

The Real But Hidden Danger of Network Timing Breach

How accurate timing can reduce network security risk and increase productivity

Security and Accurate Time Synchronization

- Establishes the correct “when” of historical (past) events
- Allows ordering of future planned events
- Enables synchronized, real-time interaction between network devices and processes

Most anyone who uses a PC no doubt sees the little clock at the bottom of the screen, dutifully ticking off the minutes of the passing day. But chances are they don't often think about the role that clock plays in the network to which the PC is attached. In fact, virtually every piece of equipment attached to a data network has a similar clock, although it may not actually show the time of day to anyone.

Until, that is, something goes wrong. Should a portion of the network go down, the clocks on every network device suddenly take on added importance. Administrators will turn to network management systems, which continuously collect log files from network devices, to try to determine what went wrong. That effort will be made far more difficult, if not futile, if the events depicted on the log files are not all working off the same, synchronized clock. If a router thinks the time is 12:05 p.m. but an application server thinks it's 12:15 p.m., good luck to the network administrator responsible for figuring out what happened and when.

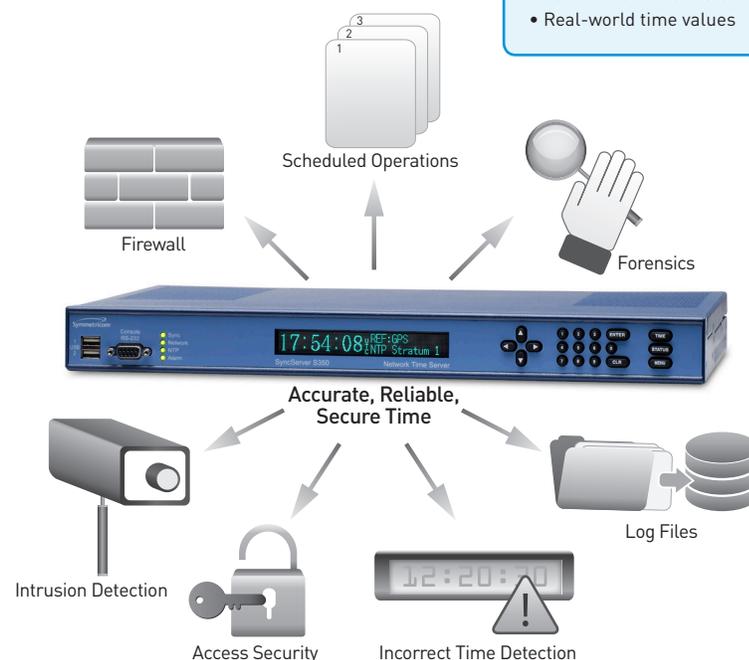
Fortunately, IT professionals typically understand the importance of time to proper network operations. Routine events such as data backup and directory synchronization, as well as potentially crucial applications such as the time stamp on a financial transaction, depend on various components agreeing on the answer to the simple question, “What time is it?”

IT pros may differ, however, on how they enable their network components to answer that question. Many organizations keep time by using one or more of the publicly available time servers that act as sources of Coordinated Universal Time,

or UTC, an internationally recognized time standard. In order to synchronize an internal clock with UTC, the Network Time Protocol (NTP) allows any system to do a simple query over the Internet to one of these time servers. However, in order to enable its systems to sync up, the organization must essentially leave a door open on its firewall through which time information can get in.

Network applications at risk due to security breach caused by inaccurate time synchronization:

- Log file accuracy, auditing and monitoring
- Network fault diagnosis and recovery
- Access security and authentication
- Scheduled operations
- False time server detection/rejection
- Real-world time values



Network security application areas that rely on secure, accurate and reliable time.

The Real But Hidden Danger of Network Timing Breach

From a security perspective, that is a big problem. Nefarious hackers know all too well that many companies use this technique to maintain network time. They even know exactly what door on the firewall will be open, since NTP is designed to use a specific port on a firewall — Port 123. In effect, then, this approach to keeping network time is like sending an invitation to an attacker, telling him how to attack your network. In this paper, we will look at why it's important for any network to keep its devices in sync and examine the security implications — and financial ramifications — of improper network time. We will also explain an approach to keeping proper time that doesn't require you to leave holes in your network defenses.

The Importance of Time

In our everyday lives, we often take time for granted, even if we never seem to have enough of it. You look at your watch to determine when it's time for a meeting, or lunch. If it's off by one or two minutes, that's not cause for alarm.

But in some instances, it's crucial that time be highly accurate. In an online stock trade, for example, that same one or two minutes could mean a dramatic difference in the price at which a block of stock is bought and sold. In that instance, time really is money.

Yet, until fairly recently, about the only legally accepted measure of time was a U.S. Postal Service postmark. If your tax return was postmarked on or before April 15, for example, it was considered on time by the Internal Revenue Service.

Increasingly, the concept of legally acceptable time stamps is being introduced into the electronic world. Since 1999, the Securities and Exchange Commission, as part of its Order Audit Trail System (OATS), has required that NASD member firms synchronize all business clocks, including those on computer systems, and that all electronic orders be time-stamped.

Similarly, the Food and Drug Administration, in its 21 CFR Part 11 guidelines for trustworthy electronic records, requires that companies use computer-generated time stamps. The regulations require procedures and controls to ensure the time stamps are accurate, including a secure synchronization mechanism.

As government gets more involved in how electronic data is handled, through regulations including the Health Insurance Portability and Accountability Act (HIPAA) and Gramm-Leach-Bliley Act (GLBA), it stands to reason that being able to prove definitively when various events happened will take on added importance.

Government mandates aside, it makes good business sense to ensure proper time because it is so important to proper network operation. In addition to the log file issue discussed above, which demonstrates the importance of time synchronization to network fault diagnosis and forensics, time plays a crucial role in scheduled operations such as data backups. If systems are out of sync, backups may not run properly, and data could be lost. Similarly, file systems such as the Network File System rely on time stamps to determine which file is the most recent. If different PCs in the same network are out of sync by several minutes or more, such file systems may well save the wrong copy of a file, causing the one with more recent changes to be discarded.

Security and Public Time Servers

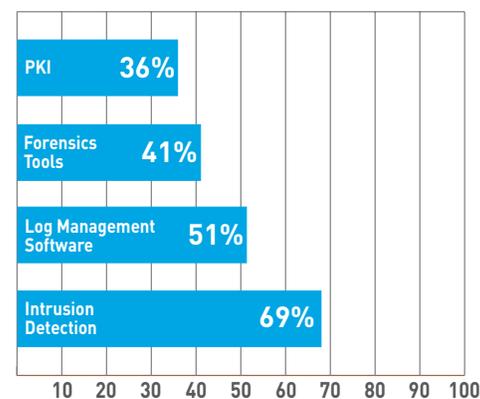
Such occurrences can be costly in terms of the lost productivity that results when employees lose files that represent even an hour's worth of work. But they often pale in comparison to the security risks related to using public, Internet-based time servers to keep systems synchronized.

The danger inherent in leaving Port 123 open in a firewall was made all too clear with the SoBig.F, the worm that hit the Internet in mid-August 2003. According to the security software firm Symantec,

SoBig.F was programmed to launch itself only on Friday or Saturday between 7 p.m. and 10 p.m. UTC time. The payload was a program that downloaded and executed files on an infected computer. Those files had a number of purposes, according to Symantec: The worm's author could use them to steal confidential information, set up servers that would relay spam email or even access a Web site controlled by the worm's author and retrieve more files of the author's choosing.

Key to the success of SoBig.F, however, was obtaining UTC time. It did so using NTP and contacting a public time server via Port 123. In other words, SoBig.F needed Port 123 to be open in order to execute. If Port 123 were closed, SoBig.F couldn't determine the accurate time, and none of the nefarious files would be downloaded.

That's not the only way that NTP and Port 123 can be responsible for a security breach, however. Should a hacker find a network that is getting its network time via NTP and Port 123, it could launch a denial of service (DoS) attack. A DoS attack essentially involves flooding an application or device with more data than it can handle. In this case, the intruder might flood the time server, sending the



Time dependent security technologies used by respondents to the 2008 CSI Computer Crime and Security Survey. 69% use Intrusion detection, 51% log management software, 41% forensics tools and 36% PKI infrastructure systems.

The Real But Hidden Danger of Network Timing Breach

NTP packets far more data than normal, possibly causing the program to crash, which means the network can no longer access the correct time.

Alternatively, the intruder could construct packets that appear to be coming from the time server an organization uses or even launch an attack from that time server. In either case, he may be able to break into the server at the user organization that hosts its NTP program. Given that NTP programs typically have system administrator rights, the intruder could then invade the rest of the network.

Out of Sync Means Insecure

Not all time-related security risks relate to NTP and public time servers, however. Poor synchronization in general can lead to unacceptable levels of risk.

Intrusion Detection and Forensics Analysis

Let's say your organization experiences a network break-in. One of the first steps to take after discovering such an event is to conduct a forensic investigation to determine exactly what resources the hacker attacked and, if possible, how he got access. Such an investigation involves a thorough examination of various network device logs, which collect information on each packet that travels through the network. By examining logs from devices such as firewalls, routers and servers, a forensic investigator can essentially follow a hacker's path to determine how he got into the network and what he did once he got in.

Such an examination is not easy. It requires an experienced professional and/or advanced tools to put the puzzle together. But if the network devices under investigation are not accurately synchronized with one another, the job becomes even more complicated. Now the investigator will have to allow for discrepancies in the log files coming from each device: Router A is three minutes slower than Server B but

two minutes ahead of Firewall C. It can be a maddening exercise, at a time when every minute counts.

Hackers, of course, know this well. As a way to cover their tracks, they may try to disrupt the time on any system they access. If the organization has no way of knowing that the time on a given machine was tampered with, it may never discover how the intruder broke in or what he did. That effectively leaves the door open for the same intruder to come back.

Access Security and Authentication

Other time-related security issues may not involve an attack of any kind. Consider authentication systems that allow users to log in to various applications, for example. Such systems, including the Kerberos-based system in Windows and systems such as RSA Security's SecurID, require some level of time synchronization between the client machine requesting access and the server that grants it. If the two aren't within an allowable time difference, access can be denied.

The simple solution may seem to be to set the allowable time differential to a value great enough that it won't likely be breached, maybe 10 or 15 minutes. But that effectively gives a potential intruder 10 or 15 minutes to try to force his way into a system — far too much time for comfort. Deploying network time synchronization allows you to shorten the allowable time differential window and increase network security.

Cost Analysis

All of these security risks come with an associated cost. The Sobig.F worm, for example, is estimated to have caused \$29.7 billion in economic damage worldwide, the most ever attributed to a single worm or virus, according to mi2g Ltd., a risk management firm in London. Much of that number comes from productivity losses. Consider the cost at your own company if

every employee who uses a PC has to stop using it for even 20 or 30 minutes while the IT department runs a scan to verify that the machine is Sobig-free.

DoS attacks can likewise result in lost dollars in any number of ways. If an intruder succeeded in taking a server down, any applications hosted on that server might be inaccessible until the root of the problem was found. If the hacker successfully changed the time on one or more network components during his break-in, that would make the forensics investigation all the more difficult, increasing the time it took to identify the cause of the attack and get the server back online. Such an investigation could easily take hours, if not days. Perhaps worse, an intruder who successfully broke into a network might steal valuable data.

The ROI From a Time Server

Consider the potential Return on Investment (ROI) of having the ability to access accurate time behind the firewall combined with software that prevents intruders from altering the time on any device.

The 2008 CSI Computer Crime and Security Survey cited an entire collection of security technologies used to protect network operations. While time plays a role in many of technologies, it plays a major role in determining the critical "when" of several of the key technologies. In particular, respondents to the survey cited intrusion detection (69%), log management software (51%), forensics tools (41%) and PKI infrastructure systems (36%) as technologies used to protect the network. Without accurate time synchronization to UTC across the network the effectiveness of these tools in the role of security becomes marginal.

While survey data can be somewhat abstract, the dollar losses are all too real. The same survey reported the average annual loss due to security incidents was just under \$300,000. This makes a case

The Real But Hidden Danger of Network Timing Breach

for the inexpensive addition of good time keeping as a defense against the direct financial loss of a security incident. To further make this point clear, even the U.S. Department of Justice Digital Forensic Analysis Methodology (2007) requires the essential time component of establishing the "When" of a computer crime. In particular, during the data analysis phase, key questions that must be asked are:

- When was it [the evidentiary data] created, accessed, modified, received, sent, viewed, deleted and launched?
- Does it show when relevant events took place?
- Time Analysis: What else happened to the system at the same time?"

The only possible way to answer these questions is if the systems were synchronized to the correct time. Harder to quantify is the financial impact due to the lost productivity that can result from out-of-sync authorization systems, when employees can't log on to their workstations or to an application they need to do their jobs. In a network where multiple machines are out of sync, this can be a common occurrence.

Network security application areas improved via secure time synchronization:

- Secure and accurate time inside the firewall
- Improved intrusion detection and forensic analysis
- Log file accuracy, auditing & monitoring
- Network fault diagnosis and recovery
- Access security and authentication
- Scheduled operations
- False time server detection/rejection
- Real-world time values

Financial firms understand the importance of having synchronized clocks, to ensure the time stamps on transactions are accurate, as NASD OATS regulations dictate. As more and more business is conducted online, more industries will have to come to grips with the importance of accurate time stamps to ensure the integrity of their business transactions.

Keeping in Sync

While the consequences of poor time synchronization can be dire for computer networks, keeping them in sync without exposing the network to security risk is a simple matter. What's required is an accurate source of time to which networked systems can refer without having to poke holes in the corporate firewall. Simply put, that means installing a time server that sits behind the firewall.

Symmetricon, Inc., for example, makes a line of time servers that synchronize to satellites that constantly orbit the earth. Each of the 30+ satellites that are used to support the Global Positioning System (GPS) is outfitted with three atomic clocks. Symmetricon time servers synchronize with the clocks on those satellites, which are accurate to approximately one-millionth of a second to UTC. Each networked system can then synchronize with the Symmetricon device as often as is deemed necessary.

Symmetricon's family of GPS NTP/PTP network time synchronization products with atomic clock accuracy ensure that time is synchronized and accurate throughout your IT enterprise network. Our NTP time servers protect network log file accuracy, security, billing systems, electronic transactions, database integrity, VoIP, and many other essential applications. Our GPS NTP Network Appliances make it easy to manage and monitor time synchronization through workstations, servers and routers, thus assuring the highest integrity throughout the network and can offer a high level of security to ensure that your business is always protected.

Even an entry-level Symmetricon time server can keep thousands of client systems in sync. For a relatively small investment, companies can be assured that out-of-sync systems won't result in far more costly productivity or data losses and that keeping accurate time won't open up their network to security risk. It's time to stop taking time for granted. For more information on Symmetricon time servers, visit:

<http://www.symmetricon.com/industry-solutions/enterprise-it/>