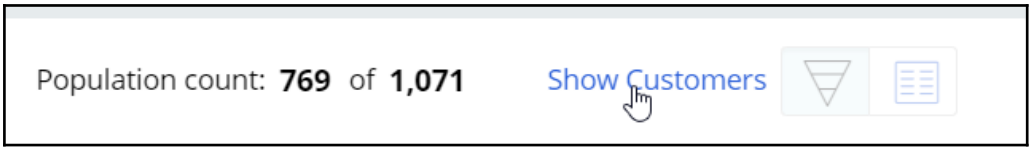

The list is shown in a separate window.

CustomerID*	ACCOUNT ID	Title	Full Name	Gender	Age	Language	MaritalStatus
14			Troy Murphy	M	26		Married
15			Barbara Stockton	F	32		Single
16			Joanna Williams	F	25		Married
17			Robert Walton	M	26		Single
18			John Smith	M	45		Single
19			Arnold Green	M	56		Single
CON-1			Layne McLaughlin	F	19		Single
CON-10			Ezekiel Funk	M	32		Single

Key segment configurations

On the **Options and Schedule** tab, there are two key settings under **Data Options** that control how data in the segment is refreshed at run-time and what level of control a segment can exert over other segments.

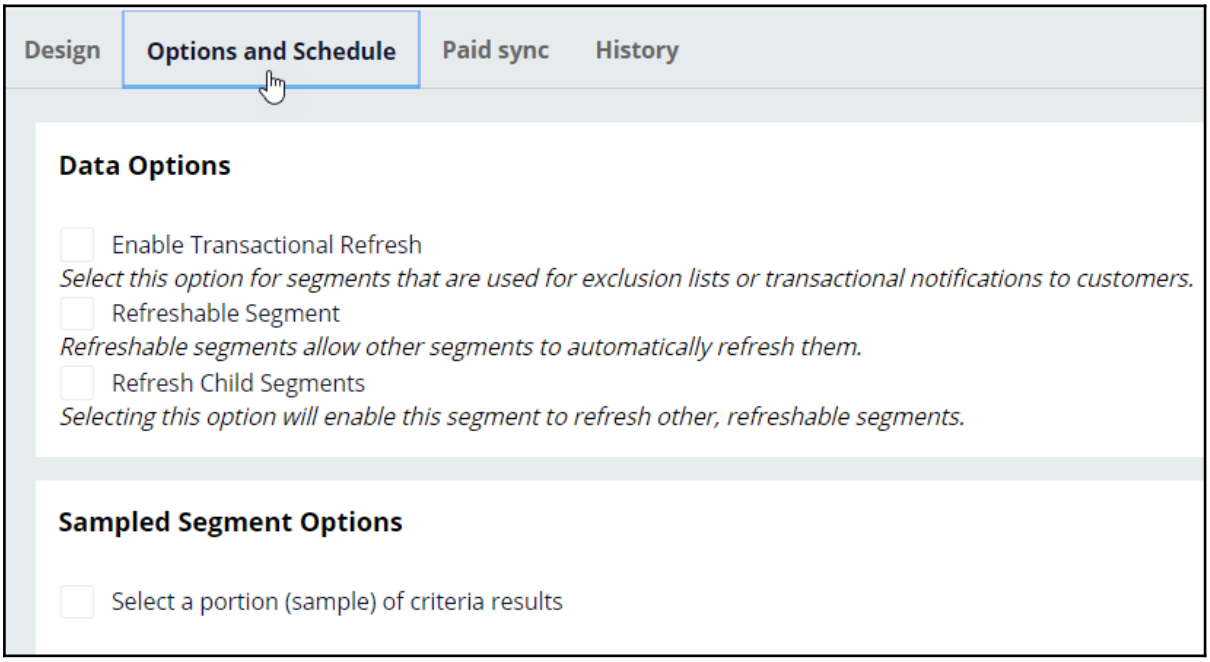

When the **Refreshable Segment** option is enabled, the segment will be refreshed when the other segments that reference it are run.

The **Refresh Child Segments** option allows you to control the cascading effect of the refresh operation. The segment can automatically refresh data in its child segments, if they are marked as refreshable.

The **Sampled Segment Options** allows you to select a random sample of a segment result to be selected as the final segment result.

Segment run history

On the **History** tab, you can view the details of previous execution runs. The history lists the runs triggered manually by a user, indicated by ***Populate Segment Requested by: <user name>.*** *It also lists* runs triggered by the outbound schedule in Next-Best-Action Designer, indicated by ***Populate Segment Requested by: Program run***.

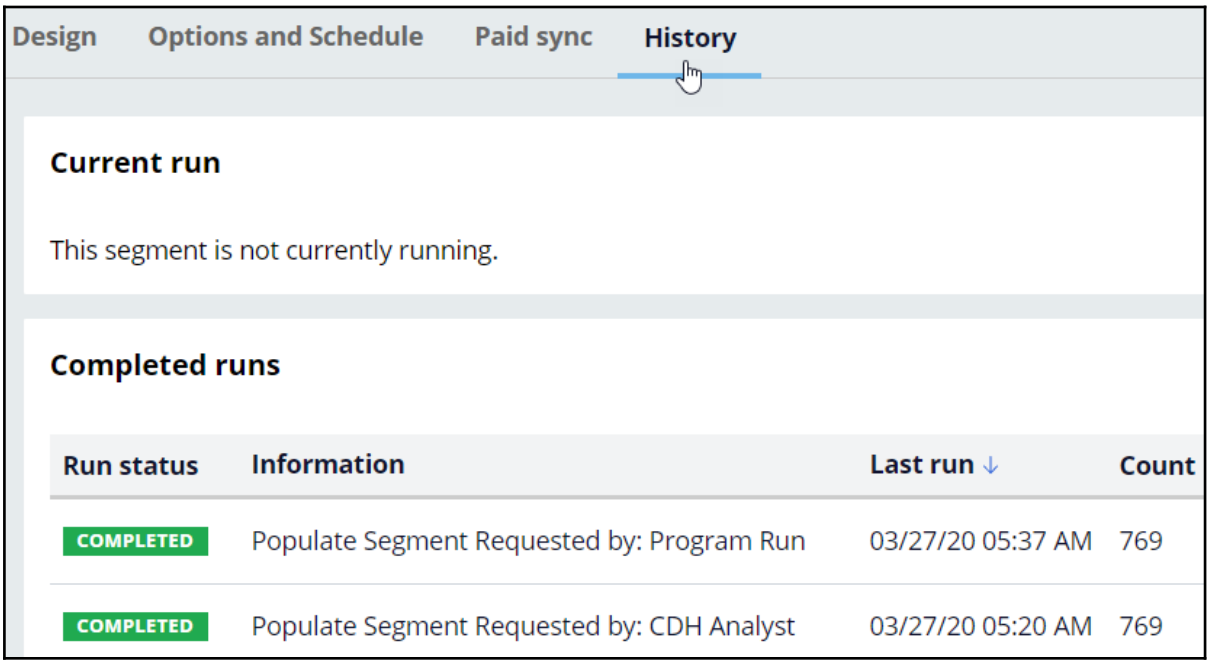

Defining an action for outbound

Description

A Next-Best-Action is selected by the brain after it considers customer engagement policies, business objectives and recommendations from the AI. An action can be presented to the customer on various inbound and outbound channels. Learn how to configure an action to make it suitable for outbound customer engagement and design a visually appealing email message using easy-to-use templates. Also, learn how to configure an action flow, which contains the steps involved in delivering an action to a customer.

Learning objectives

- Create an email treatment
- Design visually appealing emails using templates
- Configure actions with re-usable action flows

Creating an email treatment

You can send offer emails to qualified customers in an always-on outbound environment using email treatments. An email treatment helps to design a visually appealing email message for customers.

Transcript

This demo shows you how to create an email treatment to present offer details in an email.

U+, a retail bank, wants to promote their Rewards credit card via email. They want to send the email to qualified customers in an always-on outbound environment.

The email is personalized and contains the offer details and a call to action. When customers click the **APPLY TODAY** call to action in the email, they are directed to the credit cards page on the bank's website to learn more about the offer and apply for it if interested.

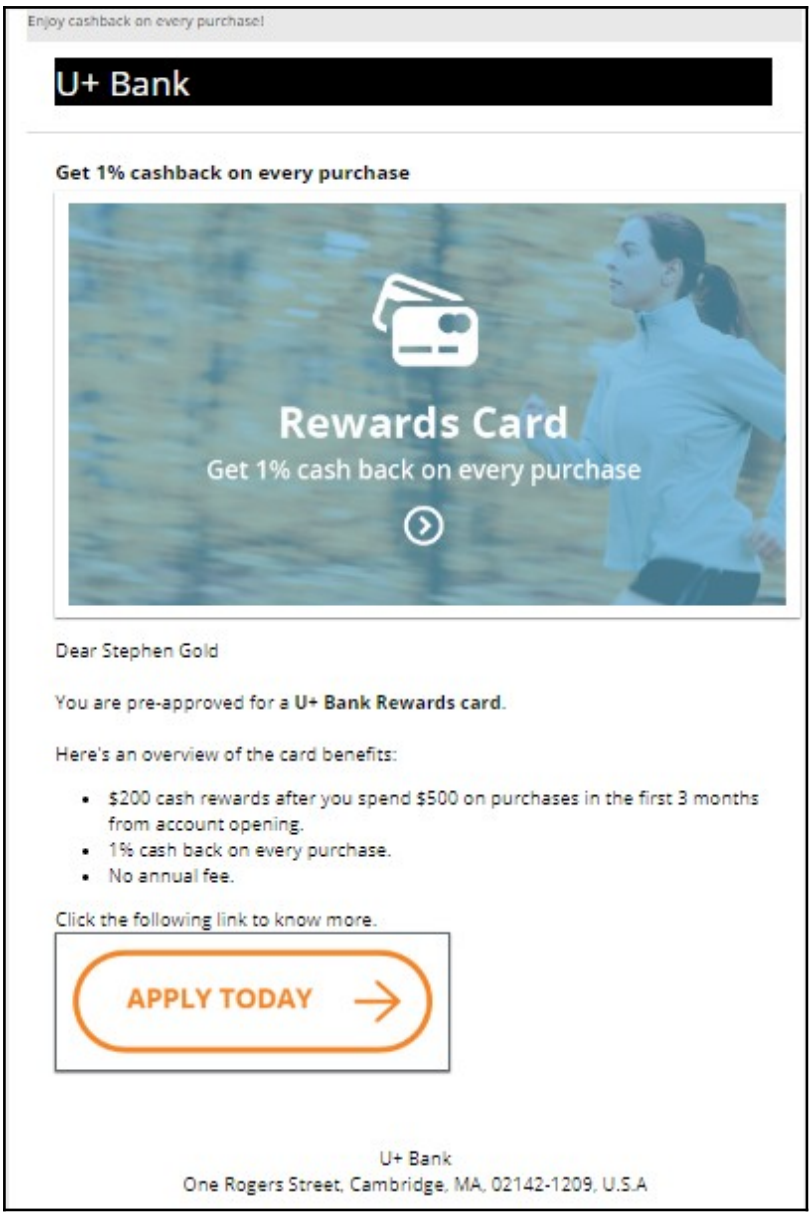

This is Pega Customer Decision Hub™ portal. To send emails in an always-on outbound environment, you need to add an email treatment to every action that you want to promote through email. An email treatment contains the message that you want to communicate to customers.

In this case, you create and add an email treatment to the Rewards card. Open the Rewards card to add new treatments to the **action**.

You can manage treatments associated with an action on the **Treatments** tab. To associate a treatment with an action, click **Add channel**.

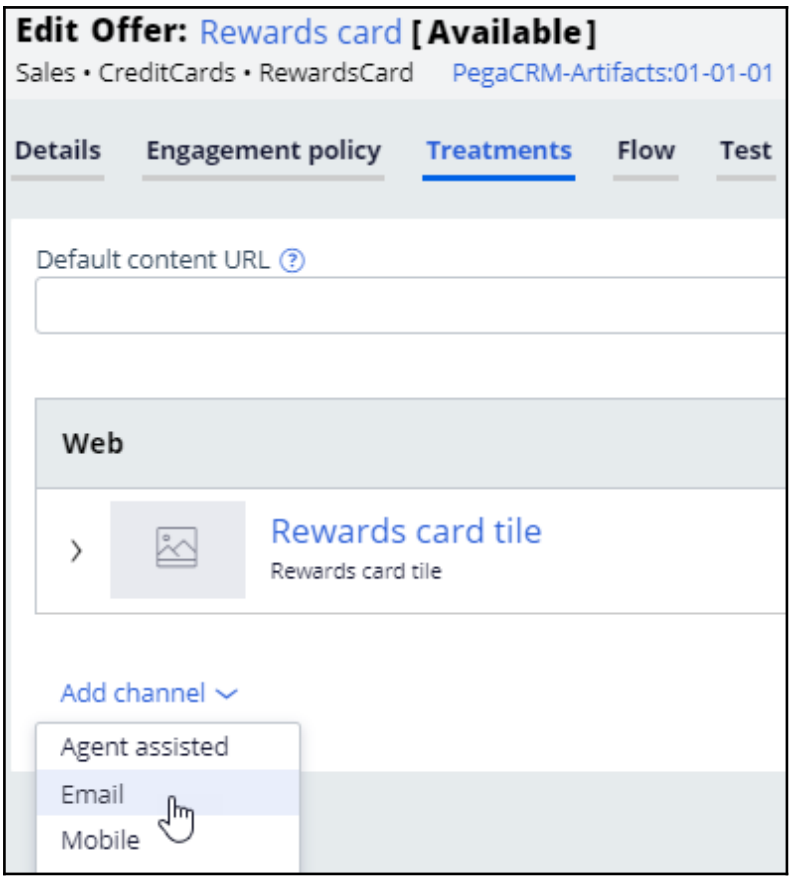

You can use an existing email treatment or create a new one. In this case, you want to create a new treatment.

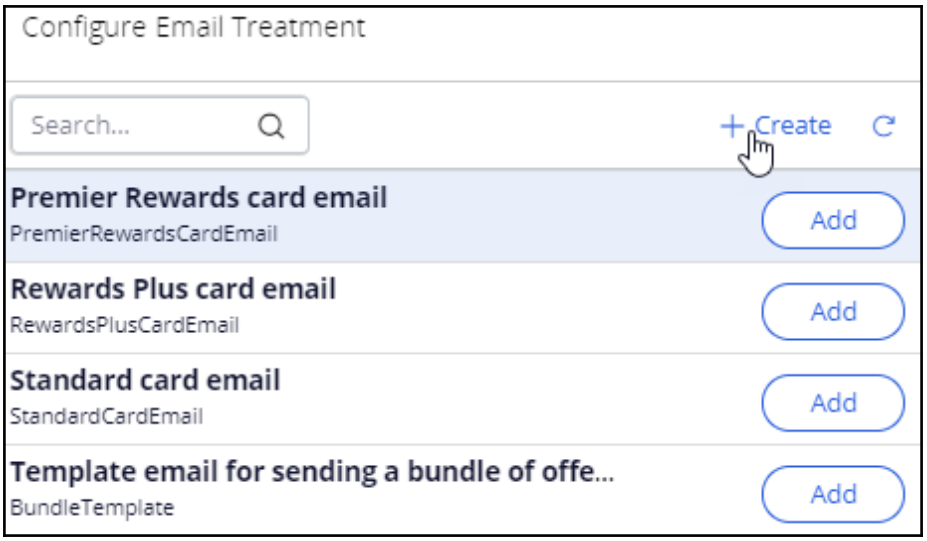

Now, enter an appropriate **short description** to create the new treatment. It is a best practice to create the treatment in the same issue/group hierarchy in which you plan to use it; this will make the treatment easier to access and manage.

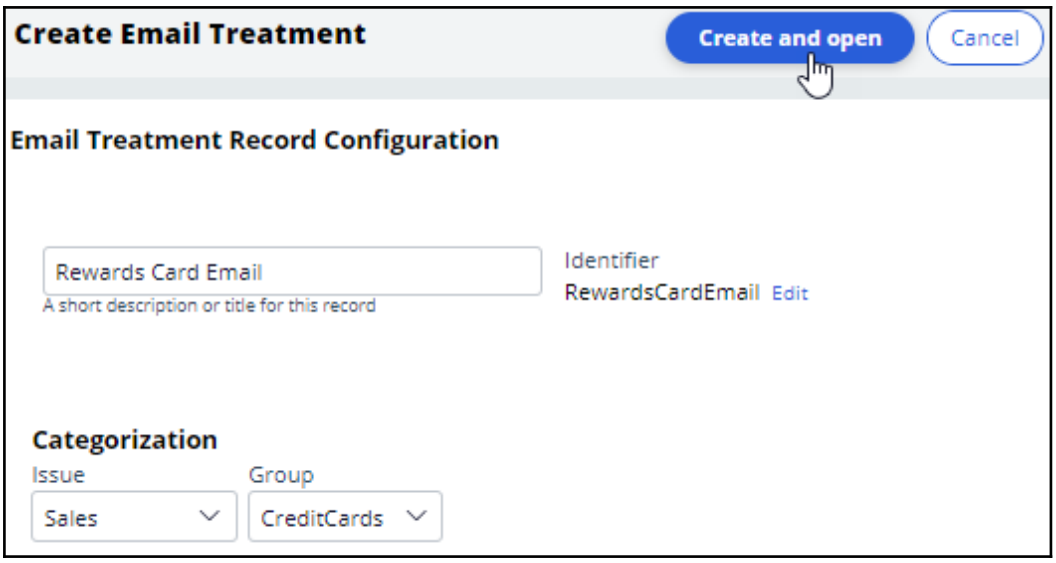

Then, enter the email content.

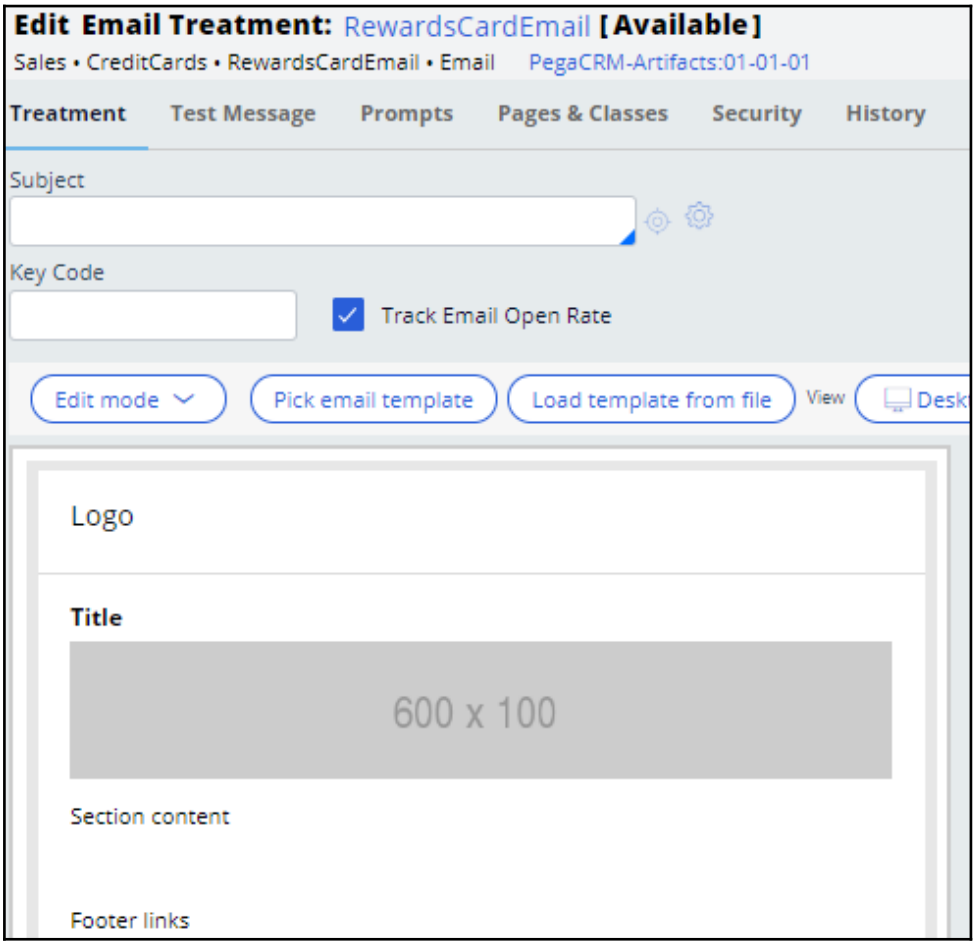

First, enter the email subject. You can include both static text and dynamic placeholders in the subject.

To access dynamic placeholders, start typing dot followed by Customer to access customer properties or OfferData to access action properties.

Dynamic placeholders allow you to personalize the subject. You can construct the placeholders with both customer and offer attribute properties.

At runtime, the properties are replaced with the real values. For example, if the customer's first name is *Stephen*, the generated subject line reads: *Hi Stephen, here's your U+ Bank card*.

Then, enter an appropriate keycode. When the **Track Email Open Rate** option is enabled, the system records an impression when a customer opens an email.

You can design the email using one of the three edit modes: Simple Rich Text Editor, Source Only, and Templated Email.

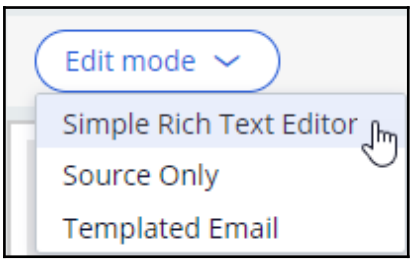

The **Simple Rich Text Editor** provides a blank page with a rich text editor, which allows you to compose an email manually.

The **Source Only** mode allows you to work with HTML content only. If the content is designed in an external tool, you can copy and paste the HTML content in this mode.

The **Templated Email** mode provides you with several easy-to-use templates for designing appealing emails; this is also the default mode.

These are the available template options.

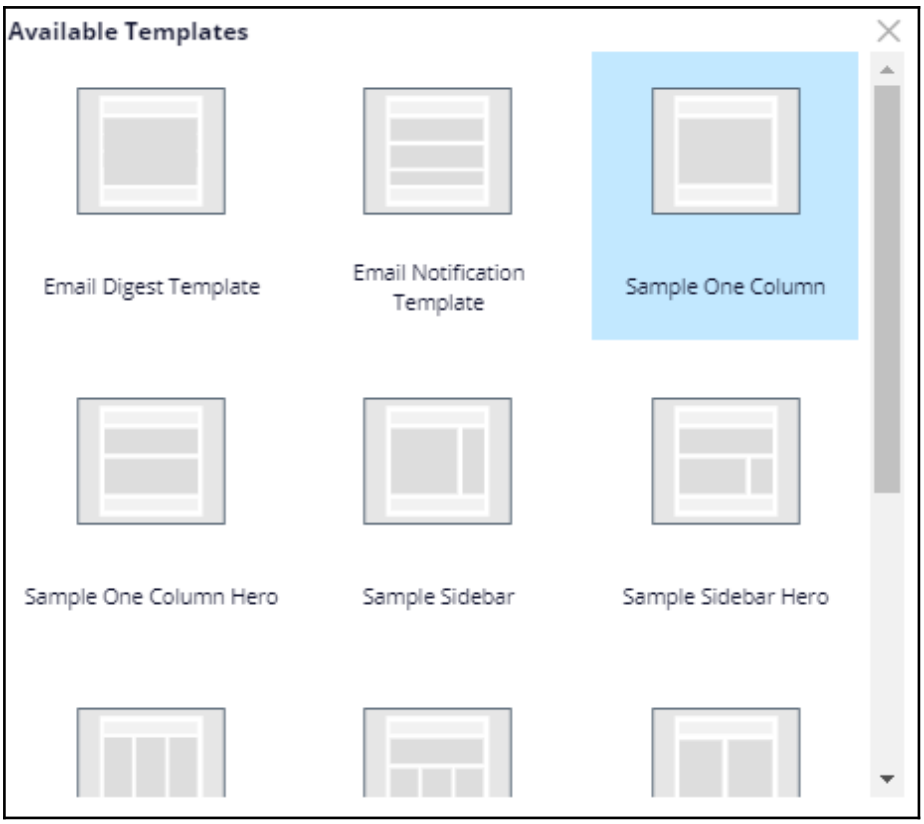

In this case, you use the default **Simple One Column** template.

The **Simple One Column** template contains four sections: Preheader, Header, Body, and Footer.

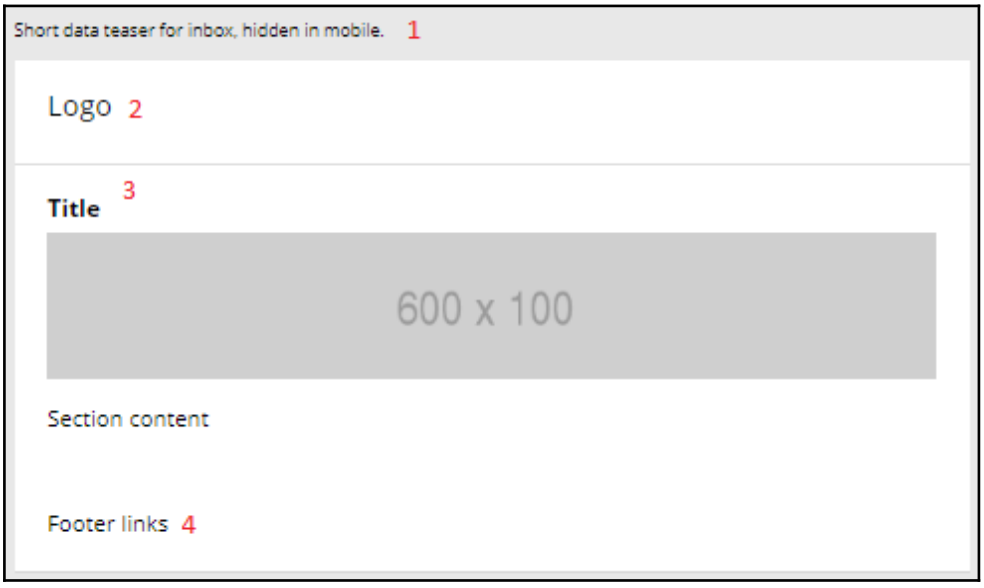

The Preheader section can contain a short **data teaser for inbox** that is intended to persuade customers to read the email further. This text is hidden on a mobile device.

The Header section can contain the organization's logo.

The Body section contains a placeholder for a title, image, text, or buttons as required. This section is where the main content goes.

The Footer section can contain contact details or links, if any.

When you click on the different placeholders or sections in the template, an editing pane is displayed on the right side of the screen.

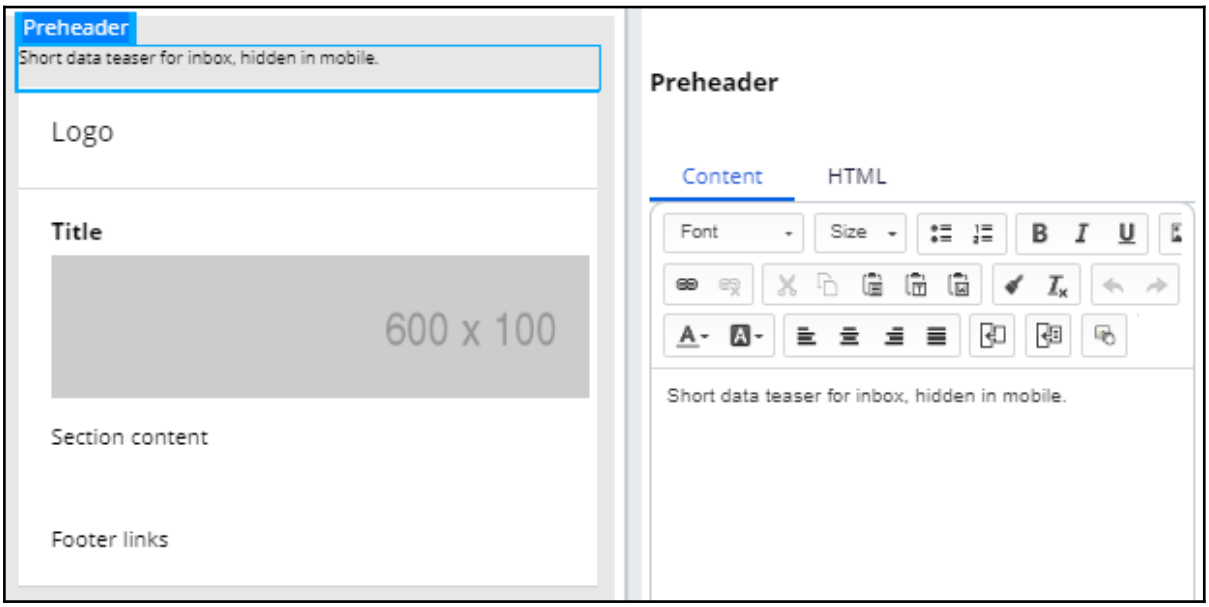

This editor also allows you to copy and paste HTML content from other tools using the **HTML** tab.

For example, here is how you can create an HTML version of the U+ Bank logo by using HTML code.

While composing the section, you can see the preview on the left.

You can also add a title to the email.

To add an image, you can open the image placeholder, which opens the **Image Properties**.

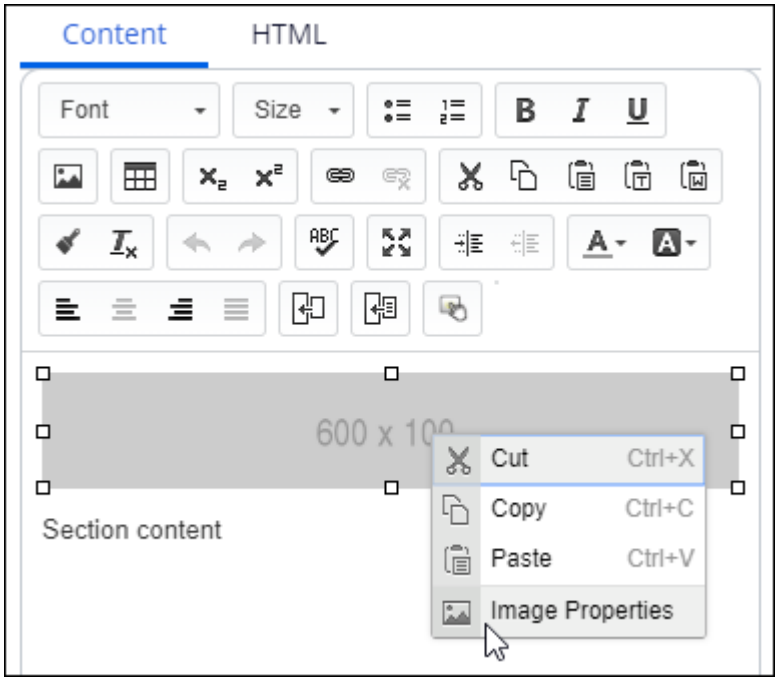

In the **Image Properties,** you can reference an image from a content server.

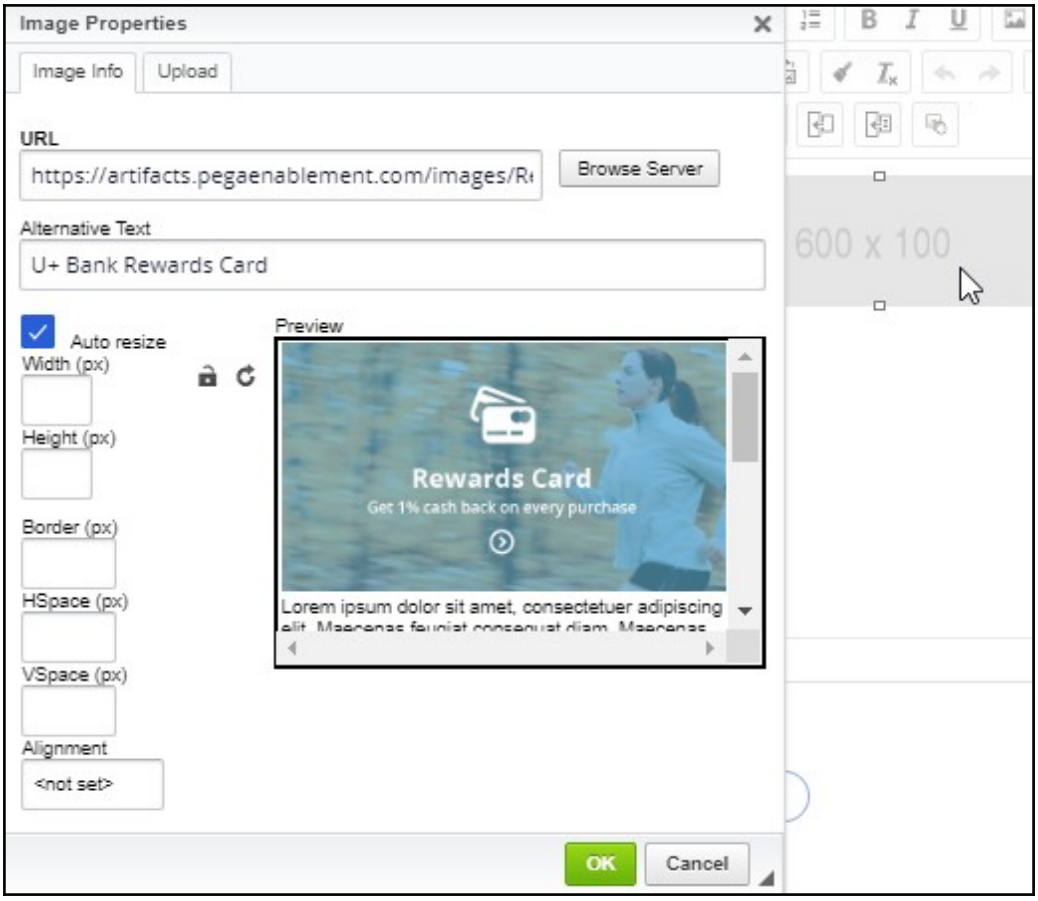

It is best practice to host an image on an external server and reference it in the email. Do not embed images directly in the email, as it makes the email bulky.

In the email content, you can add static content and dynamic placeholders to personalize the message.

For example, the property **.Customer.pyFirstName** can be added as a placeholder. The customer's real first name will replace the placeholder at the time of email delivery.

```
Dear <<.Customer.pyFullName>>
```

To add a customer or offer attribute as a dynamic placeholder, use the **Insert Property** button.

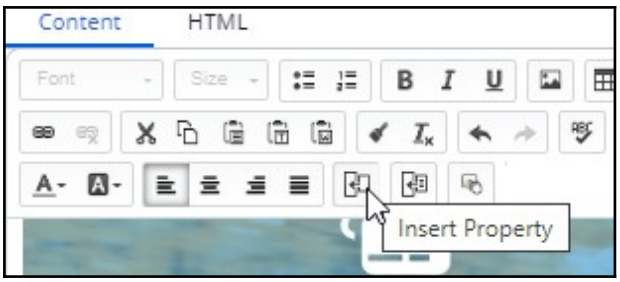

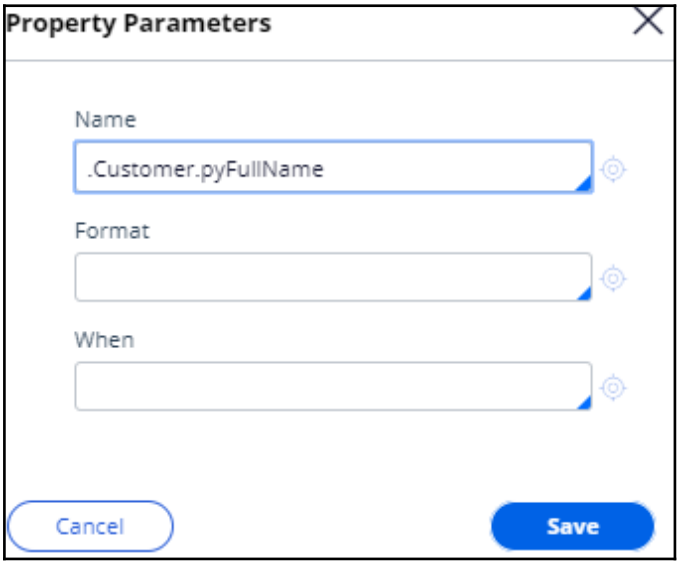

Any property that you add in an email treatment has an equivalent Pega markup tag (**pega:reference)** to reference customer or offer attributes. After you insert the property, you can view the tag (which is automatically added) in the **HTML** tab. When you design the email in an external application, you can use these tags to add dynamic placeholders, so the personalization works when you import the HTML into the email treatment.

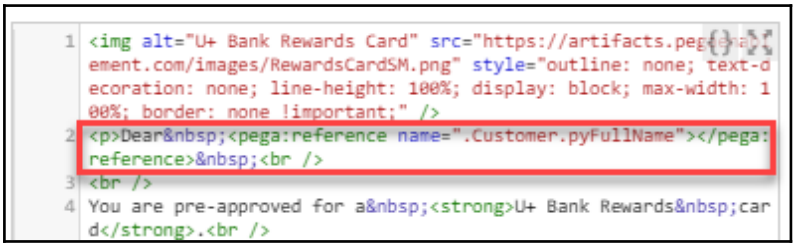

Note that to add an offer attribute, you can add a dynamic placeholder, which is then replaced with the offer attribute at the time of email delivery. For example, you can use <<.OfferData.pyLabel>> to reference the offer attribute.

In this case, you may want to replace U+ Bank card names with an offer attribute by using the <<.OfferData.pyLabel>> placeholder.

You can also add links or buttons that the customer can click to respond to an offer or view more details. In this case, you add an **APPLY TODAY** call to action button.

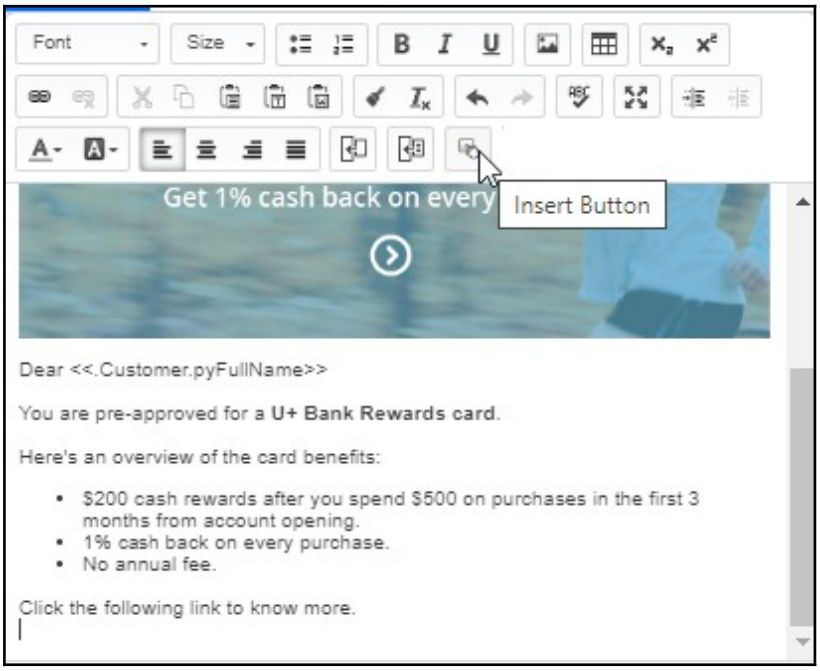

In the **Button Parameters**, the **Function** drop-down menu lists the possible functions that you can assign to a button.

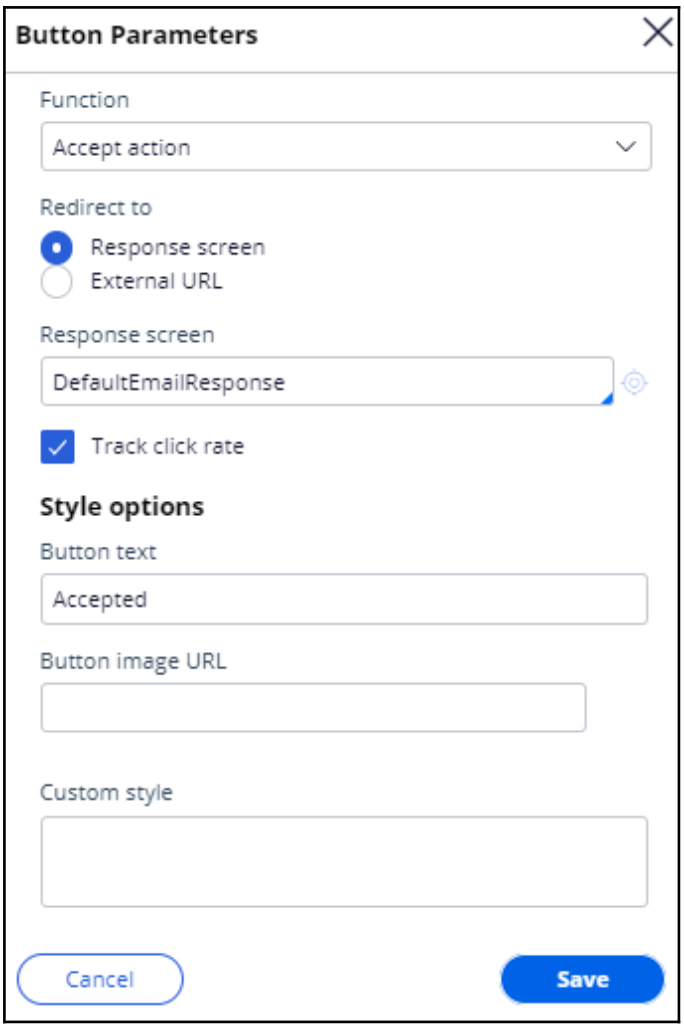

Now, select **Open external URL** and provide an external URL that opens when the customer clicks the button. Then, provide the button image URL.

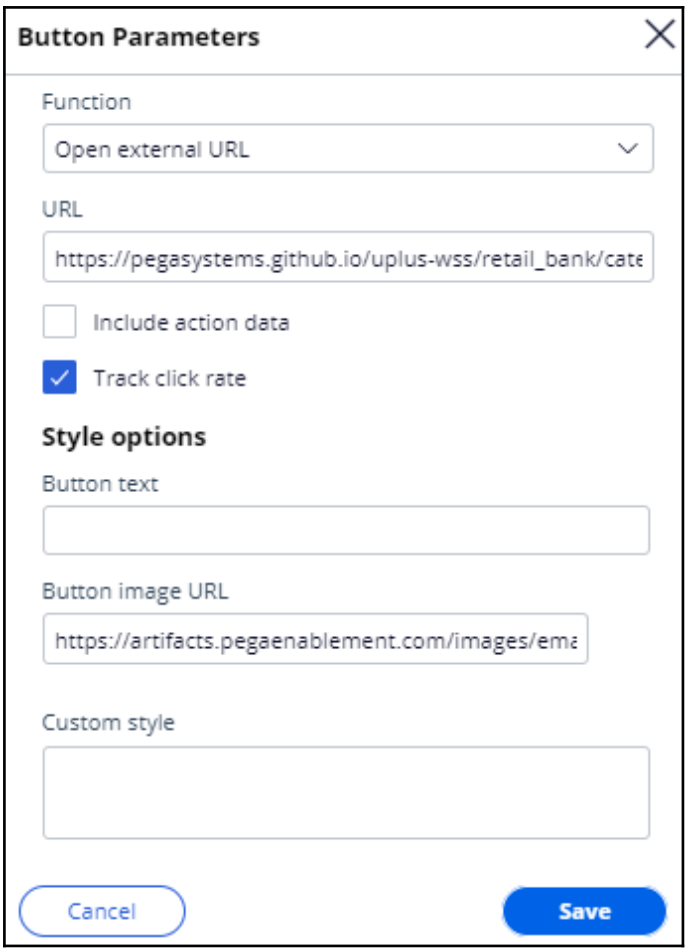

A composed email looks like this.

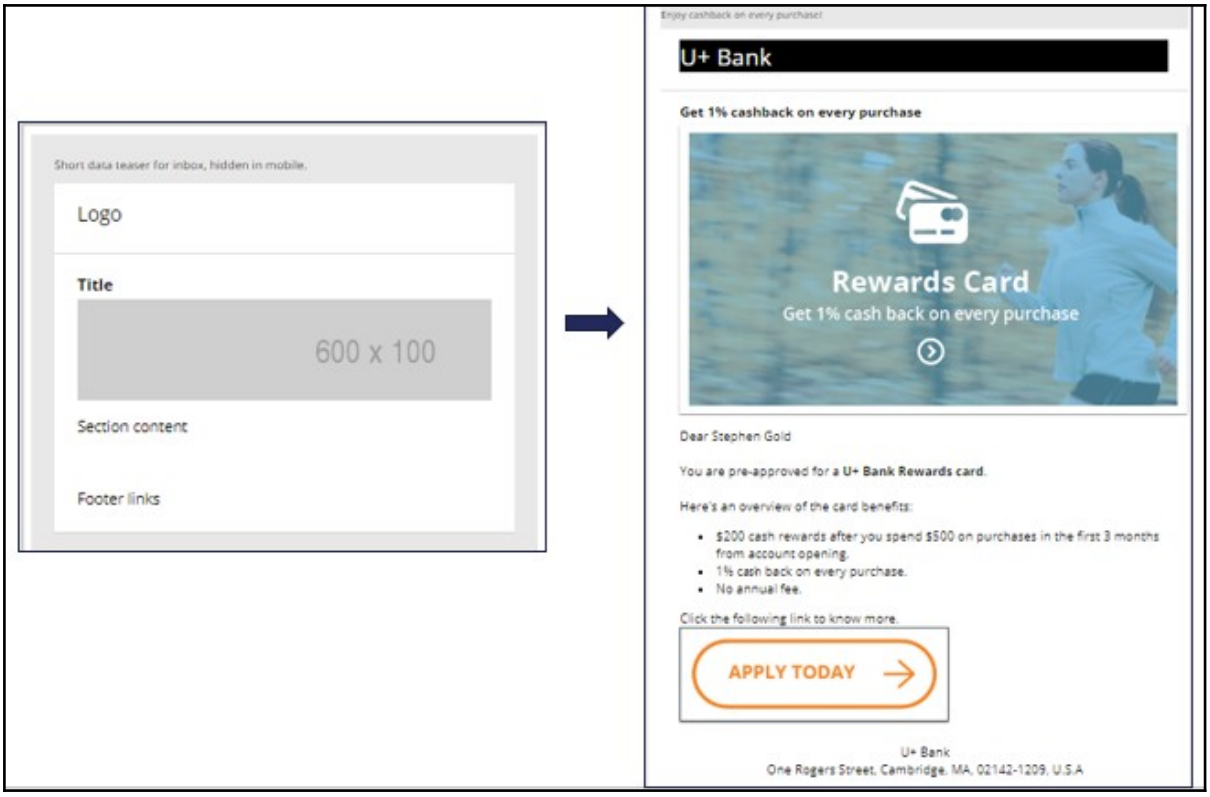

Now, save the treatment and add it to the Rewards card action.

You can send the email to test accounts or to yourself to test the visualization and rendering of the content on various email clients.

To test the treatment, switch to the **Test Message** tab. In this tab, you can specify the target email addresses and the email subject specifically for the test.

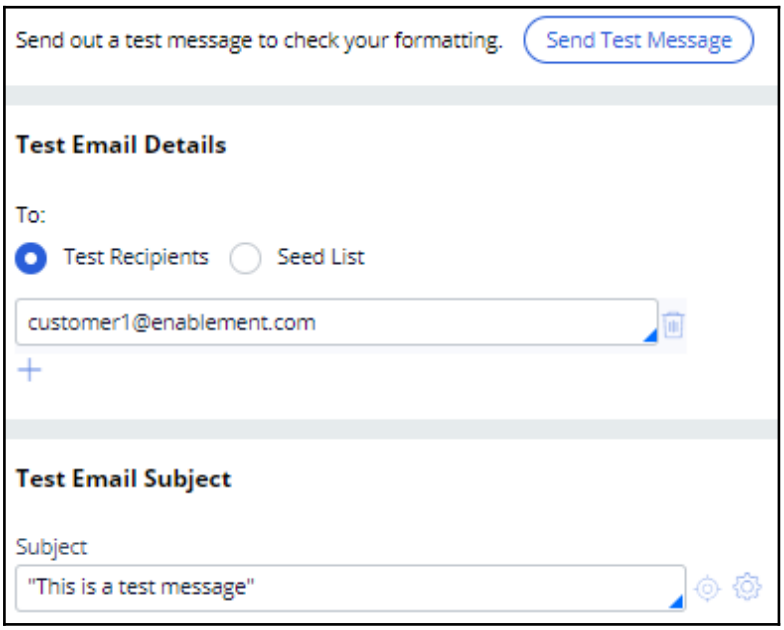

To test the treatment, click **Send Test Message**. Once the test is complete, you can check the emails in the email client.

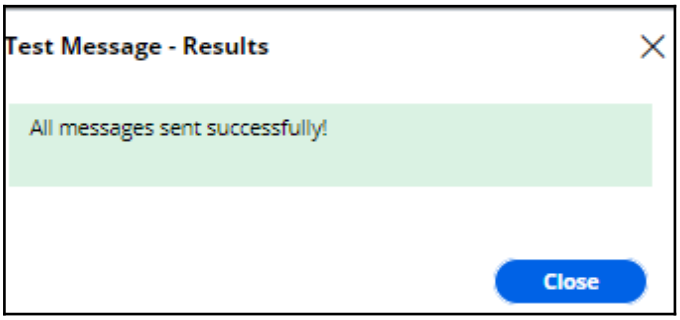

Note that none of the attribute placeholders are populated when you manually enter the recipients' email address. However, you can test the functionality of the buttons and links.

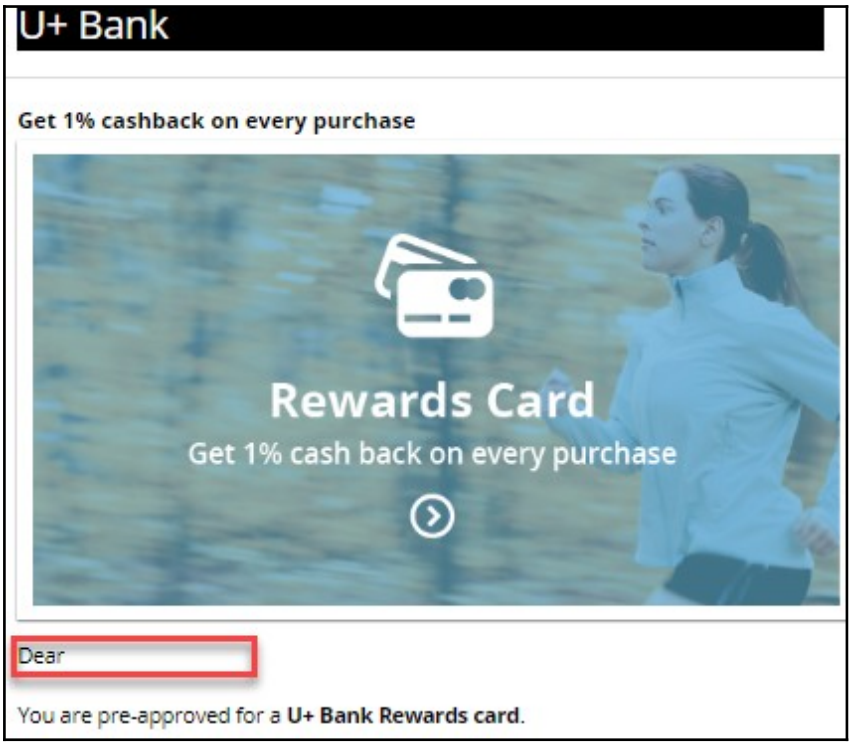

To populate the attribute placeholders, you can use a **seed list** as test input.

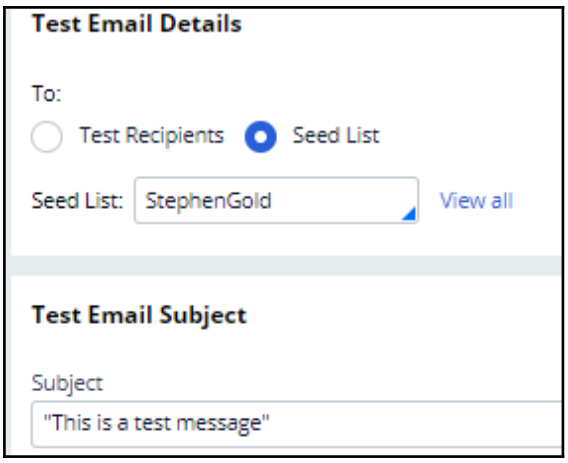

A seed list contains test customer data. It contains values for customer attributes, including email addresses.

When using a seed list to test the Treatment, the dynamic placeholders are replaced with customer information from the seed list.

Note that the customer name placeholder is replaced with the data from the seed list.

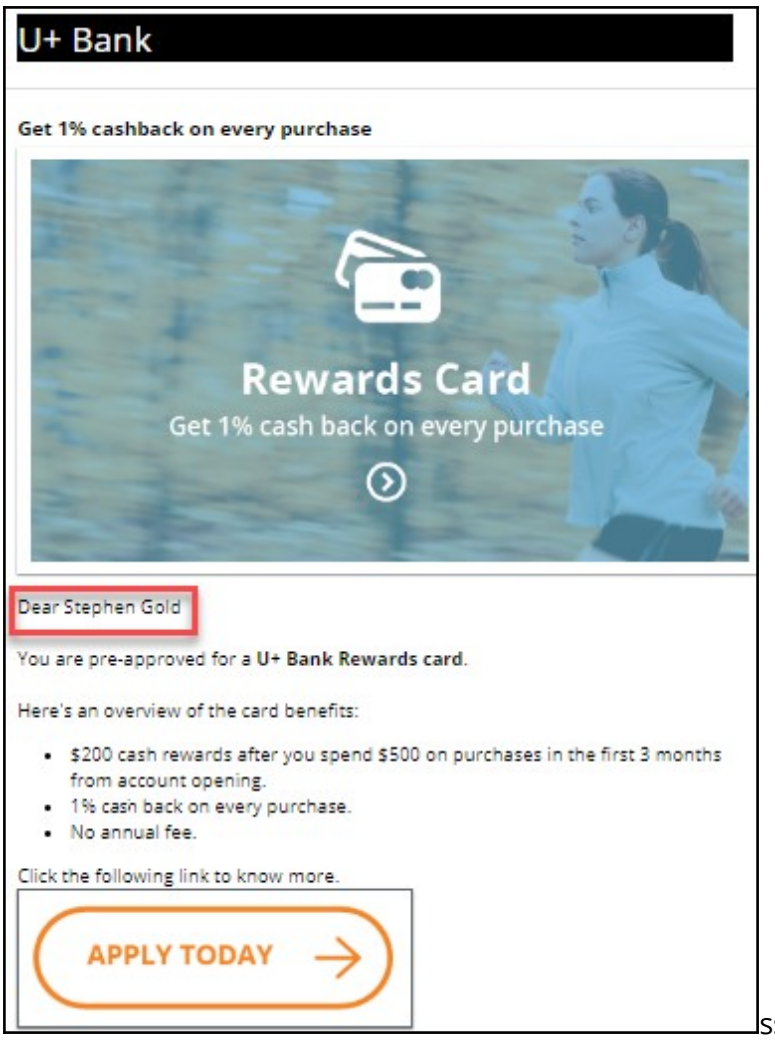

ss

This demo has concluded. What did it show you?

- How to add an email treatment to an action
- How to use dynamic placeholders to populate the customer and offer attributes in an email treatment
- How to use HTML code to generate content in an email treatment
- How to use an image from an external server in an email treatment
- How to add a button to an email treatment
- How to send emails to test accounts to test the visualization and rendering of the email content
- How to use seed lists to test the dynamic placeholders

Defining the action flow

Business scenario

U+, a retail bank, has designed a new email treatment to promote the Rewards Card offer. Now, an action flow must be created to ensure that the email is delivered when this action is selected for a customer.

Action Flow

The action flow is a flow diagram that contains the sequence of steps carried out during the lifecycle of an action. These steps include delivering the action message via an outbound channel, waiting for a response, updating the Interaction History, sending a follow-up message, and so on. In a flow, each of these steps is represented by a variety of shapes. Best practice is to keep the flow short and simple, preferably not more than 2 or 3 steps per channel. Every additional step taken in the action flow is a missed opportunity to let the Pega brain decide the Next-Best-Action.

A new action contains only the default **Start** and **End** shapes. To send emails in an always-on outbound environment, you need to add a **Send Email** shape to the action flow.

When there are several actions that need the same flow pattern, you can create that flow pattern once and reuse it for multiple actions using the **Dynamic template** option.

In the **Configure Dynamic Template** window, select the action that contains the desired flow pattern.

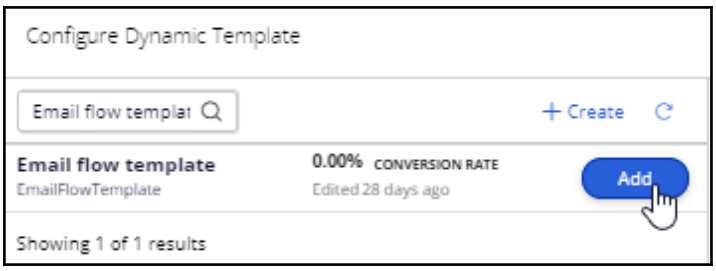

When you use the **Dynamic template** to configure the flow, you cannot change the flow in the referenced action. To edit the flow, you must edit the original action in which the flow was defined.

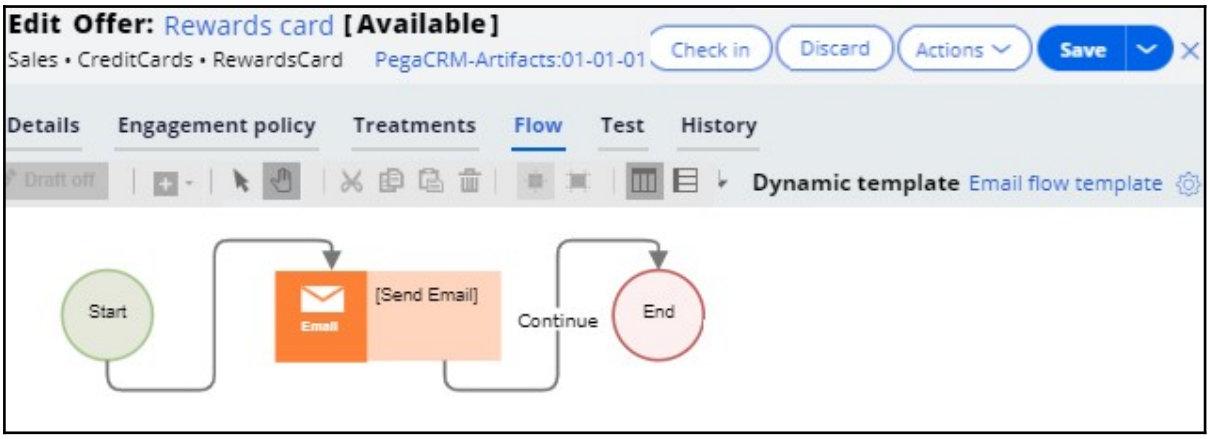

Basic email flow template

Let's look at a basic flow pattern for email delivery.

A flow template is just a normal action with the flow configured. In the following example, an action is used as a flow template within other actions to enable them to send emails. Thus, the **Availability** of the action is set to ***Never*** to prevent it from being presented to customers.

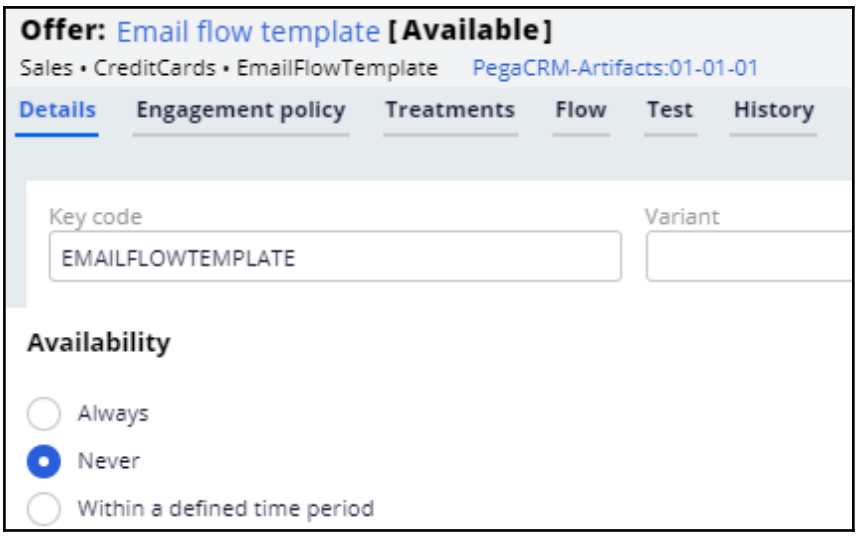

The **Send Email** shape in the flow typically references the email treatment that contains the email content to be sent.

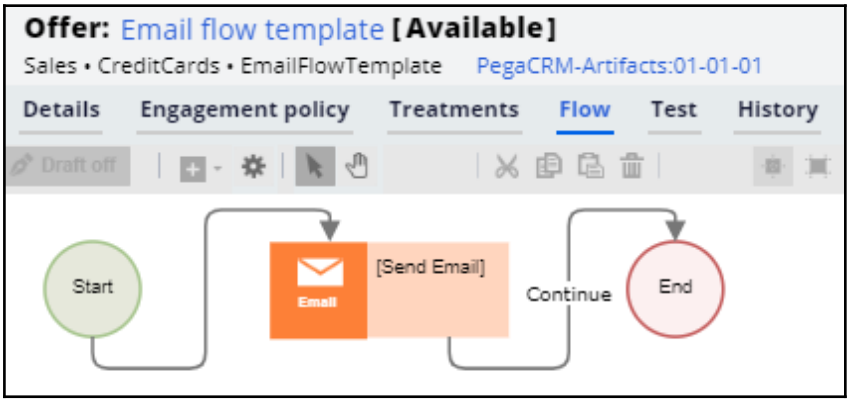

However, when you want the flow to be used as a template, you want the treatment belonging to the individual action to be used rather than a standard email treatment. In the **Send Email** shape, set the **Treatment** option to **Use treatment from strategy** to ensure that the treatment from each action, as selected by the Next-Best-Action decision, is used to send emails, instead of a standard treatment.

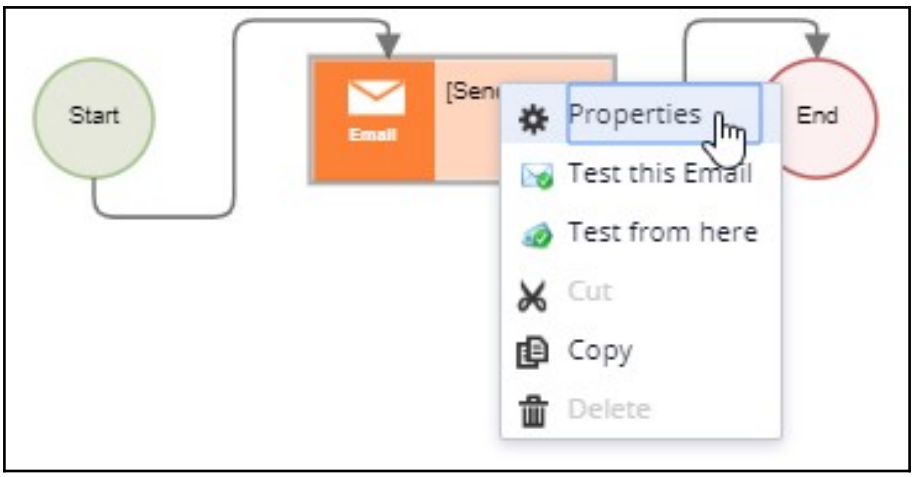

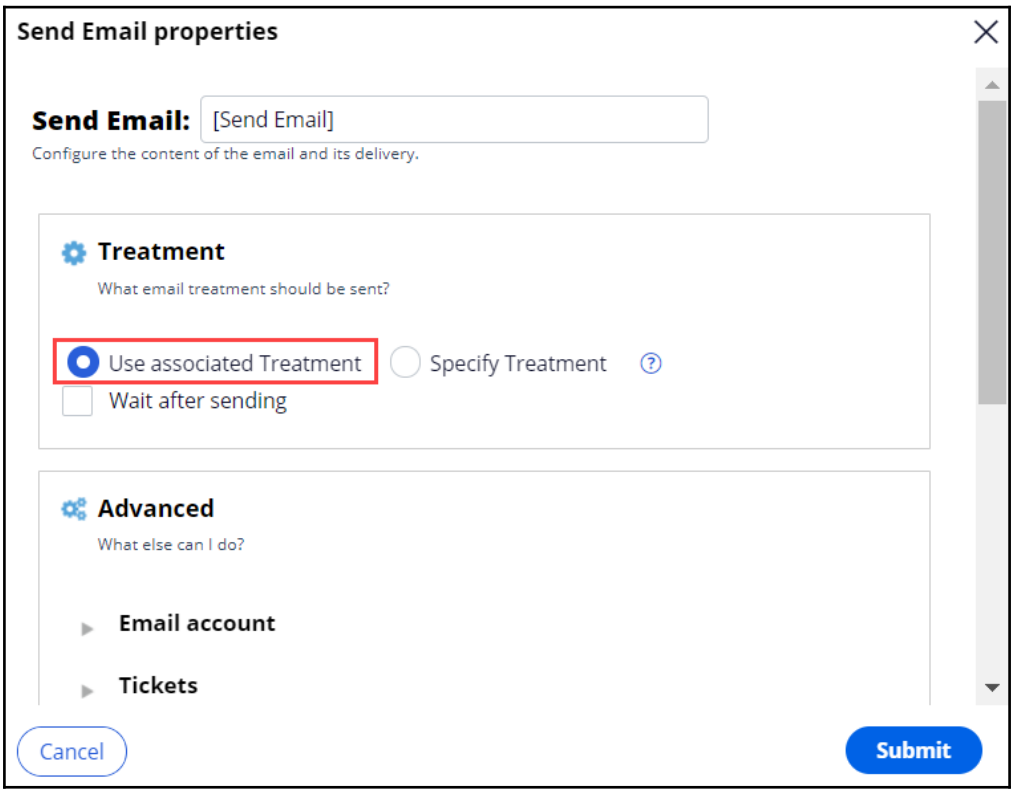

This ensures that the email treatment is independent of the flow template associated with the action. For example, when the Rewards Card uses the **Email flow template** as its **Dynamic flow** template, the Rewards Card email treatment will be used for sending the email.

Avoiding overexposure of actions on outbound

Description

Contacting a customer too frequently can have a negative impact on the customer's attitude toward further actions by your company. Learn how to define contact limits and use suppression policies to pause actions for a given time period based on past customer behavior.

Learning objectives

- List the types of contact policies
- Explain how to avoid overexposure of actions using contact limits and suppression policies

Contact policy types

Types of constraints

Too many contact attempts over a short period of time can have a negative impact on a customer's attitude toward further actions by your company. To maximize the lifetime value of every customer relationship, organizations must prevent outreach fatigue by optimizing the number of actions taken.

In Pega Customer Decision Hub™, there are two types of contact limits that allow you to limit actions over a given period of time.

1. Customer contact limits
2. Suppression policies

Customer contact limits

Customer contact limits allow you to limit the number of actions that a customer can receive over a given period of time on a specific channel. These customer contact limits prevent an action from reaching a customer on a specific channel, irrespective of past responses to that action by the customer.

For example, you can decide that you do not want your customers to receive more than two promotional emails per month.

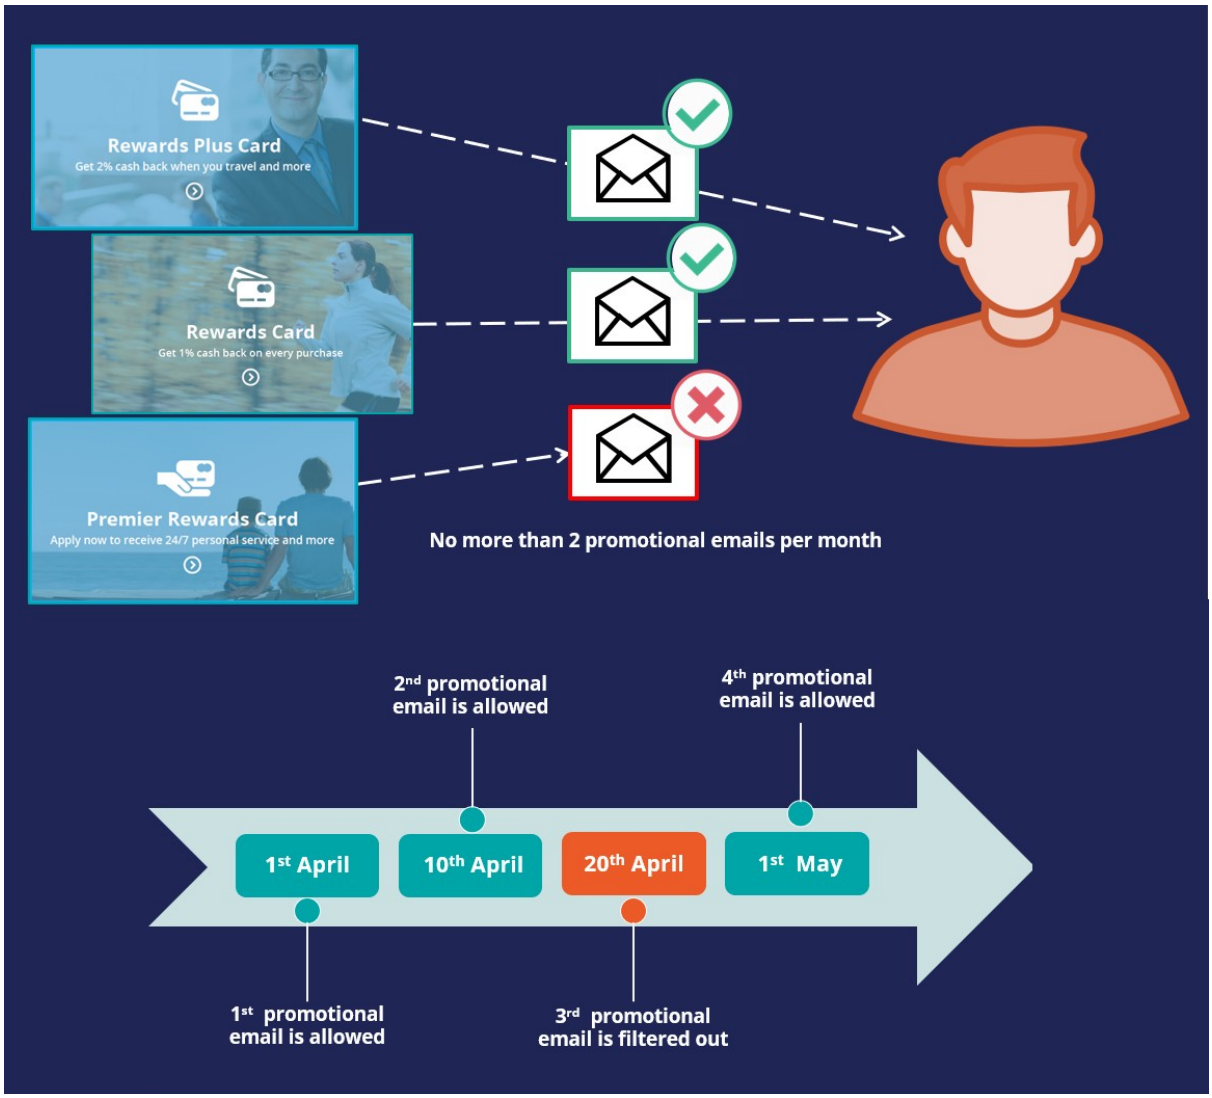

Suppression policies

Suppression policies determine when and for how long an action or group of actions should not be shown to a customer. These suppression policies put an action on hold after a specific number of outcomes are recorded for some or all channels.

Suppressing or pausing an action prevents oversaturation by limiting the number of times a customer is exposed to the same action.

For example, if a customer ignores a credit card promotional offer email twice in the previous 30 days, then the same credit card offer should not be sent to the customer again for the next 60 days.

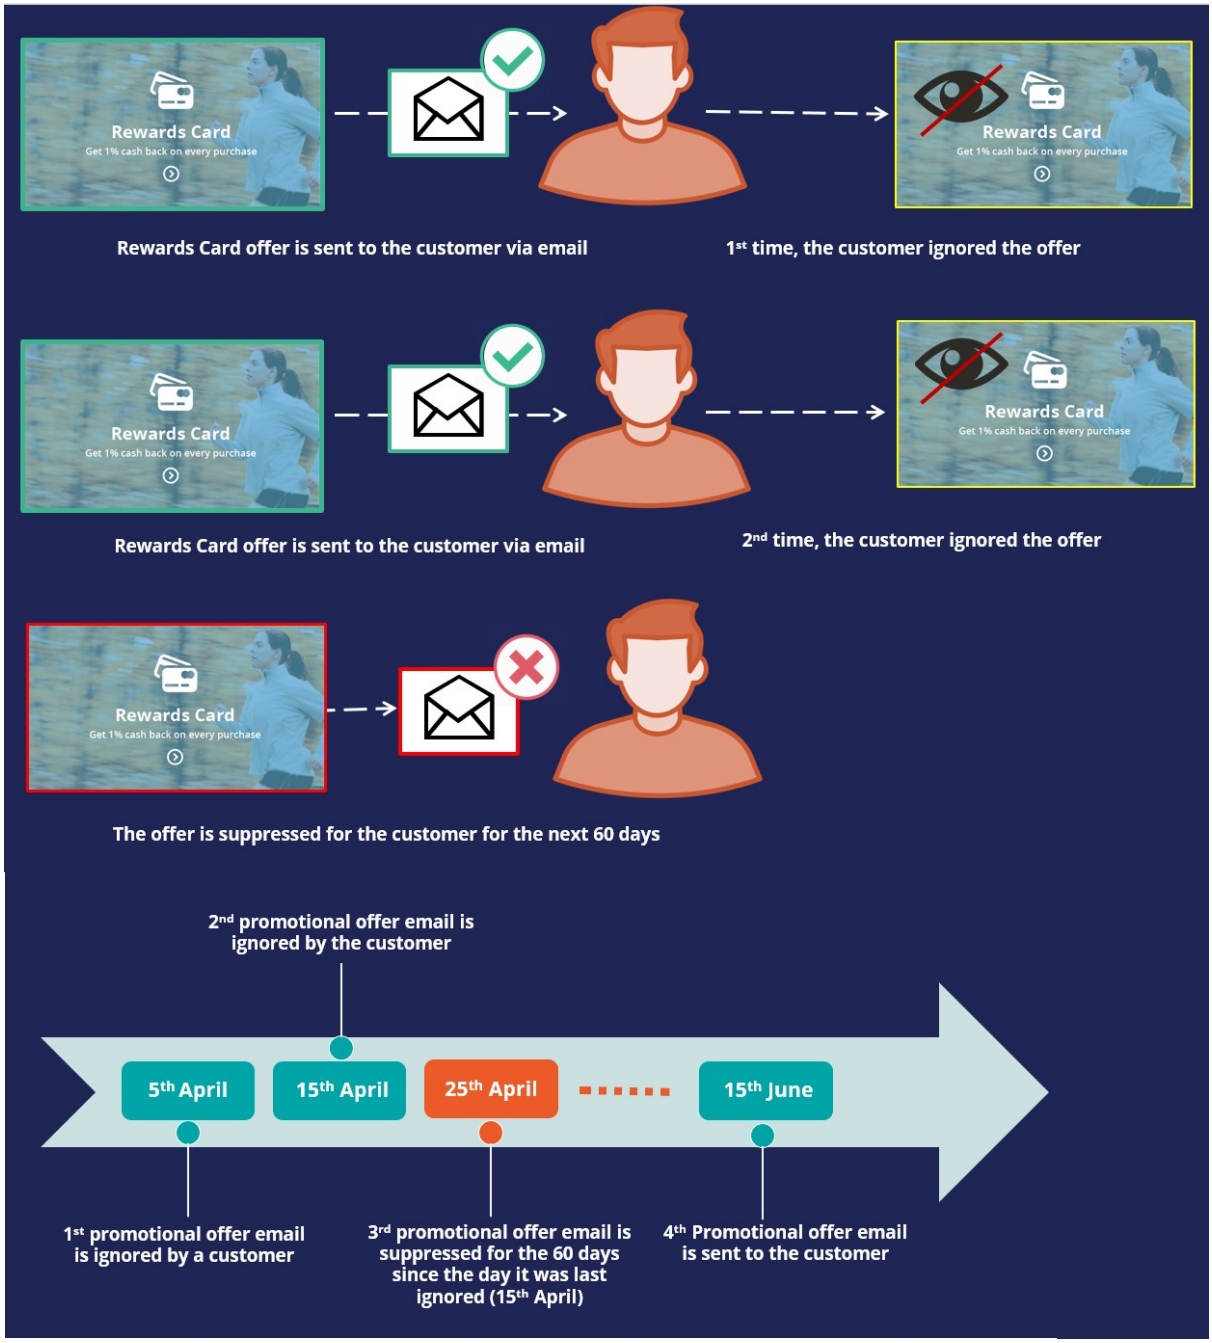
Rewards Card
Get 1% cash back on every purchase
Rewards Card offer is sent to the customer via email
1st time, the customer ignored the offer
Rewards Card
Get 1% cash back on every purchase
Rewards Card offer is sent to the customer via email
2nd time, the customer ignored the offer
Rewards Card
Get 1% cash back on every purchase
The offer is suppressed for the customer for the next 60 days
2nd promotional offer email is ignored by the customer
5th April
15th April
25th April
15th June
1st promotional offer email is ignored by a customer
3rd promotional offer email is suppressed for the 60 days since the day it was last ignored (15th April)
4th Promotional offer email is sent to the customer

Avoid sending the same action for a time period

Business scenario

U+ Bank's marketing department currently promotes various credit card offers by sending emails to qualified customers.

Now the bank wants to implement a new set of requirements: they want to limit the number of emails sent to their customers to two per month. Also, the bank wants to avoid sending a credit card offer via email for 60 days if one was sent in the past 30 days.

Updating constraints

On the Next-Best-Action Designer **Constraints** tab, you configure outbound channel limits and suppression policies.

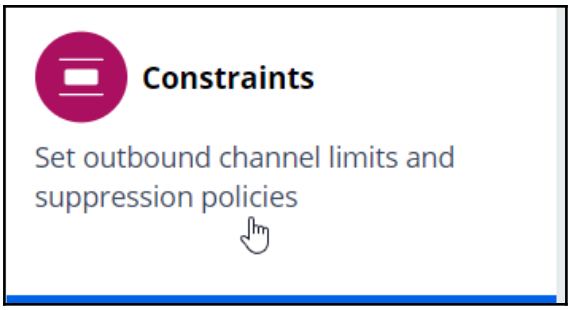

Edit **Constraints** to modify the limits.

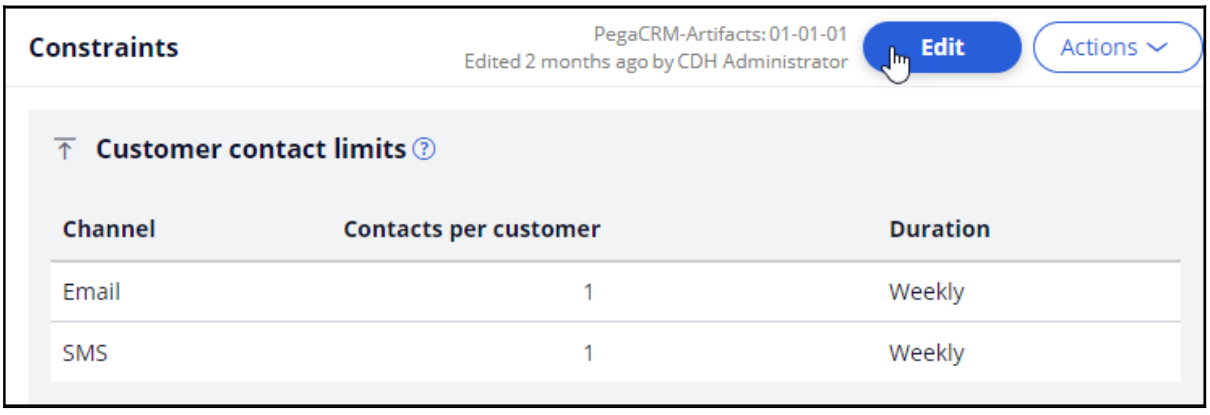

Updating customer contact limits

To limit the number of outbound interactions that a customer can receive over a given period of time on a specific channel, you need to set customer contact limits.

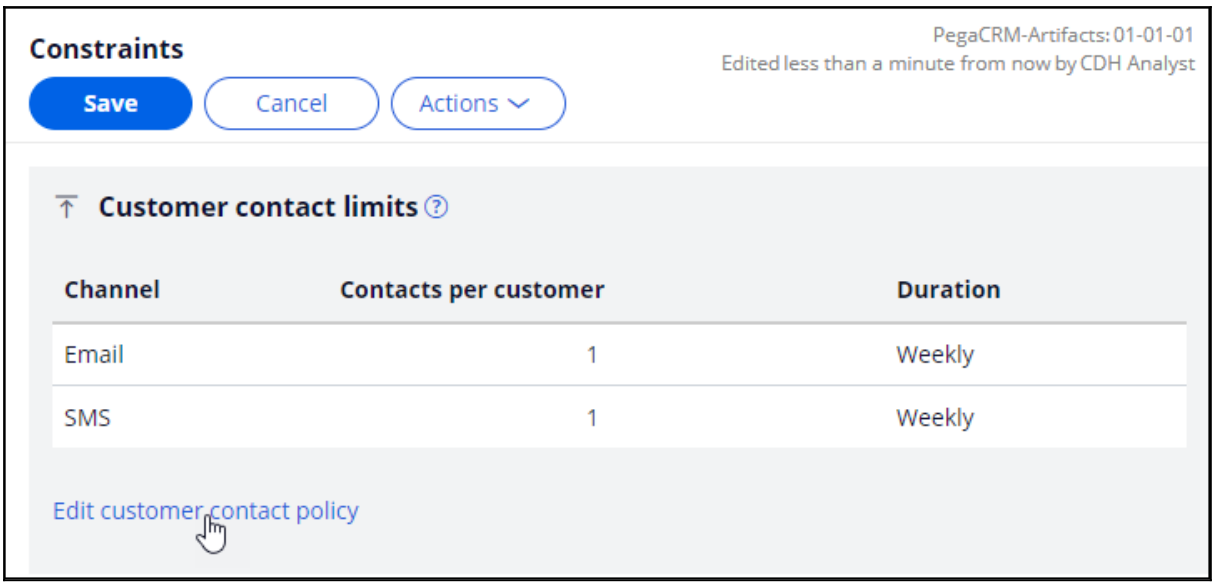

In this scenario, the bank wants their customers to receive no more than two emails per month. So, you need to update the existing customer contact policy for the email channel to two emails per month.

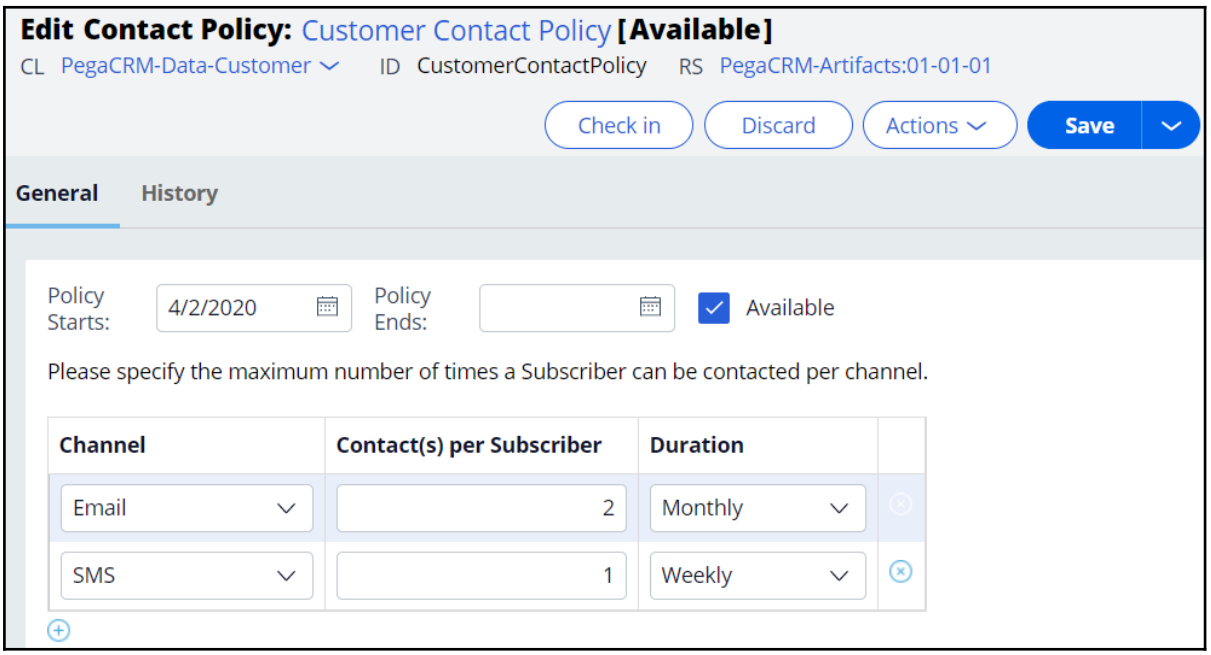

In the contact policy, you can update current limits or add new channel limits.

Note: You cannot repeat a channel in a contact policy.

Defining suppression policies

To suppress a single action or a group of actions for a limited time period, you need to define suppression policies.

In this scenario, the bank wants to avoid sending promotional emails related to a credit card offer for 60 days if one was sent in the past 30 days. So, you need to add a contact policy to put an action on hold for 60 days.

When an email is sent, a record is added to the Interaction History with the outcome value ***Pending***.

Provide a name for the new contact policy and then select the type of outcome it will track; in this case it will track outcomes with the value ***Pending***.

You can specify whether the responses are tracked for one specific action, or for all actions in the group. Track the responses at the action level, since you want to suppress only credit card promotions in email.

You can select the time period over which the responses should be tracked. In this case, responses should be tracked for the past 30 days.

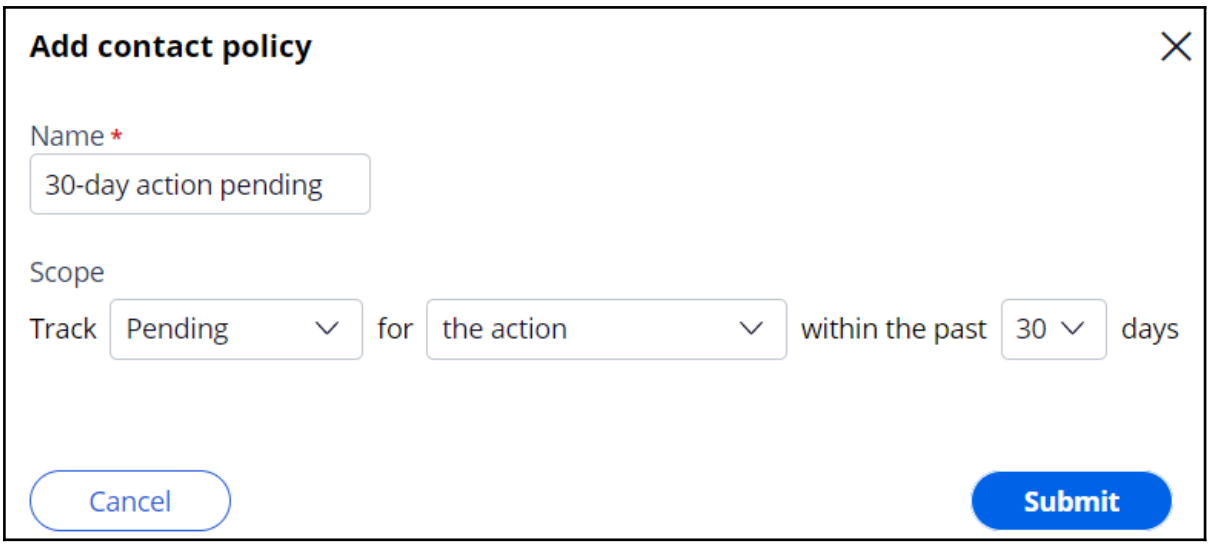

The business requirement is to suppress the action for 60 days if an email was sent in the past 30 days. In other words, if there is even one pending response for the email treatment over the past 30-day period, the action must be suppressed. So, you need to enter the number of responses required to fulfill the suppression criteria as 1.

Then you need to select the channel for which the responses are tracked and enter the number of days for which an action should be paused when the suppression criteria are met.

Associating a contact policy with an issue and a group

On the **Engagement policy** tab, you can associate contact policies with an issue and group. The contact policies are reusable across all issues and groups.

As the bank wants to suppress promotional emails related to credit card offers, you need to edit the **CreditCards** group to add the newly defined contact policy to it.

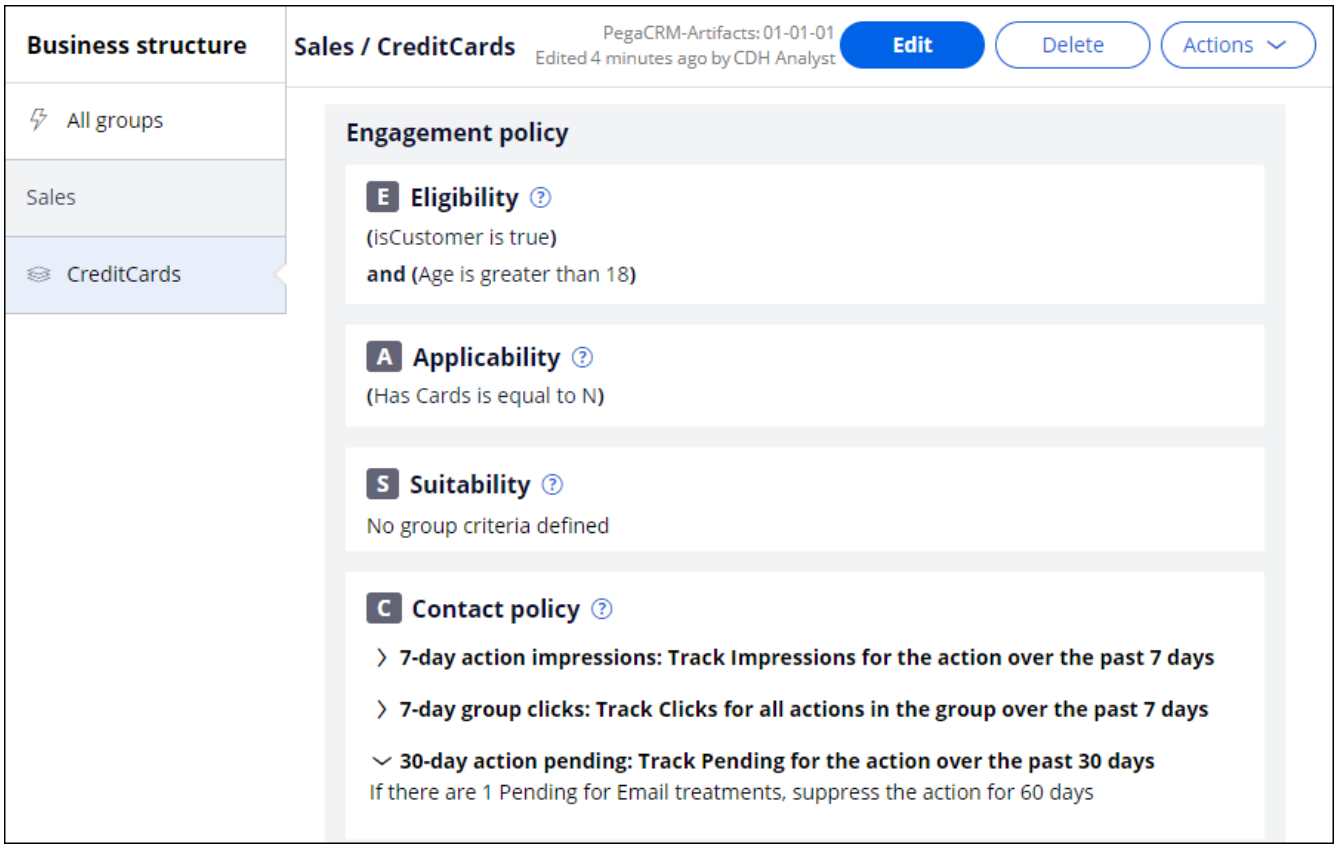
Business structure
All groups
Sales
CreditCards
Sales / CreditCards
PegaCRM-Artifacts: 01-01-01
Edited 4 minutes ago by CDH Analyst
Edit
Delete
Actions
Engagement policy
E Eligibility
(isCustomer is true)
and (Age is greater than 18)
A Applicability
(Has Cards is equal to N)
S Suitability
No group criteria defined
C Contact policy
7-day action impressions: Track Impressions for the action over the past 7 days
7-day group clicks: Track Clicks for all actions in the group over the past 7 days
30-day action pending: Track Pending for the action over the past 30 days
If there are 1 Pending for Email treatments, suppress the action for 60 days

Limiting action volume on outbound

Description

In outbound customer engagement, you need to cater to different kinds of operational limitations. These may include limits on the number of email communications that can be sent in a day, or the total number of communications that can be sent regarding a particular offer. Learn how to define volume constrains to limit the volume of actions in an outbound run.

Learning objectives

- Limit outbound action volumes using volume constraints

Volume constraints

What is volume constraint?

In Pega Customer Decision Hub™ you can limit the volume of actions delivered to customers via always-on outbound.

Volume constraints can be set at an action, channel, and/or property level.

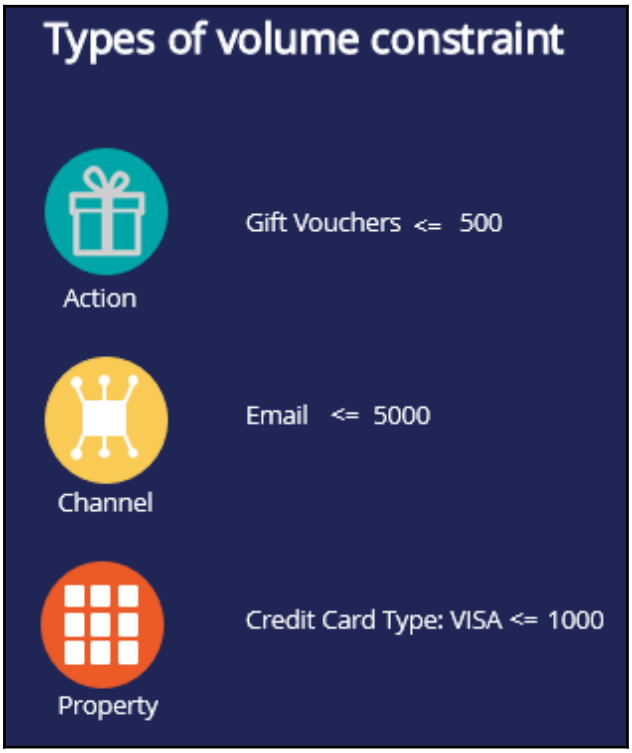

Action constraint: It limits the volume of actions that can be sent to customers.

An action constraint is useful in a situation where you have a limited number of offers that you can send to customers, for example, 500 gift vouchers. Configuring an action constraint would ensure that you do not offer more offers than are available.

Channel constraint: It limits the volume of actions that can be delivered via a certain channel.

You can use channel constraint for channel capacity planning. For example, assume that you have a limited call handling capacity in your inbound call center, for example 1000 calls/day. When you send 50,000 emails on a day on a certain regulatory update, let's say you except 2%-3% follow-up calls a day from your customers to your contact center, which

is about 1000-1500 calls. So, you need to limit the outbound email volume to 35,000 emails a day to match the inbound call center capacity.

Property constraint: By default, actions are grouped under the issue / group business hierarchy. Using a property constraint, you can limit the volume of a set of actions across different issues and groups by grouping them using a property value.
For example, you can group all credit card offers by setting a property, CreditCardType, to VISA. Then you can define a constraint to limit the action volume where CreditCardType is VISA.

Volume constraint modes

Within a volume constraint, the combination of constraints can be defined for one or more actions with one or more channel types.

In situations where a customer is eligible for multiple actions, volume constraint can be applied in two different ways depending on the business requirement. That is, you may specify whether volume constraints should be applied to the actions as a group, or whether each action should be evaluated separately.

When volume constraint is evaluated separately, each action is checked for the customer separately against applicable channel and action constraints. A customer receives only one action, the one with the top priority, even though the limits for the other low priority actions, are not reached. The top action is selected based on factors like propensity, action value and business levers defined in Next-Best-Action Designer Arbitration.

For example, let's consider that a customer is eligible for Standard card, Rewards card, and Rewards Plus card credit card actions and Standard card is the top priority card. Also,

assume that the remaining volume for all three actions are above zero. When you set volume constraints to check each action, the customer will be presented only with the top priority action, Standard card.
When volume constraint is applied for actions as a group, the highest priority action is selected to apply constraints on the channels and actions, without removing the other actions for a customer. That is, a customer receives all the actions he/she qualifies for, provided that the limit for the top priority action is not reached. The other low priority actions are selected even though the limits for those actions are reached.

For example, let's consider that a customer is eligible for Standard card, Rewards card, and Rewards Plus card actions and Standard card is the top priority action. Also, the remaining volume for Standard card is above zero, but for the other two actions it has reached zero. When you set volume constraints to check actions as a group, the customer receives all the actions.

Evaluating actions as a group allows Pega Customer Decision Hub to present all eligible actions to a customer.

Configuring a volume constraint

Business scenario

U+ bank has been running always-on outbound and now they would like to apply some limitations on the outbound volume. They have the following three different kinds of operational constraints to implement.

- The IT department has enforced a limit of 1000 emails per week to avoid being classified as spammers. Hence, the marketing team wants to adhere to the limits.
- The bank would like to restrict the number of Standard card offer emails to only 100 per week due to a business requirement.
- The bank would also like to limit offer emails related to credit cards of type “Master” to only 400 per week due to a logistical issue with their fulfillment vendor.

Creating a volume constraint

A constraint is an optional configuration element of always-on outbound. When configured, it will limit the volume of actions that will ultimately be initiated.

To limit action volume on the email channel, you need to configure a volume constraint for the always-on outbound.

The volume constraint landing page can be accessed from **Intelligence** -> **Volume Constraints** from the navigation menu.

You can add a new volume constraint from this page.

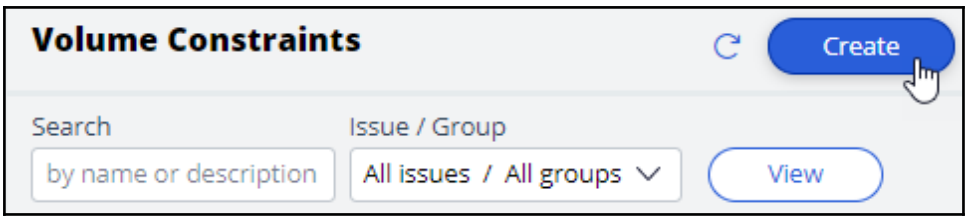

Provide a short description for the new volume constraint and the identifier is automatically populated when you tab out.

Note: The volume constraint must be created at the top level for it to be used in the Next-Best-Action Designer outbound schedule.

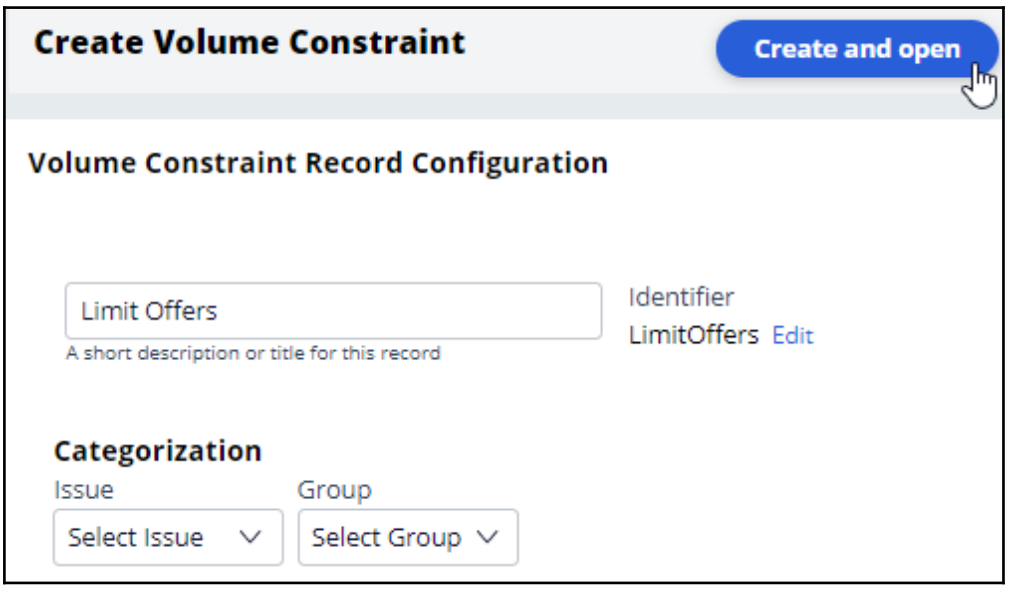

Configuring a volume constraint

The **General** tab is where you configure the different types of volume constraints.

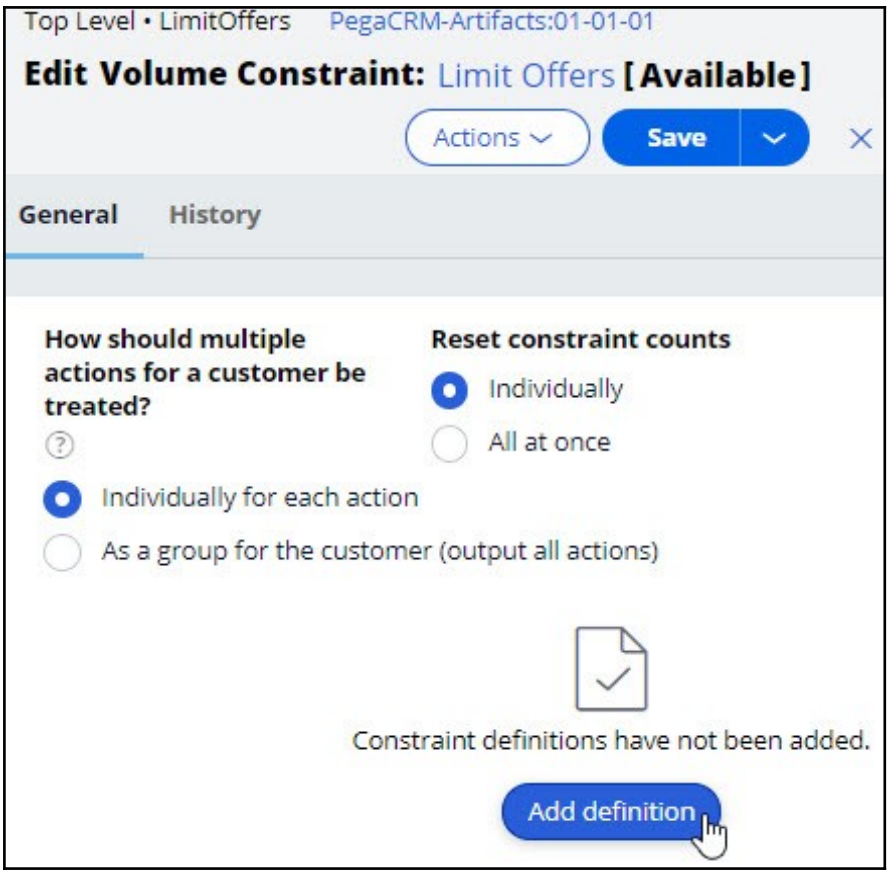

In this scenario the bank wants to implement three types of volume constraints. As a first step, add a channel constraint that limits the number of outbound messages to 1000/week on the email channel.

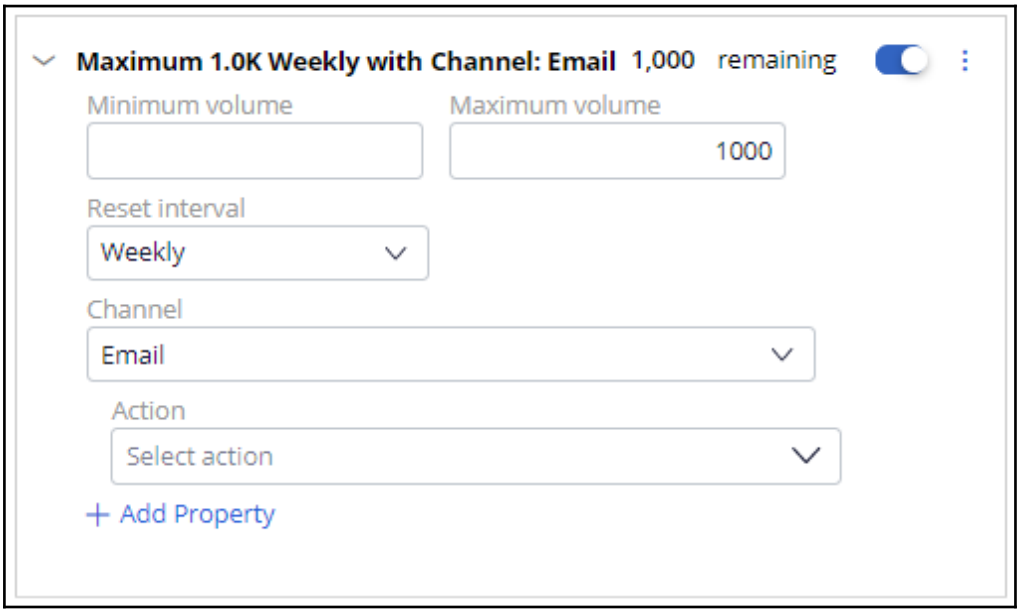

The **Rest interval** option determines how frequently the volume limit is reset to the full quota. There are a number of reset interval options available, including: *Manually*, which will allow you to reset the constraint limits manually using the **Reset** button.

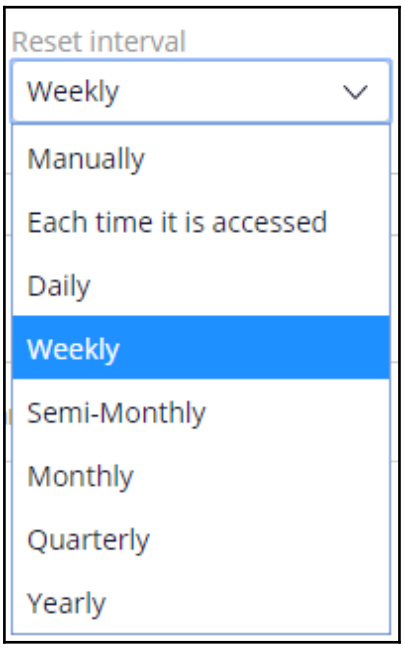

The other options in this list can be used when you want the system to automatically reset the constraint. For example, *Weekly* means that the system will reset the limits the first time they are applied in a particular week, Sunday being considered the first day of the week.

The remaining quota of the volume constraints, if any, in a given reset interval will be used in the following scheduled run. For example, if only 300 Rewards card emails can be sent in a week, the scheduled run might send 250 emails based on action prioritization. The remaining 50 emails will be sent in the same week during the next scheduled run.

In this case, configure the **Reset interval** as *Weekly* for both channel and action constraints.

Next, add an action constraint that will restrict the Standard card offer emails to only 100 per week.

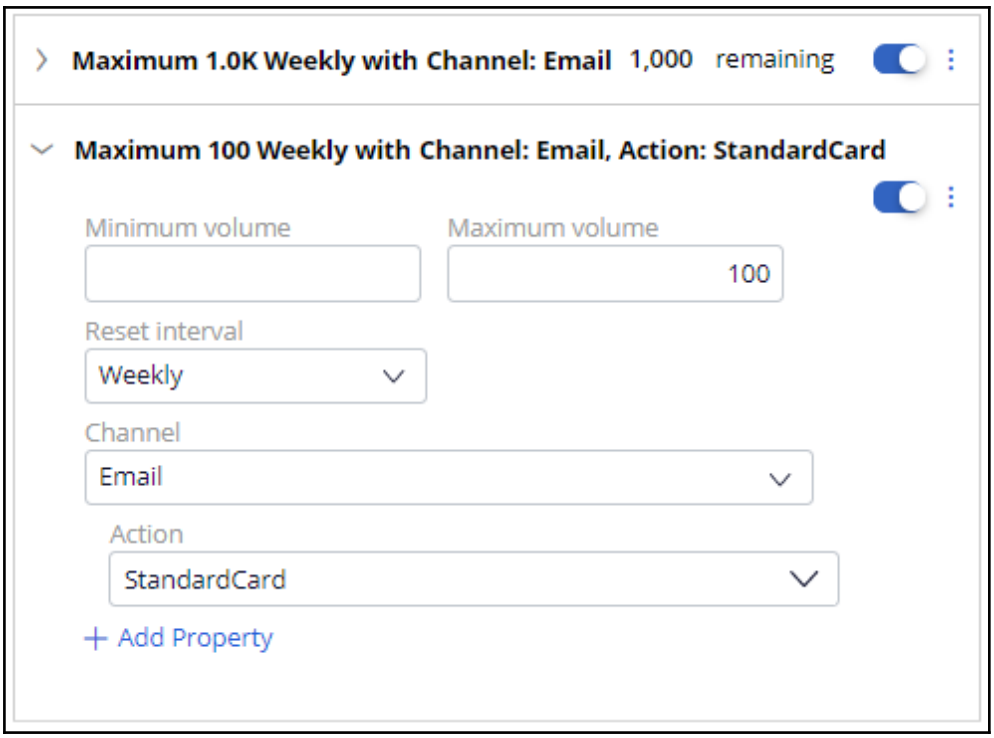

Finally, as per the requirements, add a property constraint that will limit offer emails related to credit cards of type “Master” to only 400 per week.

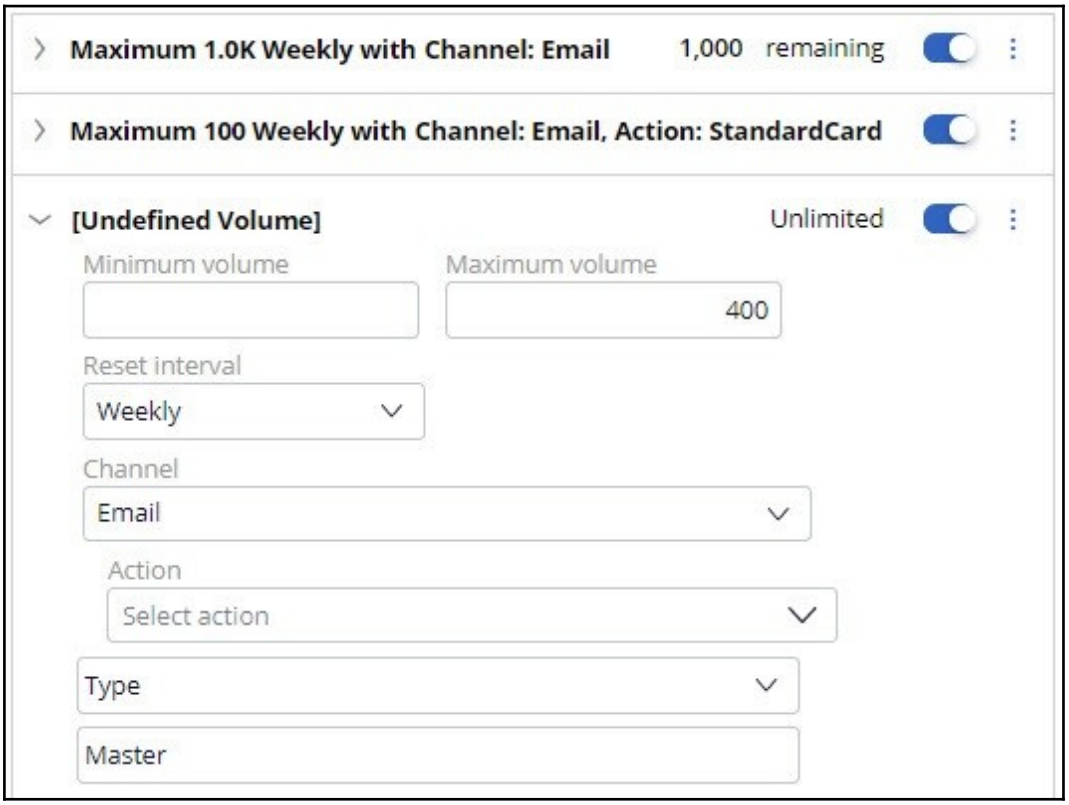

Configuring volume constraint in the Next-Best-Action Designer

To implement this business requirement, you need to configure the Next-Best-Action Designer to perform always-on outbound customer engagement with the volume constraint configured.

On the **Channels** tab, configure the **Primary Schedule** to achieve this.

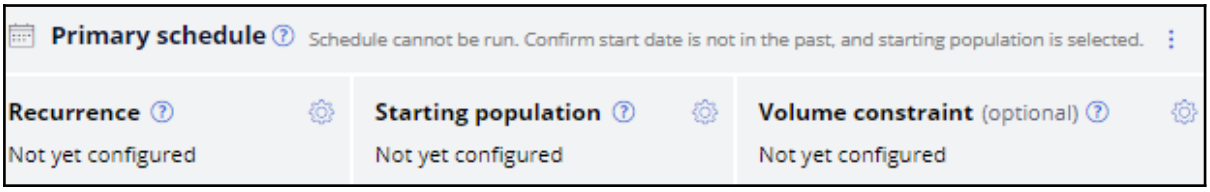

The schedule has two mandatory configurations, **Recurrence** and **Starting population**. To include the volume limits on outbound, you must configure the **Volume constraint**.

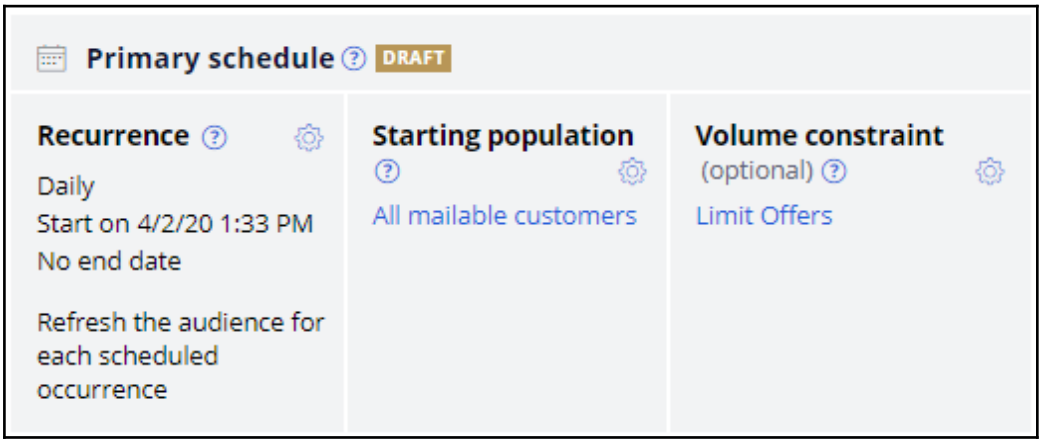

With that, all the necessary configurations for this scenario are complete.

Run the schedule. Once the primary schedule has run, you can check the remaining volume limits for the constraints. The remaining quota of the volume constraints, if any, in a given interval will be used in the following scheduled run.

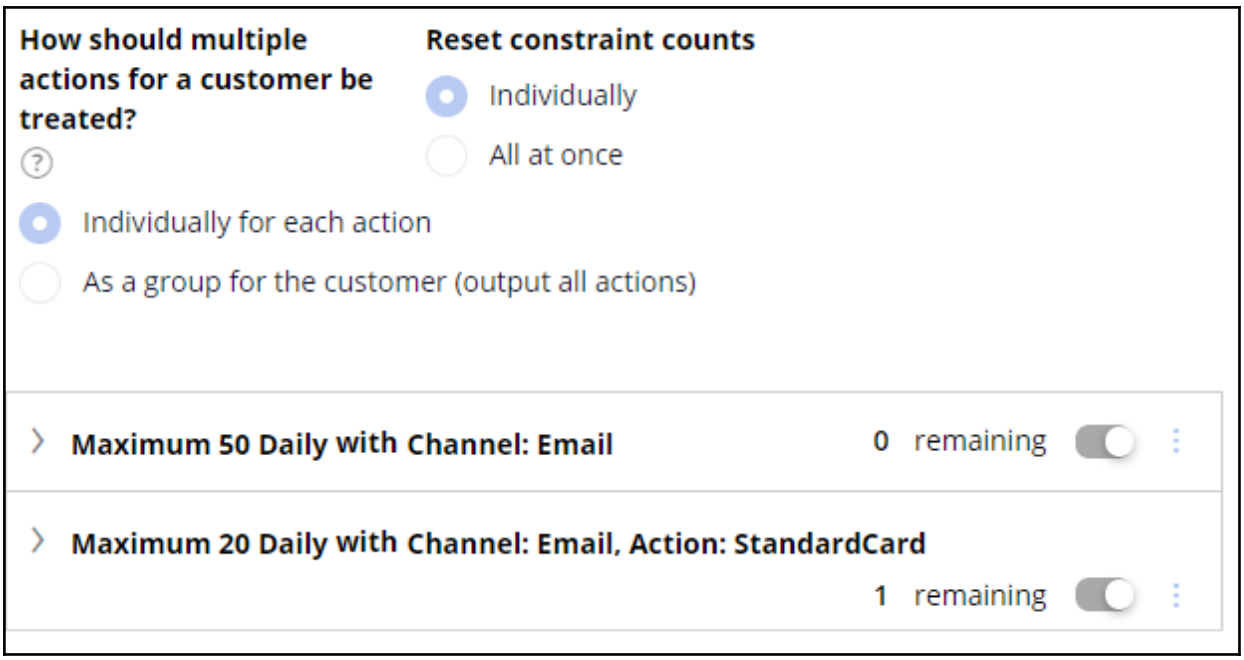
How should multiple actions for a customer be treated?
Reset constraint counts
Individually
All at once
Individually for each action
As a group for the customer (output all actions)
Maximum 50 Daily with Channel: Email
0 remaining
Maximum 20 Daily with Channel: Email, Action: StandardCard
1 remaining

Sharing action details with third-party distributors

Description

Sending outbound communications via email and SMS has its challenges, including various operational, security and legal requirements. Sometimes, it's easier to let an external email service provider handle message delivery for you. Learn how to share action details with an external service provider using files and database templates.

Learning objectives

- Explain outbound file and database templates
- Write action details to files or database tables to send to third-party email service vendors

File and database templates

Why do you need outbound templates?

In outbound customer engagement, Pega Customer Decision Hub sends action messages to customers directly via email or SMS. That is, Pega Customer Decision Hub directly communicates with the email or SMS servers to send out outbound messages to qualified customers.

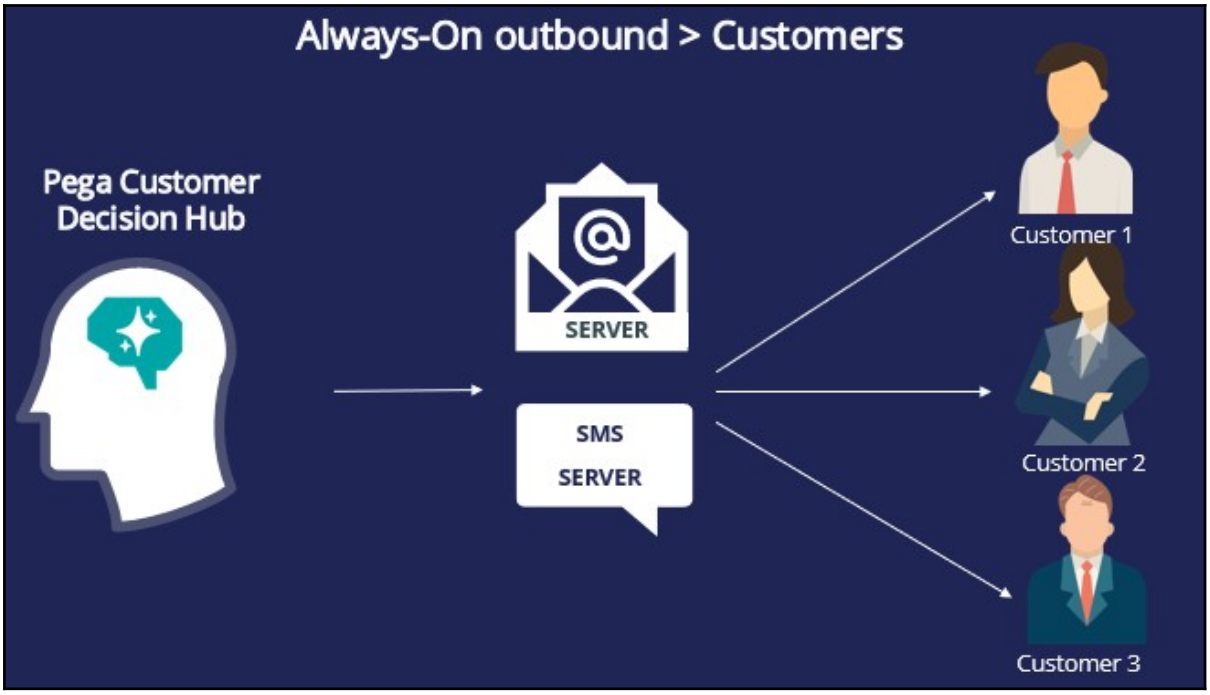

However, sometimes the necessary infrastructure to send action messages directly to customers may not be available. In such cases, you need to use a third-party email/SMS distributor. You need a medium to transfer the action and customer details, to a third-party distributor, which is typically a file or a database table.

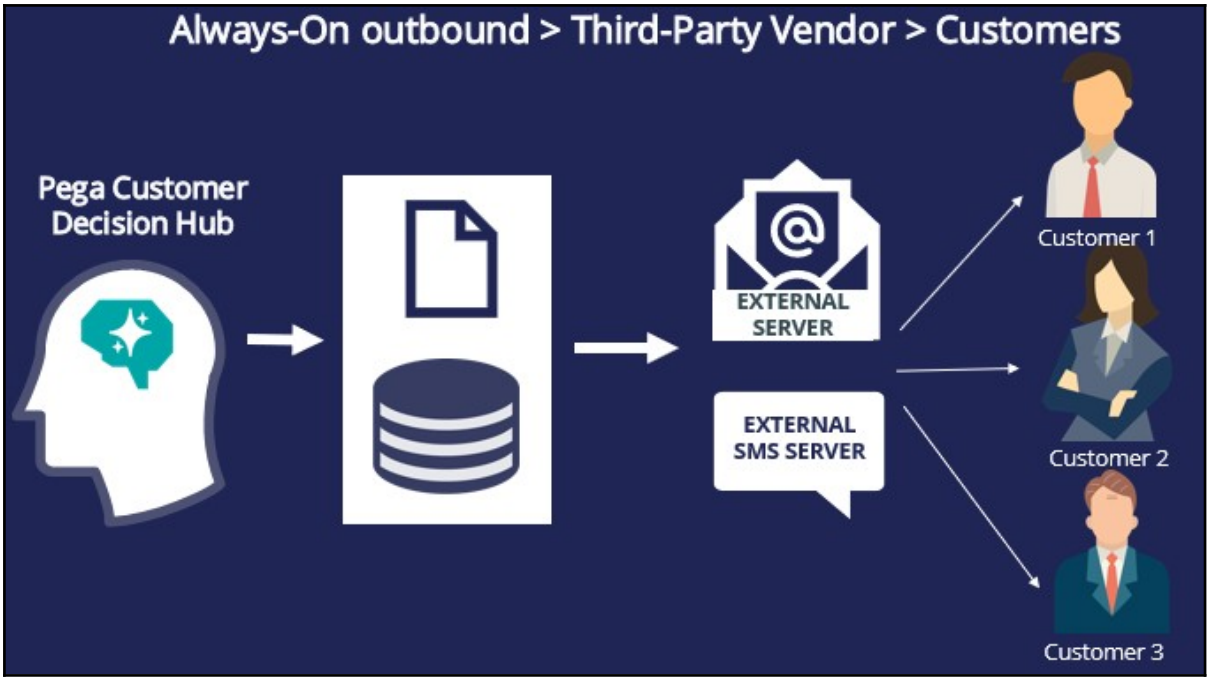

Pega Customer Decision Hub enables you to write the file in one of the following cloud-based destinations as well as the local filesystem.

- JFrog Artifactory
- Amazon S3
- Microsoft Azure

Writing action details to a file or database

Pega Customer Decision Hub enables you to easily write action details to a file on the local system or on a cloud-based destination such as Amazon S3 or to a database table, which can be sent to a third-party vendor or application for distribution. To achieve this, you need to configure an action flow to include a Send shape with the file or database template option enabled.

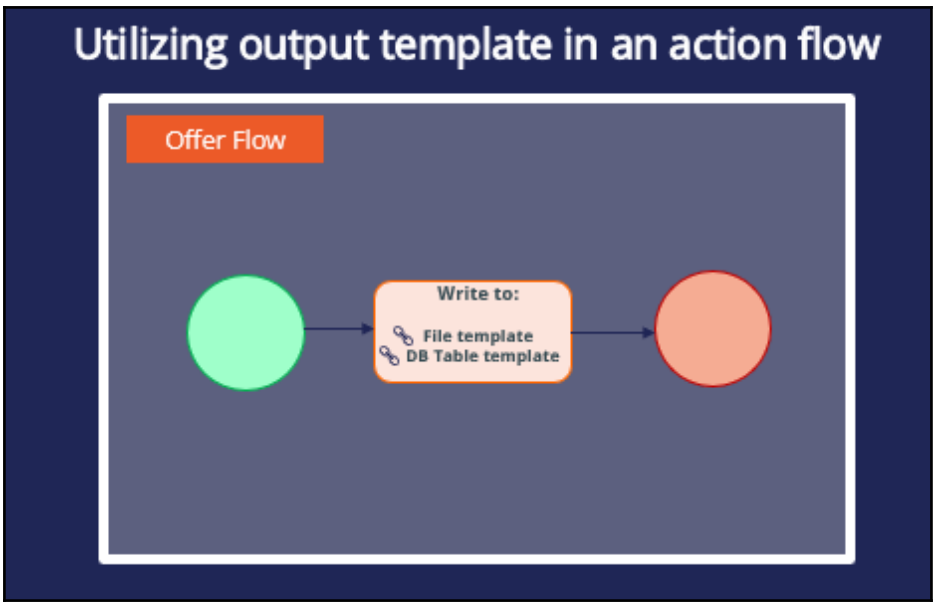

The four available Send shapes that can be used for this purpose are: **Send Email**, **Send SMS**, **Send Multi,** and **Send Generic**. Any one of these shapes can be used to write action details to a file or a database table. However, when you know that the action messages will be sent eventually via Email or SMS by the third-party vendor, it is better to use the corresponding shape, as this will clearly illustrate the purpose of the component and the channel will be recorded as an Email or SMS in the Interaction History.

If you are unsure of the communication channel, or if it is a channel that Pega Customer Decision Hub cannot directly integrate with, use the Send Generic shape.

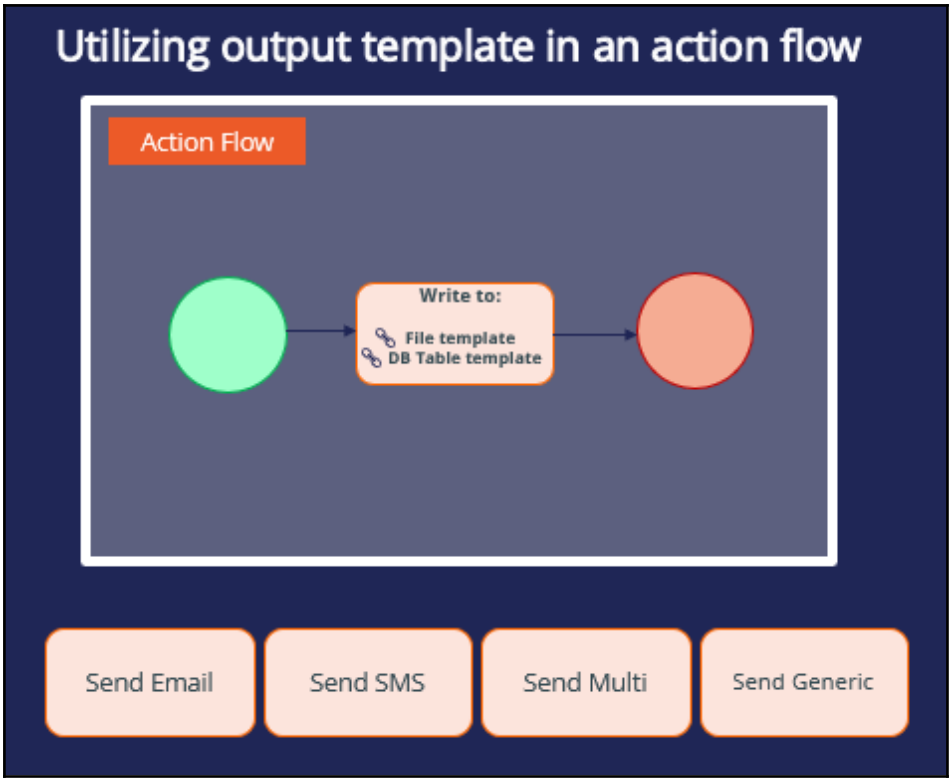

Multiple action details in a single location

Details of multiple actions can be written to the same file or database table. To do this, you need to configure the same file/database table template in multiple action flows. This functionality is useful when *the* ***Next-Best-Action outbound schedule*** *includes multiple actions. This approach consolidates multiple action details into a single location.*

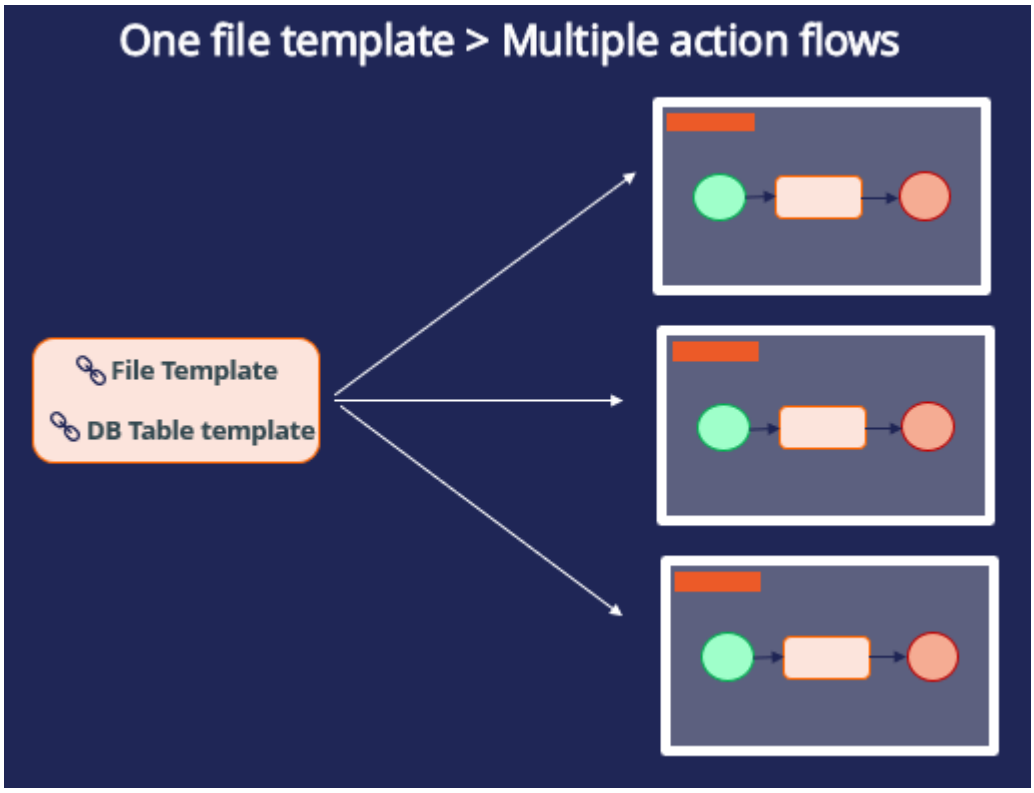
One file template > Multiple action flows
File Template
DB Table template

Writing action details to a file or database

Business scenario

U+ Bank, a retail bank, sends credit card offers to eligible customers by email. However, to comply with the security and spamming regulations, the bank wants to use a third-party email delivery service instead of using the in-house email server. As a result, the offer details need to be written to a file that can be shared with a third-party email distributor.

Creating a file template

Pega Customer Decision Hub™ enables you to easily write action details to a file that can be sent to a third-party vendor or application for distribution.

To implement this business requirement, you need to configure a file template. The file template landing page can be accessed from the navigation menu, **Content** -> **Treatments**.

You can create a new file template from this page.

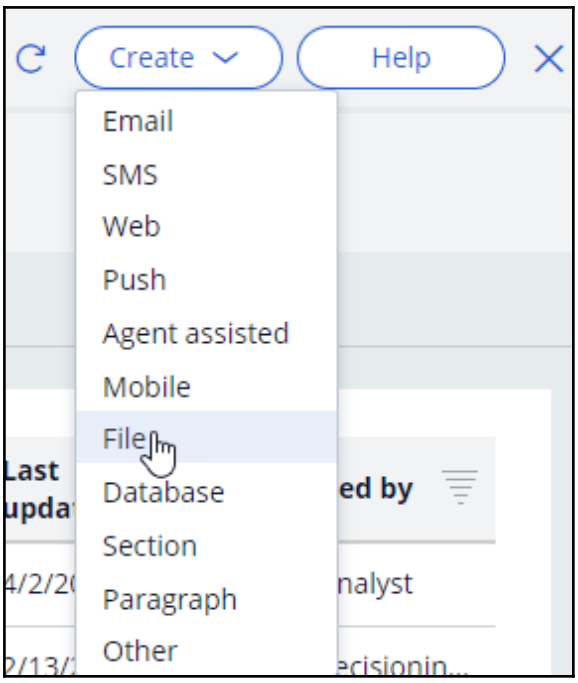

Provide a short description for the new file template and the identifier is automatically populated when you tab out.

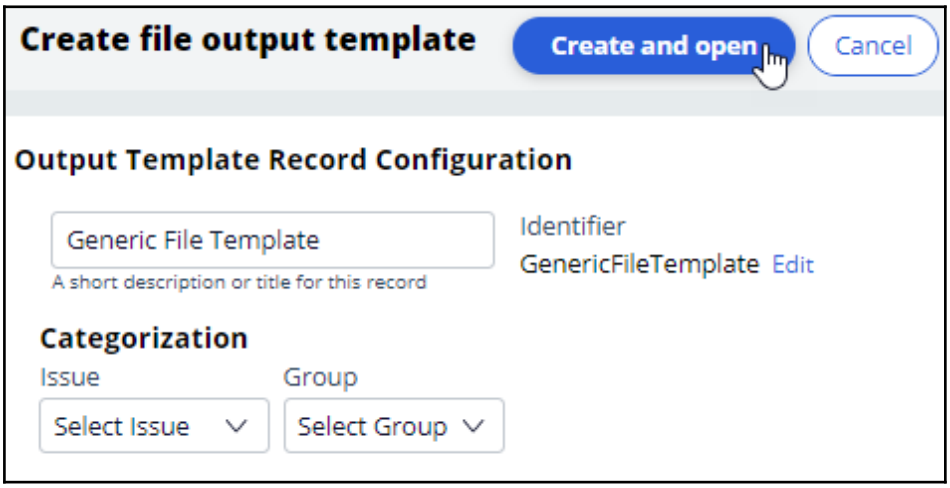

Creating a repository

When configuring a file template, provide details on an existing repository or to create a new repository where the file needs to be stored.

In this case, you will create a new repository.

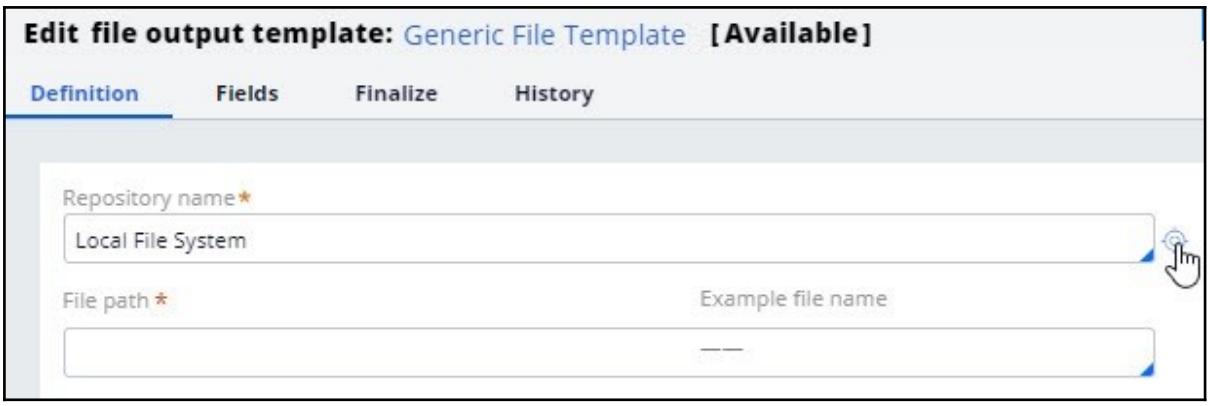

Provide a short description and a name for the repository to create one.

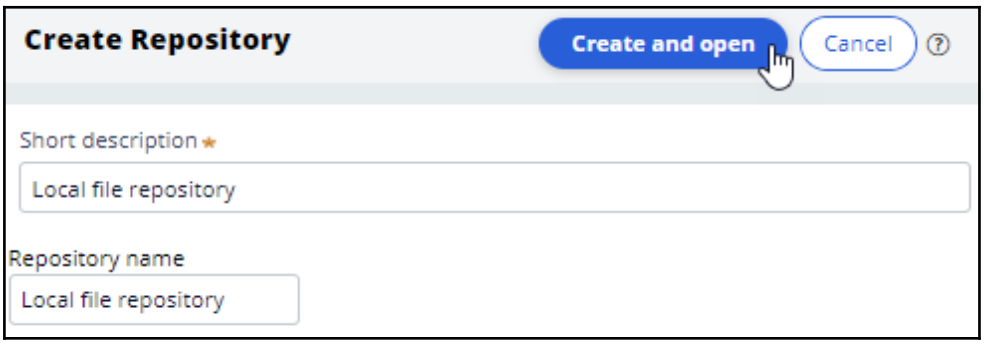

Select a **Repository Type** where you would like to store the file.

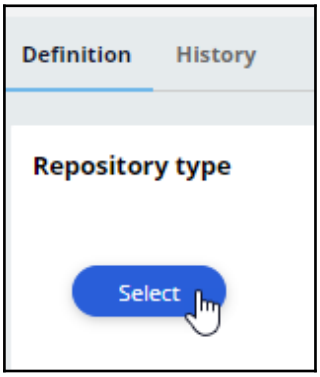

Pega Customer Decision Hub supports various cloud-based destinations as repository types such as **JFrog Artifactory**, **Amazon S3**, and **Microsoft Azure** in additional to the local file system.

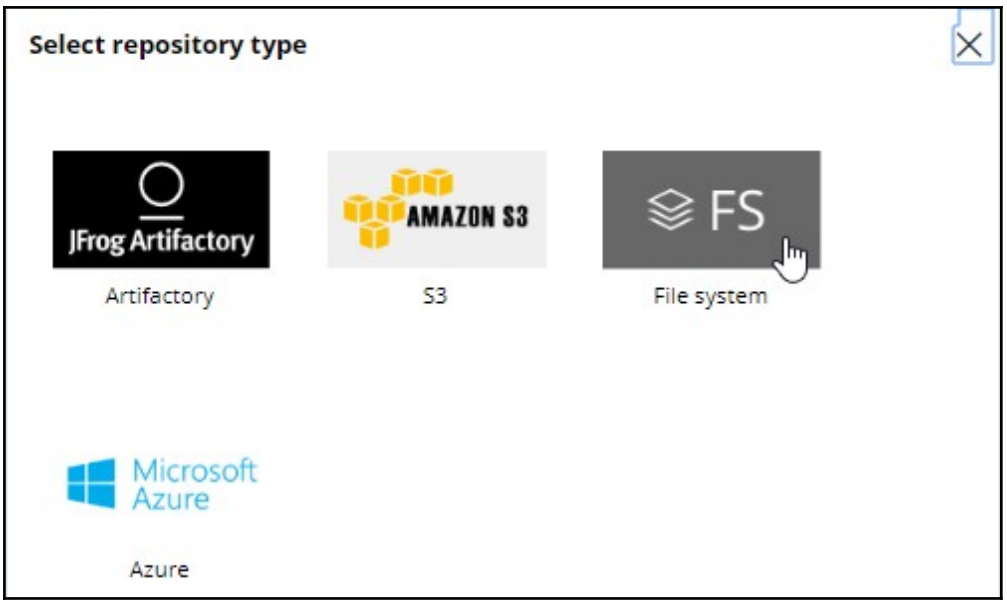

When you select the **File system** repository type, in the **Resource Path**, provide a directory path in the local file system to where you want to store the file.

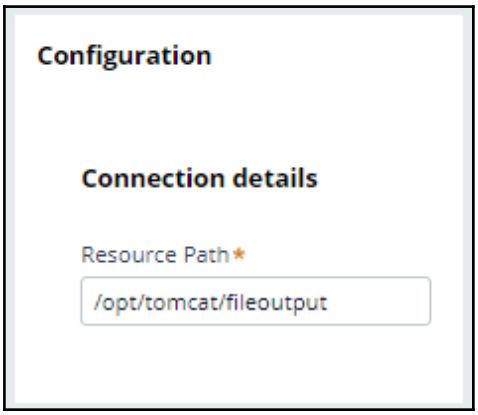

Configuring a file template

In the **Edit file output template** page, continue configuring the file template.

Select the repository and provide a path where the file needs to be stored. The location is a sub-directory under the directory specified in the file system repository. The format for the file path and name is folder/filename.

For example, if the repository path is */opt/tomcat/fileoutput* and the file path is *actions/output/mailing-list.csv,* then the file will be created under */opt/tomcat/fileoutput/actions/output.*

Output template finalization

Finalization is the technical process of writing data from an internal staging area to the final file. Finalized data is ready to be consumed by a different system. Finalization can be triggered manually or automatically by the system based on a schedule. Select the **Finalize template at the end of each outbound run option to trigger the finalization after each run.**

By finalizing a template at the end of each run, you ensure that all data in the file is the result of a single run. That is, the final output includes all data related to the run that triggered the finalization. If you want the data from multiple outbound runs to be written to the same file, leave this option unselected. When you leave this option unselected, you have to configure a separate finalization schedule in the **Finalize** tab or finalize the file manually in the **Treatments** > **File** landing page.

If you chose to finalize at the end of each run, select **Include run ID in the file name** to include the run ID for easier identification of the files generated in each run.

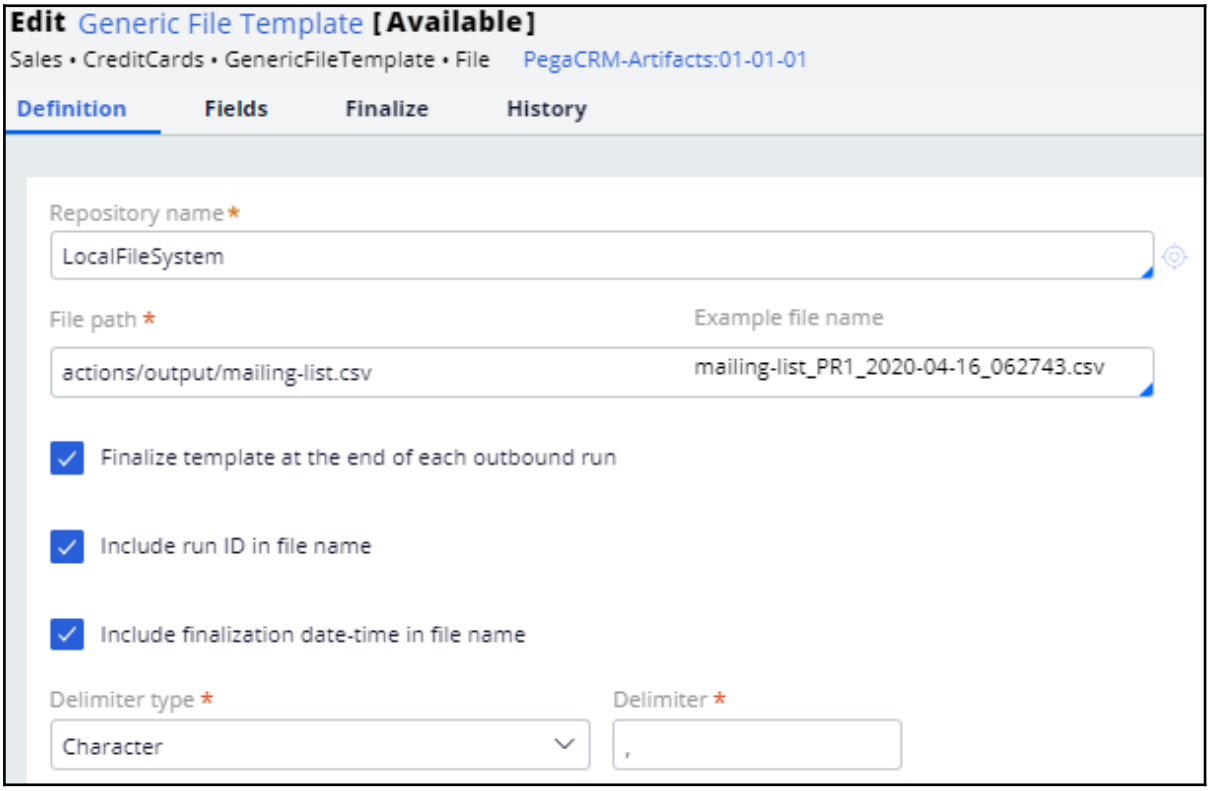

In the **Fields** tab, you can configure the information elements that you want to output.

That is, you can configure the customer and action details that you want to share with the third-party email distributor. In the **Name** field you enter a name for the field. In

the **Content** field, you select the corresponding property. Use the *.Customer.* construct to access customer properties and *.OfferData.* construct to access action properties.

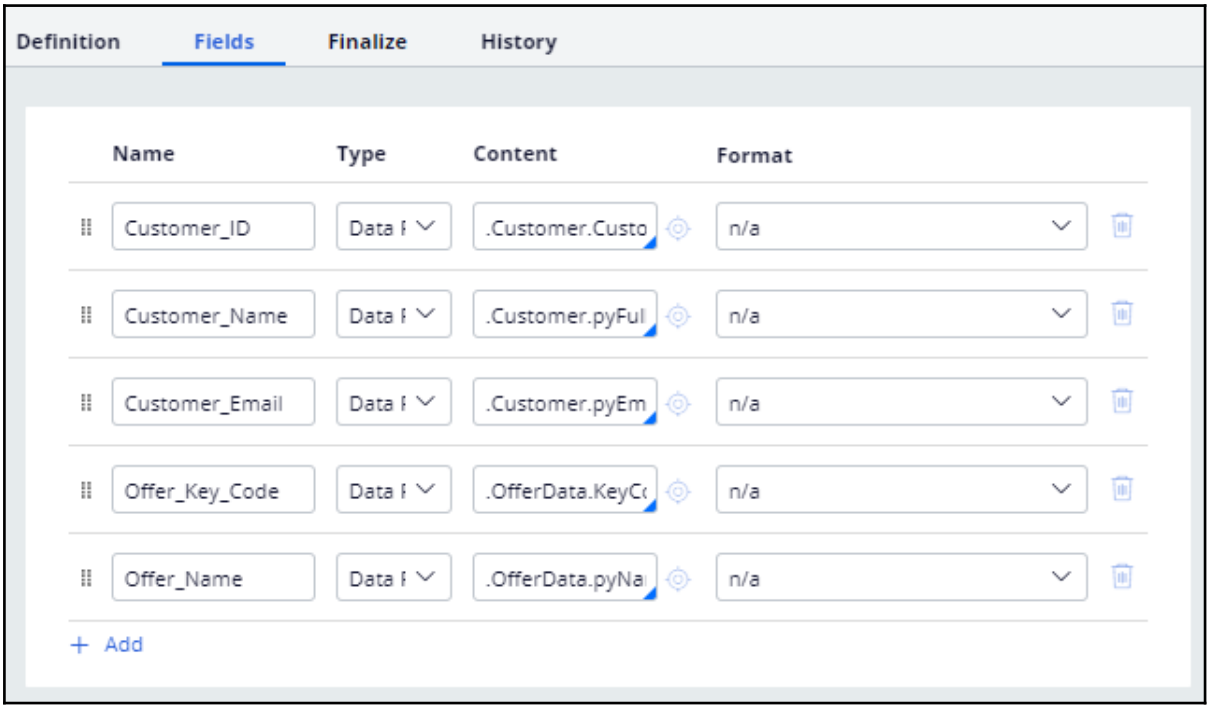

In the **Finalize** tab, you can configure further finalization options.

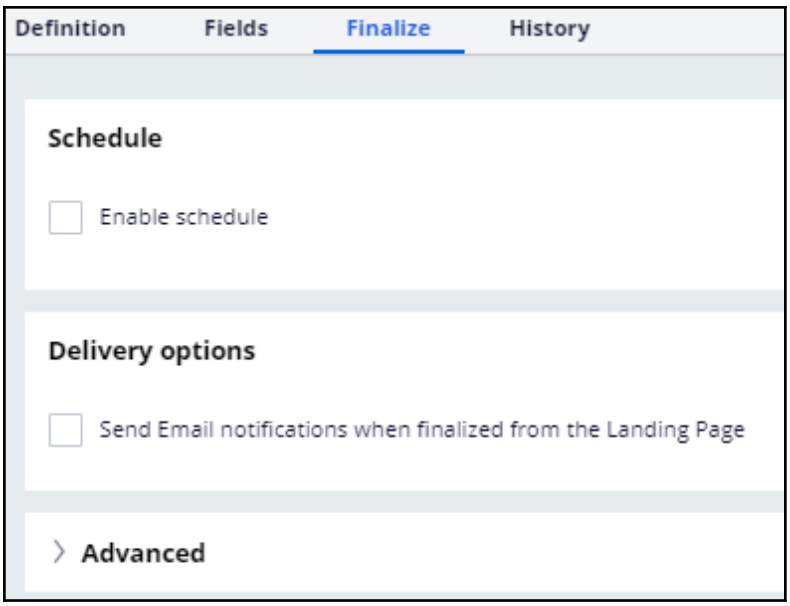

The **Advanced** section can be used for additional processing during the finalization process. For example, you can use it to initiate the file transfer to a particular location.

Using the file template in an action flow

You need to configure one of the **Send*** shapes, that is, Send Email, Send SMS, Send Generic or Send Multi shapes, in the action flow to reference the file template, so that the action details are written to the output file during an outbound run.

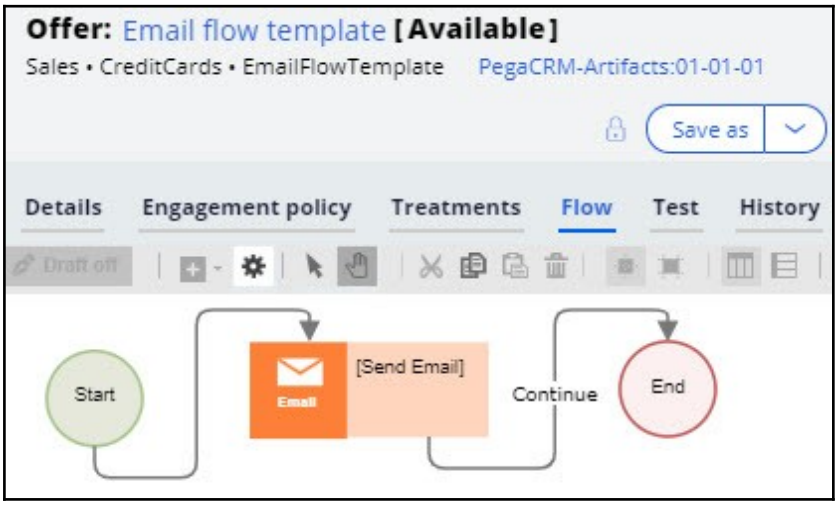

To write action details to a file template, select the **Specify Treatment** option. Now, select a treatment and reference the file template in the **Write To File** option.

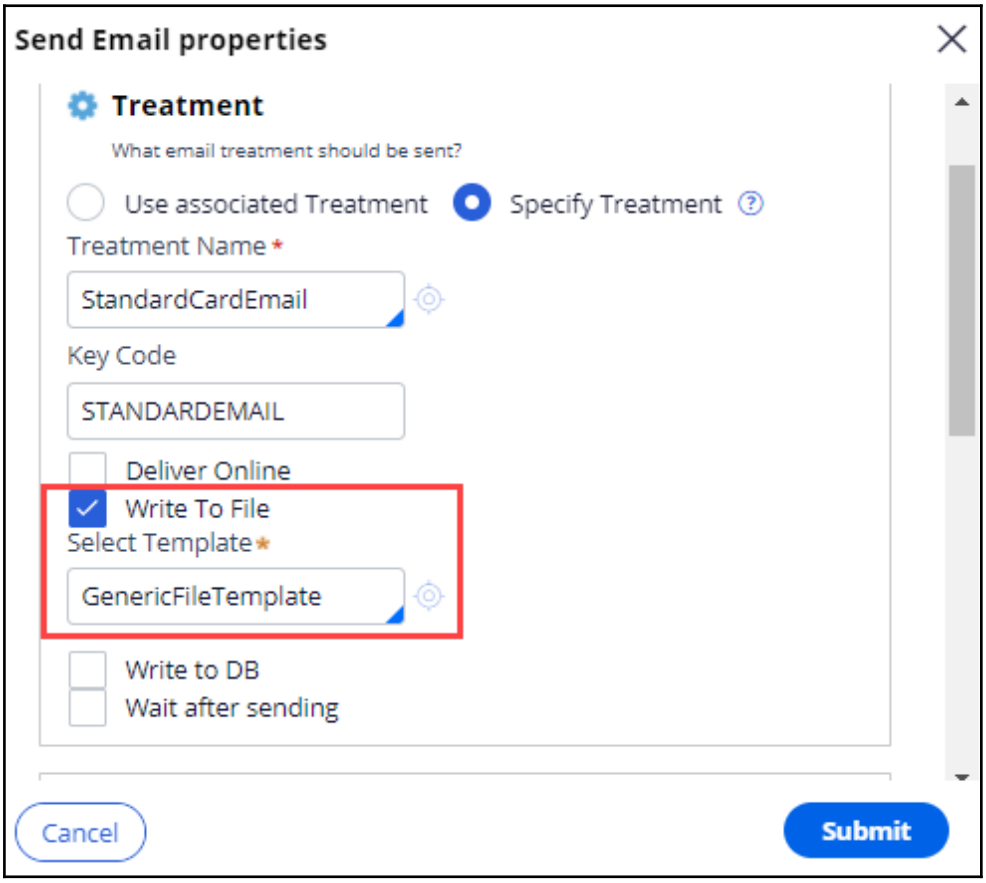

To write action details to a file, run an outbound schedule. The file with action details will be stored in the specified repository.

```
[root@ip-172-31-20-32 output]# wc -l mailing-list_PR12_2020-04-16_122456.csv
42 mailing-list_PR12_2020-04-16_122456.csv
[root@ip-172-31-20-32 output]# more mailing-list_PR12_2020-04-16_122456.csv
Customer_ID,Customer_Name,Customer_Email,Offer_Key_Code,Offer_Name
CON-894,Adrienne Marks,customer1@enablement.com,REWARDSPLUSCARD,RewardsPlusCard
CON-752,Angus Schultz,customer1@enablement.com,STANDARDCARD,StandardCard
CON-539,Antwon Graham,customer1@enablement.com,REWARDSPLUSCARD,RewardsPlusCard
CON-220,Freeman Schinner,customer1@enablement.com,REWARDSPLUSCARD,RewardsPlusCard
CON-306,Deondre Predovic,customer1@enablement.com,STANDARDCARD,StandardCard
CON-245,Jacey Hilpert,customer1@enablement.com,STANDARDCARD,StandardCard
CON-880,Lourdes Rohan,customer1@enablement.com,REWARDSPLUSCARD,RewardsPlusCard
CON-613,Roscoe Skiles,customer1@enablement.com,REWARDSPLUSCARD,RewardsPlusCard
CON-123,Mose Macejkovic,customer1@enablement.com,STANDARDCARD,StandardCard
CON-986,Delta Toy,customer1@enablement.com,STANDARDCARD,StandardCard
CON-79,Madge Renner,customer1@enablement.com,STANDARDCARD,StandardCard
CON-347,Pearlie Collier,customer1@enablement.com,REWARDSPLUSCARD,RewardsPlusCard
CON-766,Nelle Davis,customer1@enablement.com,REWARDSPLUSCARD,RewardsPlusCard
CON-367,Eleanora Larson,customer1@enablement.com,REWARDSPLUSCARD,RewardsPlusCard
CON-140,Ali Emmerich,customer1@enablement.com,REWARDSPLUSCARD,RewardsPlusCard
CON-338,Melvina Ledner,customer1@enablement.com,REWARDSPLUSCARD,RewardsPlusCard
CON-107,Pete Kerluke,customer1@enablement.com,REWARDSPLUSCARD,RewardsPlusCard
CON-916,Haskell Cummerata,customer1@enablement.com,REWARDSPLUSCARD,RewardsPlusCard
CON-141,Reva Koelpin,customer1@enablement.com,STANDARDCARD,StandardCard
CON-186,Clay Leffler,customer1@enablement.com,STANDARDCARD,StandardCard
CON-610,Claude Runolfsdottir,customer1@enablement.com,STANDARDCARD,StandardCard
CON-719,Aliza Doyle,customer1@enablement.com,STANDARDCARD,StandardCard
CON-688,Gwen Mertz,customer1@enablement.com,STANDARDCARD,StandardCard
CON-830,George Boyle,customer1@enablement.com,STANDARDCARD,StandardCard
CON-336,Kale Franecki,customer1@enablement.com,STANDARDCARD,StandardCard
CON-479,Elna Hessel,customer1@enablement.com,REWARDSPLUSCARD,RewardsPlusCard
CON-405,Glenna Emard,customer1@enablement.com,STANDARDCARD,StandardCard
CON-111,Dorcas Schumm,customer1@enablement.com,REWARDSPLUSCARD,RewardsPlusCard
CON-508,Madelyn Carter,customer1@enablement.com,REWARDSPLUSCARD,RewardsPlusCard
CON-804,Ahmad Gleason,customer1@enablement.com,REWARDSPLUSCARD,RewardsPlusCard
CON-939,Garrison Brekke,customer1@enablement.com,REWARDSPLUSCARD,RewardsPlusCard
CON-118,Joanny Pfeffer,customer1@enablement.com,STANDARDCARD,StandardCard
CON-519,Destini Vandervort,customer1@enablement.com,REWARDSPLUSCARD,RewardsPlusCard
CON-587,Caden Dibbert,customer1@enablement.com,REWARDSPLUSCARD,RewardsPlusCard
CON-528,Kaleb Aufderhar,customer1@enablement.com,REWARDSPLUSCARD,RewardsPlusCard
CON-163,Oran Beer,customer1@enablement.com,REWARDSPLUSCARD,RewardsPlusCard
CON-493,Trace Zboncak,customer1@enablement.com,REWARDSPLUSCARD,RewardsPlusCard
CON-654,Antonio Walsh,customer1@enablement.com,REWARDSPLUSCARD,RewardsPlusCard
CON-518,Germaine Schamberger,customer1@enablement.com,REWARDSPLUSCARD,RewardsPlusCard
CON-571,Charlie Hoppe,customer1@enablement.com,REWARDSPLUSCARD,RewardsPlusCard
CON-938,Cristina Satterfield,customer1@enablement.com,REWARDSPLUSCARD,RewardsPlusCard
[root@ip-172-31-20-32 output]#
```

The file contains the fields as configured in the file template.

Writing action details to a database

Similar to writing details to a file, you can make the outbound run write details to a database table using a database template. The process of creating a database template is

similar to creating a file template. Instead of the repository, file path, and file name configurations, you must provide the **Database** name and **Table** name.

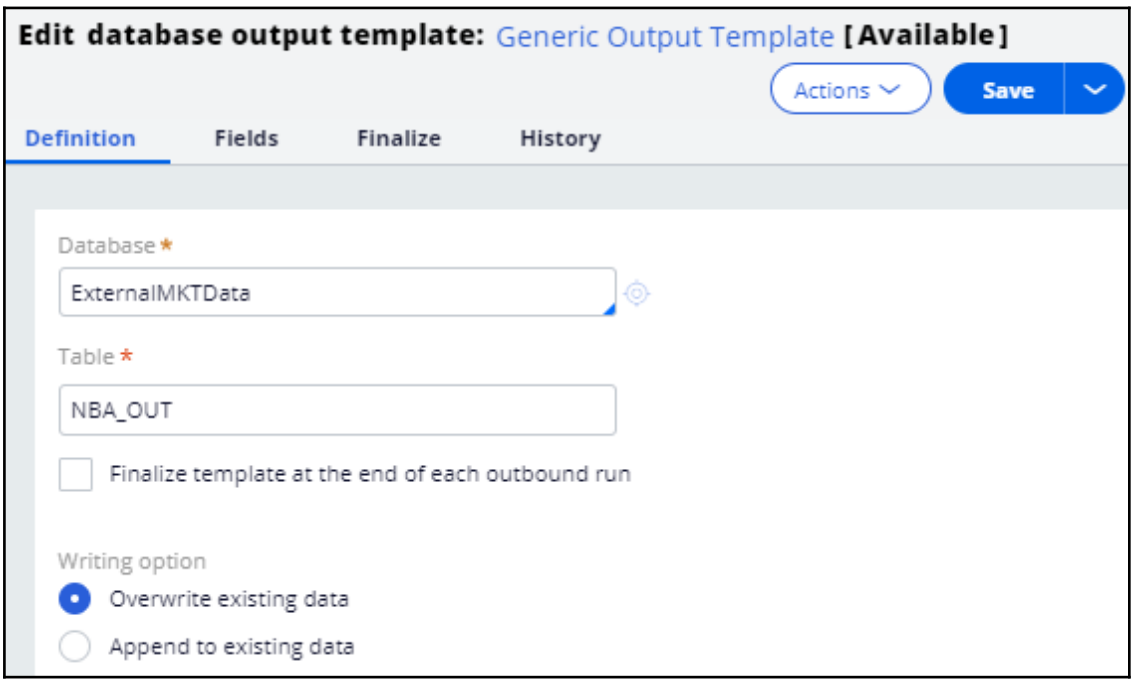

Bypassing action flow processing

In the always-on outbound scenario, you can bypass the action flow processing and write action details to a database table directly with a database template. Use this feature when you do not need to perform multiple steps in the life cycle of an action (such as send an action message, wait for a response, or take appropriate action upon response).

To use this feature, go to **Next-Best-Action Designer**, **Channel**. Under the **Primary schedule** section, click the gear icon next to **Recurrence** to access the outbound schedule configuration. Enable the **Write results using a database template** option to add a database template.

> **Note**: Here, you can configure only database templates that have the **Writing option** set to **Append to existing data**.

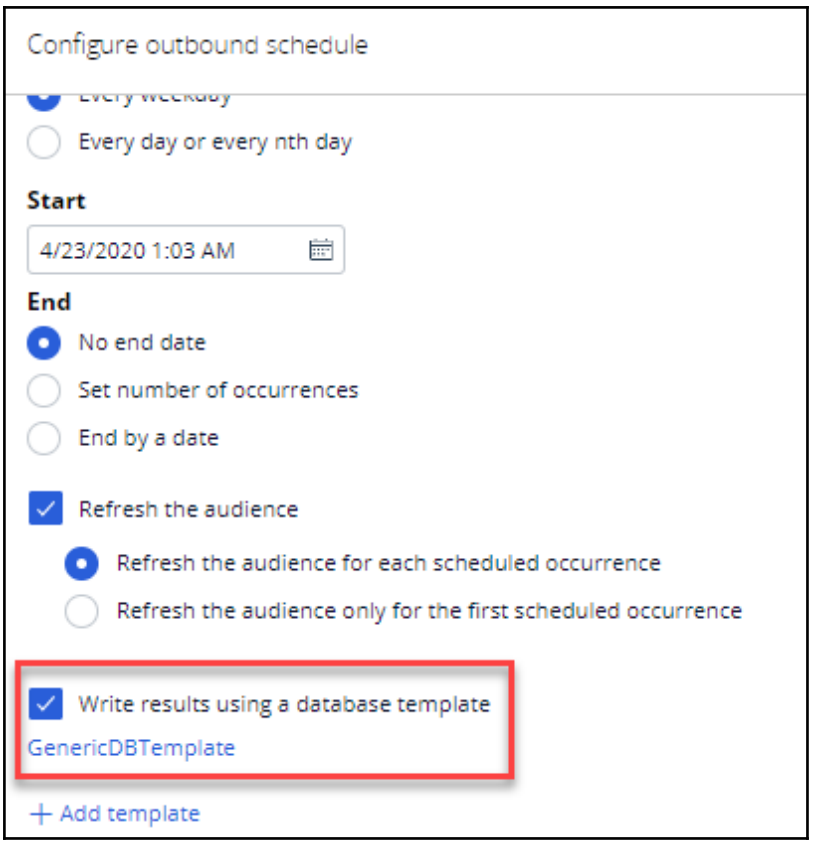
Configure outbound schedule
Every day or every nth day
Start
4/23/2020 1:03 AM
End
No end date
Set number of occurrences
End by a date
Refresh the audience
Refresh the audience for each scheduled occurrence
Refresh the audience only for the first scheduled occurrence
Write results using a database template
GenericDBTemplate
Add template

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.

Learning objectives

- Explain the purpose of simulation testing
- Define the different types of simulations

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.

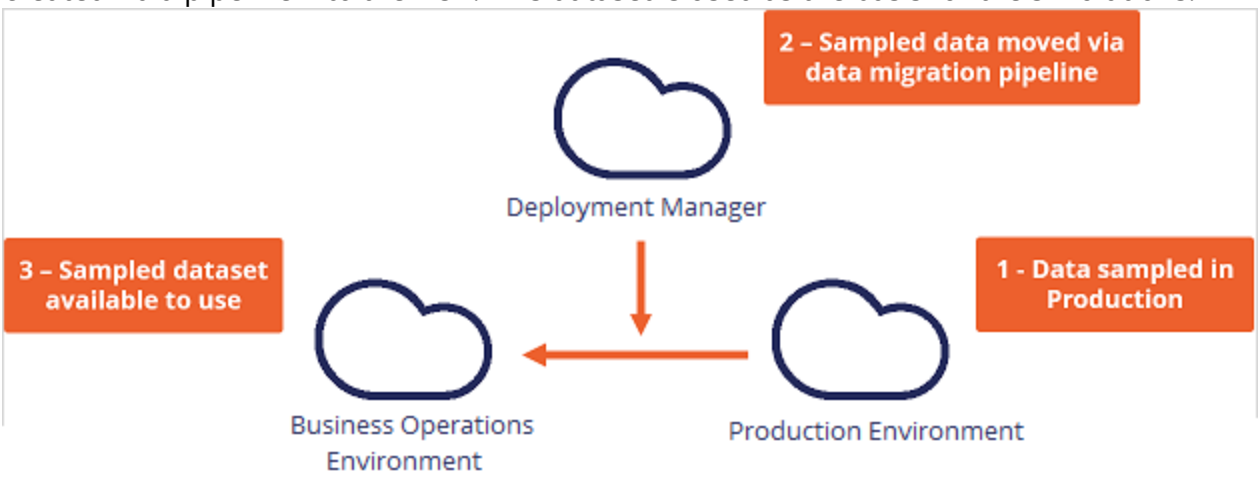

Let's see how each of these simulations can be used in a business-as-usual use case.

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.

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.

Analyzing customer distribution using Pega Value Finder

A Value Finder simulation allows you to engage more empathetically with customers by identifying and profiling “under-served” customers. The business can 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+, 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 check if there are any unserved or underserved customers 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.

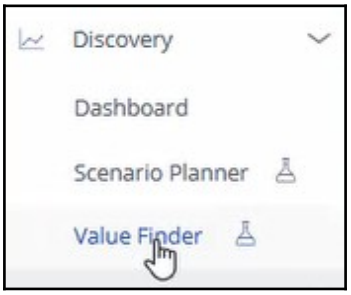

To create a Value Finder simulation run, select the issue and group in which you want to find opportunities. Then, select an audience on which to perform the simulation run. The audience is a list of potential target customers. You can modify the simulation name as required to identify the specific runs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

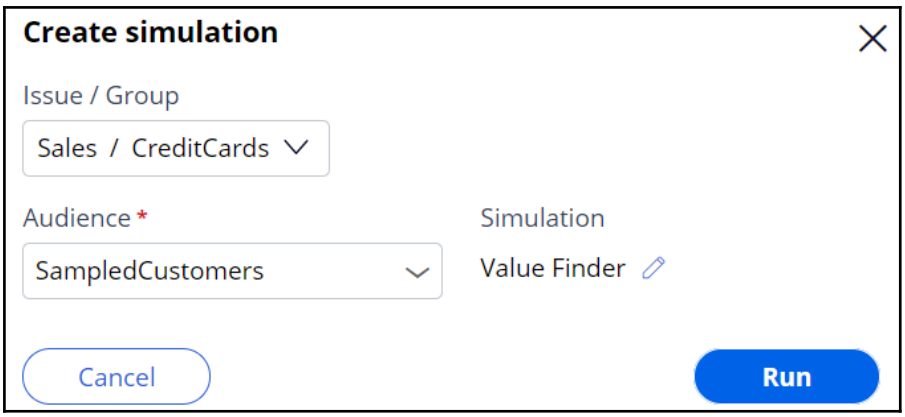

Once the simulation run is complete, the value finder indicates several opportunities.

The pie chart displays the cumulative numbers of customers without actions, under-served customers, and well-served customers after arbitration. The pie chart shows the results after eligibility, applicability, suitability conditions, and arbitration are applied. The **Without actions** category indicates the number of customers who received no actions. The **Under-served** category indicates the number of customers who received low propensity actions. The **Well-served** category indicates the customers who received high propensity actions.

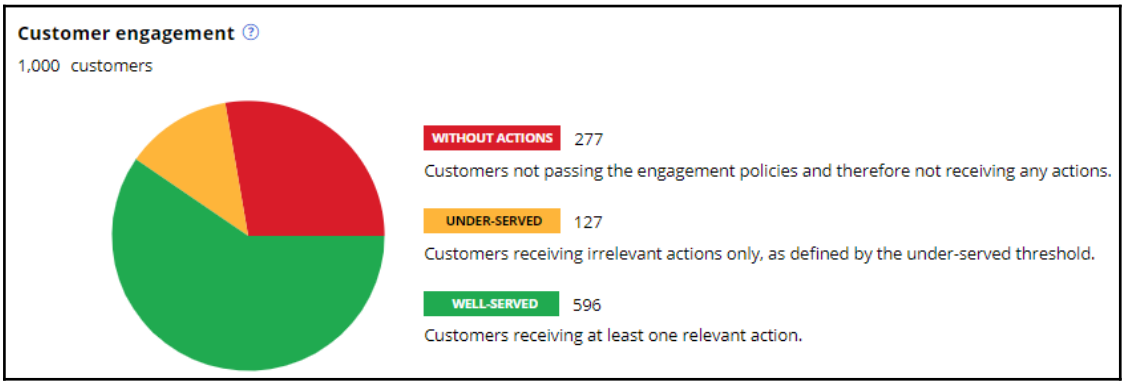

Value Finder identifies several opportunities to improve the Next-Best-Action Strategy.

Value Finder provides three different types of recommendations:

- Loosen a strictly defined engagement policy condition type: Eligibility, Suitability, and Applicability
- Shift to more empathetic arbitration settings
- Create more relevant actions or treatments

These results show how many customers the business loses after a particular Next-Best-Action stage.

In this case, the first opportunity shows the number of customers who are unable to receive actions due to suitability conditions. The second opportunity represents a group of customers that are underserved due to eligibility conditions. These customers receive less relevant actions with low propensity. The third opportunity shows the number of customers who are unable to receive actions due to eligibility conditions.

Notice that 258 customers have no actions because the suitability condition may be too strict, 201 customers are underserved because they have no relevant actions or treatments, and 19 customers have no actions because the eligibility condition may be too strict.

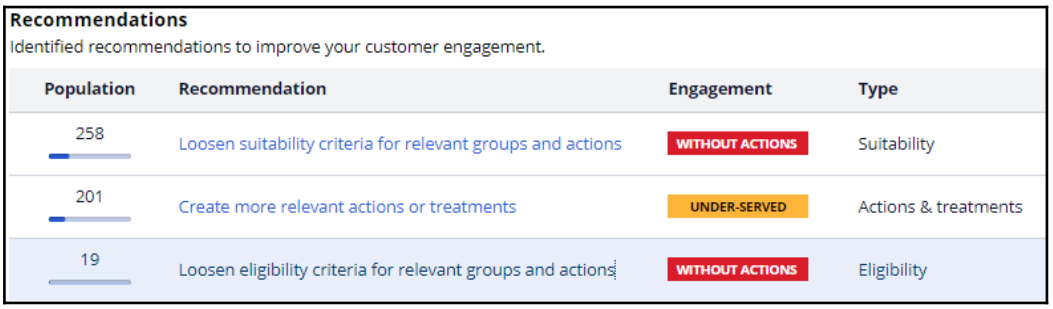

Value Finder identifies customers as underserved if the propensity of every action and treatment available to the customer is below the threshold propensity. To provide a convenient starting point, Value Finder chooses the initial propensity threshold value so that 1 out of 20 customers is defined as under-served. In this case, this results in an underserved threshold of 6.3%.

If the business wants to present offers with a higher propensity, this threshold value can be changed.

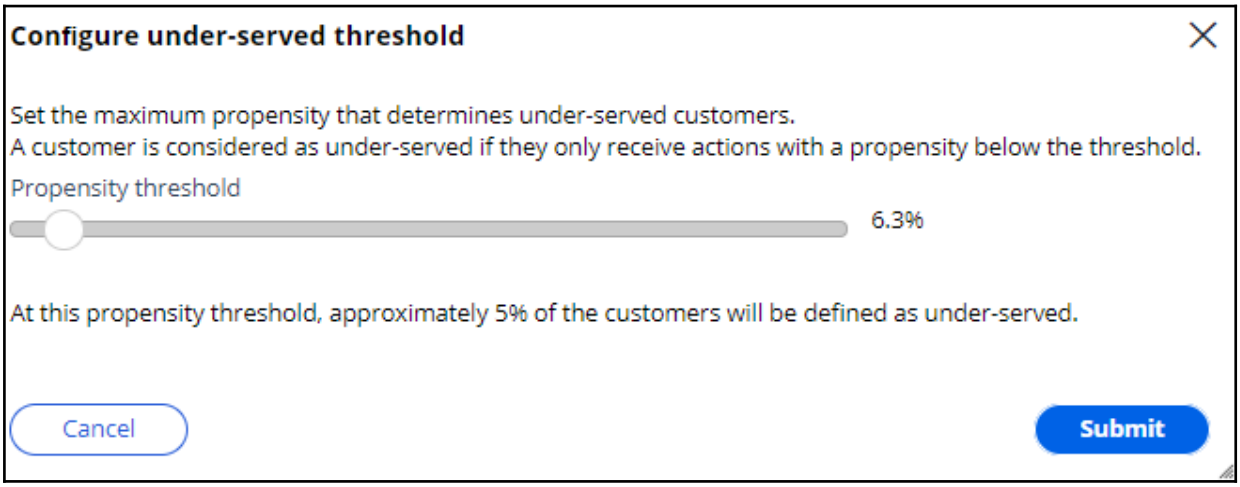

Details are provided for the **Create more relevant actions and treatments** recommendation to show how eligibility conditions contribute to the issue. Although Value Finder provides information for all engagement policy conditions, you must look at the Eligibility details to create more actions and treatments.

You can view Applicability and Suitability details for more insights.

Value Finder supplies the following information about each group:

Under-served customers: The number of underserved customers in the group.

Description: Group characteristics such as Age or LifeCycle Period.

Accuracy: The number of under-served customers in the group divided by the total number of customers in the group. If the accuracy is 100%, this means that all customers in the group are underserved. If the accuracy is lower, for example, 94.1%, then 94.1% of the customers in the group are underserved. The remaining 5.9% have at least one action above the propensity threshold.

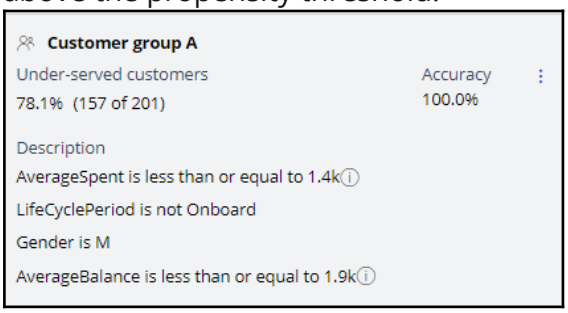

You can also manage the group by removing a field from the description. Typically, the system lists all potential fields available for the Next-Best-Action strategy. You can decide to remove a field from the list if it is not a required field. For example, if the business does not want to categorize customers based on age to avoid discrimination, they can remove the Age field from the list. If a field is removed, Value Finder then recalculates the values for **Under-served customers** and **Accuracy**.

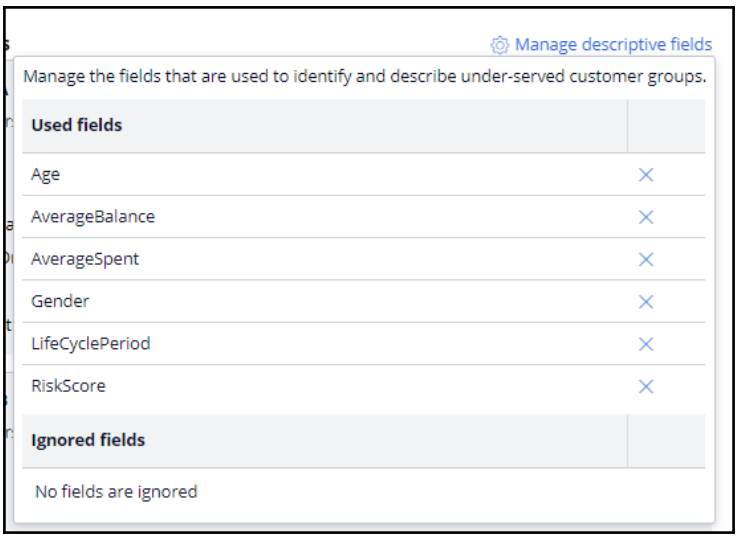

You can also save or export the customer groups that are identified as underserved as audiences.

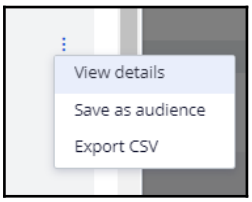

You can then run distribution tests to get more insight into the current actions that these audiences receive. Use Value Finder recommendations and your distribution test results as feedback for business stakeholders. This information can inspire them to create new actions and treatments that are relevant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

In this specific scenario, 258 customers are presented with no actions due to the suitability condition. You can further analyze and view the suggested recommendation to provide more customers with actions. Use this information from Value Finder to run a funnel filtration on an audience simulation to detect which level the suitability condition is preventing customers from receiving an action. (whether it is the suitability condition at the group level or the action level). You can also look at the no actions due to the eligibility category to view the suggested recommendation.

This demo has concluded. What did it show you?

- How to configure and run a value finder simulation
- What the top opportunities identified by Value Finder are
- How to interpret the customer group in the under-served customer category
- How to interpret the without actions customer category

- How to interpret the under-served customer category

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.

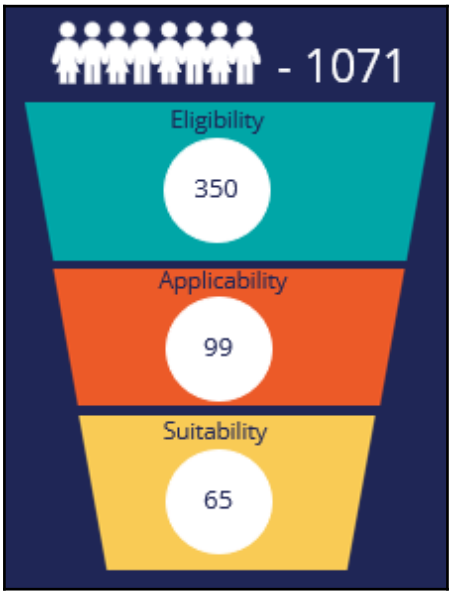

The filtering that happens in this simulation is similar to a funnel filtration for every engagement policy condition type. However, within an engagement policy condition type, each of the criterion is applied separately to the corresponding input audience.

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.

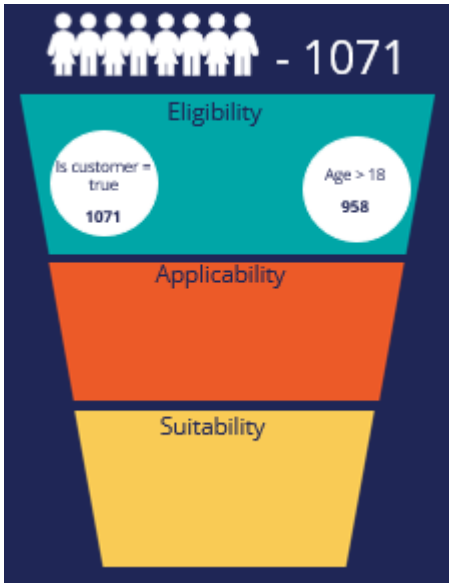

Since both conditions need to be met, the intersection of these two conditions, 958, is the final number of customers who pass through at the group level. Now the system checks for individual action-level eligibility conditions that also need to be met.

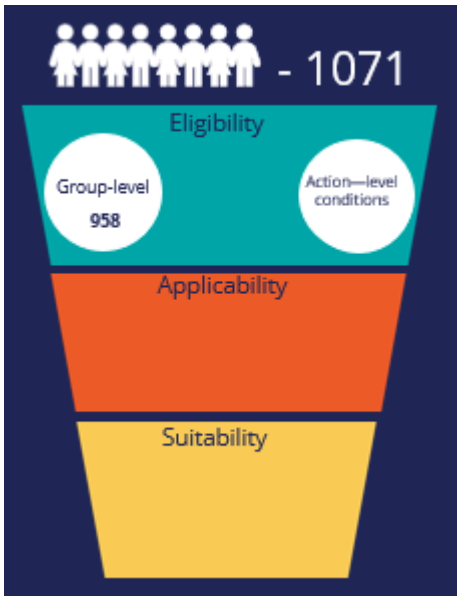

The intersection of the group-level and action-level conditions form the final eligibility-level output population of 350.

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.

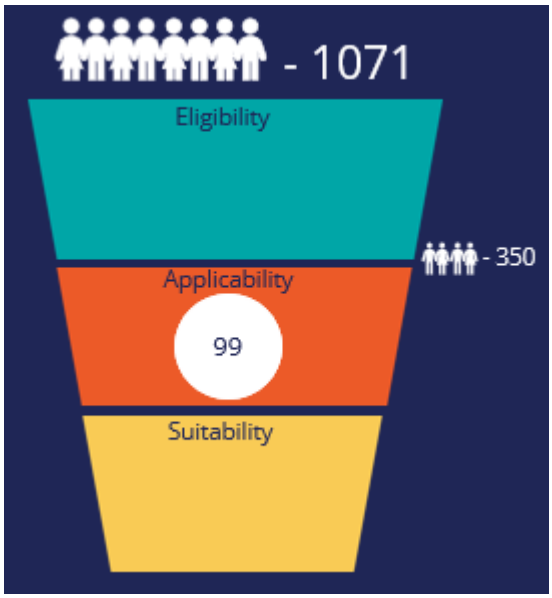

The output population of the applicability level, 99 customers, is passed on as the input to the suitability level.

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

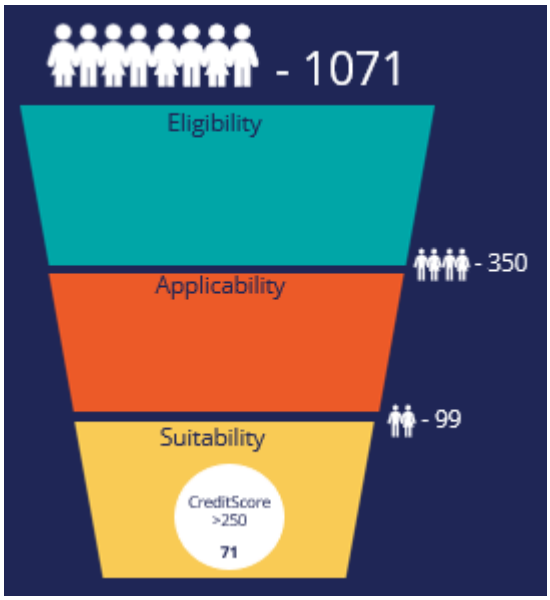

Now the system checks for individual action-level suitability conditions, and the intersection of both group-level and action-level conditions forms the final suitability-level output population, which is 65.

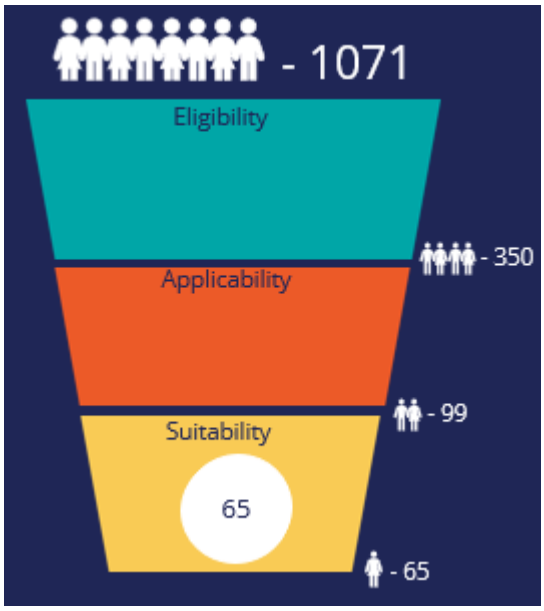

In summary, with the audience simulation test you can check how many potential actions are filtered out by each component of the engagement policy and discover if a particular criterion is too broad or too narrow for your requirements.

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.

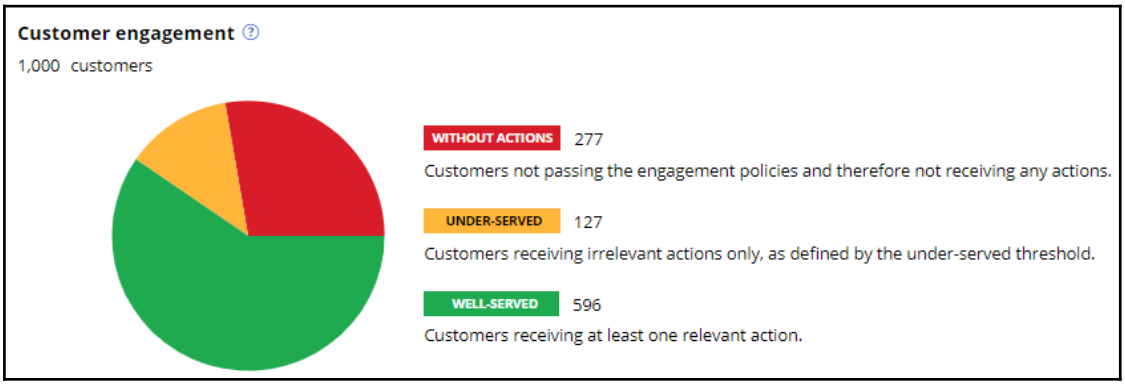

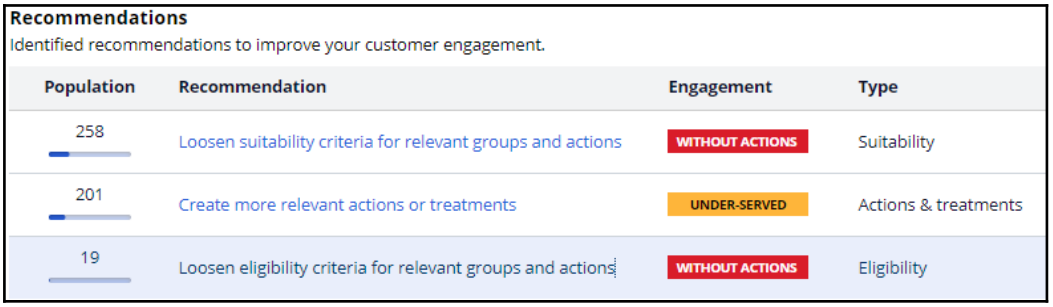

The bank would now like to run an audience simulation to investigate why these customers have no actions presented to them and correct the condition that causes the deviation.

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.

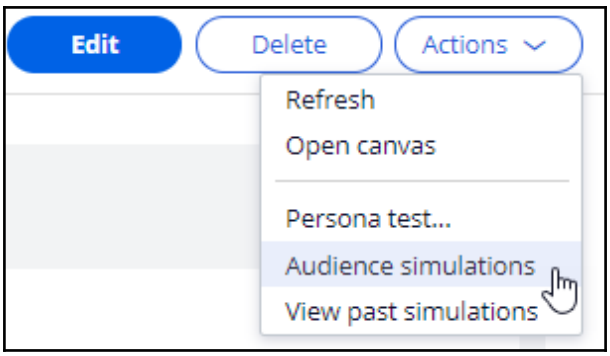

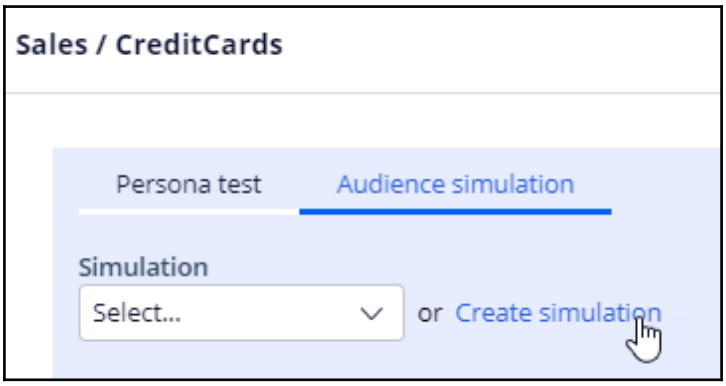

To create an audience simulation run, select an audience on which you would like to do the simulation. The audience is a list of potential target customers. You can modify the simulation name as required to easily identify the specific runs. Then, choose the scope of the run. That is, you can choose 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.

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.

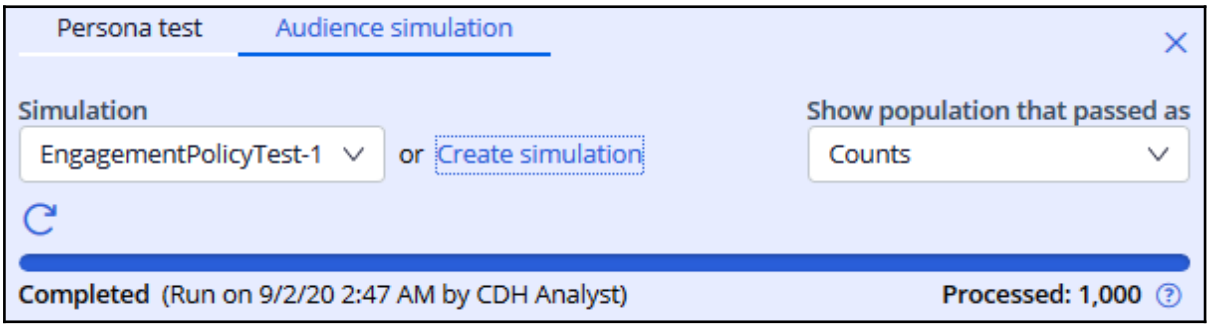

For each component of the engagement policy, the simulation test shows a numerical or percentage value of the audience that will receive the action based on current criteria. For example, if the result of a criterion that checks if the action is active returns 100%, the component did not filter out any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

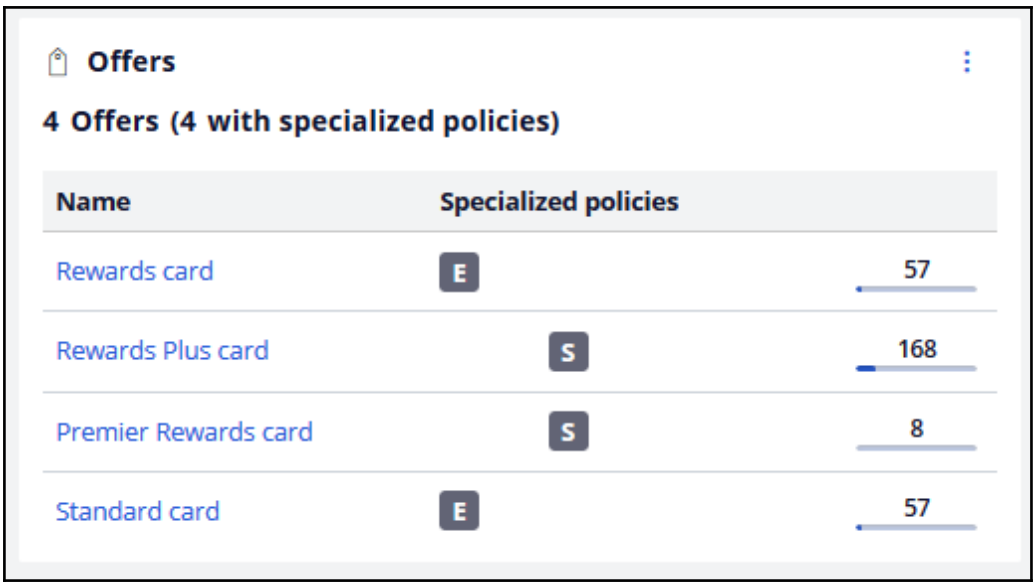

You can choose the same audience simulation to view the action-level filtration details and investigate which engagement policy condition is causing the current outcome.

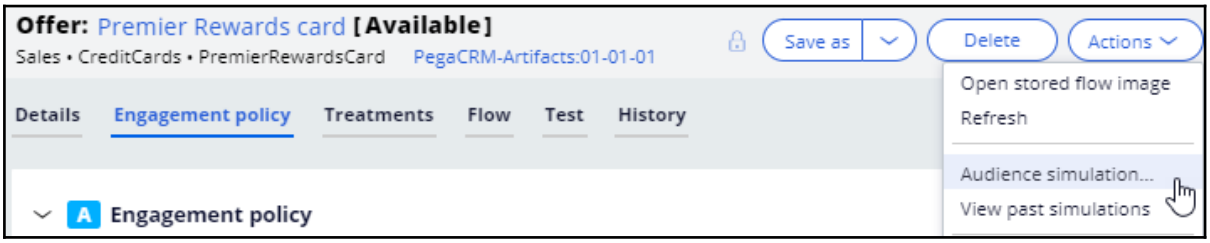

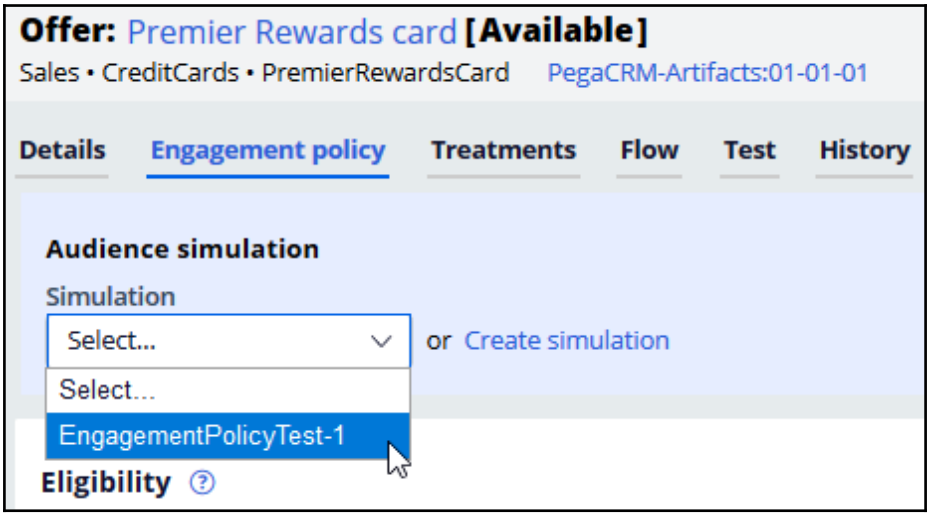

These numbers show how the engagement policies that are inherited from the CreditCards group are filtered. At the bottom, you can also view the final number of customers who qualify for the specific offers once the action-level engagement policy is applied.

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

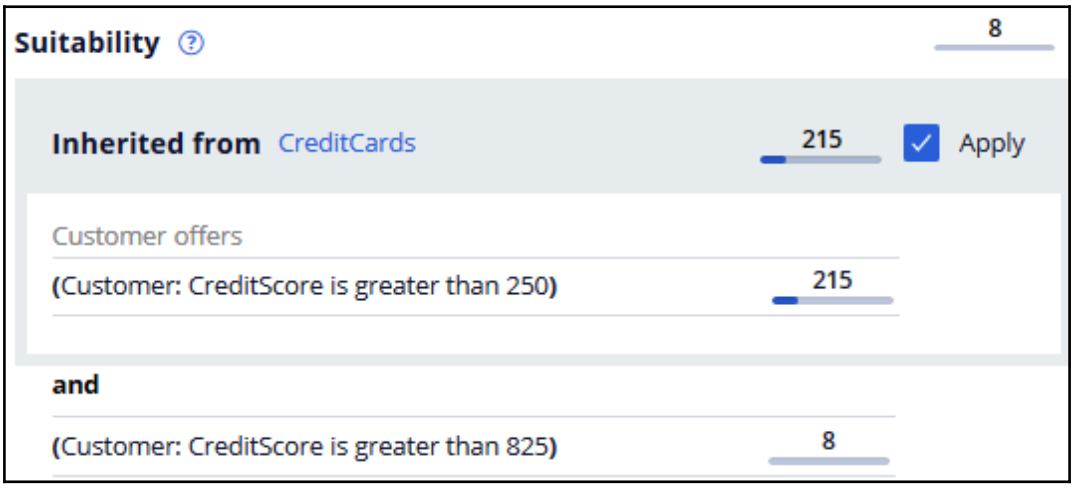

This is because the configured engagement policy criterion is too narrow. The results of this simulation show us that from the 215 customers who passed the eligibility and applicability conditions, all qualify for the group-level condition, but only eight customers qualify for the action-level condition. 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.

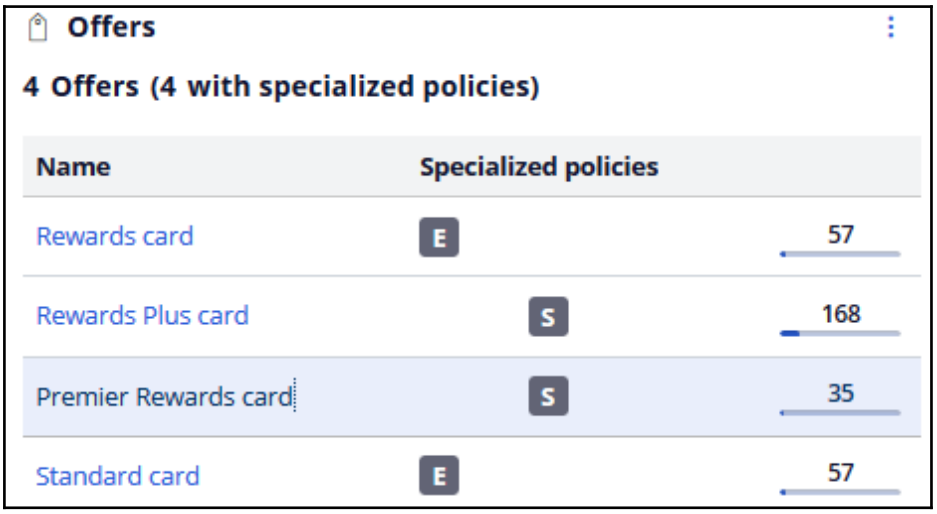

Let’s now run an audience simulation with engagement policies and arbitration in scope.

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.

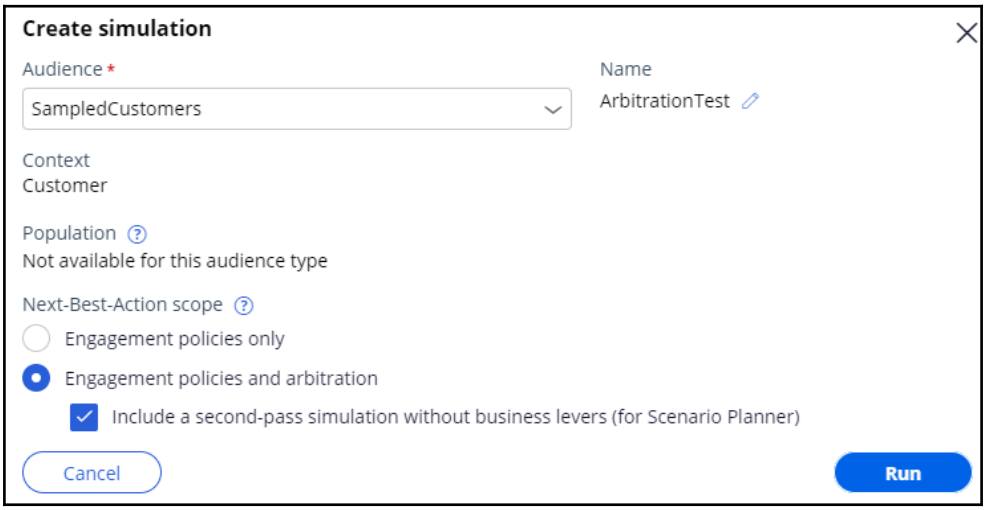

Note that the simulation results are different. This is because the simulation includes arbitration, adaptive analytics, treatment and channel processing, and constraints. That is, this simulation result shows the number of customers who will receive an action as their top action. Thus, the numbers at 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.

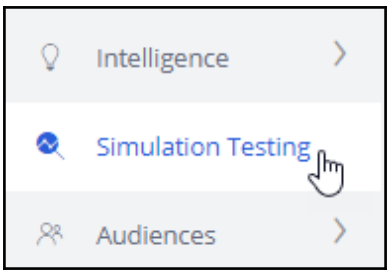

This is the auto-created funnel simulation. Alternatively, you can also create a simulation run from here and show the simulation results in the **Engagement policies > Audience Simulation**.

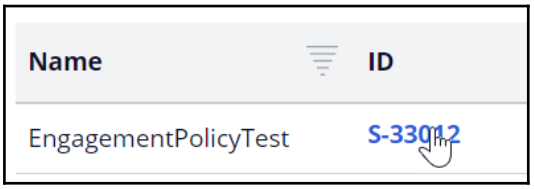

To view the funnel simulation, open the simulation.

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.

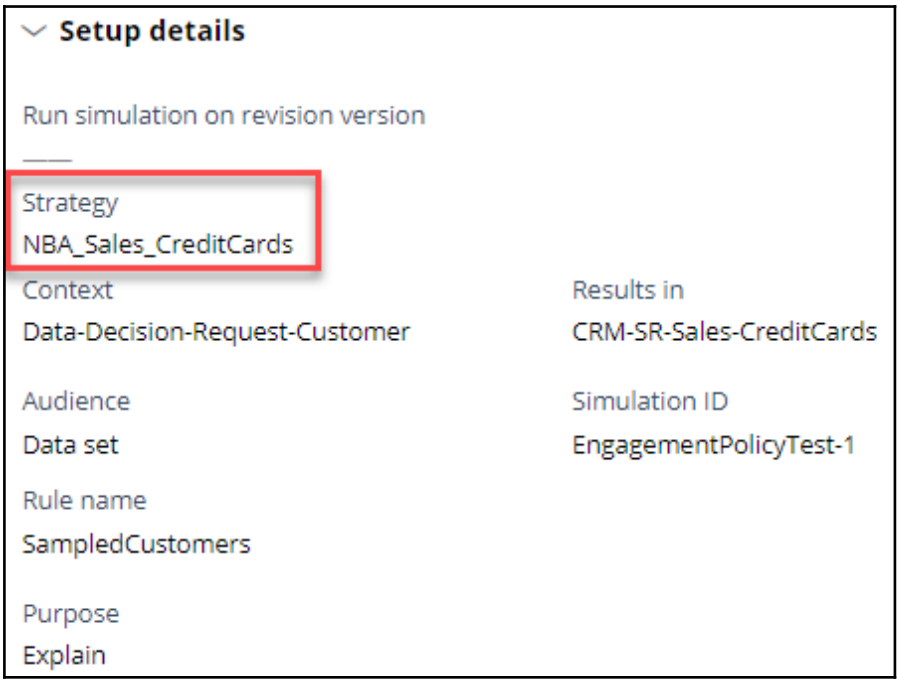

When a funnel simulation is created, a set of reports is added to the run results.

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.

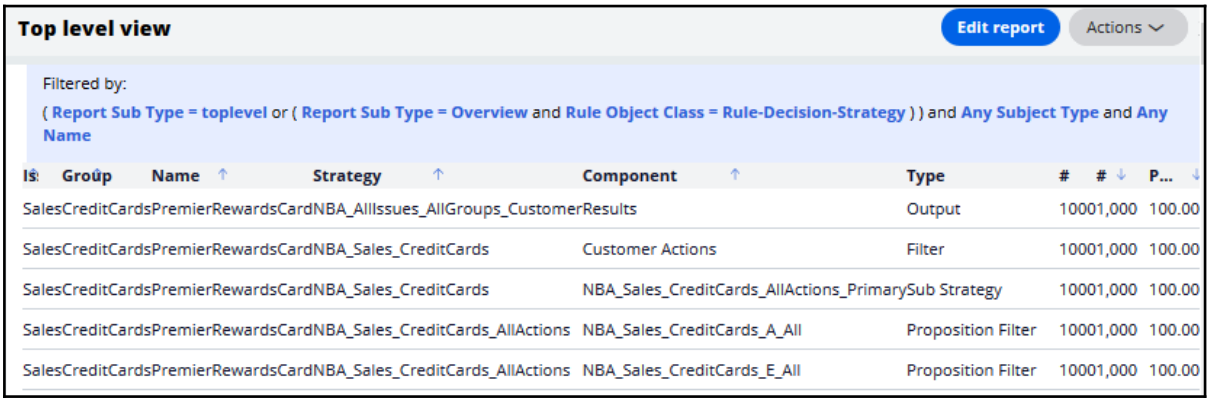

If you do not want to re-run the simulation in the future, you can wrap up the simulation. Any attached reports remain available for review.

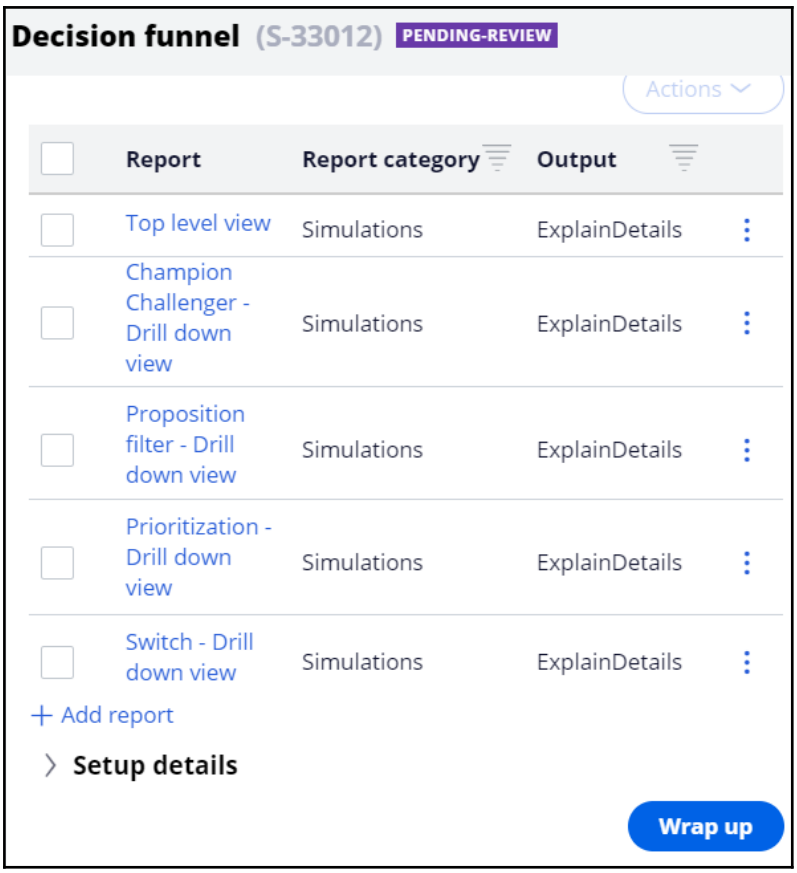

All the resolved simulations are accessible from the **Include resolved simulations** section.

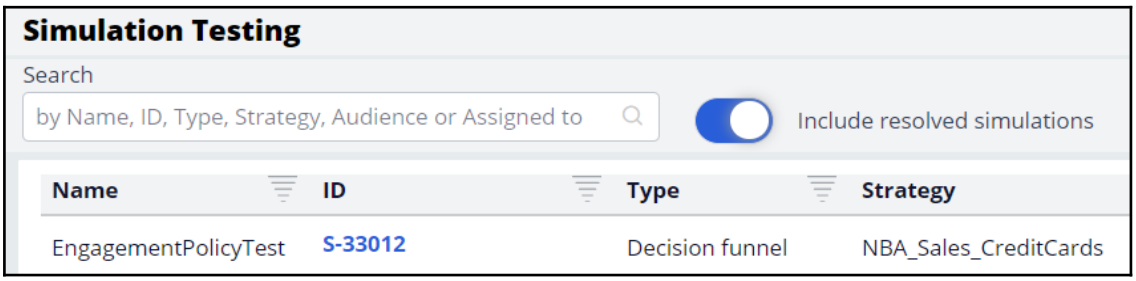

This demo has concluded. What did it show you?

- 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.

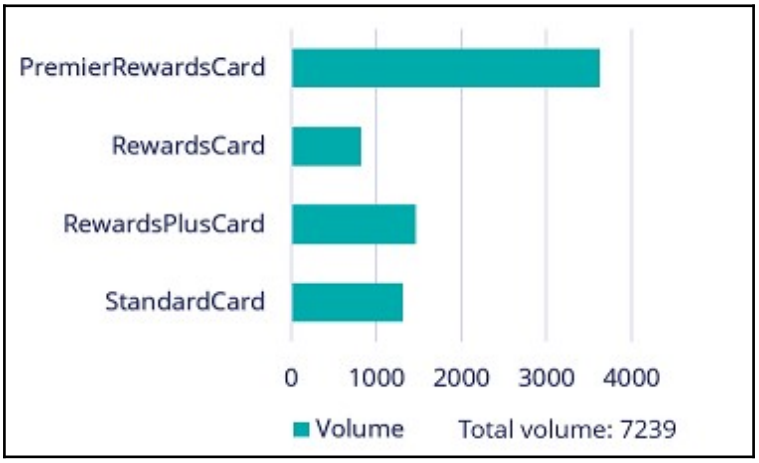

The Action priority report shows the average priority of each action. Here you can see that the RewardsPlus card offer has the highest average priority, meaning the expectation is that this will be the best performing action. If an action has a low average priority, the chance of this action being offered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

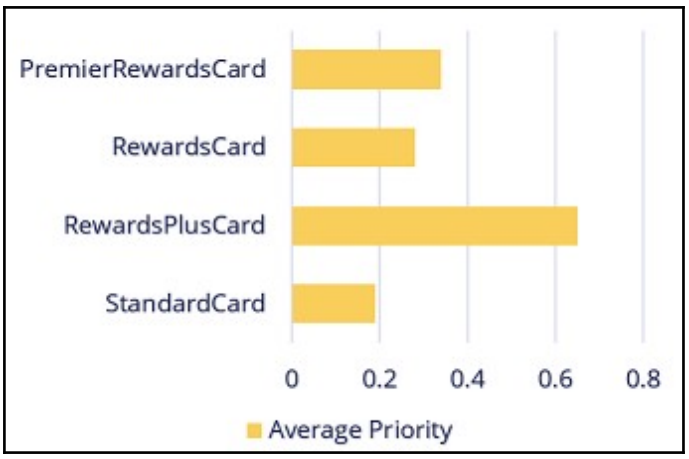
PremierRewardsCard
RewardsCard
RewardsPlusCard
StandardCard
0
0.2
0.4
0.6
0.8
Average Priority

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**.

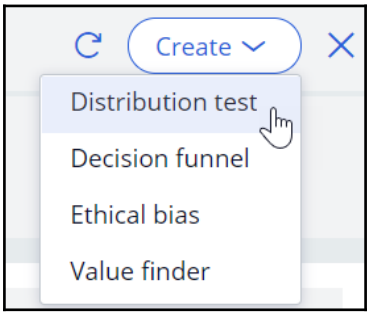

To create a simulation, select the top-level strategy on which you would like to run the simulation. Use the top-level next-best-action strategy, as this simulation pulls in data from all relevant actions and treatments.

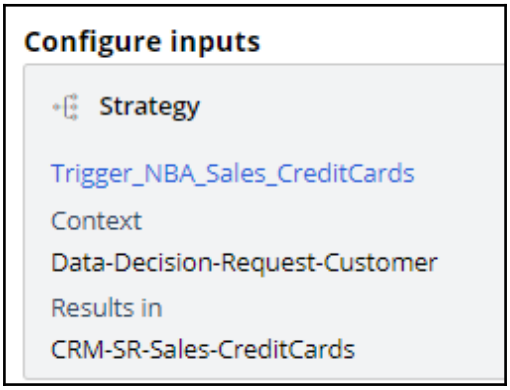

Then, select the input population on which you want to execute the simulation.

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.

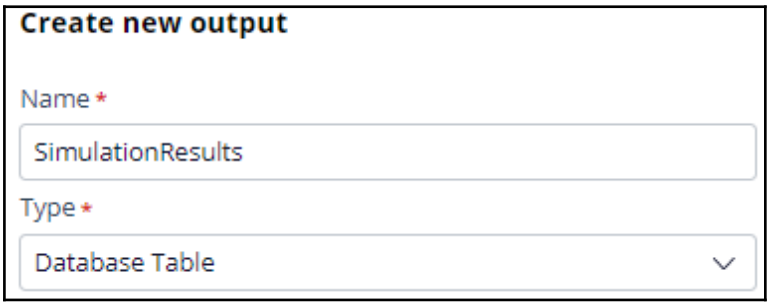

When one or more output destinations is a database table, you need to select which reports you want to have available for each destination. Select the output database table, then the report category, and finally one or more reports. Note that there are several out-of-the-box reports available. A decisioning 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**.

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.

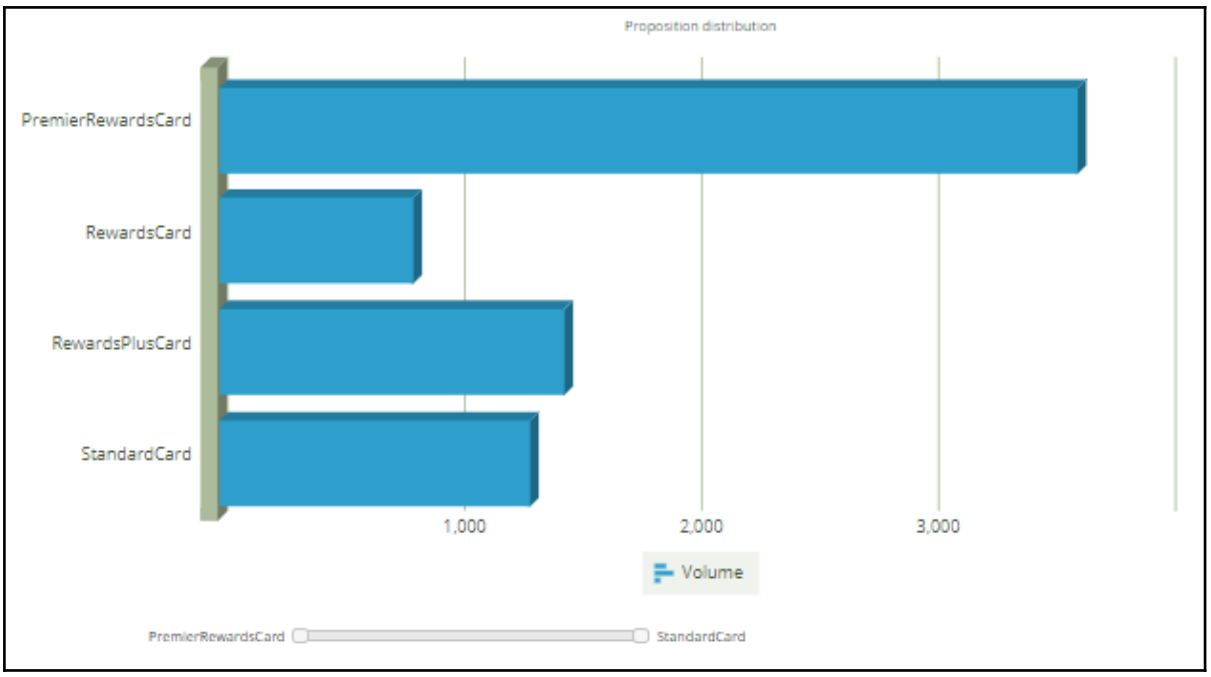

In contrast, the Action distribution report you would see in an audience or decision funnel simulation is only the number of customers who qualify for an offer.

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.

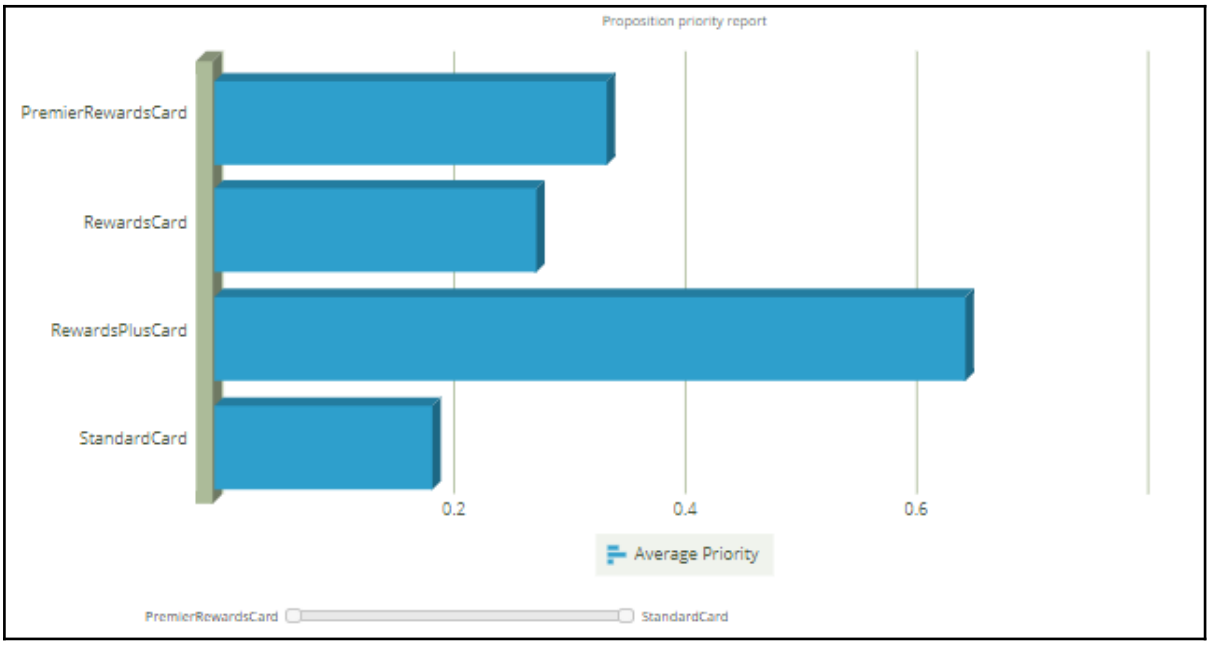

This demo has concluded. What did it show you?

- 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.

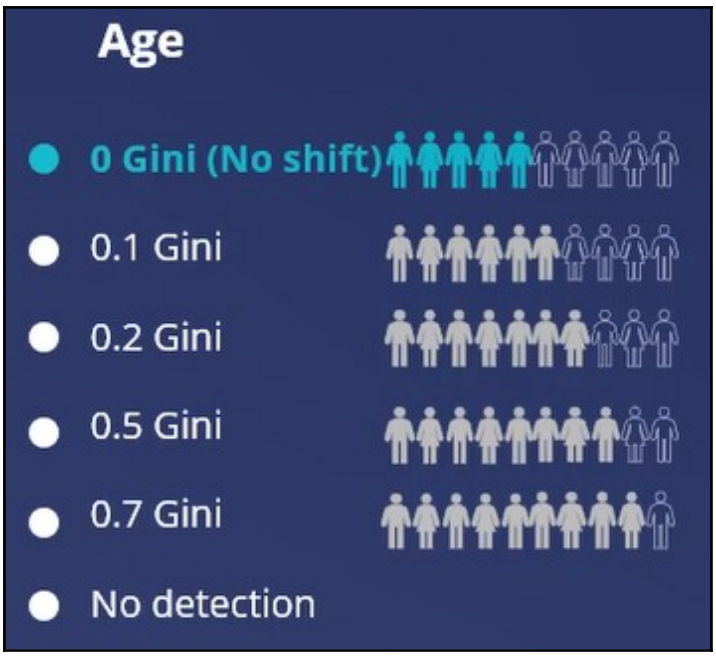

As Gender is a nominal value, a rate ratio is used to determine bias. A rate ratio is used to determine bias for categorical fields by comparing the number of customers who were selected for an action to those not selected for an action, and correlating that to the selected bias field.

A rate ratio of 1 represents perfect distribution equality. You can select a warning threshold between 1, warn if any bias is detected that significantly deviates from a rate ratio of 1, 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.

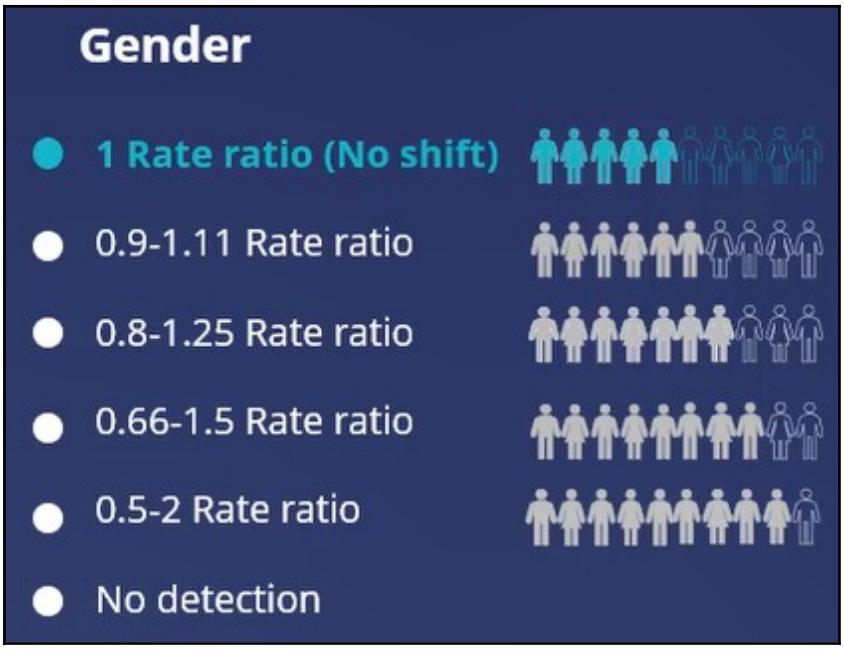

Even though no bias is ideal, it is not practical to set the threshold at that level as that would block any actions from reaching customers. Hence, in real business scenarios, a little bias is often allowed to prevent customers from receiving no actions.

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, specify it as an ordinal number. Age is an example of an ordinal number.

Next, add gender as a bias field.

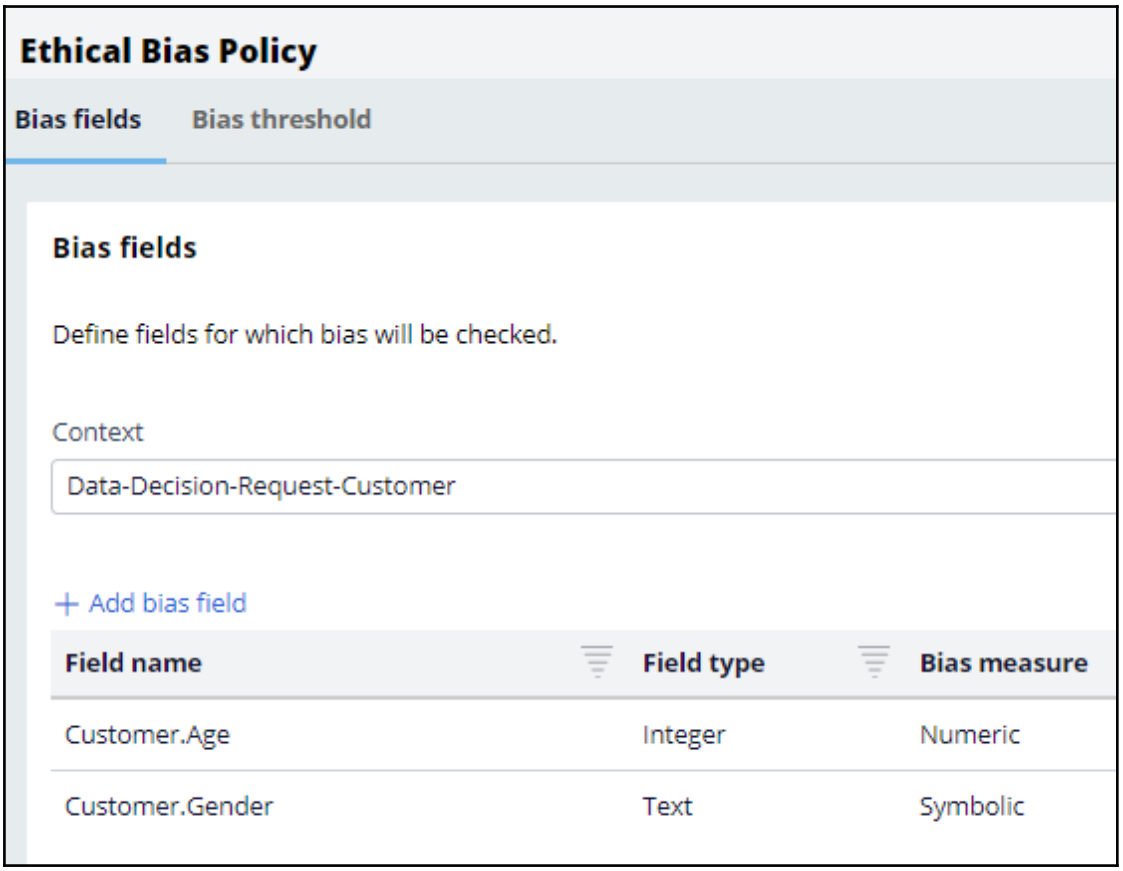

Now, on the **Bias threshold** tab, review and configure the bias threshold settings for each property you selected. This configuration is done at the business issue level. This means, for every business issue, you can decide how much bias to allow. For instance, for the risk issue, you might want to disallow any bias, whereas for sales, 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.

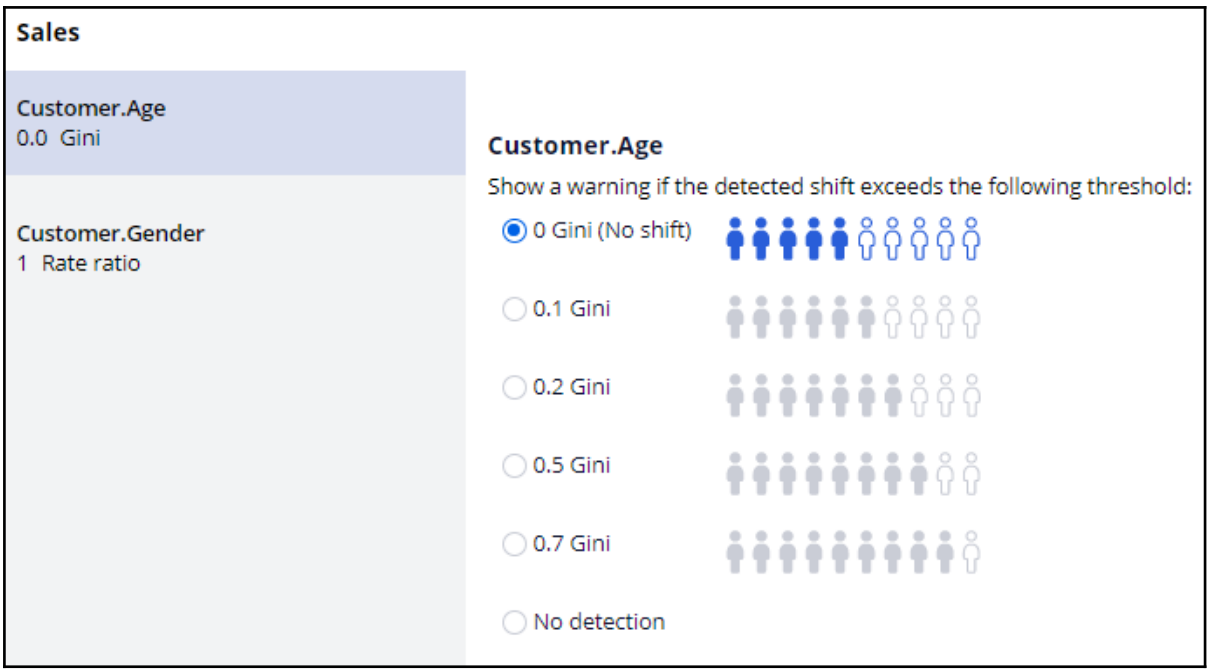

The bank also does not want to discriminate on gender, so select the appropriate threshold value. Save the changes.

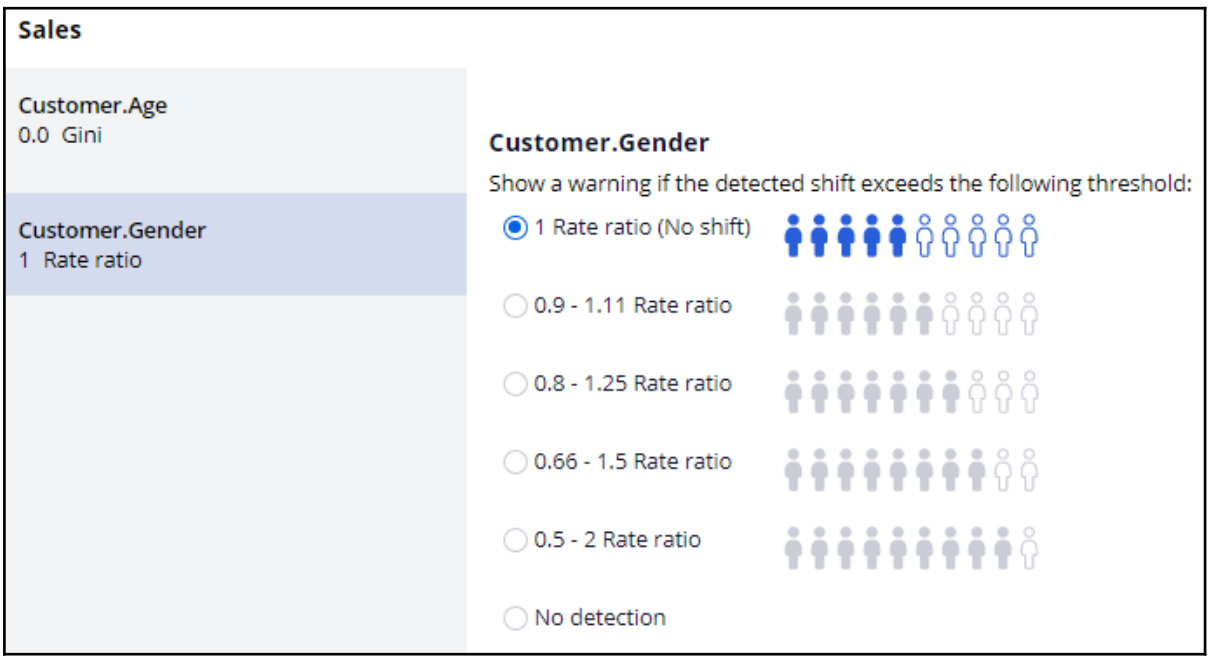

Navigate to the **Simulations** landing page to create and run an **Ethical bias** simulation.

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.

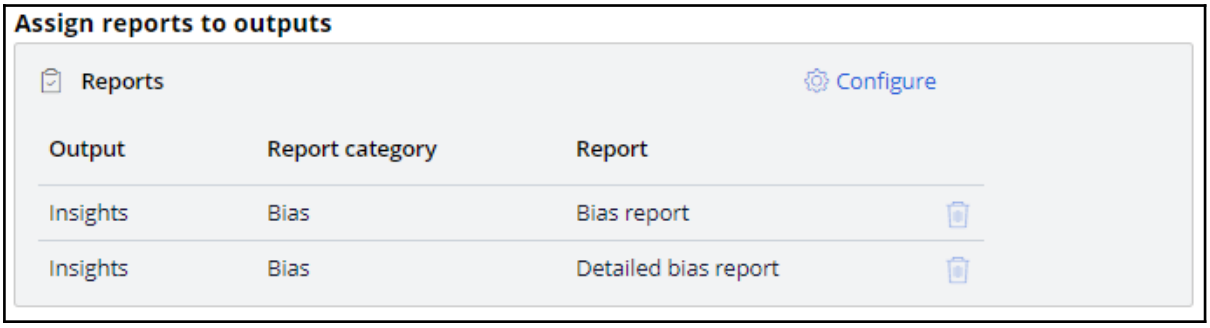

Run the simulation. Once the run is complete, there will be an indication if any unwanted bias is identified. In this scenario, although business did not want any bias, it seems some bias has been detected.

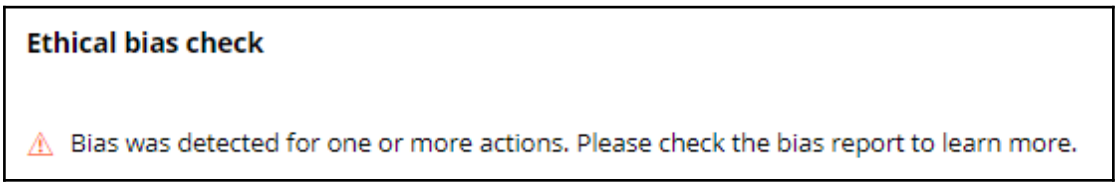

Open the generated reports and view the information in detail to understand where the bias was detected. You can sort the report in the bias detected column. Notice that there is bias detected on age. You can check the **Bias value** against the **Confidence interval** to check if the bias value is within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

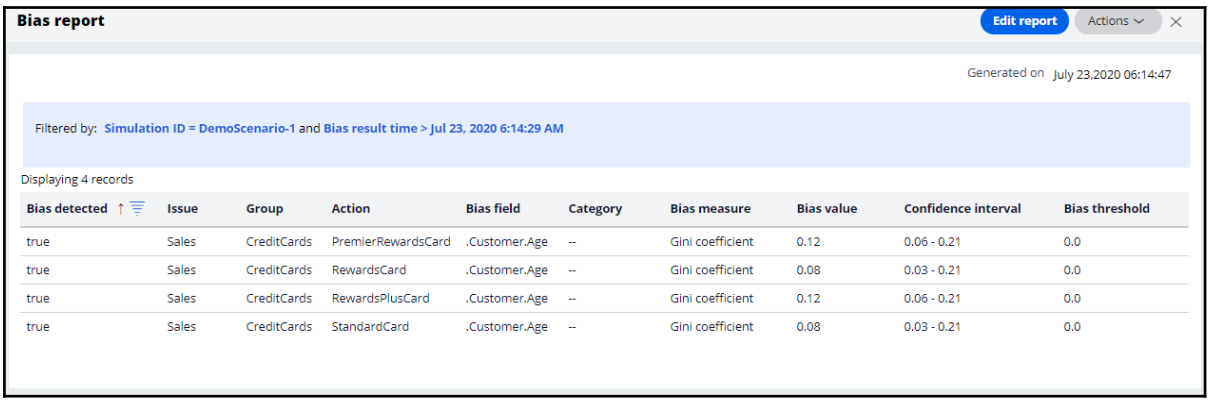

Navigate to Next-Best-Action Designer to view the engagement policy conditions. Note that there is an eligibility condition that uses age. Credit cards are valid only for customers with an age greater than 18. This is a hard eligibility rule that cannot be ignored. It is likely that bias has been detected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

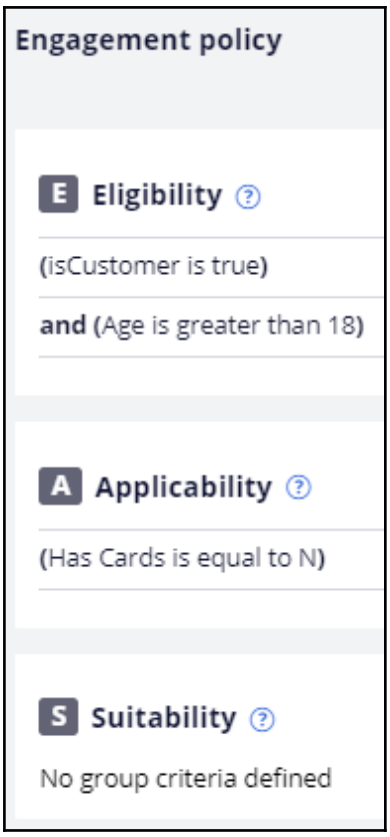

Based on the bank's policies and regulations, you must review the bias threshold and decide how much bias the business will allow.

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.

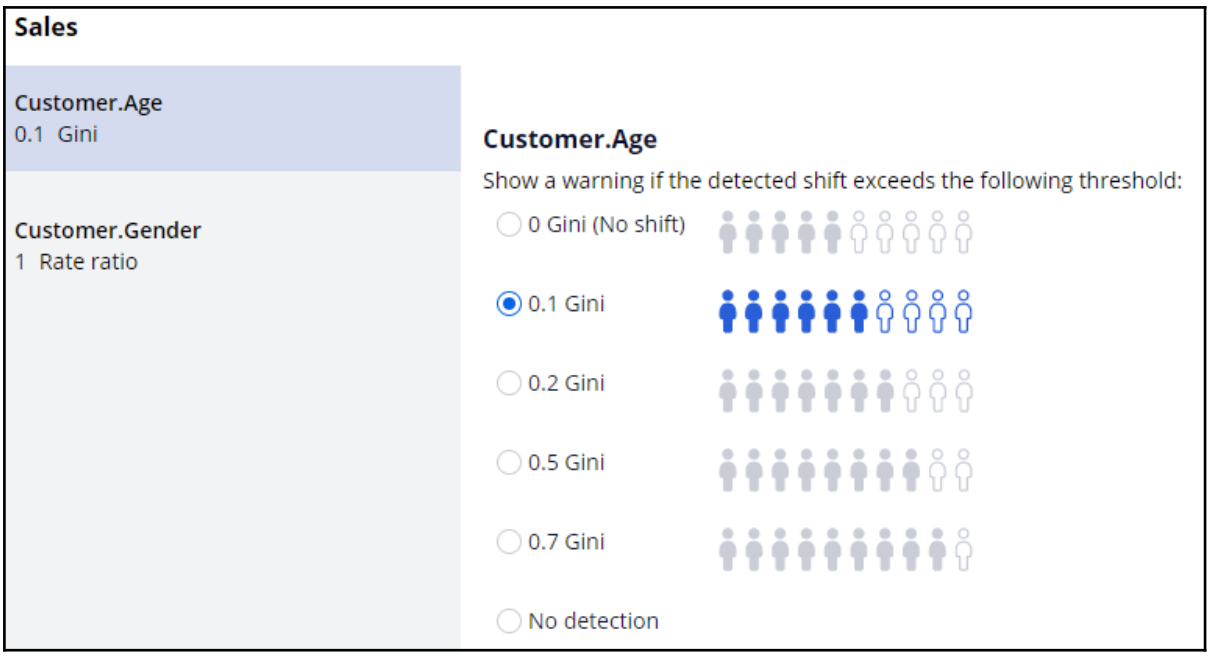

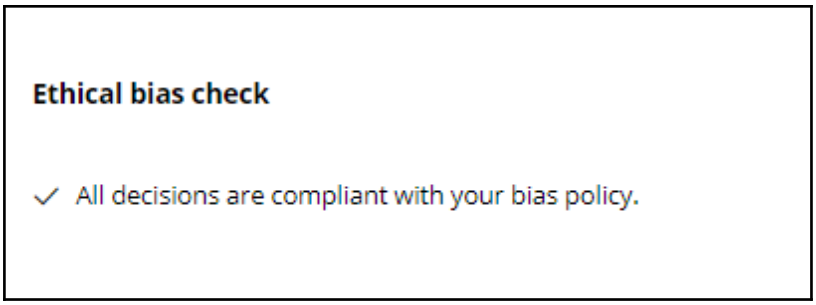

Since the age bias threshold was increased, no bias should be detected.

Filtered by: **Simulation ID = DemoScenario-1** and **Bias result time > Jul 23, 2020 6:17:01 AM**

Displaying 4 records

Bias detected	Issue[1]	Group[2]	Action[3]	Bias field[4]	Category
false	Sales	CreditCards	PremierRewardsCard	.Customer.Age	--
false	Sales	CreditCards	RewardsCard	.Customer.Age	--
false	Sales	CreditCards	RewardsPlusCard	.Customer.Age	--
false	Sales	CreditCards	StandardCard	.Customer.Age	--

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.

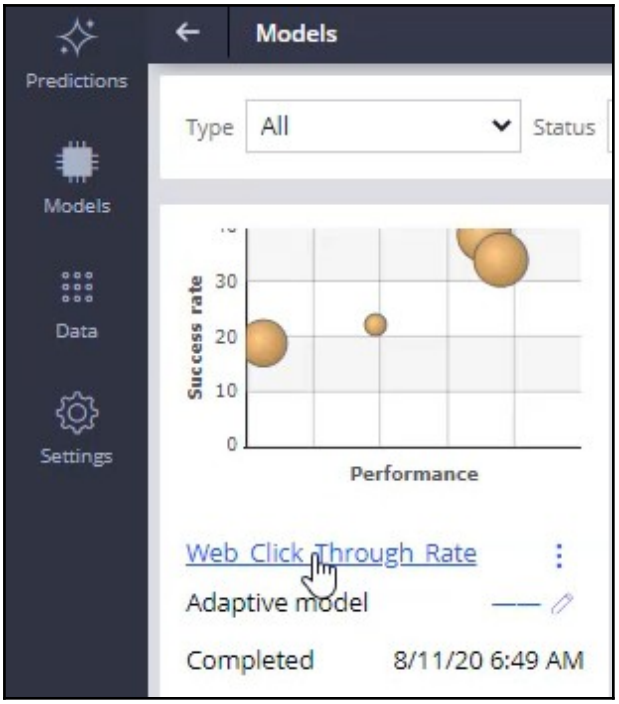

Now, open the Premier rewards card model report to view the predictors.

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.

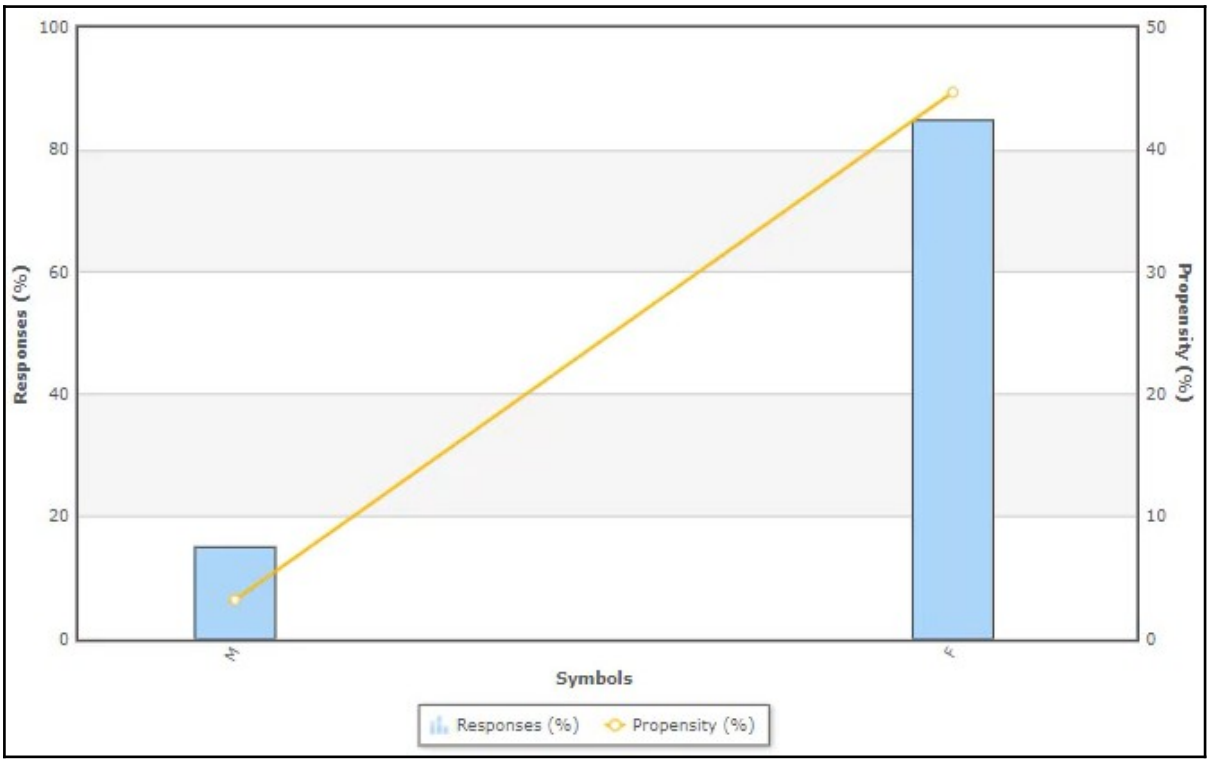

If you do not want the adaptive model to learn based on gender, ensure gender is not included as a predictor.

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.

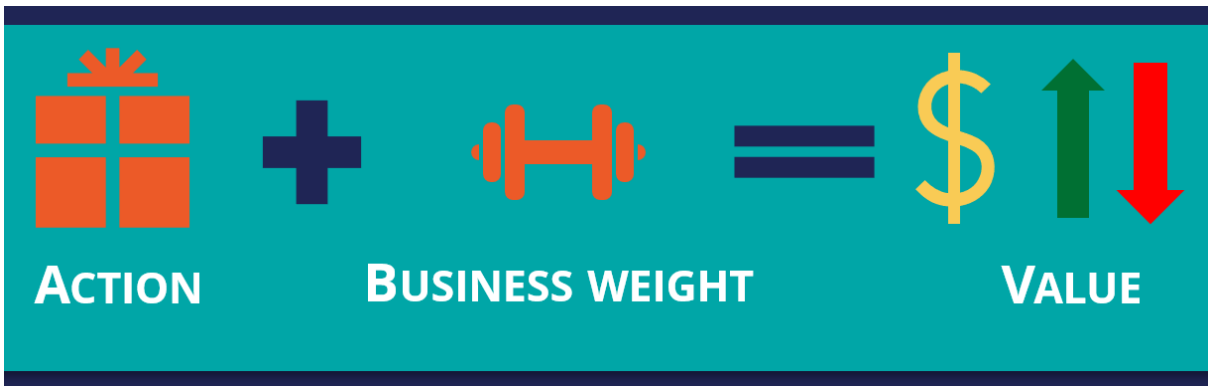

The Scenario Planner landing page shows three useful Key Performance Indicators: projected value capture, responses from Interaction History (including accepts, sends, and impressions), audience simulation counts (or projected reach), and projected accepts based on adaptive models.

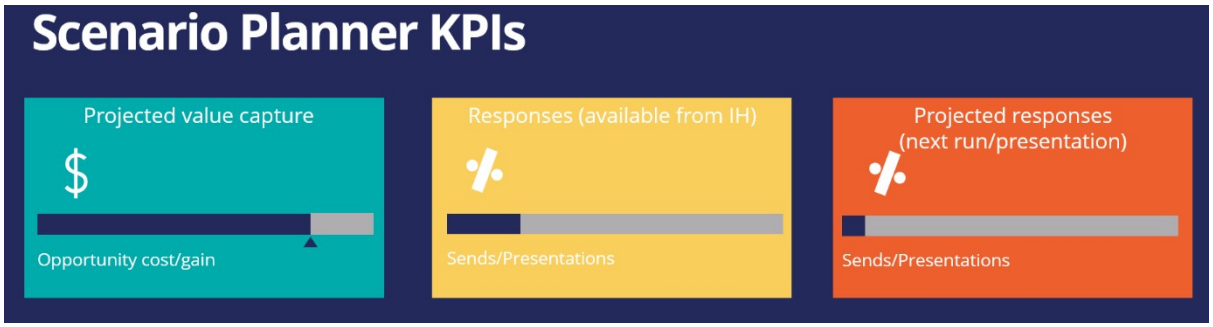

Depending on the selected filters, the Scenario Planner view can display the results for a specific channel, action, group, or issue.

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.

	Top 3 actions	Action Propensity	Action value	Expected action value	AI order
Customer1	Action1	0.09	51	4.59	1
	Action3	0.08	46	3.68	2
	Action2	0.07	49	3.43	3

With the action propensity (p) and the action value (v), the expected action value is computed using the formula p*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*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.

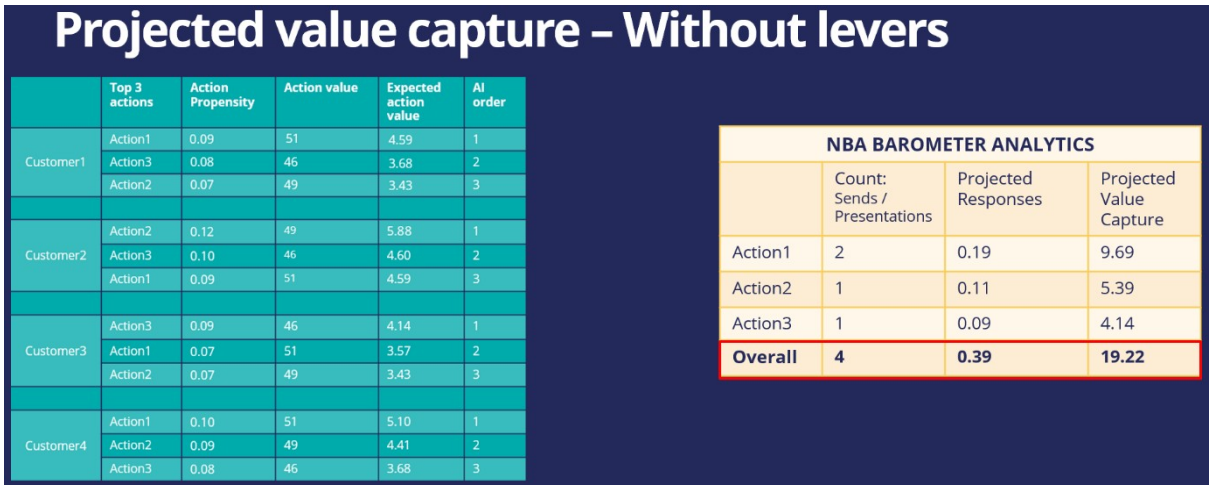

Now, consider another situation where the business has assigned levers to the actions to boost them based on the business requirement. For example, the business has assigned a business lever of 0%, 30%, and 10% to Action1, Action3, and Action2, respectively.

Applying levers affects the expected value and the action ranking. The adjusted expected value of the action is calculated using the formula, p*v*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.

Projected value capture – With levers

	Top 3 actions	Action Propensity	Action value	Expected action value	AI order	Lever	Adjusted expected value	New order
Customer1	Action1	0.09	51	4.59	1	+ 0%	4.59	2
	Action3	0.08	46	3.68	2	+ 30%	4.78	1
	Action2	0.07	49	3.43	3	+ 10%	3.77	3
Customer2	Action2	0.12	49	5.88	1	+ 10%	6.46	1
	Action3	0.10	46	4.60	2	+30%	5.98	2
	Action1	0.09	51	4.59	3	+ 0%	4.59	3
Customer3	Action3	0.09	46	4.14	1	+ 30%	5.38	1
	Action1	0.07	51	3.57	2	+ 0%	3.57	3
	Action2	0.07	49	3.43	3	+ 10%	3.77	2
Customer4	Action1	0.10	51	5.10	1	+ 0%	5.10	1
	Action2	0.09	49	4.41	2	+ 10%	4.85	2
	Action3	0.08	46	3.68	3	+ 30%	4.78	3

NBA BAROMETER ANALYTICS			
	Count: Sends / Presentations	Projected Responses	Projected Value Capture
Action1	1	0.10	5.10
Action2	1	0.12	5.88
Action3	2	0.17	7.82
Overall	**4**	**0.39**	**18.80**

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.

Projected value capture (without levers) = 19.22

NBA BAROMETER ANALYTICS			
	Count: Sends / Presentations	Projected Responses	Projected Value Capture
Action1	2	0.19	9.69
Action2	1	0.11	5.39
Action3	1	0.09	4.14
Overall	**4**	**0.39**	**19.22**

Projected value capture (with levers) = 18.80

NBA BAROMETER ANALYTICS			
	Count: Sends / Presentations	Projected Responses	Projected Value Capture
Action1	1	0.10	5.10
Action2	1	0.12	5.88
Action3	2	0.17	7.82
Overall	**4**	**0.39**	**18.80**

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.

To analyze the simulations using Scenario Planner, you first need to run an audience simulation using the engagement policies and arbitration as the scope, with second-pass simulation enabled.

In the navigation pane click **Discovery** to open the Scenario Planner and analyze one or more simulations.

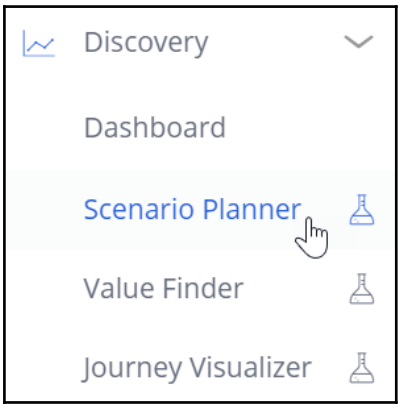

Select the scope of the run. The scope can be at the all issues level or for a particular issue and group. In this case, select the **Sales** issue and the **CreditCards** group. Now, select the audience simulation run that you executed earlier.

You can also select a simulation without levers or another audience 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.

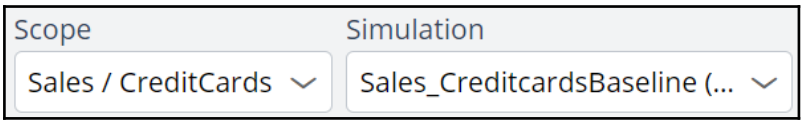

A Scenario Planner simulation displays metrics that show the projected value capture based on the configurations within Next-Best-Action Designer, the actual responses from the interaction history, and the projected reach and responses calculated from the results of the audience simulation.

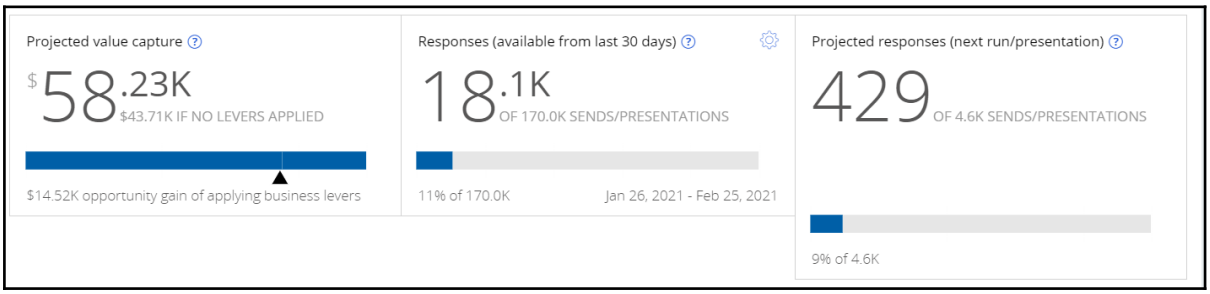

In this tile, you see the projected value that represents the estimated value that the upcoming next-best-actions (outbound run and inbound impressions) bring with the currently-selected configuration and the opportunity cost or gain of applying the levers. Prioritizing a specific action by increasing 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.

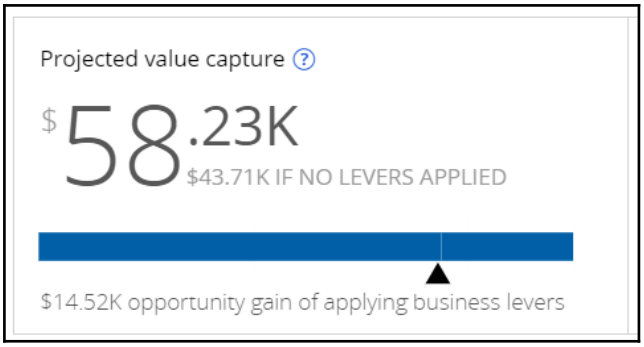

The projected value capture is $43.71K without levers, represented by the blue bar until the pointer. In addition, there is a $14.52K increase in value due to the implementation of business levers, amounting to a total of $58.23K projected value capture.

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.

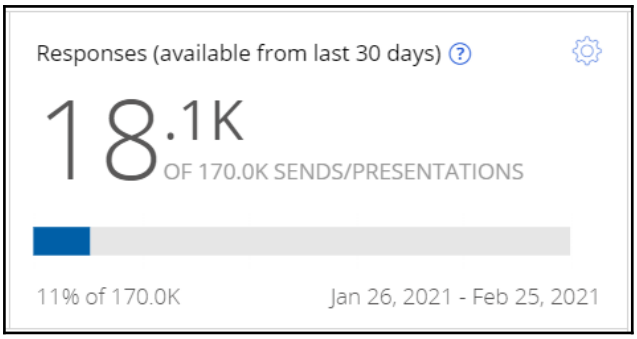

You can change the timeframe for historical data from the last 30 days to the last 60, 90, 180, or 365 days. You can also aggregate data from the last full quarter up to a full year.

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 audience simulation run results.

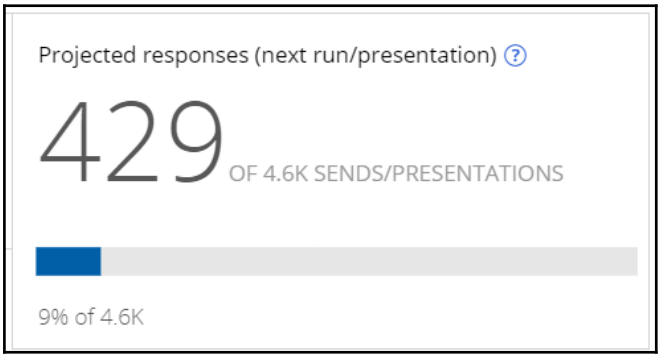

The following table displays the results per the selected simulation, issue, group, action, and channel. In this case, you see the results for the actions in the CreditCards group.

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 audience simulation without levers and compare the results.

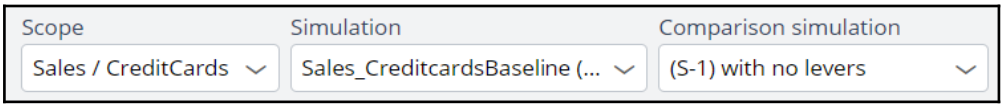

Note that in comparison mode, the Scenario Planner indicates the projected value and reach in green, as there is an increase when the lever is applied.

You can click the individual section menu to view the projected data comparison.

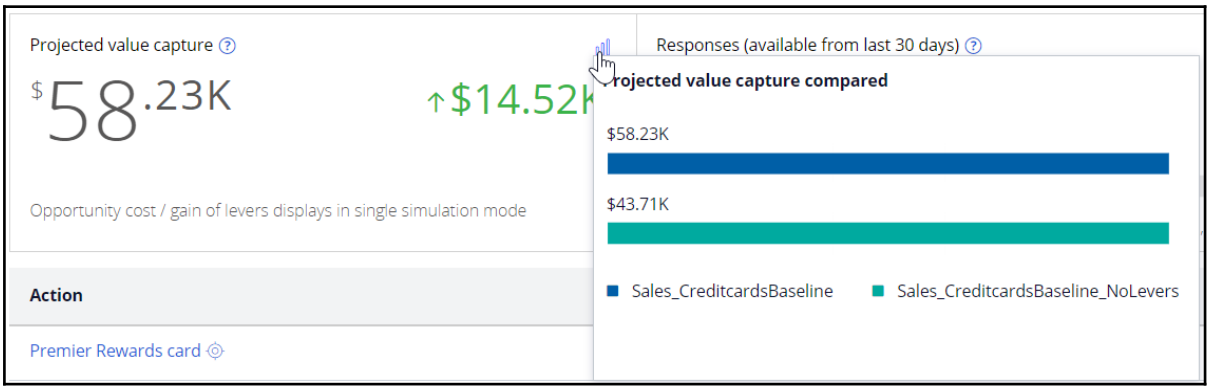

Similarly, in the table, the projected reach and responses are indicated per action with a green and red arrow to show which action is presented more and which action is presented less due to the business lever implementation. This information gives the business an idea

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.

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