



High Volume Replicator



Job Juggler User Guide

Version 3.6
May 2004



Job Juggler User Guide


P S B / High Volume Replication bv
Haaksbergweg 45
1101 BR Amsterdam
The Netherlands

Telephone: (+31) 20 312 75 12
Fax: (+31) 20 312 75 09
<http://www.HighVolumeReplicator.com>

Copyright © 2004 P S B / High Volume Replication bv
Printed in The Netherlands.
All rights reserved.

Table of Contents

Introduction	Intended Readers	7
	Structure of the Documentation	8
	Conventions	9
Job Juggler User Guide	Introduction	11
	1. Job Juggler Concepts	13
	1.1. Job Space	13
	1.2. Attributes	13
	1.3. Attribute Precedence	13
	1.4. Job States	13
	1.5. Job Output Redirection	14
	2. Job Navigator	15
	2.1. Introduction	15
	2.2. Starting the Job Navigator	15
	2.3. Browser Window	17
	2.4. Job Navigator Window	19
	2.5. Detail Windows	25
	2.6. Controlling Jobs	27
	2.7. Disabling and Suspending jobs	28
	2.8. Job Navigator Links	29
	Index	



Introduction

Intended Readers

The information in this Job Juggler User Guide is intended for the following audience:

- ❖ Application developers.
- ❖ Maintenance staff who need to change and or reset jobs run within the Job Juggler.

This user guide is written for a general audience. There is no special knowledge required to understand the content of this user guide.

Structure of the Documentation

This Manual This user guide explains what the Jog Juggler is and how to use its user interface Job Navigator.









Documentation Set This is the structure of the complete documentation set:

Manual	Describes
<i>HVR & Job Juggler Installation & Operations Guide</i>	Software installation steps
<i>HVR User Guide</i>	How to use the High Volume Replicator
<i>HVR Reference Manual</i>	Each component of HVR
<i>Job Juggler User Guide</i>	How to use Job Juggler
<i>Job Juggler Reference Manual</i>	Each component of HVR's scheduler

Conventions


This documentation uses some set patterns in explanations and descriptions. This supports the understanding of the text and improves the recognition of recurring parts. Instructions always follow a distinctive step pattern; every necessary action is described separately in this pattern. Often the result of an action is given.

This documentation uses the following conventions:

Convention	Indicates
bold	Computer 'words'
<i>italics</i>	Manual titles AND Variable 'computer' words
monospace	Examples
"...."	Paragraph titles
Select LTS menu > Enter Data > Entity > Organization	Select the menu LTS menu in the menu bar; Select the Enter Data menu option in the LTS menu; Select the menu option Entity in the Enter Data submenu; Select the menu option Organization in the Entity submenu.
Numbered steps	Actions the user must take; in a table with the action in the left column and the results in the right column; the table is marked by horizontal lines
	References to other parts of the HVR documentation
	Note, warning or recommendation
	Feature only supported on Oracle DBMS
	Feature only supported on Ingres DBMS
	Feature only supported on Microsoft SQL Server
	Feature only supported on Unix
	Feature only supported on Microsoft Windows
	Feature only supported on OpenVMS

File pathnames

All pathnames are shown using Unix conventions, e.g. using a forward slash '/' as file pathname delimiter. For the Microsoft Windows platform this can be understood as a backward slash '\'; generally HVR converts between forward and backwards slashes as appropriate, so the two can be used interchangeably. For OpenVMS a pathname beginning with a variable name can be understood to begin with a logical of the same name, e.g. `$HVR_HOME/lib/msg.dict` corresponds to `HVR_HOME:[lib]msg.dict`.



Job Juggler User Guide

Introduction

In order to understand the functionality of this program, the first section describes basic Job Juggler information. The second section describes the Job Navigator.

For detailed information on the Job Juggler, please refer to the Job Juggler Reference Manual.

1. Job Juggler Concepts

1.1. Job Space

Job Juggler uses a concept of ‘Job Space’, a two-dimensional area containing jobs and queues. A job is a task scheduled in the Job Juggler. A queue is a group of jobs. A queue may contain both jobs and queues. In Job Space, jobs are represented as points (defined by X and Y coordinates) and queues are represented as boxes (defined by 4 coordinates minimum X, maximum X, minimum Y and maximum Y). All jobs and queues are contained within the largest queue, the so-called **SYSTEM** queue.

1.2. Attributes

Attributes are properties which can be defined on an object in Job Space (jobs or queues). There are different types of attributes; different types allow quotas, time-outs and retry to be configured, as well as allowing output redirection and interactive links for display in the Job Navigator. Some attributes are known as control attributes, others as environment attributes. Control attributes include attributes for quota, retry, and time-out, which directly change scheduling. Environment attributes tell the job juggler how to run a job. An example is the **set** attribute, which is used to define an environment variable for a job; they have a less direct effect on scheduling.

Some attributes can be defined for specific periods or all periods. A specific period is typically a part of the day. For example, you can define a time-out attribute that operates differently at night, than during day time.



For more information on Attributes, see [3. “Attribute Types”](#) of the Job Juggler Reference Manual.

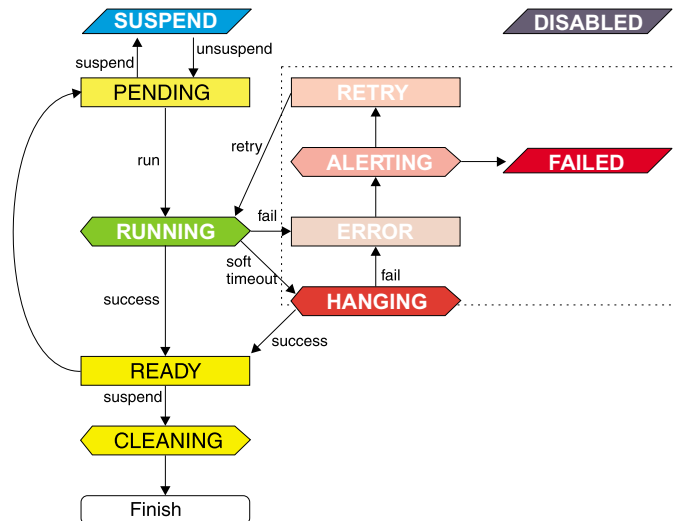
1.3. Attribute Precedence

Since the same attribute with different values can be assigned to, for example, a job and the queue that contains the job, Job Juggler uses a hierarchy when determining the effects of any given attribute. Specifically, attributes take precedence in Job Space from the smallest object to the largest (i.e. a job attribute takes precedence over the same attribute type assigned to a queue).

1.4. Job States

The Job Juggler schedules jobs. Each job performs a certain task. At any moment a job is in a certain state. For instance, when a job is waiting to be run, it is in state **PENDING**; when a job is running, it is in state **RUNNING**.

Job State Changes



Jobs can be either acyclic or cyclic. Acyclic jobs will only run once, whereas cyclic jobs will rerun repeatedly. When a cyclic job runs, it goes from state **PENDING** to **RUNNING** and then from state **RUNNING** to **READY**, from where it automatically changes to state **PENDING**. In this state it waits to receive a signal (trigger) in order to run again. When an acyclic job runs, it goes from state **PENDING** to **RUNNING**, then to **READY** and disappears.

If for some reason a job fails to run successfully it will move to state **ERROR**, and will then send an alert in state **ALERTING**. If the job is allowed to retry it will go to state **RETRY** and will eventually run again. Otherwise the job will be marked **FAILED**, where it will remain until its state is manually changed. If a job stays in state **RUNNING** for too long it may be marked with state **HANGING**; if it finishes successfully it will just become **READY**, otherwise it becomes state **ERROR**.

1.5. Job Output Redirection

Jobs under the control of Job Juggler have their output (**Stdout** and **Stderr**) intercepted by the Job Juggler server. The Job Juggler Server analyzes it and redirects it into multiple log files as required. Typically, there are 3 log files: critical only, critical and errors only, and all output. What constitutes 'critical' is user-defined (Job Juggler matches 'patterns' defined in a configuration file). These logfiles allow an operator to choose between viewing the error messages only, or viewing the error messages in the context of all the other messages.

In addition, output can be grouped in the same way across an entire queue, or the entire Job Space. This means an operator can choose to just see the output from a particular job, or that output in the context of the output from other jobs in a queue, or the output in the context from all the other jobs in Job Space.

2. Job Navigator

2.1. Introduction

Job Navigator is the intuitive user-friendly graphical user interface to the Job Juggler. This interface offers the user all the advantages of a graphical user interface such as:

- ❖ Drag and drop
- ❖ Visual overviews
- ❖ Zoom in and out
- ❖ Pop-up windows
- ❖ Mouse selections and clicks

Job Navigator visualizes by means of animating 'Job Space'. It can be used for monitoring jobs, diagnosing the status of a job, troubleshooting job failure or inactivity and resetting jobs and queues.

In general, everybody is allowed to see Job Navigator. However, making changes with Job Navigator is password protected and limited to only certain people.

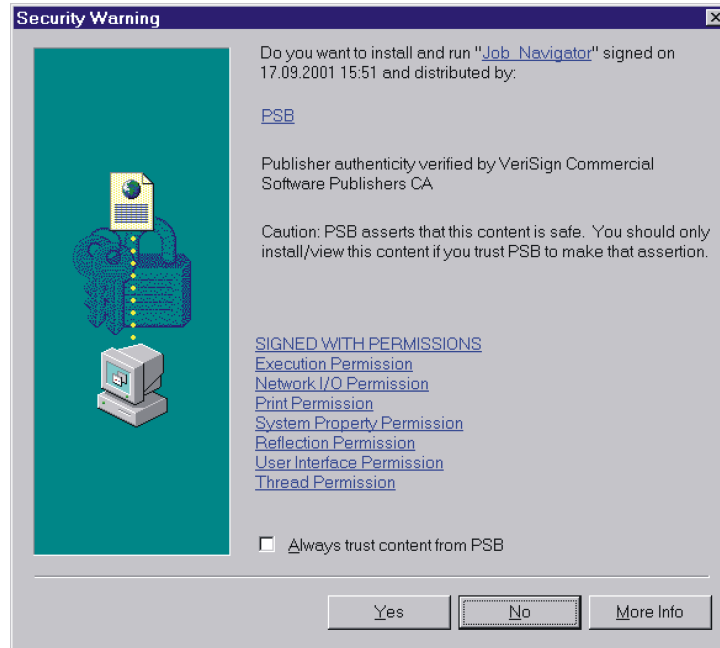
The layout of Job Navigator consists of a menu bar, buttons, toolbar and Job Space. This last element displays in animated form the scheduled processes. The functionality of these elements will be described in the next paragraphs.

2.2. Starting the Job Navigator

Job Navigator is a Java applet, which means that it runs inside an Internet browser. In order to start Job Navigator you have to open your Internet browser (e.g. Internet Explorer or Netscape Navigator) and type in the URL. The URL will refer to a machine from which the Job Navigator applet will be fetched. Typically, this machine will be the same machine on which the Job Juggler is running. The exact URL will depend on the location of the Job Navigator applet and how the web server on that machine is configured. This information is available after installation of the Job Juggler.

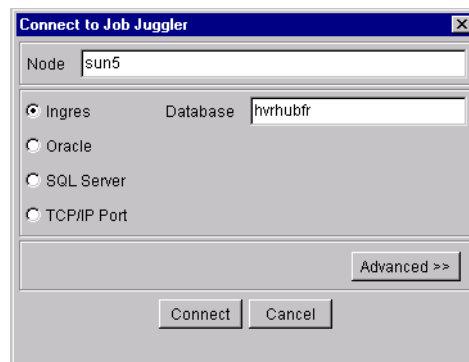
After the URL has been typed in, the browser fetches the Job Navigator applet, which in turn will connect to the Job Juggler. On Internet Explorer, a standard applet security warning appears containing information about necessary special permissions for running Job Navigator.

Internet Explorer applet
security pop-up



After you have clicked **Yes** in order to accept the security warning, the Job Navigator graphical user interface starts. The following pop-up appears:

Job Navigator connection
pop-up



In this pop-up you have to fill in the necessary information in order to establish the database connection.

In field **Node** you fill in the name of the hub machine where the Job Juggler is running.

Depending on the type of database, there are different fields that have to be filled in.

Platform	Field	Description
Ingres	Database	The name of the control database on the hub machine to which you would like to connect.
Oracle	Oracle SID	The Oracle instance identifier of the control machine.
	User (or schema)	The user name.
	TNS	Tick checkbox and fill in the TNS alias, if the Job Juggler connects to Oracle using a TNS alias.
SQL Server	Database	The name of the database.
	Database server	Tick checkbox if the Job Juggler connects to a Microsoft SQL Server database on another machine.
TCP/IP Port	Port Number	TCP/IP Port Number on which the Job Juggler Server is listening

After you have typed in the necessary information to establish the database connection, click **Connect**.

Once connection has been established, there will be two windows open: the original browser window containing the applet bar and the new Job Navigator window.

Job Navigator is directly connected to the Job Juggler on the hub machine. Therefore, Job Juggler has to be running in order to make a connection.

Menu File By accessing menu bar option **File** you can **Disconnect** a hub machine and reconnect to a different machine with **Connect**.

2.3. Browser Window

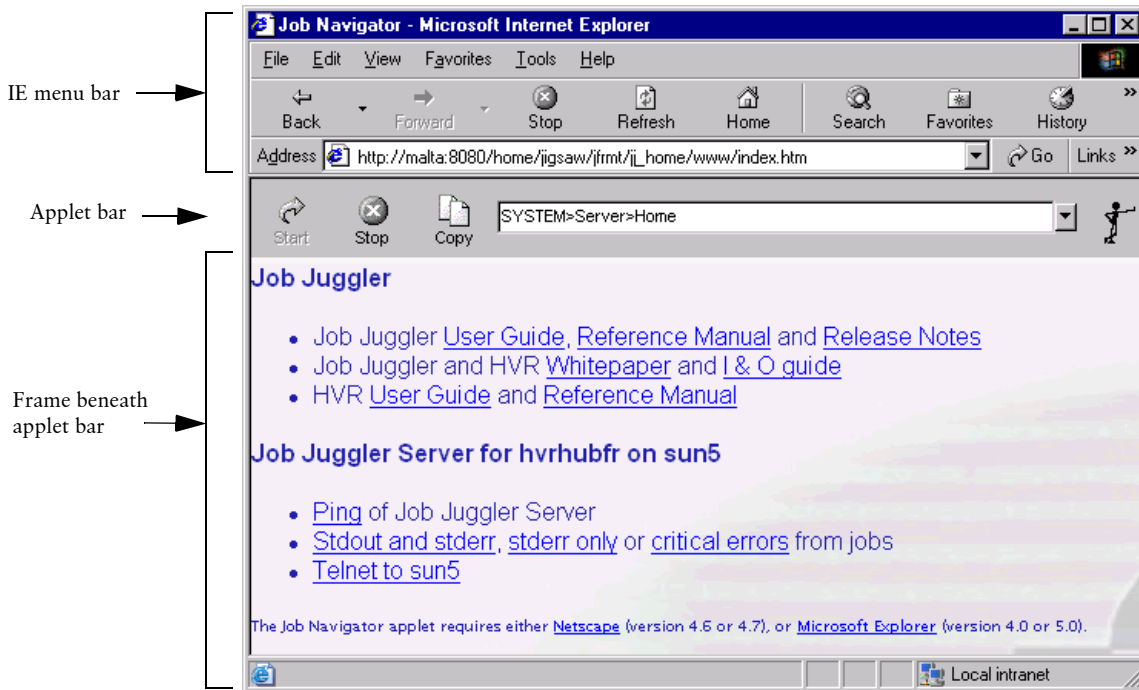
After connection has been established the original browser window looks as follows:

All windows that have been opened contain a warning text (**Warning: Applet Window**) in the status bar. This warning can be ignored by users, it is merely telling you that Java code is running.

Applet Bar The applet bar appears in the original browser window and is used to control the Job Navigator window. By clicking on **Stop**, the Job Navigator will be stopped. By clicking on **Start**, the Job Navigator will be restarted. The applet bar also contains a pull-down menu, which displays the last link that has been accessed in Job Navigator. By clicking on the drop-down button on the right side of the address bar you will see a list of recently visited links, which can be revisited by selecting the link. If you



want the information that is shown in the frame underneath the applet



bar to be shown or used in another window you can use **Copy** to transport that information to another window. Note that if the information in the frame to be copied was generated dynamically by a Job Juggler command, the copy button will cause that command to be re-run.

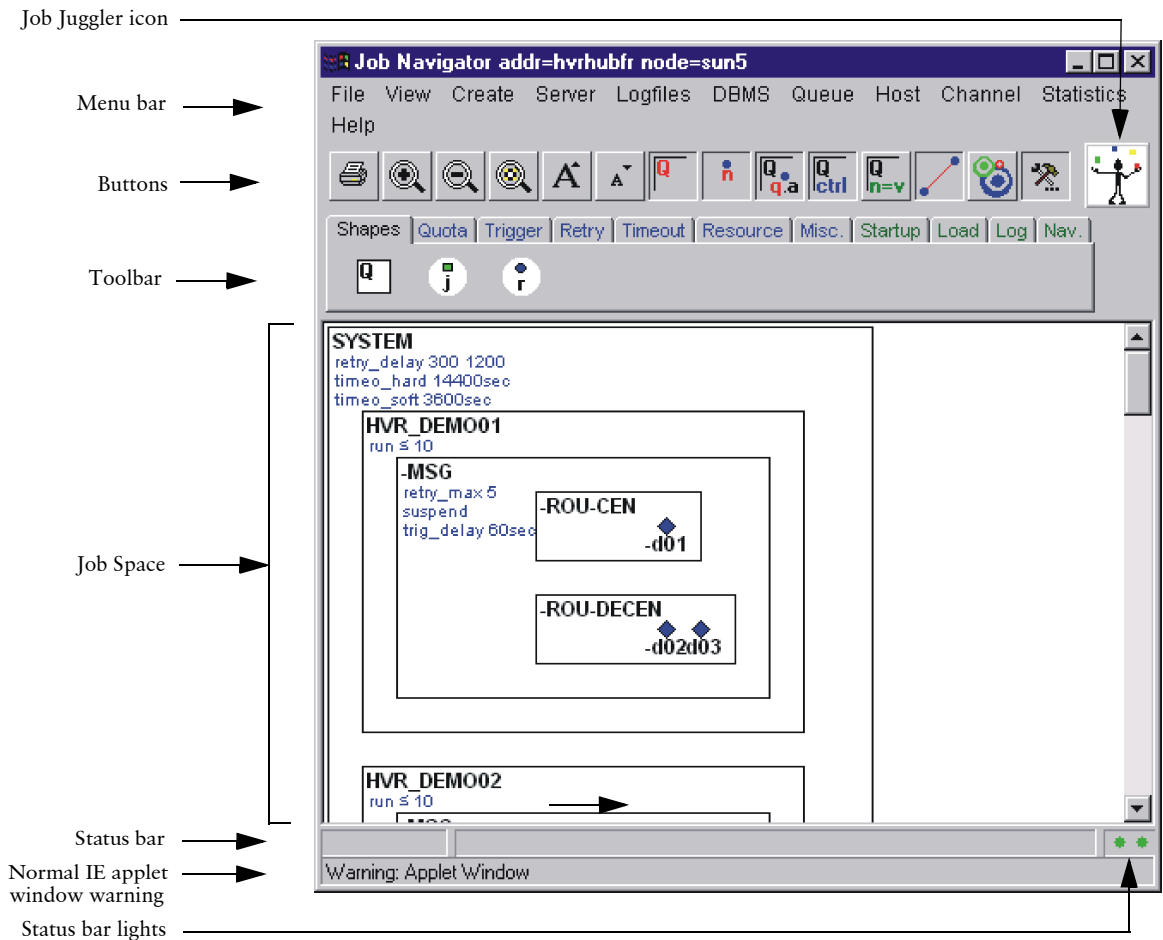
The frame beneath the applet bar is used by Job Navigator to display links and the output form various commands, such as changes made in the job detail window.



For more information on links, see 2.8. “Job Navigator Links”.

2.4. Job Navigator Window

The general layout of the Job Navigator consists of a menu bar, Job Juggler icon, Job Navigator buttons, Toolbar and Job Space.


















Job Juggler Icon

In the top-right corner of the Job Navigator window you can see the Job Juggler icon. This icon indicates the state of activity of all the replication processes monitored by the Job Navigator. If the Job Juggler icon is dark blue, there is high activity. A light blue background represents medium activity and a white background represents no recent activity.

Job Navigator Buttons

Job Navigator buttons allow you to control the way objects are displayed in Job Space such as what is shown and what is hidden, and the scale of the display. The following buttons are available:

Button	Description
	Print the visible part of Job Space.
	Zoom in to increase the scale of Job Space.
	Zoom out to decrease the scale of Job Space.
	Zoom in on the selected object.
	Bigger font: increases the font of the objects and attributes in Job Space.

Button	Description
	Smaller font: decreases the font of the objects and attributes in Job Space.
	Hide or show the names of queues and network segments.
	Hide or show the names of jobs and resources.
	Abbreviate names in Job Space. This button abbreviates the full names of queues and jobs by omitting elements of the name, which occur in enclosing shapes. For example, if a job's full name is <code>hvr_demo01-msg-rou-d03</code> but that job is enclosed by a queue whose full name is <code>HVR_DEMO01-MSG</code> then this button will abbreviate the job's name into <code>-rou-d03</code> . Likewise if queue <code>HVR_DEMO01</code> was enclosed by a queue named <code>HVR_DEMO01-MSG</code> , then this button would cause the name of that queue to be abbreviated to <code>-MSG</code> .
	Hide or show the control attributes: the quotas, schedule times, retry and time-out attributes.
	Hide or show the environment attributes: the startup, load balancing, log output, navigation and miscellaneous attributes.
	Hide or show the connections between queues or jobs and resources.
	Hide canvas padding, hides or displays white space behind text so that text becomes less or more legible.
	Hide or show colored activity ripples of a job. The color indicates the state of activity of the job. The ripple is a kind of circle moving away from the job symbol.
	Hide and show the toolbar, for creating new shapes and attributes.


Toolbar

The toolbar is used to create new shapes and attributes. It consists of various tab sheets from which you can drag and drop shapes, control attributes and environment attributes.

In order to create a queue, job or resource, just drag the icon from the shape tab to the required location in Job Space. In order to create an attribute, drag the attribute from the toolbar and drop it on a queue, job or resource. A detail window appears in which you can specify arguments. When the new shape or attribute is saved, it will be displayed in Job Space.

When dragging and dropping attributes, the display bar displays text, indicating the name of the attribute you are moving as well as the job or queue, which you are currently above. If you are dropping an attribute or shape on an invalid part in Job Space, the attribute will be displayed in a pale color and there will be no message in the status bar - thereby indicating that the action is invalid.

The following shapes and attributes are available:



↑ Create a queue.
 ↑ Create a job.
 ↑ Create a resource.



↑ Create **suspend** attribute to bring jobs into SUSPEND state.
 ↑ Create **quota_run** attribute to set max. running jobs in queue.
 ↑ Create **quota_children** to set max. child processes.
 ↑ Create **quota_speed** to set max. n spawns per $secs$ seconds.



For more information see 3.1. “Queue Quotas” of the Job Juggler Reference Manual.



↑ Create **trig_delay** to trigger $secs$ after last successful run.
 ↑ Create **trig_crono** at moment $crono$, trigger job.
 ↑ Create **trig_at** to trigger job at tim .
 ↑ Create **period** to set jobs $period$ during window $crono$.



For more information see 3.2. “Trigger Attributes” Of the Job Juggler Reference Manual.




↑ Create **retry_max** to allow n retries of failed job.
 ↑ Create **retry_delay** to initial retry delay $isecs$, max $fsecs$.




For more information see 3.3. “Retry Attributes” of the Job Juggler Reference Manual.



↑ Create **timeo_soft** to send a time-out message after $secs$ running.
 ↑ Create **timeo_hard** to kill job after running too long.
 ↑ Create **timeo_aux** to kill alerting or cleaning too long.

 For more information see 3.4. “Time-out Attributes” of the Job Juggler Reference Manual.



↑ Create **res_require** for jobs that need resource *name*.
 ↑ Create **res_unlock_delay** for secs of delay before re-using resource.
 ↑ Create **res_share** so jobs may share the same resource.

 For more information see 3.5. “Resource Attributes” the Job Juggler Reference Manual.



↑ Create **breakpoint** to hold job when it reaches *STATE*.
 ↑ Create **params** to write job parameters into file.
 ↑ Create **secondary_server** for jobs to be run by a different Job Juggler.

 For more information see 3.6. “Miscellaneous Attributes” of the Job Juggler Reference Manual.




↑ Create **set** to set variable *NM* to value *val*.
 ↑ Create **path** to add *dir* to command invocation path.
 ↑ Create **arg** to set argument *n* for running job.
 ↑ Create **cmd_run** to set command to run job.
 ↑ Create **cmd_alert** to set command for alert action.
 ↑ Create **cmd_clean** to set command for clean action.

 For more information see 3.7. “Startup Attributes” of the Job Juggler Reference Manual.




↑ Create **max_sessions** to set max. control sessions.
 ↑ Create **tick_duration** to limit tick to *min* & *max* milliseconds.
 ↑ Create **nap_duration** to limit nap to *min* & *max* milliseconds.
 ↑ Create **timeo_db** to warn if database slave takes too long.

 For more information see 3.8. “Load Balancing” of the Job Juggler Reference Manual.



Create `log_cc` to carbon-copy severity *SEV* to *file*.
 Create `log_to` to redirect severity *SEV* output to *file*.

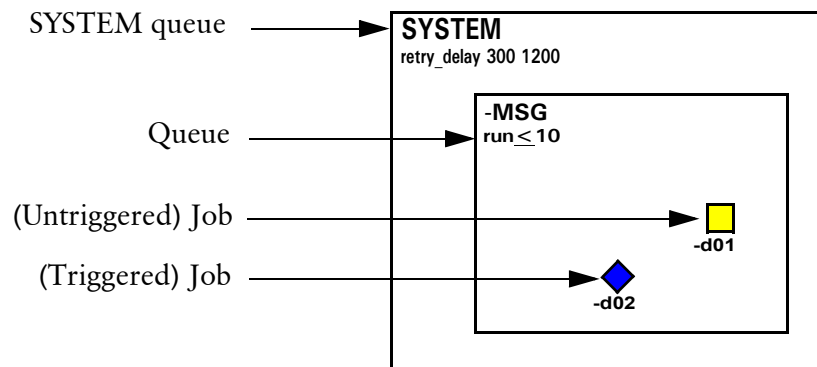
For more information see 3.9. “Log Output Attributes” of the Job Juggler Reference Manual.



Create `nav_applet` to set Job Navigator link to JAVA applet *appl*.
 Create `nav_include` to set file containing Job Navigator links.
 Create `nav_url` to give Job Navigator a link to *url*.

For more information see 3.10. “Job Navigator Attributes” of the Job Juggler Reference Manual.

You can highlight certain attributes in Job Space, by double clicking on an attribute type in the toolbar. Obviously, for this to happen, control attributes or environment attributes must be visible. Therefore, buttons [Hide control attributes](#) and [Hide environment attributes](#) must be selected in the toolbar.



Job Space The main section of the Job Navigator is Job Space. In it you will see queues represented as boxes and jobs represented as small square icons. If jobs have been triggered, they have the shape of a diamond instead of squares. The color of the job indicates in the job’s state. For instance, yellow means that a job is **PENDING**, green means that a job is **RUNNING**, and red means that a job has **FAILED**.

Job Space is animated in Job Navigator. Whenever there is activity going on in the Job Juggler, it is immediately visualized.

With the use of horizontal and vertical scroll bars it is possible to scroll to certain queues and jobs within Job Space. Another possibility is the use of the zoom in and out buttons.

Besides jobs changing their colors when their state changes, queues also change color. These are called activity colors.

The type of color provides you with information on how recently activity occurred in that queue. Queues displayed in dark blue represent a queue containing jobs that have recently changed state (recent activity). Pale blue queues represent less recent activity. Queues just shown with black lines represent no recent activity.




The Job Navigator icon also changes color depending on the activity status of the most recent activity in Job Space.

Status Bar Lights

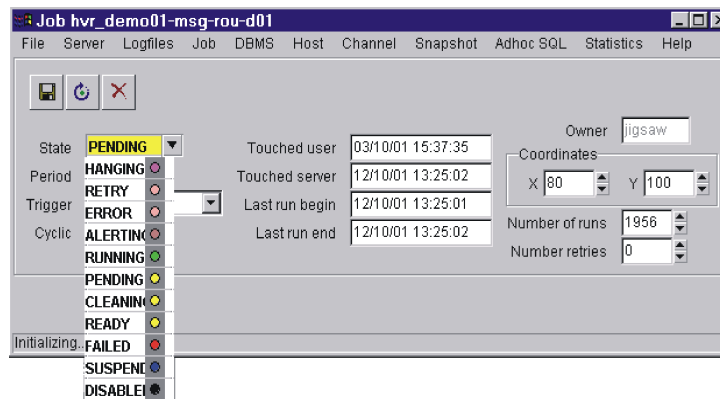
In the status bar there are two lights indicating job juggler activity. The status bar lights change color as data is transferred between the Job Navigator and the Job Juggler.

2.5. Detail Windows

By double clicking on a job, queue or attribute you access their details. A new dialog window opens; the title bar of the dialog screen displays its type and name. By default, all details windows contain the following three buttons:

-  Save the changes that have been made.
-  Refresh the status of the shape.
-  Delete job or queue.

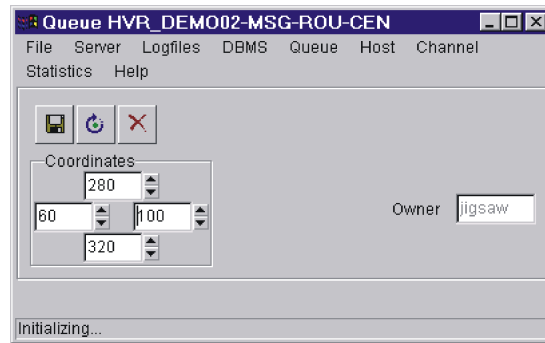
Job Detail Window



If you double click on a job the dialog window displays the details of a job. The menu bar of the dialog window contains links that are related to the selected job. The dialog window also contains the following fields:

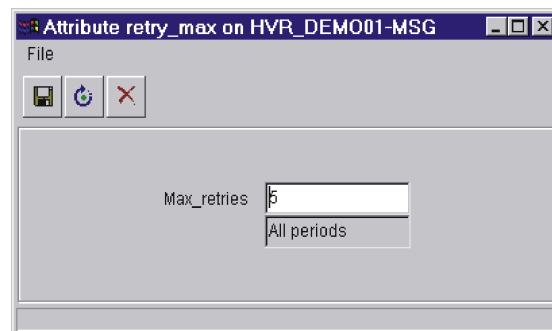
Field	Description
State	The state a job is in at the moment.
Period	The period in which the job is active (feature which is seldom used).
Trigger	Indicates whether a shape has been set to run or not. A triggered job is shown in Job Space as a diamond. If it has not been triggered, it will not run automatically. A non-triggered job is shown as a square.
Cyclic	If a job is cyclic it will change to state PENDING again when it has finished its activity. If a job is a-cyclic (checkbox is not ticked), the job disappears from Job Space after it has run successfully.
Touched user	The last time a user changed the job. Typically, this is when the job was created.
Touched server	The last time the Job Juggler server last changed the job. Typically, this is the last time the job changed from state RUNNING to another state.
Last run begin	The date and time when the job last started running.
Last run end	The date and time when the job last finished running.
Owner	The person who created the job.
Coordinates	The coordinates of the job in Job Space.
Number of runs	The number of times the job has run.
Number of retries	If it has failed, the number of times the job has been retried without successfully completing.

Queue Detail Window



In the dialog box of a queue, it is only possible to change the **Coordinates** of the queue. Furthermore, it contains the field **Owner**, which holds the name of the creator of the queue.

Attribute Detail Window



If you double click on an attribute in Job Space you will get a pop-up containing details on the selected attribute. In this pop-up you can view and change settings depending on the type of attribute. For instance, for a retry attribute you can change the maximum number retries that are allowed.

2.6. Controlling Jobs

Under normal circumstances, all of the jobs and queues in Job Space will have been generated in the Job Juggler by [hvrload](#). Therefore, a user does not normally create HVR jobs or queues. Modification only occurs in special circumstances. This can be done by accessing the details of jobs and queues in the Job Navigator.

Triggering jobs When you want to trigger a job, or the jobs in a queue you can double click on a job and reset field [Trigger](#) of the detail window, or you can set a trigger by accessing the links section. In order to trigger a single job, click on a job and select link [Job > Trigger](#). In order to trigger all the jobs within a queue, click on a queue and select link [Queue > Trigger](#).

Suspending jobs To suspend jobs in a queue, simply right-click on a queue and select link [Queue > Suspend](#). Another way to suspend jobs in a queue is to select the Quota tab from the toolbar and drag and drop the [suspend](#) attribute onto a queue. When you want to delete a [suspend](#) attribute, select the attribute and press the delete button on your keyboard.



For more information see [2.7. “Disabling and Suspending jobs”](#).

Changing schedule attribute In order to change scheduling of a job or queue, drag and drop an attribute from the Quota, Trigger, Retry, or Timeout tabs of the toolbar onto the job or queue. It is also possible to double click on an existing scheduling attribute and make changes in the detail window that appears.

Changing job states In order to change the state of a job, double click on the job that you want to change. A detail window appears in which you can change the state of a job by selecting another job state in field [State](#). After you have clicked [SAVE](#) the job state will have changed.

Other Job Space changes In the event you need to add a job or queue, access the shapes toolbar and drag and drop the job or queue into the preferred location in Job space.

You can delete a job, queue or attribute by accessing the detail window. Click in the detail window on the button [Delete](#). Another way of deletion is to select the item you wish to delete and press the delete button on your keyboard.

When you want to move a job or resize a queue you have to access the detail window where you can change its coordinates. It is also possible to select a shape and move or resize it by dragging.

After you have deleted or resized an item in Job Space, a confirmation pop-up appears requesting you to confirm the changes you are about to make. Every change that has been made in Job Space through the Job Navigator will be recorded in a special logfile [jj.ctrl](#) on the hub machine. This file can be viewed from the Job Navigator by using link [Logfiles > Control log](#).

2.7. Disabling and Suspending jobs

There are two ways to prevented jobs in the Job Juggler from running;

- ❖ Abruptly, by explicitly setting the state of an individual job to **DISABLED**. Once in **DISABLED** state a job is immune to triggering, retry and the presence or absence of suspend attribute; it can only be reactivated by explicitly resetting its state (e.g. back to **PENDING**).
- ❖ Gracefully, by adding a **suspend** attribute to a queue containing the job. **PENDING** jobs will immediately assume state **SUSPEND**, and jobs which are **RUNNING** or **HANGING** will continue running but assume state **SUSPEND** when they finish. Jobs in other states (**DISABLED**, **FAILED** and **RETRY**) ignore the suspend attribute, although job retrying is inhibited.

The **SUSPEND** state means a job has been halted due to the existence of a suspend attribute. If you try to set a running job's state to **SUSPEND** the Job Juggler may kill it and force it back to **SUSPEND** but the job will then spring back to **PENDING** state when it sees it is not affected by a **suspend** attribute.

There are 3 methods of adding a **suspend** attribute:

1. In Job Navigator by drag & drop from attribute toolbar.
2. In Job Navigator using the **Queue>Suspend** right-click menu option.
3. Using command line utility **jjsuspend**

Methods 2 and 3 also allow the definition of a suspend timeout (using option **Queue>Suspend>Allow 1 minute timeout** or the **jjsuspend -t** option); if running jobs do not finish within that amount of time they are killed and forced into **SUSPEND** state. This suspend timeout feature differs from the timeout of the **jjtrigger** command because if a triggered job does not complete in time it is left running and only **jjtrigger** stops.

2.8. Job Navigator Links

Job Navigator is used for three things: (1) animating scheduling for spotting errors, (2) changing scheduling at run time (add or remove attributes) and (3) providing context sensitive links for a job or a queue. These links can be used to diagnose problems. This section describes how to access links and describes the links that are commonly used.

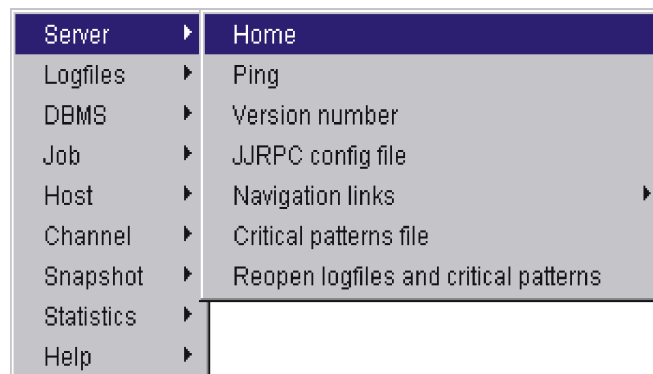
There are three ways of accessing links: by right clicking a job or queue, by double clicking on a job or queue and by using the menu bar of Job Navigator.

Right click: By right-clicking on a job or a queue a sub menu appears containing links for various purposes. Note that the sub menu for jobs is different than the sub menu for queues. This is because jobs and queues serve different functions.

Detail window: Another method of obtaining links is by double clicking on a job or queue. A detail window appears containing detailed information on that job or queue. In this detail window you can select links by choosing the menu bar items.

Menu bar of Job Navigator: The menu bar of the Job Navigator also contains links. These are the links for the system queue. So, if you right click, or double click, on the **SYSTEM** queue you will get the same links as the links displayed in the menu bar of Job Navigator.

Sub menu after right clicking on a job



When you choose a link, the output of the URL is sent to the frame underneath the applet bar (Internet browser window).

New Browser Window

Normally, if you choose a link in Job Navigator it will be displayed in the frame underneath the applet bar (browser window). However, when you keep the **SHIFT** button pressed as you click on a link, the output of that link will be shown in a new browser window.

Log Files

By selecting a log file you have to keep in mind that this is a continuous tail. So as it runs, there is continually information added to the log file. When viewing a log file using the Job Navigator, new lines will be automatically appended to the bottom of the frame. If the log contains any error messages, they will be shown in red. Some log information also contains hyperlinks. By clicking on such a link you look into that particular file.

If a log is a large file, it only shows the last (most recent) part of the file. This can be identified by a statement in the top of the screen such as “[Skipped to last 32k of 102k \(view entire file\)](#)”. By clicking on the link ‘[entire file](#)’ you can access the complete log file.

Output from link tail of job’s logfile, displayed in browser window

```

Tail of /home/jigsaw/jig_gsja/jj_config/log/hubdb/hvr_demo01-
Skipped to last 32K of 33K file (view entire file)

'hvr_demo01-msg-rou-d01' has been in state 'RUNNING' longer than the 2 \
seconds allowed by attribute 'timeo_soft' defined on object \
'hvr_demo01-msg-rou-d01'. It will be marked as HANGING but will not be \
interrupted. The line which is currently being executed is unknown.
Oct 12 11:55:00: hvr_demo01-msg-rou-d01[eof]: Finished. (elapsed=4.36s)
Oct 12 11:55:12: hvr_demo01-msg-rou-d01[eof]: Finished. (elapsed=1.08s)
Oct 12 11:55:23: hvr_demo01-msg-rou-d01[eof]: Finished. (elapsed=1.09s)
Oct 12 11:55:34: hvr_demo01-msg-rou-d01[eof]: Finished. (elapsed=1.09s)
Oct 12 11:55:45: hvr_demo01-msg-rou-d01[eof]: Finished. (elapsed=1.07s)
Oct 12 11:55:56: hvr_demo01-msg-rou-d01[eof]: Finished. (elapsed=1.10s)
Oct 12 11:56:07: hvr_demo01-msg-rou-d01[eof]: Finished. (elapsed=1.09s)
Oct 12 11:56:18: hvr_demo01-msg-rou-d01[eof]: Finished. (elapsed=1.08s)
Oct 12 11:56:29: hvr_demo01-msg-rou-d01[eof]: Finished. (elapsed=1.22s)
Oct 12 11:56:40: hvr_demo01-msg-rou-d01[eof]: Finished. (elapsed=1.08s)
Oct 12 11:56:51: hvr_demo01-msg-rou-d01[eof]: Finished. (elapsed=1.27s)
Oct 12 11:57:02: hvr_demo01-msg-rou-d01[eof]: Finished. (elapsed=1.10s)
Oct 12 11:57:13: hvr_demo01-msg-rou-d01[eof]: Finished. (elapsed=1.09s)
Oct 12 11:57:24: hvr_demo01-msg-rou-d01[eof]: Finished. (elapsed=1.10s)
Oct 12 11:57:35: hvr_demo01-msg-rou-d01[eof]: Finished. (elapsed=1.10s)
Oct 12 11:57:46: hvr_demo01-msg-rou-d01[eof]: Finished. (elapsed=1.10s)
Oct 12 11:57:57: hvr_demo01-msg-rou-d01[eof]: Finished. (elapsed=1.08s)
Oct 12 11:58:08: hvr_demo01-msg-rou-d01[eof]: Finished. (elapsed=1.08s)

```

Links to Job Juggler

Here are some common Job Juggler links:

Server > Home: shows you the home page for the Job Juggler server. This contains general links.

Server > Ping: provides a so-called ping. It means that it checks whether the Job Juggler server is running.

Server > Version number: provides information on what version of the Job Juggler is currently being used.

Logfiles > SYSTEM stderr & stdout: provides log files of activity that reached the **SYSTEM** queue; both error and non-error messages are included.

Job > Stdout & Stderr: provides all output from a particular job.

Job > Trigger: can force the job to be run immediately, providing quotas are sufficient and no **suspend** attribute exists.

Job > Delete: will provide you the possibility to delete a job.

DBMS > Ingres errlog_log on [machine]: provides access to error logs of the DBMS on a given machine.

Host > Telnet to [machine]: provides the opportunity to telnet (obtain access) onto a machine.

Queue > Trigger: allows you to force the jobs within a queue to be run immediately.

Queue > Suspend: will suspend all jobs within a queue; jobs which are currently running will be suspended when they finish running.

Help: provides links to the various online Help documentation, such as this document.

Other possible links vary according to the specific job or queues selected.



For more information see 2.3. “Browser Window”.

Links to HVR

When you are using Job Juggler in combination with the High Volume Replicator there are a number of specific links for jobs and queues, of which the commonly used links are described here.

Job > Jot script displays the jot script.

Output of Jot script

```

File: /home/jigsaw/jfrja/hvr_config/jotgen/hubdb/hvr_demo01/hvr_
open rou db=hubdb /* HVR hub */
open d03 class=ingres db=hvrd03_a role=hvr_integrate
t= rou.StateHub.hvr_demo01.d03.tog
commit
if rou.StateHub.hvr_demo01.d03.busy != 0 || $1 == '-R' then
  commit
  echo "Recovery: After actions for rou to d03 were interrupted..."
  if $t != d03.StateInt.hvr_demo01.tmir then
    commit
    echo "Recovery: Resending buffers from rou to d03..."
    copy rou:dm01_product d03i+!$t | mask | pack | to d03 | pack -u | coer\
[]d03.Describe(dm01_product__it) | write -b dm01_product__it
    copy rou:dm01_order_d03i+!$t | mask | pack | to d03 | pack -u | coer [\
]d03.Describe(dm01_order__it) | write -b dm01_order__it
    d03.StateInt.hvr_demo01.tmir= $t
    commit
  endif
  commit
  call rou.TruncCommit(dm01_product_d03i+!$t dm01_order_d03i+!$t )
  if $1 != '-i' then
    call d03.Integrate(FAIL ALL TRUNC hvr_demo01 dm01_product__it \
dm01_order__it)
  endif

```

Channel > [channel] channel definition catalogs > as tables displays the content of the catalogs for a channel definition.

Output of channel definition catalogs, as tables

The screenshot shows a web browser window with the following content:

Address: http://jamaica:8080/home/jigsaw/irja/rl_home/cgi-sys/jnav.cgi?jot=jot_arch_def&HVR_HOME=/home/jigsaw/irja/hvr_home&l

jot jot_arch_def hubdb hvr_demo01

Catalog hvr_channel

chn_name	chn_descrip
hvr_demo01	Simple reference channel.

Catalog hvr_buffer

chn_name	buf_code	base_table	addr_type	seq_major	seq_minor
hvr_demo01	dm01_order	dm01_order	ToAny	1	2
hvr_demo01	dm01_product	dm01_product	ToAny	1	1

Catalog hvr_buf_column

chn_name	buf_code	col_sequence	col_name	col_key_sequence	col_datatype	col_length	col_nullable	col_long

Catalog hvr_loc_group

chn_name	grp_suffix	grp_id	single_member	file_class_only	grp_descrip
hvr_demo01	CEN	1	1	0	Headquarters
hvr_demo01	DECEN	2	0	0	Decentral

Catalog hvr_action

chn_name	grp_suffix	buf_code	act_type	act_arg1	act_arg2
hvr_demo01	CEN	*	StdCap		
hvr_demo01	DECEN	*	StdInt		

Snapshot > compare [loc1] with [loc2] > Checksums only displays the output of a checksum comparison of two databases, indicating whether they are synchronized.

Snapshot > compare [loc1] with [loc2] > Row-by-row displays the output of a row-by-row comparison of two databases, indicating whether a row needs an insert, update or deletion.

Snapshot > Refresh [loc1] with [loc2] > Bulk refresh synchronizes two databases by overwriting the entire tables of *loc1* with the tables of *loc2*.

Snapshot > Refresh [loc1] with [loc2] > Row-by-row synchronizes two databases by merging their contents row by row.

Adhoc SQL > List buffer objects in [location] displays all buffer objects in a certain database location for the current channel.

Adhoc SQL > Show # rows in buffer tables in [location] displays the number of rows present in each buffer for the current channel.

Adhoc SQL > Move contents of fail buffers in [location] into integrate buffers causes the contents of all fail buffers for this channel to be moved into the integrate buffers.

Adhoc SQL > Empty contents of fail buffers in [location] causes the contents of the current channel's fail buffers to be cleared.

Channel > Hvrload_log on [machine] the output log of hvrload commands run with option **-v**.

Index

A

action 9, 20, 22
 ad-hoc query 32
 ad-hoc sql 32
 ALERT 14, 22
 ALERTING 14, 21
 arg 22
 attribute 13, 19, 20, 21, 23, 25, 26, 27, 28, 29, 30

B

breakpoint 22
 buffer 32
 buffer object 32
 buffer table 32

C

channel 31, 32
 channel definition 31, 32
 CLEANING 21
 cmd_alert 22
 cmd_clean 22
 cmd_run 22
 compare 32
 critical 14
 crono pattern 21

D

database location 32
 date 25
 Define 13
 delete 25, 27, 30
 DISABLED 28

E

error 14, 29, 30

H

HANGING 14, 28
 hub 16, 17, 27
 HVR_HOME 9
 hvrload 27, 32

I

Ingres 9, 17, 30
 insert 32
 installation 8, 15
 integration 32
 integration buffer 32
 Internet Explorer 15, 16

J

JAVA 15, 17, 23
 jjsuspend 28
 jjtrigger 28
 job 3, 7, 8, 11, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31
 Job Juggler 3, 7, 8, 11, 13, 14, 15, 16, 17, 18, 19, 21, 22, 23, 24, 25, 27, 28, 30, 31
 Job Navigator 8, 11, 13, 15, 16, 17, 18, 19, 23, 24, 27, 28,

29

job space 13, 14, 15, 19, 20, 23, 24, 25, 26, 27
 jot 31
 jot script 31

L

layout 15, 19
 legible 20
 log_cc 23
 log_to 23
 long 14, 21, 22

M

max_sessions 22
 message channel 20, 21

N

nap 22
 nav_applet 23
 nav_include 23
 nav_url 23
 Netscape 15
 network 20
 network segment 20

O

Oracle 9, 17
 output 13, 14, 18, 20, 23, 29, 30, 31, 32

P

params 22
 PATH 22
 path 22
 PENDING 13, 14, 23, 25, 28
 period 13, 21, 25
 port number 17

Q

queue 13, 14, 15, 20, 21, 23, 24, 25, 26, 27, 28, 29, 30, 31
 quota 13, 20, 21, 27, 30
 quota_children 21
 quota_run 21
 quota_speed 21

R

READY 14
 refresh 25, 32
 resource 20, 21, 22
 RETRY 13, 14, 20, 21, 26, 27, 28
 retry_delay 21
 retry_max 21

S

scheduling 13, 27, 29
 Secondary_server 22
 segment 20
 severity 23
 snapshot 32
 SQL 9, 17, 32

SQL Server 9, 17
suspend 21, 27, 28, 30

T

-t option 28
task 13
TCP/IP 17
tick 17, 22
timeo_db 22
timeo_hard 21
timeo_soft 21
timeout 13, 20, 21, 27, 28
transport 18
trig_at 21
trig_crono 21
trig_delay 21

U

Unix 9
update 32

V

-v option 32

W

Windows 9, 15, 17, 25

