W H I T E P A P E R

Introduction to *NTP*

Time is not just an extraneous network service. Accurate time is essential to determining the order in which events occur and is a fundamental aspect of transaction integrity, logging/auditing, troubleshooting and forensics. Accurate, reliable time is necessary for financial and legal transactions, transportation and distribution systems, database management and many other applications involving widely distributed resources.

NETWORK TIME PROTOCOL (NTP)

By far, the most widely used and accepted method for maintaining accurate time across entire networks is an implementation of the Network Time Protocol (NTP). NTP is one of the oldest, most-used protocols on the Internet and is built on the Internet Protocol (IP) and User Datagram Protocol (UDP). It is specifically designed to maintain time accuracy and reliability, even when used over typical Internet paths involving multiple gateways and unreliable networks. Simple Network Time Protocol (SNTP) is a simpler version of NTP which is compatible with NTP.

OPERATION

NTP is not based on the principle of synchronizing machines with each other. It is based on the principles of having all machines get as close as possible to the correct time - Universal Coordinated Time (UTC). A basic NTP network is composed of a time server and clients (workstations, routers, other servers, etc.). The function of a time server is to provide accurate time to the clients.

The individual clients run a small program as a background task that periodically queries the server for a precise UTC time reference. These queries are performed at designated time intervals (generally about every 15 minutes) in order to maintain the required synchronization accuracy for the network. The basic operation of the NTP is time stamping of data packets transferred between the server and the client.

The order of operations is:

1. The client stamps the time when he sends an NTP request packet to the server.
3. The server stamps the time when the NTP request packet is received from the client.
3. The server stamps the time when he sends the NTP reply packet back to the client.
4. The client stamps the time when this NTP reply packet is received.

So, an NTP packet consists of 4 timestamps. The client uses these timestamps to determine the difference between its internal time and the UTC time reference and adjusts its local time to coincide with the reference. The client can also determine the network latency and apply a correction factor when it adjusts its internal time. The ability to remove the network latency results in a more precise level of synchronization, often in the few millisecond range. However, not all the latency can be removed unless the paths to and from the server are symmetrical.

HIERARCHY

The NTP protocol has a hierarchical design in order to prevent large numbers of clients from accessing the same primary time sources. This hierarchy should be adhered to, and a large number of clients should not be configured to overload a busy stratum 1 time server. In addition, networks should be designed to minimize the number of servers that interact with public NTP servers. At the top of the hierarchy is what is accepted as the actual time - usually UTC. Each timeserver is assigned a "stratum" level that corresponds with its distance from an accurate time source. Stratum 1 servers have direct access to a UTC time source (GPS, CDMA, WWV). Stratum 2 servers receive their synchronization from stratum 1 servers. Stratum 3 servers receive time from stratum 2 servers and so on.

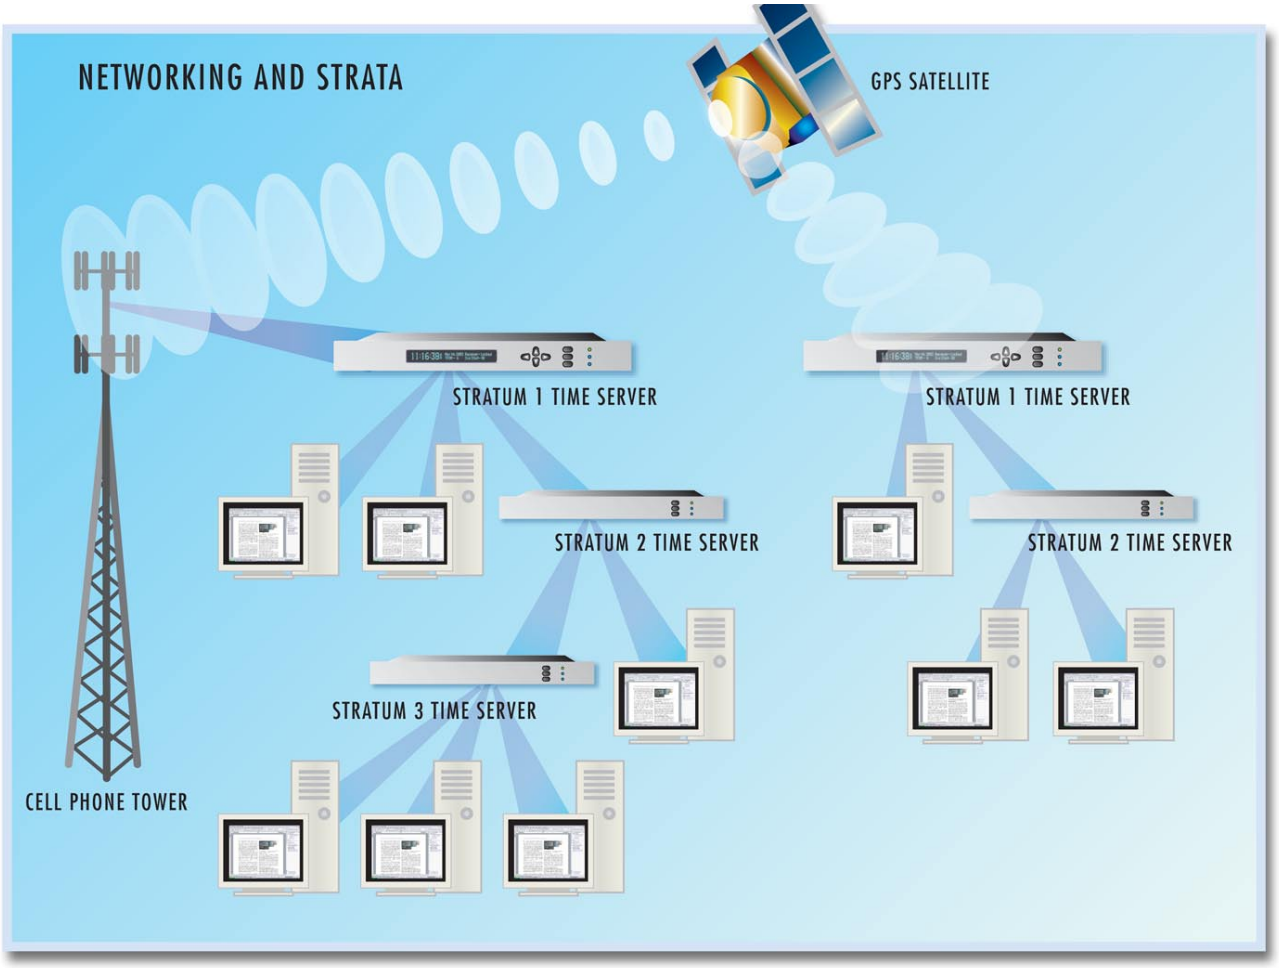

Administrators who are building an NTP-based time infrastructure have four choices:

1. Obtain an NTP Server appliance to use as a stratum 1 server. This is the easiest choice for providing an accurate, reliable, secure and autonomous UTC-synchronized network.

2. Obtain an external time source such as a GPS or CDMA reference to create a stratum 1 server. This external time reference is then connected to an existing server to create a stratum 1 time server. Although this method is more difficult to setup and configure it will provide an accurate, reliable, secure and autonomous UTC-synchronized network.

3. Synchronize an internal NTP server from publicly available servers on the Internet, making it a stratum 2 or 3 server. However, as with any externally provided service, it is also an entry point for attackers. In addition, obtaining time from the Internet is less accurate. For secure environments where synchronized time is critical, the use of a public time server would not be appropriate.

4. Designate a machine as the time authority, using its internal clock as the arbitrary time source. However, as this time source wanders all of the NTP clients connected to it will wander with it. While the primary clock could be manually adjusted to the true time occasionally, this would cause all of the clients to jump when the server adjusts. If a clock is ever adjusted to shift more than 17 minutes, all of the NTP client software will abort due to the sudden time shift. This option can still provide a synchronized network and may be acceptable in a few rare cases, but in any sort of large installation it is critical to keep the clocks synchronized with some maintained time standard.

ACCURACY

The degree of server accuracy is dependent upon the choice made from the list above. The degree of client synchronization to a server is dependent primarily on network latency. Anything that adds latency, such as hubs, switches, routers, or network traffic, will reduce accuracy. Under good conditions on a LAN without too many sources of network delay, synchronization accuracy is typically within the ½ -2 millisecond range. In other words, all clients on the network will be synchronized to UTC and to each other to within ½ -2 milliseconds. The synchronization accuracy on a WAN is typically within the range of 10-100 milliseconds. For the Internet, synchronization accuracy is unpredictable, so special attention is needed when configuring a client to use public NTP servers. *The final achievable accuracy of each client depends on the accuracy of the time server used, the network latency, and the symmetry of the network paths to and from the time server.* Obviously, no client can be more accurate than its server.

SECURITY

NTP is founded on the UDP protocol. As such, it is highly susceptible to IP spoofing. The NTP protocol uses UDP port 123. Blocking this port at the firewall is a minimum requirement for network perimeter security. This prohibits a client from obtaining time from a public NTP server on the Internet. Security is maximized when the time server is installed within the network firewall. The time server acquires time from the GPS, WWV, or CDMA system, via an antenna, with no threat to network security. It then distributes the time to the clients over the network within the firewall.

To enhance security from threats contained within the firewall, the time server should use the access control and authentication facilities in NTP to restrict access to the service. If possible, only authenticated NTP packets should be accepted. The server should also accept packets from only known, approved sources. For communicating with the time server for status and control it is best to use a secure protocol such as Secure Shell (SSH) and/or SNMP v3 (encrypted SNMP). (SNMP v1 and v2c are not secure). In addition, security is further enhanced if risky protocols such as FTP, Telnet, Time and Daytime can be disabled.

RESOURCE REQUIREMENTS

NTP requires little resource overhead, allowing it to be deployed on servers hosting other services, even if the servers are heavily loaded. The bandwidth requirements are also minimal, as NTP packets are only 60 bytes long. An NTP client/server transaction requires 2 packets per transaction. When first started, transactions occur about once per minute, gradually changing to once every 15 minutes. Obviously the network resources required to support NTP are minimal.

CLIENT SOFTWARE

In addition to adding a network time server to your network, the NTP client software must be installed on each of the workstations or servers that will interface with the time server. This software is operating system/platform specific. In most cases the client software is already resident in the operating system of the workstation, server, router or firewall. In other cases the client software is available as free, public-domain software or as low-cost shareware. For information on NTP client software go to www.ntp.org or www.endruntechnologies.com/ntp-client.htm.

NTP VERSIONS

The first version of NTP was released in 1985, but is no longer in common usage. The specifications for version 3 were released in 1992 (RFC 1305). The finalized specifications for version 4 are in-progress and have yet to be released but implementations of NTP v4 are widely available. NTP versions can interoperate and most commercial NTP Time Servers run NTP v4. Some of the newer features include advanced cryptographic authentication.

CONCLUSION

The need for synchronized time is critical for today's network environments. Using NTP is an excellent way to keep a large number of network nodes in close synchronization. NTP requires a minimum of network overhead and can also maintain a high level of synchronization accuracy and security. In addition, an effectively designed NTP infrastructure is relatively easy to implement, making NTP ideal for both small and large enterprise networks.

EndRun TECHNOLOGIES
"Smarter Timing Solutions"
Santa Rosa, CA, USA
TEL 1-877-749-3878
FAX 707-573-8619
www.endruntechnologies.com