



UNICORE UFTPD SERVER

UNICORE Team

| | |
|--------------------|------------|
| Document Version: | 1.0.0 |
| Component Version: | 2.1.1 |
| Date: | 26 09 2014 |

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 UFTP server | 4 |
| 2.4 | Configuration parameters | 4 |
| 2.5 | Protecting the Command socket | 5 |
| 2.6 | Firewall configuration | 7 |
| 3 | UNICORE Integration | 7 |
| 3.1 | Enabling the UFTP service | 7 |
| 3.2 | Enabling the UFTP protocol | 7 |
| 3.3 | Configuration options | 8 |
| 3.4 | UFTP servers with multiple interfaces | 8 |
| 3.5 | Enabling data encryption | 9 |
| 3.6 | Limiting bandwidth per transfer | 9 |
| 3.7 | Disabling SSL on the command port | 9 |
| 3.8 | Enabling "local" UFTP mode on the UNICORE/X server | 9 |
| 4 | Testing the UFTP server without UNICORE | 10 |
| 5 | Performance measurement hints | 10 |

This is the UFTP 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

1 UNICORE UFTP

UFTP is a data streaming library and file transfer tool.

It is integrated into UNICORE, allowing to transfer data from client to server (and vice versa), as well as providing data staging between UFTP-enabled UNICORE sites. UFTP can also be used independently from UNICORE, requiring a authentication server and a standalone UFTP client.

A full UFTP server installation consists of two parts

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

This manual covers the UFTP 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
- written in Java

1.2 How does UFTP work

The server part, 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 user's client machine or another UFTP enabled UNICORE 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
 - the client's IP address
 - the source/target file name
 - whether to send or receive data
 - a "secret", i.e. a string the client will send to authenticate itself
 - how many data connections will be opened
 - the user and group id for who to create the file (in case of send mode)
 - an optional key to encrypt/decrypt the data
- 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, the requested number of data connections from the client can be opened. Firewall transversal will be negotiated using a pseudo FTP protocol.
- the file is sent/received using the requested number of data connections
- to access the requested file, *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.

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 1.7 (or later) runtime is required
- the server "listen" port needs to be accessible through your firewalls, declaring it an "FTP" port
- 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)

We recommend using UFTP2 with UNICORE 7.0 or later. Older server versions will work, but new features like multi-file transfers or compression will not be available.

NOTE ON PATHS

The UNICORE UFTP 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 UFTP files are installed under a single directory. Path prefixes are as follows, where INST is the directory where UFTP 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 ↵

UFTPd_MEM : the maximum memory 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 ↵

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 ↵

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 users (rather, 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, SSL should ALWAYS be enabled.

2.5.1 SSL setup

IMPORTANT SECURITY NOTE

Without SSL enabled, users logged in to the UFTPD server can potentially read or write files with root privileges.

As usual for SSL, you need a keystore containing the UFTPD server's certificate, 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/udtpd-ssl.conf` file.

- `javax.net.ssl.keyStore` Keystore file containing the UFTPD key
- `javax.net.ssl.keyStorePassword` Keystore password
- `javax.net.ssl.keyStoreType` (optional) JKS or PKCS12, default is JKS
- `javax.net.ssl.trustStore` Truststore file containing the trusted CAs
- `javax.net.ssl.trustStorePassword` Truststore password

If the `javax.net.ssl.keyStore` property is NOT set, SSL will be disabled.

If the `javax.net.ssl.keyStorePassword` property is not set, the keystore password will be queried from the command line. If you want to use this, you will need to adapt the start script, avoiding the use of "nohup".

Please refer also to the UNICORE FAQ for keystore related questions.

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).

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
```


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
```

3 UNICORE Integration

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

3.1 Enabling the UFTP service

In the `wsrflite.xml` file, add a service definition for the UFTP filetransfer service:

```
<service name="FileTransferUFTP" wsrf="true" persistent="true">
  <interface class="de.fzj.unicore.uas.fts.FileTransfer" />
  <implementation class="de.fzj.unicore.uas.fts. ←
    FileTransferHomeImpl"/>
</service>
```

3.2 Enabling the UFTP protocol

For those storages where you want UFTP enabled, you need to add it to the list of supported protocols. This is done in the UNICORE/X main config file `uas.config`.

```
# supported protocols
coreServices.sms.protocols=BFT UFTP RBYTEIO SBYTEIO
```

Please check the UNICORE/X manual for further details.

3.3 Configuration options

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). In the `uas.config` file, set the following properties:

```
#
# 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=...

#
# Optional parameters
#

# How many parallel streams to use per file transfer
coreServices.uftp.streams=2

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

# File read/write buffer in kbytes
coreServices.uftp.bufferSize=128
```

3.4 UFTP servers with multiple interfaces

If your UFTP 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.

3.5 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.

3.6 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
```

3.7 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
```

3.8 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 Perl 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. To enable local mode, edit `uas.config` and set

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

4 Testing the UFTP server without UNICORE

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 UFTP distribution contains two scripts that allow you to test the UFP functionality without using UNICORE. 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
unicore-uftpd-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

To run performance tests, it is possible to read and write from/to pseudo files in `/dev`. Since UFTP needs to know the number of bytes to transfer, this can be given using a "pseudo" filename. For example, you could read from `/dev/zero_<number_of_bytes>` and write to `/dev/null`.

UFTP will treat any filename that starts with `/dev` and contains an underscore `"_"` character in this manner.

For example, to transfer a gigabyte of zeros to `/dev/null`,

```
. /etc/unicore/uftpd/uftpd.conf
unicore-uftpd-job.sh -c localhost -f /dev/zero_1000000000 -s true - ↵
  x my_secret -n 2 -u unicore -g unicore
uftpd.sh -r -f /dev/null -L $SERVER_PORT -l $SERVER_HOST -x ↵
  my_secret -n 2
```