



# UNICORE COMMANDLINE CLIENT: USER MANUAL

UNICORE Team

---

Document Version:	1.0.0
Component Version:	7.13.0
Date:	07 10 2019

---

This work is co-funded by the EC EMI project under the FP7 Collaborative Projects Grant Agreement Nr. INFSO-RI-261611.



## Contents

<b>1</b>	<b>Overview</b>	<b>1</b>
<b>2</b>	<b>Installation and configuration</b>	<b>1</b>
2.1	Prerequisites . . . . .	1
2.2	Download . . . . .	1
2.3	Installation and configuration . . . . .	2
2.4	Preferences file . . . . .	2
2.5	Logging . . . . .	3
2.6	Installing UCC extensions . . . . .	3
2.7	Testing the installation . . . . .	3
<b>3</b>	<b>Getting started with UCC</b>	<b>3</b>
3.1	Getting help . . . . .	4
3.2	Connecting . . . . .	4
3.3	List available sites . . . . .	4
3.4	Running your first job . . . . .	4
3.5	Listing your jobs . . . . .	5
<b>4</b>	<b>Common options to UCC</b>	<b>5</b>
4.1	User attributes and VOs . . . . .	5
4.2	Configuration file . . . . .	6
4.3	Credential and truststore options . . . . .	7
4.4	Trust store examples . . . . .	12
4.5	Using Unity . . . . .	12
4.6	Support for oidc-agent . . . . .	13
4.7	Using MyProxy . . . . .	14
4.8	Client options . . . . .	14
4.9	Other options . . . . .	17

<b>5</b>	<b>SAML PUSH support</b>	<b>17</b>
5.1	Introduction . . . . .	17
5.2	Basic usage . . . . .	17
5.3	Attribute filtering . . . . .	18
5.4	Rules for multiple filters . . . . .	19
<b>6</b>	<b>Running jobs</b>	<b>20</b>
6.1	Introduction . . . . .	20
6.2	Options overview . . . . .	21
6.3	Resource brokering . . . . .	21
6.4	Processing jobs asynchronously . . . . .	22
6.5	Executing a command . . . . .	24
<b>7</b>	<b>Job description format</b>	<b>24</b>
7.1	Site name . . . . .	25
7.2	Specifying the application or executable . . . . .	25
7.3	Arguments and Environment settings . . . . .	26
7.4	Application parameters . . . . .	26
7.5	Job data management . . . . .	27
7.6	Resources . . . . .	31
7.7	Execution environments . . . . .	32
7.8	Miscellaneous options . . . . .	32
<b>8</b>	<b>Data management functions</b>	<b>33</b>
8.1	Specifying remote locations . . . . .	34
8.2	Data movement . . . . .	35
8.3	General commands . . . . .	36
8.4	Finding data . . . . .	37
8.5	Deprecated commands . . . . .	38
8.6	Using the StorageFactory service . . . . .	38
<b>9</b>	<b>Metadata management functions</b>	<b>39</b>
9.1	Basics . . . . .	39
9.2	Available commands . . . . .	39

<b>10 Workflows</b>	<b>41</b>
10.1 Introduction . . . . .	41
10.2 Command overview . . . . .	41
10.3 Basic use . . . . .	41
10.4 Managing workflow data . . . . .	42
<b>11 Batch processing</b>	<b>45</b>
11.1 Options . . . . .	45
11.2 Performance tuning options . . . . .	45
11.3 Resource selection in batch mode . . . . .	46
<b>12 OGSA-BES functions</b>	<b>47</b>
12.1 OGSA-BES Setup . . . . .	47
12.2 Running and monitoring OGSA-BES jobs . . . . .	48
12.3 Get OGSA-BES job outputs . . . . .	49
12.4 Enabling username/password authentication . . . . .	49
<b>13 The UCC shell</b>	<b>49</b>
13.1 Exiting the shell . . . . .	50
13.2 Changing property settings . . . . .	50
<b>14 Sharing resources</b>	<b>50</b>
14.1 Editing ACLs . . . . .	51
14.2 Permission levels . . . . .	51
<b>15 UCC for site administrators</b>	<b>51</b>
15.1 Security considerations . . . . .	52
15.2 Admin commands . . . . .	52
15.3 Listing jobs, sites, . . . . .	55
15.4 Low-level operations . . . . .	56
<b>16 Scripting</b>	<b>56</b>
16.1 Script context . . . . .	56
16.2 Examples . . . . .	57

<b>17 Frequently asked questions</b>	<b>59</b>
17.1 Configuration . . . . .	59
17.2 Usage . . . . .	60

## 1 Overview

The UNICORE Commandline client (UCC) is a full-featured client for the UNICORE middleware. UCC has client commands for all the UNICORE basic services, the OGSA-BES interface and the UNICORE workflow system.

It offers the following functions

- Job submission and management for both UNICORE native and OGSA-BES interfaces
- Batch mode job submission and processing with many performance tuning options
- Data movement (upload, download, server-to-server copy, etc) using the UNICORE storage management functions and available data transfer protocols
- Storage functions (ls, mkdir, ...) including creation of storage instances via storage factories
- Full UNICORE workflow system support, including the possibility to run single jobs through the resource brokering system
- Support for the UNICORE metadata system
- Support for sharing UNICORE resources via ACLs
- Information about the available services is provided via the "system-info" command
- Various utilities like a "shell" mode, the ability to generate SAML trust delegations, low-level WSRF operations and others
- Extensibility through custom commands and the possibility to run scripts written in the Groovy programming language
- Built-in help

For more information about UNICORE visit <http://www.unicore.eu>.

## 2 Installation and configuration

### 2.1 Prerequisites

To run UCC, you need a Java runtime version 8 or later (OpenJDK, Oracle or IBM). You might need to install the "unlimited cryptography extensions" for the Oracle or IBM Java.

### 2.2 Download

You can get the latest version from the SourceForge [UNICORE download page](#).

## 2.3 Installation and configuration

To install, unpack the distribution in a directory of your choice. It's a good idea to add the bin/ directory to your PATH variable,

```
export PATH=$PATH:<UCC_HOME>/bin
```

where UCC\_HOME is the directory you installed UCC in.

---

### Note

**Windows only** Please do not install UCC into a directory containing spaces such as "Program files".

Also avoid long path names, this can lead to errors due to the Windows limit on command line length.

Setting environment variables can be done (as administrator) using the Control panel→System→Extras panel.

---

Though you can specify your keystore location and other parameters on the commandline, it is easiest to place this information in a file, so that you do not have to key in this information repeatedly.

## 2.4 Preferences file

UCC checks by default whether the file <userhome>/ucc/preferences exists, and reads it.

A minimal example that specifies keystore, password and your preferred UNICORE registry URL would look as follows:

```
credential.path=<your keystore>
credential.password=<your password>
truststore.type=keystore
truststore.keystorePath=<your keystore>
truststore.keystorePassword=<your password>
client.serverHostnameChecking=NONE
registry=<your registry>
```

Please refer to Section 4 for a full description of available options.

---

### Note

If you are worried about security, and do not want specify the password: UCC will ask for it if it is not given in the preferences or on the commandline.

---

---

**Note**

**Windows only** The preferences are usually searched in the "c:\Users\\.ucc" folder.

To create the .ucc folder, you might have to use the command prompt "mkdir" command.

When specifying paths in the preferences file, the backslash | character needs to be written using an extra backslash ||

---

For example, if you are using a local UNICORE installation for testing, you could use

```
registry=https://localhost:8080/DEMO-SITE/services/Registry?res= ↵  
default_registry
```

---

**Note**

If you wish to change the default property file location, you can set a Java VM property in the UCC start script, for example by editing the command that starts UCC

```
java .... -Ducc.preferences=<preferences location> ....
```

---

## 2.5 Logging

UCC writes some messages to the console, more if you choose the verbose mode (-v option). If you need real logging (e.g. when using the batch mode), you can edit the <UCC\_HOME>/conf/logging.properties file, which configures the Log4J logging infrastructure used in UNICORE.

## 2.6 Installing UCC extensions

UCC can be extended with additional commands. It is enough to copy the libraries (.jar files) of the extension into a directory that is scanned by UCC: in general these are the UCC lib and the \${HOME}/.ucc/lib directory.

## 2.7 Testing the installation

To test your UCC installation and to get information about the services available in the UNICORE system you're connecting to, do

```
ucc system-info -l -v
```

## 3 Getting started with UCC

Assuming you have successfully installed UCC, this section shows how to get going quickly.



### 3.1 Getting help

Calling UCC with the "-h" option will show the available options. To get a list of available commands, type

```
ucc -h
```

To get help on a specific command, type

```
ucc <command> -h
```

See also [here](#) for a list of common options.

### 3.2 Connecting

First, contact UNICORE and make sure you have access to some target systems.

```
ucc connect [options]
```

### 3.3 List available sites

Then, list the sites available to you using

```
ucc list-sites [options]
```

### 3.4 Running your first job

The UCC distribution contains samples that you can run. Let's run the "date" sample. The "-v" switch prints more info so you can see what's going on.

```
ucc run [options] -v [UCC_HOME]/samples/date.u
```

---

**Note**

Look for UCC samples in the `/usr/share/doc/unicore/ucc/samples` directory,

---

This will run "date" on a randomly chosen site, and retrieve the output. To run on a particular site, use the "-s" option to specify a particular target system.

### 3.5 Listing your jobs

The command

```
ucc list-jobs [options]
```

will print a list of jobs (actually their addresses) with their respective status (RUNNING, SUCCESSFUL, etc)

## 4 Common options to UCC

The following table lists the options understood by most UCC commands. Most commands have additional options. You can always get a summary of all available options for a command by calling UCC with the "-h" or "--help" option, for example

```
ucc batch --help
```

Since it is not possible to give all the required options on the commandline, it is mandatory to create a preferences file containing e.g. your settings for keystore, registry etc.

Table 1: Common options for the UCC

Option (short and long form)	Description
-c, --configuration <Properties_file>	Properties file containing your preferences. By default, a file <i>userhome/.ucc/preferences</i> is checked.
-k, --authentication <auth>	Authentication method to use (default: X509)
-o, --output <Output_dir>	Directory for any output produced (default is the current directory)
-r, --registry <List_of_Registry_URLs>	The comma-separated list of URLs of UNICORE registries
-v, --verbose	Verbose mode
-h, --help	Print help message
-y, --with-timing	Timing mode

### 4.1 User attributes and VOs

If you have multiple user IDs or are a member of multiple Unix Groups on the target system, you may wish to control the user attributes that are used when invoking UCC. Also, when using UCC with infrastructures that run a SAML-enabled VO server such as UVOS, there are

additional options for specifying the VO server, your VO etc. Please see [?] section for details about the VO support.

Here is a list of options related to user attributes and VOs.

Table 2: Security and VO options for the UCC

Option (short and long form)	Description
-U, --user	User ID to use remotely (if you have multiple)
-Z, --preference	Select from your remote attributes (e.g. xlogin)
-J, --VO	VO server URL
-G, --voGroup	VO group
-A, --attributeAssertion	File containing a VO attribute assertion
-I, --includeAttributes	Filter VO attributes
-Q, --excludeAttributes	Filter VO attributes

## 4.2 Configuration file

By default, UCC checks for the existence of a file <userhome/.ucc/preferences> and reads settings from there. As shown above, you can use a different file by specifying it on the commandline using the "-c" option.

The configuration file can contain default settings for many commandline options, which are given in the form <option name>=<value> where <option name> is the long form of the option. The property values may contain variables in the form \${VAR\_X}, which are automatically replaced with the environmental variable values with the same name. Additionally a special variable \${UCC\_CONFIG} is recognized and is replaced with the absolute path of your configuration file.

For example, to set your keystore, truststore and registry, the file would contain the following settings

```
credential.path=<your keystore>
credential.password=XXXXXXX
truststore.type=keystore
truststore.keystorePath=<your keystore>
truststore.keystorePassword=XXXXXX
registry=https://localhost:8080/DEMO-SITE/services/Registry?res= ↵
default_registry
```

**Note**

To protect your passwords, you should make the file non-readable by others, for example on Unix using a command such as `chmod 600 preferences`

**Note**

If the credential and/or truststore passwords are not given in the properties file, they will be queried interactively.

### 4.3 Credential and truststore options

In general you need a keystore containing your identity in order to use UNICORE, as well as a truststore file (or directory) containing trusted certificates. (Note that there may be other options available for authentication, try `ucc help-auth` to find out)

A full list of options related to credential and truststore management is available in the following table. You can also get them via the online help using

```
ucc help-auth
```

Table 3: Credential properties

Property name	Type	Default value / mandatory	Description
<code>credential.path</code>	filesystem path	<i>mandatory to be set</i>	Credential location. In case of <i>jks</i> , <i>pkcs12</i> and <i>pem</i> store it is the only location required. In case when credential is provided in two files, it is the certificate file path.
<code>credential.format</code>	[jks, pkcs12, der, pem]	-	Format of the credential. It is guessed when not given. Note that <i>pem</i> might be either a PEM keystore with certificates and keys (in PEM format) or a pair of PEM files (one with certificate and second with private key).
<code>credential.password</code>	string	-	Password required to load the credential.

Table 3: (continued)

Property name	Type	Default value / mandatory	Description
credential.keyPath	string	-	Location of the private key if stored separately from the main credential (applicable for <i>pem</i> and <i>der</i> types only),
credential.keyPass	string	-	Private key password, which might be needed only for <i>jks</i> or <i>pkcs12</i> , if key is encrypted with different password then the main credential password.
credential.keyAlias	string	-	Keystore alias of the key entry to be used. Can be ignored if the keystore contains only one key entry. Only applicable for <i>jks</i> and <i>pkcs12</i> .

Table 4: Truststore properties

Property name	Type	Default value / mandatory	Description
truststore.allowProxy	[ALLOW, DENY]	ALLOW	Controls whether proxy certificates are supported.
truststore.type	[keystore, openssl, directory]	<i>mandatory to be set</i>	The truststore type.
truststore.updateInterval	integer number	600	How often the truststore should be reloaded, in seconds. Set to negative value to disable refreshing at runtime. ( <i>runtime updateable</i> )
--- Directory type settings ---			
truststore.directoryTimeout	integer number	Timeout	Connection timeout for fetching the remote CA certificates in seconds.

Table 4: (continued)

Property name	Type	Default value / mandatory	Description
truststore.directory	file system path	Path	Directory where CA certificates should be cached, after downloading them from a remote source. Can be left undefined if no disk cache should be used. Note that directory should be secured, i.e. normal users should not be allowed to write to it.
truststore.directory	PEM/DER	PEM	For directory truststore controls whether certificates are encoded in PEM or DER. Note that the PEM file can contain arbitrary number of concatenated, PEM-encoded certificates.
truststore.directory	list of locations. *-properties with a common prefix	*	List of CA certificates locations. Can contain URLs, local files and wildcard expressions. <i>(runtime updateable)</i>
<i>--- Keystore type settings ---</i>			
truststore.keystore	string	-	The keystore type (jks, pkcs12) in case of truststore of keystore type.
truststore.keystore	string	-	The password of the keystore type truststore.
truststore.keystore	string	-	The keystore path in case of truststore of keystore type.
<i>--- Openssl type settings ---</i>			
truststore.openssl	boolean	false	In case of openssl truststore, specifies whether the trust store is in openssl 1.0.0+ format (true) or older openssl 0.x format (false)

Table 4: (continued)

Property name	Type	Default value / mandatory	Description
truststore.openssl	GLOBAL_EUGRIDPMA_GLOBUS_REQUIRE	GRIDPMA_GLOBUS_REQUIRE	Location of openssl truststore, controls which (and in which order) namespace checking rules should be applied. The REQUIRE settings will cause that all configured namespace definitions files must be present for each trusted CA certificate (otherwise checking will fail). The AND settings will cause to check both existing namespace files. Otherwise NONE, found is checked (in the order defined by the property).
truststore.openssl	filesystem path	/etc/grid	Directory to be used for openssl truststore.
--- Revocation settings ---			
truststore.crlConn	integer number	15	Connection timeout for fetching the remote CRLs in seconds (not used for Openssl truststores).
truststore.crlDisk	filesystem path	-	Directory where CRLs should be cached, after downloading them from remote source. Can be left undefined if no disk cache should be used. Note that directory should be secured, i.e. normal users should not be allowed to write to it. Not used for Openssl truststores.
truststore.crlLocal	list of .* properties with a common prefix	-	List of CRLs locations. Can contain URLs, local files and wildcard expressions. Not used for Openssl truststores. ( <i>runtime updateable</i> )

Table 4: (continued)

Property name	Type	Default value / mandatory	Description
truststore.crlMode	[REQUIRE, IF_VALID, IGNORE]	IF_VALID	General CRL handling mode. The IF_VALID setting turns on CRL checking only in case the CRL is present.
truststore.crlUpdate	integer number	600	How often CRLs should be updated, in seconds. Set to negative value to disable refreshing at runtime. ( <i>runtime updateable</i> )
truststore.ocspCache	integer number	3600	For how long the OCSP responses should be locally cached in seconds (this is a maximum value, responses won't be cached after expiration)
truststore.ocspDisk	filesystem path	-	If this property is defined then OCSP responses will be cached on disk in the defined folder.
truststore.ocspLocal	list of responders properties with a common prefix	.<NUMBER>	Optional list of local OCSP responders
truststore.ocspMode	[REQUIRE, IF_AVAILABLE, IGNORE]	IF_AVAILABLE	General OCSP checking mode. REQUIRE should not be used unless it is guaranteed that for all certificates an OCSP responder is defined.
truststore.ocspTimeout	integer number	10000	Timeout for OCSP connections in miliseconds.
truststore.revocation	[CRL, OCSP, OCSP_CRL]	OCSP_CRL	Controls overall revocation sources order
truststore.revocationAll	[true, false]	false	Controls whether all defined revocation sources should be always checked, even if the first one already confirmed that a checked certificate is not revoked.



## 4.4 Trust store examples

Here are some examples for commonly used trust store configurations.

Directory trust store with a minimal set of options

```
truststore.type=directory
truststore.directoryLocations.1=/trust/dir/*.pem
```

Directory trust store with more options

```
truststore.type=directory
truststore.allowProxy=DENY
truststore.updateInterval=1234
truststore.directoryLocations.1=/trust/dir/*.pem
truststore.directoryLocations.2=http://caserver/ca.pem
truststore.directoryEncoding=PEM
truststore.directoryConnectionTimeout=100
truststore.directoryDiskCachePath=/tmp
truststore.crlLocations.1=/trust/dir/*.crl
truststore.crlLocations.2=http://caserver/crl.pem
truststore.crlUpdateInterval=400
truststore.crlMode=REQUIRE
truststore.crlConnectionTimeout=200
truststore.crlDiskCachePath=/tmp
```

Java keystore used as a trust store:

```
truststore.type=keystore
truststore.keystorePath=/some/dir/truststore.jks
truststore.keystoreFormat=JKS
truststore.keystorePassword=xxxxxxx
```

OpenSSL trust store

```
truststore.type=openssl
truststore.opensslPath=/truststores/openssl
truststore.opensslNsMode=EUGRIDPMA_GLOBUS_REQUIRE
truststore.allowProxy=ALLOW
truststore.updateInterval=1234
truststore.crlMode=IF_VALID
```

## 4.5 Using Unity

If your UNICORE installation is using the Unity identity management service (see <http://www.unity-idm.eu>), you can setup UCC to use Unity. In this case you do not need a private key, only a truststore. UCC is configured for Unity using the following properties

```

authenticationMethod=unity
unity.address=https://<host>:<port>/unicore-soapidp/saml2unicoreidp ←
  -soap/AuthenticationService
unity.username=<your Unity username>
unity.password=<your Unity password>

```

At minimum, you must specify the "authenticationMethod" and "unity.address" parameters. Password and username are optional: if you do not specify them, they will be queried interactively.

In some installations, you can use a so-called OIDC Bearer token instead of username and password. In this case specify only unity.address and the token, like this:

```

authenticationMethod=unity
unity.address=https://<host>:<port>/unicore-soapidp/saml2unicoreidp ←
  -soap-oidc/AuthenticationService
oauth2.bearerToken=<your oauth token>

```

## 4.6 Support for oidc-agent

For infrastructures using OIDC authentication, UCC supports the *oidc-agent* tool that allows to interact with common OIDC servers to retrieve new access tokens.

Please visit <https://github.com/indigo-dc/oidc-agent> for more information.

To use oidc-agent, UCC has the following properties

Table 5: Options for oidc-agent

Property name	Type	Default value / mandatory	Description
oidc-agent.account	string	<i>mandatory to be set</i>	Account short name.
oidc-agent.lifetime	integer >= 1	-	Minimum lifetime of the issued access token.
oidc-agent.scope	string	-	OpenID scope(s) to request.

Note that you'll also need to configure a Unity server that will accept the OIDC token. Your config file would require at least:

```

authenticationMethod=oidc-agent
unity.address=https://<host>:<port>/unicore-soapidp/saml2unicoreidp ←
  -soap-oidc/AuthenticationService

```

```
oidc-agent.account=<oidc-agent account to be used>
```

## 4.7 Using MyProxy

UCC can retrieve a short lived certificate from a MyProxy server. To setup UCC for this, use the following properties

```
authenticationMethod=MYPROXY
myproxy.host=myproxy.teragrid.org
myproxy.port=7512
myproxy.username=<your MyProxy username>
myproxy.password=<your MyProxy password>
myproxy.expires=12
```

At minimum, you must specify the "authenticationMethod" and "myproxy.host" parameters. The default port of 7512 will be used if not otherwise specified. Password and username are optional: if you do not specify them, they will be queried interactively. The lifetime (in hours) of the short lived certificate can be set with the "myproxy.expires" option, if not set the lifetime is 12 hours.

## 4.8 Client options

The configuration file may also contain low-level options, for example if you need to specify connection timeouts, http proxies etc.

Table 6: Client options

Property name	Type	Default value / mandatory	Description
client.digitalSignatureEnabled	[true, false]	true	Controls whether signing of key web service requests should be performed.
client.httpAuthnEnabled	[true, false]	false	Whether HTTP basic authentication should be used.
client.httpPassword	string	empty string	Password for use with HTTP basic authentication (if enabled).
client.httpUser	string	empty string	Username for use with HTTP basic authentication (if enabled).

Table 6: (continued)

Property name	Type	Default value / mandatory	Description
<code>client.inHandlers</code>	string	<i>empty string</i>	Space separated list of additional handler class names for handling incoming WS messages
<code>client.maxWsCallRetries</code>	integer number	3	Controls how many times the client should try to call a failing web service. Note that only the transient failure reasons cause the retry. Note that value of 0 enables unlimited number of retries, while value of 1 means that only one call is tried.
<code>client.messageLogging</code>	[true, false]	false	Controls whether messages should be logged (at INFO level).
<code>client.outHandlers</code>	string	<i>empty string</i>	Space separated list of additional handler class names for handling outgoing WS messages
<code>client.securitySessions</code>	[true, false]	true	Controls whether security sessions should be enabled.
<code>client.serverHostCertificateChecking</code>	[NONE, WARN, FAIL]	WARN	Controls whether server's hostname should be checked for matching its certificate subject. This verification prevents man-in-the-middle attacks. If enabled WARN will only print warning in log, FAIL will close the connection.
<code>client.sslAuthnEnabled</code>	[true, false]	true	Controls whether SSL authentication of the client should be performed.
<code>client.sslEnabled</code>	[true, false]	true	Controls whether the SSL/TLS connection mode is enabled.
<code>client.wsCallRetryDelay</code>	integer number	10000	Amount of milliseconds to wait before retry of a failed web service call.

Table 6: (continued)

Property name	Type	Default value / mandatory	Description
<i>--- HTTP client settings ---</i>			
client.http.allowChunking	{true, false}	true	If set to false, then the client will not use HTTP 1.1 data chunking.
client.http.connectionClose	{true, false}	false	If set to true then the client will send connection close header, so the server will close the socket.
client.http.connectionTimeout	integer number	20000	Timeout for the connection establishing (ms)
client.http.maxPerHostConnections	integer number	6	How many connections per host can be made. Note: this is a limit for a single client object instance.
client.http.maxRedirects	integer number	3	Maximum number of allowed HTTP redirects.
client.http.maxTotalConnections	integer number	20	How many connections in total can be made. Note: this is a limit for a single client object instance.
client.http.socketTimeout	integer number	0	Socket timeout (ms)
<i>--- HTTP proxy settings ---</i>			
client.http.nonProxyHosts	string	-	Space (single) separated list of hosts, for which the HTTP proxy should not be used.
client.http.proxy.Password	string	-	Relevant only when using HTTP proxy: defines password for authentication to the proxy.
client.http.proxy.Username	string	-	Relevant only when using HTTP proxy: defines username for authentication to the proxy.
client.http.proxy.Hostname	string	-	If set then the HTTP proxy will be used, with this hostname.

Table 6: (continued)

Property name	Type	Default value / mandatory	Description
client.http.proxyPort	integer number	-	HTTP proxy port. If not defined then system property is consulted, and as a final fallback 80 is used.
client.http.proxyType	string	HTTP	HTTP proxy type: HTTP or SOCKS.

## 4.9 Other options

The following table lists other options, that are more rarely used.

Table 7: Other options for the UCC

Property name	Description
contact-registry	Do not attempt to contact the registry, even if one is configured

## 5 SAML PUSH support

### 5.1 Introduction

UCC supports the "SAML PUSH" mode of authentication. In this mode, the user is authenticated by an attribute assertion signed by a trusted third party. The third party is a Virtual Organisation (VO), SAML 2.0 server such as UVOS or VOMS.

### 5.2 Basic usage

First, retrieve an attribute assertion from VO server, scoped for a particular group you are in, and save it into a file.

```
ucc save-attributes -J https://uvos.example.com:2443 -G /vo/group -O assertion.xml
```

If the VO server has been registered in the registry, you may use "auto" instead of server's URL (this is possible only in case of UVOS).

```
ucc save-attributes -J auto -G /vo/group -O assertion.xml
```

Then, you may use this assertion to get access to the server. Most UCC commands support this, for example you may connect to the server using this assertion.

```
ucc connect -A assertion.xml
```

You also may combine those two steps into one:

```
ucc connect -J https://uvos.example.com:2443 -G /vo/group
```

or assuming that UVOS server is in the registry:

```
ucc connect -J auto -G /vo/group
```

However note that using <<<-A>>> option is faster: attributes are read from a file, while in the combined scenario the attributes are fetched from a remote VO server prior to invocation of the intended UCC command.

### 5.3 Attribute filtering

You might want to obtain an assertion with only handful of attributes and their values. This is done by using attribute filters.

UCC can display a list of all attributes. Names, scopes, values and descriptions (if available) will be displayed.

```
ucc list-attributes -J auto
```

The last column (F) contains a letter Y next to all attribute values, that pass through attribute filters. There are two kinds of filters: inclusive filters (which specify what attributes should pass) and exclusive filters (which specify what attributes should be filtered out). If not specified, the default inclusive filter approves of all attributes and their values, and the default exclusive filter does not reject anything.

In this example an inclusive filter choose only those attributes, whose name contains the word "xlogin".

```
ucc list-attributes -J auto -I ".*xlogin.*"
```

The value for the <<<-I>>> option is a list of semicolon-separated Java regular expressions. The expressions that contain the equal sign are called name-value filters, those which do not are called name filters.

For example, you may choose all attributes with names that contain the word "login", and the "admin" value of all attributes containing the word "role" in their name.

```
ucc list-attributes -J auto -I ".*xlogin.*;.*role.*=admin"
```

You may also exclude attributes. For example, the following filter chooses all attributes but those ending with an "a".

```
ucc list-attributes -J auto -Q ".*a"
```

If you are content with the results of filtering, you can obtain an assertion only with the filtered attributes.

```
ucc save-attributes -J auto -G /vo/group -O assertion.xml -I ".* ↵  
xlogin.*;.*role.*=admin" -Q ".*a"
```

You can also use it directly.

```
ucc connect -J auto -G /vo/group -I ".*xlogin.*;.*role.*=admin" -Q ↵  
".*a"
```

## 5.4 Rules for multiple filters

Each attribute is considered as a set of name-value pairs, containing the name of attribute and the value of attribute.

A pair matches a name filter (i.e. a filter defined using a regular expression without equals sign) if the name of the attribute matches the regular expression.

A pair matches a name-value filter (i.e. a filter defined using a regular expression with equals sign) if the name of the attribute matches the left hand of the regular expression, and the value matches the right hand.

For each pair, the following steps are performed:

- if there is a matching exclusive name-value filter, the pair is rejected
- otherwise, if there is a matching inclusive name-value filter, the pair passes
- otherwise, if there is a matching exclusive name filter, the pair is rejected
- otherwise, if there is a matching inclusive name filter, or there are no inclusive name filters, the pair passes
- otherwise, the pair is rejected

The result of filtering are the attributes, containing only those values which passed (after filtering, empty attributes are discarded).



## 6 Running jobs

### 6.1 Introduction

The UCC can run jobs specified in a simple job description format Section 7 . In the following it is assumed that you have UCC installed Section 2 and tried some examples Section 3 .

For example, assume the file "myjob.u" looks as follows

```
{
  ApplicationName="Date",
  ApplicationVersion="1.0"
}
```

To run this through UCC, issue the following command

```
ucc run myjob.u
```

This will submit the job, wait for completion, download the stdout and stderr files, and place them in your default output directory. The run command has many options, to see all the possibilities use the built-in help:

```
ucc run -h
```

#### 6.1.1 Controlling the output location and file names

Output files will be placed in the directory given by the "-o" option, if not given, the current directory is used. Also, file names will be put into a subdirectory named as the job id, to prevent accidental overwriting of existing files. This behaviour can be changed using the "-b" option. When "-b" is given on the command line, no subdirectory will be created.

#### 6.1.2 Specifying the site

In the example above, a random site will be chosen to execute the job. To control it, you can use the "-s" option. This will accept the name of a target system. The target systems available to you can be listed by

```
ucc list-sites
```

#### 6.1.3 Accessing a job's working directory

Using the UCC's data management functions, the job working directory can be accessed at any time after job submission. Please see section Section 8 for details.

## 6.2 Options overview

The following options are available when running jobs (see also the general options overview in Section 4).

Table 8: Job submission options for UCC

Option (Short and long form)	Description
-a, --asynchronous	Run asynchronously
-b, --brief	Do not create a sub-directory for output files
-B, --broker	Select the type of resource broker to use (LOCAL or SERVORCH)
-d, --dryRun	Only show candidate sites, but do not submit the job
-s, --sitename <SITE>	Site where the job shall be run
-S, --schedule <Time>	Schedule the submission of the job at the given time
-j, --jsdl	Tell UCC that the job file is a JSDL document
-o, --output <Output_dir>	Directory for any output produced (default is the current directory)
-O, --stdout <stdout_name>	specify a name for the exported standard out (by default: <i>stdout</i> )
-E, --stderr <stderr_name>	specify a name for the exported standard error (by default: <i>stderr</i> )

## 6.3 Resource brokering

If no site is specified upon submission, UCC will select a matching site, where the requirements (resources, application and execution environments) are met.

There are two types of brokers available, which can be selected using the "-B" or "--broker" option.

- LOCAL : brokering is done by UCC itself
- SERVORCH (default) : brokering is done by the "service orchestrator" component of the UNICORE Workflow system.

The SERVORCH broker requires an accessible UNICORE workflow system in version 6.6.0 or later. If this is not available, UCC will fall back to the LOCAL broker. As usual, you can set this option in your UCC preferences file.

## 6.4 Processing jobs asynchronously

In case of long-running jobs, you will want to run the job asynchronously, i.e. just submit the job, stage in any files and start it, in order to get the results later. UCC supports this, of course. The basic idea is that when submitting a job in asynchronous mode, a job descriptor file is written that contains the job's address, and any information about export files.

### 6.4.1 Asynchronous submission

Use the "-a" flag when submitting a job

```
ucc run -a <job file>
```

This will submit the job, stage-in any local files, start the job and exit. A job descriptor file (ending in ".job") will be written to your configured output directory.

### 6.4.2 Get the status of particular jobs

The command

```
ucc job-status <job_desc> <job_desc_2> ...
```

will retrieve the status of the given jobs. If not given on the command line, a job ID will be read from the console.

### 6.4.3 Download results

To get stdout, stderr and other files you have marked for export in your [job description](#), do

```
ucc get-output -o <outdir> <job_desc>
```

Here, the option "-o" specifies the directory where to put the output, by default the current directory is used. As before, the job address can also be read from the console.

### 6.4.4 Referencing a job by its EPR (Endpoint reference)

In case you want to check on a job not submitted through UCC, or in case you do not have the job descriptor file any more, you can also refer to a job given its EPR. For example, the "list-jobs" command will produce a list of all job EPRs that you can access.

Note that in this case UCC will only retrieve stdout and stderr files. To download other result files, you'll have to use the datamovement functions described in [Section 8](#).

### 6.4.5 Uploading and executing an executable

To upload and execute a file on a remote server, you might need a small helper script to make the uploaded file executable and run it:

```
#!/bin/sh
chmod +x myapp
./myapp
```

Your ucc job description would then look as follows

```
{
  Executable: "/bin/sh",
  Arguments: ["helper.sh"],
  Imports: [
    {From: "helper.sh", To: "helper.sh"},
    {From: "myapp", To: "myapp"},
  ],
}
```

---

**Note**

Recent servers do not require this trick, and you should be able to directly execute your uploaded executable without requiring a helper script.

---

### 6.4.6 Scheduling job submission to the batch system

Sometimes a user wishes to control the time when a job is submitted to the batch queue, for example because she knows that a certain queue will be empty at that time.

---

**Note**

This feature only works with server release 6.4.0 or higher.

---

To schedule a job, you can either use the "-S" option to the ucc "run" command:

```
ucc run -S "12:24" ...
```

Alternatively, you can specify the start time in your job file using the "Not before" key word

```
{
  Not before: "12:30",
}
```

In both cases, the specified start time can be given in the brief "HH:mm" (hours and minutes) format shown above, or in the full ISO 8601 format including year, date, time and time zone:

```
{  
  
  Not before: "2011-12-24T12:30:00+0200",  
  
}
```

## 6.5 Executing a command

If you just want to execute a simple command remotely (i.e. without data staging, resource specifications etc), you can use the "exec" command.

This will run the given command remotely (similarly to "ssh"), and print the output to the console. You can specify the site with the "-s" option. If you do not specify the site, a random site will be chosen.

UNICORE will run the command on the login node, it will not be submitted to the batch system.

For example, try

```
ucc exec /bin/date
```

Watch out to properly escape any arguments, in order not to interfere with the arguments to UCC.

## 7 Job description format

UCC uses a simple format that allows you to specify the application or executable you want to run, arguments and environment settings, any files to stage in from remote servers or the local machine and any result files to stage out.

A number of sample files can be found in the "samples" directory of your UCC distribution. (on Linux, check also /usr/share/unicore/ucc/samples)

The format used is called **JSON**, and contains comma-separated key-value mappings, where the values can be simple strings, or lists of values, or maps. String values should be placed in "quotes". Comments are (inofficially) possible using the "#" hash character, as in Unix shell scrips.

Each JSON file must begin and end with curly braces "{ ... }". Several complete job samples can be found in the "samples" directory of the distribution.

---

### Note

Note: quotes "" are needed around the keys and values in case special characters (like : or /) appear, if in doubt use quotes!

---

To view an example job showing most of the available options, simply run

```
ucc run -H
```

(most of the options shown are not mandatory, of course)

---

**Note**

You may alternatively specify jobs in the JSDL format that is used internally in UNICORE. To do this, run UCC with the "-j" option.

---

Usually, a UCC job file describe a single batch job on the target system. However there is a UNICORE feature called "parameter sweep" which leads to the creation of multiple batch jobs from a single "template" job. UCC can also create these "sweep jobs", as described in the relevant parts of the job description. Note that a "sweep job" still is treated as a single job by UNICORE. Sweep jobs are very useful if you need to run jobs that are highly similar, and only differ by a parameter setting or even by a different input file.

## 7.1 Site name

You can (optionally) specify on which site (if available) the job should be run.

```
Site: "DEMO-SITE",
```

If you do not specify anything UCC will select a site that will match your requirements (at least those that UCC checks for).

## 7.2 Specifying the application or executable

You can specify a UNICORE application by name and version, or using a (machine dependent) path to an executable file.

```
#using application name and version
{
  ApplicationName: "Date",
  ApplicationVersion: "1.0",
}
```

Note the comma-separation and the curly braces. To call an executable,

```
#using an executable
{
  Executable: "/bin/date",
}
```

### 7.3 Arguments and Environment settings

Arguments and environment settings are specified using a list of String values. Here is an example.

```
{
  Executable: "/bin/ls",
  Arguments: ["-l", "-t"],
  Environment: { PATH: "/bin", FOO: "bar" } ,
}
```

#### 7.3.1 Argument sweeps

To create a sweep over an Argument setting by replacing the value by a sweep specification. This can be either a simple list:

```
Arguments: [
  { Values: ["-o 1", "-o 2", "-o 3"] },
],
```

or a range:

```
Arguments: {
  "-o", { From: 1, To: 3, Step: 1 },
},
```

where the From, To and Step parameters are floating point or integer numbers.

### 7.4 Application parameters

In UNICORE, parameters for applications are often transferred in the form of environment variables. For example, the POVRay application has a large set of parameters to specify image width, height and many more. In UCC, you can specify these parameters in a very simple way using the "Parameters" keyword:

```
{
  ApplicationName: POVRay,

  Parameters: {
    WIDTH: 640,
    HEIGHT: 480,
    DEBUG: "",
  },
}
```

```
}
```

Note that an "empty" parameter (which does not have a value) needs to be written with an explicit empty string due to the limitations of the JSON syntax.

#### 7.4.1 Parameter sweeps

You can sweep over application parameters by replacing the parameter value by a sweep specification. The replacement can be either a simple list:

```
Parameters: {  
  WIDTH: { Values: [240, 480, 960] },  
},
```

or a range:

```
Parameters: {  
  WIDTH: { From: 240, To: 960, Step: 240 },  
},
```

where the From, To and Step parameters are floating point or integer numbers.

## 7.5 Job data management

In general your job will require data files, either from your client machine, or from some remote location. An important concept in UNICORE is the job's workspace (also called *Uspace*, which is the default location into which files are placed. The same applies to result files: by default, files will be downloaded from the job's workspace.

However, other remote storage locations are supported, too.

The remote location can be given as a full UNICORE URI, or using the more user friendly (but slower) "u6://" notation. Read more on remote locations in Section 8.

Local files can be given as an absolute or relative path; in the latter case the configured output directory will be used as base directory.

### 7.5.1 Importing files into the job workspace

To import files from your local computer or from remote sites to the job's working directory on the remote UNICORE server, there's the "Imports" keyword. Here is an example Import section that specifies three imports:



```
{
  Imports: [
    # import a local file into the job workspace
    { From: "/work/data/fileName", To: "uspaceFileName" },
    # import a set of PDF files into the Uspace
    { From: "/work/data/pdf/*.pdf", To: "/" },
    # import a remote file from a UNICORE storage
    { From: "u6://DEMO-SITE/Home/testfile", To: "otherUspaceFile" ←
      },
  ]
}
```

If for some reason it may happen that the local file does not exist, and you want the job to run anyway, there is a flag "FailOnError" that can be set to "false" :

```
Imports: [
  # do not fail on errors for this import:
  { From: "/work/data/fileName", To: "uspaceFileName", ←
    FailOnError: "false", },
]
```

Recent UNICORE servers (7.3 and later) support optimized input data handling. When multiple jobs or workflows use the same input data files, it may be possible to just create a symbolic link instead of physically copying the file. To tell UNICORE that this is OK, you must set a "ReadOnly" flag on the import:

```
Imports: [
  # tell UNICORE that the file can be sym-linked, if possible
  { From: "/work/data/fileName", To: "uspaceFileName", ReadOnly: ←
    "true", },
]
```

---

**Note**

UCC supports simple wild cards ("\*" and "?") for importing exporting LOCAL files, and for remote files if the server is version 7.x or later. Wildcards do not work for server-to-server imports and exports on 6.x servers.

---

### 7.5.2 Importing files into other storage locations

A UNICORE site may supports multiple storages (for example, a TMP or SCRATCH directory). To instruct the server to stage-in a file into such a storage, the "Filesystem" tag may be used.

For example to stage-in a file into SCRATCH space, the following Imports definition can be used:

```
Imports: [
  # import a file from a remote storage into the SCRATCH space on ←
  the target resource
  { From: "u6://DEMO-SITE/Home/work/data/fileName", To: "fileName ←
    ", Filesystem: "SCRATCH", },
```

### Sweeping over a stage-in file

You can also sweep over files, i.e. create multiple batch jobs that differ by one imported file. To achieve this, replace the "From" parameter by list of values, for example:

```
Imports: [
  { From: [ "u6://DEMO-SITE/Home/work/data/file1",
            "u6://DEMO-SITE/Home/work/data/file2",
            "u6://DEMO-SITE/Home/work/data/file3",
          ],
    To: "fileName", },
```

Note that only a single stage-in can be swept over in this way, and that this will not work with files imported from your local client machine.

### 7.5.3 Exporting result files from the job workspace

To export files from the job's working directory to your local machine or to some remote storage, use the "Exports" keyword. Here is an example Exports section that specifies two exports:

```
{
  Exports: [
    #this exports all png files to a local directory
    { From: "*.png", To: "/home/me/images/" },

    #this exports a single file to a to local directory
    #failure of this data transfer will be ignored
    { From: "error.log", To: "/home/me/logs/error.log", FailOnError ←
      : "false", },

    #this exports to a UNICORE storage
    { From: "stdout", To: "u6://DEMO-SITE/Home/results/myjob/stdout ←
      " },

  ]
```

```
}
```

As a special case, UCC also supports downloading files from other UNICORE storages using the Exports keyword:

```
{
  Exports: [
    #this exports a file from a UNICORE storage
    { From: "u6://DEMO-SITE/Work/somefile", To: "/home/me/somefile" ←
      },
  ]
}
```

The protocol to be used for imports and exports can be chosen using the "Preferred Protocols" entry, containing a space-separated list of protocols:

```
{
  Preferred protocols: "BFT RBYTEIO",
}
```

If not specified, BFT will be used.

#### 7.5.4 Specifying credentials for data staging

Some data staging protocols supported by UNICORE require credentials such as username and password. To pass username and password to the server, the syntax is as follows

```
{
  Imports: [
    { From: "ftp://someserver:25/some/file", To: "input_data",
      Credentials: { Username: "myname", Password: "mypassword" },
    },
  ]
}
```

and similarly for exports.

Servers 7.9.0 and later also support OAuth Bearer token for HTTP data transfers.

```
{
  Imports: [
    { From: "https://someserver/some/file", To: "input_data",
      Credentials: { BearerToken: "some_token" },
    },
  ]
}
```

You can leave the token value empty, `BearerToken: ""`, if the server already has your token by some other means.

### 7.5.5 Redirecting standard input

If you want to have your application or executable read its standard input from a file, you can use the following

```
Stdin: filename,
```

then the standard input will come from the file named "filename" in the job working directory.

## 7.6 Resources

A job definition can have a `Resources` section specifying the resources to request on the remote system. For example

```
Resources: {  
  
    # memory per node (bytes, you may use the common "K", "M" or "G ←  
    ")  
    Memory: 2G ,  
  
    # walltime (seconds, use "min", "h", or "d" for other units)  
    Runtime: 86400 ,  
  
    # Total number of requested CPUs  
    CPUs: 64 ,  
  
    # you may optionally give the number of nodes  
    #Nodes: 2 ,  
    # together with the CPUs per node  
    #CPUsPerNode: 32,  
  
    # Custom resources (site-dependent!)  
    StackLimitPerThread : 262144,  
  
    # Operating system  
    Operating system: LINUX, #MACOS, WINNT, ...  
  
    # Resource reservation reference  
    Reservation: job1234,  
  
}
```

Note that you can also specify a reservation reference if your batch system supports this and you have made a resource reservation.

## 7.7 Execution environments

To run a job in a special execution environments (as supported by the server), you can use the following syntax.

```
Execution environment: {
  Name: ...,
  Arguments: {
    ArgName1: "value1", ArgName2: "value2", ...
  },
  Options: [ ... ],
  Precommands: [ ... ],
  Postcommands: [ ... ],
  User precommand: "...",
  RunUserPrecommandOnLoginNode: "false",
  User postcommand: "...",
  RunUserPostcommandOnLoginNode: "true",
},
```

The user pre and post commands can be run either separately on the login node (default) or be placed in the job script. This is achieved using the "RunUserPrecommandOnLoginNode" and "RunUserPostcommandOnLoginNode" directives, which can be set to true or false.

## 7.8 Miscellaneous options

### 7.8.1 Selecting the remote login and/or group

In case you have multiple logins or Unix groups on the remote site mapped to the same credential, you can select the user name and/or group to use as follows

```
User name: yourlogin,
Group: yourgroup,
```

Hint: you can get a list of your logins/groups on the site by executing

```
ucc list-sites -s SITENAME -l
```

### 7.8.2 Specifying a project

If the system you're submitting to requires a project name for accounting purposes, you can specify the account (or project) you want to charge the job to using the "Project" tag:

```
Project: "my_project",
```

### 7.8.3 Specifying the user email for batch system notifications

Some batch systems support sending email upon completion of jobs. To specify your email, use

```
User email: foo@bar.org
```

Hint: if you want to explicitly switch off the email notification, use "NONE" as email value. This might be necessary because older UNICORE server versions try to use the email address from your certificate (if present).

### 7.8.4 Specifying the job name

The job name can be set simply by

```
Name: Test job
```

### 7.8.5 Specifying the status check interval for batch mode

Once a job is started, it is often not useful to check its status every few seconds, because the job might be running several minutes or more. Especially in batch mode it can reduce the load on the servers if the update interval is chosen longer. This can be achieved by using the following setting (this only affects batch mode!):

```
Update interval: 60, #only check once a minute (default is one ←  
second)
```

### 7.8.6 Specifying the "lifetime" of the job

If you want to specify a lifetime of the job, and not rely on the server default, you can use the lifetime attribute:

```
Lifetime: 12h, #sec, min, h, d
```

## 8 Data management functions

UCC offers access to all the data management functions in UNICORE. You can upload or download data from a remote server, initiate a server-to-server transfer, create directories and so on.

## 8.1 Specifying remote locations

Remote locations can be specified in two ways. The first way is to use a URI that includes protocol, storage server and filename, for example

```
BFT:https://mygateway:8080/SITE/services/StorageManagement?res= ↵
default_storage#/file
```

which specifies a file named "/file" on the storage instance "https://mygateway:8080/SITE/services/StorageManagement?res=" using the BFT protocol.

Paths are relative to the storage root, not the root of the actual file system.

This explicit format is sometimes inconvenient, so you can use a shorter, more intuitive format. This is also a URI, but you need to know only the name of the virtual site (target system), and the storage or job id. For example

```
unicore6://SITE/Home/file?protocol=PROTOCOL
```

or shorter

```
u6://SITE/Home/file?protocol=PROTOCOL
```

This will resolve the current user's "Home" storage at the target system named "SITE". Note that if you do not specify the protocol, the BFT protocol will be used as default.

You can also refer to a job's working directory on a given site. For this, you will need the unique ID of that job, which you can get for example using the *list-jobs* command. For example,

```
u6://SITE/1f3bc2e2-d814-406e-811d-e533f8f7a93b/outfile
```

refers to the file "outfile" in the working directory of the given job on the "SITE" target system. And

```
ucc ls -l u6://SITE/1f3bc2e2-d814-406e-811d-e533f8f7a93b/
```

will list the job's working directory.

It is also possible to refer to storage services that are registered in the registry using their name, for example

```
u6://SHARE/myfiles/a_file
```

can be used to refer to the shared storage named "SHARE" if it is registered in the registry.

Though convenient, the method using "unicore6://" is much slower, and will generate some network traffic. If you do a lot of operations on the same resource, you should use the *resolve* command to find out the URI of the resource, and use that later.

### 8.1.1 The resolve command

This will figure out the "real" address for a "unicore6://" URL as defined above.

```
ucc resolve u6://SHARE/
```

## 8.2 Data movement

### 8.2.1 cp

The *cp* command is a generic command for copying source file(s) to a target destination, where source and target can be remote locations or files on the local machine. Wild card characters \* and ? are supported.

Examples for client-server transfers:

```
ucc cp data/*.pdf u6://SHARE/pdfs
ucc cp u6://SHARE/test.pdf .
ucc cp u6://SHARE/test1 u6://SHARE/test3 data/
```

The "-R" option allows to choose whether subdirectories are to be copied, too.

The "-X" option allows to resume a previous transfer. Missing data will be appended to an existing target file (if the chosen protocol supports it).

Examples for server-server transfers:

```
ucc cp u6://SHARE/*.pdf u6://Demo-SITE/Home/
```

For server-to-server transfers, the *cp* command supports several additional options.

The "-S" option allows to schedule a transfer for a certain time. For example

```
ucc cp -S "23:00" u6://SHARE/*.pdf u6://Demo-SITE/Home/
```

The format is simply "HH:mm" (hours and minutes). Alternatively you can give the time in the full ISO 8601 format including year, date, time and time zone:

```
ucc cp -S "2011-12-24T12:30:00+0200" ...
```

Another useful option is "-a" which will execute the server-server transfer asynchronously, i.e. the client will not wait for the transfer to finish.

### 8.2.2 copy-file-status

This will print the status of the given data transfer. As argument, it expects a file name containing the transfer reference, or directly the reference.

Example (for Unix) which captures the reference into a shell variable:



```
export ID=$(ucc cp -a u6://OTHER-SITE/Home/test.txt u6://DEMO-SITE/ ↵  
    Home/test.txt)  
ucc copy-file-status $ID
```

### 8.2.3 Specifying the file transfer protocol

To use a different protocol from the default BFT, you can use the "-P" option to specify your preferred protocol. UCC will try to match them with the capabilities of the storage and use the first match. Your preferred protocol can also be listed in your preferences file using the "protocols" key:

```
protocols=UFTP
```

---

#### Note

If necessary, you can specify additional filetransfer options in your preferences file as well. For example, to use the UFTP protocol you may need to specify the client host address and the number of parallel streams explicitly:

```
uftp.client.host=your_client_ip_address  
uftp.streams=2  
# encrypt data (at the cost of performance)  
uftp.encryption=true  
# compress data  
uftp.compression=true
```

Use the special value "all" to enable all available client IP addresses for UFTP.

```
uftp.client.host=all
```

You can also override the UFTP server host, which can be useful in case the UFTP server is accessible via multiple network interfaces:

```
uftp.server.host=myhost.com
```

UCC will try to use reasonable defaults for any missing parameters.

---

## 8.3 General commands

### 8.3.1 mkdir

This will create a directory (including required parent directories) remotely.

Example

```
ucc mkdir u6://DEMO-SITE/Home/testdirectory/data/pdfs
```

### 8.3.2 rm

This will remove a file or directory remotely. By default, UCC will ask for a confirmation. Use the "--quiet" or "-q" option to disable this confirmation (e.g. when using this command in scripts).

Example

```
ucc rm u6://DEMO-SITE/Home/testdirectory/data/pdfs
```

### 8.3.3 rename

This will rename/move a remote file (on the same storage).

Examples

```
ucc rename u6://DEMO-SITE/Home/testdirectory/data/foo1.pdf / ↔  
testdirectory/data/foo2.pdf
```

will rename the file "foo1.pdf" to "foo2.pdf", leaving it in the same directory.

```
ucc rename u6://DEMO-SITE/Home/testdirectory/data /testdirectory/ ↔  
data_new
```

will rename the "data" directory.

### 8.3.4 stat

This command shows full information on a certain file or directory. Add the "-m" flag to also print user-defined metadata.

Example

```
ucc stat -m u6://DEMO-SITE/Home/testdirectory/foo.txt
```

## 8.4 Finding data

### 8.4.1 ls

This will list a remote directory. Useful options are: "-l" (detailed output), "-H" (human-friendly) and "-R" (recurse). Example:

```
ucc ls u6://DEMO-SITE/Home -l -H
```

If the storage supports metadata, you can get the metadata of a single file using "ls -l -m":

```
ucc ls u6://DEMO-SITE/Home/.bashrc -l -m
```

### 8.4.2 find

This command is similar to the well-known Unix utility, however much less powerful. It allows to do recursive listings and retrieve files matching certain conditions. Currently only "name match" is available. For example to get all PDF files on a storage,

```
ucc find -r -l u6://DEMO-SITE/Home/ -N .pdf
```

The *find* command is currently implemented synchronously, and may thus run into a network timeout when it takes too long. This limitation will be overcome in future versions of this command.

## 8.5 Deprecated commands

While the old *get-file*, *put-file* and *copy-file* are still available, you should use the simpler *cp* instead.

## 8.6 Using the StorageFactory service

UNICORE sites may allow users to dynamically create storage resources, which even can be linked to special back-end systems like Apache HDFS, iRODS, or cloud storage like Amazon S3.

You can find out if there are sites supporting this "StorageFactory" service either by running the *system-info -l* command, or better using

```
$> ucc create-storage -i
```

This will list the available StorageFactory services and also show which types of storage are supported and how much space is left on each of them.

UCC supports creating storages via the *create-storage* command. The simple

```
$> ucc create-storage
```

will create a new storage resource using the default storage type at some site.

Usually you want to control at least where the storage is created. Additionally, the type of storage and some parameters can be passed to UCC.

As an example, creating a storage of type "S3" would look like this

```
$> ucc create-storage -t S3 accessKey=... secretKey=...
```

You can also read parameters from a file. Say you have your S3 keys in a file *s3.properties*, then you can use the following syntax

```
$> ucc create-storage -t S3 @s3.properties
```

You can also mix this with the normal *key=value* syntax, or mix it like this:

```
$> ucc create-storage -t S3 accessKey=@s3.accessKey secretKey=@s3. ↵  
secretKey
```

The last version *key=@file* causes just the value to be read from the named file.

## 9 Metadata management functions

UCC offers a simple interface to access the metadata management service in UNICORE.

### 9.1 Basics

The metadata functions are all accessed via a single UCC command `metadata`. The actual operation to be performed is given with the "-C" (i.e. "command") option.

The storage to be operated upon is given using the "-s" option, alternatively the "-m" option can be used to directly give the metadata service URL.

In addition to the URL, the name of the target file on the storage is required.

Metadata is represented in JSON format. The metadata operations usually read metadata from a file (or write results to file), which is specified using the "-f" option.

In the following examples, `<STORAGE>` denotes the URL of a storage capable of handling metadata.

### 9.2 Available commands

#### 9.2.1 creating metadata

To create metadata, a file in JSON format is required containing key-value pairs. For example, edit the file "meta.json" to contain:

```
{  
  foo: bar  
}
```

Say we have a file "test" on our storage, then you can create metadata as follows

```
ucc metadata -C create -f meta.json -s <STORAGE> /test
```

If you now look at the file with "ls -l -m",

```
ucc ls -l -m <STORAGE>#/test
```

you should get something like this:

```
-rw-          3344 2011-06-27 22:32 /test
{
  "foo": "bar",
  "resourceName": "/test"
}
```

### 9.2.2 reading metadata

Apart from the "ls -l -m" used above, there is also an explicit "read" command, which can write the metadata to a file as well.

```
ucc metadata -C read -s <STORAGE> /test -f out.json
```

The "-f" option is optional.

### 9.2.3 updating metadata

Using update, the given metadata is merged with any existing metadata. Say we have a file x.json containing:

```
{
  x: y
}
```

we can append this to the existing metadata

```
ucc metadata -C update -s <STORAGE> /test -f x.json
```

Check that the metadata has indeed been appended.

### 9.2.4 deleting metadata

Explicitly deleting is also possible:

```
ucc metadata -C delete -s <STORAGE> /test
```

Check that the metadata has indeed been deleted.

### 9.2.5 searching

Searching requires a search string (according to the rules of Apache Lucene), and is triggered by the "search" command:

```
ucc metadata -C search -q "foo" -s <STORAGE> /
```

### 9.2.6 triggering metadata extraction

To trigger the extraction of metadata on the server, use the "start-extract" command:

```
ucc metadata -C start-extract -s <STORAGE> /
```

In this case the "/" denotes the base path from which to start the extraction process. The extraction process is asynchronous, so a "Task" service address will be returned which can be used to monitor the extraction process using the "wsrf getproperties" command.

## 10 Workflows

### 10.1 Introduction

UCC supports the UNICORE workflow system and allows to submit workflows to the workflow engine or single jobs to the service orchestrator (broker).

The workflows are executed server-side, and UCC is used only for submitting, managing data and getting results.

### 10.2 Command overview

The following commands are provided. More details and examples follow below.

- `workflow-submit` : submit a workflow file
- `workflow-control` : abort or resume a running workflow
- `list-workflows` : list information about workflows
- `broker-run` : use the service orchestrator to find matching sites or run single jobs

### 10.3 Basic use

To check the availability of workflow services, issue the following command

```
ucc system-info -l
```

This should show at least an accessible workflow engine and service orchestrator.

The distribution contains some example workflow files in the <[?]> directory that you can edit and submit.

```
ucc workflow-submit yourworkflow.swf
```

which will submit the workflow and print the address of the workflow to standard output. To get the workflow status,

```
ucc list-workflows <workflow_address>
```

To list all your workflows, you can use the `<[?]>` command without an explicit workflow address

```
ucc list-workflows -l
```

## 10.4 Managing workflow data

During workflow execution, data files will be produced that the workflow system will move to a storage location that is accessible for the individual jobs. Usually this location is created automatically by UCC before the workflow is submitted, using a special UNICORE service called a storage factory. If you want to influence this decision, UCC allows to select the storage factory to be used via the `-f` option to the `workflow-submit` command:

```
ucc workflow-submit -f <factory-url> <workflow-file>
```

You can check the available factories with the `system-info` command. If not specified, UCC will use the first storage factory it finds in the registry.

In general, storages that are dynamically created will be deleted when the workflow is deleted. To persistently store data, you need to make sure to export important result files to a persistent location (e.g. your Home on UNICORE site, or a persistent storage).

Alternatively you can directly specify a storage URL, either using the convenient `"u6://..."` notation, or as a real network URL:

```
ucc workflow-submit -S u6://MY-SITE/Home <workflow-file>
```

```
ucc workflow-submit -S https://my-gateway/SITE/services/ ↵  
StorageManagement?res=myuser-Home <workflow-file>
```

### 10.4.1 Importing local data for use by a workflow

If you have local files that need to be imported before starting the workflow, you have to specify this using a normal UCC job file that contains only an `"Imports"` section:

```
{  
  #stage-in specification  
  
  Imports: [  
    {From: local-file.sh, To: "c9m:${WORKFLOW_ID}/input.sh"}  
  ],  
}
```

When submitting the workflow, add the "-u <filename>" option to specify the imports file.

This will cause UCC to copy the local file "local-file.sh" to the workflow storage space. You can refer to this file in your workflow using the "global" name "c9m:...", say in a script activity:

```
....
  <jsd1:DataStaging>
  <jsd1:CreationFlag>overwrite</jds1:CreationFlag>
  <jsd1:FileName>input.sh</jds1:FileName>
  <jsd1:Source>
    <jsd1:URI>c9m:${WORKFLOW_ID}/input.sh</jds1:URI>
  </jds1:Source>
</jds1:DataStaging>
....
```

The workflow system will resolve the name at runtime and your file will be used. This allows you to group your files by workflow ID.

#### 10.4.2 Downloading output files

You can use the usual `get-file` command to download files using the "global IDs" used by the workflow engine. Hint: the `workflow-info` command will list the files that are produced by the workflow.

#### 10.4.3 Workflow templates

If the workflow's `Documentation` section contains `Argument` definitions, the corresponding replacement will be done by reading parameter values from the `.u` file. These so-called workflow templates can be a very simple and safe way to make adjustments in complex workflows # before submission. As an example, consider the following simple example workflow

```
<s:Workflow xmlns:s="http://www.chemomentum.org/workflow/simple"
  xmlns:jds1="http://schemas.ggf.org/jds1/2005/11/jds1"
  xmlns:jds1-p="http://schemas.ggf.org/jds1/2005/11/jds1- ↵
    posix"
  xmlns:jds1-u="http://www.unicore.eu/unicore/jds1- ↵
    extensions">

  <s:Documentation>
  <!-- template parameter definitions -->
  <jds1-u:Argument>
    <jds1-u:Name>UC_PRECISION</jds1-u:Name>
    <jds1-u:ArgumentMetadata>
      <jds1-u:Type>choice</jds1-u:Type>
      <jds1-u:ValidValue>low</jds1-u:ValidValue>
      <jds1-u:ValidValue>medium</jds1-u:ValidValue>
      <jds1-u:ValidValue>high</jds1-u:ValidValue>
```



```

    </jsdl-u:ArgumentMetadata>
  </jsdl-u:Argument>
</s:Documentation>

<s:Activity Id="date" Name="JSDL">
  <s:JSDL>
    <jsdl:JobDescription>
      <jsdl:Application>
        <jsdl:ApplicationName>ComputeIt</jsdl:ApplicationName>
        <jsdl-p:POSIXApplication>
          <jsdl-p:Argument>--precision ${UC_PRECISION}</jsdl-p: ↵
            Argument>
        </jsdl-p:POSIXApplication>
      </jsdl:Application>
    </jsdl:JobDescription>
  </s:JSDL>
</s:Activity>
</s:Workflow>

```

This defines a parameter `UC_PRECISION` which is used as application argument in a job. The `ArgumentMetadata` follow the same specification as application metadata in the UNICORE IDB. The UCC will now try to take the value for `UC_PRECISION` from the supplied `.u` file. As an example, the `.u` file could contain

```

{
  "UC_PRECISION" : "high",
}

```

Before submission, UCC will then replace this value, resulting in the workflow

```

<s:Workflow xmlns:s="http://www.chemomentum.org/workflow/simple"
  xmlns:jsdl="http://schemas.ggf.org/jsdl/2005/11/jsdl"
  xmlns:jsdl-p="http://schemas.ggf.org/jsdl/2005/11/jsdl- ↵
    posix"
  xmlns:jsdl-u="http://www.unicore.eu/unicore/jsdl- ↵
    extensions">

  <s:Documentation>
  <!-- ... -->
  </s:Documentation>

  <s:Activity Id="date" Name="JSDL">
    <s:JSDL>
      <jsdl:JobDescription>
        <jsdl:Application>
          <jsdl:ApplicationName>ComputeIt</jsdl:ApplicationName>
          <jsdl-p:POSIXApplication>
            <jsdl-p:Argument>--precision high</jsdl-p:Argument>
          </jsdl-p:POSIXApplication>
        </jsdl:Application>
      </jsdl:JobDescription>
    </s:JSDL>
  </s:Activity>
</s:Workflow>

```

```
</jsdl:Application>
  </jsdl:JobDescription>
</s:JSDL>
</s:Activity>

</s:Workflow>
```

As this is just string replacement, variables can be used anywhere in the workflow, including application definitions, file staging etc.

## 11 Batch processing

The *batch* command allows you to run many jobs without having to start UCC each time. You can control how many jobs should go to which site. This allows efficient job processing, while putting some load on the client machine. If you need to take the client offline, you should consider using the workflow system instead, which also allows efficient high-throughput processing.

Assume you have a bunch of jobs in UCC's job description format (Section 7) stored in a directory *jobs*. The output should go to a directory *out*. You can run them all through UCC using a single invocation as follows:

```
ucc batch -i jobs -o out
```

As job files, UCC will accept files ending in ".u", ".jsdl" or ".xml".

### 11.1 Options

You can run in "follow" mode, where UCC will watch the input directory, and will process new files as they arrive.

```
ucc batch -f -i jobs -o out
```

UCC can also process JSDL files, to batch-process these, use the "-j" option:

```
ucc batch -j -i jobs -o out
```

### 11.2 Performance tuning options

Getting the most performance out of UCC and the UNICORE installation can be a challenging task. Sending too many jobs to a site might decrease throughput, sometimes the client machine can be the limiting factor, etc.

You should experiment a bit to get the best performance for your specific setup. UCC has many options available for tuning. Here is an overview.

Table 9: Tuning options for the UCC batch mode

Option (short and long form)	Description
-K, --keep	Do not delete finished jobs on the server. By default, finished jobs are destroyed.
-m, --max <MaxRunningJobs>	Limit on jobs submitted by UCC at one time (default: 100)
-t, --threads <NumThreads>	Number of threads to be used for processing (default: 4)
-u, --update <UpdateInterval>	Minimum time in milliseconds between status requests on a single job (Default: 1000)
-R, --noResourceCheck	Do not check if the necessary application is available on the target system (will increase performance a bit)
-X, --noFetchOutcome	Do not fetch standard output and error
-S, --submitOnly	Only submit the jobs, do not wait for them to finish
-M, --maxNewJobs	Limit the number of job submissions (default: 100)
-s, --sitename	Specify which site to use
-W, --siteWeights	Specify a file containing site weights
-j, --jsdl	Assume jobs are in JSDL format instead of the default JSON .u files

### 11.3 Resource selection in batch mode

By default, the UCC batch mode will select a random site for running a job. You can modify the selection in different ways.

- using the "-s" option or a "Site: <sitename>," entry in the job file, you can specify the site directly
- use the "-W" option to specify a file containing site weights.

Say you have two sites where one site is a big cluster and the other a small cluster. To send more jobs to the big cluster, you can use the site weights file,

```
#example site weights file for use with "ucc batch -W ..."

BIG-CLUSTER = 100
SMALL-CLUSTER = 10

#send no jobs to this site
BAD-CLUSTER = 0

# set default weight (for any sites not specified here)
UCC_DEFAULT_SITE_WEIGHT = 10
```

This would tell UCC to send 10 times more jobs to the "BIG-CLUSTER" site, and send no jobs to the "BAD-CLUSTER". All other sites would get weight "10", i.e. the same as "SMALL-CLUSTER".

## 12 OGSA-BES functions

Assuming you have successfully installed UCC Section 2, this section shows you how to manage and monitor jobs on OGSA-BES services using UCC. The set of commands not only supports the UNICORE implementation, but may also work with implementations in other Grid middlewares compliant with OGF's [OGSA-BES specification](#).

### 12.1 OGSA-BES Setup

In UNICORE style, users are required to provide a Registry URL inside the preferences file. For BES users it is not always the case that an endpoint is advertised via a UNICORE Registry. Therefore, the configuration options allow user to modify this behaviour.

```
contact-registry=[true|false]
```

Users who wish to disable UCC calling the registry can set the "contact-registry" option to false. By default the "contact-registry" option is true.

When setting "contact-registry" to false, OGSA-BES users must provide at least one BESFactory URL using the following format.

```
bes.1=https://site1.com/services/BESFactory
bes.2=https://site2.com/services/BESFactory
bes.3=https://site3.com/services/BESFactory
bes.4=file:///tmp/bes-jugene.xml
bes.5=/tmp/bes-juropa.xml
...
```

If the "contact-registry" option is set to false and no OGSA-BES URL is specified, UCC will report an error. To use an XML endpoint reference (EPR) read from a file for contacting a BESFactory service, the contents of a EPR file must validate against the WS-Addressing's endpoint reference schema. See below the contents of the sample endpoint reference file,

```
<wsa:EndpointReference xmlns:wsa="http://www.w3.org/2005/08/ ↵
  addressing">
  <wsa:Address>https://localhost:8080/DEMO-SITE/services/BESFactory ↵
    ?res=default_bes_factory</wsa:Address>
</wsa:EndpointReference>
```

In the above XML snippet, under the "Address" tag, you must specify the URL of a target BESFactory service.

For the sake of convenience, here is an XML infoset representation taken from the [WS-Addressing specification](#):

```
<EndpointReference>
  <Address>xs:anyURI</Address>
  <ReferenceParameters>xs:any*</ReferenceParameters> ?
  <Metadata>xs:any*</Metadata>?
</EndpointReference>
```

## 12.2 Running and monitoring OGSA-BES jobs

UCC provides an easy to use command for submitting jobs on OGSA-BES complaint endpoints. To send a job read from a JSDL file,

```
ucc bes-submit-job -j hellompi.xml -s bes.3 -v
```

Alternatively, the job can be submitted using a BESFactory URL or endpoint reference file path.

```
ucc bes-submit-job -j hellompi.xml -s https://example3.com/services ←
/BESFactory -v
```

or

```
ucc bes-submit-job -j hellompi.xml -s file:///tmp/bes-jugene.xml -v
```

The JSON job description Section 7 can also be used, although only a subset of JSON constructs are supported for the OGSA-BES extensions.

Users can fetch the job status by specifying the descriptor (.job) file. This file is automatically generated after a successful execution of "bes-submit-job" command. Example:

Fetch job status example:

```
ucc bes-job-status jobid.job
```

Job can be terminated using a job descriptor file:

```
ucc bes-terminate-job jobid.job
```

To list BESFactory properties:

```
ucc bes-list-att -s bes.1
```

The above command will result in BESFactory's properties without jobs information. To see the list of the user's jobs on a BESFactory

```
ucc bes-list-job -s bes.1
```

### 12.3 Get OGSA-BES job outputs

The UCC OGSA-BES extension allow users to fetch job's output through the descriptor (.job) file. The command is called "bes-get-output". It supports downloading standard error and output files. The common usage scenario is, as soon as user issues the command, UCC will attempt to download the output files. Before downloading these files, UCC probes the job's status to check whether the job has successfully finished. If it is, then the output files will be fetched instantly, otherwise UCC waits until the job approaches completion.

Get job output example:

```
ucc bes-get-output jobid.job
```

Note that, at the moment the bes-get-output command only supports UNICORE based BES implementations. However, it can be used against other BES services which implement UNICORE specific server extensions such as BES-Activity and UNICORE's native StorageManagementService.

### 12.4 Enabling username/password authentication

Some BES implementations support authentication using username and password. To add username and password to the messages sent by UCC, the UCC preferences file must contain the following settings

```
#
# setup username/password
#
uas.security.out.handler.classname=de.fzj.unicore.bes.security. ↵
    UsernameOutHandler
de.fzj.unicore.bes.security.UsernameOutHandler.wsUserName=< ↵
    your_username>
de.fzj.unicore.bes.security.UsernameOutHandler.wsPassword=< ↵
    your_password>
de.fzj.unicore.bes.security.UsernameOutHandler. ↵
    wsActivateUsernameProfile=true
```

## 13 The UCC shell

If you want to run a larger number of UCC commands, the overhead of starting the Java VM or checking the registry may become annoying. For this scenario, UCC offers a "shell" that allows the user to enter UCC commands interactively.

It is usually started by

```
ucc shell
```

If you want to process a list of commands from a file instead of typing them, you can start the shell like this

```
ucc shell -f commandsfile
```

or on Unix you can use the redirection features

```
ucc shell < commandsfile
```

### 13.1 Exiting the shell

To exit, type `exit` or press CTRL-D

### 13.2 Changing property settings

To change a property setting in shell mode, you can use the `set` command. Without additional arguments, current properties are listed:

```
ucc>set
registry=https://...
output=/tmp
...
```

To set one or more properties, add space separated `key=value` strings:

```
ucc>set output=/work registry=https://....
```

You can also clear a property (set it to null) by using `unset`

```
ucc>unset registry
```

## 14 Sharing resources

Accessing UNICORE resources (jobs, storages, ...) is usually only possible when you "own" the resource or when there are special server-side policies in place that allow you access.

Starting with server version 7.3, UNICORE supports ACLs on a per-service instance basis. This means, that you can give other users access to your target systems, jobs, storages,

For example, you might have access to an S3 cloud storage via UNICORE, and you want to securely share data on this resource. Or, you want to allow others to check job status, or even allow them to abort jobs.

Note that to access actual **files** the permissions on file system level still need to match. Usually this is achieved by using Unix groups.

## 14.1 Editing ACLs

The ACLs are managed via the "share" command. Use the basic

```
ucc share <URL>
```

to show the current ACL for the given resource, where "URL" is the full WSRF service URL of the resource, e.g.

```
ucc share https://localhost:8080/DEMO-SITE/services/ ↔  
StorageManagement?res=demo-Home
```

To add an ACL entry, use

```
ucc share ACE1 ACE2 ... <URL>
```

where "ACE" is an access control entry expressed in a simple format:

```
[read|modify]:[DN|VO|GROUP|UID]:[value]
```

For example to give "modify" permission to a user whose UNIX user id on the target system is "test", you would use

```
ucc share modify:UID:test <URL>
```

To delete entries, use the "-d" option

```
ucc share -d modify:UID:test <URL>
```

To delete **all** entries, use the "-b" option

```
ucc share -b <URL>
```

## 14.2 Permission levels

The permissions controlled by ACLs are as follows

- read : access resource properties
- modify : perform actions e.g. job submission or creating a file export

Only the owner of a resource can edit the ACL or destroy the resource.

## 15 UCC for site administrators

UCC can be used for administrative and user support tasks, like checking server status, or getting the full details of a user job.



## 15.1 Security considerations

Usually, each UNICORE user has only access to his or her own resources (such as jobs). For administrative use, you will need to acquire administrator privileges. There are two ways to achieve this.

- create dedicated user credentials (e.g. a certificate) and map them to the role "admin" (in the XUADB, or whatever attribute source you are using). This method is recommended if you want to remotely administrate UNICORE/X.
- use the server keystore (of the UNICORE/X server you want to administrate) as UCC keystore. This will also give you administrator privileges. For this you will need to be logged on to the UNICORE/X server.

## 15.2 Admin commands

UCC has dedicated commands for accessing the "AdminService" of a UNICORE/X container. To get started, try

```
ucc admin-info -l
```

UCC will try to access the admin service on each available UNICORE/X server. For each server, a list of statistical and performance data will be listed.

It will also list the available admin commands for each server, with a short description of their parameters. For example, here is a sample output:

```
https://localhost:8080/DEMO-SITE/services/AdminService?res= ↵
  default_admin
  Services:
    TargetSystemFactoryService[1]
    ...
  Monitors:
    use.externalConnectionStatus.REST_UnitySAMLAuthenticator: OK
    use.security.overview: ServerIdentity: CN=Demo UNICORE/X,O= ↵
      UNICORE,C=EU;Expires: Thu Sep 09 12:01:19 CEST 2032; ↵
      IssuedBy: CN=Demo CA,O=UNICORE,C=EU
    ....
  Metrics:
    use.externalConnectionStatus.REST_UnitySAMLAuthenticator: OK
    use.rest.callFrequency: 0.016677196376660174
    ...
  Available commands:
    ShowJobDetails : parameters: jobID, [xnjsReference]
    ShowServerUsageOverview : parameters: [clientDN]
    ToggleResourceAvailability : 'resources' - comma separated list ↵
      of IDs
    ToggleJobSubmission : parameters: [message]
    ToggleBESJobSubmission :
```

To invoke a command, the "admin-runcommand" is used. It can take optional parameters.

### 15.2.1 Disabling/enabling job submission

For example, it is possible to disable/enable job submission to the server, using the *ToggleJobSubmission* command, which can take an optional message:

```
ucc admin-runcommand ToggleJobSubmission message="Maintenance"
```

The service will reply:

```
SUCCESS, service reply: OK - job submission is disabled
```

If a user now tries to submit, she will receive an error message on submission. Running the command again will re-enable the service.

```
ucc admin-runcommand ToggleJobSubmission message="Maintenance"
SUCCESS, service reply: OK - job submission is now enabled
```

### 15.2.2 Getting job details

To get the full job details (for example in user support), try

```
ucc admin-runcommand ShowJobDetails jobID=<unique_jobid>
```

for example

```
ucc admin-runcommand ShowJobDetails jobID=cdfdafc5-0274-464d-ac4a ↵
-463f46c942fa
SUCCESS, service reply: Job information for cdfdafc5-0274-464d-ac4a ↵
-463f46c942fa
{Info=Action ID      : cdfdafc5-0274-464d-ac4a-463f46c942fa
Action type         : JSDL
Status              : DONE (trans.: none)
Result              : SUCCESSFUL [Success.]
Owner               : CN=Demo User,O=UNICORE,C=EU
Exec. Definition: <JobDefinition xmlns="http://schemas.ggf.org/jsdl ↵
/2005/11/jsdl">
<JobDescription>
<JobIdentification>
<JobName>Date</JobName>
</JobIdentification>
<Application>
<jsdl:POSIXApplication xmlns:jsdl="http://schemas.ggf.org/ ↵
jsdl/2005/11/jsdl-posix">
<jsdl:Executable>/bin/date</jsdl:Executable>
</jsdl:POSIXApplication>
```

```

    </Application>
    <Resources/>
  </JobDescription>
</JobDefinition>
Orig. Definition: <JobDefinition xmlns="http://schemas.ggf.org/jsdl ↵
  /2005/11/jsdl">
  <JobDescription>
    <JobIdentification>
      <JobName>Date</JobName>
    </JobIdentification>
    <Application>
      <ApplicationName>Date</ApplicationName>
      <POSIXApplication xmlns="http://schemas.ggf.org/jsdl/2005/11/ ↵
        jsdl-posix"/>
    </Application>
  <Resources/>
</JobDescription>
</JobDefinition>
Processing context: de.fzj.unicore.xnjs.ems. ↵
  ProcessingContext@4308cff0
Application Info: AppInfo for Date 1.0
Job log:
Thu Sep 08 09:25:03 CEST 2016: Created with ID cdfdafc5-0274-464d- ↵
  ac4a-463f46c942fa
Thu Sep 08 09:25:03 CEST 2016: Created with type 'JSDL'
Thu Sep 08 09:25:03 CEST 2016: Client: Name: CN=Demo User,O=UNICORE ↵
  ,C=EU
Xlogin: uid: [schuller], gids: [addingOSgroups: true]
Role: user: role from attribute source
Security tokens: User name: CN=Demo User,O=UNICORE,C=EU
Consignor DN: CN=Demo User,O=UNICORE,C=EU
Delegation to consignor status: true, core delegation status: true
Message signature status: UNCHECKED
Client's original IP: 127.0.0.1
Thu Sep 08 09:25:04 CEST 2016: Using default execution environment.
Thu Sep 08 09:25:04 CEST 2016: No staging in needed.
Thu Sep 08 09:25:04 CEST 2016: Status set to READY.
Thu Sep 08 09:25:04 CEST 2016: Status set to PENDING.
Thu Sep 08 09:25:04 CEST 2016: Incarnated resources: [CPUsPerNode ↵
  =1.0, MemoryPerNode=2.68435456E8, ArraySize=1, Nodes=1.0, ↵
  ArrayLimit=64, CPUTime=3600.0]
Thu Sep 08 09:25:04 CEST 2016: Command is:
#!/bin/sh
# ....
chmod u+x /bin/date 2> /dev/null
rm -f /opt/unicore-servers/FILESPACE/cdfdafc5-0274-464d-ac4a-463 ↵
  f46c942fa//UNICORE_SCRIPT_EXIT_CODE
/bin/date

echo $? > /opt/unicore-servers/FILESPACE/cdfdafc5-0274-464d-ac4a ↵

```

```

-463f46c942fa//UNICORE_SCRIPT_EXIT_CODE

Thu Sep 08 09:25:04 CEST 2016: TSI reply: submission OK.
Thu Sep 08 09:25:04 CEST 2016: Submitted to classic TSI as [ ↔
schuller NONE] with BSSID=3269519504309
Thu Sep 08 09:25:13 CEST 2016: Exit code 0
Thu Sep 08 09:25:13 CEST 2016: Job completed on BSS.
Thu Sep 08 09:25:14 CEST 2016: Status set to DONE.
Thu Sep 08 09:25:14 CEST 2016: Result: Success.
Thu Sep 08 09:25:14 CEST 2016: Total: 10.01 sec., Stage-in: 0.05 ↔
sec., Stage-out: 0.00 sec., Datamovement: 0 %}

```

Thus you can get a full view of what the user submitted and what was executed.

### 15.3 Listing jobs, sites, ...

You can also use all normal UCC commands to access the server. Note however that due to the authentication and authorisation system in UNICORE, this may not always work as expected: the "admin" user might not have the required Unix permissions to access files, list directories etc.

The UCC commands that list server-side things (list-jobs etc) accept a filtering option, that can be used to limit the results of the operation. Filtering works on the XML resource properties of the resource in question.

Filtering is enabled by the "-f" or "--filter" option of the form

```
-f XMLNAME OPERATOR VALUE
```

where XMLNAME is the name of an XML Element from the WSRF resource properties document.

For example, to list all running jobs:

```
ucc list-jobs -f Status equals RUNNING
```

To list all jobs submitted on Nov 13, 2007:

```
ucc list-jobs -f SubmissionTime contains 2007-11-13
```

etc.

Table 10: Filtering options

Operator (long and short form)	Description
equals, eq	String equality (ignoring case)

Table 10: (continued)

<b>Operator (long and short form)</b>	<b>Description</b>
<code>notequals, neq</code>	String inequality (ignoring case)
<code>contains, c</code>	Substring match
<code>notcontains, nc</code>	substring non-match
<code>greaterthan, gt</code>	Lexical comparison
<code>lessthan, lt</code>	Lexical comparison

## 15.4 Low-level operations

UCC supports several low-level operations using the "wsrf" command.

To destroy a resource,

```
ucc wsrf destroy <Address>
```

To get a complete property listing (i.e. print the XML resource property document)

```
ucc wsrf getproperties <Address>
```

To extend the lifetime of a resource

```
ucc wsrf extend <Address> <Days>
```

These commands can be abbreviated, e.g. `+ucc wsrf d <Address>`

## 16 Scripting

UCC can execute Groovy scripts. Groovy (<http://groovy.codehaus.org>) is a dynamic scripting language similar to Python or Ruby, but very closely integrated with Java. The scripting facility can be used for automation tasks or implementation of custom commands, but it needs a bit of insight into how UNICORE and UCC work.

### 16.1 Script context

Your Groovy scripts can access some predefined variables that are summarized in the following table

Table 11: Variables accessible for scripts

variable	description	Java type
registry	A preconfigured client for accessing the registry	de.fzj.unicore.uas.client.IRegistryQuery
configurationProvider	Security configuration provider (keystore, etc)	de.fzj.unicore.ucc.authn.UCCConfigurationProvider
registryURL	the URL of the registry	java.lang.String
messageWriter	for writing messages to the user	de.fzj.unicore.ucc.MessageWriter
commandLine	the command line	org.apache.commons.cli.CommandLine
properties	defaults from the user's properties file	java.util.Properties

## 16.2 Examples

Some example Groovy scripts can be found in the samples/ directory of the UCC distribution.

Here is a script that will delete all your finished jobs (use at your own risk):

### Groovy example: delete all your jobs

```

/*
 * remove all jobs
 */

//import UNICORE/X client classes
import de.fzj.unicore.uas.client.*;
import eu.unicore.security.wsutil.client.authn. ←
    DelegationSpecification

//helper: kills a job with the given status (SUCCESSFUL, RUNNING, ←
    ...)

def kill(job, statuscode){
    if (job.status.toString() == statuscode) job.destroy()
}

//iterate over TSSs and remove all jobs
def lister = new de.fzj.unicore.uas.lookup.SiteLister(registry, ←
    configurationProvider)

lister.each {
    it.jobs.each{
        messageWriter.message "Job at "+it.address.stringValue

```

```
        securityProperties = configurationProvider. ←
            getClientConfiguration(it,DelegationSpecification. ←
                STANDARD)
        kill(new JobClient(it, securityProperties),"SUCCESSFUL")
    }
}
```

### Groovy example: list available storages

```
/*
 * list available storages
 */
import de.fzj.unicore.uas.client.*
import eu.unicore.security.wsutil.client.authn. ←
    DelegationSpecification
import javax.xml.namespace.QName

//porttype of storage service
def SMSSPORT=new QName("http://unigrids.org/2006/04/services/sms"," ←
    StorageManagement")

//method to extract storage name from a storage client
def findName(epr){
    securityProperties = configurationProvider.getClientConfiguration ←
        (epr,DelegationSpecification.STANDARD)
    sms=new StorageClient(epr, securityProperties)
    return sms.resourcePropertiesDocument.storageProperties. ←
        fileSystem.name
}

//list storages from registry
registry.listAccessibleServices(SMSSPORT).each {
    name=findName(it)
    messageWriter.message "Storage <"+name+"> at "+it.address. ←
        stringValue
}

//list storages attached to target systems
def lister = new de.fzj.unicore.uas.lookup.SiteLister(registry, ←
    configurationProvider)

lister.each {
    it.storages.each {
        name=findName(it)
        messageWriter.message "Storage <"+name+"> at "+it. ←
            address.stringValue
    }
}
```

## 17 Frequently asked questions

### 17.1 Configuration

#### 17.1.1 Do I really have to store my password in the preferences file? Isn't this insecure?

Putting the password in a file or giving it as a commandline parameter can be considered insecure. The file could be read by others, and the commandline parameters may be visible in for example in the output of the *ps* command. Thus, UCC will simply ask for the password in case you did not specify it.

#### 17.1.2 How can I enable more detailed logging?

UCC uses log4j, by default the configuration is done in `<UCC_HOME>/conf/logging.properties`. You can edit this file and increase the logging levels, choose to log to a file or to the console, etc.

#### 17.1.3 How can I set the HTTP connection timeout?

In your properties file, set

```
client.http.connection.timeout=<timeout in milliseconds>
```

#### 17.1.4 How can I log the SOAP messages sent and received by UCC?

In your properties file, set

```
#log messages  
client.messageLogging=true
```

which will log the messages on INFO level to the log file.

#### 17.1.5 How can I generate a proxy cert and add it to my message in order to use e.g. GridFTP?

In your properties file, add

```
#enable proxy cert out handler  
client.outHandlers=de.fzj.unicore.uas.security.ProxyCertOutHandler
```

which will add a handler that creates a proxy cert and adds it to the message.



## 17.2 Usage

### 17.2.1 Can I use multiple registries with UCC?

Yes. Simply use a comma-separated list of URLs for the "-c" option. However, you may only use a single key/truststore, so all registries (and sites listed in them) must accept the same security credentials.

### 17.2.2 Can I upload and execute my own executable?

Yes. Check Section 6.

### 17.2.3 Can I use UCC to list the contents of the registry?

Using the *wstrf* command, and the UNIX *grep* utility, this is very easy, for example

```
ucc wstrf getproperties https://localhost:8080/DEMO-SITE/services/ ↵  
Registry?res=default_registry | grep Address
```

will list the addresses of all services registered in the registry.

### 17.2.4 How can I use plain JSDL files instead of a .u JSON file for job submission?

Add the "-j" option when submitting a job.

### 17.2.5 I get strange errors related to security

Please read the general UNICORE FAQ on [www.unicore.eu](http://www.unicore.eu)[the UNICORE website] which contains descriptions of many common errors.

### 17.2.6 Are the JSDL documents and workflow documents validated?

The JSDL documents passed to the *run* command and to *submit-workflow* are validated, and any errors are logged. If you wish UCC to stop in case of validation errors, you need to set a property

```
ucc.validation.fail_on_errors=true
```