Backup & Incident Monitoring

```
         (___)
         ( O )
  /-------\ /
 / |     ||V
*  ||----||
   ^^    ^^
```

System and Network Administration

Revision 2 (2020/21)

Pierre-Philipp Braun <pbraun@nethence.com>

Table of contents

- Backup & Compression
- Migrations & Disaster Recovery
- Incident Monitoring
- Outgoing Email
- Backup Tips & Tricks

Backup systems

Features for pre-sales…

- Compression
- De-duplication
- Policy (full vs differential vs incremental)
- Encryption (if uploading/synching on remote & untrusted environment)

Backup types

Full vs. differential vs. incremental

	full	diff	incr
d1	full	full	full
d2	full	diff-d1	incr-d1
d3	full	diff-d1	incr-d2

Backup policies

You define what you want e.g.

```
Sun -- full backup
Mon -- differential
Tue -- differential
Wed -- differential
Thu -- full backup
Fri -- differential
Sat -- differential
```

Note incremental backups are fine only if the product deals with those for you.

Infrastructure components

- Backup server
- Backup agents
- No agent and SAN/NAS
- No agent and VMM e.g. VM snapshots (& file-system freeze)

aka `Quiesce Guest File System` during a snapshot with ESXi or vCenter

Products?…

==> THE COMPETITION (AND WITH GUI)

- Window Server Backup
- Timeshift (works fine on Ubuntu – only local?)
- NetBackup
- HP DataProtector
- Acronis
- Veeam Backup (against VMware snapshots)

LAB industrialize Timeshift on Ubuntu workstations

==> OPEN SOURCE ASSETS

De-duplication capable? (LAB)

- Amanda
- Bacula
- Duplicity (librsync)

De-duplication *there is*

- Attic/**Borg** (recommended)
- bup
- Restic

DIY BACKUP

- Easy backup script in a daily cron job
- DIY scheduled FTP or RSYNC/SSH upload
- –or– DIY scheduled remote SSH call & retrieve

Does any argument remain against that method?

UNIX TIME

- UNIX Epoch time: 1 Jan 1970
- leap seconds ignored
- one day = 86 400 seconds

TIME-ZONE POLICY

- according to physical location
- –vs.– all servers around the world on the main time-zone?

What if you have a PRA or some CDN?…

==> Obviously you’re gonna use the same time-zone…

UNIX ASSETS

Repeated full backups

- `tar` – backups as archive tarballs (recommended)
- `dump` as traditionally defined in `fstab` // LAB try it out
- `cpio` e.g. `initramfs`
- `cp -a`
- Afio // LAB try it out

LAB benchmark a few of those DIY backup systems one against eacher other

Repeated differential synchronization but in the end there's only one backup

- `rsync`
- `psync` (parallelized rsync-like clone)

LAB is there a way to script a differential backup system?

REMOTE STORAGE vs. BACKUP MASTER

whatever packaged as a product or scripted

Initiated by nodes

- Send it in some node-based chroot service
  (FTP, RSYNC/SSH, RSYNCD?)

Initiated by backup master server

- Schedule the job from the backup server

What scenario is best?...

==> Let’s have a closer look.

Initiated by nodes - a closer look

Threats

- A compromized node could remove previous backups…
- Avoid your backups to be reached by an attacker

Mitigations are over-complicated and error-prone

- send it in some **node-based** (not customer account based) chroot
- –or– send it in some `+t` upload folder?_
- –or– play with folder & file restrictive umasks?
- –or– encryption at rest AND different symmetric key for every node?

Initiated by backup master server - a closer look

Threats

- Backup server compromized? Your company is dead.

Brutal mitigation

- The backup manager does the job

Resulting thread

- But beware, backup server has full access to your infrastructure…

LAB PoC your own scripted backup-server against a few nodes * well defined archive naming (node name, date, …) * also beware of the brutal max-age folder clean-up

Definitive answer is…

==> Backup master server w/o encryption at rest

- Backup server does not listen on the network
- Dedicated VLAN? Anyhow IDS/IPS should not get crazy with that
- (Distinguish data leaks from nightly backups)

File/stream compression

Data compressors

- `compress` / `uncompress` (BSD) - adaptive Lempel-Ziv coding
- `gzip` / `gunzip` (GNU) - DEFLATE algorithm (LZ77 + Huffman coding)
- `bzip2` / `bunzip2` - Burrows–Wheeler transform
- `xz` / `unxz` (LZMA) - Lempel–Ziv–Markov chain algorithm
- Super fast LZO & LZ4 (LZ77) lossless (low CPU usage)
- Fat-transfert-optimized RZIP (900MB de-duplicated chunks)

LAB what algos are WinRAR / UnZIP / 7-Zip archivers using?

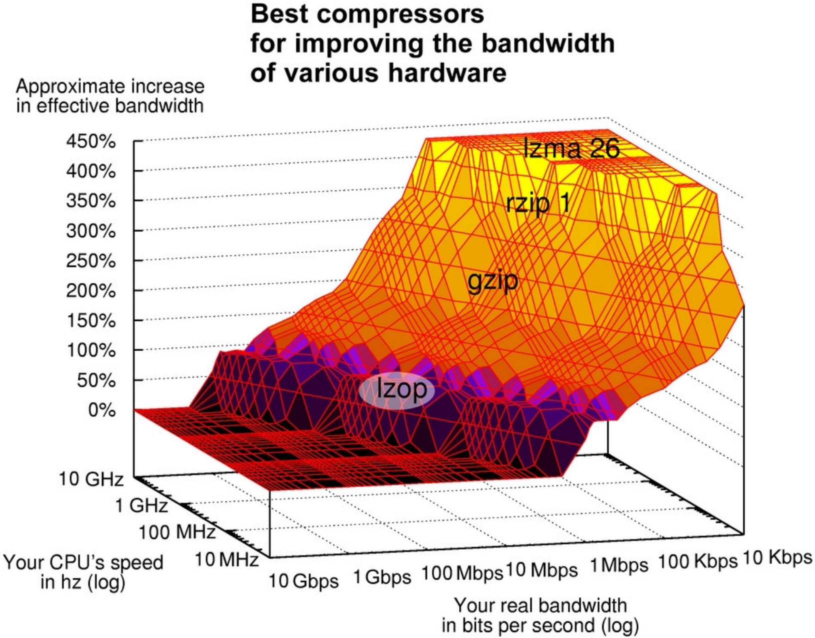

Best compressors
for improving the bandwidth
of various hardware

// linuxjournal.com

Keep file permissions

NOT needed for creating a software tarball

Create / extract an archive

```
tar -czpf
# -p, --preserve-permissions, --same-permissions


tar -xzpf
```

BONUS QUESTION does this apply to creation also or just extraction? (as manual implies)

WHAT IS A SPARSE-FILE?

- got that feature e.g. in EXT4
- a sequence of zeros got skipped at the file-system level

e.g. RAW virtual disks (if not block device nor QCOW2)

SPARSE-FILE-CAPABLE ARCHIVE

Handle those files properly. Got that feature with TAR also. Apparently the flag is only necessary while *creating* the archive.

```
tar czSf

# -S, --sparse

# --hole-detection=seek (default)

# --hole-detection=raw

# --sparse-version=1.0
```

RELATIVE PATH

```
cd /var/www/
tar -czpf html.tar.gz html/
```

vs.

```
tar -czpf html.tar.gz html/ -C /data/www/
```

Exclude from archive

e.g. loose videos and CGI chroot device files

```
tar -cJpf html.tar.xz \
    --exclude ”*.avi” \
    --exclude ”www/dev/*” \
    -C /var/www/ html/
```

Restore a TAR-based backup

extract the archive tarball in a specific folder

```
tar xzf /var/backup/html.tar.gz -C /var/tmp/
```

visualize the changes since then

```
diff -rbu /var/tmp/html/ /var/www/html/
```

and eventually rollback (rename and move at once)

```
mv /var/www/html/ /var/www/html.damn/ \
    && mv /var/tmp/html/ /var/www/html/
```

Mirror / synchronize it

Compression on transit only (not at rest)

- `rsync -z` (zlib ~ gzip)

RSYNC USAGE

Archive mode

```
rsync -avz --delete <source> <dest>
#-rlptgoD (no -H,-A,-X)
```

The trailing-slash

- The trailing-slash `/` is VERY important while defining the source directory
- Without it, it sends the full directory to the destination
- With it, it sends precisely its content to the destination

Both directions

- Both source and/or destination can be local or remote
- Choose your direction wisely

POSSIBLY REVERSED (BACKUP SRV IS THE MASTER)

```
rsync -avz --delete root@target:/var/backup/ \
    /path/to/backup/folder/
```

- Infrastructure architecture looks good
- But this is assuming the target server nodes already have local backups
- Which is still not ideal (local jobs have to be processed beforehand on the node)
- And this is just a mirror, far from being a backup policy

DATABASE BACKUP

Why not backup the folders directly?…

==> Database is mounted - it has its own storage format

- Oracle –> RMAN
- MySQL / MariaDB –> `mysqldumb`

BASH/KSH SCRIPTING

We want STDOUT and STDERR by email!

```
vi /etc/cron.daily/DAILY

#!/bin/bash
tar czpf /var/backup/`date +%s`.foldername.tar.gz html \
    -C /var/www/
#upload through lftp or rsync...

chmod +x /etc/cron.daily/DAILY
```

LAB // compare with manually defined cron job and where goes `stdout` vs `stderr`

So let us consider that we are doing full-backups every night at 01:00

…*Any problem with this plan? How to solve it?*

==> Would eat your local storage… There is a need to clean-up

```
find /var/backup/ -type f -maxdepth 1 -mtime +10 \
    -exec rm -f {} \;
    #xargs rm -f
```

- Note `-maxdepth 1` to wipe only files from that precise directory
- Sub-directory with older files will remain

```
         (____)
         ( O )
  /-------\ /
 / |     ||V
*  ||----||
   ^^    ^^
```

// Questions on backup and compression?

Migrations

Let’s say you’ve got obsolete servers in production, with various CPU architectures, and you want to consolidate it all.

How to proceed?…

==> Two kinds of migrations

the long hard road of app & db migration to change architecture

- DB / application / data migration (recommended)
- P2V & V2V
  - if it’s too much a mess to rebuild…
  - and when possible…

DB upgrades & application migrations

Most important is the database

- Oracle upgrades easier by means of export/import
- PostgreSQL migrations == backup by means of `pg_dumpall`
  - note there's also `pg_basebackup` for setting up replicas
- MariaDB migrations == backup by means of `mysqldump`

The rest of the app is usually static with eventually some upload folder for user files to synchronize. Same goes for Docker instances.

VMM migrations

Any idea what P2V and V2V means?…

==> Physical to Virtual

==> Virtual to Virtual

P2V & V2V products on-premises

- VMware Converter
- Novell Platespin

LAB // are there any other X2V products since then?

DIY P2V & V2V

- Convert the virtual disk
  - RAW vs QCOW2 vs VMDK vs …
- Configuration file can mostly be rebuild from scratch
- ESXi vs KVM vs PVHVM vs PV vs PVH?
  - devices and network devices may change

V2V « on the cloud »

- GCP[1]
- AWS[2]
- anything else (SCW, on-premises, ...)
  - DIY guest snapshot then rescue mode and DD the virtual disk over SSH...

[1] https://cloud.google.com/migrate/compute-engine/

[2] https://aws.amazon.com/cloudendure-migration/

// Questions on migrations?

Disaster Recovery

What's the difference with HA?...

==> It’s a process, or a very slow replication-based HA in best case scenario

==> It’s not a cluster, or at least not the same one

Principles

Supposedly on…

- different datacenters
- different IP range and back-bone
- but it can also name a slow and manual application HA system

Application DR - RS/6000 local DR example

some kind of slow and manual HA

- replication: an rsync script once an hour between two identical machines running AIX
- in case node A goes down –> manual trigger and node B takes over
- no cluster: nothing is shared, there's just a regular sync process
- ok for an application
- NOT ok for data, which needs to live elsewhere

VMware vSphere Replication

- copies vdisks (data protection)
- *can start the VM on the other side?*
- *probably NOT live migration capable*

VMware Site Recovery Manager

- integrates with storage solutions
  - vSphere Replication
  - vSphere Virtual Volumes (vVols)
  - third-parity vendors…
- configure a recovery plan (define policies)
- automate « the execution of the recovery »

Storage DR

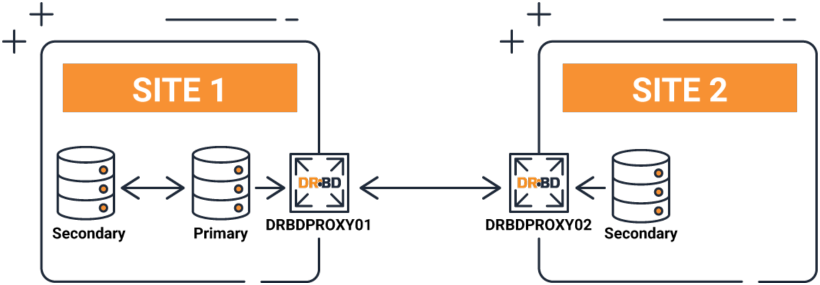

// linbit.com

// Questions on disaster recovery?

Incident Monitoring (status alerts)

```
\|/          (__)
     `\------(oo)
       ||    (__)
       ||w--||     \|/
   \|/
```

THE DASH-BOARD

- Big screen in operations room
    - large-scale hosting
    - IT outsourcing
    - any company with critical servers & services
- Viewing alerts live on dashboard
- Viewing alerts live on host/services view
- Getting alerts by email/SMS

Any monitoring products in mind?…

==> OPEN SOURCE ASSETS

- Nagios Core // LAB manage to setup performance graphs w/o XI
- Centreon (Nagios fork?)
- Munin
- Monit agent
  - sends alerts on its own
  - collects and sends data to M/Monit
- Zabbix
- Sentry // LAB

==> THE COMPETITION

- Nagios XI // LAB grab trial version
  - study and discuss the business model
  - and check if some parts closed-source
- M/Monit helps store data and visualize

Got more proprietary products on this front?

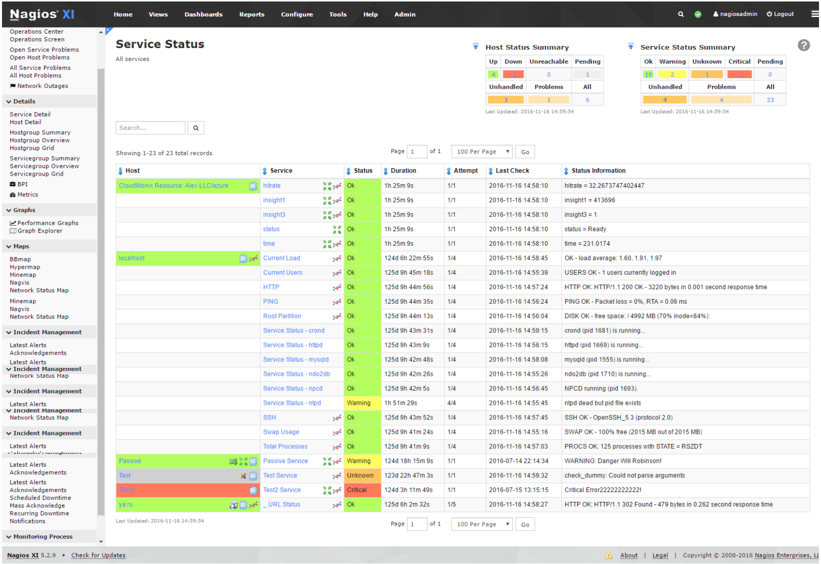

Service Status

All services

Host Status Summary

Up	Down	Unreachable	Pending
4	1	0	1

Unhandled	Problems	All
1	1	6

Last Updated: 2016-11-16 14:59:54

Service Status Summary

Ok	Warning	Unknown	Critical	Pending
19	2	1	1	0

Unhandled	Problems	All
4	4	23

Last Updated: 2016-11-16 14:59:54

Showing 1-23 of 23 total records

Page 1 of 1 100 Per Page

Host	Service	Status	Duration	Attempt	Last Check	Status Information
CloudMonix Resource: Alex LLC/azure	hitrate	Ok	1h 25m 9s	1/1	2016-11-16 14:58:10	hitrate = 32.2673747402447
	insight1	Ok	1h 25m 9s	1/1	2016-11-16 14:58:10	insight1 = 413696
	insight3	Ok	1h 25m 9s	1/1	2016-11-16 14:58:10	insight3 = 1
	status	Ok	1h 25m 9s	1/1	2016-11-16 14:58:10	status = Ready
	time	Ok	1h 25m 9s	1/1	2016-11-16 14:58:10	time = 231.0174
localhost	Current Load	Ok	124d 6h 22m 55s	1/4	2016-11-16 14:58:45	OK - load average: 1.60, 1.91, 1.97
	Current Users	Ok	125d 9h 45m 18s	1/4	2016-11-16 14:55:39	USERS OK - 1 users currently logged in
	HTTP	Ok	125d 9h 44m 56s	1/4	2016-11-16 14:57:24	HTTP OK: HTTP/1.1 200 OK - 3220 bytes in 0.001 second response time
	PING	Ok	125d 9h 44m 35s	1/4	2016-11-16 14:56:24	PING OK - Packet loss = 0%, RTA = 0.06 ms
	Root Partition	Ok	125d 9h 44m 13s	1/4	2016-11-16 14:56:04	DISK OK - free space: / 4992 MB (70% inode=84%):
	Service Status - crond	Ok	125d 9h 43m 31s	1/4	2016-11-16 14:59:15	crond (pid 1681) is running...
	Service Status - httpd	Ok	125d 9h 43m 9s	1/4	2016-11-16 14:56:15	httpd (pid 1669) is running...
	Service Status - mysqld	Ok	125d 9h 42m 48s	1/4	2016-11-16 14:58:08	mysqld (pid 1555) is running...
	Service Status - ndo2db	Ok	125d 9h 42m 26s	1/4	2016-11-16 14:55:26	ndo2db (pid 1710) is running...
	Service Status - npcd	Ok	125d 9h 42m 5s	1/4	2016-11-16 14:56:45	NPCD running (pid 1693).
	Service Status - ntpd	Warning	1h 51m 29s	4/4	2016-11-16 14:55:45	ntpd dead but pid file exists
	SSH	Ok	125d 9h 43m 52s	1/4	2016-11-16 14:57:45	SSH OK - OpenSSH_5.3 (protocol 2.0)
	Swap Usage	Ok	125d 9h 41m 24s	1/4	2016-11-16 14:55:16	SWAP OK - 100% free (2015 MB out of 2015 MB)
	Total Processes	Ok	125d 9h 41m 9s	1/4	2016-11-16 14:57:03	PROCS OK: 125 processes with STATE = RSZDT
Passive	Passive Service	Warning	124d 18h 15m 9s	1/1	2016-07-14 22:14:34	WARNING: Danger Will Robinson!
Test	Test Service	Unknown	123d 22h 47m 3s	1/1	2016-11-16 14:59:32	check_dummy: Could not parse arguments
Test2	Test2 Service	Critical	124d 3h 11m 49s	1/1	2016-07-15 13:15:15	Critical Error22222222222!
ya.ru	_ URL Status	Ok	125d 6h 2m 32s	1/5	2016-11-16 14:58:27	HTTP OK: HTTP/1.1 302 Found - 479 bytes in 0.262 second response time

Last Updated: 2016-11-16 14:59:54

Page 1 of 1 100 Per Page

Nagios XI host/services

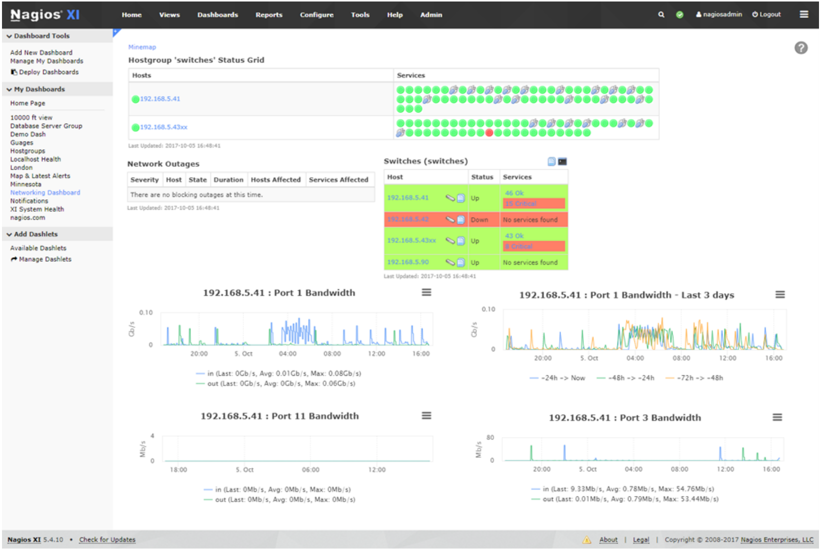

Dashboard Tools

- Add New Dashboard
- Manage My Dashboards
- Deploy Dashboards

My Dashboards

- Home Page
- 10000 ft view
- Database Server Group
- Demo Dash
- Guages
- Hostgroups
- Localhost Health
- London
- Map & Latest Alerts
- Minnesota
- Networking Dashboard
- Notifications
- XI System Health
- nagios.com

Add Dashlets

- Available Dashlets
- Manage Dashlets

Minemap

Hostgroup 'switches' Status Grid

Hosts	Services
192.168.5.41	
192.168.5.43xx	

Last Updated: 2017-10-05 16:48:41

Network Outages

Severity	Host	State	Duration	Hosts Affected	Services Affected
There are no blocking outages at this time.					

Last Updated: 2017-10-05 16:48:41

Switches (switches)

Host	Status	Services
192.168.5.41	Up	46 Ok 15 Critical
192.168.5.42	Down	No services found
192.168.5.43xx	Up	43 Ok 8 Critical
192.168.5.90	Up	No services found

Last Updated: 2017-10-05 16:48:41

192.168.5.41 : Port 1 Bandwidth

in (Last: 0Gb/s, Avg: 0.01Gb/s, Max: 0.08Gb/s)
out (Last: 0Gb/s, Avg: 0Gb/s, Max: 0.06Gb/s)

192.168.5.41 : Port 1 Bandwidth - Last 3 days

-24h -> Now   -48h -> -24h   -72h -> -48h

192.168.5.41 : Port 11 Bandwidth

in (Last: 0Mb/s, Avg: 0Mb/s, Max: 0Mb/s)
out (Last: 0Mb/s, Avg: 0Mb/s, Max: 0Mb/s)

192.168.5.41 : Port 3 Bandwidth

in (Last: 9.33Mb/s, Avg: 0.78Mb/s, Max: 54.76Mb/s)
out (Last: 0.01Mb/s, Avg: 0.79Mb/s, Max: 53.44Mb/s)

Nagios XI 5.4.10 • Check for Updates

About | Legal | Copyright © 2008-2017 Nagios Enterprises, LLC

Nagios XI hostgroups

FROM-THE-DIY-DEPT

- DIY alerting with ClusterIt as cron jobs
- Jobs can be scheduled from the backup server (which may have all the necessary SSH accesses already)

TYPES OF CHECKS

- Remote/network checks & metrics
- Local/agent checks & metrics
- Hypervisor/host metrics
- SNMP

REMOTE ALERTS

Viewing and receiving alerts on…

- Host absence (no ping response)
- Services down
- Services too slow
- Web pages down
- Web pages too slow

SYSTEM/VMM & BMC ALERTS

Viewing and receiving alerts on ***status & thresholds***

SYSTEM/VMM

- RAID *optimal*
- NIC negociated speed e.g. `1000baseT-FD`
- LACP…
- File-system usage e.g. close to 90%

SYSTEM/VMM or BMC

- Temperature
- Fan status and RPM

BMC-only

- Energy-waste (Watt / Voltages)

Viewing and receiving alerts on ***timed thresholds***

VMM performance bottlenecks

- Constant CPU 100%
- Constant RAM 100%
- Constant DISK I/O 100%
- Bandwidth usage
    - Per network link RX 100% during 15 minutes…
    - Per network link TX 100% during 2,5 hours…

About network *TX* overload, that should rather be for IDS/IPS data leak prevention.

SNMP ALERTS

covered by another lecture: SNE/NETWORK/SNMP

```
\|/          (__)
     `\------(oo)
       ||    (__)
       ||w--||     \|/
   \|/
```

// Questions on incident monitoring?

OUTGOING EMAIL

- considering a DIY backup server
- –or– considering a DIY monitoring station
- –or– any other kind of *script-in-a-cron-job* output

Where does its `stdout` *and* `stderr` *go locally?...*

==> `/var/mail/USER` (BSD)

==> `/var/mail/spool/USER` (GNU)

How to read those email stored locally?…

==>

```
cat
less
mail
alpine
mutt
```

Otherwise how to get the alerts posted to a real email address?…

==> EMAIL ALIASES

GNU/Linux

```
vi /etc/aliases


root:   TRUE-EMAIL@example.net


newaliases
ls -lF /etc/aliases.db
```

BSD & Sendmail

```
vi /etc/mail/aliases
(idem)
ls -lF /etc/mail.aliases.db
```

Will the server be able to send to `TRUE-EMAIL@example.net`?...

SMTP CLIENT vs. OUTBOUND MTA

assuming a server

- got an smtp relay on the internal network
- –or– authenticate through SASL
- –or– outbound MTA with a public IP
- –or– PTR and SPF trickery behind a NAT

SMTP RELAY

aka smarthost

e.g. with Postfix

```
vi /etc/postfix/main.cf

relayhost = 10.1.1.253
    smtpd_tls_security_level = encrypt
smtp_tls_security_level = encrypt

postfix reload
```

e.g. with DragonFlyBSD Mail Agent (DMA)

```
hostname --long
vi /etc/dma/dma.conf

    MAILNAME FQDN-WITH-VALID-PTR-HERE
    SECURETRANSFER
    STARTTLS
```

PUBLIC IP OR NAT

Need good

- PTR & SPF DNS records
- EHLO
- MAIL FROM (sender)

What happens if you're behind a NAT?…

==> Use the PTR of the gateway as hostname for the MTA.

How to test that outgoing email works anyhow?…

THIS IS AN ACCEPTANCE TESTING EXAMPLE

```
apt install bsd-mailx
#apt install mailutils
#apt install s-nail

#tail -F /var/log/maillog
tail -F /var/log/mail.log

date | mailx -s `uname -n` root
```

What network specifications are required to deploy a host system?…

IP, netmask …?

==> FULL NETWORK SPECS FOR A NEW HOST

- HOSTNAME & DOMAIN
- IP/NETMASK/GATEWAY
- DNS validating resolver (DNSSEC)
- SMTP RELAY
- NTP (usually same as DNS or domain controller)
- (SNMP community and trap destination)
- (MONITORING SRV)

Don’t forget to re-generate SSH host keys in case you’re dealing with guest templates.

More for workstations

- HTTP_PROXY
- – and/or – push a CA and a client cert in there

BACKUP TIPS & TRICKS

PACK IT UP

```
tar -cZf archive.tar.Z archive/
tar -czf archive.tar.gz archive/
tar -cjf archive.tar.bz2 archive/
tar -cJf archive.tar.xz archive/
tar -I lz4 -cpf archive.tar.lz4 archive/
```

LAB benchmark, compare and discuss speed/compression ratios

TIMEZONE APPLIED

check configured time-zone

```
ls -lF /etc/localtime
```

UNIX TIME APPLIED

quick and dirty timestamp

```
date +%s
```

force UTC (and leap seconds?)

```
date -u +%s
```

as of Feb 2021 there’s three second gap

```
1613451956
1613451953
```

INCIDENT MONITORING TIPS & TRICKS

How to check file-system usage manually?…

==> File-system usage

```
slack2# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/root        66G   54G  9.5G  85% /
devtmpfs         62G     0   62G   0% /dev
tmpfs            62G  900K   62G   1% /run
tmpfs            62G     0   62G   0% /dev/shm
cgroup_root      62G     0   62G   0% /sys/fs/cgroup
/dev/sdb1       299G   94G  191G  33% /data
cgmfs           100K     0  100K   0% /run/cgmanager/fs
```

try to standardize things across platforms

```
   -P, --portability
          use the POSIX output format
```

```
slack2# df -P
Filesystem    1024-blocks    Used Available Capacity Mounted on
/dev/root         69075456 55685528   9858000      85% /
devtmpfs          64948268        0  64948268       0% /dev
tmpfs             64951772      900  64950872       1% /run
tmpfs             64951772       0  64951772       0% /dev/shm
cgroup_root     64951772      0 64951772     0% /sys/fs/cgroup
/dev/sdb1        313296192 97542500 200061100      33% /data
cgmfs              100      0     100      0% /run/cgmanager/fs
```

And what about shells?…

==> KSH93 & BASH are pretty much compatible and offer loads of scripting features beyond POSIX

DIY alerting - File-system usage

Prints output only if there is a problem…

```
vi /root/report/diskusage.bash

#!/bin/bash

tmp=`df -P | sed 1d | grep -vE '^udev|tmpfs|^cgroup|^rpool/ROOT/'`

echo ”$tmp” | while read line; do
        percent=`echo $line | awk '{print $5}' | sed 's/%//'`
        (( percent > 89 )) && echo $line
        unset percent
done; unset line

chmod +x /root/report/diskusage.bash
```

May be executed in a loop for live display -or- put it in a cron job

```
crontab -e
```

```
*/5 * * * * /usr/pkg/bin/dsh -e -g linux -s /root/report/diskusage.bash
```