Blue Teaming

Offensive Technologies

Revision 4 (2025/26)

Pierre-Philipp Braun <pbraun@nethence.com>

Table of contents

- Infrastructure Hardening
- Intrusion Detection Systems
- Bayesian Inference
- AI-assisted Heuristics
- Honey Potting

Infrastructure Hardening

know your infrastructure

Blue loves red

- what you *think* your infrastructure is
- `!=` What *is it in fact* and this is how the attacker will get to know it

==> *pentesting MUST occur anyhow* so you get to know about your own infrastructure

Game is on

who's gonna win?...

==> who ever understand the infra at best

- hardware – servers **and devices**
- services / apps
- systems
- network architecture & routing

how to make sure it will be us defending, not them?...

Security basics

seen in SNA/SECURITY lecture – in a nutshell:

- keep your system up-to-date
- beware of what services are listening
- no weak passwords

what else as for engineers?...

Scan yourself

- exhaustive **host n services discovery**
- ideally distributed e.g. with Scan My Ass https://codeberg.org/elge/sma/

public addresses

- your cloud and datacenter IPs
- your various office locations and warehouses
- your VPN users…

possible locations to distribute/scan from:

==> need remote hosts beyond your cloud…

internal addresses

possible locations to distribute/scan from:

- the backup system
- the monitoring system
- the IDS
- the SIEM

Usability vs. security

- make the users happy e.g. deploy ssl certs for them (in case of private CA)
- make the devs happy e.g. give them access to all logs (eventually obfuscate tokens)
- make everybody happy e.g. nice SSO authentication model (however need to maintain and clean-up - check with HR)

Performance vs. security

you may prefer performance for those

- grid computing, mining, …
- storage clusters (what do when it’s convergent however?)
- ilsolated network segments (w/o inbound ports)
- trusted env like your own datacenter cage

you may prefer security for those

- DMZ & front-facing HA clusters
- workstations & large user VLANs with 1000+ users
- untrusted env like foreign clouds (and eventually even local) clouds…

improved performance

as long as attackers cannot execute anything remotely

eventually disable Spectre, Meltdown & friends mitigations as Linux kernel argument[1]

```
noibrs noibpb nopti nospectre_v2 nospectre_v1 l1tf=off
    nospec_store_bypass_disable no_stf_barrier mds=off
    tsx=on tsx_async_abort=off mitigations=off
```

or just

```
mitigations=off
```

- NOT for hosting (that goes through virtualization)
- NOT for workstations (javascript does execute some stuff)

[1] Make Linux Fast Again https://make-linux-fast-again.com/

System & network best practices

consistent infrastructures

Customize all layers (hardware, system, …)

- remember previous mention of *An architecture a day…*?

Fine-tune all things (daemons, alerts, …)

Upgrade all firmwares

- preferably open-source
- or from a specific vendor you are in business with

Customize bare-metal & firmwares

- enterprise-class
- low-cost clustered bare-metal
- what CPU exactly, what micro-code version?
- what firmwares exactly?
- what chips and features are in there – Intel ME[2]?

LAB // more on micro-code versions upgrades, and what distro package has those

[2] me_cleaner https://github.com/corna/me_cleaner

Abstraction layers

too many layers to upgrade

- bare-metal vs. virtualized vs. containers
- hardware abstraction is cool
- use containers only if you’re ready to upgrade images and restart instances…

Customize systems

kiss and keep control

- *aka sysprep, post-install, system tuning, customization, optimization (can be automated)*
- linux vs. BSD vs. exotic
- consider kernel and userland
- what libc is in there – wanna try musl?
- what booting process, partition table and volume manager is there?

(more on BSD systems in SNE/ES/OS)

Fine-tune daemons

optimize what's listening and rest in peace…

ssh daemon hardening

assuming public network

Truly useful

- ip4 vs. ip6 & what interface to bind to?
- specific user group – `AllowGroups wheel` / `root`
- –or– specific users – `AllowUsers root user1`
- –and/or– specific IP ranges – `AllowUsers root@CLIENT-IP gollum@CLIENT2 *@CIDR`
- no passwords, never ever
- host key ED25519

and just to read the logs in peace…

- alternate SSH port, ideally NOT top 1000

ssh client usage

- passphrases are still recommended
- ssh-agent is fine

Auditing tools

rootkit detection + hardening helpers

```
lynis       -- reports on system configs
rkhunter    -- search rootkit
chkrootkit  -- search rootkit
#tiger      -- brute force?
```

colorful log reading for deep-dive RCA & forensics

```
lnav
```

Monitoring dashboards

- CPU RAM DISKIO TX/RX –> check for DoS attacks against resources
- mount point space usage –> *idem* (log flood or thin-provisioning saturation)
- network TX/RX –> exflitration alert

and know what is ***considered normal*** *using heuristics (incl. your logs)*

System reports

- systems talk esp. BSD – setup outgoing email
- will tell when ever a config file changes
- will tell whenever an automated update failed

Will tell whatever you ask for (see tips & tricks)…

// Questions on infrastructure hardening?

host-based log-aware IPS

- mandatory for public IPs
- mandatory for internal network…

dedicated daemon reads logs and deals with system-firewall

- sshguard (not only ssh…)
- denyhosts
- fail2ban

the right way (no additional daemon)

- blacklistd netbsd-only

Intrusion Detection Systems

IDS / IPS

- **D**etection (passive) – *just an alert*
- **P**revention (active) – *alert* + *blocked*
- *ideally hybrid with manual validation (enable active) for next time it happens // LAB*

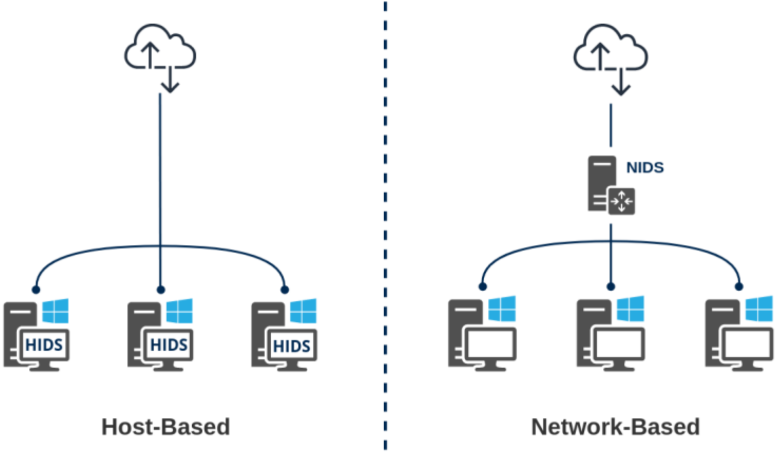
NIDS
HIDS
HIDS
HIDS
Host-Based
Network-Based

Note: corporate HIDS fleet are generally managed by a centralized mgmt interface e.g. with Kaspersky

Host vs. network based

Host IDS (easier IPS)

- extra host application
- works locally
- **active, blocks known attacks**

Network IDS (harder IPS)

- sniffing the network
- **usually passive**

Host IDS products

- OSSEC
- Prelude SIEM vs. OSS (hybrid host/network)
- Modern anti-virus (*not* Clam-AV)

Network IDS products

- Suricata (got IPS/online mode)
- (OPNsense incl. Suricata)
- Snort
- Zeek (formerly Bro)
- Prelude SIEM vs. Prelude OSS (community edition)

…some of those are essentially generate logs and alerts (need some UI?)

LAB // setup some Dashboard against a FOSS IDS

Intrusion detection

Static rule-sets

- signature-based – *character strings, binary size/checksum*
- known app-level exploit attempts – *just to catch those who try*
- protocol exploit attempts – *unprobable values in headers, known protocol attacks*
- stateful protocol analysis – *keep track of some connections*

Hint: enable as much community rules as possible – and eventually pay for a few more

LAB // check the rules and validate a stateful detection

Inference & heuristics

- anomaly-based – *classify what is normal or not*

BONUS QUESTION vs. LAB // can an IDS detect MITM attacks? e.g. does it check SSL certificate chains? e.g. does it check for SMTP downgrade attacks?

Network architecture - IDS locations

What's your target traffic?

Where to put those in a network architecture?...

monitoring floors vs. routing vs. internet

Possible locations for an IDS

- switch uplink – *attempt to catch lateral movements*
- internal router / ACL – *between network segments & VLANs*
- public gateway / firewall – *internet traffic (most important)*

==> IDS goes there – *passive only*

- port mirror *aka* SPAN/RSPAN on a switch – *monitoring the uplink*
- port mirror on an internal router? – *otherwise embedded*
- port mirror on a router/gateway? – *otherwise embedded*

BONUS QUESTION // does some Cisco switch with embedded IDS exist?

==> IPS goes there – *need to be on the path*

only routers & firewalls

- *not* on a switch – *unless the feature exists nowadays?*
- internal router / ACL
- public gateway / firewall

Consider gateway monitoring e.g. NAT

- remember we monitor only the traffic going through us

Will malicious traffic between node1 talking to node2 in the company be catched?…

==> of course not, it’s not passing through the gateway

Now consider internal router monitoring

Will malicious activity within a VLAN be catched?…

==> nope, it would have to cross a VLAN to another

Now consider switch uplink monitoring

Will malicious activity between floor-neighbors be catched?…

==> also not, unless they’re connected to different switches (and without stacking)

Detection use-cases

During exploit attacks

- external attacker & DDoS

Persistent malware & covert channels

- malicious insider maintains access + evades network
- malware is a rootkit/backdoor

Other kinds of covert channels

- consultant (or spy) reaches his internal-network station from home
- *seen in NETWORK/VPN and OT/COVERT –> Other Tunnels*

Worms

- malware spreads around

LAB // check the ruleset against some known *backdoor* and validate its detection (warning: isolate env in case this is the unmodified malware/worm)

Evasion techniques

- obfuscation vs. encryption
- self-modifying & polymorphic malware – *not sure there are mitigations for this*

Obfuscation e.g. tricky URLs

```
http://victim/cgi/../../winnt/system32/cmd.exe?/c+dir+
    c:\/ (root) ./ (current dir) ../ (parent of current dir)

http://victim/cgi/%252E%252E%252F%252E%252E%252Fwinnt/system32/
    cmd.exe?/c+dir+c:\

http%3A%2F%2Fvictim%2Fcgi%2F..%2F..%2Fwinnt%2Fsystem32%2F
    cmd.exe%3F%2Fc%2Bdir%2Bc%3A%5C
```

LAB // test this against popular IDSen

Obfuscation e.g. XOR, encoders, crypters, packers

Some packers get noticed

LAB // check the ruleset against some known *packer* and validate its detection (warning: isolate env in case this is the unmodified malware/worm)

Encryption

- Easy evasion by means of encryption
- E.g. SSL is authenticated AND end-to-end

How to work around that situation?...

IDS with SSL interception

assuming public gateway

- SSL covert channels cannot be easily identified
- the only way is to terminate the SSL tunnels

==> block anything encrypted and proxy/intercept SSL

- the only way to the public network goes through (transparent) proxy

LAB // plug the IDS to an SSL interception engine

How to SSL intercept?...

==> your PKIX CA in da place

Clients need to trust Gateway/IPS's CA which signs-on-the-fly

- deploy CA certificate in user's browsers & systems trust stores

Any idea why "on-the-fly"?...

==> remember an SSL certificate is a binding between a CN/SAN and a key pair

Hence we need to generate-and-sign certs for every requested domain

```
www.google.com
somethingelse.fr
```

E.g. some user wants to reach `www.google.com`

```
CLIENT
--> asks for www.google.com

MITM SSL PROXY intercepts (either as defined or transparent)
--> CA creates and signs www.google.com.crt

CLIENT
--> verifies the chain of trust against its CA store

MITM SSL PROXY
--> forward-proxy delivers and relays traffic
```

Idem for `somethingelse.fr`, etc.

Note: forward proxy products (& interception feature) are discussed in NETWORK/LBS-PROXY

Technology Intelligence LAB // are on-the-fly certs cached and how?

Note: there’s also a blacklist feature e.g. with Squid+SquidGuard, to simply block a few websites…

Back to casual IDS (w/o SSL interception)

- we don’t get into the encrypted channels
- but still, we’re sniffing, and we got a lot…

What data can we collect anyways?…

==> play Eve… *(Eavesdropping)*

1. look at the **plain-text** (first rule from some agency…)
2. network-level **meta-data** (supposedly Zeek is good for that) // LAB
3. app-level **meta-data** (cross-join with GAFAM utmost-plausible…)

LAB // possible to grab meta-data with Zeek against a specific target node?

Network-level meta-data

- *IP accounting, NetFlow and IPFIX are done elsewhere*
- Server Name Indication (SNI) is in clear-text
- various things e.g. who's doing SSH, …

LAB // how to combine and merge IP accounting with upper-layer meta-data from the IDS?

LAB // otherwise simply do IP accounting from the IDS itself (e.g. Zeek has that feature)

App-level meta-data

- geo-localization (GPS vs. Wifi…)
- contact lists
- *see CCF/SURVEILLANCE*

PROJECT // how to combine and merge network-level with app-level meta-data?

Traffic inspection limitations

- no way to differnciate an HTTPS connection from VPN/SSL

Any idea why?...

==> OSI layer 6 – Presentation

- that’s a tunnel anyhow
- we don’t see what’s inside

IDS architectures

let's make a cluster!

- analysis unit (A) – *the main IDS instance*
- monitor (M) – *like an snmp agent…*

LAB // can we do that with Suricata and friends? Try to setup a separate Monitor from the main FOSS IDS instance e.g. with Suricata

- for sure we've got something similar with Zeek (cluster-capable)
- also we can do RSPAN with Cisco (remote SPAN)

LAB // cluster with Suricata possible?

Collaborative Intrusion Detection (CIDS)

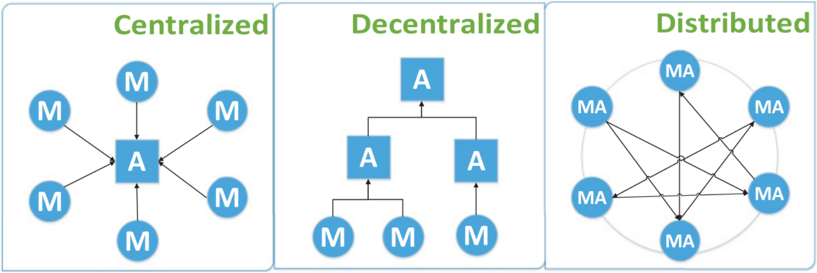

BONUS QUESTION // anything concrete on that front? what products?

IDS tuning

- community and commercial rules are not enough
- profile for company A doesn’t fit company B
- daily tuning is required
- many false-positives by design (alert doesn’t mean unsafe)
- unknown amount of false-negatives (**no alert doesn’t mean safe**) **–** we are not aware of all the bad things in the world

Rules' limitations – detects well-known attacks

- will not detect targetted and specific attacks
- unless you **create specific**, highly effective **detection rules**

Loads of tuning… – *about 3 monthes part-time for the auditor to fine-tune*

- need to update corporate security policy
- what is NOT allowed? DropBox, ToR, Torrent, Facebook, …
- so you can get rid of false-positives

also layer 3-4 tuning…

ACL & firewall vs. fine-grained layer 3-4 IDS rules

- some kind of a passive firewall
- some kind of a firewall honey-pot

Disowning the IDS

tips & tricks for the attacker

- forging fake alerts and confuse the auditor
- exclude IDS's IP from network scans

e.g.

```
   nmap -T5 --exclude 10.1.1.253 ...
masscan -T5 --exclude 10.1.1.253 ...
```

Owning the IDS

Snort 2 DCE/RPC Preprocessor Buffer Overflow

Snort 2.6.1, 2.7 Beta 1

SourceFire IDS 4.1, 4.5 and 4.6

```
msf > use exploit/multi/ids/snort_dce_rpc
```

What will attacker do once the IDS is compromized?...

==>

- IDS becomes a stepping stone for lateral movements
- IDS becomes a malicious monitoring point

Attacker covers his tracks

```
if(ip.source == attacker) drop alert
```

// Questions on IDS / IPS?

Bayesian Inference

Bayes-powered anti-spam story

- A Bayesian Approach to Filtering Junk E-mail (Jul 1998)
- A Plan for Spam (Aug 2002)
- Spam Detection (Sep 2002)
- Better Bayesian Filtering (Jan 2003)
- A Statistical Approach to the Spam Problem (Mar 2003)

Bayes-powered anti-spam products

- Bayesian Mail Filter (BMF)
- Bogofilter
- SpamAssassin (various ways of **scoring**)
- Quick Spam Filter (QSF)
- DSPAM, SpamProbe, ifile, CRM114, Annoyance Filter, SpamBayes

Classified as	NON-SPAM	SPAM
HAM	negative	**false-positive**
SPAM	**false-negative**	positive

Mail is NON-SPAM

- seen as HAM – negative – *everything went fine*
- seen as SPAM – **false-positive** – *filter is too agressive*

Mail is SPAM

- seen as HAM – **false-negative** – *filter is too lazy*
- seen as SPAM – positive – *everything went fine*

How to differenciate spam from ham?…

==>

- white-list of words
- black-list of words
- e.g. implemented with BMF against two Sleepy Cat Berkely `.db` files

Conditional probability

aka statistical inference

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}$$

`A` and `B` are events / conditions

`P(A|B)` is probability observing `A` given `B`

`P(B|A)` is probability observing `B` given `A` *applied to SPAM evaluation*

$$P(H|E) = \frac{P(E|H)P(H)}{P(E)}$$

`A` becomes `H` – the *hypothesis*

`B` becomes `E` – the *evidence*

`P(H)` becomes the *prior probability*

`P(H|E)` becomes the *posterior probability*

`P(E|H)` becomes the *likelihood*

LAB // what's `E`, simply the email body? and what's `H`, SPAM or HAM?

Bayes’ theorem

$$P(A \mid B) = \frac{P(B \mid A) \cdot P(A)}{\sum_{i=1}^{n} P(A_i) \cdot P(B \mid A_i)}$$

LAB // find what algo is in use with BMF, Bogo or QSF

Ok so we got Bayesian anti-spam – *which we don't really use it anyhow, as most of the job is done thanks to RBL and protocol checks*

What about classifying normal vs. anormal network activity?...

- *not sure it's in Suricata or any IDS*

What about analyzing logs?...

- *not sure there's a log server matching our requirement on heuristics*
- *see LIA/MONITORING –> logsrv*

LAB // does Graylog vs. ELK has the feature?

PROJECT // maybe DIY on top of sysklogd / rsyslog centralized logs?

// Questions on Bayesian inference?

AI-assisted Heuristics

- bayesian inference
- scoring

Can we imagine something better based on that?...

What modern tech do we have at hand? Any idea?...

==>

- artificial intelligence
- machine learning & training datasets
- deep structured learning & artificial neural network (ANN)

AI

- always over-estimated
- failure of « systèmes expert »

ML & datasets

anything familiar here?…

==> similar algorithms

we’ve seen

- (statistical inference)
- (Bayes’ theorem)

here comes

- Bayesian Detection Rate
- Classifier Adjusted Density Estimation (CADE)

Bayesian Detection Rate

$$P(I \mid A) = \frac{P(I) \cdot P(A \mid I)}{P(I) \cdot P(A \mid I) + P(\neg I) \cdot P(A \mid \neg I)}$$

// VILHELM GUSTAVSSON, KTH Royal Institute of Technology

Friedland, Gentzel, and Jensen (SDM 2014)

- refers to Hastie et al.
- Classifier Adjusted Density Estimation
- (CADE) approach for outlier detection

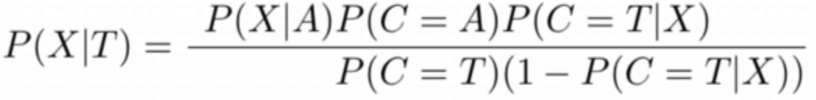

$$P(X|T) = \frac{P(X|A)P(C = A)P(C = T|X)}{P(C = T)(1 - P(C = T|X))}$$

Friedland et al. / Hastie et al.

DSL

- layers of connected nodes
- try to mimic human brain

neural network types

```
Artificial Neural Network (ANN)
Convolutional Neural Network (CNN)
Recurrent Neural Networks (RNN)
```

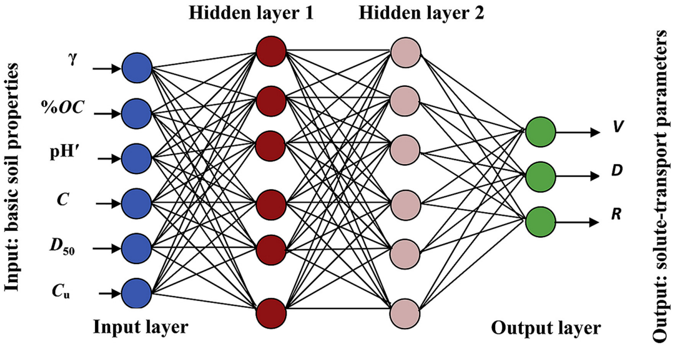

// M.A.Mojida et al.

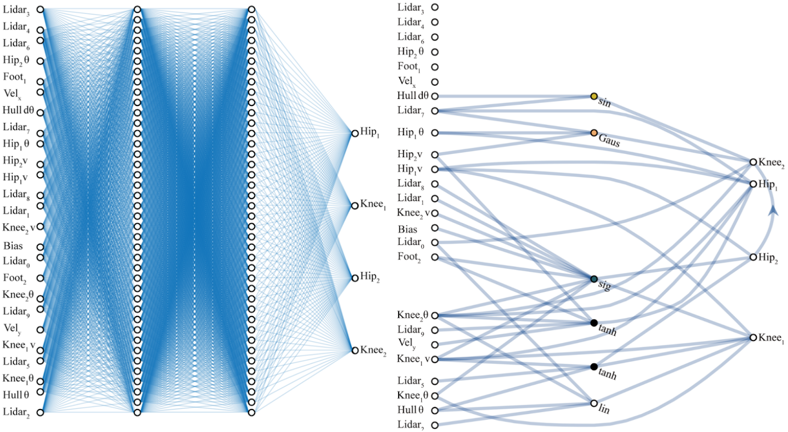

// ai.googleblog.com

// Questions on ai-assisted heuristics?

HONEY POTTING

What is a honey pot?…

==> Active prevention

- Looks like an interesting real system
- Contains fake data
- Interactive: trick the attacker
- Access to honey pot is always suspect

*So what is **the purpose** of a Honeypot?…*

==> Identify attacks and attackers

- A sensor where nothing should happen
- No noise hence see who’s there more easily
- Trace attacker’s activity (see what he’s looking for)
- Eventually a full network, infrastructure with data and traffic

What if an attacker leverages your Honeypot to reach back to your network?...

==> Not a new attack vector (avoid shooting yourself in the foot)

- The pot network should be isolated from the true network
- While being reachable from the true DMZ and/or true internal network
- …otherwise it is just gift as an attack vector and becomes a pivoting end-point

Attacks to identify

- Service/Network honeypot
- Spam honeypot (open relay on purpose to catch spammers –> RBL)
- Malware honeypot (vunlerable APIs)
- Database honeypot (audit & learn SQL injections)
- Spider honeypot (detect web crawlers and advertising networks)

Honeypot products

- SSH/ Cowrie, Kippo
- HTTP/ Glastopf, Nodepot
- Wordpress/ Formidable Honeypot, Blackhole for Bad Bots, Wordpot
- DB/ MongoDB-HoneyProxy, ElasticHoney, HoneyMysql
- Email/ Honeymail, Mailoney, SpamHAT
- Directories/ DCEPT, Canarytokens
- WebAppSec/ OWASP Honeypot
- Other/ HoneyNTP, Honeypot-ftp, Miniprint
- …

All-in-one honeypot products

- Honeydrive, MHN, Labrea
- Dionaea + LibEmu, Honeyd, T-Pot

LAB // Labrea

DIY honeypot

some vulnerable service or box

- **don’t shoot yourself in the foot**
- e.g. containment and analysis using Cuckoo[3]
- preferably behind yet another *and* isolated NAT network (and beware of NAT pivot)
- be ready to read all the necessary logs and receive an alert *in time*

[3] <https://cuckoosandbox.org/>

Spam traps

same concept but applied to an unused email address…

- publish an unused email address here and there on some web pages
- ideally hidden for the normal/human user
- wait for emails to arrive…

// Questions on honeypotting?

Tips & Tricks

lynis reports

e.g. Ubuntu/bionic with NGINX and Jitsi Meet

```
lynis audit system
```

gives (shows up in red)

```
  - Installed compiler(s)                        [ FOUND ]
 - net.ipv4.conf.all.accept_redirects (exp: 0)       [ DIFFEREN
...
 - Checking nginx                                [ FOUND ]
     - Parsing configuration options
     - SSL configured                              [ YES ]
        - Insecure protocols found                 [ YES ]
  - Checking for empty ruleset                  [ WARNING ]
    - Postfix banner                           [ WARNING ]
- Accounts without password                     [ WARNING ]
  - Permissions for directory: /etc/sudoers.d        [ WARNING ]
```

cron jobs

GNU/Linux

```
/etc/cron.hourly/
/etc/cron.daily/
/etc/cron.weekly/
/etc/cron.monthly/
```

BSD

```
vi /etc/daily
vi /etc/weekly
vi /etc/monthly
```

Manually

```
crontab -e
```

```
15 3 * * * /root/DAILY 2>&1
```

```
vi /root/DAILY
```

Daily cron job tuning example

Useful behind a NAT

```
echo WHAT IS MY IP
echo
curl -s ifconfig.me; echo
echo
```

Useful for a standalone server with bad monitoring

```
echo SERVICE STATUS
echo
/root/STATUS
echo
```

GNU/Linux specific

```
echo Who\'s who
echo
w
echo

echo Top 10 processes
echo
LINES=17 top –b –n1 –w # top 10
echo

echo Process tree excl. kernel
echo
ps --pid 2 --ppid 2 --deselect ufww
echo

echo Listening services
echo
netstat –ltupe
```

BSD specific

```
echo Who\'s who
echo
w -w
echo

echo Processes
echo
top -b 10
ps auxww | sort # by user
echo

echo Listening services
echo
sockstat -4 -l
sockstat -6 -l
echo

echo Active connections
```

LAB HINTS

a few hints

IDS/IPS outcomes

- Performing attacks on real-life systems and applications
- Detection of the attacks you’ve performed (if possible)
- Prevention of the attacks you have performed
- Honneypot w/o compromising yourself (don’t offer a pivot)

IDS/IPS hint

An easy way to check if your IDS works, even before going for covert channel detection

```
curl http://testmyids.com/
curl -A BlackSun http://testmyids.com/noexist
```

Non RCE vulns

- Non-authenticated SSL MITM
- Attacker's private CA SSL MITM
- Curveball + SSL MITM
- Network pivoting / route fuzzing
- SSH MITM (not through spoofing, needs to be persistent)

Some tools

VPN pivoting

- SSH SOCKS / tunnel / reverse-tunnel
- ProxyChains

MITM

- HonSSH (SSH MITM)
- DIY Postfix (can give conf example)

R&D

- passive IDS but possibly active: manual validation, for next time it happens… (avoid false-positives)

LAB alternative

funky saucers

- find out which high-end Cisco-or-friends firmwares would do either IDS and/or SSL interception
- and run it with GNS3 or EVE-NG
- and validate IDS and/or SSL interception

// Questions on the lab assignment?

This is the end