



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: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Backup Rotations – A Final Defense

Defense in depth is an important strategy in protecting and securing your network infrastructure. However, many organizations are just beginning to create a more secure environment and have a limited amount of defensive lines presently installed. With this being the case, the last line of defense that the organization has, and possibly the most important because in many organizations it is all they have, is an effective backup strategy.

In the GSEC Paper “How to Implement an Effective Backup Solution: A Companies True Story” by Wanda Jackson we were presented with the different types of backups (full, incremental, differential).

However, even if you are doing backups and you have verified your backups, can you recover the information that is needed in the event of an emergency? If an intrusion incident occurs, will your organization be able to revert to an earlier version or recover with your present backup rotation strategy. Will you have already reused the media that has the data that someone wishes to recover? Or will all your backups be compromised and you have no fall back point? What we will look at in this paper are the different ways to rotate your backup media to aid in recovering either your system in the event of an intrusion or data in the event of an accident or other deletion or failure in the system or disaster.

What you choose for a backup rotation strategy will be based upon balancing three specific areas against your business needs and goals. The three areas are Retention Length, Availability and Integrity.

Defining several terms will assist in evaluating what will be an appropriate strategy for protecting your organization’s data. Retention time is defined as: How long do you want to keep your media before you reuse it? Also included in this the implied concept of how far back into the past will you want or will you need to go to recover what you need?

Availability is defined as: How often over a period of time can you recover files? For example, how often over a two week period could a file be recovered? Integrity is defined as: Can you recover the file that is needed? For this paper Backup Media will be defined as the type and amount of media required to complete a backup in a given night. For example, if your system requires two tapes to complete a backup those two tapes are a media set.

Aleen Frisch in Essential System Administration wrote “.....it’s best to have five sets of tapes that you reuse each week; if you can afford it, you might even have 20 sets that you rotate through every four weeks” (Frisch. p.472).

Let’s look at some of the different rotation strategies that we can apply to protecting our data.

There are several rotation strategies that can be used. Each strategy has its own benefits and costs. Below is a list of the strategies and some variations to those strategies that will be discussed.

1. Father/Son
 - a. Basics
 - b. 6 Tape Strategy
 - c. 10 Tape Strategy
2. Grandfather/Father/Son
 - a. Basics
 - b. 10 Tape Strategy
 - c. 19 / 24 Tape Strategy
3. Tower of Hanoi
4. Incremental Rotation

Father/Son

A basic Father/Son rotation consists of four tapes used daily and two tapes used on successive Fridays. Some combination of full/incremental/differential backups will be used on the Monday–Thursday series, and a full backup will be run on Fridays. (to start the process an initial full backup should be made).

| | | | | | | | | | |
|-------|--------|--------|--------|--------------|-------|-------|-------|--------|---------------|
| Mon | Tues | Wed | Thurs | Fri | Mon | Tues | Wed | Thurs | Fri |
| Tape1 | Tape 2 | Tape 3 | Tape 4 | Tape5 | Tape1 | Tape2 | Tape3 | Tape 4 | Tape 6 |

Recovery in this scenario is limited to a maximum of six days. Your daily file recovery is also limited to a maximum of six days. It also is putting extensive wear on the main weekday media. So your retention length is short but your availability over the short term is high.

A modification of the above strategy is to use ten tapes. This increases your daily file recovery length to ten days and increases your maximum recovery length to ten days.

| | | | | |
|--------|--------|--------|--------|----------------|
| Mon | Tues | Wed | Thurs | Fri |
| Tape 1 | Tape 2 | Tape 3 | Tape 4 | Tape5 |
| | | | | |
| Mon | Tues | Wed | Thurs | Fri |
| Tape 6 | Tape 7 | Tape 8 | Tape 9 | Tape 10 |

This will also decrease the wear on the Monday through Thursday tapes.

Grandfather/Father/Son

A second method, the Grandfather/Father/Son, works in much the same way with four rotating daily media. Each successive Friday uses a different backup media, three tapes for three successive Fridays. An additional three media set for three consecutive monthly backups is added to the rotation, this is the Grandfather set.

The rotation strategy in its simplest form uses only 10 tapes. The daily tapes are reused each week. The three Friday tapes are rotated through the month and on the fourth Friday of the month one of the monthly (grandfather) tapes is used. This allows for a maximum possible recovery back to the third month, about 90 days after you have completed the full rotation and are one month into the second rotation. But the short-term daily backup is limited to only six days back. However, by the end of the third month's rotation you have now created an image of the end of five consecutive weeks (one of which is the monthly at the end of month two) and an additional monthly backup from the first month. This increases your ability to fail/fall back to a known good state in the event of an incident. With these three extra monthly tapes you now have six good fall back points in addition to the four days of the week.

| | | | | |
|--------|--------|--------|--------|-----------|
| Mon | Tues | Wed | Thurs | Fri |
| Tape 1 | Tape 2 | Tape 3 | Tape 4 | Weekly 1 |
| | | | | |
| Mon | Tues | Wed | Thurs | Fri |
| Tape 1 | Tape 2 | Tape 3 | Tape 4 | Weekly 2 |
| | | | | |
| Mon | Tues | Wed | Thurs | Fri |
| Tape 1 | Tape 2 | Tape 3 | Tape 4 | Weekly 3 |
| | | | | |
| Mon | Tues | Wed | Thurs | Fri |
| Tape 1 | Tape 2 | Tape 3 | Tape 4 | Monthly 1 |
| | | | | |

A more extended rotation could use nineteen backup media. With nineteen media sets you will have the same short term recovery timeline as with the 10 tape rotation, about six days, however, you will have increased your monthly (grandfather) to twelve and will now have a full year worth of monthly backups.

With the increase of four additional tapes you can create a deeper short-term recovery and reduce wear on your tapes by creates a Mon2-Thurs2 set of tapes. This creates immediate short-term recovery of eleven days but still leaves one backup media for each month as a recovery point.

Tower of Hanoi

A third method is the Tower of Hanoi. It is a more complex method of rotation. The Tower of Hanoi rotation is taken from the mathematical game of the same name.

At this point a short description of the game is probably in order. The Tower of Hanoi is a game that has three posts in a row. The left most post has an N number of disks in increasing size from the bottom to the top and the other two posts are empty. The object is to move the N number of disks from post 1 to post 3, but never putting a larger disk on top of a smaller disk. You are also only allowed to move one disk at a time. Let's look at how you solve the Tower of Hanoi.

“It is well known that the optimal solution of the Towers of Hanoi with N disks requires 2^N-1 moves.” (Art of Prolog)

We will look at just the following three scenarios 3,4, and 5 backup media sets.

The table below is a rotation for N=3 disks labeled A B C and for Posts labeled 1, 2 and 3, where 1 is the left-most post and 3 is the right most post and is also the destination for the disks. A is the smallest disk and C is the largest disk.

For the first set we will explain the full run. (see first table below)

A is moved from Post 1 to Post3, then B is moved to post 2. We then move A to post 2 on top of disk B. We now have C on post one and A on top of B on Post 2. We then move C to post 3, and move A from on top of B on post 2 to post one. This allows us to move disk B to post three and finally A to post 3.

| | | | | | | | | | |
|------|---|---|---|---|---|---|---|--|--|
| Disk | A | B | A | C | A | B | A | | |
| Post | 3 | 2 | 2 | 3 | 1 | 3 | 3 | | |

In the above example three pieces of backup media are used. The longest recovery date will be seven days (C being the oldest media in the run). However, the short-term file recovery is limited to only two days.

If we take this scenario another step further by adding another tape to the rotation.

| | | | | | | | | | |
|------|---|---|---|---|---|---|---|---|---|
| Disk | A | B | A | C | A | B | A | D | A |
| Post | 2 | 3 | 3 | 2 | 1 | 2 | 2 | 3 | 3 |
| | | | | | | | | | |
| Disk | B | A | C | A | B | A | | | |
| Post | 1 | 1 | 3 | 2 | 3 | 3 | | | |

By increasing our backup media by one we now have increased our long-term storage to a maximum of fourteen days. Our short-term storage is still two days, but also includes a fourth

day back. This then yields recoverable days of 1,2,4,8 (or 15 depending where you are in the rotation cycle).

If we add just one additional backup media we end up with the following scenario. Applying the formula $2^n - 1$, where $n=5$ we end up with a thirty one day rotation.

| | | | | | | | | | |
|------|---|---|---|---|---|---|---|---|---|
| Disk | A | B | A | C | A | B | A | D | A |
| Post | 3 | 2 | 2 | 3 | 1 | 3 | 3 | 2 | 2 |
| | | | | | | | | | |
| Disk | B | A | C | A | B | A | E | A | B |
| | 1 | 1 | 2 | 3 | 2 | 2 | 3 | 1 | 3 |
| | | | | | | | | | |
| Disk | A | C | A | B | A | D | A | B | A |
| Post | 3 | 1 | 2 | 1 | 1 | 3 | 3 | 2 | 2 |
| | | | | | | | | | |
| Disk | C | A | B | A | | | | | |
| Post | 3 | 1 | 3 | 3 | | | | | |

This now creates a maximum long-term file recovery of thirty one days and the following short-term recovery: Days 1,2,4,8,16.

The Tower of Hanoi strategy is an effective method of backup for creating the longest possible recovery situation with a limited number of backup media. In the Grandfather/Father/Son with 10 tapes we had a situation where we had effective short-term recovery of eleven days but we were limited to 90 days of long-term recovery. However, with a ten media set scenario, for the Towers of Hanoi we would have a short-term backup of only two days. But we would have ten fall back points with the oldest tape being 1023 days old. Rotation is as follows: 1,2,4,8,16,32,64,128,256,512 and 1023 at the furthest end of the rotation.

The biggest two drawbacks of the Tower of Hanoi rotation is the wear and tear on the more daily backup media, the ABCD media sets, and the reduced short-term file recovery. The benefit is that you can, with a very small number of media sets, create an exceptionally long-term recovery strategy.

In this scenario an automated backup log is absolutely required because of the complexity of the rotation and as additional tapes are added the rotation will become increasingly cumbersome to do manually. We would not want our backup operators going crazy.

Incremental Rotation

The incremental rotation is a sets of media labeled from one to N, where N is the last media set. The initial backup is done on media set one and continues forward through the first week.

| | | | | |
|--------|--------|--------|--------|--------|
| Mon | Tues | Wed | Thurs | Fri |
| Tape 1 | Tape 2 | Tape 3 | Tape 4 | Tape 5 |
| | | | | |
| Mon | Tues | Wed | Thurs | Fri |
| Tape 2 | Tape 3 | Tape 4 | Tape 5 | Tape 6 |

At the beginning of the next week we pull the first used set from the previous week out of the rotation and store. Then add the next higher media set number to the end of the rotation. This yields a short-term recovery time of five or six days depending upon when the failure occurs. With twelve media sets you will have seven weekly backup sets before you start the rotation over again.

The benefit of this rotation is a reduced wear on media. The media sets all rotate in and out of the cycle equally. However, you will have to use more media sets to achieve the same long-term recovery of the GFS strategy. However, it is fairly easy to implement and with only ten tapes you have a deeper recovery than is possible with just the Father/Son strategy (five weekly sets back as opposed to only one in the Father/Son).

Conclusions

It is definitely important to be able to recover files over the short-term. However, it is also necessary to be able to recover some files, or even an OS drive in the event of either intrusion or loss of file(s). I have seen instances and have had instances where having three and eight-month old backups were the difference between a successful recover and failure (and maybe keeping your job!!).

As we have seen there are several different backup rotation strategies. The one that works the best for you may not be what works best for another site. However, with the above information you should be able to find a strategy that, with some tuning will allow you to meet the recovery needs of your organization.

A parting thought from one of the books that I read through during this process.

It may take a week or a month to realize that a file has been deleted. Therefore, you should keep some backup media for a week, some for a month, and some for several months. After all, tape is cheap, and **rm** is forever. Keeping a yearly or a biannual backup forever is a very small investment in the event that it should ever be needed again (Garfinkel. p.108).

Citations / References

Feidler and Hunter, Revised by Ben Smith. “Unix System V Release 4 Administration 2ed.” Hayden Books, 1991.

Frisch, A Elen . “Essential System Administration 2nd Edition.” O’Reilly and Associates, 1995. p 472.

Garfinkel, Simson and Spafford, Gene. ”*Practical Unix Security.*” O’Reilly & Associates, 1991. p.108.

Jackson, Wanda. “How to Implement an Effective Backup Solution: A Companies True Story .”
<http://www.sans.org/infosecFAQ/sysadmin/backup.htm>

Kern, Johnson, Hawkins, Law with William Kennedy. “Managing the New Enterprise.” SunSoft Press A Prentice Hall Title, 1996.

Sterling, Leon / Shapiro, Ehud . “The Art of Prolog.” MIT Press Series in Logic Programming, Massachusetts Institute of Technology, 1986 p. 65.

Seagate Web Site. “Tape Rotation Schemes.”
URL: <http://www.seagate.com/products/tapesales/backup.a2g1.html>

Mateyaschuk, Jennifer. “Backup Plans Become Critical.” Information Week. January 11, 1999.
URL: <http://www.informationweek.com/716/16iubkp.htm>

NIST Web Site
URL: <http://hissa.nist.gov/dads/HTML/towershanoi.html>

Amarillo Datasafe Web Site. Backup Rotation Methods.
URL: http://www.amaonline.com/dlps/backup_rotation.htm

Landscape Web Site. “Tape Backup Strategy.”
URL: <http://www.landscape.com.au/support.backup.htm>

Koller, Mike. “Service Takes Collective Approach to Backup.” Internet Week. July 10, 2000.
URL: <Http://www.internetwk.com/story/INW20000710S0008>

Unix Systems Independent Learning (USAIL) Backups. 12 May 1999.
URL: www.uwsg.indiana.edu/usail/index/backup.html

University of Melbourne – Academic and Corporate Services – IT. 25 June 2001.
URL: <http://www.acs.unimelb.edu.au/backups/strategy.html>

© SANS Institute 2000 - 2005, Author retains full rights.