CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

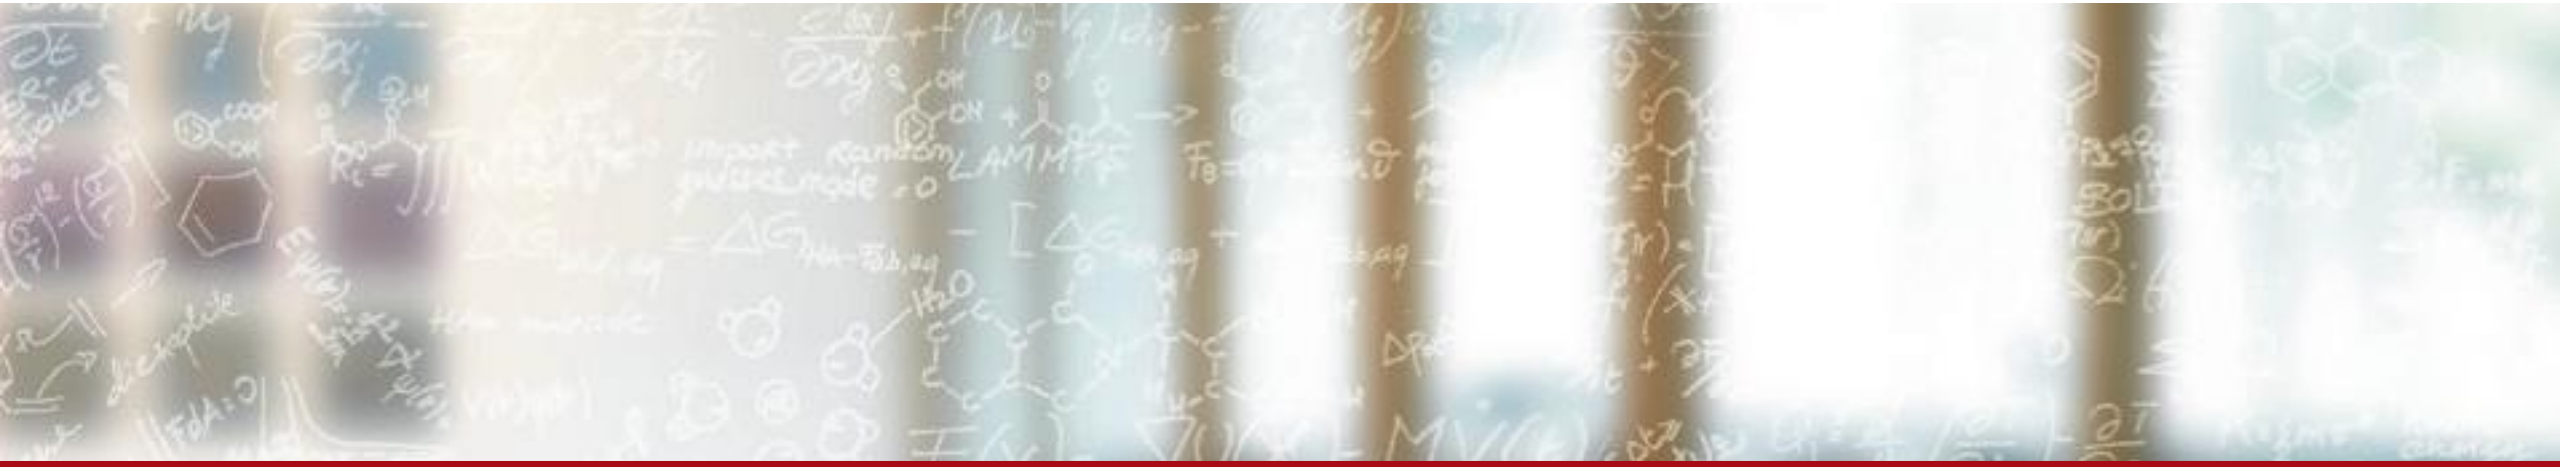

400 Gbps Observability

PacketFest '25

Fabio Zambrino, CSCS

May 8th, 2025

TPL: Green

Diverse Infrastructure

- The flagship system has more than 4'000 nodes
- vCluster technology
  - dedicated login and compute nodes
- Internet access
- Different software stack used by the users
- Multitude of technologies to support the HPC infra

How “traditional” Science is Done

The data sources

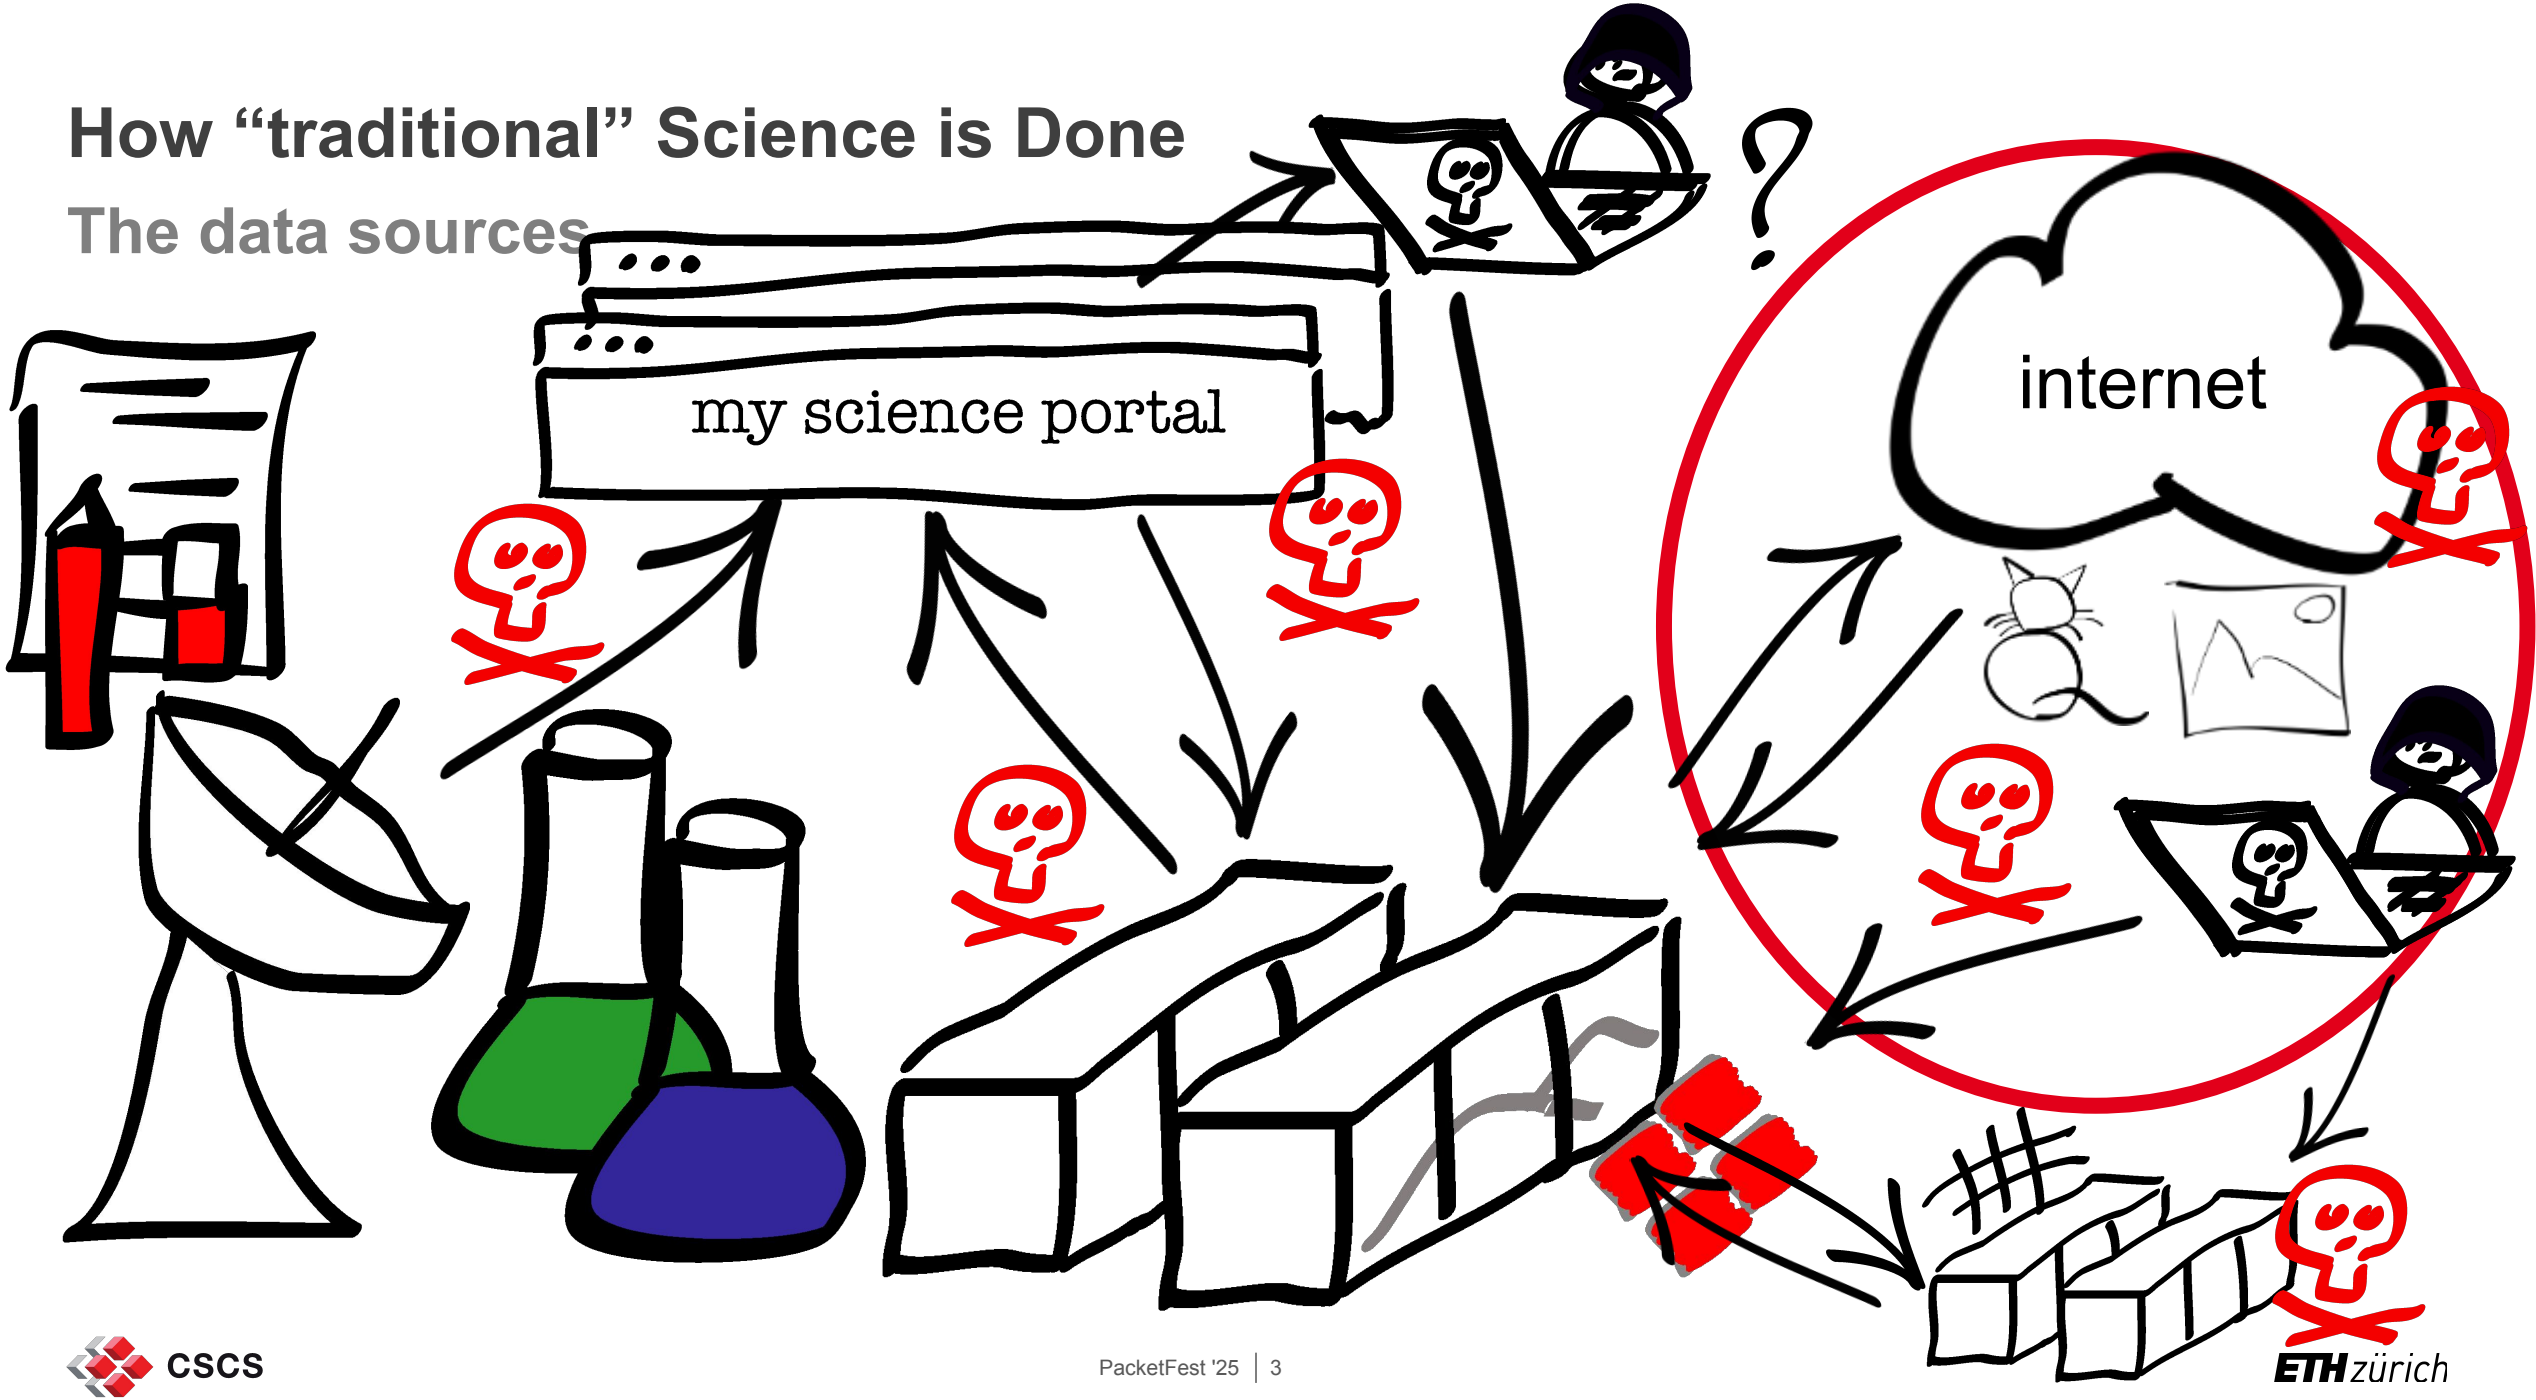

Challenges

- Lowest impact possible on performances

Challenges

- Lowest impact possible on performances

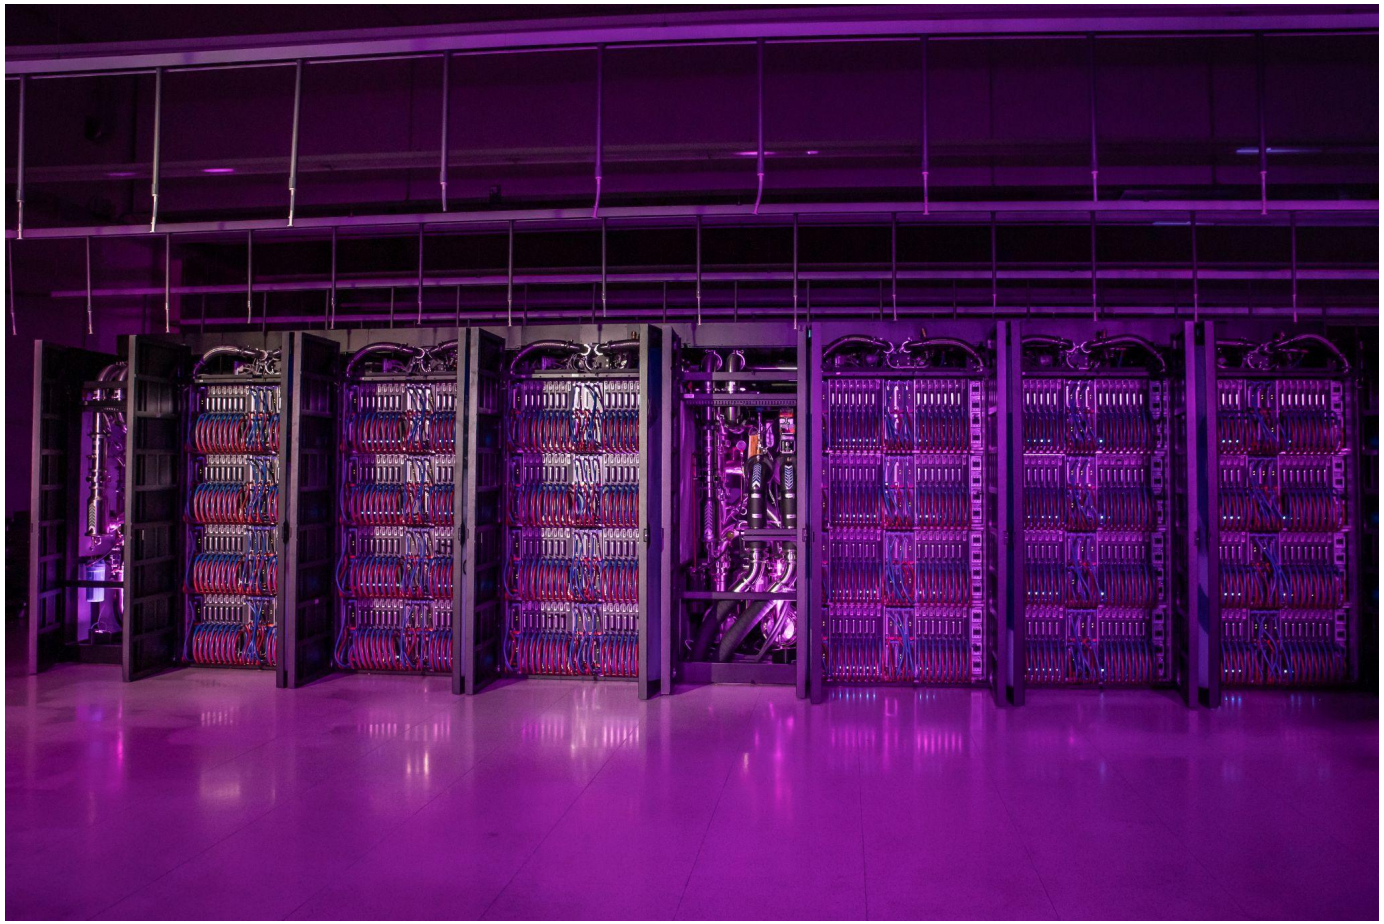

Challenges

- Lowest impact possible on performances
- Very broad scope to monitor

Challenges

- Lowest impact possible on performances
- Very broad scope to monitor

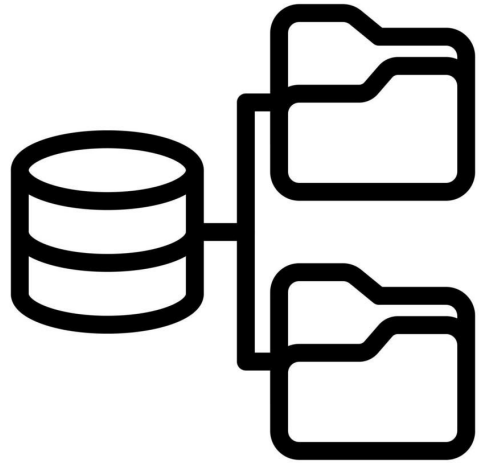

Knowledge Base

vmware®

Challenges

- Lowest impact possible on performances
- Very broad scope to monitor
- Supply chain attacks

Challenges

- Lowest impact possible on performances
- Very broad scope to monitor
- Supply chain attacks
- Data collection
  - Which data should we collect
  - How to collect the data
  - Where to centrally store the data
  - …

Challenges

- Lowest impact possible on performances
- Very broad scope to monitor
- Supply chain attacks
- Data collection
  - Which data should we collect
  - How to collect the data
  - Where to centrally store the data
  - …
- Data storage and retention
  - For how long should we store the collected data
  - e.g. Network traffic collected with ZEEK
    - raw text logs, compressed ~ 250GB/day
    - indexed logs in elasticsearch ~ 2 TB/day

Inc

Wh

Wh

- I
- H

Wh

- Increase in the phising campaings
- M
- I
- -

SWI swissinfo.ch The Swiss voice in the world since 1935

Login Search Menu

News >
Science >

Hackers target Zurich university with 'professional' cyberattack

Popular Stories

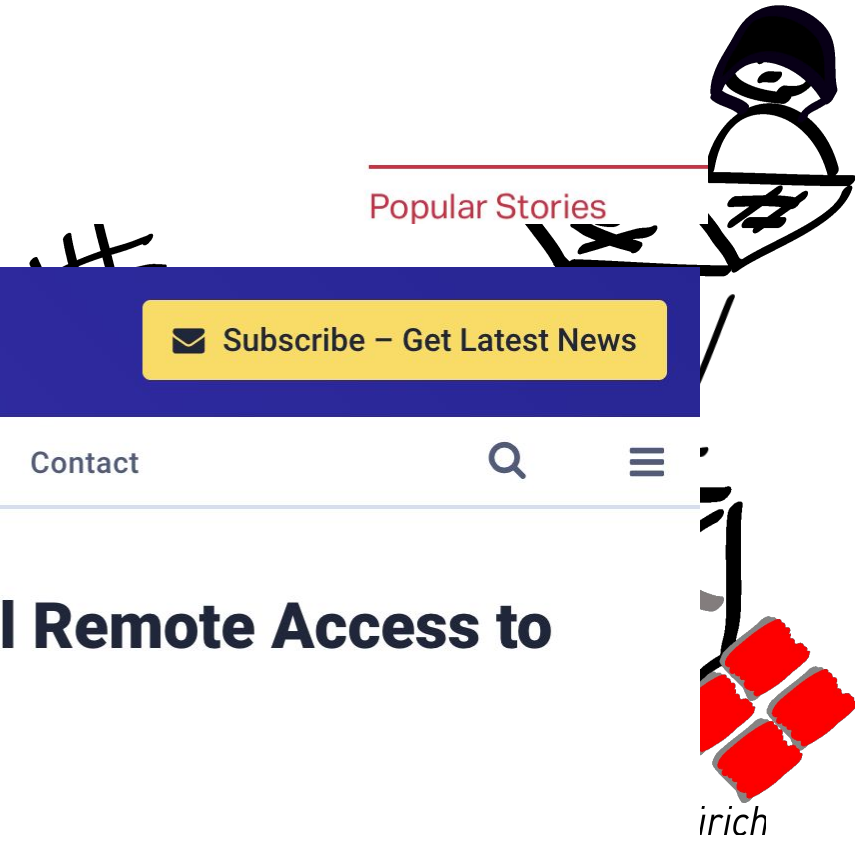

The Hacker News

Subscribe – Get Latest News

Home Data Breaches Cyber Attacks Vulnerabilities Webinars Expert Insights Contact

New Linux Malware 'Auto-Color' Grants Hackers Full Remote Access to Compromised Systems

Feb 26, 2025 Ravie Lakshmanan Linux / Endpoint Security

irich

Weaponizi

We ca

What'

- Dat
- The

What

- Use
- We
- Security
- AI workflc

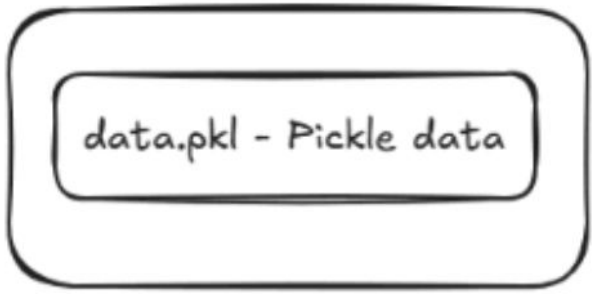

The Pickle format is well-known to be a **dangerous** serialization format, since in addition to serialized data, it may contain serialized code which will be automatically executed when the Pickled/Serialized file is loaded.

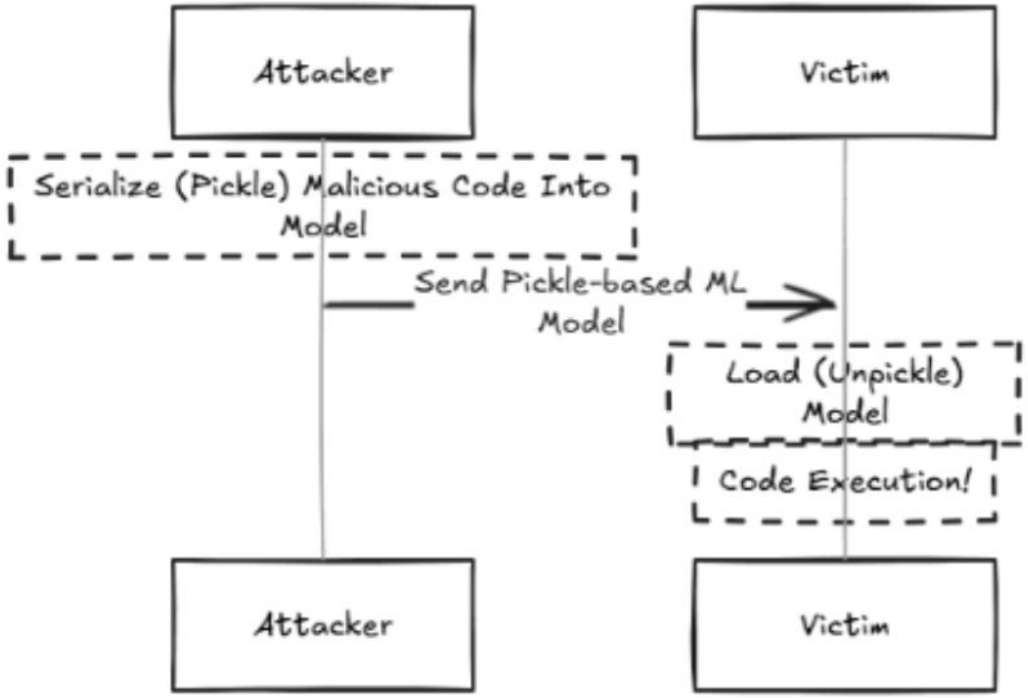

The balance

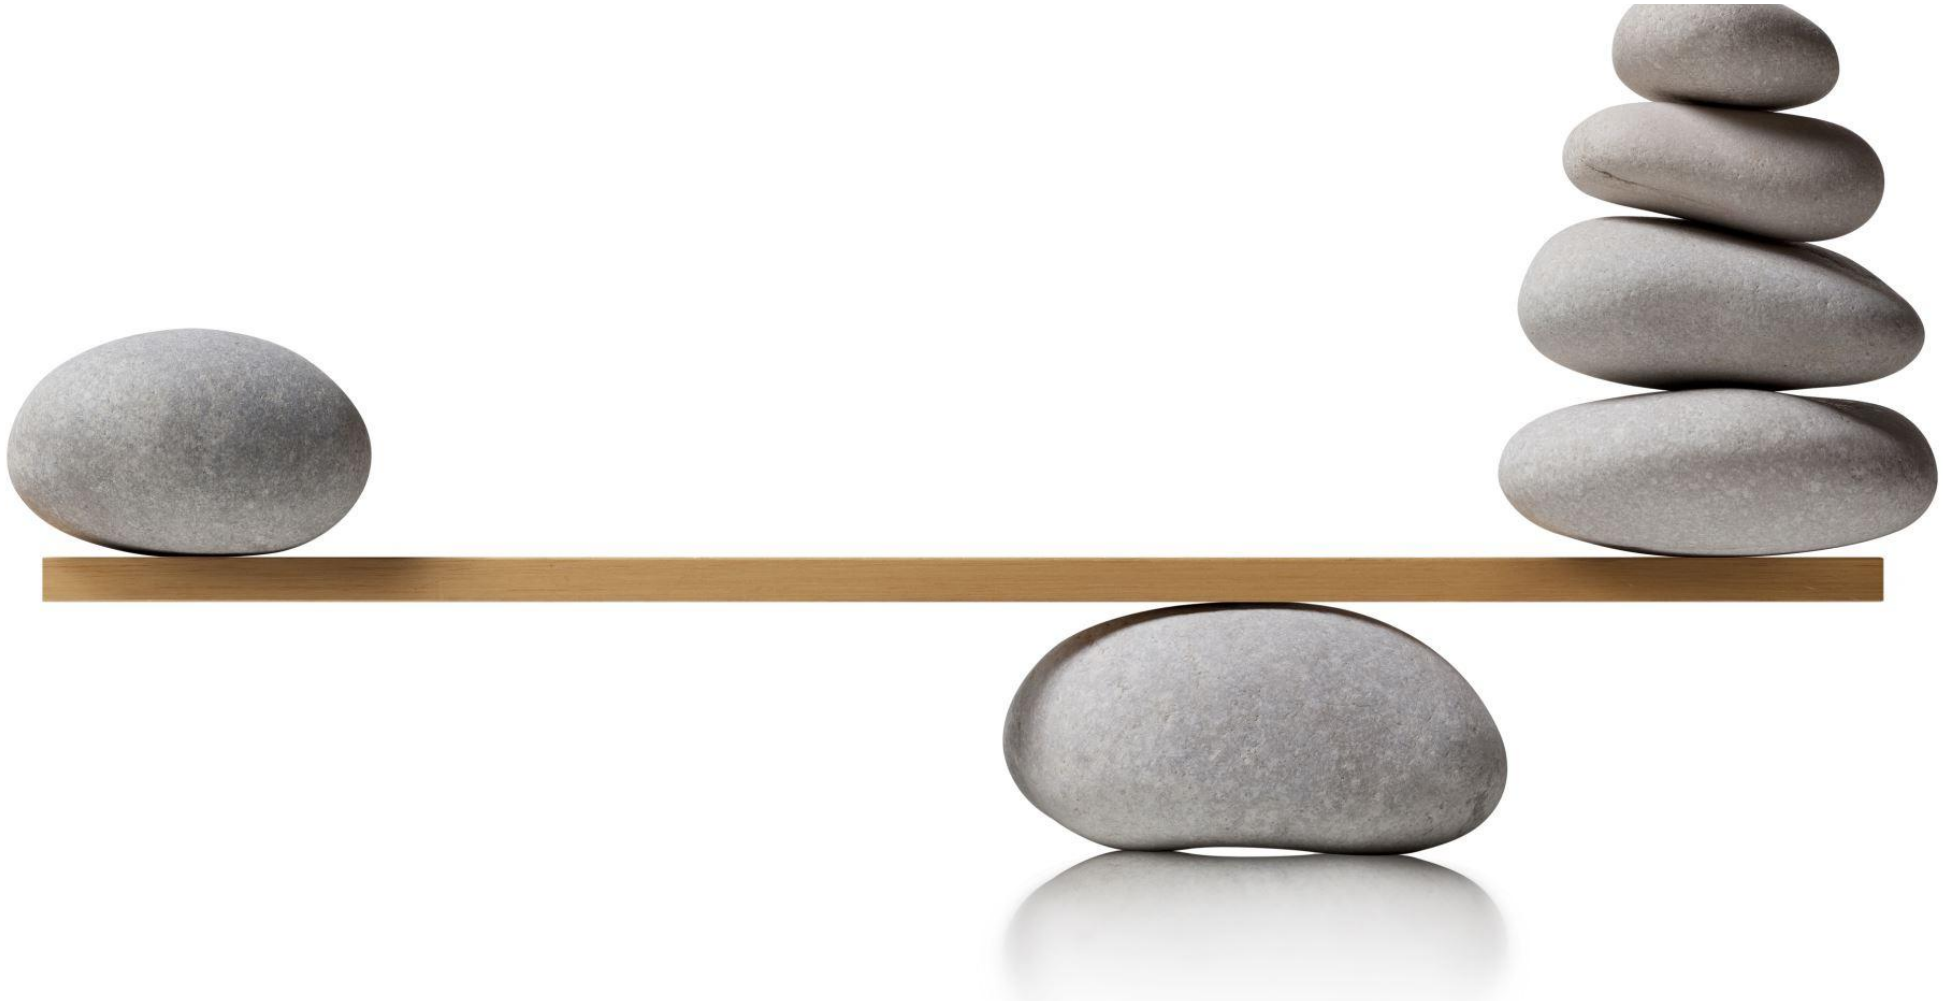

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

What are we doing?

High level view - Network Security Monitoring Stack

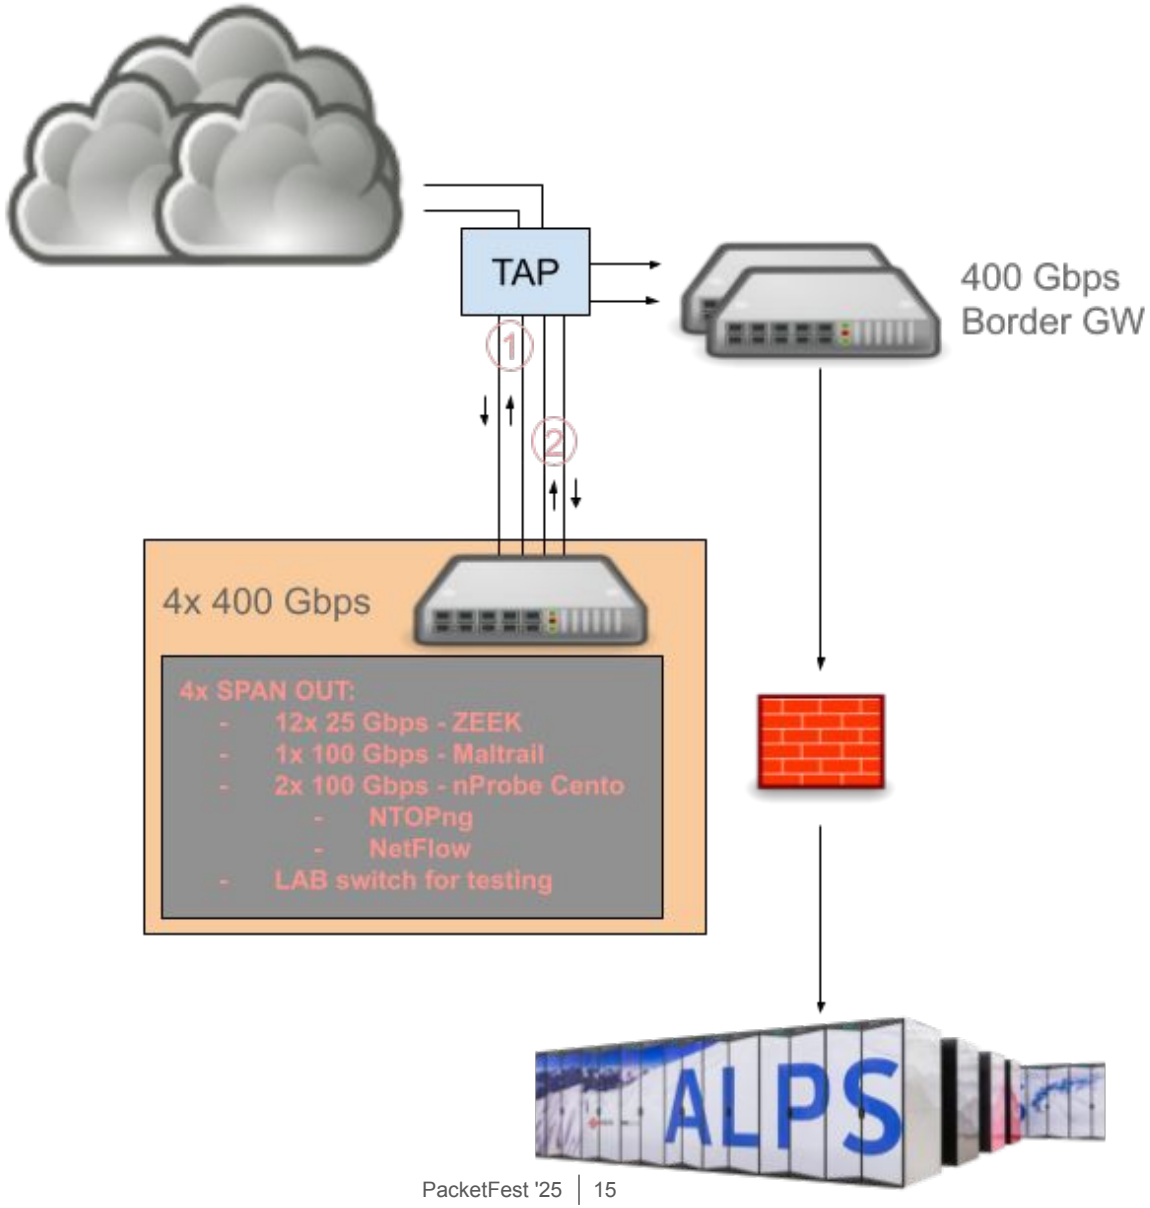

Current NTOP deployment

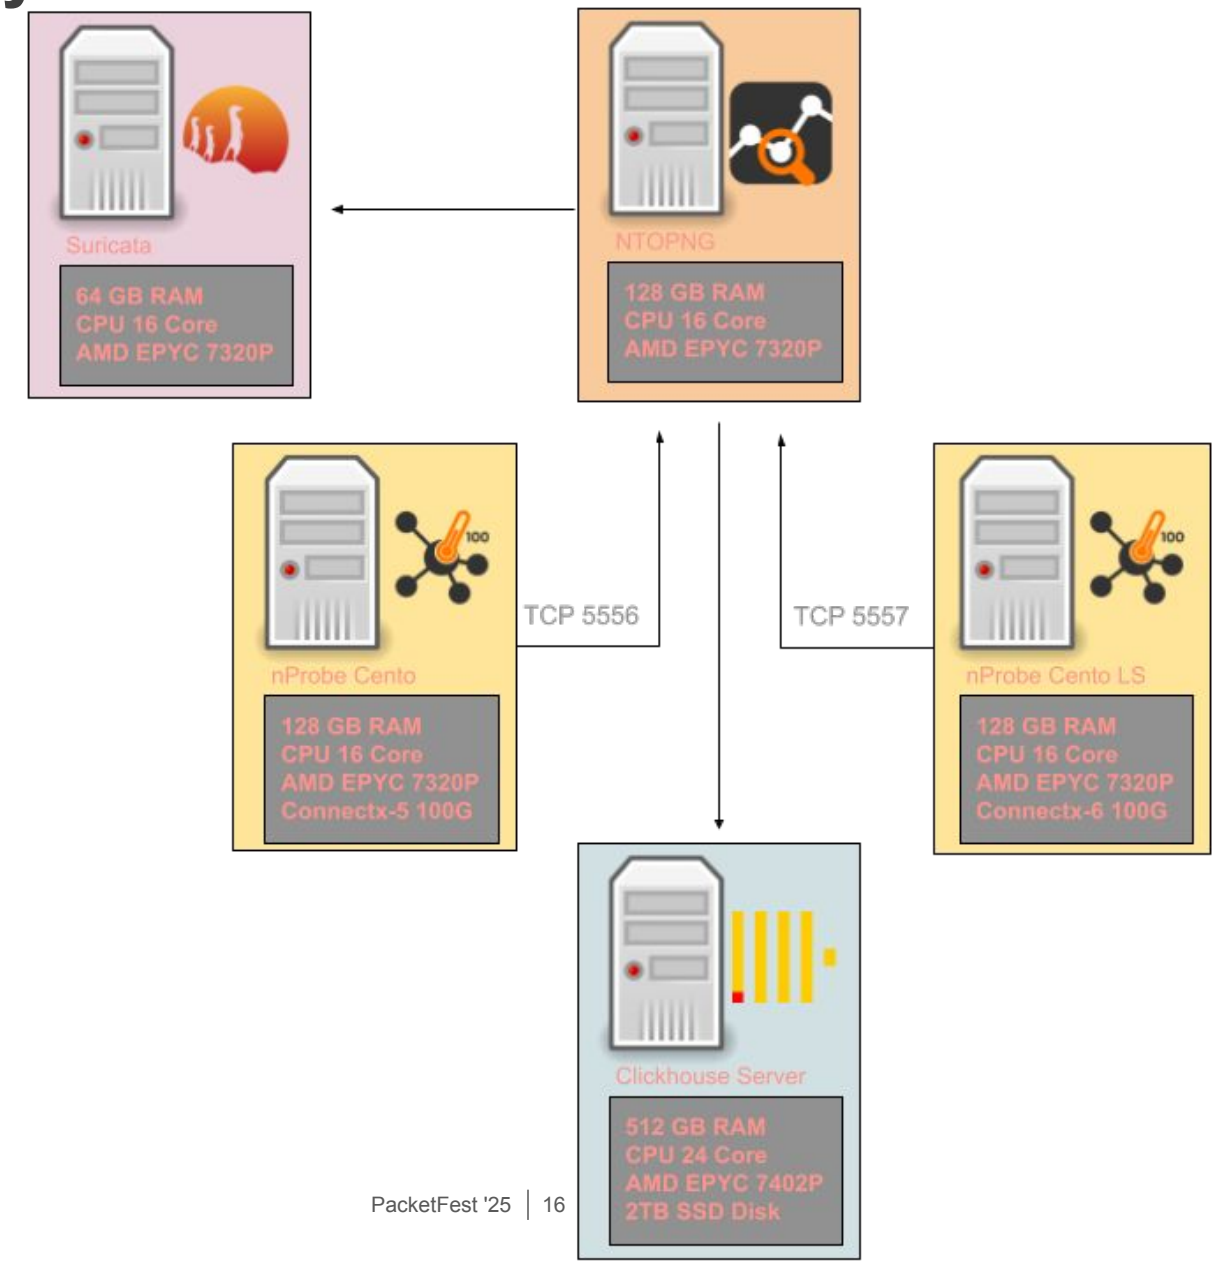

Current ZEEK deployment

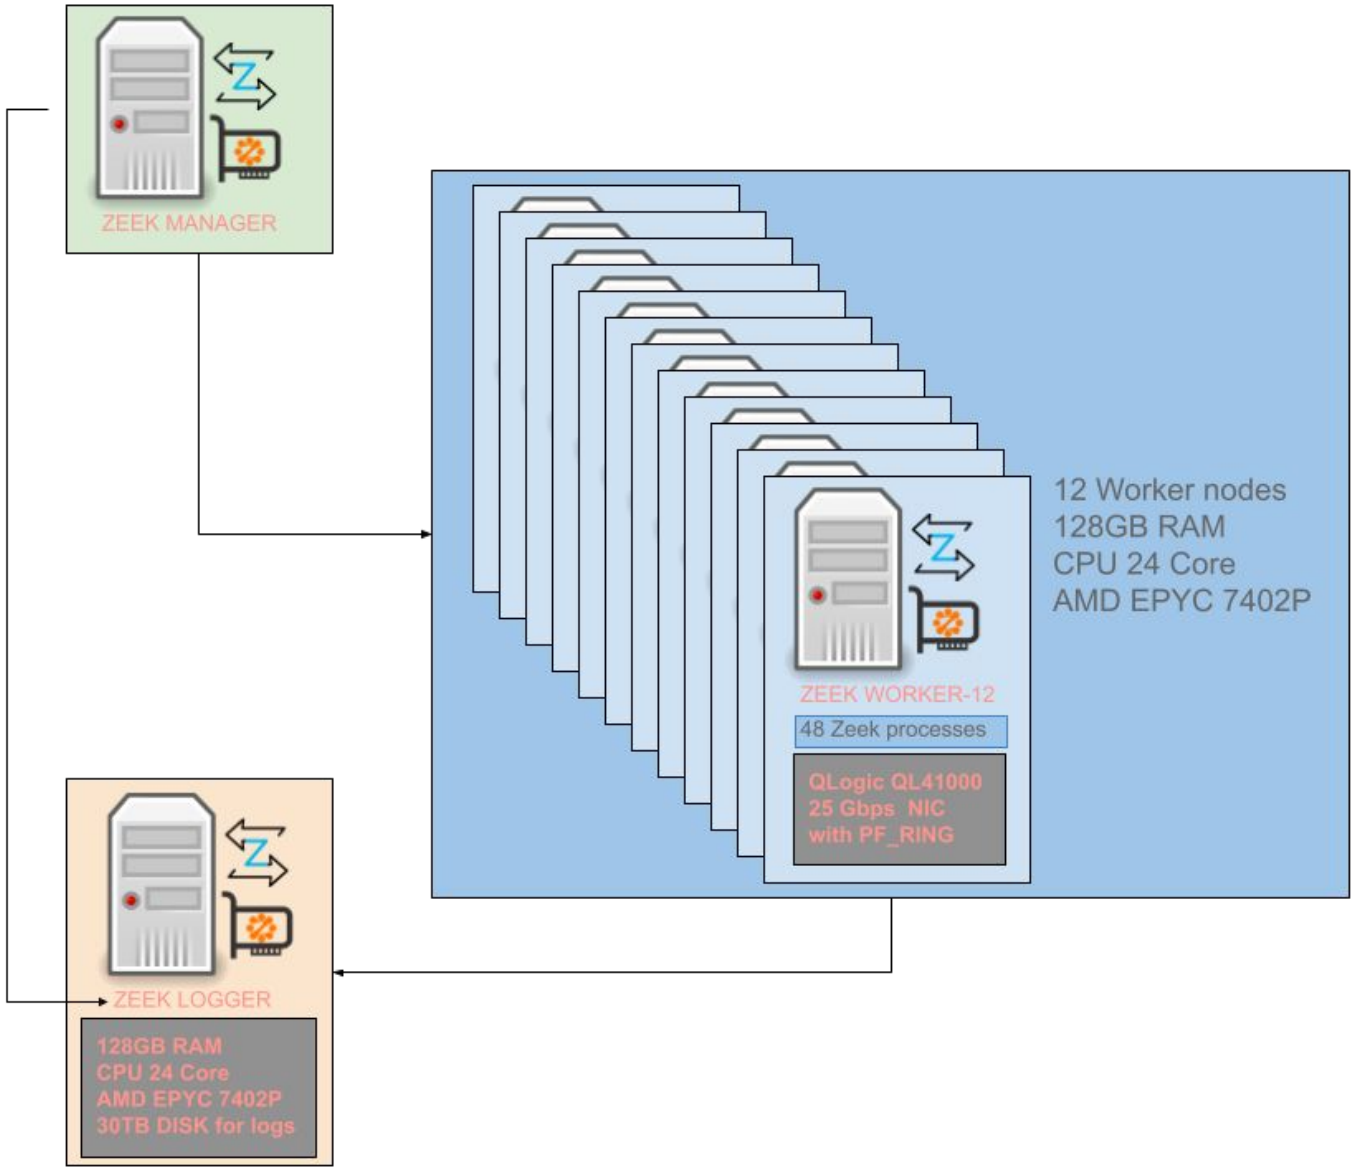

Security Monitoring Stack

A small slice of our reality

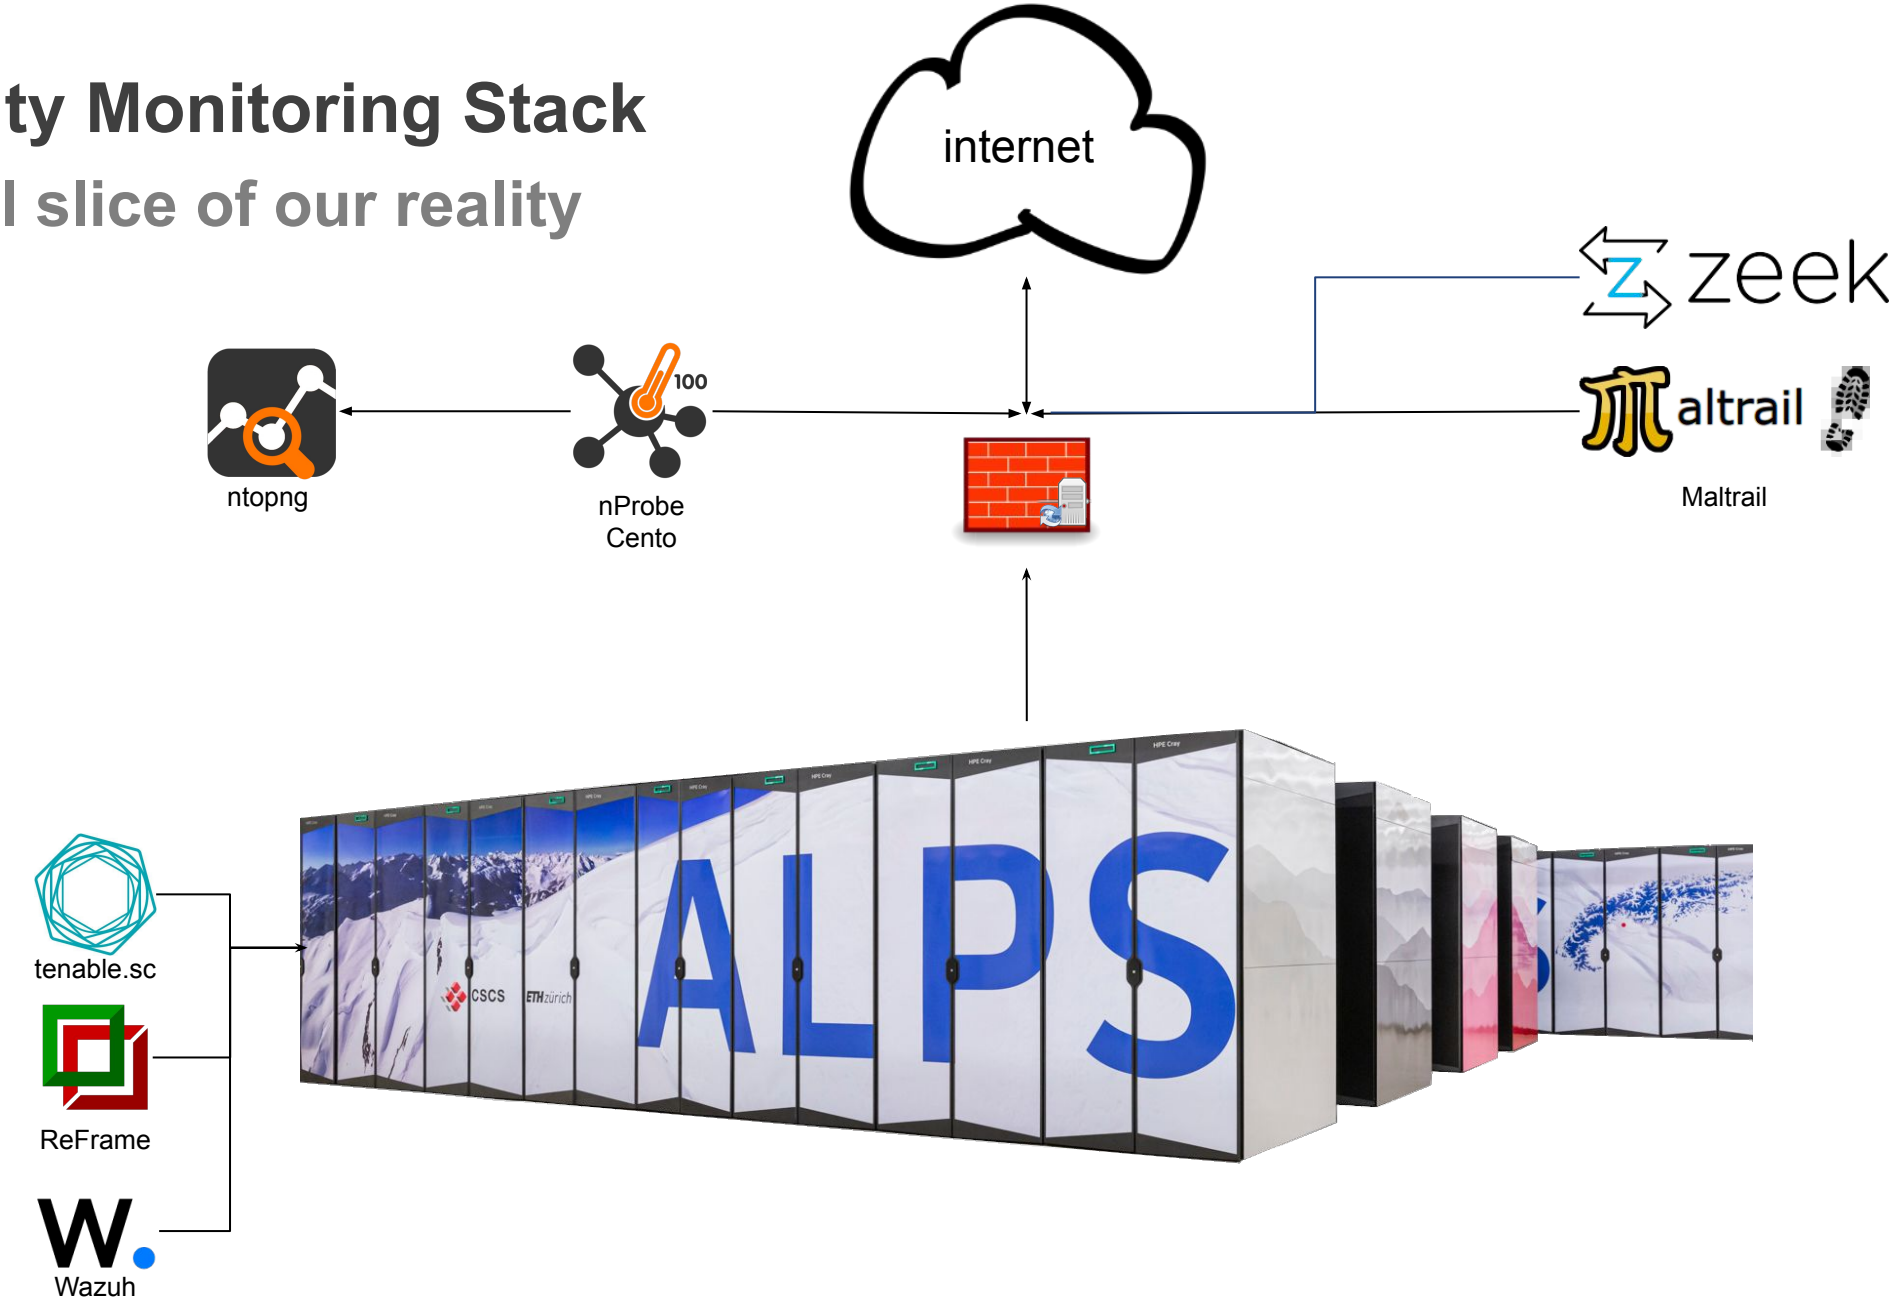

The tooling

Our current capability

- Generate reports to management about security threats and posture of our

CONFIDENTIAL//FOR OFFICIAL USE ONLY

CSCS
Centro Svizzero di
Swiss National Su

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

ETH Data Classification Level: CONFIDENTIAL

System Health Check

System:

Dashboard ISO

Generated on March 30, 202

Summary | Check Results | Disclaimer

Security Tests Profile

nid001293	root
Hostname	User

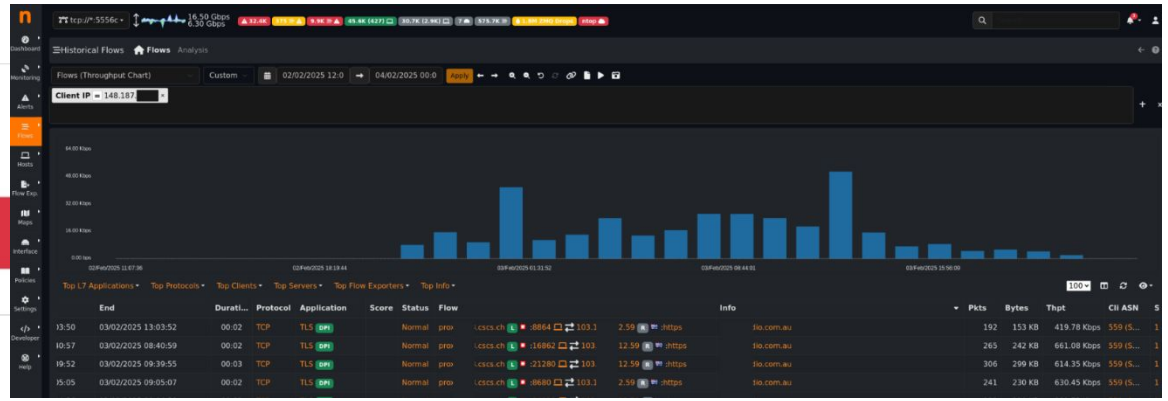

Wrong usage of ALPS infrastructure

An example of real threat

When we became the “bad guys”

- User crawled the internet and gray listed CSCS IP addresses

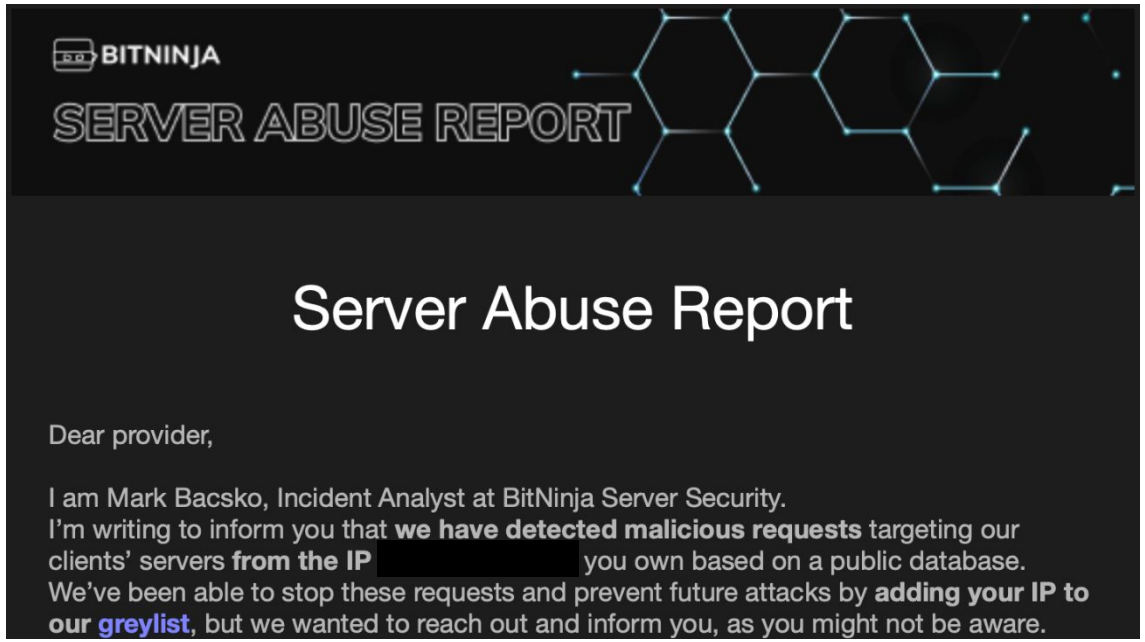

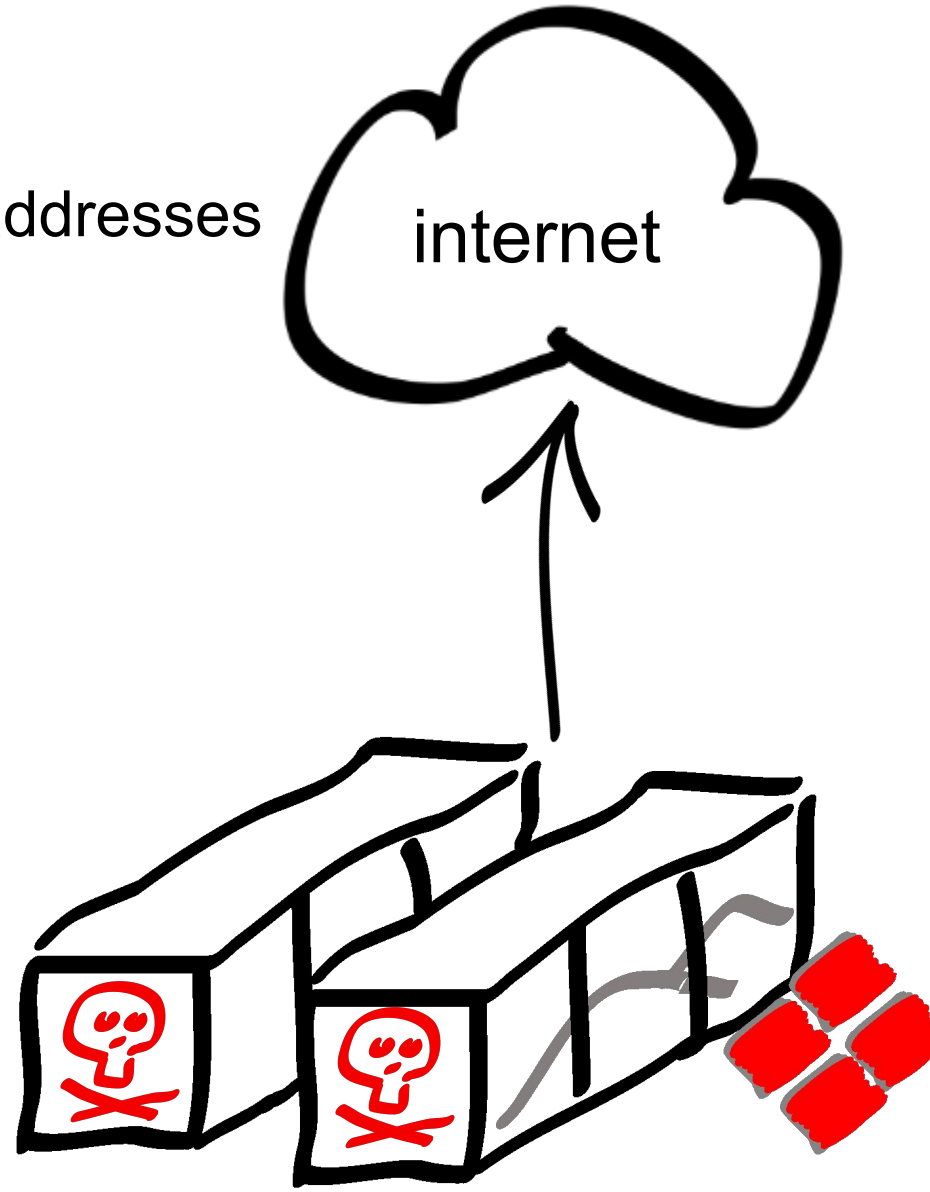

Identification - Firewall alert

We receive an alert from the firewall because of an unusual very high number of requests from ALPS network towards internet

Identification - Maltrail alert

Maltrail send an alert triggered by multiple connections towards known malicious websites

```
"2025-02-02 22:16:31.043301" maltrail 148.187.XXX.XXX 5826 81.XXX.XXX.12 443
TCP IPORT 81.XXX.XXX.12:443 "cobaltstrike-2 (malware)" (static)
```

1. https://github.com/stamparm/maltrail

Incident response - evidence collection

- Check on Maltrail logs to see the destination IP/URL
  - This helps to make some queries on the Firewall and Proxy
- Check on Firewall and Proxy to identify the responsible node(s)
- Once the node(s) is identified, search the start time


```
Feb  2 22:02:51 148.187.XXX.XXX date=2025-02-02 time=21:59:28
…
sessionid=547100635 srcip=172.XXX.XXX.XXX srcport=44162 srccountry="Reserved" srcintf="Ext"
srcintfrole="undefined" dstip=88.208.XXX.XXX dstport=80 dstcountry="United Kingdom" dstintf="Ext"
dstintfrole="undefined" proto=6 httpmethod="GET" service="HTTP" hostname="XXX.com"
…
url="http://XXX.com/wp-content/uploads/2013/06/Homeopathy-on-holiday-log.jpg" sentbyte=275 rcvdbyte=0
direction="outgoing" msg="URL belongs to a category with warnings enabled" ratemethod="domain" cat=26
catdesc="Malicious Websites" crscore=30 craction=4194304 crlevel="high"
```

Incident response - evidence collection

Check on NTOPng to find additional information on all the captured flows of the node in the identified time range

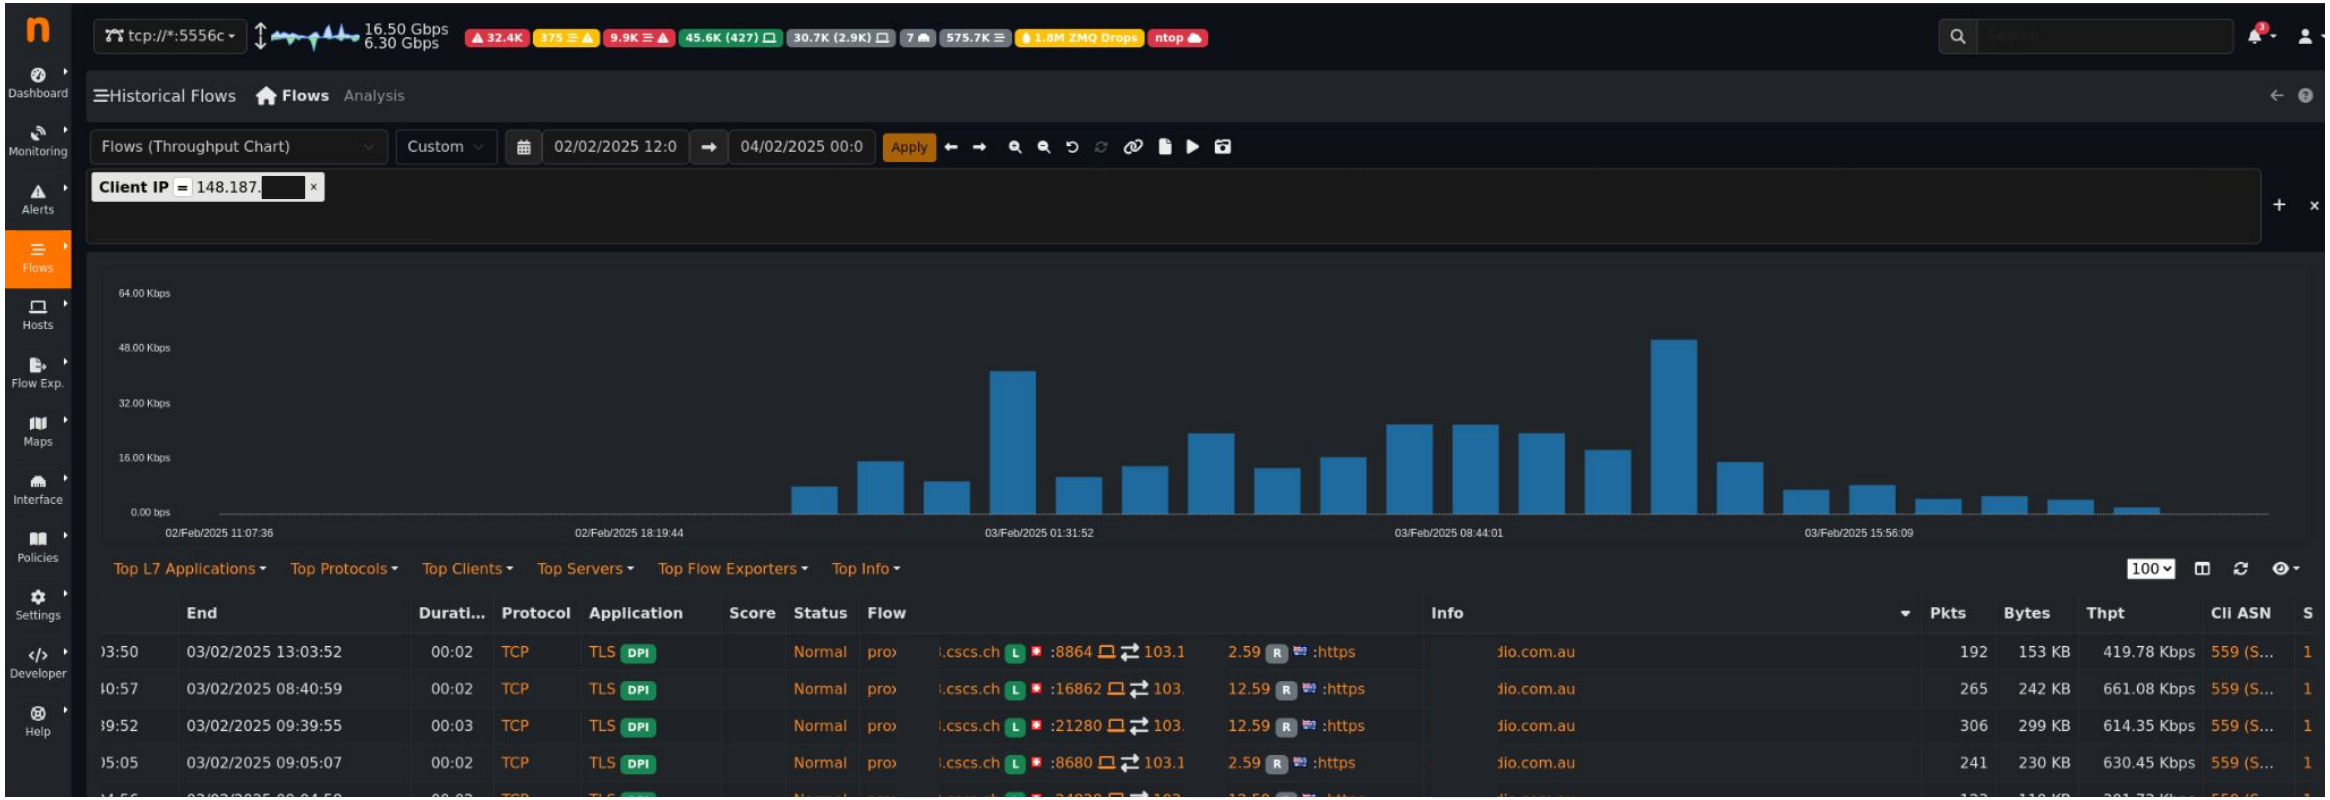

1. https://www.ntop.org/

Incident response - evidence collection

At this point we check on the node what is running and we identify the user

```
$ sacct -S 2025-02-02T21:00 -o start,end,user,jobid,jobname -N nid00XXXX
              Start                 End      User         JobID    JobName
------------------- ------------------- --------- ------------ ----------
2025-02-02T20:16:57 2025-02-02T21:03:56     user1 127157       voxDataGen
2025-02-02T20:16:57 2025-02-02T21:03:56           127157.batch batch
2025-02-02T20:16:57 2025-02-02T21:03:56           127157.exte+ extern
2025-02-02T21:10:56 2025-02-03T21:11:22     user2 127380       install_c+
2025-02-02T21:10:56 2025-02-03T21:11:23           127380.batch batch
2025-02-02T21:10:56 2025-02-03T21:11:24           127380.exte+ extern
2025-02-03T21:22:21 2025-02-03T21:27:06     user3 133124       run0_dpre+
2025-02-03T21:22:21 2025-02-03T21:27:06           133124.exte+ extern
2025-02-03T21:23:07 2025-02-03T21:27:06           133124.1     bash
```

Incident response - evidence collection

At this point we check on the node what is running and we identify the user

```
$ sacct -S 2025-02-02T21:00 -o start,end,user,jobid,jobname -N nid00XXXX
              Start                 End      User         JobID    JobName
------------------- ------------------- --------- ------------ ----------
2025-02-02T20:16:57 2025-02-02T21:03:56     user1 127157       voxDataGen
2025-02-02T20:16:57 2025-02-02T21:03:56           127157.batch batch
2025-02-02T20:16:57 2025-02-02T21:03:56           127157.exte+ extern
2025-02-02T21:10:56 2025-02-03T21:11:22     user2 127380       install_c+
2025-02-02T21:10:56 2025-02-03T21:11:23           127380.batch batch
2025-02-02T21:10:56 2025-02-03T21:11:24           127380.exte+ extern
2025-02-03T21:22:21 2025-02-03T21:27:06     user3 133124       run0_dpre+
2025-02-03T21:22:21 2025-02-03T21:27:06           133124.exte+ extern
2025-02-03T21:23:07 2025-02-03T21:27:06           133124.1     bash
```

Incident response - evidence collection

```
$ sacct -j 127380 --format=User,JobID,Jobname,state,start,end,elapsed,nnodes,ncpus,nodelist
 User     JobID        JobName         State      Start               End                 Elapsed    NNodes NCPUS NodeList
--------- ------------ --------------- ---------- ------------------- ------------------- ---------- ------ ----- ----------
user2     127380       install_coyo    TIMEOUT    2025-02-02T21:10:56 2025-02-03T21:11:22 1-00:00:26    1   288   nid00XXXX
          127380.batch batch           CANCELLED  2025-02-02T21:10:56 2025-02-03T21:11:23 1-00:00:27    1   288   nid00XXXX
          127380.exte+ extern          COMPLETED  2025-02-02T21:10:56 2025-02-03T21:11:24 1-00:00:28    1   288   nid00XXXX
```

Incident response - containment

- We proceed with node isolation
  - network isolation
  - drain the node in slurm
- Contact the Service Manager and the user

Incident response - evidence collection

- We collect and store relevant logs and files from the node(s)
  - /etc/passwd
  - /etc/group
  - /var/log/messages
  - /var/log/audit/audit.log
  - /var/log/zypper.log
  - /var/log/fakerootuidsync.log
  - /var/log/fabricmanager.log
  - /var/log/cray-lldp.log
  - /var/log/wtmp
  - /var/log/btmp
  - /var/log/sssd/sssd_pam.log
  - /var/log/sssd/sssd_ssh.log
  - /var/log/munge/munged.log
  - the user logs of the job
  - data downloaded during the job

Conclusions

- Currently we need to improve the visibility on the nodes
- The data collection must be tuned - we have many sources
- We will develop and implement a new tool to enhance the visibility of software running on the machines at a lower level with eBPF technology (stay tuned…)
- We are always looking for new ideas and brainstorming with other teams

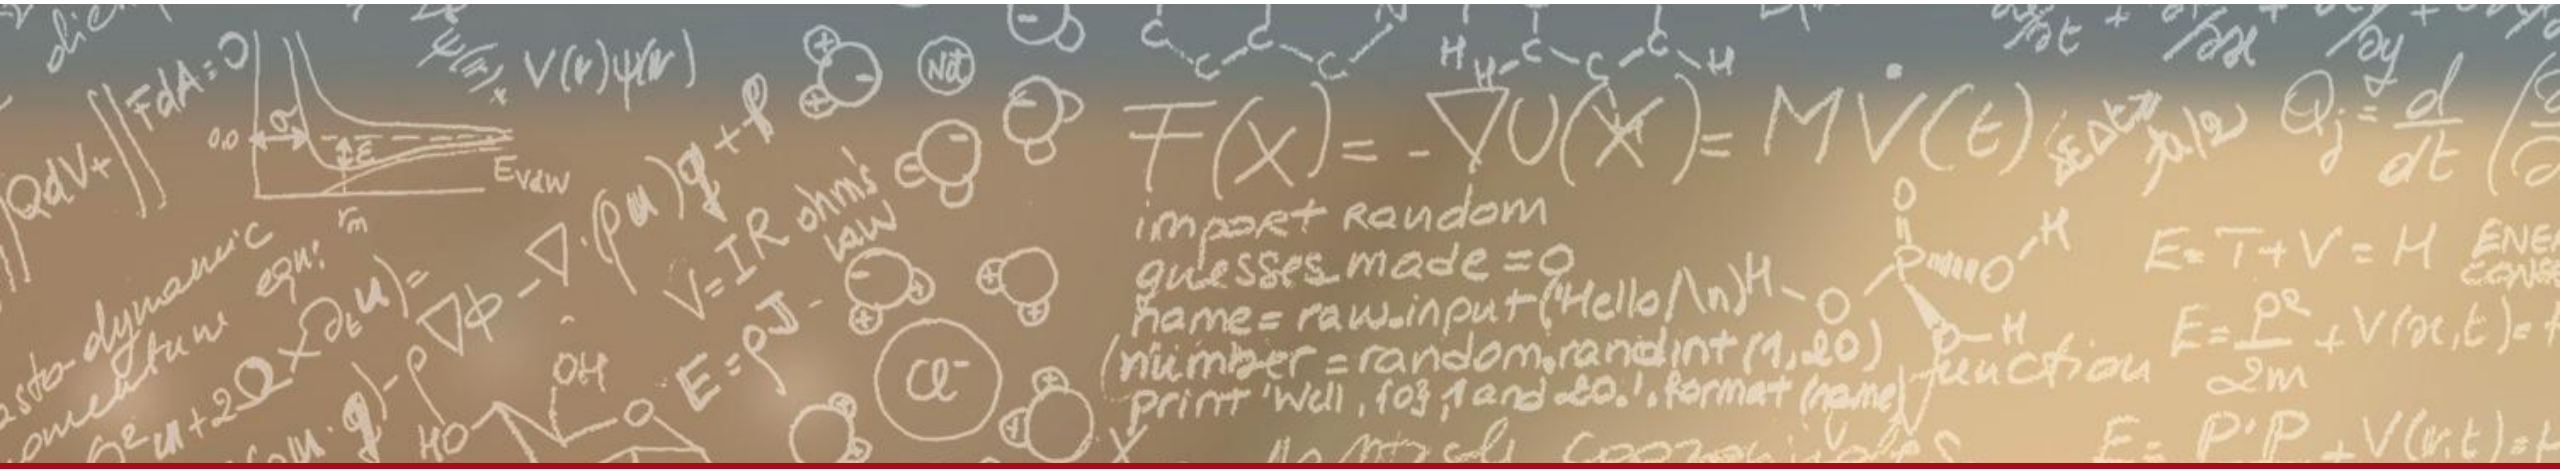

Thank you for your attention.

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

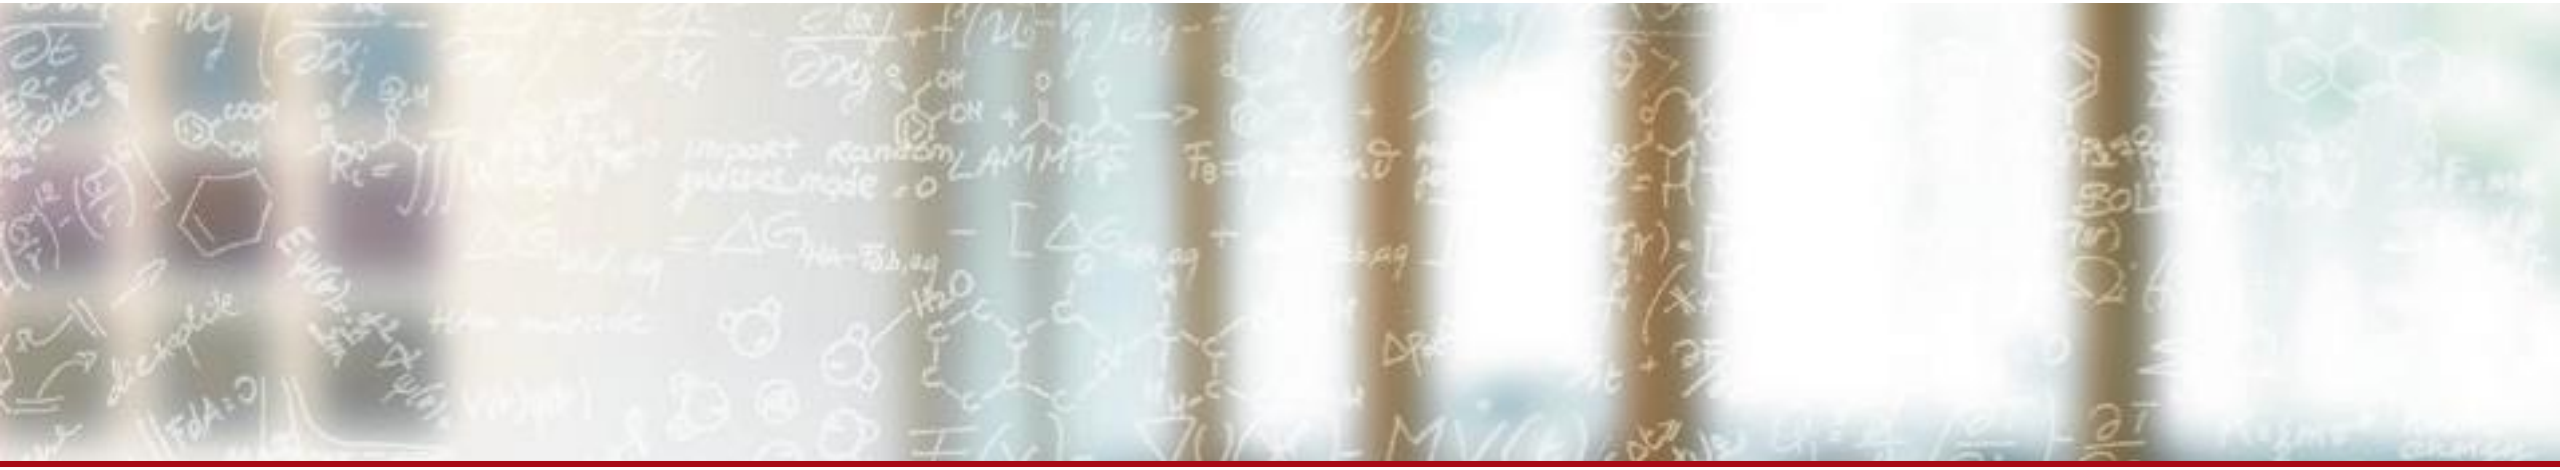

400 Gbps Observability

PacketFest '25

Fabio Zambrino, CSCS

May 8th, 2025

TPL: Green