



User Manual

v3.2.4

1 Introduction

LogFaces is a standalone solution for collecting, storing, dispatching and filtering log data. LogFaces is a "box" which resides next to your system and lifts the logging burden off you. Built one level above the conventional logging API's such as [Apache Logging Services](#) it easily integrates with existing systems, including those built on J2EE, native C++, or Microsoft© .NET Framework. If your system is already using those API's, you can use logFaces for storing log, real-time log monitoring, evaluating operational discrepancies, receiving scheduled email reports with log files and more...

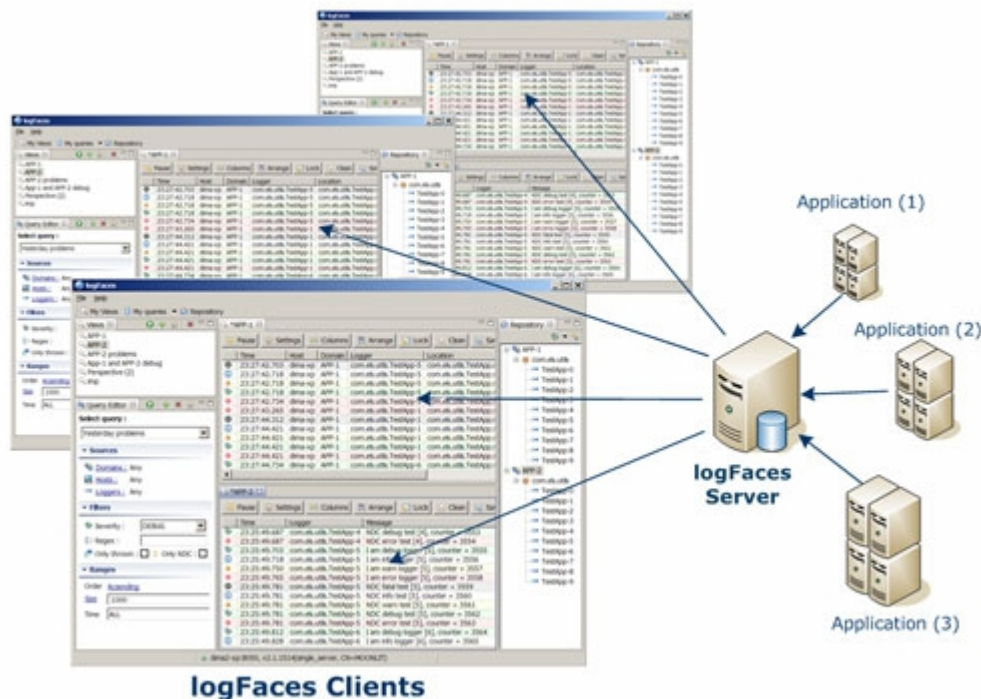


Figure 1-1.1 LogFaces Architecture

There are three players - your system, logFaces Server and logFaces Client. It works like this: your system/s send log data to the logFaces Server which routes this data to interested clients or **does something with this data on its way**. For example; saving some of it to database or discarding it, splitting and filtering it for various client listeners. The only requirement for any system to be operable with logFaces is to produce log in a format understood by the logFaces Server which collects this log. The current version of logFaces supports [Apache log4xxx API's](#) as schema of log statements. But technically logFaces will not be limited only to those formats. Today log4xxx API's are the de-facto standards in software development and most systems already support them. In such cases the integration with logFaces is straightforward and is only a matter of simple configuration.

LogFaces can be used in a compact mode (or **Server Mode** as we call it) where applications send log data directly to the Client. This mode is mostly dedicated for those who want to use logFaces as a real time viewer for single user. It fits better smaller projects due to its simplicity, however it doesn't have persistence capabilities, reporting and data mining features – it's a plain real time viewer of log data.

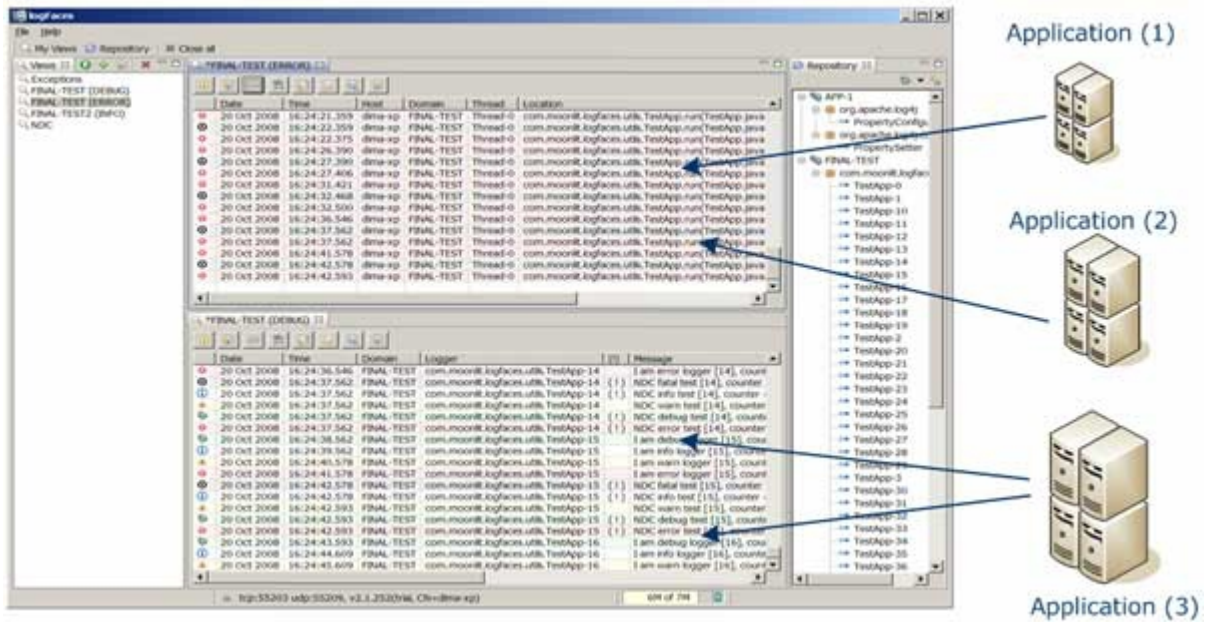


Figure 1-1.2 logFaces Client in server mode

2 Getting started with logFaces Server

This section will guide you through the quick process of installation and configuration. Java™ Runtime Environment (JRE) 1.6 or later must be installed before proceeding with installation on any operating system. LogFaces Server will start after the installation with default settings and trial license for **10 days evaluation**.

2.1 Installing logFaces Server

On Windows, download and run the installer which will walk you through the process. During installation you will be asked to register and/or run logFaces as Windows service. Linux and Solaris distributions come as **tar.gz** archives, just unzip the archive and you're ready to go.

2.2 Using silent installer options

Running installer exe with -h flag will display list of options you can use. Amongst them is -q option for running in unattended mode. It is also possible to use a response file from previous installation and re-use it, response files located in .install4j directory and named 'response.varfile'. For example, to run installation silently you can do "**ifs.windows.win32.x86.exe -q -varfile response.varfile**". This will silently install the server using all options specified in response file.

2.3 Running logFaces Server on Windows

If you selected to run logFaces as windows service, the service named LFS will be registered on your computer automatically and will run every time you start your computer. Use conventional service commands to start/stop the service by doing so in command prompt - "**net start ifs**" or "**net stop ifs**"

You can also run the server as a console application by using **/bin/run.bat**.

2.4 Running logFaces Server on Linux/Solaris

- In order to start the server in the terminal execute command **./bin/ifs console**
- In order to start the server as daemon execute command **./bin/ifs start**
- In order to stop the server process execute command **./bin/ifs stop**
- In order to check the server process status execute command **./bin/ifs status**

2.5 Server directory tree

Once installed and ran for the first time, the server will explode its default configuration and you will find a bunch of new folders under installation directory. The table below is a brief summary of folders and their purpose. Normally you won't need all the technical details, but it's a good idea to familiarize yourself with some internals, some of the files are referred to throughout this manual.

Path	Updated with version increments	Description
/.install4j	yes	Visible only on windows, created by installers.
/bin	yes	Binaries for bootstrapping the server, depending on the platform the content will be different.
/admin	yes	Admin web application and resources
/conf	no	Server configuration files
/db	no	Relevant only when embedded database is used
/doc	yes	Holds release notes and other documents
/legal	no	Holds end user license agreements
/lib	yes	Binary distribution libraries and dependencies
/log	no	Holds server internal logs
/overflow	no	Contain overflow files when server overloads
/temp	no	Temporally files

Note that folders which are created and updated by the server (those which are not part of version updates), can be located elsewhere and not necessarily under installation directory. We call these folders **artifacts** and sometimes it's a good idea to keep them separately, for example for the backup purposes or switching from one setup to another. By default all artifacts reside under installation directory, this is the argument for running the server, and you can change it in `/bin/lfs.conf` like this:

wrapper.app.parameter.1=MY_DIR

MY_DIR can be an empty folder; when you run the server for the first time, it will create all default artifacts automatically. Alternatively, you can move other artifacts into MY_DIR or point the above parameter to another location with different artifacts in there - this is what we often do during tests. Once again, all this is not required, in most cases the default directory structure is good for nearly all circumstances.

2.6 Integrating with applications

To work with logFaces, your application needs to be configured by adding several elements to its logging configuration file. Provided that your system is based on log4xxx API's, you should be having log4xxx configuration file/s, which usually come as property files or XML configuration files. We add an appender to your configuration file which knows how to communicate with logFaces Server. Sections below will show how to setup such appenders in various situations. As of this writing we provide our own appenders for log4j and logback frameworks while .Net, PHP and C++ can be freely obtained from Apache downloads and don't require any change to work with logFaces.

2.6.1 log4j TCP configuration

We offer a fairly standard yet sophisticated log4j appender with asynchronous TCP socket connection and fail over mechanism. When your application emits log statements, they will not be sent to server at the expense of the calling thread. They will be queued and sent to logFaces server by the background thread. In addition to queuing, the appender also knows how to fail over to another host or save logs to a local file. You can specify how often to retry the connection, how many times to retry and to what host to switch to when all retries are exhausted. This is an example of log4j XML configuration:

```
<appender name="LFS" class="com.moonlit.logfaces.appenders.AsyncSocketAppender">
  <param name="remoteHost" value="host1,host2,host3" />
  <param name="port" value="55200" />
  <param name="locationInfo" value="true" />
  <param name="threshold" value="ALL" />
  <param name="application" value="My application" />
  <param name="reconnectionDelay" value="5000" />
  <param name="offerTimeout" value="0" />
  <param name="queueSize" value="100" />
  <appender-ref ref="CONSOLE" />
</appender>
```

This is an example of log4j property file configuration :

```
log4j.appender.LFS = com.moonlit.logfaces.appenders.AsyncSocketAppender
log4j.appender.LFS.application = APP-1
log4j.appender.LFS.remoteHost = host1,host2,host3
log4j.appender.LFS.port = 55200
log4j.appender.LFS.locationInfo = true
log4j.appender.LFS.threshold = ALL
log4j.appender.LFS.reconnectionDelay = 5000
log4j.appender.LFS.offerTimeout = 0
log4j.appender.LFS.queueSize = 100
log4j.appender.LFS.backupFile = "c:/lfs-backup.log"
```

(note what log4j property configurations don't support appender references, so for the backup we are using built-in RollingFileAppender and direct it to 'backupFile' attribute)

Property	Description	Default	Mandatory
application	Identifies application under this name. All logs coming through this appender will be stamped with this name, which can later be used on client.	-	no
remoteHost	Comma separated list of logFaces servers. If more than one host specified, the appender will automatically fail over to the next host when current host becomes unavailable. Switching hosts is done in the loop. If only one host specified, the retries will be done indefinitely with this host.	-	yes
port	Port where logFaces server will accept the connection from this appender.	55200	no
locationInfo	Specifies whether to include location data, such as class name, method name and line numbers.	FALSE	no
reconnectionDelay	Rate of reconnection retries in milliseconds.	5000	no
nofRetries	How many times to retry before dropping current host and switching to the next one. If only one host specified in remoteHost attribute, the retries will go indefinitely to the same host.	3	no
queueSize	Size of the event queue. The larger the size, the less likely the data will get lost when connection is lost, because events will be re-transmitted to the server when connection recovers. However, queue size affects JVM heap memory, so be considerate.	500	no
offerTimeout	How long to wait (in ms) while offering event to the appender queue. When server is slower than application and queue gets full, the caller has an option to wait before giving up. Queue can typically get full when server is down or when server can't consume log data in the rate of this appender. WARNING: Use with care as it will slow down the calling thread when queue fills up.	0	no
appender-ref	Reference to an appender to use when logFaces server is not reachable on any host. This is a backup delegate appender. When specified and server is unreachable with full queue, every log event will be delegated to the referenced appender. When not specified, the events will be discarded. Used only with XML based configuration.	-	no
backupFile	Used with log4j properties based configuration for setting up a backup file. When server becomes unreachable, the logs will be sent to this file. Its XML layout will let you easily import the data in server later.		no

2.6.2 log4j UDP appenders

Should your choice of transport be a UDP, there is a fairly simple appender you can use. Keep in mind that UDP is unreliable protocol and it's not recommended to be used with unusually heavy loaded producers of logs. Having mentioned that, there is still one benefit for using UDP – it's extremely cheap on resources. Below are examples of configurations:

XML style:

```
<appender name="LFSU" class="com.moonlit.logfaces.appenders.LFUDPAppender">
  <param name="application" value="MyApp" />
  <param name="remoteHost" value="localhost" />
  <param name="port" value="55201" />
  <param name="locationInfo" value="true" />
</appender>
```

Or properties style:

```
log4j.appender.LFSU = com.moonlit.logfaces.appenders.LFUDPAppender
log4j.appender.LFSU.application = MyApp
log4j.appender.LFSU.remoteHost = localhost
log4j.appender.LFSU.port = 55201
log4j.appender.LFSU.locationInfo = true
```

2.6.3 logback appenders

If you are using [logback](#) framework in your applications, you can easily integrate with logFaces.

Below is an example logback configuration :

```
<configuration>
  <contextName>MYAPPLICATION</contextName>
  <appender name="CONSOLE" class="ch.qos.logback.core.ConsoleAppender">
    <layout class="ch.qos.logback.classic.PatternLayout">
      <Pattern>%d{HH:mm:ss.SSS} [%-5level] %logger{36} | %msg%n</Pattern>
    </layout>
  </appender>
  <appender name="LFS" class="com.moonlit.logfaces.appenders.logback.LogfacesAppender">
    <remoteHost>host1,host2</remoteHost>
    <port>55200</port>
    <locationInfo>true</locationInfo>
    <application>${CONTEXT_NAME}</application>
    <reconnectionDelay>1000</reconnectionDelay>
    <offerTimeout>0</offerTimeout>
    <queueSize>200</queueSize>
    <appender-ref ref="CONSOLE" />
    <delegateMarker>true</delegateMarker>
  </appender>
  <root level="trace">
    <appender-ref ref="CONSOLE" />
    <appender-ref ref="LFS" />
  </root>
</configuration>
```


Note how **contextName** can be referenced to specify the application name. When LFS appender will unable to work with server, it will fall back to CONSOLE appender referenced by **appender-ref** – normally you would want to use some rolling file appender instead of just console.

Meaning of the attributes are identical to our log4j appender described above except "**delegateMarker**" option which is specific to logback. If set to true, the appender will automatically copy logback MARKER into event context which will then appear on server as special MDC property named "**marker**". This will allow you to filter and query logs by the markers set in your application. For example, if you do this in your code :

```
Marker ADMIN = MarkerFactory.getMarker("SYS-ADMIN");  
logger.trace(ADMIN, "this event is for the attention of sysadmin");
```

Then you will be able easily to fish our all events marked with 'SYS-ADMIN' token. Make sure to specify MDC mapping on server so that it contains "**marker**" property, see Context section under Administration [22] for more details.

Make sure to place **lfsappenders.jar** into the class path of your application, it can be found either in /lib directory of server installation or from our [download page](#). Logback dependency jars must be in the class path as well, make sure you grab them from the authors web site.

2.6.4 log4php appenders

[log4php](#) comes out of the box with its own [socket appender](#) and can be directed to logFaces server, using conventional configuration file. Below is an example of such configuration for both TCP and UDP transports. You will normally need to use one of them:

```
<configuration xmlns="http://logging.apache.org/log4php/">
  <appender name="lfs-tcp" class="LoggerAppenderSocket">
    <param name="remoteHost" value="localhost" />
    <param name="port" value="55200" />
    <param name="timeout" value="2" />
    <layout class="LoggerLayoutXml">
      <param name="locationInfo" value="true" />
      <param name="log4jNamespace" value="true" />
    </layout>
  </appender>
  <appender name="lfs-udp" class="LoggerAppenderSocket">
    <param name="remoteHost" value="udp://localhost" />
    <param name="port" value="55201" />
    <layout class="LoggerLayoutXml">
      <param name="locationInfo" value="true" />
      <param name="log4jNamespace" value="true" />
    </layout>
  </appender>
</root>
  <level value="TRACE" />
  <appender_ref ref="lfs-tcp" />
  <appender_ref ref="lfs-udp" />
</root>
</configuration>
```

Note the timeout parameter in lfs-tcp appender, if you choose to use TCP version beware that every log statement will cause new socket connection to open against logFaces server. Being very reliable, the TCP appender may cause delays when your server is down or its network is slow. Having very long timeouts will cause your page visitors to wait if log server is unavailable. UDP appender comes to the rescue, but UDP is inherently unreliable protocol so there is always a chance that some of the datagrams will get lost. The conclusion – choose wisely upon your needs and circumstances.

Below is an example of what normally happens in PHP code. Note how we use MDC '**application**' property to let logFaces server know which application the logs are coming from:

```
<?php
```

```
// 1. use the path where you unpacked log4php
include('/development/testphp/php/Logger.php');

// 2. point to configuration file which must include logFaces appenders
Logger::configure('/development/testphp/log4php.xml');

// 3. fetch a logger, any name is OK, best to use names which will
// be easy to use in logFaces hierarchy like this
// logFaces will split it into package-like notation for easier traceability
$log = Logger::getLogger('com.mycompany.myproject.mypage');

// 4. specify application (domain) name, this is optional.
// it will allow logFaces server to associate logs with 'application' token
// if not specified, the logFaces will use the host name of the originating logs
LoggerMDC::put("application", "my-product");

// 5. typical logging stuff...
$log->trace("this is a trace message");
$log->debug("this is an info message");
$log->info("this is an info message");
$log->warn("this is a warning message");
$log->error("this is an error message");
$log->fatal("this is a fatal message");
```

```
?>
```

2.6.5 log4net appenders

If your system is .Net based and using [Apache log4net API](#) for logging, you can use its out of the box [UdpAppender](#) :

```
<log4net>
  <appender name="logFaces" type="log4net.Appender.UdpAppender">
    <param name="RemoteAddress" value="10.0.0.110" />
    <param name="RemotePort" value="55201" />
    <param name="Encoding" value="UTF-8" />
    <layout type="log4net.Layout.XmlLayoutSchemaLog4j, log4net">
      <locationInfo value="true" />
    </layout>
  </appender>

  <root>
    <level value="ALL" />
    <appender-ref ref="logFaces" />
  </root>
</log4net>
```

As mentioned earlier, logFaces can listen for TCP and/or UDP. In this example, we use UDP appender - make sure that RemotePort attribute in this example corresponds to the one configured in logFaces.

2.6.6 NLog appenders

If your system is .Net based and using [NLog logging platform](#), you can use either TCP or UDP or even both of them depending on your needs. Here is the configuration sample :

```
<?xml version="1.0" encoding="utf-8" ?>
<nlog xmlns="http://www.nlog-project.org/schemas/NLog.xsd"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">

  <targets>
    <target name="console" xsi:type="Console" />
    <target name="logfile" xsi:type="File" fileName="file.txt" />
    <target name="lfs-tcp" xsi:type="Network" address="tcp://localhost:55200" />
      layout="${log4jxmlevent:includeMdc=true:appInfo=MYAPP}"
    <target name="lfs-udp" xsi:type="Network" address="udp://localhost:55201" />
      layout="${log4jxmlevent:includeMdc=true:appInfo=MYAPP}"
  </targets>

  <rules>
    <logger name="com.package1.*" minlevel="Error" writeTo="lfs-tcp" />
    <logger name="com.package2.*" minlevel="Trace" writeTo="lfs-udp" />
  </rules>
</nlog>
```

Note how we use '**appInfo**' attribute to identify your application name, and '**includeMDC**' to make sure mapped diagnostic context is transmitted to the server (optional).

2.6.7 log4cxx appenders

If your system is based on C++ and using [Apache Log4cxx API](#) for logging, you should configure it by adding XMLSocketAppender included in log4cxx API itself. Here is a snippet of configuration example:

```
log4j.appender.stdout=org.apache.log4j.ConsoleAppender
log4j.appender.stdout.layout=org.apache.log4j.PatternLayout
log4j.appender.stdout.layout.ConversionPattern= %-5p %d{HH:mm:ss} %-20c{1} %X{stam} | %m%n

log4j.appender.LFS=org.apache.log4j.net.XMLSocketAppender
log4j.appender.LFS.RemoteHost=10.200.1.110
log4j.appender.LFS.Port=55200

log4j.rootLogger=debug, stdout, LFS, FILE
```

The meaning of the attributes is identical to the previous Java^{*} example. However, note that XMLSocketAppender doesn't (yet?) provide "Application" and "LocationInfo" attributes. This is not a problem for logFaces – those loggers which don't correspond to any logging domain will be automatically grouped in logFaces under name "**Default Domain**". Unfortunately, until those attributes are supported by the underlying API's, we will have to add some code when initializing the logger in the application. The code snippet is shown below, what we do is simply getting into a root logger, digging out the LFS appender from there and manually set the missing attributes of the layout like this:

```
// this is a workaround for XMLSocketAppender to allow routing of properties
// over the network, we manually setup the layout
LoggerPtr root = Logger::getRootLogger();
AppenderPtr app = root->getAppender(LOG4CXX_STR("LFS"));
if(app != NULL){
    LayoutPtr layout = app->getLayout();
    if(layout != NULL){
        layout->setOption(LOG4CXX_STR("locationinfo"), LOG4CXX_STR("true"));
        layout->setOption(LOG4CXX_STR("properties"), LOG4CXX_STR("true"));
        MDC::put("application", "WSC");
    }
}
```

In any case, those missing attributes are not a show stoppers, your application can still work with logFaces out of the box with those limitations.

IMPORTANT:

The MDC (message diagnostic context) works only in the context of the current thread. In case you have several threads in your application you should add **MDC::put("application", "xxx")** call in the beginning of every thread. Otherwise, the log statements coming from those threads will be orphaned and server will automatically put them under "Default Domain" which might be a bit confusing. Future versions of logFaces will include proper appender to avoid those workarounds.

2.7 Administration

Most of logFaces Server configuration is done remotely using conventional browser. To access Administration Console, open your browser and navigate to this URL:

<http://your-logfaces-server-host:8050>

The shortcut is also created for you during the installation in Windows Start menu. Of course, you can access the console from any other host as long as your network configuration allows the access. Access to logFaces Administration Console is secured by user name and password, which are both defaulted to "admin" during installation and can later be changed. Administration Console is grouped into several tabs where you can easily modify needed functionality – connectivity, database, access control, context variables, e-mail, notifications, licensing and monitor status.

2.7.1 Connectivity

In client connectivity section you specify how server is accessed by rich clients and browsers. This includes binding address, port number, SSL and credentials for administrator to access the administration application.

Administration Console
v3.1.3 (build #2185)

Connectivity [Security](#) [Database](#) [Context](#) [Email](#) [Notifications](#) [License](#) [Status](#)

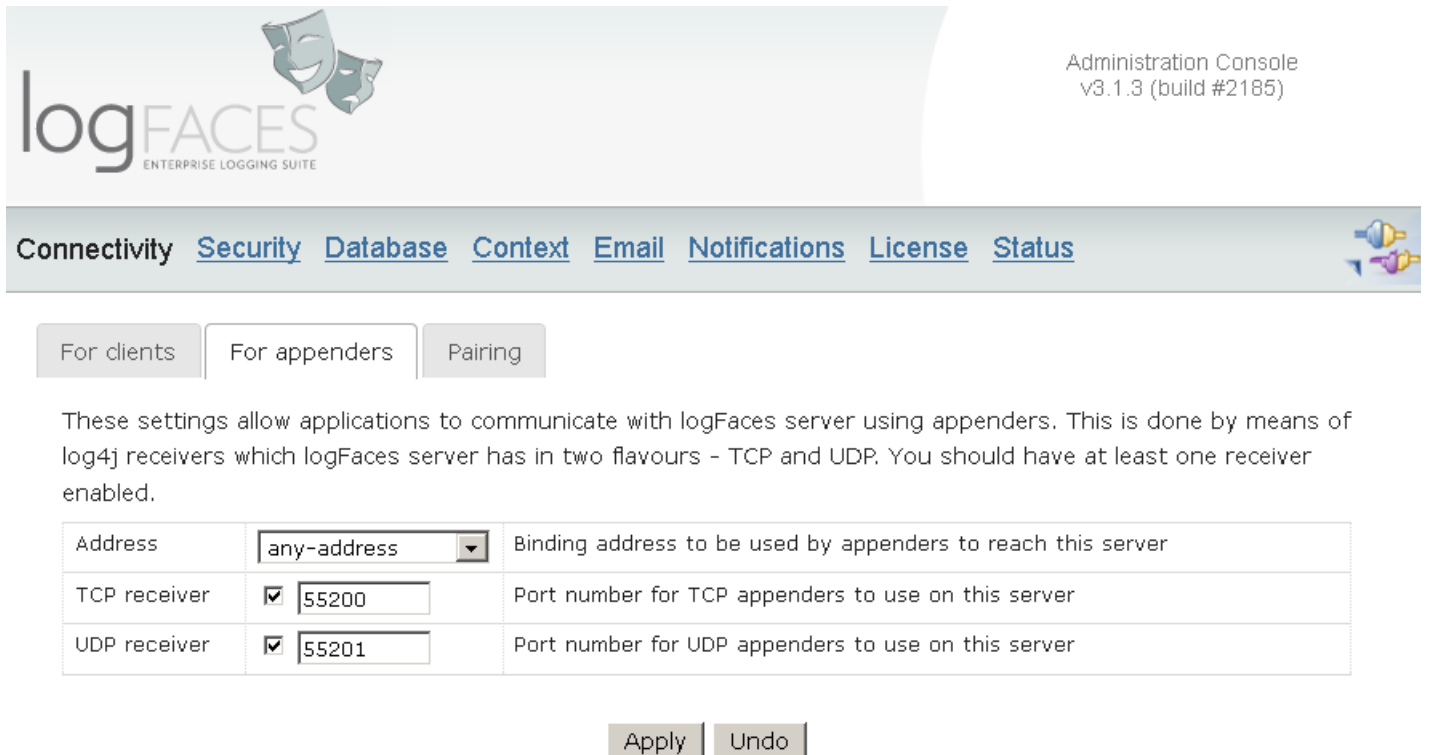
For clients For appenders Pairing

Here you specify how logFaces server is exposed to rich clients and administration. Note that changing any of these parameters will require server restart.

Address	<input type="text" value="any-address"/>	Binding address to be used by external clients to reach this server
Port	<input type="text" value="8050"/>	Listening port number
Secure	<input type="text" value="No"/>	Access to this server from outside will be strictly over SSL
Admin user	<input type="text" value="admin"/>	User name for administrator login
Admin password	<input type="password" value="•••••"/>	Password for administrator login

Figure 2.1: External connectivity

In appenders connectivity section you specify how server gets exposed to applications. There are two types of receivers available – TCP and UDP, you can enable both or either of them depending on what appenders you use in your apps.



Administration Console
v3.1.3 (build #2185)

Connectivity [Security](#) [Database](#) [Context](#) [Email](#) [Notifications](#) [License](#) [Status](#)

For clients For appenders Pairing

These settings allow applications to communicate with logFaces server using appenders. This is done by means of log4j receivers which logFaces server has in two flavours - TCP and UDP. You should have at least one receiver enabled.

Address	<input type="text" value="any-address"/>	Binding address to be used by appenders to reach this server
TCP receiver	<input checked="" type="checkbox"/> <input type="text" value="55200"/>	Port number for TCP appenders to use on this server
UDP receiver	<input checked="" type="checkbox"/> <input type="text" value="55201"/>	Port number for UDP appenders to use on this server

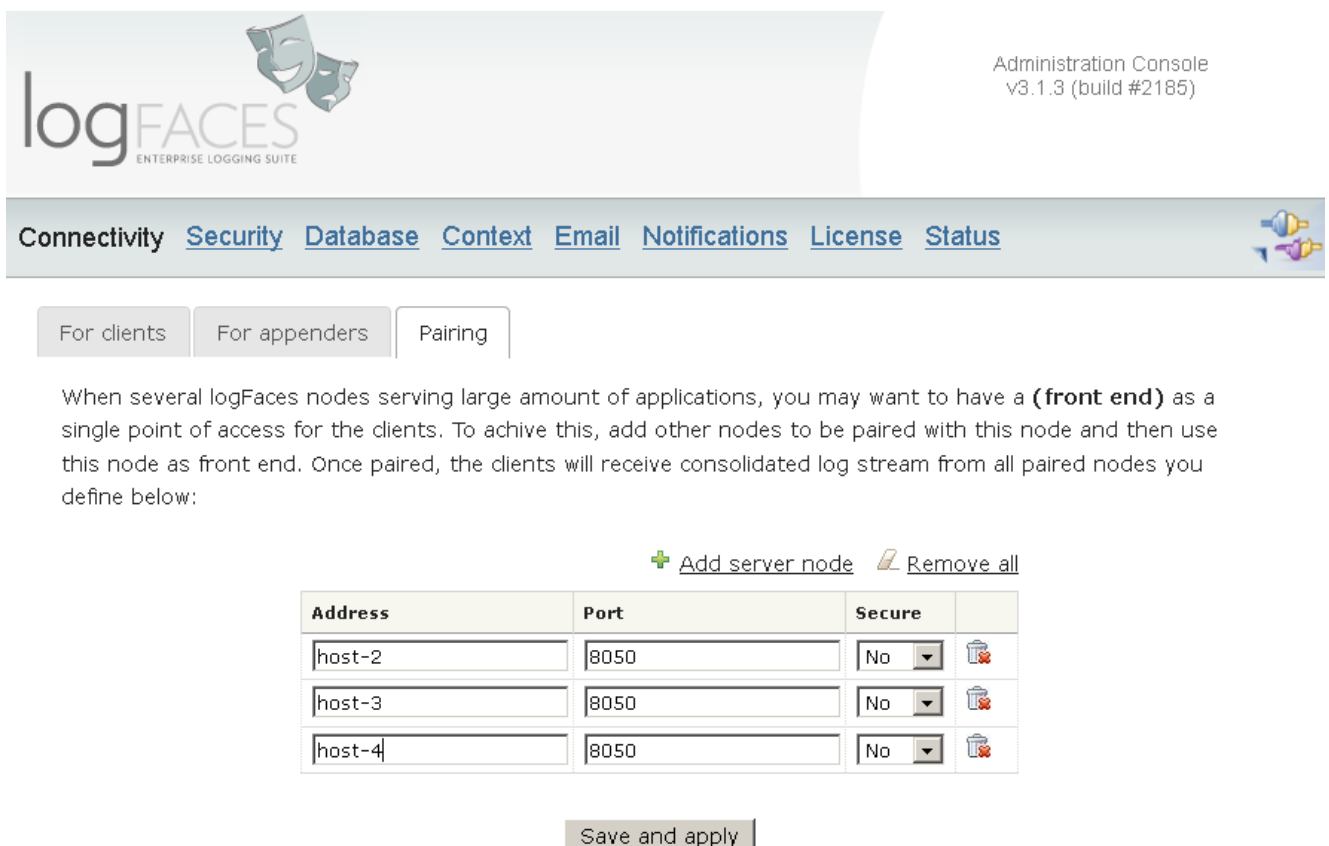
Figure 2.2: Internal connectivity

Note the binding address in both clients and appenders tabs. If server is installed on a computer with several network cards, you can bind server sockets to a particular address or a host name. This is a good idea if you need physical separation of the transport. When “any address” is specified, the server will pick the default binding, this will allow access to server by any address it supports. Otherwise, the access will be strictly by the address you select.

2.7.2 Pairing

Pairing is a mechanism for providing single access point to the data when several logFaces nodes are deployed as a single system. Multiple nodes are used for two purposes: a) for splitting the inflow when working with large amount of applications and b) for providing fail over. In either case, you deploy several nodes which do actual data processing against the apps (**back end**) and another node for user access only (**front end**). This way, the users don't have to know which node to connect to in order to receive the logs, they always connect to a single **front end** node which is paired with **back end** nodes by specifying their connection end points – host, port and SSL option. Note that this configuration assumes that all nodes, back and front, are sharing the same database.

The example below demonstrates how **this** node (host-1) being the front end is paired with 4 back end nodes at host-2, host-3 and host-4. The front end will be accessed by end users and whatever they request will be served by the relevant node from the back end or any combination of them. So, the end users never aware of the back end nodes existence, the front end delegates everything to its pairs.



Administration Console
v3.1.3 (build #2185)

Connectivity [Security](#) [Database](#) [Context](#) [Email](#) [Notifications](#) [License](#) [Status](#)

For clients For appenders **Pairing**

When several logFaces nodes serving large amount of applications, you may want to have a (**front end**) as a single point of access for the clients. To achieve this, add other nodes to be paired with this node and then use this node as front end. Once paired, the clients will receive consolidated log stream from all paired nodes you define below:

[+ Add server node](#) [Remove all](#)

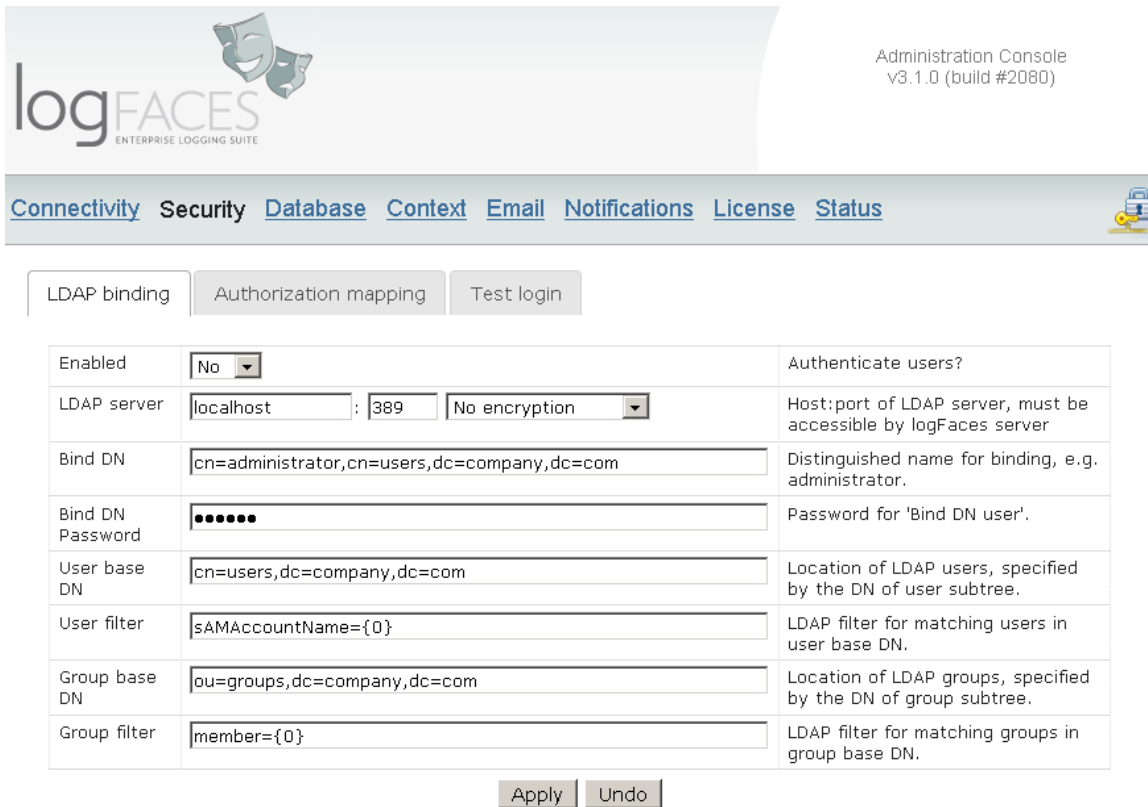
Address	Port	Secure	
host-2	8050	No	
host-3	8050	No	
host-4	8050	No	

Save and apply

Figure 2.3: Pairing example

2.7.3 Security

logFaces server can be integrated with external LDAP server for authentication and authorization of clients. Both features are optional. When enabled, any client trying to connect to logFaces server will be prompted to log-in. User credentials will then be delegated to the LDAP server which will do actual authentication. If authorization is enabled as well, the logged in client will only have access to those applications you specify.



logFaces ENTERPRISE LOGGING SUITE

Administration Console v3.1.0 (build #2080)

Connectivity Security Database Context Email Notifications License Status

LDAP binding Authorization mapping Test login

Enabled	<input type="checkbox"/> No	Authenticate users?
LDAP server	localhost : 389 No encryption	Host:port of LDAP server, must be accessible by logFaces server
Bind DN	cn=administrator,cn=users,dc=company,dc=com	Distinguished name for binding, e.g. administrator.
Bind DN Password	••••••	Password for 'Bind DN user'.
User base DN	cn=users,dc=company,dc=com	Location of LDAP users, specified by the DN of user subtree.
User filter	§AMAccountName={0}	LDAP filter for matching users in user base DN.
Group base DN	ou=groups,dc=company,dc=com	Location of LDAP groups, specified by the DN of group subtree.
Group filter	member={0}	LDAP filter for matching groups in group base DN.

Apply Undo

Figure 2.4 LDAP binding

LDAP server - host and port should be pointing to your LDAP server, make sure that it can be accessed from logFaces server host.

Encryption – use it if your LDAP server works over SSL. If your LDAP server is using well known root CA, the communication should work straight away. Otherwise you will want to introduce the trust store with that certificate into logFaces JVM so that it can trust your LDAP server. One of the ways of doing this is by setting the following JVM system properties in /bin/lfs.conf file :

```
wrapper.java.additional.x=-Djavax.net.ssl.trustStore="truststore_file"
wrapper.java.additional.x=-Djavax.net.ssl.trustStorePassword="truststore_password"
```

It will instruct logFaces JVM to use this particular store with that password for getting the certificate.

Bind DN - distinguished name for binding to the LDAP server, logFaces will use this DN in order to gain an access to user base. Usually those credentials are obtained from LDAP server administrator and must have permissions for walking user base tree.

Bind DN Password - corresponding password for the binding user.

User base DN - distinguished name corresponding to the location of users to be authenticated.

User filter - LDAP filter for matching users in user base DN. This parameter gives a very sophisticated way to match users in the user base. The default value **attr={0}** will match any user whose user ID is mapped to the attribute named '**attr**'. This attribute name varies in different LDAP implementations, for example in Apache DS this is normally '**uid**' while in MS Active Directory it show as **sAMAccountName**. Note the **{0}** parameter – it must be present all the time to match the actual user ID supplied by the user. When you want to do more complex matching of users, you can specify fairly complex LDAP filters in this field – please refer to LDAP documentation for the syntax details. Here is an example, the filter below will only match users from SALES organization unit

```
(&(ou=SALES)(uid={0}))
```

So, even when user is part of user base (**uid={0}**), it will only be attempted for authentication when she belongs to SALES unit. This way, having fairly large user base DN you filter out only relevant users for accessing logFaces.

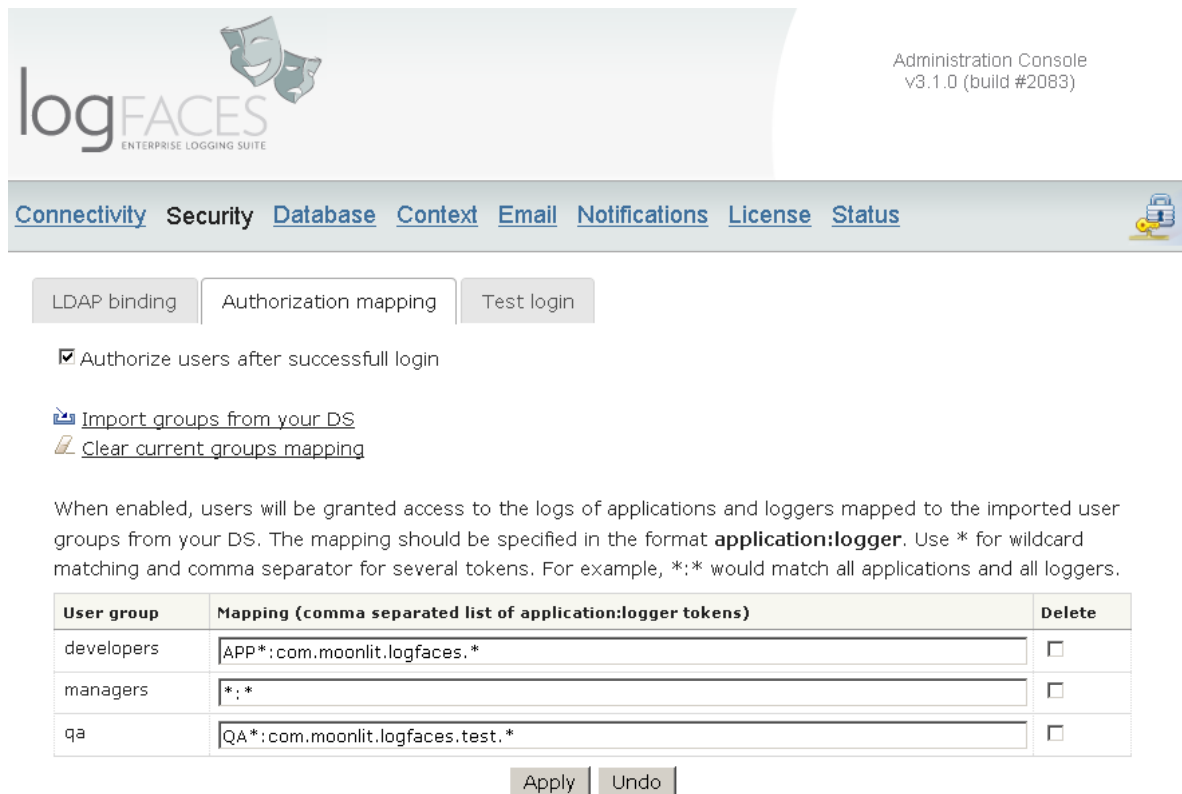
Group base DN - location of user groups sub-tree. Groups will be used for authorization, if you don't need authorization - leave this field with default value.

Group filter - LDAP filter for matching groups in group base DN. This parameter is very similar to the user filter except that it acts on user groups. The default value is again **attr{0}** which will match group members mapped to the attribute named **attr**. Usually this attribute is named '**member**' in most LDAP implementations. Below is an example of filter which will grant authorities matching the group description (USA) and user ID matching the one provided by actual user during login.

```
(&(description=USA)(member={0}))
```

If you don't need this flexibility, just leave group filter to its default **member={0}**. Once you specify LDAP settings and click Apply button, you should see a response list of users that will be available to pass the authentication.

After settings up the LDAP binding are accepted, you can also enable authorization :



logFACES ENTERPRISE LOGGING SUITE

Administration Console
v3.1.0 (build #2083)

[Connectivity](#) [Security](#) [Database](#) [Context](#) [Email](#) [Notifications](#) [License](#) [Status](#)

LDAP binding | Authorization mapping | Test login

Authorize users after successful login

[Import groups from your DS](#)
[Clear current groups mapping](#)

When enabled, users will be granted access to the logs of applications and loggers mapped to the imported user groups from your DS. The mapping should be specified in the format **application:logger**. Use * for wildcard matching and comma separator for several tokens. For example, *:* would match all applications and all loggers.

User group	Mapping (comma separated list of application:logger tokens)	Delete
developers	APP*:com.moonlit.logfaces.*	<input type="checkbox"/>
managers	*:*	<input type="checkbox"/>
qa	QA*:com.moonlit.logfaces.test.*	<input type="checkbox"/>

Apply Undo

Figure 2.5: Authorization mapping

Clicking on “Import groups from your DS” link will bring up a list of user groups available in your LDAP tree as specified by **Group base DN** attribute. For each user group you can assign a list of applications and loggers to be accessible by the members of the group. It goes as a comma separated list of pairs like this:

APP1:LOGGER1, APP2:LOGGER2

where tokens can use * wild cards, for example APP1*:com.company1.*, APP2*:com.company2.*. In the figure above, users from 'developers' group will be granted access to all logs relevant to applications starting from APP and loggers starting from com.moonlit.logfaces while managers will have access to everything. If you remove the group, its members will be denied access entirely.

Note that both authentication and authorization are optional. When authentication is disabled, authorization, of course, is not applicable. However, you can enable authentication while allowing all authenticated users to access any application logs by disabling authorization.

Finally, you can test all these settings in “Test login” page by entering actual user name and password. The server response will help you to verify that security is setup as expected.

2.7.4 Database

In Database section we specify how logFaces server should work with your database.

logFACES
ENTERPRISE LOGGING SUITE

Administration Console
v3.1.0 (build #2083)

[Connectivity](#) [Security](#) **Database** [Context](#) [Email](#) [Notifications](#) [License](#) [Status](#)

General options **Commit criteria**

Day capacity	1 Month	Keep this amount of days of log data and remove older records
Batch commit size	100	Number of records to buffer for batch commits
Number of commit failures	10	How many failed commits should trigger recovery mechanism
Recovery attempts rate	1	Rate of database reconnection attempts in minutes
Reconnection attempts	30	How many reconnection attempts to make before disabling database
Maintenance schedule	0 15 10 ? * 6#3	Cron job expression for maintenance (empty = no maint)

Apply Undo

Figure 2.6: Database settings

Database **capacity** is specified in days of log. If you specify "1 week" for example, then latest week of data will always be available. As time goes, older records are automatically removed while new ones are appended. You should carefully specify this value according to your needs; it affects overall performance as well as disk space usage. When using MongoDB with capped collections, this field will be automatically set to 'unlimited' to let MongoDB manage the storage size automatically.

Batch commit size is the size of the buffer used to insert log statements into database as a batch. The smaller the buffer the more frequently commits will be performed. Depending on the data inflow intensity, the buffer should be adjusted in such way that it does less frequent commits. On the other hand, large commit buffer size could be stressful for the database. Optimal sizes are usually in range of 50 - 500. You should use higher number if your system has frequent spikes of log data, this will greatly improve the performance of server overall. Half full commit buffers will be committed with a timer job running every minute.

Number of commit failures specifies how many commits can fail in a row to trigger recovery mechanism. This mechanism is designed specifically for situations when database goes down for maintenance or temporarily unavailable for some other reasons.

Recovery attempts rate specifies how frequently to try reconnection with database during recovery.

Number of reconnection attempts specifies how many times to try before giving up on database and switching to a router mode. In the screen shot above, the recovery will run for 30 minutes trying to reconnect every minute. If during this time database comes back, everything will continue as normal. Note that during recovery process, incoming log statements are persisted on local disk and flushed into database when it comes back. When database is unavailable for a long time while application log keeps coming, there could be quite large amount of those backed up records.

Maintenance schedule is an optional [cron expression](#) which will trigger a database maintenance job. Depending on the database the task is different. For example, in case of embedded database, its storage will be compacted and indexes rebuilt, other external SQL databases will have their indexes rebuild. In case of MongoDB there will be a 'repair' task which is very important for releasing the disk space to the OS since MongoDB doesn't release its storage by default and grows indefinitely. Note, however, that when capped collections are used, there is no need for any maintenance and this field will be hidden. No action will be taken if cron expression is not specified.

Commit criteria lets you specify which events the server should persist into the database. Criteria is a simple collection of Boolean rules which you can manipulate to achieve a fine tuned filter.

logFACES ENTERPRISE LOGGING SUITE

Administration Console v3.1.0 (build #2083)

[Connectivity](#) [Security](#) [Database](#) [Context](#) [Email](#) [Notifications](#) [License](#) [Status](#)

General options **Commit criteria**

& loggerLevel higher or equals TRACE

and domainName is not LFS

or

& loggerLevel higher or equals WARN

and domainName is LFS

[Add rule set](#)

Apply

Figure 2.7: Commit criteria

The example above will insure that everything except LFS application is persisted, however if LFS application emits some warnings (or above) events, we still want them saved.

2.7.5 Context

One of the advanced features in many logging systems is diagnostic context attached by the application to logging events. In log4j and its other flavors there is MDC – Mapped Diagnostic Context. You can read more about it [here](#). To provide convenient integration with MDC's, logFaces lets you map your application context properties in such way so that they could later be used in queries and other displays. You can specify up to 5 different context variables which will be automatically extracted from logging events and stored in database.

logFACES
ENTERPRISE LOGGING SUITE

Administration Console
v3.1.0 (build #2083)

[Connectivity](#) [Security](#) [Database](#) [Context](#) [Email](#) [Notifications](#) [License](#) [Status](#)

Mapped Diagnostic Context Custom logger levels

Here you can map your application MDC names with logFaces server. At run time, the server will look for those properties in logging events and automatically map them to corresponding database columns. Those names will also appear on client for doing real-time filtering and database queries.

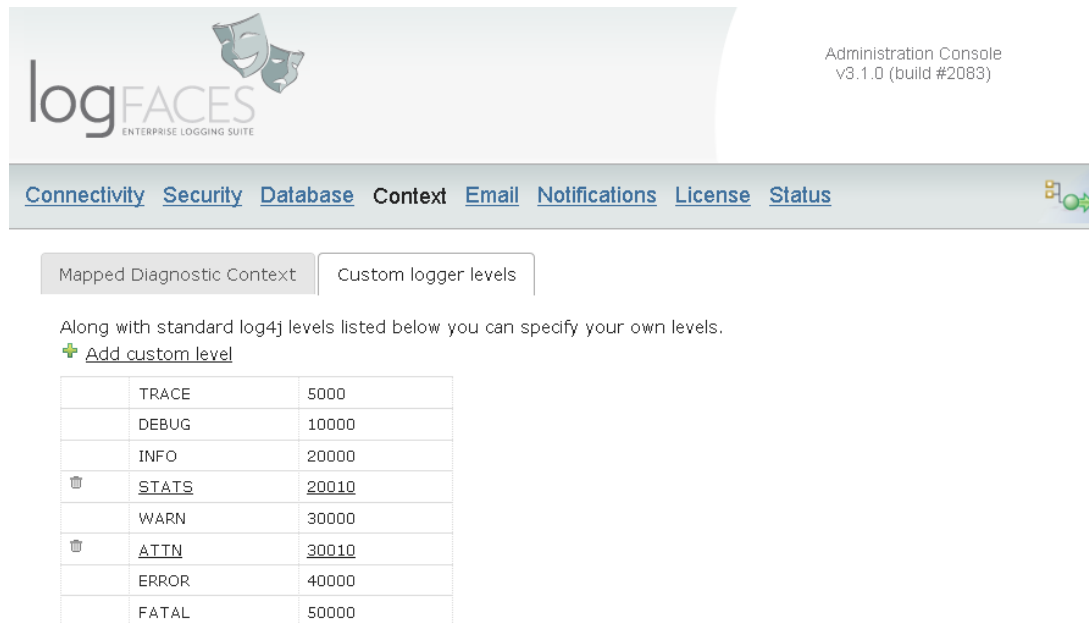
mdc1	userID
mdc2	sessionID
mdc3	targetID
mdc4	
mdc5	

Apply Undo

Figure 2.8: MDC mapping

You can modify those names any times during run time, but it's best to setup the mapping as earlier as possible. For example, if you initially worked with sessionID as first MDC parameter and then decided to rename it to sessionName, you won't be able to extract previously saved values as sessionID because it will be named differently.

Often there is a need to use custom severity levels in addition (or even instead of) default levels provided by log4j. If this is your case, there is a way to specify custom levels in logFaces server :



logFACES
ENTERPRISE LOGGING SUITE

Administration Console
v3.1.0 (build #2083)

Connectivity Security Database Context Email Notifications License Status

Mapped Diagnostic Context Custom logger levels

Along with standard log4j levels listed below you can specify your own levels.
 + Add custom level

	TRACE	5000
	DEBUG	10000
	INFO	20000
+	<u>STATS</u>	<u>20010</u>
	WARN	30000
+	<u>ATTN</u>	<u>30010</u>
	ERROR	40000
	FATAL	50000

Figure 2.9: Custom levels

Note how **STATS** and **ATTN** levels are defined – these are custom levels in addition to defaults. Once you add custom levels, they become available throughout the system – in filters, reports, triggers, etc. Those levels will also appear on client side so that users can utilize them in queries and displays.

2.7.6 Email

Email settings will allow logFaces server to send e-mails when required by reports and triggers. Here you define outgoing SMTP properties and also can verify that these settings are correct by sending test email to some recipient.

Administration Console
v3.1.0 (build #2083)

[Connectivity](#) [Security](#) [Database](#) [Context](#) [Email](#) [Notifications](#) [License](#) [Status](#)

Those are SMTP settings to be used for sending out e-mail notifications and reports. If you are intending to use reports and triggers, make sure that outgoing email is setup properly.

[Test outgoing email](#)

Server	<input type="text" value="smtp.gmail.com"/>	Host name or IP address of SMTP server
Port	<input type="text" value="587"/>	SMTP server port
Account name	<input type="text" value="me@foo.com"/>	Account name for authentication (leave blank if not required)
Password	<input type="password" value="••••••••"/>	Account password for authentication (leave blank if not required)
Use secure connection :	<input type="text" value="TLS"/>	Specify what security connection SMTP server requires

Figure 2.10: Outgoing e-mail settings

Click on "Test outgoing e-mail" link to verify that logFaces can send e-mails successfully. Should anything go wrong, you will be shown an error describing the cause. If everything was correct, you will receive an acknowledging email.

2.7.7 Reports

Reports are custom log files that server periodically generates according to the schedule and criteria query. Reports are then emailed to the recipients of your choice. Reports are organized in a list where you can see the overall information. Reports can be enabled or disabled – the rightmost icon indicates that second report is disabled in the example below. Disabled report stays in the system but doesn't actually do anything until you enable it.

Administration Console
v3.1.0 (build #2083)

[Connectivity](#) [Security](#) [Database](#) [Context](#) [Email](#) [Notifications](#) [License](#) [Status](#)

Scheduled reports Real-time triggers

Reports are scheduled queries which server will execute and deliver results to email recipients. Along with delivery options you can specify complex query to cover certain period of time. One of the typical examples of reports usage is to schedule daily list of errors, exceptions or other interesting items. Fired periodically such report will bring you fresh information daily.

[+ Add new report](#)

	Report name	Cron expression	Next fire time	
●	logFaces issues	0 0 0 ? * *	19-12-2011 00:00:00	
●	logFaces daily	0 0 0 ? * *		

Figure 2.11: Report list

The list also shows the cron expression which drives the report schedule, its closest fire time and links to manage each report individually.

Note that reports can also be individually tested (click on icon), - this is quite useful because it lets you receive real data instantly without waiting for the complex cron expression to trigger the report.

Each report comes with a bunch of parameters explained below as well as its criteria query. The query will be executed when report trigger fires. The results of the query will be then packaged into a log file, zipped if necessary and sent over to your recipients.

The example below demonstrates a typical report – it will fire every midnight and if there is anything in the query, it will send an email to our support. It covers past 24 hours and flags high email priority.

Look at the query it does – we want WARN+ events coming from com.moonlit.logfaces package.

Administration Console v3.1.4 (build #2242)

Report Editor

Delivery options Database query

Enabled	<input checked="" type="checkbox"/>	When report disabled, it won't be fired but will stay in the list
Report name	My report	For management purposes
Cron expression	0 0 0 ? * *	Specifies when report will be fired and how it will repeat
Time range to cover	24.0	Hours to cover since trigger time (fractions allowed)
Email to	me@gmail.com	Email report to these addresses, use ; for separation
Email subject	Problems in \${hostName}	Accepts <code>\${variable}</code> notation for dynamic text creation
Mail priority	highest	Email message will be tagged with this priority
Zip threshold	100	When attachment is larger than specified in KB, it will be zipped
Log file layout	%-5p %d{dd-MMM-yyy}	Attached log file layout (see log4j spec for details)

Save Cancel

Figure 2.12: Report delivery options

Administration Console v3.1.4 (build #2242)

Report Editor

Delivery options Database query

This query will be executed when report is fired:

& loggerLevel higher or equals WARN

and loggerName is com.moonlit.logfaces

Add rule set

Save Cancel

Figure 2.13: Report query

Enabled - when unchecked, the report will stay in the system but will never fire. You can enable it any time and it will fire at the next schedule slot.

Report name – this field is used only for the management purposes, give it a friendly name so that you could easily find the report in the list.

Cron expression – is an expression which specifies when and how to fire the report. Cron expressions are very flexible and used to make fairly complex scheduling rules. You can get more information about cron expressions [here](#).

Time range to cover – this specifies the time range which report query will cover, the count begins from the actual trigger time backwards. For example, if you want to cover single day and your report is fired daily, specify 24 hours.

E-mail to - list of recipients to receive the e-mail (use semi column as separator)

E-mail subject – this text will be used in email subject when report is dispatched. Note that you can use `${variable}` notation here where variable could be **domainName**, **hostName**, **loggerName**, **message** and any of the mapped MDC names. When report is built, this variable will be substituted with the corresponding value taken from the first log statement in the report.

Mail priority - e-mails can be flagged with standard e-mail priorities (highest, high, normal, low, lowest).

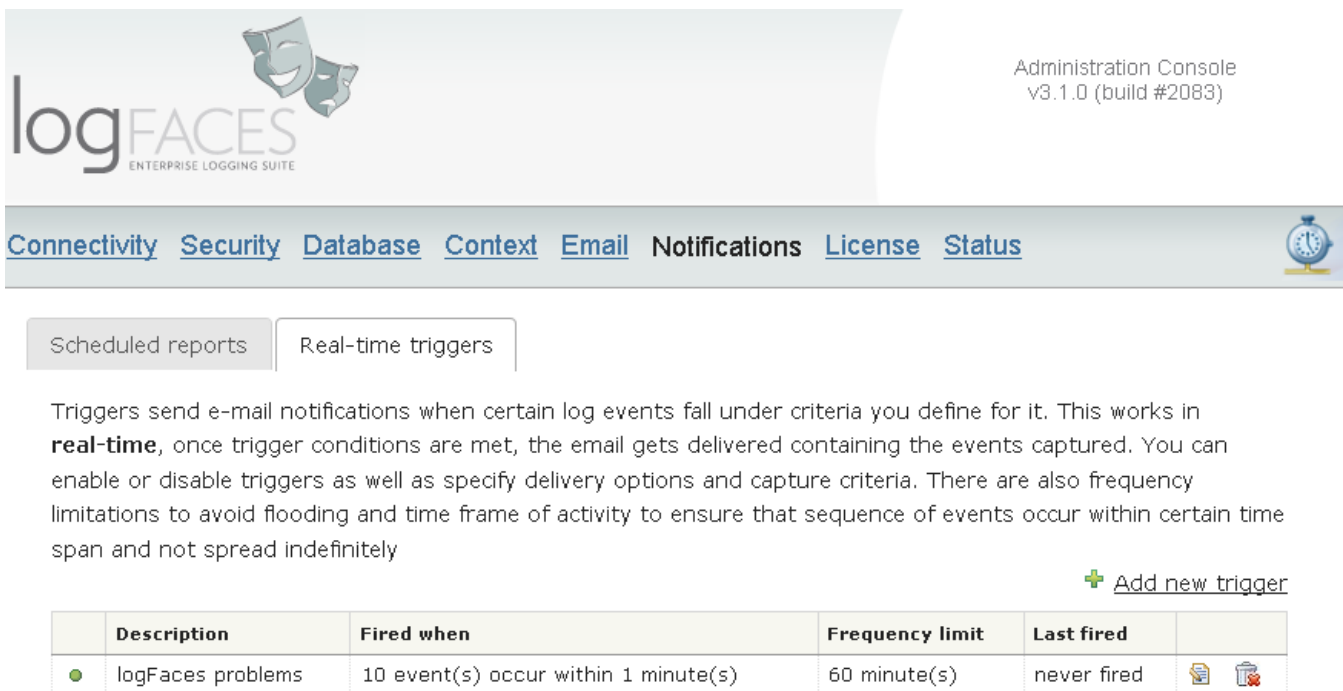
Zip attachments - specify a maximum size of log file in KB; if attachment file will be larger than specified, it will be automatically zipped.

Layout – specifies how to layout the text in the log files. LogFaces is using log4j formatting rules; you can find more details [here](#)

2.7.8 Triggers

Triggers are similar to reports except that they are not scheduled but fired immediately when certain conditions met. Conditions are based on the log data going through the server. By specifying criteria you will be able to detect very particular log statements from very particular sources. In addition to this, you can also specify how many of such events to capture and within what time span they should be in order to fire the trigger.

Like reports, triggers are listed to give you an overall view of what triggers are there, what is enabled and when and how they get fired.



Administration Console
v3.1.0 (build #2083)

[Connectivity](#) [Security](#) [Database](#) [Context](#) [Email](#) [Notifications](#) [License](#) [Status](#)

Scheduled reports | Real-time triggers

Triggers send e-mail notifications when certain log events fall under criteria you define for it. This works in **real-time**, once trigger conditions are met, the email gets delivered containing the events captured. You can enable or disable triggers as well as specify delivery options and capture criteria. There are also frequency limitations to avoid flooding and time frame of activity to ensure that sequence of events occur within certain time span and not spread indefinitely

[+ Add new trigger](#)



	Description	Fired when	Frequency limit	Last fired	
●	logFaces problems	10 event(s) occur within 1 minute(s)	60 minute(s)	never fired	 

Figure 2.14: List of triggers

The example below demonstrates a trigger intended to react to exceptions from a certain application. Note that this trigger is setup to react only to 10 consecutive exceptions occurring withing 1 minute time span. It will not react if there is say only 9 exceptions registered in one minute. This way we don't want to jump on every exception, we want to see a bunch of issues coming our way to take action. Also, we don't want flood of e-mails, it's enough to get an e-mail at least once in an hour.

Administration Console
v3.1.4 (build #2242)

logFACES
ENTERPRISE LOGGING SUITE

Trigger Editor

Delivery options | Capture criteria

Enabled	<input checked="" type="checkbox"/>	Trigger will not fire unless enabled.
Name	logFaces problems	For management purposes only . Will show in emails and attachment.
Recipients	me@gmail.com	E-mails will be sent to those recipients (comma separated)
Email subject	logFaces problems	Accepts <code>{variable}</code> notation for dynamic text creation
Email priority	normal	Emails will be sent with this priority for convinence
Log file layout	%-5p %d{dd-MMM-yy}	Attached log file layout (see log4j spec for details)
Counter	10	Will fire when this number of log events captured in sequence
Time frame	1	Will fire only when events are captured within this time frame (minutes)
Frequency limit	60	Will NOT fire more often than specified (minutes).

Save | Cancel

Figure 2.15: Trigger example

Trigger Editor

Delivery options | Capture criteria

The trigger will capture log statements which match this criteria:

&	thrown	is	true	🗑
and	domainName	is	LFS	🗑

[Add rule set](#)

Save | Cancel

Figure 2.16: Trigger query

Enabled – the trigger will not fire if disabled, but will stay in the system

Name - textual information which will be sent out with an email

Recipients - to whom the trigger may concern - comma separated list of e-mail recipients

Email priority - e-mail priority (highest, high, normal, low, lowest)

Email subject - this text will be used in email subject when report is dispatched. Note that you can use `${variable}` notation here where variable could be **domainName**, **hostName**, **loggerName**, **message** and any of the mapped MDC names. When report is built, this variable will be substituted with the corresponding value taken from the first log statement in the report.

Layout - the layout of log file which will be attached to the email.

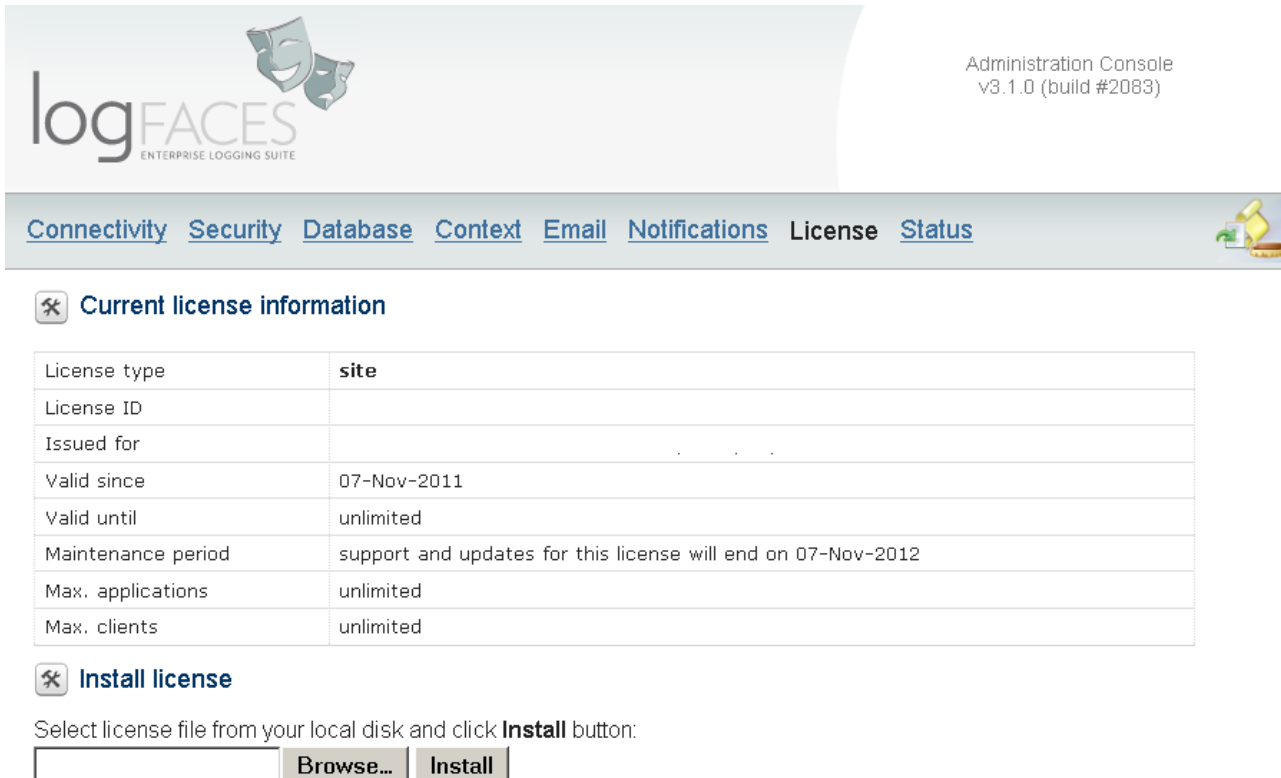
Counter - the trigger will fire only when at least this many events are trapped by the criteria.

Time frame - the trigger will fire only when events are captured within this time frame in minutes.

Frequency limit - the trigger will not fire more often than specified by this value in minutes.

2.7.9 License

License tab displays currently installed license information as well as allows you to install new license file. When you install logFaces Server for a first time, it automatically activates one time trial evaluation period for 30 days. If you decide to purchase a license, the license file should be submitted through this form:



logFACES
ENTERPRISE LOGGING SUITE

Administration Console
v3.1.0 (build #2083)

[Connectivity](#) [Security](#) [Database](#) [Context](#) [Email](#) [Notifications](#) [License](#) [Status](#)

Current license information

License type	site
License ID	
Issued for	
Valid since	07-Nov-2011
Valid until	unlimited
Maintenance period	support and updates for this license will end on 07-Nov-2012
Max. applications	unlimited
Max. clients	unlimited

Install license

Select license file from your local disk and click **Install** button:

Browse... **Install**

Figure 2-2.17 License Panel

What happens when evaluation license expires? LogFaces Server will shutdown its engine and only allow Administration Console access; applications will not be able to use the server and clients won't be able to connect to it. When you install new license, the engine should be started manually. This can be done by simply restarting the service from command prompt or control panel, or from the Status panel link – see the next section.

2.7.10 Status

Status tab contains useful health monitoring information of the Server and allows basic instrumentation tasks.

Property	Value	Actions
Server version	3.1.3.2185	check for updates
Server license	138c28db-c225-4dbb-ac4f-3c1deeab5636	
Operating system	Windows XP (5.1), x86, 1.7.0 (UTF-8)	
Engine status	running (003:47:48)	stop restart
Last errors	none	
Internal log	Quiet (enable verbose)	download log file
JVM heap memory	max: 494 MB, total:247 MB, free: 219 MB	run garbage collection
Number of threads	39	fetch stack traces
Inflow rate (events/sec)	last: 0, average: 0, maximum: 0	reset stats
Database throughput (events/sec)	last: 65, average: 65, overload: 0.00	
Connections	users: 1, apps: 0	
Connected to database	yes	
Number of records	2 (update count)	re-create database
Database storage size	1 MB	
Database product	Apache Derby, 10.7.1.1 - (1040133)	
Database driver	Apache Derby Embedded JDBC Driver, 10.7.1.1 - (1040133)	
Next database maint time	not scheduled	do maintenance now

Figure 2-2.18 Status panel

Most of the information in this tab should be self explanatory for technical people. We will just mention the important instrumentation actions which are available from the links on the right side of the table:

- **Check for updates** will try to detect updates we regularly post on our web site. If update is available, the new version number and package size will be displayed. Note that this operation requires live internet connection. Updates can also be automatically downloaded and installed – when new version is detected, you will get a link to activate the installer. The process is fully automated but will take your server offline for a few moments.

- **Engine start/stop**; sometimes it's required to put the server down without actually shutting the process down. One of the typical uses of this option is when trial license expires. In such case, the logFaces Server will start so that you would be able install proper license, but its engine will be down and no logging will be taken from applications.
- **JVM restart** will do full restart of server JVM, this normally takes few seconds and may be required in some situation.
- **Run garbage collection**; explicitly call garbage collection now
- **Fetch stack traces** will download full dump of all threads currently running on server.
- **Last errors** is a list of latest errors encountered by server, you can browse through them to see if anything went wrong lately, or simply reset them.
- **Internal log** can be tuned to verbose or silent mode. You can also download an internal log file for inspection. This file can be sent to our [support team](#).
- **Inflow rate** represents the throughput of the server on the network side, it indicates the amount of logs flowing into logFaces from the appenders in second.
- **Database throughput** indicates how much data your database can commit per second. You should keep an eye on this metric to be below **Inflow rate** most of the time. When database throughput is significantly lower than inflow rate for a long time, the data will be stored in local disk storage – normally this is an expensive operation and may result in higher than usual CPU and IO use. **Overload** is the percentage ratio of total number of events went through an overflow buffer on local disk to a total committed. This ratio is very important for detecting the database bottle neck. When overload gets too high, it will be emphasized in red color. The default threshold is set to 10% but you can adjust it in environment properties (see paragraph on advanced configuration). You will also see a flag icon indicating that currently server handles its internal overflow cache trying to push it into the database.
- **Re-create database** allows to remove all database records; be careful with this operation, it is not recoverable and can't be undone

- **Update records counter** will re-count total number of log records stored in your database. Because counting with some databases is very expensive operation, this action is set for explicit user request.
- **Database maintenance** – will start the database maintenance job whether its currently scheduled or not. Depending on the database type, the task is different. For example, in case of embedded database, its storage will be compacted and indexes rebuilt, other external SQL databases will have their indexes rebuild. In case of MongoDB there is optional 'repair' task and optional index rebuilding.

2.8 Advanced setup

There are some settings which can not be configured through the Administration Console but can be configured manually in several configuration files. These settings should only be modified when server is down. Please review the settings carefully before doing any manual change. It's a good idea to keep a backup of the files you are intending to modify.

2.8.1 Environment - /conf/environment.properties

Environment properties are fed into server JVM upon start up, prefix is **com.moonlit.logfaces**

Property	Mandatory	Description
.config.server	yes	Points to a main configuration file
.config.mongodb	no	If set to true then server will work with MongoDB defined in /conf/mongodb.properties file. Default is false.
.config.hibernate	yes	Points to hibernate configuration file
.config.schema	yes	Points to schema file which will be created in database.
.config.jobs	yes	Points to jobs configuration file
.config.reports	yes	Points to reports configuration file
.url.revision	no	URL for checking software updates
.url.downloads	no	URL for update downloads
.url.updates	no	URL to version update descriptor
.url.notes	no	URL to version release notes
.monitoring.highThreadCount	no	Maximum number of threads the server should be able to sustain, if during run time the amount of threads will be higher, the server will issue an internal warning. Default is 300.
.monitoring.lowMemoryThreshold	no	Minimum of free heap memory specified as a percentage of maximum heap memory. When free memory will go below this value, the server will issue a warning which will appear in Admin. console. Default is 2.
.monitoring.overloadThreshold	no	Threshold for detecting database overload – ratio between overflow and total commits. Default is 10%.
.monitoring.overflowSize	no	Size of the local overflow queue. When specified size is reached, the overflow queue will refuse to queue more events and will raise system error. Default is 500000 events.
.security.keyStore	no	File name for SSL certificate key store
.security.trustStore	no	File name for SSL key trust store
.security.keyPass	no	Password for key store (plain or obfuscated)
.security.trustPass	no	Password for trust store (plain or obfuscated)

Table 2.1: Environment properties

2.8.2 Main configuration file - /conf/lfs.xml

This file is updated by Administration console, there should be no need to modify it manually in normal circumstances.

Element or attribute	Description	Default
remoting/host	Host or IP address for socket binding	any-address
remoting/port	Port number for external client connections	8050
remoting/ssl	If true, the SSL will be enforced	false
remoting/sessionTimeout	Timeout in minutes after which inactive clients are kicked off.	2 minutes
email/emailServer	Outgoing e-mail server name to use when sending e-mail from logFaces	
email/emailAccount	Outgoing e-mail account name to use when sending e-mail from logFaces	
email/emailPassword	Outgoing e-mail account password to use when sending e-mail from logFaces	
hub/host	Host or IP address for socket binding	any-address
hub/ssl	If true, the SSL will be enforced	false
hub/receiver/name	Name of log receiver – XMLTCP and/or XMLUDP	
hub/receiver/class	Receiver class, internal use only, don't modify.	
hub/receiver/port	Receiver port number	TCP: 55200 UDP: 55201
hub/receiver/enabled	Enables/disables receiver	
database/commitBuffer	Size of commit buffer used to insert log events into database in a batch operation.	100
database/dayCapacity	Capacity of data storage in days	3
database/recoveryRate	When unable to commit to database, this parameter specifies how often to retry re-connection in minutes.	1
database/nofRecoveryAttempts	How many retries to make in order to re-connect to database after failure before giving up.	5

Table 2.2: Main configuration file

2.8.3 How do I work with external SQL databases?

You have to obtain relevant database driver from your database vendor and place the jars in `/lib/dbdrivers` directory on server. **Our installation only include drivers which permitted by publisher's license.** We do our tests for Oracle, MySQL, SQL Server, DB2 and PostgreSQL, but theoretically there shouldn't be a problem to work with other relational databases as well. It's only a matter of configuration and database driver you wish to use.

Look at `/conf/environment.properties` file – you should see these two properties pointing out to hibernate configuration file and database schema :

`com.moonlit.logfaces.config.hibernate=${ifs.home}/conf/hibernate.properties`

`com.moonlit.logfaces.config.schema=${ifs.home}/conf/ifs.sql`

You can modify these references by pointing to different files, but make sure the files are correct. Our distribution contains both hibernate examples and database schema for all databases we support at this moment. The example above will use embedded database driver settings and Derby schema. If, for example, you would like to use PostgreSQL, modify `environment.properties` as follows :

`com.moonlit.logfaces.config.hibernate=${ifs.home}/conf/hibernate-postgree.properties`

`com.moonlit.logfaces.config.schema=${ifs.home}/conf/ifs-postgree.sql`

Then modify hibernate properties file to point to your database. LogFaces Server uses [Hibernate](#) ORM framework; it is recommended to have good knowledge of this framework before you decide to make re-configuration.

Once you prepared the configuration, simply restart the logFaces server. During the start up new schema will be published automatically. After the start up, open **Administration Console** and navigate to the **Status** tab. If everything went well, you should see that engine is started, database connection is OK and there are versions of database and driver in the table. If something goes wrong and there are problems with database connection, you will see red marks and last error numbers. To see what happened, click on "show last errors" or "download log file" link; the problem is usually related to a configuration error, typo or perhaps your database is not responding as logFaces expects. If you're unable to figure out the problem yourself, submit this log file or an error code to our support site and we will try to help.

Another, perhaps easier, option is to run the server in console mode and see that there are no

exception thrown during it's start up. On Windows you can run server in console mode from `/bin/run.bat` on. On Linux, run the following command: `./lfs console` while in `/bin` directory.

2.8.4 Can I modify SQL database schema?

Yes, you can do that to a certain extent. Database schema file is referenced in `/conf/environment.properties` through property named `com.moonlit.logfaces.config.schema`.

You can not modify the structure of the tables or column names because they're mapped through hibernate mapping in the code. But you can adjust column size constraints, modify indexes or add some additional statements as long as they don't break the mapping.

After you changed the schema file you need to re-create the database. This can be done either from administration console status tab, or by using some other external tool. Note that when using embedded database, there is no other choice but using the admin. console.

IMPORTANT: If you care for the existing data in the database, make sure to back it up before doing any changes to the schema. One of the options is to use our backup utility described next.

2.8.5 Can I use my own PK generators with SQL databases?


Yes. By default PKs are generated with simple native generator defined in hibernate mapping files which are not exposed for general usage. If you want to generate your own PKs, the mappings need to be altered and server needs to know which mapping to use. The procedure to follow:

1. extract `Event.hbm.xml` and `Repository.hbm.xml` from `/lib/lfs.jar`
2. modify the 'id' column generator section as you need, this is standard hibernate format
3. create separate jar file and place modified mappings along with your own classes
4. place this jar under `/lib` directory on server
5. enforce those classpath resources by changing `/conf/environment.properties`:

```
com.moonlit.logfaces.resources.eventMapping = classpath:/your/Event.hbm.xml  
com.moonlit.logfaces.resources.repoMapping = classpath:/your/Repository.hbm.xml
```
6. run server in console mode (`/bin/lfs.bat`) and watch the log output. If there is something wrong in the mapping or something is missing, the database layer will not start and there will be many errors.

2.8.6 How do I work with MongoDB?

Starting from version 3.0 logFaces server can work with MongoDB – one of the NoSQL tools available today. To enable this you need to modify `/conf/environment.properties` file and specify `com.moonlit.logfaces.config.mongodb=true`. Once enabled, the server will ignore all other settings related to hibernate and schema – logFaces can work either in SQL or NoSQL modes exclusively. Next step - start the server and go to admin/status page. It's very likely that logFaces won't find your database out of its defaults, so you will see that database is down. Then go to database tab and adjust MongoDB parameters, most importantly the connection end points :



[Connectivity](#) [Security](#) [Database](#) [Context](#) [Email](#) [Notifications](#) [License](#) [Status](#)

General options | **MongoDB options** | Commit criteria

[+ New database instance](#)

State	connected	Current connection to MongoDB
Database name	<input type="text" value="logfaces"/>	logFaces database name in MongoDB space
Connection(s)	<input type="text" value="localhost:27017"/>	Comma separated list of host:port pairs
Secured	<input type="text" value="No"/>	Database access requires authentication
Write concern	<input type="text" value="NORMAL"/>	Controls the driver write behavior
Error check rate	<input type="text" value="20"/>	How often to force getLastError call
Slave OK	<input type="text" value="No"/>	Propagate queries to replica slaves
Sort queries	<input type="text" value="Yes"/>	Sort result sets by time
Storage type	<input type="text" value="Capped collection"/>	Type of collection for storing log data
Storage size	<input type="text" value="100"/>	Size of capped collection in MB

Once applied, the server will try to connect to MongoDB daemon and create its database and collections there. If this will be successful, you could then use MongoDB shell and verify that database is created as you named it and it has two collections there - "log" and "repo". The former is where you log data will be stored, the later is where the repository info will be stored – host names, app names, etc. Basically that's all there is to it. You can always get back to this page later and adjust the settings. In most cases the changes applied instantly.

Note that parameters on this page come from and saved to `/conf/mongodb.properties` file on your server. You will rarely need to edit them manually, but it's a good reference to understand the meaning of these parameters. Table below is a complete list of properties we use to integrate with MongoDB, it will help you to understand how to use connections, replica sets, indexes and other important features:

Property name	Description	Default
.connection	Comma separated list of host:port pairs for specifying replica sets.	localhost:27017, localhost:27018, localhost:27019
.user	User name for authentication. Leave blank if your database doesn't require security.	
.password	Password for authentication. Leave blank if your database doesn't require security.	
.writeConcern	By default driver sends data to the MongoDB and if it is received it is OK. Sometimes you want more than that. It might be very important that content is stored in one node or maybe even in multiple nodes. By providing a writeConcern you can specify how much safety you want. NORMAL is the default write and forget. SAFE only waits for the master by using the getLastError. REPLICAS_SAFE does the same but also waits for the slaves - be careful in a test environment without slaves, it will wait and wait. FSYNC_SAFE force fsync to disk on all operations.	NORMAL
.getLastErrorRate	Enforces getLastError call on every X inserts. This parameter may have an affect on heavy loads and improve the database write throughput. When 0 is specified, getLastError will never be called. Otherwise getLastError will be forced on every X inserts specified here.	20
.slaveOK	When querying a replica set, requests routed only to the master by default. To permit queries against slaves set this parameter to true. Only affective when replica sets are used.	false
.dbname	Name of the database logFaces will create	logfaces
.capped	When set to 'true' logFaces will force the capped 'log' collection and will try to convert it to capped when it's not, and will fail if it can't be converted.	false
.cappedSize	Mandatory size of capped collection in MB. Affective when 'capped' is set to true	50 MB
.ttl	When set to true logFaces will try to convert the collection to TTL unless it's already TTL.	false
ttlDays	Number of days for TTL collection. Affective only when TTL collection is enabled	7

2.8.7 MongoDB schema and indexes

When started, the server will automatically create two collections named '**log**' and '**repo**'. Log collection stores log events in such way that each event corresponds to exactly one MongoDB document. This collection grows rapidly as data gets committed. Repo collection is a helping repository to keep names of applications, hosts and loggers – this information is used mostly by clients to assist with selection of items for queries and hardly ever grows once the system stabilizes.

Documents stored in **log** collection have attributes corresponding to the fields of actual logging event. In order to save storage space those names are shrank to a minimum, so when you look into collection, you will see extremely short names. Below is the mapping of these names:

```
{
  "_id" : ObjectId("51abeff2e0fd7bab0f2cf20c"), // object ID (generated by mongodb)
  "t" : ISODate("2013-06-03T01:22:12.794Z"), // loggerTimeStamp (creation time)
  "q" : NumberLong(3256236), // sequenceNumber (generated by lfs)
  "p" : 5000, // loggerLevel (severity level)
  "r" : "355941742@qtp-2017211435-7749", // threadName (originating thread name)
  "m" : "This is the message", // message (message body)
  "h" : "my host", // hostName (originating host name)
  "a" : "my app", // domainName (originating app name)
  "w" : false, // thrown (exception true/false)
  "g" : "com.myapp.logger", // loggerName (originating logger name)
  "f" : "Myclass.java", // locFileName (location file name)
  "e" : "MyMethod", // locMethodName (location method name)
  "l" : "489", // locLineNumber (location line)
  "c" : "com.myapp.logger", // locClassName (location class)
  "p_targetID" : "1360911352", // custom MDC mapped to targetID
  "p_sessionID" : "sid4454.col080" // custom MDC mapped to sessionID
}
```

Because logFaces allows usage of any of the above attributes in queries, you will need to carefully select [compound indexes](#). One important thing to keep in mind when dealing with compound indexes is that the **order of the attributes** in index is crucial to index performance. For example indexes (t,p) and (p,t) may have completely different performance depending on the data store. First one will seek records in specified time range (t), and then reduce to severity level (p). Second one works in reverse – first seeking level (p) and then reduce by time (t). It is often desirable to have several compound indexes so that you cover as many frequently used queries as possible. Keep in mind that the aim of selecting right indexes is to **compromise** between storage size (indexes are greedy) and the end user satisfaction with the speed of queries. Use MongoDB console tools to examine [performance](#) details.

Nearly all queries generated by clients contain loggerTimeStamp (t) attribute. It means that in most cases you will want to include this attribute and position it correctly within the index.

2.8.8 MongoDB capped collections

From MongoDB site:

“Capped collections are fixed sized collections that have a very high performance auto-FIFO age-out feature (age out is based on insertion order). They are a bit like the “RRD” (Round Robin Database) concept if you are familiar with that. In addition, capped collections automatically, with high performance, maintain insertion order for the objects in the collection; this is very powerful for certain use cases such as logging.”

RRD doesn't ensure data storage by time (for example 3 days of data), it ensures that storage will not grow bigger in size than specified. However, having such tremendous performance advantage, capped collections could be a preference for some. If you are one of those, this is how you should setup logFaces to utilize this feature – go to admin/database/MongoDB tab, check Capped Collection option and specify size in MB. The server will automatically do the relevant conversion of 'log' collection into capped. It is also possible to convert to a capped collection directly from MongoDB shell. If you do this, make sure to modify **mongodb.properties** file accordingly before restarting the server!

When capped collections used, the database day capacity automatically becomes '**unlimited**' and logFaces will not manage the data store size. Maintenance of such database is also irrelevant since MongoDB storage size will not grow from the specified. Thus you will not see '**capacity**' and '**maintenance cron**' options in admin database page.

Note that switching from capped collections back to regular collections is not possible without manually backing up your data and re-inserting it back into a newly created regular collection. It is also not possible (as of this writing) to extend the size of capped collections without doing manual re-inserts. Again, make sure to modify **mongodb.properties** file before restarting the server, otherwise it will convert everything back to capped. In other words, **mongodb.properties** must always be adjusted manually **if you manually administer your database** otherwise it may override your changes.

2.8.9 MongoDB TTL collections

TTL stands for “Time To Live”. From MongoDB documentation:

“Implemented as a special index type, TTL collections make it possible to store data in MongoDB and have the `mongod` automatically remove data after a specified period of time. This is ideal for some types of information like machine generated event data, logs, and session information that only need to persist in a database for a limited period of time ”

To enable this feature with logFaces server, open administration page and modify “**storage type**” option to “**TTL collection**” as shown below. You will also need to specify the “**number of days**” for this collection to operate with.

The screenshot shows the 'MongoDB options' tab in the logFaces administration interface. The interface includes a navigation bar with links for Connectivity, Security, Database, Context, Email, Notifications, License, and Status. Below the navigation bar are three tabs: General options, MongoDB options (selected), and Commit criteria. A '+ New database instance' link is visible in the top right of the options area.

State	connected	Current connection to MongoDB
Database name	logfaces	logFaces database name in MongoDB space
Connection(s)	localhost:27017	Comma separated list of host:port pairs
Secured	No	Database access requires authentication
Write concern	NORMAL	Controls the driver write behavior
Error check rate	20	How often to force getLastError call
Slave OK	No	Propagate queries to replica slaves
Sort queries	Yes	Sort result sets by time
Storage type	TTL collection	Type of collection for storing log data
Number of days	7	Time to live of log data in days

At the bottom of the form are 'Apply' and 'Undo' buttons. Two blue arrows point to the 'Storage type' and 'Number of days' fields.

Figure 2.19: Setting TTL collection

Note that when TTL collection is used, the server will automatically set “Day capacity” to unlimited and will not remove older records from database on its own – this duty will solely be in hands of MongoDB.

2.8.10 MongoDB custom `_id` creation

By default, logFaces uses the `_id` generated by MongoDB driver, which is a standard [ObjectId](#). To override this behavior and create your own `_id` object follow the procedure below:

1. Implement `com.moonlit.logfaces.server.core.OIDGenerator` interface. The interface definition and other dependencies can be found in `/lib/lfs-core.jar`. Implemented class will be instantiated by reflection, so the default constructor must be present.
2. Place your implementation on server under `/lib` directory.
3. Tell the server to use your implementation instead of the defaults by changing this property in `/conf/environment.properties`:

```
com.moonlit.logfaces.resources.mongoIdGenerator = com.your.id.Generator
```
4. Run the server, feed some data and see that your implementation works as expected. If server fails to find or instantiate your implementation, it will fall back to the defaults silently.

2.8.11 How do I tune the server for the best performance?

First of all, to determine whether the server under-performs, you need some metrics. You will find them in admin status page, below is an example of the most interesting parameters:



Engine status	running (094:47:40)	stop engine
Last errors	4	show , reset
Internal log	Verbose (silence to quiet)	download log file
JVM heap memory	max:508 MB, total:254 MB, free:176 MB	run garbage collection
Number of threads	40	fetch stack traces
Inflow rate (events/sec)	last: 1, average: 1, maximum: 518	 refresh  reset
Database throughput (events/sec)	last: 300, average: 1122, overload: 2.52%	
Active clients / appenders	0 / 4	

Figure 2.20: Performance metrics

The most common performance factor is the memory factor. Note that free heap memory should never get close to about 2-5% of the maximum allocated for the server JVM; when this happens the memory monitor will issue a warning in “last errors”. Once JVM memory gets drained, its behavior gets unpredictable and eventually may result in a complete fault. For the details read section named “[How to increase server JVM memory](#)”.

Another common performance factor is the balance between **inflow rate** and **database throughput** rate. The inflow rate is the number of log events arriving to the server per second from appenders. Database throughput is the number of events your database commits per second. You will see these numbers changing quite a lot and depend on many factors. To name just a few - network speed, database performance, log statement sizes, etc. When inflow rate is higher than database throughput, the server will activate an overflow mechanism which delegates residuals to a local disk storage while database is busy committing. Residuals are then flushed into database whenever database permits. In general, this mechanism is designed to prevent intermittent load spikes and can't be effective when the inflow is permanently higher than the database throughput. When local storage grows large, the CPU and IO usage may get much higher than normal. The larger the local storage grows, the harder it becomes to get the data from there. To prevent this from happening on regular basis you should watch the amount of “overload” - the percentage of total inflow which went through the local storage before it reached your database.

Try to adjust commit buffer size in admin database page. Note that larger commit buffer sizes not necessarily improve the database throughput, it very much depends on the database setup. Usually, most databases start getting sluggish when they significantly grow in size – try to control this by

reducing day capacity or persisting less information.

Indexes - they play very significant role in insert rates, try to tune your indexes to fit best your queries. Removing unnecessary indexes can significantly improve the overall performance – you can modify indexes in schema templates for the database you use, or alternatively – manage them with external database tools.

Of course the amount of server connections and threads play their role too. Each remote client or appender will result in an additional thread. Depending on the hardware and OS, these numbers should be taken into account.

When you see that server obviously under-performs, consider to either reduce the amount of load, or add another logFaces node to split the load.

2.8.12 How do I backup my database storage?

You can perform backup of entire database storage into a binary file. This backup file can later be re-imported into another system. We provide a script to fully automate this process - `/bin/backup.bat` on Windows or `/bin/backup.sh` on Linuxes. Generally this is a good idea to stop the server during backup process, but this is only required for when you use embedded databases, external databases can be backed up even while server is running. To **import** backed up data into another system, go to client **Tools** menu and select "Import data into database". You will be prompted to select the backup file. Depending on the amount of data, the process may take some time to complete.

Note that this is not the most efficient way to do the backup and shouldn't be used with large volumes of data.

2.8.13 How can I automate monitoring of logFaces server logs?

logFaces server itself has its own internal log which you can monitor by means of triggers or reports. The application name used by logFaces server is "**LFS**" and root package of server classes is named "**com.moonlit.logfaces**". When specifying criteria you can use either of those parameters to detect problems in real-time by means of triggers, or obtain scheduled summaries by means of reports. Generally it's a good idea to have this set up before you contact our support, we normally will ask for the internal logs before getting into the details.

2.8.14 How do I increase server JVM memory?

Open `\bin\lfs.conf` file – this is a bootstrap configuration file.

JVM memory is setup with those attributes:

```
wrapper.java.initmemory=256
```

```
wrapper.java.maxmemory=512
```

Those values are default, if you experience extensive memory usage, try to increase maxmemory property value. Note that values are in mega bytes.

2.8.15 How do I use NTLM authentication with MS SQL Server on Windows?

1. Download latest jTDS driver from here <http://jtds.sourceforge.net/>
2. Place **jtds-xxx.jar** and relevant **ntlmauth.dll** under **/lib/dbdrivers** in your server installation
3. Modify **/bin/lfs.conf** so that JVM loads ntlmauth.dll by adding the library path property:
wrapper.java.library.path.2=..lib\dbdrivers
4. Comment out user name and password in hibernate properties file – this will ensure that NTLM authentication is used. Make sure that connection URL specifies correct domain.
5. If logFaces server runs as windows service, you will have to make sure that LFS service starts with proper user account and not default system account - modify these in service properties. Otherwise the driver will use currently logged in user credentials to connect to database.

2.8.16 How do I make logFaces win32 service depend on other services?

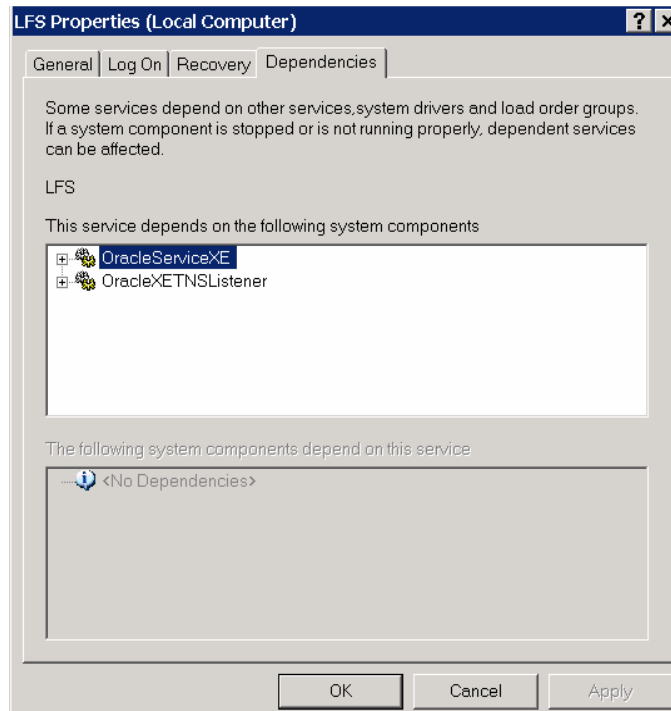
In Windows it's often required to have logFaces service dependent on other services during start up. For example, if you run database server on the same machine as logFaces server, you may want to make sure that it starts only after database successfully starts. To achieve that, you must specify service dependencies.

Open **/bin/lfs.conf** file and modify **wrapper.ntservice.dependency.xxx** properties accordingly. For example:

```
wrapper.ntservice.dependency.1 = SomeService1
```

```
wrapper.ntservice.dependency.2 = SomeService2
```

Make sure to preserve proper numbering order and specify correct names of dependent services. Then you will need to re-register LFS service. Make sure the service is not running; execute **/bin/uninstallservices.bat** and then **/bin/installservices.bat** – this will re-register service configuration in Windows registry. To verify that everything went OK, check out service properties in your Windows administration tools. For example, in case of Oracle, the dependencies should look similar to what is shown below. If everything looks OK, go ahead and restart your computer to verify that dependencies actually work.



2.8.17 How do I enable SSL for clients and browsers.

By default the access to the server from clients and browsers is based on plain HTTP connection. To use secure (SSL) based connection follow these steps:

1. Obtain public a private keys from trusted CA authority. Usually they come in a PEM format – text files with bunch of encrypted numbers. Note that it's no possible to use self signed certificates, so before you proceed you must have both these files handy.
2. The keys must be now imported into the trusted store which will then be used by logFaces server to present its identity and encrypt the traffic for the outside clients. To do this we need to do the following steps:
 - a) First we need to convert the keys into combined PKCS12 format, one of the tools which can be used for this is [OpenSSL](#) , here is the command you should execute:

```
openssl pkcs12 -inkey private.key -in mycert.crt -export -out mykeys.pkcs12
```

This will combine the private (**private.key**) and public (**mycert.crt**) keys into a single file named **mykeys.pkcs12**

- b) Now import the **mykeys.pkcs12** into the key store using conventional JDK keytool utility.

```
keytool -importkeystore -srckeystore mykeys.pkcs12 -srcstoretype PKCS12 -destkeystore mykeystore
```

You will be prompted for passwords – remember them for the next steps.

- c) Created **mykeystore** file should now be introduced to the logFaces server to be used. Open the `/conf/environment.properties` file and fill in the following properties:

```
com.moonlit.logfaces.security.keyStore = path_to_mykeystore
com.moonlit.logfaces.security.trustStore = path_to_mykeystore
com.moonlit.logfaces.security.keyPass = OBF:1x881v2h1t371shw1shmlt2z1v1x1x8m
com.moonlit.logfaces.security.trustPass = OBF:1x881v2h1t371shw1shmlt2z1v1x1x8m
```

Note that passwords can be entered in clear text or obfuscated like in the example above. To obfuscate your passwords you can use the existing Jetty utilities located in `/lib/jetty` folder. The command line below will obfuscate a text string:

```
java -cp jetty.jar;jetty-util.jar org.mortbay.jetty.security.Password my-password
```

3. At this point you can start the server. If everything is specified correctly, the clients and browsers will have to use SSL connection to communicate with the server. Please note that once configured to use SSL, the server will enforce only SSL connection, there is no other port for non-secure access – it's either SSL or not. If you can't get connected, try running the server in console mode or look in its internal log.

Note that `/conf/environment.properties` file has default presets for file names and passwords, so you might as well choose to use those names in the first place. The obfuscated passwords are '**logfaces**'.

3 Getting started with logFaces client

logFaces client application is available for Windows, Linux and Mac OSX for 32 and 64 bit architectures. Think of logFaces client as a pair of glasses you wear when looking into the log stream pouring through the log server. The client will assist you to make sense out of this stream and convert large amount of log data into a meaningful piece of work.

3.1 Installing

On Windows, download and run the installer which will walk you through the process. Linux distributions come as **tar.gz** archives, just unzip the archive and run **logfaces** executable file. Mac OS X distributions come in a form of Mac OS X folder, open it and deploy to a location of your choice.

3.2 Modes of operation



Figure 3.1: Client Mode




Figure 3.2: Server Mode

LogFaces Client can work in two modes - **Client Mode** or **Server Mode**. You select the mode during application start-up. In Client Mode the application connects to and works with one logFaces server instance. In Server Mode the application acts as actual log server but without database.

In **Client Mode** we specify logFaces server **host** name, **port** number and SSL option. These options will be remembered for the next time you run the client, but you can also indicate to use those settings automatically in the next time and not asking again. It is possible to modify these settings in the File/Preferences menu later. The communication between logFaces Client and Server is one way (from client to server) and is HTTP(s) based, so normally there shouldn't be a problem with firewalls. Of course, the access through the given port should be allowed by your network administrator.

In **Server Mode** the application runs with embedded **compact** version of logFaces Server. This is a

limited version of server and client combined into single application which provides **only real time viewing** of log data. You can use it when you don't need database and other features available in standalone logFaces Server. In order to run in **Server Mode** we need to specify at least one of the ports which will be used by the application to receive log data from appenders. You can specify either TCP, or UDP or both, just make sure those ports are available and your application appenders are configured accordingly.

Both modes look very much alike from user experience point of view, except that in Server Mode there are no database and querying features.

*Note that in order to run in **Server Mode** you need to install the license on the computer where you run it. Note also that this is not the same license type as installed on the logFaces Server. This license needs to be purchased and installed separately unless you hold OEM or Site license.*

3.3 Workspace

When using several server nodes, you often need to switch quickly from one system to the other. This is done by means of Workspace which stores connection parameters, real-time perspectives, queries, counters, and all other settings. Workspaces can be exported to and imported from a text files. Additionally, workspaces can be stored on server and shared amongst team members, this way you do the setup only once and let others to import workspace with all the settings. File menu contains all workspace related actions :




Figure 3.3: Workspace actions

3.4 User interface concept

The screenshot displays the logFaces client user interface. The main window is titled "logFaces" and contains several panels:

- Top Panel:** Includes a menu bar (File, Tools, Window, Help) and a toolbar with options like "Pause", "Modify", "Columns", "Pack", "Lock", "Clear", "To log file", "Text edit", and "Find...".
- Left Panel:** Contains "My Perspectives" and "My Queries" sections. "My Perspectives" lists various tracing and problem categories like "App-1 tracing", "Data layer", and "Global problems". "My Queries" shows a list of query filters such as "ALL", "Administration traces", and "Data layer issues".
- Center Panel:** A large table displaying log entries with columns for Time, Host, Domain, SESSION_ID, DEVICE_ID, and Message. The messages include events like "Null returned for the key 51", "New list length 2", and "File deleted: file:///C:/lighform/groups/interface/xxcopy/t".
- Right Panel:** Features "Domain filter" and "Package filter" input fields, and a tree view showing a hierarchy of application components like "App-2" with sub-items like "com.lighformr", "Monitorec", and "DeviceAg".
- Bottom Left Panel:** Shows "Distributions" settings, including "Informative levels", "All applications", "All hosts", and a "30 min sampling rate". It also includes a calendar for "August 2012" and a histogram titled "Log events count" showing event frequency over time.
- Bottom Center Panel:** A "Message" view showing a detailed log entry: "Null returned for the key 51". Below it is a "Counters" view with a bar chart showing metrics like "Potential problems", "Thrown exceptions", "Errors", "Warnings", "User access", "Database hits", "Support calls", and "Access denied".
- Bottom Right Panel:** An "Event details" view showing a list of attributes: Time, Sequence #, Severity, Domain, Host, Logger, Thread, Class, Method, and File. A tooltip is visible over the "File" attribute, displaying an exception message: "Exception in App-4 [DIMA2-XP] com.moonlit.logfaces.test.emitter.LogEmitter5 java.lang.Exception this is ERROR event (607)".

Figure 3.4: logFaces client user interface

logFaces is designed with two major goals in mind – real-time **monitoring** and **drilling** into the log history. The user interface you see above is built for achieving those tasks. The client is an Eclipse RCP application and retains most of the Eclipse paradigms – there is a working area in a middle and it's surrounded by auxiliary views which drive the work flow.

Real time data and query results are displayed in a middle part within data tables and each piece of information has its own tab. You can move those tabs around to have several views displayed together.

Along with the data, there are structural pieces of information like names of loggers, applications, hosts, exceptions, event details – they occupy auxiliary views around the main area. Auxiliary views also can be hidden, shown or re-arranged as you see fit.

The following sections explain how to do the monitoring and drilling as well as the usage of auxiliary views.

3.5 Data tables

Log data is displayed in tables - flexible to manipulate and easy to navigate with. Once you have data (real time or historical) you will want to customize the way it looks or navigate from it somewhere else. Table headers are special – they allow additional filters on existing table content. For example, if you want to focus on particular Thread, click on the corresponding header and select that thread. Several column filters can be combined. Tables can also be displayed with specified columns so that you only see relevant details. The content can be searched and annotated.




Figure 3.5: Data tables

Note the colored bars on the sides – **problematic** events are color-mapped on right side bar while **tagged** events are color-mapped on the left side bar. Hovering the mouse over those bars gives a quick jump to a particular event without too much scrolling around.

The content of data tables can be instantly converted into textual log files – click on “To log file” button or even open the text in an editor by clicking on “Text editor”. The text layout can be specified in preferences.

3.6 Monitoring tasks

Real time monitoring is a unique feature in logFaces as it was designed for incredible volumes of data going through the server. To serve many users, logFaces server pushes logs to each client separately and according to its very own criteria filter. Those filters you will find throughout the system - including queries, administration tasks like reports, trigger and database. This is how real time perspective is defined:

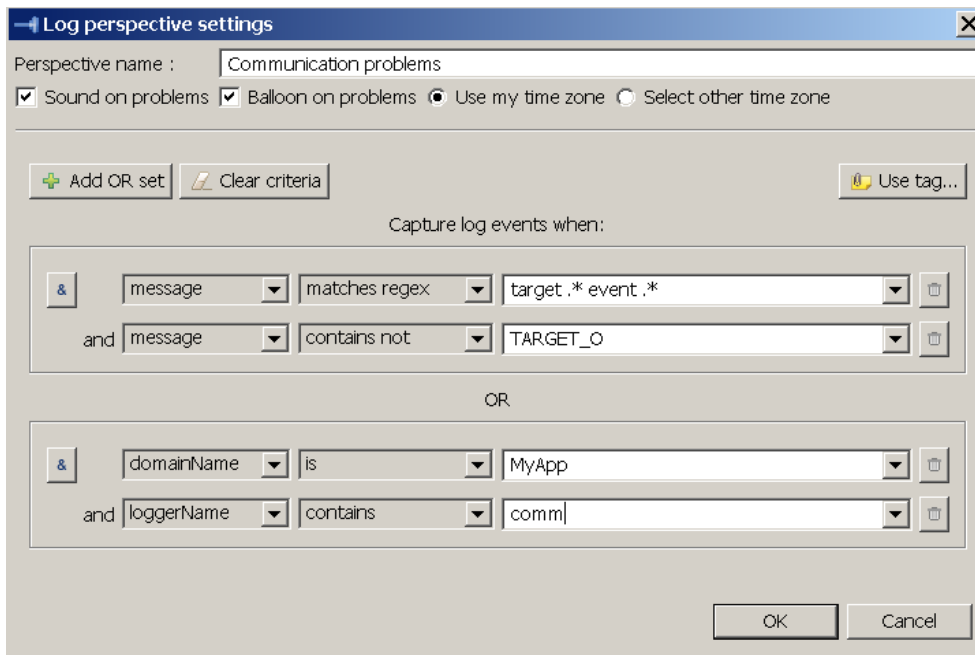


Figure 3.6: Perspective settings

Along with criteria filters there are other options – sound, balloon pop-ups and time zone presentation. Note that perspectives are named and can be stored for later use in “My perspectives” axillary view. Perspectives can be created from several contexts:

- From “My perspectives” view where you will specify the criteria yourself
- From most other axillary views where perspectives can be launched instantly using the selected context. For example, right clicking on a host name within “Hosts” view you can create a perspective which will listen for selected host logs.
- From data tables (query results, or other perspectives) where you can “Follow” some interesting item by right clicking in the table rows.

Perspectives tool bar gives you full control of each perspective separately where you can pause, resume, setup columns, lock, clear or do text searching.

3.7 Data mining tasks

Real-time monitoring alone is not good enough unless accompanied with data mining. You will find numerous ways of getting the historical data – there is an instant fetch, context drilling, range lookups to peer into all sorts of time spans – all made simple and easy to use. Database queries are based on the same idea of criteria filters:

Figure 3.7: Database query

Along with criteria you specify a time range to cover, result set limitation, order display and time zone for presentation. Note that queries can be named – you may want to keep some of the queries you use often in “My Queries” view and spin them out any time. Since queries are based on predefined time ranges, they will always work no matter when you run them. For example query with “1 day” time range will give you data spanning from this moment and back 24 hours.

Queries can be launched from various contexts in data tables and most axillary views. Those queries which you decide to keep in “My Queries” can be instantly launched from main tool bar drop down menu.

3.8 Analytical tasks

Few clicks away you will find convenient tools to assist you with basic qualitative or quantitative analysis of log data. Click on main tool bar buttons “**Distributions**” or “**Slices**” to bring up the corresponding view. Combined use of both those views should help you get best understanding of the underlying log data. This is particularly true when dealing with very large volumes of logs.

3.8.1 Distributions (time line)

This view visualizes log statements in a time line and tells you **when** certain events took place, in what quantities and categories. There are three different modes in which time lines can be used – Daily, Query and Context. Switch the mode by selecting corresponding toggle button on the upper right corner of the view.

Daily distributions display particular day of logs per applications, hosts, severity levels and thrown exceptions. Few clicks of a mouse and you will get your data nicely spread in a time line where you can see occurrences of relevant logs statement. The colors will tell you the full story. Clicking on any of the bars will instantly fetch the actual logs into a separate data table.




Figure 3.8: Daily distributions

Query distributions work in a similar way but visualize particular query results in time so that you see what happened when. Selecting one of your prepared queries will draw the chart instead of delivering the actual logs. For example, on a picture below we can say that most of so-called 'disconnections' took place at around three PM. And then by clicking on an individual bar you will get actual logs contributed to this chart.




Figure 3.9: Query distributions




Figure 3.10: Query distributions

Context distributions, unlike the other two, work strictly in the context of currently active data table and will graphically visualize textual content of this table. The example below illustrates distribution of tags where bars are colored by the color of the corresponding tag counted. You may pick any of the available categories – levels, hosts, app names, exceptions, tags, threads and event MDC items.




Figure 3.11: Context distributions by TAGS




Figure 3.12: Context distribution by COMPONENT

You can also modify the rate of the distributions to make bars wider or narrower in time. Clicking on each bar will filter out actual log statement according to the category used and time point selected.

3.8.2 Slices

Unlike distributions which show time line of logs, **Slices** displays categorized quantities and their correlations within the category contributors. Slices view is also linked to currently displayed data table and will automatically update when data table gets changed or switched. Below are several examples:

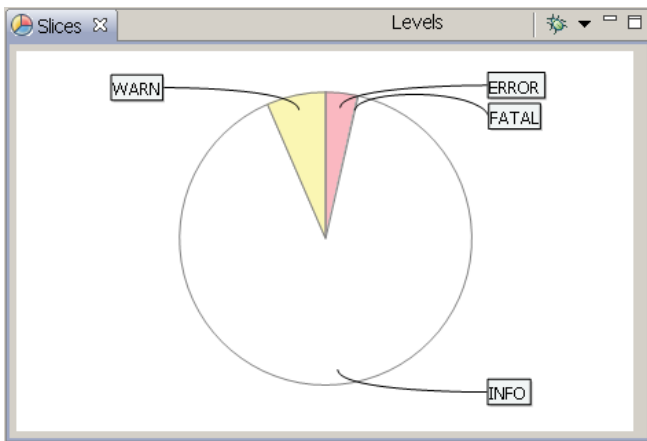


Figure 3.14: Sliced by severity levels




Figure 3.13: Sliced by tags




Figure 3.15: Sliced by components

3.8.3 Counters

Counters is another form of looking at logs – by counting certain events in correlation to each other. Counting is done in **real time** and can be setup separately per client and reset at any time. From server point of view Counters are yet another log perspectives but they displayed differently:




Figure 3.16: Counters view

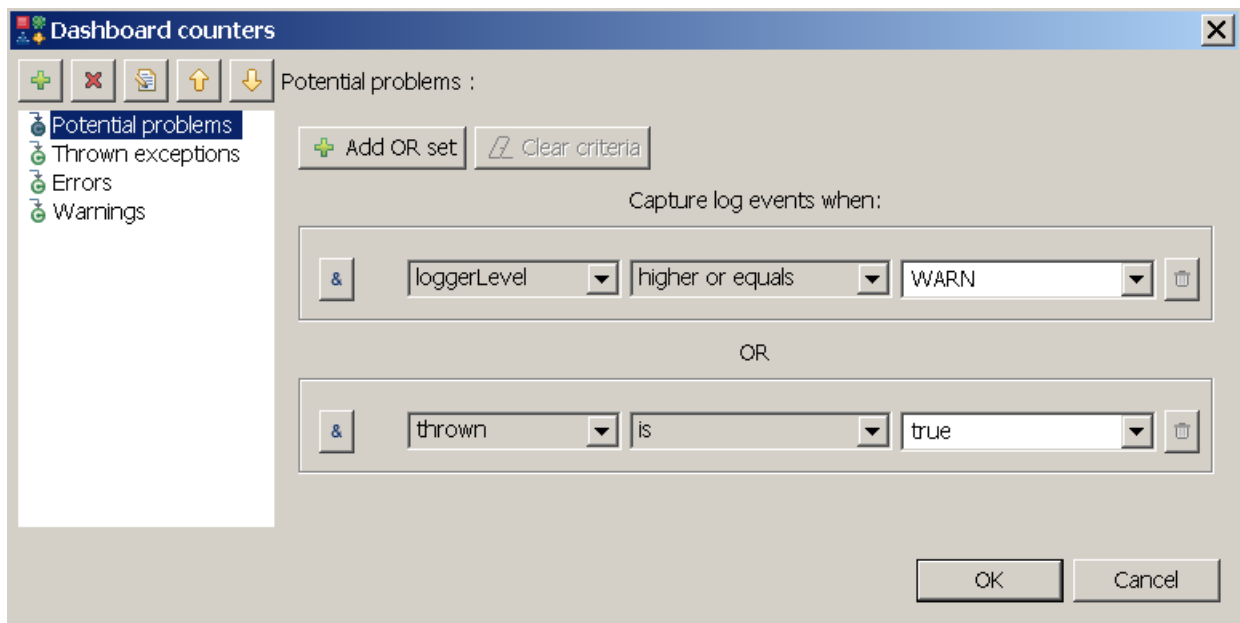


Figure 3.17: Counter setup

Each counter has its own criteria filter and name. Whenever criteria is met the counter gets incremented and its bar grows horizontally. Clicking on the bars will fetch the contributing log statements into ordinary data table where you can further work with the logs.

3.9 Repositories

Domains, hosts and components are jump starts into real-time monitoring and data mining tasks. Those views provide structural information about your log data and underlying system. These are actual names of applications, hosts, packages and classes. This information is accumulated by server continuously and never gets removed. Use refresh button to update the content of the views, they aren't updated automatically.

Use right click context menu to spin off real time perspectives or queries. Selected items will be used to construct criteria filter automatically.

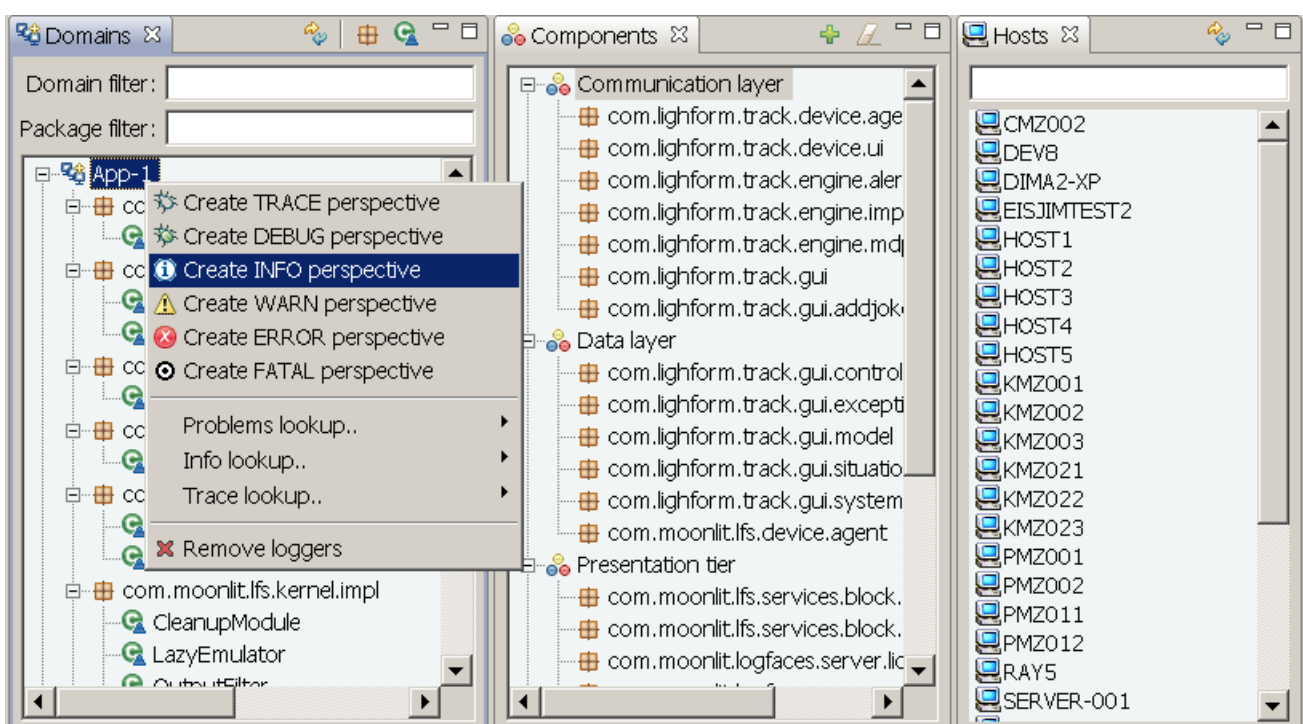


Figure 3.18: Repositories

One note about **Components**. This view is special because components are artifacts, server doesn't have a knowledge of them. Component is a collection of packages grouped together to allow you additional abstraction level of the underlying system. For example, you can make a “Presentation layer” component holding dozens of packages and since components are named, you will be able to use those names in filtering and queries.

3.10 Messages, exceptions and details

Whenever something is selected in data tables, the additional details are also shown in separate views for convenience. Every log event has a message text; when it's too large or spans several lines the “**Message**” view will show it nicely so that you could select and copy text into a clipboard.

When event carries an exception, the “**Exception**” view will be displayed with full stack trace where you further jump into sources.

“**Event Details**” view has a summary of all information known about the log statement. Everything is colored by event severity so that you will easily know what it's all about. For example, below we have an error selected in data table, so the “Message”, “Exception” and “Details” views are painted with the red background.




Figure 3.19: Using auxiliary views

You can also spin queries from selections and even real-time perspectives to follow particular patterns – use the mouse right click wherever possible.

3.11 Source code correlation

Having appenders to transmit location information allows instant access to the origin of the log statements. Most appenders have a property named “**locationInfo**”; when set to “true”, the appender transmits file name, class name, method name and a line number of originating log statement.




Figure 3.20: Accessing the source code

In order to be able to use this information we need to map the locations of source files – this can be done in File/Preferences/Source mapping. Whenever the source files are requested, the specified locations are scanned to find the corresponding source file. If source file is found, it will be displayed in separate window pointing to exact location in the code where log statement is coming from.

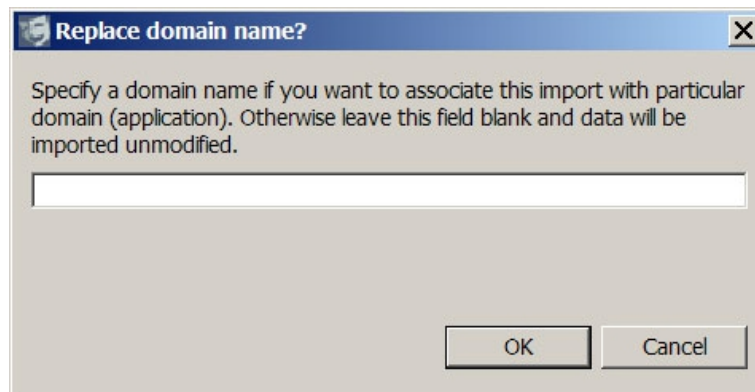
Note that you can also open sources in external editor instead of built in viewer like in the example above. This can be configured in File/Preferences/Files section.

3.12 Import and Export

Log data can be exported from one server and imported into the other, we use binary data image for those operations. You can export query results directly into binary file and import it later on somewhere else.

To export all data from the database, go to Tools menu and select "**Export data from database**". You will be prompted for the file name. This operation is identical to what is done by Backup utility.

To import data into database, go to Tools menu and select "**Import data into database**". You will be prompted to select a file exported previously from logFaces. Before the import actually starts, you will be asked to specify a Domain name to which imported data should be associated:



If you leave this field blank, the data will be imported without any modifications. Otherwise, the Domain name of all events will be replaced to the value provided. This is very convenient when you exchange data between different logFaces servers.

Please take into account, that depending on the dataset size and the speed of the database, those operations may take up to several minutes while taking considerable amount of CPU and database I/O.

3.13 Command Line Arguments

3.13.1 logfaces.ini file

Client installation directory should contain a file named **logfaces.ini**. If it doesn't exist – make it manually (it should be in the same folder as **logfaces.exe**). The following parameters define JVM heap sizes, make sure those lines are present as shown :

```
-vmargs
-Xms64m
-Xmx256m
-XX:MaxPermSize=128m
```

3.13.2 Launching queries on start-up

It is possible to launch the client so that it automatically runs a query before showing up. To do this, use **-query** argument which is followed by **semi-column** separated query parameters. None of the parameters is mandatory (see defaults below). For example, the following command line will display first 100 exceptions with ERROR and above severity level :

logfaces.exe -query name=Problems;thrown=true;loggerLevel=error

Here is the full list of parameters which can be used:

Parameter	Description	Default
<code>name</code>	Query name, will be displayed on the editor tab	My query
<code>fromTime</code>	Include data from this time (UTC numeric long in msec)	0
<code>toTime</code>	Include data until this time (UTC numeric long in msec)	MAX_LONG
<code>domains</code>	Comma separated list of application names to match	-
<code>hosts</code>	Comma separated list of host names to match	-
<code>loggers</code>	Comma separated list of logger names to match	-
<code>exception</code>	Text to match in stack traces	-
<code>matchMessage</code>	Text to match in event message	-
<code>onlyThrown</code>	True for including only thrown exceptions	false
<code>level</code>	Log4j level, numeric or strings (INFO, ERROR, etc)	TRACE
<code>size</code>	Maximum numbers of events to fetch	100
<code>timeZone</code>	Which time zone to use for display	JVM time zone

Table 3.1: Command line parameters

3.14 Preferences

Global preferences are accessible through the File menu:

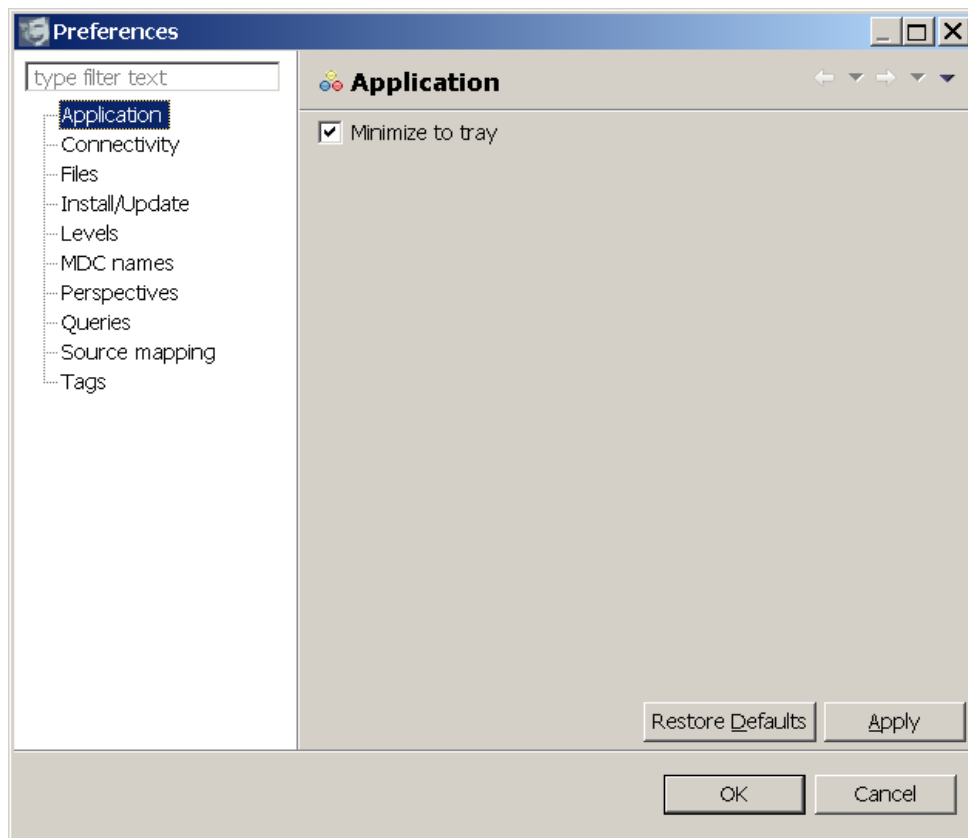


Figure 3.21: Application preferences

In **Application** section you can specify whether the client main window should be minimized to tray or stay minimized in the task bar so that you see the title. When application is minimized to tray, double clicking on its icon will restore its main window.

In **Connectivity** section you can specify the server connection endpoints which apply to a currently used workspace.

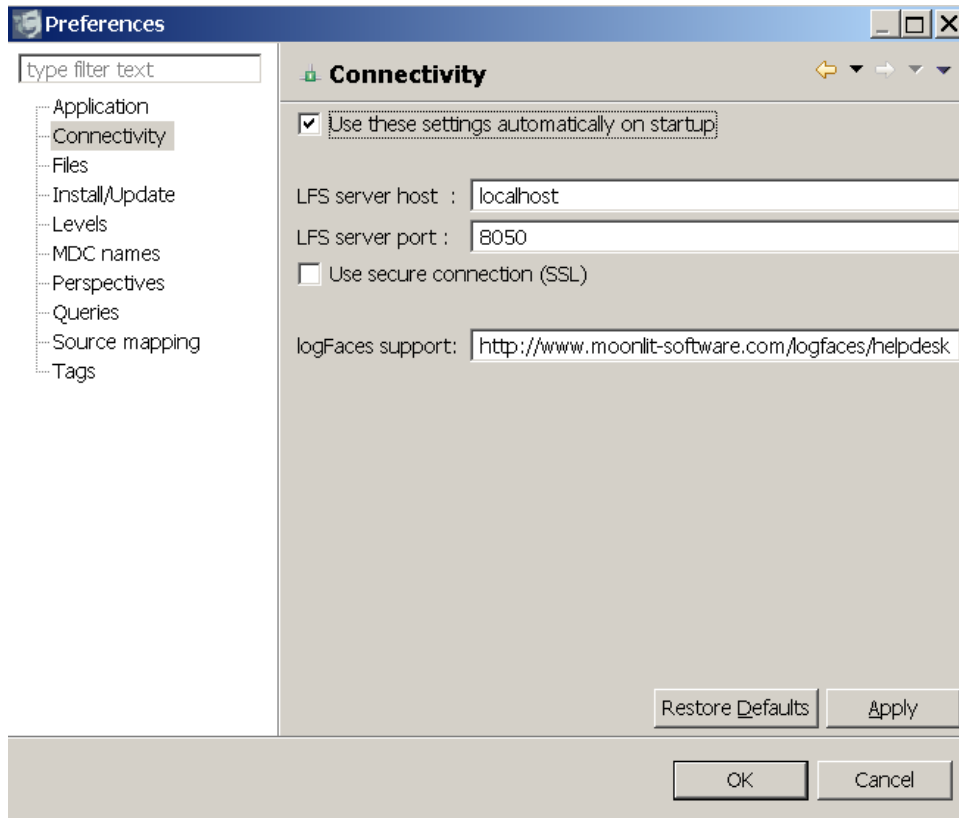


Figure 3-3.22 Connectivity Preferences

Support site URL is used for submitting bugs, feature requests or questions – you will be taken to this URL automatically when selecting **Help/Report bug or request feature** menu.

In **Files** section you specify the layout format of log files and external editor to be used throughout the application. The layout format will be applied to all file related operations when saving data to a file or copying logs into clipboard. The format is specified in Apache Log4j documentation, you might want to [read about it here](#).

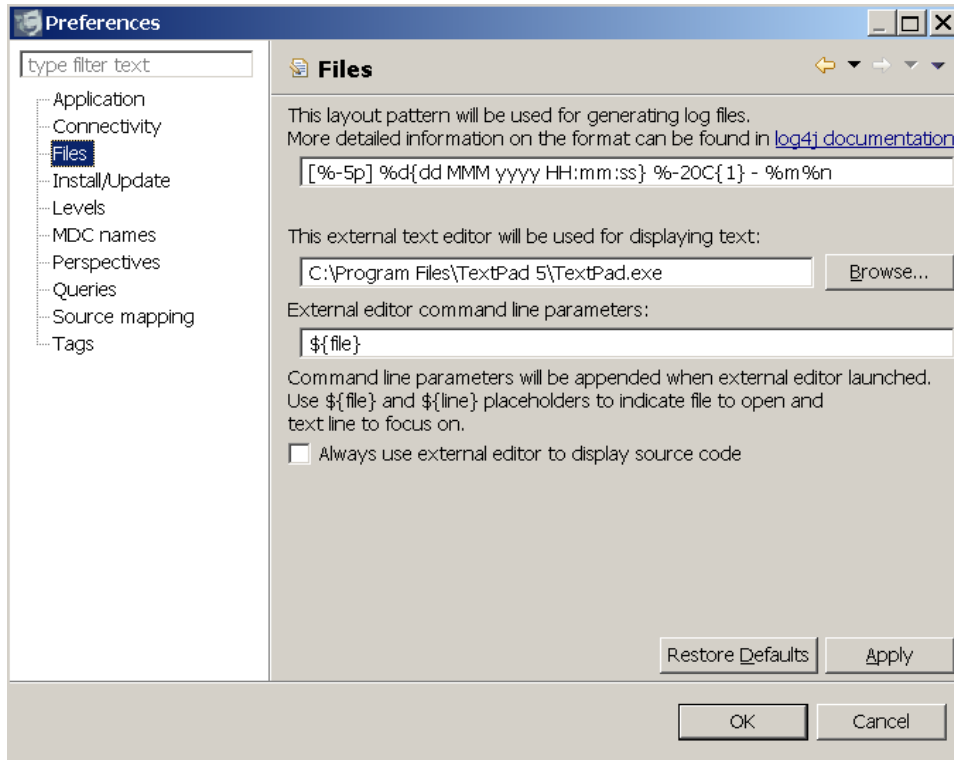


Figure 3-3.23 Files Preferences

Note that external text editor can also be used for displaying the source files. By default sources are displayed in internal source viewer, but you can change this and have your own editor opened whenever sources are displayed. You can also use `${file}` and `${line}` variables to construct a command line for your editor.

In **Install/Update** section you specify how to obtain software updates. There are several options to specify the update policy, for example, you can request to check for the updates every time the application is run and notify when they're available for installation:

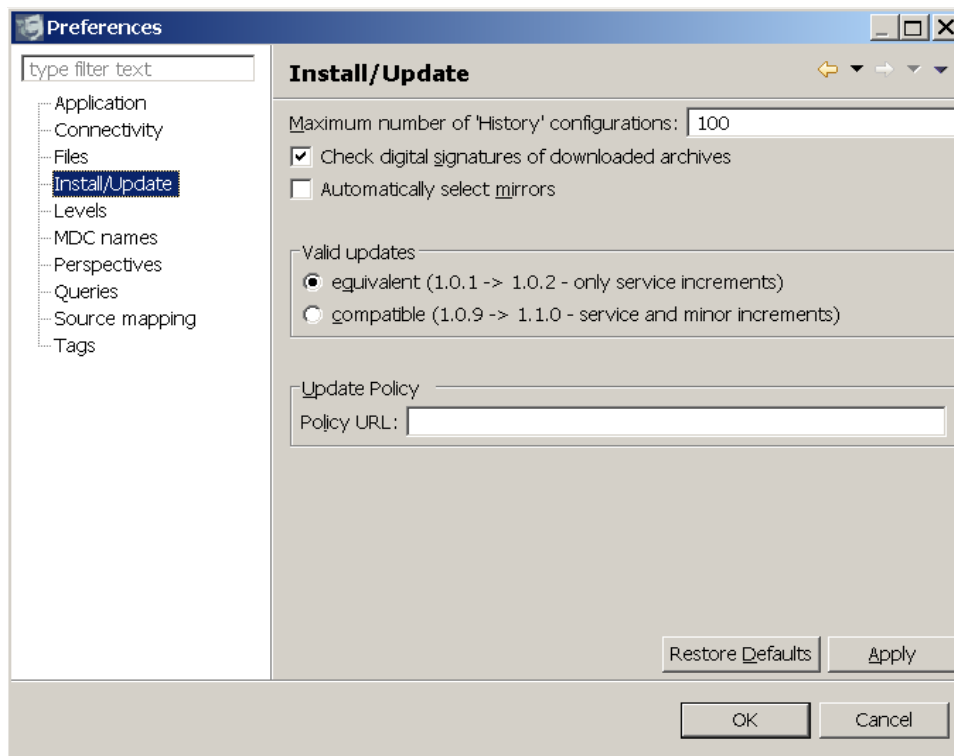


Figure 3-3.24 Update Preferences

You may as well disable the automatic updates and do it manually from the **Help menu** on the main menu bar. When new version of the software will be released, the application will display notification dialog and ask your permission to install the updates. Normally this will require consequent restart of the application.

In **Levels** section you can customize the way log statements look – background or foreground colors and image icons. Note that some axillary views are using the colors you define here.

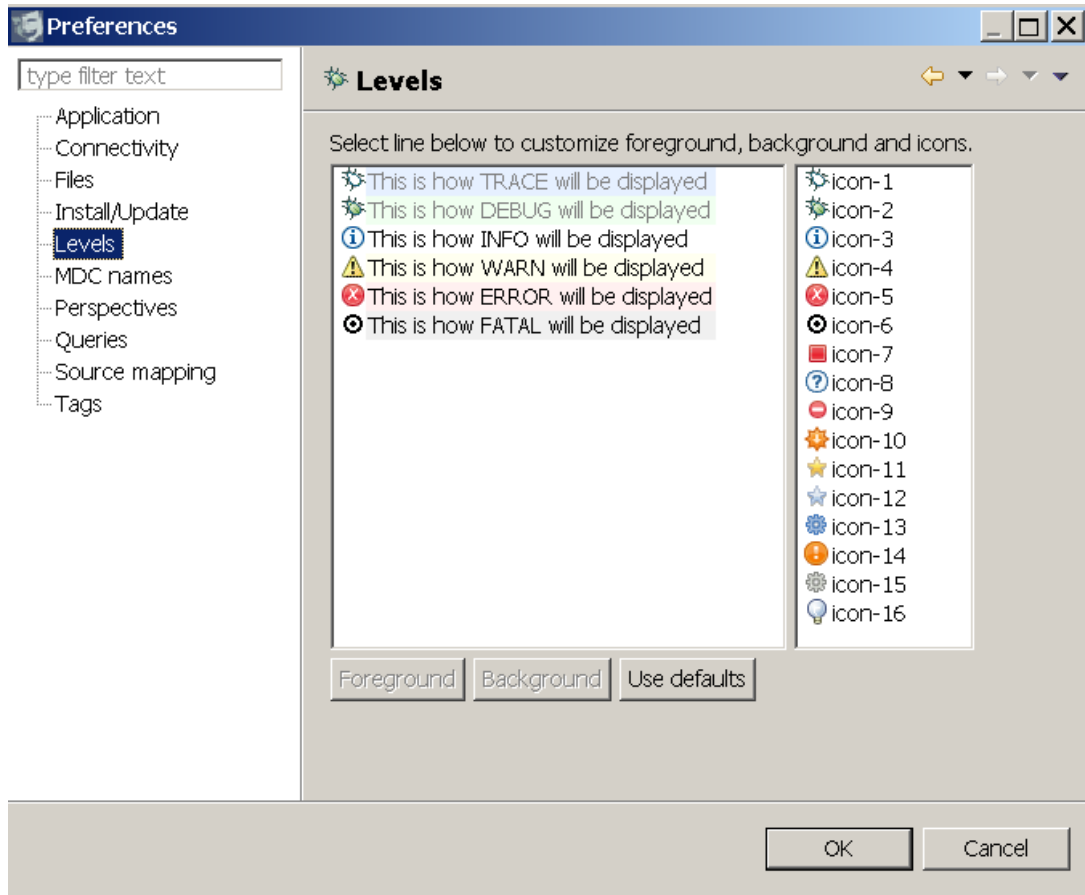


Figure 3.25: Levels

In **MDC names** section you will find current MDC names mapped from applications to the logFaces server. MDC mapping is modified on server side, so in case you need to change the mapping, you will have to use server administration.

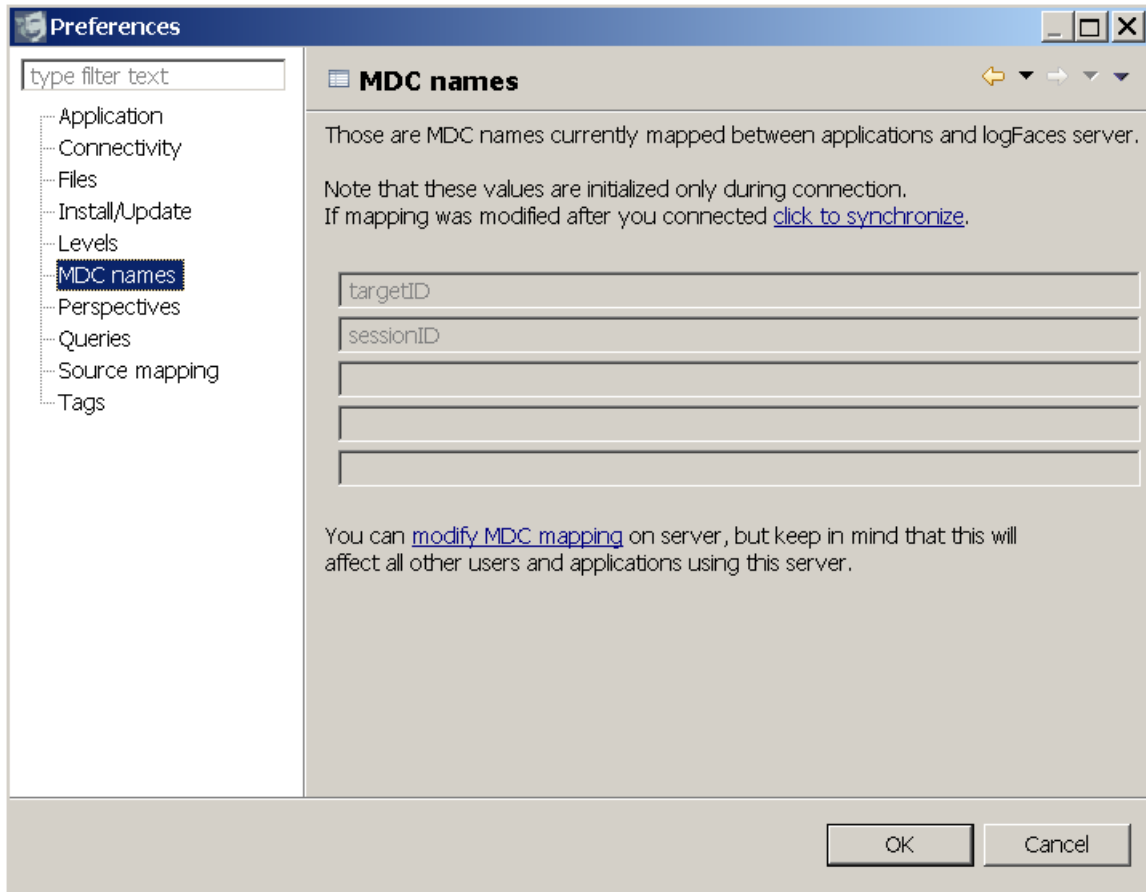


Figure 3.26: MDC names mapping

When MDC mapping has some names, you will be able to use those names in queries and filters. For example, when SESSION_ID name is defined, it will appear in data tables column and you will be able to search events based on this name as well as trace real time data and filter data tables containing particular value of this attribute.

In **Perspectives** section there are ways to customize real-time perspective views. When real-time view is opened, there is an option to grab historical data from database before the view gets live – this will give you the context for current log events coming into the view.

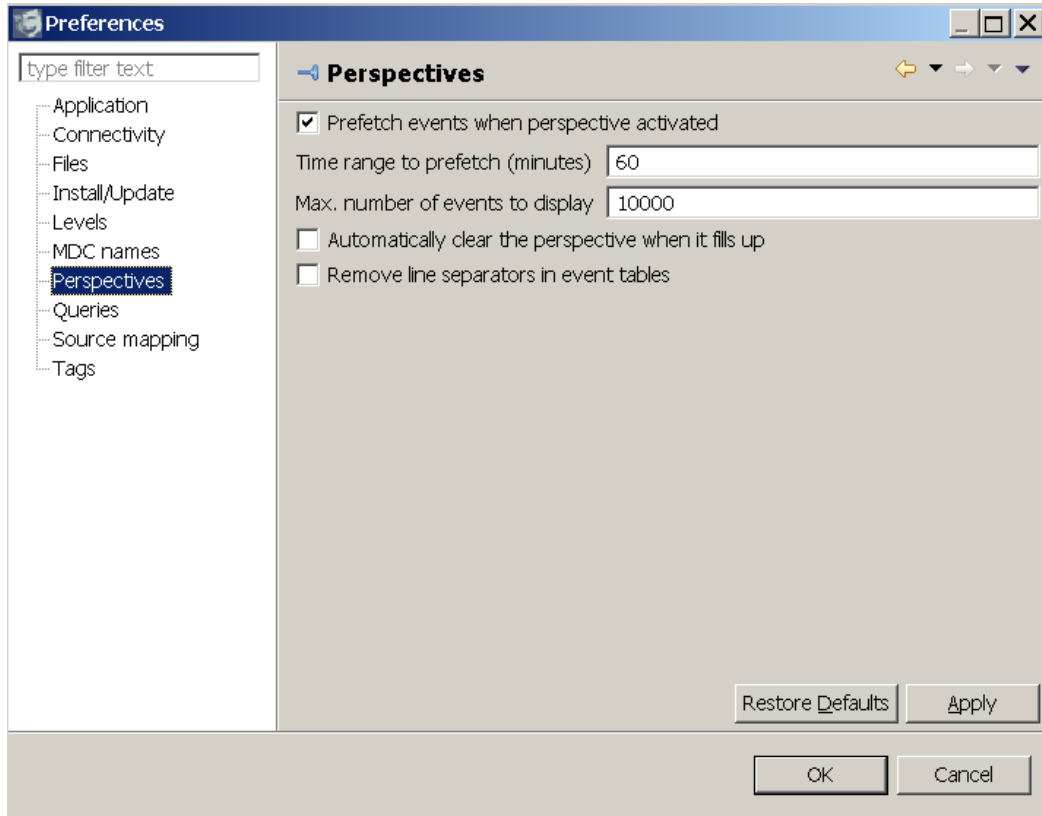


Figure 3.27: Perspectives

You can limit the size of the view to a certain amount of events, when view gets full it starts rotating by removing older events. You may want to clear up the view entirely when it fills up. Some messages may have End Of Line separators – this can be optionally removed to fit the message nicely into the table.

In **Queries** section you can specify the way queries are executed against the server. These are global settings and applied to all queries throughout the client.

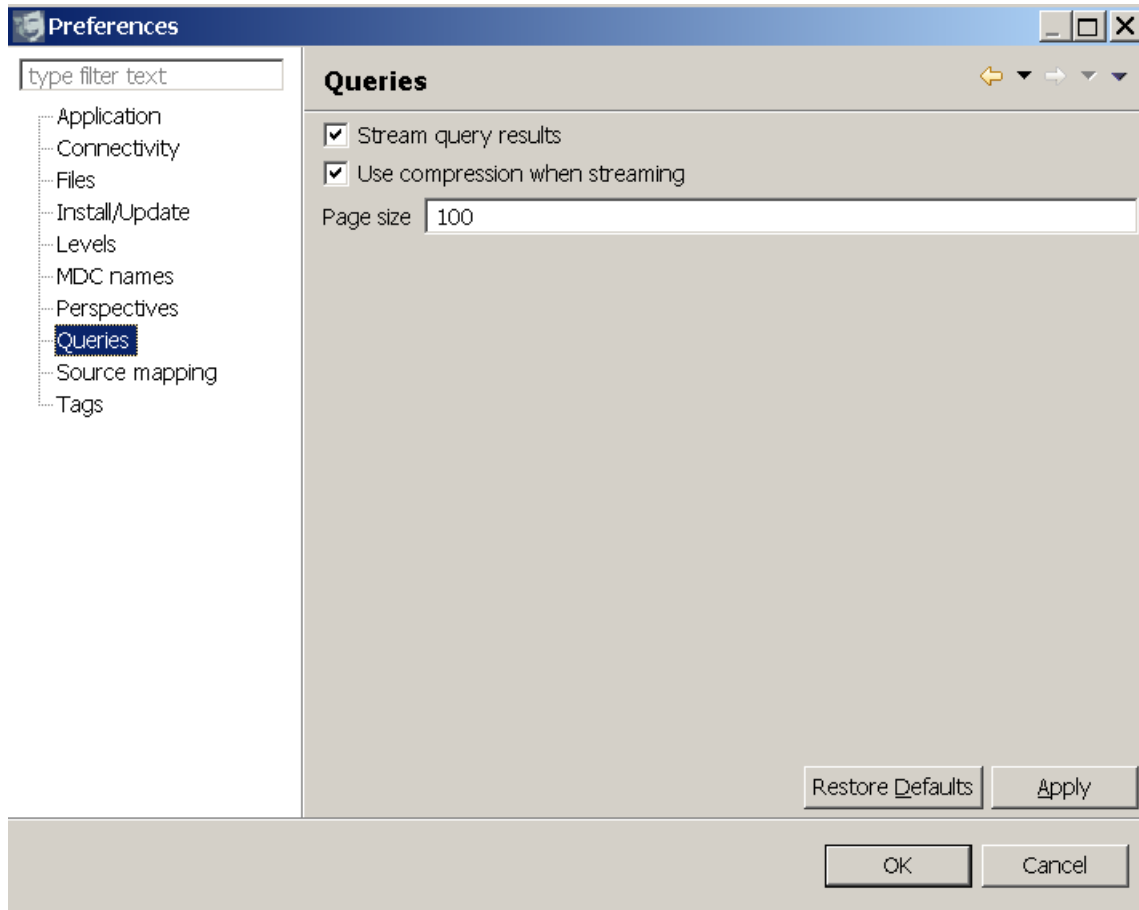


Figure 3.28: Queries

If you ever used logFaces over the public internet or VPN, you might have noticed that large result sets are slow to get to the client even with a well tuned database. It may get particularly painful with huge queries containing exception dumps, which is often what people do. This issue is addressed by two options you can try here - “Stream Query Results” and “Use Compression”.

The slower the client-server connection, the more vivid will be this improvement. Now, instead of paging the result sets and hitting the server on every page, we do streaming which takes a single server hit and delivers everything to the client in one shot. The stream may get compressed. Plus it will bypass many conversions while streaming simple JSON text.

In **Source Mapping** section you can specify where to look for actual source files. References to source files (if enabled in your appenders) will be available for every log event as well as exception stack traces. So, in order to jump into a source code directly from logFaces client, this mapping is essential.

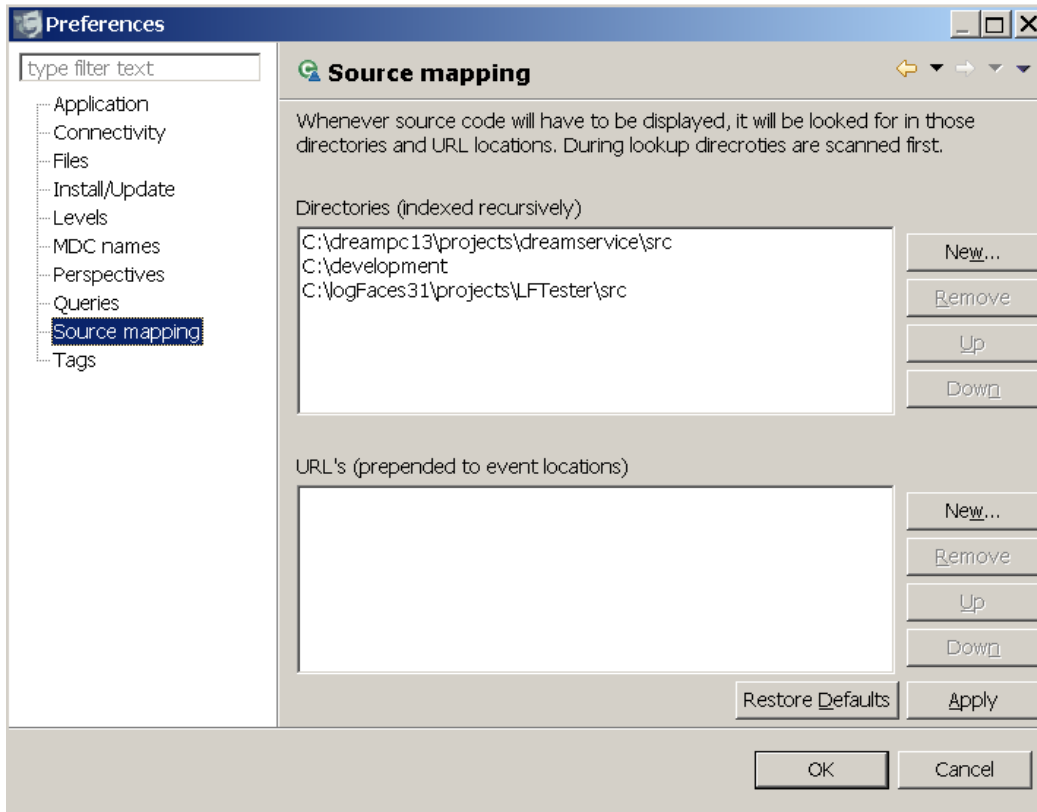


Figure 3.29: Source code mapping

When source file is being resolved, client first tries to locate the file in directories in the order you specify. If not found, the client will try to locate the file in one of the URL's. The URL must be HTTP based and contain the base for file locations, client will append relative file path to the URL during look up.

Tags add domain specific indications to a dull technical log stream. You create a tag by giving it a short friendly name, color and matching criteria. Logs matching the tag criteria get tagged before they get displayed. The criteria is the same bunch of rules we use in queries and filters, nothing new here. Tags are a very powerful tool - they participate in view filters, queries and analytical charts.

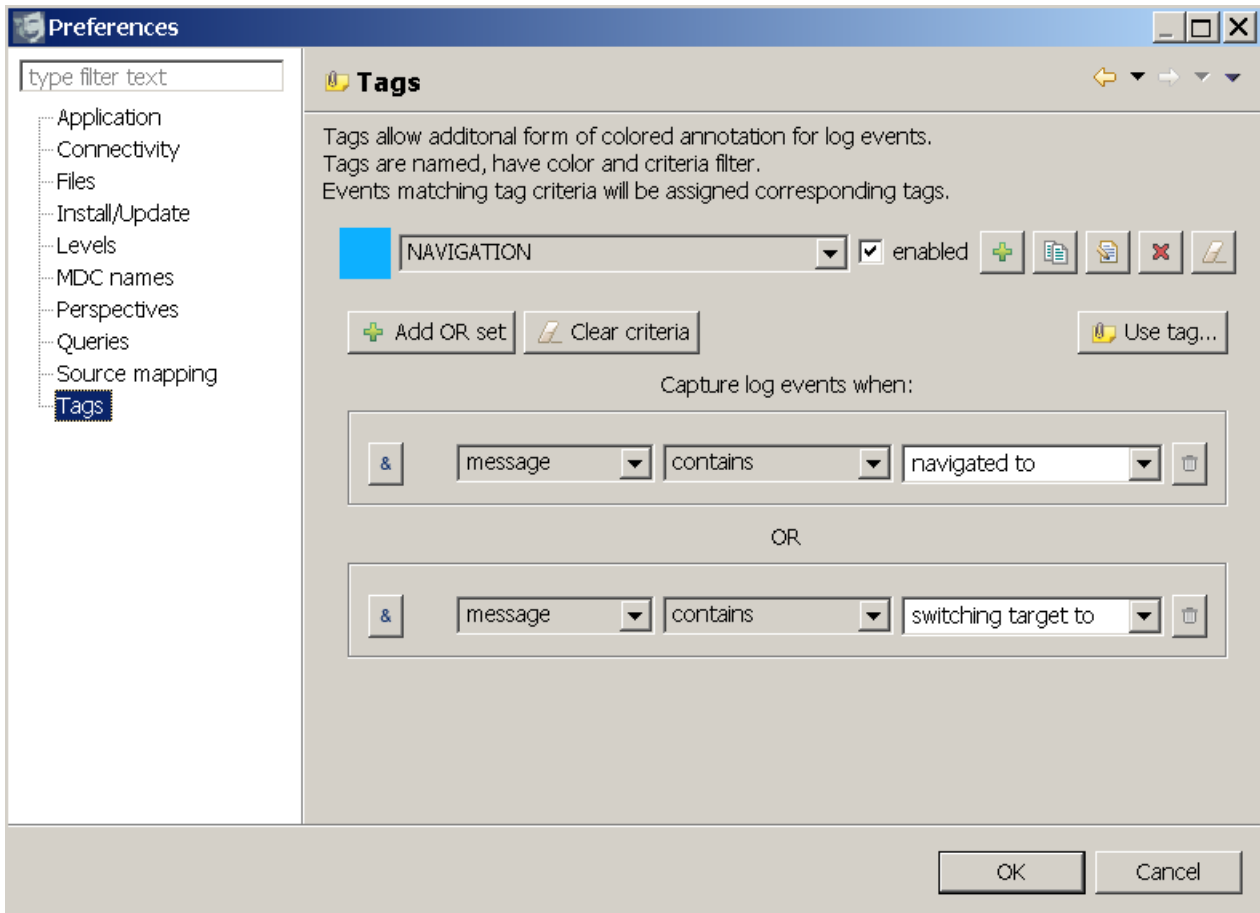


Figure 3.30: Tags

3.15 Status bar

Status bar displays the following information:

- current connection state
- current connection end point
- current version of logFaces server
- current license
- number of database records (click on the icon to refresh the counter)
- current RAM used on client versus maximum RAM allocated plus manual garbage collector.



Figure 3-3.31 Status Bar