



UNICORE UFTPD SERVER

UNICORE Team

Document Version:	1.0.0
Component Version:	2.8.0
Date:	08 07 2020

Contents

1 UNICORE UFTP	1
1.1 UFTP features	1
1.2 How does UFTP work	2
2 Installation and use	3
2.1 Prerequisites	3
2.2 C library for switching user ID	3
2.3 Starting and stopping the UFTPD server	4
2.4 Configuration parameters	4
2.5 Protecting the Command socket	6
2.6 Firewall configuration	12
2.7 Logging	12
3 UNICORE Integration	13
3.1 Configuring the UFTP service	14
3.2 UFTPD servers with multiple interfaces	14
3.3 Enabling data encryption	15
3.4 Limiting bandwidth per transfer	15
3.5 Disabling SSL on the command port	15
3.6 Enabling "local" UFTP mode on the UNICORE/X server	15
4 Testing the UFTPD server	16
5 Performance measurement hints	16

This is the UFTPD user manual providing information on running and using the UNICORE UFTP server *uftpd*. Please note also the following places for getting more information:

UNICORE Website: <http://www.unicore.eu>

Support list: unicore-support@lists.sf.net

Developer's list: unicore-devel@lists.sf.net

UFTP issue tracker: <https://sourceforge.net/p/unicore/uftp-issues>

1 UNICORE UFTP

UFTP is a data streaming library and file transfer tool.

UFTP is best used using the *uftp* client-side application (available from <https://sourceforge.net/projects/unicore/files/Clients/UFTP-Client>), but is easily integrated into custom applications due to its FTP compliance.

It can be also integrated into UNICORE, allowing to transfer data from client to server (and vice versa), as well as providing data staging and third-party transfer between UFTP-enabled UNICORE sites.

A full UFTP server installation consists of two parts

- the "uftpd" file server
- either a UNICORE/X server, or a standalone authentication service.

This manual covers the UFTP file server "uftpd".

1.1 UFTP features

- dynamic firewall port opening using a pseudo FTP connection. UFTPD requires only a single open port.
- parallel input/output streams based on code from the JPARSS library, Copyright (c) 2001 Southeastern Universities Research Association, Thomas Jefferson National Accelerator Facility
- optional encryption of the data streams using a symmetric key algorithm
- optional compression of the data streams (using *gzip*)
- command port protected by SSL
- partial reads/writes to a file. If supported by the filesystem, multiple UFTP processes can thus read/write a file in parallel (striping)

- supports efficient synchronization of single local and remote files using the rsync algorithm
- integrated into UNICORE clients for fast file upload and download
- integrated with UNICORE servers for fast data staging and server-to-server file transfers
- standalone (non-UNICORE) client available
- data upload/download with *curl* or *ftp* possible in conjunction with the auth server
- written in Java

1.2 How does UFTP work

The UFTP file server, called *uftp*, listens on two ports (which may be on two different network interfaces):

- the command port receives control commands
- the listen port accepts data connections from clients.

The *uftp* server is "controlled" (usually by UNICORE/X) via the command port, and receives/sends data directly from/to a client machine (which can be an actual user client machine or another server). Data connections are made to the "listen" port, which has to be accessible from external machines. Firewalls have to treat the "listen" port as an FTP port.

A UFTP file transfer works as follows:

- the UNICORE/X server sends a request to the command port. This request notifies the UFTPD server about the upcoming transfer and contains the following information
 - a "secret", i.e. a one-time password which the client will use to authenticate itself
 - the user and group id which *uftp* should use to access files
 - an optional key to encrypt/decrypt the data
 - the client's IP address
- the UFTPD server will now accept an incoming connection from the announced IP address, provided the supplied "secret" matches the expectation.
- if everything is OK, an FTP session is created, and the client can use the FTP protocol to open data connections, list files, transfer data etc. Firewall transversal will be automatic using passive FTP (see below for alternatives like using a fixed port range).
- to access requested files, *uftp* attempts to switch its user id to the requested one prior to reading/writing the file. This uses a C library which is accessed from Java via the Java native interface (JNI). See also the installation section below.

[NOTE] .IMPORTANT SECURITY NOTE

The UNICORE UFTPD server is running with root privileges. Make sure to read and understand the section below on protecting the command socket. Otherwise, users logged on to the UFTPD machine can possibly read and write other user's files.

2 Installation and use

2.1 Prerequisites

- Java 8 (or later) runtime is required
- the server "listen" port needs to be accessible through your firewalls, declaring it an "FTP" port (FTP connection tracking). Alternatively a fixed range of open ports can be configured and used
- the UFTPD server needs access to the target file systems
- a server certificate for the UFTPD server is **STRONGLY** recommended for production use (see the section on SSL below)

A functional UFTP installation requires either a full UNICORE/X server or the *auth server*.

NOTE ON PATHS

The UNICORE UFTPD server is distributed either as a platform independent and portable tar.gz or zip bundle, or as an installable, platform dependent package such as RPM. Depending on the installation package, the paths to various files are different. If installing using distribution-specific package the following paths are used:

```
CONF=/etc/unicore/uftpd
SBIN=/usr/sbin
BIN=/usr/bin
LOG=/var/log/unicore/uftpd
LIB=/usr/share/unicore/uftpd/lib
```

If installing using the portable bundle, all UFTPD files are installed under a single directory. Path prefixes are as follows, where INST is the directory where UFTPD was installed:

```
CONF=INST/conf
SBIN=INST/bin
BIN=INST/bin
LOG=INST/log
LIB=INST/lib
```

These variables (CONF, SBIN, BIN and LOG) are used throughout the rest of this manual.

2.2 C library for switching user ID

It may be required to re-compile the libuftp-unix.so library on your system. This library uses the Java Native Interface (JNI). The version supplied with the distribution has been compiled on a 64bit Linux system. The folder LIB/native contains the required headers and C source files, as well as an exemplary makefile. Please edit the makefile, the following information is required:

- the base directory of our Java installation (JAVA_HOME)
- the location of platform-specific include files
- the location of the `uftp-<version>.jar` file (LIB)

Then, run "make install" to build the library, which will compile the code and install the library into the LIB folder. If any problems occur during this procedure, please consult UNICORE support.

2.3 Starting and stopping the UFTPD server

In the SBIN directory, start/stop and status scripts are provided:

- `unicore-uftp-start.sh` starts the server
- `unicore-uftp-stop.sh` stops the server
- `unicore-uftp-status.sh` checks the server status

The parameters such as server host/port, control host/port, and others are configured in the `CONF/uftp.conf` file

In a production scenario with multiple users, the uftp server needs to be started as root. This is necessary to be able to set the correct file permissions.

2.4 Configuration parameters

The following variables can be defined in the configuration file (`uftp.conf`):

```
SERVER_HOST      : the interface where the server listens for ←  
  client data  
                  connections  
  
SERVER_PORT      : the port where the server listens for ←  
  client data  
                  connections  
  
ADVERTISE_HOST  : Advertise this server as having the ←  
  following IP in the  
                  control connection. This is useful if the ←  
                  server is behind  
                  a NAT firewall and the public address is ←  
                  different from  
                  SERVER_HOST.
```

CMD_HOST : the interface where the server listens for control commands ↵

CMD_PORT : the port where the server listens for control commands ↵

SSL_CONF : File containing SSL settings for the command port ↵

ACL : File containing the list of server DNSs that are allowed access to the command port ↵

UFTPd_MEM : the maximum memory (Java heap size) allocated to the UFTPd server ↵

MAX_CONNECTIONS : the maximum number of concurrent control connections per client IP ↵

MAX_STREAMS : the maximum number of parallel TCP streams per connection ↵

BUFFER_SIZE : the size of the buffer (in kilobytes) for reading/writing local files ↵

PORT_RANGE : (optional) server-side port range in the form 'lower:upper' that will be used for data connections. By default, any free ports will be used. ↵
Example: set to '50000:50500' to limit the port range. ↵
NOTE ports in this range *must not* be used by other services! ↵

DISABLE_IP_CHECK : (optional) in some situations, the client IP can be different from the one that was sent to the UFTPd server. This will lead to rejected transfers. Setting this variable to a non-zero value will disable the IP check. Only the one-time password will be checked. ↵

```
UFTPD_KEYFILES      : (optional) list of files (relative to ↵
    current user's $HOME) where uftpd will
                        read public keys for authentication. ↵
                        List is separated by ":". This ↵
                        defaults
                        to ".ssh/authorized_keys"

UFTPD_NO_WRITE      : (optional) ":"-separated list of file ↵
    name patters that uftpd should
                        not write to
```

As usual if you set the `SERVER_HOST` to be "0.0.0.0", the server will bind to all the available network interfaces.

If possible, use an "internal" interface for the Command socket. If that is not possible, make sure the Command socket is protected by a firewall!

We **STRONGLY** recommend enabling SSL for the Command socket. Please refer to the next section.

2.5 Protecting the Command socket

Using SSL for the Command port ensures that only trusted parties (i.e. trusted UNICORE servers) can issue commands to the UFTPD server. To further limit the set of trusted users, an access control list (ACL) file is used.

In production settings where users can log in to the UFTPD server machine, SSL **MUST** be enabled to prevent unauthorized data access!

[NOTE] .IMPORTANT SECURITY NOTE

Without SSL enabled, users logged in to the UFTPD server can easily create exploits to read or write files with arbitrary user privileges (except *root*).

2.5.1 SSL setup

To setup SSL, you need a keystore containing the UFTPD server's credential, and a truststore containing certificate authorities that should be trusted. Keystore and truststore can be the same file.

The following properties can be set in the `CONF/uftpd-ssl.conf` file.

Table 1: 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.
<code>credential.keyPath</code>	string	-	Location of the private key if stored separately from the main credential (applicable for <i>pem</i> and <i>der</i> types only),
<code>credential.keyPassword</code>	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.
<code>credential.keyAlias</code>	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 2: 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.
truststore.directoryPath	filesystem 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.directoryEncoding	[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.directoryLocations	list of properties with a common prefix	*	List of CA certificates locations. Can contain URLs, local files and wildcard expressions. (<i>runtime updateable</i>)
--- Keystore type settings ---			

Table 2: (continued)

Property name	Type	Default value / mandatory	Description
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)
truststore.openssl	string	GLOBAL, EUGRIDPMA, GLOBUS, EUGRIDPMA, GLOBUS_REQUIRE, EUGRIDPMA_GLOBUS_REQUIRE, GLOBUS_REQUIRE, EUGRIDPMA_REQUIRE, EUGRIDPMA_AND_GLOBUS_REQUIRE, GLOBUS_REQUIRE_IGNORE]	In case 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 REQUIRE is checked (in the order defined by the property).
truststore.openssl	file system path	/etc/grid	Directory to be used for openssl truststore.
<i>--- Revocation settings ---</i>			
truststore.crlConn	integer	15	Connection timeout for fetching the remote CRLs in seconds (not used for Openssl truststores).

Table 2: (continued)

Property name	Type	Default value / mandatory	Description
truststore.crlDiskfileSystemPath	file system 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.crlLocalList	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>)
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.crlUpdateInterval	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.ocspCacheTime	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.ocspDiskfileSystemPath	file system path	-	If this property is defined then OCSP responses will be cached on disk in the defined folder.
truststore.ocspLocalList	list of responders *. <NUMBER> properties with a common prefix	.<NUMBER>	Optional list of local OCSP responders

Table 2: (continued)

Property name	Type	Default value / mandatory	Description
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 milliseconds.
truststore.revocationOrder	[CRL, OCSP, OCSP_CRL]	OCSP_CRL	Controls overall revocation sources order
truststore.revocationCheckAll	[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.

If the `credential.path` property is NOT set, SSL will be disabled.

[NOTE] .Backwards compatibility to previous versions

Existing configuration files with the `javax.net.ssl.*` properties used in UFTPD < 2.6 are still supported

2.5.2 ACL setup

The access control list contains the distinguished names of those certificates that should be allowed access.

The "ACL" setting in `CONF/uftpd.conf` is used to specify the location of the ACL file

```
export ACL=conf/uftpd.acl
```

The default ACL contains the certificate DN of the UNICORE/X server from the UNICORE core server bundle. In production, you need to replace this by the actual DNs of your UNICORE/X server(s) and UFTP Authentication server(s).

The ACL entries are expected in RFC2253 format. To get the name from a certificate in the correct format using `openssl`, you can use the following OpenSSL command:

```
$> openssl x509 -in your_server.pem -noout -subject -nameopt ↵  
RFC2253
```

The ACL file can be updated at runtime.

2.6 Firewall configuration

Note

Please consult the firewall documentation on how to enable an "FTP" service on your firewall (or operating system).

With Linux iptables, you may use rules similar to the following:

```
iptables -A INPUT -p tcp -m tcp --dport $SERVER_PORT -j ACCEPT  
iptables -A INPUT -p tcp -m helper --helper ftp-$SERVER_PORT -j ↵  
ACCEPT
```

where `$SERVER_PORT` is the `SERVER_PORT` defined in `uftpd.conf`. The first rule allows anyone to access port `$SERVER_PORT`. The second rule activates the iptables connection tracking FTP module on port `$SERVER_PORT`.

On some operating systems it may be required to load additional kernel modules to enable connection tracking, for example on CentOS:

```
modprobe nf_conntrack_ipv4  
modprobe nf_conntrack_ftp ports=$SERVER_PORT
```

2.7 Logging

UFTPD uses log4j, the same logging system as other UNICORE components. Logging is configured in the `CONF/logging.properties` file.

Note

You can change the logging configuration at runtime by editing the `logging.properties` file. The new configuration will take effect a few seconds after the file has been modified.

By default, log files are written to the the `LOG` directory.

For more info on controlling the logging we refer to the log4j documentation:

- [PatternLayout](#)

- [RollingFileAppender](#)
- [DailyRollingFileAppender](#)

Log4j supports a very wide range of logging options, such as date based or size based file rollover, logging different things to different files and much more. For full information on Log4j we refer to the publicly available documentation, for example the [Log4j manual](#).

2.7.1 Logger categories, names and levels

Logger names are hierarchical. Prefixes are used (e.g. "uftp.server") to which the Java class name is appended.

The logging output produced can be controlled in a fine-grained manner. Log levels in Log4j are (in increasing level of severity) TRACE, DEBUG, INFO, WARN, ERROR, and FATAL.

For example, to track the UFTP server's communication with clients in detail, you can set

```
log4j.logger.uftp.server=DEBUG
```

Here is a table of logger categories

Log category	Description
uftp	All UFTP logging
uftp.security	Security
uftp.server	Server code
uftp.client	Client code

Note

Please take care to not set the global level to TRACE or DEBUG for long times, as this may produce a lot of output.

2.7.2 Usage logging

Often it is desirable to keep track of the usage of your UFTP server. The server has a special logger category which logs information about finished jobs at INFO level. If you wish to enable this, set the level to INFO or higher:

```
log4j.logger.uftp.server.USAGE=INFO
```

3 UNICORE Integration

In UNICORE, UFTP can be used for uploading/downloading data from a server, as well as server-to-server transfers, where one UNICORE/X instance acts as UFTP client.

To enable a UNICORE/X server for the UFTP filetransfer the following settings have to be made.

3.1 Configuring the UFTP service

The UNICORE/X server needs to know the UFTPD server settings (i.e. host and port of both the command and the main listener socket). If the `coreServices.uftp.server.host` property is set, UFTP will be enabled and advertised to clients.

In the `uas.config` file, the following properties are mandatory for UFTP:

```
#
# REQUIRED parameters
#

# Listener (pseudo-FTP) socket
coreServices.uftp.server.host=...
coreServices.uftp.server.port=...

# Command socket
coreServices.uftp.command.host=...
coreServices.uftp.command.port=...

#
# Full path to the 'uftp.sh' client executable
# on the TSI login node
#
coreServices.uftp.client.executable=...
```

There are further configuration options to tune how UFTP behaves

```
#
# Optional parameters
#

# How many parallel streams to use per file transfer
coreServices.uftp.streams=1

# Limit the maximum number of streams per file transfer
coreServices.uftp.streamsLimit=4
```

3.2 UFTPD servers with multiple interfaces

If your UFTPD server is on a machine with multiple network interfaces, you can give a comma-separated list of host names, like so:

```
coreServices.uftp.server.host=net1.domain.org,net2.domain.org
```


The UFTP client in this case will try to connect to the hosts in the specified order, and will use the first "working" one.

3.3 Enabling data encryption

If you wish to encrypt the data sent/received by UFTPD (in data staging or server-to-server transfers), the following property can be used:

```
# enable data encryption
coreServices.uftp.encryption=true
```

This will by default encrypt data with a symmetric key using the Blowfish algorithm. This costs some performance due to the additional CPU load. Encryption only works in single-stream mode. Users can override this setting.

3.4 Limiting bandwidth per transfer

It is possible to limit the bandwidth that is consumed by a single UFTP transfer. It is given in bytes per second.

```
# limit transfer rate (bytes/second)
coreServices.uftp.rateLimit=10000000000
```

3.5 Disabling SSL on the command port

While not recommended, it may be sometimes useful to disable SSL for communicating with the UFTPD, e.g. while setting up and testing. To do this add a property in `uas.config` (no UNICORE/X restart is required)

```
coreServices.uftp.command.sslDisable=true
```

3.6 Enabling "local" UFTP mode on the UNICORE/X server

In case the UNICORE/X server has direct access to the target file system, and you're not using the UNICORE TSI, it can be an interesting option to run the UFTP filetransfer client code directly in the UNICORE/X server instead of passing it to the TSI. This means more load in the UNICORE/X process. It is also a way to use UFTP to/from a UNICORE/X server running on Windows. To enable local mode, edit `uas.config` and set

```
# enable local client mode
coreServices.uftp.client.local=true
```

4 Testing the UFTPD server

Testing as described in this section works only if SSL is not enabled. Therefore, you should run these tests as a non-root user. Enable SSL and restart the UFTPD server with root privileges once you are finished with these tests.

The UFTPD distribution contains two scripts that allow you to test the UFP functionality without using any external servers or clients. Making a data transfer involves two steps:

- invoke `uftp-job.sh` to "announce" an upcoming transfer to the UFTPD server
- invoke `uftp.sh` to initiate the actual transfer

Note, in case you installed from an RPM or DEB package, these files are located in `/usr/bin`.

The following shell commands "transfer" the file `.bashrc` to the `/tmp` directory.

Assuming you installed from RPM/DEB:

```
. /etc/unicore/uftpd/uftpd.conf
uftp-job.sh -c localhost -f ~/.bashrc -s true -x my_secret -n 2 -u ↵
    unicore -g unicore
uftp.sh -r -f /tmp/test -L $SERVER_PORT -l $SERVER_HOST -x ↵
    my_secret -n 2
```

This should create a file `/tmp/test` identical to `~/.bashrc`. Check the console output and the UFTPD log file `LOG/uftpd.log` in case of errors.

After the transfer finished, check that indeed

```
md5sum /tmp/test ~/.bashrc
```

gives the correct checksum for the newly created file.

It is also possible to enable encryption "manually", by appending `"-E <key>"` to the commands above, where "key" is a sequence of 12 characters (really a base64-encoded 64 bit key).

5 Performance measurement hints

You should use the `uftp` client to run performance tests, which contains many options such as the number of concurrent FTP connections, and can use `/dev/null` and `/dev/zero` as data source/sink.