

**NAME**

guasi\_create, guasi\_free, guasi\_submit, guasi\_fetch, guasi\_req\_cancel, guasi\_req\_info, guasi\_req\_free  
 guasi\_\_open, guasi\_\_read, guasi\_\_write, guasi\_\_pread, guasi\_\_pwrite, guasi\_\_sendfile, guasi\_\_opendir,  
 guasi\_\_readdir, guasi\_\_stat, guasi\_\_fstat, guasi\_\_recv, guasi\_\_recvfrom, guasi\_\_send, guasi\_\_sendto,  
 guasi\_\_accept, guasi\_\_connect, guasi\_\_fsync

**SYNOPSIS****Core API**

```
#include <guasi.h>
```

```
guasi_t guasi_create(int min_threads, int max_threads, int max_priority);
void guasi_free(guasi_t hctx);
guasi_req_t guasi_submit(guasi_t hctx, void *priv, void *asid, int prio,
                        guasi_syscall_t proc, guasi_cleanup_t free, int nparams, ...);
int guasi_fetch(guasi_t hctx, guasi_req_t *reqs, int minreqs,
               int maxreqs, long timeo);
int guasi_req_cancel(guasi_req_t hreq);
int guasi_req_info(guasi_req_t hreq, struct guasi_reqinfo *rinf);
void guasi_req_free(guasi_req_t hreq);
```

**POSIX wrappers:**

```
#include <guasi_syscalls.h>
```

```
guasi_req_t guasi__open(guasi_t hctx, void *priv, void *asid, int prio,
                       char const *name, long flags, long mode);
guasi_req_t guasi__read(guasi_t hctx, void *priv, void *asid, int prio,
                       int fd, void *buf, size_t size);
guasi_req_t guasi__write(guasi_t hctx, void *priv, void *asid, int prio,
                        int fd, void const *buf, size_t size);
guasi_req_t guasi__pread(guasi_t hctx, void *priv, void *asid, int prio,
                        int fd, void *buf, size_t size, off_t off);
guasi_req_t guasi__pwrite(guasi_t hctx, void *priv, void *asid, int prio,
                        int fd, void const *buf, size_t size, off_t off);
guasi_req_t guasi__sendfile(guasi_t hctx, void *priv, void *asid, int prio,
                           int ofd, int ifd, off_t *off, size_t cnt);
guasi_req_t guasi__opendir(guasi_t hctx, void *priv, void *asid, int prio,
                          char const *name);
guasi_req_t guasi__readdir(guasi_t hctx, void *priv, void *asid, int prio,
                          DIR *dir);
guasi_req_t guasi__stat(guasi_t hctx, void *priv, void *asid, int prio,
                       char const *name, struct stat *sbuf);
guasi_req_t guasi__fstat(guasi_t hctx, void *priv, void *asid, int prio,
                       int fd, struct stat *sbuf);
guasi_req_t guasi__recv(guasi_t hctx, void *priv, void *asid, int prio,
                      int fd, void *buf, size_t size, int flags,
                      long timeo);
guasi_req_t guasi__recvfrom(guasi_t hctx, void *priv, void *asid, int prio,
                          int fd, void *buf, size_t size, int flags,
                          struct sockaddr *from, socklen_t *fromlen, long timeo);
guasi_req_t guasi__send(guasi_t hctx, void *priv, void *asid, int prio,
                      int fd, void const *buf, size_t size, int flags,
                      long timeo);
```

```

guasi_req_t guasi__sendto(guasi_t hctx, void *priv, void *asid, int prio,
                          int fd, void const *buf, size_t size, int flags,
                          struct sockaddr const *to, socklen_t tolen, long timeo);
guasi_req_t guasi__accept(guasi_t hctx, void *priv, void *asid, int prio,
                          int fd, struct sockaddr *addr, socklen_t *addrlen,
                          long timeo);
guasi_req_t guasi__connect(guasi_t hctx, void *priv, void *asid, int prio,
                           int fd, struct sockaddr const *addr, socklen_t addrlen,
                           long timeo);
guasi_req_t guasi__fsync(guasi_t hctx, void *priv, void *asid, int prio,
                         int fd);

```

## DESCRIPTION

The **guasi** library implements a thread based generic asynchronous execution engine, to be used to give otherwise synchronous calls an asynchronous behaviour. It can be used to wrap any synchronous call, so that it can be scheduled for execution, and whose result can be fetched at later time (hence not blocking the submitter thread). The **guasi** library can be used as complement to standard event retrieval interfaces like **poll(2)**, **select(2)** or **epoll(4)**. The **guasi** library is generic, by meaning that any (not only system calls) otherwise synchronous function can be made asynchronous. Functions submitted with **guasi\_submit(3)** will/might execute in parallel, so proper care must be taken to handle this behaviour. Care must be taken to select the proper stack size, according to what the asynchronous functions will require. The default **PTHREAD\_STACK\_MIN** value is more than enough when executing simple system call wrappers, but it may not be enough if the function submitted with **guasi\_submit(3)** uses a lot of local (stack) allocations.

## Structures and Types

The following structures are defined:

### **guasi\_param\_t**

The **guasi\_param\_t** type is the basic type used to return asynchronous operation results, and to pass parameters to asynchronous callbacks. Its size guarantee that a pointer can be safely passed through it.

### **guasi\_syscall\_t**

The **guasi\_syscall\_t** type is a function pointer type, that is used to pass asynchronous functions to **guasi\_submit**. An asynchronous function looks like:

```

guasi_param_t async_proc(void *priv, guasi_param_t const *params) {
    guasi_param_t result;

    ...
    return result;
}

```

### **guasi\_cleanup\_t**

The **guasi\_cleanup\_t** type is a function pointer type, that is used to pass to **guasi\_submit** to give the caller to perform a cleanup of the parameters passed to **guasi\_submit** itself. A cleanup function looks like:

```

void cleanup_proc(void *priv, guasi_param_t const *params) {
    ...
}

```

```
    }
```

### struct guasi\_reqinfo

```
struct guasi_reqinfo {
    void *priv;
    void *asid;
    guasi_param_t result;
    long error;
    int status;
};
```

The **struct guasi\_reqinfo** describes the status of a request submitted to a **guasi** handle. The **guasi\_req\_info** function can be used to retrieve such information from a request handle (**guasi\_req\_t**). The *priv* member is the same *priv* parameter that is passed to the **guasi\_submit** function when the request was submitted. The *asid* is the request identifier, passed to **guasi\_submit**. The *result* member is the result of the asynchronous operation, that is in turn, the value returned by the *proc* parameter of the **guasi\_submit** function. The *error* member is the value or the C library **errno** after the *proc* asynchronous function execution. The *status* is the current status of the request. Possible values for the *status* member are:

#### **GUASI\_STATUS\_PENDING**

The request has still to be dequeued.

#### **GUASI\_STATUS\_ACTIVE**

The request is executing.

#### **GUASI\_STATUS\_COMPLETE**

The request has completed.

#### **GUASI\_STATUS\_CANCELED**

The request has been canceled.

The *result* and *error* members are undefined until the request *status* reaches the **GUASI\_STATUS\_COMPLETE** value.

### Functions (Core API)

The following functions are defined:

```
guasi_t guasi_create(int min_threads, int max_threads, int max_priority);
```

The **guasi\_create** function creates a **guasi** handle to be used as a gateway for all the following **guasi** operations. The *min\_threads* parameter specifies the minimum number of threads to be used in the **guasi** thread pool, and the *max\_threads* parameter specifies the maximum number of threads (zero means unlimited). Since **guasi** requests can have different priorities, the *max\_priority* parameter specifies the number of priorities that the new **guasi** handle must support. The function returns the new **guasi** handle if succeeded, or **NULL** if failed.

```
void guasi_free(guasi_t hctx);
```

The **guasi\_free** function frees all the resources associated with the **guasi** handle passed in *hctx*.

```
guasi_req_t guasi_submit(guasi_t hctx, void *priv, void *asid, int prio,
guasi_syscall_t proc, guasi_cleanup_t free, int nparams, ...);
```

The **guasi\_submit** function submits a new request to the **guasi** handle passed in *hctx*. The *priv* parameter is an opaque value that is passed as is to the *proc* asynchronous function. The *free* function, if not **NULL**, will be called to give the caller the chance to cleanup the parameters. The *asid* is a cookie that identifies the request, and is returned by the **guasi\_req\_info** function, inside the **struct guasi\_reqinfo** structure. The priority of the request is passed in the *prio* parameter, that should range from 0 to (*max\_priority* - 1). The *nparams* parameter specifies the number of arguments that follows and that will be passed to the *proc* asynchronous function. The function returns a request handle if succeeded, or **NULL** if failed.

```
int guasi_fetch(guasi_t hctx, guasi_req_t *reqs, int minreqs,
int maxreqs, long timeo);
```

The **guasi\_fetch** function retrieves completed requests from the **guasi** handle passed in *hctx*. Not less than *minreqs*, and up to *maxreqs* requests are retrieved, and stored inside the *reqs* pointer. The *timeo* represent a maximum time (in milliseconds) to wait for some request to complete. If *timeo* is negative, the wait for completed requests will not have a time limit. The function returns the number of completed requests, or a negative number in case of error.

```
int guasi_req_cancel(guasi_req_t hreq);
```

The **guasi\_req\_cancel** function cancels a pending request specified in the *hreq* parameter. The function returns the current status of the request, or a negative number in case of error. A canceled request will be available at the next **guasi\_fetch()** call, with **guasi\_req\_info()** returning the proper status. If the reported status is **GUASI\_STATUS\_CANCELED**, it means that the request has been canceled before having the chance to be executed. If the reported status is **GUASI\_STATUS\_COMPLETE**, the caller should check the *error* member of the **struct guasi\_reqinfo** structure to see if the request has fully completed before we canceled, or if it has been interrupted while executing (in which case *error* should hold the **EINTR** value - at least for system call wrappers).

```
int guasi_req_info(guasi_req_t hreq, struct guasi_reqinfo *rinf);
```

The **guasi\_req\_info** function retrieves information about the request passed in the *hreq* parameter. The request information will be stored inside the *rinf* **struct guasi\_reqinfo** pointer. The function returns 0 in case of success, or a negative number in case of error.

```
void guasi_req_free(guasi_req_t hreq);
```

The **guasi\_req\_free** function frees all the resources associated with the request handle passed in *hreq*. A request can be freed only after it has been returned by the **guasi\_fetch()** function.

### Functions (POSIX wrappers)

All the POSIX wrappers functions implemented by **guasi**, map to the corresponding POSIX function in terms of parameters and result type. Some networking functions can have an extra *timeo* parameter, that can be used to specify a timeout (in milliseconds) for the operation.

## PERFORMANCE

One software that allows you to compare **guasi** performance with other kinds of asynchronous I/O providers, is **FIO**:

<http://freshmeat.net/projects/fio/>

Performance results of **guasi** against **libaio** using **FIO** are reported in this page:

<http://www.xmailserver.org/guasi-libaio-fio.html>

## NOTES

The following notes applies to the **guasi** library usage:

### Linux and NPTL

When using **guasi** on a Linux system using NPTL (thread library), the maximum number of threads is limited to about 350 on a 32 bit system with a 3G/1G split. This because NPTL reserves a 8MB VM area (default stack reservation) for each thread stack. Using **ulimit**(sh) it is possible to reduce the stack area to a lower value, hence allowing an higher number of threads to be created. The command is:

```
$ ulimit -s NN
```

Where *NN* is the number of KB to be reserved for the stack. Something around 64 should be enough for most applications.

## LICENSE

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2.1 of the License, or (at your option) any later version. A copy of the license is available at:

<http://www.gnu.org/copyleft/lesser.html>

## AUTHOR

Developed by Davide Libenzi <[davidel@xmailserver.org](mailto:davidel@xmailserver.org)>

## AVAILABILITY

The latest version of the **guasi** library can be found at:

<http://www.xmailserver.org/guasi-lib.html>

## BUGS

There are no known bugs. Bug reports and comments to Davide Libenzi <[davidel@xmailserver.org](mailto:davidel@xmailserver.org)>