

SPnode



Contents

1	Introducing SPnode	3
2	SPnode Core concepts	3
2.1	Planning items	3
2.2	Sequences	4
2.2.1	Activation state	4
2.2.2	Sequence current status	5
2.3	Time and time ranges	5
3	The SPnode Scheduler	5
3.1	Presentation of the SPnode Scheduler panel user interface	5
3.1.1	The time ranges banner in the SPnode Scheduler	6
3.1.2	Time bar	8
3.1.3	Sequences section	8
3.1.4	Visual characteristics of items	10
4	The SPnode Supervisor	14
4.1	Presentation of SPnode Supervisor panel user interface	14
4.2	Monitoring information	14
4.2.1	Default Sequences display	14
4.2.2	Sequence activation state display	15
4.2.3	Past Sequence error notifications	15
4.2.4	Sequence details	16
4.3	The SPnode Visibilities	17
4.3.1	Presentation of the SPnode Visibilities panel user interface	18

4.3.2 Types of visibility 18

1 Introducing SPnode

What is the system used for and in what environment is it displayed?

SPnode is a time scheduling and task execution system. The system consists of two main components: the **SPnode Scheduler** and the **SPnode Supervisor**, which work together to provide a complete, real-time overview of your schedule and its execution.

- The **SPnode Scheduler** is the central element of this system. It is specially designed to visualise your planning and its execution on timelines. As an SPnode user, you can configure planning scenarios, organise resources, and automate various actions according to a precise schedule. The SPnode Scheduler is presented in The SPnode Scheduler panel section.
- The **SPnode Supervisor** plays a crucial role in monitoring the execution of your scheduling scenarios in real time. It provides an overview of activities in progress, their status and their results. By using the SPnode Supervisor, you can quickly detect problems and react in the event of an incident. The panel elements and features are described in The SPnode Supervisor panel section.

These two components of the SPnode system are integrated into Grafana panels. **Grafana** is an open source data visualisation platform that offers advanced features for creating interactive dashboards. This means that you access SPnode via a Grafana dashboard in a web browser. This integration into Grafana allows you to take advantage of the user-friendly interface and advanced data visualisation features to monitor, manage and adjust your planning operations.

In short, SPnode is a complete time scheduling management system that gives you precise control over how your tasks are organised and executed. With SPnode, you have a powerful tool for managing time, automating recurring operations, optimising the use of resources and ensuring that your activities are carried out efficiently.

2 SPnode Core concepts

Familiarise yourself with the concepts.

2.1 Planning items

Items in SPnode are scheduling elements. There are two main types: **EVENT items** and **ACTIVITY items**.

- An **EVENT item** refers to a time event or a specific period in your schedule. It does not lead to any particular action, acting instead as a passive indicator of time.
- An **ACTIVITY item** is an active element, representing a scheduled execution at a precise moment and for a specific duration. It is associated with a concrete action taking place at a given moment.

An **ACTIVITY item** can have different execution status. When an **ACTIVITY item** is scheduled and not yet executed, its status is **SCHEDULED**. Other possible status for an **ACTIVITY item** are (in order of severity, from lowest to highest): **NONE**, **WAITING**, **SKIPPED**, **RUNNING**, **KILLED**, **OK**, **WARN**, **ERROR**, **SYSTEM_ERROR**, **NO_AGENT** and **LOST**. For an **EVENT item**, only the **SCHEDULED** status is assigned to it when it is scheduled in a Sequence.

Each item has specific characteristics that make it unique, such as an identifier, a name, a type, a start date, an end date, etc. In addition, they can be linked to each other by conditions, which means that an item can depend on the execution or progress of other items, so introducing logical relationships between them.

For more information on items please refer to the SPnode API documentation, or see Visual characteristics of items.

Dependencies

There may be dependency relationships between EVENT and ACTIVITY items, influencing their execution or triggering other specific processing.

Conditions at the origin of these dependencies

These dependencies are based on conditions, which are Boolean expressions written in Python 2.7.3 and specified in the `beforeCondition` and `afterCondition` properties of items.

In these conditions, items can be referenced by their name. For example, the condition `myActivity.getStatus() == OK` checks the status of the item with the name `myActivity`. This referencing between items creates dependencies, and distinguishes two types of items: **conditioning items** and **conditioned items**.

Conditioning and conditioned items

An item is considered **conditioning** for another item when it is referenced in the other item's `beforeCondition` and `afterCondition` properties. A conditional item can influence the start of execution of the other item (if it is an ACTIVITY item) and the deactivation of the sequence on which the other item is scheduled (whatever the type of item). An item can also be a condition for several other items, known as conditioned items.

An item is **conditioned** by another item when it references another item in its `beforeCondition` and `afterCondition` properties. The start of execution of a conditioned item (if it is an ACTIVITY item) and the deactivation of the sequence on which the conditioned item is scheduled can be influenced by the conditioning item. An item can be conditioned by several other conditioning items.

It is also possible for an item to be both conditioning and conditioned.

Specific visual indicators in the interface indicate the existence of dependencies between items.

2.2 Sequences

Understand Sequences and their hierarchy.

A sequence is an active or inactive line of execution, sequential and autonomous, on which **items** are scheduled. SPnode interfaces can contain and display several Sequences at once. Although autonomous, these execution lines can also be dependent on each other, through items.

Sequences can be independent or hierarchical by belonging to groups of Sequences.

2.2.1 Activation state

A sequence can be in two different states: **Activated** or **Deactivated**. When a sequence is **activated**, this means that the activities scheduled on this sequence can be executed, as soon as the SPnode time bar reaches the date on which these activities will actually be executed.

A new sequence is activated by default.

When a sequence is **deactivated**, this means that the next activities scheduled on this sequence will not be executed. It is important to note that running **ACTIVITY** on a sequence will not be interrupted at the moment the sequence is deactivated, but subsequent executions on this sequence will be ignored.

2.2.2 Sequence current status

The current status of the sequence is defined by the status the last run activity. The last run activity is defined according to the following order of priority :

1. the activity currently running
2. the last activity executed (successfully or not)

The possible current status of the Sequence are (in order of severity, from lowest to highest) : **NONE**, **SCHEDULED**, **WAITING**, **SKIPPED**, **RUNNING**, **KILLED**, **OK**, **WARN**, **ERROR**, **SYSTEM_ERROR**, **NO_AGENT** and **LOST**. When no execution has yet occurred on the Sequence, or when the Sequence contains only **EVENT** items, the Sequence's current status is considered **SCHEDULED**.

2.3 Time and time ranges

Understand what SPnode time is and familiarise yourself with time ranges.

SPnode time is like a clock that you can set yourself, marking the precise moment when executions occur. Imagine a vertical bar that moves from left to right as the seconds tick by, giving you a visual representation of the past, present and future.

In the SPnode Scheduler, there are two important concepts related to time ranges:

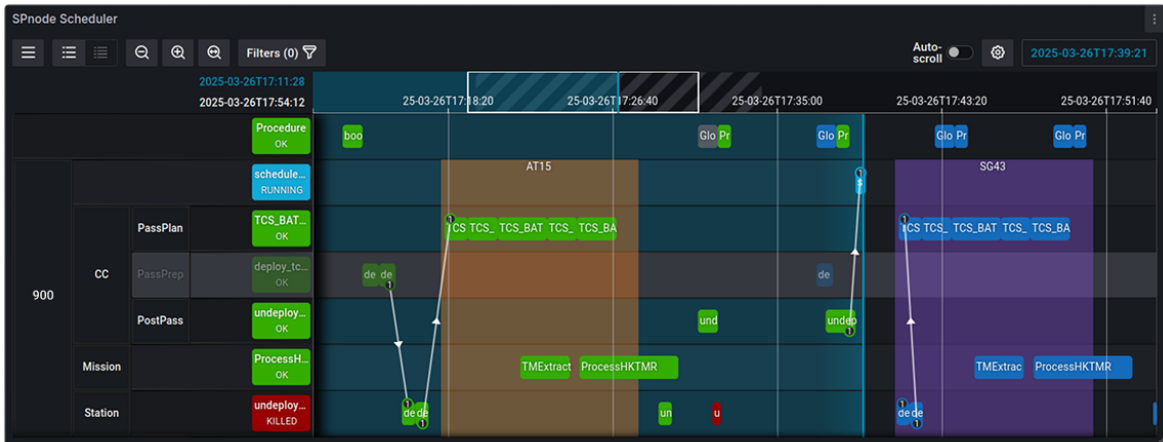
- First, the time range configured in the Grafana dashboard acts as a frame. It limits what you can see in the SPnode Scheduler.
- Secondly, there's the time range selection, which acts as a window within this frame, showing the period of time you are looking at. It allows you to move around in time to observe past, present or future events.

For a detailed explanation of SPnode time and time ranges, please refer to the sections **Time bar** and **The time ranges banner** in the SPnode Scheduler.

3 The SPnode Scheduler

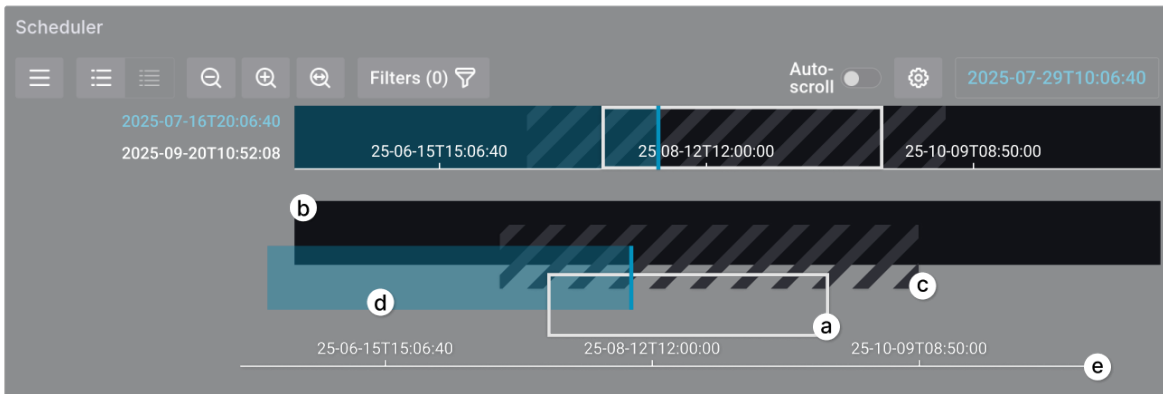
3.1 Presentation of the SPnode Scheduler panel user interface

Explore the interface for a better understanding.



3.1.1 The time ranges banner in the SPnode Scheduler

The SPnode Scheduler takes into account two essential time ranges for its operation.



3.1.1.1 The Grafana dashboard time range

As its name suggests, this range corresponds to the time configured in the Grafana dashboard. On the SPnode Scheduler interface, this range is represented by a rectangle in the background of the ranges banner (b).

It is important to note that this range limits the viewing of Sequence content, which means that you cannot view items outside this configured range from the Grafana dashboard.

Hatching (c) may appear in this rectangle to indicate where the schedules are in time. It extends from the first item scheduled in SPnode to the last, but does not inform you of time slots where no items are scheduled between these two points.

This rectangle contains two time indicators: 1. A vertical bar (d) representing SPnode time moves from left to right, showing visually where executions are taking place in relation to the time range configured from the Grafana dashboard. A coloured veil accompanies it, and is visible to its left. This indicates the time passed (see explanations in the section Time bar).


- It also includes a horizontal scale graduated by time values (e). These graduations represent points in time currently visible in the SPnode Scheduler, and belonging to the time range selection. Lines run vertically through the Sequences at each graduation point.

3.1.1.2 The time range selection

The time range selection in the SPnode Scheduler corresponds to the visible period of time.

In the interface, this range is symbolised by a rectangle with a border, without fill colour, and located in the foreground of the ranges banner (a).

All the scheduled elements you see in Sequences are located within this range. In other words, the start of this range always corresponds to the start of what is displayed in a Sequence, and the end of this range always corresponds to the end of what is displayed in a Sequence.



By default, the bounds of this range are the same as those of the Grafana dashboard time range.

The rectangle representing the range lets you move around in time. You can configure this range from the SPnode Scheduler in different ways (see Configure SPnode time and displays according to time).

3.1.1.3 Text labels representing the bounds of the time range selection

To locate the positions of the two bounds of the selection range, there are two text labels to the left of the Scheduler time ranges.



- The colour of these labels depends on the time position of the bounds relative to the SPnode time :
- If a bound of the range is to the right of the SPnode time position, i.e. in the **future**, then its inscription is **white** (in Grafana's dark theme mode) or black (in light theme mode).
 - If a bound of the range is to the left of the SPnode time position, i.e. in the **past**, then its inscription takes the colour of time (**blue** or **yellow-orange**).
 - If a bound of the range is **not included** in the Grafana dashboard time range, its inscription is **grey**.

It is important to note that when the bounds of the time range selection correspond to the bounds of the Grafana dashboard time range, the text inscriptions reflect the time values of the bounds of the Grafana dashboard time range.

3.1.2 Time bar

The time bar, associated with the SPnode time value (configured or not), visually illustrates the precise time (T) at which the executions are taking place.

As a reminder, the textual representation of the SPnode time value is displayed at the top right of the Scheduler interface (see Top toolbar > *Knowing the SPnode time value and refocusing the display in relation to it*).

When no time shift configuration is applied (see Top toolbar > Modify the SPnode time), the time bar is blue. On the other hand, when a time offset is applied, the time bar takes on a yellow-orange tint.




Figure 1: Time bars over sequences

Like the coloured veil to its left, this time bar hovers over the Sequences. The role of this veil is to indicate directly whether you are in a period when executions have already taken place or whether they are still to come.

The time bar and its associated veil are also visible in the time ranges banner.



Figure 2: Time bars in time ranges banner

3.1.3 Sequences section

Remember that Sequences are autonomous or dependent execution lines, which can be active or inactive, on which items are scheduled. In the SPnode Scheduler interface, a sequence can be displayed independently of the others, or organised into hierarchical groups.

3.1.3.1 Hierarchical and alphabetical ordering

A Sequence has a name, which begins with the “/” character and is composed of a series of words separated by “/” characters. The name represents the unique identifier of the sequence, and it is used to classify the Sequences into groups. Each word corresponds to a group.

Two Sequences with the same root in their name belong to the same group. For example, two sequences with the names “/My/SeqA” et “/My/SeqB” belong to the same “My” group because they have the same root “/My”.

Sequences and groups of Sequences are ordered vertically in **ascending alphabetical order**.





Figure 3: Hierarchical and alphabetical ordering example

Let’s take the example of two sequences with the names “/My/SeqA” et “/My/SeqB”. These two Sequences will belong to the same “My” group, which in turn will contain the “SeqA” and “SeqB” groups. As explained above, the “SeqA” group will be placed vertically above the “SeqB” group by alphabetical comparison.

3.1.3.2 Activation state



When a Sequence is **deactivated**, the Sequence appears greyed out.

When a Sequence is **activated**, the Sequence retains its normal appearance, without being greyed out.

A group (e.g. “My”) is considered activated when at least one of its children is activated (e.g. “SeqA”). When all its children are deactivated, the group appears as deactivated.

3.1.3.3 Sequence current activity

On the right-hand side of the hierarchical grouping of Sequences, you can locate visual indicators that provide information about the last execution of each Sequence.

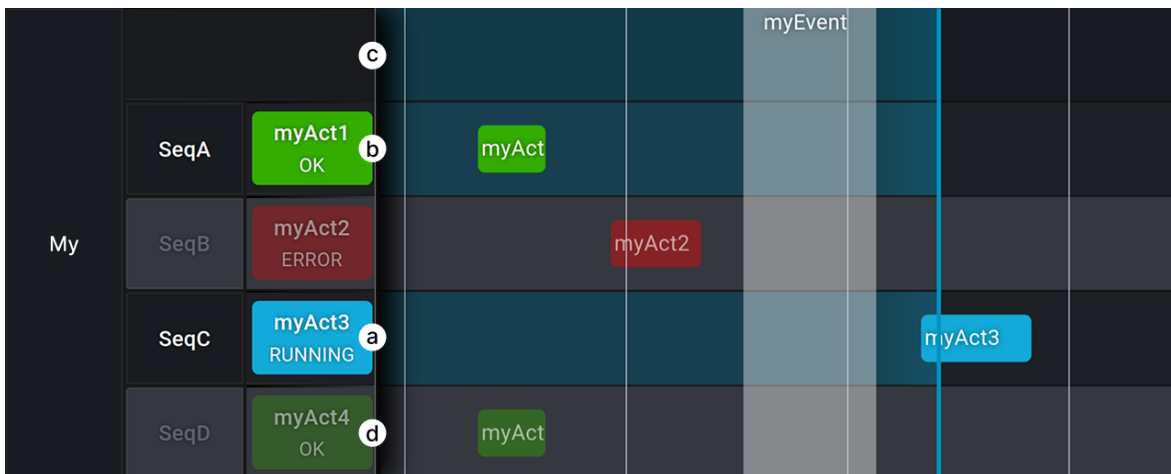


Figure 4: Information from the last ACTIVITY item executed on each sequence

This visual indicator is displayed on a Sequence as soon as at least one execution of an ACTIVITY item is in progress or has already occurred on this Sequence (a, b & d). Otherwise, or when the Sequence only contains EVENT items, it remains invisible (c).

The colour of this indicator represents the last known execution status of an ACTIVITY item in the Sequence. The meaning of these colours is explained in the next section (see Visual characteristics of items > The visual representation of an ACTIVITY item). It is important to note that the opacity of the indicator colour varies according to the activation state of the Sequence. For example, the colour of a status displayed on a deactivated Sequence will be less intense (d), whereas the status displayed on an activated Sequence will be fully visible (a & b).

Within this indicator, you will find the name of the ACTIVITY item being executed, as well as the last status transcribed verbatim. All the characteristics of an item can be consulted via this indicator (see Manage planning and execution > Display item information) for a detailed analysis or simple consultation.

By double-clicking on this indicator, you can refocus the display on the item.



Figure 5: Double-click on last run item indicator to refocus on the item

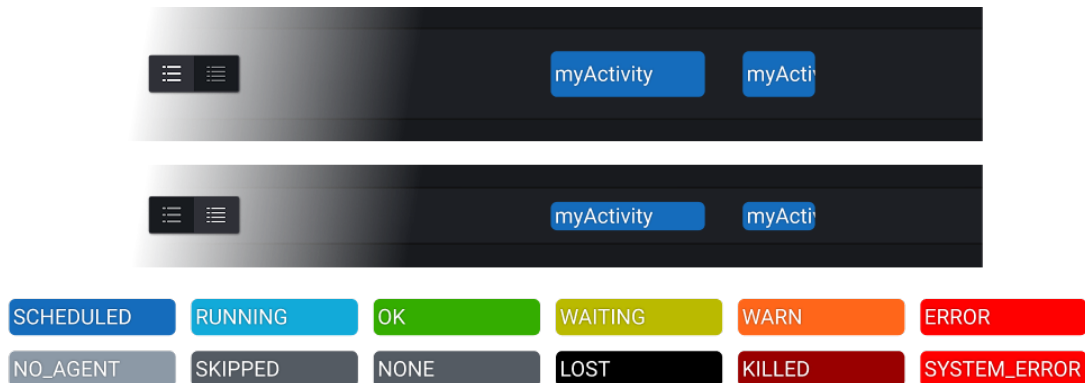
To make it easier to locate the item (when scheduling is busy, for example), a white border appears on the item temporarily for 4 seconds.

Please note that this action causes a time shift by modifying the time ranges.

3.1.4 Visual characteristics of items

3.1.4.1 The visual representation of an ACTIVITY item

An ACTIVITY item is visually represented as a rectangle with slightly rounded corners.



The height of the rectangle is 50% of the height of the Sequence on which it is located, and its width varies according to the duration of the item (`expectedDuration` or `realEndDate - realStartDate`) and the zoom level in the SPnode Scheduler.

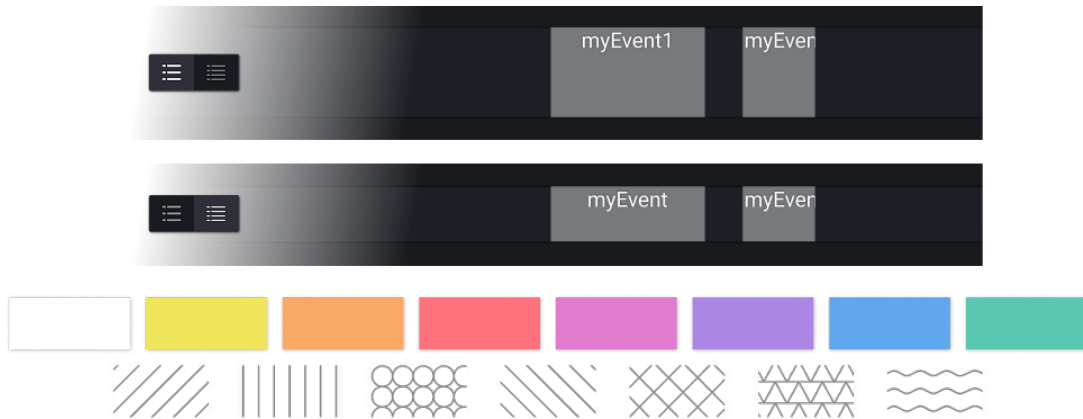
Its color reflects the status of the item, and changes according to the activation or deactivation of the corresponding Sequence.

The text, representing the item's name (`name` property), is aligned with the rectangle's left border. If the width of the rectangle is insufficient to display the text, it is truncated on the right, leaving only the beginning of the name visible. Vertically, the text is centered within the rectangle.

Unlike an `EVENT` item, an `ACTIVITY` item is in the foreground of the Sequence. An `ACTIVITY` item can overlap an `EVENT` item, but two `ACTIVITY` items in the same Sequence cannot be superimposed.

3.1.4.2 The visual representation of an `EVENT` item

An `EVENT` item is visually represented as a semi-transparent rectangle.




The height of the rectangle is 100% of the height of the Sequence on which it is located, or the sum of the heights of several Sequences when it is visually crossing (see `Visible on one Sequence or traversing several other Sequences`). Its width varies according to the duration of the item (`expectedDuration` or `realEndDate - realStartDate`) and the zoom level in the SPnode Scheduler.

The default item color is white when the system is in “dark” mode or black when the system is in “light” mode, but can be customized by choosing a color and/or pattern from those available (see `Custom EVENT item appearances`).

The text, displayed at the top of the item, represents the item's name (`name` property). On the horizontal axis, the text is centered relative to the item. If the width of the item becomes less than that of the text, the text aligns with the left border of the item and is trimmed to the right. This reveals the beginning of the item name.

Unlike an `ACTIVITY` item, an `EVENT` item is placed in the background of a sequence. `EVENT` items can overlap. The semi-transparent appearance of these visual elements allows you to glimpse these possible overlaps.

3.1.4.2.1 Visible on one Sequence or traversing several other Sequences



An EVENT item is attached to a specific Sequence via its `sequence` property, making it visible on that sequence. For example, the item “myEvent” is scheduled on the Sequence “/900/CC”, and therefore appears on this sequence.


However, this item can also be visually traversed on other Sequences if their name begins with the name of the Sequence to which the item belongs. Let’s take an example: if an EVENT item is attached to Sequence “/900”, it will be visible on its belonging Sequence “/900” as well as on any other Sequence whose name begins with “/900”. In this way, the “myTraversingEvent2” item, programmed on the “/900” Sequence, is visible and traversing on the “/900”, “/900/CC”, “/900/Mission” and “/900/Station” Sequences.

Particular case: EVENT items belonging to the “/” Sequence will be visually traversed on all other Sequences, as all Sequences have, by default, a name beginning with the “/” character. This is the case for the EVENT item “myTraversingEvent1”, for example.

If a Sequence is masked by a filter (see Filter settings in the SPnode Scheduler and Supervisor panels) and contains an EVENT item that visually traverses other Sequences that remain visible after filtering, this EVENT item will also be masked and will therefore not be visible on other Sequences.

3.1.4.3 Round markers on conditioning/conditioned items

In the SPnode Scheduler, you will see special round markers on certain items, marking dependencies. The markers allow you to quickly identify the presence of dependencies attached to your existing items within your planning.



- A round marker showing that an item is **conditioning** is **positioned at the bottom right** of it. It indicates an outgoing dependency.


- A round marker showing that an item is **conditioned** is **positioned at the top left** of it. It indicates an incoming dependency.

The placement of these markers follows the convention of reading text, starting at the top left and ending at the bottom right. Here, an incoming dependency in a conditioned item is signified by a marker positioned at the top left, while an outgoing dependency in a conditioning item is signified by a marker positioned at the bottom right.

The size of the markers varies according to the type of item (ACTIVITY or EVENT), the width of the item, and the height of the sequence on which the item is scheduled. The wider the ACTIVITY item, the larger the marker. On the other hand, if the item is of type EVENT, or if the width of the ACTIVITY item is small, or if the height of the sequence is reduced (see Top toolbar > 2. *Sequence display height*), then the marker is smaller.

When the marker is large enough, a number appears inside. This indicates the number of incoming dependencies for a conditioned item, or outgoing dependencies for a conditioning item. The smallest marker you can see is represented by a dot.


You can right-click on a marker to access information about conditioning or conditioned items.



To find out more about interactions and access to information, please refer to section Manage planning and execution > Display item information.

3.1.4.4 Lines linking conditioning/conditioned items

Lines can also be drawn between these markers to help visualize the relationships between these items.



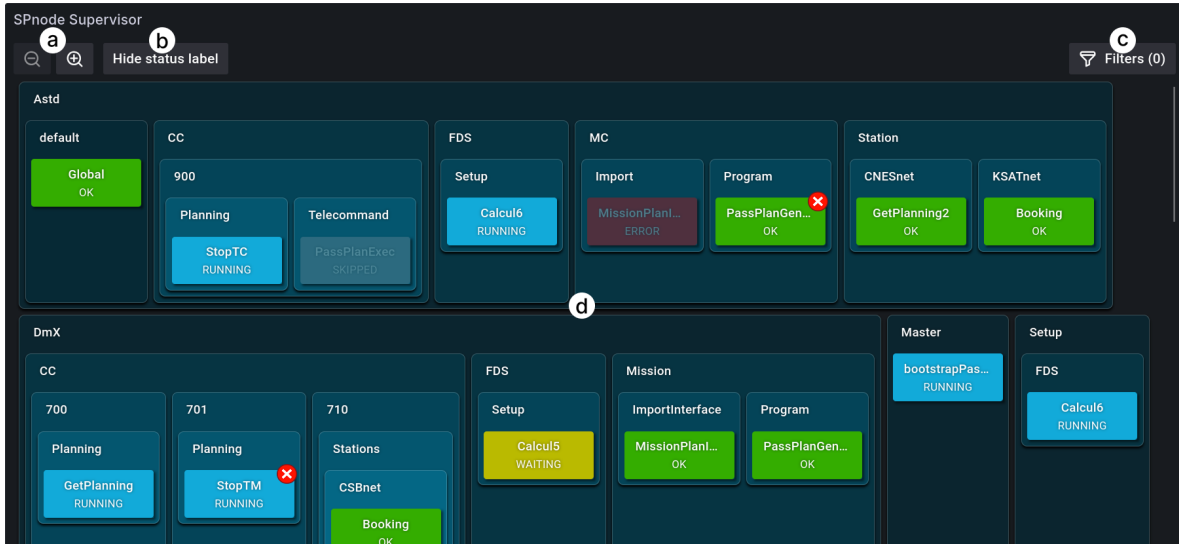
A line is drawn from the marker positioned on the conditioning item to the marker positioned on the conditioned item.

An arrow appears in the middle of the drawn line to indicate the direction of the dependency.

You can choose to show or hide them (see Manage planning and execution > Show or hide dependencies between items).

4 The SPnode Supervisor

4.1 Presentation of SPnode Supervisor panel user interface



a. *Display Size.*

The display size of the user interface elements can be modified through the display size buttons from small to large.

b. *Status Label Display.*

This option provides the ability to hide/show the status label at the panel level. Hiding status label will allow you to display more Sequences in the same space but suppose that you are familiar with the status color code.

c. *Temporary filters*

This allows you to define further filter settings besides other permanent filter settings.

d. *Sequences Overview*

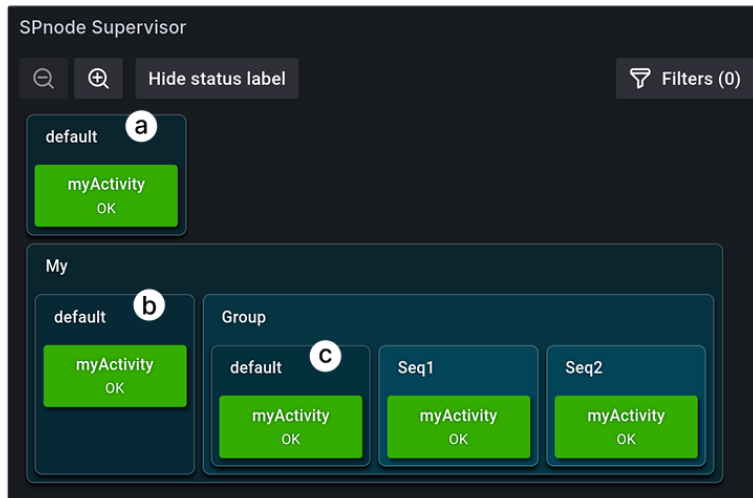
This is the main display area of the panel is dedicated to the Sequences status display. All monitoring takes place in this section, this overview reproduces the actual Sequence tree organization and shows the current status for each Sequence.

4.2 Monitoring information

4.2.1 Default Sequences display

Default Sequences offer a way to group items concerning a whole subset of Sequences. With default Sequences, any level of Sequence can nest activities. Sequence containers marked as default refers to the Sequence root of the Sequence they are contained in.

Considering this organization of sequence :

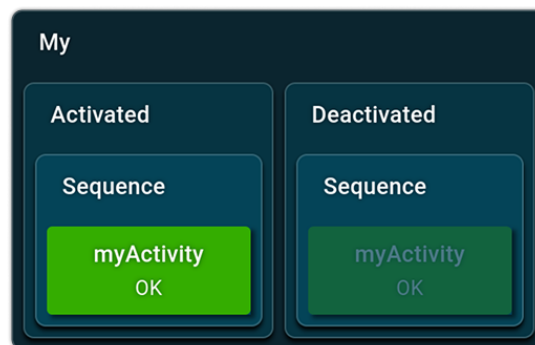


The *default* Sequence **a** refers to the root of the Sequence tree and its corresponding Sequence name would be “/”.

The *default* Sequence **b** contained within the “/My/...” Sequence refers to root of **My** Sequence and its corresponding Sequence path would be “/My”. It is the same for Sequence **c**, which corresponds to Sequence “/My/Group”.

4.2.2 Sequence activation state display


Sequences activation state can be quickly identified through a disabled visual effect. Activities tag for deactivated Sequences are darkened.



4.2.3 Past Sequence error notifications

Non-blocking Sequence errors that occurred in the past are potentially not visible via current status.

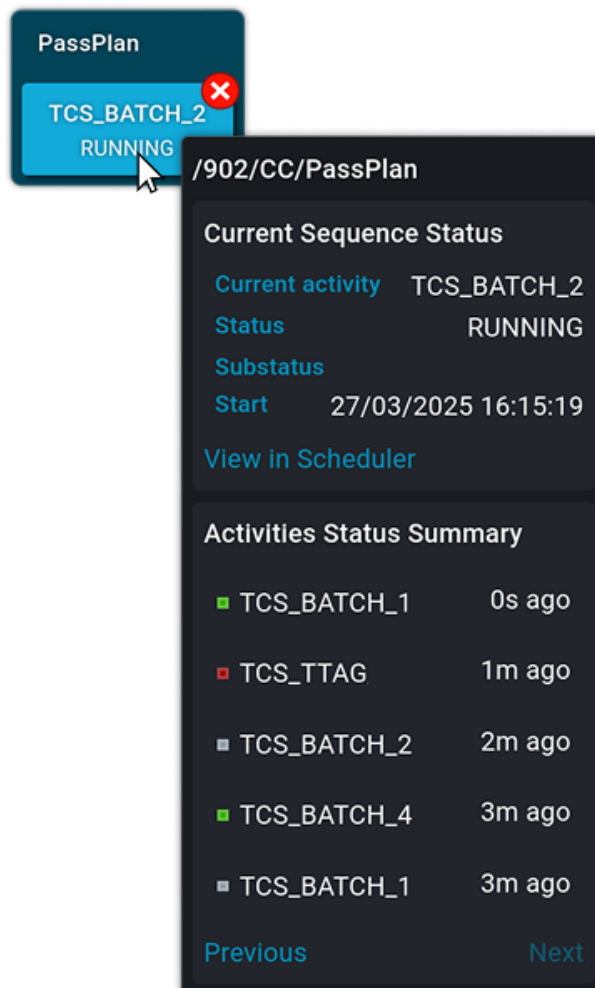
To identify Sequences that may have encountered easily, all errors that occurred on a Sequence **within the timerange defined in Grafana** will be notified by a error badge on the right upper corner.



It is possible to retrieve which activities cause the error by opening the Sequence see Sequence Details section.

4.2.4 Sequence details

By clicking on an activity, a popup appears it is possible to view more detailed information about this activity.



In the **Current Sequence Status** section, you can find information about the current activity

In the **Activity Status Summary**, a summary of the activities executed within the defined Grafana timerange is splayed.

Each activity has a *View in Scheduler* associated link which provide a direct link to the related activity in a linked SPnode Scheduler panel (see Linking to a SPnode Scheduler panel for more detail about linking Supervisor panel to a Scheduler panel).

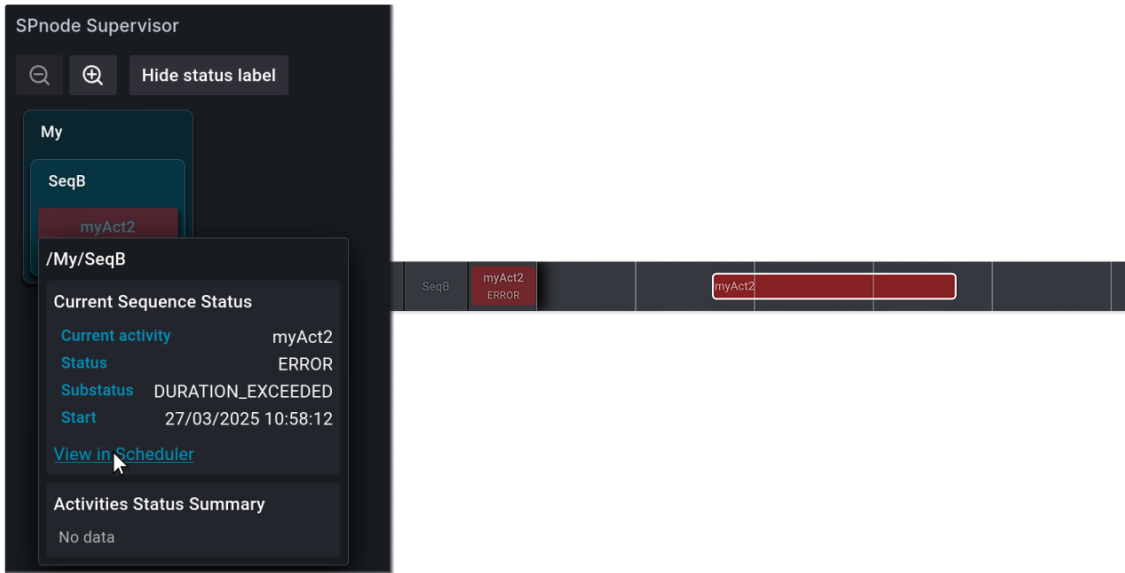
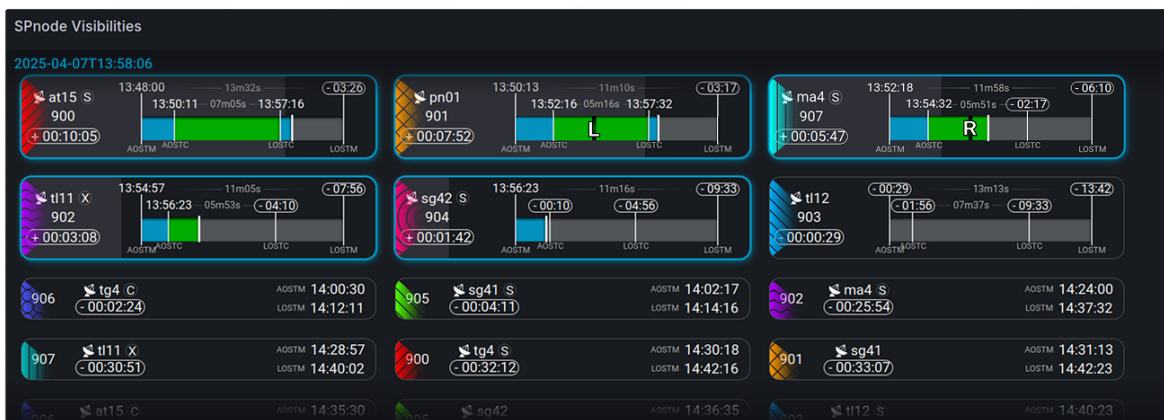


Figure 6: View in Scheduler function

Keep in mind that only activities that occurred within the specified Grafana timerange will be displayed in the activity history.

4.3 The SPnode Visibilities



4.3.1 Presentation of the SPnode Visibilities panel user interface

The SPnode Visibilities panel displays satellite visibilities in real time. Elements are organized chronologically, according to AOSTM dates (in ISO8601 format).

The interface reads like a page in a book, from top to bottom and left to right.

The elements displayed represent different types of visibility, which can be :

- Past
- In progress
- To come
- To come and detailed

Each type of visibility has a distinct visual presentation.


4.3.2 Types of visibility

Past visibility



Past visibility looks greyish. Only satellite name, station antenna name and AOSTM and LOSTM times are displayed.

Visibility in progress



Visibility in progress is visually highlighted by a thick, luminous blue border.

It displays :

- On the left: the name of the antenna station, the band (X, S, etc.), the name of the satellite and a positive counter for the time elapsed since the start of visibility (AOSTM).
- Right: a graph showing the AOS (Acquisition Of Signal) and LOS (Loss Of Signal) events, the times between AOSTM-LOSTM and AOSTC-LOSTC, the polarity change event (L, R, ...) if any, and the progress of visibility. On the graph, events are positioned proportionally according to their datetime.

Events that have not yet occurred are displayed with a negative counter (e.g. LOSTC and LOSTM in the figure). When an event occurs, its time is indicated (for example, AOSTM and AOSTC).


Visibility to come



An upcoming visibility is displayed as “active” (not grayed out, unlike a past visibility), with the following information:

- Satellite name, station antenna, band and a countdown to the start of visibility.
- AOSTM and LOSTM times.

Visibility to come and detailed



A detailed upcoming visibility is similar to an upcoming visibility, but with the addition : - A graph detailing AOS and LOS events appears a few moments before visibility begins, depending on the configuration chosen by the user. This delay can be set in the panel options (see Panel options > Graph display delay).

In the example above, the graph appears 30 seconds before the start of visibility, due to the “PT30S” configuration.