



This PDF is generated from authoritative online content, and is provided for convenience only. This PDF cannot be used for legal purposes. For authoritative understanding of what is and is not supported, always use the online content. To copy code samples, always use the online content.

AI Core Services

9.0.013.01

12/24/2025

9.0.013.01

AI Core Services Release Notes

9.x AI Core Services is part of 9.x starting in **9.0.006.05**.

Release Date	Release Type	Restrictions	AIX	Linux	Solaris	Windows
12/18/18	General	Under Shipping Control		X		

What's New

This is a general release for this component. This version was first released as an Update on 10/25/2018. For availability of this release, contact your Genesys representative.

This release includes the following new features and enhancements:

- AI Core Services (AICS) now requires HTTPS communication by default. Note that this change also requires you to adjust your configuration for Agent State Connector and the URS instance implementing the URS Strategy Subroutines component. For instructions, contact your Genesys representative. (PRR-1940)
- Dataset handling has been made significantly faster by means of the following improvements:
 - For the initial data upload, this release introduces the Minio container. The Minio container is installed by default during the AICS deployment process and requires no user configuration. For more information, see The Minio Container in the AICS [High-Level Architecture](#) section of the *Genesys Predictive Routing Deployment and Operations Guide*. To take advantage of this improvement, you must also configure the new **S3_ENDPOINT** environment variable.
 - The Dataset import to MongoDB now uses a multithreaded process. A Dataset is now split into smaller chunks (minimum size, 50,000 rows), which MongoDB reads simultaneously. To speed up the Dataset import, Genesys recommends that you first create and initialize the schema by uploading a small sample Dataset (10-100 rows). Then append the rest of the data in chunks of up to 1 million rows. Appending data consisting of 100 features/1 million rows (708 MB) to the synced schema now requires no more than 40 minutes, including calculation of

Helpful Links

Releases Info

- [List of 9.0.x Releases](#)
- [9.0.x Known Issues](#)

Product Documentation

- [Genesys Predictive Routing](#)

Genesys Products

- [List of Release Notes](#)

cardinalities.

- A new environment variable, `HOST_DOMAIN`, now enables you to specify the public IP address or the name of the host where GPR is deployed or, for HA environments, the IP address of the load balancer. This variable also affects how URLs are formed in the emails sent by the GPR application, in which the hostname part of the URL is taken from the `HOST_DOMAIN` value. For instructions on how to configure this variable, see [Set Values for Environment Variables](#) in the *Genesys Predictive Routing Deployment and Operations Guide*.
- The NGINX container has been removed from AICS. NGINX is an optional load balancer that had been provided for use only in test environments.
- The *Sizing Guide* for Genesys Predictive Routing (GPR) has been entirely reworked and expanded. It now provides a simplified, comprehensive set of sizing guidelines that generates hardware requirements for all GPR components. See [Sizing for Premise Deployments](#) for a link to the worksheet and instructions for its use.
- The explanation for how to use Composite Predictors has been revised and clarified. See [Composite Predictors](#) for more information.
- AICS now checks whether a Dataset is connected to a Predictor before deleting the Dataset. This change is to prevent accidental deletion of a Dataset that is used by a Predictor, in which case you cannot make updates to the Predictor in future.
- The GPR web application and GPR API now use the same process to create Agent and Customer Profile schemas. As a result, the workflow in the web application has changed slightly to correspond with the workflow used in the API. To create Agent and Customer Profile schemas, you now do the following:
 1. Upload a small dataset sample to establish the schema structure. This dataset is used only to discover the fields and infer their datatypes. No data is uploaded.
 2. Make any necessary changes to the schema, such as correcting datatypes and setting the ID field.
 3. Accept and sync the schema.
 4. Upload (or append) data. AICS loads the data to the schema and updates cardinalities.

Previously, in the web application, AICS would discover the schema and upload the data in one step.

For a detailed discussion of the procedure for creating the Agent Profile schema, see [Configuring Agent Profiles](#) in the *Genesys Predictive Routing Help*.

- The `LOG_LEVEL` environment variable has been added to the **tango.env** configuration file. By default, it is set to `INFO`, which is a minimal logging level, adequate for most circumstances. If you do need to increase the log level, set the `LOG_LEVEL` variable to `DEBUG` and then restart GPR. Note that setting `LOG_LEVEL` to `DEBUG` considerably increases log files sizes.
- This release upgrades AICS to Python 3.6 from Python 2.7. Python 2.7 will not be maintained later than 2020. See the [Upgrade Notes](#) section below for the procedure to upgrade your existing Models to be compatible with the new Python version.
- AICS now performs automatic cleanup processes which should maintain an adequate amount of free disk space. Releases prior to 9.0.013.01 require you to manually perform the clean-up procedures. For instructions, see [Clean Up Disk Space](#).
- Memory handling for MongoDB was improved In this release. By default, MongoDB consumes all available RAM on a server, allocating half for its cache. GPR now restricts MongoDB to 8 GB RAM, of which 4 GB is used for the cache.
- The [Lift Estimation report](#) has been improved in the following ways:
 - You can now export analysis results as a CSV file.

- You can toggle the report display between graph and tabular formats. Graph values can be seen in a table, if the results show a positive lift the corresponding values are highlighted.
- The order in which the report sub-tabs appear has changed to show the Aggregated display first, followed by tabs to display the results for each feature.
- The display for the analysis reports has been improved in the following ways:
 - The name of the user who generated the report now appears on the report thumbnail and the report display.
 - Additional report metadata appears, such as the ID for the associated Dataset, the date range, and (when relevant) the name of the associated Predictor.
- The line chart on the **Dataset Trend** tab (accessed from the top navigation bar) has been replaced with a bar chart, which more accurately depicts the distribution of values per unit of time.
- The display on the **Dataset Schema** tab (accessed from the left-hand **Settings** navigation bar) has been improved in the following ways:
 - You can now choose to view only the fields you set to **Visible**.
 - Additional information about the Dataset appears, such as the filename of the file from which you imported the Dataset data and the Dataset ID.
 - The check box to the left of each Dataset row no longer appears after the Dataset is synchronized. They are not used after synchronization.
- The display on the window containing a table listing all Datasets (accessed from the left-hand **Settings** navigation bar) has been improved in the following ways:
 - The columns were rearranged for better usability.
 - A new column shows the status indicators for each Dataset.
 - A new column displays the filename of the CSV file used to create the Dataset.
- When you click the gear icon on the top navigation bar to open the **Settings** left-hand navigation bar, the display opens with the **Customer Profile** tab active.
- This release includes multiple improvements and additions to the GPR API, described in the following section:

For complete information on accessing and using the GPR API, see the following reference: [*Predictive Routing API Reference*](#) (Requires a password for access. Please contact your Genesys representative if you need to view this document.)

 - The new `apply_sync` and `accept_sync` commands enable you to synchronize and accept a Dataset. These commands, together with the existing Dataset functionality, provide the ability to upload, sync, and accept a schema, as well as set the timestamp field and make sure the datatypes are correct. Use the following commands to create a Dataset using the GPR API:
 1. The following command synchronizes the schema: `apply_sync`
 2. The following commands enable you to choose the timestamp column and change datatypes: `name`, `schema`, `visible_fields`
 3. The following command accepts the schema: `accept_sync`
 4. The following commands enable you to check Dataset status, using the following fields to access specific information: `sync_status`, `sync_progress`, `accept_progress`
 - The API now enables you to check the status of various jobs, including Predictor creation and Model creation.

- The new `import_progress` command enables you to track the progress of a Dataset upload.
- You can now purge all data from a Dataset using the new `clear` command. This command retains any existing Predictors and Models, but clears out old Dataset data. You can then add new data, without having to create new Predictors.
- You can now use the GPR API to check the status of Generate and Purge Data jobs using the `check_status` command.
- GPR now supports nested queries in dictionaries during scoring using the `action_filters` parameter. For example, you could use the following query: `\{'action_filters': 'skills.skill1 > 123'\}`.
In addition, you can now change which character to use as a separator denoting nesting in dictionaries. This might be necessary if a separator character is part of a field name because it will interfere with nested queries. To change the nested query separator character, make a PUT request to the predictors endpoint with the following parameter: `\{'nested_query_char': '.'\}`.

Resolved Issues

This release contains the following resolved issues:

In Model configuration, the Test/Train data split was improved to ensure that the percent of rows actually allocated for Train vs. Test are in sync with the percentages you configure. The web application also validates the percentages you specify for the split, and does not allow invalid percentages, such as specifying either 0% or 100% for the training section. In addition, the way the end-date for the data period used to generate a Predictor was calculated has been modified to ensure that all data from the specified Dataset is used. (PRR-3564, PRR-3444)

When you delete a Predictor, GPR now also deletes all associated data that was used to train that Predictor. This ensures that out-of-date Predictor training data is not retained after the Predictor is deleted. (PRR-3484)

When a new user account is opened, the email with instructions for activating the account and resetting the password to a secure value now arrives. To enable this correction, configure the new environment variable, `HOST_DOMAIN`, which enables you to specify the public IP address or host name of the host where GPR is deployed or, for HA environments, the IP address of the load balancer. For instructions on how to configure this variable, see [Set Values for Environment Variables](#) in the *Genesys Predictive Routing Deployment and Operations Guide*. (PRR-3445, PRR-1564)

You can now export a report that has non-ASCII characters in the report name. Previously, in this scenario GPR generated a `UnicodeEncodeError...` message. (PRR-3404)

If you run a Feature Analysis report on a dataset that uses a low-cardinality field as the target metric, the report now correctly displays the features ordered by rank. Previously, the sub-reports were displayed in alphabetical order. (PRR-3356)

Non-ASCII characters are now properly displayed in the pop-up window that displays the actual cardinality values in Agent Profile and Customer Profile schemas if the feature type is "dictionary". (PRR-3334)

If you run a Feature Analysis report using the GPR API and it fails for any reason, GPR now generates an error message and removes the failed job. Previously in this scenario, the spinning icon in the web application that indicates report generation is underway continued to spin indefinitely. (PRR-3324)

If your environment contains a number of predictors based on large datasets, you no longer encounter an out-of-memory error message when you try to open the Predictors Settings page. Previously, for example, having six predictors, each based on a dataset with 500 columns, triggered the following error message: Sort operation used more than the maximum 33554432 bytes of RAM. Add an index, or specify a smaller limit. (PRR-2945)

GPR now correctly scores agents when you use a Composite Predictor. Previously, no score results were returned when you used a Composite Predictor for scoring. (PRR-2924)

The **Agent Profile** window (accessible from the left-hand **Settings** navigation bar) no longer becomes unresponsive if the Agent Profile schema contains a significant number of high-cardinality fields. (PRR-2871)

Agent features now appear correctly on the **Agents** tab (accessed from the top navigation bar) of the GPR web application when you create the Agent Profile schema using the GPR API. (PRR-2748)

When you use the GPR API to set up a Predictor, you can specify a score expression, such as `p_score*100`, that manipulates the value to be returned for a scoring request. The scoring response includes min, max, median, and mean scores, which now correctly take into account the score expression you configured for the Predictor. (PRR-2468)

AICS now correctly handles score requests based on Composite Predictors. Previously, such scoring requests failed and generated the following error message: Actions data collection is not indexed. (PRR-2461)

Upgrade Notes

In release 9.0.013.01, GPR uses an updated version of Python. Use the following procedure to upgrade to release 9.0.013.01.

1. Deploy the IP for AI Core Services release 9.0.013.01, following the instructions in the *Genesys Predictive Routing Deployment and Operations Guide* for a **Single-Server** or **high availability** environment, as appropriate.
 2. Open a terminal window and go to the **Tango** container using the following command:
-

```
docker exec -it tango /bin/bash
```

3. Switch to the **gpr** directory using the following command:

```
cd src/gpr
```

4. Run the following upgrade script:

```
MODE=prod python3.6 py3_upgrade_script.py
```

5. Restart GPR using the following commands:

```
exit
```

This command exits the Tango container.

```
cd IP<version_number>
```

```
bash scripts/restart.sh
```

6. After performing this upgrade, you must retrain all existing Models. Instructions for training Models are available in [Configuring, Training, and Testing Models](#) in the *Genesys Predictive Routing Help*.

Your AICS deployment should now be complete.

Downgrading from 9.0.013.01

If you experience problems that force you to change back to the older version of Python, use the following procedure to remove GPR release 9.0.013.01 and the updated version of Python.

Important

Use this procedure only if you are downgrading from release 9.0.013.01.

1. Open a terminal window and go to the **Tango** container using the following command:

```
docker exec -it tango /bin/bash
```

2. Switch to the **gpr** directory using the following command:

```
cd src/gpr
```

3. Run the following downgrade script:

```
MODE=prod python3.6 py3_upgrade_script.py down
```

4. Re-install your previous release, following the instructions in the *Genesys Predictive Routing Deployment and Operations Guide* for a [Single-Server](#) or [high availability](#) environment, as appropriate.