

A Tutorial for GNUnet 0.10.x (C version)
Christian Grothoff Bart Polot Matthias Wachs
March 7, 2017

This tutorial explains how to install GNUnet on a GNU/Linux system and gives an introduction on how GNUnet can be used to develop a Peer-to-Peer application. Detailed installation instructions for various operating systems and a detailed list of all dependencies can be found on our website at <https://gnunet.org/installation>.

Please read this tutorial carefully since every single step is important and do not hesitate to contact the GNUnet team if you have any questions or problems! Check here how to contact the GNUnet team: https://gnunet.org/contact_information

1 Installing GNUnet

First of all you have to install a current version of GNUnet. You can download a tarball of a stable version from GNU FTP mirrors or obtain the latest development version from our Git repository.

Most of the time you should prefer to download the stable version since with the latest development version things can be broken, functionality can be changed or tests can fail. You should only use the development version if you know that you require a certain feature or a certain issue has been fixed since the last release.

1.1 Obtaining a stable version

You can download the latest stable version of GNUnet from GNU FTP mirrors:

```
ftp://ftp.gnu.org/gnu/gnunet/gnunet-0.10.x.tar.gz
```

You should also download the signature file and verify the integrity of the tarball.

```
ftp://ftp.gnu.org/gnu/gnunet/gnunet-0.10.x.tar.gz.sig
```

To verify the signature you should first import the GPG key used to sign the tarball

```
1 $ gpg --keyserver keys.gnupg.net --recv-keys 48426C7E
```

And use this key to verify the tarball's signature

```
1 $ gpg --verify gnunet-0.10.x.tar.gz.sig gnunet-0.10.x.tar.gz
```

After successfully verifying the integrity you can extract the tarball using

```
1 $ tar xvzf gnunet-0.10.x.tar.gz
2 $ mv gnunet-0.10.x gnunet # we will use the directory "gnunet" in the remainder of this document
3 $ cd gnunet
```

However, please note that stable versions can be very outdated, as a developer you are strongly encouraged to use the version from <https://gnunet.org/git/>.

1.2 Installing Build Tool Chain and Dependencies

To successfully compile GNUnet you need the tools to build GNUnet and the required dependencies. Please have a look at <https://gnunet.org/dependencies> for a list of required dependencies and https://gnunet.org/generic_installation for specific instructions for your operating system.

Please check the notes at the end of the configure process about required dependencies.

For GNUnet bootstrapping support and the http(s) plugin you should install `libgnurl`. For the filesharing service you should install at least one of the datastore backends `mysql`, `sqlite` or `postgresql`.

1.3 Obtaining the latest version from Git

The latest development version can be obtained from our Git repository. To obtain the code you need Git installed and checkout the repository using:

```
1 $ git clone https://gnunet.org/git/gnunet
```

After cloning the repository you have to execute

```
1 $ cd gnunet
2 $ ./bootstrap
```

The remainder of this tutorial assumes that you have Git Master checked out.

1.4 Compiling and Installing GNUnet

First, you need to install at least `libgnupgerror` version 1.27¹ and `libgcrypt` version 1.7.6².

```
1 $ wget ftp://ftp.gnupg.org/gcrypt/libgpg-error/libgpg-error-1.27.tar.bz2
2 $ tar xf libgpg-error-1.27.tar.bz2
3 $ cd libgpg-error-1.27
4 $ ./configure
5 $ sudo make install
6 $ cd ..
```

```
1 $ wget ftp://ftp.gnupg.org/gcrypt/libgcrypt/libgcrypt-1.7.6.tar.bz2
2 $ tar xf libgcrypt-1.7.6.tar.bz2
3 $ cd libgcrypt-1.7.6
4 $ ./configure
5 $ sudo make install
6 $ cd ..
```

Assuming all dependencies are installed, the following commands will compile and install GNUnet in your home directory. You can specify the directory where GNUnet will be installed by changing the `--prefix` value when calling `./configure`. If you do not specify a prefix, GNUnet is installed in the directory `/usr/local`. When developing new applications you may want to enable verbose logging by adding `--enable-logging=verbose`:

```
1 $ ./configure --prefix=$PREFIX --enable-logging
2 $ make
3 $ make install
```

¹<ftp://ftp.gnupg.org/gcrypt/libgpg-error/libgpg-error-1.27.tar.bz2>

²<ftp://ftp.gnupg.org/gcrypt/libgcrypt/libgcrypt-1.7.6.tar.bz2>

After installing GNUnet you have to add your GNUnet installation to your path environmental variable. In addition you have to create the `.config` directory in your home directory where GNUnet stores its data and an empty GNUnet configuration file:

```
1 $ export PATH=$PATH:$PREFIX/bin
2 $ echo export PATH=$PREFIX/bin:\\$PATH >> ~/.bashrc
3 $ mkdir ~/.config/
4 $ touch ~/.config/gnunet.conf
```

1.5 Common Issues - Check your GNUnet installation

You should check your installation to ensure that installing GNUnet was successful up to this point. You should be able to access GNUnet's binaries and run GNUnet's self check.

```
1 $ which gnunet-arm
```

should return `$PREFIX/bin/gnunet-arm`. It should be located in your GNUnet installation and the output should not be empty. If you see an output like:

```
1 $ which gnunet-arm
2 $
```

check your `PATH` variable to ensure GNUnet's `bin` directory is included. GNUnet provides tests for all of its subcomponents. Run

```
1 $ make check
```

to execute tests for all components. `make check` traverses all subdirectories in `src`. For every subdirectory you should get a message like this:

```
make[2]: Entering directory '/home/$USER/gnunet/contrib'
PASS: test_gnunet_prefix
=====
1 test passed
=====
```

2 Background: GNUnet Architecture

GNUnet is organized in layers and services. Each service is composed of a main service implementation and a client library for other programs to use the service's functionality, described by an API. This approach is shown in figure 1a. Some services provide an additional command line tool to enable the user to interact with the service.

Very often it is other GNUnet services that will use these APIs to build the higher layers of GNUnet on top of the lower ones. Each layer expands or extends the functionality of the service below (for instance, to build a mesh on top of a DHT). See figure 1b for an illustration of this approach.

The main service implementation runs as a standalone process in the operating system and the client code runs as part of the client program, so crashes of a client do not affect the service process or other clients. The service and the clients communicate via a message protocol to be defined and implemented by the programmer.




Figure 1: GNUnet's layered system architecture

3 First Steps with GNUnet

3.1 Configure your peer

First of all we need to configure your peer. Each peer is started with a configuration containing settings for GNUnet itself and its services. This configuration is based on the default configuration shipped with GNUnet and can be modified. The default configuration is located in the `$PREFIX/share/gnunet/config.d` directory. When starting a peer, you can specify a customized configuration using the the `-c` command line switch when starting the ARM service and all other services. When using a modified configuration the default values are loaded and only values specified in the configuration file will replace the default values.

Since we want to start additional peers later, we need some modifications from the default configuration. We need to create a separate service home and a file containing our modifications for this peer:

```
1 $ mkdir ~/gnunet1/
2 $ touch peer1.conf
```

Now add the following lines to `peer1.conf` to use this directory. For simplified usage we want to prevent the peer to connect to the GNUnet network since this could lead to confusing output. This modifications will replace the default settings:

```
[PATHS]
GNUNET_HOME = ~/gnunet1/ # Use this directory to store GNUnet data
[hostlist]
SERVERS = # prevent bootstrapping
```

3.2 Start a peer

Each GNUnet instance (called peer) has an identity (*peer ID*) based on a cryptographic public private key pair. The peer ID is the printable hash of the public key.

GNUnet services are controlled by a master service the so called *Automatic Restart Manager* (ARM). ARM starts, stops and even restarts services automatically or on demand when a client connects. You interact with the ARM service using the `gnunet-arm` tool. GNUnet can then be started with `gnunet-arm -s` and stopped with `gnunet-arm -e`. An additional service not automatically started can be started using `gnunet-arm -i <service name>` and stopped using `gnunet-arm -k <servicename>`.

Once you have started your peer, you can use many other GNUnet commands to interact with it. For example, you can run:

```
1 $ gnunet-peerinfo -s
```

to obtain the public key of your peer. You should see an output containing the peer ID similar to:

```
1 I am peer `0PA02UVRKQTS2C .. JL5Q78F6H0B1ACPV1CJI59MEQUMQCC5G'.
```

3.3 Monitor a peer

In this section, we will monitor the behaviour of our peer's DHT service with respect to a specific key. First we will start GNUnet and then start the DHT service and use the DHT monitor tool to monitor the PUT and GET commands we issue using the `gnunet-dht-put` and `gnunet-dht-get` commands. Using the "monitor" line given below, you can observe the behavior of your own peer's DHT with respect to the specified KEY:

```
1 $ gnunet-arm -c ~/peer1.conf -s # start gnunet with all default services
2 $ gnunet-arm -c ~/peer1.conf -i dht # start DHT service
3 $ cd ~/gnunet/src/dht;
4 $ ./gnunet-dht-monitor -c ~/peer1.conf -k KEY
```

Now open a separate terminal and change again to the `gnunet/src/dht` directory:

```
1 $ cd ~/gnunet/src/dht
2 $ ./gnunet-dht-put -c ~/peer1.conf -k KEY -d VALUE # put VALUE under KEY in the DHT
3 $ ./gnunet/src/dht/gnunet-dht-get -c ~/peer1.conf -k KEY # get key KEY from the DHT
4 $ gnunet-statistics -c ~/peer1.conf # print statistics about current GNUnet state
5 $ gnunet-statistics -c ~/peer1.conf -s dht # print statistics about DHT service
```

3.4 Starting Two Peers by Hand

3.4.1 Setup a second peer

We will now start a second peer on your machine. For the second peer, you will need to manually create a modified configuration file to avoid conflicts with ports and directories. A peers configuration file is by default located in `./gnunet/gnunet.conf`. This file is typically very short or even empty as only the differences to the defaults need to be specified. The defaults are located in many files in the `$PREFIX/share/gnunet/config.d` directory.

To configure the second peer, use the files `$PREFIX/share/gnunet/config.d` as a template for your main configuration file:

```
1 $ cat $PREFIX/share/gnunet/config.d/*.conf > peer2.conf
```

Now you have to edit `peer2.conf` and change:

- `SERVICEHOME` under `PATHS`
- Every (uncommented) value for “`PORT`” (add 10000) in any section (the option may be commented out if `PORT` is prefixed by “`#`”, in this case, UNIX domain sockets are used and the `PORT` option does not need to be touched)
- Every value for “`UNIXPATH`” in any section (e.g. by adding a “`-p2`” suffix)

to a fresh, unique value. Make sure that the `PORT` numbers stay below 65536. From now on, whenever you interact with the second peer, you need to specify `-c peer2.conf` as an additional command line argument.

Now, generate the 2nd peer’s private key:

```
1 $ gnunet-peerinfo -s -c peer2.conf
```

This may take a while, generate entropy using your keyboard or mouse as needed. Also, make sure the output is different from the `gnunet-peerinfo` output for the first peer (otherwise you made an error in the configuration).

3.4.2 Start the second peer and connect the peers

Then, you can start a second peer using:

```
1 $ gnunet-arm -c peer2.conf -s
2 $ gnunet-arm -c peer2.conf -i dht
3 $ ~/gnunet/src/dht/gnunet-dht-put -c peer2.conf -k KEY -d VALUE
4 $ ~/gnunet/src/dht/gnunet-dht-get -c peer2.conf -k KEY
```

If you want the two peers to connect, you have multiple options:

- UDP neighbour discovery (automatic)
- Setup a bootstrap server
- Connect manually

To setup peer 1 as bootstrapping server change the configuration of the first one to be a hostlist server by adding the following lines to `peer1.conf` to enable bootstrapping server:

```
[hostlist]
OPTIONS = -p
```

Then change `peer2.conf` and replace the “`SERVERS`” line in the “[`hostlist`]” section with “`http://localhost:8080/`”. Restart both peers using:

```
1 $ gnunet-arm -c peer1.conf -e # stop first peer
2 $ gnunet-arm -c peer1.conf -s # start first peer
3 $ gnunet-arm -c peer2.conf -s # start second peer
```

Note that if you start your peers without changing these settings, they will use the “global” hostlist servers of the GNUnet P2P network and likely connect to those peers. At that point, debugging might become tricky as you’re going to be connected to many more peers and would likely observe traffic and behaviors that are not explicitly controlled by you.

3.4.3 How to connect manually

If you want to use the `peerinfo` tool to connect your peers, you should:

- Set `FORCESTART = NO` in section `hostlist` (to not connect to the global GNUnet)

- Start both peers running `gnunet-arm -c peer1.conf -s` and `gnunet-arm -c peer2.conf -s`
- Get HELLO message of the first peer running `gnunet-peerinfo -c peer1.conf -g`
- Give the output to the second peer by running `gnunet-peerinfo -c peer2.conf -p '<output>'`

Check that they are connected using `gnunet-core -c peer1.conf`, which should give you the other peer's peer identity:

```
1 $ gnunet-core -c peer1.conf
2 Peer `9TVUCS8P5A7ILLBGO6 [...shortened...] 1KNBJ4NGCHP3JPVULDG'
```

3.5 Starting Peers Using the Testbed Service

GNUnet's testbed service is used for testing scenarios where a number of peers are to be started. The testbed can manage peers on a single host or on multiple hosts in a distributed fashion. On a single affordable computer, it should be possible to run around tens of peers without drastically increasing the load on the system.

The testbed service can be access through its API `include/gnunet_testbed_service.h`. The API provides many routines for managing a group of peers. It also provides a helper function `GNUNET_TESTBED_test_run()` to quickly setup a minimalistic testing environment on a single host.

This function takes a configuration file which will be used as a template configuration for the peers. The testbed takes care of modifying relevant options in the peers' configuration such as `SERVICEHOME`, `PORT`, `UNIXPATH` to unique values so that peers run without running into conflicts. It also checks and assigns the ports in configurations only if they are free.

Additionally, the testbed service also reads its options from the same configuration file. Various available options and details about them can be found in the testbed default configuration file `src/testbed/testbed.conf`.

With the testbed API, a sample test case can be structured as follows:

```
1 #include <unistd.h>
2 #include <gnunet/platform.h>
3 #include <gnunet/gnunet_util_lib.h>
4 #include <gnunet/gnunet_testbed_service.h>
5 #include <gnunet/gnunet_dht_service.h>
6
7 #define NUM_PEERS 20
8
9 static struct GNUNET_TESTBED_Operation *dht_op;
10
11 static struct GNUNET_DHT_Handle *dht_handle;
12
13
14 struct MyContext
15 {
16     int ht_len;
17 } ctxt;
18
19
20 static int result;
21
22
23 static void
24 shutdown_task (void *cls)
25 {
26     if (NULL != dht_op)
27     {
28         GNUNET_TESTBED_operation_done (dht_op);
29         dht_op = NULL;
```

```

30     dht_handle = NULL;
31 }
32 result = GNUNET_OK;
33 }
34
35
36 static void
37 service_connect_comp (void *cls,
38                      struct GNUNET_TESTBED_Operation *op,
39                      void *ca_result,
40                      const char *emsg)
41 {
42     GNUNET_assert (op == dht_op);
43     dht_handle = ca_result;
44     // Do work here...
45     GNUNET_SCHEDULER_shutdown ();
46 }
47
48
49 static void *
50 dht_ca (void *cls, const struct GNUNET_CONFIGURATION_Handle *cfg)
51 {
52     struct MyContext *ctxt = cls;
53
54     dht_handle = GNUNET_DHT_connect (cfg, ctxt->ht_len);
55     return dht_handle;
56 }
57
58
59 static void
60 dht_da (void *cls, void *op_result)
61 {
62     struct MyContext *ctxt = cls;
63
64     GNUNET_DHT_disconnect ((struct GNUNET_DHT_Handle *) op_result);
65     dht_handle = NULL;
66 }
67
68
69 static void
70 test_master (void *cls,
71             struct GNUNET_TESTBED_RunHandle *h,
72             unsigned int num_peers,
73             struct GNUNET_TESTBED_Peer **peers,
74             unsigned int links_succeeded,
75             unsigned int links_failed)
76 {
77     ctxt.ht_len = 10;
78     dht_op = GNUNET_TESTBED_service_connect
79         (NULL, peers[0], "dht",
80          &service_connect_comp, NULL,
81          &dht_ca, &dht_da, &ctxt);
82     GNUNET_SCHEDULER_add_shutdown (&shutdown_task, NULL);
83 }
84
85
86 int
87 main (int argc, char **argv)

```



```

88 {
89     int ret;
90
91     result = GNUNET_SYSERR;
92     ret = GNUNET_TESTBED_test_run
93         ("awesome-test", "template.conf",
94         NUM_PEERS, 0LL,
95         NULL, NULL, &test_master, NULL);
96     if ( (GNUNET_OK != ret) || (GNUNET_OK != result) )
97         return 1;
98     return 0;
99 }

```

The source code for the above listing can be found at https://gnunet.org/git/gnunet.git/tree/doc/testbed_test.c or in the `doc/` folder of your repository check-out. After installing GNUnet, the above source code can be compiled as:

```

1 $ export CPPFLAGS="-I/path/to/gnunet/headers"
2 $ export LDFLAGS="-L/path/to/gnunet/libraries"
3 $ gcc $CPPFLAGS $LDFLAGS -o testbed-test testbed_test.c -lgnunettestbed -lgnunetdht -lgnunetutil

```

The `CPPFLAGS` and `LDFLAGS` are necessary if GNUnet is installed into a different directory other than `/usr/local`.

All of testbed API's peer management functions treat management actions as operations and return operation handles. It is expected that the operations begin immediately, but they may get delayed (to balance out load on the system). The program using the API then has to take care of marking the operation as “done” so that its associated resources can be freed immediately and other waiting operations can be executed. Operations will be canceled if they are marked as “done” before their completion.

An operation is treated as completed when it succeeds or fails. Completion of an operation is either conveyed as events through *controller event callback* or through respective operation completion callbacks. In functions which support completion notification through both controller event callback and operation completion callback, first the controller event callback will be called. If the operation is not marked as done in that callback or if the callback is given as `NULL` when creating the operation, the operation completion callback will be called. The API documentation shows which event are to be expected in the controller event notifications. It also documents any exceptional behaviour.

Once the peers are started, test cases often need to connect some of the peers' services. Normally, opening a connect to a peer's service requires the peer's configuration. While using testbed, the testbed automatically generates per-peer configuration. Accessing those configurations directly through file system is discouraged as their locations are dynamically created and will be different among various runs of testbed. To make access to these configurations easy, testbed API provides the function `GNUNET_TESTBED_service_connect()`. This function fetches the configuration of a given peer and calls the *Connect Adapter*. In the example code, it is the `dht_ca`. A connect adapter is expected to open the connection to the needed service by using the provided configuration and return the created service connection handle. Successful connection to the needed service is signaled through `service_connect_comp_cb`.

A dual to connect adapter is the *Disconnect Adapter*. This callback is called after the connect adapter has been called when the operation from `GNUNET_TESTBED_service_connect()` is marked as “done”. It has to disconnect from the service with the provided service handle (`op_result`).

Exercise: Find out how many peers you can run on your system.

Exercise: Find out how to create a 2D torus topology by changing the options in the configuration file.^a Then use the DHT API to store and retrieve values in the network.

^aSee <https://gnunet.org/content/supported-topologies>

4 Developing Applications

4.1 gnunet-ext

To develop a new peer-to-peer application or to extend GNUnet we provide a template build system for writing GNUnet extensions in C. It can be obtained as follows:

```
1 $ git clone https://gnunet.org/git/gnunet-ext
2 $ cd gnunet-ext/
3 $ ./bootstrap
4 $ ./configure --prefix=$PREFIX --with-gnunet=$PREFIX
5 $ make
6 $ make install
7 $ make check
```

The GNUnet ext template includes examples and a working buildsystem for a new GNUnet service. A common GNUnet service consists of the following parts which will be discussed in detail in the remainder of this document. The functionality of a GNUnet service is implemented in:

- the GNUnet service (gnunet-ext/src/ext/gnunet-service-ext.c)
- the client API (gnunet-ext/src/ext/ext_api.c)
- the client application using the service API (gnunet-ext/src/ext/gnunet-ext.c)

The interfaces for these entities are defined in:

- client API interface (gnunet-ext/src/ext/ext.h)
- the service interface (gnunet-ext/src/include/gnunet_service_SERVICE.h)
- the P2P protocol (gnunet-ext/src/include/gnunet_protocols_ext.h)

In addition the ext systems provides:

- a test testing the API (gnunet-ext/src/ext/test_ext_api.c)
- a configuration template for the service (gnunet-ext/src/ext/ext.conf.in)

4.2 Adapting the Template

The first step for writing any extension with a new service is to ensure that the `ext.conf.in` file contains entries for the `UNIXPATH`, `PORT` and `BINARY` for the service in a section named after the service.

If you want to adapt the template rename the `ext.conf.in` to match your services name, you have to modify the `AC_OUTPUT` section in `configure.ac` in the `gnunet-ext` root.

5 Writing a Client Application

When writing any client application (for example, a command-line tool), the basic structure is to start with the `GNUNET_PROGRAM_run` function. This function will parse command-line options, setup the scheduler and then invoke the `run` function (with the remaining non-option arguments) and a handle to the parsed configuration (and the configuration file name that was used, which is typically not needed):

```

1 #include <gnunet/platform.h>
2 #include <gnunet/gnunet_util_lib.h>
3
4 static int ret;
5
6 static void
7 run (void *cls,
8     char *const *args,
9     const char *cfgfile,
10    const struct GNUNET_CONFIGURATION_Handle *cfg)
11 {
12     // main code here
13     ret = 0;
14 }
15
16 int
17 main (int argc, char *const *argv)
18 {
19     static const struct GNUNET_GETOPT_CommandLineOption options[] = {
20         GNUNET_GETOPT_OPTION_END
21     };
22     return (GNUNET_OK ==
23           GNUNET_PROGRAM_run (argc,
24                               argv,
25                               "binary-name",
26                               gettext_noop ("binary description text"),
27                               options, &run, NULL)) ? ret : 1;
28 }

```

5.1 Handling command-line options

Options can then be added easily by adding global variables and expanding the `options` array. For example, the following would add a string-option and a binary flag (defaulting to `NULL` and `GNUNET_NO` respectively):

```

1 static char *string_option;
2 static int a_flag;
3
4 // ...
5 static const struct GNUNET_GETOPT_CommandLineOption options[] = {
6     {'s', "name", "SOMESTRING",
7     gettext_noop ("text describing the string_option NAME"), 1,
8     &GNUNET_GETOPT_set_string, &string_option},
9     {'f', "flag", NULL,
10    gettext_noop ("text describing the flag option"), 0,
11    &GNUNET_GETOPT_set_one, &a_flag},
12    GNUNET_GETOPT_OPTION_END
13 };
14 string_option = NULL;
15 a_flag = GNUNET_SYSERR;
16 // ...

```

Issues such as displaying some helpful text describing options using the `-help` argument and error handling are taken care of when using this approach. Other `GNUNET_GETOPT_-`functions can be used to obtain integer value options, increment counters, etc. You can even write custom option parsers for special circumstances not covered by the available handlers. To

check if an argument was specified by the user you initialize the variable with a specific value (e.g. NULL for a string and GNUNET_SYSERR for a integer) and check after parsing happened if the values were modified.

Inside the `run` method, the program would perform the application-specific logic, which typically involves initializing and using some client library to interact with the service. The client library is supposed to implement the IPC whereas the service provides more persistent P2P functions.

Exercise: Add a few command-line options and print them inside of `run`. What happens if the user gives invalid arguments?

5.2 Writing a Client Library

The first and most important step in writing a client library is to decide on an API for the library. Typical API calls include connecting to the service, performing application-specific requests and cleaning up. Many examples for such service APIs can be found in the `gnunet/src/include/gnunet*_service.h` files.

Then, a client-service protocol needs to be designed. This typically involves defining various message formats in a header that will be included by both the service and the client library (but is otherwise not shared and hence located within the service's directory and not installed by `make install`). Each message must start with a `struct GNUNET_MessageHeader` and must be shorter than 64k. By convention, all fields in IPC (and P2P) messages must be in big-endian format (and thus should be read using `ntohl` and similar functions and written using `htonl` and similar functions). Unique message types must be defined for each message struct in the `gnunet_protocols.h` header (or an extension-specific include file).

5.2.1 Connecting to the Service

Before a client library can implement the application-specific protocol with the service, a connection must be created:

```
1 struct GNUNET_MQ_MessageHandlers handlers[] = {
2     // ...
3     GNUNET_MQ_handler_end ()
4 };
5 struct GNUNET_MQ_Handle *mq;
6
7 mq = GNUNET_CLIENT_connect (cfg, "service-name", handlers, &error_cb, NULL);
```

As a result a `GNUNET_MQ_Handle` is returned which can be used henceforth to transmit messages to the service. The complete MQ API can be found in `gnunet_mq_lib.h`. The `handlers` array in the example above is incomplete. Here is where you will define which messages you expect to receive from the service, and which functions handle them. The `error_cb` is a function that is to be called whenever there are errors communicating with the service.

5.2.2 Sending messages

In GNUnet, messages are always sent beginning with a `struct GNUNET_MessageHeader` in big endian format. This header defines the size and the type of the message, the payload follows after this header.

```
1 struct GNUNET_MessageHeader
2 {
3     uint16_t size GNUNET_PACKED;
4     uint16_t type GNUNET_PACKED;
5 };
```

Existing message types are defined in `gnunet_protocols.h`. A common way to create a message is with an envelope:

```

1 struct GNUNET_MQ_Envelope *env;
2 struct GNUNET_MessageHeader *msg;
3
4 env = GNUNET_MQ_msg_extra (msg, payload_size, GNUNET_MY_MESSAGE_TYPE);
5 memcpy (&msg[1], &payload, payload_size);
6 // Send message via message queue 'mq'
7 GNUNET_mq_send (mq, env);

```

Exercise: Define a message struct that includes a 32-bit unsigned integer in addition to the standard GNUnet MessageHeader. Add a C struct and define a fresh protocol number for your message. (Protocol numbers in gnunet-ext are defined in gnunet-ext/src/include/gnunet_protocols_ext.h)

Exercise: Find out how you can determine the number of messages in a message queue.

Exercise: Find out how you can determine when a message you have queued was actually transmitted.

Exercise: Define a helper function to transmit a 32-bit unsigned integer (as payload) to a service using some given client handle.

5.2.3 Receiving Replies from the Service

Clients can receive messages from the service using the handlers specified in the `handlers` array we specified when connecting to the service. Entries in the array are usually created using one of two macros, depending on whether the message is fixed size or variable size. Variable size messages are managed using two callbacks, one to check that the message is well-formed, the other to actually process the message. Fixed size messages are fully checked by the MQ-logic, and thus only need to provide the handler to process the message. Note that the prefixes `check_` and `handle_` are mandatory.

```

1 static void
2 handle_fix (void *cls, const struct MyMessage *msg)
3 {
4     // process 'msg'
5 }
6
7 static int
8 check_var (void *cls, const struct MyVarMessage *msg)
9 {
10    // check 'msg' is well-formed
11    return GNUNET_OK;
12 }
13
14 static void
15 handle_var (void *cls, const struct MyVarMessage *msg)
16 {
17     // process 'msg'
18 }
19
20 struct GNUNET_MQ_MessageHandler handlers[] = {
21     GNUNET_MQ_hd_fixed_size (fix,
22                             GNUNET_MESSAGE_TYPE_MY_FIX,
23                             struct MyMessage,
24                             NULL),
25     GNUNET_MQ_hd_fixed_size (var,
26                             GNUNET_MESSAGE_TYPE_MY_VAR,
27                             struct MyVarMessage,
28                             NULL),

```

```

29 |
30 | GNUNET_MQ_handler_end ()
31 | };

```

Exercise: Expand your helper function to receive a response message (for example, containing just the `struct GNUnet MessageHeader` without any payload). Upon receiving the service's response, you should call a callback provided to your helper function's API.

Exercise: Figure out where you can pass values to the closures (`cls`).

5.3 Writing a user interface

Given a client library, all it takes to access a service now is to combine calls to the client library with parsing command-line options.

Exercise: Call your client API from your `run()` method in your client application to send a request to the service. For example, send a 32-bit integer value based on a number given at the command-line to the service.

6 Writing a Service

Before you can test the client you've written so far, you'll need to also implement the corresponding service.

6.1 Code Placement

New services are placed in their own subdirectory under `gnUNET/src`. This subdirectory should contain the API implementation file `SERVICE_api.c`, the description of the client-service protocol `SERVICE.h` and P2P protocol `SERVICE_protocol.h`, the implementation of the service itself `gnUNET-service-SERVICE.h` and several files for tests, including test code and configuration files.

6.2 Starting a Service

The key API definition for creating a service is the `GNUNET_SERVICE_MAIN` macro:

```

1 GNUNET_SERVICE_MAIN
2 ("service-name",
3  GNUNET_SERVICE_OPTION_NONE,
4  &run,
5  &client_connect_cb,
6  &client_disconnect_cb,
7  NULL,
8  GNUNET_MQ_hd_fixed_size (...),
9  GNUNET_MQ_hd_var_size (...),
10 GNUNET_MQ_handler_end ());

```

In addition to the service name and flags, the macro takes three functions, typically called `run`, `client_connect_cb` and `client_disconnect_cb` as well as an array of message handlers that will be called for incoming messages from clients.

A minimal version of the three central service functions would look like this:

```

1 static void
2 run (void *cls,
3      const struct GNUNET_CONFIGURATION_Handle *c,
4      struct GNUNET_SERVICE_Handle *service)

```

```

5 | {
6 | }
7 |
8 | static void *
9 | client_connect_cb (void *cls,
10 |     struct GNUNET_SERVICE_Client *c,
11 |     struct GNUNET_MQ_Handle *mq)
12 | {
13 |     return c;
14 | }
15 |
16 | static void
17 | client_disconnect_cb (void *cls,
18 |     struct GNUNET_SERVICE_Client *c,
19 |     void *internal_cls)
20 | {
21 |     GNUNET_assert (c == internal_cls);
22 | }

```

Exercise: Write a stub service that processes no messages at all in your code. Create a default configuration for it, integrate it with the build system and start the service from `gnunet-service-arm` using `gnunet-arm -i NAME`.

Exercise: Figure out how to set the closure (`cls`) for handlers of a service.

Exercise: Figure out how to send messages from the service back to the client.

Each handler function in the service **must** eventually (possibly in some asynchronous continuation) call `GNUNET_SERVICE_client_continue()`. Only after this call additional messages from the same client may be processed. This way, the service can throttle processing messages from the same client.

Exercise: Change the service to “handle” the message from your client (for now, by printing a message). What happens if you forget to call `GNUNET_SERVICE_client_continue()`?

7 Interacting directly with other Peers using the CORE Service

FIXME: This section still needs to be updated to the latest API!

One of the most important services in GNUnet is the CORE service managing connections between peers and handling encryption between peers.

One of the first things any service that extends the P2P protocol typically does is connect to the CORE service using:

```

1 | #include <gnunet/gnunet_core_service.h>
2 |
3 | struct GNUNET_CORE_Handle *
4 | GNUNET_CORE_connect (const struct GNUNET_CONFIGURATION_Handle *cfg,
5 |     void *cls,
6 |     GNUNET_CORE_StartupCallback init,
7 |     GNUNET_CORE_ConnectEventHandler connects,
8 |     GNUNET_CORE_DisconnectEventHandler disconnects,
9 |     const struct GNUNET_MQ_MessageHandler *handlers);

```

7.1 New P2P connections

Before any traffic with a different peer can be exchanged, the peer must be known to the service. This is notified by the `CORE connects` callback, which communicates the identity of the new peer to the service:

```

1 void *
2 connects (void *cls,
3         const struct GNUNET_PeerIdentity *peer,
4         struct GNUNET_MQ_Handle *mq)
5 {
6     return mq;
7 }

```

Note that whatever you return from `connects` is given as the `cls` argument to the message handlers for messages from the respective peer.

Exercise: Create a service that connects to the CORE. Then start (and connect) two peers and print a message once your connect callback is invoked.

7.2 Receiving P2P Messages

To receive messages from CORE, you pass the desired handlers to the `GNUNET_CORE_connect()` function, just as we showed for services.

It is your responsibility to process messages fast enough or to implement flow control. If an application does not process CORE messages fast enough, CORE will randomly drop messages to not keep a very long queue in memory.

Exercise: Start one peer with a new service that has a message handler and start a second peer that only has your “old” service without message handlers. Which “connect” handlers are invoked when the two peers are connected? Why?

7.3 Sending P2P Messages

You can transmit messages to other peers using the `mq` you were given during the `connect` callback. Note that the `mq` automatically is released upon `disconnect` and that you must not use it afterwards.

It is your responsibility to not over-fill the message queue, GUNet will send the messages roughly in the order given as soon as possible.

Exercise: Write a service that upon connect sends messages as fast as possible to the other peer (the other peer should run a service that “processes” those messages). How fast is the transmission? Count using the STATISTICS service on both ends. Are messages lost? How can you transmit messages faster? What happens if you stop the peer that is receiving your messages?

7.4 End of P2P connections

If a message handler returns `GNUNET_SYSERR`, the remote peer shuts down or there is an unrecoverable network disconnection, CORE notifies the service that the peer disconnected. After this notification no more messages will be received from the peer and the service is no longer allowed to send messages to the peer. The disconnect callback looks like the following:

```

1 void
2 disconnects (void *cls,
3             const struct GNUNET_PeerIdentity *peer)
4 {
5     /* Remove peer's identity from known peers */
6     /* Make sure no messages are sent to peer from now on */
7 }

```

Exercise: Fix your service to handle peer disconnects.

8 Storing peer-specific data using the PEERSTORE service

GNUnet's PEERSTORE service offers a persistance for arbitrary peer-specific data. Other GNUnet services can use the PEERSTORE to store, retrieve and monitor data records. Each data record stored with PEERSTORE contains the following fields:

- subsystem: Name of the subsystem responsible for the record.
- peerid: Identity of the peer this record is related to.
- key: a key string identifying the record.
- value: binary record value.
- expiry: record expiry date.

The first step is to start a connection to the PEERSTORE service:

```
1 #include "gnunet_peerstore_service.h"
2
3 peerstore_handle = GNUNET_PEERSTORE_connect (cfg);
```

The service handle `peerstore_handle` will be needed for all subsequent PEERSTORE operations.

8.1 Storing records

To store a new record, use the following function:

```
1 struct GNUNET_PEERSTORE_StoreContext *
2 GNUNET_PEERSTORE_store (struct GNUNET_PEERSTORE_Handle *h,
3                          const char *sub_system,
4                          const struct GNUNET_PeerIdentity *peer,
5                          const char *key,
6                          const void *value,
7                          size_t size,
8                          struct GNUNET_TIME_Absolute expiry,
9                          enum GNUNET_PEERSTORE_StoreOption options,
10                         GNUNET_PEERSTORE_Continuation cont,
11                         void *cont_cls);
```

The `options` parameter can either be `GNUNET_PEERSTORE_STOREOPTION_MULTIPLE` which means that multiple values can be stored under the same key combination (subsystem, peerid, key), or `GNUNET_PEERSTORE_STOREOPTION_REPLACE` which means that PEERSTORE will replace any existing values under the given key combination (subsystem, peerid, key) with the new given value.

The continuation function `cont` will be called after the store request is successfully sent to the PEERSTORE service. This does not guarantee that the record is successfully stored, only that it was received by the service.

The `GNUNET_PEERSTORE_store` function returns a handle to the store operation. This handle can be used to cancel the store operation only before the continuation function is called:

```
1 void
2 GNUNET_PEERSTORE_store_cancel (struct GNUNET_PEERSTORE_StoreContext *sc);
```

8.2 Retrieving records

To retrieve stored records, use the following function:

```
1 struct GNUNET_PEERSTORE_IterateContext *
2 GNUNET_PEERSTORE_iterate (struct GNUNET_PEERSTORE_Handle *h,
3     const char *sub_system,
4     const struct GNUNET_PeerIdentity *peer,
5     const char *key,
6     struct GNUNET_TIME_Relative timeout,
7     GNUNET_PEERSTORE_Processor callback,
8     void *callback_cls);
```

The values of peer and key can be NULL. This allows the iteration over values stored under any of the following key combinations:

- (subsystem)
- (subsystem, peerid)
- (subsystem, key)
- (subsystem, peerid, key)

The callback function will be called once with each retrieved record and once more with a NULL record to signal the end of results.

The GNUNET_PEERSTORE_iterate function returns a handle to the iterate operation. This handle can be used to cancel the iterate operation only before the callback function is called with a NULL record.

8.3 Monitoring records

PEERSTORE offers the functionality of monitoring for new records stored under a specific key combination (subsystem, peerid, key). To start the monitoring, use the following function:

```
1 struct GNUNET_PEERSTORE_WatchContext *
2 GNUNET_PEERSTORE_watch (struct GNUNET_PEERSTORE_Handle *h,
3     const char *sub_system,
4     const struct GNUNET_PeerIdentity *peer,
5     const char *key,
6     GNUNET_PEERSTORE_Processor callback,
7     void *callback_cls);
```

Whenever a new record is stored under the given key combination, the callback function will be called with this new record. This will continue until the connection to the PEERSTORE service is broken or the watch operation is canceled:

```
1 void
2 GNUNET_PEERSTORE_watch_cancel (struct GNUNET_PEERSTORE_WatchContext *wc);
```

8.4 Disconnecting from PEERSTORE

When the connection to the PEERSTORE service is no longer needed, disconnect using the following function:

```
1 void
2 GNUNET_PEERSTORE_disconnect (struct GNUNET_PEERSTORE_Handle *h, int sync_first);
```

If the sync_first flag is set to GNUNET_YES, the API will delay the disconnection until all store requests are received by the PEERSTORE service. Otherwise, it will disconnect immediately.

9 Using the DHT

The DHT allows to store data so other peers in the P2P network can access it and retrieve data stored by any peers in the network. This section will explain how to use the DHT. Of course, the first thing to do is to connect to the DHT service:

```
1 dht_handle = GNUNET_DHT_connect (cfg, parallel_requests);
```

The second parameter indicates how many requests in parallel to expect. It is not a hard limit, but a good approximation will make the DHT more efficient.

9.1 Storing data in the DHT

Since the DHT is a dynamic environment (peers join and leave frequently) the data that we put in the DHT does not stay there indefinitely. It is important to “refresh” the data periodically by simply storing it again, in order to make sure other peers can access it.

The put API call offers a callback to signal that the PUT request has been sent. This does not guarantee that the data is accessible to others peers, or even that it has been stored, only that the service has requested to a neighboring peer the retransmission of the PUT request towards its final destination. Currently there is no feedback about whether or not the data has been successfully stored or where it has been stored. In order to improve the availability of the data and to compensate for possible errors, peers leaving and other unfavorable events, just make several PUT requests!

```
1 static void
2 message_sent_cont (void *cls, const struct GNUNET_SCHEDULER_TaskContext *tc)
3 {
4     // Request has left local node
5 }
6
7 struct GNUNET_DHT_PutHandle *
8 GNUNET_DHT_put (struct GNUNET_DHT_Handle *handle,
9                 const struct GNUNET_HashCode *key,
10                uint32_t desired_replication_level,
11                enum GNUNET_DHT_RouteOption options,
12                enum GNUNET_BLOCK_Type type, size_t size, const void *data,
13                struct GNUNET_TIME_Absolute exp,
14                struct GNUNET_TIME_Relative timeout,
15                GNUNET_DHT_PutContinuation cont, void *cont_cls)
```

Exercise: Store a value in the DHT periodically to make sure it is available over time. You might consider using the function `GNUNET_SCHEDULER_add_delayed` and call `GNUNET_DHT_put` from inside a helper function.

9.2 Obtaining data from the DHT

As we saw in the previous example, the DHT works in an asynchronous mode. Each request to the DHT is executed “in the background” and the API calls return immediately. In order to receive results from the DHT, the API provides a callback. Once started, the request runs in the service, the service will try to get as many results as possible (filtering out duplicates) until the timeout expires or we explicitly stop the request. It is possible to give a “forever” timeout with `GNUNET_TIME_UNIT_FOREVER_REL`.

If we give a route option `GNUNET_DHT_RO_RECORD_ROUTE` the callback will get a list of all the peers the data has travelled, both on the PUT path and on the GET path.

```
1 static void
2 get_result_iterator (void *cls, struct GNUNET_TIME_Absolute expiration,
3                     const struct GNUNET_HashCode *key,
```

```

4     const struct GNUNET_PeerIdentity *get_path,
5     unsigned int get_path_length,
6     const struct GNUNET_PeerIdentity *put_path,
7     unsigned int put_path_length,
8     enum GNUNET_BLOCK_Type type, size_t size, const void *data)
9 {
10  // Optionally:
11  GNUNET_DHT_get_stop (get_handle);
12 }
13
14 get_handle =
15     GNUNET_DHT_get_start (dht_handle,
16                          block_type,
17                          &key,
18                          replication,
19                          GNUNET_DHT_RO_NONE,
20                          NULL,
21                          0,
22                          &get_result_iterator,
23                          cls)

```

Exercise: Store a value in the DHT and after a while retrieve it. Show the IDs of all the peers the requests have gone through. In order to convert a peer ID to a string, use the function `GNUNET_i2s`. Pay attention to the route option parameters in both calls!

9.3 Implementing a block plugin

In order to store data in the DHT, it is necessary to provide a block plugin. The DHT uses the block plugin to ensure that only well-formed requests and replies are transmitted over the network.

The block plugin should be put in a file `plugin_block_SERVICE.c` in the service's respective directory. The mandatory functions that need to be implemented for a block plugin are described in the following sections.

9.3.1 Validating requests and replies

The evaluate function should validate a reply or a request. It returns a `GNUNET_BLOCK_EvaluationResult`, which is an enumeration. All possible answers are in `gnunet_block_lib.h`. The function will be called with a `reply_block` argument of `NULL` for requests. Note that depending on how `evaluate` is called, only some of the possible return values are valid. The specific meaning of the `xquery` argument is application-specific. Applications that do not use an extended query should check that the `xquery_size` is zero. The block group is typically used to filter duplicate replies.

```

1 static enum GNUNET_BLOCK_EvaluationResult
2 block_plugin_SERVICE_evaluate (void *cls,
3                                enum GNUNET_BLOCK_Type type,
4                                struct GNUNET_BlockGroup *bg,
5                                const GNUNET_HashCode *query,
6                                const void *xquery,
7                                size_t xquery_size,
8                                const void *reply_block,
9                                size_t reply_block_size)
10 {
11  // Verify type, block and bg
12 }

```

Note that it is mandatory to detect duplicate replies in this function and return the respective status code. Duplicate detection is typically done using the Bloom filter block group provided by `libgnunetblockgroup.so`. Failure to do so may cause replies to circle in the network.

9.3.2 Deriving a key from a reply

The DHT can operate more efficiently if it is possible to derive a key from the value of the corresponding block. The `get_key` function is used to obtain the key of a block — for example, by means of hashing. If deriving the key is not possible, the function should simply return `GNUNET_SYSERR` (the DHT will still work just fine with such blocks).

```
1 static int
2 block_plugin_SERVICE_get_key (void *cls, enum GNUNET_BLOCK_Type type,
3                               const void *block, size_t block_size,
4                               struct GNUNET_HashCode *key)
5 {
6     // Store the key in the key argument, return GNUNET_OK on success.
7 }
```

9.3.3 Initialization of the plugin

The plugin is realized as a shared C library. The library must export an initialization function which should initialize the plugin. The initialization function specifies what block types the plugin cares about and returns a struct with the functions that are to be used for validation and obtaining keys (the ones just defined above).

```
1 void *
2 libgnunet_plugin_block_SERVICE_init (void *cls)
3 {
4     static enum GNUNET_BLOCK_Type types[] =
5     {
6         GNUNET_BLOCK_TYPE_SERVICE_BLOCKTYPE,
7         GNUNET_BLOCK_TYPE_ANY
8     };
9     struct GNUNET_BLOCK_PluginFunctions *api;
10
11     api = GNUNET_new (struct GNUNET_BLOCK_PluginFunctions);
12     api->evaluate = &block_plugin_SERVICE_evaluate;
13     api->get_key = &block_plugin_SERVICE_get_key;
14     api->types = types;
15     return api;
16 }
```

9.3.4 Shutdown of the plugin

Following GUNet's general plugin API concept, the plugin must export a second function for cleaning up. It usually does very little.

```
1 void *
2 libgnunet_plugin_block_SERVICE_done (void *cls)
3 {
4     struct GNUNET_TRANSPORT_PluginFunctions *api = cls;
5
6     GNUNET_free (api);
7     return NULL;
8 }
```

9.3.5 Integration of the plugin with the build system

In order to compile the plugin, the `Makefile.am` file for the service `SERVICE` should contain a rule similar to this:

```
1 pluginindir = $(libdir)/gnunet
2
3 plugin_LTLIBRARIES = \
4 libgnunet_plugin_block_ext.la
5 libgnunet_plugin_block_ext_la_SOURCES = \
6 plugin_block_ext.c
7 libgnunet_plugin_block_ext_la_LIBADD = \
8 $(prefix)/lib/libgnunethello.la \
9 $(prefix)/lib/libgnunetblock.la \
10 $(prefix)/lib/libgnunetutil.la
11 libgnunet_plugin_block_ext_la_LDFLAGS = \
12 $(GN_PLUGIN_LDFLAGS)
13 libgnunet_plugin_block_ext_la_DEPENDENCIES = \
14 $(prefix)/lib/libgnunetblock.la
```

Exercise: Write a block plugin that accepts all queries and all replies but prints information about queries and replies when the respective validation hooks are called.

9.4 Monitoring the DHT

It is possible to monitor the functioning of the local DHT service. When monitoring the DHT, the service will alert the monitoring program of any events, both started locally or received for routing from another peer. There are three different types of events possible: a GET request, a PUT request or a response (a reply to a GET).

Since the different events have different associated data, the API gets 3 different callbacks (one for each message type) and optional type and key parameters, to allow for filtering of messages. When an event happens, the appropriate callback is called with all the information about the event.

```
1 static void
2 get_callback (void *cls,
3               enum GNUNET_DHT_RouteOption options,
4               enum GNUNET_BLOCK_Type type,
5               uint32_t hop_count,
6               uint32_t desired_replication_level,
7               unsigned int path_length,
8               const struct GNUNET_PeerIdentity *path,
9               const struct GNUNET_HashCode *key)
10 {
11 }
12
13
14 static void
15 get_resp_callback (void *cls,
16                   enum GNUNET_BLOCK_Type type,
17                   const struct GNUNET_PeerIdentity *get_path,
18                   unsigned int get_path_length,
19                   const struct GNUNET_PeerIdentity *put_path,
20                   unsigned int put_path_length,
21                   struct GNUNET_TIME_Absolute exp,
22                   const struct GNUNET_HashCode *key,
23                   const void *data,
24                   size_t size)
25 {
```

```

26 }
27
28
29 static void
30 put_callback (void *cls,
31               enum GNUNET_DHT_RouteOption options,
32               enum GNUNET_BLOCK_Type type,
33               uint32_t hop_count,
34               uint32_t desired_replication_level,
35               unsigned int path_length,
36               const struct GNUNET_PeerIdentity *path,
37               struct GNUNET_TIME_Absolute exp,
38               const struct GNUNET_HashCode *key,
39               const void *data,
40               size_t size)
41 {
42 }
43
44
45 monitor_handle = GNUNET_DHT_monitor_start (dht_handle,
46                                           block_type,
47                                           key,
48                                           &get_callback,
49                                           &get_resp_callback,
50                                           &put_callback,
51                                           cls);

```

10 Debugging with gnunet-arm

Even if services are managed by `gnunet-arm`, you can start them with `gdb` or `valgrind`. For example, you could add the following lines to your configuration file to start the DHT service in a `gdb` session in a fresh `xterm`:

```

[dht]
PREFIX=xterm -e gdb --args

```

Alternatively, you can stop a service that was started via ARM and run it manually:

```

1 $ gnunet-arm -k dht
2 $ gdb --args gnunet-service-dht -L DEBUG
3 $ valgrind gnunet-service-dht -L DEBUG

```

Assuming other services are well-written, they will automatically re-integrate the restarted service with the peer.

GNUnet provides a powerful logging mechanism providing log levels `ERROR`, `WARNING`, `INFO` and `DEBUG`. The current log level is configured using the `$GNUNET_FORCE_LOG` environmental variable. The `DEBUG` level is only available if `--enable-logging=verbose` was used when running `configure`. More details about logging can be found under <https://gnunet.org/logging>.

You should also probably enable the creation of core files, by setting `ulimit`, and echoing 1 into `/proc/sys/kernel/core_uses_pid`. Then you can investigate the core dumps with `gdb`, which is often the fastest method to find simple errors.

Exercise: Add a memory leak to your service and obtain a trace pointing to the leak using `valgrind` while running the service from `gnunet-service-arm`.