

RCE User Guide

Build 8.3.0.0202001221326_RC

Table of Contents

1. Preface	1
1. Abstract	1
2. Intended Audience	1
3. License Information	1
4. Compatible Operating Systems	1
4.1. Support of 32 Bit Operating Systems	2
4.2. Known Issues	2
4.2.1. KDE on Red Hat Enterprise Linux 7	2
4.2.2. KDE with Oxygen	2
4.2.3. Jython scripts are executed sequentially	2
4.2.4. 32-bit Java is not supported	2
2. Setup	3
1. Installation on Windows	3
1.1. Prerequisites	3
1.2. Installation	3
1.3. Digital Signatures and Download Verification	3
1.4. Starting RCE as a GUI Client	4
1.5. Starting a Non-GUI ("Headless") Instance	5
1.6. Installation as a Service on a Windows Server	5
1.6.1. Installation and Service Management	5
1.6.2. Service Configuration	6
2. Configuration	6
2.1. Configuration Locations and Files	7
2.2. Configuration Parameters	7
2.3. Configuration UI	13
2.3.1. Remote Access: SSH account configuration	13
2.3.2. Mail: SMTP server configuration	13
3. Usage	15
1. Graphical User Interface	15
2. Workflows	17
2.1. Rationale	18
2.2. Getting Started	18
2.3. Workflow Components	18
2.4. Coupling Workflow Components	19
2.5. Execution Scheduling of Workflow Components	20
2.6. (Nested) Loops	21
2.7. Fault-tolerant Loops	22
2.8. Manual Tool Result Verification	23
3. Commands	23
3.1. Command Line Parameters	23
3.2. Profile Selection UI	25
3.3. Command Shell	26
3.3.1. Configuration Placeholder Value Files	28
4. Integration of External Tools	30
4.1. Basic Concepts	30
4.2. Directory Structure for Integrated Tools	30
4.3. Copying of Integrated Tools	32
4.3.1. Publishing a Component on a Server Instance in Headless Mode	32
4.3.2. Tool Execution Return Codes	32
4.4. Integration of CPACS Tools	33
4.4.1. Additional concepts of CPACS Tool Integration	33
4.4.2. Integrate a CPACS Tool into a Client Instance	35
4.4.3. Integrate a CPACS Tool into a Server Instance in Headless Mode	37
5. Remote Tool and Workflow Access	37
5.1. Basic Concepts	37

5.2. Tool vs. Workflow Execution	37
5.3. Setting up the Single Tool Execution Example	37
5.4. Setting up the Workflow Execution Example/Template	39
5.5. Building Your Own Remote Access Workflow	40
6. Connecting RCE instances via the RCE network or via SSH connections	41
6.1. RCE Network Connections	41
6.2. SSH Remote Access Connections	42
6.2.1. Configuring an RCE instance as an SSH server	42
6.2.2. Configuring an RCE instance as an SSH client	42

List of Figures

2.1. Configuration tool for SSH account and SMTP server configuration	13
3.1. Workbench with different views and the workflow editor opened	15
3.2. Connection Editor	16
3.3. Network View	16
3.4. Workflow Data Browser	17
3.5. Workflow Console	17
3.6. Profile Selection UI	26
3.7. Run process of an user-integrated CPACS Tool	35

List of Tables

2.1. General	7
2.2. BackgroundMonitoring	8
2.3. Network	8
2.4. Publishing	9
2.5. Component identifiers used in configuration group 'publishing'	9
2.6. ComponentSettings	10
2.7. thirdPartyIntegration	10
2.8. sshServer	10
2.9. Possible roles for SSH accounts	11
2.10. sshRemoteAccess	12
2.11. smtpServer	12
3.1. Data Type Conversion Table	19
3.2. Inputs of Optimizer	21
3.3. Outputs of Optimizer	21
3.4. Inputs of Design of Experiments	22
3.5. Outputs of Design of Experiments	22
3.6. Inputs of Parametric Study	22
3.7. Outputs of Parametric Study	22
3.8. Inputs of Converger	22
3.9. Outputs of Converger	22
3.10. Command line arguments for RCE	24
3.11. Shell Commands	27
3.12. Components and their configuration placeholders	29
3.13. Connection types	41

Chapter 1. Preface

This chapter gives an introduction to RCE.

1. Abstract

RCE (Remote Component Environment) is an open source software to help engineers, scientists and others to create, manage and execute complex calculation and simulation workflows. A workflow in RCE consists of components with predefined inputs and outputs connected to each other. A component can be a simulation tool, a tool for data access, or a user-defined script. Connections define which data flows from one component to another. There are predefined components with common functionalities, like an optimizer or a cluster component. Additionally, users can integrate their own tools. RCE instances can be connected between each other. Components can be executed locally or on remote instances of RCE (if the component is configured to allow this). Using these building blocks, use cases for complex distributed applications can be solved with RCE.

2. Intended Audience

The intended audience of this document are engineers, scientists, and everybody else interested in developing automated workflows with RCE.

3. License Information

RCE is published under the Eclipse Public Licence (EPL) 1.0. It is based on Eclipse RCP 3.7.2 (Indigo SR2), which is also published under the Eclipse Public Licence (EPL) 1.0. RCE also makes use of various libraries which may not be covered by the EPL; for detailed information, see the file "THIRD_PARTY" in the root folder of an RCE installation. (To review this file without installing RCE, open the RCE release .zip file.)

For downloads and further information, please visit <http://www.rcenvironment.de>.

4. Compatible Operating Systems

RCE releases are provided for Windows and Linux. It is regularly tested on

- Windows 7
- Windows 10
- Windows Server 2012 R2
- SUSE Linux Enterprise Desktop ("SLED") 12 SP2
- Ubuntu 16.04 LTS
- Red Hat Enterprise Linux 6

4.1. Support of 32 Bit Operating Systems

Starting with release 8.0.0, RCE is only shipped for 64 bit systems. If you still require 32 bit packages, you can continue to use previous RCE releases, but there will be no standard feature or bugfix updates for them.

4.2. Known Issues

4.2.1. KDE on Red Hat Enterprise Linux 7

On Red Hat Enterprise Linux 7 with KDE 4, RCE (like any other Eclipse-based application) can cause a segmentation fault at startup. If you encounter such an issue, you can try choosing a different GTK2 theme:

1. Open the **System Settings** application (systemsettings).
2. Go to **Application Appearance**
3. Open **GTK** page
4. Switch the GTK2 theme to "Raleigh" or "Adwaita" and click on **Apply**

4.2.2. KDE with Oxygen

On Unix Systems using KDE as desktop environment and Oxygen as theme it can happen that RCE crashes when certain GUI elements are shown. It is a known issue in the theme Oxygen and happens on other Eclipse-based applications as well. If you encounter such an issue, please choose a different theme like "Raleigh" or "Adwaita".

4.2.3. Jython scripts are executed sequentially

The Script component can use Jython for the evaluation of scripts and the pre- and postprocessing of integrated tools always uses Jython. Due to a known bug in the Jython implementation it is not possible to execute several Jython instances in parallel. Therefore, the execution will be done sequentially. If several Script components should be executed in parallel, Python should be used instead.

4.2.4. 32-bit Java is not supported

Running RCE with a 32-bit Java Runtime Environment doesn't work. On some operating systems an error dialog will be displayed in this case, on some other systems nothing will happen at all. Therefore, always make sure a 64-bit Java Runtime Environment is used to run RCE.

Chapter 2. Setup

This section describes the installation and configuration of RCE.

1. Installation on Windows

1.1. Prerequisites

To run RCE on a system, the only prerequisite is an installed Java Runtime Environment (JRE), version 7 or above. If you don't already have one on your machine, you can download it from

```
http://www.java.com/download/
```

and install it. Starting with RCE 8.x only 64 bit packages of RCE are made available. Therefore, please make sure to install the 64 bit version of the JRE.

1.2. Installation

On Windows, a single zip file is provided to set up client and server installations, which can be downloaded from

```
https://software.dlir.de/updates/rce/8.x/products/standard/releases/latest/zip/
```

Simply extract the zip file to a location on your file system.

Note

IMPORTANT: Please note that due to restrictions of the standard Windows file system (NTFS), you should choose a destination path that is as short as possible. Long paths are known to cause problems, as some RCE files may not be properly extracted from the zip file.

1.3. Digital Signatures and Download Verification

Starting with RCE 6.x, we provide digital signatures for our releases. These are intended to verify the integrity of the downloaded files, as *any* unsigned software can be tampered with while downloading it over an unsecured (HTTP) connection. This is especially important when installing RCE from a user account with administrator privileges, for example as a Windows service

For each type of release files (.zip, .deb, and .rpm files), a SHA256SUM.asc file is provided next to the actual files. This file contains checksums for each release file, and these checksums are signed with our release key. It is named "RCE 6.x-8.x Automatic Signing Key", has the id 9DC1CE34, and is valid until Nov 25, 2018. Its full fingerprint is 258B C129 EDA2 389D 3ECD 2DE6 BA88 0CB3 9DC1 CE34.

On Windows, you need to install gpg4win (<http://www.gpg4win.org/>) to verify the digital signature. You can use the Get-FileHash cmdlet of the Windows PowerShell to verify the checksum. The cmdlet is available since version 4.0 of the Windows PowerShell. Windows PowerShell 4.0 is already built-in to some Windows versions (Windows Server 2012 R2, Windows 8.1) or can manually be installed (Windows Server 2008 R2, Windows Server 2012, Windows 7). To perform the actual verification:

- Execute


```
gpg --keyserver hkp://pgp.mit.edu:80 --recv-key 258BC129EDA2389D3ECD2DE6BA880CB39DC1CE34
```

to import the signing key. This only needs to be done once per key, e.g. once for all RCE 8.x releases.

Note

If you see a message containing the text "RCE 6.x-8.x Automatic Signing Key <rcenvironment-builds@lists.sourceforge.net>", the key import was successful. Sometimes, this step fails with the message "key [...] not found on keyserver; gpg: no valid OpenPGP data found." If this happens, just repeat the command a few times until it works. This issue may also be caused by an outdated gpg version (gpg version 2.0.9 on SUSE Linux Enterprise Desktop 11 is known to be problematic, while gpg version 2.0.22 works). If you cannot access key servers at all for some reason, you can also download the key manually from

```
https://github.com/rcenvironment/rce-signing/blob/master/rce\_8.x\_signing\_key.asc
```

After the download has completed, you can import the key by executing

```
gpg --import rce_8.x_signing_key.asc
```

from the command line.

- Download the SHA256SUMS.asc file from the same location as the installation package. Run

```
gpg -v --yes SHA256SUMS.asc
```

in the location where you saved it to; this verifies the digital signature. Inspect the output to see if it is correct; you should find the text "Good signature from "RCE 6.x-8.x Automatic Signing Key <rcenvironment-builds@lists.sourceforge.net>" (or a similar translation).

Note

When following these steps, it is normal to receive a warning about the fact that the owner of this key cannot be verified. If you have received this user guide from a trustworthy source (e.g. an official RCE project site secured with HTTPS), you can assume that the key is correct, as the command used to import the key has already verified the key's integrity. Alternatively, you can fetch the key's fingerprint from a trustworthy source (e.g. a secure intranet page, or the official @rcenvironment Twitter feed accessed via HTTPS) and compare it with the one shown in the command's output. If they match, you can trust that you are using the authentic key.

- Adapt the following command to the RCE version you downloaded previously and run

```
Get-FileHash -Algorithm sha256 rce-7.1.4.201609280904-standard-win32.x86_64.zip
```

in the same folder from a PowerShell. The cmdlet will compute the hash for the specified file and print it to the screen. You need to compare the calculated hash with the hash in the SHA256SUMS file. If both are equal, it is verified that the actual download matches what has been digitally signed for the release.

1.4. Starting RCE as a GUI Client

To use RCE with a graphical user interface (GUI), simply start the "rce" executable from Windows Explorer. Optionally, create a desktop icon from the right-click menu using the "send to > desktop" option.

Once your RCE instance has started, you can open the configuration file with the menu option "Help > Open Configuration File". Edit the file, save it, and then restart RCE using the "File > Restart" menu option to apply the changes. There are configuration templates and other information available via the "Help > Open Configuration Information" option. The available configuration settings are described in the configuration chapter.

1.5. Starting a Non-GUI ("Headless") Instance

RCE can also be run from the command line without a graphical user interface (which is called "headless" mode), which uses less system resources and is therefore recommended when the GUI is not needed.

To run a headless RCE instance, open a command prompt and run the command

```
rce --headless -console
```

While RCE is running, you can enter various console commands described in Section 3, "Commands"; note that you need to prefix all RCE commands with "rce " here. To perform a clean shutdown, for example, type `rce stop` and press enter.

1.6. Installation as a Service on a Windows Server

For ad-hoc or temporary RCE network setups, running a headless RCE from the command line is perfectly fine. For more permanent installations, however, we recommend installing RCE as a Windows service instead. This has the advantage that RCE automatically shuts down when the server is shut down, and automatically restarts when the server does.

1.6.1. Installation and Service Management

Executing the following steps will install RCE as service. An RCE service will start automatically on system boot and stop before system shutdown.

1. Navigate to the `extras\windows_service` folder inside your installation folder of RCE named `rce`
2. Open the file `install_as_service.bat` and adjust these settings:
 - Set `RCE_ROOT_PATH` to the location of your RCE installation (typically, the location which you are editing these files in). Note that RCE must already be present (and unpacked) in this location; the daemon installer does not copy any files there.
 - Set `RCE_SERVICE_USER` to the name of the existing user account that the RCE service should run as.

Note

This user account will be used to run RCE and all local tools that are invoked by it. Make sure that this user has appropriate system permissions to run these tools' executables, and read/write all related files.

3. Execute with administrator rights

```
install_as_service.bat
```

by double-click or on the command line

4. Open the Windows Service Management Console (e.g. execute "services.msc" on the command line)
5. Look for entry "RCE Service", do a right-click and click "Properties"
6. Click on the "Log On" tab and fill in the corresponding password for the user account set to the `RCE_SERVICE_USER` variable. By clicking the "Apply" button a message should inform you that the privilege to "Log on as service" was added to the account.

Note

To add the "Log on as a service" right to an account on your local computer manually:

- a. Open Local Security Policy.
 - b. In the console tree, double-click **Local Policies**, and then click **User Rights Assignments**.
 - c. In the details pane, double-click **Log on as a service**.
 - d. Click **Add User or Group**, and then add the appropriate account to the list of accounts that possess the **Log on as a service** right.
-

7. Close the properties dialog by clicking "Ok"

To manually start the RCE service:

1. Open the Windows Service Management Console (e.g. execute "services.msc" on the command line)
2. Look for entry "RCE Service", do a right-click and click "Start"

To stop the RCE service:

1. Open the Windows Service Management Console (e.g. execute `services.msc` on the command line)
2. Look for entry "RCE Service", do a right-click and click "Stop"

To uninstall the RCE service:

1. Navigate to the `extras\windows_service` folder inside your installation folder of RCE named `rce`
2. Open the file `uninstall_service.bat` and set the absolute path of the variable `RCE_ROOT_PATH` to the location of your RCE installation
3. Execute with administrator rights

```
uninstall_service.bat
```

by double-click or on the command line

1.6.2. Service Configuration

After installation, the service instance will be automatically started. This will create a default configuration file if it does not exist yet.

To configure the service instance, locate the RCE configuration file in the service user's home directory - by default, it is located at "C:\users\<user id>\.rce\default\configuration.json". Edit this file, and restart the service to apply the changes.

Note

The need to restart the service is temporary; future versions of RCE will apply configuration changes as soon as the configuration file is changed.

2. Configuration

This section describes the configuration of RCE. Configuration is done within one single configuration file. It is located in the profile directory. From the graphical user interface, you can easily access it from the tool bar or *Help* menu. Note: To apply changes you need to restart RCE. The format of

the configuration file is JSON. See <http://www.json.org/> for the format definition. Also refer to the example configuration files in the installation data directory.

2.1. Configuration Locations and Files

Starting with RCE 6.0.0, all user data is strictly separated from the RCE installation itself. Each set of user data is contained in a so-called "profile". Each profile defines what is called an RCE "instance". Each profile (and therefore, each instance) belongs to exactly one user, and each user can have multiple profiles. The default profile is located within the user's "home" directory ("C:\users\<user id>\\" on Windows), in the ".rce/default" sub-folder.

Note

Note that ".rce" is a hidden directory; you may need to set operating-specific options to see hidden files and directories.

All manual configuration takes place in the profile's central configuration file, "configuration.json". As of RCE 8.0.0, most configuration settings only take effect on startup, so you need to restart RCE after editing it. (This will be changed in a future release.) This applies to all types of installations.

2.2. Configuration Parameters

Configuration parameters are grouped within the configuration file. Below are lists of the configuration parameters. There is one list per group.

Table 2.1. General

<i>Configuration key</i>	<i>Comment</i>	<i>Default value</i>
instanceName	<p>The name of the instance that will be shown to all users in the RCE network. The following placeholders can be used within the instance name:</p> <ul style="list-style-type: none"> • <code>\${hostName}</code> is resolved to the local system's host name. • <code>\${systemUser}</code> is resolved to the Java "user.name" property. • <code>\${profileName}</code> is resolved to the last part of the current profile's file system path. • <code>\${version}</code> is resolved to the build id. <p>Example: "Default instance started by \"<code>\${systemUser}</code>\" on <code>\${hostName}</code>".</p>	"<unnamed instance>"
isWorkflowHost	If set to <i>true</i> , the local instance can be used as a <i>workflow host</i> by other RCE instances. I.e., the workflow controller can be set to this instance and the workflow data is stored there as well.	false
isRelay	If set to <i>true</i> , the local node will merge all connected nodes into a single network, and forward messages between them. This behaviour is transitive; if a relay node connects to another relay node, both networks will effectively merge into one.	false

<i>Configuration key</i>	<i>Comment</i>	<i>Default value</i>
	If set to <i>false</i> (the default value), the local node can connect to multiple networks at once without causing them to merge.	
tempDirectory	Can be used to override the default path where RCE stores temporary files. Useful if there is little space in the default temp file location. Must be an absolute path to an existing directory, and the path must not contain spaces (to prevent problems with tools accessing such directories). The placeholder <code>\${systemUser}</code> can be used for path construction, e.g. <code>"/tmp/custom-temp-directory/\${systemUser}"</code>	An "rce-temp" subdirectory within the user or system temp directory.
enableDeprecatedInputTab	If set to <i>true</i> the tab 'Inputs' is enabled again in the properties view of running workflows. It is disabled by default to due robustness and memory issues. It is recommended to use the 'Workflow Data Browser' to see inputs received and outputs sent.	false

Table 2.2. BackgroundMonitoring

<i>Configuration key</i>	<i>Comment</i>	<i>Default value</i>
enabledIds	Comma-separated list of identifiers referring to certain kind of monitoring data that should be logged continuously in the background. Currently, only 'basic_system_data' is supported.	
intervalSeconds	Logging interval	10

Note

IMPORTANT: When setting up a network of RCE instances, keep in mind that the RCE network traffic is currently *not encrypted*. This means that it is *not secure* to expose RCE server ports to untrusted networks like the internet. When setting up RCE connections between different locations, make sure that they either connect across a secure network (e.g. your institution's internal network), or that the connection is secured by other means, like SSH tunneling or a VPN. Alternatively, you can set up SSH connections in RCE instead of the standard RCE connections.

Table 2.3. Network

<i>Configuration key</i>	<i>Comment</i>	<i>Default value</i>
requestTimeoutMsec	The timeout (in milliseconds) for network requests that are made by the local node. If this time expires before a response is received, the request fails.	40000
forwardingTimeoutMsec	The timeout (in milliseconds) for network requests that are forwarded by the local node on behalf of another node. If this time expires before a response is received, an error response is sent back to the node that made the request.	35000
connections	A map of all connections that the local instance tries to establish on startup. This allows the local instance to act as a client. For each connection a unique identifier (id) must be given.	<code>{}</code> (an empty map in JSON format)
connections/[id]/host	IP address of the host to connect to. Host names and IPv4 addresses are permitted.	-

<i>Configuration key</i>	<i>Comment</i>	<i>Default value</i>
connections/[id]/port	Port number of the remote RCE instance.	-
connections/[id]/connectOnStartup	If set to <i>true</i> , the connection is immediately established on startup.	true
connections/[id]/autoRetryInitialDelay	The initial delay, in seconds, to wait after a failed or broken connection before a reconnect attempt is made. This configuration must be present to enable the auto-reconnect feature.	-
connections/[id]/autoRetryDelayMultiplier	A decimal-point value ≥ 1 that the delay time is multiplied with after each consecutive connection failure. This provides an "exponential backoff" feature that reduces the frequency of connection attempts over time. This configuration must be present to enable the auto-reconnect feature.	-
connections/[id]/autoRetryMaximumDelay	Defines an upper limit for the delay time, even when applying the multiplier would create a higher value. This can be used to maintain a minimum frequency for retrying the connection. This configuration must be present to enable the auto-reconnect feature.	-
serverPorts	A map of all server ports that the local instance registers for other instances to connect to. This allows the local instance to act as a server. For each server port a unique identifier (id) must be given.	{ } (an empty map in JSON format)
serverPorts/[id]/ip	IP address to which the local instance should be bound.	-
serverPorts/[id]/port	Port number to which other instances connect to.	-
ipFilter	Allows to limit the incoming connections to a set of IP addresses.	-
ipFilter/enabled	If set to <i>true</i> , the ip filter active.	false
ipFilter/allowedIPs	List of IP addresses, which are allowed to connect to the instance.	[] (an empty list in JSON format)

Table 2.4. Publishing

<i>Configuration key</i>	<i>Comment</i>	<i>Default value</i>
components	List of workflow components, which are made available to other instances within the RCE network. Workflow components are represented by their identifiers (see table below). Note: That doesn't affect integrated tools.	[] (an empty list in JSON format)

Table 2.5. Component identifiers used in configuration group 'publishing'

Component name	Component identifier
Cluster	de.rcenvironment.cluster
Converger	de.rcenvironment.converger
CPACS VampZero Initializer	de.rcenvironment.vampzeroinitializer
CPACS Writer	de.rcenvironment.cpacswriter
Database	de.rcenvironment.database

Component name	Component identifier
Design of Experiments	de.rcenvironment.doe.v2
Evaluation Memory	de.rcenvironment.evaluationmemory
Excel	de.rcenvironment.excel
Input Provider	de.rcenvironment.inputprovider
Joiner	de.rcenvironment.joiner
Optimizer	de.rcenvironment.optimizer
Output Writer	de.rcenvironment.outputwriter
Parametric Study	de.rcenvironment.parametricstudy
Script	de.rcenvironment.script
Switch	de.rcenvironment.switchcmp
TiGL Viewer	de.rcenvironment.tigviewer
XML Loader	de.rcenvironment.xmlloader
XML Merger	de.rcenvironment.xmlmerger

Table 2.6. ComponentSettings

Configuration key	Comment	Default value
de.rcenvironment.cluster	Configuration of the cluster workflow component.	-
de.rcenvironment.cluster/maxChannels	Maximum number of channels, which are allowed to be opened in parallel to the cluster server.	8

Table 2.7. thirdPartyIntegration

Configuration key	Comment	Default value
tigviewer	Configuration of the external TiGL Viewer application integration. This needs to be configured to enable RCE's TiGL Viewer view and thus, the TiGL Viewer workflow component. Note:TiGL Viewer must be downloaded and installed separately.	-
tigviewer/binaryPath	The path to the TiGL Viewer executable file. Must be an absolute path.	-
tigviewer/startupTimeoutSeconds	The timeout in seconds, to wait for the external TiGL viewer application to start and determine its process id.	10
tigviewer/embedWindow	If set to <i>false</i> , the external TiGL Viewer application Window will not be embeded into RCE's TiGL Viewer view.	true

Table 2.8. sshServer

Configuration key	Comment	Default value
enabled	If set to <i>true</i> the local instance acts as an SSH server.	false
ip (<i>deprecated alias: "host"</i>)	The host's ip address to bind to. If you want to make the SSH server accessible from remote, you should set this to the IP of the machine's external network interface. Alternatively, you can set this to "0.0.0.0" to listen on all available IPv4	127.0.0.1

<i>Configuration key</i>	<i>Comment</i>	<i>Default value</i>
	addresses, if this is appropriate in your network setup.	
port	The port number to which SSH clients can connect to.	-
idleTimeoutSeconds	The time to keep an idle SSH connection alive, in seconds. For typical SSH usage, the default value is usually sufficient. Higher values are, for example, needed when invoking long-running tools using the SSH Remote Access feature.	10
accounts	A map of accounts. For each account a unique identifier (account name) must be given.	{ } (an empty map in JSON format)
[account name]/passwordHash	The hashed password for the account, if password authentication is used. If the SSH account is configured using the configuration UI, the hash is automatically computed and stored here.	-
[account name]/password (<i>deprecated</i>)	The password for the account. SSH passwords can also be configured as plain text, which is however not recommended. To prevent misuse of the configured login data, any configuration file with SSH accounts <i>must</i> be secured against unauthorized reading (e.g. by setting restrictive filesystem permissions). A more secure alternative is to just store the password hash.	-
[account name]/publicKey	The public key for the account, if keyfile authentication is used. Only RSA keys in the OpenSSH format are supported. The public key has to be entered here in the OpenSSH format (a string starting with "ssh-rsa", like it is used for example in authorized_keys files). Only applicable on RCE version 7.1 or newer.	-
[account name]/role	The role of the account. See next table for a list of the possible roles.	-
[account name]/enabled	If set to <i>true</i> , the account is enabled.	true

Table 2.9. Possible roles for SSH accounts

<i>Role name</i>	<i>Allowed commands</i>
remote_access_user (Standard role for using SSH remote access tools and workflows)	ra sysmon
remote_access_compatibility (backwards compatibility alias for remote_access_user)	ra sysmon
remote_access_admin	ra ra-admin sysmon components
workflow_observer	components net info sysmon wf list wf details
workflow_admin	components net info sysmon wf
local_admin	cn components mail net restart shutdown stop stats tasks
instance_management_admin	im net info
instance_management_delegate_user	cn components net restart shutdown stop stats tasks wf ra-admin

<i>Role name</i>	<i>Allowed commands</i>
developer	<all>

Table 2.10. sshRemoteAccess

<i>Configuration key</i>	<i>Comment</i>	<i>Default value</i>
sshConnections	A map of SSH connections. This allows the local instance to act as a SSH remote access client. For each connection a unique identifier (id) must be given.	{ } (an empty map in JSON format)
sshConnections/[id]/displayName	The name for the connection that will be shown in the network view.	-
sshConnections/[id]/host	The remote RCE instance to connect to. Host names and IPv4 addresses are permitted.	-
sshConnections/[id]/port	Port number of the remote RCE instance.	-
sshConnections/[id]/loginName	The login name for authentication.	-
sshConnections/[id]/keyfileLocation	Path to the private key file, if keyfile authentication is used. Only RSA keys in the OpenSSH format are supported.	-
sshConnections/[id]/noPassphrase	This option should only be set if a private key that requires no passphrase is used for authentication. If set to <i>true</i> , RCE does not ask for a passphrase before connecting.	false

Table 2.11. smtpServer

<i>Configuration key</i>	<i>Comment</i>	<i>Default value</i>
host	The IP address or hostname of the SMTP server, which should be used for mail delivery.	-
port	Port number of the SMTP server.	-
encryption	Can either be "explicit" or "implicit". Select "implicit" if you want to connect to the SMTP server using SSL/TLS. Select "explicit" if you want to connect to the SMTP server using STARTTLS. Unencrypted connections are not permitted.	-
username	The login name for authentication.	-
password	The obfuscated password for authentication. Plaintext password cannot be used here. To create the obfuscated password from the plaintext password, you need to use the Configuration UI described in Section 2.3, "Configuration UI"	-
sender	Email address, which should be displayed as the sender in the sent email.	-


Note

The used SMTP server needs to be configured using the Configuration UI described in Section 2.3.2, "Mail: SMTP server configuration", since the password needs to be obfuscated.

2.3. Configuration UI

If you want to configure SSH accounts with passphrases or you want to configure e-mail support for the instance, you need to use the Configuration UI. You can access the interactive tool by executing RCE from the command line with the option "rce --configure" or by using the "Launch Configuration UI" script in the "extras" folder of your RCE installation directory.

Figure 2.1. Configuration tool for SSH account and SMTP server configuration



2.3.1. Remote Access: SSH account configuration

If the RCE instance shall act as a SSH server, you can configure SSH accounts using the Configuration UI, which encrypts the SSH passwords before storing them in the configuration file.

Note

All SSH accounts configured with this tool initially have the role "remote_access_user", which allows to execute commands needed for remote access on tools and workflows. If you want to change the role of an SSH account, you can do this by editing the configuration file manually (see table "Possible roles for SSH accounts").

2.3.2. Mail: SMTP server configuration

If you want to send e-mails from a RCE instance, you need to configure an SMTP server. RCE does not send e-mails directly to the recipient, but instead sends the e-mails to an SMTP server, which delivers them to the recipient. You need to use the Configuration UI to configure such an SMTP server, since

the password used for authentication needs to be obfuscated before it is stored in the configuration file. The SMTP server parameters that need to be configured are described in more detail in Table 2.11, “smtpServer”

Note

Due to a known bug on Windows system with a German keyboard layout, the Configuration UI inserts the characters "q@" into a text field if you want to insert the @ sign. You can manually remove the additional character "q".

Chapter 3. Usage


This chapter describes the main usage concepts.

1. Graphical User Interface

This section introduces the Graphical User Interface (GUI).

The GUI of RCE is composed of different views and editors (besides standard GUI elements such as the menu bar, status bar, etc.). Views can be (re-)arranged by the user. They can even be closed and opened again. Some views are closed by default, but can be opened as desired. To open a view go to: Window->Show view.

Figure 3.1. Workbench with different views and the workflow editor opened




Left hand side:

- *Project Explorer:* View to manage projects. All relevant data including workflow files needs to be organized in projects.
- *Workflow List:* Lists all active workflows and allows to manage them (stop, pause, resume, dispose).


Right hand side and center:

- *Workflow Editor:* Core view of RCE used to build and configure workflows.
- *Palette:* Lists all available workflow components. If RCE runs in a distributed environment this includes local as well as remote workflow components. At the top, it also provides actions for connecting workflow components. The appearing connection editor dialog will look like this. Additionally, connections of the workflow are shown in the *Properties* view at the bottom, if the background of the workflow editor is selected.

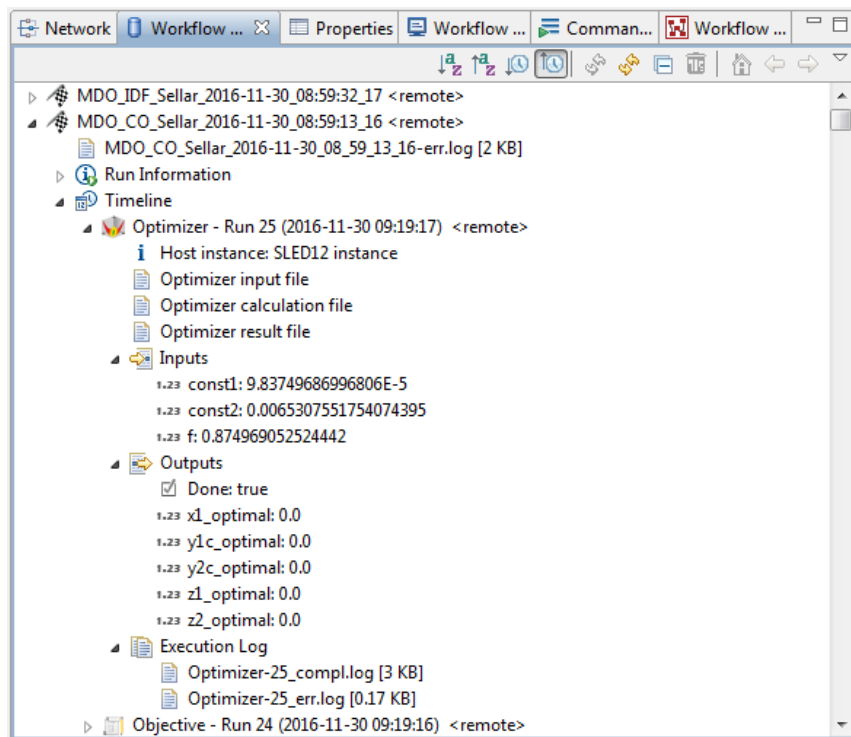
Figure 3.2. Connection Editor

Bottom:


- *Log*: Shows all log output of RCE, e.g. error messages during workflow execution.
- *Network View*: Shows all RCE nodes of the distributed RCE network and their published workflow components. It also shows the outgoing connections of the own RCE instance and allows to manage them (start, stop, etc.). Furthermore, you are able to see monitoring data like CPU or RAM usage for each instance.

Figure 3.3. Network View

- *Workflow Data Browser*: Shows workflow related result data.

Figure 3.4. Workflow Data Browser

- *Properties*: Allows configuration of workflow components (e.g. Inputs/Outputs) if they are selected in the workflow editor. View adapts to selected workflow component.
- *Workflow Console*: Shows all native console line output of integrated tools during workflow execution. Provides full text search.

Figure 3.5. Workflow Console

2. Workflows

This section describes the basics of workflows in RCE.

2.1. Rationale

RCE is designed to execute automated, distributed workflows. Workflows consist of so called workflow components which can be coupled among each other. Loops are supported, even multi-nested ones.

2.2. Getting Started

To get started with workflows in RCE it is recommended to both read the following sections about workflows and walk through the example workflows provided in RCE. The sections here refer to the workflow examples where it is useful and vice versa.

Workflows in RCE are encapsulated in so called projects. To create the workflow examples project go to: File->New->Workflow Examples Project. A dialog appears. Leave the default project name or give it a name of your choice and confirm by clicking 'Finish'. In the 'Project Explorer' on the left-hand side, the newly created project is shown. The example workflows are grouped in sub folders. It is recommended to walk through the workflows following the prefix starting with '01_01_Hello_World.wf'.

2.3. Workflow Components

Workflow components are either tools that are integrated by users or are provided by RCE supplying multi-purpose functionality. The following list shows workflow components provided by RCE grouped by purpose (workflow components that are deprecated (i.e., they are removed soon) or that are not recommended to use anymore are left out):

- *Data*:Database
- *Data Flow*:Input Provider, Output Writer, Joiner, Switch
- *Evaluation*:Optimizer, Design of Experiments, Parametric Study, Converger, Evaluation Memory
- *Execution*:Script, Cluster, Excel
- *XML*:XML Loader, XML Merger, XML Values
- *CPACS*:TiGL Viewer, VAMPzero Initializer

The example workflows in sub folder '02_Component Groups' introduce some of the workflow components provided. Additionally, there is a documentation for each workflow component available in RCE. To access it, you can either rightclick on a component in a workflow and select "Open Help" or press "F1". A help view opens on the right-hand side. Moreover there is an entry "Help Contents" in the "Help" menu where you can navigate to the component help you require.

The XML and CPACS components are able to read or extract data from an Xml-file via dynamic in- or outputs. The XPathChooser is a feature that provides help selecting the item, which shall be read or extracted. Add an in- or output and press the "XPath choosing..." button to open a window, where you can select the Xml-file, which contains the item that shall be selected. After choosing the file, the XPathChooser opens containing a tree, symbolizing the Xml-File. By selecting an element, the text below is updated and displays the current path. The last two columns are used to choose attributes. The attribute name can be selected via the column "Attributes". In the column "Values", the proper value can be selected. Use a double-click on an element to expand or fold the tree. The chosen XPath will be written in the text field of the window, in which the XPathChooser has been opened originally. Using this text field, new paths can be created. Add a slash and the name of the node that shall be created to the existing path. The new path will be added during the workflow run.

New XPaths can only be generated within the inputs tab. Using the outputs tab will cause an error.

2.4. Coupling Workflow Components

A workflow component can send data to other workflow components. Therefore, a so called connection needs to be created between the sending workflow component and the receiving one. For that purpose, workflow components can have so called inputs and outputs. A connection is always created between an output and an input. You can think of a connection as a directed data channel. Data is sent as atomic packages which are not related to each other (there is no data streaming between workflow components). Supported data types are:

Primitive data types:

- *Short Text*: A short text (up to 140 characters)
- *Integer*: Integer number
- *Float*: Floating point number
- *Boolean*: Boolean value (true or false)

Referenced data types (The actual data is stored in RCE's data management and only a reference is transferred):

- *File*: File
- *Directory*: Directory

Other data types

- *Small Table*: Table restricted to values of type Short Text, Integer, Float, Boolean (primitive data types) (up to 100.000 cells)
- *Vector*: one-dimensional "Small Table" (one column) restricted to values of type Float
- *Matrix*: Small Table restricted to values of type Float

Not all of the workflow components support all of the data types listed. A connection can be created between an output and an input if:

- The data type of the output is the same as or convertible to the data type of the input.
- The input is not already connected to another output.

Note that data from an output can be sent to multiple inputs, but an input can just receive data from a single output.

The following table shows which data types are convertible to which other types:

Table 3.1. Data Type Conversion Table

To From	Boolean	Integer	Float	Vector	Matrix	Small Table	Short Text	File	Directory
Boolean		x	x	x	x	x			
Integer			x	x	x	x			
Float				x	x	x			
Vector					x	x			
Matrix						x			
Small Table									
Short Text						x			

File									
Directory									

2.5. Execution Scheduling of Workflow Components

The execution of workflows is data-driven. As soon as all of the desired input data is available, a workflow component will be executed. Which input data is desired, is defined by the component developer (for RCE's default workflow components), the tool integrator, and/or the workflow creator. The workflow component developer and tool integrator decide which options are allowed for a particular workflow component. The workflow creator can choose between those options at workflow design time. The following options exist:

Input handling:

- *Constant*: The value won't be consumed during execution and will be reused in the next iteration (if there is any loop in the workflow). The workflow will fail if there is more than one value received, except for nested loops: All inputs of type *Constant* are resetted within nested loops, after the nested loop has been finished.
- *Single (Consumed)*: The input value will be consumed during execution and won't be reused in the next iteration (if there is any loop in the workflow). Queuing of input values is not allowed. If another value is received before the current one was consumed, the workflow will fail. This can guard against workflow design errors. E.g., an optimizer must not receive more than one value at one single input within one iteration.
- *Queue (Consumed)*: The input value will be consumed during execution and won't be reused in the next iteration (if there is any loop in the workflow). Queuing of input values is allowed.

Execution constraint:

- *Required*: The input value is required for execution. Thus, the input must be connected to an output.
- *Required if connected*: The input value is not required for execution (e.g., if a default value will be used as fall back within the component). Thus, the input doesn't need to be connected to an output. But if it is connected to an output, it will be handled as an input of type *Required*.
- *Not required*: The input value is not required for execution. Thus, the input doesn't need to be connected to an output. If it is connected to an output, the input value will be passed to the component if there is a value available at the time of execution. Values at inputs of type *Not required* cannot trigger component execution except if it is the only input defined for a component. Note: With this option, non-deterministic workflows can be easily created. Use this option carefully. If in doubt, leave it out.

Note: With RCE 6.0.0 the scheduling options changed. Below is the migration path:

- *Initial* was migrated to *Constant* and *Required*.
- *Required* was migrated to *Single (Consumed)* and *Required*.
- *Optional* was migrated to *Single (Consumed)* and *Required if connected*.

If you encounter any problems with workflows created before RCE 6.0.0, it is likely, that it affects the migration to *Single (Consumed)* instead of to *Queue (Consumed)*. We decided to migrate conservatively to not hide any existing workflow design errors. So, if queuing of input values is allowed for an input, just change the input handling option to *Queue (Consumed)* after the workflow was updated. Another issue can affect the migration of *Optional*. If it affects an input of the script component, check the option, which let the script component execute on each new value at any of its inputs. Also check *Not required* as an alternative execution constraint option.

2.6. (Nested) Loops

Workflow components can be coupled to loops. A loop must always contain a so called driver workflow component. Driver workflow components (group "Evaluation") are: Optimizer, Design of Experiments, Parametric Study, Converger (see the example workflow "02_02_Evaluation_Drivers"). The responsibilities of a driver workflow component in a loop are:

- Send values to the loop and receive the result values.
- Finish the loop based on some certain criteria.

If a loop contains another loop, we speak of the latter as a nested loop. A nested loop can contain again another loop and so on. To create workflows with nested loops (see example workflows in "03_Workflow_Logic"), some certain concepts behind nested loops must be understood:

- Loop level: If a loop contains another loop, that loop is considered as a nested loop with a lower loop level. From the perspective of the nested loop, the other loop is considered as a loop with an upper loop level.
- If a driver workflow component is part of a nested loop, you need to check the checkbox in the "Nested Loop" configuration tab
- Data exchange between loops of different loop levels is only allowed via a driver workflow component. Thereby, only particular inputs and outputs of driver workflow components are allowed to be connected to inputs and outputs of the next upper loop level and particular ones to inputs and outputs of the same loop level. For example, if a 'same loop level' output is connected to a loop with an upper loop level, the workflow won't succeed or might even get stuck. Below you find tables of inputs and outputs for each driver workflow component and whether they must be connected to the same loop level or to the next upper loop level. (Note: In the inputs and outputs tables of driver workflow components (in 'Inputs/Outputs' properties tab), the loop level requirement is present in a particular column for each input and output.)

Table 3.2. Inputs of Optimizer

Input	Loop Level
* - lower bounds - start value	To next upper loop level
* - upper bounds - start value	To next upper loop level
* - start value	To next upper loop level
* (Objective functions)	To same loop level
* (Constraints)	To same loop level
d*.d* (Gradients)	To same loop level

Table 3.3. Outputs of Optimizer

Output	Loop Level
*_optimal	To next upper loop level
Done	To next upper loop level
* (Design variables)	To same loop level
Gradient request	To same loop level
Iteration	To same loop level

Table 3.4. Inputs of Design of Experiments

Input	Loop Level
*_start	To next upper loop level
*	To same loop level

Table 3.5. Outputs of Design of Experiments

Output	Loop Level
Done	To same loop level
*	To same loop level

Table 3.6. Inputs of Parametric Study

Input	Loop Level
*_start	To next upper loop level
*	To same loop level

Table 3.7. Outputs of Parametric Study

Output	Loop Level
Done	To same loop level
*	To same loop level

Table 3.8. Inputs of Converger

Input	Loop Level
*_start	To next upper loop level
*	To same loop level

Table 3.9. Outputs of Converger

Output	Loop Level
Converged	To next upper loop level
Converged absolute	To next upper loop level
Converged relative	To next upper loop level
*_converged	To next upper loop level
Done	To same loop level
*	To same loop level

2.7. Fault-tolerant Loops

Workflow components of a loop can fail. There are two kind of failures:

- A workflow component fails gracefully, i.e. it couldn't compute any results for the inputs received but works normally. In this case, it sends a value of type "not-a-value" with the specified cause to its outputs which finally are received by the driver workflow components as results.
- A workflow component fails, i.e. it crashes for an unexpected reason. In this case, the workflow engine sends values of type "not-a-value" with the specified cause as results to the driver workflow component.

In the "Fault Tolerance" configuration tab of workflow driver components, it can be configured how to handle failures in loops, for both kind of failures separately.

2.8. Manual Tool Result Verification

After the execution of an integrated tool, the results are sent via outputs to the next workflow component (e.g. to the next integrated tool). By default, this is done in an automated manner without any user interaction. If the data should be verified by a person responsible for the tool before they are sent further, manual verification of tool results must be enabled in the tool integration wizard in the 'Verification' tab of the 'Inputs and Outputs' page.

In case manual verification of tool results is enabled, the results are hold after each tool execution and the corresponding workflow component remains in state "Waiting for approval". Then, there are two options:

- Approve tool results: The tool results are sent via the outputs to the next workflow component and the workflow continues normally.
- Reject tool results: The tool results are not sent via the outputs to the next workflow component and the workflow is cancelled.

To apply one of the options, a so called verification key is required. The verification key is generated by RCE after each tool execution and is written to a file on the file system of the machine which executed the tool. (The location is specified in the 'Verification' tab of the 'Inputs and Outputs' page in the tool integration wizard.) Optionally, the verification key can also be sent via e-mail if e-mail support is configured for the RCE instance where the tool is integrated. (E-mail support can only be configured using the Configuration UI as described in Section 2.3.2, "Mail: SMTP server configuration") E-mail delivery can be enabled and the recipients can be defined in the 'Verification' tab of the 'Inputs and Outputs' page in the tool integration wizard.

Once the verification key is known (either from the file or an e-mail), perform following steps to approve or reject the tools results:

- Start an RCE instance with a graphical user interface. (Your tool must be available, i.e. it must appear in the palette of the workflow editor.)
- In the menu bar at the top, go to Run -> Verify tool results...
- A dialog appears that guides you through the verification process.

3. Commands

This section introduces the list of commands available for the command line and the interactive shell.

3.1. Command Line Parameters

General syntax

```
> rce --[RCE arguments] -[RCP arguments] -[VM arguments]
```


Table 3.10. Command line arguments for RCE

Argument	Type	Description
profile "<profile id or path>"	RCE	Sets a custom profile folder to use. If only an id (any valid directory name) is given, the profile directory "<user home>/.rce/id" is used. Alternatively, a full filesystem path can be specified.
profile	RCE	If the profile argument is specified without a profile id or path, RCE launches the Profile Selection UI, which allows to select a profile folder for the startup as described in Section 3.2, "Profile Selection UI".
batch "<command string>"	RCE	Behaves like the "exec" command, but also implies the "--headless" option and always shuts down RCE after execution.
headless	RCE	Starts RCE in a headless modus without GUI. It will remain in the OSGi console and waits for user input.
exec "<command string>"	RCE	<p>Executes one or more shell commands defined by <command string>. For the list of available commands, refer to the command shell documentation. This argument is usually used together with --headless to run RCE in batch mode. Multiple commands can be chained within <command string> by separating them with ";" (note the spaces); each command is completed before the next is started.</p> <p>You can use the "stop" command at the end of the command sequence to shut down RCE after the other commands have been executed. However, any error during execution of these commands will cancel the sequence, and prevent the "stop" command from being executed. To ensure shut down at the end of the command sequence, use the --batch option instead of "--exec".</p> <p>As an example, <code>rce --headless --exec "wf run example.wf ; stop"</code> will execute the "example.wf" workflow in headless mode and then shut down RCE. However, if the workflow fails to start, RCE will keep running, as the "stop" command is never executed. To attempt execution of the workflow file, but then always shut down regardless of the outcome, use <code>rce --batch "wf run example.wf"</code> instead.</p>
configure	RCE	Starts the RCE Configuration UI (Section 2.3, "Configuration UI") which can be used to configure SSH accounts with passphrases or to configure e-mail support for the RCE instance.
data @noDefault	RCP	Set the default workspace location to empty
consoleLog	RCP	Logs everything for log files on the console as well.
console	RCP	Runs RCE with an additional OSGi console window, which allows you to execute RCE shell commands. See the Command Shell documentation for more information.

Argument	Type	Description
<i>Deprecated:</i> console <port>	RCP	Specify the port that will be used to listen for telnet connections. (<i>NOTE:</i> this access is insecure; configure SSH access instead)
clean	RCP	Cleans before startup
vmargs	VM	Standard JVM arguments
Dde.rcenvironment.rce.configuration.dir=<insert-config-path>	VM	Sets the configuration directory
Drce.network.overrideNodeId =<some-id>	VM	Sets the local node id, overriding any stored value. This is mostly used for automated testing. Example: "-Drce.network.overrideNodeId=a96db8fa762d59f2d2782f3e5e9662d4"
Dcommunication.uploadBlockSize=<block size in bytes>	VM	Sets the block size to use when uploading data to a remote node. This is useful for very slow connections (less than about 10 kb/s) to avoid timeouts. The default value is 262144 (256 kb). Example: "-Dcommunication.uploadBlockSize=131072" - sets the upload block size to 128kb (half the normal size)

3.2. Profile Selection UI

During startup of the instance, the Profile Selection UI allows to select a profile folder which should be used for the current run of RCE. Furthermore it allows to specify a default profile for future runs. You can access the Profile Selection UI by executing RCE from the command line with the option "rce --profile".

Figure 3.6. Profile Selection UI

If the first option "Select a profile and start RCE" is chosen, a list of available profiles is presented. On selection of one of these profiles, RCE is started using this profile.

If the second option "Select the default profile for future runs" is chosen, a list of available profiles is presented. On selection of one of these profiles, RCE will not be started using this profile, but instead the selected profile will be marked as the default profile for future runs. This selection can be temporarily overwritten again by using the '-profile "<profile id or path>"' option. The default profile setting will be stored for the current user and the current installation location of RCE. Different users on the same machine can therefore configure different default profiles. Furthermore, different installations of RCE can have different default profiles configured.

Note

The Profile Selection UI will only display profiles if they have been started once with RCE 7.0 or newer.

3.3. Command Shell

RCE provides an integrated shell (sometimes referred to as "console") for executing commands. It can be accessed in three different ways:

- Start RCE with the "-console" command-line option, or add "-console" to the rce.ini file before starting; this will open an OSGi console window. Due to the nature of an OSGi console, all RCE commands must be prefixed with "rce". For example, type "rce help" to show the available commands.
- *Deprecated:* Start RCE with the "-console <port>" command-line option; this will accept telnet OSGi console sessions on that port. As with the "-console" option, RCE commands must be prefixed with "rce" (for example, type "rce help").

Note that this option is *insecure*, as there is no authentication nor encryption, so it should only be used in fully trusted networks. Whenever possible, use the SSH console (see below) instead.

- Configure SSH access. To do so, refer to Section Configuration Parameters. After RCE has started, you can access the shell on the configured port with a standard SSH client. On Windows systems, the "putty" software works well as a client.

As this option creates a pure RCE shell (as opposed to the OSGi consoles created above), you can enter RCE commands without a prefix - for example, just type "help" to list the available commands. Note that to avoid confusion, adding a "rce" prefix still works, but it is not necessary.

The following table lists some shell commands; more documentation coming soon.

Table 3.11. Shell Commands

Argument	Description
help	Lists all available commands.
components	Short form of "components list".
components list	Lists components published by reachable RCE nodes.
cn	Short form of "cn list".
cn add <target> ["<description>"]	Adds a new network connection. (Example: cn add activemq-tcp:rceserver.example.com:20001 "Our RCE Server")
cn list	Lists all network connections, including ids and connection states.
cn start <id>	Starts/Connects a READY or DISCONNECTED connection (use "cn list" to get the id).
cn stop <id>	Stops/Disconnects an ESTABLISHED connection (use "cn list" to get the id).
mail <recipient> <subject> <body>	Sends an email to the specified recipient.
net	Short form of "net info".
net filter	Shows the status of the IP whitelist filter.
net filter reload	Reloads the IP whitelist configuration.
net info	Lists all reachable RCE nodes.
restart	Restarts RCE.
shutdown	Shuts down the local RCE instance.
stop	Shuts down the local RCE instance (alias of "shutdown").
version	Shows version information.
wf	Short form of "wf list"
wf list	Lists all current workflows, their states and execution ids.
wf details <workflow execution id>	Shows details about one workflow.
wf cancel <workflow execution id>	Cancels a running or paused workflow.
wf delete <workflow execution id>	Deletes a finished, cancelled or failed workflow from the data management and disposes it.
wf pause <workflow execution id>	Pause a running workflow.
wf resume <workflow execution id>	Resume a paused workflow.
wf run [--delete <unfinished always never>] [--compact-	Executes the given workflow file and waits until it has completed. Workflow file paths containing spaces must be enclosed in double quotes ("...").

Argument	Description
output] [-p <placeholder value file>] <workflow file>	<p>The "--delete" option defines the deletion behavior after workflow completion. Deleting a workflow deletes all files in the data management of it and releases certain resources that may or may not be used after it has finished, for example data to be visualized in component's runtime views. The default of this setting is "onfinished": The workflow is deleted if it terminates in state "Finished" (which means normal completion without errors), otherwise it is left unchanged for inspection.</p> <p>The "--compact-output" option reduces this command's output as much as possible, which is intended to simplify scripted calls of this command. The first line printed will either be the workflow's assigned id if the start was successful, or a text starting with "Error " if the workflow could not be started. If (and only if) the start was successful, a second line will be printed once the workflow has terminated. The pattern of this second line is "<workflow id>: <final state>".</p> <p>The "-p" option can be used to define a placeholder value file (see below).</p>
wf verify [-- delete <onfinished always never>] [--pr <parallel runs>] [-- sr <serial runs>] [- p <placeholder value file>] --basedir <directory> <workflow file> [<workflow file> ...]	<p>Runs several workflows and creates a summary of which ones failed and succeeded.</p> <p>The "--pr" option defines how often the workflow is started in parallel. The "--sr" options defines how often the workflow is started in serial. E.g. "--pr 5 --sr 3" starts the workflow three times with five in parallel. If "*" is used with the "--basedir" option or multiple workflow filenames are passed, "--pr" and "--sr" are applied for each of the workflows.</p> <p>For the "--delete" and "-p" options refer to "wf run" above.</p> <p>The "--basedir <directory>" parameter specifies the directory containing the workflow files. File paths containing spaces must be enclosed in double quotes ("...").</p> <p>The second parameter defines the workflow's filenames. Using "*" as workflow file runs all non-backup workflows in the basedir. Workflow file paths containing spaces must be enclosed in double quotes ("...").</p>

3.3.1. Configuration Placeholder Value Files

Some workflow components use placeholders for configuration values. The values for the placeholders are defined at workflow start. When executing workflows from the command line (e.g. in headless or batch mode), the placeholder's values must be defined in a file, which will be passed to the command with the -p option. Placeholder value files have following format:

```
{
<component id>/<component version> : {
<configuration placeholder id> : <configuration value>
}
<component id>/<component version>/<component instance name> : {
<configuration placeholder id> : <configuration value>
}
}
```

Note

Every id and every value must be enclosed in double quotes ("...").

The component `id` is the id string of a component (e.g. `de.rcenvironment.script`), the component `version` is the version of the component that is used in the workflow (e.g. 3.4).

There are two ways of defining values for configuration placeholders: per component type and per component instance. When defined per component type, the id and version must be specified (e.g. `"de.rcenvironment.script/3.4"`). When defined per component instance the component id, component version, and the name of the component in the workflow must be specified (e.g. `"de.rcenvironment.script/3.4/MyScript"`). In both cases, the configuration placeholder id, which is the name of the configuration placeholder, and the actual configuration value must be specified.

Component instance values override component type values.

Note

It is possible to mix component type and component instance values.

Below is an example placeholder value file, which defines one placeholder value (component type) for the script component, one placeholder value (component type) for the input provider component and a placeholder value (component instance) for a specified input provider component of the workflow:

```
{
  "de.rcenvironment.script/3.4": {
    "pythonExecutionPath": "C:/Python/python.exe"
  },
  "de.rcenvironment.inputprovider/3.2": {
    "inputFile": "C:/input/globalInputFile.txt"
  },
  "de.rcenvironment.inputprovider/3.2/Provider 1" : {
    "inputFile": "C:/input/Provider1.txt"
  }
}
```

The following table lists components and their configuration placeholders.

Table 3.12. Components and their configuration placeholders

Component	Component id and version	Configuration placeholders
Cluster	de.rcenvironment.cluster/3.2	authuser - user name authphrase - password (base64 encoded)
Input Provider	de.rcenvironment.inputprovider/3.2	<output name> - value of output

Component	Component id and version	Configuration placeholders
Output Writer	de.rcenvironment.outputwriter/2.0	targetRootFolder - path to target root folder
Script	de.rcenvironment.script/3.4	pythonExecutionPath - path to the Python executable (only required if Python is set as script language)

4. Integration of External Tools

4.1. Basic Concepts

The Tool Integration concept of RCE is used to integrate external tools for calculations, simulations and so on into RCE and use them in a workflow as a component. The tools must fulfill these requirements:

- The external tool must be callable via command line
- It must have a non-interactive mode which is called via command line
- Input for the tool must be provided through command line arguments or files

If these requirements are fulfilled, a configuration file can be created that is used for the integration.

If you use RCE with a graphical user interface this can be done with the help of a wizard which guides you through the settings. This wizard can be found in the menu *File -> Integrate Tool....* Required fields are marked with an asterisk (*). When the wizard is finished and everything is correct, the integrated tool will automatically show up in the Workflow Editor palette.


Note

The wizard has a dynamic help, which is shown by clicking on the question mark on the bottom left or by pressing F1. It will guide you through every page of the wizard.


4.2. Directory Structure for Integrated Tools

When executing an integrated tool, a certain directory structure is created in the chosen working directory. This structure depends on the options you have chosen in the integration wizard. The two options that matter are "Use a new working directory each run" and "Tool copying behavior".


„Use a new working directory on each run“ **not** selected
 „Do not copy tool“ selected




„Use a new working directory on each run“ **not** selected
 „Copy tool to working directory once“ selected



„Use a new working directory on each run“ selected
 „Copy tool to working directory once“ selected



„Use a new working directory on each run“ selected
 „Copy tool to working directory on each run“ selected



Root Working Directory: This is the directory you choose in the "Tool Integration Wizard" as "Working Directory" on the "Launch Settings" page.

Config Directory: In this directory, the configuration file that may be created by the tool integration will be created by default. The configuration files can be created from the properties that are defined for the tool on the "Tool Properties" page.

Input Directory: All inputs of type "File" and "Directory" will be copied here. They will have a subdirectory that is named exactly as the input's name (e.g. the input "x" of type "File" will be put into "Input Directory/x/filename").

Output Directory: All outputs of type "File" and "Directory" can be written into this directory. After that, you can use the placeholder for this directory to assign these outputs to RCE outputs in the post execution script. To write the output directory into an output "x" of type "Directory" the following line in the post execution script would be required: `${out:x} = "${dir:output}"`

Tool Directory: This is the directory, where the actual tool is located. If the tool should not be copied, it will be exactly the same directory that you choose, otherwise it will be the same as the chosen directory but copied to the working directory.

Working Directory: A working directory is always the location, where all the other directories will be created. If the option "Use a new working directory on each run" is disabled, this will always be the same as the "Root Working Directory". Otherwise, a new directory is created each run (the name will be the run number) and is the working directory for the run.

4.3. Copying of Integrated Tools

When a component is created in the integration wizard, a configuration file is created.

All configuration files from the tool integration are stored in the directory

```
<profile folder>/integration/tools/
```

In this directory, there is a separation between different kinds of integration realized through one subdirectory for each. The `common` folder is always existent.

In these subdirectory, the integrated tools are stored, again separated through a subdirectory for each. The name of the directory equals the name of integration of the tool.

If an integrated tool should be copied to another RCE instance or another machine, the directory of the tool must be copied, containing a `configuration.json` and some optional files. It must be put in the equivalent integration type directory of the target RCE instance. After that, RCE automatically reads the new folder and if everything is valid, the tool will be integrated right away.

Note

If you want to delete a tool folder that contains some documentation, this can cause an error. If you have this problem, delete the documentation folder at first (it must be empty), afterwards you can delete the tool folder.

4.3.1. Publishing a Component on a Server Instance in Headless Mode

An integrated tool on a server is not necessarily visible for other users in the same network. Publishing means to make it available for all others, but they are not allowed to change the tool's integration. If a component should be published, the `published.conf` file in the integration kind subdirectory must be edited. In this file, the tool's directory name (or absolute path) must be written in a single line. After saving the file, RCE reads the file automatically and the component will be published in the network.

Note

Using a graphical editor to edit the `published.conf` or a tool configuration file **might** not work as expected. It is possible, that the tool or the publishing status is not updated when saving the tool. This problem did not occur with headless editors.

4.3.2. Tool Execution Return Codes

The tools are executed by using a command line call on the operating system via the "Execution Script". When the tool finished executing (with or without error), its exit code is handed back to the execution script and can be analyzed in this script. If in the script nothing else is done, the exit code is handed back to RCE. When there is an exit code that is not "0", RCE assumes that the tool crashed and thus lets the component crash without executing the "Post Script". Using the option "Exit codes other than 0 is not an error" can prevent the component to crash immediately. With this option enabled, the post script will be executed in any way and the exit code from the tool execution can be read by using the placeholder from "Additional Properties". In this case, the post script can run any post processing and either not fail the component, so the workflow runs as normal, or let the component crash after some debugging information was written using the Script API `"RCE.fail('reason')"`.

4.4. Integration of CPACS Tools

4.4.1. Additional concepts of CPACS Tool Integration

Extending the common Tool Integration concept, the CPACS Tool Integration has some additional features.

- **Parameter Input Mapping (optional):** Substitutes single values in the incoming CPACS content, based on an XPath configured at workflow design time as a dynamic input of the component
- **Input Mapping:** Generates the tool input XML file as a subset of the incoming CPACS file XML structure, specified by a mapping file
- **Tool Specific Input Mapping (optional):** Adds tool specific data to the tool input file, based on a mapping file and a data XML file
- **Output Mapping:** Merges the content of the tool output XML file into the origin incoming CPACS file, based on a mapping file
- **Parameter Output Mapping (optional):** Generates output values as single values of the CPACS result file, based on an XPath configured at workflow design time as a dynamic output of the component
- **Execution option to only run on changed input:** If enabled, the integrated tool will only run on changed input. Therefore the content of the generated tool input file is compared to the last runs content. Additionally the data of the static input channels are compared to the previous ones.

All the features listed above can be configured in the tool integration wizard on the dedicated *CPACS Tool Properties* page.

The mappings can be specified by XML or XSLT as shown in the following examples. RCE differentiates between these methods in accordance to the corresponding file extension (.xml or .xsl).

Example for an input or tools pecific XML mapping :

```
<?xml version="1.0" encoding="UTF-8"?>
<map:mappings xmlns:map="http://www.rcenvironment.de/2015/mapping" xmlns:xsl="http://www.w3.org/1999/XSL/Transform">

  <map:mapping>
    <map:source>/cpacs/vehicles/aircraft/model/reference/area</map:source>
    <map:target>/toolInput/data/var1</map:target>
  </map:mapping>

</map:mappings>
```

And input or tool specific XSLT mapping:

```
<?xml version="1.0" encoding="utf-8"?>
<xsl:stylesheet version="1.0" xmlns:xsl="http://www.w3.org/1999/XSL/Transform" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:noNamespaceSchemaLocation="cpacs_schema.xsd">
  <xsl:output method="xml" media-type="text/xml" />
  <xsl:template match="/">
    <toolInput>
      <data>
        <var1>
          <xsl:value-of select="/cpacs/vehicles/aircraft/model/reference/area" />
        </var1>
      </data>
    </toolInput>
  </xsl:template>
</xsl:stylesheet>
```

Example of an output XML mapping:

```
<?xml version="1.0" encoding="UTF-8"?>
```

```
<map:mappings xmlns:map="http://www.rcenvironment.de/2015/mapping" xmlns:xsl="http://www.w3.org/1999/XSL/Transform">

  <map:mapping mode="delete">
    <map:source>/toolOutput/data/result1</map:source>
    <map:target>/cpacs/vehicles/aircraft/model/reference/area</map:target>
  </map:mapping>

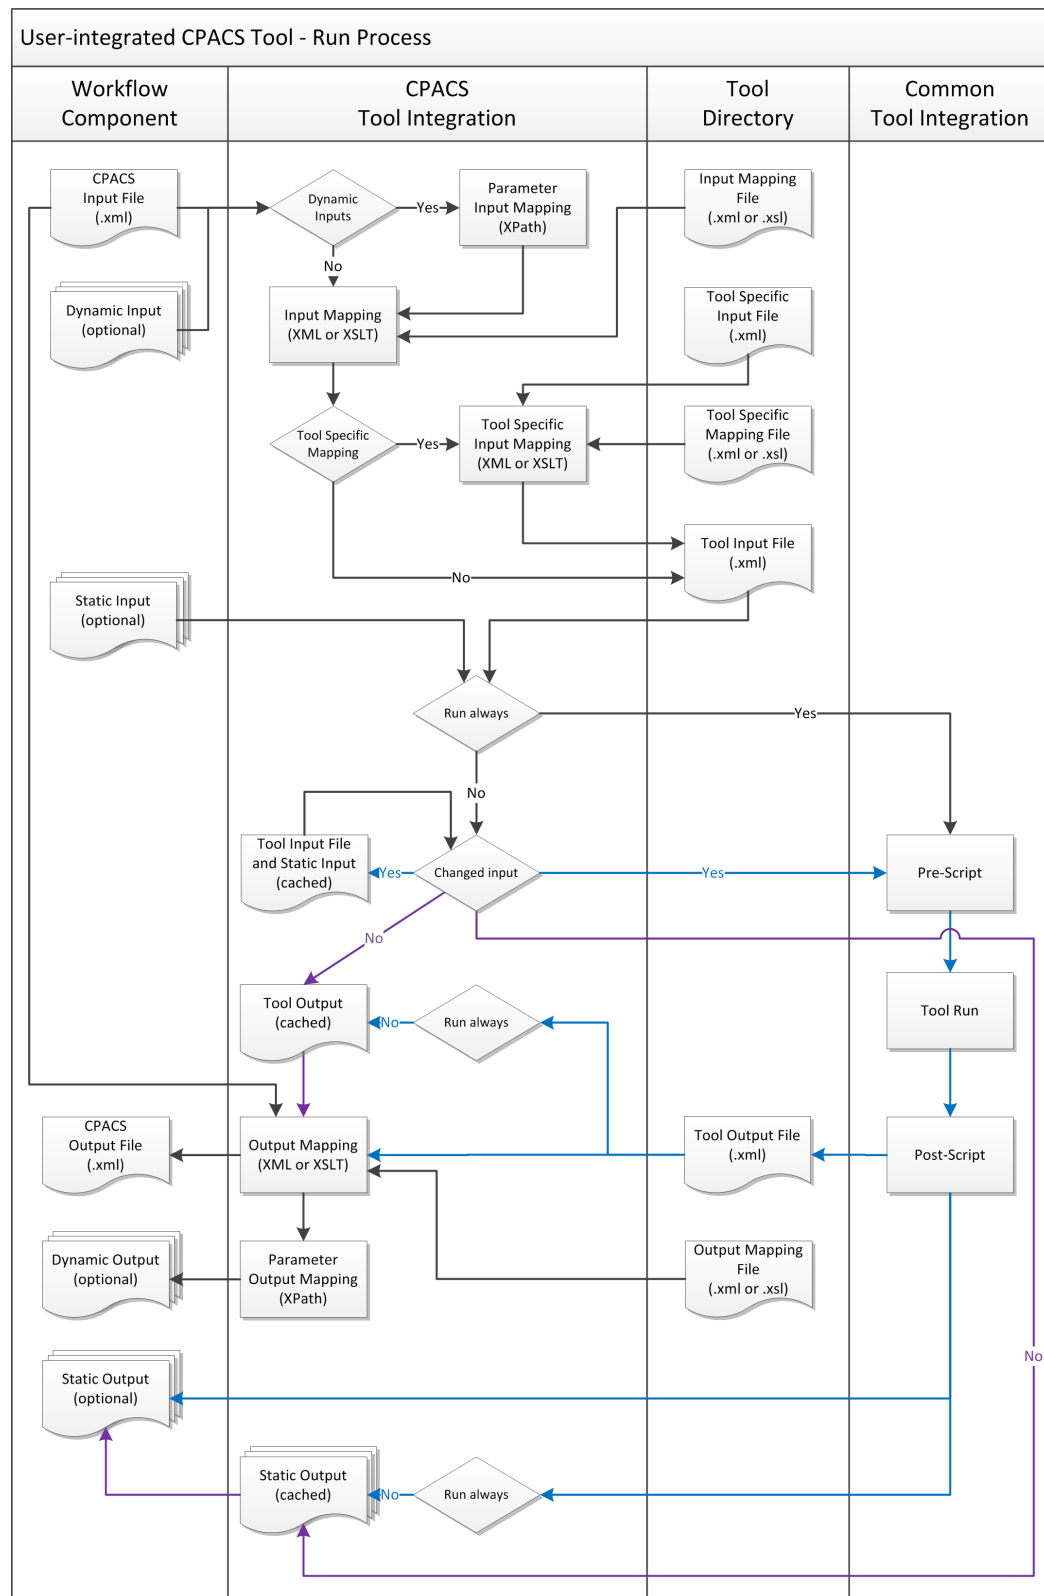
</map:mappings>
```

And output XSLT mapping:

```
<?xml version="1.0" encoding="UTF-8"?>
<xsl:stylesheet version="1.0" xmlns:xsl="http://www.w3.org/1999/XSL/Transform" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" exclude-result-prefixes="xsi">
  <xsl:output method="xml" version="1.0" encoding="UTF-8" indent="yes"/>
  <!--Define Variable for toolOutput.xml-->
  <xsl:variable name="toolOutputFile" select="'./ToolOutput/toolOutput.xml'"/>
  <!--Copy complete Source (e.g. CPACInitial.xml) to Result (e.g. CPACResult.xml)-->
  <xsl:template match="@* | node()">
    <xsl:copy>
      <xsl:apply-templates select="@* | node()"/>
    </xsl:copy>
  </xsl:template>
  <!--Modify a value of an existing node-->
  <xsl:template match="/cpacs/vehicles/aircraft/model/reference/area">
    <area>
      <xsl:value-of select="document($toolOutputFile)/toolOutput/data/result1"/>
    </area>
  </xsl:template>
</xsl:stylesheet>
```

Please ensure to use the proper namespace for map (xmlns:map="http://www.rcenvironment.de/2015/mapping") in XML mapping files.

The figure below illustrates how the additional features are used in the run process of an user-integrated CPACS tool.

Figure 3.7. Run process of an user-integrated CPACS Tool

4.4.2. Integrate a CPACS Tool into a Client Instance

1. Start RCE as Client
2. Open the *Tool Integration Wizard* by clicking the *Integrate Tool...* in the *File* menu.

Note

You will always find further help by clicking the ? on the bottom left corner on each page of the wizard or by pressing *F1*.

3. Choose the option *Create a new tool configuration from a template*.

Note

The CPACS templates delivered with RCE are designed to match the conventions of the old CPACS tool wrapper (respectively ModelCenter tool wrapper). Most of the properties are preconfigured and do not need to be changed.

4. Select one of the *CPACS* templates.
Click *Next*.
5. Fill in the *Tool Description* page.
Click *Next*.
6. On the *Inputs and Outputs* page you will find preconfigured static in- and outputs, that will match the old tool wrapper conventions. If your tool needs additional in- or outputs, feel free to configure.
Click *Next*.
7. Skip the page *Tool Properties* by clicking *Next* since it is not relevant for tools that match the conventions of the old CPACS tool wrapper.
8. Add a launch setting for the tool by clicking the *Add* button on the *Launch Settings* page. Configure the path of the CPACS tool and fill in a version, click *OK*. If you would like to allow users of your tool to choose that the temp directory won't be deleted at all after workflow execution, check the property *Never delete working directory(ies)*. Not to delete the working directory can be very useful for users for debugging purposes, at least if they have access to the server's file system. But this option can result in disc space issues as the amount required grows continuously with each workflow execution. It is recommended to check that option during integrating the tool and uncheck it before publishing the tool. If you like to publish your tool to the entire network right after the integration, check *Publish component after integration*. Note that you can publish your tool at any later point in time.
Click *Next*.
9. The *CPACS Tool Properties* are preconfigured to match the folder structure defined for the old CPACS tool wrapper. In most cases you do not have to change this configuration. If you are using XSLT mapping, please select the corresponding mapping files. If your tool does not work with static tool specific input, please deselect this property.
Click *Next*.
10. In the *Execution command(s)* tab on the *Execution* page, you need to define your execution command itself as well as optional pre and post commands. Commands will be processed sequentially line by line. An example for a typical Windows command including pre and post commands will look like the following:

```
rem pre-command
pre.bat

rem tool-execution
YourTool.exe ToolInput/toolInput.xml ToolOutput/toolOutput.xml

rem post-command
post.bat
```

11. Click *Save and activate* and your tool will appear immediately in the palette and is be ready to use.
12. If not already done, do not forget to publish your tool after testing it locally. Therefore, edit your tool configuration by clicking *Edit Tool* in the *File* menu, navigate to the page *Launch Configuration* and check the property *Publish component after integration*. Click *Save and activate*. To check if

your tool is successfully published to the RCE network open the tab *Network View* at the bottom and checkout *Published Components* after expanding the entry of your RCE instance.

4.4.3. Integrate a CPACS Tool into a Server Instance in Headless Mode

The way to integrate a CPACS tool on a server running RCE in headless mode is as follows: Perform the steps to integrate a CPACS tool on a client instance and make sure that the path of the CPACS tool configured on the *Launch Settings* page (step 8) matches the absolute tool path on your server system. Afterwards, you will find the configuration files inside your rce profile folder at the following location:

```
/integration/tools/cpacs/[YourToolName]
```

Copy the folder [YourToolName] to the same location inside the profile folder running with your headless server instance. Edit the file `published.conf` in the cpacs directory by adding your tool's directory name as a single line. If the server instance is already running, your tool will be available immediately after saving the file (cf. section "Publishing a component in headless mode").

5. Remote Tool and Workflow Access

RCE provides an interface that allows external applications to access and run single tools or complete workflows within RCE instances. This allows existing applications to make use of RCE's features (like network distribution, data management, logging, ...) when integrating the application itself is not possible or desirable. For example, applications that are based on frequent user interaction are not a good fit for being run as part of an automated workflow, but may still want to use some of RCE's features.

This section describes how to publish tools or workflows such that they can be used via Remote Access. It will guide you through the creation of a simple example, which you can expand to build your own solutions.

5.1. Basic Concepts

The basic idea of the Remote Access interface is that the external application opens an SSH connection to RCE, and initiates the tool or workflow execution via text commands. Input and output data, as well as log files, are transferred using SCP. For using the integrated tools and workflows remotely, a standalone remote access tool is currently under development and will be released soon. The integrated tools can also be used in a remote RCE instance via SSH connections.

5.2. Tool vs. Workflow Execution

In RCE 5.0.0, the "remote access" feature was able to invoke a single integrated tool. Starting with RCE 5.1.0, arbitrary user-defined workflows can be executed. The following sections describe each approach separately.

5.3. Setting up the Single Tool Execution Example

These steps will guide you through the configuration of an integrated tool. Using this example, you will be able to easily integrate your own text-based tools into RCE and invoke them using the Remote Access interface.

- Setting RCE up as a "Remote Access Server".
 - Download and unpack/install an RCE distribution (version 7.0.0 or higher).
 - Start RCE and select the "Help > Configuration Information" menu option. From the list of configuration examples, double-click "Remote Access Server" to open it. Mark all text (Ctrl-A) and copy it to the clipboard (Ctrl-C).
 - Select the "Help > Open Configuration File" menu option, select all text (Ctrl-A) and paste the copied text from the example configuration.
 - Save the file, and select "File > Restart" to apply the new configuration. The Remote Access SSH interface is now running, with a pre-configured SSH account of "ra_demo", password "ra_demo".

Note

IMPORTANT: "ra_demo" is just an example account for testing within a secure network. Create a new account with a better password (for example using the configuration UI described in the "Configuration" section) before using the Remote Access interface in production. Publishing tools using this default account is a potential security risk.

- Define an example tool using the following steps.
 - Select the "File > Integrate Tool" menu option.
 - In the dialog that has opened, select "Create a new configuration from a template" and choose "Remote Tool Access" from the list. Click "Next >".
 - On the "Tool Description" page, enter a name for the example tool (e.g. "test"). Click "Next >".
 - You can leave the next two wizard pages as they are. Skip them by clicking "Next >" twice.
 - You should be on the "Launch Settings" page now. Click "Add" on the right-hand side. On the dialog that opens, you normally enter a version number describing your tool's release number and the path where it is installed. For this example, just enter an arbitrary version and select an arbitrary directory (for example, a temporary files directory on your computer). Click "Ok", check "Publish component after integration" and then "Next >" to proceed.

Note

After completing this Remote Access example, refer to the "Tool Integration" section on how to integrate your actual tools.

- On the "Execution" page, enable the command section that matches your operating system. This is where you define the commands that execute your tool. For this example, just leave the commands as they are. Click "Save and Activate" to complete the wizard.
- You now have an example tool that can be accessed with the "Remote Access" feature. How to setup a client to access this tool see the "Configuring an RCE instance as an SSH client" section
- To get an impression of how this feature interacts with existing RCE features, you can examine several areas within the RCE instance.
 - Open the "Workflow List" view in RCE and watch it while the "run-tool" script is executing. After a short preparation time where the input data is uploaded, you will see the automatically generated workflow containing the tool being executed. It will disappear automatically if it finishes successfully; if it fails, it will remain in the list for review.
 - Open the "Workflow Console" view; if the tool produced any output, it should be visible there.
 - Open the "Workflow Data Browser" and refresh it; there should be a "Remote_Tool_Access-..." workflow entry matching the remote tool run. When you expand this entry, you should see the

uploaded content of the input folder, the generated output folder, any generated text output (in the "Execution Log" folder), and the exit code of the tool process (also in the "Execution Log" folder).

5.4. Setting up the Workflow Execution Example/Template

These steps will guide you through the creation of a remote-executable workflow, and will show you how to invoke it using the provided example scripts.

- Download and unpack/install an RCE distribution (version 7.0.0 or higher).
- Run RCE and select the "Help > Open Configuration Information" menu option. From the list of configuration examples, double-click "Remote Access Server" to open it. Mark all text (Ctrl-A) and copy it to the clipboard (Ctrl-C).
- Select the "Help > Open Configuration File" menu option, select all text (Ctrl-A) and paste the copied text from the example configuration.
- Save the file, and select "File > Restart" to apply the new configuration. The Remote Access SSH interface is now running, with a pre-configured SSH account of "ra_demo", password "ra_demo".

Note

IMPORTANT: "ra_demo" is just an example account for testing within a secure network. Create a new account with a better password (for example using the configuration UI described in the "Configuration" section) before using the Remote Access interface in production. Publishing tools using this default account is a potential security risk.

- As a first example we are going to execute the unmodified "Remote_Workflow_Access_Template" workflow file in the Workflow Examples Project. If you haven't created this project already, right-click in the Project Explorer on the left side, and choose "New > Workflow Examples Project", and choose a name for it. The template file is contained within the project folder. To get an impression of the basic setup, open the template workflow file. You will see an Input Provider on the left side with two outputs: one Directory and one ShortText. On the right side, there is an Output Provider with a single Directory input. These are the points where the Remote Workflow Access feature sends the provided input files (as a Directory) and parameters (as a ShortText) into your workflow, and collects the final output files (as a Directory).
- As a security measure, you need to explicitly publish your workflow to allow remote access to it. This is done via a console command at this time; future RCE versions will most likely add a option to do this from the GUI. To issue this command, open the "Command Console" view (if it is not already visible) by selecting "Windows > Show View > Other" from the menu, and then double-clicking "Command Console" in the "RCE" section.
- Right-click your workflow file in the "Project Explorer" and select the "Properties" entry at the bottom of the popup menu. Mark the value of the "location" entry with the mouse, and press Ctrl+C to copy the full path to the workflow file to the clipboard.

Note

This step demonstrates how to get the path of a workflow file in the current workspace, but you can use workflow files that are located anywhere on your system.

- To make the workflow available for remote execution, enter the command `ra-admin publish-wf "<workflow file>" <id>` in the command window. Press Ctrl-V in place of `<workflow file>` to insert the path to your workflow file there. For `<id>`, choose a string (without whitespace) that callers can use to execute the workflow. Press "enter" to execute the command. The workflow file will be inspected, and you will either see a message describing what is missing,

or a message that the workflow was successfully published. Fix any errors until the workflow is published.

Note

Starting with RCE 6.2.0, published workflows are persistent by default, so they will still be available after the local RCE instance is restarted. Use the `ra-admin unpublish-wf <id>` command to remove a published workflow from remote access.

To publish a workflow for the current RCE instance's life-time only, use the `-t` option: `ra-admin publish-wf -t "<workflow file>" <id>`.

- If some of the workflow's components use placeholders for configuration values, you can use the `-p` option to specify a placeholder values file. The structure of placeholder value files is explained in Section 3.3.1, "Configuration Placeholder Value Files". Placeholder files can be used with both persistent and non-persistent workflows (see above).

Example: `ra-admin publish-wf -p myPlaceholderValues.json myWorkflowFile.wf myPublishId`

- You now have a workflow file that can be executed using the "Remote Access" feature.
- To get an impression of how this feature interacts with existing RCE features, you can examine several areas within the RCE instance.
 - Open the "Workflow List" view in RCE and watch it while the "run-wf" script is executing. After a short preparation time where the input data is uploaded, you will see the workflow being executed. It will disappear automatically if it finishes successfully; if it fails, it will remain in the list for review. You can also double-click on a running workflow to monitor its execution.
 - Open the "Workflow Console" view; if the tool produced any output, it should be visible there.
 - Open the "Workflow Data Browser" and refresh it; there should be an entry for the Remote Access workflow. When you expand this entry, you should see the uploaded content of the input folder, the generated output folder, any generated text output (in the "Execution Log" folder), and the exit code of the tool process (also in the "Execution Log" folder).

5.5. Building Your Own Remote Access Workflow

After running the example/template workflow as described in the previous section, you can proceed to building your own actual workflow.

As described above, open the "Remote_Workflow_Access_Template" workflow file. You will see an Input Provider on the left side with two outputs: one Directory and one ShortText. On the right side, there is an Output Provider with a single Directory input. These are the points where the Remote Workflow Access feature sends the provided input files (as a Directory) and parameters (as a ShortText) into your workflow, and collects the final output files (as a Directory). You should leave these two components unchanged, and place your workflow logic between them. The Script component in the middle is just a placeholder - unless you need a Script component anyway, you can just delete it.

There are two basic approaches to building your workflow:

- Either you start with the template, and build your workflow between the two standard components. This is straight-forward, but means that you cannot test run the workflow from the RCE GUI (as the Input Provider will fail), but have to use the Remote Access feature to test it.
- The other approach is to build your workflow normally, where you add an Input Provider and Output Writer with the same outputs and inputs as the ones in the template (but with standard configuration values). You can then test (and if needed, modify) your workflow from the GUI until it behaves as it

should. Then, mark all components *except* the Input Provider and Output writer in your workflow, and select "Copy" from the right-click menu. Switch to the template file, click an empty area, and select "Paste" from the right-click menu. Then, connect the two template components (Input Provider and Output Writer) as in your original workflow.

Note

(Advanced Usage) You can also build your workflow in the template file, add your own Input Provider and Output Writer, and use the new "Enable/Disable Component" feature to toggle between them for testing and Remote Access usage. As this requires some helper components to work, this is not recommended for your first example, but may be a useful trick to keep in mind.

After you have finished building your workflow, the process of publishing and executing it is the same as described above for the unmodified template file. Please note that publishing your workflow for remote execution automatically creates an (invisible) copy of it. Modifications you make to your workflow file are not published right away. Once you have made the changes you want to publish, run the same "ra-admin publish-wf" command again to update the published version.

Note

Tip: To repeat a previous command, presse the "up arrow" key in the Command Console window.

6. Connecting RCE instances via the RCE network or via SSH connections

RCE provides two possibilities to connect your RCE instance to other RCE instances and to use the user-integrated tools and components published on those instances: The RCE network connections and SSH connections. RCE connections are meant to be used only in a trusted network (e.g. your institution's internal network). The RCE network traffic is currently *not encrypted*. This means that it is *not secure* to expose RCE server ports to untrusted networks like the internet. In the case that it is not possible or not secure to use RCE connections, SSH connections provide a more secure alternative. However, not all tools can be used via an SSH connection in the current version. The following table compares the two connection types:

Table 3.13. Connection types

Connection type	Access to all published tools	Supporting all RCE network functions (e.g. documentation, remote monitoring...)	Encrypted traffic	Using login name and password or RSA key	Symmetrical
RCE connections	yes	yes	no	no	yes
SSH connections	no	no	yes	yes	no

6.1. RCE Network Connections

RCE connections are meant to be used only in a trusted network (e.g. your institution's internal network). To build up a network of RCE instances, at least one of the instances has to be configured as a server (see the "Configuration" section or the sample configuration file "Relay server" for details).

On the client side, RCE network connections can be added in the "network" view by clicking "Add network connection" and entering the hostname and port of an RCE server instance. The connections

are shown in the "RCE Network"->"Connections" subtree. They can also be edited, connected and disconnected in the network view. However, the changes made here are not saved in the configuration yet, i.e. they will be lost when RCE is closed or restarted. To permanently add connections, you can edit the configuration file (see the "Configuration" section for details).

In the "RCE Network" -> "Instances" subtree all RCE instances in the network are listed. When expanding the entry for an instance, you can see monitoring data like CPU or RAM usage for this instance, and the published components and tools of this instance (if it has any).

The published components and tools of the other instances in your network are also shown in the palette of the Workflow Editor. From there, you can use them in your workflows just like your local components and tools. When you start a workflow, in the "Execute Workflow" wizard there is an overview which component will be run on which RCE instance. If a component is available on several instances, you can choose here on which instance it should be run. In the same wizard, you can also choose another instance as the "Controller Target Instance", which means that the workflow execution will be controlled by this instance (see the section "Configuration Parameters" for more information). This can be useful when you start a long-running workflow where all components are run on remote instances and you do not want to keep your local computer connected all the time.

6.2. SSH Remote Access Connections

SSH connections provide a more secure alternative to the standard RCE connections and can be used to access tools remotely that were integrated using the "Remote Tool Access" template, which determines certain inputs and outputs. The published tools meeting the requirements of the "Remote Tool Access" template are shown in the palette of the client's Workflow Editor in the group "SSH Remote Access Tools" (this may take a few seconds after connecting, as the tool list is fetched from the remote hosts every few seconds). From there, you can use them in your workflows just like your local components and tools. Differently from tools accessed by RCE network connections, in this case the component is shown to be executed on your local instance in the Workflow Execution wizard.

Also workflows that were published on the remote instance (for information about the publishing see section "Remote Tool and Workflow Access") are shown as components in the palette of the client's Workflow Editor in the group "SSH Remote Access Workflows" (if the client runs RCE 7.1 or newer). These remote workflows can be added to workflows and executed like local components/tools.

6.2.1. Configuring an RCE instance as an SSH server

The RCE instance that publishes the tool, or another instance connected to it by RCE network connections, has to be configured as an RCE remote access server (see the "Configuration" section or the sample configuration file "Remote access server" for details).

Note

When configuring an SSH account using a key file, note that for using key files, both server and client have to run RCE 7.1 or newer and that only RSA-Keys in the OpenSSH format are supported by RCE. You can generate such keys on Linux e.g. using the "ssh-keygen" tool (which creates the correct format by default). On Windows, you can use the software "puttygen" to generate a key (choose the type "SSH-2 RSA") and convert it to the OpenSSH format ("Conversions"->"Export OpenSSH key" in puttygen).

6.2.2. Configuring an RCE instance as an SSH client

On the client side, SSH connections can be added in the "network" view by clicking "Add SSH Remote Access Connection". In the following dialog, enter the hostname and port of an RCE instance that provides an SSH server as well as the user name and the authentication data of an SSH account configured on this instance. Depending on the SSH account, you have to authenticate using a passphrase or by an RSA private key file. If your private key is protected by a passphrase, select

the authentication type "Keyfile with passphrase protection", else select "Keyfile without passphrase protection".

The connections are shown in the "SSH Remote Access"-"SSH Remote Access Connections" subtree. They can also be edited, connected and disconnected in the "network" view. It is possible to store passphrases using the Eclipse Secure Storage Mechanism. However, the changes made here are not saved in the configuration yet, i.e. they will be lost when RCE is closed or restarted. To permanently add SSH connections, you can edit the configuration file (see the "Configuration" section for details). However, it is not possible yet to store the passphrases for such connections permanently.