Full credit is given to the above companies including the OS that this PDF file was generated!

Red Hat Enterprise Linux Release 9.2 Manual Pages on 'chronyc.1' command

$ man chronyc.1

CHRONYC(1) User manual CHRONYC(1)

NAME

chronyc - command-line interface for chrony daemon

SYNOPSIS

chronyc [OPTION]... [COMMAND]...

DESCRIPTION

chronyc is a command-line interface program which can be used to monitor chronyd's performance and to change various operating parameters whilst it is running.

If no commands are specified on the command line, chronyc will expect input from the user. The prompt chronyc> will be displayed when it is being run from a terminal. If chronyc's input or output are redirected from or to a file, the prompt will not be shown.

There are two ways chronyc can access chronyd. One is the Internet Protocol (IPv4 or IPv6) and the other is a Unix domain socket, which is accessible locally by the root or chrony user. By default, chronyc first tries to connect to the Unix domain socket. The compiled-in default path is /run/chrony/chronyd.sock. If that fails (e.g. because chronyc is running under a non-root user), it will try to connect to 127.0.0.1 and then ::1.

Only the following monitoring commands, which do not affect the behaviour of chronyd, are allowed from the network: activity, manual list, rtcdata, smoothing, sourcename, sources, sourcestats, tracking,

waitsync. The set of hosts from which chronyd will accept these commands can be configured with the cmdallow directive in the chronyd's configuration file or the cmdallow command in chronyc. By default, the commands are accepted only from localhost (127.0.0.1 or ::1).

All other commands are allowed only through the Unix domain socket. When sent over the network, chronyd will respond with a ?Not authorised? error, even if it is from localhost.

Having full access to chronyd via chronyc is more or less equivalent to being able to modify the chronyd's configuration file and restart it.

OPTIONS

-4

With this option hostnames will be resolved only to IPv4 addresses.

-6

With this option hostnames will be resolved only to IPv6 addresses.

-n

This option disables resolving of IP addresses to hostnames, e.g. to avoid slow DNS lookups. Long addresses will not be truncated to fit into the column.

-N

This option enables printing of original hostnames or IP addresses of NTP sources that were specified in the configuration file, or chronyc commands. Without the -n and -N option, the printed hostnames are obtained from reverse DNS lookups and can be different from the specified hostnames.

-c

This option enables printing of reports in a comma-separated values (CSV) format. Reverse DNS lookups will be disabled, time will be printed as number of seconds since the epoch, and values in seconds will not be converted to other units.

-d

This option enables printing of debugging messages if chronyc was compiled with debugging support.

-m

Normally, all arguments on the command line are interpreted as one command. With this option multiple commands can be specified. Each argument will be interpreted as a whole command.

-h host

This option specifies the host to be contacted by chronyc. It can be specified with a hostname, IP address, or path to the local Unix domain socket. Multiple values can be specified as a comma-separated list to provide a fallback.

The default value is /run/chrony/chronyd.sock,127.0.0.1,::1, i.e. the host where chronyc is being run. First, it tries to connect to the Unix domain socket and if that fails (e.g. due to running under a non-root user), it will try to connect to 127.0.0.1 and then ::1.

-p port

This option allows the user to specify the UDP port number which the target chronyd is using for its monitoring connections. This defaults to 323; there would rarely be a need to change this.

-f file

This option is ignored and is provided only for compatibility.

-a

This option is ignored and is provided only for compatibility.

-v, --version

With this option chronyc displays its version number on the terminal and exits.

--help

With this option chronyc displays a help message on the terminal and exits.

COMMANDS

This section describes each of the commands available within the chronyc program.

System clock

tracking

The tracking command displays parameters about the system?s clock performance. An example of the output is shown below.

Reference ID : CB00710F (foo.example.net)

Stratum : 3

Ref time (UTC) : Fri Jan 27 09:49:17 2017

System time : 0.000006523 seconds slow of NTP time

Last offset : -0.000006747 seconds

RMS offset : 0.000035822 seconds

Frequency : 3.225 ppm slow

Residual freq : -0.000 ppm

Skew : 0.129 ppm

Root delay : 0.013639022 seconds

Root dispersion : 0.001100737 seconds

Update interval : 64.2 seconds

Leap status : Normal

The fields are explained as follows:

Reference ID

This is the reference ID and name (or IP address) of the server to which the computer is currently synchronised. For IPv4 addresses, the reference ID is equal to the address and for IPv6 addresses it is the first 32 bits of the MD5 sum of the address.

If the reference ID is 7F7F0101 and there is no name or IP address, it means the computer is not synchronised to any external source and that you have the local mode operating (via the local command in chronyc, or the local directive in the configuration file).

The reference ID is printed as a hexadecimal number. Note that in older versions it used to be printed in quad-dotted notation and could be confused with an IPv4 address.

Stratum

The stratum indicates how many hops away from a computer with an attached reference clock we are. Such a computer is a stratum-1 computer, so the computer in the example is two hops away (i.e. foo.example.net is a stratum-2 and is synchronised

from a stratum-1).

Ref time

This is the time (UTC) at which the last measurement from the reference source was processed.

System time

This is the current offset between the NTP clock and system clock. The NTP clock is a software (virtual) clock maintained by chronyd, which is synchronised to the configured time sources and provides time to NTP clients. The system clock is synchronised to the NTP clock. To avoid steps in the system time, which might have adverse consequences for certain applications, the system clock is normally corrected only by speeding up or slowing down (up to the rate configured by the maxslewrate directive). If the offset is too large, this correction will take a very long time. A step can be forced by the makestep command, or the makestep directive in the configuration file.

Note that all other offsets reported by chronyc and most offsets in the log files are relative to the NTP clock, not the system clock.

Last offset

This is the estimated local offset on the last clock update. A positive value indicates the local time (as previously estimated true time) was ahead of the time sources.

RMS offset

This is a long-term average of the offset value.

Frequency

The ?frequency? is the rate by which the system?s clock would be wrong if chronyd was not correcting it. It is expressed in ppm (parts per million). For example, a value of 1 ppm would mean that when the system?s clock thinks it has advanced 1 second, it has actually advanced by 1.000001 seconds relative to true time.

Residual freq

This shows the ?residual frequency? for the currently selected reference source. This reflects any difference between what the measurements from the reference source indicate the frequency should be and the frequency currently being used.

The reason this is not always zero is that a smoothing procedure is applied to the frequency. Each time a measurement from the reference source is obtained and a new residual frequency computed, the estimated accuracy of this residual is compared with the estimated accuracy (see ?skew? next) of the existing frequency value. A weighted average is computed for the new frequency, with weights depending on these accuracies. If the measurements from the reference source follow a consistent trend, the residual will be driven to zero over time.

Skew

This is the estimated error bound on the frequency.

Root delay

This is the total of the network path delays to the stratum-1 computer from which the computer is ultimately synchronised.

Root dispersion

This is the total dispersion accumulated through all the computers back to the stratum-1 computer from which the computer is ultimately synchronised. Dispersion is due to system clock resolution, statistical measurement variations, etc.

An absolute bound on the computer?s clock accuracy (assuming the stratum-1 computer is correct) is given by:

clock_error <= |system_time_offset| + root_dispersion + (0.5 * root_delay)

Update interval

This is the interval between the last two clock updates.

Leap status

This is the leap status, which can be Normal, Insert second,

Delete second or Not synchronised.

makestep, makestep threshold limit

Normally chronyd will cause the system to gradually correct any time offset, by slowing down or speeding up the clock as required. In certain situations, the system clock might be so far adrift that this slewing process would take a very long time to correct the system clock.

The makestep command can be used in this situation. There are two forms of the command. The first form has no parameters. It tells chronyd to cancel any remaining correction that was being slewed and jump the system clock by the equivalent amount, making it correct immediately.

The second form configures the automatic stepping, similarly to the makestep directive. It has two parameters, stepping threshold (in seconds) and number of future clock updates for which the threshold will be active. This can be used with the burst command to quickly make a new measurement and correct the clock by stepping if needed, without waiting for chronyd to complete the measurement and update the clock.

makestep 0.1 1

burst 1/2

BE WARNED: Certain software will be seriously affected by such jumps in the system time. (That is the reason why chronyd uses slewing normally.)

maxupdateskew skew-in-ppm

This command has the same effect as the maxupdateskew directive in the configuration file.

waitsync [max-tries [max-correction [max-skew [interval]]]]

The waitsync command waits for chronyd to synchronise.

Up to four optional arguments can be specified. The first is the maximum number of tries before giving up and returning a non-zero error code. When 0 is specified, or there are no arguments, the number of tries will not be limited.

The second and third arguments are the maximum allowed remaining correction of the system clock and the maximum allowed skew (in ppm) as reported by the tracking command in the System time and Skew fields. If not specified or zero, the value will not be checked.

The fourth argument is the interval specified in seconds in which the check is repeated. The interval is 10 seconds by default.

An example is:

```
waitsync 60 0.01
```

which will wait up to about 10 minutes (60 times 10 seconds) for chronyd to synchronise to a source and the remaining correction to be less than 10 milliseconds.

Time sources

sources [-a] [-v]

This command displays information about the current time sources that chronyd is accessing.

If the -a option is specified, all sources are displayed, including those that do not have a known address yet. Such sources have an identifier in the format ID#XXXXXXXXXX, which can be used in other commands expecting a source address.

The -v option enables a verbose output. In this case, extra caption lines are shown as a reminder of the meanings of the columns.

```
MS Name/IP address         Stratum Poll Reach LastRx Last sample
===============================================================================
#* GPS0                     0   4   377    11   -479ns[ -621ns] +/-  134ns
^? foo.example.net           2   6   377    23   -923us[ -924us] +/-   43ms
^+ bar.example.net           1   6   377    21  -2629us[-2619us] +/-   86ms
```

The columns are as follows:

M

This indicates the mode of the source. ^ means a server, = means a peer and # indicates a locally connected reference clock.

S

This column indicates the selection state of the source.

? * indicates the best source which is currently selected for synchronisation.

? + indicates other sources selected for synchronisation, which are combined with the best source.

? - indicates a source which is considered to be selectable for synchronisation, but not currently selected.

? x indicates a source which chronyd thinks is a falseticker (i.e. its time is inconsistent with a majority of other sources, or sources specified with the trust option).

? ~ indicates a source whose time appears to have too much variability.

? ? indicates a source which is not considered to be selectable for synchronisation for other reasons (e.g. unreachable, not synchronised, or does not have enough measurements).

The selectdata command can be used to get more details about the selection state.

Name/IP address

This shows the name or the IP address of the source, or reference ID for reference clocks.

Stratum

This shows the stratum of the source, as reported in its most recently received sample. Stratum 1 indicates a computer with a locally attached reference clock. A computer that is synchronised to a stratum 1 computer is at stratum 2. A computer that is synchronised to a stratum 2 computer is at stratum 3, and so on.

Poll

This shows the rate at which the source is being polled, as a base-2 logarithm of the interval in seconds. Thus, a value of 6 would indicate that a measurement is being made every 64 seconds. chronyd automatically varies the polling rate in

response to prevailing conditions.

Reach

This shows the source?s reachability register printed as an octal number. The register has 8 bits and is updated on every received or missed packet from the source. A value of 377 indicates that a valid reply was received for all from the last eight transmissions.

LastRx

This column shows how long ago the last good sample (which is shown in the next column) was received from the source. Measurements that failed some tests are ignored. This is normally in seconds. The letters m, h, d or y indicate minutes, hours, days, or years.

Last sample

This column shows the offset between the local clock and the source at the last measurement. The number in the square brackets shows the actual measured offset. This can be suffixed by ns (indicating nanoseconds), us (indicating microseconds), ms (indicating milliseconds), or s (indicating seconds). The number to the left of the square brackets shows the original measurement, adjusted to allow for any slews applied to the local clock since. The number following the +/- indicator shows the margin of error in the measurement. Positive offsets indicate that the local clock is ahead of the source.

sourcestats [-a] [-v]

The sourcestats command displays information about the drift rate and offset estimation process for each of the sources currently being examined by chronyd.

If the -a option is specified, all sources are displayed, including those that do not have a known address yet. Such sources have an identifier in the format ID#XXXXXXXXXX, which can be used in other commands expecting a source address.

The -v option enables a verbose output. In this case, extra caption

lines are shown as a reminder of the meanings of the columns.

An example report is:

```
   Name/IP Address            NP  NR  Span  Frequency  Freq Skew  Offset  Std Dev
   ==============================================================================
   foo.example.net            11   5   46m     -0.001      0.045      1us    25us
```

The columns are as follows:

Name/IP Address

This is the name or IP address of the NTP server (or peer) or reference ID of the reference clock to which the rest of the line relates.

NP

This is the number of sample points currently being retained for the server. The drift rate and current offset are estimated by performing a linear regression through these points.

NR

This is the number of runs of residuals having the same sign following the last regression. If this number starts to become too small relative to the number of samples, it indicates that a straight line is no longer a good fit to the data. If the number of runs is too low, chronyd discards older samples and re-runs the regression until the number of runs becomes acceptable.

Span

This is the interval between the oldest and newest samples. If no unit is shown the value is in seconds. In the example, the interval is 46 minutes.

Frequency

This is the estimated residual frequency for the server, in parts per million. In this case, the computer?s clock is estimated to be running 1 part in 10^9 slow relative to the server.

Freq Skew

This is the estimated error bounds on Freq (again in parts per

million).

Offset

This is the estimated offset of the source.

Std Dev

This is the estimated sample standard deviation.

selectdata [-a] [-v]

The selectdata command displays information specific to the selection of time sources. If the -a option is specified, all sources are displayed, including those that do not have a known address yet. With the -v option, extra caption lines are shown as a reminder of the meanings of the columns.

An example of the output is shown below.

```
S Name/IP Address        Auth COpts EOpts Last Score     Interval  Leap
=============================================================================
D foo.example.net         Y ----- --TR-    4   1.0   -61ms   +62ms  N
* bar.example.net         N ----- -----    0   1.0 -6846us +7305us  N
+ baz.example.net         N ----- -----   10   1.0 -7381us +7355us  N
```

The columns are as follows:

S

This column indicates the state of the source after the last source selection. It is similar to the state reported by the sources command, but more states are reported.

The following states indicate the source is not considered selectable for synchronisation:

- N - has the noselect option.
- s - is not synchronised.
- M - does not have enough measurements.
- d - has a root distance larger than the maximum distance (configured by the maxdistance directive).
- ~ - has a jitter larger than the maximum jitter (configured by the maxjitter directive).
- w - waits for other sources to get out of the M state.
- S - has older measurements than other sources.

- O - has a stratum equal or larger than the orphan stratum (configured by the local directive).
- T - does not fully agree with sources that have the trust option.
- x - does not agree with other sources (falseticker).

The following states indicate the source is considered selectable, but it is not currently used for synchronisation:

- W - waits for other sources to be selectable (required by the minsources directive, or the require option of another source).
- P - another selectable source is preferred due to the prefer option.
- U - waits for a new measurement (after selecting a different best source).
- D - has, or recently had, a root distance which is too large to be combined with other sources (configured by the combinelimit directive).

The following states indicate the source is used for synchronisation of the local clock:

- + - combined with the best source.
- * - selected as the best source to update the reference data (e.g. root delay, root dispersion).

Name/IP address

This column shows the name or IP address of the source if it is an NTP server, or the reference ID if it is a reference clock.

Auth

This column indicites whether an authentication mechanism is enabled for the source. Y means yes and N means no.

COpts

This column displays the configured selection options of the source.

- N indicates the noselect option.
- P indicates the prefer option.

- ? T indicates the trust option.
- ? R indicates the require option.

EOpts

This column displays the current effective selection options of the source, which can be different from the configured options due to the authentication selection mode (configured by the authselmode directive). The symbols are the same as in the COpts column.

Last

This column displays how long ago was the last measurement of the source made when the selection was performed.

Score

This column displays the current score against the source in the * state. The scoring system avoids frequent reselection when multiple sources have a similar root distance. A value larger than 1 indicates this source was better than the * source in recent selections. If the score reaches 10, the best source will be reselected and the scores will be reset to 1.

Interval

This column displays the lower and upper endpoint of the interval which was expected to contain the true offset of the local clock considering the root distance at the time of the selection.

Leap

This column displays the current leap status of the source.

- ? N indicates the normal status (no leap second).
- ? + indicates that a leap second will be inserted at the end of the month.
- ? - indicates that a leap second will be deleted at the end of the month.
- ? ? indicates the unknown status (i.e. no valid measurement was made).

To avoid excessive switching between sources, chronyd can stay synchronised to a source even when it is not currently the best one among the available sources.

The reselect command can be used to force chronyd to reselect the best synchronisation source.

reselectdist distance

The reselectdist command sets the reselection distance. It is equivalent to the reselectdist directive in the configuration file.

NTP sources

activity

This command reports the number of servers and peers that are online and offline. If the auto_offline option is used in specifying some of the servers or peers, the activity command can be useful for detecting when all of them have entered the offline state after the network link has been disconnected.

The report shows the number of servers and peers in 5 states:

online

the server or peer is currently online (i.e. assumed by chronyd to be reachable)

offline

the server or peer is currently offline (i.e. assumed by chronyd to be unreachable, and no measurements from it will be attempted.)

burst_online

a burst command has been initiated for the server or peer and is being performed; after the burst is complete, the server or peer will be returned to the online state.

burst_offline

a burst command has been initiated for the server or peer and is being performed; after the burst is complete, the server or peer will be returned to the offline state.

unresolved

the name of the server or peer was not resolved to an address

yet; this source is not visible in the sources and sourcestats reports.

authdata [-a]

The authdata command displays information specific to authentication of NTP sources. If the -a option is specified, all sources are displayed, including those that do not have a known address yet. An example of the output is shown below.

```
Name/IP address            Mode KeyID Type KLen Last Atmp  NAK Cook CLen
==============================================================================
foo.example.net            NTS     1   15  256 135m    0    0    8  100
bar.example.net             SK    30   13  128    -    0    0    0    0
baz.example.net              -     0    0    0    -    0    0    0    0
```

The columns are as follows:

Name/IP address

This column shows the name or the IP address of the source.

Mode

This column shows which mechanism authenticates NTP packets received from the source. NTS means Network Time Security, SK means a symmetric key, and - means authentication is disabled.

KeyID

This column shows an identifier of the key used for authentication. With a symmetric key, it is the ID from the key file. With NTS, it is a number starting at zero and incremented by one with each successful key establishment using the NTS-KE protocol, i.e. it shows how many times the key establishment was performed with this source.

Type

This columns shows an identifier of the algorithm used for authentication. With a symmetric key, it is the hash function or cipher specified in the key file. With NTS, it is an authenticated encryption with associated data (AEAD) algorithm, which is negotiated in the NTS-KE protocol. The following values can be reported:

- ? 1: MD5
- ? 2: SHA1
- ? 3: SHA256
- ? 4: SHA384
- ? 5: SHA512
- ? 6: SHA3-224
- ? 7: SHA3-256
- ? 8: SHA3-384
- ? 9: SHA3-512
- ? 10: TIGER
- ? 11: WHIRLPOOL
- ? 13: AES128
- ? 14: AES256
- ? 15: AEAD-AES-SIV-CMAC-256

KLen

This column shows the length of the key in bits.

Last

This column shows how long ago the last successful key establishment was performed. It is in seconds, or letters m, h, d or y indicate minutes, hours, days, or years.

Atmp

This column shows the number of attempts to perform the key establishment since the last successful key establishment. A number larger than 1 indicates a problem with the network or server.

NAK

This column shows whether an NTS NAK was received since the last request. A NAK indicates that authentication failed on the server side due to chronyd using a cookie which is no longer valid and that it needs to perform the key establishment again in order to get new cookies.

Cook

This column shows the number of NTS cookies that chronyd

currently has. If the key establishment was successful, a number smaller than 8 indicates a problem with the network or server.

CLen

This column shows the length in bytes of the NTS cookie which will be used in the next request.

ntpdata [address]

The ntpdata command displays the last valid measurement and other NTP-specific information about the specified NTP source, or all NTP sources (with a known address) if no address was specified. An example of the output is shown below.

```
Remote address  : 203.0.113.15 (CB00710F)
Remote port     : 123
Local address   : 203.0.113.74 (CB00714A)
Leap status     : Normal
Version         : 4
Mode            : Server
Stratum         : 1
Poll interval   : 10 (1024 seconds)
Precision       : -24 (0.000000060 seconds)
Root delay      : 0.000015 seconds
Root dispersion : 0.000015 seconds
Reference ID    : 47505300 (GPS)
Reference time  : Fri Nov 25 15:22:12 2016
Offset          : -0.000060878 seconds
Peer delay      : 0.000175634 seconds
Peer dispersion : 0.000000681 seconds
Response time   : 0.000053050 seconds
Jitter asymmetry: +0.00
NTP tests       : 111 111 1111
Interleaved     : No
Authenticated   : No
TX timestamping : Kernel
```

RX timestamping : Kernel

Total TX : 24

Total RX : 24

Total valid RX : 24

Total good RX : 22

The fields are explained as follows:

Remote address

The IP address of the NTP server or peer, and the corresponding reference ID.

Remote port

The UDP port number to which the request was sent. The standard NTP port is 123.

Local address

The local IP address which received the response, and the corresponding reference ID.

Leap status, Version, Mode, Stratum, Poll interval, Precision, Root delay, Root dispersion, Reference ID, Reference time

The NTP values from the last valid response.

Offset, Peer delay, Peer dispersion

The measured values.

Response time

The time the server or peer spent in processing of the request and waiting before sending the response.

Jitter asymmetry

The estimated asymmetry of network jitter on the path to the source. The asymmetry can be between -0.5 and 0.5. A negative value means the delay of packets sent to the source is more variable than the delay of packets sent from the source back.

NTP tests

Results of RFC 5905 tests 1 through 3, 5 through 7, and tests for maximum delay, delay ratio, delay dev ratio (or delay quantile), and synchronisation loop.

Interleaved

This shows if the response was in the interleaved mode.

Authenticated

This shows if the response was authenticated.

TX timestamping

The source of the local transmit timestamp. Valid values are Daemon, Kernel, and Hardware.

RX timestamping

The source of the local receive timestamp.

Total TX

The number of packets sent to the source.

Total RX

The number of all packets received from the source.

Total valid RX

The number of packets which passed the first two groups of NTP tests.

Total good RX

The number of packets which passed all three groups of NTP tests, i.e. the NTP measurement was accepted.

add peer name [option]...

The add peer command allows a new NTP peer to be added whilst chronyd is running.

Following the words add peer, the syntax of the following parameters and options is identical to that for the peer directive in the configuration file.

An example of using this command is shown below.

```
add peer foo.example.net minpoll 6 maxpoll 10 key 25
```

add pool name [option]...

The add pool command allows a pool of NTP servers to be added whilst chronyd is running.

Following the words add pool, the syntax of the following parameters and options is identical to that for the pool directive in the configuration file.

An example of using this command is shown below:

add pool foo.example.net maxsources 3 iburst

add server name [option]...

The add server command allows a new NTP server to be added whilst chronyd is running.

Following the words add server, the syntax of the following parameters and options is identical to that for the server directive in the configuration file.

An example of using this command is shown below:

add server foo.example.net minpoll 6 maxpoll 10 key 25

delete address

The delete command allows an NTP server or peer to be removed from the current set of sources.

burst good/max [mask/masked-address], burst good/max [masked-address/masked-bits], burst good/max [address]

The burst command tells chronyd to make a set of measurements to each of its NTP sources over a short duration (rather than the usual periodic measurements that it makes). After such a burst, chronyd will revert to the previous state for each source. This might be either online, if the source was being periodically measured in the normal way, or offline, if the source had been indicated as being offline. (A source can be switched between the online and offline states with the online and offline commands.)

The mask and masked-address arguments are optional, in which case chronyd will initiate a burst for all of its currently defined sources.

The arguments have the following meaning and format:

good

This defines the number of good measurements that chronyd will want to obtain from each source. A measurement is good if it passes certain tests, for example, the round trip time to the source must be acceptable. (This allows chronyd to reject measurements that are likely to be bogus.)

max

This defines the maximum number of measurements that chronyd will attempt to make, even if the required number of good measurements has not been obtained.

mask

This is an IP address with which the IP address of each of chronyd's sources is to be masked.

masked-address

This is an IP address. If the masked IP address of a source matches this value then the burst command is applied to that source.

masked-bits

This can be used with masked-address for CIDR notation, which is a shorter alternative to the form with mask.

address

This is an IP address or a hostname. The burst command is applied only to that source.

If no mask or masked-address arguments are provided, every source will be matched.

An example of the two-argument form of the command is:

```
burst 2/10
```

This will cause chronyd to attempt to get two good measurements from each source, stopping after two have been obtained, but in no event will it try more than ten probes to the source.

Examples of the four-argument form of the command are:

```
burst 2/10 255.255.0.0/1.2.0.0
burst 2/10 2001:db8:789a::/48
```

In the first case, the two out of ten sampling will only be applied to sources whose IPv4 addresses are of the form 1.2.x.y, where x and y are arbitrary. In the second case, the sampling will be applied to sources whose IPv6 addresses have first 48 bits equal to 2001:db8:789a.

Example of the three-argument form of the command is:

```
burst 2/10 foo.example.net
```

maxdelay address delay

This allows the maxdelay option for one of the sources to be modified, in the same way as specifying the maxdelay option for the server directive in the configuration file.

maxdelaydevratio address ratio

This allows the maxdelaydevratio option for one of the sources to be modified, in the same way as specifying the maxdelaydevratio option for the server directive in the configuration file.

maxdelayratio address ratio

This allows the maxdelayratio option for one of the sources to be modified, in the same way as specifying the maxdelayratio option for the server directive in the configuration file.

maxpoll address maxpoll

The maxpoll command is used to modify the maximum polling interval for one of the current set of sources. It is equivalent to the maxpoll option in the server directive in the configuration file. Note that the new maximum polling interval only takes effect after the next measurement has been made.

minpoll address minpoll

The minpoll command is used to modify the minimum polling interval for one of the current set of sources. It is equivalent to the minpoll option in the server directive in the configuration file. Note that the new minimum polling interval only takes effect after the next measurement has been made.

minstratum address minstratum

The minstratum command is used to modify the minimum stratum for one of the current set of sources. It is equivalent to the minstratum option in the server directive in the configuration file.

offline [address], offline [masked-address/masked-bits], offline [mask/masked-address]

The offline command is used to warn chronyd that the network connection to a particular host or hosts is about to be lost, e.g.

on computers with intermittent connection to their time sources. Another case where offline could be used is where a computer serves time to a local group of computers, and has a permanent connection to true time servers outside the organisation. However, the external connection is heavily loaded at certain times of the day and the measurements obtained are less reliable at those times. In this case, it is probably most useful to determine the gain or loss rate during the quiet periods and let the whole network coast through the loaded periods. The offline and online commands can be used to achieve this.

There are four forms of the offline command. The first form is a wildcard, meaning all sources (including sources that do not have a known address yet). The second form allows an IP address mask and a masked address to be specified. The third form uses CIDR notation. The fourth form uses an IP address or a hostname. These forms are illustrated below.

```
offline
offline 255.255.255.0/1.2.3.0
offline 2001:db8:789a::/48
offline foo.example.net
```

The second form means that the offline command is to be applied to any source whose IPv4 address is in the 1.2.3 subnet. (The host?s address is logically and-ed with the mask, and if the result matches the masked-address the host is processed.) The third form means that the command is to be applied to all sources whose IPv6 addresses have their first 48 bits equal to 2001:db8:789a. The fourth form means that the command is to be applied only to that one source.

The wildcard form of the address is equivalent to:

```
offline 0.0.0.0/0.0.0.0
offline ::/0
```

online [address], online [masked-address/masked-bits], online [mask/masked-address]

The online command is opposite in function to the offline command.

It is used to advise chronyd that network connectivity to a

particular source or sources has been restored.

The syntax is identical to that of the offline command.

onoffline

The onoffline command tells chronyd to switch all sources that have

a known address to the online or offline status according to the

current network configuration. A source is considered online if it

is possible to send requests to it, i.e. a network route to the

source is present.

polltarget address polltarget

The polltarget command is used to modify the poll target for one of

the current set of sources. It is equivalent to the polltarget

option in the server directive in the configuration file.

refresh

The refresh command can be used to force chronyd to resolve the

names of configured sources to IP addresses again, e.g. after

suspending and resuming the machine in a different network.

Sources that stop responding will be replaced with newly resolved

addresses automatically after 8 polling intervals, but this command

can still be useful to replace them immediately and not wait until

they are marked as unreachable.

reload sources

The reload sources command causes chronyd to re-read all *.sources

files from the directories specified by the sourcedir directive.

sourcename address

The sourcename command prints the original hostname or address that

was specified for an NTP source in the configuration file, or the

add command. This command is an alternative to the -N option, which

can be useful in scripts.

Note that different NTP sources can share the same name, e.g.

servers from a pool.

manual on, manual off, manual delete index, manual list, manual reset

The manual command enables and disables use of the settime command, and is used to modify the behaviour of the manual clock driver.

The on form of the command enables use of the settime command.

The off form of the command disables use of the settime command.

The list form of the command lists all the samples currently stored in chronyd. The output is illustrated below.

```
210 n_samples = 1
#    Date  Time(UTC)    Slewed   Original   Residual
==============================================
 0 27Jan99 22:09:20       0.00       0.97       0.00
```

The columns are as as follows:

1. The sample index (used for the manual delete command).
2. The date and time of the sample.
3. The system clock error when the timestamp was entered, adjusted to allow for changes made to the system clock since.
4. The system clock error when the timestamp was entered, as it originally was (without allowing for changes to the system clock since).
5. The regression residual at this point, in seconds. This allows ?outliers? to be easily spotted, so that they can be deleted using the manual delete command.

The delete form of the command deletes a single sample. The parameter is the index of the sample, as shown in the first column of the output from manual list. Following deletion of the data point, the current error and drift rate are re-estimated from the remaining data points and the system clock trimmed if necessary. This option is intended to allow ?outliers? to be discarded, i.e. samples where the administrator realises they have entered a very poor timestamp.

The reset form of the command deletes all samples at once. The system clock is left running as it was before the command was entered.

settime time

The settime command allows the current time to be entered manually, if this option has been configured into chronyd. (It can be configured either with the manual directive in the configuration file, or with the manual command of chronyc.)

It should be noted that the computer?s sense of time will only be as accurate as the reference you use for providing this input (e.g. your watch), as well as how well you can time the press of the return key.

Providing your computer?s time zone is set up properly, you will be able to enter a local time (rather than UTC).

The response to a successful settime command indicates the amount that the computer?s clock was wrong. It should be apparent from this if you have entered the time wrongly, e.g. with the wrong time zone.

The rate of drift of the system clock is estimated by a regression process using the entered measurement and all previous measurements entered during the present run of chronyd. However, the entered measurement is used for adjusting the current clock offset (rather than the estimated intercept from the regression, which is ignored). Contrast what happens with the manual delete command, where the intercept is used to set the current offset (since there is no measurement that has just been entered in that case).

The time is parsed by the public domain getdate algorithm. Consequently, you can only specify time to the nearest second.

Examples of inputs that are valid are shown below:

settime 16:30

settime 16:30:05

settime Nov 21, 2015 16:30:05

For a full description of getdate, see the getdate documentation (bundled, for example, with the source for GNU tar).

NTP access

accheck address

This command allows you to check whether client NTP access is allowed from a particular host.

Examples of use, showing a named host and a numeric IP address, are as follows:

```
accheck foo.example.net
accheck 1.2.3.4
accheck 2001:db8::1
```

This command can be used to examine the effect of a series of allow, allow all, deny, and deny all commands specified either via chronyc, or in chronyd's configuration file.

clients [-p packets] [-k] [-r]

This command shows a list of clients that have accessed the server, through the NTP, command, or NTS-KE port. It does not include accesses over the Unix domain command socket.

The -p option specifies the minimum number of received NTP or command packets, or accepted NTS-KE connections, needed to include a client in the list. The default value is 0, i.e. all clients are reported. With the -k option the last four columns will show the NTS-KE accesses instead of command accesses. If the -r option is specified, chronyd will reset the counters of received and dropped packets or connections after reporting the current values.

An example of the output is:

```
Hostname                  NTP   Drop Int IntL Last     Cmd   Drop Int  Last
===============================================================================
localhost                   2     0   2   -   133      15     0  -1    7
foo.example.net            12     0   6   -    23       0     0   -    -
```

Each row shows the data for a single host. Only hosts that have passed the host access checks (set with the allow, deny, cmdallow and cmddeny commands or configuration file directives) are logged.

The intervals are displayed as a power of 2 in seconds.

The columns are as follows:

1. The hostname of the client.
2. The number of NTP packets received from the client.

3. The number of NTP packets dropped to limit the response rate.
4. The average interval between NTP packets.
5. The average interval between NTP packets after limiting the response rate.
6. Time since the last NTP packet was received
7. The number of command packets or NTS-KE connections received/accepted from the client.
8. The number of command packets or NTS-KE connections dropped to limit the response rate.
9. The average interval between command packets or NTS-KE connections.
10. Time since the last command packet or NTS-KE connection was received/accepted.

serverstats

The serverstats command displays NTP and command server statistics.

An example of the output is shown below.

```
NTP packets received       : 1598
NTP packets dropped        : 8
Command packets received   : 19
Command packets dropped    : 0
Client log records dropped : 0
NTS-KE connections accepted: 3
NTS-KE connections dropped : 0
Authenticated NTP packets  : 189
Interleaved NTP packets    : 43
NTP timestamps held        : 44
NTP timestamp span         : 120
```

The fields have the following meaning:

NTP packets received

The number of valid NTP requests received by the server.

NTP packets dropped

The number of NTP requests dropped by the server due to rate limiting (configured by the ratelimit directive).

Command packets received

The number of command requests received by the server.

Command packets dropped

The number of command requests dropped by the server due to rate limiting (configured by the cmdratelimit directive).

Client log records dropped

The number of client log records dropped by the server to limit the memory use (configured by the clientloglimit directive).

NTS-KE connections accepted

The number of NTS-KE connections accepted by the server.

NTS-KE connections dropped

The number of NTS-KE connections dropped by the server due to rate limiting (configured by the ntsratelimit directive).

Authenticated NTP packets

The number of received NTP requests that were authenticated (with a symmetric key or NTS).

Interleaved NTP packets

The number of received NTP requests that were detected to be in the interleaved mode.

NTP timestamps held

The number of pairs of receive and transmit timestamps that the server is currently holding in memory for clients using the interleaved mode.

NTP timestamp span

The interval (in seconds) covered by the currently held NTP timestamps.

Note that the numbers reported by this overflow to zero after 4294967295 (32-bit values).

allow [all] [subnet]

The effect of the allow command is identical to the allow directive in the configuration file.

The syntax is illustrated in the following examples:

allow 1.2.3.4

allow all 3.4.5.0/24

allow 2001:db8:789a::/48

allow 0/0

allow ::/0

allow

allow all

deny [all] [subnet]

The effect of the allow command is identical to the deny directive in the configuration file.

The syntax is illustrated in the following examples:

deny 1.2.3.4

deny all 3.4.5.0/24

deny 2001:db8:789a::/48

deny 0/0

deny ::/0

deny

deny all

local [option]..., local off

The local command allows chronyd to be told that it is to appear as a reference source, even if it is not itself properly synchronised to an external source. (This can be used on isolated networks, to allow one computer to be a master time server with the other computers slaving to it.)

The first form enables the local reference mode on the host. The syntax is identical to the local directive in the configuration file.

The second form disables the local reference mode.

smoothing

The smoothing command displays the current state of the NTP server time smoothing, which can be enabled with the smoothtime directive.

An example of the output is shown below.

Active : Yes

Offset : +1.000268817 seconds

Frequency : -0.142859 ppm

Wander : -0.010000 ppm per second

Last update : 17.8 seconds ago

Remaining time : 19988.4 seconds

The fields are explained as follows:

Active

This shows if the server time smoothing is currently active. Possible values are Yes and No. If the leaponly option is included in the smoothtime directive, (leap second only) will be shown on the line.

Offset

This is the current offset applied to the time sent to NTP clients. Positive value means the clients are getting time that?s ahead of true time.

Frequency

The current frequency offset of the served time. Negative value means the time observed by clients is running slower than true time.

Wander

The current frequency wander of the served time. Negative value means the time observed by clients is slowing down.

Last update

This field shows how long ago the time smoothing process was updated, e.g. chronyd accumulated a new measurement.

Remaining time

The time it would take for the smoothing process to get to zero offset and frequency if there were no more updates.

smoothtime activate, smoothtime reset

The smoothtime command can be used to activate or reset the server time smoothing process if it is configured with the smoothtime directive.

Monitoring access

cmdaccheck address

This command is similar to the accheck command, except that it is used to check whether monitoring access is permitted from a named host.

Examples of use are as follows:

```
cmdaccheck foo.example.net
cmdaccheck 1.2.3.4
cmdaccheck 2001:db8::1
```

cmdallow [all] [subnet]

This is similar to the allow command, except that it is used to allow particular hosts or subnets to use chronyc to monitor with chronyd on the current host.

cmddeny [all] [subnet]

This is similar to the deny command, except that it is used to allow particular hosts or subnets to use chronyc to monitor chronyd on the current host.

Real-time clock (RTC)

rtcdata

The rtcdata command displays the current RTC parameters.

An example output is shown below.

```
RTC ref time (GMT) : Sat May 30 07:25:56 2015
Number of samples  : 10
Number of runs     : 5
Sample span period :  549
RTC is fast by     :    -1.632736 seconds
RTC gains time at  :  -107.623 ppm
```

The fields have the following meaning:

RTC ref time (GMT)

This is the RTC reading the last time its error was measured.

Number of samples

This is the number of previous measurements being used to determine the RTC gain or loss rate.

Number of runs

This is the number of runs of residuals of the same sign

following the regression fit for (RTC error) versus (RTC time).
A value which is small indicates that the measurements are not well approximated by a linear model, and that the algorithm will tend to delete the older measurements to improve the fit.

Sample span period
This is the period that the measurements span (from the oldest to the newest). Without a unit the value is in seconds; suffixes m for minutes, h for hours, d for days or y for years can be used.

RTC is fast by
This is the estimate of how many seconds fast the RTC when it thought the time was at the reference time (above). If this value is large, you might (or might not) want to use the trimrtc command to bring the RTC into line with the system clock. (Note, a large error will not affect chronyd's operation, unless it becomes so big as to start causing rounding errors.)

RTC gains time at
This is the amount of time gained (positive) or lost (negative) by the real time clock for each second that it ticks. It is measured in parts per million. So if the value shown was +1, suppose the RTC was exactly right when it crosses a particular second boundary. Then it would be 1 microsecond fast when it crosses its next second boundary.

trimrtc
The trimrtc command is used to correct the system?s real-time clock (RTC) to the main system clock. It has no effect if the error between the two clocks is currently estimated at less than a second.

The command takes no arguments. It performs the following steps (if the RTC is more than 1 second away from the system clock):

1. Remember the currently estimated gain or loss rate of the RTC and flush the previous measurements.

2. Step the real-time clock to bring it within a second of the system clock.

3. Make several measurements to accurately determine the new offset between the RTC and the system clock (i.e. the remaining fraction of a second error).

4. Save the RTC parameters to the RTC file (specified with the rtcfile directive in the configuration file).

The last step is done as a precaution against the computer suffering a power failure before either the daemon exits or the writertc command is issued.

chronyd will still work perfectly well both whilst operating and across machine reboots even if the trimrtc command is never used (and the RTC is allowed to drift away from true time). The trimrtc command is provided as a method by which it can be corrected, in a manner compatible with chronyd using it to maintain accurate time across machine reboots.

The trimrtc command can be executed automatically by chronyd with the rtcautotrim directive in the configuration file.

writertc

The writertc command writes the currently estimated error and gain or loss rate parameters for the RTC to the RTC file (specified with the rtcfile directive). This information is also written automatically when chronyd is killed (by the SIGHUP, SIGINT, SIGQUIT or SIGTERM signals) or when the trimrtc command is issued.

Other daemon commands

cyclelogs

The cyclelogs command causes all of chronyd's open log files to be closed and re-opened. This allows them to be renamed so that they can be periodically purged. An example of how to do this is shown below.

```
# mv /var/log/chrony/measurements.log /var/log/chrony/measurements1.log
# chronyc cyclelogs
# rm /var/log/chrony/measurements1.log
```

dump

The dump command causes chronyd to write its current history of measurements for each of its sources to dump files in the directory specified in the configuration file by the dumpdir directive and also write server NTS keys and client NTS cookies to the directory specified by the ntsdumpdir directive. Note that chronyd does this automatically when it exits. This command is mainly useful for inspection whilst chronyd is running.

rekey

The rekey command causes chronyd to re-read the key file specified in the configuration file by the keyfile directive. It also re-reads the server NTS keys if ntsdumpdir is specified and automatic rotation is disabled in the configuration file.

reset sources

The reset sources command causes chronyd to drop all measurements and switch to the unsynchronised state. This command can help chronyd with recovery when the measurements are known to be no longer valid or accurate, e.g. due to moving the computer to a different network, or resuming the computer from a low-power state (which resets the system clock). chronyd will drop the measurements automatically when it detects the clock has made an unexpected jump, but the detection is not completely reliable.

shutdown

The shutdown command causes chronyd to exit. This is equivalent to sending the process the SIGTERM signal.

Client commands

dns option

The dns command configures how hostnames and IP addresses are resolved in chronyc. IP addresses can be resolved to hostnames when printing results of sources, sourcestats, tracking and clients commands. Hostnames are resolved in commands that take an address as argument.

There are five options:

dns -n

Disables resolving IP addresses to hostnames. Raw IP addresses will be displayed.

dns +n

Enables resolving IP addresses to hostnames. This is the default unless chronyc was started with -n option.

dns -4

Resolves hostnames only to IPv4 addresses.

dns -6

Resolves hostnames only to IPv6 addresses.

dns -46

Resolves hostnames to both address families. This is the default behaviour unless chronyc was started with the -4 or -6 option.

timeout timeout

The timeout command sets the initial timeout for chronyc requests in milliseconds. If no response is received from chronyd, the timeout is doubled and the request is resent. The maximum number of retries is configured with the retries command.

By default, the timeout is 1000 milliseconds.

retries retries

The retries command sets the maximum number of retries for chronyc requests before giving up. The response timeout is controlled by the timeout command.

The default is 2.

keygen [id [type [bits]]]

The keygen command generates a key that can be added to the key file (specified with the keyfile directive) to allow NTP authentication between server and client, or peers. The key is generated from the /dev/urandom device and it is printed to standard output.

The command has three optional arguments. The first argument is the key number (by default 1), which will be specified with the key

option of the server or peer directives in the configuration file.

The second argument is the name of the hash function or cipher (by default SHA1, or MD5 if SHA1 is not available). The third argument is the length of the key in bits if a hash function was selected, between 80 and 4096 bits (by default 160 bits).

An example is:

```
keygen 73 SHA1 256
```

which generates a 256-bit SHA1 key with number 73. The printed line should then be securely transferred and added to the key files on both server and client, or peers. A different key should be generated for each client or peer.

An example using the AES128 cipher is:

```
keygen 151 AES128
```

exit, quit

The exit and quit commands exit from chronyc and return the user to the shell.

help

The help command displays a summary of the commands and their arguments.

SEE ALSO

chrony.conf(5), chronyd(8)

BUGS

For instructions on how to report bugs, please visit https://chrony.tuxfamily.org/.

AUTHORS

chrony was written by Richard Curnow, Miroslav Lichvar, and others.