



Interested in learning more about security?

SANS Institute InfoSec Reading Room

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

Systems Survivability

The goal of this paper is to introduce the concept of systems survivability, and explore methods in which this broad concept can be utilized. Although much of this paper is based on personal experience in a healthcare setting, the ideas contained herein can easily be applied to any organization. However, due to the nature of this paper, it is not a one-stop solution to any systems problem; rather the intent of this paper is to (hopefully) make the reader think about how this concept can be applied to systems under thei...

Copyright SANS Institute
Author Retains Full Rights

AD

An advertisement banner for Rational AppScan. On the left, the word "Rational." is in white on a blue background, with the IBM logo below it. To the right, the text reads "TAKE BACK CONTROL OF YOUR APPLICATION SECURITY" in bold, followed by "»»» DOWNLOAD A TRIAL VERSION OF RATIONAL APPSCAN" in a smaller font. On the far right, there is a small image of a man in a white shirt and tie, holding a red object.

Systems Survivability

by

Douglas Lancaster, B.Ed., CAN, A+

**Submitted in part for the
GIAC Security Essentials Certification**

15 January 2002

© SANS Institute 2002. Author retains full rights.

Introduction

GOD, grant me the serenity to accept the things I cannot change,
Courage to change the things I can,
And the wisdom to know the difference.
Reinhold Neibuhr-1926

The goal of this paper is to introduce the concept of systems survivability, and explore methods in which this broad concept can be utilized. Although much of this paper is based on personal experience in a healthcare setting, the ideas contained herein can easily be applied to any organization. However, due to the nature of this paper, it is not a “one-stop solution” to any systems problem; rather the intent of this paper is to (hopefully) make the reader think about how this concept can be applied to systems under their control. As the Serenity Prayer suggests, there is often much that is beyond our control – that which is within our control is the main focus of this paper.

What is a System?

Systems survivability is more of a philosophy than an actual plan of action. The primary purpose is to ensure that a system continues to operate in adverse conditions without any severe disruptions. However, any discussion of systems survivability must start by defining systems. Peter G. Neumann defines a system as “operating systems, dedicated application systems, systems of systems, and networks” (2000:ix). Although these aspects are present in most organizations, the definition does not encompass the complete realm of systems.

In generic terms, a system is nothing more than a series of steps required to complete a specific goal. In business terms, the series of steps contribute toward the overall operation, and hence the success or failure, of the organization. Essentially, the complete business operation can be construed as one large system, with a series of subsystems all required to intermesh in order to keep the business operating. If one of these subsystems breaks down, it can have a profound effect on the success of the business, however the failure of one subsystem should not cause the system as a whole to fail.

Thus, systems survivability goes beyond the operating systems, the network infrastructure, or the computer applications. Although these are all important aspects, and generally the ones over which the readers of this paper will have influence, they are only parts of the whole system.

Systems in Action

In a hospital setting, the complete system revolves around the ultimate goal – patient care. The goal of the primary system, therefore, is to take in unhealthy patients, give them proper medical care, and release them when they are healthy. Within this primary system, there are a number of subsystems such as payroll, pharmacy, and nursing, to name just a few, and all must interact effectively in order to achieve the ultimate goal. If there is a fault in any one of these subsystems, the overall goal can be adversely affected as each subsystem relies to a certain extent on other subsystems.

Many of these subsystems rely on computer programs in order to function. From this, it is easy to elevate the computer to lofty heights – after all, without the computer, would not the other systems be doomed to failure? This, unfortunately, is the attitude of some network administrators. Computers are merely tools, and in the hands of trained professionals, they can be very valuable assets to a business. Likewise, in the hands of untrained individuals, they can also be dangerous weapons. By seeing the computers as such, it is often easier to gain perspective on their importance in the operation of the business as a whole.

Holistically, any subsystem should be viewed as a spoke in a bicycle wheel. Each spoke must be kept under a specific amount of tension in order for the wheel to spin true – too much or too little, and the wheel will wobble. Likewise, a subsystem must be able to fail without affecting the operation of the complete system, just as breaking a single spoke will not cause the bicycle wheel to stop rolling. When viewed in this light, the subsystem made up of what is often referred to as Information Technology is an integral but not an essential part of systems survivability.

Potential Systems Problems and Solutions

Legacy Systems

Many systems, including computer networks, are rarely built as single entities, but often built upon an existing framework, or through the gradual improvement and replacement of parts over time. Few if any network professionals would state that their network could not be improved if they could completely gut it and start from scratch (and with an unlimited budget).

A computer network is constantly growing and evolving. Unfortunately, the growth of the network is rarely planned in advance, and when new computers are added to a network, the old ones are often given to another user. To make matters worse, the number of additional users are often attached using temporary measures, such as small network hubs, and these temporary measures often become permanent. As a result, the technology upon which the network is based is suddenly stretched beyond its capabilities, and the computers that are in service are often limited in their capabilities. An additional problem is hardware that is no longer supported by the manufacturer, and that can not easily be repaired or replaced. In the healthcare system, many specialized instruments send information to the computer via a card that requires an ISA slot. These cards are no longer manufactured, and ISA slots are rarely seen in new computers. The result of this is that a number of older workstations must be kept working for this specific purpose, long after they should have been retired.

This problem, to a certain extent, is present in most organizations, and it is a difficult problem to overcome. While it is often difficult to predict network growth, it is necessary to analyze the growth rates of previous years, and attempt to predict future growth. While it is essential to prepare for future growth through the purchase of extra equipment such as network switches, overestimating the growth may be just as expensive as underestimating as there will be a surplus of unused, and rapidly devaluating, equipment. A more sensible approach would be to ensure that the budget grows as rapidly as does the number of users.

Another problem is the required support for legacy software. Computer programs, like the computers that the programs run on, also have an effective life span. However, due to the large initial cost involved in purchasing the necessary licenses, companies will often utilize software even after it should have been retired. An example of this is the large number of DOS-based programs that are still being used by many organizations. Essentially, the attitude is that “if it isn’t broke, don’t fix it”, however it was this attitude that resulted in the Y2K crisis, as legacy programs were being utilized long after they should have been replaced. Unfortunately, the only solution for legacy software is to replace it with a more current version, or a suitable substitute.

Disaster and Contingency Planning

Since September 11, one of the most overused buzzwords in the media is the term, “disaster planning”. Despite this, many organizations only give a cursory glance to disaster planning. Disaster planning involves much more than ensuring that backups are done on a regular basis, but unfortunately this is often the level of thought that is put into this area. There are many very good books and articles on creating plans and policies, so this paper will not go into great detail in this area but will offer forth some suggestions that are often missed.

In order to create an effective disaster plan, all departments must work together. All departments will be affected by a disaster, so interdepartmental cooperation is essential. The plan should clearly spell out the chain of command, and the duties of subordinates, along with a well organized flowchart indicating the appropriate responses to various scenarios. Those involved in the top levels of the chain of command should be familiar with the disaster plan, and should conduct mock disasters. As with the logic behind a fire drill, the time not to read the disaster plan is during the disaster. Mock disasters not only make people more familiar with the disaster plan, but they are also very useful in evaluating the plan. Often, inadequacies in the plan only become evident during the drills.

Other areas that are often not covered, or covered inadequately, in disaster plans include contingency planning, again with other departments. What happens when the power goes out? Although most organizations have Uninterruptable Power Supplies attached to their servers, and occasionally some of the desktop computers, most organizations are not prepared for an extended power outage. The disaster plan should have a level of redundancy in place for such eventualities. Examples of this would include hard copies of important documents and contact lists, stored on and off site, and paper forms that can be completed by hand and later entered into a database.

The goal of any disaster plan is to allow the organization to continue to function as normally as possible under the circumstances. Although it would be impossible to predict or prepare for all disasters, this would be the ultimate goal. With a proper disaster plan, it is possible for an organization to survive many problems that would otherwise cause the system to collapse.

Components Failure

Although there has been great advances in the field of computers in the past two decades, some things have not changed. Most computers still rely on a few pieces of hardware with moving

parts, such as hard drives and cooling fans. When either of these parts stops working, the computer that relies on it will also soon stop working. In an attempt to allay the public's fears, most hard drive manufacturers include in their advertising a statistic called Mean Time Between Failure, or MTBF.

Although the MTBF does provide some information, one can not rely on the MTBF to predict systems survivability. The first problem with MTBF is often referred to as the Complex Systems theory. Essentially, this theory states that the more moving parts in a system, the greater the opportunity for failure, or in other words, an airplane with two engines is twice as likely to have an engine failure than an airplane with one engine. Although MTBF is often used as a selling point, it is important to understand exactly what the MTBF really means.

MTBFs for drives are typically between 200,000 and 500,000 hours—22 to 57 years. Don't confuse these numbers with design life, which is typically 5 years. The MTBF doesn't tell you how long an average drive will last. It tells you how often, on average, a drive will fail if you replace each drive at the end of its design life, even though it's still working. In other words, if a given model drive has a 200,000-hour MTBF, and you have 200,000 drives, one drive, on average, will fail every hour. And if you have 200 drives, one will fail, on average, every 1,000 hours. So the MTBF doesn't tell you anything about how long the drive is designed to last. Rather, it tells you how likely it is to fail during its design lifetime.

(<http://www.zdnet.com.au/shopping/buyerguides/story/0,2000023936,20219778-8,00.htm>)

A second problem is that MTBF is based on exponential distribution. The period of hours indicated by the MTBF is the mean, which in mathematical terms, states that one half of the failures should occur on either side of the MTBF. In essence, the MTBF predicts that the actual number of failures over time, if plotted on a graph, should result in a bell curve. Most people are familiar with bell curves being used in grading classes – the principal is exactly the same. The MTBF predicts that very few components will fail early, and very few will last considerably longer than the mean, with most of the predicted failures occurring close to the mean, and evenly distributed on either side of the mean.

In actuality, failures can occur randomly, and the mean can be affected by premature failures. Rather than relying on MTBF as an indicator, one can again take a lesson from aviation, and utilize the Maintenance Free Operating Period instead of the Mean Time Between Failures. As the name suggests, the MFOP is essentially the length of time that a device will function before it will require periodic maintenance. The goal of MFOP is not to predict failures, but to replace components prior to their failure. Although there have been attempts to mathematically predict the MFOP, the more useful method would be to keep a database of actual failure rates, and share this information with others. Once enough data has been collected, it will then be possible to predict the length of time that a system will operate before it requires maintenance.

To maintain a specified MFOP probability it may be required to replace the components prematurely (or insert inspection). Discarding components before

they fail will inevitably cost money (however, it reduces the cost due to unscheduled maintenance). Besides the increase in the number of spares of these particular components, there is likely to be an increase in the number of LRI (line replaceable item) removals. In general, each component will have different failure distribution parameters and may have different ages.

(Maintenance free operating period - an alternative measure to MTBF and failure rate for specifying reliability ? U. Dinesh Kumar and J. Knezevic, J. Crocker <http://www.ex.ac.uk/mirce/mfop1.htm>)

In mission critical areas, such as aircraft engines, from whence this idea originates, replacing components prematurely would be somewhat expensive, as they would not be usable elsewhere. In the typical business environment, this would have less of an impact, as a component removed from a server could be used in a less critical environment, such as a desktop computer.

In an ideal situation, however, a preventative maintenance schedule would be applied, in which any component that was nearing its lifespan would be replaced. This rarely happens in real life, however, which makes redundancy of components essential. Although there is greater likelihood of one cooling fan failing in a server with three power supplies, the probability of the server surviving a failing cooling fan increases as well. Redundancy is a science unto itself – multiple levels of redundancy could almost eliminate failure, however it would also be cost prohibitive. Thus, it becomes necessary to balance the probability of failure and the cost of redundancy, only instituting levels of redundancy where necessary.

Patches and Service Packs

This paper opened with the Serenity Prayer. This could be, in view of some of the service packs released by certain software companies, be changed to the IT Prayer. To paraphrase Shakespeare, to patch or not to patch, that is the question. Almost any article on computer security will mention the need to keep up to date with the patches and service packs. It is widely accepted that many of the computer viruses and worms would not be as widespread or do as much damage if those in control of computer systems were more vigilant in applying patches.

The problem is that cure is often worse than the disease. There have been many cases in which a recommended patch causes more problems than it cures. An example of this is Service Pack 2 for Microsoft Office, which cripples Microsoft Outlook so various files can not be attached to e-mail messages (<http://www.woodyswatch.com/wowmm/archtemplate.asp?v2-n07>). Although this does rectify some problems, it also puts unacceptable limits on what legitimate users can attach to their e-mail messages. Of course, the preferred solution to dealing with e-mail attachments is to educate the users, and use a good virus scanner and filter at the gateway, the e-mail server and the desktop.

Because of problems such as these, many network administrators will take a “wait and see” attitude toward applying a patch. Only after the patch or service pack has been declared safe by their peers will they load it on their systems. While this does eliminate some potential for harm, the network administrators must be vigilant to ensure that they do not neglect to apply the patch. A preferable solution to this would be to apply the patch to a test system that is not connected to

the network, and stress that system to see if the patch does what it is supposed to do.

Because of some of these problems, applying patches and service packs is often viewed as a balancing act. As shown by the number of computers infected by the Code Red Worm, a failure to apply patches promptly can have disastrous consequences. The Code Red Worm would have been rendered harmless if the systems that it was designed to attack, Microsoft IIS v.4.0 and 5.0, had been properly patched (<http://www.nipc.gov/warnings/alerts/2001/01-016.htm>).

An Example of Systems Failure

Peter Neumann states that systems failures are often the result of either problems with security or reliability (2000:29). The following example of a catastrophic systems failure is due to reliability problems, and a dependence on technology.

In January, 1998, over 100 millimeters (or approximately 4 inches) of freezing rain fell on much of the Canadian provinces of Ontario and Quebec. As a result of this ice storm, a large number of utility towers, poles and lines collapsed. In a country that prides itself on its infrastructure, 1.6 million people were without electricity, some for as long as a month, in the middle of a cold Canadian winter.

Nearly one quarter of the country's dairy cows - 274,000 - were in the affected areas and many could not be milked because farmers depend on mechanized milking.

Cows that are not milked regularly become vulnerable to mastitis, an infection of the udder. Dairy cows that survived the power outages may never regain their pre-storm productivity.

Milk processing plants were shut down and more than 10 million litres of milk had to be dumped. However, 1.5 million litres were processed in American facilities and returned to Canada.

(http://www.canoe.ca/CNEWSIceStorm/icestorm_dec15_cp.html)

Many people reading this may be asking why, in a paper dealing primarily with Information Technology, would the author be talking about dairy cows. The reason is simple. This is a graphic example of complete systems failure. Due to the increasing size of dairy farms over the past few decades, it is no longer feasible to milk cows by hand. With increases in milking technology, it was possible to milk more cows in a less time with fewer people. However, this also lead to a total dependence on this technology. The result was a single point of failure – no electricity equals no milking. This analogy can easily be applied to many industries that rely on computers and computer-based technologies.

As an interesting sidebar to this story, a Quebec government official stated that only fifty municipalities of the 1,400 in the province of Quebec had workable disaster plans during the crisis. (http://www.canoe.ca/CNEWSIceStorm/feb3_quebec.html)

Learning from the Mistakes of Others

Being human, we are all prone to err. Keeping a journal of our mistakes and the solutions is one of the best methods of preventing the same mistake from happening again, not just in the network environment, but in life as well. Although we all learn from our mistakes, it is usually preferable to allow others to make the mistakes and then learn from them.

In an article published in the TechRepublic website, Matthew Villano outlines the events of September 11, as seen through the eyes of the CIO of the Greater New York Chapter of the Red Cross. Like many of us, Leslie Hunt would often spend time thinking of ways to improve her networks and was doing so the morning of September 11.

The article describes her actions over the course of that day. As she added more computers for the volunteers, network stability became a problem and the servers would often crash. The organization's website went down several times, and then the e-mail server got hit by a virus. She describes her actions as taking it one step at a time and focusing on the positives instead of the negatives. Over the next week, Hunt and her staff began evaluating what went wrong, and also what went right. From this was laid the foundation for a two-year plan to improve the network. (<http://www.techrepublic.com/article.jhtml?id=r00520011211gcn02.htm>)

What can be learned from Leslie Hunt's experience? First, although there was a disaster plan in place, it was not adequate for the magnitude of the events that occurred. Although her office was indirectly affected, due to the nature of the Red Cross, a worse-case scenario should have covered large-scale disasters. Second, during a disaster is not the time to be applying patches or trying to eliminate a virus. This is an example of patches not being applied in an appropriate or timely fashion. Although a plan to improve the network has now been developed and is in the implementation stage, this is an example of a reactive response to a problem. The article also mentions a common problem with many IT departments; lack of funding. Often it does take something such as what befell Leslie Hunt and the GNY Chapter of the Red Cross for the funding to be made available, however many of the problems, such as the virus attack, could have been prevented even without additional funds.

Conclusion

As the Serenity Prayer states, much of what the intended audience of this paper does from day to day involves change – applying service packs and patches, closing holes as soon as they appear or are discovered, and keeping that which needs security secure. In doing so, often we lose sight of the reason that we are doing what we do. The computer network is only a part of a much larger system, but it is the part that we can control, and hopefully improve. It is hoped that through reading this paper, the reader may take a proactive stance rather than a reactive response to systems survivability. The quickest reaction will only help limit the damage once it is done; a proactive stance will hopefully prevent the damage from occurring.

Bibliography

Bueckert, Dennis

1998 Ice Storm damage tallied. December 15, 1998, Associated Press.
(http://www.canoe.ca/CNEWSIceStorm/icestorm_dec15_cp.html)

Canadian Press

1998 *Quebec plans for the next disaster*. February 3, 1998, Canadian Press.
(http://www.canoe.ca/CNEWSIceStorm/feb3_quebec.html)

Kumar, U. Dinesh, J. Knezevic, and J. Crocker

Maintenance free operating period - an alternative measure to MTBF and failure rate for specifying reliability? University of Exeter, Exeter, EX4 4QF
(<http://www.ex.ac.uk/mirce/mfop1.htm>)

Leonhard, Woody

Woody's Office for Mere Morals Archives.
(<http://www.woodyswatch.com/wowmm/archtemplate.asp?v2-n07>)

National Infrastructure Protection Center

2001 *Alert 01-016 "Code Red Worm"*. July 29, 2001
(<http://www.nipc.gov/warnings/alerts/2001/01-016.htm>)

Neumann, Peter G.

2000 *Practical Architectures for Survivable Systems and Networks*. SRI International, Menlo Park, California

Villano, Matthew

2001 *9/11: A lesson in crisis control*. TechRepublic, CIO.
(<http://www.techrepublic.com/article.jhtml?id=r00520011211gcn02.htm>)

Ziff-Davis Net

MTBF versus Design Life. Australia
(<http://www.zdnet.com.au/shopping/buyerguides/story/0,2000023936,20219778-8,00.htm>)



Upcoming SANS Training

[Click Here for a full list of all Upcoming SANS Events by Location](#)

Hong Kong Advanced Forensics Seminar	Hong Kong, Hong Kong	Nov 09, 2009 - Nov 14, 2009	Live Event
SANS Sydney 2009	Sydney, Australia	Nov 09, 2009 - Nov 14, 2009	Live Event
SANS Vancouver 2009	Vancouver,	Nov 14, 2009 - Nov 19, 2009	Live Event
SecurityByte 2009	New Delhi, India	Nov 17, 2009 - Nov 20, 2009	Live Event
SANS Geneva CISSP at HEG 2009 Autumn	Geneva, Switzerland	Nov 23, 2009 - Nov 28, 2009	Live Event
SANS London 2009	London, United Kingdom	Nov 28, 2009 - Dec 06, 2009	Live Event
SANS WhatWorks in Incident Detection Summit 2009	Washington, DC	Dec 09, 2009 - Dec 10, 2009	Live Event
SANS CDI East 2009	Washington, DC	Dec 11, 2009 - Dec 18, 2009	Live Event
SANS WhatWorks in Data Leakage Prevention and Encryption Summit 2010	New Orleans, LA	Jan 07, 2010 - Jan 12, 2010	Live Event
SANS Security East 2010	New Orleans, LA	Jan 10, 2010 - Jan 18, 2010	Live Event
SANS AppSec 2010 and WhatWorks in AppSec Summit	San Francisco, CA	Jan 29, 2010 - Feb 05, 2010	Live Event
SANS San Francisco 2009	OnlineCA	Nov 09, 2009 - Nov 14, 2009	Live Event
SANS OnDemand	Books & MP3s Only	Anytime	Self Paced