

**NAME**

**mminfo** – NetWorker media database reporting command

**SYNOPSIS**

**mminfo** [ **-avV** ] [ **-o order** ] [ **-s server** ] [ **report** ] [ **query** ] [ *volname...* ]

< **report** >: [ **-m** | **-p** | **-B** | **-S** | **-X** | **-r reportspec** ]

< **query** >: [ **-c client** ] [ **-N name** ] [ **-t time** ] [ **-q queriespec** ]

**DESCRIPTION**

The **mminfo** command reports information about NetWorker's media and save sets. The **mminfo** command can produce several different reports depending on the flags specified. Several built-in reports can be specified using short-hand flags. Custom reports can also be specified. The default report, along with the built-in reports printed by the use of the **-v**, **-V**, **-m**, **-p**, **-S**, **-B**, and **-X** flags, are described first below. The custom query and report generators, using the **-q queriespec** and **-r reportspec** options, are described in the **CUSTOM QUERIES AND REPORTS** section. Other options are described in the **OPTIONS** section.

Without any options, **mminfo** displays information about the save sets that completed properly during the last twenty four hours, and are still contained in an on-line file index (browsable save sets). The following information is printed for each save set: the containing volume name, the client's name, the creation date, the size saved on that volume, the save set level, and the save set name. The size field is displayed in Bytes (B), KiloBytes (KB), MegaBytes (MB), GigaBytes (GB), TeraBytes (TB), PetaBytes (PB), or ExaBytes (EB). The save set level will display 'full', 'incr', 'migration' or 1 through 9, for full, incremental, migration save sets, level 1 through 9, respectively. The level is only kept for scheduled saves and file migration; save sets generated by explicitly running the **save(8)** command (called *ad hoc* saves) do not have an associated level.

Specifying the **-v** flag causes three additional fields to be displayed: the creation time, the internal save set identifier (ssid), and two flags. One character is used per flag.

The first flag indicates which part of the save set is on the volume. When the save is completely contained on the volume, a **c** is displayed. An **h** is displayed when the save set spans volumes and the head is contained on this volume. The remaining sections will be on other volumes. An **m** is displayed when the save set spans volumes and a middle section is contained on this volume. The head and tail sections will be on different volumes. There may be more than one middle section. A **t** is displayed when the tail section of a spanning save set is contained on this volume. Again, the other sections will be on other volumes.

The second flag indicates the status of the save set. A **b** indicates that the save set is in the on-line index and is browsable via the **recover(8)** command. An **r** indicates that the save set is not in the on-line index and is recoverable via the **scanner(8)** command. An **E** indicates that the save set has been marked eligible for recycling and may be over-written at any time. An **S** denotes that the save set was scanned in (or rolled in). Rolled in save sets are not subject to the standard index management procedures and will remain in the file index until the user manually purges the save set. An **a** indicates that the save was aborted before completion. Aborted save sets are removed from the on-line file index by **nsrck(8)**. An **i** indicates that the save is still in progress. The **-v** flag prints aborted, purged, and incomplete save sets in addition to the complete, browsable save sets printed by default.

The **-V** flag displays even more detail than the **-v** flag, and is generally used for debugging. This format also displays information (i.e. media file number and record number) that can be used to speed the operation of the **scanner(8)** command. Rather than displaying one line per save set per volume, three lines are displayed each time a section of a save set occurs within a file on a volume. A single save set will have multiple index entries if it starts in one file on a volume and ends in another. This report contains all of the information reported via the **-v** flag, but, because of the additional detail, some of this information is reordered. The first line will contain the volume name, the client's name, the size saved in that section, the save set level, and the save set name. The size field lists the number of bytes that are contained in the section, rather than the total amount of the save set on this volume. The second line contains the following fields: the internal save set identifier (ssid), the save time in seconds since 00:00:00 GMT, Jan 1, 1970, the

creation data and time of day, the internal save set identifier (ssid), the browse time, and the retention time. The third line contains: the offset of the first and last bytes of the save set contained within section, the media file number, the first record within the media file containing data for this save set, the internal volume identifier (valid), the total size of the save set, and the flags, described in the `-v` paragraph above, indicating which part of the save set is contained in this media file (**c**, **h**, **m**, or **t**) and the save set's status (**b**, **r**, **a**, or **i**).

The `-p` flag causes **mminfo** to display a report on the browse and retention times for save sets. Each line of the report displays the save set creation date, and the stored browse and retention dates ('undef' is displayed when connecting to a downrev server), the save set identifier, the client's name, and the save set's name. The `-v` and `-V` options have no effect on the columns included in this report.

The `-m` flag causes **mminfo** to display the name of each volume in the media database, the number of bytes written to it, the percent of space used (or the word 'full' indicating that the volume is filled to capacity), the number of bytes read, the expiration date, the number of times the volume has been mounted, and the volume's capacity. Volumes that are recyclable (see **nsrim(8)**) are flagged by an **E** in the 1st column (meaning **E**ligible for recycling). If a volume has been marked as manually-recyclable, an **M** is displayed instead of the **E**. If a volume is both manually-recyclable and eligible for recycling, an **X** will be displayed. Archive and migration volumes are flagged by an **A**, also in the 1st column. If the volume is not an archive or migration volume, and is not recyclable, no flag appears.

Specifying the `-v` flag with the `-m` flag causes three additional fields to be displayed: the internal volume identifier (valid), the number of the next file to be written, and the type of media.

Using a `-V` flag with the `-m` adds a column of flags to the output. There are currently two possible flags. The **d** flag is set if the volume is currently being written (**d**irty). The **r** flag is set if the volume is marked as **r**ead-only. If neither condition is present, the flags column will be empty.

The `-S` flag displays a long, multi-line save set report, which is used for debugging. The number of lines varies per save set. Due to the length, there are no column headers. Instead, each attribute of the save set is displayed in a 'name=value' manner, except the client and save set name, which are displayed as 'client:name', and the extended attributes, described below. The first line of each multi-line group starts on the left margin and includes the save set identifier (ssid), save time as both a date/time string and seconds since 00:00:00 GMT, Jan 1, 1970, and the client and save set names. Subsequent lines for this save set are indented. If the save set is part of a save set series (a 'continued save set') and is not the first in the series, the save set identifier of the previous save set in the series is shown on the second line by itself. The next line displays the level, the save set flags (in 'ssflags' format, as described in the table in the **CUSTOM QUERIES AND REPORTS** section), the save set size in bytes, the number of files in the save set, and the save set insertion date. The next line displays the save set's create, completion, browse and retention (aka expiration) dates. The string 'undef' for any of the values on these two lines generally means an older server that does not store these values is being queried. If the client identifier is set, it is printed on the next line. If the save set has extended attributes (such as the group to which the save set was a part or the archive annotation), they are printed next, at most one attribute per line. The format of each extended attribute is "name: values;". The clones or instances of the save set are shown last (every save set has at least once instance). The first line of each clone shows the clone identifier, the date and time the instance was created, and the per-clone flags (in 'clflags' format from the **CUSTOM QUERIES AND REPORTS** table). For each instance, each section of that instance is shown as a fragment line. The fragment line shows the offset of that fragment from the beginning of the save set, the volume identifier (valid) containing the fragment, the media file and record numbers of start of the fragment, an absolute positioning identifier (unused by existing servers), and the date of last access of the fragment. The `-v` and `-V` options have no effect on this report.

The `-X` flag prepares a save set summary report instead of one or more lines per save set. Note that the entire media database must be examined to resolve this query, making it very slow and expensive. The summary lists the total number of save sets, and breaks the total down into several overlapping categories summarizing the save set types. The recent save set usage, if appropriate to the query, is also printed. The categories are the number of fulls, the number of incrementals, The number of other non-full, non-incremental saves, the number of ad hoc, archive, migration, empty and purged save sets, the number of index

save sets, and finally, the number of incomplete save sets. For recent usage, the number of save sets per day are shown, up to a week ago, along with a summary of the week's save sets and, if applicable, a summary of the month's save sets. For each line, the number of files (saved in the time interval specified), number of save sets, total size, and average size per save set, and average size per file, are listed. The percentage of the amount saved for incrementals v.s. fulls and the percentage of browsable files are also printed, when appropriate. The **-v** and **-V** options have no effect on the summary report.

The **-B** flag performs a canned query to output, in a convenient format, the list of bootstraps generated in the previous five weeks. In this format, there is one line of output for each matched save set. Each line contains the save date and time, save level, save set identifier (ssid), starting file number, starting record number, and the volume. The equivalent query is described below in the **EXAMPLES** section. The **-v** and **-V** options have no effect on the bootstrap display.

## OPTIONS

- a** Causes queries to apply to all complete, browsable save sets, not just those in the last 24 hours. This option is implied by the **-c**, **-N**, **-q**, **-m**, and **-o** options, described below. When combined with a media-only report (**-m** or a custom report showing only media information), **-a** applies to all volumes, not just those with complete and browsable save sets.
- c client**  
Restricts the reported information to the media and/or save sets pertaining to the specified *client*.
- m** Display a media report instead of the default save set report (in other words, a report about the media containing save sets, not the save sets themselves).
- N name**  
Restricts the reported information to the media and/or save sets pertaining to the specified save set *name*.
- o order**  
Sort the output in the specified order. Before displaying the save sets, they are sorted by various fields. Numeric fields are sorted least to greatest, other fields are sorted alphabetically. *order* may be any combination of the letters **celmontR**, representing **c**lient, **e**xpiration date, **l**ength, **m**edia name, **n**ame of save set, **o**ffset on media (file and record number), **t**ime, and **R**everse, respectively. The default sorting order for save set reports is **mocntl**. The offset fields (file and record) are only considered when the **-V** option has been selected and for custom reports that show save set section (fragment) information. When applied to **-m** media-only reports, the length is the amount used on the volume, the expiration date applies to the volume, the time is the last time the media was accessed, and the other order flags are ignored.
- q queryspec**  
Adds the given query constraint to the list of constraints on the current query. Multiple **-q** options may be given. See the **CUSTOM QUERIES AND REPORTS** section below for the syntax of the *queryspec*.
- r reportspec**  
Appends the given report specification to the list of attributes to be displayed for the current query. Multiple **-r** options may be given. See the **CUSTOM QUERIES AND REPORTS** section below for the syntax of the *reportspec*.
- s server**  
Display volume and save set information from the NetWorker system on *server*. See **nsr(8)** for a description of server selection. The default is the current system.
- t time** Restricts the reported information to the media and/or save sets pertaining to the save sets created on or after *time*. See **nsr\_getdate(3)** for a description of the recognized time formats. The default is 'yesterday'.
- v** Turn on the verbose display reports, described above.
- B** Run the canned query to report bootstraps which have been generated in the past five weeks, as

described above. This option is used by **savegrp**(8) when saving the server's index and bootstrap.

- S Displays a long, multi-line save set report, as described above.
- V Displays additional verbose report output, as described above.
- X Prepare a summary report, as described above.

### CUSTOM QUERIES AND REPORTS

The custom query and report options of **mminfo** allow one to generate media and save set reports matching complex constraints without resorting to pipelines and scripts. This section describes the syntax of custom query and report specifications, and gives some simple examples. Further examples are shown in the **EXAMPLES** section, below.

The custom query option, **-q *queryspec***, is an extension to the short-hand query options, such as **-c *client***, which allow you to make queries based on almost any media or save set attribute in the database, and allow various comparisons in addition to the simple equality comparison provided by the short-hand options. The format of a *queryspec* is

```
[!] name [ comp value ] [ , ... ]
```

where *name* is the name of a database attribute, listed in the table below, *comp* is a valid comparator for the attribute, from the set '>', '>=', '=', '<=', '<', and *value* is the value being compared. Leading and trailing spaces can be used to separate the individual components of the specification. The comparator and value must be specified for all but flag attributes. Generally numeric attributes allow all five comparators, and character string attributes generally only allow equality. When comparing flags, whose values are normally 'true' and 'false', one may alternatively use the '[ ! ] *name*' syntax. The '!*name*' form is equivalent to '*name=false*', and '*name*' by itself is equivalent to '*name=true*'. The comparisons in the specification are separated by commas. If a time or a string contains commas, you must quote the value with single or double quotes. Quotes are escaped within a string by repeating them. The following is a valid string comparison:

```
name="Joe's daily, ""hot"" Save Set"
```

Note that command line shells also interpret quotes, so you will need to enclose the entire query within quotes, and quote the single value inside the query, possibly with a different kind of quote, depending on the shell. Except for multiple character string values, explained below, all of the specified constraints must match a given save set and/or media volume before a line will be printed in the report. Multiple **-q** options may be specified, and may be combined with the short-hand query constraints **-c**, **-N** and **-t**. The order of the above query constraints is unimportant.

Numeric constraints, except for identifiers (volume, save set and clone identifiers) allow ranges to be specified, and all character string constraints allow multiple possible values to be specified. Note that times and levels are considered to be numeric values, not character strings. The upper and lower bounds of a numeric range are specified as two separate constraints. For example,

```
%used>20,%used<80
```

matches volumes that are between 20% and 80% used. Each possible value of a given character string attribute is specified as a separate equality constraint. For example,

```
client=pegasus,client=avalon
```

matches save sets from the client 'pegasus' or the client 'avalon'.

The custom report option, **-r *reportspec***, allows one to specify exactly which media and save set attributes should be shown in the report, the order of the columns, the column widths, and where line breaks should be placed. The format of a *reportspec* is

```
name [ (width) ] [ , name [ (width) ] ... ]
```

where *name* is the name of a database attribute, listed below, and the optional *width*, enclosed in parentheses, specifies how wide the column should be. Leading and trailing spaces are ignored. The default column width depends on the attribute; default widths are also shown in the table, below. Multiple **-r** options

may be specified. The order of the columns in the report will be left to right, and correspond to the order of the attribute names specified. Each line of output will contain all of the data requested (you can cause line breaks within a logical line by using the *newline* attribute name). If a value does not fit in the requested column width, subsequent values in the line will be shifted to the right (values are truncated at 256 characters).

The table below lists all of the recognized attribute names, their valid range of query values (or 'NA' for attributes that are only valid for report specifications), their default column width in characters (or 'NA' for flag attributes that are only valid for query specifications), and a short description.

Numeric attributes (shown as *number* in the valid range column of the table) can be specified using any of the comparators listed above, and can be used in range comparisons.

The *=id* attributes are used for various identifiers (volume identifier, save set identifier, and so on) and only allow equality comparisons. In most cases, if the column is narrow (less than 50 characters), only the short ID shown, which corresponds to the ID used by downrev servers. If the column is wide enough, the full ID is shown. Client identifiers always display as full IDs, and clone identifiers always display as short IDs.

*Flag* attributes have the values 'true' or 'false', only apply as query constraints, and have corresponding flag summary strings for report specifications.

*Time* attributes are specified in **nsr\_getdate(3)** format and are otherwise treated as numeric attributes (note that you will need to quote times that contain commas). The special time 'forever', when used as an expiration date, means a save set or volume will never expire. The special time 'undef' is displayed when the time is undefined. When output, times are displayed as *MM/DD/YY HH:MM:SS* for numeric month, day year (last two digits), hours, minutes, and seconds, respectively. If the column is very narrow (less than 14 characters), only the date is shown, columns less than 17 characters wide drop the seconds, columns 17 or more characters wide print the full date.

*Size* and *kbsize* attributes may have a scale factor appended to them, 'KB' for kilobytes, 'MB' for MegaBytes, 'GB' for GigaBytes, 'TB' for TeraBytes, 'PB' for PetaBytes, or 'EB' for ExaBytes. The default scale (when no scale is explicitly specified) on query constraints for attributes is bytes; the default for *kbsize* attributes is kilobytes. The scale varies in reports, depending on the actual value.

*String* attributes may be any arbitrary character string, enclosed in quotes if necessary, as described above in the query syntax paragraph.

attribute name	value range	width	description
space	NA	1	White space before the next column.
newline	NA	1	Line break(s) within a logical line. <i>Width</i> is actually the number of newlines desired.
volume	string	15	The volume name.
volid	=id	11	The unique volume identifier.
barcode	string	15	The volume barcode, when set.
family	string	4	The media family (eg. tape, disk).
type	string	7	The media type (eg. 8mm, optical).
volflags	NA	5	Volume summary flags, <b>d</b> and <b>r</b> , for dirty (in use) and read-only.
state	NA	3	Volume state summary, <b>E</b> , <b>M</b> , <b>X</b> and <b>A</b> , meaning eligible for recycling, manually-recyclable, both, and archive or migration volumes, respectively.
full	flag	NA	Matches full volumes.
inuse	flag	NA	Matches in-use (dirty) volumes.
volrecycle	flag	NA	Matches recyclable volumes.
readonly	flag	NA	Matches read-only volumes.

manual	flag	NA	Matches manually-recyclable volumes.
pool	string	15	The pool containing the volume.
location	string	15	The volume's location.
capacity	size	8	The volume's estimated capacity.
written	kbsize	7	Kbytes written to volume.
%used	number or 'full'	5	Estimated percentage used, or 'full' for volumes marked as full.
read	kbsize	8	Kbytes read (recovered) from the volume.
next	number	5	Next media file for writing.
nrec	number	5	Next media record for writing.
volaccess	time	9	Last time volume was accessed (written to or labeled). Old servers do not provide this value reliably.
volexp	time	9	Volume expiration date.
volretent	time	9	The time the last save set on this volume will expire.
olabel	time	9	The first time the volume was labeled.
labeled	time	9	The most recent time the media volume was (re)labeled.
mounts	number	6	Number of times the volume was mounted explicitly (reboots do not count).
recycled	number	4	Number of times the volume was relabeled.
avail	NA	3	Summary of volume availability, current valid values, <b>n</b> meaning nearline (ie. in a jukebox), and <b>ov</b> meaning the volume is being managed by SmartMedia.
near	flag	NA	Matches nearline volumes.
smartmedia	flag	NA	Matches volumes managed by SmartMedia.
metric	number	6	Volume speed and desirability metric (unused by existing servers).
savesets	NA	6	Number of save sets on a volume.
volattrs	NA	31	The extended volume attributes.
name	string	31	The save set name.
savetime	time	9	The save time (on the client).
nsavetime	NA	11	The save time, printed as seconds since 00:00:00 GMT, Jan 1, 1970.
sscreate	time	9	The creation time (on the server). If the client and server clocks are out of sync, this time may be different from the savetime.
ssid	=id	11	The unique save set identifier.
level	0..9, full, incr, migration or manual	5	The backup level. Manual backups are printed as blank column values in reports.
client	string	11	The client name for the save set.
attrs	NA	31	The extended save set attributes.
pssid	=id	11	When part of a save set series, the previous save set identifier in the series, zero for the first or only save set in a series.

ssflags	NA	7	The save set flags summary, one or more characters in the set <b>CvrSENRIIF</b> , for continued, valid, purged (recoverable), scanned-in (rolled-in), eligible for recycling, NDMP generated, raw, incomplete, in-progress and finished (ended), respectively.
continued	flag	NA	Matches continued save sets.
recoverable	flag	NA	Matches recoverable (purged) save sets.
ssrecycle	flag	NA	Matches recyclable save sets.
incomplete	flag	NA	Matches incomplete save sets.
rolledin	flag	NA	Matches rolled-in save sets.
ndmp	flag	NA	Matches NDMP save sets.
raw	flag	NA	Matches raw save sets.
valid	flag	NA	Matches valid save sets. All save sets are marked 'valid' by current servers.
sumflags	NA	3	Per-volume save set summary flags, as described for the <b>-v</b> report.
fragflags	NA	3	Per-section save set summary flags, as described for the <b>-V</b> report.
totalsize	number	11	The total save set size.
nfiles	number	5	The number of the client's files in the save set.
ssbrowse	time	9	The save set's browse time. This is the time limit that the save set will remain browsable. 'undef' is displayed when connected to a downrev server.
ssretent	time	9	The save set's retention time (aka expiration time). This is the time limit that the save set will be remain in the media database.
ssinsert	time	9	The save set's insertion time. This is the time the save set was most recently introduced into the database (e.g by a backup or by running <b>scanner(8)</b> ).
sscomp	time	9	The save set's completion time. This is the time the save set backup was completed.
clientid	=id	9	The globally unique client identifier for the host that was backed up in this save set.
copies	number	6	The number of copies (instances or clones) of the save set, all with the same save time and save set identifier.
cloneid	=id	11	The clone identifier of one copy.
clonetime	time	9	The time a copy was made.
clflags	NA	5	The clone flags summary, from the set <b>ais</b> for aborted, incomplete and suspect (read error), respectively.
suspect	flag	NA	Matches suspect save set copies, copies that had errors during file recovery.
annotation	string	31	The (archive) save set's annotation. In a <b>queryspec</b> , the string is a regular expression in the form used by <b>grep(1)</b> .
group	string	12	The group of this save set. This is the

group that backed up this save set.

first	number	11	The offset of the first byte of the save set contained within the section.
last	NA	11	The calculated offset of the last byte of the save set contained within the current section.
fragsize	NA	7	The calculated size of the current section of the save set.
sumsize	NA	7	The calculated total size of all of the sections of the save set on this volume.
mediafile	number	5	The media file number containing the current section of the save set.
mediarec	number	5	The media record number where the first bytes of the save set are found within the current media file.
mediamark	number	5	The absolute positioning data for the current section (not used by existing servers).
ssaccess	time	9	The last time this section of the save set was access (for backup or recover).

#### EXAMPLES

In the following examples, the equivalent short-hand and custom versions of the report are shown, when a short-hand option exists for a given report or query.

Display all bootstraps generated in the previous five weeks, as reported by **savegrp(8)**:

```
mminfo -B
mminfo -N bootstrap -t '5 weeks ago' -ot
-r 'savetime(17),space,level(6),ssid'
-r 'mediafile(6),mediarec(8),space(3),volume'
```

Display information about all of the volumes:

```
mminfo -m
mminfo -a -r 'state,volume,written,%used,read,space,volexp'
-r 'mounts(5),space(2),capacity'
```

Display media information from volumes **mars.001** and **mars.002**:

```
mminfo -m mars.001 mars.002
mminfo -m -q 'volume=mars.001,volume=mars.002'
```

Display all save sets named **/usr**:

```
mminfo -N /usr
mminfo -q name=/usr
```

Display save sets named **/usr**, generated by client **venus**, in the past week:

```
mminfo -N /usr -c venus
mminfo -q 'name=/usr,client=venus'
```

Display save sets named **/usr**, generated by client **venus**, on volume **mars.001**:

```
mminfo -N /usr -c venus mars.001
mminfo -q 'name=/usr,client=venus,volume=mars.001'
```

Display a media report of all volumes written on in the past week:

```
mminfo -m -t 'last week'
mminfo -m -q 'savetime>=last week'
```



Display a media report of all non-full volumes, showing the percent-used, pool and location of each volume:

```
mminfo -a -r 'volume,%used,pool,location' -q '!full'
```

Display a media report similar to the `-m` report but showing the barcode instead of the volume label:

```
mminfo -a -r 'state,barcode,written,%used,read,space,volexp'
-r 'mounts(5),space(2),capacity'
```

Display a verbose list of the instances of all save sets with more than one copy, sorted by save time and client name:

```
mminfo -otc -v -q 'copies>1'
```

Display all archive save sets with an annotation of "project data" for the past four months.

```
mminfo -q'annotation=project data'
-r'volume,client,savetime,sumsize,ssid,name,annotation"
-t'four months ago'
```

## FILES

`/nsr/mm/mmvolume` The save set and media volume databases (actually accessed by `nsrmmdbd(8)`).

## SEE ALSO

`grep(1)`, `nsr_getdate(3)`, `nsr_layout(5)`, `nsradmin(8)`, `nsrmmdbd(8)`, `recover(8)`, `savegrp(8)`, `scanner(8)`.

## DIAGNOSTICS

### no matches found for the query

No save sets or volumes were found in the database that matched all of the constraints of the query.

### invalid volume name '*volname*'

The volume name given is not in a valid format. Note that volume names may not begin with a dash. Queries that match no volumes will return the error 'no matches found for the query'.

### only one of `-m`, `-B`, `-S`, `-X` or `-r` may be specified

Only one report can be generated at a time. Use separate runs of `mminfo` to obtain multiple reports.

### invalid sorting order specifier, choose from '*celmnotR*'

Only letters from *celmnotR* may be used with the `-o` option.

### only one `-o` allowed

Only one sorting order may be specified.

### only one `-s` allowed

Only one server can be queried at one time. Use multiple runs of `mminfo` to obtain reports from multiple servers.

### Out of Memory

The query exhausted available memory. Try issuing it again, using the sorting order `-om`, or make the query more restrictive (eg. list specific volumes, clients, and/or save set names).

### invalid value specified for '*attribute*'

The value specified is either out of range (eg. a negative number for a value that can only take positive numbers), the wrong type (an alphabetic string value specified for a numeric attribute), or just poorly formatted (eg. non-blank characters between a close quote and the next comma or a missing close quote).

### value of '*attribute*' is too long

The value specified for *attribute* is longer than the maximum accepted value. Query attributes must have values less than 65 characters long.

### non-overlapping range specified for '*attribute*'

The range specified for *attribute* is a non-overlapping numeric range, and cannot possibly match

any save set or volume in the database.

**unknown query constraint:** *attribute*

The given query *attribute* is not valid. See the **CUSTOM QUERIES AND REPORTS** table for a list of all valid attribute names.

**need a value for query constraint** '*attribute*'

The *attribute* is not a flag, and must be specified in the '*name comparator value*' format.

**constraint** '*attribute*' **is only valid for reports**

The *attribute* specified for a query may only be used in report (**-r**) specifications. Calculated values, flag summaries, save set extended attributes, and formatting tools (**space** and **newline**) may not be used in queries.

**invalid comparator for query constraint** '*attribute*'

The comparator used is not valid for the given *attribute*. See the **CUSTOM QUERIES AND REPORTS** section for a list of the valid comparators for *attribute*.

**query constraint** '*attribute*' **specified more than once**

The given attribute was specified more than once with the same comparator, and is not a string attribute (string attributes can match one of several specific values).

**unknown report constraint:** *attribute*

The given report *attribute* is not valid, see the **CUSTOM QUERIES AND REPORTS** table for a list of all valid attribute names.

**constraint** '*attribute*' **is only valid for queries**

The *attribute* specified for a report is a flag matching attribute and may only be used in query (**-q**) specifications. See the **CUSTOM QUERIES AND REPORTS** table for the appropriate flag summary attribute that one may use in reports of a given flag.

**column width of** '*attribute*' **is invalid**

The width specified for *attribute* is out of range. Column widths must be positive numbers less than 256.

**missing close parenthesis after report constraint** '*attribute*'

The width of *attribute* is missing a close parenthesis.

**missing comma after report constraint** '*attribute*'

There are non-blank characters after the width specification for *attribute* without any comma preceding them.

**No data requested, no report generated**

The given report specification contains only formatting, no data attribute names.

## LIMITATIONS

You cannot specify save set extended attributes as query constraints.

You cannot list several possible equality matches for numbers, only for strings.

Some queries, namely those that are not highly selective (few query constraints) and use a sorting order where the volume name is not the primary sort key, still require **mminfo** to retrieve the entire database before printing any of it. Such queries use up large amounts of memory in **mminfo**, but not, as was the case with older versions, in **nsrmmdbd**.

You cannot make a report that shows save set or media instances and a summary without running **mminfo** at least twice.

You cannot specify query constraints that compare database attributes with each other.

You cannot make a report that uses **-B** flag with **-c** flag.