Google

Linux CVE workgroup

LPC 2025

melotti@google.com

Dec 2025

Context

- Since Feb ‘24, Linux is its own CNA
- Linux Kernel CVEs, What Has Caused So Many to Suddenly Show Up? - Greg Kroah-Hartman @ CNCF
- CVE volume has significantly increased
  - Before: ~20/month
  - After: ~60/**week**
  - 1300+ filed from past commits

Challenge

- How do we triage them all?
- Different challenges for different players in the community
- For a cloud provider:
  - Can’t reboot machines
  - Decision in {ignore, (expedite) backport, livepatch}
- Applicability depends on threat model, exploitability, permissions, etc.
- Is there space for collaboration?

Outcome from LPC 2024

- Very high interest demonstrated by many different entities
  - Google, Microsoft, Canonical, Oracle, Red Hat, SUSE, Amazon...
  - Intention to collaborate on community-powered analysis
- CNA team is open to add our analyses to CVE records
- Threat models can be different, but the baseline analysis stays the same
- No engagement should be demanded to maintainers

Linux CVE workgroup

- Idea: share basic and indisputed CVE analysis with the community
  - Following a threat-model agnostic template
  - Everyone is already doing this work independently!
    - → Reduce toil
    - → Improve accuracy
- Goal: become an established source of vulnerability assessments
- https://github.com/cloud-lts/linux-cve-analysis
- Bi-weekly meetings
- #linux-cve-workgroup IRC channel on libera.chat

Example – CVE-2024-46800

```
reachability: Local
memory_corruption: yes
bug_class: Use-After-Free
impact: LPE
privileges_required: no
notes: An exploit has been submitted as part of kCTF (exp192). Privileged
  capabilities (CAP_NET_ADMIN) in user namespace are sufficient to exploit the
  bug.
author: Google
```

Example – CVE-2024-46800

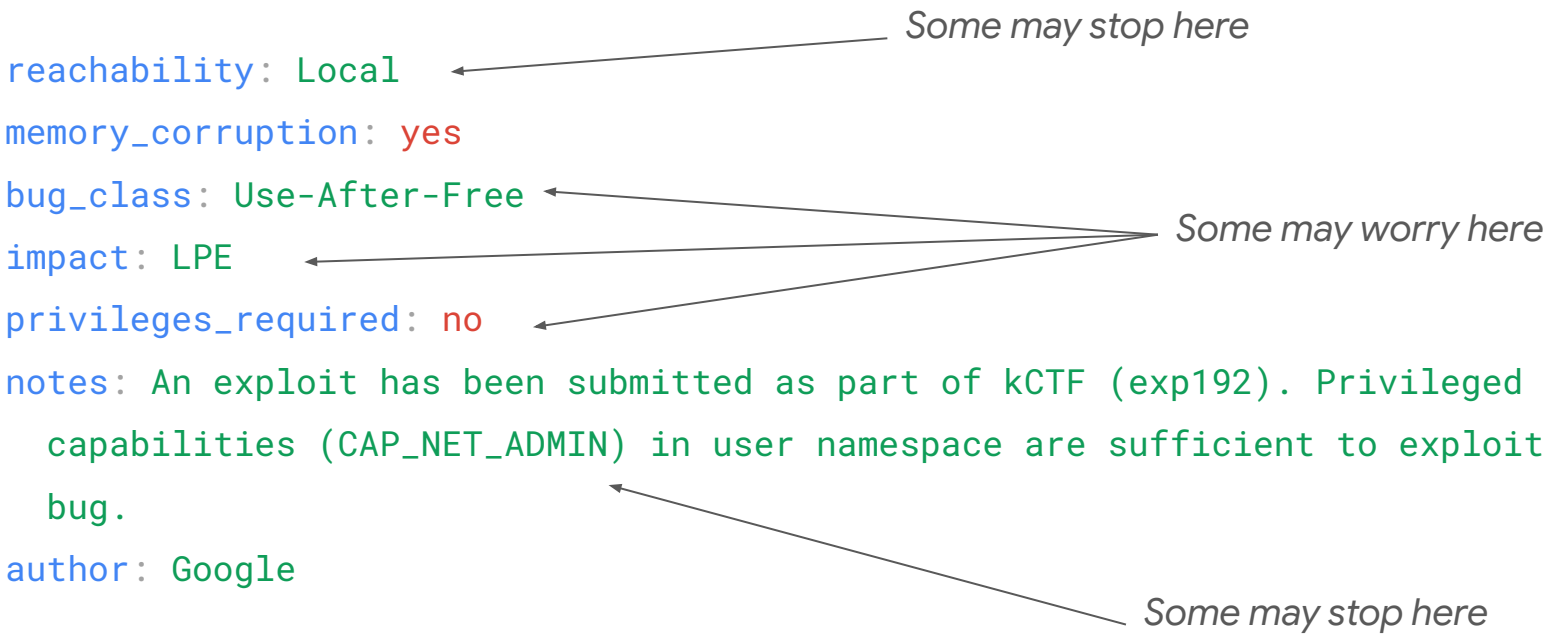

Why not CVSS?

- ✅ Industry-standard for vulnerability severity scoring
- ✅ Some organizations need to provide it as part of their advisories
- ❌ Fails to properly describe a kernel vulnerability
  - What's the difference between a WARN_ON and a null ptr deref?
  - What's the difference between a local process and a VM?
- ❌ Specific to threat models, environments, configurations
  - Some CVEs already have different CVSS scores

→ Idea: a template that is a **subset** of CVSS

What’s working?

- Collaboration is recognized as essential in this space
- Lots of important stakeholders actively participate
- 704 CVEs analyzed
- Active discussions over progressive improvements
    - How to leverage LLMs
    - Focusing only on high-severity issues
    - Etc.

What can be improved

- **Latency**
  - Analyses pushed in batches some time after CVEs were published
  - In the meantime, everyone had to triage them already
- **Throughput**
  - Only a small subset of CVEs is analyzed by the group
- Alignment challenges
- Our template vs CVSS

What’s next?

- My 2c: we have to integrate this into our internal processes
- AI to the rescue?
  - AI-powered analysis are becoming better and better, to the point that we should start using and sharing them
  - LLMs can also help with subtasks (e.g., translate internal analysis with upstream-able one)
  - Still, there will always be humans in the loop