

NAME

hash, hash32, hash32_buf, hash32_str, hash32_strn, hash32_stre, hash32_strne, jenkins_hash, jenkins_hash32, murmur3_32_hash, murmur3_32_hash32 - general kernel hashing functions

SYNOPSIS

```
#include <sys/hash.h>
```

uint32_t

```
hash32_buf(const void *buf, size_t len, uint32_t hash);
```

uint32_t

```
hash32_str(const void *buf, uint32_t hash);
```

uint32_t

```
hash32_strn(const void *buf, size_t len, uint32_t hash);
```

uint32_t

```
hash32_stre(const void *buf, int end, const char **ep, uint32_t hash);
```

uint32_t

```
hash32_strne(const void *buf, size_t len, int end, const char **ep, uint32_t hash);
```

uint32_t

```
jenkins_hash(const void *buf, size_t len, uint32_t hash);
```

uint32_t

```
jenkins_hash32(const uint32_t *buf, size_t count, uint32_t hash);
```

uint32_t

```
murmur3_32_hash(const void *buf, size_t len, uint32_t hash);
```

uint32_t

```
murmur3_32_hash32(const uint32_t *buf, size_t count, uint32_t hash);
```

DESCRIPTION

The **hash32()** functions are used to give a consistent and general interface to a decent hashing algorithm within the kernel. These functions can be used to hash ASCII NUL terminated strings, as well as blocks of memory.

A *len* argument is the length of the buffer in bytes. A *count* argument is the length of the buffer in 32-bit

words.

The **hash32_buf()** function is used as a general buffer hashing function. The argument *buf* is used to pass in the location, and *len* is the length of the buffer in bytes. The argument *hash* is used to extend an existing hash, or is passed the initial value HASHINIT to start a new hash.

The **hash32_str()** function is used to hash a NUL terminated string passed in *buf* with initial hash value given in *hash*.

The **hash32_strn()** function is like the **hash32_str()** function, except it also takes a *len* argument, which is the maximal length of the expected string.

The **hash32_stre()** and **hash32_strne()** functions are helper functions used by the kernel to hash pathname components. These functions have the additional termination condition of terminating when they find a character given by *end* in the string to be hashed. If the argument *ep* is not NULL, it is set to the point in the buffer at which the hash function terminated hashing.

The **jenkins_hash()** function has same semantics as the **hash32_buf()**, but provides more advanced hashing algorithm with better distribution.

The **jenkins_hash32()** uses same hashing algorithm as the **jenkins_hash()** function, but works only on *uint32_t* sized arrays, thus is simpler and faster. It accepts an array of *uint32_t* values in its first argument and size of this array in the second argument.

The **murmur3_32_hash()** and **murmur3_32_hash32()** functions are similar to **jenkins_hash()** and **jenkins_hash32()**, but implement the 32-bit version of MurmurHash3.

RETURN VALUES

The **hash32()** functions return a 32 bit hash value of the buffer or string.

EXAMPLES

```
LIST_HEAD(head, cache) *hashtbl = NULL;
u_long mask = 0;

void
sample_init(void)
{
    hashtbl = hashinit(numwanted, type, flags, &mask);
}
```

```
void
sample_use(char *str, int len)
{
    uint32_t hash;

    hash = hash32_str(str, HASHINIT);
    hash = hash32_buf(&len, sizeof(len), hash);
    hashtbl[hash & mask] = len;
}
```

SEE ALSO

free(9), hashinit(9), malloc(9)

LIMITATIONS

The **hash32()** functions are only 32 bit functions. They will prove to give poor 64 bit performance, especially for the top 32 bits. At the current time, this is not seen as a great limitation, as these hash values are usually used to index into an array. Should these hash values be used for other means, this limitation should be revisited.

HISTORY

The **hash** functions first appeared in NetBSD 1.6. The current implementation of **hash32** functions was first committed to OpenBSD 3.2, and later imported to FreeBSD 6.1. The **jenkins_hash** functions were added in FreeBSD 10.0. The **murmur3_32_hash** functions were added in FreeBSD 10.1.

AUTHORS

The **hash32** functions were written by Tobias Weingartner. The **jenkins_hash** functions were written by Bob Jenkins. The **murmur3_32_hash** functions were written by Dag-Erling Smørgrav <des@FreeBSD.org>.