

**NAME**

libarchive — functions for reading and writing streaming archives

**OVERVIEW**

The **libarchive** library provides a flexible interface for reading and writing archives in various formats such as tar and cpio. **libarchive** also supports reading and writing archives compressed using various compression filters such as gzip and bzip2. The library is inherently stream-oriented; readers serially iterate through the archive, writers serially add things to the archive. In particular, note that there is currently no built-in support for random access nor for in-place modification.

When reading an archive, the library automatically detects the format and the compression. The library currently has read support for:

- old-style tar archives,
- most variants of the POSIX “ustar” format,
- the POSIX “pax interchange” format,
- GNU-format tar archives,
- most common cpio archive formats,
- 7-Zip archives,
- ar archives (including GNU/SysV and BSD extensions),
- Microsoft CAB archives,
- ISO9660 CD images (including RockRidge and Joliet extensions),
- LHA archives,
- mtree file tree descriptions,
- RAR and most RAR5 archives,
- WARC archives,
- XAR archives,
- Zip archives.

The library automatically detects archives compressed with *compress*(1), *bzip2*(1), *grzip*(1), *gzip*(1), *lrzip*(1), *lz4*(1), *lzip*(1), *lzop*(1), *xz*(1), or *zstd*(1) and decompresses them transparently. Decompression of some formats requires external decompressor utilities. It can similarly detect and decode archives processed with *uuencode*(1) or which have an *rpm*(1) header.

When writing an archive, you can specify the compression to be used and the format to use. The library can write

- POSIX-standard “ustar” archives,
- POSIX “pax interchange format” archives,
- cpio archives,
- 7-Zip archives,
- ar archives,
- two different variants of shar archives,
- ISO9660 CD images,
- mtree file tree descriptions,
- XAR archives,
- Zip archive.

Pax interchange format is an extension of the tar archive format that eliminates essentially all of the limitations of historic tar formats in a standard fashion that is supported by POSIX-compliant *pax*(1) implementations on many systems as well as several newer implementations of *tar*(1). Note that the default write format will suppress the pax extended attributes for most entries; explicitly requesting pax format will enable those attributes for all entries.

The read and write APIs are accessed through the **archive\_read\_XXX()** functions and the **archive\_write\_XXX()** functions, respectively, and either can be used independently of the other.

The rest of this manual page provides an overview of the library operation. More detailed information can be found in the individual manual pages for each API or utility function.

## READING AN ARCHIVE

See *archive\_read(3)*.

## WRITING AN ARCHIVE

See *archive\_write(3)*.

## WRITING ENTRIES TO DISK

The *archive\_write\_disk(3)* API allows you to write *archive\_entry(3)* objects to disk using the same API used by *archive\_write(3)*. The *archive\_write\_disk(3)* API is used internally by **archive\_read\_extract()**; using it directly can provide greater control over how entries get written to disk. This API also makes it possible to share code between archive-to-archive copy and archive-to-disk extraction operations.

## READING ENTRIES FROM DISK

The *archive\_read\_disk(3)* supports for populating *archive\_entry(3)* objects from information in the filesystem. This includes the information accessible from the *stat(2)* system call as well as ACLs, extended attributes, and other metadata. The *archive\_read\_disk(3)* API also supports iterating over directory trees, which allows directories of files to be read using an API compatible with the *archive\_read(3)* API.

## DESCRIPTION

Detailed descriptions of each function are provided by the corresponding manual pages.

All of the functions utilize an opaque struct archive datatype that provides access to the archive contents.

The struct *archive\_entry* structure contains a complete description of a single archive entry. It uses an opaque interface that is fully documented in *archive\_entry(3)*.

Users familiar with historic formats should be aware that the newer variants have eliminated most restrictions on the length of textual fields. Clients should not assume that filenames, link names, user names, or group names are limited in length. In particular, pax interchange format can easily accommodate pathnames in arbitrary character sets that exceed *PATH\_MAX*.

## RETURN VALUES

Most functions return *ARCHIVE\_OK* (zero) on success, non-zero on error. The return value indicates the general severity of the error, ranging from *ARCHIVE\_WARN*, which indicates a minor problem that should probably be reported to the user, to *ARCHIVE\_FATAL*, which indicates a serious problem that will prevent any further operations on this archive. On error, the **archive\_errno()** function can be used to retrieve a numeric error code (see *errno(2)*). The **archive\_error\_string()** returns a textual error message suitable for display.

**archive\_read\_new()** and **archive\_write\_new()** return pointers to an allocated and initialized struct archive object.

**archive\_read\_data()** and **archive\_write\_data()** return a count of the number of bytes actually read or written. A value of zero indicates the end of the data for this entry. A negative value indicates an error, in which case the **archive\_errno()** and **archive\_error\_string()** functions can be used to obtain more information.

## ENVIRONMENT

There are character set conversions within the *archive\_entry(3)* functions that are impacted by the currently-selected locale.

## SEE ALSO

*tar(1)*, *archive\_entry(3)*, *archive\_read(3)*, *archive\_util(3)*, *archive\_write(3)*, *tar(5)*

## HISTORY

The **libarchive** library first appeared in FreeBSD 5.3.

## AUTHORS

The **libarchive** library was originally written by Tim Kientzle <kientzle@acm.org>.

**BUGS**

Some archive formats support information that is not supported by struct archive\_entry. Such information cannot be fully archived or restored using this library. This includes, for example, comments, character sets, or the arbitrary key/value pairs that can appear in pax interchange format archives.

Conversely, of course, not all of the information that can be stored in an struct archive\_entry is supported by all formats. For example, cpio formats do not support nanosecond timestamps; old tar formats do not support large device numbers.

The ISO9660 reader cannot yet read all ISO9660 images; it should learn how to seek.

The AR writer requires the client program to use two passes, unlike all other libarchive writers.