



UNICORE GATEWAY

UNICORE Team

Document Version:	1.1.0
Component Version:	8.0.0
Date:	12 11 2019

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

Contents

1	Introduction	1
2	Installation	2
2.1	Installation from the core server bundle	2
2.2	Installation from a Linux package (rpm or deb)	2
3	Upgrading	3
4	Configuration	3
4.1	Java and environment settings: startup.properties	3
4.2	Configuring sites: connections.properties	3
4.3	Main server settings: gateway.properties	4
4.4	Certificate-less end-user access	4
4.5	security.properties	11
4.6	Logging	16
5	Using Apache httpd as a frontend	18
6	Using the Gateway for failover and/or loadbalancing of UNICORE sites	19
6.1	Configuration	19
6.2	Available Strategies	20
7	Gateway failover and migration	21
7.1	Gateway's migration	21
7.2	Failover and loadbalancing of the Gateway	21
8	Building the Gateway from source	21

This user manual provides information on running and using the UNICORE Gateway. Please note also the following places for getting more information:

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

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

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

1 Introduction

The Gateway is the entry point into a UNICORE site, routing HTTPS traffic to servers like UNICORE/X. It is installed in front of any networking firewall. It (optionally) authenticates incoming messages and forwards them to their intended destination. The Gateway receives the reply and sends it back to the client. In this way, only a single open port in a site's firewall has to be configured.

LIMITATIONS

This forwarding process only works for "most" HTTP requests, and is not a complete HTTP reverse proxy implementation. For example, it is not possible to run a full, complex web application like the UNICORE Portal "behind" the Gateway. Check the respective components's manual whether it can be run behind the Gateway.

In effect, traffic to a *virtual* URL, e.g. *https://mygateway:8088/Alpha* is forwarded to the real URL, e.g. *https://host1:7777*.

The mappings of virtual URL to real URL for the available sites are listed in a configuration file `connections.properties`. Additionally, the Gateway supports dynamic registration of sites.

The second functionality of the Gateway is (optional) authentication of incoming requests. Connections to the Gateway are made using SSL, so the Gateway can be configured to check whether the caller presents a certificate issued by a trusted authority. Information about the client is forwarded to services behind the Gateway in UNICORE proprietary format (as a SOAP or HTTP header).

The Gateway will forward the IP address of the client to the back-end server.

Last not least, the Gateway can be configured as a HTTP load balancer.

IMPORTANT NOTE ON PATHS

The UNICORE Gateway is distributed either as a platform independent and portable bundle (as a part of the UNICORE core server package) or as an installable, platform dependent package format such as RPM.

Depending on the installation method, the paths to various Gateway files are different. If installing using a distribution-specific package the following paths are used:

```
CONF=/etc/unicore/gateway
BIN=/usr/sbin
LOG=/var/log/unicore/gateway
```

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

```
CONF=INST/conf
BIN=INST/bin
LOG=INST/log
```

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

2 Installation

The UNICORE Gateway is distributed in the following formats:

1. As a part of platform independent installation bundle called UNICORE core server bundle. The UNICORE core server bundle is provided in two forms: one with a graphical installer and one with a command line installer.
2. As a binary, platform-specific package available currently for RedHat (Centos) and Debian platforms. Those packages are not tested on all possible platforms, but should work without any problems with other versions of similar distributions, e.g. SL6, Centos, or Fedora.

2.1 Installation from the core server bundle

Download the core server bundle from the UNICORE project website.

If you use the graphical installer, follow the on-screen instructions and do not forget to enable the Gateway checkbox when prompted.

If you use the console installer, please review the README file available after extracting the bundle. You don't have to change any defaults as the Gateway is installed by default.

2.2 Installation from a Linux package (rpm or deb)

Use your distribution's package manager to install.

3 Upgrading

The general update procedure is presented below, with possible variations:

1. Stop the old Gateway.
2. Update the server package. This step mostly applies for RPM/DEB managed installations. For Quickstart installation it is enough to replace the *.jar files with the new ones.
3. Start the newly installed Gateway.
4. Verify log file and fix any problems reported.

4 Configuration

The Gateway is configured using a set of configuration files, which reside in the CONF subdirectory.

4.1 Java and environment settings: startup.properties

This file contains settings related to the Java VM, such as the Java command to use, memory settings, library paths etc.

4.2 Configuring sites: connections.properties

This is a simple list connecting the names of sites and their physical addresses. An example is:

```
DEMO-SITE = https://localhost:7777
REGISTRY = https://localhost:7778
```

If this file is modified, the Gateway will re-read it at runtime, so there is no need to restart the Gateway in order to add or remove sites.

Optionally administrator can enable a possibility for dynamic site registration at runtime, see Section 4.4.2 for details. Then this file should contain only the static entries (or none if all sites register dynamically).

Further options for back-end sites configuration are presented in Section 6.

4.3 Main server settings: gateway.properties

Use the `gateway.hostname` property to configure the network interface and port the Gateway will listen on. You can also select between `https` and `http` protocol, though in almost all cases `https` will be used.

Example:

```
gateway.hostname = https://192.168.100.123:8080
```

Note

If you set the host to `0.0.0.0`, the Gateway will listen on all network interfaces of the host machine, else it will listen only on the specified one.

If the scheme of the hostname URL is set to `https`, the Gateway uses the configuration data from `security.properties` to configure the HTTPS settings.

4.4 Certificate-less end-user access

With UNICORE 7, it is possible that end-users do not have client certificates. To enable them to connect, the Gateway needs to accept TLS connections without a client certificate. To configure this, set the following in `gateway.properties`

```
gateway.httpServer.requireClientAuthn=false
```

4.4.1 Scalability settings

To fine-tune the operational parameters of the embedded Jetty server, you can set advanced HTTP server parameters. See [?informaltable] for details. Among others you can use the non-blocking IO connector offered by Jetty, which will scale up to higher numbers of concurrent connections than the default connector.

The Gateway acts as a `https` client for the VSites behind it. The number of concurrent calls is limited, and controlled by two parameters:

```
# maximum total number of concurrent calls to Vsites
gateway.client.maxTotal=100
# total number of concurrent calls per site
gateway.client.maxPerService=20
```

You can also control the limit on the maximum SOAP header size which is allowed by the Gateway. **Typically you don't have to touch this parameter.** However if your clients do produce very big SOAP headers and the Gateway blocks them, you can increase the limit. Note that such a giant SOAP header usually means that the client is not behaving in a sane way, e.g. is trying to perform a DoS attack.

```
# maximum size of an accepted SOAP header, in bytes
gateway.soapMaxHeader=102400
```

Note that Gateway may consume this amount of memory (plus some extra amount for other data) for each opened connection. Therefore, this value multiplied by the number of maximum allowed connections, should be significantly lower, then the total memory available for the Gateway.

4.4.2 Dynamic registration of Vsites

Dynamic registration is controlled by three properties in `CONF/gateway.properties` file:

```
gateway.registration.enable=true
gateway.registration.secret=<your key>
```

If set to true, the Gateway will accept dynamic registrations which are made by sending a HTTP POST request to the URL `/VSITE_REGISTRATION_REQUEST`. This request must contain a parameter "secret" which matches the value configured in the `gateway.properties` file.

Filters can be set to forbid access of certain hosts, or to require certain strings in the Vsite addresses. For example:

```
gateway.registration.deny=foo.org example.org
```

will deny registration if the remote hostname contains *foo.org* or *example.org*. Conversely,

```
gateway.registration.allow=mydomain.org
```

will only accept registrations if the remote address contains *mydomain.org*. These two (deny and allow) can be combined.

4.4.3 Web interface ("monkey page")

For testing and simple monitoring purposes, the Gateway displays a website showing detailed site information (the details view can be disabled). Once the Gateway is running, open up a browser and navigate to https://<gateway_host>:8080 (or whichever URL the gateway is running on). If the Gateway is configured to do SSL authentication, you will need to import a suitable client certificate into your web browser.

A HTML form for testing the dynamic registration is available as well, by clicking the link in the footer of the main Gateway page.

To disable the Vsite details page, set

```
gateway.disableWebpage=true
```

4.4.4 Main options reference

Property name	Type	Default value / mandatory	Description
gateway.hostname	string	<i>mandatory to be set</i>	external gateway bind address
gateway.registrationAllowed	string	-	Space separated list of allowed hosts for dynamic registration.
gateway.registrationDeny	string	-	Space separated list of denied hosts for dynamic registration.
gateway.registrationEnabled	{true, false}	false	Whether dynamic registration of sites is enabled.
gateway.registrationSecret	string	-	Required secret for dynamic registration.
<i>--- Passing Consignor info ---</i>			
gateway.consignorTolerance	Integer >= 0	30	The validity time of the authenticated client information passed to backend sites will start that many seconds before the real authentication. It is used to mask time synchronization problems between machines.
gateway.consignorValidity	Integer >= 0	60	What is the validity time of the authenticated client information passed to backend sites. Increase it if there machines clocks are not synhronized.
gateway.signConsignor	{true, false}	false	Controls whether information about the authenticated client (the consignor) passed to backend sites should be signed, or not. Signing is slower, but is required when sites may be reached directly, bypassing the Gateway.
<i>--- Gateway → Site client ---</i>			

Property name	Type	Default value / mandatory	Description
gateway.client.chunked	[true, false]	true	Controls whether chunked passing of HTTP requests to backend sites is supported.
gateway.client.connectionTimeout	integer number	30000	Connection timeout, used when connecting to backend sites.
gateway.client.expectContinue	[true, false]	true	Controls whether the HTTP expect-continue mechanism is enaled on connections to backend sites.
gateway.client.gzip	[true, false]	true	Controls whether support for compression is announced to backend sites.
gateway.client.keepAlive	[true, false]	true	Whether to keep alive the connections to backend sites.
gateway.client.maxConnectionsPerBackendSite	integer number	20	Maximum allowed number of connections per backend site.
gateway.client.maxConnectionsTotal	integer number	100	Maximum total number of connections to backend sites allowed.
gateway.client.socketTimeout	integer number	30000	Connection timeout, used when connecting to backend sites.
--- Advanced ---			
gateway.disableWebStatusPage	[true, false]	false	Whether the (so called monkey) status web page should be disabled.
gateway.externalHost	string	<i>not set</i>	External address of the gateway, when it is accessible through a frontend server as Apache HTTP.
gateway.soapMaxHeaderSize	integer [1024 — 1024000000]	102400	Size in bytes of the accepted SOAP header. In the most cases you don't need to change it.

Property name	Type	Default value / mandatory	Description
gateway.httpServer	string	S_allowedHeaders	CORS: comma separated list of allowed HTTP headers (default: any)
gateway.httpServer	string	S_allowedMethods	CORS: comma separated list of allowed HTTP verbs.
gateway.httpServer	string	S_allowedOrigins	CORS: allowed script origins.
gateway.httpServer	[true, false]	chainPreflight	CORS: whether preflight OPTION requests are chained (passed on) to the resource or handled via the CORS filter.
gateway.httpServer	string	S_exposedHeaders	CORS: comma separated list of HTTP headers that are allowed to be exposed to the client.
gateway.httpServer	string	disabledCipherSuites	Space separated list of SSL cipher suites to be disabled. Names of the ciphers must adhere to the standard Java cipher names, available here: http://docs.oracle.com/javase/8/docs/technotes-guides/security-SunProviders.html#SupportedCipherSuites
gateway.httpServer	[true, false]	CORS	false Control whether Cross-Origin Resource Sharing is enabled. Enable to allow e.g. accessing REST services from client-side JavaScript.

Property name	Type	Default value / mandatory	Description
gateway.httpServer[true, false] stricts	[true, false]	false	Control whether HTTP strict transport security is enabled. It is a good and strongly suggested security mechanism for all production sites. At the same time it can not be used with self-signed or not issued by a generally trusted CA server certificates, as with HSTS a user can't opt in to enter such site.
gateway.httpServer[true, false] random	[true, false]	false	Use insecure, but fast pseudo random generator to generate session ids instead of secure generator for SSL sockets.
gateway.httpServer[true, false] enable	[true, false]	false	Controls whether to enable compression of HTTP responses.
gateway.httpServer integer number gzipSize	integer number	51200	Specifies the minimal size of message that should be compressed.
gateway.httpServer integer number maxConnections	integer number	200	If the number of connections exceeds this amount, then the connector is put into a special <i>low on resources</i> state. Existing connections will be closed faster. Note that the server will also go to the low on resources mode if there are no available threads in the pool. You can set this to 0 to disable the connections limit (and have only thread pool size governed limit). If set to a negative number then the <i>low on resources</i> mode won't be used at all.

Property name	Type	Default value / mandatory	Description	
gateway.httpServer	integer	ResourceMaxIdleTime	In low resource conditions, time (in ms.) before an idle connection will time out.	
gateway.httpServer	integer	IdleTime	200000	Time (in ms.) before an idle connection will time out. It should be large enough not to expire connections with slow clients, values below 30s are getting quite risky.
gateway.httpServer	integer	threads	255	Maximum number of threads to have in the thread pool for processing HTTP connections. Note that this number will be increased with few additional threads to handle connectors.
gateway.httpServer	integer	threads	1	Minimum number of threads to have in the thread pool for processing HTTP connections. Note that this number will be increased with few additional threads to handle connectors.
gateway.httpServer	[true, false]	ClientAuthn		Controls whether the SSL socket requires client-side authentication.
gateway.httpServer	[true, false]		true	DEPRECATED, no effect
gateway.httpServer	[true, false]	ClientAuthn		Controls whether the SSL socket accepts (but does not require) client-side authentication.
gateway.httpServer	string	AllowedOrigins	http://localhost	URI origin that is allowed to embed web interface inside a (i)frame. Meaningful only if the xFrameOptions is set to <i>allowFrom</i> . The value should be in the form: <i>http[s]://host[:port]</i>

Property name	Type	Default value / mandatory	Description
gateway.httpServer	[deny, sameOrigin, allowFrom, allow]	deny	Defines whether a clickjacking prevention should be turned on, by insertion of the X-Frame-Options HTTP header. The <i>allow</i> value disables the feature. See the RFC 7034 for details. Note that for the <i>allowFrom</i> you should define also the <i>xFrameAllowed</i> option and it is not fully supported by all the browsers.

4.5 security.properties

In the `security.properties` file, the trust store and Gateway credential is configured.

Property name	Type	Default value / mandatory	Description
gateway.credential	file system 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.
gateway.credential	[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).
gateway.credential	string	-	Password required to load the credential.

Property name	Type	Default value / mandatory	Description
gateway.credentialPath	string	-	Location of the private key if stored separately from the main credential (applicable for <i>pem</i> and <i>der</i> types only),
gateway.credentialPassword	string	-	Private key password, which might be needed only for <i>jks</i> or <i>pkcs12</i> , if key is encrypted with different password than the main credential password.
gateway.credentialAlias	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> .

Property name	Type	Default value / mandatory	Description
gateway.truststore	[ALLOW, DENY]	ALLOW	Controls whether proxy certificates are supported.
gateway.truststore	[keystore, openssl, directory]	<i>mandatory to be set</i>	The truststore type.
gateway.truststoreInterval	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 ---			
gateway.truststoreConnectionTimeout	integer number	5	Connection timeout for fetching the remote CA certificates in seconds.

Property name	Type	Default value / mandatory	Description
gateway.truststore.filesystem.path	fileSystemPath	DiskCachePath	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.
gateway.truststore.[PEM,DER]	PropertyEncoding	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.
gateway.truststore.listofpropertieswithacommonprefix	listofPropertyLocations.*		List of CA certificates locations. Can contain URLs, local files and wildcard expressions. <i>(runtime updateable)</i>
<i>--- Keystore type settings ---</i>			
gateway.truststore.string	storeFormat		The keystore type (jks, pkcs12) in case of truststore of keystore type.
gateway.truststore.string	storePassword		The password of the keystore type truststore.
gateway.truststore.string	storePath		The keystore path in case of truststore of keystore type.
<i>--- Openssl type settings ---</i>			
gateway.truststore.[true,false]	NewStoreFormat	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)

Property name	Type	Default value / mandatory	Description
gateway.truststore	[GLOBUS_EU-GRIDPMA, GLOBUS_EU-GRIDPMA_GLOBUS, GLOBUS_EU-GRIDPMA_GLOBUS_REQUIRE, GLOBUS_EU-GRIDPMA_GLOBUS_REQUIRE_IGNORE]	GRIDPMA	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, found is checked (in the order defined by the property).
gateway.truststore	filesystem path	/etc/grid	Directory to be used for openssl truststore.
--- Revocation settings ---			
gateway.truststore	integer number	5	Connection timeout for fetching the remote CRLs in seconds (not used for Openssl truststores).
gateway.truststore	filesystem path	CachePath	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.
gateway.truststore	list of properties with a common prefix	Locations.*	List of CRLs locations. Can contain URLs, local files and wildcard expressions. Not used for Openssl truststores. (<i>runtime updateable</i>)

Property name	Type	Default value / mandatory	Description
gateway.truststore	[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.
gateway.truststore	integer number	Interval	How often CRLs should be updated, in seconds. Set to negative value to disable refreshing at runtime. (<i>runtime updateable</i>)
gateway.truststore	integer number	13600	For how long the OCSP responses should be locally cached in seconds (this is a maximum value, responses won't be cached after expiration)
gateway.truststore	filesystem path	cache	If this property is defined then OCSP responses will be cached on disk in the defined folder.
gateway.truststore	list of properties with a common prefix	LocalResponders.<	Optional list of local OCSP responders
gateway.truststore	[REQUIRE, IF_AVAILABLE, IGNORE]	IF_AVAILABLE	General OCSP ckecking mode. REQUIRE should not be used unless it is guaranteed that for all certificates an OCSP responder is defined.
gateway.truststore	integer number	10000	Timeout for OCSP connections in miliseconds.
gateway.truststore	[CRL, OCSP, OCSP_CRL]	OCSP_CRL	Controls overall revocation sources order
gateway.truststore	[true, false]	UseAll	Controls whether all defined revocation sources should be always checked, even if the first one already confirmed that a checked certificate is not revoked.

4.5.1 Proxy certificate support

The UNICORE Gateway optionally accepts proxy certificates as used by other Grid middleware systems. In general, we think proxies are a bad idea, but for interoperability purposes, proxies support can be enabled. If enabled, the clients using proxies are authenticated as the initial issuer of the presented proxy certificates chain. See the above reference properties table for the actual setting.

4.6 Logging

UNICORE uses the Log4j logging framework. It is configured using a config file. By default, this file is found in components configuration directory and is named `logging.properties`. The config file is specified with a Java property `log4j.configuration` (which is set in startup script).

Several libraries used by UNICORE also use the Java utils logging facility (the output is two-lines per log entry). For convenience its configuration is also controlled in the same `logging.properties` file and is directed to the same destination as the main Log4j output.

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 LOGS directory.

The following example config file configures logging so that log files are rotated daily.

```
# Set root logger level to INFO and its only appender to A1.
log4j.rootLogger=INFO, A1

# A1 is set to be a rolling file appender with default params
log4j.appender.A1=org.apache.log4j.DailyRollingFileAppender
log4j.appender.A1.File=logs/uas.log

#configure daily rollover: once per day the uas.log will be copied
#to a file named e.g. uas.log.2008-12-24
log4j.appender.A1.DatePattern='.'yyyy-MM-dd

# A1 uses the PatternLayout
log4j.appender.A1.layout=org.apache.log4j.PatternLayout
log4j.appender.A1.layout.ConversionPattern=%d [%t] %-5p %c{1} %x - ←
    %m%n
```

Note

In Log4j, the log rotation frequency is controlled by the `DatePattern`. Check <http://logging.apache.org/log4j/1.2/apidocs/org/apache/log4j/DailyRollingFileAppender.html> for the details.

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](#).

4.6.1 Logger categories, names and levels

Logger names are hierarchical. In UNICORE, prefixes are used (e.g. "unicore.security") to which the Java class name is appended. For example, the XUADB connector in UNICORE/X logs to the "unicore.security.XUADBAuthoriser" logger.

Therefore the logging output produced can be controlled in a fine-grained manner. Log levels in Log4j are (in increasing level of severity):

TRACE on this level *huge* pieces of unprocessed information are dumped, # DEBUG on this level UNICORE logs (hopefully) admin-friendly, verbose information, useful for hunting problems, # INFO standard information, not much output, # WARN warnings are logged when something went wrong (so it should be investigated), but recovery was possible, # ERROR something went wrong and operation probably failed, # FATAL something went really wrong - this is used very rarely for critical situations like server failure.

For example, to debug a security problem in the UNICORE security layer, you can set:

```
log4j.logger.unicore.security=DEBUG
```

If you are just interested in details of credentials handling, but not everything related to security you can use the following:

```
log4j.logger.unicore.security=INFO
log4j.logger.unicore.security.CredentialProperties=DEBUG
```

so the XUADBAuthoriser will log on DEBUG level, while the other security components log on INFO level.

Note

(so the full category is printed) and turn on the general DEBUG logging for a while (on uniconore). Then interesting events can be seen and subsequently the logging configuration can be fine tuned to only show them.

4.6.2 Logging

The most important, root log categories used by the Gateway's logging are:

unicore.gateway	General Gateway logging
unicore.connections	Log IPs of clients, and the DN after the SSL handshake
unicore.httpserver	HTTP processing, Jetty server
unicore.security	Certificate details and other security

The Gateway uses the so called MDC (Mapped Diagnostic Context) to provide additional information on the client which is served. In the MDC the client's IP address and client's Distinguished Name is stored. You can control whether to attach MDC to each line by the `%X{entry} log4j` pattern entry. As the entry you can use: `clientIP` or `clientName`. For example:

```
log4j.appender.A1.layout.ConversionPattern=%d [%t] [%X{clientIP} %X ←
  {clientName}] %-5p %c{1} - %m%n
```

5 Using Apache httpd as a frontend

You may wish to use the Apache webserver (httpd) as a frontend for the Gateway (e.g. for security or fault-tolerance reasons).

Requirements

- Apache httpd
- mod_proxy for Apache httpd

External references

- https://wiki.eclipse.org/Jetty/Howto/Configure_mod_proxy

6 Using the Gateway for failover and/or loadbalancing of UNICORE sites

The Gateway can be used as a simple failover solution and/or loadbalancer to achieve high availability and/or higher scalability of UNICORE/X sites without additional tools.

A site definition (in `CONF/connections.properties`) can be extended, so that multiple physical servers are used for a single virtual site.

An example for such a so-called multi-site declaration in the `connections.properties` file looks as follows:

```
#declare a multisite with two physical servers

MYSITE=multisite:vsites=https://localhost:7788 https://localhost ↵
:7789
```

This will tell the Gateway that the virtual site "MYSITE" is indeed a multi-site with the two given physical sites.

6.1 Configuration

Configuration options for the multi-site can be passed in two ways. On the one hand they can go into the `connections.properties` file, by putting them in the multi-site definition, separated by ";" characters:

```
#declare a multisite with parameters

MYSITE=multisite:param1=value1;param2=value2;param3=value3;...
```

The following general parameters exist

vsites	List of physical sites
strategy	Class name of the site selection strategy to use (see below)
config	Name of a file containing additional parameters

Using the "config" option, all the parameters can be placed in a separate file for enhanced readability. For example you could define in `connections.properties`:

```
#declare a multisite with parameters read from a separate file

MYSITE=multisite:config=conf/mysite-cluster.properties
```

and give the details in the file "conf/mysite-cluster.properties":

```
#example multisite configuration
vsites=https://localhost:7788 https://localhost:7789

#check site health at most every 5 seconds
strategy.healthcheck.interval=5000
```

6.2 Available Strategies

A selection strategy is used to decide where a client request will be routed. By default, the strategy is "Primary with fallback", i.e. the request will go to the first site if it is available, otherwise it will go to the second site.

Primary with fallback

This strategy is suitable for a high-availability scenario, where a secondary site takes over the work in case the primary one goes down for maintenance or due to a problem. This is the default strategy, so nothing needs to be configured to enable it. If you want to explicitly enable it anyway, set

```
strategy=primaryWithFallback
```

The strategy will select from the first two defined physical sites. The first, primary one will be used if it is available, else the second one. Health check is done on each request, but not more frequently as specified by the "strategy.healthcheck.interval" parameter. By default, this parameter is set to 5000 milliseconds.

Changes to the site health will be logged at "INFO" level, so you can see when the sites go up or down.

Round robin

This strategy is suitable for a load-balancing scenario, where a random site will be chosen from the available ones. To enable it, set

```
strategy=roundRobin
```

Changes to the site health will be logged at "INFO" level, so you can see when the sites go up or down.

It is very important to be aware that this strategy requires that all backend sites used in the pool, share a common persistence. It is because Gateway does not track clients, so particular client requests may land at different sites. This is typically solved by using a non-default, shared database for sites, such as MySQL.

Note

Currently loadbalancing of target sites is an experimental feature and is not yet fully functional. It will be improved in future UNICORE versions.

Custom strategy

You can implement and use your own failover strategy, in this case, use the name of the Java class as strategy name:

```
strategy=your_class_name
```

7 Gateway failover and migration

The Section 6 covered usage of the Gateway to provide failover of backend services. However it may be needed to guarantee high-availability for the Gateway itself or to move it to other machine in case of the original one's failure.

7.1 Gateway's migration

The Gateway does not store any state information, therefore its migration is easy. It is enough to install the Gateway at the target machine (or even to simply copy it in the case of installation from the core server bundle) and to make sure that the original Gateway's configuration is preserved.

If the new machine uses a different address, it needs to be reflected in the server's configuration file (the listen address). Also, the configuration of sites behind the Gateway must be updated accordingly.

7.2 Failover and loadbalancing of the Gateway

Gateway itself doesn't provide any features related to its own redundancy. However as it is stateless, the standard redundancy solutions can be used.

The simplest solution is to use Round Robin DNS, where DNS server routes the Gateway's DNS address to a pool of real IP addresses. While easy to set up this solution has a significant drawback: DNS server doesn't care about machines being down.

The much better choice is to use the Linux-HA software suite, often known under the name of its principal component, the *heartbeat*. For details see <http://www.linux-ha.org>

Additionally a more advanced HTTP-aware software can be used, such as HA-Proxy (<http://haproxy.1wt.eu>). Currently Gateway and UNICORE don't maintain HTTP sessions so usage of the HTTP-aware load-balancer is not strictly required, but such solutions generally provide more general purpose features.

8 Building the Gateway from source

To checkout the latest version of the Gateway source code, do

```
svn co http://svn.code.sf.net/p/unicore/svn/gateway/trunk gateway
```

You will need to install Maven from <http://maven.apache.org> Compile using

```
mvn install
```

which compiles the code and runs the tests.

The file "README-building.txt" contains instructions for building distributable packages.