May 2026

Release Notes

Version 17.0

Table of Contents

1. General

1.1 Get answers faster with the Aurora Support Assistant in the Aurora interface

The Aurora interface now includes **the Aurora Support Assistant**, an AI-driven chat assistant that helps you find accurate answers to Aurora-related questions. It reduces the need to manually search through documentation by providing direct, context-aware responses within the platform.

The Support Assistant is trained on the Aurora Help System, Knowledge Base, and official resources. It supports a wide range of queries, including platform concepts, configuration guidance, and technical behavior, helping you access reliable information during your workflow.

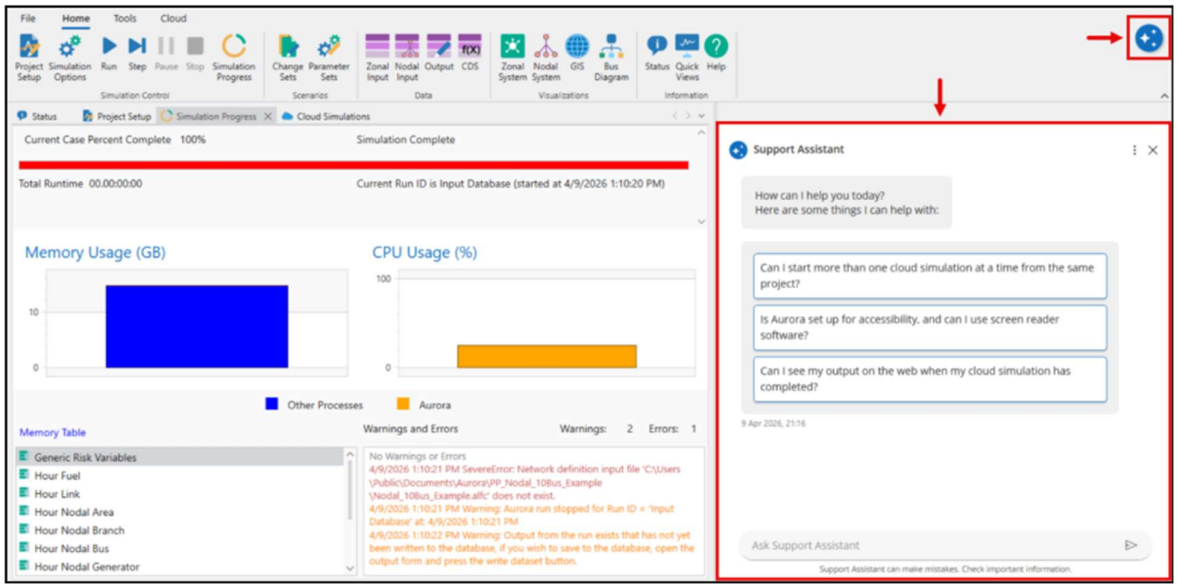

You can use the Support Assistant to ask a wide range of Aurora-related questions, such as:

- Which database types does Aurora support?
- How is inflation applied?
- What is the purpose of aggregate area tables?
- Where can you configure parallelization settings?
- What are the hardware recommendations for Aurora?
- How does emissions logic work?

The Support Assistant is accessible directly from the **Aurora interface** and remains available across the platform, ensuring that you can retrieve information whenever needed.

➤ Key benefits

- **Instant access to information:** Retrieves answers from Aurora documentation in real time.
- **Reduced search effort:** Minimizes time spent navigating help resources.
- **Integrated workflow:** Access support without leaving the Aurora interface.

For more information, refer to the detailed **Support Assistant** help article

---

1.2 Merge RMT records into the Resource Table from the input ribbon

A new **Merge RMT** button is available in the **input ribbon** to merge RMT records into the active Resource Table. This feature streamlines the process of creating and updating resource records based on RMT data.

Aurora automatically creates new resource records and updates existing ones during the merge. The result is saved as a new change set in the applied state.

Refer to the following image for an example:

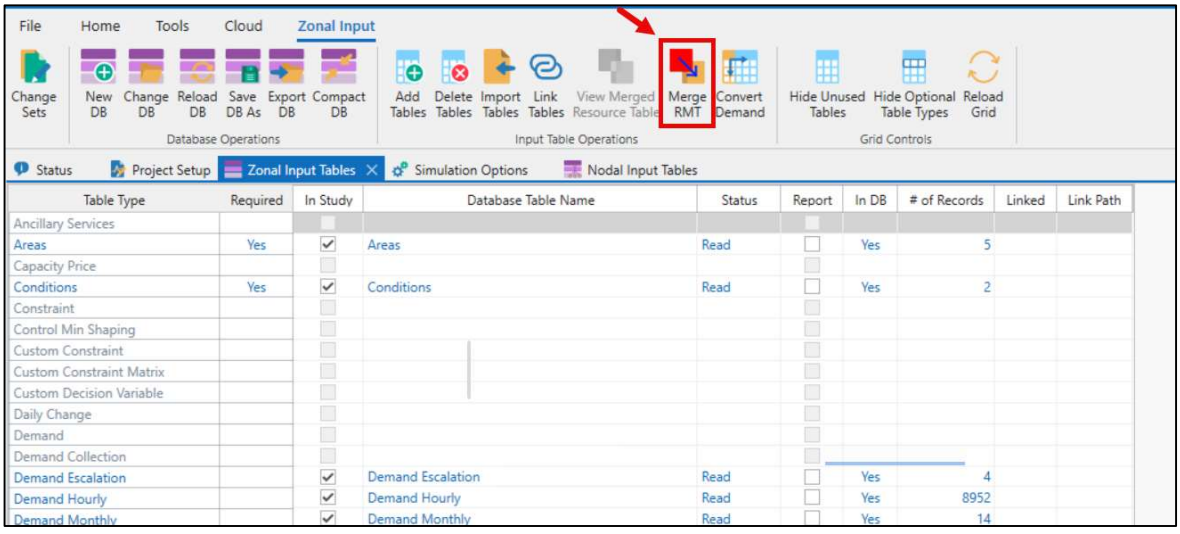

For scripting users this functionality can also be accessed through a new scripting command using the AC.MergeRMTIntoResourceTable function.

1.3 Apply a dedicated parameter set for final dispatch after LT studies

Aurora now supports a **Post LT Parameter** Set that applies a separate configuration specifically to the final dispatch run following a long-term (LT) study.

This parameter set allows overriding project-level settings, such as dispatch options, input tables, commitment settings, and output configurations, only for the final dispatch run. It does not affect the iterative LT optimization phase.

The Post LT Parameter Set is configured at the study case level alongside the standard Parameter Set used at the start of the simulation. When defined, it overrides Post LT dispatch settings specified in the General Information table.

Refer to the following image for an example:

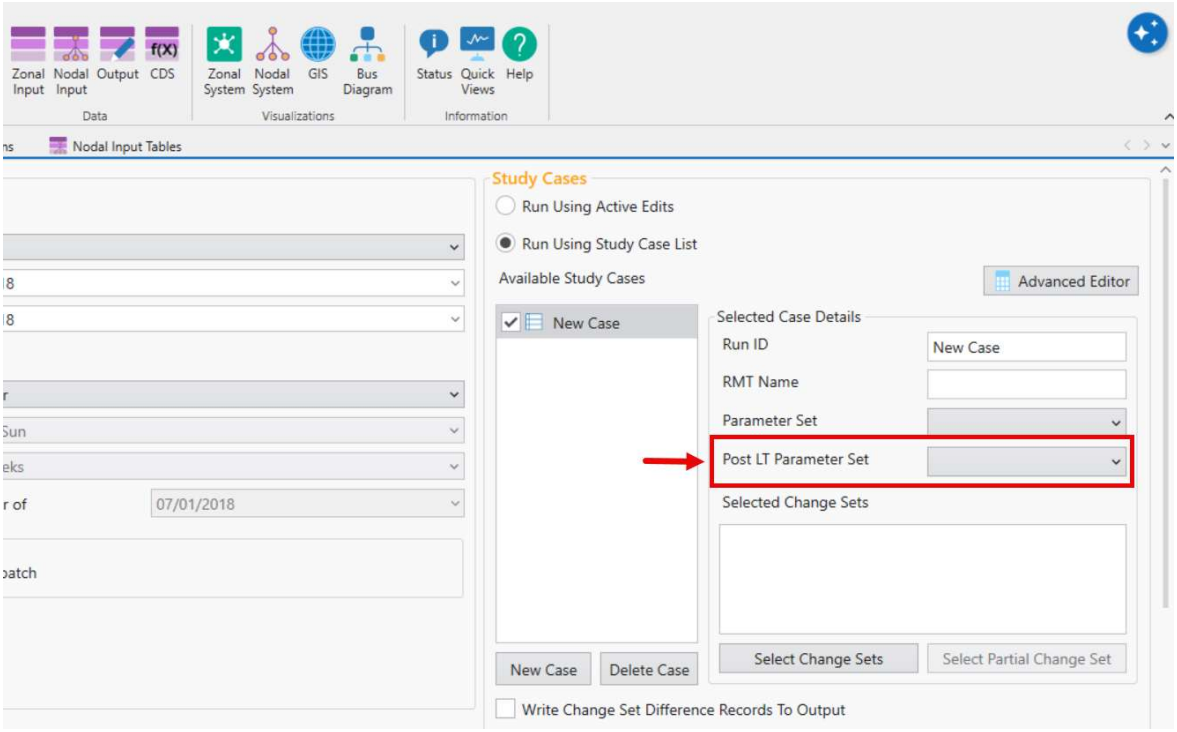

➤ Key Benefits

- Separation of LT optimization and final dispatch configuration
- Targeted control over final run behavior
- Ability to override project settings without impacting earlier iterations
- Flexible configuration at the study case level

1.4 Option to use Table Type Prefixes for Input to Output Write

Aurora now allows control over whether table type prefixes are included when writing input tables to the output database.

This option is available as a **checkbox** in the **Output Database Options** under **Simulation** Options. It enables customization of output table naming based on table type inclusion or exclusion.

Refer to the following image for an example:

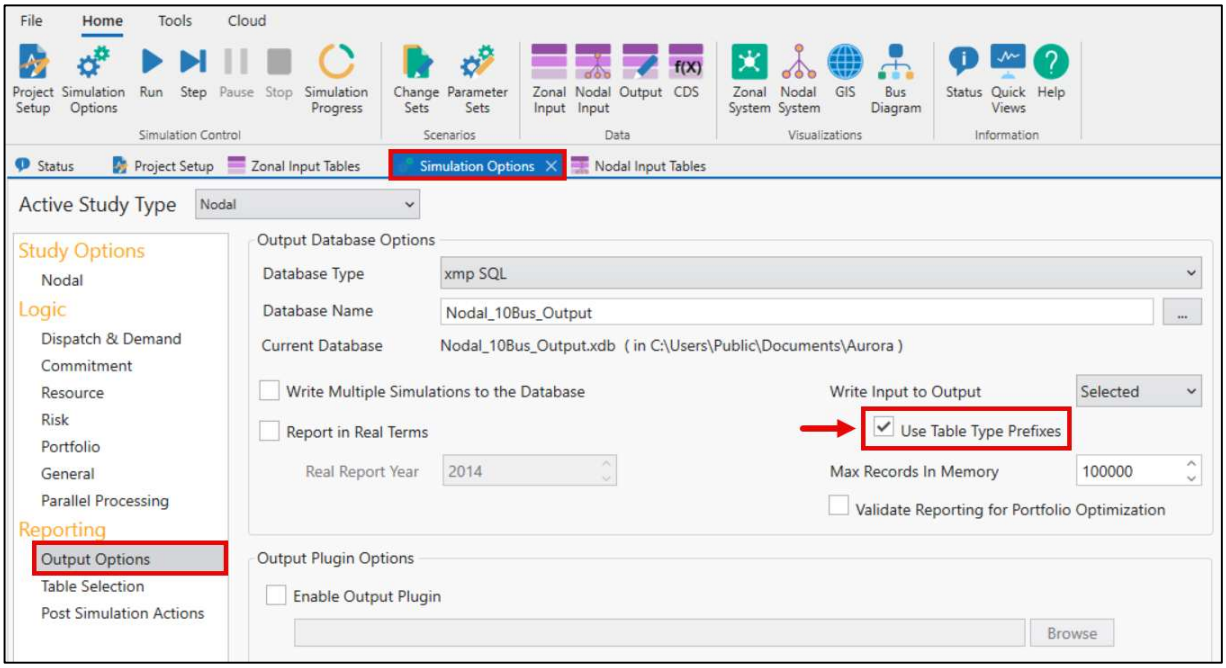

➤ Key Benefits

- Flexible control over output table naming conventions.
- Ability to simplify or standardize output structures.

1.5 Reduced installer size by removing Chromium dependencies

Aurora **no longer** depends on **Chromium browser** components for displaying the **Welcome Page** and help content. The Welcome Page now uses lightweight internal controls, while help content opens in the system's default browser.

This change removes Chromium library dependencies, reduces installer size, eliminates cache folders in working directories, and improves application startup performance.

➤ Key Benefits

- Reduced installer size due to removal of Chromium libraries.
- Faster application startup.
- Elimination of Chromium cache folders in working directories.

---

2. Cloud

2.1 Automatically download and combine cloud simulation outputs

Aurora now supports **automatic download** of **cloud simulation outputs**. When enabled, the system monitors selected cloud simulations and downloads the output database once the simulation completes successfully.

A new **"Automatically download output after Run"** checkbox is available in the **Run Cloud Simulation panel**, allowing this behavior to be configured at the time of submission.

The feature supports downloading **single** or **multiple simulation outputs,** combining multiple output databases when applicable, converting outputs to the selected format, and downloading associated log files. This reduces manual effort required after simulation completion and ensures outputs are readily available for analysis.

Refer to the following image for an example:

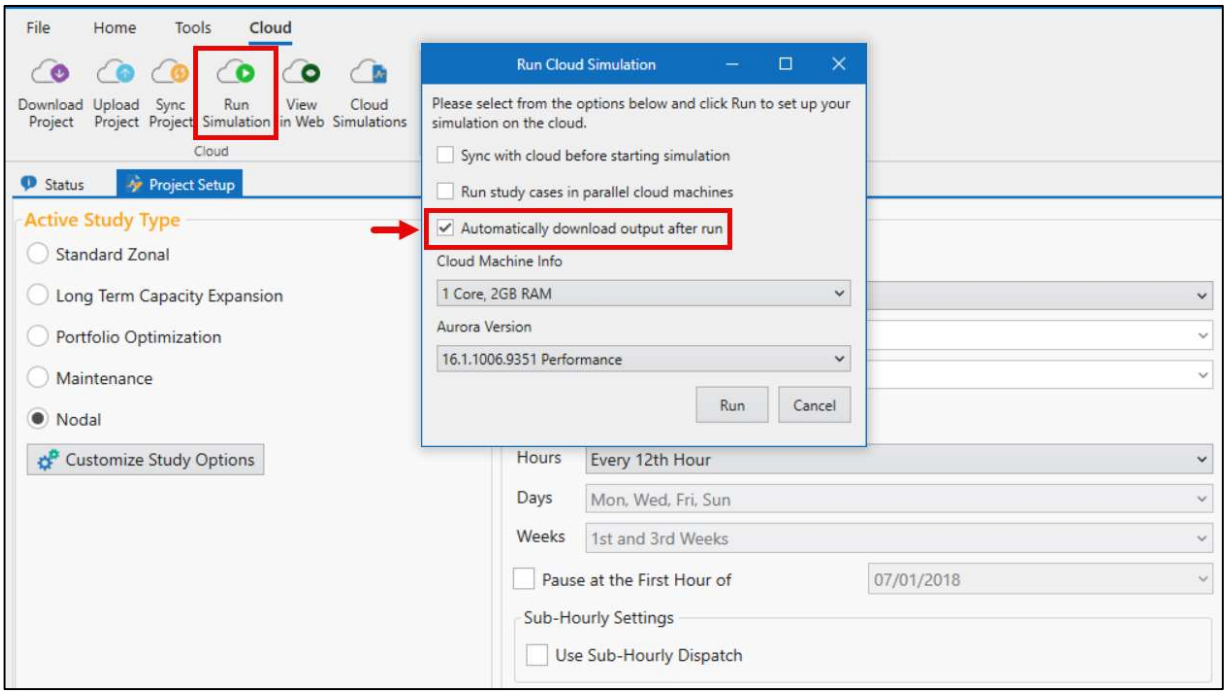

➤ Key Benefits

- Automatic retrieval of simulation outputs upon completion.
- Support for multiple simulations with combined output handling.
- Automatic download of log files.
- Reduced manual post-processing effort.

➤ How it works

1. Log in to Aurora using your credentials.
2. Navigate to the **Cloud** tab.
3. Click **Run Simulation**.
4. In the **Run Cloud Simulation panel**, select **Automatically download output after Run**.

When the simulation runs:

- Aurora monitors the simulation in the background.
- Upon completion, the output database is downloaded automatically.
- If multiple simulations are selected, Aurora can combine the downloaded output databases into a single result.
- If the configured output format is not Xmp SQL, Aurora converts the downloaded database to the selected format.

2.2 Track cloud uploads and downloads with real-time progress indicators

Aurora now provides visual progress indicators for **cloud uploads** and **downloads.** Progress bars display the status of each operation, including scenarios where multiple output databases are downloaded simultaneously.

Refer to the following images for an example:

➥ **For Upload:**

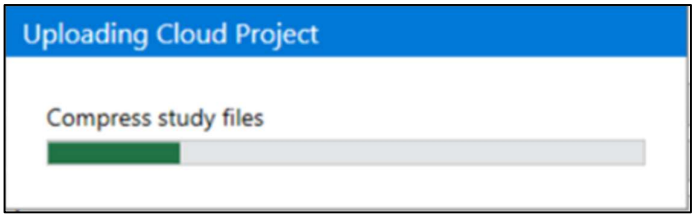

➥ **For Download:**

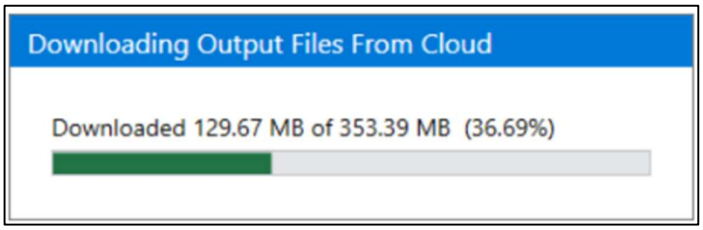

This update also expands the **Status Window** to display detailed, real-time progress messages during cloud study creation and upload.

The Status Window now shows each step as it occurs, including:

- creating local study files
- compressing study files

- uploading files to the cloud
- extracting files in the cloud
- performing the initial commit
- confirming successful study creation

This provides clear visibility into the full process and helps identify the current execution stage.

Refer to the following image for an example:

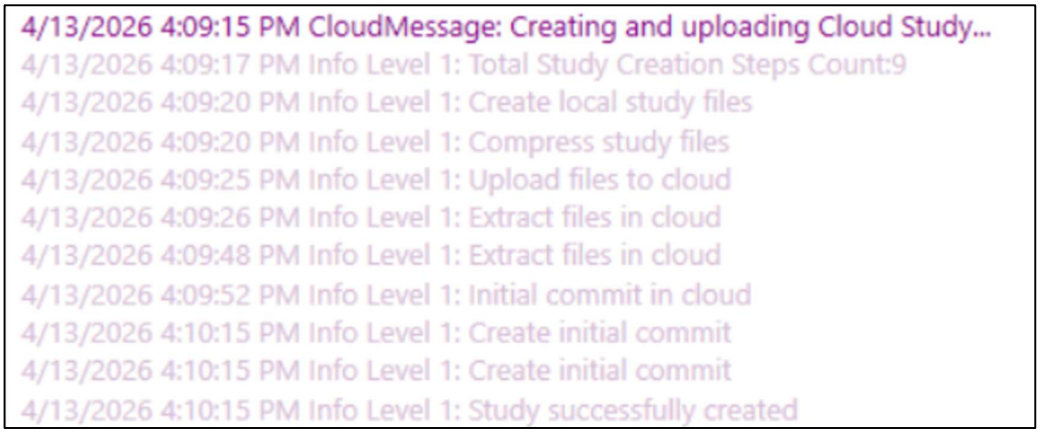

➤ Key Benefits

- Real-time visibility into upload and download progress.
- Easier identification of current processing stage.
- Better transparency in cloud workflows.

2.3 Re-run cloud simulations using the previous configuration

Aurora now allows **re-running cloud simulations** directly from the Cloud Simulations window. This option uses the **configuration of a previous simulation** to create and enqueue a new run without modifying the original simulation record.

The system retrieves key parameters such as Study ID, Model Name, Changeset ID, and Simulation ID, and rebuilds the simulation request using the Cloud SDK. A new simulation is created and submitted while preserving the original run.

➤ How it works

1. Select one or more simulations in the Cloud Simulations window.

2. Right-click and choose **Re-Run Simulation**.

After this:

- Aurora retrieves previous simulation parameters.
- A new simulation is created and enqueued.
- A confirmation message indicates the simulation has started.

Refer to the following image for an example:

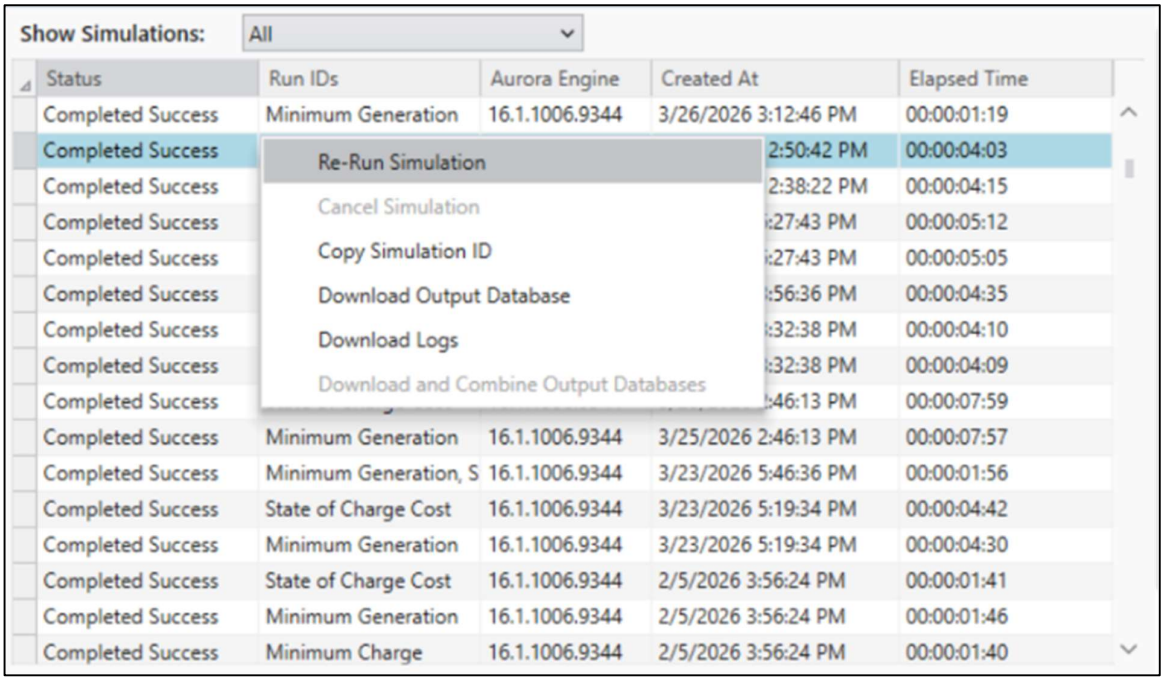

➤ Key Benefits

- Eliminates the need to manually reconfigure simulation parameters.
- Reuses existing simulation configurations.
- Supports single or multiple simulation selections.
- Speeds up repeated simulation workflows.

---

3. Scripting

3.1 Added new scripting functions for input and output database operations

Aurora introduces new scripting functions to support working with input tables and output databases. These functions enable deletion and import of input tables, saving output databases in different formats, and merging RMT records into the Resources table.

The additions expand automation capabilities and improve flexibility when managing data programmatically.

The following functions are now available:

- **ADS.DeleteInputDatabaseTable**: Deletes a specific input table from the input database
- **ADS.ImportInputDatabaseTable**: Imports a table into the input database
- **AC.SaveOutputDataBaseAs**: Saves the output database as a new file or converts it to a different format
- **AC.MergeRMTIntoResourcesTable**: Merges records from the RMT table into the Resources table

These functions enable more flexible and automated data workflows within Aurora scripting environments.

---

4. Logic

4.1 Apply control-min shaping using the new Control Min Shaping table

A new input table called the **Control Min Shaping** table is introduced. This table is an expanded feature in Aurora that allows users to shape time-varying input parameters (e.g., resource capacity) Rather than manually building an hourly time series to represent a shaped value, the Control Min Shaping method takes three descriptive statistics (minimum, average, and peak)

plus an optional hourly shape profile, and automatically derives a full set of hourly (or sub-hourly) values that honor all three targets simultaneously.

The technique is analogous to how Aurora shapes demand using the Control Min method and brings that same flexibility to any model field that supports a time series reference. Each row in the table defines a named shaping configuration. Once defined, a row is referenced from any compatible model field using the prefix **cm_** followed by the ID (e.g., **cm_MyShapeID**). This works exactly like other time series prefixes in Aurora such as **ts_**, **cds_**, **mn_**, etc.

4.2 Model ancillary services using multi-segment price-quantity demand curves

Aurora now supports modeling ancillary services using multiple price-quantity segments instead of a single flat price.

The newly added **Ancillary Service Demand Curve (ASDC)** feature enables you to model ancillary services using multiple price–quantity segments instead of a single fixed requirement and price.

This enables defining different prices for different reserve quantities.

A new column, **MW_Price**, is added to the **Ancillary Product Table.** When used with **MW_Requirement** and **Segment Definitions (sd_[Name]),** it defines demand curve segments.

The feature activates automatically when both columns reference Segment Definitions; otherwise, standard single-price behavior applies.

Refer to the following image for an example:

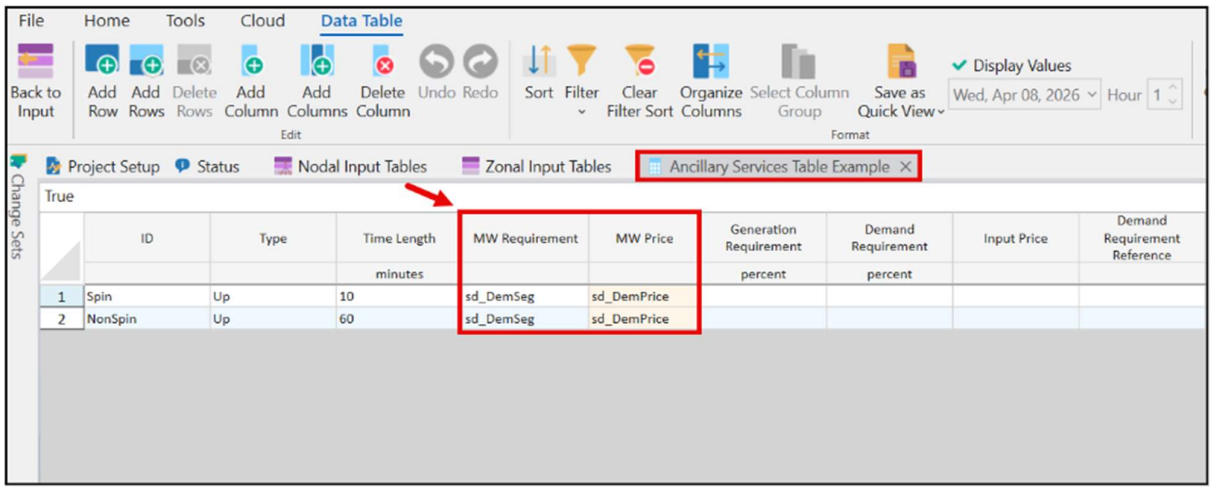

➤ Use Case

Consider a spinning reserve requirement:

Segment	MW Requirement (Cumulative)	Price ($/MW)
Segment 1	100 MW	50
Segment 2	200 MW	30
Segment 3	300 MW	10

How it works:

- The first 100 MW is valued at **$50/MW** (highest priority)
- The next 100 MW (100–200 MW) is valued at **$30/MW**
- The final 100 MW (200–300 MW) is valued at **$10/MW**

If the system only needs 150 MW:

- 100 MW is procured at $50
- 50 MW is procured at $30

This reflects how real markets value scarce reserves more highly than surplus reserves.

➤ Key Benefits

- **Realistic Market Representation**: Models reserve procurement based on marginal value instead of fixed pricing.
- **Flexible Demand Modeling**: Supports multiple price–quantity segments for each ancillary product.
- **Improved Economic Accuracy**: Reflects diminishing value of reserves as quantity increases.

For more information, refer to the detailed **Ancillary Service Demand Curve** help article.

4.3 Enable storage spill handling in commitment optimization

Aurora now includes storage spill functionality within Commitment Optimization. A new boolean column, **Allow Storage Spill**, is available in Resource, New Resource, and Fuel input tables.

When enabled, a spill variable is introduced into daily storage constraints. Spill values are reported in the **Storage_Spill_MWh** output column.

The feature applies only to storage resources with defined inflows. Resources without inflows will ignore this setting. This feature is beneficial when modeling hydro units as storage in Aurora.

4.4 Enforce Min Generation for Storage

A new Enforce Min Generation option is available in the Storage table. When set to True, this option forces the storage resource to always generate and prohibits charging in every modeled hour.

4.5 Control partial build decisions for new resources in LT studies

In a Long-Term (LT) optimization study, new resource candidates are typically treated as all-or-nothing decisions—the optimizer either builds the full unit or does not build it at all. The Allow Partial Build and Min Partial Build Ratio inputs provide finer control over this behavior. When **Allow Partial Build** is enabled for a new resource candidate, the optimizer is permitted to acquire a fractional portion of that resource's capacity rather than requiring the full unit. The **Min Partial Build Ratio** sets a floor on how small that fraction can be. If the optimizer decides to build the resource at all, it must acquire at least this minimum share (for example, a ratio of 0.25 means the resource must be built to at least 25% of its full capacity or not build). This feature is especially useful for modeling resources where incremental or phased development is realistic, such as wind and solar farms, allowing the optimizer to right-size capacity additions rather than forcing full-block investment decisions.

➤ Key Benefits

- Enables fractional resource build decisions
- Provides control over minimum build thresholds

- Improves flexibility in LT optimization outcomes

4.6 Uplift Payment Eligible flag for Resources

Gain greater control over uplift revenue calculations with this new setting: when enabled (default), resources remain eligible for uplift, while setting it to False explicitly excludes them. This provides added flexibility to refine how uplift is allocated across resources.

4.7 Add additional wheeling cost components using new input columns

Aurora introduces two new input columns, **Wheeling Cost Mod1** and **Wheeling Cost Mod2**, in the Link input table. These columns allow defining additional wheeling cost components that are included in the total wheeling cost calculation.

Both columns support time series inputs, enabling dynamic adjustment of wheeling costs over time. This provides greater flexibility in modeling cost structures associated with transmission flows.

5. Outputs

5.1 View detailed wheeling cost components in output reporting

Aurora extends wheeling cost reporting to reflect the newly introduced modifier inputs. The total wheeling cost is now broken into distinct components for improved transparency.

The output includes:

- **Wheeling_Cost_Base**
- **Wheeling_Cost_Aux1**
- **Wheeling_Cost_Aux2**

These components sum to the total wheeling cost and are aggregated across all reporting periods, providing a clearer breakdown of cost contributions.

5.2 Output Column for Days with Loss of Load in Zone and Pool

Introducing a new way to track reliability: this metric shows the number of days in a period when at least one loss of load event (LOLE) occurred in the Zone/Pool. Each day is counted only once regardless of multiple events. It provides a clear, high-level view of system stress and is available for non-hourly reporting.

5.3 Marginal Resource ID Column in Zone and Pool

This new output reports the input ID of the marginal resource in the Zone and Pool output making it easier to query the data and link up between multiple tables.

6. Supplemental Parameters

6.1 Auto Write Output For Stopped Simulations

Introducing automatic output recovery: when enabled, Aurora will write any remaining results to the database if a simulation is stopped. This helps preserve valuable data and ensures greater continuity in your analysis.

6.2 Skip Min Iters For Lowest Cost RMT

This new LT supplemental parameter is used in conjunction with **Choose Lowest Total Cost for RMT** where Aurora selects the lowest-cost iteration to determine final build decisions. With this setting enabled, only iterations that meet the minimum iteration requirement are considered—ensuring more stable, mature results drive the outcome.

6.3 Remove LT Rounding From Std Zonal

This new option has been introduced to control how rounding is applied to Daily LDC and storage constraint bounds in Standard Zonal studies when using the RMT table.

- **Default (False):** Rounding is applied, aligning Standard Zonal results with Final Long-Term (LT) run outputs. This ensures consistency between study types.

- **True:** Rounding is not applied, preserving legacy behavior and maintaining consistency with results from previous versions.

This enhancement gives users flexibility to choose between result consistency with LT studies or backward compatibility with earlier model behavior.

7. Solver

7.1 Updated Mosek solver to version 11.1.6

Aurora now includes an updated version of the Mosek solver. The solver has been upgraded to version 11.1.6 as part of this release.

7.2 Updated Gurobi solver to version 12.0.3

Aurora now includes an updated version of the Gurobi solver. The solver has been upgraded to version 12.0.3 as part of this release.

8. Support Information

All updates to Aurora are available in the Client Portal, https://portal.energyexemplar.com/.

Enter your username and password to select the desired update. Contact support@energyexemplar.com if you do not know the username and password. Be sure to review the Aurora Enhancement Log before downloading a new version of the software from the website.