CS615 - Aspects of System Administration

SMTP, Backup and Disaster Recovery

Department of Computer Science
Stevens Institute of Technology
Jan Schaumann
jschauma@stevens-tech.edu
http://www.cs.stevens-tech.edu/~jschauma/615A/

Email... still popular

Bad news, everybody: Slack has not yet replaced email.

Email... still popular

Bad news, everybody: Slack has not yet replaced email.

- **4.6 billion** - number of email accounts.
- **269 billion** - Average number of email messages per day.
  That's 3.1 million emails *per second*.
- **121** - Average number of emails an office worker receives.
- **42** - Percentage of Americans that check their email in the bathroom.
- **18** - Percentage of Americans that check their email while driving.
- **>70** - Percentage of emails that are Spam.

Sending...

```
# tcpdump -i xennet0 -w /tmp/t.out port not 22 2>/dev/null &
# mail -s "CS615 - SMTP Exercise" jschauma@stevens.edu -f jschauma@stevens.edu
Hello,

SMTP is simple.

-Jan
.
EOT
# fg
tcpdump -i xennet0 -w /tmp/t.out port not 22 2>/dev/null
^C
```

Sending...

```
# tail -5 /var/log/maillog
Apr  4 15:42:33 ip-10-235-167-232 postfix/pickup[848]: 2A17275438:
        uid=0 from=<jschauma@stevens.edu>
Apr  4 15:42:33 ip-10-235-167-232 postfix/cleanup[765]: 2A17275438:
        message-id=<20160404154233.2A17275438@ip-10-235-167-232.ec2.internal>
Apr  4 15:42:33 ip-10-235-167-232 postfix/qmgr[876]: 2A17275438:
        from=<jschauma@stevens.edu>, size=380, nrcpt=1 (queue active)
Apr  4 15:42:33 ip-10-235-167-232 postfix/smtp[1124]: 2A17275438:
        to=<jschauma@stevens.edu>, relay=spamfilter01.stevens.edu[155.246.14.37]:25,
        delay=0.62, delays=0.04/0.01/0.03/0.54, dsn=2.0.0,
        status=sent (250 Ok: queued as 688CD6F4001)
Apr  4 15:42:33 ip-10-235-167-232 postfix/qmgr[876]: 2A17275438: removed
```

Sending...

```
# tcpdump -t -r /tmp/t.out port 53
IP 10.235.167.232.65498 > 172.16.0.23.domain: 61195+ MX? stevens.edu. (29)
IP 172.16.0.23.domain > 10.235.167.232.65498: 61195 2/0/0
        MX spamfilter01.stevens.edu. 10,
        MX spamfilter02.stevens.edu. 20 (87)
IP 10.235.167.232.65497 > 172.16.0.23.domain: 1949+
        A? spamfilter01.stevens.edu. (42)
IP 172.16.0.23.domain > 10.235.167.232.65497: 1949 1/0/0 A 155.246.14.37 (58)
IP 10.235.167.232.65496 > 172.16.0.23.domain: 39922+
        AAAA? spamfilter01.stevens.edu. (42)
IP 172.16.0.23.domain > 10.235.167.232.65496: 39922 0/1/0 (113)
IP 10.235.167.232.65495 > 172.16.0.23.domain: 26844+
        A? spamfilter02.stevens.edu. (42)
IP 172.16.0.23.domain > 10.235.167.232.65495: 26844 1/0/0 A 155.246.248.24 (58)
IP 10.235.167.232.65494 > 172.16.0.23.domain: 1439+
        AAAA? spamfilter02.stevens.edu. (42)
IP 172.16.0.23.domain > 10.235.167.232.65494: 1439 0/1/0 (113)
```

Sending...

```
# host -t mx stevens.edu
stevens.edu mail is handled by 20 spamfilter02.stevens.edu.
stevens.edu mail is handled by 10 spamfilter01.stevens.edu.
# host spamfilter01.stevens.edu.
spamfilter01.stevens.edu has address 155.246.14.37
# host spamfilter02.stevens.edu.
spamfilter02.stevens.edu has address 155.246.248.24
#
```

Sending...

```
IP 10.235.167.232.65524 > 155.246.14.37.smtp: Flags [S], seq 1528496417,
IP 155.246.14.37.smtp > 10.235.167.232.65524: Flags [S.], seq 3510048077, ack 1528496
IP 10.235.167.232.65524 > 155.246.14.37.smtp: Flags [.], ack 1,
IP 155.246.14.37.smtp > 10.235.167.232.65524: Flags [P.], seq 1:72, ack 1, length 71
IP 10.235.167.232.65524 > 155.246.14.37.smtp: Flags [P.], seq 1:38, ack 72, length 37
IP 155.246.14.37.smtp > 10.235.167.232.65524: Flags [.], ack 38,
IP 155.246.14.37.smtp > 10.235.167.232.65524: Flags [P.], seq 72:244, ack 38, length
IP 10.235.167.232.65524 > 155.246.14.37.smtp: Flags [P.], seq 38:119, ack 244, length
IP 155.246.14.37.smtp > 10.235.167.232.65524: Flags [P.], seq 244:282, ack 119, lengt
IP 10.235.167.232.65524 > 155.246.14.37.smtp: Flags [.], ack 282,
IP 155.246.14.37.smtp > 10.235.167.232.65524: Flags [P.], seq 282:369, ack 119, lengt
IP 10.235.167.232.65524 > 155.246.14.37.smtp: Flags [P.], seq 119:508, ack 369, lengt
IP 155.246.14.37.smtp > 10.235.167.232.65524: Flags [.], ack 508
IP 155.246.14.37.smtp > 10.235.167.232.65524: Flags [P.], seq 369:400, ack 508, lengt
IP 155.246.14.37.smtp > 10.235.167.232.65524: Flags [FP.], seq 400:499, ack 508, leng
IP 10.235.167.232.65524 > 155.246.14.37.smtp: Flags [.], ack 500
IP 10.235.167.232.65524 > 155.246.14.37.smtp: Flags [F.], seq 508, ack 500
IP 155.246.14.37.smtp > 10.235.167.232.65524: Flags [.], ack 509
```

Sending...

```
$ telnet 155.246.14.37 25
Trying 155.246.14.37...
Connected to spamfilter01.stevens.edu.
Escape character is '^]'.
220 spamfilter01.stevens.edu ESMTP (fe32969a29a5f461e53bf93b18c8fdb5)
EHLO ip-10-235-167-232.ec2.internal
250-spamfilter01.stevens.edu Hello ec2-54-205-68-41.compute-1.amazonaws.c
        pleased to meet you
250-SIZE 50000000
250-PIPELINING
250-8BITMIME
250 HELP
MAIL FROM:<jschauma@stevens.edu> SIZE=380
250 Sender <jschauma@stevens.edu> OK
RCPT TO:<jschauma@stevens.edu>
250 Recipient <jschauma@stevens.edu> OK
```

Sending...

```
DATA
354 Start mail input; end with <CRLF>.<CRLF>
Received: by ip-10-235-167-232.ec2.internal (Postfix, from userid 0)
        id 2A17275438; Mon,  4 Apr 2016 15:42:33 +0000 (UTC)
To: jschauma@stevens.edu
Subject: CS615 - SMTP Exercise
Message-Id: <20160404154233.2A17275438@ip-10-235-167-232.ec2.internal>
Date: Mon,  4 Apr 2016 15:42:33 +0000 (UTC)
From: jschauma@stevens.edu (Charlie Root)

Hello,

SMTP is simple.

-Jan
.
250 Ok: queued as 6A9C76F4004
```

SMTP Codes

SMTP codes consist of three digits in five classes:

- **1xx** – Mail server has accepted the command, but does not yet take any action. A confirmation message is required.
- **2xx** – Mail server has completed the task successfully without errors.
- **3xx** – Mail server has understood the request, but requires further information to complete it.
- **4xx** – Mail server has encountered a temporary failure. If the command is repeated without any change, it might be completed. Try again, it may help!
- **5xx** – Mail server has encountered a fatal error. Your request can't be processed.

Receiving...

```
Date: Mon, 4 Apr 2016 15:42:33 +0000
From: jschauma@stevens.edu (Charlie Root)
To: Jan Schaumann <jschauma@stevens.edu>
Subject: CS615 - SMTP Exercise

Hello,

SMTP is simple.

-Jan
```

Receiving...

```
From jschauma@stevens.edu  Mon Apr  4 11:42:35 2016
Received: by panix.netmeister.org (Postfix, from userid 1004)
        id 6B0F56513D; Mon,  4 Apr 2016 11:42:35 -0400 (EDT)
Received: from nexus.stevens.edu (nexus.stevens.edu [155.246.14.12])
        by panix.netmeister.org (Postfix) with ESMTP id 2AD596513B
Received: from exchng02.campus.stevens-tech.edu (exchng02.campus.stevens-tech.edu [15
        by nexus.stevens.edu (Postfix) with ESMTPS id 11E3817F825
Received: from exchng04.campus.stevens-tech.edu (2002:9bf6:f826::9bf6:f826) by
        exchng02.campus.stevens-tech.edu (2002:9bf6:e17::9bf6:e17) with Microsoft
        SMTP Server (TLS) id 15.0.1104.5; Mon, 4 Apr 2016 11:42:34 -0400
Received: from exchng03.campus.stevens-tech.edu (155.246.248.36) by
        exchng04.campus.stevens-tech.edu (155.246.248.39) with Microsoft SMTP Server
        (TLS) id 15.0.1104.5; Mon, 4 Apr 2016 11:42:34 -0400
Received: from exchng03.campus.stevens-tech.edu ([::1]) by
        exchng03.campus.stevens-tech.edu ([fe80::599a:f128:d1b3:4ce7%12]) with
        Microsoft SMTP Server id 15.00.1104.000; Mon, 4 Apr 2016 11:42:34 -0400
From: Jan Schaumann <jschauma@stevens.edu>
To: Jan Schaumann <jschauma@stevens.edu>
Subject: CS615 - SMTP Exercise
Date: Mon, 4 Apr 2016 15:42:33 +0000
Message-ID: <1b1399e9c44b494f99e9d0030f0fa74b@exchng03.campus.stevens-tech.edu>
x-barracuda-apparent-source-ip: 54.205.68.41
x-ms-exchange-parent-message-id: <20160404154233.2A17275438@ip-10-235-167-232.ec2.int
Resent-Message-Id: <20160404154235.11E3817F825@nexus.stevens.edu>
```

Receiving...

```
$ tail -f /var/log/maillog
postfix/smtpd[552]: connect from nexus.stevens.edu[155.246.14.12]
postfix/smtpd[552]: 2AD596513B: client=nexus.stevens.edu[155.246.14.12]
postfix/cleanup[6975]: 2AD596513B:
        message-id=<1b1399e9c44b494f99e9d0030f0fa74b@exchng03.campus.stevens-tech.edu
postfix/cleanup[6975]: 2AD596513B:
        resent-message-id=<20160404154235.11E3817F825@nexus.stevens.edu>
postfix/smtpd[552]: disconnect from nexus.stevens.edu[155.246.14.12]
postfix/qmgr[24644]: 2AD596513B: from=<jschauma@stevens.edu>, size=3116, nrcpt=1 (que
spamd[24829]: spamd: connection from localhost [::1]:58531 to port 783, fd 5
spamd: processing message <1b1399e9c44b494f99e9d0030f0fa74b@exchng03.campus.stevens-t
        aka <20160404154235.11E3817F825@nexus.stevens.edu> for spamd:1004
postfix/pipe[8278]: 2AD596513B: to=<jschauma@netmeister.org>, relay=spamassassin, del
        delays=0.05/0.01/0/0.23, dsn=2.0.0, status=sent (delivered via spamassassin s
postfix/qmgr[24644]: 2AD596513B: removed
postfix/pickup[16387]: 6B0F56513D: uid=1004 from=<jschauma@stevens.edu>
postfix/cleanup[6975]: 6B0F56513D:
        message-id=<1b1399e9c44b494f99e9d0030f0fa74b@exchng03.campus.stevens-tech.edu
postfix/cleanup[6975]: 6B0F56513D:
        resent-message-id=<20160404154235.11E3817F825@nexus.stevens.edu>
postfix/qmgr[24644]: 6B0F56513D: from=<jschauma@stevens.edu>, size= 3442, nrcpt=1 (qu
postfix/local[2264]: 6B0F56513D: to=<jschauma@netmeister.org>, relay=local,
        delay=0.15, delays=0.08/0.01/0/0.06, dsn=2.0.0, status=sent (delivered to com
        /usr/pkg/bin/procmail)
```

Receiving...

```
IP 155.246.14.12.37547 > 166.84.7.99.25: Flags [S], seq 2435028943
IP 166.84.7.99.25 > 155.246.14.12.37547: Flags [S.], seq 4137148947, ack 2435028944
IP 155.246.14.12.37547 > 166.84.7.99.25: Flags [.], ack 1
IP 166.84.7.99.51798 > 166.84.67.2.53: 29782+ PTR? 12.14.246.155.in-addr.arpa. (44)
IP 166.84.67.2.53 > 166.84.7.99.51798: 29782 1/2/2 PTR nexus.stevens.edu. (160)
IP 166.84.7.99.51797 > 166.84.67.2.53: 17361+ A? nexus.stevens.edu. (35)
IP 166.84.67.2.53 > 166.84.7.99.51797: 17361 1/3/3 A 155.246.14.12 (173)
```

Receiving...

```
IP 166.84.7.99.25 > 155.246.14.12.37547: Flags [P.], seq 1:41, ack 1
IP 155.246.14.12.37547 > 166.84.7.99.25: Flags [.], ack 41
IP 155.246.14.12.37547 > 166.84.7.99.25: Flags [P.], seq 1:25, ack 41
IP 166.84.7.99.25 > 155.246.14.12.37547: Flags [P.], seq 41:174, ack 25
IP 155.246.14.12.37547 > 166.84.7.99.25: Flags [P.], seq 25:157, ack 174
IP 166.84.7.99.51796 > 166.84.67.2.53: 41349+ MX? stevens.edu. (29)
IP 166.84.67.2.53 > 166.84.7.99.51796: 41349 2/3/5 MX spamfilter02.stevens.edu.
      20, MX spamfilter01.stevens.edu. 10 (241)
IP 166.84.7.99.51795 > 166.84.67.2.53: 51732+ A? nexus.stevens.edu. (35)
IP 166.84.67.2.53 > 166.84.7.99.51795: 51732 1/3/3 A 155.246.14.12 (173)
IP 166.84.7.99.51794 > 166.84.67.2.53: 64776+ A? 12.14.246.155.sbl.spamhaus.org. (48)
IP 166.84.67.2.53 > 166.84.7.99.51794: 64776 NXDomain 0/1/0 (112)
```

Receiving...

```
IP 166.84.7.99.25 > 155.246.14.12.37547: Flags [P.], seq 174:239, ack 157
IP 155.246.14.12.37547 > 166.84.7.99.25: Flags [.], seq 157:1605, ack 239
IP 155.246.14.12.37547 > 166.84.7.99.25: Flags [.], seq 1605:3053, ack 239
IP 155.246.14.12.37547 > 166.84.7.99.25: Flags [P.], seq 3053:3080, ack 239
IP 166.84.7.99.25 > 155.246.14.12.37547: Flags [.], ack 3053, win 4016,
IP 166.84.7.99.25 > 155.246.14.12.37547: Flags [P.], seq 239:290, ack 3080
IP 166.84.7.99.25 > 155.246.14.12.37547: Flags [F.], seq 290, ack 3080
IP 155.246.14.12.37547 > 166.84.7.99.25: Flags [F.], seq 3080, ack 290
IP 166.84.7.99.25 > 155.246.14.12.37547: Flags [.], ack 3081, win 4197
IP 155.246.14.12.37547 > 166.84.7.99.25: Flags [.], ack 291, win 123
```

Receiving...

```
IP 166.84.7.99.50935 > 166.84.67.2.53: 12067 [1au] A? 12.14.246.155.psbl.surriel.com. (59)
IP 166.84.7.99.50935 > 166.84.67.2.53: 39395 [1au] A? 12.14.246.155.bb.barracudacentral.org. (66)
IP 166.84.7.99.50935 > 166.84.67.2.53: 17365 [1au] A? 12.14.246.155.bl.score.senderscore.com. (67)
IP 166.84.7.99.50935 > 166.84.67.2.53: 38635 [1au] TXT? 37.248.246.155.bl.spamcop.net. (58)
IP 166.84.7.99.50935 > 166.84.67.2.53: 10907 [1au] TXT? 21.14.246.155.bl.spamcop.net. (57)
IP 166.84.7.99.50935 > 166.84.67.2.53: 11246 [1au] TXT? 12.14.246.155.bl.spamcop.net. (57)
IP 166.84.67.2.53 > 166.84.7.99.50935: 39395 0/2/7 (206)
IP 166.84.67.2.53 > 166.84.7.99.50935: 12067 0/4/6 (247)
IP 166.84.67.2.53 > 166.84.7.99.50935: 17365 0/4/5 (254)
IP 166.84.67.2.53 > 166.84.7.99.50935: 38635 0/8/9 (354)
IP 166.84.67.2.53 > 166.84.7.99.50935: 13318 0/8/9 (403)
IP 166.84.67.2.53 > 166.84.7.99.50935: 10907 0/8/9 (353)
IP 166.84.7.99.50935 > 166.84.67.2.53: 8582 [1au] A? 37.248.246.155.dnsbl.sorbs.net. (59)
IP 166.84.7.99.50935 > 166.84.67.2.53: 61203 [1au] A? 21.14.246.155.dnsbl.sorbs.net. (58)
IP 166.84.7.99.50935 > 166.84.67.2.53: 59454 [1au] A? 12.14.246.155.dnsbl.sorbs.net. (58)
IP 166.84.7.99.50935 > 166.84.67.2.53: 7289 [1au] A? 12.14.246.155.iadb.isipp.com. (57)
IP 166.84.7.99.50935 > 166.84.67.2.53: 63596 [1au] TXT? 12.14.246.155.sa-trusted.bondedsender.org. (70)
IP 166.84.7.99.50935 > 166.84.67.2.53: 18732 [1au] TXT? 12.14.246.155.sa-accredit.habeas.com.
(65)
IP 166.84.7.99.50935 > 166.84.67.2.53: 43600 [1au] A? 12.14.246.155.bl.mailspike.net. (59)
IP 166.84.67.2.53 > 166.84.7.99.50935: 49122 0/19/80 (1472)
IP 166.84.67.2.53 > 166.84.7.99.50935: 6461 0/6/10 (360)
IP 166.84.67.2.53 > 166.84.7.99.50935: 54452 0/13/14 (607)
IP 166.84.67.2.53 > 166.84.7.99.50935: 8582 0/13/14 (558)
IP 166.84.67.2.53 > 166.84.7.99.50935: 61203 0/13/14 (557)
IP 166.84.67.2.53 > 166.84.7.99.50935: 59454 0/13/14 (557)
IP 166.84.67.2.53 > 166.84.7.99.50935: 7289 0/5/7 (288)
IP 166.84.67.2.53 > 166.84.7.99.50935: 63596 0/18/19 (879)
IP 166.84.67.2.53 > 166.84.7.99.50935: 18732 0/15/16 (742)
IP 166.84.67.2.53 > 166.84.7.99.50935: 43600 0/2/3 (126)
IP 166.84.7.99.50935 > 166.84.67.2.53: 2649 [1au] TXT? nexus.stevens.edu. (46)
IP 166.84.67.2.53 > 166.84.7.99.50935: 2649 0/3/4 (168)
```

Relaying mail

```
$ telnet spamfilter01.stevens.edu 25
Trying 155.246.14.37...
Connected to spamfilter01.stevens.edu.
Escape character is '^]'.
220 spamfilter01.stevens.edu ESMTP (fe32969a29a5f461e53bf93b18c8fdb5)
ehlo localhost
250-spamfilter01.stevens.edu Hello ec2-54-205-68-41.compute-1.amazonaws.com
        [54.205.68.41], pleased to meet you
250-SIZE 50000000
250-PIPELINING
250-8BITMIME
250 HELP
MAIL FROM:<obama@whitehouse.gov>
250 Sender <obama@whitehouse.gov> OK
RCPT TO:<comey@fbi.gov>
550 No such domain at this location
```

Authenticity

```
MAIL FROM:<obama@whitehouse.gov>
250 Sender <obama@whitehouse.gov> OK
RCPT TO:<jschauma@stevens-tech.edu>
250 Recipient <jschauma@stevens-tech.edu> OK
DATA
354 Start mail input; end with <CRLF>.<CRLF>
From: "Barack Obama" <obama@whitehouse.gov>
Subject: Friday

Yo,

Party at my house.
BYOB.

-B
.
250 Ok: queued as 61D4D6F4003
```

Authenticity

```
MAIL FROM:<president@stevens.edu>
250 Sender <president@stevens.edu> OK
RCPT TO:<jschauma@stevens-tech.edu>
250 Recipient <jschauma@stevens-tech.edu> OK
DATA
354 Start mail input; end with <CRLF>.<CRLF>
From: "Barack Obama" <obama@whitehouse.gov>
Subject: Friday

Yo,

Party at my house.
BYOB.

-B
.
250 Ok: queued as 61D4D6F4003
```

Authenticity

```
Date: Mon,  4 Apr 2016 15:35:12 -0400 (EDT)
From: Barack Obama <obama@whitehouse.gov>
To: undisclosed-recipients: ;
Subject: Friday

Yo,

Party at my house.
BYOB.

-B
```

Authenticity

```
From president@stevens.edu  Mon Apr  4 15:35:12 2016
Return-Path: <president@stevens.edu>
Date: Mon,  4 Apr 2016 15:35:12 -0400 (EDT)
From: Barack Obama <obama@whitehouse.gov>
To: undisclosed-recipients: ;
Subject: Friday

Yo,

Party at my house.
BYOB.

-B
```

SMTP is a Simple Mail Transfer Protocol.

- TCP port 25
- DNS MX records
- mail may be relayed or processed by many servers in transit
- transport is in clear text
- STARTTLS may provide (opportunistic) transport encryption
- SPAM controls may include DNS lookups, bayesian scoring, ...
- authenticity not guaranteed

The Mail System

Divided into:

- *Mail User Agent* or MUA, such as `mutt(1)`, *Mail.app*, *Outlook*, a browser (ugh) ...
- *Mail Transfer Agent* or MTA, such as *postfix*, *sendmail*, *qmail*, ...
- *Mail Delivery Agent* or MDA, such as *procmail*
- *Access Agent* providing access via *POP*, *IMAP* etc.

Anatomy of an email message

An email consists of:

- mandatory headers (such as ”From ”, ”Delivered-To: ”, ...)
- optional headers (such as ”From: ”, ”To: ”, ”Subject: ”, ...)
- the (optional) body of the message

Service Considerations

- outsourcing versus in-house
- privacy considerations
- spam protections
- phishing protections
- mail delivery cannons for notifications vs. spam lists
- high volume traffic demands fine-tuned systems
- high volume traffic implications on logging

See also:

- `http://is.gd/JQp1zM`
- `http://is.gd/cXyrwX`
- `http://is.gd/o6Y5f8`

Hooray!

5 Minute Break

Know a Unix Command

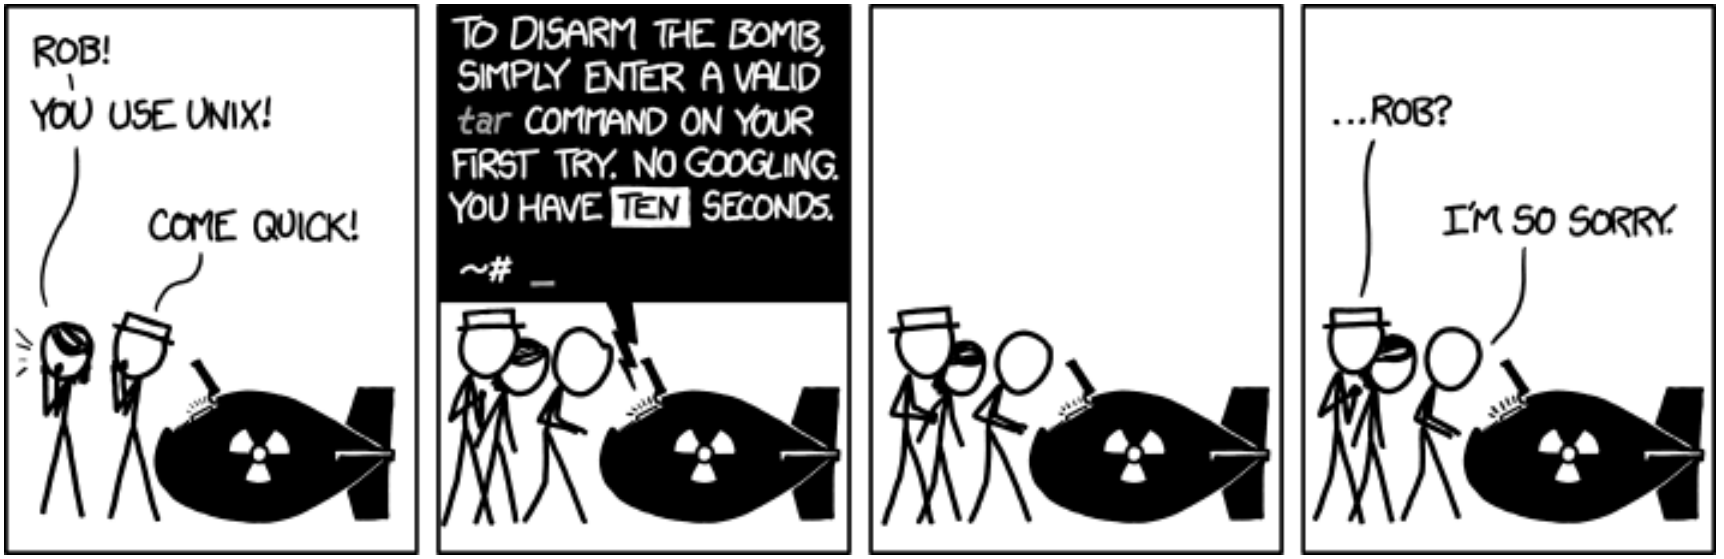

https://www.xkcd.com/1168/

https://www.cs.stevens.edu/~jschauma/615/tar.html

Backups

- backup vs. restore

Backups

- backup vs. restore
- backup devices and media

Backups

- backup vs. restore
- backup devices and media

Backups

- backup vs. restore
- backup devices and media

Backups

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

Backups

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

Backups

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

Backups and Restore Basics

When do we need backups?

Backups and Restore Basics

When do we need backups?

- disaster recovery: off-site storage of sensitive data
- long-term storage requirements
- recover from data loss

Backups and Restore Basics

When do we need backups?

- disaster recovery: off-site storage of sensitive data
- long-term storage requirements
- recover from data loss due to

Backups and Restore Basics

When do we need backups?

- disaster recovery: off-site storage of sensitive data
- long-term storage requirements
- recover from data loss due to

Backups and Restore Basics

When do we need backups?

- disaster recovery: off-site storage of sensitive data
- long-term storage requirements
- recover from data loss due to

Backups and Restore Basics

When do we need backups?

- disaster recovery: off-site storage of sensitive data
- long-term storage requirements
- recover from data loss due to

Backups and Restore Basics

When do we need backups?

- disaster recovery: off-site storage of sensitive data
- long-term storage requirements
- recover from data loss due to

Backups and Restore Basics

When do we need backups?

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

Backups and Restore Basics

When do we need backups?

- disaster recovery: off-site storage of sensitive data
- long-term storage requirements
- recover from 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.

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

Filesystem backup

`dump(8)` / `restore(8)`

- in use since ~1975
- full filesystem level backups
- direct interaction with tape devices
- integration with `/etc/fstab`
- efficient incremental backups

Filesystem backup

- start an EC2 instance
- create a full filesystem backup using `dump(8)`
- add e.g. the 'apache' package
- create an incremental backup
- delete the 'apache' package
- restore all files from the incremental backup using the `restore(8)` command
- verify that the 'apache' package is fully installed

Filesystem backup

```
ssh ec2-instance "sudo dump -u -0 -f - /" | bzip2 -c -9 >tmp/ec2.0.bz2
  DUMP: Date of this level 0 dump: Thu Mar 26 19:25:06 2015
  DUMP: Dumping /dev/xvda1 (/) to standard output
  DUMP: Label: cloudimg-rootfs
  DUMP: Writing 10 Kilobyte records
  DUMP: mapping (Pass I) [regular files]
  DUMP: mapping (Pass II) [directories]
  DUMP: estimated 823759 blocks.
  DUMP: Volume 1 started with block 1 at: Thu Mar 26 19:25:07 2015
  DUMP: dumping (Pass III) [directories]
  DUMP: dumping (Pass IV) [regular files]
  DUMP: Volume 1 completed at: Thu Mar 26 19:28:21 2015
  DUMP: Volume 1 820690 blocks (801.46MB)
  DUMP: Volume 1 took 0:03:14
  DUMP: Volume 1 transfer rate: 4230 kB/s
  DUMP: 820690 blocks (801.46MB)
  DUMP: finished in 194 seconds, throughput 4230 kBytes/sec
  DUMP: Date of this level 0 dump: Thu Mar 26 19:25:06 2015
  DUMP: Date this dump completed:  Thu Mar 26 19:28:21 2015
  DUMP: Average transfer rate: 4230 kB/s
  DUMP: DUMP IS DONE
```

Filesystem backup

```
$ cat /var/lib/dumpdates
/dev/xvda1 0 Thu Mar 26 19:25:06 2015 +0000
$ sudo apt-get install apache2
[...]
$ ssh ec2-instance "sudo dump -u -i -f - /" | bzip2 -c -9 >tmp/ec2.1.bz2
  DUMP: Date of this level i dump: Thu Mar 26 19:56:45 2015
  DUMP: Date of last level 0 dump: Thu Mar 26 19:48:58 2015
  DUMP: Dumping /dev/xvda1 (/) to standard output
  DUMP: Writing 10 Kilobyte records
  DUMP: mapping (Pass I) [regular files]
  DUMP: mapping (Pass II) [directories]
  DUMP: estimated 39644 blocks.
  DUMP: Volume 1 started with block 1 at: Thu Mar 26 19:56:50 2015
  DUMP: dumping (Pass III) [directories]
  DUMP: dumping (Pass IV) [regular files]
  DUMP: Volume 1 completed at: Thu Mar 26 19:56:56 2015
  DUMP: Volume 1 39820 blocks (38.89MB)
  DUMP: Volume 1 transfer rate: 6636 kB/s
  DUMP: 39820 blocks (38.89MB)
  DUMP: finished in 5 seconds, throughput 7964 kBytes/sec
  DUMP: Date of this level i dump: Thu Mar 26 19:56:45 2015
  DUMP: Date this dump completed:  Thu Mar 26 19:56:56 2015
  DUMP: Average transfer rate: 6636 kB/s
  DUMP: DUMP IS DONE
```

Filesystem backup

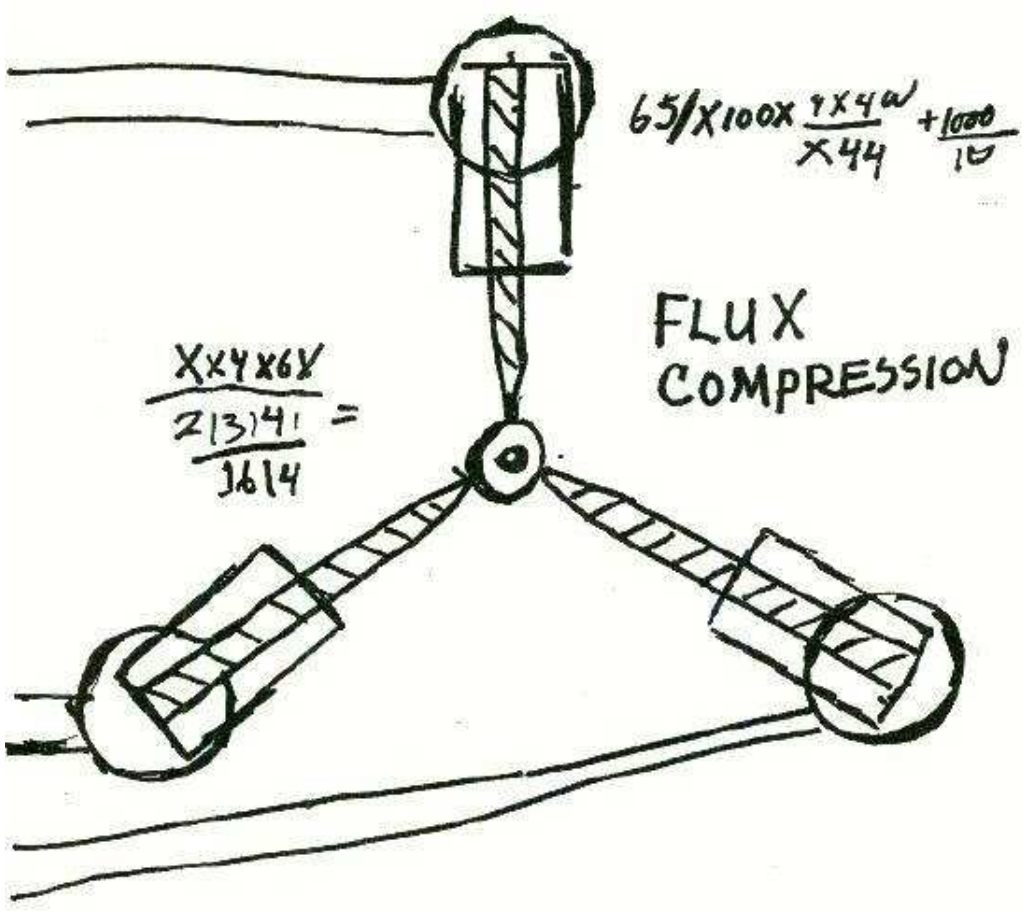

Filesystem backup

Filesystem backup

Filesystem backup

Example: Mac OS X “Time Machine”:

- automatically creates a full backup (equivalent of a ”level 0 dump”) to separate device or NAS, recording (specifically) last-modified date of all directories
- every hour, creates a full copy via *hardlinks* (hence no additional disk space consumed) for files that have not changed, new copy of files that have changed
- changed files are determined by inspecting last-modified date of directories (cheaper than doing comparison of all files’ last-modified date or data)
- saves hourly backups for 24 hours, daily backups for the past month, and weekly backups for everything older than a month.

Filesystem backup

Example: WAFL (Write Anywhere File Layout)

- used by NetApp's "Data ONTAP" OS
- a snapshot is a read-only copy of a file system (cheap and near instantaneous, due to CoW)
- uses regular snapshots ("consistency points", every 10 seconds) to allow for speedy recovery from crashes

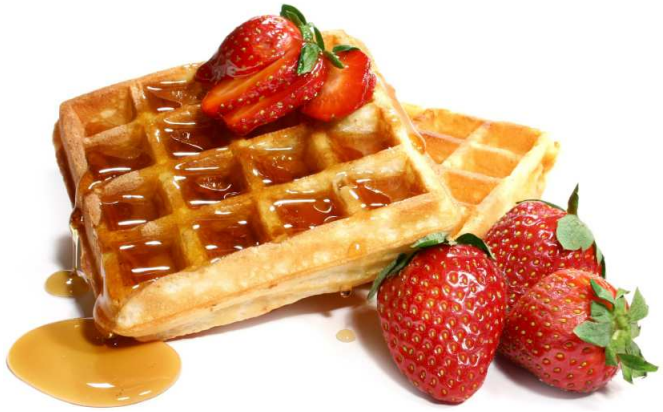

Filesystem backup

Example: WAFL (Write Anywhere File Layout)

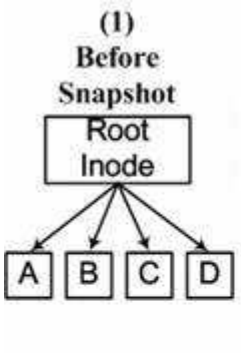

Filesystem backup

Example: WAFL (Write Anywhere File Layout)

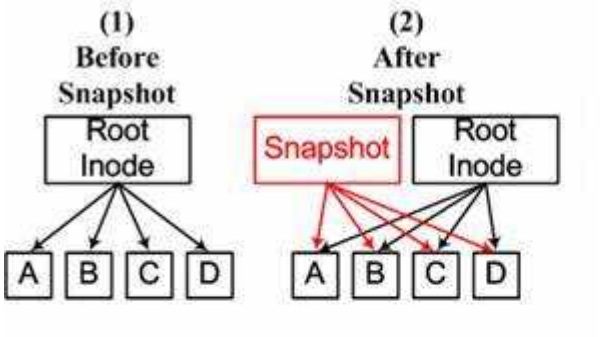

Filesystem backup

Example: WAFL (Write Anywhere File Layout)

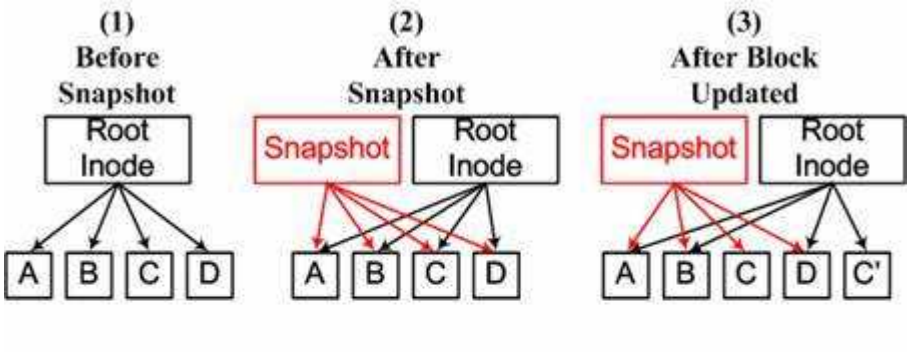

Filesystem backup

Example: WAFL (Write Anywhere File Layout)

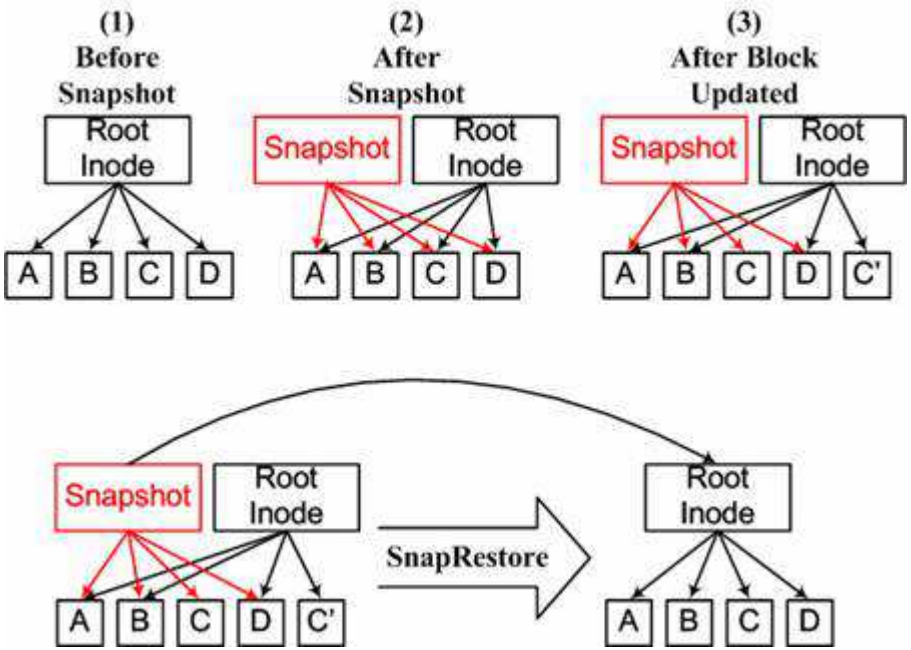

Filesystem backup

Example: ZFS snapshots

- ZFS uses a copy-on-write transactional object model (new data does not overwrite existing data, instead modifications are written to a new location with existing data being referenced), similar to WAFL
- a snapshot is a read-only copy of a file system (cheap and near instantaneous, due to CoW)
- initially consumes no additional disk space; the writable filesystem is made available as a “clone”
- conceptually provides a branched view of the filesystem; normally only the “active” filesystem is writable

ZFS Snapshots

```
$ pwd
/home/jschauma
$ ls -l .z*
ls: cannot access .z*: No such file or directory
$
```

ZFS Snapshots

```
$ pwd
/home/jschauma
$ ls -l .z*
ls: cannot access .z*: No such file or directory
$ ls -lid .zfs
1 dr-xr-xr-x 3 root root 3 Jan 10  2013 .zfs
$
```

ZFS Snapshots

```
$ pwd
/home/jschauma
$ ls -l .z*
ls: cannot access .z*: No such file or directory
$ ls -lid .zfs
1 dr-xr-xr-x 3 root root 3 Jan 10  2013 .zfs
$ ls -lai .zfs/snapshot
total 13
2 dr-xr-xr-x  4 root     root       4 Feb 28 21:00 .
1 dr-xr-xr-x  3 root     root       3 Jan 10  2013 ..
4 drwx--x--x 37 jschauma professor 88 Feb 24 22:32 amanda-_export_home_jschauma-0
4 drwx--x--x 37 jschauma professor 88 Feb 26 11:47 amanda-_export_home_jschauma-1
$
```

ZFS Snapshots

```
$ pwd
/home/jschauma
$ ls -l .z*
ls: cannot access .z*: No such file or directory
$ ls -lid .zfs
1 dr-xr-xr-x 3 root root 3 Jan 10  2013 .zfs
$ ls -lai .zfs/snapshot
total 13
2 dr-xr-xr-x  4 root     root       4 Feb 28 21:00 .
1 dr-xr-xr-x  3 root     root       3 Jan 10  2013 ..
4 drwx--x--x 37 jschauma professor 88 Feb 24 22:32 amanda-_export_home_jschauma-0
4 drwx--x--x 37 jschauma professor 88 Feb 26 11:47 amanda-_export_home_jschauma-1
$ cd .zfs/snapshot
$ echo foo > amanda-_export_home_jschauma-0/oink
-ksh: amanda-_export_home_jschauma-0/oink: cannot create [Read-only file system]
$ ls -laid . /
2 dr-xr-xr-x  4 root root    4 Feb 28 21:00 .
2 drwxr-xr-x 26 root root 4096 Jan 27 11:44 /
```

ZFS Snapshots

```
$ pwd
/home/jschauma/.zfs/snapshot
$ ls -lai amanda-_export_home_jschauma-0 >/tmp/a
$ ls -lai amanda-_export_home_jschauma-1 >/tmp/b
$ diff -bu /tmp/[ab]
--- /tmp/a 2014-03-01 22:55:49.000000000 -0500
+++ /tmp/b 2014-03-01 22:55:59.000000000 -0500
@@ -35,7 +35,7 @@
 57723 drwx------  3 jschauma professor         6 Dec 31 15:08 .subversion
 49431 -rw-------  1 jschauma professor         6 Dec 22 12:25 .sws.pid
    20 drwx------  2 jschauma professor         3 Jan 26 10:30 .vim
-61768 -rw-------  1 jschauma professor     14538 Feb 24 22:32 .viminfo
+61775 -rw-------  1 jschauma professor     14557 Feb 26 09:23 .viminfo
   173 -rw-------  1 jschauma professor      4355 Sep 17  2012 .vimrc
 45744 -rw-r--r--  1 jschauma professor         0 Jul 28  2013 .xsession-errors
    21 drwxr-xr-x  3 jschauma professor         6 Apr  4  2010 CS615A
$
```

Summary

- backups are most commonly done as incrementals of a filesystem, mountpoint, or directory hierarchy
- consider (long-term) storage:
  - media and location
  - increased storage requirements
  - privacy and safety of the data
- self-service restores and filesystem snapshots
- backups need to be:
  - regular, frequent, automated
  - invisible
  - verifiable

Summary

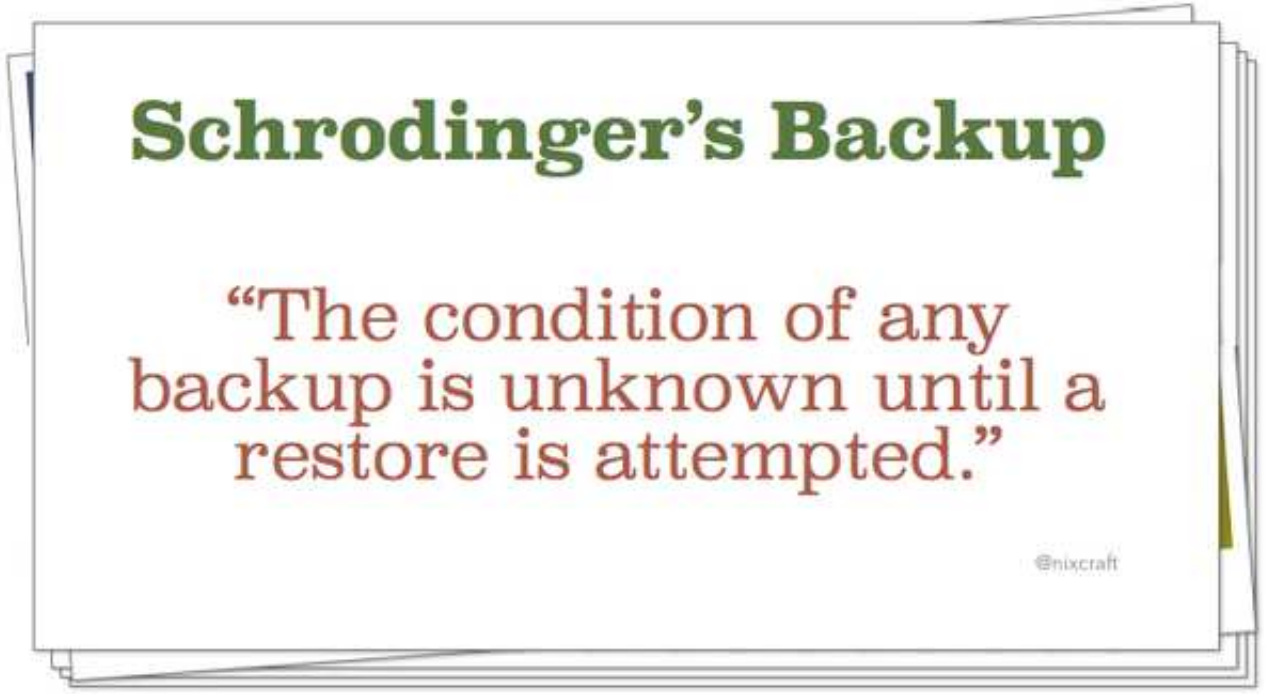
Schrodinger's Backup
"The condition of any backup is unknown until a restore is attempted."
@nixcraft

Reading

Hurricane Sandy

- http://is.gd/aaxzvI
- http://is.gd/Y75pEA
- http://is.gd/32Az7y
- http://is.gd/FhAuFZ

Reading

Backups with `dump(8)` and `restore(8)`:

- `dump(8)` and `restore(8)`
- `http://is.gd/bXG9of`

Filesystem snapshots:

- `https://en.wikipedia.org/wiki/Snapshot_(computer_storage)`
- `https://en.wikipedia.org/wiki/Time_Machine_(Apple_software)`
- `http://comet.lehman.cuny.edu/jung/cmp426697/WAFL.pdf`
- `http://www.cs.tau.ac.il/~ohadrode/slides/WAFL.pdf`

Book: `http://www.oreilly.com/catalog/unixbr/`