

CS615 - Aspects of System Administration

Backup, Restore and Disaster Recovery

Department of Computer Science
Stevens Institute of Technology
Jan Schaumann

`jschauma@cs.stevens-tech.edu`

`http://www.cs.stevens-tech.edu/~jschauma/615/`

Related topics

- backup vs. restore

Related topics

- backup vs. restore
- backup devices and media



Related topics

- backup vs. restore
- backup devices and media



Related topics

- backup vs. restore
- backup devices and media



Related topics

- backup vs. restore
- backup devices and media
- filesystem considerations

Related topics

- backup vs. restore
- backup devices and media
- filesystem considerations
- backup strategies

Related topics

- backup vs. restore
- backup devices and media
- filesystem considerations
- backup strategies
- planning for disasters

Backups and Restore Basics

Why are backups needed?

Backups and Restore Basics

Why are backups needed?

- off-site storage of sensitive data
- long-term data storage
- data loss

Backups and Restore Basics

Why are backups needed?

- off-site storage of sensitive data
- long-term data storage
- data loss due to



Backups and Restore Basics

Why are backups needed?

- off-site storage of sensitive data
- long-term data storage
- data loss due to



Backups and Restore Basics

Why are backups needed?

- off-site storage of sensitive data
- long-term data storage
- data loss due to



Backups and Restore Basics

Why are backups needed?

- off-site storage of sensitive data
- long-term data storage
- data loss due to



Backups and Restore Basics

Why are backups needed?

- off-site storage of sensitive data
- long-term data storage
- data loss due to



Backups and Restore Basics

Why are backups needed?

- off-site storage of sensitive data
- long-term data storage
- data loss due to



Backups and Restore Basics

Why are backups needed?

- off-site storage of sensitive data
- long-term data storage
- data loss due to



Backups and Restore Basics

Why are backups needed?

- off-site storage of sensitive data
- long-term data storage
- data loss due to
 - equipment failure
 - bozotic users
 - natural disaster
 - security breach
 - software bugs

Backups and Restore Basics

Why are backups needed?

- off-site storage of sensitive data
- long-term data storage
- data loss due to
 - equipment failure
 - bozotic users
 - natural disaster
 - security breach
 - software bugs

Think of your backups as *insurance*: you invest and pay for it, hoping you will never need it.

Backup and Restore Basics

Simply adding disks to a system is easy – determining an efficient backup and restore system for these disks is *hard*.

Terminology:

- the *backup of a file system* is often called a dump
- a *full backup* is referred to as a level 0 dump
- subsequent *partial dumps* are referred to as incremental dumps
- while most things we consider hold for *any* backup media, we'll simply refer to backup tapes

Developing a Backup and Restore System

Determine the desired end-result:

- minimum set of data considered absolutely crucial

Developing a Backup and Restore System

Determine the desired end-result:

- minimum set of data considered absolutely crucial
- time frames and downtime

Developing a Backup and Restore System

Determine the desired end-result:

- minimum set of data considered absolutely crucial
- time frames and downtime
- length of archival

Developing a Backup and Restore System

Determine the desired end-result:

- minimum set of data considered absolutely crucial
- time frames and downtime
- length of archival
- how are files restored

Developing a Backup and Restore System

Determine the desired end-result:

- minimum set of data considered absolutely crucial
- time frames and downtime
- length of archival
- how are files restored

Then, tune these guidelines with less and less abstraction:

- service level agreement
- policy
- schedule
- software configuration

Key Reasons for Restores

Three key reasons for restores: *Accidental File Deletion*, *Disk Failure* and *Archival*.

1. Accidental File Deletion

- ability to restore a file within a certain time frame
- restore time, including
 - actual time spent restoring
 - waiting until resources permit the restore
 - staff availability
- self-service restore

Key Reasons for Restores

2. Disk Failure

- loss of entire file system
- leads to downtime
- RAID may help
- takes long time to restore

3. Archival

- *full* set of level 0 backups
- separate set from regular backups
- usually stored off-site
- store for long time

Example Data-Recovery Service Level Agreement

1. users should be able to get back a file with a granularity of one day for the last 7 days
2. users should be able to get back a file with a granularity of one month for the last year
3. disk failures should be restored in four hours
4. disk failure must not lead to data loss of more than two business days
5. archival of all content every quarter
6. archives are stored “forever”

Backup Schedule

SLA dictates one-day granularity. So:

- perform daily (incremental) backups
- at least one full dump per month to be stored for 12 months
- a full dump every 3 months, stored appropriately (offsite?)

Backup Cycle

Do the math:

Level-1 only

| Day Number | Daily | 7-day | 14-day | 21-day |
|---------------------|----------------|-------------------|-------------------|------------------|
| 1 | 500 | 500 | 500 | 500 |
| 2 | 500 | 25 | 25 | 25 |
| 3 | 500 | 26.25 | 26.25 | 26.25 |
| 4 | 500 | 26.31 | 26.31 | 26.31 |
| 5 | 500 | 26.32 | 26.32 | 26.32 |
| 6 | 500 | 26.32 | 26.32 | 26.32 |
| 7 | 500 | 500 | 26.32 | 26.32 |
| 8 | 500 | 25 | 26.32 | 26.32 |
| 9 | 500 | 26.25 | 26.32 | 26.32 |
| 10 | 500 | 26.31 | 26.32 | 26.32 |
| 11 | 500 | 26.32 | 26.32 | 26.32 |
| 12 | 500 | 26.32 | 26.32 | 26.32 |
| 13 | 500 | 26.32 | 26.32 | 26.32 |
| 14 | 500 | 500 | 500 | 26.32 |
| 15 | 500 | 25 | 25 | 26.32 |
| ... | ... | ... | ... | ... |
| 15-day total | 7500 GB | 1786.73 GB | 1340.76 GB | 867.08 GB |

Backup Cycle

Do the math:

Multiple levels:

| Day Number | L0 | L1 + L2 | L1 + L2 + L3 | L3 + L2 + L1 |
|--------------------|----------------|------------------|------------------|------------------|
| 1 | 500 | 500.00 | 500.00 | 500.00 |
| 2 | 500 | 25.00 | 25.00 | 25.00 |
| 3 | 500 | 1.25 | 1.25 | 26.25 |
| 4 | 500 | 26.87 | 0.06 | 26.31 |
| 5 | 500 | 1.34 | 25.00 | 1.31 |
| 6 | 500 | 25.06 | 1.25 | 1.38 |
| 7 | 500 | 1.25 | 0.06 | 25.07 |
| ... | ... | ... | ... | ... |
| 7-day total | 3500 GB | 580.79 GB | 552.63 GB | 605.33 GB |

Reading

- `dump(8)` and `restore(8)`

Book: <http://www.oreilly.com/catalog/unixbr/>