NAME

rrd-ntp – description of the NTP monitoring graphs

DESCRIPTION

The *ntp.xml* file in the **rrdmon** package describes a couple of data channels retrieved from the NTP daemon **ntpd**(8) that can help to assess the use of the NTP time synchronization service. The following sections describe the graphs.

NTP STRATUM AND ROOT DISPERSION

This graph contains a visual representation of the current stratum and the root dispersion. The *stratum* is the number of network hops required to reach a reference clock. In normal operation, a stratum value between 2 and 5 is expected. After startup, the NTP daemon **ntpd**(8) has the default stratum 16 associated with it, indicating that is not yet synchronized.

The *root dispersion* value, measured in milliseconds, indicates the variance of the accumulated round trip time back to the root servers, also called the *root/*delay*,* see the section NTP SYSINFO below.

NTP SYSINFO

This graph combines a few performance numbers from the output of the *sysinfo* command of the **ntpq**(8) query client.

The root delay is accumulated round trip time to the stratum 1 server. This is an approximate upper bound of the time error of this server.

The clock jitter and wander are weighted RMS averages of the local clocks offset derived by the local clock module.

NTP PEER

This graph shows some basic performance information of the peer server currently used to synchronize with. The light blue background shows the *poll intervall*. In normal operation, the interval starts at 1 second and is then double as synchronization is established. The maximum value is 1024. If no responses are received from the peer or if the precision is in doubt, the poll interval is reduced.

The blue *delay* curve shows the time it takes the peer to respond to the request. A large delay makes precise time determination difficult.

The green *jitter* value shows the variance of the various time comparisons of local and peer clock derived from the NTP requests.

NTP OFFSETS

The time offset on the other hand is the actual time difference in milliseconds between the local clock and the reference clock. It is displayed as a dark red curve.

The dark blue curve is the PLL offset value that is used to correct the PLL frequency (see next section).

NTP KERNEL STATISTICS

The NTP daemon **ntpd**(8) keeps a PLL to coerce the kernel clock locked to the reference clock. This PLL frequency value indicates how much the clock frequency of the local clock needs to change to keep the reference clock and the local clock in sync. The PLL frequency is controlled by the offset displayed in the NTP OFFSETS graphs. If the PLL offset increases, the PLL frequency is adapted to bring the clock back into sync with the reference. A large PLL frequency value indicates that uncontrolled, the local clock

deviates from standard time quickly.

From the performance of the PLL, the NTP daemon **ntpd**(8) derives an estimation for the error, which is displayed as a blue curve.

NTP I/O STATISTICS

The NTP I/O statistics graph reports the number of packets sent and received by the daemon.

NTP ADDRESSES

The NTP daemon remembers addresses of clients that have talked to it. This graph shows the peak number of addresses remembered as a violett background and the current numer of addresses as a dark blue line. The Intranet servers will have very large numbers of addresses, but infrastructure servers on the path to the Internet or the local NTP daemon running an every machine will only have small numbers, usuall not more than about 3 addresses.

SEE ALSO

rrdsetup(8), **rrdupdate**(8)

AUTHOR

Andreas Mueller, andreas.mueller@othello.ch