

1 Workflow Design Rules

When you create or run workflows you should follow some basic rules:

- Rules for algorithms
- Rules to set up correct schedules to avoid timing conflicts
- Rules for backup and restore

1.1 Rules for Algorithms

When you design your synchronizations, you must be sure not to produce conflicts. This section provides some rules to avoid problems in this area.

1.1.1 Single Master System Design

If you follow the rules in this chapter, your system will be very robust. Conflicts cannot occur per design.

- Define only one master for entry creation and deletion.
- Run only workflows with non-overlapping execution times. If this is not possible, be sure that the workflows do not influence each other.
- If workflows create or use the same files (for example one workflow creates an LDIF file, the other reads it), be sure that the execution times do not overlap.

1.1.2 Multi Master System Design

The rules in this chapter help to control a multi master system. They help to minimize conflicts.

- If you need to set up several masters for entry creation and deletion, define unique criteria to distinguish the corresponding entries. Two examples are:
 - Create different trees in your directory that belong to the different masters.
 - Set up attributes or object classes, which allow to distinguish between the entries of the different masters, for example by filtering expressions.

- If you cannot follow the previous guideline, be sure that workflows that create or delete entries do not run at the same time (avoid overlapping execution time – see next chapter).
- If workflows create or use the same files (for example one workflow creates an LDIF file, the other reads it), be sure that the execution times do not overlap.

1.2 Rules for Correct Set Up of Schedules

A lot of issues have to be regarded when scheduling workflows. This chapter explains the rules and their reasons. Set up schedules so that

- they do not have unrealistic values (for example an interval of 0 which means that the workflow would run always – the scheduler ignores schedules with this setting).
- their possible start times cannot overlap
- their execution times cannot overlap

DirXmetahub has some inbuilt features, which handle conflicting situations, but only when the parameters are set correctly.

Start of the same workflow twice

The DirXmetahub server does not allow starting the same workflow twice (for example from the scheduler and from the DirXmetahub Manager by hand). Both workflows would try to open, read and write the same files in the work area, which cannot work properly.

A message is written into the server logging.

Start of workflows after start of server

At the time when the DirXmetahub server is started, all workflows are started where the start time lies in the past and the deviation time is not over.

In this case it can happen that a workflow was aborted previously (during the stop sequence of the DirXmetahub server). MQSeries keeps the related start message and resends it after startup of the DirXmetahub server. The scheduler at the same time tries to start the same workflow. This will fail, because the same workflow was already started (a message is written into the logging of the server). Thus it is assured that the workflow is only started once.

1.2.1 Basic Rules for Parameters in the Central Configuration Object

Rule 1: Polling time

- The **Polling Time** in the central configuration object **must be greater than 0**. Otherwise, the scheduler will assume a default of 5 seconds.

Reason: a timer based on the set polling time activates the scheduler. A polling time of 0 would result in heavy system load. Therefore we do not allow setting this value.

Rule 2: Time Interval

- The **Time Interval** in the central configuration object **must be greater than 0**. Otherwise, the scheduler will assume a default of 1 day (86400 seconds).

Reason: this value determines the interval, the status tracker removes status entries from the database. A Time Interval of 0 would result in heavy system load. Therefore we do not allow setting this value.

Rule 3: Schedule Sync Interval

- The **Schedule Sync Interval** of the central configuration object **must be greater than 0**. Otherwise, a default of 1 hour (3600 seconds) will be assumed.

Reason: normally the DirXmetahub Manager informs the scheduler automatically when schedule information is changed. If this fails, this mechanism reads the schedules regularly from the configuration database. A Schedule Sync Interval of 0 would result in heavy system load (the schedule reread would happen always). Therefore we do not allow setting this value.

1.2.2 Basic Rules for Parameters in a Schedule Object**Rule 4: Time Interval**

The **Time Interval** in a schedule object must be

- **greater than 0**.

Reason: this value determines the interval, when this workflow has to be started again. A Time Interval of 0 does not make sense. Therefore we do not allow setting this value.

- and **greater than the polling time** (see central configuration object). The scheduler will ignore schedules that do not satisfy this condition. Related error messages will be written into the logs (event log or log files).

Reason: A time interval smaller than the polling time could result in missed starts of workflows.

- and **greater than the sum of deviation + timeout + 5 * polling time**. The **time buffer** of $5 * \text{polling time}$ is recommended because multiple Server components (scheduler, workflow engine, agent controller, messaging service) are involved in abortion of a workflow in case of exceeded timeout. Here, timeout is the sum of all job timeouts of related workflow, multiplied with $(1 + \text{latency-factor} / 100)$.

Example: Suppose that a workflow consists of 2 jobs running in a sequence, both with a job-timeout of 10 minutes, and the latency factor is 20. Then timeout (of the workflow) is 24 minutes.

Reason: The schedule might miss start times when this condition is not fulfilled.

Rule 5: Deviation

- The **Deviation** of the schedule **must be greater than the polling time** (see central configuration object).

Reason: When this condition is not fulfilled, the scheduler could miss workflow start times.

1.2.3 Combined Rules

These rules must be considered when several workflows, which could influence each other shall not be started simultaneously or may not have overlapping execution times.

To give full control to the DirXmetahub server (for example to abort agents after the timeout has been reached), you must set the **Abort Execution allowed** flag in the agent object in the expert view.


Rule 6: Non overlapping start times for workflows

Suppose you do not want the scheduler to start a set of workflows simultaneously. Then define related schedule parameters **start time**, **deviation** and **interval** so that the intervals of possible start times do not overlap.

Example1: Non overlapping start times

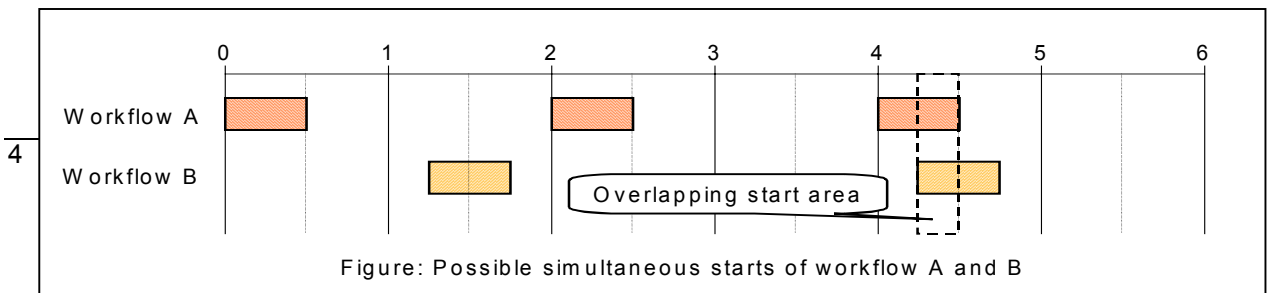
Schedule	Start Time	Interval	Deviation
A	01.01.2001 00:00:00	2 h	30 min
B	01.01.2001 01:15:00	2 h	30 min

Then the scheduler will not start the related workflows simultaneously (see figure below).



Example 2: Overlapping start times

If we change the interval of workflow B to 3 h, then overlapping starts would be possible at 4:15 to 4:30, 10:15 to 10:30 and so on.



We must also take into account that after a restart of the DirXmetahub server all workflows which are in their deviation time will be started simultaneously at the same time (the related messages are waiting in the MQSeries server). Let's assume that the server is restarted at 4:20 then both workflows are started immediately.

Rule 7: Non overlapping execution of workflows

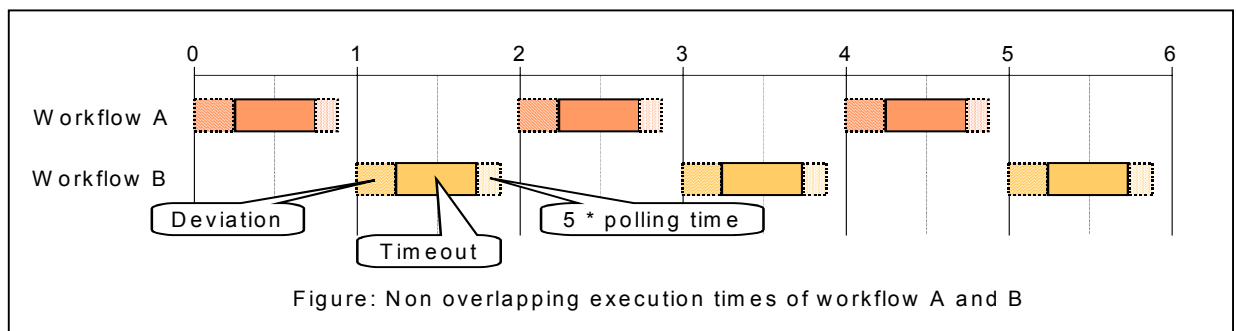
Suppose you want to schedule workflows so that their execution does not overlap (what may be convenient for resource-consumptive workflows). Then you must define the related schedule parameters **start time**, **deviation**, **interval** and **polling time** so that the intervals of possible execution times of a workflow do not overlap.

Example1: Non overlapping execution times

Schedule	Start Time	Interval	Deviation	Timeout
A	01.01.2001 00:00:00	2 h	15 min	30 min
B	01.01.2001 01:00:00	2 h	15 min	30 min

We additionally suppose that the polling time is set to 5 seconds.

As shown in the next figure, the workflow execution cannot overlap.



In normal operation, the upper limit of a life time of a workflow instance in the DirXmetahub Server will be **timeout + polling time**, because there may be a delay of polling time before scheduler and workflow engine really detect a timeout.

The example shows that a workflow can only be active in time intervals between **start time<n>** and **start time<n> + deviation + timeout + 5*polling time**, where $\text{start time}\langle n \rangle = \text{start time} + n * \text{interval}$ (n being an integer).

1.3 Rules for Backup and Restore

If you want to backup and restore your LDAP server or the configuration part of it, we recommend the following procedure:

- Deactivate all schedules to assure that no new workflows are started (uncheck the **Active** flag in each active schedule).
- Check with **Get Server State** on each DirXmetahub server in your environment that no workflows are still running (only status tracker and scheduler threads may be visible in the details view of this feature).
- Stop all DirXmetahub servers.
- Perform a backup or restore of your server with the native methods of your LDAP directory (for DirX you will normally use DirXmanage) or perform an **Export Configuration** or **Import Configuration** command from the DirXmetahub Manager's expert view.
- After a restore, restart all DirXmetahub Managers (because the content of the configuration database could be changed completely and the caches have to be cleared).
- Restart all DirXmetahub servers.
- Activate all schedules.

Normal operation is performed again.

After a restore of a complete database or only the configuration tree in the database, you must be aware of the following problems:

- If you have restored an older version of your configuration tree, of course status entries in the monitor view of the DirXmetahub Manager could be missing. Thus you cannot access the related status entry file information in your status area in the file system any more.
- If your database comes from a different environment (for example from a test environment), some settings in your database could not be correct. You have to adjust it by hand. Candidates are the name of the MQSeries queue manager, settings in the service objects and others, which are different in the two environments. Therefore we do not recommend exchanging whole databases or configuration trees. Use instead the **Export data** and **Import data** feature of the DirXmetahub Manager to exchange parts of the configuration database (whole scenario, workflow, job or connected directory objects).
- Do not try to use the **Import data** method with data with data, which was created with a previous version of DirXmetahub V6 if this was not explicitly allowed (see release notes). Changes in the data structure could not fit any longer with the new version of DirXmetahub.

DirXmetahub V 6.0

Workflow Design Rules
Edition October 2001



Table of Contents

1	Workflow Design Rules	1
1.1	Rules for Algorithms	1
1.1.1	Single Master System Design.....	1
1.1.2	Multi Master System Design	1
1.2	Rules for Correct Set Up of Schedules.....	2
1.2.1	Basic Rules for Parameters in the Central Configuration Object	2
1.2.2	Basic Rules for Parameters in a Schedule Object.....	3
1.2.3	Combined Rules	4
1.3	Rules for Backup and Restore.....	6
	Table of Contents	i

Infoline:

Tel: +49 (89) 636-48878

Fax: +49 (89) 636-47168

E-Mail: directory@icn.siemens.de

Support:

<http://www.siemens.com/directory>

**Comments
Suggestions
Corrections
Courses**

The User Documentation
Department would like to know
your opinion on this manual.
Your feedback helps us to
optimize our documentation to
suit your individual needs.

All product names quoted are trademarks or registered trademarks of the manufacturers concerned.

© Copyright 1997-2001 Siemens AG

Distribution and reproduction not permitted without the consent of Siemens.

How to use the softbook

1 Structure

The softbook contains several chapters, plus possibly glossary, index, title sheet, table of contents, copyright and notes on how to use the softbook. The order in which they are listed here is the order in which they appear, i.e. the notes are at the end.

Changes and additions made after the copy deadline are documented as "current information" and can be found after the keywords in the softbook.

2 Printing the softbook

For best results print to a PostScript printer.

You can print either the contents of the entire softbook (maximal 255 pages at a time) or particular pages. You can find help for printing problems in the readme file located in the program folder of Acrobat reader.

If you want to compile your own printout of the book, it is advisable to place the title sheet and the table of contents, neither of which have any page numbers, in front of the first chapter.

3 Navigating

You can navigate through the softbook e.g. by paging up and down one page at a time or by retracing your steps through the document (previous view).

In the following some navigational structures are described.

a Bookmarks

The bookmarks palette in the navigation pane contains the structure of the softbook in the form of a visual table of contents. The texts and numbering appearing on the bookmarks correspond to the chapter headings. When you open the softbook, initially only the first-level headings are displayed.

The bookmarks are used to jump direct to the individual chapters and sections of the softbook.

b Index

The page numbers quoted after the keywords are in general linked to the corresponding pages. A click on the sensitive area will take you to the page containing the keyword.

4 Full-Text Index

This online documentation set also provides a full-text index generated by Acrobat Catalog®. To use this index you need *Adobe Acrobat Reader plus Search*.

The full-text index includes all online manuals of DirX or DirXmetahub. The corresponding index is attached automatically when an online manual is opened. All word options (Case sensitive, Sounds Like, and Word Stemming) were enabled when the index was built. There were no numbers or stopwords excluded from the index.