



# Global Information Assurance Certification Paper

Copyright SANS Institute  
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

## Interested in learning more?

Check out the list of upcoming events offering  
"Security Essentials Bootcamp Style (Security 401)"  
at <http://www.giac.org/registration/gsec>

# Logging – From Your Fingertips to the Desktop and beyond.....

Steven Leikeim

Feb 12, 2004

GIAC Securities Essentials Certification (GSEC)

Practical Assignment version 1.4b, option 1

## Abstract

Logging is a process by which we collect and analyze information about our activities. This information can be used in various ways to give us feedback on activities that occur on our systems. Methods of logging in a computing environment include, starting at our fingertips, with a personal log book and extending from there with electronic, console and graphical logs. Concentrating on log books the format, maintenance and use of the logging alternatives are discussed along with their strengths and limitations.

Our logging processes need to be complete and consistent to ensure the validity of the logs and any data or summaries that we generate from them. Summarization of logs is a valuable activity for electronic logs as the quantity of data can be overwhelming otherwise.

log *n*, ....

3. **a:** the record of the rate of a ship's speed or of her daily progress; also: the full nautical record of a ship's voyage **b:** the full record of a flight by an aircraft
4. any of various records of performance <a computer log><sup>1</sup>

logbook *n*.

1. The official record book of a ship or an aircraft.
2. A record book with periodic entries.<sup>2</sup>

## Background

### *What do we mean by logging?*

Historically, log books have been used to record travel by ships and aircraft. People have also kept journals and diaries for personal and business matters. As commonly used in the computer industry, logging is the process of recording selected items which are expected to be of potential interest. Typically, computers log some "normal" activities (e.g. logins and printer activity) and try to log exceptions (e.g. login failure, disk full, hardware failures, etc). We may ourselves have a personal log (typically in the form of a daytimer or planner) which might contain meeting information, to-do lists and records of work activities. We need to keep in mind that logs are NOT a complete record of events that have occurred, but only a snapshot of particular events or times. When evaluating logs in the context of security, we need to be cognizant of and work within the limitations of the logs that we have.

### *Why do we log things? (Why is logging important)*

There are a large variety of methods we can use to log information. Some of the ones we will commonly see are log books, electronic logs, printing consoles and graphical logs. Often, we log events because they are deemed to be "interesting" in some fashion. For system logs, it's probably due to someone in the software development process deciding that certain events should be recorded or that we have specifically requested the logging of certain events. In personal logs it could be that something exceptional (or just different) happened, an expected event occurred, or a particular accomplishment or achievement was completed. Recording "normal" activities of the system can be an important step in defining a baseline level against which you can compare activities to determine if they are "interesting" or not. We need a baseline to be able to differentiate between periodic events (e.g. network traffic spikes due to backups) from truly "interesting" events, such as intrusion attempts. These logs may also record

---

<sup>1</sup> Webster's Ninth New Collegiate Dictionary, p. 702

<sup>2</sup> Dictionary.com

hardware and software failure events that are important. There may be regulatory or procedural requirements for logging (e.g. legislation, corporate policy, certification requirements or convention). In all cases, the best practice to follow is to ensure that logs are consistent and that all related items are logged in a similar fashion (e.g. logging of logins should be similar on all platforms in an organization.)

### ***How do we log?***

In most instances, logging is built into our computing systems and devices and we simply have to properly collect and store the logs. Activities not directly related to computing may also have logging requirements of their own. (e.g. regulatory or certification processes.) As security practitioners, we should also keep a personal log of security activities that have an impact on us. In any case, for any logging that occurs, we need a clear policy on what we are logging, and how. We also need to be consistent in our logging. In particular, if we have a set procedure for maintaining a particular log, that procedure must be properly followed at all times. Otherwise, we run the risk of not being able to make any statements regarding our logs at all!! Included in our logging should be procedures that can be used to periodically verify, and even certify, that the data logged is correct and valid. Additionally, we may have the capability to log data that is sent to a console or controlling terminal of a system or device. Events may be recorded here that are not recorded elsewhere. In a posting to the unisog mailing list ([unisog-help@sans.org](mailto:unisog-help@sans.org)) Valdis Kletnieks makes the following statement

4) There's an interesting legal facet about logs when trying to use them as evidence - logs that you have been keeping for a long time, and making business/planning decisions based on them, will be given greater weight than logs you suddenly started keeping in the middle of an incident. (I wish I had a citation for this one).<sup>3</sup>

While he does not have a citation for this, nor was I able to find one, this is a reasonable frame of mind to be in when setting up your logging policies.

Sometimes it is necessary to change the presentation of a log for a particular purpose. For instance, management generally does not want to see all of the details of operations, but instead wants a summary. Graphing data from a log file can show trends and make exceptional occurrences obvious. Of course, when we present information detailing a specific event or occurrence that we are interested in, we need to remove irrelevant details from our presentation. It's important to note that anytime we are working with log files, that we work only with copies of the files and leave the original logs untouched. It may be important to later show how our summary in fact is correct based on the original log files.

---

<sup>3</sup> Kletnieks

## ***What should we be logging?***

We should log anything of interest of course!! What this means depends strongly on how we define “interesting”. Certainly we must log items that we are required to, whether by law or convention. Beyond that we need to be logging any exceptional events as they could be indicative of problems or of attempts to circumvent local policies. In order to accurately place events in time, or to place them in a proper context, we should log events that occur with reasonable frequency so as to provide a proper context to evaluate against. On many systems it is possible to turn on very detailed logging. If you enable auditing, particularly filesystem auditing, these logs can be VERY large. These logs can be useful in the proper circumstances but normally we do not need to log at that level. As an example, the UNIX syslog facility specifies several levels from debug to emerg<sup>4</sup>. If we only log “emerg” options, we will have difficulty placing these events in the normal course of the system activities. Logging at the “debug” level is likely to produce much more log detail than we are likely to be interested in. Therefore, debug level is likely only useful when debugging specific programs in a limited timeframe

## ***Where should we be logging?***

We should keep logs at every level from a personal log to our desktop machine and on all machines we are responsible for. Where possible, there may be advantages to centralizing logging. In his paper “Logging and Reporting: A view from the top”<sup>5</sup>, Rick Hislop covers the issues of log collection and reporting with a centralized server. In “Importance of Understanding Logs from an Information Security Standpoint”<sup>6</sup>, Stewart Allen discusses the issues of some of the problems we face with logs and again, options for central collection and storage of logs. Please refer to these papers for a more detailed discussion of these issues. Centralized logging can benefit from normal log items since those items will be properly logged in sequence, with all applicable timestamps, and will help place other entries in their proper chronological relationship. As system clocks can drift<sup>7</sup>, the central log will keep the proper order regardless of the amount an individual clock drifts. This can be useful in case we want to compare logs from two, or more, machines whose clocks have drifted apart.

---

<sup>4</sup> SunOS Reference manual - syslog.conf(4)

<sup>5</sup> Hislop

<sup>6</sup> Allen

<sup>7</sup> NTP.org

# Personal log

## *Format*

As security practitioners, it is essential to keep a personal log of our activities. Written log books are used in a variety of industries such as engineering<sup>8</sup>, security<sup>9</sup> and inventors<sup>10</sup>. Many of the suggestions for log books in these industries are applicable to us as security practitioners. Our personal log book should contain ALL of our security related activities. Queries, requests, investigations, even maintenance and patching ... everything! The exact format of your personal log is a matter of preference, taking into account, of course, any relevant policies or conventions. For our purposes, a good security log book should have most, if not all, of the following attributes:

- It should be a good quality bound book, preferably hard cover.
- The pages should all be pre-numbered.
- All entries should be legible and made in ink.
- All entries must be dated.
- Deletions and/or corrections should only be made by striking out the deleted item with a single line.
- Use every page in the book and use every part of every page.
- A security policy regarding the personal log book should be maintained and followed at all times.
- Small items can be pasted or stapled into the log book.
- Normal or periodic items should be included in the log book.
- Keep the book handy at all times.

Your personal log book should be a good quality bound book. A bound book makes it more difficult to remove (or replace) a page in the book undetected. Generally, a good quality book will be more resistant to falling apart while you are using it. The advantage of a hard cover is that it makes the book more durable. You also find that there are less problems with curling of pages or ripping.

Pre-numbered pages give you a reference point in your book to aid in referring to entries in your log book. As well, this helps to make the log book more resistant to tampering, as it is easy to see if pages are missing. If you use a book that is not pre-numbered, the best thing to do is to number all of the pages yourself before you start to use the book. While this may not look as nice as pre-numbered pages, if you do this consistently, it has all of the advantages of the pre-numbered book.

---

<sup>8</sup> Engineering Design Course

<sup>9</sup> Greater Victoria Security

<sup>10</sup> Bellis

All entries in the log book should be legible and made in permanent ink. Since your log book is a permanent record of your activities, you need to ensure that the record IS permanent. Entries need to be very legible so that others who may view your log book can understand what is written there. This is particularly true if access to your log book is required while you are unavailable or if your log book needs to be used by a 3<sup>rd</sup> party. (e.g. certifications, legal proceedings.) Using pencils or erasable ink should be discouraged as these can be modified or erased with relative ease. You also need to be careful with felt-tip style and fountain pens as these may bleed through the paper rendering both sides of the page illegible. Within these guidelines, the choice of pen style and ink colour are personal preferences. The best choice is one you are comfortable with and that you will use. If you have multiple people using a log book, it could be appropriate to have each person use a different ink colour for their entries. If the work you are doing is sensitive it may even be necessary to have an entry added to your log book authorizing or verifying work performed. Your local security policies may also have a requirement for authorization or verification entries.

All entries in the log book should be dated. In many cases it may also be appropriate to log the time of an entry. This shows when you were working on an issue and can be used to determine an activity timeline if necessary. Additionally, if you go to a prior entry in your log book and add a note, the date will clearly show when an issue was started and completed. It would be best to limit additions to prior entries to either an indication of when the issue was completed or a pointer to the page(s) where further work was done on an issue. In either case, your note should indicate the page number where you completed or did further work on the issue.

Deletions and/or corrections should only be made by striking out the deleted item with a single line. This is important to show that you are not attempting to hide any entries made in your log book. White-out should not be used, nor should anything be pasted over top of written entries. It is possible that you may need to use your log book in legal proceedings. Ensuring that everything recorded in your log book is visible, even items that should not be there, should help show that your log book is a complete and valid record of your activities.

Use every page in the book and use every part of every page. It's important that items recorded in your log book are consistent and not subject to later changes. If you fill pages as you use your log book, you can say with some conviction that your activities did take place in the order they appear in your log book. This is not to say that everything you write needs to be jammed together. White space is an effective way to separate material and you do need some white space to keep entries readable. The important consideration is to be consistent in how you make entries in your log book. Certainly, there will still be small places where you can make small notes, primarily as references to activities recorded on other pages. You may want to use a book that has lined pages on one side and blank or graph paper on the other for your log book. In this case, you may want to set

your policy to record your log proper on the lined paper and use the other side for diagrams or as scratch working space. It's still part of the log and your policy regarding this log should cover this space. If it is necessary to leave a large amount of white space within the log (say to always start a new incident on a fresh page) then you should Z or X out the space that you are leaving to show that you did intend to leave that space blank.

You need to have a personal policy regarding maintenance of your log book. This policy needs to be followed at all times. The main purpose for this policy is to lay down what kinds of activities you will record in your log book and how you will record them. What you are looking for here is to keep your log book as consistent, comprehensive and relevant as possible.

Small items can be pasted or stapled into the log book. Sometimes, you may receive small items which you would like to keep in your log book. Business cards of people you work with on a project is one example. Printed graphs, small excerpts of other material or other output is another. Certainly, you will want to record in your log book why this item is included. This is only really suitable for small items. If you try to include large amounts of extra material in your log book, it will quickly fill up and become hard to use. If you want to make a note of large items in your log book, you will also need to have a place to store the item so that you can refer to it when you need to. Of course, all of this information needs to be recorded in your log book.

Normal or periodic items should be included in the log book. It is important to record normal activities in your log book, some of which may be periodic in nature. These entries will help place the, hopefully, not so frequent entries of security issues in some context, and will show any pro-active, rather than strictly reactive, measures that you are taking. For example, patches to operating systems, periodic preventative maintenance, changes to security measures, completion of certain important activities like full system dumps are some of these activities.

Keep your log book handy at all times. It should be easy to access so that you will record important issues and activities as they occur. If your log book is difficult to access, you are likely to put off recording your activities until a "convenient" time and will likely miss some details. It's important to record activities, particularly in a security context, as they happen. Your log book should have enough details in it so that you can reconstruct what you did if necessary, or explain to someone else how you got the results you did. It should even be possible for someone else, with a basic understanding of the issue, to be able to understand what you did by reading your log book.



## **Benefits**

Written logs have a number of benefits, specifically:

- Difficult to modify or erase.
- Easy to share with others locally.
- Inaccessible to remote parties.
- Very portable.
- Always available and easy to add to.

As a rule, if you follow the procedures listed above, it will not be easy for entries in your log book to be erased or modified. Pages removed from the log book will be visible due to missing pages numbers and, usually, remnants of the removed pages remaining in the log book. Entries made in ink are difficult to modify without damaging the pages or resorting to other, obvious methods of hiding entries.

Log books are an easy way to share information with local people. No special readers, converters nor programs are required for this. All that's required is that both parties understand the same language and that the entries are neat enough to not cause any problems due to legibility.

Log books are generally inaccessible to remote parties. This means that in order for others to see what is in your log book, they must have physical access to it. This also means it is impossible for remote parties to modify your log book in any way.

Log books are very portable. They are easy to carry with you when dealing with security issues. In fact, it would be reasonable to always have your log book with you as you deal with security issues. This also gives added security to your log book as no-one can access to your log book without your knowledge if it is never out of your possession.

Log books are always available for use. Power failures, viruses, worms, hackers all have no impact on your log book. As long as you have empty pages and access to a working writing instrument, you can record in your log. While we want to always be making log entries in ink, if all you have access to is a pencil, that will have to do. In addition to making our own entries, if we have a need for someone else to make an entry in our log book, this is easily accomplished and we can easily identify these entries. In some cases, it may be suitable to have someone else add a verification to our log book that certain activities did occur and/or that proper procedures were followed.

## **Limitations**

Some of the limitations of a log book, in no particular order are:

- Difficult to backup.
- Subject to environmental damage.

- Difficult to search.
- Difficult to share.
- Fairly easy to lose.
- Can expose private information.

It is very difficult to make a backup copy of a log book. Of course, you could make a photocopy of the individual pages, but if you lost the original, you have lost the benefits of using a hard bound book. Mainly you would lose the ability to easily detect if pages have been tampered with.

Books are subject to environmental damage. Certainly fire and water will destroy, or render unusable, a log book. Even spilling your coffee on your log book will cause potentially serious damage. If pages are removed from your log book, unless you have a backup copy, you will not have access to the information on those pages.

Written log books are difficult to search through. You can, of course, create an index for entries you recorded in the log, but the only way to do a general search not related to any indexes you create is to go through the log book page by page.

Log books can be difficult to share. If you are working with someone on a project or issue, you may want to record everything in once place. This could be a problem with a log book if both of you want to record in it at the same time, or if you are working a long distance apart. An additional complexity is that it could be more difficult to determine who made particular entries or additions. This is more true of items added to the log (e.g. business cards) as you can usually tell people's entries apart by the style of their handwriting or printing.

Due to their size and relative portability, log books can be lost, misplaced or even stolen. This is a serious problem as your log book may have the only record of some of your activities or thoughts.

You may, at times, record sensitive information, even passwords, in your log book. This could result in possible security issues if someone else were to gain access to your log book.

## **Summary**

Log books are a valuable tool for security professionals. They are open to recording any information necessary and are available wherever and whenever required. While this method of logging has some limitations, mainly due to the size and construction of the log book itself, most of these limitations can be minimized by simply keeping the log book with us at all times. If we are not able to have it with us all the time, keeping it locked in a safe location will keep it from being accessed or modified by others.

## System (electronic) logs

### *Format*

Most of the computer equipment that we use can log information electronically. Servers certainly do, printers most often do, but so can Ethernet switches (Extreme Networks Summit series<sup>11</sup>), disk arrays (Sun T3<sup>12</sup>), and many other devices. As noted in “Logging and Reporting: A view from the top”<sup>13</sup>, and “Importance of Understanding Logs from an Information Security Standpoint”<sup>14</sup>, in most cases the data logged and the format are dependent on the operating system and/or application being run. Normally, it is possible to collect and store logged entries even if the default is local to the device/system itself. Often the ability to “send” logs to another machine for processing and storage is available.

When we log information, it is best to log all base information available to us. (e.g. login/out records, privilege change, significant events and operational events of important services (email/web)). In addition centralized logging can offer advantages, particularly in terms of relating events that may happen across multiple systems.

We do need to be careful when selecting items to log as certain types of logging (e.g. auditing and debugging log levels) can generate HUGE logs very quickly. However, these may be required in some instances so our security policies should cover this potential requirement.

In “Importance of Understanding Logs from an Information Security Standpoint”<sup>15</sup>, Stewart Allen covers management of logs, not only of collecting, but of archiving and securing your logs. When we are collecting and storing logs, we need to ensure that our logs are as complete as possible and that they are as usable as possible. We should be able to select log information fairly quickly (at an initial level) to avoid having to wade through unnecessary entries. If you cannot archive logs quickly and need some disk space, compression can be an option. Most logs contain lots of repeated data which compresses well. (Compression rates in excess of 90% have been observed on web server logs!)

Mentioned in Stewart’s paper is encrypting log files to safeguard them. While this is certainly true, you then have additional CRITICAL information to maintain with your logs, that being the key to decrypt them and the proper software for the decryption. Should that key, or the software, become unavailable, logs secured by that key or software will not be accessible to anyone, even you! Thus key management then becomes a critical part of your storage management policy

---

<sup>11</sup> Extremeware, p 18-9

<sup>12</sup> Sun T3 Disk Admin Guide – pp 4-9 – 4-16

<sup>13</sup> Hislop

<sup>14</sup> Allen

<sup>15</sup> Allen

relating to your archived logs. If you store the keys too close to the logs, that can make it easy for an adversary to gain access to those logs. If the keys are stored too far away, then the keys are more likely to be discarded as “no longer required”.

The best procedure from the standpoint of stability of the logs is to not change them at all if it can be avoided. (i.e. no compression or encryption.) Your security policy may require one or both of these, but logs that remain unchanged at all times, except for transfer to archive media, are probably more likely to be seen to be valid logs than logs that have additional operations done on them as there is always the possibility of corruption of the log when you make some change to it.

## **Benefits**

Electronic logs have a number of benefits, namely:

- Support for automatic collection.
- Easily searchable and can be quickly and easily summarized.
- Inclusion of normal data and from multiple sources can help establish timelines of events.
- No specific limits on size.
- Even lack of logged information could be used for detection of system problems.
- Can be used for resource allocation/charging.
- Backup copies can be maintained.
- Centralized logging machines can collect entries from virtually anywhere.

Electronic logs have the major benefit that they are generally collected automatically. In fact, most systems have at least some logging turned on in default installations.

Electronic logs are easily searchable with even the most basic of tools. Programs to monitor log files are available (e.g. `swatch`<sup>16</sup>). You can even write your own program to search your logs or convert them to another format (e.g. SQL database) for storage and retrieval if you desire. With the correct tools, logs can be quickly and easily summarized in any format desired. The tool `Logwatch`<sup>17</sup> is useful to generate a summary of many system log files.

Including data from “normal” activities and from multiple machines (centralized logging) can help establish timelines of events. Using data from normal activities, such as login/logout events, can help place other events logged on a machine by having the context of known events to work with. Centrally logging data from multiple machines can help with coordination of events that affect multiple machines and can help to establish when events occurred even if one or more machines have drifted from the correct time. Having centralized logs can also

---

<sup>16</sup> <http://swatch.sourceforge.net>

<sup>17</sup> Bauer

assist in looking at activities across a large number of machines as a single search or analysis operation can apply to all of the machines the log covers with little more effort than the same operation on a single machine.

Electronic logs generally do not have hard limits on the allowed size of the logs. Even if devices, applications and operating systems have a limit to the size of logs, configurable or not, logs can always be copied to another location, or another media for archiving. As long as we take care to avoid modifying the log in transit to other machines or media, we can be assured of the validity of the log.

With electronic logs, even a LACK of logged information can be useful to us. If we have a machine that generally logs a fair amount of data each day to our central logging machine, that machine suddenly not logging anything may be an indication that the machine has been compromised and had logging turned off. There could also be other reasons of course, but if this machine is important to us, we should investigate and find out why the logging has stopped. We do need to be careful about saying that an event did NOT occur based on a lack of logged data. It is certainly easy to say that an event did occur based on a record of that event being logged. It is not so easy to say that an event did not occur based on no record being logged as it is possible for a record to be lost due to events such as network failure, software failure, disk errors or full logs. To be able to say that an event did not occur with some authority, you would need to base that conclusion on other logged entries being consistent with the event you are claiming did not occur.

Logs can be used for resource allocation and charging. Keeping track of limited resources can give us the capability to decide where resources are being over or under utilized and re-allocate or purchase additional resources as appropriate.<sup>18</sup> Logs of usage (e.g. printing, disk space, CPU, license usage, etc) can be used to generate charging information.

You can make backup copies of logs. These backup copies can, of course, reside anywhere on any media you desire. Certainly, this makes an attackers work harder if multiple copies of a log are maintained, particularly if some of the copies are on read-only media. If we are searching through, or summarizing logs, especially if we are developing tools for these jobs, we need to be working with copies of the log files in case something goes wrong.

Centralized logging gives us the capability to collect log entries from virtually anywhere. Certainly, it makes a lot of sense to centrally collect logs from local machines. Collecting logs from machines around the world is not only possible, but necessary if that's the extent of your responsibility. We still have to keep logs on the machine(s) themselves, even if we centrally log the same entries. This will help ensure that we have some copy of our logs if one of our machines is compromised, even if that machine is the central logging machine.

---

<sup>18</sup> Hislop

## **Limitations**

Some of the limitations of electronic logs are:

- Easily and undetectably modifiable.
- What if the original and backup do not match?
- Large log files can be difficult to sort.
- Collection can be affected by system problems or hostile activity.

Given that it is very easy to record and store electronic logs, it's not surprising that it is also very easy to delete or modify these logs as well. If the person modifying the logs is at all careful, any modifications made can be almost totally undetectable or, at least, very difficult to trace.

Having a backup copy of a log helps to alleviate the previous point. However, if one of the copies is modified, how do you tell which one is the correct version? Unless you can be absolutely sure which one was modified you can't be entirely confident of what changes were made. (Just because a machine is "more secure" doesn't mean it's not the one where the log was changed.) Certainly, having multiple copies of logs will help determine what is correct, but if more than one copy of the entries we are questioning has been damaged or modified, all that we might be able to ascertain is that we are under attack or that we are experiencing some failure within our logging infrastructure.

If we let our log files accumulate for some time before we rotate them, or if we are simply logging too many things to one file (e.g. debugging information) we run the risk of having log files that are difficult to sort through, even with tools to help us. If it's too difficult for tools to deal with, it's almost certain that the log would be beyond sorting through manually. While consistent, constant log rotation will certainly help with dealing with logs over a short term, this problem could still arise if we are looking for activity over an extended time which we might accomplish by merging copies of the rotated logs for processing.

Failures related to logging machine(s) can result in significant information loss. These failures could be due to system problems (e.g. disk failure, network outage) or hostile attacks (e.g. Denial of Service (DoS), attempts to flood the log). In "Importance of Understanding Logs from an Information Security Standpoint"<sup>19</sup>, Stewart Allen covers this issue and suggests having multiple centralized logging machines for each class of log styles (specifically UNIX, NT and other devices). He further suggests log replication across these servers to provide greater redundancy.

---

<sup>19</sup> Allen

## **Summary**

Electronic logs are probably the easiest type of logs to keep as most systems keep some electronic logs by default. Our largest concern regarding electronic logs is in collecting them and storing them. Ideally we would store logs centrally, and use them to report on our system activities. We can use these logs to determine normal activity, and to indicate when abnormal activity, such as an intrusion attempt is taking place. If we are going to keep logs for a long period of time, then we need to have policies in place regarding storage and retention of our logs so that the logs are available and usable if and when we need them.

## **Printed logs (consoles)**

### **Format**

Many computer systems (particularly UNIX servers) will often send important information to a “console” device which is generally a terminal of some description. This is also true for other devices which may be in use such as switches, disk arrays, tape libraries, routers and firewalls. In some cases, the information sent to this terminal is NOT recorded elsewhere. Most of these messages fall into one of the following categories:

- Monitor/BIOS messages.
- Kernel and/or panic messages
- Indications of hardware problems or failures, particularly disk problems.

Using a printing terminal, you can keep a permanent record of any messages sent to this console. As well, any system maintenance or recovery activities performed on this machine, which may be performed in a mode where logging is turned off, will be recorded. Alternatives to a printing terminal include terminal servers, some of which can give remote access to the console, and a program like “conserver”<sup>20,21</sup> which uses terminal ports on a computer as a terminal server and includes the ability to log any console output in an electronic format.

### **Benefits**

Logging console data to a printing terminal has the benefit that you instantly have a permanent printed record of that data. Since it will be on paper already, it is possible for someone to annotate the record with any additional information desired and/or who is performing the maintenance or recovery activities. While this method can record data which is not recorded elsewhere, it will also record some data that is recorded in other locations, which can be useful.

### **Limitations**

Printed console logs have many of the same limitations as written log books, namely:

---

<sup>20</sup> Fine

<sup>21</sup> [conserver.com](http://conserver.com)

- Difficult to backup.
- Subject to environmental damage.
- Difficult to search.
- Can expose private information

In addition to these printed console logs have some additional limitations:

- Subject to terminal device errors.
- Fragility of output.

As with written log books, it is difficult to backup a printed console log. However, most of the messages recorded here will probably be recorded elsewhere on the system (e.g. syslog). Additionally, the size and format (generally continuous form paper) make making a copy of the console log difficult.

Environmental damage is an issue with printed console output. Damage by fire and water are of course issues. Since your console will be typically in your computer room, it's possible you might have both at the same time if you had a fire there and the sprinkler system went off.

Unless the information was recorded elsewhere, it is difficult to search through a console log for information. In most cases, you are probably interested in what was recorded at or around a particular time and that can be a fairly easy search as the console log will be in strictly chronological in order.

Private information can be exposed in a console log. This can be either due to typing in your password when the system is not expecting it, thus recording it on the paper, or as a logged copy of a login attempt where the user typed their password instead of their username.

Errors on the console terminal can cause the console output to not be recorded. In extreme cases, this could even affect the operation of the attached server. Problem indications that could arise with a printing terminal are: out of paper or ribbon, paper jam and general device errors. Since printing terminals are mainly mechanical devices, there is lots of potential for failure.

The output generated from a printing terminal can be quite fragile. The quality of the paper may not be as good as that in your log book for instance, which could result in faster deterioration of the log. Additionally, the continuous form output can be difficult to page through, particularly in the middle of a large printout, resulting in torn or separated pages if you are not careful.

## **Summary**

Printed (console) logs can be an effective way to collect information that is not otherwise recorded elsewhere. When using printed logs, we have to be careful that our devices for printing these logs are working properly, or use an alternative



means of collecting this data. While these records are difficult to search through, our likely interest is probably going to be chronological in nature which, given how the data is recorded, will work very well.

## Graphical (Performance) logs

### Format

While not necessarily logs in the purest sense, graphical (often performance) logs can be an effective tool to monitor systems for performance, utilization and availability. Graphs can be “real-time” (e.g. CPU or network utilization) or summaries of past system activities (e.g. daily virus blocking statistics, license usage). Graphs can often show system changes much more clearly than looking through raw logs. Areas where this is particularly evident are in rapid changes in a monitored resource, periodic bursts of activity and gradual but constant long-term change in utilization of a resource (e.g. disk usage).

One tool, which is very useful in collecting and displaying graphs, is MRTG<sup>22</sup>. While originally developed for graphing network traffic, extensions have been developed to summarize virtually any sort of numeric information<sup>23</sup>.




Figure 1 - Network traffic - frodo.bagend.nsd

As we can see from Figure 1, fairly constant network activity was observed until 7:00 at which time a sudden increase in traffic was noted. We are unable to tell from this graph just what that activity is, but it can point us in a direction for investigation.





Figure 2 - Network traffic - weekly - frodo.bagend.nsd

<sup>22</sup> Oetiker

<sup>23</sup> Oetiker, MRTG Companion sites

From Figure 2 we can see a longer term graph network traffic of the same machine. We can see here that, while there are large changes in the network traffic, they appear to be consistent in time of day and relative network impact. From this we can deduce that these spikes are probably due to periodic, probably scheduled, activity relating to frodo.bagend.nsd, such as backups over the network. This is something we would have to discover via further investigation. The anomalous spike on Monday afternoon seems unrelated to the periodic activity and again would require further investigation.



**Figure 3 - Disk usage (%) - frodo.bagend.nsd**

Figure 3 shows us the disk utilization of our machine over a period of a year. We can see clearly that the usage is slowly increasing over time and, if this continues, we will probably need to be concerned about it at some point in time. We can also see that there was a large increase in usage during June (possibly due to a project) but that requirement seems to only have been there for a short duration. To find the underlying cause of the steadily increasing disk usage, we must investigate further. It could be due to log files growing, growing web browser caches, or simply due to a constant but gradual increase in the number of users utilizing this resource.

## **Benefits**

Using graphs to view logs (or summaries from logs) allows us to easily see trends and rapid changes in activity or usage. They can also allow us to select an area for further investigation by looking at a collection of graphs and quickly spotting the graph, or graphs, which are likely to be worthy of further investigation. When looking at problems, you almost always have to look elsewhere for more information about a problem as graphs can only tell you so much.

## **Limitations**

Probably the main limitation of using graphs for logging is that most graphs tend to be transient. That is, we only create them when we are trying to find or solve a problem (e.g. CPU monitoring) or we are using something like MRTG<sup>24</sup> where the data collected for the graphs is periodically consolidated and eventually,

---

<sup>24</sup> Oetiker

discarded. If we want to log and maintain this data for longer periods, then we will need to use different tools for this purpose.

Graphs can tell you that a particular change has, or did, occur, but cannot tell you WHY that occurred. Further investigation is almost always required. In some cases this might point to areas requiring more detailed logging to record the events that we are interested in. If your graph does not cover a long enough time, you might not have enough data to determine that the activity you are seeing is expected periodic activity and not an anomaly to be concerned about.

## **Summary**

We can use graphs to quickly see trends and changes in a resource activity. Graphs can also be used in groups (e.g. traffic summaries for each port of an Ethernet switch) to quickly view and select areas of interest. Sometimes that will be the most active, sometimes the least. At times, the most variable or the one with the most obvious changes will be one we are interested in. Sometimes, you will not sure what you're looking for until you see it.

## **Conclusions**

As we have seen, logs can be generated in a variety of formats. The most important considerations when logging are completeness and consistency. That is always log as much as is practical and always log all of the data you've decided to log. Our logging must be consistent with applicable laws, policies (corporate and personal) and operational needs. If 3rd parties use our logs, we need to ensure that their requirements are also met. As we use different methods to log, we need to be aware of the strengths and limitations of each method and make our choices accordingly.

When summarizing or working with logs, it is important to work on copies of the log so as to not corrupt the original log. We must do everything possible to preserve the original logs without modification.

We can use all of the logging methods available to us to have multiple records of important events. Having multiple records will make it more difficult to someone to erase all traces of an important event.

Ensuring that enough "normal" data is logged will help us identify exceptions and unusual situations by defining a baseline. The "normal" entries will also serve to accurately place anomalies in time and possibly by source.

## Bibliography

Merriam-Webster Inc., Webster's Ninth New Collegiate Dictionary. Markham, Ontario: Thomas Allen & Son Limited, 1986

Dictionary.com. <http://dictionary.com> (12 Feb 2004)

Sun Microsystems, Inc. SunOS Reference Manual, Section 4. File Formats. Mountain View, CA. August, 1994

Government of Alberta, Canada. "Freedom of Information and Protection of Privacy". URL: <http://www3.gov.ab.ca/foip> (12 Feb 2004)

Hislop, Rick. "Logging and Reporting: A view from the top". August 17, 2003. URL: [http://www.giac.org/practical/GSEC/Rick\\_Hislop\\_GSEC.pdf](http://www.giac.org/practical/GSEC/Rick_Hislop_GSEC.pdf) (12 Feb 2004)

Allen, Stewart. "Importance of Understanding Logs from an Information Security Standpoint". 2001. URL: <http://www.sans.org/rr/papers/33/200.pdf> (12 Feb 2004)

NTP.org. "ntp.org: Home of the Network Time Protocol". November 13, 2003. URL: <http://www.ntp.org> (12 Feb 2004)

Brusse-Gendre, T. "Engineering Records". Engineering Design. 2002. URL: <http://www.ucalgary.ca/~design/toolbox/toolbox-records2.htm> (12 Feb 2004)

Greater Victoria Security. "Greater Victoria Security – The Services We Offer". URL: <http://www.gvssecurity.ca/services.html> (12 Feb 2004)

Bellis, Mary. "Keep an Inventors Log Book". URL: [http://inventors.about.com/cs/logbook/ht/Log\\_book.htm](http://inventors.about.com/cs/logbook/ht/Log_book.htm) (12 Feb 2004)

Extreme Networks, Inc. ExtremeWare Software User Guide, Software Version 6.1. Santa Clara, CA: Extreme Networks, Inc. April 2000

Sun Microsystems, Inc. Sun StorEdge T3 Disk Tray Administrator's Guide. Palo Alto, CA: Sun Microsystems, Inc. July 2000

"SWATCH – The Simple WATCHer of Logfiles", URL: <http://swatch.sourceforge.net> (12 Feb 2004)

Bauer, Kirk. "Logwatch". URL: <http://www.logwatch.org>. (12 Feb 2004)

Valdis Kletnieks. “[unisog] IDS vs. Privacy – unisog message #4226”. 2 Feb 2004. Retrieve via mail to [unisog-get.4226@sans.org](mailto:unisog-get.4226@sans.org) (no subject nor body required.)

Oetiker, Tobias and Rand, Dave. “MRTG: The Multi Router Traffic Grapher”. 2004. URL: <http://people.ee.ethz.ch/~oetiker/webtools/mrtg>. (12 Feb 2004)

Oetiker, Tobias. “MRTG Companion Sites”. URL: <http://people.ee.ethz.ch/~oetiker/webtools/mrtg/links.html>. (12 Feb 2004)

Fine, Thomas A, and Romig, Steven M. “A Console Server”. Conference Proceedings – Large Installation Systems Administration IV October 1990 (1990): 97-100

Stansell, Bryan. “Conserver”. URL: <http://www.conserver.com>. (12 Feb 1004)

The Institute of Internal Auditors - UK and Ireland. “NOTES ON COMPLETING A LOGBOOK FOR THE AWARD OF QICA”. URL: <http://www.iaa.org.uk/images/qicalogbook.doc>. (12 Feb 2004)

Penner, Russell. “Securing the Network in a K-12 Public School Environment”. October 25, 2003. URL: [http://www.giac.org/practical/GSEC/Russell\\_Penner\\_GSEC.pdf](http://www.giac.org/practical/GSEC/Russell_Penner_GSEC.pdf). (12 Feb 2004)

Posey, Brien M. “Keeping track of your network”. May 26, 2000. URL: <http://networking.earthweb.com/netsysm/article.php/623431>. (12 Feb 2004)

# Upcoming Training

Click Here to  
**{Get CERTIFIED!}**



SANS Stockholm 2017	Stockholm, Sweden	May 29, 2017 - Jun 03, 2017	Live Event
SANS San Francisco Summer 2017	San Francisco, CA	Jun 05, 2017 - Jun 10, 2017	Live Event
SANS Houston 2017	Houston, TX	Jun 05, 2017 - Jun 10, 2017	Live Event
Security Operations Center Summit & Training	Washington, DC	Jun 05, 2017 - Jun 12, 2017	Live Event
Community SANS Ottawa SEC401	Ottawa, ON	Jun 05, 2017 - Jun 10, 2017	Community SANS
SANS Rocky Mountain 2017	Denver, CO	Jun 12, 2017 - Jun 17, 2017	Live Event
SANS Charlotte 2017	Charlotte, NC	Jun 12, 2017 - Jun 17, 2017	Live Event
SANS Rocky Mountain 2017 - SEC401: Security Essentials Bootcamp Style	Denver, CO	Jun 12, 2017 - Jun 17, 2017	vLive
Community SANS Portland SEC401	Portland, OR	Jun 12, 2017 - Jun 17, 2017	Community SANS
SANS Secure Europe 2017	Amsterdam, Netherlands	Jun 12, 2017 - Jun 20, 2017	Live Event
SANS Minneapolis 2017	Minneapolis, MN	Jun 19, 2017 - Jun 24, 2017	Live Event
SANS Columbia, MD 2017	Columbia, MD	Jun 26, 2017 - Jul 01, 2017	Live Event
SANS Cyber Defence Canberra 2017	Canberra, Australia	Jun 26, 2017 - Jul 08, 2017	Live Event
SANS Paris 2017	Paris, France	Jun 26, 2017 - Jul 01, 2017	Live Event
SANS London July 2017	London, United Kingdom	Jul 03, 2017 - Jul 08, 2017	Live Event
Cyber Defence Japan 2017	Tokyo, Japan	Jul 05, 2017 - Jul 15, 2017	Live Event
SANS Cyber Defence Singapore 2017	Singapore, Singapore	Jul 10, 2017 - Jul 15, 2017	Live Event
Community SANS Minneapolis SEC401	Minneapolis, MN	Jul 10, 2017 - Jul 15, 2017	Community SANS
SANS Los Angeles - Long Beach 2017	Long Beach, CA	Jul 10, 2017 - Jul 15, 2017	Live Event
Community SANS Phoenix SEC401	Phoenix, AZ	Jul 10, 2017 - Jul 15, 2017	Community SANS
SANS Munich Summer 2017	Munich, Germany	Jul 10, 2017 - Jul 15, 2017	Live Event
Mentor Session - SEC401	Macon, GA	Jul 12, 2017 - Aug 23, 2017	Mentor
Mentor Session - SEC401	Ventura, CA	Jul 12, 2017 - Sep 13, 2017	Mentor
Community SANS Atlanta SEC401	Atlanta, GA	Jul 17, 2017 - Jul 22, 2017	Community SANS
Community SANS Colorado Springs SEC401	Colorado Springs, CO	Jul 17, 2017 - Jul 22, 2017	Community SANS
SANSFIRE 2017	Washington, DC	Jul 22, 2017 - Jul 29, 2017	Live Event
SANSFIRE 2017 - SEC401: Security Essentials Bootcamp Style	Washington, DC	Jul 24, 2017 - Jul 29, 2017	vLive
Community SANS Charleston SEC401	Charleston, SC	Jul 24, 2017 - Jul 29, 2017	Community SANS
Community SANS Fort Lauderdale SEC401	Fort Lauderdale, FL	Jul 31, 2017 - Aug 05, 2017	Community SANS
SANS San Antonio 2017	San Antonio, TX	Aug 06, 2017 - Aug 11, 2017	Live Event
SANS Prague 2017	Prague, Czech Republic	Aug 07, 2017 - Aug 12, 2017	Live Event