Symmetricom

WHITE PAPER

Five Dangers of Poor Network Timekeeping

Easy and cost effective solutions to avoid networks fall out of sync

Introduction

Most organizations today rely on networks of computers, all of which rely on clocks. So what happens when the clocks in these computers don't agree with each other – or with the correct time? What happens to the processes running on these networks? And, what happens to the organizations?

A time bomb is ticking away in the heart of most IT infrastructures – the same infrastructure on which organizations depend to produce products, buy from suppliers, sell to customers, prepare financial reports, and communicate both internally and externally, in short do all the things modern organizations do. When the clocks in a network fall out of sync – with each other or with the correct time – bad things start to happen. Processes fail. Data is lost. Security is compromised. Legal exposure increases. And organizations lose credibility with customers and business partners which eventually can lead to decreased revenue and profit.

Why is this time bomb allowed to exist? Because most people mistakenly assume that computer clocks are inherently accurate. They also don't fully appreciate the consequences when computers are made interdependent and when the clocks in these computers don't agree. Finally, they don't realize that solutions to the network synchronization issues are inexpensive and easy to implement. These solutions, called time servers, typically cost under $5,000 and can support a network consisting of thousands of computers. They are also virtually self-administering. In other words, there is absolutely no reason why any of the dangers discussed in this paper need ever threaten any company's operations again.

Five Dangers – An Overview

The negative consequences of running out-ofsync computers in a network fall into five main categories:

Operations Failure — Automated events – like data backups or order processing – simply don't occur or else they break. This happens because a trigger set to go off at a certain time does not or because tasks that are supposed to be carried out on different computers can not be executed in the proper sequence.

Data Loss — This occurs because system software like directory services erroneously saves an out-of-date version of a file instead of the latest version.

Security Holes — These occur both directly and indirectly and as a result of poor timekeeping. For example, most systems set time using an industry standard protocol called NTP (Network Time Protocol) that opens the firewall to hackers. Other security lapses occur because administrators cannot retrace hacker activities since log files are inaccurate. A third example is when security applications (like badge readers), designed to protect company assets, stop working.

Legal Liability — In a commercial dispute, there is no way to prove that transactions took place when alleged or that digital signatures on contracts are authentic.

Loss of Credibility — While any of these dangers can mean financial loss – so can the mere inability to demonstrate a competent business process. All business practices, and audits of business practices, involve time. So failure in this one area throws all other processes into question.

A major factor in each of these areas is the fact that time on a computer is usually measured in milliseconds or fractions of milliseconds. Given the fact that there might be thousands of events occurring simultaneously on any network (however large you want to define "network"), it makes timekeeping just that much more important, and difficult. Without an accurate, reliable source of time, there is simply no intuitive way to tell whether an organization is about to run into trouble.

Operational Failure

These failures cover a broad range of activities that touch almost every aspect of a company. Problems fall into three main areas:

- Automated tasks
- Network consolidated tasks
- Interdependent application tasks

Five Dangers of Poor Network Timekeeping

Automated area tasks are those like data backups that run overnight. These are typically multi-stage events, each of which occurs (or should occur) at a scheduled time. If one event is triggered out-of-sequence with the others, the entire process may fail. Furthermore, since these tasks often occur at off-hours, the likelihood is that the failure may not be discovered or corrected until the next day.

Some network-based tasks save resources by allowing a single machine to perform a common service that might otherwise have to be performed on multiple machines. Others, like time synchronization, are inherently network-centric and are optimally performed on a common machine. Either way, the process represents a single point of failure. Directory services utilize a common time source to schedule the order in which events occur. Should the server with the central time source be out of sync with the clocks of the various workstations and servers over which it has supervisory control, requests from their users and applications on those machines may not be recognized as valid.

A similar situation applies to distributed computing "middleware". Middleware is the "glue" tying together processes running on multiple machines so that they behave as a single coherent application – for example sales order administration where billing, point-of-sale, inventory control, and other systems all interoperate. In the case of IBM's DCE, if any of the clocks running the various processes are more than five minutes out of sync with DCE's Distributed Time Service, those processes will fail. That could be the billing system, a point-of-sale register, or virtually any other part of an infrastructure.

Computer events that require accurate, or accurately synchronized, time:

- Manufacturing process control
- Communications network time-of-day configuration
- Computer maintenance
- Funds transfers or purchases
- Database file time stamps (i.e. NFS, UNIX "make" process)
- Determining fault sequences via SNMP event traps
- Time stamping telephone and radio dispatch call records
- Employee time cards
- Measuring packet transit times
- Tracing intruder steps
- Time-dependent security processes (i.e. Kerberos authentication)
- Packet time-to-live stamps

Applications don't have to be part of a distributed computing environment, however, to be interdependent. In fact, commercial trading partners probably would not want to be locked into sharing a middleware layer just so they could do business. And they don't have to. There are many other ways applications can talk to each other – such as by using XML. Whatever the method, however, the demand for synchronicity remains high – such as when a parts manufacturer supplies "just in time" inventory to a carmaker. Each transaction, along with its various components, is time stamped, typically within a tolerance of one second. When a supplier's bill of materials or price, for example, are not received by a customer's waiting application within the required time window the application may simply move on to process another transaction. This could occur, for example, if a time stamp mistakenly indicates that information was sent before it was requested or arrived before it was sent. Another example is e-mail. If the user's e-mail program is configured to display messages sorted by time sent, a message with the wrong time could be easily overlooked.

Data Loss

One way operations can fail is by losing data. Unlike other kinds of breakdowns, however, data loss may go undetected until long after the fact – creating even more damage because people and applications relied on data they were led to believe was accurate. A prime example is a network file system, a network-consolidated resource that keeps track of date and times when files were created, last modified, last accessed, and last archived. If one of the machines in the networks submits a file stamped with a time earlier than the file maintained on the central server, the server may simply assume the file is old and discard it, along with any changes.

Another example is software development. Time stamps are used to indicate which pieces of a software program are the latest version – and should therefore become part of the production copy. If (as is often the case) different software modules are written on different machines, it is possible for versions to be listed on the central file system in a chronological sequence other than the one in which they were actually written. That means older modules may be erroneously substituted for newer modules. At best, this means that the production product won't have the latest features; or worse that the software will simply stop working – or work erratically – and that programmers might waste weeks looking for bugs that are not really there.

Five Dangers of Poor Network Timekeeping

Security Holes

How organizations keep track of time has a major impact on the overall security of the organization's IT infrastructure, for two reasons. First, the mechanisms used to keep track of time are among the most vulnerable to exploitation by a hacker. Second, time stamps are critical evidence for retracing a hacker's movements inside a target system – and therefore of hardening the system against future attacks.

With respect to time-related vulnerabilities, one the most common involves the Network Time Protocol (NTP). This program, present on virtually all computers, allows systems to synchronize their clocks with a time source over a UDP / IP network such as the Internet or a corporate local area network. A potential problem arises if this time source is located beyond the corporate firewall. If it is that means there must be "hole" left open in the firewall (specifically port 123) to allow packets containing the time information through. (Even if the time source is not outside the firewall, that does not automatically mean that port 123 is closed, only that it is not needed. The system administrator must still make sure unused ports are closed.) One way to exploit this opening is to crash the NTP program itself. This can be done (on several variations of both the Unix and Linux® operating systems) by sending too much data in an NTP packet. The result is a denial of time services and (depending on what else is happening on the network) potentially a crash of the network itself.

A second way to exploit NTP is to construct a packet which doesn't crash the NTP program, but instead uses that program to take over the target machine – using the same privileges as the NTP program itself (typically system administrator-level).

Even if the organization blocks all access to port 123 except from the external time source, that still leaves open the possibility that a hacker could attack the network from there.

A more insidious effect of weak timekeeping is that it damages the ability to investigate security breaches and other kinds of system problems. Hackers, for example, will often exploit backdoors, and proxy computers when mounting an attack – both to hide their tracks and to exploit whatever opportunities (like NTP system privileges) the hacker encounters along the way. Finding these stopping off points is critical for shutting the door to future attacks – and requires precise measurement of time in order to reconstruct the exact sequence of events. Log files and application time stamps obviously become essential pieces of evidence.

Of course, this is the same kind of evidence used for investigating system problems generally – not just hacker break-ins. Since network log files usually consist of time stamps from different machines, administrators can use them to reconstruct the events leading up to an incident occurring anywhere on the network. Performance related statistical information can also be collected and analyzed, allowing administrators to identify process bottlenecks and other opportunities for system optimization. All this of obviously depends on whether time stamps are synchronized to the correct time.

Finally, there is the performance of security systems themselves to consider – including firewalls, access card readers, and digital certificate authentication systems. Like many other systems an organization owns, these too can be compromised by weak network timekeeping. Take digital certificate authentication systems; these are used to check certificates used to authorize payments, sign contracts and carry out other sensitive business that requires proof of identity. As a security precaution, certificates are issued with a validity period and must periodically be renewed. If the authentication system clock is out of sync, an expired certificate may be accepted, potentially allowing for security breaches. A similar problem exists with firewalls, which may be opened temporarily during certain parts of the day – for example to perform maintenance or file uploading on remote servers. If their systems' clocks are not set correctly, these firewalls may be left open (or allowed to open) arbitrarily.

A reverse example is access card readers. Here an out-of-sync system may fail to recognize a legitimate card. That can occur because the card and the reader use the current time to generate an entry code. If their clocks are out of sync, so will be the codes and the cards will not work.

Legal Liability

Keeping accurate time on a network is more than just a technical issue – it is also a legal one. That's because time is used as a basis for making contracts. In the real world, people receive receipts, sign agreements, and audit performance. These documents, signatures, and transactions all include time references that make them legally binding. Recently both the United States and the European Union have passed laws making digitally signed documents legal. As contracts executed in cyberspace become more commonplace, parties to an online agreement or transaction will increasingly be called upon to prove that what was alleged to have occurred actually did occur, and when. Nowhere is that more true than in the brokerage business. Take the National Association of Security Dealers

Five Dangers of Poor Network Timekeeping

(NASD), a network consisting of 5,500 members and 82,000 branch offices. NASD requires its members to time stamp each stock trade to within an accuracy of three seconds. Furthermore, members must be able to prove that the time in the time stamps came from a recognized time source, specifically the National Institute of Standards and Technology (NIST). For transactions to be legal their time stamps must be accurate, and the accuracy must be proven. In digital commerce, merely synchronizing your own network clocks won't be good enough; they must also be synchronized with an external Coordinated Universal Time (UTC) source.

Loss of Credibility

Behind all of these dangers looms perhaps the greatest danger of all – loss of credibility in the marketplace. This seems obvious: the more operations fail, the more data that is lost, the more that security breaches occur, or the more legal liability that is incurred, the harder it will be to attract business. Companies will be so busy fighting fires they won't have time left to focus on customers – many of whom will likely be upset anyway because they were the ones affected by the failed operation, lost data, security breach or illegal violation.

But none of these lapses need to actually cause harm in order for a business to sustain a significant competitive loss due to inadequate timekeeping. Companies lose even when they just look uncompetitive. Time synchronization is an obvious priority for any audit of a company's operations – whether conducted by its own accounting firm or a large potential customer wishing to investigate the operational fitness of its commercial partners. Such audits are increasingly popular among supply chain partners, given the degree to which different companies' systems can become interdependent. Companies want to make sure they are not exposing themselves to security issues or system meltdowns caused by others' inattention to timekeeping. It's one thing when a poorly written web site crashes a PC at home. It's another when the same thing happens to a just-in-time parts ordering system at the plant.

The good news in all of this is that timekeeping solutions can be easy to implement, inexpensive, and extremely effective at stopping any and all of these threats. With respect to questions from potential customers or business partners, sometimes the only thing required to prove competency is to say that a particular timekeeping solution is already in place – analogous to having a security company's sticker on the windows of your home. This brings up the topic of the rest of this paper, which is this: what are the requirements for a good timekeeping solution and how do organizations satisfy them?

Providing Good Network Time

There are two features that make good network timekeeping easy to recognize: 1) that the clocks on the various computers are synchronized by a time server, and 2) removing the need to go through the firewall to acquire time from a time source. There are other benefits that can come from using a time server as well, depending on the features of the server that is installed. Such benefits might be missing if the network were simply getting its time (as is often the case) via the Internet from a public NTP time source. Here is a summary of the most important time server features:

Accurate Time Source — Obviously, the most important benefit a time source can provide is accurate time. There are three main contributors to accuracy: the time source, the availability of the time source, and the reliability of the time server to maintain accurate time once it has received the time from its source.

By definition "accurate" time is time that agrees with Coordinated Universal Time (UTC), the internationally recognized time standard. UTC is available from the National Measurement Institutes (NMIs) of various countries, such as NIST in the United States. UTC can be received from NIST via a couple of ways that include:

- A dialup modem to NIST's Automated Computer Time Service (ACTS)
- Internet via a NIST NTP Server (www.time.nist.gov)
- Radio Broadcast (WWVB)

UTC is also provided by the United States Naval Observatory (USNO) via the Global Positioning System (GPS).

UTC is available from a number of sources over the Internet. Internet-based time sources, however, introduce the security issues discussed earlier. There is also the issue of latency – i.e., delays between when the time packets leave the source and when they arrive at your network. Minimizing the latency improves the synchronization accuracy.

Redundant Time Sources — A better quality server is one that can receive time from multiple sources, not just one. A redundant time source means the time server can always switch to a different source should the need arise — such as when a company moves or redeploys network assets in a corporate restructuring. Moreover, mounting an antenna to receive GPS signals is not always practical, and signal interruptions are possible. Receiving time from GPS satellites requires mounting a coffee cupsize antenna with an unobstructed view of at least half the sky. Although GPS works with a 180° view, a 360° view greatly enhances reliability and reduces how long it takes to acquire GPS time.

Five Dangers of Poor Network Timekeeping

Reliable Time Synchronization — Of course, once UTC has been acquired, the time server becomes the time source for the network. The way the time server works is as follows: Each of the computers in the network makes requests to the time server for an accurate time stamp. By comparing their local clocks to the time server clock, and accounting for network delays, the local clocks are able to set their clock to match the time server.

The key factor that affects the time server's reliability is the accuracy of its own internal clock. The more accurate the clock, the longer the server can go between UTC resets – and the greater the accuracy between those resets. A rubidium atomic clock (the type used in some of the newer GPS satellites) is the most accurate clock to be found in commercially available network time servers. These can maintain an accuracy of within 1 millionth of a second per day, well within the tolerance of most time-dependent software. (By comparison, the Windows default authentication protocol (MIT Kerberos version 5) requires that network domain controllers operate within a time difference of 5000 milliseconds in order for it to authorize logon attempts between the controllers.)

Secure Time Source — Enhanced security is an obvious byproduct of running behind the firewall.

Ease-of-Use — The terms "plug & play" and "set & forget" should both apply to a network time server. Network configuration simply requires plugging the server into the network over a standard Ethernet cable. Setting the time is a one-time operation: you merely plug in the GPS antenna or dialup modem. The server automatically acquires the GPS signals or performs the required dial-out connection. System administrators should no longer be forced to serve as expensive timekeepers and can be freed to focus on other issues.

Cost Efficiency — Time servers are (or at least should be) among the most cost-effective purchases that increase the reliability, performance, and security of a network. A single time server costing less than $5,000 can service hundreds of thousands of computers on a network. Compared to other "enterprise" investments, this cost is virtually zero on a per–CPU basis. It is also far below the cost of many time-related operational or security mishaps.

Summary

When most executives today talk about time management or operating on Internet time, they may not think about the reality behind those expressions. There is, in fact, a real Internet time. There is also a real need to maintain correct time within the computer network – which is, after all, the digital equivalent of your organization's nervous system. Given the importance of time it should not be surprising that there are actual consequences for letting network time go unmanaged, and that these consequences can be serious. These consequences matter at both the technical and business level. What is perhaps most surprising is that, for such a potentially dangerous issue, there is available a simple, effective, and extremely low cost remedy. As organizations, and processes become even more highly synchronized, the importance of network timekeeping will only grow – and so will the application of time servers.

Symmetricom®

2300 Orchard Parkway
San Jose, California 95131-1017
tel: 408.433.0910 fax: 408.428.6960
www.symmetricom.com

WP/5DPNT/042313