pForecast User Manual

Software as a service solution to digitalize, simplify, and standardize how production forecasts are generated and utilized.

www.pForecast.com

Published: 29-04-24

Contents

Table of Figures

5

6

Overview

[Powersim Software](#) is a Norwegian company based in Bergen. Powersim Software provides a family of business tools tailored to forecast and analyze complex dynamic problems. pForecast, which has been developed for predicting oil and gas production, is one of these tools. pForecast is an uncertainty-centric software that digitalizes, simplifies, and standardizes how production forecasts are generated and utilized. pForecast performs a full lifetime simulation of the production and injection forecast, including historical data, in keeping with the industry's ever-increasing need for agility.

This manual covers the essential steps to get started with the pForecast software. For users seeking specialized support, please reach out to us at:

Email: pForecast@powersim.no

Phone: +47 55 60 65 00

Introduction

pForecast is a SaaS solution with a user-friendly and interactive interface, offering a simple and consistent methodology for production forecasting. It performs deterministic and stochastic analyses, and generates unbiased production forecasts taking the involved uncertainties into consideration for the entire field. pForecast consolidates the forecasting and acts as a common framework across assets.

pForecast, which uses Monte Carlo simulations at its core, improves analysis quality as uncertainties are an integral part of the forecasting, and not added as a deviation on a deterministic run. The software applies to both short and long-term forecasting.

pForecast's hierarchical structures consist of three levels: asset, forecast, and scenario. This manual shows how to configure each level and generate a comprehensive production forecast. A short description of each chapter is given in Table 1.

Table 1: Structure of the pForecast user manual.

Chapter	Description
1	Accessing pForecast
2	Asset Configuration
3	Forecast Configuration
4	Scenario Configuration
5	Running the Simulation
6	Corporate Analysis
7	Incremental Profiles
8	Corner Cases
9	Glossary

1. Accessing pForecast

pForecast is accessible through your web browser. The users need to request access to pForecast via their company IT support. A designated "Company Administrator" in your company will be defined as administrator when pForecast is introduced. The Company Administrator can assign different roles at the corporate level and on the asset level.

1.1. Company Administrator

In *Company Properties*, the Company Administrator can add new end users to the administrative level, either as an additional Company Administrator or as a Company Analyst.

For each role in the Company Properties, there are associated access rights:

1. **Company Administrator:** Allowed to administrate company, e.g., create assets and assign users to roles.
2. **Company Analyst:** Allowed to work with corporate roll-ups.

Please note that you will need to have at least two assets defined, each of them with at least one saved stochastic scenario, before you can run a corporate roll-up analysis, see Chapter 6.

Figure 1.1 shows where you can assign roles for the company.

Figure 1.1 Company Administrator properties.

1.2. Asset

You will at least need to define one asset that mirrors the structure in your organization model. Several users can use and add forecasts and scenarios for the same asset, depending on the settings of access rights in the next sub chapter.

Figure 1.2 shows how to create new assets for the end users.
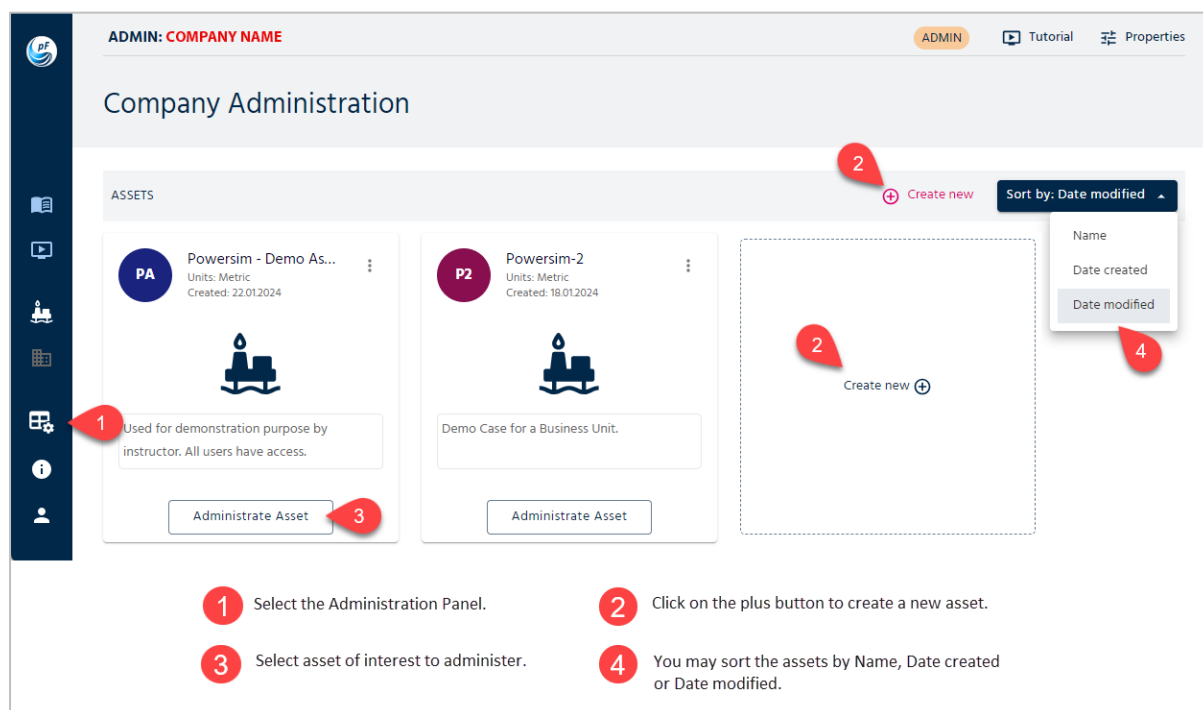
Figure 1.2: Creating the assets.

1.3. Asset roles

Once the asset is created, the Company Administrator can assign user roles for each asset. After selecting the asset, choose the Asset Roles tab and you will see the page shown in Figure 1.3, where you can assign the roles.

For each role, there are associated access rights:

1. Administrator: Allowed to administrate an asset, e.g., assign users to roles, and set up connectivity properties for assets
2. Contributor: Allowed to work with forecasts
3. Editor: Allowed to work with forecasts and flag forecasts for review
4. QC: Quality Controller is allowed to work with forecasts, flag forecasts for review, and approve/reject forecasts
5. Subscriber: Allowed to view forecasts
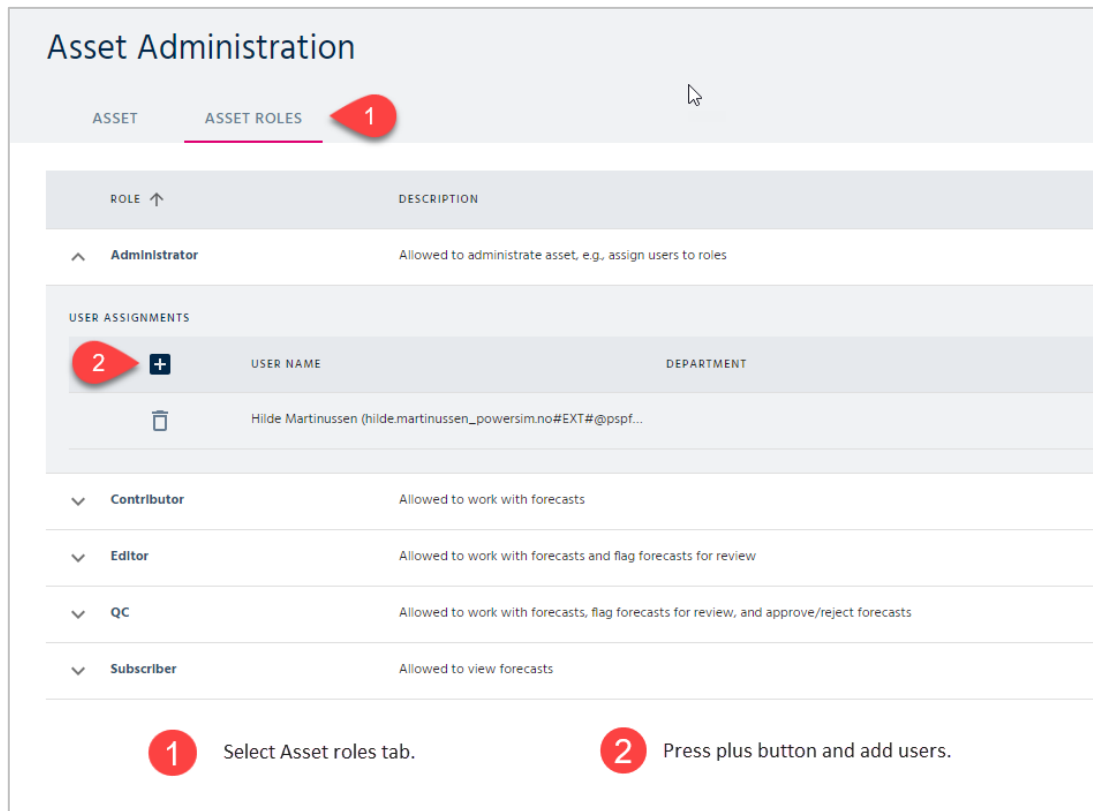
Figure 1.3: Defining users' roles.

1.4. Connectivity

pForecast can be configured to allow for scheduled imports of data from other systems and databases. If so, you will find a Connectivity tab under Asset Administration where you can control how input data will be imported into pForecast. Currently only historical production data is supported in this context.

As an administrator you can control how often and for which facilities historical production data should be uploaded to pForecast. This allows you to configure each of your assets with their own customised set of input data.

Connectivity must be established through your IT department. pForecast currently supports only one custom connectivity solution where connectivity settings can be fully controlled. This is a solution for Equinor's Production Data Mart (PDM), allowing import of historical production data as described in chapter 1.4.1. pForecast also supports connections to Aker BP's Cognite Data Fusion (CDF) for import of historical production data. As of yet, for this solution connectivity settings cannot be user controlled, and historical production data imports are scheduled to happen every night for all assets.

For a full-fledged connectivity solution, we have:

- The import is handled automatically, i.e., the communication between pForecast and the database happens in a scheduled background process.
- Imports can be scheduled on monthly, weekly, or daily basis. In all cases the import uses a time of day given as Coordinated Universal Time (UTC). This is a standard used to set all time zones around the world. So, for instance, Norway is in the time zone UTC plus one, meaning that we are one hour ahead of the reading on a UTC clock (except during daylight savings, when we are two hours ahead).

11

- You may also import data instantly, by selecting the "Import now" option.

Figure 1.4 shows all the elements you may define to set up a daily schedule for importing historical data.

Figure 1.4 Daily import of historical data.

If you choose a daily interval of 2-4, please note that the selected day interval is realigned to the beginning of each month, possibly resulting in a smaller (reduced) interval in the transition between two months.

When you select to only import data for specific facilities, the switch called "Limit import to the following facilities" is automatically toggled on. You cannot toggle it on manually. You can however toggle it off if there are checked facilities in the list, effectively unchecking all facilities, meaning that you choose to import historical data for all facilities.

In fig 1.4 the options to schedule a weekly import of historical data is shown. The settings which are equal to the daily settings explained in Figure 1.4 are not included.

Figure 1.5 Weekly import of historical data.

12

When you push the "Import now!" button, a dialog window will appear, as depicted in Figure 1.6. You have the flexibility to adjust these input parameters to values different from your regularly scheduled imports.

Figure 1.6 Input parameters for the Import now function.

If historical data are updated only once a month, you may utilize the monthly import choice as illustrated in Figure 1.7. You may only select one day within the month. To import more frequently, see weekly or daily import described above.
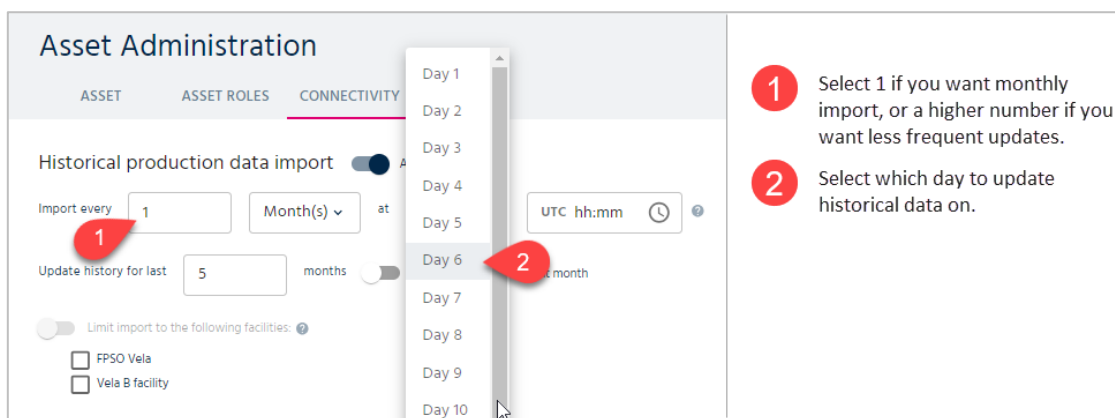
Figure 1.7 Import of monthly historical data.

When configuring pForecast for your company, one can decide to allow for an end user option to include values from the current month. If this feature is configured in, you can simply choose to import monthly data for the current month or not, using is the toggle illustrated in Figure 1.8
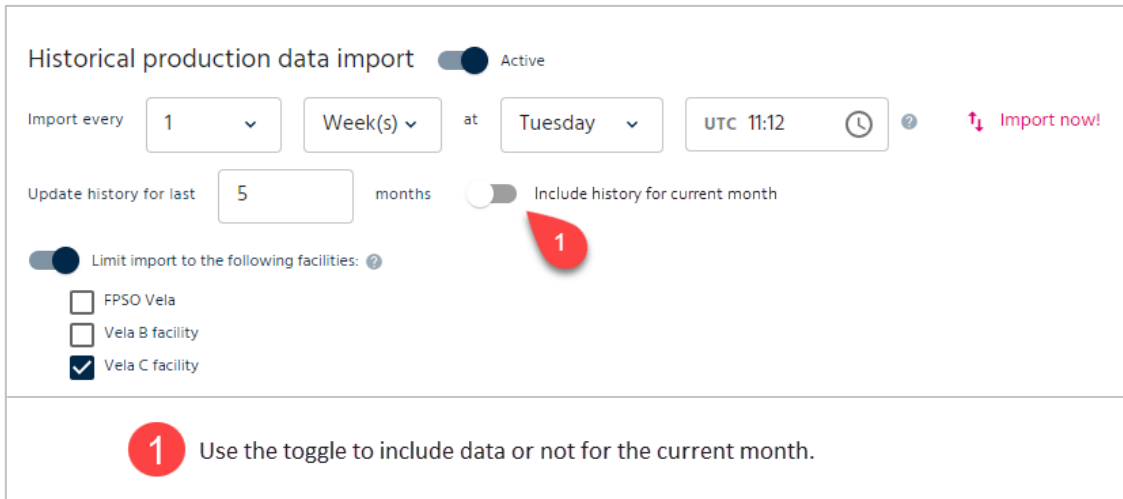
Figure 1.8 Optional setting for connectivity.

1.4.1. Configuration using PDM

If pForecast is set up to import historical production data using the bespoke Production Data Mart (PDM) solution for Equinor, there is one user-configurable property in addition to the generic ones described above.

You may select a geographical area for the asset. Setting this PDM area property will limit access to this asset and its historical production history to users who have access to the corresponding area in PDM.

Note: Once the property is specified, it cannot be changed or reset, i.e., the dropdown will be read only. If Figure 1.9 you can see an example of this setting.
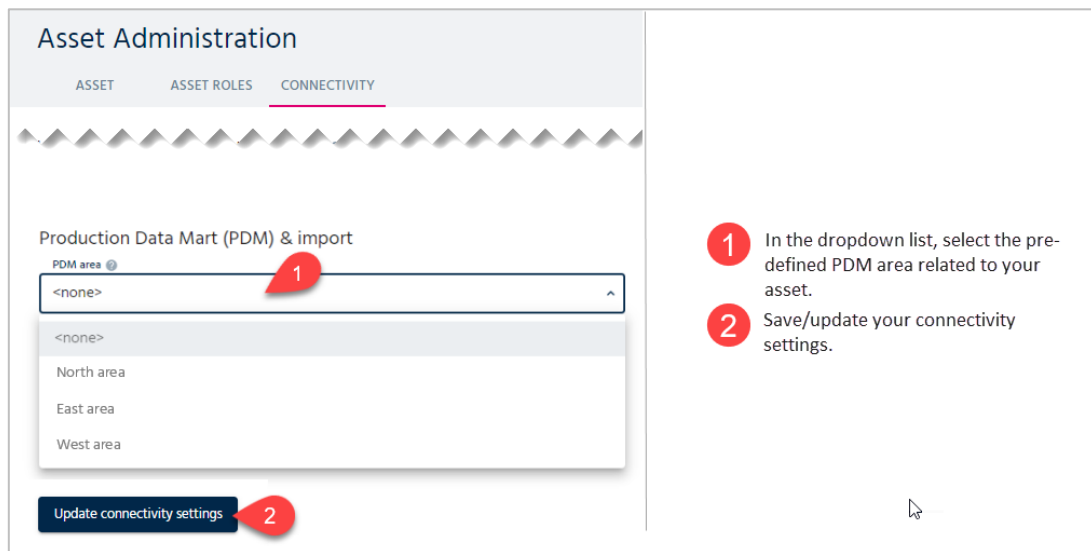
Figure 1.9 Import using the PDM connection/connectivity.

Once an import has been initiated, the administrator may follow the progress and result of the import job under the "Previous & ongoing imports" section. Figure 1.10 illustrates what type of information you can see regarding the import jobs for an asset.

Figure 1.10 Previous and ongoing imports.

2. Asset Configuration

An asset or a business unit in the pForecast context is a group of fields that share data. When a user is given access to an asset, the user can be given read-only or read-and-edit access. The access provided applies to the entire asset. The first time the users start pForecast, they need to select the asset they are working on. The left-hand side navigation bar has a button called "Select asset." You can choose or switch between different assets you have access to by clicking on this button. See Figure 2.1.

Figure 2.1: Asset selection.

When the asset of interest is selected, the user is presented with the following screen (Figure 2.2), where it is possible to configure the asset by pushing the properties button in the top-right corner.
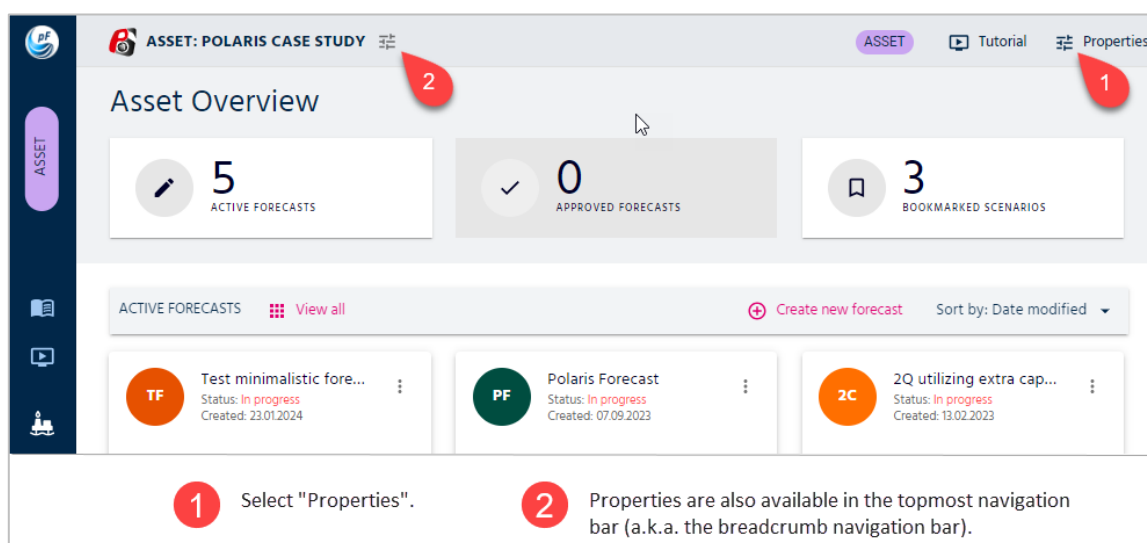
Figure 2.2: Asset configuration.

16

Asset configuration is for setting up static input that does not usually change or require frequent updating. There are five main areas to choose from: Facility, Field, Reporting Schemes, Ownership Groups, and Gas Lift Curve Sets. For the first three areas, there are additional parameters that can be set, as detailed in the following structure:

Facility:

- Clusters
- Wells
- Fuel and Flare

Field:

- PVT
- Voidage
- RNB
- Sales Conversion

Reporting Schemes:

- Tags
- Reporting Groups

In the following sections, you will find an explanation of how to configure these parameters.

2.1. Facility

The user can select an existing facility or create a new one by pressing the "Create new" card. When a facility is selected, it is possible to edit or delete it. It is worth noting that if the facility is deleted, all associated data is deleted. Figure 2.3 depicts how to make a new facility.

Figure 2.3: Facility configuration.

In the oil and gas industry, it is common to find multiple production facilities within a single asset. This practice serves several purposes, driven by the following factors:

- Large field size: Some oil and gas fields are characterized by their significant size and abundant reserves. To efficiently extract and process the hydrocarbons from such fields, it becomes necessary to establish multiple production facilities. These facilities encompass drilling platforms, processing plants, storage tanks, and other crucial infrastructure elements. By distributing the production capacity across various facilities, operators can optimize the extraction process and ensure that the asset reaches its full potential.
- Geographic dispersion: In certain instances, an oil and gas asset comprises multiple fields that are spread across a vast geographic area. Each field may possess distinct characteristics, such as varying reservoir pressures or different types of hydrocarbons. To address these unique requirements and maximize operational efficiency, multiple facilities are strategically deployed. This approach enables operators to tailor production and processing operations according to the specific needs of each field, minimizing logistical challenges and enhancing overall performance.
- Expansion and tie-in fields: Oil and gas assets often undergo development in stages. Initially, a production facility is designed and constructed to handle the output from primary fields. However, as new fields are discovered or tie-in fields are integrated into the asset following the initial development, the existing facility's capacity may prove insufficient. Consequently, additional facilities are built, or existing ones are expanded to accommodate the increased production volume. This adaptive approach allows for seamless integration of new fields into the asset's infrastructure.

Within the pForecast software, users have the capability to simulate the segregation of phases within a single facility, exporting specific portions of production while redirecting the remaining volumes to another facility. Additionally, the system allows for the modeling of dependencies in facility uptimes, which arise when multiple facilities share power units or processing plants.

Figure 2.4 Setting up a hierarchical structure for a multi-facility asset.

Figure 2.4 shows how, under the Facility Properties, you can set up a two-level hierarchical structure for a multi-facility asset, in order to model separation of phases at child facilities as well as dependencies on PE and shutdowns in the hierarchy.

2.1.1. Clusters

Clusters are used to group wells and to constrain the production and injection to the capacity of that group. Clusters can be routed to other clusters; there can be up to four levels of clusters in a facility. Later, when configuring forecasts, you can attach constraints to the clusters defined here. When the "Clusters" tab is selected, it is possible to create, edit, and delete clusters. When the plus icon is clicked, the cluster dialog box appears. It is possible to give the cluster a name. The name must be unique. It is also possible to specify an existing parent cluster. If no parent is specified, the cluster will be a top-level cluster. Figure 2.5 shows how clusters can be created and structured.

Figure 2.5: Cluster configuration.

2.1.2. Wells

The most crucial feature of wells in the pForecast software is to hold the production history. pForecast has two notions of a well. Firstly, it is a physical well, as detailed here; secondly, it is a forecast well that connects to the physical well. The forecast well connects to a wellbore of a given physical well, and each well must have at least one wellbore to get the historical production over to the forecast. Wells can be created, edited, and deleted. Figure 2.6 depicts how historical production and injection data can be uploaded from Excel.

Figure 2.6: Well configuration and uploading of historical production data from Excel.

Figure 2.7 shows the Excel template for uploading historical production per well.

Figure 2.7: Excel template for uploading historical production per well.

2.1.2.1. Import historical production data from CDF

Historical production data can also be imported from Cognite Data Fusion (CDF). Once the connection is set up, pForecast will collect updated historical data from CDF daily.

Figure 2.8 shows how to configure for import of historical production data from CDF.

Figure 2.8: Import historical production data from CDF.

After importing production data through either Excel sheets or CDF, for quality control, you can open the "History details" menu for each well and observe the production and injection profiles for different fluids (see Figure 2.9).

Figure 2.9: Quality control of imported historical data per well.

2.1.3. Fuel and Flare

Both fuel and flare are calculated at the facility level as linear functions of produced and injected volumes. The fuel and flare calculated per facility are distributed to the contributing wells for sales calculations based on one of the following methods:

a. Liquid volume (use as the default)
b. Gas volume
c. Oil volume
d. Sum of produced oil, gas in oil equivalents and water
e. Sum of produced oil, gas in oil equivalents, produced water, injected gas in oil equivalents, and injected water

The below formula will be used to distribute the fuel based on the liquid volume method:

$$V_{\text{fuel from well i}} = V_{\text{total fuel for facility}} \cdot \frac{V_{\text{liquid from well i}}}{V_{\text{total liquid facility}}}$$

The other methods follow this pattern.

In pForecast, fuel and flare input data can be imported through an Excel file or edited manually using the embedded table. See Figure 2.10 for more details.

Figure 2.10: Uploading fuel and flare input data.

Figure 2.11 shows the available Excel template in pForecast for uploading fuel and flare input data.

Figure 2.11: Excel template for uploading fuel and flare input data.

2.2. Fields

The user can select an existing field or create a new one on the asset level by pressing the create card. You can define as many fields as needed. When a field is selected, it is possible to edit or delete it. If one of the fields is deleted, all associated data is deleted. Figure 2.12 illustrates how to create a new field.

Figure 2.12: Creating a new field on the asset level.

2.2.1. PVT

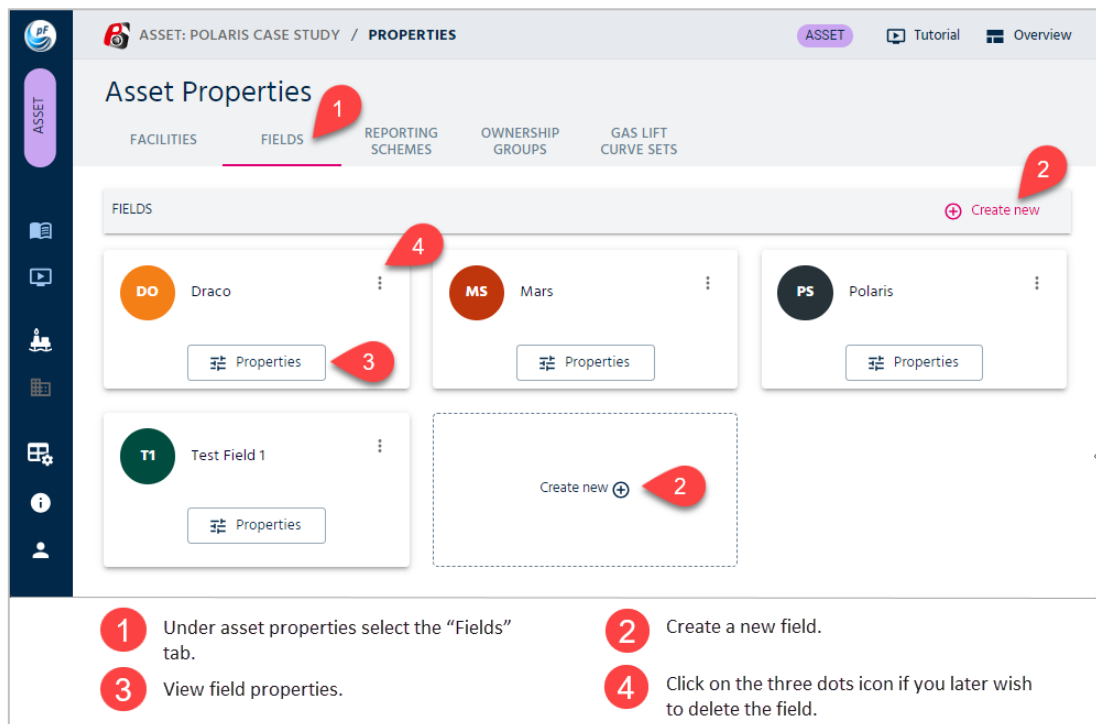Voidage groups use PVT properties. Constant PVT properties are assumed to be sufficiently accurate to model injection to achieve a reasonable voidage replacement. PVT properties can be edited by selecting the PVT Properties tab. PVT properties can be created, edited, and deleted. When the plus icon is pressed, PVT properties can be set, and the following parameters are entered: formation volume factors for oil and gas and solution gas oil ratio. The formation volume factor for water is assumed to be unity. Figure 2.13 indicates how to apply PVT properties.
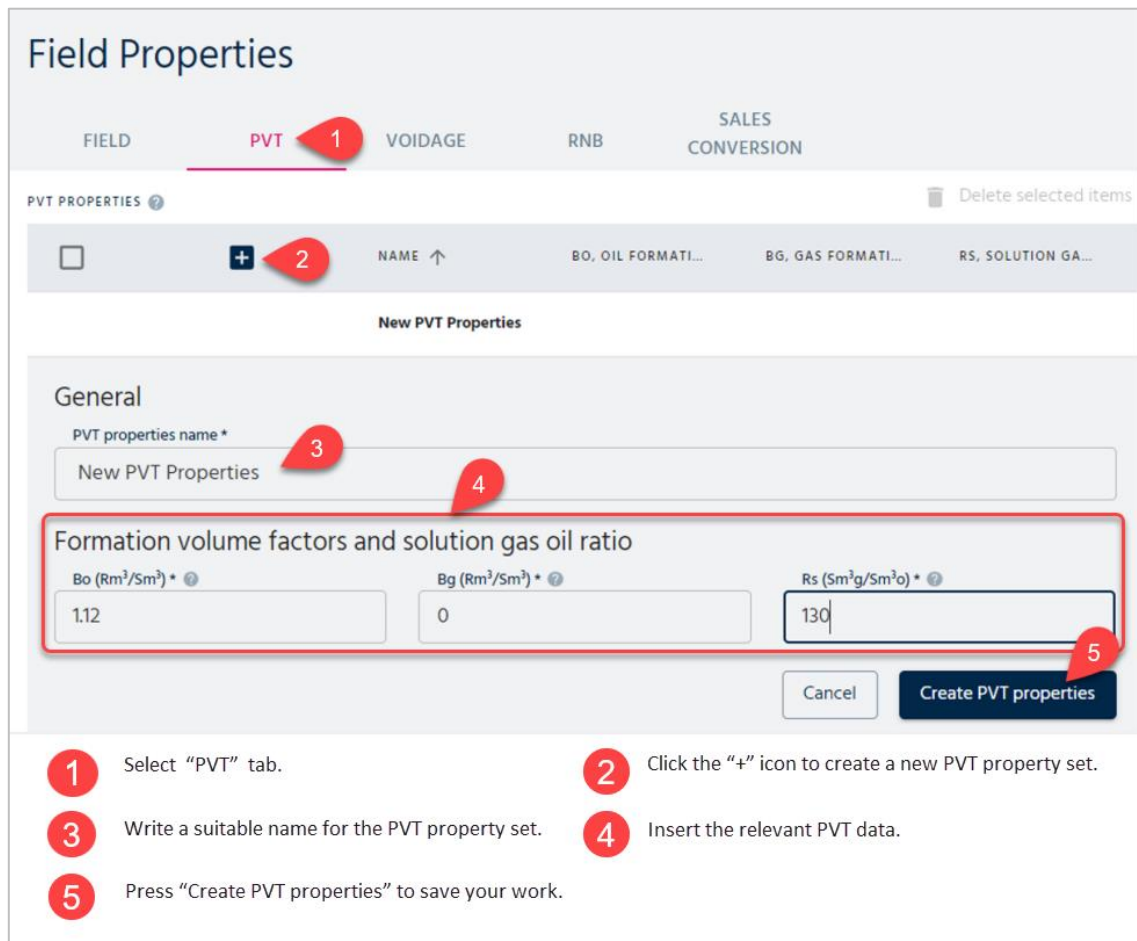
Figure 2.13: Applying PVT properties for voidage calculations.

2.2.2. Voidage

Voidage groups are used to model required injection volumes. The groups are created here and used later when a scenario is built. When the voidage tab is selected, it is possible to create a new one and edit or delete the selected voidage group. In the field properties, the following page appears when the "Voidage" tab is pressed. A unique name must be entered, and a set of PVT properties, defined in the previous step, must be selected. By applying the target voidage factor as a fraction, pForecast will steer the injection rates to reach the target voidage factor. See Figure 2.14 for more details.

Figure 2.14: Defining the voidage group.

2.2.3. RNB

A structure of RNB profiles and projects can be defined for a field. These profiles and projects are used to aid in reporting data to the Revised National Budget (RNB) in Norway each autumn, 15 October. An RNB profile is a set of profiles given to the authorities. The Norwegian Petroleum Directorate (NPD) allows up to 15 profiles with annual reports on recoverable resources in classes 0 to 5. Projects in resource classes 6 and 7, which are reported separately in the RNB reports, are to be placed in profile 0. The reader is referred to the NPD website for more information about resource classification. Each profile should be given a unique name and a profile number.

When the RNB Profiles tab is selected, it is possible to create a new RNB profile or to edit or delete the selected profile. See Figure 2.15 for more details.

Figure 2.15: RNB profiles and projects.

2.2.4. Sales Conversion

For a field, the user can define yearly or monthly entries, each containing various parameters related to conversion from production to sales figures. If a set of parameters is provided for both a month and the year of that month, the month data will be used in calculations. In the case that there is no provided entry for a month, the defaults for the various parameters will be used. See Figure 2.16 for more details.

Figure 2.16: Sales conversion factors.

If your sales conversion factors vary over time, we recommend downloading the Excel template and populating it with your values. Figure 2.17 shows this template partially filled in, listing some of the factors you may include.

Figure 2.17: Excel template for uploading sales conversion parameters.

2.3. Reporting Schemes

A reporting scheme is used to define how the simulation results are aggregated and output. pForecast calculates the production rates per well, and reporting schemes enable the users to define what groups of wells they wish to output production rates for. It is possible to define multiple reporting schemes.

Each forecast well is assigned one 'tag' that describes a property of the well. In fact, a well can only be assigned one tag per reporting scheme. The tags must therefore be disjoint, meaning the properties do not overlap. Examples of properties used to define tags are which reservoir the well produces from, which production license it belongs to, if the well is a producer or injector, and so on and so forth.

On the other hand, tags are collected into one or more report groups. The report groups will be used to output results containing the sum of production from all wells assigned tags belonging to the report group. To sum up, the steps for configuring a reporting scheme in pForecast software are:

1. Create a new reporting scheme.
2. Define the tags that will be assigned to the wells.
3. Define the report groups that will be output in the results and assign well tags to the report group.
4. Assign tags to the specific wells in the forecast. This step is performed in the forecast configuration.

In pForecast, reporting schemes can be defined either by populating the provided Excel file and uploading it or by creating them manually. Figure 2.18 demonstrates the process of creating a new reporting scheme, Figure 2.19 showcases the creation of tags, and Figure 2.20 explains the creation of report groups. On the other hand, Figure 2.21 demonstrates the method of importing an Excel file to accomplish the same task.


Figure 2.18: Creating a new reporting scheme.

Figure 2.19: Creating tags.

Figure 2.20: Creating report groups.

Figure 2.21: Creating report groups by importing an Excel file.

2.4. Ownership Groups

For an asset, the user can define a set of ownership groups consisting of ownership shares for sales oil, gas, and natural gas liquid (NGL).

A forecast well, defined under forecast properties, can subsequently connect to an ownership group. The various ownership shares are used when calculating net sales values. In other words, wells are tagged to ownership groups to characterize the company's ownership share of the produced volumes. It is also possible to define a default ownership group; in that case, all wells that have not been specifically tagged to an ownership group will be assigned to this group. Suppose there is no default ownership group for the asset, and a well is not assigned to a specific ownership group. In that case, pForecast will assume 100% ownership share as the default value when generating net sales profiles. In that case, the net sales and the gross sales are equal. Figure 2.22 shows how to create an ownership group step by step.

Figure 2.22: Creating ownership groups.

2.5. Gas Lift Curve Sets

For an asset, the user can define gas lift curve sets, consisting of curves that give the relationship between the fraction of gas lift potential and the fraction of production potential for given water cuts. A production profile, which is defined under Forecast Properties, can subsequently connect to a gas lift curve set. The gas lift curves are used in connection with constraint calculations involving gas lift. Figure 2.23 shows how to specify gas lift curves.

Figure 2.23: Generating gas lift curves.

3. Forecast Configuration

Forecast configuration provides future projections for three main input data domains: subsurface, facility operations, and drilling. The following sections give details on creating a forecast and configuring previously mentioned input data.

3.1. Create a Forecast

The first step is to create a new forecast. It is worth mentioning that if the user selects to duplicate a forecast, the original one is kept, and changes in the duplicated forecast are not reflected in the original. It is also possible to edit or delete the selected forecast. See Figure 3.1 and Figure 3.2 for more details.
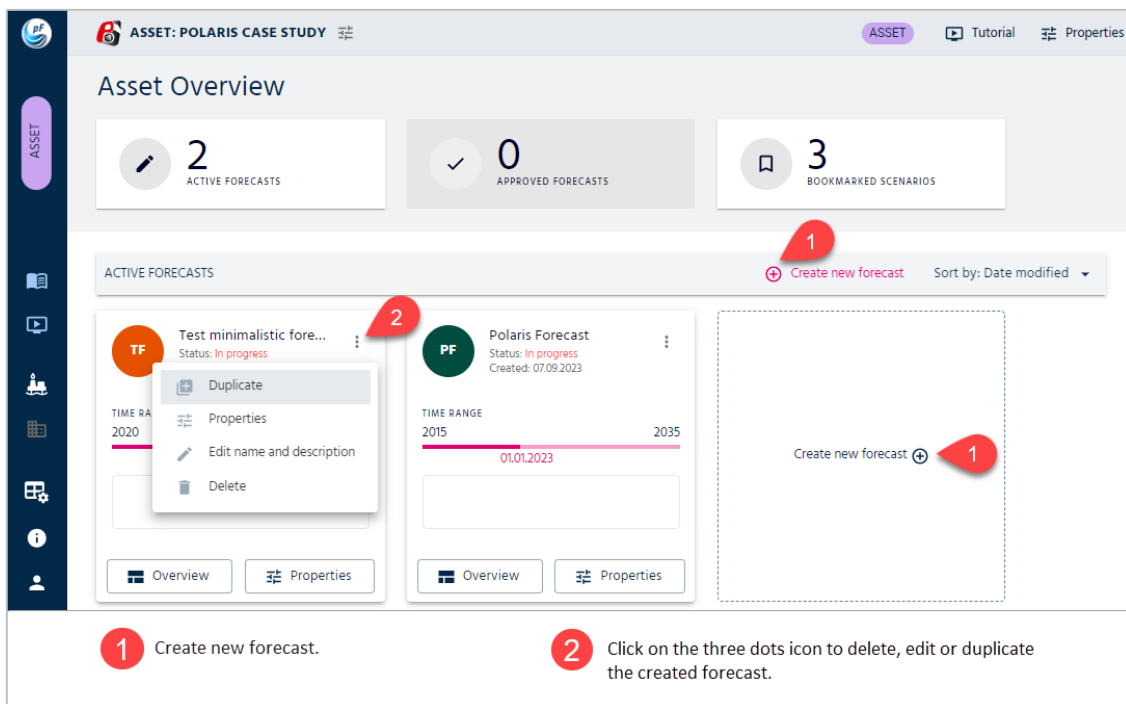

Figure 3.1: Creating a new forecast (part 1).

After creating a new forecast, the following page will be shown, in which the user can provide high-level forecast properties and time ranges:
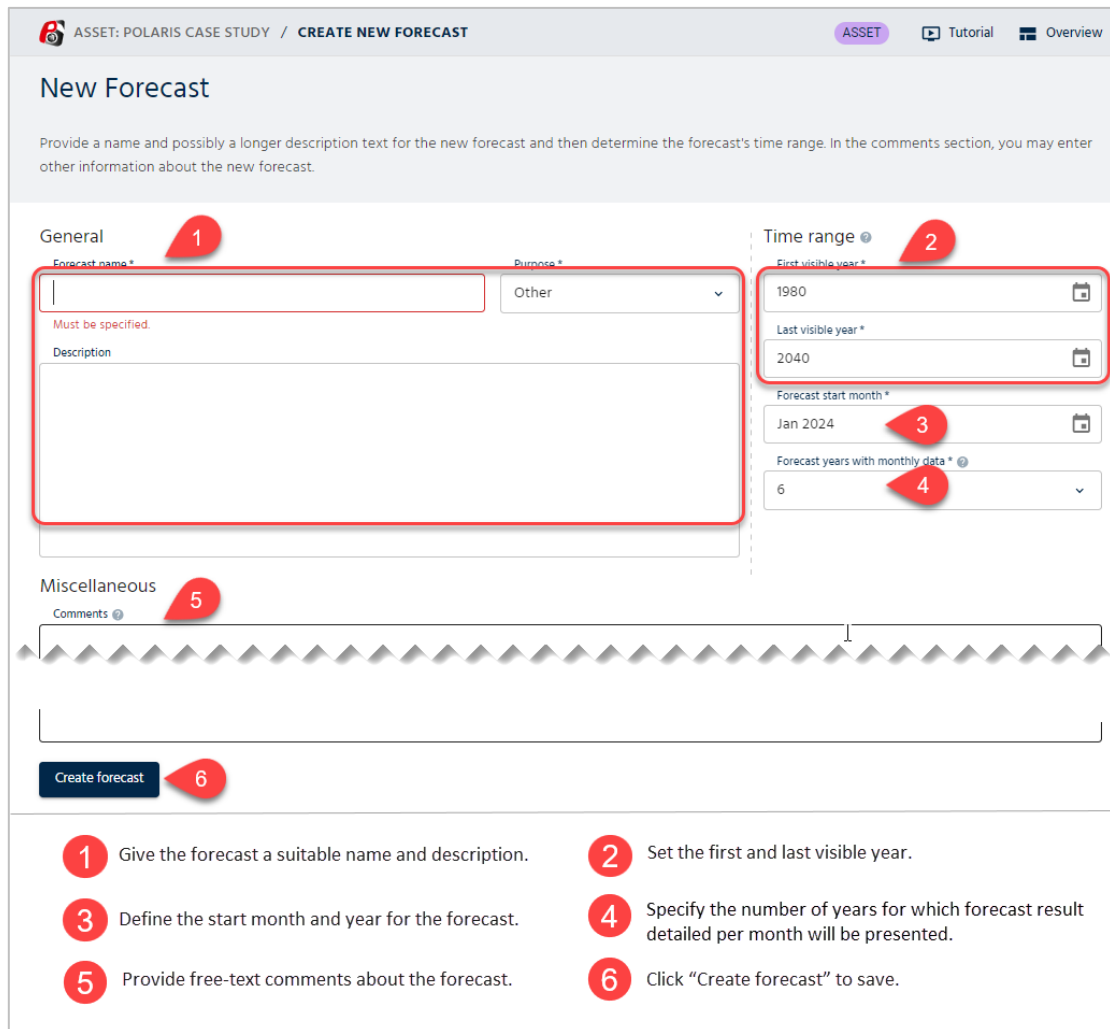
Figure 3.2: Creating a new forecast (part 2).

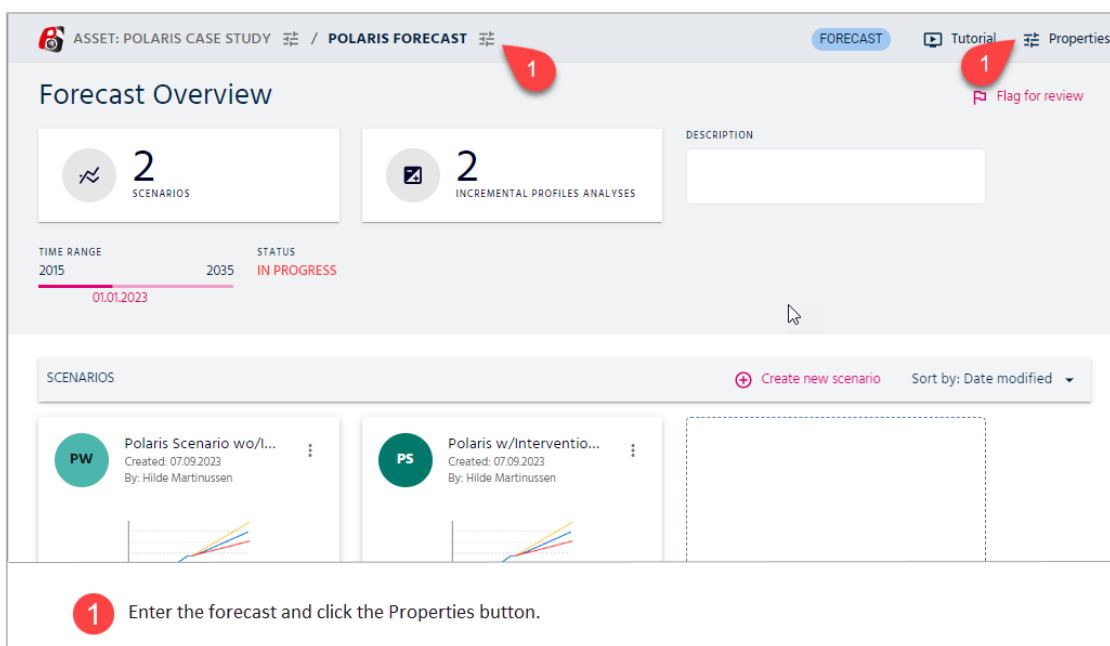After making a new forecast, the user can easily access all properties, see Figure 3.3.

Figure 3.3: Access to the forecast properties.

3.2. Subsurface Input and Configuration

3.2.1. Production Potentials

For a forecast, the user can define a set of production profiles to be used by forecast wells of type producers to establish their production potential. The potential is given as a rate per stream day, where a stream day is 24 hours when running at full capacity under optimal conditions. pForecast supports six types of production profiles:

- Stream day rate vs. month
- Stream day rate vs. delta month
- Fractions vs. delta month (see section 3.2.1.1)
- Decline curve
- Multi-segmented decline curve
- Stream day rate vs. accumulated volume

The actual production (given as a rate per calendar day) can be different from what the input production potential profile would suggest since the production is affected by various factors such as production efficiency, shutdown periods, and capacity limitations. In addition, potential profiles can be set to be volume-based. The production volumes can be quite different than the input production potential profile would suggest. To adjust for this, the user can specify that the production profile is to be calculated using a volume-based profile rather than a time-based. In order to have the correct relationship between time and produced volume, the user should enter an average production efficiency applied in the establishment of the production profile. In addition to using volume-based potentials for production estimation, it is possible to cut the production based on a given produced volume, a given cut-off rate, a time limit when drilling a new well starts, or any combination of these.

The production potential for the well is imported from an Excel file, and for the first time, a new profile will be created and appear in the list. When production potential is uploaded for a profile that already exists, the new data will overwrite the old data.

After importing, a list of the production profiles is shown. A unique name for each production profile is shown together with the type of production profile. It is also specified whether the profile is generic or not. Different forecast wells can reuse a generic production potential profile. The list also shows how many times the production potential profile is used by forecast wells, noted by use count. The dates for the creation and last modification of the production potential profile are also shown.

Figure 3.4 and Figure 3.5 display how to upload and configure well production potentials.

Figure 3.4: Uploading well production potentials.

After uploading the well potentials, click the pen icon to enter the menu for potential configuration. More details are shown in Figure 3.5.

Figure 3.5: Configuration of well production potentials.

3.2.1.1. Distribution types in pForecast

pForecast supports several types of distributions. For all of them, up to three parameters are used when generating samples from the distribution. A low value, a high value, and a mid/deterministic value (the value used in a deterministic analysis).

Triangular (P10, P90) - The Triangular distribution is a continuous probability distribution with lower limit, upper limit and mode. There are two variants of the Triangular distribution in pForecast. In this variant, the three defining parameters are calculated from the provided low/P10 value (10% probability of non-exceedance), the provided high/P90 value (90% probability of non-exceedance), and the mid value that is used either as a most likely value (mode) or an expected value.

Normal - The Normal (or Gaussian) distribution is a continuous probability distribution given by a mean (expected value) and a standard deviation. In pForecast, these two parameters are calculated from a provided low value (10% probability of non-exceedance), and a provided high value (90% probability of non-exceedance). The mean is the average of the low and high value, while the standard deviation is the difference between high and low multiplied with 0.39. Note that for this distribution the provided mid value is not used; it should however not deviate too much from the calculated mean.

Pert - The Pert distribution is a continuous probability distribution defined by a minimum value, a most likely value and a maximum value. It is a smoother shape alternative to the Triangular distribution. In pForecast, the provided low value and high values give the minimum and maximum values directly, while the mid value is either the most likely value or an expected value that is used to derive the most likely value.

Log normal - A Log normal distribution is a continuous probability distribution of a random variable whose logarithm is normally distributed. Like the Normal distribution, it is defined by a mean and a standard deviation. In pForecast, these two parameters are calculated from a provided low value (10% probability of non-exceedance), and a provided high value (90% probability of non-exceedance). The mean is 0.5 multiplied with the natural logarithm of the product between the low and high value, while the standard deviation is 0.39 multiplied with the natural logarithm of the quotient of the high value and the low value. Note that for this distribution the provided mid value is not used; it should however not deviate too much from the calculated mean.

Uniform - A Uniform distribution is a continuous probability distribution that defines an arbitrary outcome that lies between certain bounds. In pForecast, these bounds are defined by subtracting one eighth of the difference between the provided high (90% probability of non-exceedance) value and the low (10% probability of non-exceedance) value from the lower value and adding one eighth to the high value. Note that for this distribution the provided mid value is not used; it should however not deviate too much from the average of the high and low values.

Triangular (P0, P100) - The Triangular distribution is a continuous probability distribution with lower limit, upper limit and mode. There are two variants of the Triangular distribution in pForecast. In this variant, the three defining parameters are given by the provided low/P0 value (the actual lower limit), the provided high/P100 value (the actual upper limit), and the mid value that is used either as a most likely value (mode) or an expected value.

3.2.1.2. Fraction Resources

The fraction resources profile type (Fractions vs delta month) is very similar to the delta monthly profile type. The values you provide per month represent the proportion of the resource volume. Typically, the sum of all your fraction monthly values should add up to one. If this condition is met, the entire resource volume will be produced during the specified relative months.

Example: let's say you have a resource of one million Sm^3. Normally it would take 400 months to produce this resource, however, for the sake of simplicity, let's assume that the total production time is five months. In that case, the five production fractions could be: [0.3, 0.3,

0.2, 0.1, 0.1]. That will give a production volume per month of [300.000, 300.000, 200.000, 100.000, 100.000]. So far, this is quite straightforward. This input could also work fine if we change our resource to 900.000 Sm^3. In that case we would get another scale on our production, only by changing the resource value: [270.000, 270.000, 180.000, 90.000, 90.000]. In many cases you want to do this, because the shape of your profile remains the same for many resource volumes. pForecast requires input about how much the resources will decrease over time. This is given in a separate Excel sheet containing resource values versus time. This could for example mean that if the well starts producing in 2027 the resource will be one million Sm^3, but if it starts producing in 2035, the resource will be 900.000 Sm^3. By giving this in the ''VsTime'' Excel sheet, pForecast will interpolate the resource volumes in the years between. The same monthly fraction values will be used regardless of the resource volume which is dependent on the production start time.

3.2.2. Forecast Wells

The user can create, edit, and delete forecast wells by selecting the "Forecast well" tab. A list of forecast wells is presented. It is worth mentioning that when the production potential for a well is uploaded for the first time, a new forecast well will automatically be created and appear under the tab "Forecast Well." The list shows the name, whether this well is already in operation (or whether it is a drilling target), and the facility the forecast well belongs to. It is also shown how many times the forecast well is used in scenarios and the dates for the creation and last modification of the forecast well.

After clicking on the "Forecast well" tab, the user should click the pen icon to enter the menu to configure the forecast well. A unique name for the well should be given. The well should be linked to a facility. The production efficiency of the well typically depends on which facility the well is connected to. The default set is used if the well is not connected to any facility. Users can specify production efficiency forecasts for all facilities and the default set. It is also possible to override the production efficiency for a particular well by choosing a specific production efficiency forecast for the well. A forecast well that is in operation is typically linked to a wellbore. The production history is taken from the given wellbore. No history is included if the well is not linked to any wellbore. The forecast well can also be linked to a cluster. A cluster is a group of wells that share the same capacity, and this capacity limit may constrain the well. Another option is to link a well to a voidage group. The voidage groups are used for calculating injection requirements. No producing wells are influenced by this, but injection can be reduced according to voidage requirements. The well can also be linked to an uncertainty group. It is also possible to assign the well a tag per reporting scheme. See Figure 3.6 for more details.

Figure 3.6: Configuration of forecast wells.

3.2.2.1.   Injectors Wells

Under the "Forecast well" tab, injector wells can also be created, and provided a target rate schedule for injectors.

The user must supply a start year and a start month for the injection control. The user then provides a mode for the injector. The mode can be one of the following:

- Water injection
- Gas injection
- Water Alternating Gas (WAG) injection.

Priority is given to all wells within the same cluster. If an injection constraint is exceeded, the higher priority wells are kept on injection. The target rate for water and gas are entered. If the well does not belong to a voidage group, the well is kept on the target rate if injection constraints permit it. If the well belongs to a voidage group, the target rate may be reduced to match the production level of the voidage group. If the injection mode is selected to be WAG, then the user must supply WAG duration in months. If the well injects gas, the gas formation volume factor for the injection gas (B_g) should be entered. Figure 3.7 shows more details about injector well configuration.

Figure 3.7: Creating injector wells.

Figure 3.8 shows the available Excel template in pForecast for uploading injection schedules.

Figure 3.8: Injection schedules template.

3.2.3. Production Correlation

The well can also be linked to an uncertainty group. An uncertainty group is used to correlate the production of a collection of wells in stochastic analyses. In fact, the group's correlation factor is used to define the strength of the wells' interdependency. It is also possible to have negative dependencies (correlations). For instance, for a group of two wells, a negative correlation would mean a relationship in which one well's production increases as the other's decreases and vice versa. Figure 3.9 shows the steps toward making an uncertainty group.
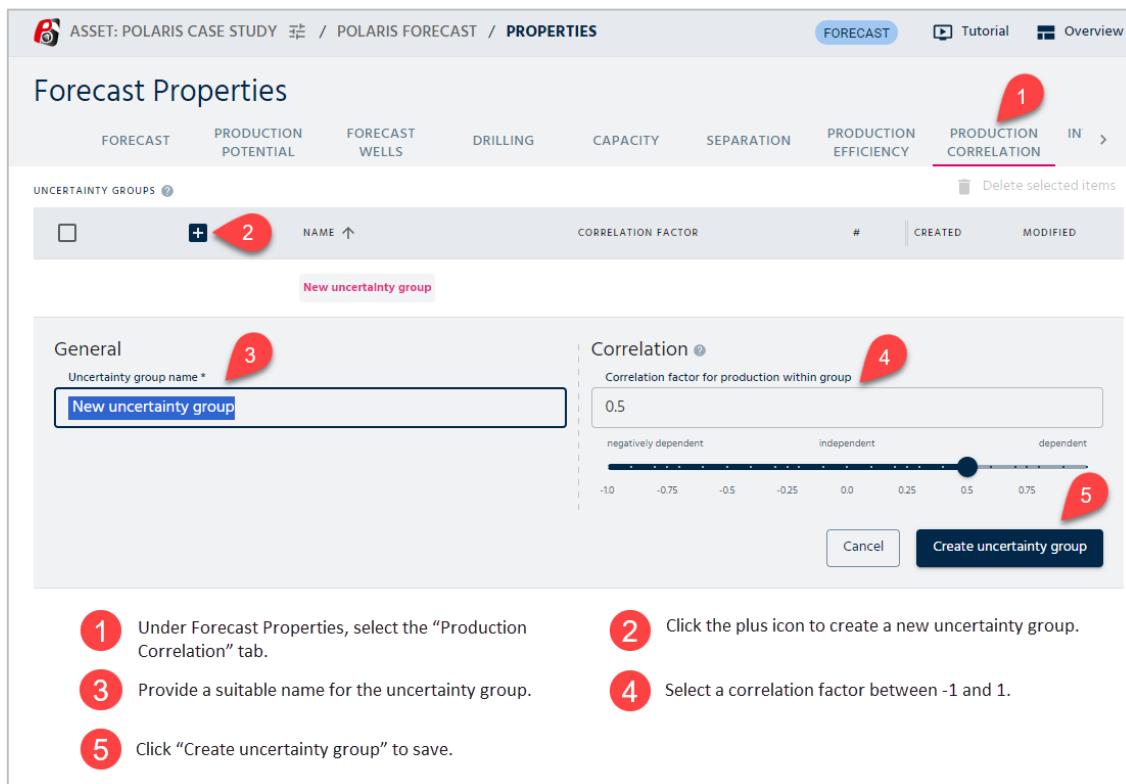
Figure 3.9: Creating uncertainty group.

After creating an uncertainty group, the user can assign production wells to the uncertainty groups under the "Forecast well" tab. The created uncertainty group will now be an option in the dropdown menu (see step 9 in Figure 3.6).

3.3. Facility Operations Input and Configuration

3.3.1. Production Efficiency

Normally, the uptime for the production facilities is close to 100%. In pForecast, in addition to most likely values for Production Efficiencies (PE), it is possible to provide upside and downside cases. Users can create, edit, and delete PE forecasts by selecting the "Production Efficiency" tab under the forecast menu. A list of PE forecasts is presented. The list shows the name and the facility the PE forecast belongs to. It is also shown how many times the PE forecast is used in scenarios (use count) and the dates for the creation and last modification of the PE forecast. Figure 3.10 illustrates how to upload the PE forecast from an Excel file.
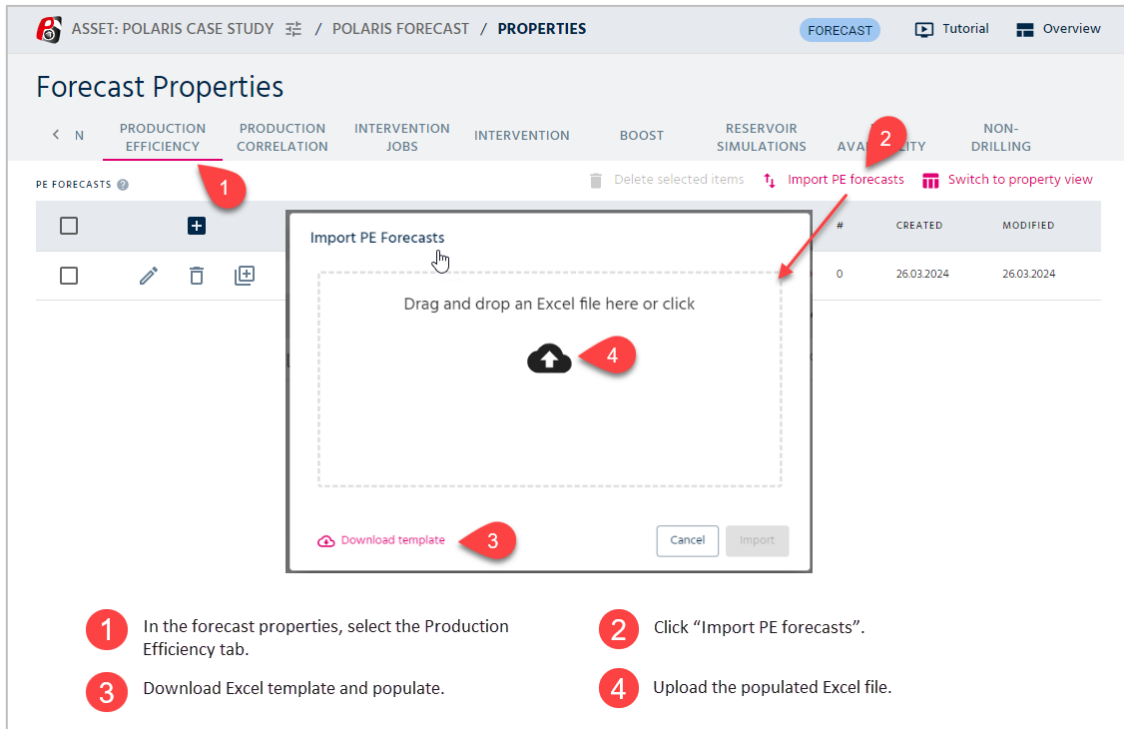
Figure 3.10: Uploading production efficiency schedule.

Figure 3.11 shows the available Excel template in pForecast for uploading PE forecasts.
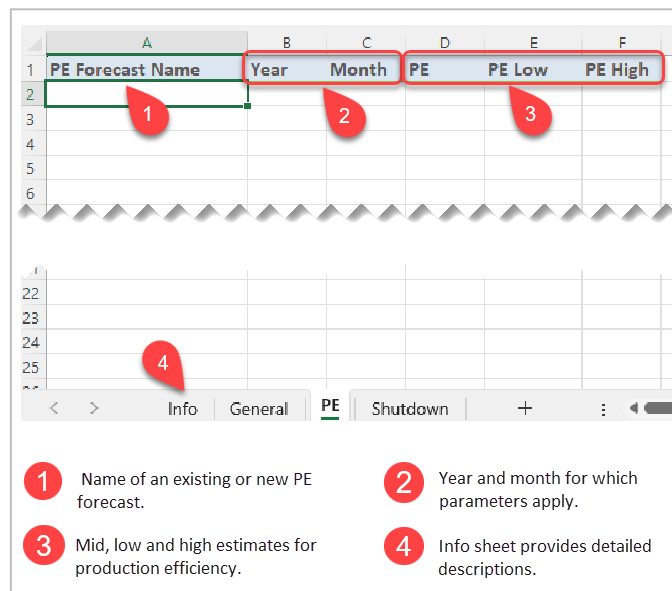

Figure 3.11: Production Efficiencies schedule template.

The user can also specify monthly and annual estimates for the production efficiency manually by filling out the embedded table. It is also possible to choose distribution type and boost factors. See Figure 3.12 for more details.

Figure 3.12: Providing PE values in tabular format.

If you have previously set up a two-level hierarchical structure for a multi-facility asset at the asset level, you can control dependencies on PE (Production Efficiency) within the hierarchical facility structure. These dependencies are established for PE forecasts of child facilities in the two-level hierarchy.

A correlation factor determines how the PE values of a child facility are related to the PE values of the parent facility. It is important to note that negative dependencies (correlations) are allowed. This implies a relationship in which one facility's PE increases as the other facility's PE decreases, and vice versa. Furthermore, you can specify that the PE values of the child facility will never exceed those of the parent facility under any circumstances. See Figure 3.13.
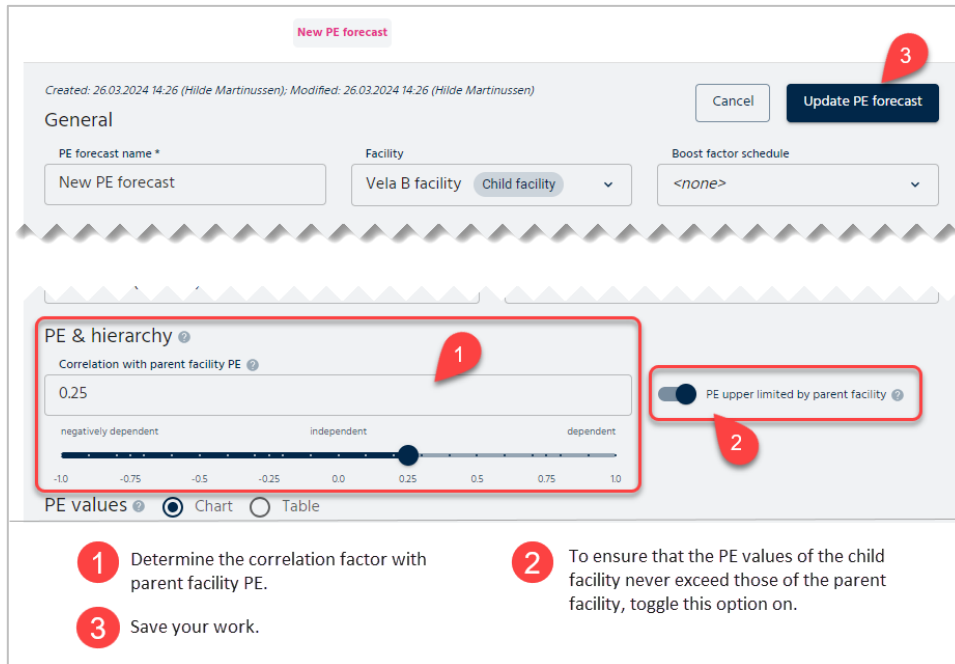
Figure 3.13: Controlling the dependencies on PE in a hierarchical facility structure.

3.3.2. Boost Factors

Boost factors model boost production after a long shut-in of a field. After such a stop, many fields experience significant changes in water cut, gas-oil ratio, or well potential until the production stabilizes after a period of production.
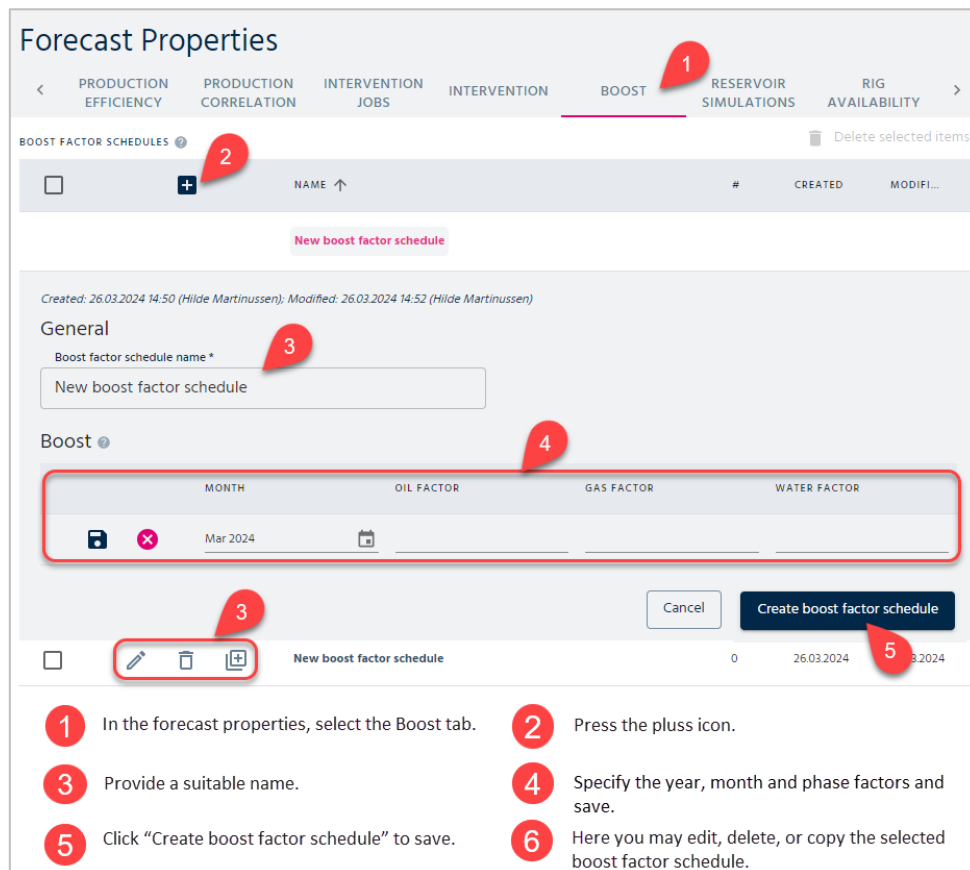

Figure 3.14: Boost factor schedule.

In order to model this, it is possible to specify a monthly boost factor in a boost factor schedule and connect this schedule to the PE forecast. Each phase rate is then multiplied by the boost factor for that phase for the year and month specified. Users can create, edit, and delete boost factors by selecting the "Boost" tab. A list of boost factor schedules is presented. The list shows the name of the boost factor. It is also shown how many times the boost factor schedule is used in scenarios and the dates for the creation and last modification of the boost factor schedule.

The boost factor schedule dialog box appears when the plus icon is clicked. A unique name for the boost factor schedule should be entered. It is then possible to add, edit or delete individual boost factors from the list in the dialog box. When a new boost factor is entered, the user must supply the year and month for the boost and the boost factors for oil, gas, and water. See Figure 3.14 for more details.

3.3.3. Planned Shutdowns

In addition to PE values, planned shutdowns can also be considered. For instance, in case of scheduled shutdowns due to facility maintenance, turnarounds (TAR) can be defined, and it is possible to assign uncertainties both to the start date and the duration of these scheduled shut-ins.
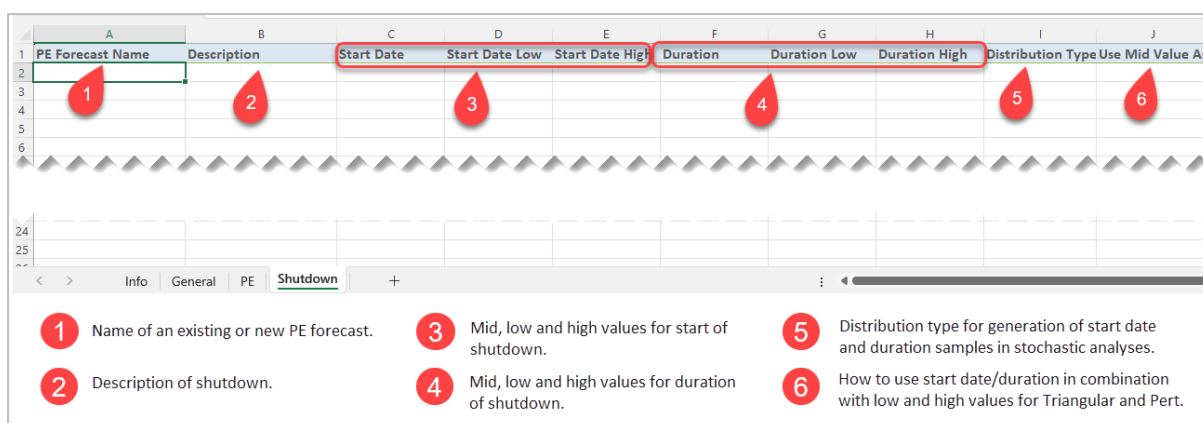
Figure 3.15: Planned shutdowns schedule template.

In pForecast planned shutdowns can be defined either by populating the provided Excel file and uploading it, see Figure 3.15, or by manually defining the schedule, see Figure 3.16.

In a two-level facility hierarchy, a child facility's PE forecast can incorporate shutdowns from the parent facility's PE forecasts. When inheritance is enabled, you will observe potential shutdowns from the parent PE forecasts mixed with the shutdowns specific to this PE forecast. For each inheritable shutdown, you have the option to choose whether to inherit it or not. If there are multiple PE forecasts for the parent facility, the inherited shutdown descriptions will include the respective PE forecast names as prefixes.

It's important to note that inheriting a shutdown from a parent PE forecast does not activate the shutdown in a scenario unless that specific PE forecast is also included for the parent facility. In other words, you can opt to inherit shutdowns from multiple PE forecasts, and the final decision regarding activation depends on the inclusion of the respective PE forecast for the parent facility. See Figure 3.17.

47

Figure 3.16: Defining planned shutdowns manually.

Figure 3.17: Inheriting planned shutdowns from parent facility.

3.3.4. Capacity Constraints

Users can add constraints by selecting the "Capacity" tab under the forecast properties section. It is also possible to copy, edit or delete selected constraints. When the plus icon is selected, the user can create a new capacity constraint. The user must enter a name and indicate from when the constraint is valid. A constraint is valid from this point in time and onwards until a new constraint is entered. Maximum values for oil, gas, water, and liquid production are entered along with maximum values for the injection of water and gas. Regularity factors for the injection can also be entered. Figure 3.18 shows more details.

Figure 3.18: Adding capacity constraints.

3.3.5. Separation

In a forecast, users can define separation forecasts, which provide monthly estimates for the phases' separation before routing to the next level of the cluster hierarchy. These forecasts help users understand how the phases progress and are distributed within the hierarchy, aiding in planning and decision-making. Figure 3.19 shows how to create a separation forecast.

Figure 3.19: Creating a separation forecast.

3.4. Drilling Input and Configuration

3.4.1. Drilling Schedules

For a forecast, the user can define a set of wells in operation and drilling targets (wells to be drilled). The two types are both referred to as forecast wells. It is worth mentioning that a forecast well is not to be confused with a physical well defined under the asset level. The latter is used to hold the production history, and an in-operation forecast well typically connects to a wellbore of a physical well in order to get the historical production over to the forecast. Figure 3.20 shows how to create the undrilled wells.

Figure 3.20: Creating a well to be drilled.

The next step is to set up drilling schedules. The list under the "Drilling" tab shows the name of the drilling schedule and which facility it belongs to. It is also shown how many times the drilling schedule is used in scenarios and the dates for the creation and last modification of the drilling schedule. A drilling schedule can be created, and the selected drilling schedule can be copied, edited, or deleted. A unique name should be entered for each drilling schedule. A description can also be given. The drilling schedule can be linked to a facility. If it is linked to a given facility, only wells from that facility can be included in the drilling schedule. The start date of the drilling schedule must be provided. In addition, the user can specify a rig availability schedule. A rig availability schedule is a forecast of how much time the rig is used for drilling. This factor is given as a yearly efficiency factor. Non-drilling periods can also be specified. These are planned periods when the rig is not drilling. These periods are provided with a fixed start date and duration, including uncertainties. Finally, the sequence of wells on the drilling schedule is shown in a list. Wells can be added or removed from the list. A selected well can be easily moved up or down in the sequence. Figure 3.21 shows how to create a drilling schedule.

Figure 3.21: Creating a drilling schedule.

3.4.2. Rig Availability Configuration

Normally, the rig cannot be utilized for drilling the entire time it is in operation. Maintenance on the rig itself and other activities on the platform may prevent drilling operations. These activities are either planned or unplanned. In pForecast, unplanned periods without drilling are entered as rig availability. While planned work can be entered as non-drilling periods with fixed start times and uncertain duration.

Let us first look at unplanned periods. A rig availability factor is an expression of how much a drilling rig is used for drilling wells. The rig availability is entered as a percentage on an annual basis and influences the drilling schedule performance. Selecting the "Rig Availability" tab presents a list of rig availabilities. The list shows the name and the facility the rig availability belongs to. It is also shown how many times the rig availability is used in scenarios (use count) and the dates for the creation and last modification of the rig availability.

Rig availability can be created, and the selected rig availability can be copied, edited, or deleted. The rig availability form appears when the plus icon is clicked. A unique name should be entered, and a facility to link to. The rig availability for each year is entered in the table as a percentage. If no rig availability is given, pForecast assumes it to be 100%. See Figure 3.22 for more details.

Figure 3.22: Rig availability configuration.

3.4.3. Non-Drilling Configuration

On the other hand, planned periods when no drilling takes place are modeled as non-drilling periods in the pForecast software. Typical examples are plant maintenance stops, major rig upgrades, or when the license owner has decided not to drill wells for an extended period. A non-drilling period is assumed to have a fixed start point in time. The stop duration is given alongside low and high estimates for the duration. The user can create, edit, and delete the non-drilling schedule by selecting the "Non-drilling" tab.

A list of non-drilling schedules is presented. The list shows the name of the schedule. It also shows how many times the schedule is used in scenarios and the dates for the creation and last modification. The non-drilling schedule form appears when the plus icon is clicked. The user should enter a unique name, description, fixed start date, and duration. The duration is entered together with low and high estimates. If the distribution type is set on triangular, the value of the duration can be assumed either as expected or as the most likely value.

See Figure 3.23 for more details.

Figure 3.23: Non-drilling schedule configuration.

3.5. Intervention Input and Configuration

A well intervention is an operation that is performed on the well during its operational lifetime. The objective for performing a well intervention is to perform maintenance on installed equipment, improve the well performance, improve the well integrity, or acquire data about the well or near-well area.

After it has been determined when the intervention job initiates and how long it lasts, we want to model the production increase. A production increment is added to the well's initial production potential. The model accounts for uncertainty related to both the base production and the increment. These are assumed to be independent draws. For reference, typical impacts of successful interventions are one or more of the following events:

- Higher oil rates
- Lower water rates or reduced water-cut
- Higher or lower gas rates
- No change (maintenance)

The impact of well interventions can be significant for a field or an asset and should therefore be reflected in its production forecast. It is common for oil companies to perform a significant number of well interventions.

An intervention operation impacts the well's production:

a) The well's production is temporarily shut in during the operation.
b) The well's performance can be improved as a result of the operation.

This is illustrated in the Figure 3.24 below, showing the concept of well intervention for one selected well in the deterministic model with a positive increment after the shut-in.

Figure 3.24 Intervention operation with production increment.

As the figure illustrates, the well has an initial base production potential. The well is shut-in temporarily for the intervention operation. When the well starts producing again, it continues producing its initial base production potential, plus an incremental production potential. The key uncertainties are related to timing of the production shut-in and the size of the incremental production potential.

The timing of the production shut-in is primarily decided by the well intervention schedule, which is described in further detail in Chapter 3.5.2. The key uncertainty driver for the timing of production shut-ins is the duration of the intervention and schedule dependencies. Figure 3.25 illustrates how the stochastic model includes uncertainty on both timing and the size of the production increment.
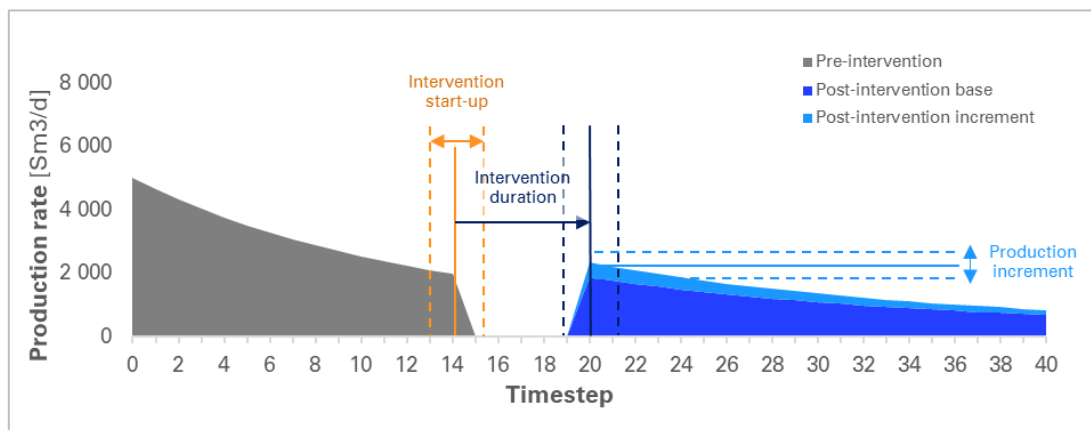
Figure 3.25 Intervention operation including uncertainties.

The improvement in the well's production potential following an intervention operation is represented as an incremental production profile. The well's resulting production potential thereby includes both its initial base production potential and an incremental production potential. Both types of production potential are prone to uncertainty and the incremental potential can be both positive and negative values. See illustrations in Figure 3.25.

Similarly, if an intervention is performed on an injection well, the resulting injection profile should consist of an initial base injection and an incremental injection.

If no incremental production or injection are provided, the increment can be assumed to be zero and the resulting profile equals the initial base profile. This is, for instance, applicable for maintenance interventions.

Running a stochastic analysis with 30 runs, Figure 3.26 shows the realizations for a selected well. Each individual run is created from a random draw according to the uncertainty settings for the intervention job on the well. Bear in mind that the result may also be affected by uncertainties applied to the previous wells in the schedule.

Figure 3.26 Intervention operation uncertainties with 30 runs.

In pForecast's result viewer, stochastic results are presented with percentiles as depicted in Figure 3.27 below. In this figure we have shown a selection with three percentiles. In chapter 5.3.3. more choices are available for viewing stochastic results from intervention operations.

Figure 3.27 Intervention operation with percentile results.

You can effectively do quality assurance on the impact of intervention campaigns by running scenarios with and without interventions and then compare the results. (Chapter 5.1.3).

In addition to the wells being intervened, nearby wells, or offset wells, may be temporarily shut in during the operation. These shut-ins are also reflected in pForecast's production forecasts.

3.5.1. Intervention Jobs

The intervention functionality in pForecast adds an opportunity to control timing of well interventions, including schedule dependencies between wells, intervention job durations and temporary shut-in of wells.

The following assumptions are made for the intervention scheduling functionality in pForecast:

- An intervention crew follows a schedule.
- One well is intervened at a time, per crew.
- The duration of an intervention job is uncertain.
- Once a well intervention job is completed, another job initiates.
- There can be multiple intervention crews working in parallel on the same asset.

For reference, the duration of a typical well intervention ranges from between 3 to 60 days. Figure 3.28 shows how to define one intervention job. The duration of an intervention job is a property of the intervention job and independent of when in the intervention schedule it occurs.

Figure 3.28 Defining an intervention job.

1. Under Forecast Properties, select the "Intervention Jobs" tab.

2. Click the "+" icon to create a new intervention job.

3. Provide a suitable name and description.

4. Select the intervened well.

5. Define duration for the job, and optionally include low and high values as well as any pre- or post-job shut-ins.

6. If low and high values are specified, select a distribution type.

7. Once non-optional settings are specified, you may save your work.

8. Additional settings will be explained with new figures.

A forecast well can have multiple interventions in its forecast period, but it is not common practice. pForecast offers an opportunity to test different intervention schedules quickly, i.e., you can define multiple schedules on the forecast level and then include various combinations of these in different scenarios of the forecast.

During execution of a well intervention, the intervened well's production (or injection) is shut in. For some jobs, the well is shut in in advance of the operation (pre-job shut-in). Production/injection starts up again after the intervention is completed, plus a "hook-up time" or "time to restart production" (post-job shut-in).

During execution of certain intervention jobs, production may be shut in on one or more nearby wells too, so called offset wells. The offset wells are not impacted by pre-job shut-ins and post-job shut-ins. The input for offset wells are illustrated in Figure 3.29

Figure 3.29 Including offset wells for an intervention job.

Each offset well has a set of properties that you may view and edit once you have added it to your intervention job's list. This is illustrated in Figure 3.30. Please note that if you change any values in the dialog window that appears, this will impact the characteristics for this well for all scenarios in this forecast.

For a job that is an intervention on a producer well, a production increment profile can be specified in order to model the incremental effect of the intervention on the production.

This profile is added to the base production profile for the well, starting from the production restart date given by the job end date and possibly a post-job shut-in.

pForecast supports two types of production increment profiles:

- Stream day rate vs delta month – Incremental potentials and gas lift rates given for a range of months relative to production restart.
- Decline curve – Incremental potentials given as an Arps formula decline curve with production restart as initial time.

Both types of increment profiles can be specified with uncertainty, i.e. with up to three values (low, mid/deterministic, and high) per parameter to be combined with a provided distribution type to create samples for the profile when performing a stochastic analysis.

Figure 3.30 View properties of an offset well.

3.5.1.1. Stream day rate vs delta month

You can see the details of a "Stream day rate vs delta month" increment profile in either chart or table form. The window opens by default showing the profile details as a chart. This is illustrated in Figure 3.31.

If the input data ranges over many months, it is most convenient to add input for the intervention profile through importing the data, something which is described in chapter 3.5.3.

Figure 3.31 Parameters for the profile details "Stream day rate vs delta month".

3.5.1.2.　Decline curve

You can see the details of a decline curve increment profile in either chart or table form. The window opens by default showing the profile details as a chart. This is illustrated in Figure 3.32.

Figure 3.32 Parameters for the profile details for "Decline curve".

Decline curve increment profiles are given by Arps formula. In the table input, you may add values for the following parameters:

- Initial rate, initial (nominal) decline and hyperbolic exponent (B factor) for the oil phase given by three values, one low value, one mid/deterministic value and one high value.
- For the gas phase, it is GOR times oil rate.
- And for the water phase, it is liquid minus oil rate.

For a decline curve you may check the resulting curves visually before submitting them, by switching to the chart output view. The table input is illustrated in Figure 3.33.

Figure 3.33 Table input for decline curve increment profiles.

There can also be slots in the intervention schedule that are not related to work on a producing or injecting well, so-called non-intervention activities. Examples of these are moving of crews and interventions on non-producing wells, e.g. plug and abandonment, rigging equipment and maintenance campaigns. Such slots will be defined as a job that is not connected to any well. An example of a non-intervention activity can be found in Figure 3.34.

Figure 3.34 Non-intervention activity added as part of an intervention project.

To view all the defined input parameters in a more compact way, you may switch to property view mode using the button just above the list of defined intervention jobs. This view is illustrated in Figure 3.35.
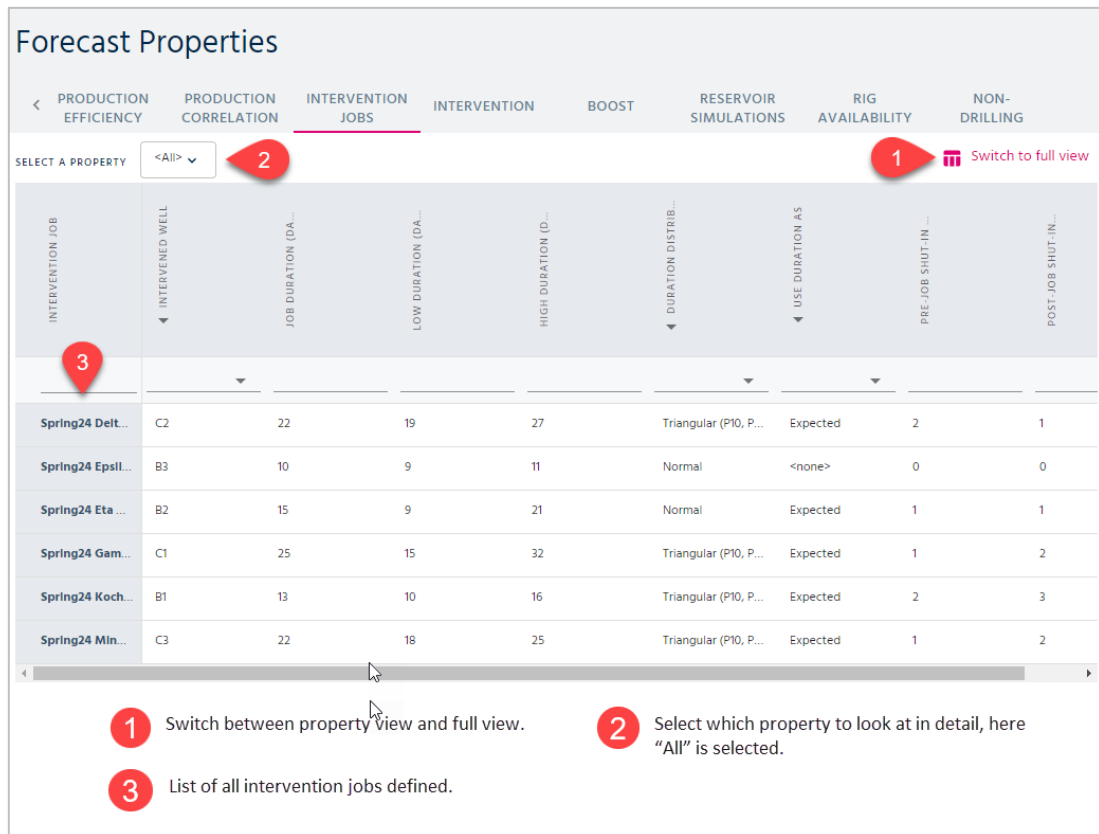
Figure 3.35 Intervention jobs shown in property view mode.

For the following three parameters you can view and control the selections in a detailed radio-button-based view:

- Intervened well
- Duration distribution
- Use duration as

We have selected the "Intervened well" property in the example shown in Figure 3.36, and this gives us an easy overview of which intervened well is connected to which job. The view also allows you to change a connection if it turns out to be incorrect. For the intervention job in question, simply click on the radio button corresponding to the new well.

Figure 3.36 Connections between intervention jobs and intervened wells.

In the same way you may see an overview of 'Duration distribution', as shown in Figure 3.37 and 'Use duration as', as shown in Figure 3.38.

Figure 3.37 Setup of "Duration distribution" property for each intervention job.

Figure 3.38 Setup of "Use duration as" property for each intervention job.

The set of intervention jobs for a forecast may be imported through pre-made Excel templates, see chapter 3.5.3. This is particularly useful in cases with extensive job lists.

3.5.2. Intervention Schedules

The well intervention schedule reflects the timing of the actual intervention operations. In addition to shutting in the well for the operation, a well may be shut in some time in advance of the operation (pre-job shut-in) and remain shut in for some time after the operation (post-job shut-in).

There may be uncertainty in the start date of an intervention job; this is related to the job being part of a schedule. For the end date we have uncertainty related to the schedule but also to the duration of the job. We may also add idle slots (non-intervention jobs) in the schedule, modeling for example moving of crews and other operations taking time but not being an intervention job per se.

The background and assumptions for modeling schedule dependencies and temporary production shut-ins in pForecast is as follows:

An intervention job is always executed within a schedule. The time it takes to perform an intervention job is prone to uncertainty. The timing of a particular intervention job is dependent on all the intervention jobs performed/lined up earlier in the schedule.

Typically, the timing uncertainty is largest for well intervention jobs late in the schedule.

For the production forecast, the timing of an intervention tells us when a well is not in production and when we can expect additional production after the intervention.

The well's production may be stopped days in advance of an intervention job to allow the pressure around the well to stabilize prior to any operations. Similarly, the well's production may be shut in some time after the intervention job due to slowness in restarting flow or additional work needed to reconnect the well to the flowline.

66

To define an intervention schedule, there must be at least one intervention job defined. Figure 3.39 illustrates how you get started with defining an intervention schedule, when setting it up through the dashboard's user interface.

Figure 3.39 Creating a new intervention schedule.

When working with interventions it is often useful to create a copy/clone of an already defined schedule. You can for example use this to analyze an alternative sequence of jobs or to see the impact of adding or subtracting maintenance items. Figure 3.40 illustrates how to create a copy/clone.

Figure 3.40 Creating a copy of an intervention schedule.

For a quality assurance check or a last-minute update to the schedule, you can access each individual job of the schedule from the schedule directly. This is illustrated in the dialog window in Figure 3.41. It is possible to update any of the properties that are presented in the view, but please note that changes will be reflected in *all instances* of this intervention job, also in other schedules where it is included.

Figure 3.41 Last-minute changes to jobs of a schedule.

The setup of intervention schedules and their contained jobs can be imported using pre-made Excel templates, see chapter 3.5.3.

3.5.3. Import Intervention

You may also import a list of both intervention schedules and intervention jobs using an Excel workbook template. Download the template as described in Figure 3.42. When using the template, you will work with the same parameters that are described as part of the user interface, see chapter 3.5.1 and chapter 3.5.2. All the parameters are listed on the Excel sheet named Info, which also lists what type of information to input for each parameter on each of the different worksheets.

Figure 3.42 Find the Excel workbook template to import intervention setup.

On the Intervention Jobs worksheet, you may add a list of intervention jobs to import into pForecast. You need to supply the information as described in Figure 3.43. Our example illustrates three different intervention jobs, the first uses a Triangular distribution to describe the uncertainty of the duration of the job. The second one uses a Normal distribution. The last job is a non-intervention job, therefore no intervened well is connected to this job.

Figure 3.43 Excel input data on the Intervention Jobs sheet.

To import the data from Excel, follow the instructions shown in Figure 3.44. The data will be validated upon import, and in our example from above, we can see that there are two

problems with the input data. pForecast offers sensible error messages that easily allow you to identify the problem at hand and rectify the problem.
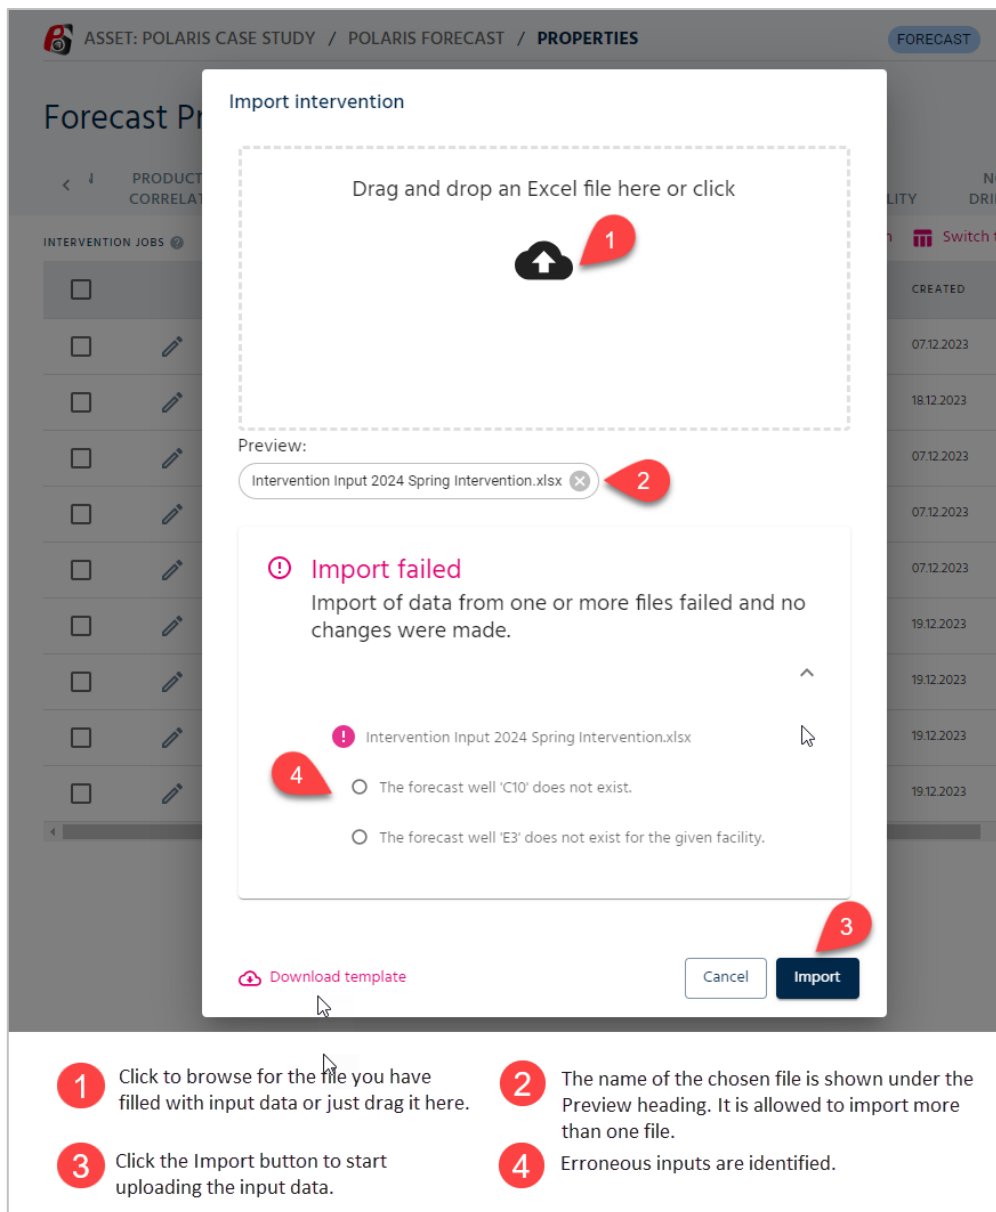

Figure 3.44 Validation of imported intervention data.

In our example the well mentioned for our first intervention job should have been well "C4", while the correct facility notation for the second job was "FPSO Vela" instead of "Vela" only. If you now try to import again, you will receive feedback that the file is successfully imported. This is illustrated in Figure 3.45. A list of which input data was added is also presented.

Figure 3.45 Successful import of intervention jobs.

On the Intervention Schedules worksheet, you may add a list of intervention schedules that can be imported into pForecast in one operation. You need to supply the information as described in Figure 3.46. Our example includes one intervention schedule, and we always need to specify which intervention job should be the first to be executed in the schedule. We have listed this new intervention job on the previous tab of intervention jobs, and the import of the schedule will then give us the familiar complete message, illustrated in Figure 3.47.


Figure 3.46 Input data for intervention schedules.

Figure 3.47 Successful import of intervention schedules.

Also on the Intervention Schedules worksheet, you may import monthly data for a production increment profile of type "Stream day rate vs delta month". Since this profile type typically contains data for many months, Excel import is the most convenient way of adding data for such production increments. See Figure 3.46.

Figure 3.48 Input data for "Stream day rate vs delta month" data.

4. Scenario Configuration

Scenario configuration is the last level of the pForecast hierarchical structure. A scenario is essentially a consumer of existing data and can also be seen as a version of a forecast. At this stage, it is possible to run different scenarios of the forecast and visualize, analyze, and compare results.

4.1. Create a Scenario

In order to create a scenario, the user should select an active forecast. After selecting the forecast, you will see the following overview page, shown in Figure 4.1, where you can create a new scenario. After clicking the "Create the new scenario" button, the page in Figure 4.2 will be shown, allowing high-level scenario configuration.

Figure 4.1: Creating a scenario (part 1).

Figure 4.2: Creating a scenario (part 2).

4.2. Scenario Properties

After making a new scenario, the user can easily access all scenario properties as illustrated in Figure 4.3, and start the configuration process.
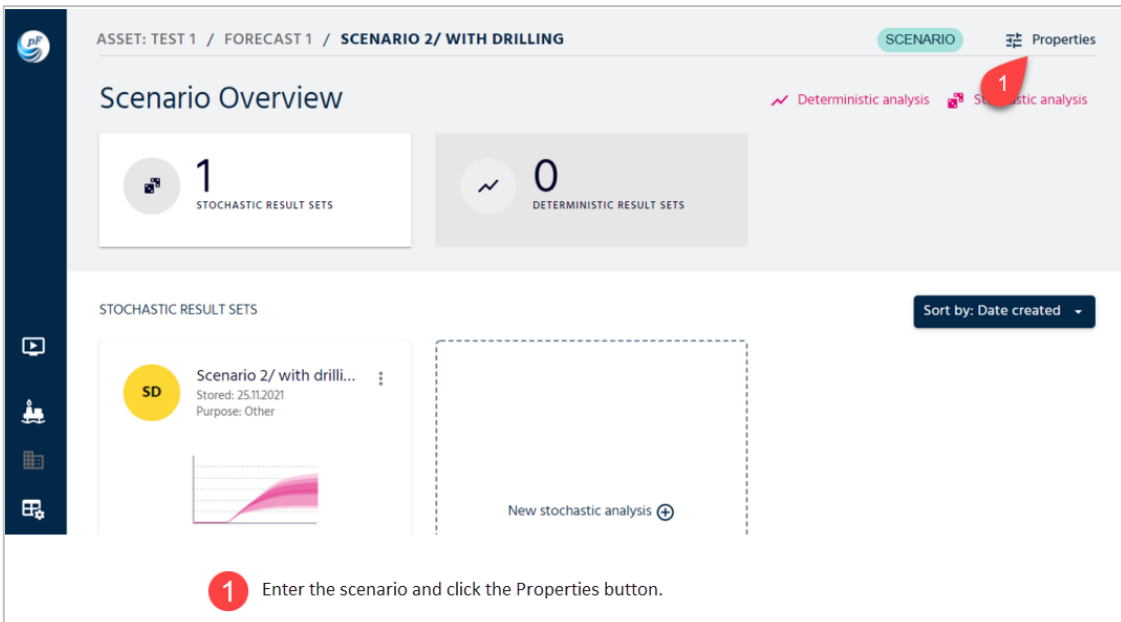

Figure 4.3: Access to the scenario properties.

4.2.1. Scenario

Under Scenario Properties, when the "Scenario" tab is selected, a list of attributes for the current scenario is shown. You can choose to set both default production efficiency (PE) and define dependencies also called correlations. A PE forecast gives monthly and annual estimates of the production efficiency as well as potential shutdowns. The PE forecast selected here (typically a PE forecast not linked to a facility) will be used for all facilities that are configured to use the default PE forecast rather than a specific one. See Figure 4.5 for more details.

If the "Use scaling factor" switch is switched on, all drilling times will be multiplied with the factor, a number between zero and 1, given in the involved drilling schedules.

Specify the dependency, correlation, between various events/characteristics within a simulation run as a factor between zero, fully independent, and on, fully dependent. These factors are only used when running stochastic analyses.

Figure 4.4 Define default PE and dependencies/correlations for scenarios.

4.2.2. Wells in operation

Under the scenario properties, when the "Wells in operation" tab is selected, a list of existing wells in the current forecast is shown. You can choose to include or exclude the wells in the current scenario. See Figure 4.5: Activation of wells in operation, a.k.a. existing wells.Figure 4.5 for more details.

Figure 4.5: Activation of wells in operation, a.k.a. existing wells.

4.2.3. Drilling

The next step is about wells to be drilled. When the "Drilling" tab is selected, you can activate previously defined drilling schedules. It is worth mentioning that if a drilling schedule is activated, all forecast wells belonging to that drilling schedule will be automatically included in the scenario. Figure 4.6 gives more details.

Figure 4.6: Activation of wells to be drilled.

4.2.4. Capacity

As previously mentioned, see section 3.3.4, a set of capacity constraints can be defined for a forecast. A constraint is given with maximum rates for production/injection as well as regularity factors for injection and is valid from a given point in time and onwards until a new constraint is activated. Each constraint is connected to a cluster in the hierarchical capacity cluster structure. Use the toggles under the "Included" column to determine which of these constraints to include in the current scenario. See Figure 4.7 for more details.

Figure 4.7 Activation of constraints.

4.2.5. Separation

For a forecast, you can define a set of separation forecasts, each of them giving monthly estimates for separation of phases prior to routing the next level of the cluster hierarchy. Use the dropdowns of the separation forecast column to determine which separation forecast to use for each involved child facility in the current scenario. See more details in Figure 4.8.

Figure 4.8 Add separation for your facility.

4.2.6. Production Efficiency

The user can select the scenario's active production efficiencies by selecting the "Production efficiency" tab under the scenario properties. The user is asked to provide a PE forecast for each facility. Different facilities may use the same PE forecast if appropriate. Each of the facilities in the model is shown in the list, and the PE forecast is chosen from the drop-down menu. See Figure 4.9 for more details.

Figure 4.9: Activation of PE forecasts.

4.2.7. Intervention

To take intervention projects/campaigns into account for a forecast, you must first define both intervention jobs and one or more containing intervention schedules for the forecast. This is described earlier in chapter 3.5.

In the Intervention tab for Scenario Properties, you are presented with a cockpit to include one or more intervention-related campaigns. You may use the toggles in the Included column to determine which of the listed interventions schedules should be included in the current scenario. In this way it is easy to evaluate two alternative schedules by switching the corresponding schedule on and off in two separate scenarios.

In our example shown in Figure 4.10 Activate Intervention schedules, we have multiple schedules to implement an upcoming intervention project for an upgrade during the spring of 2024. Each schedule has a slightly different setup, and we will run different scenarios to compare the outcome of each schedule. Note that a specific intervention job can only be present once in a scenario. This means that in our case where the alternative schedules contain some of the same jobs, only one of the alternative schedules can be active in a scenario. In general, as long as there are no common intervention jobs, it is allowed with multiple active schedules in the same scenario.

Figure 4.10 Activate Intervention schedules for scenarios.

When you open the details of an intervention schedule, a pop-up window will appear as illustrated in Figure 4.11. The sequence (the execution order) of jobs in the intervention schedule is shown as a list.

Figure 4.11 View sequence of an intervention schedule.

4.2.8. Reporting

As mentioned earlier, see chapter 2.3, a set of reporting schemes can be defined for an asset. In order to configure reporting for a given scenario, the user should select the "Reporting" tab under the scenario properties. You can choose which of these schemes to consider for the reports of the current scenario. In order to add, edit or delete group schemes, you should go to the asset level and perform any necessary changes there. Figure 4.12 gives more details.

Figure 4.12: Activation of reporting schemes.

5. Running the Simulation

pForecast can be run in two different modes, deterministic and stochastic. In the deterministic mode, it is possible to review the reference case without considering any uncertainties. In the stochastic mode, the Monte Carlo approach is used to perform uncertainty analyses.

5.1. Running a deterministic simulation

Running a scenario in deterministic mode allows performing an analysis of a reference case without considering uncertainties. Figure 5.1 shows how to enter the deterministic analysis mode, and Figure 5.2 shows how to start running a deterministic scenario.
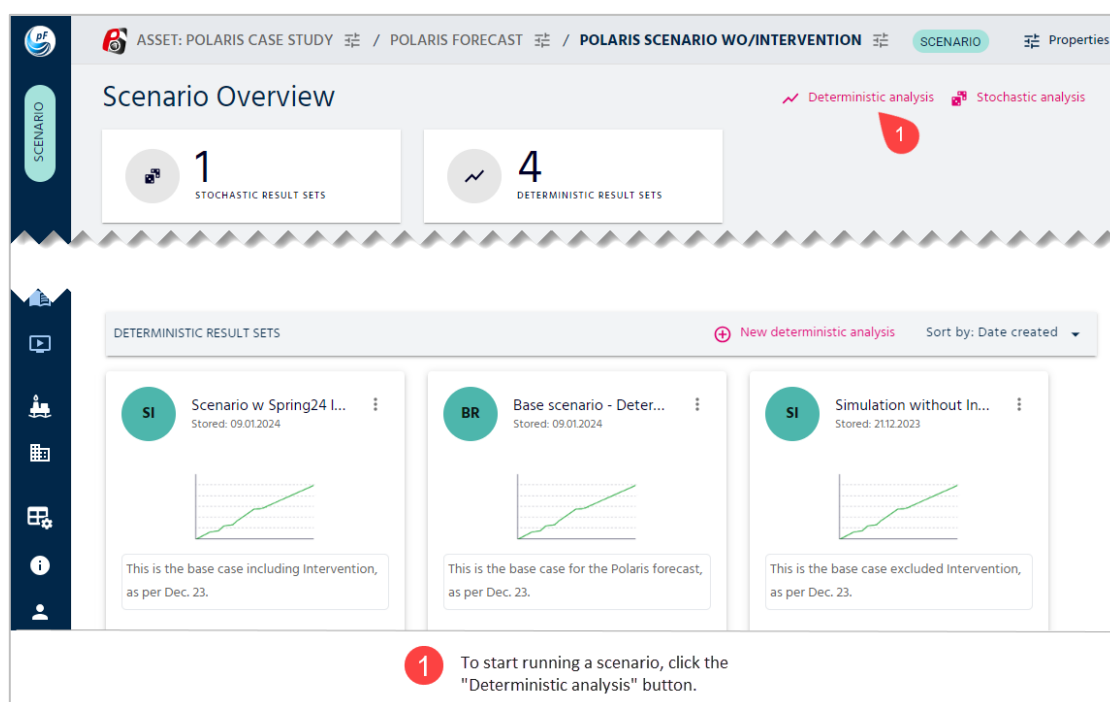
Figure 5.1 Running a deterministic simulation (part one).

Figure 5.2 Running a deterministic simulation (part two).

When the simulation is complete, the user is presented directly with the production profile results.

5.1.1. Profiles

In pForecast, the user can choose between different report groups for viewing results. It is possible to present both calendar day rate, volume and accumulated volume, as well as production efficiency (PE). Figure 5.3 shows how to navigate the result viewer under the "Profiles" tab.
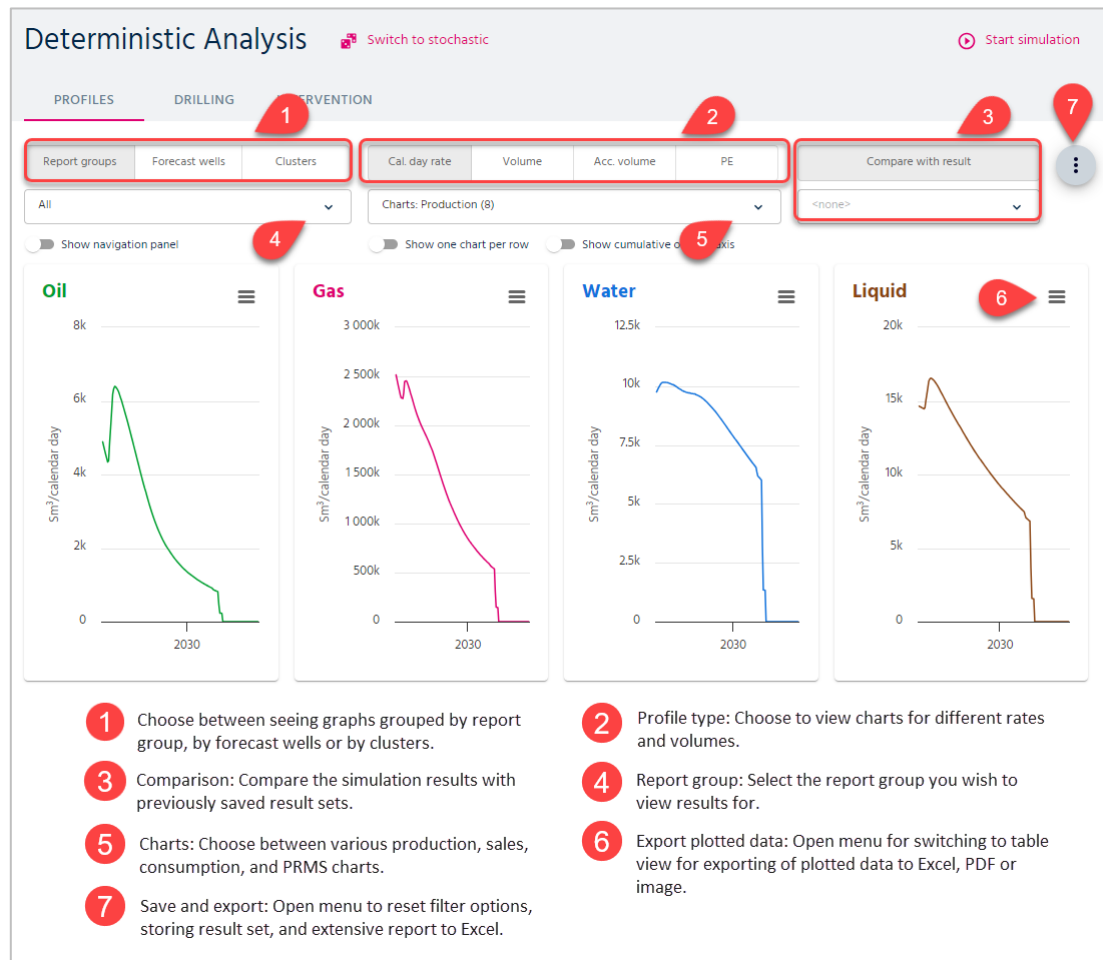
Figure 5.3: The pForecast result viewer for a deterministic analysis.

In addition to charts for the three basic phases (oil, gas, and water), you can also select other charts like GOR (gas oil ratio) and WCT (water cut) and injection. Figure 5.4 shows the selection of available charts in the pForecast result viewer for deterministic analyses.

Figure 5.4: Available charts for deterministic analysis in pForecast.

5.1.2. Drilling

For Exploration and Production (E&P) companies, it is essential to consider the time related to drilling, the time before a well is put into production. To meet this need, pForecast considers the drilling programs involved, automatically generates resulting Gantt charts for drilling schedules and shows how drilling targets propagate.
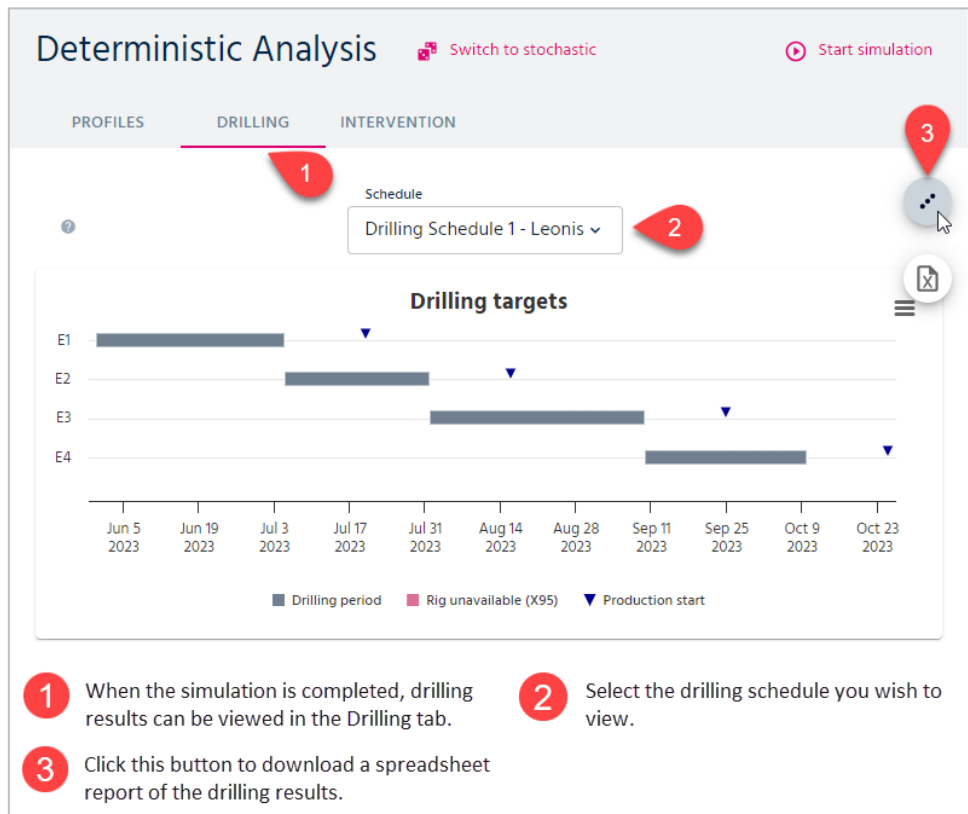
Figure 5.5: Gantt charts for drilling schedules in a deterministic analysis.

You can view the expected drilling schedules by selecting the Drilling tab under Deterministic Analysis. If there is more than one drilling schedule, you select between them by using the schedule dropdown list. Start of drilling, end of drilling, and start of production are visualized for each well. Figure 5.5 provides more details.

5.1.3. Intervention

You can see the expected intervention schedules by selecting the Intervention tab under Deterministic Analysis.

In Figure 5.6, you will see intervention activities for interventions on producers or injectors indicated with dark blue lines. Non-intervention activities (crew movements, etc.) are shown using slanted blue lines.

Diamonds and triangles indicate production shut-in for intervention jobs, starting at the diamond icon and ending at the triangle icon.

Figure 5.6 Gantt charts for intervention schedules in a deterministic analysis.

5.2. Deterministic result sets

If you save a result set for your simulation, you can find these data in the Scenario Overview window as shown in Figure 5.7.

When you open this result set, you may analyze the data for each well, on a yearly or monthly basis. It is possible to zoom in to study details of a graph by using the slider underneath the graph. This is illustrated in Figure 5.8.
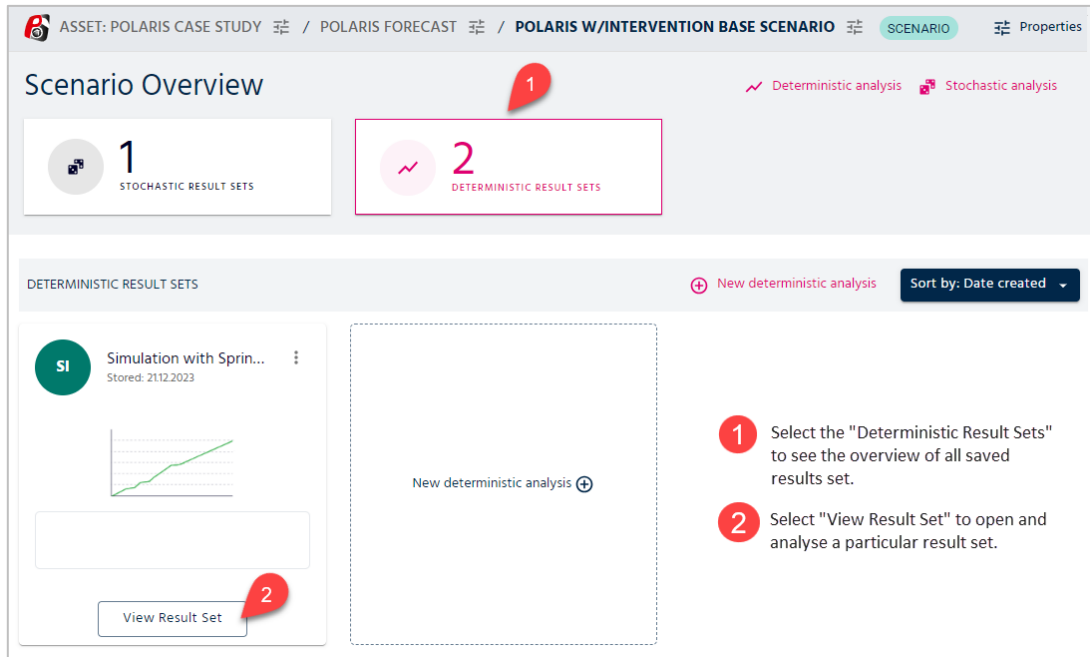
Figure 5.7 Saved result sets for a deterministic analysis of a scenario.

Saving results sets are very useful, because you may compare one result set towards another simulation as described in chapter 5.1.1, shown in Figure 5.3 point 3.

Figure 5.8 Result set for a deterministic analysis.

5.3. Running a stochastic simulation

In pForecast there is also a second mode, the stochastic mode, where a Monte Carlo approach is used to perform an uncertainty analysis. Figure 5.9 and Figure 5.10 show how to run a stochastic simulation.
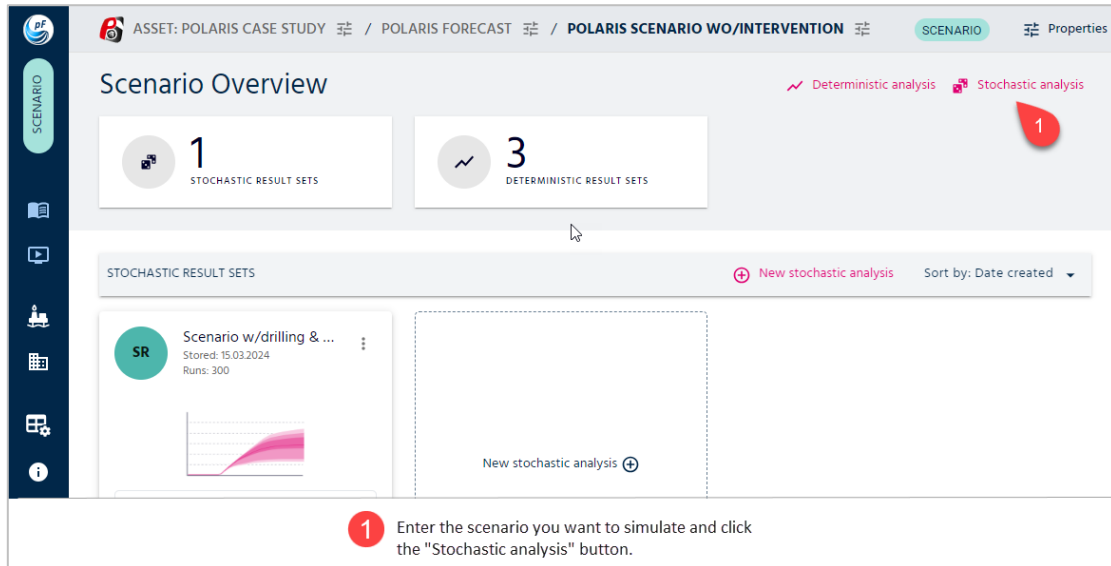
Figure 5.9: Running a stochastic simulation (part one).

After clicking the "Stochastic analysis" button, you will see the following page (Figure 5.10), where you can adjust the number of runs and start the simulation. The user can run between 10 to 2000 simulations and derive production profiles for different scenarios. It is recommended to run at least 200 runs for reporting purposes.
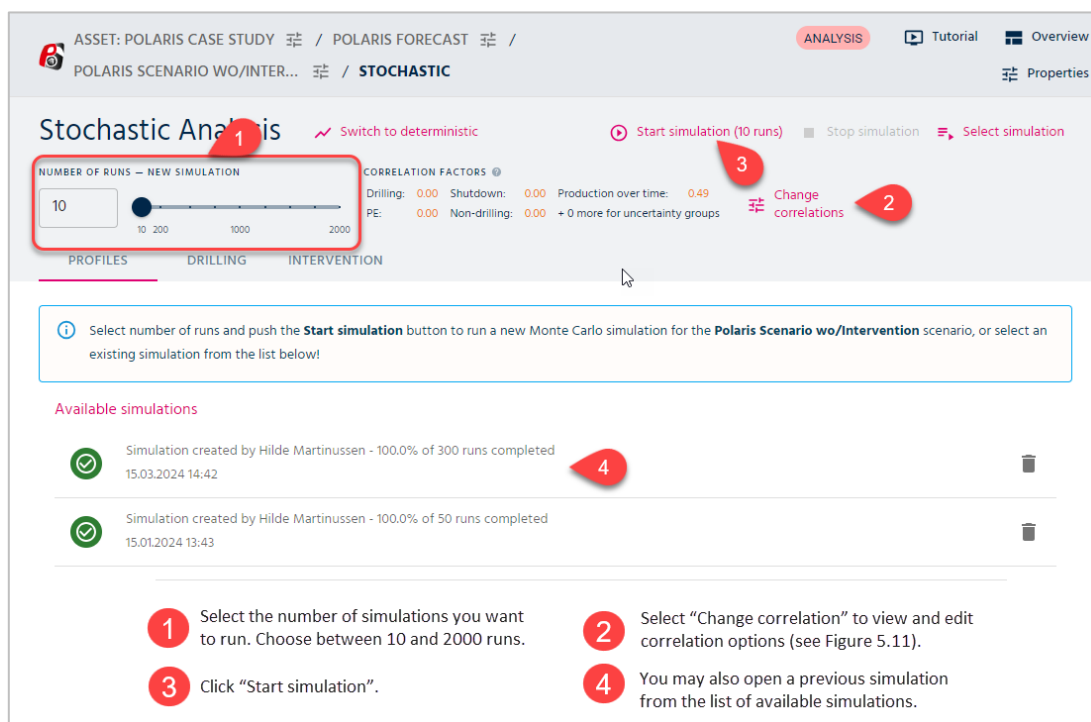
Figure 5.10: Running a stochastic simulation (part two).

Before running the stochastic simulation, the user can specify how a series of events/characteristics in the simulation relate to each other. Figure 5.11 shows how to set up the dependencies within a simulation run as a factor between zero (fully independent) and one (entirely dependent). For uncertainty groups, it is also possible to set correlation factors between minus one and zero to model negative dependencies.
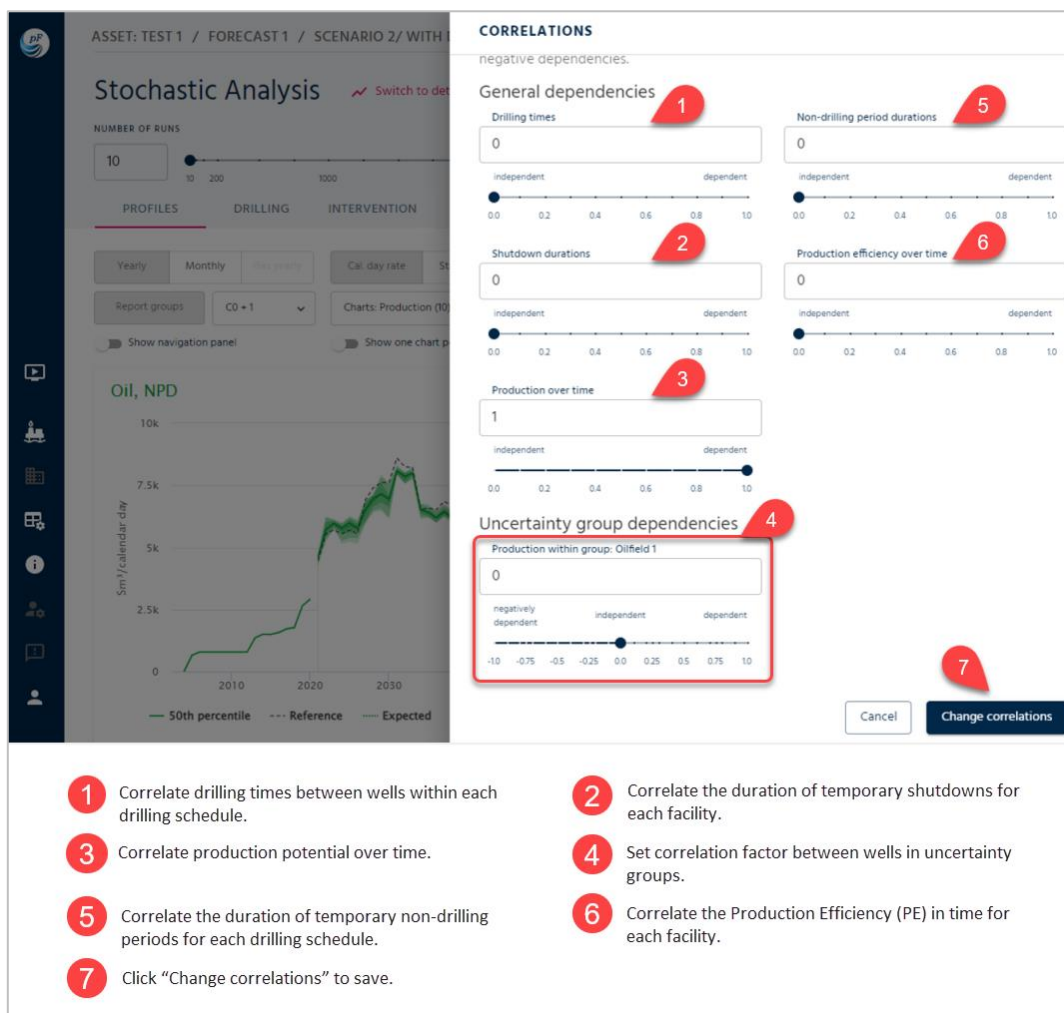
Figure 5.11: Description of correlation options.

When the simulation is complete, you are presented with production profile results showing a reference case as well as uncertainty estimates.

5.3.1. Profiles

In the result viewer for a stochastic analysis, you can choose between yearly or monthly plot granularity. Results are shown by report groups. It is possible to present both stream day rates, assuming 100 percent uptime, and calendar day rates accounting for planned and unplanned downtime. Figure 5.12 shows how to navigate the result viewer under the Profiles tab of Stochastic Analysis.
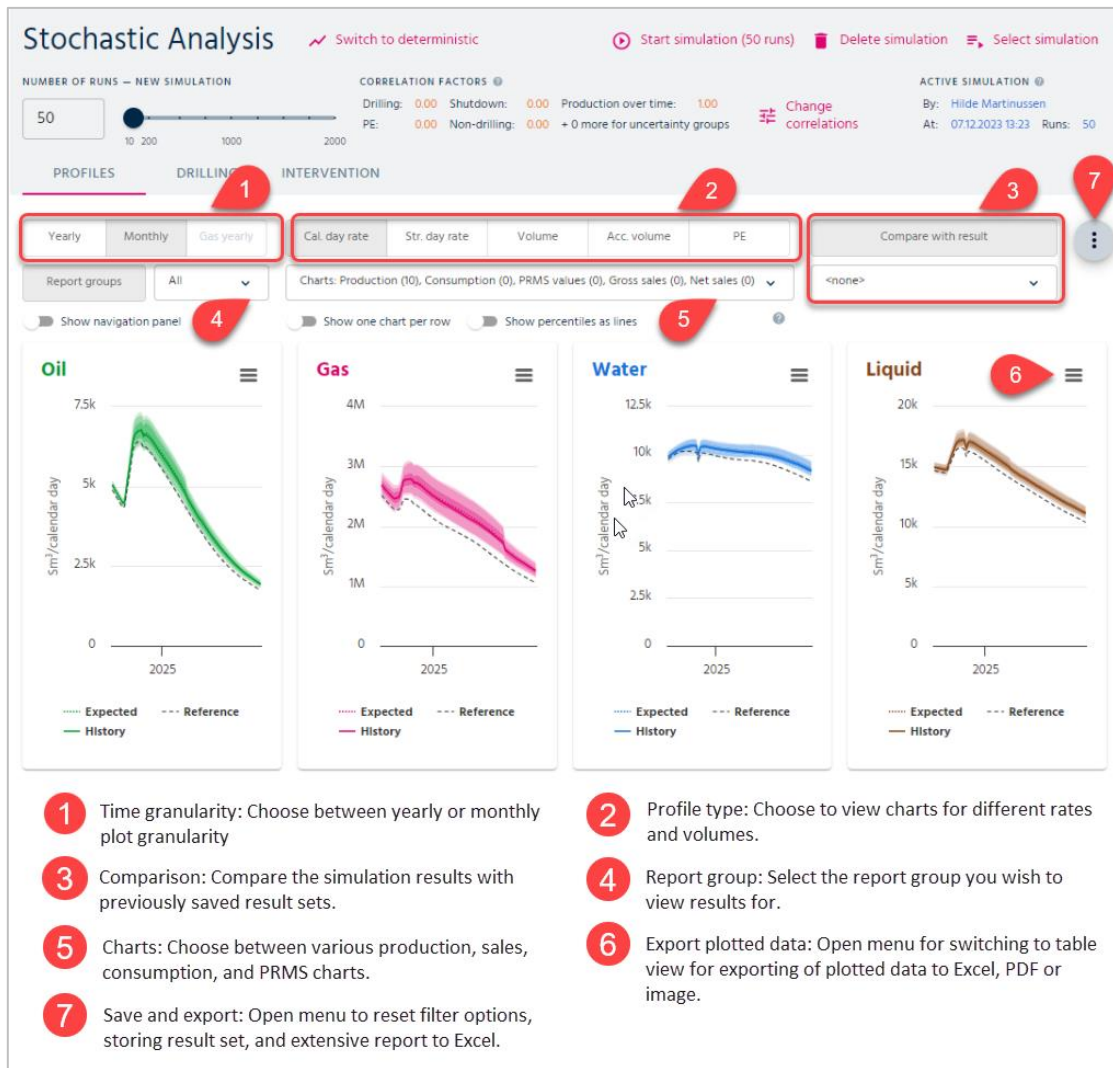
Figure 5.12: pForecast result viewer for a stochastic analysis.

In addition to production profiles, the user can easily generate other profiles, i.e., injection, fuel, and flare, as well as gross and net sales. Figure 5.13 shows an overview of available profiles in the pForecast software.

Figure 5.13 Available profiles in pForecast.

5.3.2. Drilling

For Exploration and Production (E&P) companies, it is essential to take into account the uncertainties related to drilling. To meet this need, pForecast considers uncertainties involved in drilling programs, automatically generates Gantt charts for drilling schedules and shows how uncertainty propagates over drilling targets. You can view the expected drilling schedules by selecting the "Drilling" tab under Stochastic Analysis. If there is more than one drilling schedule, you can select one of them by using the schedule dropdown list. The early start of drilling, expected start of drilling, expected end of drilling, expected start of production, and the late end of drilling are visualized for each well.

Figure 5.14 provides more details.
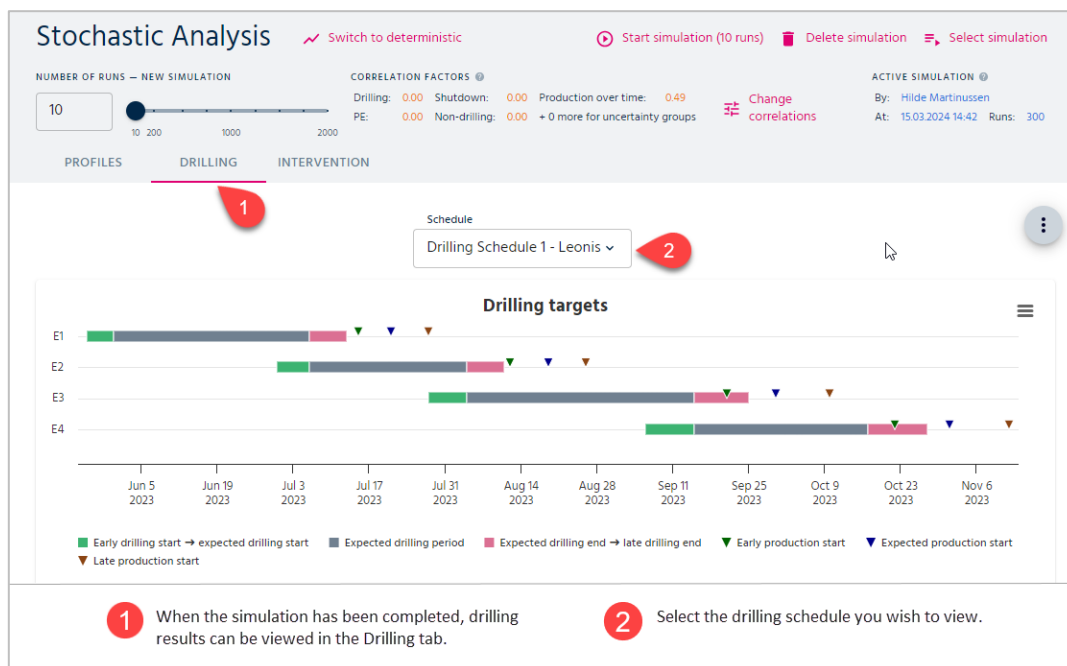
Figure 5.14: Gantt charts for drilling schedules in a stochastic analysis.

5.3.3. Intervention

You can see the expected intervention schedules by selecting the Intervention tab under Stochastic Analysis.

In Figure 5.15, you will see intervention activities for interventions on producers or injectors indicated with blue lines with three different shades (representing uncertainty estimates). Non-intervention activities (crew movements, etc.) are shown using slanted blue lines with the same three shades.

Diamonds and triangles indicate production shut-in for intervention jobs, starting at the diamond icon and ending at the triangle icon. The colors of the diamond and triangle icons indicate different uncertainty estimates for shut-in start and end.
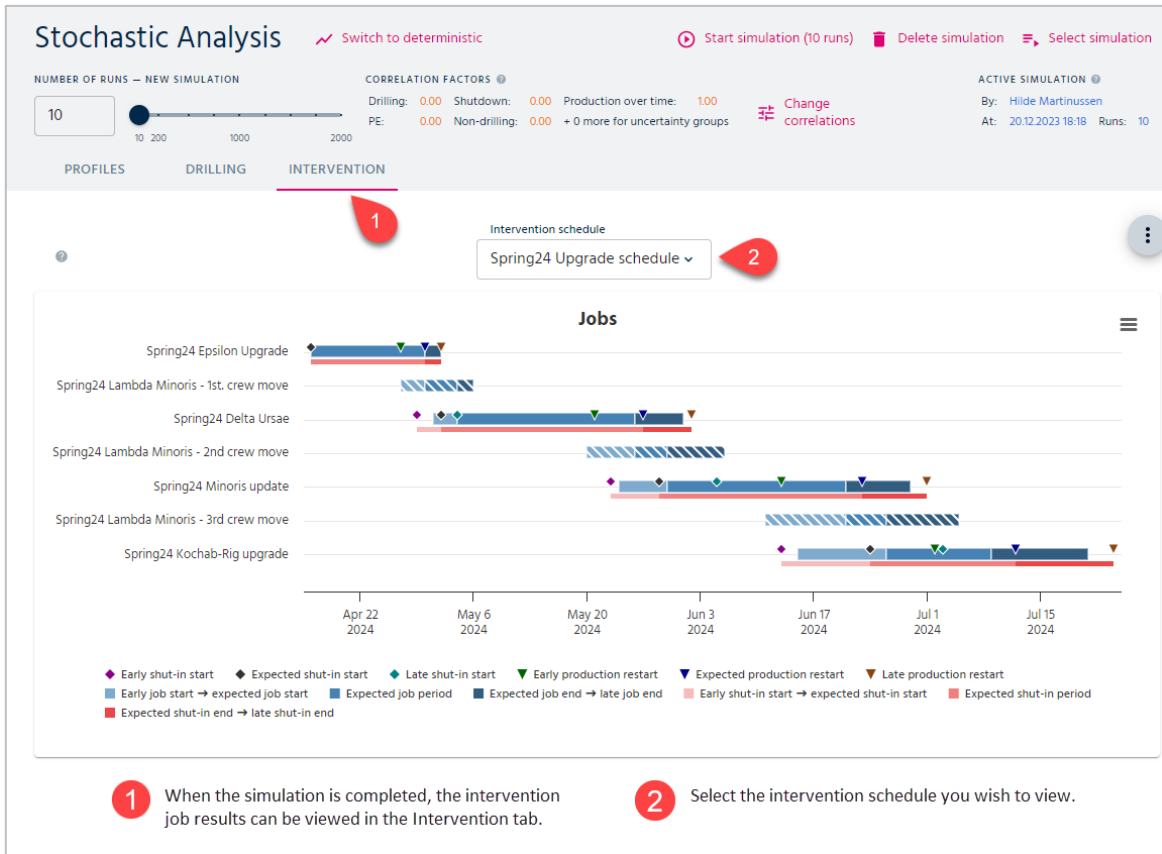
Figure 5.15 Gantt charts for intervention schedules in a stochastic analysis.

5.4. Stochastic result sets

If you save a result set of a stochastic analysis, you can find these data in the Scenario Overview window as shown in Figure 5.16.
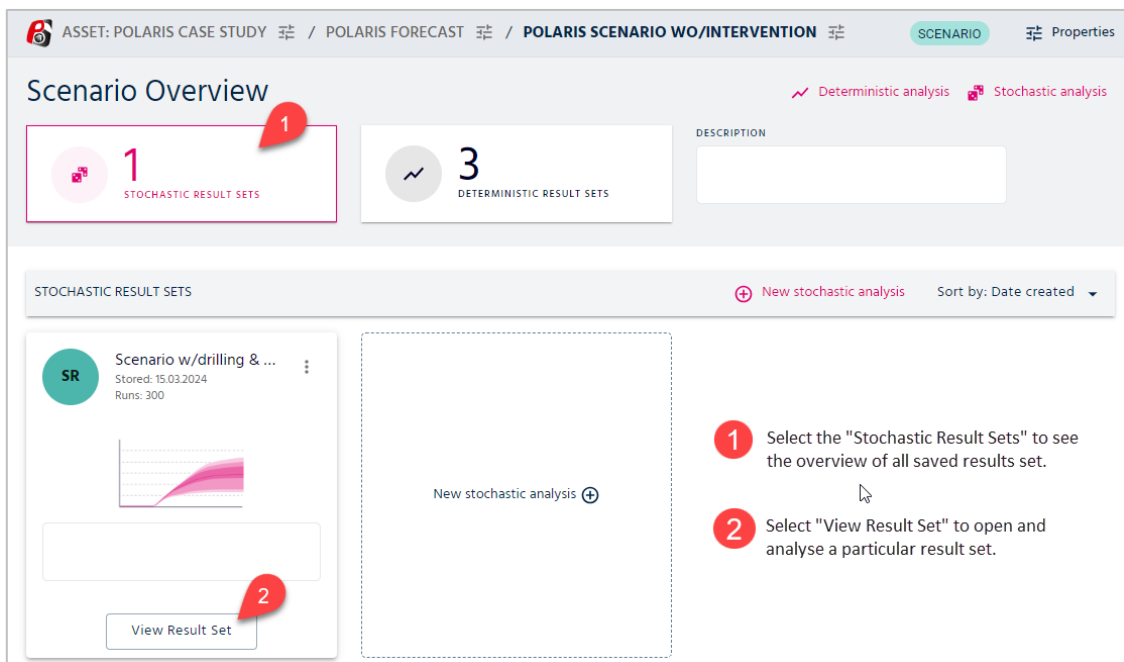

Figure 5.16 Saved result sets for a stochastic analysis.

When you open a saved result set, you will find the same information and same tabs as described in sections 5.3.1 - 5.3.3.

A great feature of pForecast is the possibility to compare one result set with another. This is illustrated in Figure 5.17. The graph shows one result set without an intervention, the other result set includes a suggested intervention and allows you to compare the benefits or drawbacks of your planned/scheduled intervention.
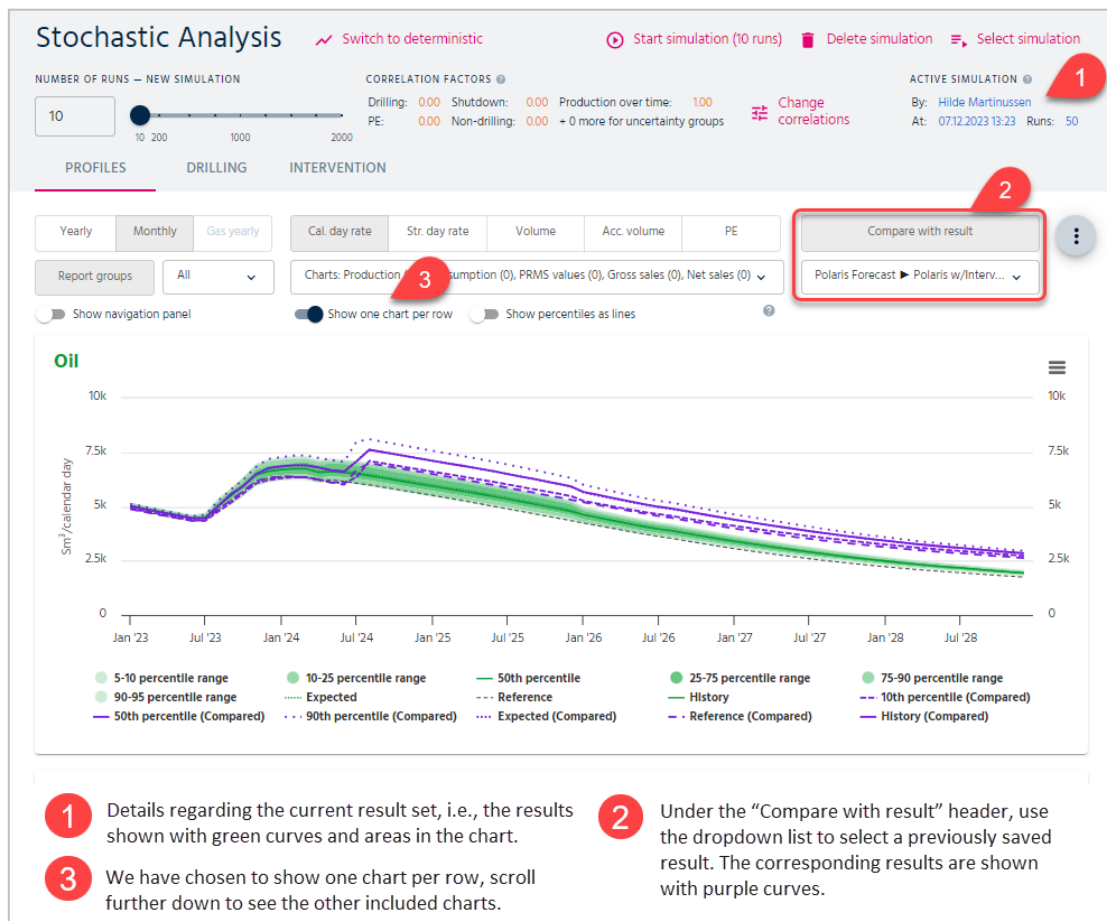

Figure 5.17 Comparing result sets.

6. Corporate Analysis

In most cases, companies own multiple assets, and it is desired to see the cumulation of values for all or a few of these assets. The pForecast software can support corporate roll-up analyses, which are cumulations of values for a selected set of assets. The users can create as many roll-ups as they wish for corporate-wide analyses and reports. Figure 6.1 through Figure 6.6 illustrate how to create a new roll-up.
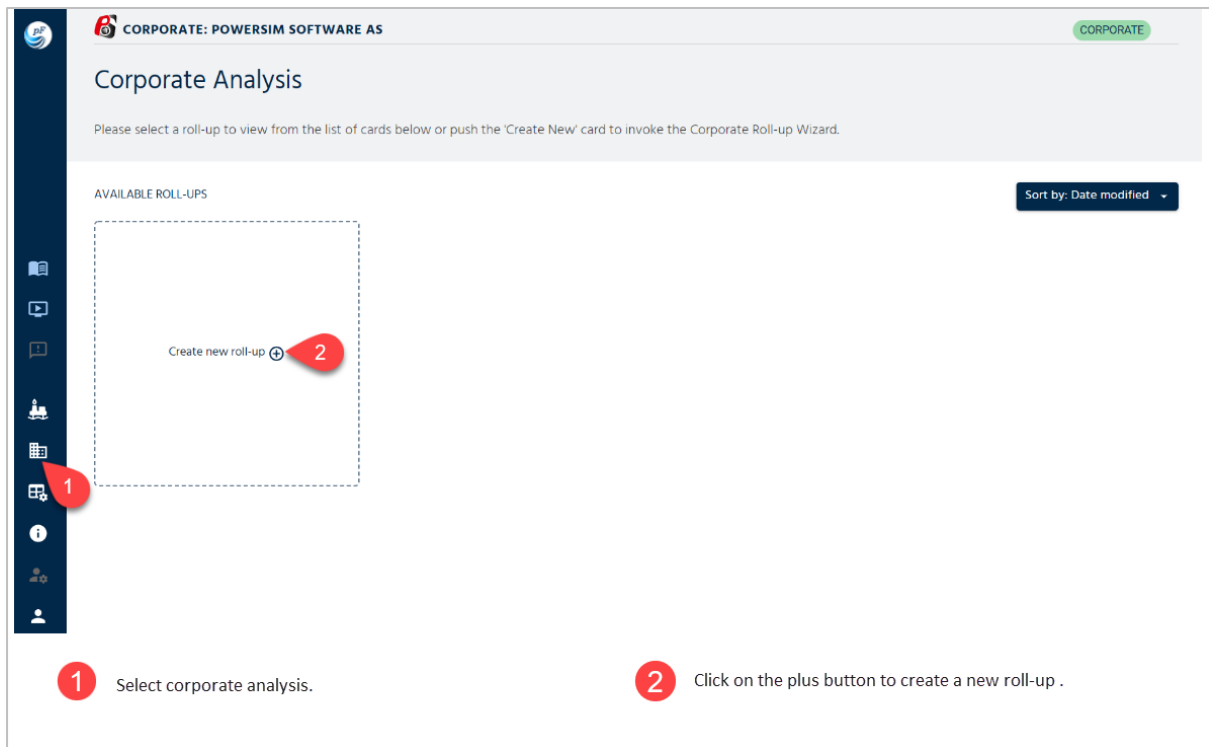

Figure 6.1: Creating a new corporate roll-up (part 1).

When creating a new roll-up, the following wizard page will be shown, in which the user can provide a name and description. It is also possible to insert the number of runs. The user can select a number between 100 and 10000. An automatic number will be used if this box is left empty. The automatic number is determined by selecting the larger value between 1000 and the number with the highest occurrence in the input (result sets). See Figure 6.2 for more details. Since corporate roll-ups consist of more than one asset, you will be asked to select the assets of interest in the next step. See Figure 6.3.

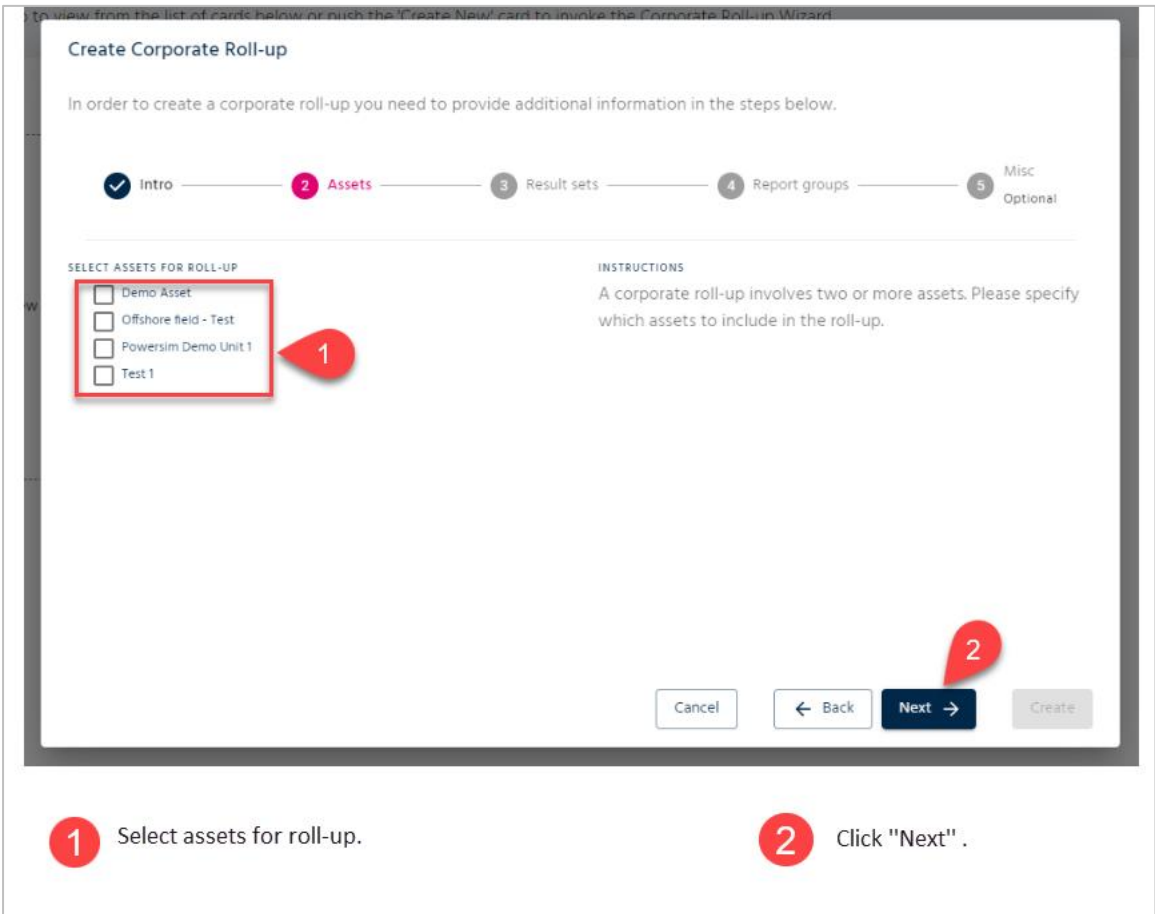Figure 6.2: Creating a new corporate roll-up (part 2).

Figure 6.3: Creating a new corporate roll-up (part 3).

In the next step, for each of the assets selected in the previous step, a stochastic result set that is to produce the results for the asset during corporate roll-up should be specified. Figure 6.4 gives more details.

Figure 6.4: Creating a new corporate roll-up (part 4).

Now for each of the assets and result sets selected in the previous step, you need to specify which report groups from the result set that are to represent all wells, existing wells, and new wells of the asset during a roll-up. See Figure 6.5.

Figure 6.5: Creating a new corporate roll-up (part 5).

In the last step and after having provided the necessary information for the roll-up, you can specify if you want to view the roll-up once the wizard closes. Figure 6.6 shows how to do so.

Figure 6.6: Creating a new corporate roll-up (part 6).

7. Incremental Profiles

Introducing new fields, new wells, or performing facility upgrades and maintenance affect the production of already producing fields. To support project investment decisions, incremental profile analysis can be performed to quantify the net production of the entire project.

Incremental profiles are based on two scenarios, a base case, and a new case, within the same forecast. In fact, the total production of the incremental profile is the difference between the base case and the new case. Figure 7.1 through Figure 7.6 illustrate how to create incremental profiles analyses.
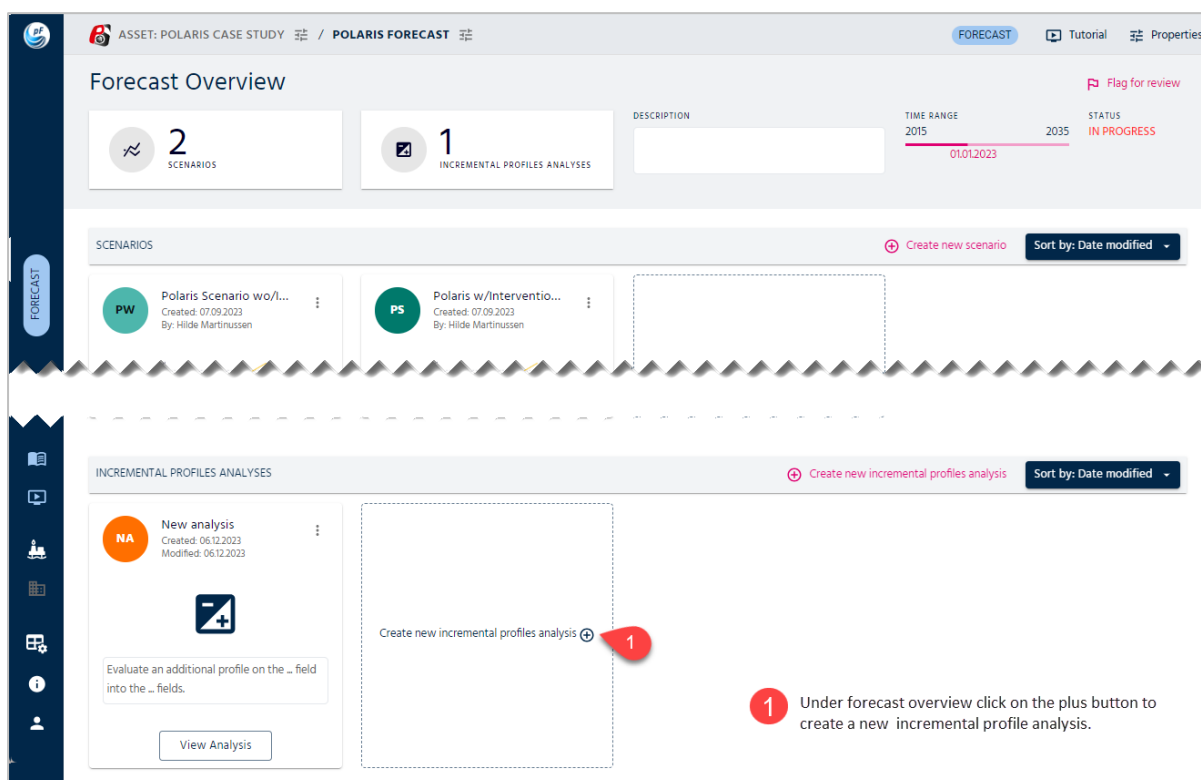

Figure 7.1: Creating a new incremental profiles analysis.

When clicking the 'Create new incremental profiles analysis' card, the following wizard page will be shown, in which the user can provide a name, description, and the number of runs for the Monte Carlo simulation behind the analysis. The user can select a number between 10 and 2000. See Figure 7.2 for more details.

Please note that if you select 50 as the number of runs, this means that your Incremental Profiles Analysis will run both scenarios 50 times, that is 100 runs all together.

Also, please note that in this type of analysis, the two stochastic simulations are run with the same draw of random values for mutual entities, which makes them mathematically viable to compare with each other. This is one of pForecast's greatest strengths.
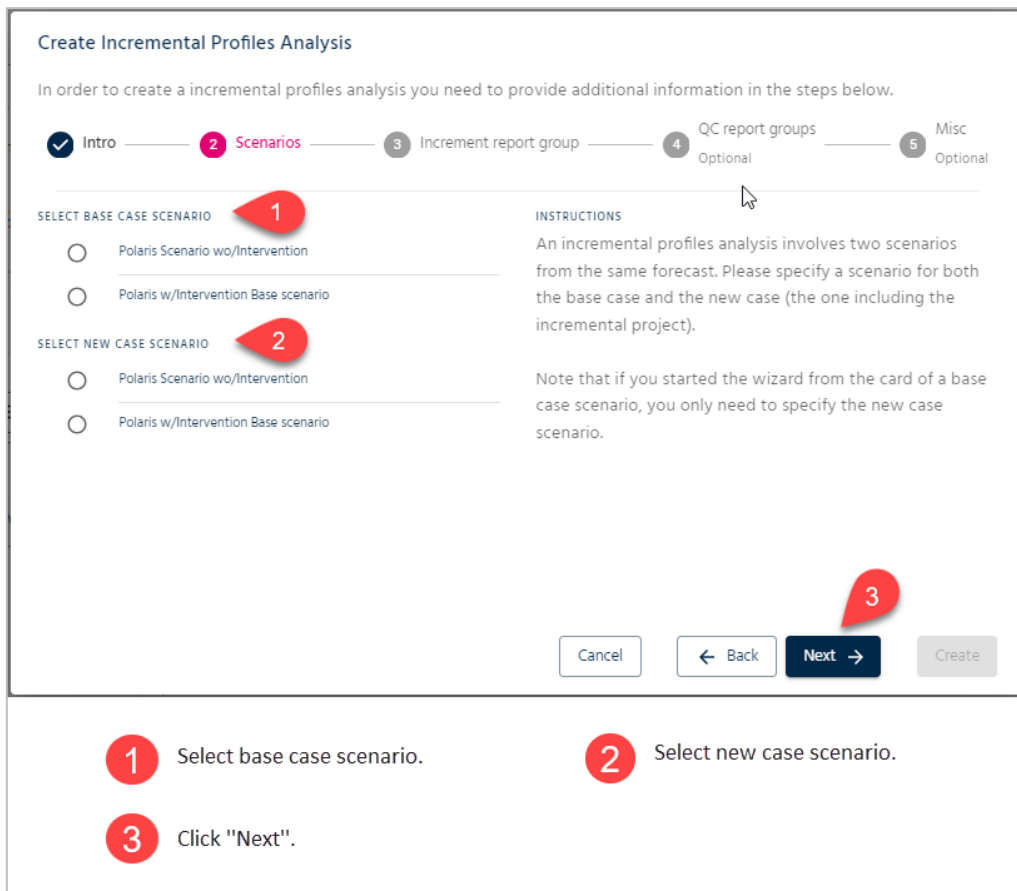
Figure 7.2: Apply a name and decide number of runs.

Since an incremental profiles analysis consists of two scenarios, in the next step, you need to select both a base and new case (see Figure 7.3). Note that the new case scenario typically include an incremental project.

Figure 7.3: Selecting scenarios for incremental profile analysis.

Next, you are required to select the report group which constitutes the basis for the increment calculations. This is normally a report group that represents the total project. In the incremental profiles analysis result viewer, the selected report group will result in three available report groups:
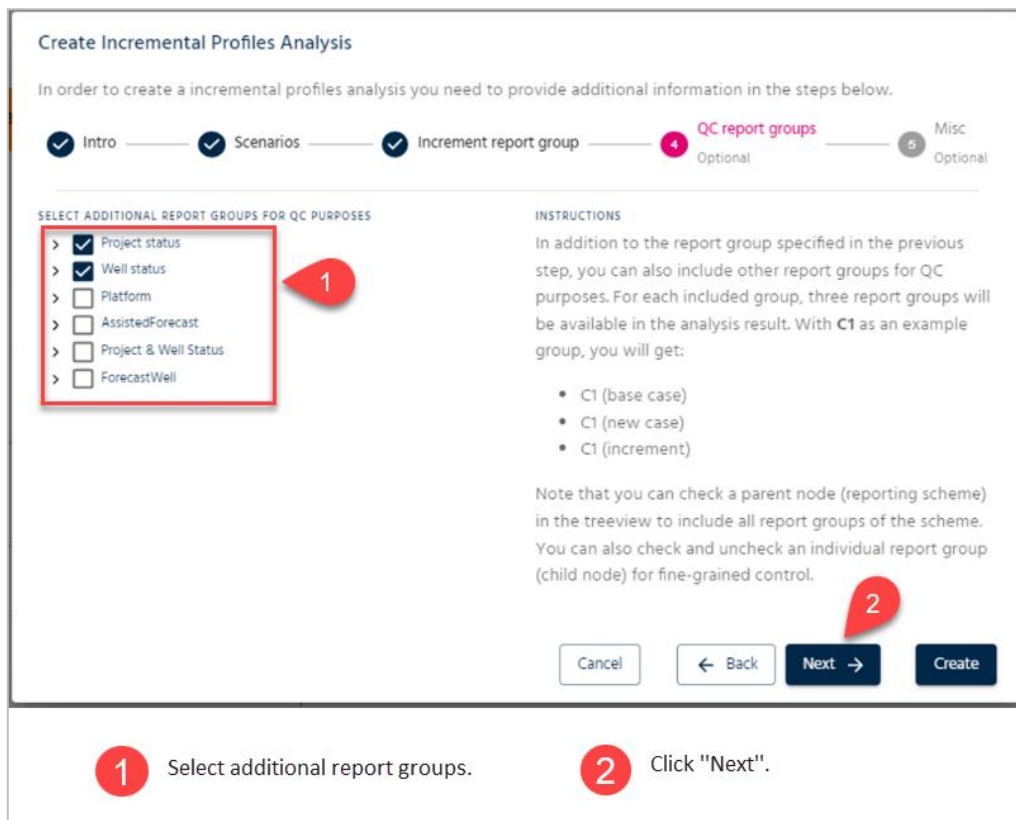
1. Report group for the base case gross value.
2. Report group for the new case gross value.
3. Report group for the increment value.

Figure 7.4 illustrates how to perform this step.

Figure 7.4: Selecting a report group for incremental profiles analysis.

In addition to the report group that you have already specified, you can also include other report groups for quality check purposes. Like with the previous step, for each selected group, three report groups will be available, base, new, and increment. Figure 7.5 shows how to select additional report groups.

Figure 7.5: Selecting additional report groups for incremental profiles analysis.

After collecting all the necessary data for the analysis of the incremental profiles, you can decide whether you want to view the incremental profiles analysis once the wizard closes. See Figure 7.6.
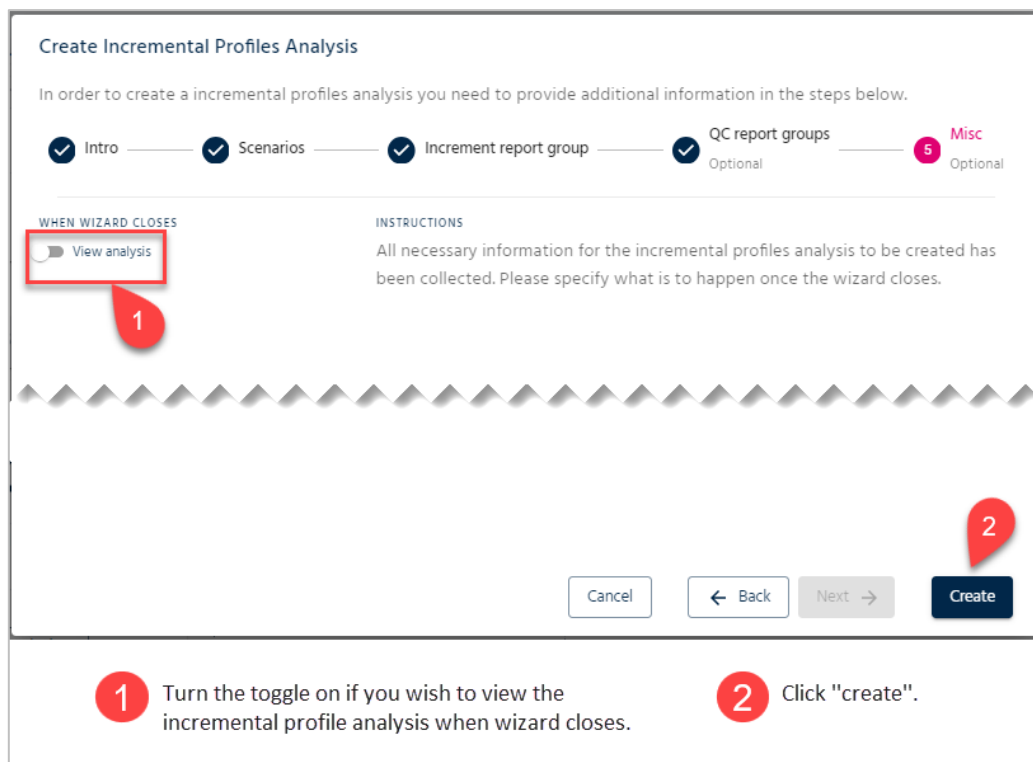
Figure 7.6: Creating a new incremental profiles analysis.

After running the incremental analysis, you will be able to see the results in the Result Viewer. You can view the base case scenario, the new case scenario, and most importantly the resulting increment.

Figure 7.7 shows accumulated volume for the increment for oil, gas, water and liquid respectively.
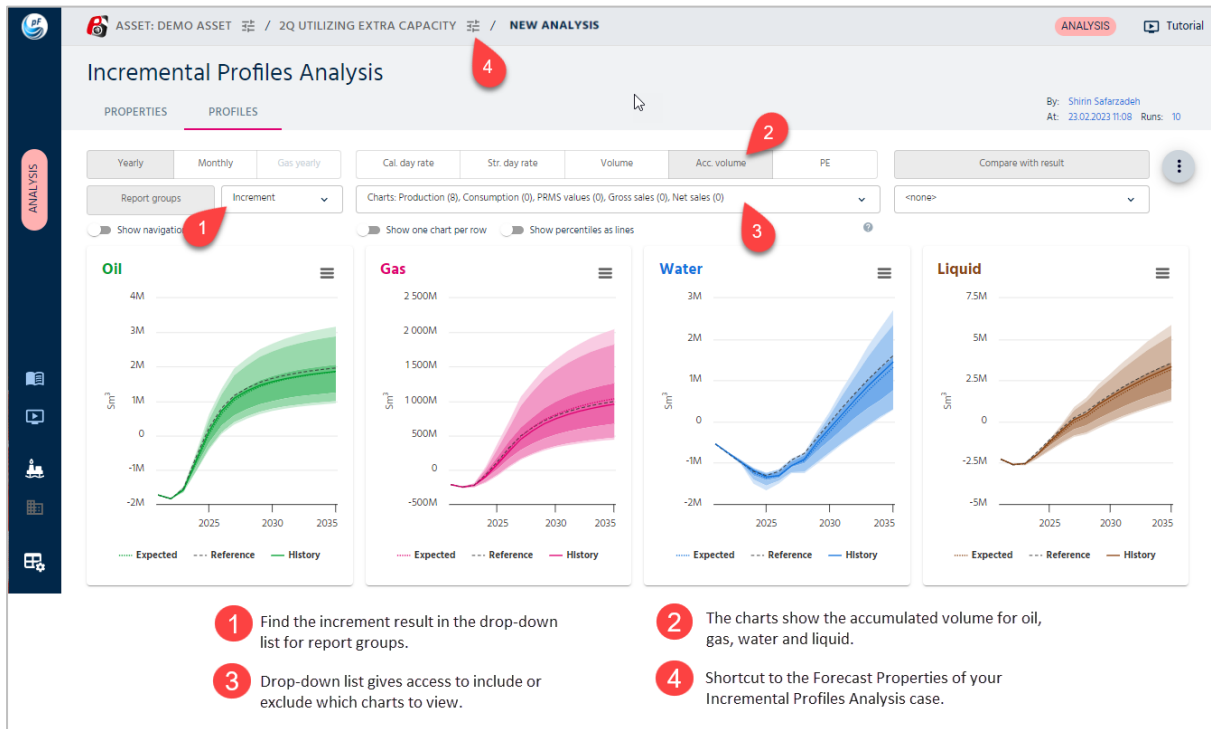
Figure 7.7 Viewing the increment's accumulated volume.

8. Corner Cases

This chapter presents how to deal with corner cases and how to combine functionalities in pForecast to model specific activities.

8.1. Slot Recovery

In case of modeling a sidetrack (Figure 8.1), pForecast can automatically shut down the existing well when the drilling of the new wellbore starts. Note that the new wellbore is secondary, and since the drilling times have uncertainty, we don't know exactly when to shut down the old wellbore. In such cases, you would normally want to keep producing as long as possible. The solution is to connect these events by selecting the drilling start of the second wellbore as the cut-off condition.

Figure 8.1: Slot recovery or sidetrack.

To model a slot recovery in pForecast, first, you need to create a new wellbore. Figure 3.20 illustrates how to do so. Next, you need to select the "Production Potential" tab and select the profile used for the existing well you want to shut down. Then, under the Cut-off section, use the ''At drilling start of well'' option. There you can choose a new wellbore; the one which is to replace the old wellbore. See Figure 8.2.

Figure 8.2: Modelling a slot recovery.

8.2. Capacity Upgrade

In connection with a capacity upgrade, it is necessary to define a shut-in period during the process of upgrading.

To model this, the old capacity should be valid before the shut-in period, while the new capacity should be valid from the end of the shut-in period. But since the duration of the shut-in period has uncertainty, the last day of the shut-in period will not be exactly known. To handle this issue, under the "Capacity" section we assign the "valid from-day" of the constraint to the start date of the shut-in period rather than to the end, even if the end would be the most correct time of change (see Figure 3.18).

Since it doesn't matter which capacity we have during the shut-in; this approach will not compromise the validity time of the constraints before and after the shut-in.

To make this work, it is important to keep a fixed time for the start of the shut-in and use the same time for the new constraint. Figure 8.3 gives more details about how to model capacity upgrades in pForecast software.

Figure 8.3: Capacity upgrade with facility shut-in during upgrade.

9. Glossary

The glossary summarizes the most used terminology in the pForecast software.

Asset: An asset in this context is a group of fields that share data.

Business Unit: Business unit and asset are used interchangeably.

Capacity limitation priorities: If capacity constraints limit your production, you have several ways of prioritizing the order in which wells will be choked. If you specify the default GOR or water cut, these values will be used instead of the calculated water cut or GOR. The priority groups are considered before GOR and water cut in a two-step prioritization approach. A low priority means being among the first wells to be choked and vice versa.

Cluster: Clusters are used to group wells and to constrain the production and injection to the capacity of that group. Clusters can be routed to other clusters; there can be up to four levels of clusters in a facility.

Corporate Analysis: Corporate roll-up analyses are cumulations of values for a selected set of assets. The users can create as many roll-ups as they wish for corporate-wide analyses and reports.

Cut-off: There are four criteria you may specify to cut the well production permanently: at a given rate, at a given volume, on a given date, and at drilling start of another well.

Field: An area in which there are one or several reservoirs. You import data for one field at a time. The field information is used for reporting purposes.

Gas Formation Volume Factor (B$_g$): This is the ratio of the gas volume at the reservoir condition to its corresponding volume at the standard condition (P = 101.325 kPa and T = 288.15 K). In pForecast, it is used for Voidage Management.

GOR: Gas-Oil ratio (Rs): This is the ratio of the volume of gas that comes out of the solution to the volume of oil at standard conditions (P = 101.325 kPa and T = 288.15 K)

Oil Formation Volume Factor (B$_o$): This is the ratio of the oil volume at the reservoir condition to its corresponding volume at the standard condition (P = 101.325 kPa and T = 288.15 K). In pForecast, it is used for Voidage Management.

PE: Production Efficiency.

PVT: This is the abbreviation for pressure, volume, and temperature.

Regularity: In pForecast, we use regularity for constraints on injection, which can be between zero and one. If you want to specify regularity for production, you should specify PE (Production Efficiency).

RNB: It is mandatory for operating companies in Norway to submit their forecasts to the Revised National Budget (RNB). In pForecast, a structure of RNB profiles and projects can be defined under the RNB tab.

Rs: Rs and GOR are used interchangeably.

TAR: A turnaround or a "TAR" is a very expensive process in which an industrial plant or refinery goes through a scheduled shutdown in order to perform maintenance on the facility. During this time, production must come to a complete stop.

Voidage factor: In pForecast, the voidage factor is a user input that means the same as Voidage Replacement Ratio (VRR). It is defined as follows:

$$VRR = \frac{Injected\ reservoir\ volume}{Produced\ reservoir\ volume}$$

Voidage group: Used to group producing wells and injector wells on the same reservoir in connection with Voidage Management. A specific voidage factor (VRR) will be assigned to each voidage group.

Volume-based profile: By turning on this toggle, the total volume of the production potential will be preserved even if there are limitations on the production. Such limitations could be production efficiency, shut-down periods, and potential reductions due to capacity limitations.

WTC: Water Cut. The ratio of water produced compared to the volume of total liquids produced.

About us

We help the petroleum industry generate the best possible forecasts for the most complex reservoirs. pForecast is a Software as a Service (SaaS) solution to digitalize, simplify, and standardize how production forecasts are generated and utilized. pForecast performs a full lifetime simulation of the production and injection forecast, including historical data, in keeping with the industry's ever-increasing need for agility. Powersim Software is the software company and developers behind the pForecast solution.

Powersim Software AS

Litleåsvegen 79, Bergen, Norway

pForecast@powersim.no

Phone: +47 55 60 65 00

 www.pforecast.com

www.powersim.com