

NAME

recover - Recover files from Incremental Dump

USAGE

recover <filename> [<filename>...<filename>]

DESCRIPTION

recover allows the retrieval of lost/scrambled files from the Incremental archives (dumps) created daily @ 1:00am (the dump begins at 1:00; it may take several hours to complete, thus a file may not actually be saved until 3:00 or 4:00am). Incremental archives contain all files which were modified/created since the EPOCH Dump. They MIGHT NOT contain files which were removed (see below).

Each machine, DEV1 and DEV2 (scd1 & scd2), has 2 dedicated disk areas (numbered 1 & 2) utilized for Incremental Dumps. Each area is used on alternate days (the EPOCH Dump day is typically skipped) and thus provide a 2 day history of file system changes (ie. every other day, a given area is overwritten with new data). Their contents are referred to as the DAY1 and DAY2 Incrementals. See below for details of the files contained on an Incremental Dump.

The <filename> argument to recover should be the entire relative path name of the file you wish to recover. That is, the leading (root) '/' must be omitted (files in the incremental archive are referenced exclusively by their path name relative to the root dir '/'). If you are not sure where a file was, enter as much of the name as possible (since recover 'greps' a listing of Incremental contents for <filename>, this will limit the number of extraneous matches you must weed through).

recover will then:

Output a 'ls -l' of the latest 2 Incremental Contents Listings (list1 & list2 - corresponding to DAY1 & DAY2). Each listing is created at the time of the Incremental. Its modification date is the completion date/time for the last successful Dump for that day.

'Grep' the listings for the string (<filename>) provided. A match may be found in either, both or neither of the listings, for the reasons outlined below. If found, the mode/uid/gid/size of the file in the archive is output. This info and/or the date of the listing may be used to determine from which 'DAY' the file should be extracted.

Prompt the user for the necessary retrieval information:

Listing number where desired version found (1 or 2 - defaults to 1)

File name on Incremental (relative path name - defaults to
<filename> arg)

New name of retrieved file (defaults to last component of
Incremental file name, above)

Path of target directory for retrieved file (defaults to .)

Background? (defaults to no - you must wait for file to be
retrieved)

The prompts are defaulted by entering a <Carriage Return> alone
as a response.

recover will then retrieve the file from the Incremental archive
corresponding to the list number entered (/etc/updfs is used).

[Directories may not be retrieved with recover. Consult a gnome
when a large number of files must be recovered.]

NOTE:

Incremental dumps are only a 2 day history of file modifica-
tions. A given change made on a Thursday may be recovered before
1:00am Sunday AND, if no more mods are made, will be available
through the following Friday. If further modifications are made,
however, (eg. on Monday), the previous change will be lost after
1:00am of the 2nd day following the new change (ie. Wednesday).
Confusing? - you bet.

If a file has not been modified/created since the EPOCH, it will
not appear on the Incremental. Thus, files which have 'vanished
mysteriously' cannot be recovered with recover unless they
were changed/created since the EPOCH and had been removed (or
vanished mysteriously) WITHIN the LAST 2 DAYS. Two days after a
changed/created file is deleted, the Incremental image of the
file systems will no longer contain the file.

The EPOCH packs must be consulted for retrieval of unmodified
files and those modified files removed more than 2 days prior.

FILES


/incrdir/list1 /incrdir/list2 /dev/rDAY1 /dev/rDAY2

DIAGNOSTICS

'Cannot mount /dev/incrdir'. This is where the list? files re-
side. There are probably too many file systems mounted already.

BUGS

Note that Incrementals are not guaranteed to be successful every
day. Perturbances in the file systems can cause the incremental
areas of disk (56k blocks each - 2 per/machine) to be filled to
capacity and not contain all the changes since the EPOCH. At-



tempts are made to ensure that the user directories are dumped first, but even this isn't guaranteed (perturbances include things like 'touch'ing every file in your directories - eg. by changing their modes/groups, etc. -or- copying an entire PR directory into your own).

Since the incremental directory, /incrdir, is a mounted file system, a failure to mount /dev/incrdir will result in an aborted Incremental Dump.

