

FraunhoferFS

Installation and Setup Guide




Table of Contents

1 Introduction.....	3
1.1 Important Notes	3
2 GUI-based Installation and Management	4
2.1 Important Notes	4
2.2 Get FhGFS-Setup.....	4
2.3 Starting Setup.....	5
2.4 Install FhGFS.....	5
2.4.1 Configuration	5
2.4.1.1 Define Roles	6
2.4.1.2 Create Configuration	6
2.4.1.3 Configure Infiniband	7
2.4.2 Download Release	7
2.4.3 Install	8
2.5 Start / Stop FhGFS Services	8
3 Manual Installation.....	9
3.1 Download and Installation of Packages.....	9
3.2 Building Kernel Modules and OpenTk Library	10
3.3 Basic Configuration and first Startup	11
3.4 Additional Notes	11
4 Support	12
5 Licensing.....	12

1 Introduction

There are two ways to install FhGFS: GUI-based or manually. In general, the GUI-based installation is recommended, but it does not provide the full flexibility of a manual installation (e.g. some configuration settings are not available and installation into an image is currently not supported by the GUI-based installation process). Chapter 2 of this document gives an overview of the GUI-based installation and describes the installation process in more detail. Chapter 3 describes the manual installation process.

1.1 Important Notes

- **Runtime-Compatibility of different Versions:** Different versions of FhGFS clients and servers are only compatible if they are part of the same release set (i.e. contained in the same release folder at <http://fhgfs.com/release>).
- **Storage Format Compatibility:** Some versions of FhGFS may introduce a new storage format, which might be incompatible with previous versions. Therefore, we strongly advise you to backup all your data before you upgrade to a new major release of FhGFS. (See release changelog for compatibility notes).
- **Compiler:** Please note that you must have a GNU C compiler (`gcc`) installed on every client node to build the FhGFS client kernel modules. If you enable Infiniband support, you will also need a GNU C++ compiler (`g++`) on the corresponding FhGFS nodes to build the FhGFS OpenTk library.
- **Required Infiniband Libraries:** To compile FhGFS with native Infiniband support, you need to have the `ibverbs` and `rdmacm` libraries and development packages of OFED 1.2 or higher installed.
- **FhGFS Installation Paths:** The FhGFS binaries and libraries will be installed to `/opt/fhgfs`. The configuration files are located in the directory `/etc/fhgfs`. Each service (including the client) installs a corresponding init script to `/etc/init.d`.

2 GUI-based Installation and Management

2.1 Important Notes

- **SSH Configuration:** As setup (the FhGFS Admon daemon) is using SSH to connect to the FhGFS nodes, it is essential that you enable password-less login using a public/private key pair.

2.2 Get FhGFS-Setup

The first step to install FhGFS is to download the FhGFS installation package from <http://www.fhgfs.com> and install it on a management host. There are several setup packages available, depending on your Linux distribution and architecture. Download the appropriate one (see Table 1), according to your management host configuration.

Linux Distribution	Architecture	Package
Suse Linux (OpenSuse, SLED, SLES)	x86	fhgfs-suse-2009.08-r1.i586.rpm
	x86_64	fhgfs-suse-2009.08-r1.x86_64.rpm
Red Hat Linux (including derivatives such as CentOS and Fedora)	x86	fhgfs-rh-2009.08-r1.i586.rpm
	x86_64	fhgfs-rh-2009.08-r1.x86_64.rpm
Debian GNU Linux (including derivatives such as Ubuntu)	x86	fhgfs_2009.08-1_i586.deb
	x86_64	fhgfs_2009.08-1_x86_64.deb

Table 1: Available installer packages

After downloading, you can install the package using your distribution's package manager (`rpm` for Suse and Red Hat, `dpkg` for Debian). The GUI-based installation will be done using a graphical Java frontend, provided by the "FhGFS Administration and Monitoring Tool" (Admon). The Admon consists of two parts: a daemon process (`fhgfs_admon`), which is installed and started automatically by the RPM, and a Java frontend, which can be run on any host with network access to your management host. For more information see the FhGFS Administration and Monitoring documentation.

2.3 Starting Setup

To start setup, download the Java frontend from <http://www.fhgfs.com/download>. This is a jar-file which can be started on Windows and Linux workstations via double-click or command line:

```
# java -jar fhgfs_admon_GUI.jar
```

After startup, you will be prompted for login information. To install FhGFS, you have to log in as "Administrator". The default password is "admin". For further information about user and password management, please refer to the FhGFS Administration and Monitoring documentation.

After starting the application, you need to define the host on which the remote daemon runs, as well as the port on which it listens to connections (default port: 8000) (Figure 1).

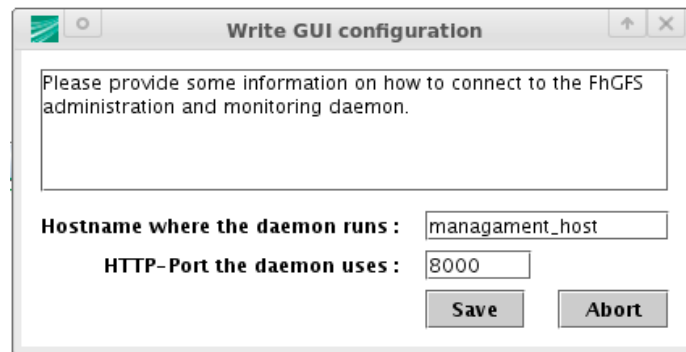


Figure 1: Connection Settings

2.4 Install FhGFS

2.4.1 Configuration

Start the configuration of your FhGFS installation by choosing "Management" → "Installation" → "Configuration" in the menu (Figure 2 shows the menu structure).

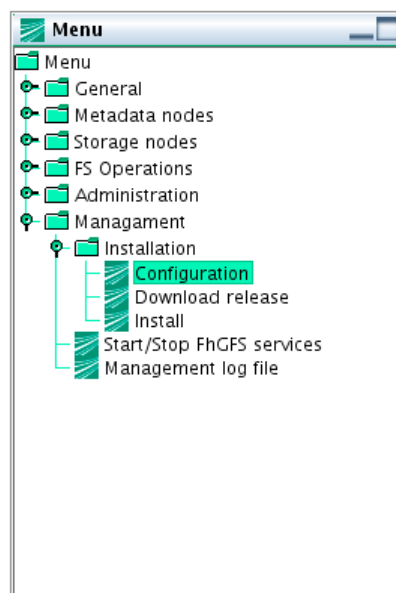


Figure 2: Menu Structure

2.4.1.1 Define Roles

First, you have to define the host names of your servers and the tasks they should perform in your file system.

You can define metadata servers, storage servers and clients (Figure 3 shows an example configuration). The location of the management daemon and the administration and monitoring service are fixed to your management host. Please refer to the FhGFS User Guide for more information on the different types of services.

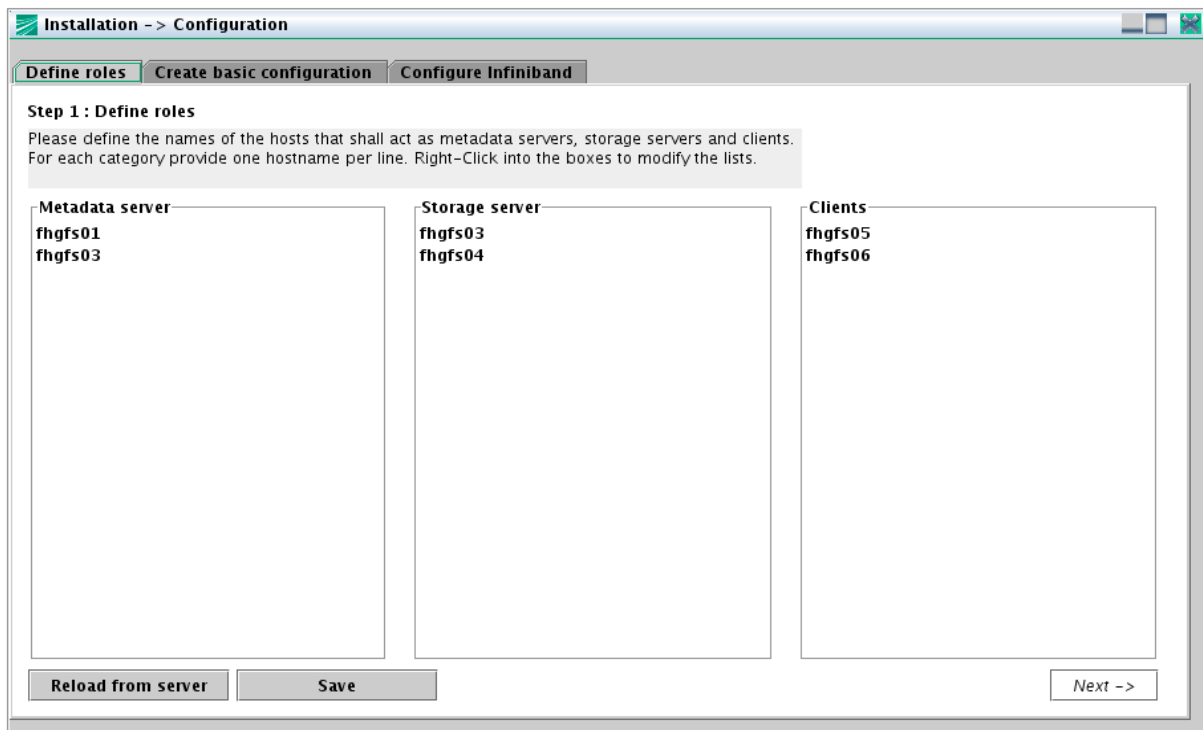


Figure 3: Example role definition

Apply your changes by clicking the "Save"-button. The installation system will automatically perform a check for your used distributions and architectures afterwards (this might take a while and you will be notified when the detection is finished).

2.4.1.2 Create Configuration

The next step is to set some basic configuration options for the FhGFS services. A description of every configuration option can be viewed by hovering the mouse over the information symbol.

The configuration dialog is filled with reasonable default values, so you do not have to change anything if you do not have any special needs.

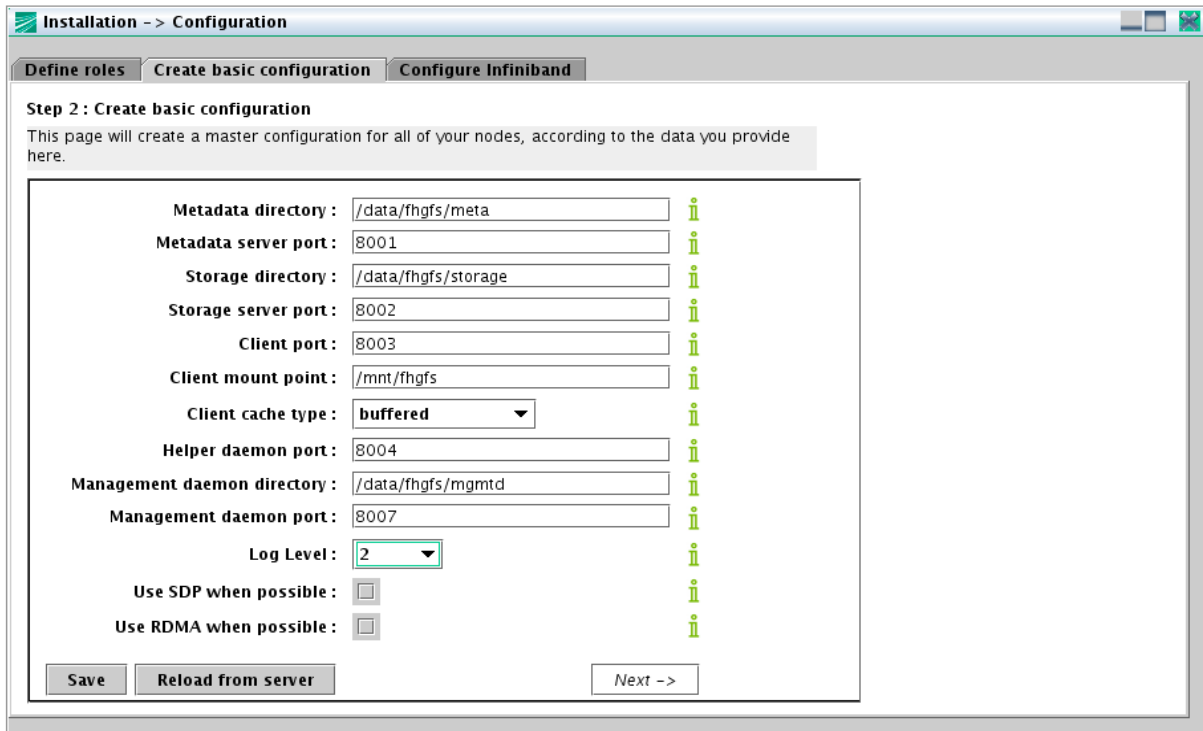


Figure 4: Basic Configuration

2.4.1.3 Configure Infiniband

If you want to enable native Infiniband (ibverbs) support for your nodes, you have to specify some additional information in this step. The checkbox on the left lets you define for each host, whether native Infiniband shall be used or not.

In order to build and link some parts of FhGFS with your installed version of the OFED software stack, you need to provide the paths to your OFED include- and library-files. If these directories are identical on all hosts, it is sufficient to specify them once and click the "Use for all" -button. You may also leave the fields completely empty, in case all include files and libraries are installed to the default locations.

If you need more information on a configuration option, try hovering the mouse over the corresponding element. In most cases, a tooltip will be shown.

2.4.2 Download Release

Before FhGFS can be installed, the latest release packages have to be downloaded. This can be done automatically. Please use "Management" → "Download release" from the menu. If you need to use a proxy server for HTTP connections to the internet, please provide its name and port.

If you do not want to use the current release or if you prefer to do this step manually, please download all packages for the distributions and architectures you use from <http://www.fhgfs.com/release> and put them into the directory `/opt/fhgfs/setup/rpms` on your management host.

2.4.3 Install

Finally, to install the packages on the configured hosts, choose “Management” → “Install” from the menu.

A summary of the data gathered about your hosts will be shown. If the information is correct, you can click the "Install"-button. The application will transfer the corresponding files to the target systems, install them and create a configuration according to your selections. Afterwards, an information dialog is shown which tells you if the installation succeeded. Please be aware that you must accept the FhGFS license agreement to perform an installation.

2.5 Start / Stop FhGFS Services

The menu item “Management” → “Start/Stop FhGFS services” (Figure 5) allows you to start and stop FhGFS services on the configured nodes from within the GUI interface.

After choosing the menu option, an overview on started and stopped services is presented and you are able to either start/stop all services or particular services on individual hosts.

You can also restart the Administration and Monitoring daemon, which is required after changing the configuration (e.g. port settings) or installing FhGFS.

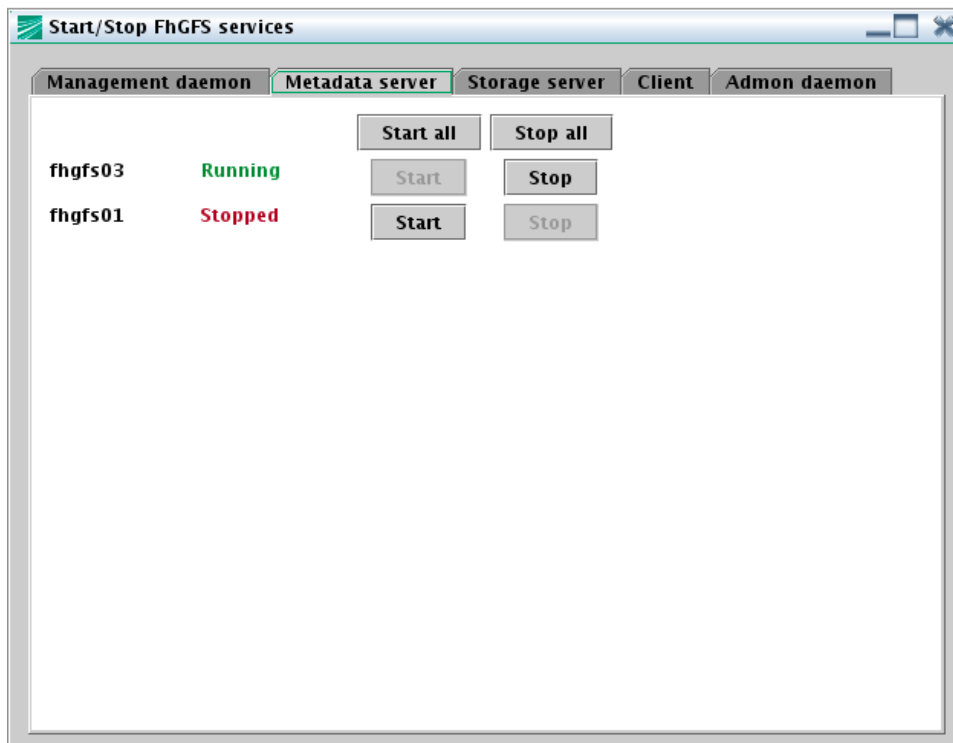


Figure 5: Start/Stop services

The software will log into each of the nodes via SSH to start/stop the corresponding services.

3 Manual Installation

3.1 Download and Installation of Packages

To install FhGFS manually, start by downloading the FhGFS packages for your architecture and distribution. FhGFS supports heterogeneous environments with different architectures and/or distributions.

Table 2 shows the roles that need to be defined to run FhGFS and the corresponding packages for each role. FhGFS has been designed to allow running any combination of services (e.g., metadata and storage server) on the same node. Additional information about the roles of FhGFS nodes can be found in the FhGFS User Guide. The latest stable packages can be downloaded from:

<http://www.fhgfs.com/release/latest/packages>

Roles	Packages
Management Server (<i>one node</i>) <ul style="list-style-type: none">➤ Manages configuration and group membership➤ Hostname or IP address must be known by other nodes at service start time	fhgfs_mgmt fhgfs_opentk_lib
Metadata Server (<i>at least one node</i>) <ul style="list-style-type: none">➤ Stores directory information and allocates file space on storage servers	fhgfs_meta fhgfs_opentk_lib
Storage Server (<i>at least one node</i>) <ul style="list-style-type: none">➤ Stores raw file contents	fhgfs_storage fhgfs_opentk_lib
Client <ul style="list-style-type: none">➤ Kernel module to mount the file system➤ Requires userspace helper daemon for logging and hostname resolution	fhgfs_client fhgfs_helperd fhgfs_opentk_lib
Admon - Administration and Monitoring Server (<i>optional</i>) <ul style="list-style-type: none">➤ Continuous monitoring of servers➤ Live statistics	fhgfs ¹ fhgfs_opentk_lib

Table 2: FhGFS Roles and corresponding Packages

After downloading, install the packages with the package manager of your distribution, e.g. rpm or dpkg. The package contents will be installed to `/opt/fhgfs`. For distributions that are not RPM- or Debian-based, you might consider using a tool like `rpm2cpio` to extract the package contents (but this method will not be further discussed here).

¹ The FhGFS Admon is included in the general FhGFS (setup) package.

After the installation, all services will be set to automatic startup at system boot. If you do not want to run FhGFS services automatically, use `inserv`, `chkconfig` or `update-rc.d` (depending on your distribution) to disable automatic startup.

Each FhGFS service (including the client) comes with a corresponding `fhgfs_...` init script in `/etc/init.d` and a configuration file in `/etc/fhgfs`. Installation does not start any of the services directly, because there are a few configuration options that need to be set before the first startup.

3.2 Building Kernel Modules and OpenTk Library

Most parts of FhGFS are pre-built and ready to use when you install the packages. However, the client kernel module of FhGFS still needs to be compiled for your specific Linux kernel version.

In addition to that, it might be necessary to build the open toolkit library (`fhgfs_opentk_lib`), which contains a part of the FhGFS communication layer. The library comes pre-built and needs to be recompiled only if you have an Infiniband network and want to enable native Infiniband support. (Note that you also might need to enable usage of the native Infiniband protocol in the FhGFS configuration files by setting the value of `connUseRDMA` for the clients, metadata and storage servers.)

Client kernel modules: To build the client kernel modules, become root and change to the module build directory on a client node: `/opt/fhgfs/src/client/fhgfs_client_module/build`

The build directory contains a makefile. Use `“make help”` to see a list of available options. By using make arguments, you can specify a kernel source directory or enable native Infiniband support. Note that it is necessary to have the kernel header files installed to build the kernel module. If you want to build the module for the currently running Linux kernel without native Infiniband support, just type `“make”` without any arguments. Otherwise use `“make FHGFS_OPENTK_IBVERBS=1”` to build with Infiniband support, using default include paths.

After a successful build of the kernel modules, use `“make install”` to copy the kernel modules to their default location (`/opt/fhgfs/client`).

Open toolkit library: To build the FhGFS open toolkit library with native Infiniband support, become root and change to the library build directory: `/opt/fhgfs/src/fhgfs_opentk_lib/build`

The build directory contains a makefile and `“make help”` will show you the available options. The make options allow you to specify an OFED include directory in case you did not install the OFED packages to the standard path. Note that you need to have the `ibverbs` and `rdmacm` libraries and development packages installed. If you want to build the library with Infiniband support using the standard paths, just type `“make FHGFS_OPENTK_IBVERBS=1”`.

After a successful build of the library, use `“make install”` to copy the library to the default location (`/opt/fhgfs/lib/fhgfs`).

3.3 Basic Configuration and first Startup

There are three configuration options that need to be changed for your specific environment:

- **Storage paths**
(Files: `/etc/fhgfs/fhgfs_mgmt.d.conf`, `fhgfs_meta.conf`, `fhgfs_storage.conf`)
- **Hostname of the FhGFS management daemon**
(Files: `/etc/fhgfs/fhgfs_meta.conf`, `fhgfs_storage.conf`, `fhgfs_client.conf`, `fhgfs_admon.conf`)
- **Client mount file**
(File: `/etc/fhgfs/fhgfs_mounts.conf`)

Storage paths: Set the value of the “store...Directory” option to configure where file system should store its internal data. Note that FhGFS stores all data in a subdirectory of an existing formatted partition, so there is no need to format a separate partition for FhGFS.

Management host: Set the value of the “sysMgmtdHost” option in the configuration files of the metadata servers, the storage servers, the clients, and the optional Admon to the IP address or hostname of the management daemon.

Client mount file: The client mount file consists of two space-separated values. The first one is the path to where you want to mount the file system; the second value is the client configuration file for this mount point. You will typically have a line like this in the file:

```
/mnt/fhgfs /etc/fhgfs/fhgfs_client.conf
```

(Note that it is also possible to enter multiple mount/config pairs, one pair per line, in this file to allow a client to access multiple distinct file systems.)

Start FhGFS services: Now, you are ready to start all FhGFS services (in arbitrary order) by using the corresponding init scripts (`/etc/init.d/fhgfs_...`).

3.4 Additional Notes

Multiple network interfaces: If your nodes are equipped with multiple network interfaces and you want to assign specific priorities or exclude some of the interfaces, read the corresponding section (“Miscellaneous” → “Network interfaces”) in the general FhGFS User Guide.

Command-line configuration tool: In addition to the FhGFS Admon GUI, there is a command-line tool (`/opt/fhgfs/sbin/fhgfs_online_cfg`) contained in the client package, which be used to query file system status information and to configure striping settings.

Admon: If you are using the Admon package, you might want to go to the “Management” → “Installation” page and enter the names of your nodes. This will allow you to use the Admon to automatically distribute configuration files and start the build processes for you.

4 Support

FhGFS is currently available free of charge as a binary package. If you have any trouble installing or running the software, please send an e-mail to support@fhgfs.com. Of course, you are also welcome to send any other kind of question, recommendation or benchmark result to this e-mail address.

Commercial support for FhGFS is also available. Support fees are based on the number of storage servers (clients are free). Contact support@fhgfs.com for additional information.

5 Licensing

FhGFS and the FhGFS Setup are licensed under the terms of the FhGFS End User License Agreement (see <http://www.fhgfs.com> for more information).

The FhGFS Administration and Monitoring system makes use of the following unmodified third party components:

- Cxxtools and Tntnet by Tommi Mäkitalo (<http://www.tntnet.org>), which are licensed under the GNU Lesser General Public License (LGPL, <http://www.gnu.org/licenses/lgpl.html>).
- Sqlite (<http://www.sqlite.org>), which was given to the public domain by the authors.