



Sun Java™ System

Web Proxy Server 4 NSAPI Programmer's Guide

2005Q2

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054
U.S.A.

Part No: 819-0910-10

Copyright © 2005 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.sun.com/patents> and one or more additional patents or pending patent applications in the U.S. and in other countries.

THIS PRODUCT CONTAINS CONFIDENTIAL INFORMATION AND TRADE SECRETS OF SUN MICROSYSTEMS, INC. USE, DISCLOSURE OR REPRODUCTION IS PROHIBITED WITHOUT THE PRIOR EXPRESS WRITTEN PERMISSION OF SUN MICROSYSTEMS, INC.

U.S. Government Rights - Commercial software. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

This distribution may include materials developed by third parties.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, Java, Solaris, JDK, Java Naming and Directory Interface, JavaMail, JavaHelp, J2SE, iPlanet, the Duke logo, the Java Coffee Cup logo, the Solaris logo, the SunTone Certified logo and the Sun ONE logo are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon architecture developed by Sun Microsystems, Inc.

Legato and the Legato logo are registered trademarks, and Legato NetWorker, are trademarks or registered trademarks of Legato Systems, Inc. The Netscape Communications Corp logo is a trademark or registered trademark of Netscape Communications Corporation.

The OPEN LOOK and Sun(TM) Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

Products covered by and information contained in this service manual are controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical biological weapons or nuclear maritime end uses or end users, whether direct or indirect, are strictly prohibited. Export or reexport to countries subject to U.S. embargo or to entities identified on U.S. export exclusion lists, including, but not limited to, the denied persons and specially designated nationals lists is strictly prohibited.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright © 2005 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, Etats-Unis. Tous droits réservés.

Sun Microsystems, Inc. détient les droits de propriété intellectuelle relatifs à la technologie incorporée dans le produit qui est décrit dans ce document. En particulier, et ce sans limitation, ces droits de propriété intellectuelle peuvent inclure un ou plusieurs des brevets américains listés à l'adresse <http://www.sun.com/patents> et un ou des brevets supplémentaires ou des applications de brevet en attente aux Etats - Unis et dans les autres pays.

CE PRODUIT CONTIENT DES INFORMATIONS CONFIDENTIELLES ET DES SECRETS COMMERCIAUX DE SUN MICROSYSTEMS, INC. SON UTILISATION, SA DIVULGATION ET SA REPRODUCTION SONT INTERDITES SANS L'AUTORISATION EXPRESSE, ECRITE ET PREALABLE DE SUN MICROSYSTEMS, INC.

Cette distribution peut comprendre des composants développés par des tierces parties.

Des parties de ce produit peuvent être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, Java, Solaris, JDK, Java Naming and Directory Interface, JavaMail, JavaHelp, J2SE, iPlanet, le logo Duke, le logo Java Coffee Cup, le logo Solaris, le logo SunTone Certified et le logo Sun[tm] ONE sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays.

Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

Legato, le logo Legato, et Legato NetWorker sont des marques de fabrique ou des marques déposées de Legato Systems, Inc. Le logo Netscape Communications Corp est une marque de fabrique ou une marque déposée de Netscape Communications Corporation.

L'interface d'utilisation graphique OPEN LOOK et Sun(TM) a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui, en outre, se conforment aux licences écrites de Sun.

Les produits qui font l'objet de ce manuel d'entretien et les informations qu'il contient sont regis par la législation américaine en matière de contrôle des exportations et peuvent être soumis au droit d'autres pays dans le domaine des exportations et importations. Les utilisations finales, ou utilisateurs finaux, pour des armes nucléaires, des missiles, des armes biologiques et chimiques ou du nucléaire maritime, directement ou indirectement, sont strictement interdites. Les exportations ou réexportations vers des pays sous embargo des Etats-Unis, ou vers des entités figurant sur les listes d'exclusion d'exportation américaines, y compris, mais de manière non exclusive, la liste de personnes qui font objet d'un ordre de ne pas participer, d'une façon directe ou indirecte, aux exportations des produits ou des services qui sont régi par la législation américaine en matière de contrôle des exportations et la liste de ressortissants spécifiquement désignés, sont rigoureusement interdites.

LA DOCUMENTATION EST FOURNIE "EN L'ETAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFACON.

Contents

About This Guide	9
Who Should Use This Guide	9
How This Guide Is Organized	10
Documentation Conventions	10
Using the Documentation	11
Contacting Sun Technical Support	12
Third-Party Web Site References	12
Feedback	12
Chapter 1 NSAPI Function Reference	13
NSAPI Functions (in Alphabetical Order)	13
C	14
cache_digest	14
cache_filename	14
cache_fn_to_dig	15
CALLOC	15
ce_free	16
ce_lookup	16
cif_write_entry	17
cinfo_find	18
condvar_init	18
condvar_notify	19
condvar_terminate	19
condvar_wait	20
crit_enter	20
crit_exit	21
crit_init	21

crit_terminate	22
D	22
daemon_atrestart	22
dns_set_hostent	23
F	24
fc_close	24
fc_open	24
filebuf_buf2sd	25
filebuf_close	26
filebuf_getc	26
filebuf_open	27
filebuf_open_nostat	28
filter_create	28
filter_find	30
filter_insert	30
filter_layer	31
filter_name	31
filter_remove	32
flush	32
FREE	33
fs_blk_size	34
fs_blks_avail	34
func_exec	35
func_find	35
func_insert	36
I	37
insert	37
L	38
log_error	38
M	39
magnus_atrestart	39
MALLOC	39
N	40
net_flush	40
net_ip2host	41
net_read	41
net_sendfile	42
net_write	44
netbuf_buf2sd	44
netbuf_close	45
netbuf_getc	45
netbuf_grab	46
netbuf_open	46

nsapi_module_init	47
NSAPI_RUNTIME_VERSION	47
NSAPI_VERSION	48
P	49
param_create	49
param_free	49
pblock_copy	50
pblock_create	50
pblock_dup	51
pblock_find	51
pblock_findlong	52
pblock_findval	52
pblock_free	53
pblock_ninsert	54
pblock_nninsert	54
pblock_nvinsert	55
pblock_pb2env	55
pblock_pblock2str	56
pblock_pinsert	56
pblock_remove	57
pblock_replace_name	57
pblock_str2pblock	58
PERM_CALLOC	59
PERM_FREE	59
PERM_MALLOC	60
PERM_REALLOC	61
PERM_STRDUP	61
prepare_nsapi_thread	62
protocol_dump822	63
protocol_finish_request	63
protocol_handle_session	64
protocol_parse_request	64
protocol_scan_headers	65
protocol_set_finfo	65
protocol_start_response	66
protocol_status	67
protocol_uri2url	68
protocol_uri2url_dynamic	68
R	69
read	69
REALLOC	70
remove	71
request_create	71

request_free	72
request_header	72
S	73
sem_grab	73
sem_init	74
sem_release	74
sem_terminate	75
sem_tgrab	75
sendfile	76
session_create	76
session_dns	77
session_free	78
session_maxdns	78
shexp_casecmp	79
shexp_cmp	79
shexp_match	80
shexp_valid	81
shmem_alloc	81
shmem_free	82
STRDUP	83
system_errmsg	83
system_fclose	84
system_flock	85
system_fopenRO	85
system_fopenRW	86
system_fopenWA	86
system_fread	87
system_fwrite	88
system_fwrite_atomic	88
system_gmtime	89
system_localtime	90
system_lseek	90
system_rename	91
system_ulock	91
system_unix2local	92
systhread_attach	92
systhread_current	93
systhread_getdata	93
systhread_init	94
systhread_newkey	94
systhread_setdata	95
systhread_sleep	95
systhread_start	96

systhread_terminate	96
systhread_timerset	97
U	98
USE_NSAPI_VERSION	98
util_can_exec	99
util_chdir2path	100
util_cookie_find	100
util_does_process_exist	101
util_env_create	101
util_env_find	102
util_env_free	102
util_env_replace	103
util_env_str	103
util_get_current_gmt	104
util_get_int_from_aux_file	104
util_get_int_from_file	105
util_get_long_from_aux_file	105
util_get_long_from_file	106
util_get_string_from_aux_file	106
util_get_string_from_file	107
util_getline	108
util_hostname	108
util_is_mozilla	109
util_is_url	109
util_itoa	110
util_later_than	110
util_make_filename	111
util_make_gmt	111
util_make_local	112
util_move_dir	112
util_move_file	113
util_parse_http_time	113
util_put_int_to_file	114
util_put_long_to_file	114
util_put_string_to_aux_file	115
util_put_string_to_file	115
util_sect_id	116
util_sh_escape	116
util_snprintf	117
util_sprintf	118
util_strcasecmp	118
util_strftime	119
util_strncasecmp	120

util_uri_check	120
util_uri_escape	121
util_uri_is_evil	121
util_uri_parse	122
util_uri_unescape	122
util_url_cmp	123
util_url_fix_host name	123
util_url_has_FQDN	124
util_vsprintf	124
util_vsprintf	125
W	126
write	126
writev	127
Chapter 2 Data Structure Reference	129
Privatization of Some Data Structures	130
Session	131
pblock	131
pb_entry	132
pb_param	132
Session->client	132
Request	133
stat	133
shmem_s	134
cinfo	134
sendfiledata	135
Filter	135
FilterContext	136
FilterLayer	136
FilterMethods	136
The CacheEntry Data Structure	137
The CacheState Data Structure	138
The ConnectMode Data Structure	139
Chapter 3 Time Formats	141
Appendix A Alphabetical List of NSAPI Functions and Macros	143
Index	153

About This Guide

This guide describes how to configure and administer the Sun Java™ System Web Proxy Server 4, formerly known as Sun ONE Web Proxy Server and iPlanet™ Web Proxy Server (and hereafter referred to as Sun™ Java System Web Proxy Server or just Proxy Server).

This guide provides a reference of the NSAPI functions you can use to define new plugins.

This preface contains information about the following topics:

- [Who Should Use This Guide](#)
- [How This Guide Is Organized](#)
- [Documentation Conventions](#)
- [Using the Documentation](#)
- [Contacting Sun Technical Support](#)
- [Third-Party Web Site References](#)
- [Feedback](#)

Who Should Use This Guide

The intended audience for this guide is the person who develops, assembles, and deploys NSAPI plugins in a corporate enterprise. This guide assumes you are familiar with the following topics:

- HTTP
- HTML

- NSAPI
- C programming
- Software development processes, including debugging and source code control

How This Guide Is Organized

The guide is divided into parts, each of which addresses specific areas and tasks. The following table lists the parts of the guide and their contents .

Table 1 Guide Organization

Chapter	Description
Chapter 1, “NSAPI Function Reference”	This chapter presents a reference of the NSAPI functions. You use NSAPI functions to define SAFs.
Chapter 2, “Data Structure Reference”	This chapter discusses some of the commonly used NSAPI data structures.
Chapter 3, “Time Formats”	This chapter lists time formats.
Appendix A, “Alphabetical List of NSAPI Functions and Macros”	This appendix provides an alphabetical list of NSAPI functions and macros.

Documentation Conventions

The following table lists the documentation conventions used in this guide.

Table 2 Documentation Conventions

Element	Usage
File and directory paths	Given in UNIX® format, with forward slashes separating directory names
Installation root directories	Indicated as <i>install_dir</i> .
<i>italic</i> text	Titles, emphasis, terms
monospace text	Code examples, file names, path names, command names, programming language keywords, properties
<i>italic monospace</i> text	Variable path names, environment variables in paths

Using the Documentation

At the time of final release, Proxy Server documentation will be available in PDF and HTML formats at:

<http://docs.sun.com/db/prod/s1.webproxys#hic>

The following table lists the tasks and concepts described in guide..

Table 3 Proxy Server Documentation

For Information About	See
Late-breaking information about the software and documentation	<i>Release Notes</i>
Performing installation and migration tasks: <ul style="list-style-type: none"> Supported platforms and environments Installing Sun Java System Web Proxy Server Migrating from version 3.6 to version 4 	<i>Installation and Migration Guide</i>
Performing administration and management tasks: <ul style="list-style-type: none"> Using the Administration and command-line interfaces Configuring server preferences Managing users and groups Monitoring and logging server activity Using certificates and public key cryptography to secure the server Controlling server access Proxying and routing URLs Caching Filtering content Using a reverse proxy Using SOCKS Tuning the Proxy Server to optimize performance 	<i>Administrator's Guide</i> (and the online Help included with the product)
Creating custom Netscape Server Application Programmer's Interface (NSAPI) plugins	<i>NSAPI Programmer's Guide</i>
Editing configuration files	<i>Administrator's Configuration File Reference</i>

Contacting Sun Technical Support

For technical questions about this product not answered in the product documentation, go to:

<http://www.sun.com/service/contacting>

Third-Party Web Site References

Sun is not responsible for the availability of third-party Web sites mentioned in this document. Sun does not endorse and is not responsible or liable for any content, advertising, products, or other materials that are available on or through such sites or resources. Sun will not be responsible or liable for any actual or alleged damage or loss caused or alleged to be caused by or in connection with use of or reliance on any such content, goods, or services that are available on or through such sites or resources.

Feedback

Sun is interested in improving its documentation and welcomes your comments and suggestions. To share your comments, go to <http://docs.sun.com> and click "Send comments." Be sure to provide the document title and part number in the online form.

NSAPI Function Reference

This chapter lists all of the public C functions and macros of the Netscape Server Applications Programming Interface (NSAPI) in alphabetic order. These are the functions you use when writing your own Server Application Functions (SAFs).

For information about the predefined SAFs used in `obj.conf`, see the Sun Java System Web Proxy Server 4 *Administrator's Configuration File Reference*.

Each function provides the name, syntax, parameters, return value, a description of what the function does, and sometimes an example of its use and a list of related functions.

For more information on data structures, see [Chapter 2, "Data Structure Reference,"](#) and also look in the `nsapi.h` header file in the `include` directory in the build for Sun Java System Web Proxy Server 4.

NSAPI Functions (in Alphabetical Order)

For an alphabetical list of function names, see [Appendix A, "Alphabetical List of NSAPI Functions and Macros."](#)

C D F I L M N P R S U W

cache_digest

The `cache_digest` function calculates the MD5 signature of a specified URL and stores the signature in a `digest` variable.

Syntax

```
#include <libproxy/cache.h>
void cache_digest(char *url, unsigned char digest[16]);
```

Returns

void

Parameters

char **url* is a string containing the cache filename of a URL.

name **digest* is an array to store the MD5 signature of the URL.

See also

[cache_fn_to_dig](#)

cache_filename

The `cache_filename` function returns the cache filename for a given URL, specified by MD5 signature.

Syntax

```
#include <libproxy/cutil.h>
char *cache_filename(unsigned char digest[16]);
```

Returns

A new string containing the cache filename.

Parameters

char **digest* is an array containing the MD5 signature of a URL.

See also

[cache_fn_to_dig](#)

cache_fn_to_dig

The `cache_fn_to_dig` function converts a cache filename of a URL into a partial MD5 digest.

Syntax

```
#include <libproxy/cutil.h>
void *cache_fn_to_dig(char *name, unsigned char digest[16]);
```

Returns

void

Parameters

char **name* is a string containing the cache filename of a URL.

name **digest* is an array to receive first 8 bits of the signature of the URL.

CALLOC

The `CALLOC` macro is a platform-independent substitute for the C library routine `calloc`. It allocates `num*size` bytes from the request's memory pool. If pooled memory has been disabled in the configuration file (with the `pool-init` built-in SAF), `PERM_CALLOC` and `CALLOC` both obtain their memory from the system heap.

Syntax

```
void *CALLOC(int size)
```

Returns

A void pointer to a block of memory.

Parameters

int *size* is the size in bytes of each element.

Example

```
char *name;
name = (char *) CALLOC(100);
```

See Also

[FREE](#), [REALLOC](#), [STRDUP](#), [PERM_MALLOC](#), [PERM_FREE](#), [PERM_REALLOC](#), [PERM_STRDUP](#)

ce_free

The `ce_free` function releases memory allocated by the `ce_lookup` function.

Syntax

```
#include <libproxy/cache.h>
void ce_free(CacheEntry *ce);
```

Returns

void

Parameters

CacheEntry **ce* is a cache entry structure to be destroyed.

See also

[ce_lookup](#)

ce_lookup

The `ce_lookup` cache entry lookup function looks up a cache entry for a specified URL.

Syntax

```
#include <libproxy/cache.h>
CacheEntry *ce_lookup(Session *sn, Request *rq, char *url, time_t ims_c);
```

Returns

- NULL if caching is not enabled
- A newly allocated CacheEntry structure, whether or not a copy existed in the cache. Within that structure, the `ce->state` field reports about the existence:

CACHE_NO signals that the document is not and will not be cached; other fields in the cache structure may be NULL

CACHE_CREATE signals that the cache file doesn't exist but may be created once the remote server is contacted. However, during the retrieval it may turn out that the document is not cacheable.

CACHE_REFRESH signals that the cache file exists, but it needs to be refreshed (an up-to-date check must be made) before it's used; note that the data may still be up-to-date, but the remote server needs to be contacted to find that out. If not, the cache file will be replaced with the new document version sent by the remote origin server.

CACHE_RETURN_FROM_CACHE signals that the cache file exists and is up-to-date based on the configuration and current parameters controlling what is considered fresh.

CACHE_RETURN_ERROR is a signal that happens only if the proxy is set to no-network mode (connect-Modenese), and the document does not exist in the cache.

Parameters

Session **sn* identifies the Session structure.

Request **rq* identifies the Request structure.

char **url* contains the name of the URL for which the cache is being sought.

time-out *misc*. is the if-modified-since time.

See also

[ce_free](#)

cif_write_entry

The `cif_write_entry` function writes a CIF entry for a specified `CacheEntry` structure. The CIF entry is stored in the cache file itself.

Syntax

```
#include <libproxy/cif.h>
int cif_write_entry(CacheEntry *ce,int new_cachefile)
```

Returns

- nonzero if the write was successful
- 0 if the write was unsuccessful

Parameters

`CacheEntry *ce` is a cache entry structure to be written to the `.cif` file.

int `new_cachefile` The values are 1 or 0.

1 if it is a new cache file;

0 if the file exists and the CIF entry is to be modified

cinfo_find

The `cinfo_find()` function uses the MIME types information to find the type, encoding, and/or language based on the extension(s) of the Universal Resource Identifier (URI) or local file name. Use this information to send headers (`rq->srvhdrs`) to the client indicating the `content-type`, `content-encoding`, and `content-language` of the data it will be receiving from the server.

The name used is everything after the last slash (/) or the whole string if no slash is found. File name extensions are not case-sensitive. The name may contain multiple extensions separated by period (.) to indicate type, encoding, or language. For example, the URI `a/b/filename.jp.txt.zip` could represent a Japanese language, text/plain type, zip encoded file.

Syntax

```
cinfo *cinfo_find(char *uri);
```

Returns

A pointer to a newly allocated `cinfo` structure if content info was found, or NULL if no content was found.

The `cinfo` structure that is allocated and returned contains pointers to the `content-type`, `content-encoding`, and `content-language`, if found. Each is a pointer into static data in the types database, or NULL if not found. Do not free these pointers. You should free the `cinfo` structure when you are done using it.

Parameters

`char *uri` is a Universal Resource Identifier (URI) or local file name. Multiple file name extensions should be separated by periods (.).

condvar_init

The `condvar_init` function is a critical-section function that initializes and returns a new condition variable associated with a specified critical-section variable. You can use the condition variable to prevent interference between two threads of execution.

Syntax

```
CONDVAR condvar_init(CRITICAL id);
```

Returns

A newly allocated condition variable (`CONDVAR`).

Parameters

`CRITICAL id` is a critical-section variable.

See Also

[condvar_notify](#), [condvar_terminate](#), [condvar_wait](#), [crit_init](#), [crit_enter](#), [crit_exit](#), [crit_terminate](#)

condvar_notify

The `condvar_notify` function is a critical-section function that awakens any threads that are blocked on the given critical-section variable. Use this function to awaken threads of execution of a given critical section. First, use `crit_enter` to gain ownership of the critical section. Then use the returned critical-section variable to call `condvar_notify` to awaken the threads. Finally, when `condvar_notify` returns, call `crit_exit` to surrender ownership of the critical section.

Syntax

```
void condvar_notify(CONDVAR cv);
```

Returns

void

Parameters

`CONDVAR cv` is a condition variable.

See Also

[condvar_init](#), [condvar_terminate](#), [condvar_wait](#), [crit_init](#), [crit_enter](#), [crit_exit](#), [crit_terminate](#)

condvar_terminate

The `condvar_terminate` function is a critical-section function that frees a condition variable. Use this function to free a previously allocated condition variable.

Warning

Terminating a condition variable that is in use can lead to unpredictable results.

Syntax

```
void condvar_terminate(CONDVAR cv);
```

Returns

void

Parameters

CONDVAR *cv* is a condition variable.

See Also

[condvar_init](#), [condvar_notify](#), [condvar_wait](#), [crit_init](#), [crit_enter](#), [crit_exit](#), [crit_terminate](#)

condvar_wait

The `condvar_wait` function is a critical-section function that blocks on a given condition variable. Use this function to wait for a critical section (specified by a condition variable argument) to become available. The calling thread is blocked until another thread calls `condvar_notify` with the same condition variable argument. The caller must have entered the critical section associated with this condition variable before calling `condvar_wait`.

Syntax

```
void condvar_wait(CONDVAR cv);
```

Returns

void

Parameters

CONDVAR *cv* is a condition variable.

See Also

[condvar_init](#), [condvar_terminate](#), [condvar_notify](#), [crit_init](#), [crit_enter](#), [crit_exit](#), [crit_terminate](#)

crit_enter

The `crit_enter` function is a critical-section function that attempts to enter a critical section. Use this function to gain ownership of a critical section. If another thread already owns the section, the calling thread is blocked until the first thread surrenders ownership by calling `crit_exit`.

Syntax

```
void crit_enter(CRITICAL crvar);
```

Returns

void

Parameters

`CRITICAL crvar` is a critical-section variable.

See Also

[crit_init](#), [crit_exit](#), [crit_terminate](#)

crit_exit

The `crit_exit` function is a critical-section function that surrenders ownership of a critical section. Use this function to surrender ownership of a critical section. If another thread is blocked waiting for the section, the block will be removed and the waiting thread will be given ownership of the section.

Syntax

```
void crit_exit(CRITICAL crvar);
```

Returns

void

Parameters

`CRITICAL crvar` is a critical-section variable.

See Also

[crit_init](#), [crit_enter](#), [crit_terminate](#)

crit_init

The `crit_init` function is a critical-section function that creates and returns a new critical-section variable (a variable of type `CRITICAL`). Use this function to obtain a new instance of a variable of type `CRITICAL` (a critical-section variable) to be used in managing the prevention of interference between two threads of execution. At the time of its creation, no thread owns the critical section.

Warning

Threads must not own or be waiting for the critical section when `crit_terminate` is called.

Syntax

```
CRITICAL crit_init(void);
```

Returns

A newly allocated critical-section variable (`CRITICAL`).

Parameters

none

See Also

[crit_enter](#), [crit_exit](#), [crit_terminate](#)

crit_terminate

The `crit_terminate` function is a critical-section function that removes a previously allocated critical-section variable (a variable of type `CRITICAL`). Use this function to release a critical-section variable previously obtained by a call to `crit_init`.

Syntax

```
void crit_terminate(CRITICAL crvar);
```

Returns

void

Parameters

`CRITICAL crvar` is a critical-section variable.

See Also

[crit_init](#), [crit_enter](#), [crit_exit](#)

daemon_atrestart

The `daemon_atrestart` function lets you register a callback function named by `fn` to be used when the server terminates. Use this function when you need a callback function to deallocate resources allocated by an initialization function. The `daemon_atrestart` function is a generalization of the `magnus_atrestart` function.

The `magnus.conf` directives `TerminateTimeout` and `ChildRestartCallback` also affect the callback of NSAPI functions.

Syntax

```
void daemon_atrestart(void (*fn)(void *), void *data);
```

Returns

void

Parameters

void (* fn) (void *) is the callback function.

void *data is the parameter passed to the callback function when the server is restarted.

Example

```
/* Register the log_close function, passing it NULL */
/* to close *a log file when the server is */
/* restarted or shutdown. */
daemon_atrestart(log_close, NULL);
NSAPI_PUBLIC void log_close(void *parameter)
{
    system_fclose(global_logfd);
}
```

dns_set_hostent

The `dns_set_hostent` function sets the DNS host entry information in the request. If this is set, the proxy won't try to resolve host information by itself, but instead it will just use this host information which was already resolved within custom DNS resolution SAF.

Syntax

```
int dns_set_hostent(struct hostent *hostent, Session *sn, Request *rq);
```

Returns

REQ_PROCEED on success or REQ_ABORTED on error.

Parameters

struct hostent *hostent is a pointer to the host entry structure.

Session *sn is the Session

Request *rq is the Request

Example

```
int my_dns_func(pblock *pb, Session *sn, Request *rq)
{
    char *host = pblock_findval("dns-host", rq->vars);
    struct hostent *hostent;
    hostent = gethostbyname(host); // replace with custom DNS
    implementation
    dns_set_hostent(hostent, sn, rq);
    return REQ_PROCEED;
}
```

fc_close

The `fc_close` function closes a file opened using `fc_open`. This function should only be called with files opened using `fc_open`.

Syntax

```
void fc_close(PRFileDesc *fd, FcHdl *hdl);
```

Returns

void

Parameters

`PRFileDesc *fd` is a valid pointer returned from a prior call to `fc_open`.

`FcHdl *hdl` is a valid pointer to a structure of type `FcHdl`. This pointer must have been initialized by a prior call to `fc_open`.

fc_open

The `fc_open` function returns a pointer to `PRFileDesc` that refers to an open file (`fileName`). The `fileName` must be the full path name of an existing file. The file is opened in read mode only. The application calling this function should not modify the currency of the file pointed to by the `PRFileDesc *` unless the `DUP_FILE_DESC` is

also passed to this function. In other words, the application (at minimum) should not issue a read operation based on this pointer that would modify the currency for the `PRFileDesc *`. If such a read operation is required (that may change the currency for the `PRFileDesc *`), then the application should call this function with the argument `DUP_FILE_DESC`.

On a successful call to this function, a valid pointer to `PRFileDesc` is returned and the handle `'FcHdl'` is properly initialized. The size information for the file is stored in the `'fileSize'` member of the handle.

Syntax

```
PRFileDesc *fc_open(const char *fileName, FcHdl *hDl, PRUint32 flags,
Session *sn, Request *rq);
```

Returns

Pointer to `PRFileDesc`, or `NULL` on failure.

Parameters

`const char *fileName` is the full path name of the file to be opened.

`FcHdl*hDl` is a valid pointer to a structure of type `FcHdl`.

`PRUint32 flags` can be `0` or `DUP_FILE_DESC`.

`Session *sn` is a pointer to the session.

`Request *rq` is a pointer to the request.

filebuf_buf2sd

The `filebuf_buf2sd` function sends a file buffer to a socket (descriptor) and returns the number of bytes sent.

Use this function to send the contents of an entire file to the client.

Syntax

```
int filebuf_buf2sd(filebuf *buf, SYS_NETFD sd);
```

Returns

The number of bytes sent to the socket if successful, or the constant `IO_ERROR` if the file buffer could not be sent.

Parameters

`filebuf *buf` is the file buffer that must already have been opened.

`SYS_NETFD sd` is the platform-independent socket descriptor. Normally this will be obtained from the `csd` (client socket descriptor) field of the `sn` (session) structure.

Example

```
if (filebuf_buf2sd(buf, sn->csd) == IO_ERROR)
    return(REQ_EXIT);
```

See Also

[filebuf_close](#), [filebuf_open](#), [filebuf_open_nostat](#), [filebuf_getc](#)

filebuf_close

The `filebuf_close` function deallocates a file buffer and closes its associated file.

Generally, use `filebuf_open` first to open a file buffer, and then `filebuf_getc` to access the information in the file. After you have finished using the file buffer, use `filebuf_close` to close it.

Syntax

```
void filebuf_close(filebuf *buf);
```

Returns

void

Parameters

`filebuf *buf` is the file buffer previously opened with `filebuf_open`.

Example

```
filebuf_close(buf);
```

See Also

[filebuf_open](#), [filebuf_open_nostat](#), [filebuf_buf2sd](#), [filebuf_getc](#)

filebuf_getc

The `filebuf_getc` function retrieves a character from the current file position and returns it as an integer. It then increments the current file position.

Use `filebuf_getc` to sequentially read characters from a buffered file.

Syntax

```
filebuf_getc(filebuf b);
```

Returns

An integer containing the character retrieved, or the constant `IO_EOF` or `IO_ERROR` upon an end of file or error.

Parameters

`filebuf b` is the name of the file buffer.

See Also

[filebuf_close](#), [filebuf_buf2sd](#), [filebuf_open](#), [filter_create](#)

filebuf_open

The `filebuf_open` function opens a new file buffer for a previously opened file. It returns a new buffer structure. Buffered files provide more efficient file access by guaranteeing the use of buffered file I/O in environments where it is not supported by the operating system.

Syntax

```
filebuf *filebuf_open(SYS_FILE fd, int sz);
```

Returns

A pointer to a new buffer structure to hold the data if successful, or `NULL` if no buffer could be opened.

Parameters

`SYS_FILE fd` is the platform-independent file descriptor of the file which has already been opened.

`int sz` is the size, in bytes, to be used for the buffer.

Example

```
filebuf *buf = filebuf_open(fd, FILE_BUFFER_SIZE);
if (!buf) {
    system_fclose(fd);
}
```

See Also

[filebuf_getc](#), [filebuf_buf2sd](#), [filebuf_close](#), [filebuf_open_nostat](#)

filebuf_open_nostat

The `filebuf_open_nostat` function opens a new file buffer for a previously opened file. It returns a new buffer structure. Buffered files provide more efficient file access by guaranteeing the use of buffered file I/O in environments where it is not supported by the operating system.

This function is the same `filebuf_open`, but is more efficient, since it does not need to call the `request_stat_path` function. It requires that the stat information be passed in.

Syntax

```
filebuf* filebuf_open_nostat(SYS_FILE fd, int sz,
    struct stat *finfo);
```

Returns

A pointer to a new buffer structure to hold the data if successful, or NULL if no buffer could be opened.

Parameters

`SYS_FILE fd` is the platform-independent file descriptor of the file that has already been opened.

`int sz` is the size, in bytes, to be used for the buffer.

`struct stat *finfo` is the file information of the file. Before calling the `filebuf_open_nostat` function, you must call the `request_stat_path` function to retrieve the file information.

Example

```
filebuf *buf = filebuf_open_nostat(fd, FILE_BUFFER_SIZE, &finfo);
if (!buf) {
    system_fclose(fd);
}
```

See Also

[filebuf_close](#), [filebuf_open](#), [filebuf_getc](#), [filebuf_buf2sd](#)

filter_create

The `filter_create` function defines a new filter.

The name parameter specifies a unique name for the filter. If a filter with the specified name already exists, it will be replaced.

Names beginning with `magnus-` or `server-` are reserved by the server.

The `order` parameter indicates the position of the filter in the filter stack by specifying what class of functionality the filter implements.

The following table describes parameters allowed order constants and their associated meanings for the `filter_create` function. The left column lists the name of the constant, the middle column describes the functionality the filter implements, and the right column lists the position the filter occupies in the filter stack.

Table 1-1 filter-create constants

Constant	Functionality Filter Implements	Position in Filter Stack
<code>FILTER_CONTENT_TRANSLATION</code>	Translates content from one form to another (for example, XSLT)	Top
<code>FILTER_CONTENT_CODING</code>	Encodes content (for example, HTTP gzip compression)	Middle
<code>FILTER_TRANSFER_CODING</code>	Encodes entity bodies for transmission (for example, HTTP chunking)	Bottom

The `methods` parameter specifies a pointer to a `FilterMethods` structure. Before calling `filter_create`, you must first initialize the `FilterMethods` structure using the `FILTER_METHODS_INITIALIZER` macro, and then assign function pointers to the individual `FilterMethods` members (for example, `insert`, `read`, `write`, and so on) that correspond to the filter methods the filter will support.

`filter_create` returns `const Filter *`, a pointer to an opaque representation of the filter. This value may be passed to `filter_insert` to insert the filter in a particular filter stack.

Syntax

```
const Filter *filter_create(const char *name, int order, const
FilterMethods *methods);
```

Returns

The `const Filter *` that identifies the filter or `NULL` if an error occurred.

Parameters

`const char *name` is the name of the filter.

`int order` is one of the order constants above.

const FilterMethods *methods contains pointers to the filter methods the filter supports.

Example

```
FilterMethods methods = FILTER_METHODS_INITIALIZER;
const Filter *filter;
/* This filter will only support the "read" filter method */
methods.read = my_input_filter_read;
/* Create the filter */
filter = filter_create("my-input-filter", FILTER_CONTENT_TRANSLATION,
&methods);
```

filter_find

The `filter_find` function finds the filter with the specified name.

Syntax

```
const Filter *filter_find(const char *name);
```

Returns

The const `Filter *` that identifies the filter, or NULL if the specified filter does not exist.

Parameters

const char *name is the name of the filter of interest.

filter_insert

The `filter_insert` function inserts a filter into a filter stack, creating a new filter layer and installing the filter at that layer. The filter layer's position in the stack is determined by the order value specified when `filter_create` was called, and any explicit ordering configured by `init-filter-order`. If a filter layer with the same order value already exists in the stack, the new layer is inserted above that layer.

Parameters may be passed to the filter using the `pb` and `data` parameters. The semantics of the `data` parameter are defined by individual filters. However, all filters must be able to handle a `data` parameter of NULL.

When possible, plugin developers should avoid calling `filter_insert` directly, and instead use the `insert-filter` SAF (applicable in `Input-class` directives).

Syntax

```
int filter_insert(SYS_NETFD sd, pblock *pb, Session *sn, Request *rq, void
*data, const Filter *filter);
```

Returns

Returns `REQ_PROCEED` if the specified filter was inserted successfully, or `REQ_NOACTION` if the specified filter was not inserted because it was not required. Any other return value indicates an error.

Parameters

`SYS_NETFD sd` is `NULL` (reserved for future use).

`pblock *pb` is a set of parameters to pass to the specified filter's init method.

`Session *sn` is the Session.

`Request *rq` is the Request.

`void *data` is filter-defined private data.

`const Filter *filter` is the filter to insert.

filter_layer

The `filter_layer` function returns the layer in a filter stack that corresponds to the specified filter.

Syntax

```
FilterLayer *filter_layer(SYS_NETFD sd, const Filter *filter);
```

Returns

The topmost `FilterLayer *` associated with the specified filter, or `NULL` if the specified filter is not part of the specified filter stack.

Parameters

`SYS_NETFD sd` is the filter stack to inspect.

`const Filter *filter` is the filter of interest.

filter_name

The `filter_name` function returns the name of the specified filter. The caller should not free the returned string.

Syntax

```
const char *filter_name(const Filter *filter);
```

Returns

The name of the specified filter, or NULL if an error occurred.

Parameters

const Filter *filter is the filter of interest.

filter_remove

The `filter_remove` function removes the specified filter from the specified filter stack, destroying a filter layer. If the specified filter was inserted into the filter stack multiple times, only that filter's topmost filter layer is destroyed.

When possible, plugin developers should avoid calling `filter_remove` directly, and instead use the remove-filter SAF (applicable in Input-, Output-, Service-, and Error-class directives).

Syntax

```
int filter_remove(SYS_NETFD sd, const Filter *filter);
```

Returns

Returns `REQ_PROCEED` if the specified filter was removed successfully or `REQ_NOACTION` if the specified filter was not part of the filter stack. Any other return value indicates an error.

Parameters

`SYS_NETFD sd` is the filter stack, `sn->csd`.

const Filter *filter is the filter to remove.

flush

The `flush` filter method is called when buffered data should be sent. Filters that buffer outgoing data should implement the `flush` filter method.

Upon receiving control, a `flush` implementation must write any buffered data to the filter layer immediately below it. Before returning success, a `flush` implementation must successfully call the `net_flush` function:

```
net_flush(layer->lower).
```

Syntax

```
int flush(FilterLayer *layer);
```

Returns

0 on success or -1 if an error occurred.

Parameters

`FilterLayer *layer` is the filter layer the filter is installed in.

Example

```
int myfilter_flush(FilterLayer *layer)
{
    MyFilterContext context = (MyFilterContext *)layer->context->data;
    if (context->buf.count) {
        int rv;
        rv = net_write(layer->lower, context->buf.data, context->buf.count);
        if (rv != context->buf.count)
            return -1; /* failed to flush data */
        context->buf.count = 0;
    }
    return net_flush(layer->lower);
}
```

See Also

[net_flush](#)

FREE

The `FREE` macro is a platform-independent substitute for the C library routine `free`. It deallocates the space previously allocated by `MALLOC`, `CALLOC`, or `STRDUP` from the request's memory pool.

Syntax

```
FREE(void *ptr);
```

Returns

void

Parameters

`void *ptr` is a `(void *)` pointer to a block of memory. If the pointer is not one created by `MALLOC`, `CALLOC`, or `STRDUP`, the behavior is undefined.

Example

```
char *name;
name = (char *) MALLOC(256);
...
FREE(name);
```

See Also

[CALLOC](#), [REALLOC](#), [STRDUP](#), [PERM_MALLOC](#), [PERM_FREE](#), [PERM_REALLOC](#), [PERM_STRDUP](#)

fs_blk_size

The `fs_blk_size` function returns the block size of the disk partition on which a specified directory resides.

Syntax

```
#include <libproxy/fs.h>
long fs_blk_size(char *root);
```

Returns

the block size, in bytes

Parameters

char **root* is the name of the directory.

See also

[fs_blks_avail](#)

fs_blks_avail

The `fs_blks_avail` function returns the number of disk blocks available on the disk partition on which a specified directory resides.

Syntax

```
#include <libproxy/fs.h>
long fs_blks_avail(char *root);
```

Returns

The number of available disk blocks

Parameters

`char *root` is the name of the directory.

See also

[fs_blk_size](#)

func_exec

The `func_exec` function executes the function named by the `fn` entry in a specified `pblock`. If the function name is not found, it logs the error and returns `REQ_ABORTED`.

You can use this function to execute a built-in Server Application Function (SAF) by identifying it in the `pblock`.

Syntax

```
int func_exec(pblock *pb, Session *sn, Request *rq);
```

Returns

The value returned by the executed function, or the constant `REQ_ABORTED` if no function was executed.

Parameters

`pblock pb` is the `pblock` containing the function name (`fn`) and parameters.

`Session *sn` is the Session.

`Request *rq` is the Request.

The `Session` and `Request` parameters are the same as the ones passed into your SAF.

See Also

[log_error](#)

func_find

The `func_find` function returns a pointer to the function specified by `name`. If the function does not exist, it returns `NULL`.

Syntax

```
FuncPtr func_find(char *name);
```

Returns

A pointer to the chosen function, suitable for dereferencing, or NULL if the function could not be found.

Parameters

char *name is the name of the function.

Example

```
/* this block of code does the same thing as func_exec */
char *afunc = pblock_findval("afunction", pb);
FuncPtr afnptr = func_find(afunc);
if (afnptr)
    return (afnptr)(pb, sn, rq);
```

See Also

[func_exec](#)

func_insert

The `func_insert` function dynamically inserts a named function into the server's table of functions. This function should only be called during the `Init` stage.

Syntax

```
FuncStruct *func_insert(char *name, FuncPtr fn);
```

Returns

Returns the `FuncStruct` structure that identifies the newly inserted function. The caller should not modify the contents of the `FuncStruct` structure.

Parameters

char *name is the name of the function.

FuncPtr fn is the pointer to the function.

Example

```
func_insert("my-service-saf", &my_service_saf);
```

See Also

[func_exec](#), [func_find](#)

insert

The `insert` filter method is called when a filter is inserted into a filter stack by the `filter_insert` function or `insert-filter` SAF (applicable in Input-class directives).

Syntax

```
int insert(FilterLayer *layer, pblock *pb);
```

Returns

Returns `REQ_PROCEED` if the filter should be inserted into the filter stack, `REQ_NOACTION` if the filter should not be inserted because it is not required, or `REQ_ABORTED` if the filter should not be inserted because of an error.

Parameters

`FilterLayer *layer` is the filter layer at which the filter is being inserted.

`pblock *pb` is the set of parameters passed to `filter_insert` or specified by the `fn="insert-filter"` directive.

Example

```
FilterMethods myfilter_methods = FILTER_METHODS_INITIALIZER;
const Filter *myfilter;
int myfilter_insert(FilterLayer *layer, pblock *pb)
{
    if (pblock_findval("dont-insert-filter", pb))
        return REQ_NOACTION;
    return REQ_PROCEED;
}
...
myfilter_methods.insert = &myfilter_insert;
myfilter = filter_create("myfilter", &myfilter_methods);
...
```

L

L

log_error

The `log_error` function creates an entry in an error log, recording the date, the severity, and a specified text.

Syntax

```
int log_error(int degree, char *func, Session *sn, Request *rq, char *fmt,
...);
```

Returns

0 if the log entry was created, or -1 if the log entry was not created.

Parameters

`int degree` specifies the severity of the error. It must be one of the following constants:

`LOG_WARN` -- warning

`LOG_MISCONFIG` -- a syntax error or permission violation

`LOG_SECURITY` -- an authentication failure or 403 error from a host

`LOG_FAILURE` -- an internal problem

`LOG_CATASTROPHE` -- a nonrecoverable server error

`LOG_INFORM` -- an informational message

`char *func` is the name of the function where the error has occurred.

`Session *sn` is the Session.

`Request *rq` is the Request.

The `Session` and `Request` parameters are the same as the ones passed into your SAF.

`char *fmt` specifies the format for the `printf` function that delivers the message.

`...` represents a sequence of parameters for the `printf` function.

Example

```
log_error(LOG_WARN, "send-file", sn, rq,
"error opening buffer from %s (%s)", path,
system_errmsg(fd));
```

See Also

[func_exec](#)

M

magnus_atrestart

NOTE Use the `daemon-atrestart` function in place of the obsolete `magnus_atrestart` function.

The `magnus_atrestart` function lets you register a callback function named by *fn* to be used when the server receives a restart signal. Use this function when you need a callback function to deallocate resources allocated by an initialization function.

Syntax

```
#include <netsite.h>
void magnus_atrestart(void (*fn)(void *), void *data);
```

Returns

void

Parameters

void (**fn*) (void *) is the callback function.

void **data* is the parameter passed to the callback function when the server is restarted.

Example

```
/* Close log file when server is restarted */
magnus_atrestart(brief_terminate, NULL);
return REQPROCEED;
```

MALLOC

The `MALLOC` macro is a platform-independent substitute for the C library routine `malloc`. It normally allocates from the request's memory pool. If pooled memory has been disabled in the configuration file (with the `pool-init` built-in SAF), `PERM_MALLOC` and `MALLOC` both obtain their memory from the system heap.

Syntax

```
void *MALLOC(int size)
```

Returns

A void pointer to a block of memory.

Parameters

`int size` is the number of bytes to allocate.

Example

```
/* Allocate 256 bytes for a name */
char *name;
name = (char *) MALLOC(256);
```

See Also

[FREE](#), [CALLOC](#), [REALLOC](#), [STRDUP](#), [PERM_MALLOC](#), [PERM_FREE](#), [PERM_CALLOC](#), [PERM_REALLOC](#), [PERM_STRDUP](#)

N

net_flush

The `net_flush` function flushes any buffered data. If you require that data be sent immediately, call `net_flush` after calling network output functions such as `net_write` or `net_sendfile`.

Syntax

```
int net_flush(SYS_NETFD sd);
```

Returns

0 on success, or a negative value if an error occurred.

Parameters

`SYS_NETFD sd` is the socket to flush.

Example

```
net_write(sn->csd, "Please wait... ", 15);
net_flush(sn->csd);
/* Perform some time-intensive operation */
...
net_write(sn->csd, "Thank you.\n", 11);
```

See Also

[net_write](#), [net_sendfile](#)

net_ip2host

The `net_ip2host` function transforms a textual IP address into a fully-qualified domain name and returns it.

NOTE This function works only if the `DNS` directive is enabled in the `obj.conf` file. For more information, see *Sun Java System Web Proxy Server 4 Administrator's Configuration File Reference*.

Syntax

```
char *net_ip2host(char *ip, int verify);
```

Returns

A new string containing the fully-qualified domain name if the transformation was accomplished, or `NULL` if the transformation was not accomplished.

Parameters

`char *ip` is the IP address as a character string in dotted-decimal notation:
`nnn.nnn.nnn.nnn`

`int verify`, if nonzero, specifies that the function should verify the fully-qualified domain name. Though this requires an extra query, you should use it when checking access control.

net_read

The `net_read` function reads bytes from a specified socket into a specified buffer. The function waits to receive data from the socket until either at least one byte is available in the socket or the specified time has elapsed.

Syntax

```
int net_read (SYS_NETFD sd, char *buf, int sz, int timeout);
```

Returns

The number of bytes read, which will not exceed the maximum size, `sz`. A negative value is returned if an error has occurred, in which case `errno` is set to the constant `ETIMEDOUT` if the operation did not complete before `timeout` seconds elapsed.

Parameters

`SYS_NETFD sd` is the platform-independent socket descriptor.

`char *buf` is the buffer to receive the bytes.

`int sz` is the maximum number of bytes to read.

`int timeout` is the number of seconds to allow for the read operation before returning. The purpose of `timeout` is not to return because not enough bytes were read in the given time, but to limit the amount of time devoted to waiting until some data arrives.

See Also

[net_write](#)

net_sendfile

The `net_sendfile` function sends the contents of a specified file to a specified a socket. Either the whole file or a fraction may be sent, and the contents of the file may optionally be preceded and/or followed by caller-specified data.

Parameters are passed to `net_sendfile` in the `sendfiledata` structure. Before invoking `net_sendfile`, the caller must initialize every `sendfiledata` structure member.

Syntax

```
int net_sendfile(SYS_NETFD sd, const sendfiledata *sfd);
```

Returns

A positive number indicates the number of bytes successfully written, including the headers, file contents, and trailers. A negative value indicates an error.

Parameters

`SYS_NETFD sd` is the socket to write to.

`const sendfiledata *sfd` identifies the data to send.

Example

The following `Service SAF` sends a file bracketed by the strings "begin" and "end."

```
#include <string.h>
#include "nsapi.h"
```

```
NSAPI_PUBLIC int service_net_sendfile(pblock *pb, Session *sn, Request *rq)
{
```

```

char *path;
SYS_FILE fd;
struct sendfiledata sfd;
int rv;

path = pblock_findval("path", rq->vars);
fd = system_fopenRO(path);
if (!fd) {
    log_error(LOG_MISCONFIG, "service-net-sendfile", sn, rq,
              "Error opening %s (%s)", path, system_errmsg());
    return REQ_ABORTED;
}

sfd.fd = fd; /* file to send */
sfd.offset = 0; /* start sending from the
beginning */
sfd.len = 0; /* send the whole file */
sfd.header = "begin"; /* header data to send before the
file */
sfd.hlen = strlen(sfd.header); /* length of header data */
sfd.trailer = "end"; /* trailer data to send after the
file */
sfd.tlen = strlen(sfd.trailer); /* length of trailer data */

/* send the headers, file, and trailers to the client */
rv = net_sendfile(sn->csd, &sfd);

system_fclose(fd);

if (rv < 0) {
    log_error(LOG_INFORM, "service-net-sendfile", sn, rq, "Error sending
%s (%s)", path, system_errmsg());
    return REQ_ABORTED;
}

return REQ_PROCEED;
}

```

See Also[net_flush](#)

net_write

The `net_write` function writes a specified number of bytes to a specified socket from a specified buffer.

Syntax

```
int net_write(SYS_NETFD sd, char *buf, int sz);
```

Returns

The number of bytes written, which may be less than the requested size if an error occurred.

Parameters

`SYS_NETFD sd` is the platform-independent socket descriptor.

`char *buf` is the buffer containing the bytes.

`int sz` is the number of bytes to write.

Example

```
if (net_write(sn->csd, FIRSTMSG, strlen(FIRSTMSG)) == IO_ERROR)
    return REQ_EXIT;
```

See Also

[net_read](#)

netbuf_buf2sd

The `netbuf_buf2sd` function sends a buffer to a socket. You can use this function to send data from IPC pipes to the client.

Syntax

```
int netbuf_buf2sd(netbuf *buf, SYS_NETFD sd, int len);
```

Returns

The number of bytes transferred to the socket, if successful, or the constant `IO_ERROR` if unsuccessful.

Parameters

`netbuf *buf` is the buffer to send.

`SYS_NETFD sd` is the platform-independent identifier of the socket.

`int len` is the length of the buffer.

See Also

[netbuf_close](#), [netbuf_getc](#), [netbuf_grab](#), [netbuf_open](#)

netbuf_close

The `netbuf_close` function deallocates a network buffer and closes its associated files. Use this function when you need to deallocate the network buffer and close the socket.

You should never close the `netbuf` parameter in a `session` structure.

Syntax

```
void netbuf_close(netbuf *buf);
```

Returns

void

Parameters

`netbuf *buf` is the buffer to close.

See Also

[netbuf_buf2sd](#), [netbuf_getc](#), [netbuf_grab](#), [netbuf_open](#)

netbuf_getc

The `netbuf_getc` function retrieves a character from the cursor position of the network buffer specified by `b`.

Syntax

```
netbuf_getc(netbuf b);
```

Returns

The integer representing the character if one was retrieved, or the constant `IO_EOF` or `IO_ERROR` for end of file or error.

Parameters

`netbuf b` is the buffer from which to retrieve one character.

See Also

[netbuf_buf2sd](#), [netbuf_close](#), [netbuf_grab](#), [netbuf_open](#)

netbuf_grab

The `netbuf_grab` function reads `sz` number of bytes from the network buffer's (`buf`) socket into the network buffer. If the buffer is not large enough it is resized. The data can be retrieved from `buf->inbuf` on success.

This function is used by the function `netbuf_buf2sd`.

Syntax

```
int netbuf_grab(netbuf *buf, int sz);
```

Returns

The number of bytes actually read (between 1 and `sz`) if the operation was successful, or the constant `IO_EOF` or `IO_ERROR` for end of file or error.

Parameters

`netbuf *buf` is the buffer to read into.

`int sz` is the number of bytes to read.

See Also

[netbuf_buf2sd](#), [netbuf_close](#), [netbuf_grab](#), [netbuf_open](#)

netbuf_open

The `netbuf_open` function opens a new network buffer and returns it. You can use `netbuf_open` to create a `netbuf` structure and start using buffered I/O on a socket.

Syntax

```
netbuf* netbuf_open(SYS_NETFD sd, int sz);
```

Returns

A pointer to a new `netbuf` structure (network buffer).

Parameters

`SYS_NETFD sd` is the platform-independent identifier of the socket.

`int sz` is the number of characters to allocate for the network buffer.

See Also

[netbuf_buf2sd](#), [netbuf_close](#), [netbuf_getc](#), [netbuf_grab](#)

nsapi_module_init

Plugin developers may define an `nsapi_module_init` function, which is a module initialization entry point that enables a plugin to create filters when it is loaded. When an NSAPI module contains an `nsapi_module_init` function, the server will call that function immediately after loading the module. The `nsapi_module_init` presents the same interface as an `Init` SAF, and it must follow the same rules.

The `nsapi_module_init` function may be used to register SAFs with `func_insert`, and create filters with `filter_create`.

Syntax

```
int nsapi_module_init(pblock *pb, Session *sn, Request *rq);
```

Returns

`REQ_PROCEED` on success, or `REQ_ABORTED` on error.

Parameters

`pblock *pb` is a set of parameters specified by the `fn="load-modules"` directive.

`Session *sn` (the Session) is `NULL`.

`Request *rq` (the Request) is `NULL`.

NSAPI_RUNTIME_VERSION

The `NSAPI_RUNTIME_VERSION` macro defines the NSAPI version available at runtime. This is the same as the highest NSAPI version supported by the server the plugin is running in. The NSAPI version is encoded as in `USE_NSAPI_VERSION`.

The value returned by the `NSAPI_RUNTIME_VERSION` macro is valid only in iPlanet™ Web Server 6.0, Netscape Enterprise Server 6.0, Sun Java System Web Server 6.1, and Sun Java System Web Proxy Server 4 and higher. That is, the server must support NSAPI 3.1 for this macro to return a valid value. Additionally, to use `NSAPI_RUNTIME_VERSION`, you must compile against an `nsapi.h` header file that supports NSAPI 3.2 or higher.

Plugin developers should not attempt to set the value of the `NSAPI_RUNTIME_VERSION` macro directly. Instead, see the `USE_NSAPI_VERSION` macro.

Syntax

```
int NSAPI_RUNTIME_VERSION
```

Example

```

NSAPI_PUBLIC int log_nsapi_runtime_version(pblock *pb, Session *sn, Request
*rq) {
    log_error(LOG_INFORM, "log-nsapi-runtime-version", sn, rq,
              "Server supports NSAPI version %d.%d\n",
              NSAPI_RUNTIME_VERSION / 100,
              NSAPI_RUNTIME_VERSION % 100);
return REQ_PROCEED;
}

```

See Also

[NSAPI_VERSION](#), [USE_NSAPI_VERSION](#)

NSAPI_VERSION

The `NSAPI_VERSION` macro defines the NSAPI version used at compile time. This value is determined by the value of the `USE_NSAPI_VERSION` macro, or, if the plugin developer did not define `USE_NSAPI_VERSION`, by the highest NSAPI version supported by the `nsapi.h` header the plugin was compiled against. The NSAPI version is encoded as in `USE_NSAPI_VERSION`.

Plugin developers should not attempt to set the value of the `NSAPI_VERSION` macro directly. Instead, see the `USE_NSAPI_VERSION` macro..

Syntax

```
int NSAPI_VERSION
```

Example**Example**

```

NSAPI_PUBLIC int log_nsapi_compile_time_version(pblock *pb, Session *sn,
Request *rq) {
    log_error(LOG_INFORM, "log-nsapi-compile-time-version", sn, rq,
              "Plugin compiled against NSAPI version %d.%d\n",
              NSAPI_VERSION / 100,
              NSAPI_VERSION % 100);
return REQ_PROCEED;
}

```

See Also

[NSAPI_RUNTIME_VERSION](#), [USE_NSAPI_VERSION](#)

P

param_create

The `param_create` function creates a `pb_param` structure containing a specified name and value. The name and value are copied. Use this function to prepare a `pb_param` structure to be used in calls to `pblock` routines such as `pblock_pinsert`.

Syntax

```
pb_param *param_create(char *name, char *value);
```

Returns

A pointer to a new `pb_param` structure.

Parameters

`char *name` is the string containing the name.

`char *value` is the string containing the value.

Example

```
pb_param *newpp = param_create("content-type", "text/plain");  
pblock_pinsert(newpp, rq->srvhdrs);
```

See Also

[param_free](#), [pblock_pinsert](#), [pblock_remove](#)

param_free

The `param_free` function frees the `pb_param` structure specified by `pp` and its associated structures. Use the `param_free` function to dispose a `pb_param` after removing it from a `pblock` with `pblock_remove`.

Syntax

```
int param_free(pb_param *pp);
```

Returns

1 if the parameter was freed or 0 if the parameter was NULL.

Parameters

`pb_param *pp` is the name-value pair stored in a `pblock`.

Example

```
if (param_free(pblock_remove("content-type", rq-srvhdrs)))
    return; /* we removed it */
```

See Also

[param_create](#), [pblock_pinsert](#), [pblock_remove](#)

pblock_copy

The `pblock_copy` function copies the entries of the source `pblock` and adds them into the destination `pblock`. Any previous entries in the destination `pblock` are left intact.

Syntax

```
void pblock_copy(pblock *src, pblock *dst);
```

Returns

void

Parameters

`pblock *src` is the source `pblock`.

`pblock *dst` is the destination `pblock`.

Names and values are newly allocated so that the original `pblock` may be freed, or the new `pblock` changed without affecting the original `pblock`.

See Also

[pblock_create](#), [pblock_dup](#), [pblock_free](#), [pblock_find](#), [pblock_findval](#), [pblock_remove](#), [pblock_nvinsert](#)

pblock_create

The `pblock_create` function creates a new `pblock`. The `pblock` maintains an internal hash table for fast name-value pair lookups.

Syntax

```
pblock *pblock_create(int n);
```

Returns

A pointer to a newly allocated `pblock`.

Parameters

`int n` is the size of the hash table (number of name-value pairs) for the pblock.

See Also

[pblock_copy](#), [pblock_dup](#), [pblock_find](#), [pblock_findval](#), [pblock_free](#), [pblock_nvinsert](#), [pblock_remove](#)

pblock_dup

The `pblock_dup` function duplicates a pblock. It is equivalent to a sequence of `pblock_create` and `pblock_copy`.

Syntax

```
pblock *pblock_dup(pblock *src);
```

Returns

A pointer to a newly allocated pblock.

Parameters

`pblock *src` is the source pblock.

See Also

[pblock_create](#), [pblock_find](#), [pblock_findval](#), [pblock_free](#), [pblock_nvinsert](#), [pblock_remove](#)

pblock_find

The `pblock_find` function finds a specified name-value pair entry in a pblock, and returns the `pb_param` structure. If you only want the value associated with the name, use the `pblock_findval` function.

This function is implemented as a macro.

Syntax

```
pb_param *pblock_find(char *name, pblock *pb);
```

Returns

A pointer to the `pb_param` structure if one was found, or NULL if name was not found.

Parameters

char *name is the name of a name-value pair.

pblock *pb is the pblock to be searched.

See Also

[pblock_copy](#), [pblock_dup](#), [pblock_findval](#), [pblock_free](#), [pblock_nvinsert](#), [pblock_remove](#)

pblock_findlong

The `pblock_findlong` function finds a specified name-value pair entry in a parameter block, and retrieves the name and structure of the parameter block. Use `pblock_findlong` if you want to retrieve the name, structure, and value of the parameter block. However, if you want only the name and structure of the parameter block, use the `pblock_find` function. Do not use these two functions in conjunction.

Syntax

```
#include <libproxy/util.h>
long pblock_findlong(char *name, pblock *pb);
```

Returns

- A long containing the value associated with the name
- -1 if no match was found

Parameters

char *name is the name of a name-value pair.

pblock *pb is the parameter block to be searched.

See also

pblock_ninsert

pblock_findval

The `pblock_findval` function finds the value of a specified name in a pblock. If you just want the `pb_param` structure of the pblock, use the `pblock_find` function.

The pointer returned is a pointer into the pblock. Do not FREE it. If you want to modify it, do a `STRDUP` and modify the copy.

Syntax

```
char *pblock_findval(char *name, pblock *pb);
```

Returns

A string containing the value associated with the name or NULL if no match was found.

Parameters

char *name is the name of a name-value pair.

pblock *pb is the pblock to be searched.

Example

see [pblock_nvinsert](#).

See Also

[pblock_create](#), [pblock_copy](#), [pblock_find](#), [pblock_free](#), [pblock_nvinsert](#), [pblock_remove](#), [request_header](#)

pblock_free

The `pblock_free` function frees a specified `pblock` and any entries inside it. If you want to save a variable in the `pblock`, remove the variable using the function `pblock_remove` and save the resulting pointer.

Syntax

```
void pblock_free(pblock *pb);
```

Returns

void

Parameters

pblock *pb is the pblock to be freed.

See Also

[pblock_copy](#), [pblock_create](#), [pblock_dup](#), [pblock_find](#), [pblock_findval](#), [pblock_nvinsert](#), [pblock_remove](#)

pblock_ninsert

The `pblock_ninsert` function creates a new parameter structure with a given name and long numeric value and inserts it into a specified parameter block. The name and value parameters are also newly allocated.

Syntax

```
#include <libproxy/util.h>
pb_param *pblock_ninsert(char *name, long value, pblock *pb);
```

Returns

The newly allocated parameter block structure

Parameters

char **name* is the name by which the name-value pair is stored.

long *value* is the long (or integer) value being inserted into the parameter block.

pblock **pb* is the parameter block into which the insertion occurs.

See also

[pblock_findlong](#)

pblock_nninsert

The `pblock_nninsert` function creates a new entry with a given name and a numeric value in the specified `pblock`. The numeric value is first converted into a string. The name and value parameters are copied.

Syntax

```
pb_param *pblock_nninsert(char *name, int value, pblock *pb);
```

Returns

A pointer to the new `pb_param` structure.

Parameters

char **name* is the name of the new entry.

int *value* is the numeric value being inserted into the `pblock`. This parameter must be an integer. If the value you assign is not a number, then instead use the function `pblock_nvinsert` to create the parameter.

pblock **pb* is the `pblock` into which the insertion occurs.

See Also

[pblock_copy](#), [pblock_create](#), [pblock_find](#), [pblock_free](#), [pblock_nvinsert](#), [pblock_remove](#), [pblock_str2pblock](#)

pblock_nvinsert

The `pblock_nvinsert` function creates a new entry with a given name and character value in the specified `pblock`. The name and value parameters are copied.

Syntax

```
pb_param *pblock_nvinsert(char *name, char *value, pblock *pb);
```

Returns

A pointer to the newly allocated `pb_param` structure.

Parameters

`char *name` is the name of the new entry.

`char *value` is the string value of the new entry.

`pblock *pb` is the `pblock` into which the insertion occurs.

Example

```
pblock_nvinsert("content-type", "text/html", rq->srvhdrs);
```

See Also

[pblock_copy](#), [pblock_create](#), [pblock_find](#), [pblock_free](#), [pblock_nvinsert](#), [pblock_remove](#), [pblock_str2pblock](#)

pblock_pb2env

The `pblock_pb2env` function copies a specified `pblock` into a specified environment. The function creates one new environment entry for each name-value pair in the `pblock`. Use this function to send `pblock` entries to a program that you are going to execute.

Syntax

```
char **pblock_pb2env(pblock *pb, char **env);
```

Returns

A pointer to the environment.

Parameters

`pblock *pb` is the pblock to be copied.

`char **env` is the environment into which the pblock is to be copied.

See Also

[pblock_copy](#), [pblock_create](#), [pblock_find](#), [pblock_free](#), [pblock_nvinsert](#), [pblock_remove](#), [pblock_str2pblock](#)

pblock_pblock2str

The `pblock_pblock2str` function copies all parameters of a specified pblock into a specified string. The function allocates additional nonheap space for the string if needed.

Use this function to stream the pblock for archival and other purposes.

Syntax

```
char *pblock_pblock2str(pblock *pb, char *str);
```

Returns

The new version of the `str` parameter. If `str` is NULL, this is a new string; otherwise, it is a reallocated string. In either case, it is allocated from the request's memory pool.

Parameters

`pblock *pb` is the pblock to be copied.

`char *str` is the string into which the pblock is to be copied. It must have been allocated by `MALLOC` or `REALLOC`, not by `PERM_MALLOC` or `PERM_REALLOC` (which allocate from the system heap).

Each name-value pair in the string is separated from its neighbor pair by a space, and is in the format *name*="value."

See Also

[pblock_copy](#), [pblock_create](#), [pblock_find](#), [pblock_free](#), [pblock_nvinsert](#), [pblock_remove](#), [pblock_str2pblock](#)

pblock_pinsert

The function `pblock_pinsert` inserts a `pb_param` structure into a pblock.

Syntax

```
void pblock_pinsert(pb_param *pp, pblock *pb);
```

Returns

void

Parameters

`pb_param *pp` is the `pb_param` structure to insert.

`pblock *pb` is the `pblock`.

See Also

[pblock_copy](#), [pblock_create](#), [pblock_find](#), [pblock_free](#), [pblock_nvinsert](#), [pblock_remove](#), [pblock_str2pblock](#)

pblock_remove

The `pblock_remove` function removes a specified name-value entry from a specified `pblock`. If you use this function, you should eventually call `param_free` to deallocate the memory used by the `pb_param` structure.

Syntax

```
pb_param *pblock_remove(char *name, pblock *pb);
```

Returns

A pointer to the named `pb_param` structure if it was found, or NULL if the named `pb_param` was not found.

Parameters

`char *name` is the name of the `pb_param` to be removed.

`pblock *pb` is the `pblock` from which the name-value entry is to be removed.

See Also

[pblock_copy](#), [pblock_create](#), [pblock_find](#), [pblock_free](#), [pblock_nvinsert](#), [param_create](#), [param_free](#)

pblock_replace_name

The `pblock_replace_name` function replaces the name of a name-value pair, retaining the value.

Syntax

```
#include <libproxy/util.h>
void pblock_replace_name(char *oname,char *nname, pblock *pb);
```

Returns

void

Parameters

char **oname* is the old name of a name-value pair.

char **nname* is the new name for the name-value pair.

pblock **pb* is the parameter block to be searched.

See also

[pblock_remove](#)

pblock_str2pblock

The `pblock_str2pblock` function scans a string for parameter pairs, adds them to a `pblock`, and returns the number of parameters added.

Syntax

```
int pblock_str2pblock(char *str, pblock *pb);
```

Returns

The number of parameter pairs added to the `pblock`, if any, or -1 if an error occurred.

Parameters

char **str* is the string to be scanned.

The name-value pairs in the string can have the format *name=value* or *name="value"*.

All backslashes (\) must be followed by a literal character. If string values are found with no unescaped = signs (no `name=`), it assumes the names 1, 2, 3, and so on, depending on the string position. For example, if `pblock_str2pblock` finds "some strings together," the function treats the strings as if they appeared in name-value pairs as 1="some" 2="strings" 3="together."

pblock **pb* is the `pblock` into which the name-value pairs are stored.

See Also

[pblock_copy](#), [pblock_create](#), [pblock_find](#), [pblock_free](#), [pblock_nvinsert](#), [pblock_remove](#), [pblock_pblock2str](#)

PERM_CALLOC

The `PERM_CALLOC` macro is a platform-independent substitute for the C library routine `calloc`. It allocates `int size` bytes of memory that persist after the request that is being processed has been completed. If pooled memory has been disabled in the configuration file (with the `pool-init` built-in SAF), `PERM_CALLOC` and `CALLOC` both obtain their memory from the system heap.

Syntax

```
void *PERM_CALLOC(int size)
```

Returns

A void pointer to a block of memory.

Parameters

`int size` is the size in bytes of each element.

Example

```
char **name;
name = (char **) PERM_CALLOC(100);
```

See Also

[PERM_FREE](#), [PERM_STRDUP](#), [PERM_MALLOC](#), [PERM_REALLOC](#), [MALLOC](#), [FREE](#), [CALLOC](#), [STRDUP](#), [REALLOC](#)

PERM_FREE

The `PERM_FREE` macro is a platform-independent substitute for the C library routine `free`. It deallocates the persistent space previously allocated by `PERM_MALLOC`, `PERM_CALLOC`, or `PERM_STRDUP`. If pooled memory has been disabled in the configuration file (with the `pool-init` built-in SAF), `PERM_FREE` and `FREE` both deallocate memory in the system heap.

Syntax

```
PERM_FREE(void *ptr);
```

Returns

void

Parameters

void *ptr is a (void *) pointer to block of memory. If the pointer is not one created by `PERM_MALLOC`, `PERM_CALLOC`, or `PERM_STRDUP`, the behavior is undefined.

Example

```
char *name;
name = (char *) PERM_MALLOC(256);
...
PERM_FREE(name);
```

See Also

[FREE](#), [MALLOC](#), [CALLOC](#), [REALLOC](#), [STRDUP](#), [PERM_MALLOC](#), [PERM_CALLOC](#), [PERM_REALLOC](#), [PERM_STRDUP](#)

PERM_MALLOC

The `PERM_MALLOC` macro is a platform-independent substitute for the C library routine `malloc`. It provides allocation of memory that persists after the request that is being processed has been completed. If pooled memory has been disabled in the configuration file (with the `pool-init` built-in SAF), `PERM_MALLOC` and `MALLOC` both obtain their memory from the system heap.

Syntax

```
void *PERM_MALLOC(int size)
```

Returns

A void pointer to a block of memory.

Parameters

int size is the number of bytes to allocate.

Example

```
/* Allocate 256 bytes for a name */
char *name;
name = (char *) PERM_MALLOC(256);
```

See Also

[PERM_FREE](#), [PERM_STRDUP](#), [PERM_CALLOC](#), [PERM_REALLOC](#), [MALLOC](#), [FREE](#), [CALLOC](#), [STRDUP](#), [REALLOC](#)

PERM_REALLOC

The `PERM_REALLOC` macro is a platform-independent substitute for the C library routine `realloc`. It changes the size of a specified memory block that was originally created by `MALLOC`, `CALLOC`, or `STRDUP`. The contents of the object remains unchanged up to the lesser of the old and new sizes. If the new size is larger, the new space is uninitialized.

Warning

Calling `PERM_REALLOC` for a block that was allocated with `MALLOC`, `CALLOC`, or `STRDUP` will not work.

Syntax

```
void *PERM_REALLOC(void *ptr, int size)
```

Returns

A void pointer to a block of memory.

Parameters

`void *ptr` a void pointer to a block of memory created by `PERM_MALLOC`, `PERM_CALLOC`, or `PERM_STRDUP`.

`int size` is the number of bytes to which the memory block should be resized.

Example

```
char *name;
name = (char *) PERM_MALLOC(256);
if (NotBigEnough())
    name = (char *) PERM_REALLOC(512);
```

See Also

[PERM_MALLOC](#), [PERM_FREE](#), [PERM_CALLOC](#), [PERM_STRDUP](#), [MALLOC](#), [FREE](#), [STRDUP](#), [CALLOC](#), [REALLOC](#)

PERM_STRDUP

The `PERM_STRDUP` macro is a platform-independent substitute for the C library routine `strdup`. It creates a new copy of a string in memory that persists after the request that is being processed has been completed. If pooled memory has been disabled in the configuration file (with the `pool-init` built-in SAF), `PERM_STRDUP` and `STRDUP` both obtain their memory from the system heap.

The `PERM_STRDUP` routine is functionally equivalent to:

```
newstr = (char *) PERM_MALLOC(strlen(str) + 1);
strcpy(newstr, str);
```

A string created with `PERM_STRDUP` should be disposed with `PERM_FREE`.

Syntax

```
char *PERM_STRDUP(char *ptr);
```

Returns

A pointer to the new string.

Parameters

`char *ptr` is a pointer to a string.

See Also

[PERM_MALLOC](#), [PERM_FREE](#), [PERM_CALLOC](#), [PERM_REALLOC](#), [MALLOC](#), [FREE](#), [STRDUP](#), [CALLOC](#), [REALLOC](#)

prepare_nsapi_thread

The `prepare_nsapi_thread` function allows threads that are not created by the server to act like server-created threads. This function must be called before any NSAPI functions are called from a thread that is not server-created.

Syntax

```
void prepare_nsapi_thread(Request *rq, Session *sn);
```

Returns

void

Parameters

`Request *rq` is the Request.

`Session *sn` is the Session.

The Request and Session parameters are the same as the ones passed into your SAF.

See Also

[protocol_start_response](#)

protocol_dump822

The `protocol_dump822` function prints headers from a specified `pblock` into a specific buffer, with a specified size and position. Use this function to serialize the headers so that they can be sent, for example, in a mail message.

Syntax

```
char *protocol_dump822(pblock *pb, char *t, int *pos, int tsz);
```

Returns

A pointer to the buffer, which will be reallocated if necessary.

The function also modifies `*pos` to the end of the headers in the buffer.

Parameters

`pblock *pb` is the `pblock` structure.

`char *t` is the buffer, allocated with `MALLOC`, `CALLOC`, or `STRDUP`.

`int *pos` is the position within the buffer at which the headers are to be dumped.

`int tsz` is the size of the buffer.

See Also

[protocol_start_response](#), [protocol_status](#)

protocol_finish_request

The `protocol_finish_request` function finishes a specified request. For HTTP, the function just closes the socket.

Syntax

```
#include <frame/protocol.h>
void protocol_finish_request(Session *sn, Request *rq);
```

Returns

void

Parameters

`Session *sn` is the Session that generated the request.

`Request *rq` is the Request to be finished.

See also

[protocol_handle_session](#), [protocol_scan_headers](#), [protocol_start_response](#), [protocol_status](#)

protocol_handle_session

The `protocol_handle_session` function processes each request generated by a specified session.

Syntax

```
#include <frame/protocol.h>
void protocol_handle_session(Session *sn);
```

Parameters

Session `*sn` is the that generated the requests.

See also

[protocol_scan_headers](#), [protocol_start_response](#), [protocol_status](#)

protocol_parse_request

Parses the first line of an HTTP request.

Syntax

```
#include <frame/protocol.h>
int protocol_parse_request(char *t, Request *rq, Session *sn);
```

Returns

- The constant `REQ_PROCEED` if the operation succeeded
- The constant `REQ_ABORTED` if the operation did not succeed

Parameters

char `*t` defines a string of length `REQ_MAX_LINE`. This is an optimization for the internal code to reduce usage of runtime stack.

Request `*rq` is the request to be parsed.

Session `*sn` is the session that generated the request.

See also

[protocol_scan_headers](#), [protocol_start_response](#), [protocol_status](#)

protocol_scan_headers

Scans HTTP headers from a specified network buffer, and places them in a specified parameter block.

Folded lines are joined and the linefeeds are removed (but not the whitespace). If there are any repeat headers, they are joined and the two field bodies are separated by a comma and space. For example, multiple mail headers are combined into one header and a comma is used to separate the field bodies.

Syntax

```
#include <frame/protocol.h>
int protocol_scan_headers(Session *sn, netbuf *buf, char *t, pblock *headers);
```

Returns

- The constant `REQ_PROCEED` if the operation succeeded
- The constant `REQ_ABORTED` if the operation did not succeed

Parameters

Session **sn* is the session that generated the request. The structure named by *sn* contains a pointer to a netbuf called *inbuf*. If the parameter *buf* is NULL, the function automatically uses *inbuf*.

Note that *sn* is an optional parameter that is used for error logs. Use NULL if you wish.

netbuf **buf* is the network buffer to be scanned for HTTP headers.

char **t* defines a string of length `REQ_MAX_LINE`. This is an optimization for the internal code to reduce usage of runtime stack.

pblock **headers* is the parameter block to receive the headers.

See also

[protocol_handle_session](#), [protocol_start_response](#), [protocol_status](#)

protocol_set_finfo

The `protocol_set_finfo` function retrieves the content-length and last-modified date from a specified `stat` structure and adds them to the response headers (`rq->srvhdrs`). Call `protocol_set_finfo` before calling `protocol_start_response`.

Syntax

```
int protocol_set_finfo(Session *sn, Request *rq, struct stat *finfo);
```

Returns

The constant `REQ_PROCEED` if the request can proceed normally, or the constant `REQ_ABORTED` if the function should treat the request normally but not send any output to the client.

Parameters

`Session *sn` is the Session.

`Request *rq` is the Request.

The `Session` and `Request` parameters are the same as the ones passed into your SAF.

`stat *finfo` is the `stat` structure for the file.

The `stat` structure contains the information about the file from the file system. You can get the `stat` structure info using `request_stat_path`.

See Also

[protocol_start_response](#), [protocol_status](#)

protocol_start_response

The `protocol_start_response` function initiates the HTTP response for a specified session and request. If the protocol version is HTTP/0.9, the function does nothing, because that version has no concept of status. If the protocol version is HTTP/1.0, the function sends a status line followed by the response headers. Use this function to set up HTTP and prepare the client and server to receive the body (or data) of the response.

Syntax

```
int protocol_start_response(Session *sn, Request *rq);
```

Returns

The constant `REQ_PROCEED` if the operation succeeded, in which case you should send the data you were preparing to send.

The constant `REQ_NOACTION` if the operation succeeded but the request method was `HEAD`, in which case no data should be sent to the client.

The constant `REQ_ABORTED` if the operation did not succeed.

Parameters

Session **sn* is the Session.

Request **rq* is the Request.

The Session and Request parameters are the same as the ones passed into your SAF.

Example

```
/* A noaction response from this function means the request was HEAD */
if (protocol_start_response(sn, rq) == REQ_NOACTION) {
    filebuf_close(groupbuf); /* close our file*/
    return REQ_PROCEED;
}
```

See Also

[protocol_status](#)

protocol_status

The `protocol_status` function sets the session status to indicate whether an error condition occurred. If the reason string is NULL, the server attempts to find a reason string for the given status code. If it finds none, it returns “Unknown reason.” The reason string is sent to the client in the HTTP response line. Use this function to set the status of the response before calling the function `protocol_start_response`.

For the complete list of valid status code constants, please refer to the file “`nsapi.h`” in the server distribution.

Syntax

```
void protocol_status(Session *sn, Request *rq, int n, char *r);
```

Returns

void, but it sets values in the Session/Request designated by `sn/rq` for the status code and the reason string.

Parameters

Session **sn* is the Session.

Request **rq* is the Request.

The Session and Request parameters are the same as the ones passed into your SAF.

int *n* is one of the status code constants above.

char **r* is the reason string.

Example

```

/* if we find extra path-info, the URL was bad so tell the */
/* browser it was not found */
if (t = pblock_findval("path-info", rq->vars)) {
    protocol_status(sn, rq, PROTOCOL_NOT_FOUND, NULL);
    log_error(LOG_WARN, "function-name", sn, rq, "%s not found",
        path);
    return REQ_ABORTED;
}

```

See Also

[protocol_start_response](#)

protocol_uri2url

The `protocol_uri2url` function takes strings containing the given URI prefix and URI suffix, and creates a newly allocated, fully qualified URL in the form `http://(server):(port)(prefix)(suffix)`. See [protocol_uri2url_dynamic](#).

If you want to omit either the URI prefix or suffix, use "" instead of NULL as the value for either parameter.

Syntax

```
char *protocol_uri2url(char *prefix, char *suffix);
```

Returns

A new string containing the URL.

Parameters

`char *prefix` is the prefix.

`char *suffix` is the suffix.

See Also

[protocol_start_response](#), [protocol_status](#), [pblock_nvinsert](#), [protocol_uri2url_dynamic](#)

protocol_uri2url_dynamic

The `protocol_uri2url` function takes strings containing the given URI prefix and URI suffix, and creates a newly allocated, fully qualified URL in the form `http://(server):(port)(prefix)(suffix)`.

If you want to omit either the URI prefix or suffix, use "" instead of NULL as the value for either parameter.

The `protocol_uri2url_dynamic` function is similar to the `protocol_uri2url` function, but should be used whenever the `session` and `request` structures are available. This ensures that the URL it constructs refers to the host that the client specified.

Syntax

```
char *protocol_uri2url(char *prefix, char *suffix, Session *sn,
Request *rq);
```

Returns

A new string containing the URL.

Parameters

`char *prefix` is the prefix.

`char *suffix` is the suffix.

`Session *sn` is the Session.

`Request *rq` is the Request.

The `Session` and `Request` parameters are the same as the ones passed into your SAF.

See Also

[protocol_start_response](#), [protocol_status](#), [protocol_uri2url_dynamic](#)

R

read

The `read` filter method is called when input data is required. Filters that modify or consume incoming data should implement the `read` filter method.

Upon receiving control, a `read` implementation should fill `buf` with up to `amount` bytes of input data. This data may be obtained by calling the `net_read` function, as shown in the example below.

Syntax

```
int read(FilterLayer *layer, void *buf, int amount, int timeout);
```

Returns

The number of bytes placed in `buf` on success, 0 if no data is available, or a negative value if an error occurred.

Parameters

`FilterLayer *layer` is the filter layer in which the filter is installed.

`void *buf` is the buffer in which data should be placed.

`int amount` is the maximum number of bytes that should be placed in the buffer.

`int timeout` is the number of seconds to allow for the `read` operation before returning. The purpose of `timeout` is not to return because not enough bytes were read in the given time, but to limit the amount of time devoted to waiting until some data arrives.

Example

```
int myfilter_read(FilterLayer *layer, void *buf, int amount, int
timeout)
{
    return net_read(layer->lower, buf, amount, timeout);
}
```

See Also

[net_read](#)

REALLOC

The `REALLOC` macro is a platform-independent substitute for the C library routine `realloc`. It changes the size of a specified memory block that was originally created by `MALLOC`, `CALLOC`, or `STRDUP`. The contents of the object remains unchanged up to the lesser of the old and new sizes. If the new size is larger, the new space is uninitialized.

Warning

Calling `REALLOC` for a block that was allocated with `PERM_MALLOC`, `PERM_CALLOC`, or `PERM_STRDUP` will not work.

Syntax

```
void *REALLOC(void *ptr, int size);
```

Returns

A pointer to the new space if the request could be satisfied.

Parameters

`void *ptr` is a (void *) pointer to a block of memory. If the pointer is not one created by `MALLOC`, `CALLOC`, or `STRDUP`, the behavior is undefined.

`int size` is the number of bytes to allocate.

Example

```
char *name;
name = (char *) MALLOC(256);
if (NotBigEnough())
    name = (char *) REALLOC(512);
```

See Also

[MALLOC](#), [FREE](#), [STRDUP](#), [CALLOC](#), [PERM_MALLOC](#), [PERM_FREE](#), [PERM_REALLOC](#), [PERM_CALLOC](#), [PERM_STRDUP](#)

remove

The `remove` filter method is called when the filter stack is destroyed, or when a filter is removed from a filter stack by the `filter_remove` function or `remove-filter` SAF (applicable in `Input`-, `Output`-, `Service`-, and `Error`-class directives).

Note that it may be too late to flush buffered data when the `remove` method is invoked. For this reason, filters that buffer outgoing data should implement the `flush` filter method.

Syntax

```
void remove(FilterLayer *layer);
```

Returns

void

Parameters

`FilterLayer *layer` is the filter layer the filter is installed in.

See Also

[flush](#)

request_create

The `request_create` function is a utility function that creates a new request structure.

Syntax

```
#include <frame/req.h>
Request *request_create(void);
```

Returns

A Request structure

Parameters

No parameter is required.

See also

[request_free](#), [request_header](#)

request_free

The `request_free` function frees a specified request structure.

Syntax

```
#include <frame/req.h>
void request_free(Request *req);
```

Returns

void

Parameters

Request **rq* is the Request structure to be freed.

See also

[request_header](#)

request_header

The `request_header` function finds an entry in the `pblock` containing the client's HTTP request headers (`rq->headers`). You must use this function rather than `pblock_findval` when accessing the client headers, since the server may begin processing the request before the headers have been completely read.

Syntax

```
int request_header(char *name, char **value, Session *sn, Request *rq);
```

Returns

A result code, `REQ_PROCEED` if the header was found, `REQ_ABORTED` if the header was not found, `REQ_EXIT` if there was an error reading from the client.

Parameters

`char *name` is the name of the header.

`char **value` is the address where the function will place the value of the specified header. If none is found, the function stores a `NULL`.

`Session *sn` is the Session.

`Request *rq` is the Request.

The `Session` and `Request` parameters are the same as the ones passed into your SAF.

See Also

[request_create](#), [request_free](#)

S

sem_grab

The `sem_grab` function requests exclusive access to a specified semaphore. If exclusive access is unavailable, the caller blocks execution until exclusive access becomes available. Use this function to ensure that only one server processor thread performs an action at a time.

Syntax

```
#include <base/sem.h>
int sem_grab(SEMAPHORE id);
```

Returns

- -1 if an error occurred
- 0 to signal success

Parameters

`SEMAPHORE id` is the unique identification number of the requested semaphore.

See also

[sem_init](#), [sem_release](#), [sem_terminate](#), [sem_tgrab](#)

sem_init

The `sem_init` function creates a semaphore with a specified name and unique identification number. Use this function to allocate a new semaphore that will be used with the functions `sem_grab` and `sem_release`. Call `sem_init` from an `init` class function to initialize a static or global variable that the other classes will later use.

Syntax

```
#include <base/sem.h>
SEMAPHORE sem_init(char *name, int number);
```

Returns

The constant `SEM_ERROR` if an error occurred.

Parameters

`SEMAPHORE *name` is the name for the requested semaphore. The filename of the semaphore should be a file accessible to the process.

`int number` is the unique identification number for the requested semaphore.

See also

[sem_grab](#), [sem_release](#), [sem_terminate](#)

sem_release

The `sem_release` function releases the process's exclusive control over a specified semaphore. Use this function to release exclusive control over a semaphore created with the function `sem_grab`.

Syntax

```
#include <base/sem.h>
int sem_release(SEMAPHORE id);
```

Returns

- -1 if an error occurred
- 0 if no error occurred

Parameters

`SEMAPHORE id` is the unique identification number of the semaphore.

See also

[sem_grab](#), [sem_init](#), [sem_terminate](#)

sem_terminate

The `sem_terminate` function deallocates the semaphore specified by *id*. You can use this function to deallocate a semaphore that was previously allocated with the function `sem_init`.

Syntax

```
#include <base/sem.h>
void sem_terminate(SEMAPHORE id);
```

Returns

void

Parameters

SEMAPHORE *id* is the unique identification number of the semaphore.

See also

[sem_grab](#), [sem_init](#), [sem_release](#)

sem_tgrab

The `sem_tgrab` function tests and requests exclusive use of a semaphore. Unlike the somewhat similar `sem_grab` function, if exclusive access is unavailable the caller is not blocked but receives a return value of -1. Use this function to ensure that only one server processor thread performs an action at a time.

Syntax

```
#include <base/sem.h>
int sem_grab(SEMAPHORE id);
```

Returns

- -1 if an error occurred or if exclusive access was not available
- 0 exclusive access was granted

Parameters

SEMAPHORE *id* is the unique identification number of the semaphore.

See also

[sem_grab](#), [sem_init](#), [sem_release](#), [sem_terminate](#)

sendfile

The `sendfile` filter method is called when the contents of a file are to be sent. Filters that modify or consume outgoing data may choose to implement the `sendfile` filter method.

If a filter implements the `write` filter method but not the `sendfile` filter method, the server will automatically translate `net_sendfile` calls to `net_write` calls. As a result, filters interested in the outgoing data stream do not need to implement the `sendfile` filter method. However, for performance reasons, it is beneficial for filters that implement the `write` filter method to also implement the `sendfile` filter method.

Syntax

```
int sendfile(FilterLayer *layer, const sendfiledata *data);
```

Returns

The number of bytes consumed, which may be less than the requested amount if an error occurred.

Parameters

`FilterLayer *layer` is the filter layer the filter is installed in.

`const sendfiledata *sfd` identifies the data to send.

Example

```
int myfilter_sendfile(FilterLayer *layer, const sendfiledata *sfd)
{
    return net_sendfile(layer->lower, sfd);
}
```

See Also

[net_sendfile](#)

session_create

The `session_create` function creates a new `Session` structure for the client with a specified socket descriptor and a specified socket address. It returns a pointer to that structure.

Syntax

```
#include <base/session.h>
Session *session_create(SYS_NETFD csd, struct sockaddr_in *sac);
```

Returns

- A pointer to the new Session if one was created
- NULL if no new Session was created

Parameters

`SYS_NETFD csd` is the platform-independent socket descriptor.

`sockaddr_in *sac` is the socket address.

See also

[session_maxdns](#)

session_dns

The `session_dns` function resolves the IP address of the client associated with a specified session into its DNS name. It returns a newly allocated string. You can use `session_dns` to change the numeric IP address into something more readable.

The `session_maxdns` function verifies that the client is who it claims to be; the `session_dns` function does not perform this verification.

NOTE This function works only if the `DNS` directive is enabled in the `obj.conf` file. For more information, see *Sun Java System Web Proxy Server 4 Administrator's Configuration File Reference*.

Syntax

```
char *session_dns(Session *sn);
```

Returns

A string containing the host name, or NULL if the DNS name cannot be found for the IP address.

Parameters

`Session *sn` is the Session.

The Session is the same as the one passed to your SAF.

session_free

The `session_free` function frees a specified Session structure. The `session_free` function does not close the client socket descriptor associated with the Session.

Syntax

```
#include <base/session.h>
void session_free(Session *sn);
```

Returns

void

Parameters

Session **sn* is the Session to be freed.

See also

[session_create](#), [session_maxdns](#)

session_maxdns

The `session_maxdns` function resolves the IP address of the client associated with a specified session into its DNS name. It returns a newly allocated string. You can use `session_maxdns` to change the numeric IP address into something more readable.

NOTE This function works only if the `DNS` directive is enabled in the `obj.conf` file. For more information, see *Sun Java System Web Proxy Server 4 Administrator's Configuration File Reference*

Syntax

```
char *session_maxdns(Session *sn);
```

Returns

A string containing the host name, or NULL if the DNS name cannot be found for the IP address.

Parameters

Session **sn* is the Session.

The Session is the same as the one passed to your SAF.

shexp_casecmp

The `shexp_casecmp` function validates a specified shell expression and compares it with a specified string. It returns one of three possible values representing match, no match, and invalid comparison. The comparison (in contrast to that of the `shexp_cmp` function) is not case-sensitive.

Use this function if you have a shell expression like `*.netscape.com` and you want to make sure that a string matches it, such as `foo.netscape.com`.

Syntax

```
int shexp_casecmp(char *str, char *exp);
```

Returns

0 if a match was found.

1 if no match was found.

-1 if the comparison resulted in an invalid expression.

Parameters

`char *str` is the string to be compared.

`char *exp` is the shell expression (wildcard pattern) to compare against.

See Also

[shexp_cmp](#), [shexp_match](#), [shexp_valid](#)

shexp_cmp

The `shexp_casecmp` function validates a specified shell expression and compares it with a specified string. It returns one of three possible values representing match, no match, and invalid comparison. The comparison (in contrast to that of the `shexp_casecmp` function) is case-sensitive.

Use this function if you have a shell expression like `*.netscape.com` and you want to make sure that a string matches it, such as `foo.netscape.com`.

Syntax

```
int shexp_cmp(char *str, char *exp);
```

Returns

0 if a match was found.

- 1 if no match was found.
- 1 if the comparison resulted in an invalid expression.

Parameters

`char *str` is the string to be compared.

`char *exp` is the shell expression (wildcard pattern) to compare against.

Example

```
/* Use wildcard match to see if this path is one we want */
char *path;
char *match = "/usr/netscape/*";
if (shexp_cmp(path, match) != 0)
    return REQ_NOACTION; /* no match */
```

See Also

[shexp_casecmp](#), [shexp_match](#), [shexp_valid](#)

shexp_match

The `shexp_match` function compares a specified prevalidated shell expression against a specified string. It returns one of three possible values representing match, no match, and invalid comparison. The comparison (in contrast to that of the `shexp_casecmp` function) is case-sensitive.

The `shexp_match` function doesn't perform validation of the shell expression; instead the function assumes that you have already called `shexp_valid`.

Use this function if you have a shell expression such as `*.netscape.com`, and you want to make sure that a string matches it, such as `foo.netscape.com`.

Syntax

```
int shexp_match(char *str, char *exp);
```

Returns

- 0 if a match was found.
- 1 if no match was found.
- 1 if the comparison resulted in an invalid expression.

Parameters

`char *str` is the string to be compared.

`char *exp` is the prevalidated shell expression (wildcard pattern) to compare against.

See Also

[shexp_casecmp](#), [shexp_cmp](#), [shexp_valid](#)

shexp_valid

The `shexp_valid` function validates a specified shell expression named by `exp`. Use this function to validate a shell expression before using the function `shexp_match` to compare the expression with a string.

Syntax

```
int shexp_valid(char *exp);
```

Returns

The constant `NON_SXP` if `exp` is a standard string.

The constant `INVALID_SXP` if `exp` is a shell expression, but invalid.

The constant `VALID_SXP` if `exp` is a valid shell expression.

Parameters

`char *exp` is the shell expression (wildcard pattern) to be validated.

See Also

[shexp_casecmp](#), [shexp_match](#), [shexp_cmp](#)

shmem_alloc

The `shmem_alloc` function allocates a region of shared memory of the given size, using the given name to avoid conflicts between multiple regions in the program. The size of the region will not be automatically increased if its boundaries are overrun; use the `shmem_realloc` function for that.

This function must be called before any daemon workers are spawned in order for the handle to the shared region to be inherited by the children.

Because of the requirement that the region must be inherited by the children, the region cannot be reallocated with a larger size when necessary.

Syntax

```
#include <base/shmem.h>
shmem_s *shmem_alloc(char *name, int size, int expose);
```

Returns

A pointer to a new shared memory region.

Parameters

char **name* is the name for the region of shared memory being created. The value of *name* must be unique to the program that calls the **shmem_alloc** function or conflicts will occur.

int *size* is the number of characters of memory to be allocated for the shared memory.

int *expose* is either zero or nonzero. If nonzero, then on systems that support it, the file that is used to create the shared memory becomes visible to other processes running on the system.

See also

[shmem_free](#)

shmem_free

The `shmem_free` function deallocates (frees) the specified region of memory.

Syntax

```
#include <base/shmem.h>
void *shmem_free(shmem_s *region);
```

Returns

void

Parameters

shmem_s **region* is a shared memory region to be released.

See also

[shmem_alloc](#)

STRDUP

The `STRDUP` macro is a platform-independent substitute for the C library routine `strdup`. It creates a new copy of a string in the request's memory pool.

The `STRDUP` routine is functionally equivalent to:

```
newstr = (char *) MALLOC(strlen(str) + 1);
strcpy(newstr, str);
```

A string created with `STRDUP` should be disposed with `FREE`.

Syntax

```
char *STRDUP(char *ptr);
```

Returns

A pointer to the new string.

Parameters

`char *ptr` is a pointer to a string.

Example

```
char *name1 = "MyName";
char *name2 = STRDUP(name1);
```

See Also

[MALLOC](#), [FREE](#), [CALLOC](#), [REALLOC](#), [PERM_MALLOC](#), [PERM_FREE](#), [PERM_CALLOC](#), [PERM_REALLOC](#), [PERM_STRDUP](#)

system_errmsg

The `system_errmsg` function returns the last error that occurred from the most recent system call. This function is implemented as a macro that returns an entry from the global array `sys_errlist`. Use this macro to help with I/O error diagnostics.

Syntax

```
char *system_errmsg(int param1);
```

Returns

A string containing the text of the latest error message that resulted from a system call. Do not `FREE` this string.

Parameters

`int param1` is reserved, and should always have the value 0.

See Also

[system_fopenRO](#), [system_fopenRW](#), [system_fopenWA](#), [system_lseek](#), [system_fread](#), [system_fwrite](#), [system_fwrite_atomic](#), [system_flock](#), [system_unlock](#), [system_fclose](#)

system_fclose

The `system_fclose` function closes a specified file descriptor. The `system_fclose` function must be called for every file descriptor opened by any of the `system_fopen` functions.

Syntax

```
int system_fclose(SYS_FILE fd);
```

Returns

0 if the close succeeded, or the constant `IO_ERROR` if the close failed.

Parameters

`SYS_FILE fd` is the platform-independent file descriptor.

Example

```
SYS_FILE logfd;  
system_fclose(logfd);
```

See Also

[system_errmsg](#), [system_fopenRO](#), [system_fopenRW](#), [system_fopenWA](#), [system_lseek](#), [system_fread](#), [system_fwrite](#), [system_fwrite_atomic](#), [system_flock](#), [system_unlock](#)

system_flock

The `system_flock` function locks the specified file against interference from other processes. Use `system_flock` if you do not want other processes to use the file you currently have open. Overusing file locking can cause performance degradation and possibly lead to deadlocks.

Syntax

```
int system_flock(SYS_FILE fd);
```

Returns

The constant `IO_OKAY` if the lock succeeded, or the constant `IO_ERROR` if the lock failed.

Parameters

`SYS_FILE fd` is the platform-independent file descriptor.

See Also

[system_errmsg](#), [system_fopenRO](#), [system_fopenRW](#), [system_fopenWA](#), [system_lseek](#), [system_fread](#), [system_fwrite](#), [system_fwrite_atomic](#), [system_unlock](#), [system_fclose](#)

system_fopenRO

The `system_fopenRO` function opens the file identified by `path` in read-only mode and returns a valid file descriptor. Use this function to open files that will not be modified by your program. In addition, you can use `system_fopenRO` to open a new file buffer structure using `filebuf_open`.

Syntax

```
SYS_FILE system_fopenRO(char *path);
```

Returns

The system-independent file descriptor (`SYS_FILE`) if the open succeeded, or 0 if the open failed.

Parameters

`char *path` is the file name.

See Also

[system_errmsg](#), [system_fopenRO](#), [system_fopenWA](#), [system_lseek](#), [system_fread](#), [system_fwrite](#), [system_fwrite_atomic](#), [system_flock](#), [system_unlock](#), [system_fclose](#)

system_fopenRW

The `system_fopenRW` function opens the file identified by `path` in read-write mode and returns a valid file descriptor. If the file already exists, `system_fopenRW` does not truncate it. Use this function to open files that will be read from and written to by your program.

Syntax

```
SYS_FILE system_fopenRW(char *path);
```

Returns

The system-independent file descriptor (`SYS_FILE`) if the open succeeded, or 0 if the open failed.

Parameters

`char *path` is the file name.

Example

```
SYS_FILE fd;
fd = system_fopenRO(pathname);
if (fd == SYS_ERROR_FD)
    break;
```

See Also

[system_errmsg](#), [system_fopenRO](#), [system_fopenWA](#), [system_lseek](#), [system_fread](#), [system_fwrite](#), [system_fwrite_atomic](#), [system_flock](#), [system_unlock](#), [system_fclose](#)

system_fopenWA

The `system_fopenWA` function opens the file identified by `path` in write-append mode and returns a valid file descriptor. Use this function to open those files to which your program will append data.

Syntax

```
SYS_FILE system_fopenWA(char *path);
```

Returns

The system-independent file descriptor (`SYS_FILE`) if the open succeeded, or 0 if the open failed.

Parameters

`char *path` is the file name.

See Also

[system_errmsg](#), [system_fopenRO](#), [system_fopenRW](#), [system_lseek](#), [system_fread](#), [system_fwrite](#), [system_fwrite_atomic](#), [system_flock](#), [system_ulock](#), [system_fclose](#)

system_fread

The `system_fread` function reads a specified number of bytes from a specified file into a specified buffer. It returns the number of bytes read. Before `system_fread` can be used, you must open the file using any of the `system_fopen` functions (except `system_fopenWA`).

Syntax

```
int system_fread(SYS_FILE fd, char *buf, int sz);
```

Returns

The number of bytes read, which may be less than the requested size if an error occurred or the end of the file was reached before that number of characters were obtained.

Parameters

`SYS_FILE fd` is the platform-independent file descriptor.

`char *buf` is the buffer to receive the bytes.

`int sz` is the number of bytes to read.

See Also

[system_errmsg](#), [system_fopenRO](#), [system_fopenRW](#), [system_fopenWA](#), [system_lseek](#), [system_fwrite](#), [system_fwrite_atomic](#), [system_flock](#), [system_ulock](#), [system_fclose](#)

system_fwrite

The `system_fwrite` function writes a specified number of bytes from a specified buffer into a specified file.

Before `system_fwrite` can be used, you must open the file using any of the `system_fopen` functions (except `system_fopenRO`).

Syntax

```
int system_fwrite(SYS_FILE fd, char *buf, int sz);
```

Returns

The constant `IO_OKAY` if the write succeeded, or the constant `IO_ERROR` if the write failed.

Parameters

`SYS_FILE fd` is the platform-independent file descriptor.

`char *buf` is the buffer containing the bytes to be written.

`int sz` is the number of bytes to write to the file.

See Also

[system_errmsg](#), [system_fopenRO](#), [system_fopenRW](#), [system_fopenWA](#), [system_lseek](#), [system_fread](#), [system_fwrite_atomic](#), [system_flock](#), [system_ulock](#), [system_fclose](#)

system_fwrite_atomic

The `system_fwrite_atomic` function writes a specified number of bytes from a specified buffer into a specified file. The function also locks the file prior to performing the write, and then unlocks it when done, thereby avoiding interference between simultaneous write actions. Before `system_fwrite_atomic` can be used, you must open the file using any of the `system_fopen` functions, except `system_fopenRO`.

Syntax

```
int system_fwrite_atomic(SYS_FILE fd, char *buf, int sz);
```

Returns

The constant `IO_OKAY` if the write/lock succeeded, or the constant `IO_ERROR` if the write/lock failed.

Parameters

`SYS_FILE fd` is the platform-independent file descriptor.

`char *buf` is the buffer containing the bytes to be written.

`int sz` is the number of bytes to write to the file.

Example

```
SYS_FILE logfd;
char *logmsg = "An error occurred.";
system_fwrite_atomic(logfd, logmsg, strlen(logmsg));
```

See Also

[system_errmsg](#), [system_fopenRO](#), [system_fopenRW](#), [system_fopenWA](#), [system_lseek](#), [system_fread](#), [system_fwrite](#), [system_flock](#), [system_unlock](#), [system_fclose](#)

system_gmtime

The `system_gmtime` function is a thread-safe version of the standard `gmtime` function. It returns the current time adjusted to Greenwich Mean Time.

Syntax

```
struct tm *system_gmtime(const time_t *tp, const struct tm *res);
```

Returns

A pointer to a calendar time (`tm`) structure containing the GMT time. Depending on your system, the pointer may point to the data item represented by the second parameter, or it may point to a statically-allocated item. For portability, do not assume either situation.

Parameters

`time_t *tp` is an arithmetic time.

`tm *res` is a pointer to a calendar time (`tm`) structure.

Example

```
time_t tp;
struct tm res, *resp;
tp = time(NULL);
resp = system_gmtime(&tp, &res);
```

See Also

[system_localtime](#), [util_strftime](#)

system_localtime

The `system_localtime` function is a thread-safe version of the standard `localtime` function. It returns the current time in the local time zone.

Syntax

```
struct tm *system_localtime(const time_t *tp, const struct tm *res);
```

Returns

A pointer to a calendar time (`tm`) structure containing the local time. Depending on your system, the pointer may point to the data item represented by the second parameter, or it may point to a statically-allocated item. For portability, do not assume either situation.

Parameters

`time_t *tp` is an arithmetic time.

`tm *res` is a pointer to a calendar time (`tm`) structure.

See Also

[system_gmtime](#), [util_strftime](#)

system_lseek

The `system_lseek` function sets the file position of a file. This affects where data from `system_fread` or `system_fwrite` is read or written.

Syntax

```
int system_lseek(SYS_FILE fd, int offset, int whence);
```

Returns

The offset, in bytes, of the new position from the beginning of the file if the operation succeeded, or `-1` if the operation failed.

Parameters

`SYS_FILE fd` is the platform-independent file descriptor.

`int offset` is a number of bytes relative to `whence`. It may be negative.

`int whence` is one of the following constants:

`SEEK_SET`, from the beginning of the file.

`SEEK_CUR`, from the current file position.

`SEEK_END`, from the end of the file.

See Also

`system_errmsg`, `system_fopenRO`, `system_fopenRW`, `system_fopenWA`,
`system_fread`, `system_fwrite`, `system_fwrite_atomic`, `system_flock`,
`system_unlock`, `system_fclose`

system_rename

The `system_rename` function renames a file. It may not work on directories if the old and new directories are on different file systems.

Syntax

```
int system_rename(char *old, char *new);
```

Returns

0 if the operation succeeded, or -1 if the operation failed.

Parameters

`char *old` is the old name of the file.

`char *new` is the new name for the file.

system_unlock

The `system_unlock` function unlocks the specified file that has been locked by the function `system_lock`. For more information about locking, see `system_flock`.

Syntax

```
int system_unlock(SYS_FILE fd);
```

Returns

The constant `IO_OKAY` if the operation succeeded, or the constant `IO_ERROR` if the operation failed.

Parameters

`SYS_FILE fd` is the platform-independent file descriptor.

See Also

`system_errmsg`, `system_fopenRO`, `system_fopenRW`, `system_fopenWA`,
`system_fread`, `system_fwrite`, `system_fwrite_atomic`, `system_flock`,
`system_fclose`

`system_errmsg`, `system_fopenRO`, `system_fopenRW`, `system_fopenWA`,
`system_fread`, `system_fwrite`, `system_fwrite_atomic`, `system_flock`,
`system_fclose`

system_unix2local

The `system_unix2local` function converts a specified UNIX-style path name to a local file system path name. Use this function when you have a file name in the UNIX format (such as one containing forward slashes), and you need to access a file on another system such as Windows. You can use `system_unix2local` to convert the UNIX file name into the format that Windows accepts. In the UNIX environment this function does nothing, but may be called for portability.

Syntax

```
char *system_unix2local(char *path, char *lp);
```

Returns

A pointer to the local file system path string.

Parameters

`char *path` is the UNIX-style path name to be converted.

`char *lp` is the local path name.

You must allocate the parameter `lp`, and it must contain enough space to hold the local path name.

See Also

`system_fclose`, `system_flock`, `system_fopenRO`, `system_fopenRW`,
`system_fopenWA`, `system_fwrite`

systhread_attach

The `systhread_attach` function makes an existing thread into a platform-independent thread.

Syntax

```
SYS_THREAD systhread_attach(void);
```

Returns

A `SYS_THREAD` pointer to the platform-independent thread.

Parameters

none

See Also

[systhread_current](#), [systhread_getdata](#), [systhread_init](#), [systhread_newkey](#), [systhread_setdata](#), [systhread_sleep](#), [systhread_start](#), [systhread_timeriset](#)

systhread_current

The `systhread_current` function returns a pointer to the current thread.

Syntax

```
SYS_THREAD systhread_current(void);
```

Returns

A `SYS_THREAD` pointer to the current thread.

Parameters

none

See Also

[systhread_getdata](#), [systhread_newkey](#), [systhread_setdata](#), [systhread_sleep](#), [systhread_start](#), [systhread_timeriset](#)

systhread_getdata

The `systhread_getdata` function gets data that is associated with a specified key in the current thread.

Syntax

```
void *systhread_getdata(int key);
```

Returns

A pointer to the data that was earlier used with the `systhread_setkey` function from the current thread, using the same value of `key` if the call succeeds. Returns NULL if the call did not succeed; for example, if the `systhread_setkey` function was never called with the specified key during this session.

Parameters

`int key` is the value associated with the stored data by a `systhread_setdata` function. Keys are assigned by the `systhread_newkey` function.

See Also

[systhread_current](#), [systhread_newkey](#), [systhread_setdata](#), [systhread_sleep](#), [systhread_start](#), [systhread_timerset](#)

systhread_init

The `systhread_init` function initializes the threading system.

Syntax

```
#include <base/systr.h>
void systhread_init(char *name);
```

Returns

void

Parameters

`char *name` is a name to be assigned to the program for debugging purposes.

See also

[systhread_attach](#), [systhread_current](#), [systhread_getdata](#), [systhread_newkey](#), [systhread_setdata](#), [systhread_sleep](#), [systhread_start](#), [systhread_terminate](#), [systhread_timerset](#)

systhread_newkey

The `systhread_newkey` function allocates a new integer key (identifier) for thread-private data. Use this key to identify a variable that you want to localize to the current thread, then use the `systhread_setdata` function to associate a value with the key.

Syntax

```
int systhread_newkey(void);
```

Returns

An integer key.

Parameters

none

See Also

[systhread_current](#), [systhread_getdata](#), [systhread_setdata](#),
[systhread_sleep](#), [systhread_start](#), [systhread_timeriset](#)

systhread_setdata

The `systhread_setdata` function associates data with a specified key number for the current thread. Keys are assigned by the `systhread_newkey` function.

Syntax

```
void systhread_setdata(int key, void *data);
```

Returns

void

Parameters

int `key` is the priority of the thread.

void `*data` is the pointer to the string of data to be associated with the value of `key`.

See Also

[systhread_current](#), [systhread_getdata](#), [systhread_newkey](#),
[systhread_sleep](#), [systhread_start](#), [systhread_timeriset](#)

systhread_sleep

The `systhread_sleep` function puts the calling thread to sleep for a given time.

Syntax

```
void systhread_sleep(int milliseconds);
```

Returns

void

Parameters

int milliseconds is the number of milliseconds the thread is to sleep.

See Also

[systhread_current](#), [systhread_getdata](#), [systhread_newkey](#),
[systhread_setdata](#), [systhread_start](#), [systhread_timerset](#)

systhread_start

The `systhread_start` function creates a thread with the given priority, allocates a stack of a specified number of bytes, and calls a specified function with a specified argument.

Syntax

```
SYS_THREAD systhread_start(int prio, int stksz,  
    void (*fn)(void *), void *arg);
```

Returns

A new `SYS_THREAD` pointer if the call succeeded, or the constant `SYS_THREAD_ERROR` if the call did not succeed.

Parameters

int prio is the priority of the thread. Priorities are system-dependent.

int stksz is the stack size in bytes. If stksz is zero (0), the function allocates a default size.

void (*fn)(void *) is the function to call.

void *arg is the argument for the fn function.

See Also

[systhread_current](#), [systhread_getdata](#), [systhread_newkey](#),
[systhread_setdata](#), [systhread_sleep](#), [systhread_timerset](#)

systhread_terminate

The `systhread_terminate` function terminates a specified thread.

Syntax

```
#include <base/systhr.h>
void systhread_terminate(SYS_THREAD thr);
```

Returns

void

Parameters

`SYS_THREAD thr` is the thread to terminate.

See also

[systhread_current](#), [systhread_getdata](#), [systhread_newkey](#),
[systhread_setdata](#), [systhread_sleep](#), [systhread_start](#), [systhread_timeriset](#)

systhread_timeriset

The `systhread_timeriset` function starts or resets the interrupt timer interval for a thread system.

Because most systems don't allow the timer interval to be changed, this should be considered a suggestion, rather than a command.

Syntax

```
void systhread_timeriset(int usec);
```

Returns

void

Parameters

`int usec` is the time, in microseconds

See Also

[systhread_current](#), [systhread_getdata](#), [systhread_newkey](#),
[systhread_setdata](#), [systhread_sleep](#), [systhread_start](#)

USE_NSAPI_VERSION

Plugin developers can define the `USE_NSAPI_VERSION` macro before including the `nsapi.h` header file to request a particular version of NSAPI. The requested NSAPI version is encoded by multiplying the major version number by 100 and then adding this to the minor version number. For example, the following code requests NSAPI 3.2 features:

```
#define USE_NSAPI_VERSION 302 /* We want NSAPI 3.2 (Web Server 6.1) */
#include "nsapi.h"
```

To develop a plugin that is compatible across multiple server versions, define `USE_NSAPI_VERSION` to the highest NSAPI version supported by all of the target server versions.

The following table lists server versions and the highest NSAPI version supported by each:

Table 1-2 NSAPI Versions Supported by Different Servers

Server Version	NSAPI Version
iPlanet Web Server 4.1	3.0
iPlanet Web Server 6.0	3.1
Netscape Enterprise Server 6.0	3.1
Netscape Enterprise Server 6.1	3.1
Sun ONE Application Server 7.0	3.1
Sun Java System Web Server 6.1	3.2
Sun Java System Web Proxy Server 4	3.3

It is an error to request a version of NSAPI higher than the highest version supported by the `nsapi.h` header that the plugin is being compiled against. Additionally, to use `USE_NSAPI_VERSION`, you must compile against an `nsapi.h` header file that supports NSAPI 3.3 or higher.

Syntax

```
int USE_NSAPI_VERSION
```

Example

The following code can be used when building a plugin designed to work with Sun Java System Web Proxy Server 4:

```
#define USE_NSAPI_VERSION 303 /* We want NSAPI 3.3 (Proxy Server 4) */
#include "nsapi.h"
```

See Also

[NSAPI_RUNTIME_VERSION](#), [NSAPI_VERSION](#)

util_can_exec

UNIX Only

The `util_can_exec` function checks that a specified file can be executed, returning either a 1 (executable) or a 0. The function checks if the file can be executed by the user with the given user and group ID.

Use this function before executing a program using the `exec` system call.

Syntax

```
int util_can_exec(struct stat *finfo, uid_t uid, gid_t gid);
```

Returns

1 if the file is executable, or 0 if the file is not executable.

Parameters

`stat *finfo` is the `stat` structure associated with a file.

`uid_t uid` is the UNIX user id.

`gid_t gid` is the UNIX group id. Together with `uid`, this determines the permissions of the UNIX user.

See Also

[util_env_create](#), [util_getline](#), [util_hostname](#)

util_chdir2path

The `util_chdir2path` function changes the current directory to a specified directory, where you will access a file.

When running under Windows, use a critical section to ensure that more than one thread does not call this function at the same time.

Use `util_chdir2path` when you want to make file access a little quicker, because you do not need to use a full path.

Syntax

```
int util_chdir2path(char *path);
```

Returns

0 if the directory was changed, or -1 if the directory could not be changed.

Parameters

`char *path` is the name of a directory.

The parameter must be a writable string because it isn't permanently modified.

util_cookie_find

The `util_cookie_find` function finds a specific cookie in a cookie string and returns its value.

Syntax

```
char *util_cookie_find(char *cookie, char *name);
```

Returns

If successful, returns a pointer to the NULL-terminated value of the cookie. Otherwise, returns NULL. This function modifies the cookie string parameter by null-terminating the name and value.

Parameters

`char *cookie` is the value of the `Cookie:` request header.

`char *name` is the name of the cookie whose value is to be retrieved.

util_does_process_exist

The `util_does_process_exist` function verifies that a given process ID is that of an executing process.

Syntax

```
#include <libproxy/util.h>
int util_does_process_exist (int pid)
```

Returns

- nonzero if the *pid* represents an executing process
- 0 if the *pid* does not represent an executing process

Parameters

int *pid* is the process ID to be tested.

See also

[util_url_fix_host_name](#), [util_uri_check](#)

util_env_create

The `util_env_create` function creates and allocates the environment specified by *env*, returning a pointer to the environment. If the parameter *env* is NULL, the function allocates a new environment. Use `util_env_create` to create an environment when executing a new program.

Syntax

```
#include <base/util.h>
char **util_env_create(char **env, int n, int *pos);
```

Returns

A pointer to an environment.

Parameters

char ***env* is the existing environment or NULL.

int *n* is the maximum number of environment entries that you want in the environment.

int **pos* is an integer that keeps track of the number of entries used in the environment.

See also

[util_env_replace](#), [util_env_str](#), [util_env_free](#), [util_env_find](#)

util_env_find

The `util_env_find` function locates the string denoted by a name in a specified environment and returns the associated value. Use this function to find an entry in an environment.

Syntax

```
char *util_env_find(char **env, char *name);
```

Returns

The value of the environment variable if it is found, or NULL if the string was not found.

Parameters

`char **env` is the environment.

`char *name` is the name of an environment variable in `env`.

See Also

[util_env_replace](#), [util_env_str](#), [util_env_free](#), [util_env_create](#)

util_env_free

The `util_env_free` function frees a specified environment. Use this function to deallocate an environment you created using the function `util_env_create`.

Syntax

```
void util_env_free(char **env);
```

Returns

void

Parameters

`char **env` is the environment to be freed.

See Also

[util_env_replace](#), [util_env_str](#), [util_env_create](#), [util_env_find](#)

util_env_replace

The `util_env_replace` function replaces the occurrence of the variable denoted by a name in a specified environment with a specified value. Use this function to change the value of a setting in an environment.

Syntax

```
void util_env_replace(char **env, char *name, char *value);
```

Returns

void

Parameters

`char **env` is the environment.

`char *name` is the name of a name-value pair.

`char *value` is the new value to be stored.

See Also

[util_env_str](#), [util_env_free](#), [util_env_find](#), [util_env_create](#)

util_env_str

The `util_env_str` function creates an environment entry and returns it. This function does not check for nonalphanumeric symbols in the name (such as the equal sign “=”). You can use this function to create a new environment entry.

Syntax

```
char *util_env_str(char *name, char *value);
```

Returns

A newly allocated string containing the name-value pair.

Parameters

`char *name` is the name of a name-value pair.

`char *value` is the new value to be stored.

See Also

[util_env_replace](#), [util_env_free](#), [util_env_create](#), [util_env_find](#)

util_get_current_gmt

The `util_get_current_gmt` function obtains the current time, represented in terms of GMT (Greenwich Mean Time).

Syntax

```
#include <libproxy/util.h>
time_t util_get_current_gmt(void);
```

Returns

the current GMT

Parameters

No parameter is required.

See also

[util_make_local](#)

util_get_int_from_aux_file

The `util_get_int_from_aux_file` function is used to get a single line from a specified file and return it in the form of an integer. This is a utility for storing single numbers in a file.

Syntax

```
#include <libproxy/cutil.h>
int util_get_int_from_file(char *root, char *name);
```

Returns

an integer from the file.

Parameters

char *root is the name of the directory containing the file to be read.

char *name is the name of the file to be read.

See also

[util_get_long_from_aux_file](#), [util_get_string_from_aux_file](#),
[util_get_int_from_file](#), [util_get_long_from_file](#),
[util_get_string_from_file](#), [util_put_int_to_file](#), [util_put_long_to_file](#),
[util_put_string_to_aux_file](#), [util_put_string_to_file](#)

util_get_int_from_file

The `util_get_int_from_file` function is used to get a single line from a specified file and return it in the form of an integer. This is a utility for storing single numbers in a file.

Syntax

```
#include <libproxy/cutil.h>
int util_get_int_from_file(char *filename);
```

Returns

- an integer from the file.
- -1 if no value was obtained from the file.

Parameters

char **filename* is the name of the file to be read.

See also

[util_get_long_from_file](#), [util_get_string_from_file](#), [util_put_int_to_file](#), [util_put_long_to_file](#), [util_put_string_to_file](#)

util_get_long_from_aux_file

The `util_get_long_from_file` function is used to get a single line from a specified file and return it in the form of a long number. This is a utility for storing single long numbers in a file.

Syntax

```
#include <libproxy/cutil.h>
long util_get_long_from_file(char *root, char *name);
```

Returns

a long integer from the file.

Parameters

char **root* is the name of the directory containing the file to be read.

char **name* is the name of the file to be read.

See also

[util_get_int_from_aux_file](#), [util_get_string_from_aux_file](#),
[util_get_int_from_file](#), [util_get_long_from_file](#),
[util_get_string_from_file](#), [util_put_int_to_file](#), [util_put_long_to_file](#),
[util_put_string_to_aux_file](#), [util_put_string_to_file](#)

util_get_long_from_file

The `util_get_long_from_file` function is used to get a single line from a specified file and return it in the form of a long number. This is a utility for storing single long numbers in a file.

Syntax

```
#include <libproxy/cutil.h>
long util_get_long_from_file(char *filename);
```

Returns

- a long integer from the file.
- -1 if no value was obtained from the file.

Parameters

char **file* is the name of the file to be read.

See also

[util_get_int_from_file](#), [util_get_string_from_file](#), [util_put_int_to_file](#),
[util_put_long_to_file](#), [util_put_string_to_file](#)

util_get_string_from_aux_file

The `util_get_string_from_aux_file` function is used to get a single line from a specified file and return it in the form of a word. This is a utility for storing single words in a file.

Syntax

```
#include <libproxy/cutil.h>
char *util_get_string_from_file(char *root, char *name, char *buf,
int maxsize);
```

Returns

a string containing the next line from the file.

Parameters

char **root* is the name of the directory containing the file to be read.

char **name* is the name of the file to be read.

char **buf* is the string to use as the file buffer.

int *maxsize* is the maximum size for the file buffer.

See also

[util_get_int_from_aux_file](#), [util_get_long_from_aux_file](#),
[util_get_int_from_file](#), [util_get_long_from_file](#),
[util_get_string_from_file](#), [util_put_int_to_file](#), [util_put_long_to_file](#),
[util_put_string_to_aux_file](#), [util_put_string_to_file](#)

util_get_string_from_file

The `util_get_string_from_file` function is used to get a single line from a specified file and return it in the form of a word. This is a utility for storing single words in a file.

Syntax

```
#include <libproxy/cutil.h>
char *util_get_string_from_file(char *filename, char *buf, int maxsize);
```

Returns

- a string containing the next line from the file.
- NULL if no string was obtained.

Parameters

char **file* is the name of the file to be read.

char **buf* is the string to use as the file buffer.

int *maxsize* is the maximum size for the file buffer.

See also

[util_get_int_from_file](#), [util_get_long_from_file](#), [util_put_int_to_file](#),
[util_put_long_to_file](#), [util_put_string_to_file](#)

util_getline

The `util_getline` function scans the specified file buffer to find a line feed or carriage return/line feed terminated string. The string is copied into the specified buffer, and NULL-terminates it. The function returns a value that indicates whether the operation stored a string in the buffer, encountered an error, or reached the end of the file.

Use this function to scan lines out of a text file, such as a configuration file.

Syntax

```
int util_getline(filebuf *buf, int lineno, int maxlen, char *line);
```

Returns

0 if successful; `line` contains the string.

1 if the end of file was reached; `line` contains the string.

-1 if an error occurred; `line` contains a description of the error.

Parameters

`filebuf *buf` is the file buffer to be scanned.

`int lineno` is used to include the line number in the error message when an error occurs. The caller is responsible for making sure the line number is accurate.

`int maxlen` is the maximum number of characters that can be written into `l`.

`char *l` is the buffer in which to store the string. The user is responsible for allocating and deallocating `line`.

See Also

[util_can_exec](#), [util_env_create](#), [util_hostname](#)

util_hostname

The `util_hostname` function retrieves the local host name and returns it as a string. If the function cannot find a fully-qualified domain name, it returns NULL. You may reallocate or free this string. Use this function to determine the name of the system you are on.

Syntax

```
char *util_hostname(void);
```

Returns

If a fully-qualified domain name was found, returns a string containing that name; otherwise, returns NULL if the fully-qualified domain name was not found.

Parameters

none

util_is_mozilla

The `util_is_mozilla` function checks whether a specified user-agent header string is a Netscape browser of at least a specified revision level, returning a 1 if it is, and 0 otherwise. It uses strings to specify the revision level to avoid ambiguities such as 1.56 > 1.5.

Syntax

```
int util_is_mozilla(char *ua, char *major, char *minor);
```

Returns

1 if the user-agent is a Netscape browser, or 0 if the user-agent is not a Netscape browser.

Parameters

char *ua is the user-agent string from the request headers.

char *major is the major release number (to the left of the decimal point).

char *minor is the minor release number (to the right of the decimal point).

See Also

[util_is_url](#), [util_later_than](#)

util_is_url

The `util_is_url` function checks whether a string is a URL, returning 1 if it is and 0 otherwise. The string is a URL if it begins with alphabetic characters followed by a colon (:).

Syntax

```
int util_is_url(char *url);
```

Returns

1 if the string specified by `url` is a URL, or 0 if the string specified by `url` is not a URL.

Parameters

`char *url` is the string to be examined.

See Also

[util_is_mozilla](#), [util_later_than](#)

util_itoa

The `util_itoa` function converts a specified integer to a string, and returns the length of the string. Use this function to create a textual representation of a number.

Syntax

```
int util_itoa(int i, char *a);
```

Returns

The length of the string created.

Parameters

`int i` is the integer to be converted.

`char *a` is the ASCII string that represents the value. The user is responsible for the allocation and deallocation of `a`, and it should be at least 32 bytes long.

util_later_than

The `util_later_than` function compares the date specified in a time structure against a date specified in a string. If the date in the string is later than or equal to the one in the time structure, the function returns 1. Use this function to handle RFC 822, RFC 850, and ctime formats.

Syntax

```
int util_later_than(struct tm *lms, char *ims);
```

Returns

1 if the date represented by `ims` is the same as or later than that represented by the `lms`, or 0 if the date represented by `ims` is earlier than that represented by the `lms`.

Parameters

`tm *lms` is the time structure containing a date.

`char *ims` is the string containing a date.

See Also

[util_strftime](#)

util_make_filename

The `util_make_filename` function concatenates a directory name and a filename into a newly created string. This can be handy when you are dealing with a number of files that all go to the same directory.

Syntax

```
#include <libproxy/cutil.h>
char *util_make_filename(char *root, char *name);
```

Returns

A new string containing the directory name concatenated with the filename.

Parameters

`char *root` is a string containing the directory name.

`char *name` is a string containing the filename.

util_make_gmt

The `util_make_gmt` function converts a given local time to GMT (Greenwich Mean Time), or obtains the current GMT.

Syntax

```
#include <libproxy/util.h>
time_t util_make_gmt(time_t t);
```

Returns

- the GMT equivalent to the local time `t`, if `t` is not 0
- the current GMT if `t` is 0

Parameters

`time_t t` is a time.

See also[util_make_local](#)

util_make_local

The `util_make_local` function converts a given GMT to local time.

Syntax

```
#include <libproxy/util.h>
time_t util_make_local(time_t t);
```

Returns

the local equivalent to the GMT *t*

Parameters

time_t *t* is a time.

See also[util_make_gmt](#)

util_move_dir

The `util_move_dir` function moves a directory, preserving permissions, creation times, and last-access times. It attempts to do this by renaming, but if that fails (for example, if the source and destination are on two different file systems), it copies the directory.

Syntax

```
#include <libproxy/util.h>
int util_move_dir (char *src, char *dst);
```

Returns

- 0 if the move failed
- nonzero if the move succeeded

Parameters

`char *src` is the fully qualified name of the source directory.

`char *dst` is the fully qualified name of the destination directory.

See also[util_move_file](#)

util_move_file

The `util_move_dir` function moves a file, preserving permissions, creation time, and last-access time. It attempts to do this by renaming, but if that fails (for example, if the source and destination are on two different file systems), it copies the file.

Syntax

```
#include <libproxy/util.h>
int util_move_file (char *src, char *dst);
```

Returns

- 0 if the move failed
- nonzero if the move succeeded

Parameters

`char *src` is the fully qualified name of the source file.

`char *dst` is the fully qualified name of the destination file.

See also[util_move_dir](#)

util_parse_http_time

The `util_parse_http_time` function converts a given HTTP time string to `time_t` format.

Syntax

```
#include <libproxy/util.h>
time_t util_parse_http_time(char *date_string);
```

Returns

the `time_t` equivalent to the GMT `t`

Parameters

`time_t t` is a time.

See also

[util_make_gmt](#)

util_put_int_to_file

The `util_put_int_to_file` function writes a single line containing an integer to a specified file.

Syntax

```
#include <libproxy/cutil.h>
int util_put_int_to_file(char *filename, int i);
```

Returns

- nonzero if the operation succeeded
- 0 if the operation failed

Parameters

char **file* is the name of the file to be written.

int *i* is the integer to write.

See also

[util_get_int_from_file](#), [util_get_long_from_file](#), [util_put_long_to_file](#),
[util_put_string_to_file](#)

util_put_long_to_file

The `util_put_long_to_file` function writes a single line containing a long integer to a specified file.

Syntax

```
#include <libproxy/cutil.h>
int util_put_long_to_file(char *filename, long l);
```

Returns

- nonzero if the operation succeeded
- 0 if the operation failed

Parameters

char **file* is the name of the file to be written.

long *l* is the long integer to write.

See also

[util_get_int_from_file](#), [util_get_long_from_file](#), [util_put_int_to_file](#),
[util_put_string_to_file](#)

util_put_string_to_aux_file

The `util_put_string_to_aux_file` function writes a single line containing a string to a file specified by directory name and file name.

Syntax

```
#include <libproxy/cutil.h>
int util_put_string_to_aux_file(char *root, char *name, char *str);
```

Returns

- non-zero if the operation succeeded
- 0 if the operation failed

Parameters

char **root* is the name of the directory where the file is to be written.

char **name* is the name of the file is to be written.

char **str* is the string to write.

See also

[util_get_int_from_file](#), [util_get_long_from_file](#), [util_put_int_to_file](#),
[util_put_long_to_file](#), [util_put_string_to_file](#)

util_put_string_to_file

The `util_put_string_to_file` function writes a single line containing a string to a specified file.

Syntax

```
#include <libproxy/cutil.h>
int util_put_string_to_file(char *filename, char *str);
```

Returns

- nonzero if the operation succeeded

- 0 if the operation failed

Parameters

char **file* is the name of the file to be read.

char **str* is the string to write.

See also

[util_get_int_from_file](#), [util_get_long_from_file](#), [util_put_int_to_file](#),
[util_put_long_to_file](#)

util_sect_id

The `util_sect_id` function creates a section ID from the section `dim` and an index.

Syntax

```
#include <libproxy/cutil.h>
void util_sect_id(int dim, int idx, char *buf);
```

Returns

- nonzero if the operation succeeded
- 0 if the operation failed

Parameters

int *dim* is the section dim.

int *idx* is the index.

char **buf* is the buffer to receive the section ID.

util_sh_escape

The `util_sh_escape` function parses a specified string and places a backslash (\) in front of any shell-special characters, returning the resultant string. Use this function to ensure that strings from clients won't cause a shell to do anything unexpected.

The shell-special characters are the space plus the following characters:

```
& ; ` ' " | * ? ~ < > ^ ( ) [ ] { } $ \ # !
```

Syntax

```
char *util_sh_escape(char *s);
```

Returns

A newly allocated string.

Parameters

char *s is the string to be parsed.

See Also

[util_uri_escape](#)

util_snprintf

The `util_snprintf` function formats a specified string, using a specified format, into a specified buffer using the `printf`-style syntax and performs bounds checking. It returns the number of characters in the formatted buffer.

For more information, see the documentation on the `printf` function for the runtime library of your compiler.

Syntax

```
int util_snprintf(char *s, int n, char *fmt, ...);
```

Returns

The number of characters formatted into the buffer.

Parameters

char *s is the buffer to receive the formatted string.

int n is the maximum number of bytes allowed to be copied.

char *fmt is the format string. The function handles only `%d` and `%s` strings; it does not handle any width or precision strings.

... represents a sequence of parameters for the `printf` function.

See Also

[util_sprintf](#), [util_vsnprintf](#), [util_vsprintf](#)

util_sprintf

The `util_sprintf` function formats a specified string, using a specified format, into a specified buffer, using the `printf`-style syntax without bounds checking. It returns the number of characters in the formatted buffer.

Because `util_sprintf` doesn't perform bounds checking, use this function only if you are certain that the string fits the buffer. Otherwise, use the function `util_snprintf`. For more information, see the documentation on the `printf` function for the runtime library of your compiler.

Syntax

```
int util_sprintf(char *s, char *fmt, ...);
```

Returns

The number of characters formatted into the buffer.

Parameters

`char *s` is the buffer to receive the formatted string.

`char *fmt` is the format string. The function handles only `%d` and `%s` strings; it does not handle any width or precision strings.

`...` represents a sequence of parameters for the `printf` function.

Example

```
char *logmsg;
int len;
logmsg = (char *) MALLOC(256);
len = util_sprintf(logmsg, "%s %s %s\n", ip, method, uri);
```

See Also

[util_snprintf](#), [util_vsnprintf](#), [util_vsprintf](#)

util_strcasecmp

The `util_strcasecmp` function performs a comparison of two alphanumeric strings and returns a -1, 0, or 1 to signal which is larger or that they are identical.

The comparison is not case-sensitive.

Syntax

```
int util_strcasecmp(const char *s1, const char *s2);
```

Returns

1 if `s1` is greater than `s2`.

0 if `s1` is equal to `s2`.

-1 if `s1` is less than `s2`.

Parameters

`char *s1` is the first string.

`char *s2` is the second string.

See Also

[util_strncasecmp](#)

util_strftime

The `util_strftime` function translates a `tm` structure, which is a structure describing a system time, into a textual representation. It is a thread-safe version of the standard `strftime` function

Syntax

```
int util_strftime(char *s, const char *format, const struct tm *t);
```

Returns

The number of characters placed into `s`, not counting the terminating NULL character.

Parameters

`char *s` is the string buffer to put the text into. There is no bounds checking, so you must make sure that your buffer is large enough for the text of the date.

`const char *format` is a format string, a bit like a `printf` string in that it consists of text with certain `%x` substrings. You may use the constant `HTTP_DATE_FMT` to create date strings in the standard Internet format. For more information, see the documentation on the `printf` function for the runtime library of your compiler. Refer to [Time Formats Chapter 3, “Time Formats,”](#) for details on time formats.

`const struct tm *t` is a pointer to a calendar time (`tm`) structure, usually created by the function `system_localtime` or `system_gmtime`.

See Also

`system_localtime`, `system_gmtime`

util_strncasecmp

The `util_strncasecmp` function performs a comparison of the first `n` characters in the alphanumeric strings and returns a `-1`, `0`, or `1` to signal which is larger or that they are identical.

The function's comparison is not case-sensitive.

Syntax

```
int util_strncasecmp(const char *s1, const char *s2, int n);
```

Returns

`1` if `s1` is greater than `s2`.

`0` if `s1` is equal to `s2`.

`-1` if `s1` is less than `s2`.

Parameters

`char *s1` is the first string.

`char *s2` is the second string.

`int n` is the number of initial characters to compare.

See Also

[util_strcmp](#)

util_uri_check

The `util_uri_check` function checks that a URI has a format conforming to the standard.

At present, the only URI it checks for is a URL. The standard format for a URL is

protocol://user:password@host:port/url-path

where *user:password*, *:password*, *:port*, or */url-path* can be omitted.

Syntax

```
#include <libproxy/util.h>
int util_uri_check (char *uri);
```

Returns

- `0` if the URI does not have the proper form.

- nonzero if the URI has the proper form.

Parameters

char **uri* is the URI to be tested.

util_uri_escape

The `util_uri_escape` function converts any special characters in the URI into the URI format (`%XX`, where `XX` is the hexadecimal equivalent of the ASCII character), and returns the escaped string. The special characters are `?:#:+&*"<>`, space, carriage return, and line feed.

Use `util_uri_escape` before sending a URI back to the client.

Syntax

```
char *util_uri_escape(char *d, char *s);
```

Returns

The string (possibly newly allocated) with escaped characters replaced.

Parameters

char **d* is a string. If *d* is not NULL, the function copies the formatted string into *d* and returns it. If *d* is NULL, the function allocates a properly sized string and copies the formatted special characters into the new string, then returns it.

The `util_uri_escape` function does not check bounds for the parameter *d*. Therefore, if *d* is not NULL, it should be at least three times as large as the string *s*.

char **s* is the string containing the original unescaped URI.

See Also

[util_uri_is_evil](#), [util_uri_parse](#), [util_uri_unescape](#)

util_uri_is_evil

The `util_uri_is_evil` function checks a specified URI for insecure path characters. Insecure path characters include `//, /./, /../` and `./, ..` (also for Windows `.`) at the end of the URI. Use this function to see if a URI requested by the client is insecure.

Syntax

```
int util_uri_is_evil(char *t);
```

Returns

1 if the URI is insecure, or 0 if the URI is OK.

Parameters

char *t is the URI to be checked.

See Also

[util_uri_escape](#), [util_uri_parse](#)

util_uri_parse

The `util_uri_parse` function converts `//`, `./`, and `/*./` into `/` in the specified URI (where `*` is any character other than `/`). You can use this function to convert a URI's bad sequences into valid ones. First use the function `util_uri_is_evil` to determine whether the function has a bad sequence.

Syntax

```
void util_uri_parse(char *uri);
```

Returns

void

Parameters

char *uri is the URI to be converted.

See Also

[util_uri_is_evil](#), [util_uri_unescape](#)

util_uri_unescape

The `util_uri_unescape` function converts the encoded characters of a URI into their ASCII equivalents. Encoded characters appear as `%XX`, where `XX` is a hexadecimal equivalent of the character.

NOTE You cannot use an embedded null in a string, because NSAPI functions assume that a null is the end of the string. Therefore, passing unicode-encoded content through an NSAPI plugin doesn't work.

Syntax

```
void util_uri_unescape(char *uri);
```

Returns

void

Parameters

char *uri is the URI to be converted.

See Also

[util_uri_escape](#), [util_uri_is_evil](#), [util_uri_parse](#)

util_url_cmp

The `util_url_cmp` function compares two URLs. It is analogous to the `strcmp()` library function of C.

Syntax

```
#include <libproxy/util.h>
int util_url_cmp (char *s1, char *s2);
```

Returns

- -1 if the first URL, *s1*, is less than the second, *s2*
- 0 if they are identical
- 1 if the first URL, *s1*, is greater than the second, *s2*

Parameters

char *s1 is the first URL to be tested.

char *s2 is the second URL to be tested.

See also

[util_url_fix_host_name](#), [util_uri_check](#)

util_url_fix_host name

The `util_url_fix_host name` function converts the host name in a URL to lowercase and removes redundant port numbers.

Syntax

```
#include <libproxy/util.h>
void util_url_fix_host name(char *url);
```

Returns

void (but changes the value of its parameter string)

The protocol specifier and the host name in the parameter string are changed to lowercase. The function also removes redundant port numbers, such as 80 for HTTP, 70 for gopher, and 21 for FTP.

Parameters

char **url* is the URL to be converted.

See also

[util_url_cmp](#), [util_uri_check](#)

util_url_has_FQDN

The `util_url_has_FQDN` function returns a value to indicate whether a specified URL references a fully qualified domain name.

Syntax

```
#include <libproxy/util.h>
int util_url_has_FQDN(char *url);
```

Returns

- 1 if the URL has a fully qualified domain name
- 0 if the URL does not have a fully qualified domain name

Parameters

char **url* is the URL to be examined.

util_vsnprintf

The `util_vsnprintf` function formats a specified string, using a specified format, into a specified buffer using the `vprintf`-style syntax and performs bounds checking. It returns the number of characters in the formatted buffer.

For more information, see the documentation on the `printf` function for the runtime library of your compiler.

Syntax

```
int util_vsnprintf(char *s, int n, register char *fmt, va_list args);
```

Returns

The number of characters formatted into the buffer.

Parameters

`char *s` is the buffer to receive the formatted string.

`int n` is the maximum number of bytes allowed to be copied.

`register char *fmt` is the format string. The function handles only `%d` and `%s` strings; it does not handle any width or precision strings.

`va_list args` is an STD argument variable obtained from a previous call to `va_start`.

See Also

[util_snprintf](#), [util_vsprintf](#)

util_vsprintf

The `util_vsprintf` function formats a specified string, using a specified format, into a specified buffer using the `vprintf`-style syntax without bounds checking. It returns the number of characters in the formatted buffer.

For more information, see the documentation on the `printf` function for the runtime library of your compiler.

Syntax

```
int util_vsprintf(char *s, register char *fmt, va_list args);
```

Returns

The number of characters formatted into the buffer.

Parameters

`char *s` is the buffer to receive the formatted string.

`register char *fmt` is the format string. The function handles only `%d` and `%s` strings; it does not handle any width or precision strings.

`va_list args` is an STD argument variable obtained from a previous call to `va_start`.

See Also

[util_snprintf](#), [util_vsnprintf](#)

W

write

The `write` filter method is called when output data is to be sent. Filters that modify or consume outgoing data should implement the `write` filter method.

Upon receiving control, a write implementation should first process the data as necessary, and then pass it on to the next filter layer; for example, by calling `net_write(layer->lower, ...)`. If the filter buffers outgoing data, it should implement the `flush` filter method.

Syntax

```
int write(FilterLayer *layer, const void *buf, int amount);
```

Returns

The number of bytes consumed, which may be less than the requested amount if an error occurred.

Parameters

`FilterLayer *layer` is the filter layer in which the filter is installed.

`const void *buf` is the buffer that contains the outgoing data.

`int amount` is the number of bytes in the buffer.

Example

```
int myfilter_write(FilterLayer *layer, const void *buf, int amount)
{
    return net_write(layer->lower, buf, amount);
}
```

See Also

[flush](#), [net_write](#), [writev](#)

writev

The `writev` filter method is called when multiple buffers of output data are to be sent. Filters that modify or consume outgoing data may choose to implement the `writev` filter method.

If a filter implements the `write` filter method but not the `writev` filter method, the server automatically translates `net_writev` calls to `net_write` calls. As a result, filters interested in the outgoing data stream do not need to implement the `writev` filter method. However, for performance reasons, it is beneficial for filters that implement the `write` filter method to also implement the `writev` filter method.

Syntax

```
int writev(FilterLayer *layer, const struct iovec *iov, int iov_size);
```

Returns

The number of bytes consumed, which may be less than the requested amount if an error occurred.

Parameters

`FilterLayer *layer` is the filter layer the filter is installed in.

`const struct iovec *iov` is an array of `iovec` structures, each of which contains outgoing data.

`int iov_size` is the number of `iovec` structures in the `iov` array.

Example

```
int myfilter_writev(FilterLayer *layer, const struct iovec *iov, int
iov_size)
{
    return net_writev(layer->lower, iov, iov_size);
}
```

See Also

[flush](#), [net_write](#), [write](#)

Data Structure Reference

NSAPI uses many data structures that are defined in the `nsapi.h` header file, which is in the directory `server-root/plugins/include`.

The NSAPI functions described in [Chapter 1, “NSAPI Function Reference,”](#) provide access to most of the data structures and data fields. Before directly accessing a data structure in `nsapi.h`, check to see if an accessor function exists for it.

For information about the privatization of some data structures in Sun Java System Web Proxy Server 4, see [“Privatization of Some Data Structures” on page 130](#)

The rest of this chapter describes public data structures in `nsapi.h`. Note that data structures in `nsapi.h` that are not described in this chapter are considered private and may change incompatibly in future releases.

This chapter has the following sections:

- [Privatization of Some Data Structures](#)
- [Session](#)
- [pblock](#)
- [pb_entry](#)
- [pb_param](#)
- [Session->client](#)
- [Request](#)
- [stat](#)
- [shmem_s](#)
- [cinfo](#)

- [sendfiledata](#)
- [Filter](#)
- [FilterContext](#)
- [FilterLayer](#)
- [FilterMethods](#)
- [The CacheEntry Data Structure](#)
- [The CacheState Data Structure](#)
- [The ConnectMode Data Structure](#)

Privatization of Some Data Structures

The data structures in `nsapi_pvt.h` are now considered to be private data structures, and you should not write code that accesses them directly. Instead, use accessor functions. We expect that very few people have written plugins that access these data structures directly, so this change should have very little impact on customer-defined plugins. Look in `nsapi_pvt.h` to see which data structures have been removed from the public domain, and to see the accessor functions you can use to access them from now on.

Plugins written for Enterprise Server 3.x that access contents of data structures defined in `nsapi_pvt.h` will not be source compatible with Sun Java System Web Proxy Server 4, that is, it will be necessary to `#include "nsapi_pvt.h"` to build such plugins from source. There is also a small chance that these programs will not be binary compatible with Sun Java System Web Proxy Server 4, because some of the data structures in `nsapi_pvt.h` have changed size. In particular, the `directive` structure is larger, which means that a plugin that indexes through the directives in a `dtable` will not work without being rebuilt (with `nsapi_pvt.h` included).

We hope that the majority of plugins do not reference the internals of data structures in `nsapi_pvt.h`, and therefore that most existing NSAPI plugins will be both binary and source compatible with Sun Java System Web Proxy Server 4.

Plugins written for iPlanet Web Proxy Server 3.6 will not be binary compatible with Proxy Server 4. These plugins will have to be recompiled and relinked using Web Proxy Server 4's NSAPI header files and libraries.

Session

A session is the time between the opening and closing of the connection between the client and the server. The `session` data structure holds variables that apply session wide, regardless of the requests being sent, as shown here:

```
typedef struct {
    /* Information about the remote client */
    pblock *client;

    /* The socket descriptor to the remote client */
    SYS_NETFD csd;

    /* The input buffer for that socket descriptor */
    netbuf *inbuf;

    /* Raw socket information about the remote */
    /* client (for internal use) */
    struct in_addr iaddr;
} Session;
```

pblock

The parameter block is the hash table that holds `pb_entry` structures. Its contents are transparent to most code. This data structure is frequently used in NSAPI; it provides the basic mechanism for packaging up parameters and values. There are many functions for creating and managing parameter blocks, and for extracting, adding, and deleting entries. See the functions whose names start with `pblock_` in [Chapter 1, “NSAPI Function Reference” on page 13](#). You should not need to write code that accesses `pblock` data fields directly.

```
typedef struct {
    int hsize;
    struct pb_entry **ht;
} pblock;
```

pb_entry

The `pb_entry` is a single element in the parameter block.

```
struct pb_entry {
    pb_param *param;
    struct pb_entry *next;
};
```

pb_param

The `pb_param` represents a name-value pair, as stored in a `pb_entry`.

```
typedef struct {
    char *name, *value;
} pb_param;
```

Session->client

The `Session->client` parameter block structure contains two entries:

- The `ip` entry is the IP address of the client machine.
- The `dns` entry is the DNS name of the remote machine. This member must be accessed through the `session_dns` function call:

```
/*
 * session_dns returns the DNS host name of the client for this
 * session and inserts it into the client pblock. Returns NULL if
 * unavailable.
 */
char *session_dns(Session *sn);
```

Request

Under HTTP protocol, there is only one request per session. The `request` structure contains the variables that apply to the request in that session (for example, the variables include the client's HTTP headers).

```
typedef struct {
    /* Server working variables */
    pblock *vars;

    /* The method, URI, and protocol revision of this request */
    block *reqpb;

    /* Protocol specific headers */
    int loadhdrs;
    pblock *headers;

    /* Server's response headers */
    int senthdrs;
    pblock *srvhdrs;

    /* The object set constructed to fulfill this request */
    httpd_objset *os;
} Request;
```

stat

When a program calls the `stat()` function for a given file, the system returns a structure that provides information about the file. The specific details of the structure should be obtained from your platform's implementation, but the basic outline of the structure is as follows:

```

struct stat {
    dev_t      st_dev;      /* device of inode */
    ino_t      st_ino;     /* inode number */
    short      st_mode;    /* mode bits */
    short      st_nlink;   /* number of links to file */
    short      st_uid;     /* owner's user id */
    short      st_gid;     /* owner's group id */
    dev_t      st_rdev;    /* for special files */
    off_t      st_size;    /* file size in characters */
    time_t     st_atime;   /* time last accessed */
    time_t     st_mtime;   /* time last modified */
    time_t     st_ctime;   /* time inode last changed*/
}

```

The elements that are most significant for server plugin API activities are `st_size`, `st_atime`, `st_mtime`, and `st_ctime`.

shmem_s

```

typedef struct {
    void      *data;      /* the data */
    HANDLE     fdmap;
    int       size;      /* the maximum length of the data */
    char      *name;     /* internal use: filename to unlink if exposed */
    SYS_FILE  fd;       /* internal use: file descriptor for region */
} shmem_s;

```

cinfo

The `cinfo` data structure records the content information for a file.

```

typedef struct {
    char    *type;
            /* Identifies what kind of data is in the file*/
    char    *encoding;
            /* encoding identifies any compression or other /*
            /* content-independent transformation that's been /*
            /* applied to the file, such as uuencode)*/
    char    *language;
            /* Identifies the language a text document is in. */
} cinfo;

```

sendfiledata

The `sendfiledata` data structure is used to pass parameters to the `net_sendfile` function. It is also passed to the `sendfile` method in an installed filter in response to a `net_sendfile` call.

```

typedef struct {
    SYS_FILE fd;           /* file to send */
    size_t offset;        /* offset in file to start sending from */
    size_t len;           /* number of bytes to send from file */
    const void *header;   /* data to send before file */
    int hlen;             /* number of bytes to send before file */
    const void *trailer;  /* data to send after file */
    int tlen;             /* number of bytes to send after file */
} sendfiledata;

```

Filter

The `Filter` data structure is an opaque representation of a filter. A `Filter` structure is created by calling `filter_create`.

```

typedef struct Filter Filter;

```

FilterContext

The `FilterContext` data structure stores context associated with a particular filter layer. Filter layers are created by calling `filter_insert`.

Filter developers may use the `data` member to store filter-specific context information.

```
typedef struct {
    pool_handle_t *pool; /* pool context was allocated from */
    Session *sn;         /* session being processed */
    Request *rq;         /* request being processed */
    void *data;          /* filter-defined private data */
} FilterContext;
```

FilterLayer

The `FilterLayer` data structure represents one layer in a filter stack. The `FilterLayer` structure identifies the filter installed at that layer and provides pointers to layer-specific context and a filter stack that represents the layer immediately below it in the filter stack.

```
typedef struct {
    Filter *filter; /* the filter at this layer in the filter stack */
    FilterContext *context; /* context for the filter */
    SYS_NETFD lower; /* access to the next filter layer in the stack */
} FilterLayer;
```

FilterMethods

The `FilterMethods` data structure is passed to `filter_create` to define the filter methods a filter supports. Each new `FilterMethods` instance must be initialized with the `FILTER_METHODS_INITIALIZER` macro. For each filter method a filter supports, the corresponding `FilterMethods` member should point to a function that implements that filter method.

```

typedef struct {
    size_t size;
    FilterInsertFunc *insert;
    FilterRemoveFunc *remove;
    FilterFlushFunc *flush;
    FilterReadFunc *read;
    FilterWriteFunc *write;
    FilterWritevFunc *writev;
    FilterSendfileFunc *sendfile;
} FilterMethods;

```

The CacheEntry Data Structure

The CacheEntry data structure holds all the information about one cache entry. It is created by the `ce_lookup` function and destroyed by the `ce_free` function. It is defined in the `libproxy/cache.h` file.

```

typedef struct _CacheEntry {
    CacheState state; /* state of the cache file; DO NOT refer to any
        * of the other fields in this C struct if state
        * is other than
        *     CACHE_REFRESH or
        *     CACHE_RETURN_FROM_CACHE
        */
    SYS_FILE fd_in; /* do not use: open cache file for reading */
    int fd_out; /* do not use: open (locked) cache file for writing */
    struct stat finfo; /* stat info for the cache file */
    unsigned char digest[CACHE_DIGEST_LEN]; /* MD5 for the URL */
    char * url_dig; /* URL used to for digest; field #8 in CIF */
    char * url_cif; /* URL read from CIF file */
    char * filename; /* Relative cache file name */
    char * dirname; /* Absolute cache directory name */
    char * absname; /* Absolute cache file path */
    char * lckname; /* Absolute locked cache file path */
    char * cifname; /* Absolute CIF path */
    int sect_idx; /* Cache section index */
    int part_idx; /* Cache partition index */
    CSect *section; /* Cache section that this file belongs to */
    CPart *partition; /* Cache partition that this file belongs to */
    int xfer_time; /* secs */ /* Field #2 in CIF */
    time_t last_modified; /* GMT */ /* Field #3 in CIF */
    time_t expires; /* GMT */ /* Field #4 in CIF */

```

```

time_t last_checked; /* GMT */ /* Field #5 in CIF */
long content_length; /* Field #6 in CIF */
char * content_type; /* Field #7 in CIF */
int is_auth; /* Authenticated data -- always do recheck */
int auth_sent; /* Client did send the Authorization header */
longmin_size; /* Min size for a cache file (in KB) */
longmax_size; /* Max size for a cache file (in KB) */
time_t last_accessed; /* GMT for proxy, local for gc */
time_t created; /* localtime (only used by gc, st_mtime) */
int removed; /* gc only; file was removed from disk */
long bytes; /* from stat(), using this we get hdr len */
long bytes_written; /* Number of bytes written to disk */
long bytes_in_media; /* real fs size taken up */
long blks; /* size in 512 byte blocks */
int category; /* Value category; bigger is better */
int cif_entry_ok; /* CIF entry found and ok */
time_t ims_c; /* GMT; Client -> proxy if-modified-since */
time_t start_time; /* Transfer start time */
int inhibit_caching; /* Bad expires/other reason not to cache */
int corrupt_cache_file; /* Cache file gone corrupt => remove */
int write_aborted; /* True if the cache file write was aborted */
int batch_update; /* We're doing batch update (no real user) */
char * cache_exclude; /* Hdrs not to write to cache (RE) */
char * cache_replace; /* Hdrs to replace with fresh ones from 304 response
(RE) */
char * cache_nomerge; /* Hdrs not to merge with the cached ones (RE) */
Session * sn;
Request * rq;
} CacheEntry;

```

The CacheState Data Structure

The CacheState data structure is actually an enumerated list of constants. Always use their names because values are subject to implementation change.

```

typedef enum {
CACHE_EXISTS_NOT = 0, /* Internal flag -- do not use! */
CACHE_EXISTS, /* Internal flag -- do not use! */
CACHE_NO, /* No caching: don't read, don't write cache */
CACHE_CREATE, /* Create cache; don't read */

```

```
CACHE_REFRESH,      /* Refresh cache; read if not modified */
CACHE_RETURN_FROM_CACHE, /* Return directly, no check */
CACHE_RETURN_ERROR /* With connect-mode=never when not in cache */
} CacheState;
```

The ConnectMode Data Structure

The ConnectMode data structure is actually an enumerated list of constants. Always use their names because values are subject to implementation change.

```
typedef enum {
CM_NORMAL = 0, /* normal -- retrieve/refresh when necessary */
CM_FAST_DEMO, /* fast -- retrieve only if not in cache already */
CM_NEVER     /* never -- never connect to network */
} ConnectMode;
```

The ConnectMode Data Structure

Time Formats

This chapter describes the format strings used for dates and times. These formats are used by the NSAPI function `util_strftime`, by some built-in SAFs such as `append-trailer`, and by server-parsed HTML (`parse-html`). The formats are similar to those used by the `strftime` C library routine, but not identical.

The following table describes the formats, listing the symbols and their meanings.

Table 3-1 Time Formats

Symbol	Meaning
%a	Abbreviated weekday name (3 chars)
%d	Day of month as decimal number (01-31)
%S	Second as decimal number (00-59)
%M	Minute as decimal number (00-59)
%H	Hour in 24-hour format (00-23)
%Y	Year with century, as decimal number, up to 2099
%b	Abbreviated month name (3 chars)
%h	Abbreviated month name (3 chars)
%T	Time "HH:MM:SS"
%X	Time "HH:MM:SS"
%A	Full weekday name
%B	Full month name
%C	"%a %b %e %H:%M:%S %Y"
%c	Date & time "%m/%d/%y %H:%M:%S"
%D	Date "%m/%d/%y"

Table 3-1 Time Formats

Symbol	Meaning
%e	Day of month as decimal number (1-31) without leading zeros
%I	Hour in 12-hour format (01-12)
%j	Day of year as decimal number (001-366)
%k	Hour in 24-hour format (0-23) without leading zeros
%l	Hour in 12-hour format (1-12) without leading zeros
%m	Month as decimal number (01-12)
%n	line feed
%p	A.M./P.M. indicator for 12-hour clock
%R	Time "%H:%M"
%r	Time "%I:%M:%S %p"
%t	tab
%U	Week of year as decimal number, with Sunday as first day of week (00-51)
%w	Weekday as decimal number (0-6; Sunday is 0)
%W	Week of year as decimal number, with Monday as first day of week (00-51)
%x	Date "%m/%d/%y"
%y	Year without century, as decimal number (00-99)
%%	Percent sign

Alphabetical List of NSAPI Functions and Macros

This appendix provides an alphabetical list for the easy lookup of NSAPI functions and macros.

C

`cache_digest`
`cache_filename`
`cache_fn_to_dig`
`CALLOC`
`ce_free`
`ce_lookup`
`cif_write_entry`
`cinfo_find`
`condvar_init`
`condvar_notify`
`condvar_terminate`
`condvar_wait`
`crit_enter`
`crit_exit`
`crit_init`

crit_terminate

D

daemon_atrestart

dns_set_hostent

F

fc_close

fc_open

filebuf_buf2sd

filebuf_close

filebuf_getc

filebuf_open

filebuf_open_nostat

filter_create

filter_find

filter_insert

filter_layer

filter_name

filter_remove

flush

FREE

fs_blk_size

fs_blks_avail

func_exec

func_find

func_insert

I

`insert`

L

`log_error`

M

`magnus_atrestart`

`MALLOC`

N

`net_flush`

`net_ip2host`

`net_read`

`net_sendfile`

`net_write`

`netbuf_buf2sd`

`netbuf_close`

`netbuf_getc`

`netbuf_grab`

`netbuf_open`

`nsapi_module_init`

`NSAPI_RUNTIME_VERSION`

`NSAPI_VERSION`

P

`param_create`

`param_free`

`pblock_copy`

`pblock_create`

`pblock_dup`

`pblock_find`

`pblock_findlong`

`pblock_findval`

`pblock_free`

`pblock_ninsert`

`pblock_nninsert`

`pblock_nvinsert`

`pblock_pb2env`

`pblock_pblock2str`

`pblock_pinsert`

`pblock_remove`

`pblock_replace_name`

`pblock_str2pblock`

`PERM_CALLOC`

`PERM_FREE`

`PERM_MALLOC`

`PERM_REALLOC`

`PERM_STRDUP`

`prepare_nsapi_thread`

`protocol_dump822`

`protocol_finish_request`

protocol_handle_session
protocol_parse_request
protocol_scan_headers
protocol_set_finfo
protocol_start_response
protocol_status
protocol_uri2url
protocol_uri2url_dynamic

R

read
REALLOC
remove
request_create
request_free
request_header

S

sem_grab
sem_init
sem_release
sem_terminate
sem_tgrab
sendfile
session_create
session_dns
session_free

session_maxdns
shexp_casecmp
shexp_cmp
shexp_match
shexp_valid
shmem_alloc
shmem_free
STRDUP
system_errmsg
system_fclose
system_flock
system_fopenRO
system_fopenRW
system_fopenWA
system_fread
system_fwrite
system_fwrite_atomic
system_gmtime
system_localtime
system_lseek
system_rename
system_unlock
system_unix2local
systhread_attach
systhread_current
systhread_getdata
systhread_init

systhread_newkey
systhread_setdata
systhread_sleep
systhread_start
systhread_terminate
systhread_timerset

U

USE_NSAPI_VERSION
util_can_exec
util_chdir2path
util_cookie_find
util_does_process_exist
util_env_create
util_env_find
util_env_free
util_env_replace
util_env_str
util_get_current_gmt
util_get_int_from_aux_file
util_get_int_from_file
util_get_long_from_aux_file
util_get_long_from_file
util_get_string_from_aux_file
util_get_string_from_file
util_getline
util_hostname

util_is_mozilla
util_is_url
util_itoa
util_later_than
util_make_filename
util_make_gmt
util_make_local
util_move_dir
util_move_file
util_parse_http_time
util_put_int_to_file
util_put_long_to_file
util_put_string_to_aux_file
util_put_string_to_file
util_sect_id
util_sh_escape
util_snprintf
util_sprintf
util_strcasecmp
util_strftime
util_strncasecmp
util_uri_check
util_uri_escape
util_uri_is_evil
util_uri_parse
util_uri_unescape
util_url_cmp

`util_url_fix_host name`

`util_url_has_FQDN`

`util_vsnprintf`

`util_vsprintf`

W

`write`

`writev`

Index

A

API functions

- cache_digest 14
- cache_filename 14
- cache_fn_to_dig 15
- CALLOC 15
- ce_free 16
- ce_lookup 16
- cif_write_entry 17
- cinfo_find 18
- condvar_init 18
- condvar_notify 19
- condvar_terminate 19
- condvar_wait 20
- crit_enter 20
- crit_exit 21
- crit_init 21
- crit_terminate 22
- daemon_atrestart 22
- fc_close 24
- filebuf_buf2sd 24, 25
- filebuf_close 26
- filebuf_getc 26
- filebuf_open 27
- filebuf_open_nostat 28
- filter_create 28
- filter_find 30
- filter_insert 30
- filter_layer 31
- filter_name 31
- filter_remove 32
- flush 32
- FREE 33
- fs_blk_size 34
- fs_blks_available 34
- func_exec 35
- func_find 35
- func_insert 36
- insert 37
- log_error 38
- magnus_atrestart 39
- MALLOC 39
- net_ip2host 41
- net_read 41
- net_write 44
- netbuf_buf2sd 44
- netbuf_close 45
- netbuf_getc 45
- netbuf_grab 46
- netbuf_open 46
- param_create 49
- param_free 49
- pblock_copy 50
- pblock_create 50
- pblock_dup 51
- pblock_find 51
- pblock_findlong 52
- pblock_findval 52
- pblock_free 53
- pblock_ninsert 54
- pblock_nninsert 54
- pblock_nvinsert 55
- pblock_pb2env 55
- pblock_pblock2str 56
- pblock_pinsert 56
- pblock_remove 57
- pblock_replace_name 57

pblock_str2pblock 58
 PERM_FREE 59
 PERM_MALLOC 59, 60, 61
 PERM_STRDUP 61
 prepare_nsapi_thread 62
 protocol_dump822 63
 protocol_set_finfo 65
 protocol_start_response 66
 protocol_status 67
 protocol_uri2url 68
 read 69
 REALLOC 70
 remove 71
 request_create 71
 request_free 72
 request_header 72
 sem_grab 73
 sem_init 74
 sem_release 74
 sem_terminate 75
 sem_tgrab 75
 sendfile 76
 session_create 76
 session_dns 77
 session_free 78
 session_maxdns 78
 shem_alloc 81
 shexp_casecmp 79
 shexp_cmp 79
 shexp_match 80
 shexp_valid 81
 shmем_free 82
 STRDUP 83
 system_errmsg 83
 system_fclose 84
 system_flock 85
 system_fopenRO 85
 system_fopenRW 86
 system_fopenWA 86
 system_fread 87
 system_fwrite 88
 system_fwrite_atomic 88
 system_gmtime 89
 system_localtime 90
 system_lseek 90
 system_rename 91
 system_ulock 90, 91
 system_unix2local 92
 systhread_attach 92
 systhread_current 93
 systhread_getdata 93
 systhread_newkey 83, 94
 systhread_setdata 95
 systhread_sleep 95
 systhread_start 96
 systhread_terminate 96
 systhread_timerset 83, 97
 util_can_exec 99
 util_chdir2path 100
 util_cookie_find 100
 util_env_create 101
 util_env_find 102
 util_env_free 102
 util_env_replace 103
 util_env_str 103
 util_get_current_gmt 104
 util_get_int_from_file 107
 util_get_long_from_file 106
 util_get_string_from_file 107
 util_getline 108
 util_hostname 108
 util_is_mozilla 109
 util_is_url 109
 util_itoa 110
 util_later_than 110
 util_make_filename 111
 util_make_gmt 111
 util_make_local 112
 util_move_dir 112
 util_move_file 113
 util_parse_http_time 113
 util_put_int_to_file 114
 util_put_long_to_file 114
 util_put_string_to_file 115
 util_sect_id 116
 util_sh_escape 116
 util_snprintf 117
 util_strcasecmp 118
 util_strftime 119
 util_strncasecmp 120
 util_uri_escape 121
 util_uri_is_evil 121
 util_uri_parse 122
 util_uri_unescape 122

- util_url_fix_hostthame 123, 124
- util_vsnprintf 124
- util_vsprintf 125
- util-cookie_find 100
- util-does_process_exist 101
- util-sprintf 118
- write 126
- writew 127

C

- cache_digest
 - API function 14
- cache_filename
 - API function 14
- cache_fn_to_dig
 - API function 15
- CALLOC API function 15
- ce 16
- ce_free
 - API function 16
- ce_lookup
 - API function 16
- cif_write_entry
 - API function 17
- cinfo NSAPI data structure 134
- cinfo_find API function 18
- client
 - getting DNS name for 132
 - getting IP address for 132
 - sessions and 131
- compatibility issues 130
- condvar_init API function 18
- condvar_notify API function 19
- condvar_terminate API function 19
- condvar_wait API function 20
- creating
 - custom NSAPI plugins 11
- crit_enter API function 20
- crit_exit API function 21
- crit_init API function 21
- crit_terminate API function 22

- custom
 - NSAPI plugins 11

D

- daemon_atrestart API function 22
- data structures
 - cinfo 134
 - compatibility issues 130
 - Filter 135
 - FilterContext 136
 - FilterLayer 136
 - FilterMethods 136
 - nsapi.h header file 129
 - nsapi_pvt.h 130
 - pb_entry 132
 - pb_param 132
 - pblock 131
 - privatization of 130
 - removed from nsapi.h 130
 - request 133
 - sendfiledata 135
 - session 131
 - Session->client 132
 - shmem_s 134
 - stat 133
- day of month 141
- DNS names
 - getting clients 132

E

- errors
 - finding most recent system error 83

F

- fc_close API function 24
- file descriptor
 - closing 84

- locking [85](#)
- opening read-only [85](#)
- opening read-write [86](#)
- opening write-append [86](#)
- reading into a buffer [87](#)
- unlocking [90, 91](#)
- writing from a buffer [88](#)
- writing without interruption [88](#)
- filebuf_buf2sd API function [24, 25](#)
- filebuf_close API function [26](#)
- filebuf_getc API function [26](#)
- filebuf_open API function [27](#)
- filebuf_open_nostat API function [28](#)
- Filter NSAPI data structure [135](#)
- filter_create API function [28](#)
- filter_find API function [30](#)
- filter_insert API function [30](#)
- filter_layer API function [31](#)
- filter_name API function [31](#)
- filter_remove API function [32](#)
- FilterContext NSAPI data structure [136](#)
- FilterLayer NSAPI data structure [136](#)
- FilterMethods NSAPI data structure [136](#)
- flush API function [32](#)
- FREE API function [33](#)
- fs_blk_size
 - API function [34](#)
- fs_blks_available
 - API function [34](#)
- func_exec API function [35](#)
- func_find API function [35](#)
- func_insert API function [36](#)

G

- GMT time
 - getting thread-safe value [89](#)

I

- insert API function [37](#)
- IP address
 - getting client's [132](#)

L

- localtime
 - getting thread-safe value [90](#)
- log_error API function [38](#)

M

- magnus_atrestart
 - API function [39](#)
- MALLOC API function [39](#)
- month name [141](#)

N

- net_ip2host API function [41](#)
- net_read API function [41](#)
- net_write API function [44](#)
- netbuf_buf2sd API function [44](#)
- netbuf_close API function [45](#)
- netbuf_getc API function [45](#)
- netbuf_grab API function [46](#)
- netbuf_open API function [46](#)
- NSAPI plugins, custom [11](#)
- nsapi.h [129](#)
- nsapi_pvt.h [130](#)

P

- param_create API function [49](#)

- param_free API function 49
- path name
 - converting UNIX-style to local 92
- pb_entry NSAPI data structure 132
- pb_param NSAPI data structure 132
- pblock
 - NSAPI data structure 131
- pblock_copy API function 50
- pblock_create API function 50
- pblock_dup API function 51
- pblock_find API function 51
- pblock_findlong
 - API function 52
- pblock_findval API function 52
- pblock_free API function 53
- pblock_nlinsert
 - API function 54
- pblock_nninsert API function 54
- pblock_nvinsert API function 55
- pblock_pb2env API function 55
- pblock_pblock2str API function 56
- pblock_pinsert API function 56
- pblock_remove API function 57
- pblock_replace_name
 - API function 57
- pblock_str2pblock API function 58
- PERM_FREE API function 59
- PERM_MALLOC API function 59, 60, 61
- PERM_STRDUP API function 61
- plugins
 - compatibility issues 130
 - private data structures 130
- prepare_nsapi_thread API function 62
- private data structures 130
- protocol_dump822 API function 63
- protocol_set_finfo API function 65
- protocol_start_response API function 66
- protocol_status API function 67
- protocol_uri2url API function 68

R

- read API function 69
- REALLOC API function 70
- remove API function 71
- request
 - NSAPI data structure 133
- request_create
 - API function 71
- request_free
 - API function 72
- request_header API function 72

S

- sem_grab
 - API function 73
- sem_init
 - API function 74
- sem_release
 - API function 74
- sem_terminate
 - API function 75
- sem_tgrab
 - API function 75
- semaphore
 - creating 74
 - deallocating 75
 - gaining exclusive access 73
 - releasing 74
 - testing for exclusive access 75
- sendfile API function 76
- sendfiledata NSAPI data structure 135
- session
 - defined 131
 - NSAPI data structure 131
 - resolving the IP address of 77, 78
- session structure
 - creating 76
 - freeing 78
- Session->client NSAPI data structure 132
- session_create
 - API function 76

session_dns API function 77
 session_free
 API function 78
 session_maxdns API function 78
 shared memory
 allocating 81
 freeing 82
 shell expression
 comparing (case-blind) to a string 79
 comparing (case-sensitive) to a string 79, 80
 validating 81
 shexp_casecmp API function 79
 shexp_cmp API function 79
 shexp_match API function 80
 shexp_valid API function 81
 shmем_alloc
 API function 81
 shmем_free
 API function 82
 shmем_s NSAPI data structure 134
 socket
 closing 45
 reading from 41
 sending a buffer to 44
 sending file buffer to 25
 writing to 44
 sprintf, see util_sprintf 118
 stat NSAPI data structure 133
 STRDUP API function 83
 string
 creating a copy of 83
 system_errmsg API function 83
 system_fclose API function 84
 system_flock API function 85
 system_fopenRO API function 85
 system_fopenRW API function 86
 system_fopenWA API function 86
 system_fread API function 87
 system_fwrite API function 88
 system_fwrite_atomic API function 88
 system_gmtime API function 89
 system_localtime API function 90
 system_lseek API function 90

system_rename API function 91
 system_unlock API function 90, 91
 system_unix2local API function 92
 systhread_attach API function 92
 systhread_current API function 93
 systhread_getdata API function 93
 systhread_newkey
 API function 83
 systhread_newkey API function 94
 systhread_setdata API function 95
 systhread_sleep API function 95
 systhread_start API function 96
 systhread_terminate
 API function 96
 systhread_timerset
 API function 83
 systhread_timerset API function 97

T

thread
 allocating a key for 83, 94
 creating 96
 getting a pointer to 93
 getting data belonging to 93
 putting to sleep 95
 setting data belonging to 95
 setting interrupt timer 83, 97
 terminating 96

U

unicode 122
 util_can_exec API function 99
 util_chdir2path API function 100
 util_cookie_find API function 100
 util_does_process_exist
 API function 101
 util_env_create
 API function 101

[util_env_find API function](#) 102
[util_env_free API function](#) 102
[util_env_replace API function](#) 103
[util_env_str API function](#) 103
[util_get_current_gmt](#)
 API function 104
[util_get_int_from_file](#)
 API function 107
[util_get_long_from_file](#)
 API function 106
[util_get_string_from_file](#)
 API function 107
[util_getline API function](#) 108
[util_hostname API function](#) 108
[util_is_mozilla API function](#) 109
[util_is_url API function](#) 109
[util_itoa API function](#) 110
[util_later_than API function](#) 110
[util_make_filename](#)
 API function 111
[util_make_gmt](#)
 API function 111
[util_make_local](#)
 API function 112
[util_move_dir](#)
 API function 112
[util_move_file](#)
 API function 113
[util_parse_http_time](#)
 API function 113
[util_put_int_to_file](#)
 API function 114
[util_put_long_to_file](#)
 API function 114
[util_put_string_to_file](#)
 API function 115
[util_sect_id](#)
 API function 116
[util_sh_escape API function](#) 116
[util_snprintf API function](#) 117
[util_sprintf API function](#) 118
[util_strcasecmp API function](#) 118
[util_strftime API function](#) 119, 141

[util_strncasecmp API function](#) 120
[util_uri_escape API function](#) 121
[util_uri_is_evil API function](#) 121
[util_uri_parse API function](#) 122
[util_uri_unescape API function](#) 122
[util_url_fix_hostname](#)
 API function 123, 124
[util_vsnprintf API function](#) 124
[util_vsprintf API function](#) 125

V

[vsnprintf, see util_vsnprintf](#) 124
[vsprintf, see util_vsprintf](#) 125

W

[weekday](#) 141
[write API function](#) 126
[writev API function](#) 127

