CIRCL
Computer Incident
Response Center
Luxembourg

AI at CIRCL

Building our own models for vulnerability intelligence

https://circl.lu/pub/ai-strategy

Cédric Bonhomme - cedric.bonhomme@circl.lu

April 28, 2026

Presented at International Committee of the Red Cross — Luxembourg

TLP:CLEAR

Contents

How CIRCL approaches AI

Our AI approach in one slide

We **enhance** existing tools rather than replace them with a chatbot.

- **Pragmatic**: AI is one tool among many, used where it actually helps analysts.
- **Open**: open data, open datasets, open models, open code.
- **Sovereign**: we train and host our models with our own hardware.
- **Small and specialised**: task-specific models, not general-purpose LLMs.
- **Strategy**: https://circl.lu/pub/ai-strategy

Vulnerability-Lookup, in 30 seconds

- Open Source project led by **CIRCL**, co-funded by the **European Union**.
- One place to **search**, **triage** and **track** vulnerabilities from 70+ sources.
- More than **1.7M security advisories** and **400k+ sightings** collected.
- Used by CSIRTs, SOCs, ENISA (EUVD), and many organisations.
- Reference implementation of the **GCVE** standards[a].

[a] https://gcve.eu

vulnerability-lookup

https://www.vulnerability-lookup.org

ail project

Building our own models

From raw data to AI datasets

- Years of experience handling large security datasets (MISP, AIL, Lookyloo, Passive DNS, BGP Ranking...).
- We turn messy, multi-source vulnerability data into **structured, actionable** datasets.
- Datasets are **publicly shared** on Hugging Face[1].

Dataset	Rows	Features
vulnerability-scores	694,871	Descriptions (en), CVSS, CPE
vulnerability-CNVD	127,854	Descriptions (zh), CVSS
Vulnerability-FSTEC	82,960	Descriptions (ru), CVSS
vulnerability-cwe-patch	39,260	Descriptions, CWE, full patch diffs

[1] https://huggingface.co/CIRCL

We train our own models

Custom models, trained **locally** on our own datasets.

- Stack: **PyTorch** + Hugging Face **Transformers**.
- Base models: **RoBERTa**, **MacBERT** (Chinese), **ruRoberta** (Russian)
- Fine-tuned on CIRCL datasets.
- Trained, versioned and published with **VulnTrain**[a].
- Paper: https://arxiv.org/abs/2507.03607

[a] https://github.com/vulnerability-lookup/VulnTrain

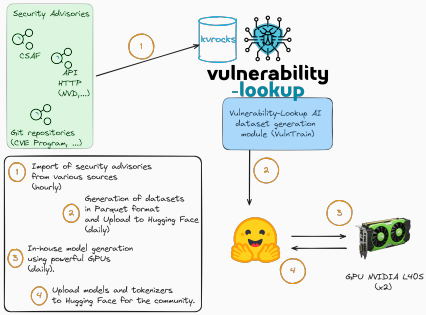

Models we have trained so far

Model	Params	Epochs	Accuracy	Train time
Severity classification (EN)[1]	125M	5	0.829	< 4h
Severity classification (ZH)[2]	102M	5	0.782	< 1h
CWE guessing[3]	125M	36–40	0.875	< 1h

All models are published openly — anyone can pull them, audit them, fine-tune them.

[1] https://huggingface.co/CIRCL/vulnerability-severity-classification-roberta-base
[2] https://huggingface.co/CIRCL/vulnerability-severity-classification-chinese-macbert-base
[3] https://huggingface.co/CIRCL/cwe-parent-vulnerability-classification-roberta-base

VLAI in action

Severity prediction — example

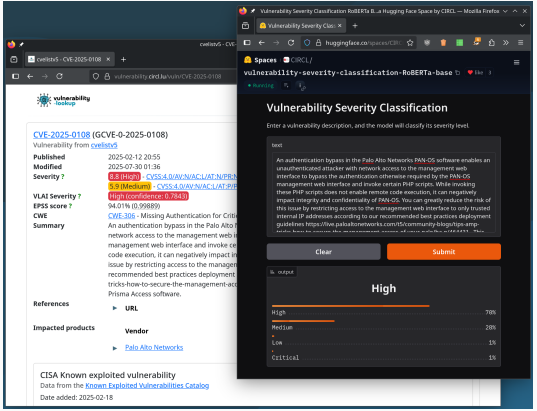

Predicted severity and confidence directly inside Vulnerability-Lookup.

Useful even when data is scarce

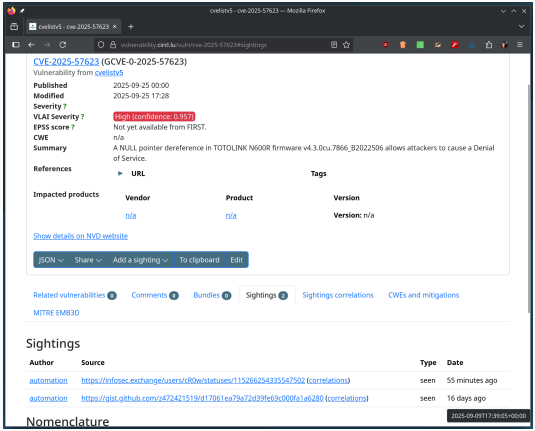

Reserved CVE, sightings only — the model still gives a useful signal.

Integration via the ML-Gateway

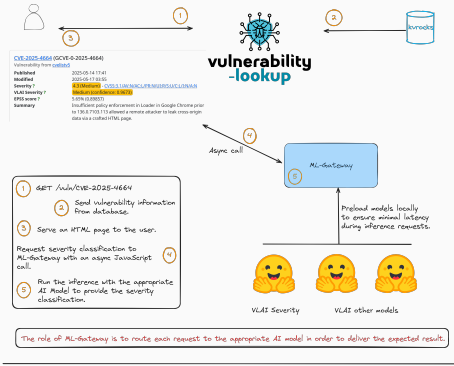

- **Optional** component, decoupled from Vulnerability-Lookup.
- Models pulled from Hugging Face, preloaded locally.
- Documented **OpenAPI** for inference.
- Self-hosted: **your data stays with you**.
- `https://github.com/vulnerability-lookup/ML-Gateway`

One HTTP call, one prediction

```
$ curl -X POST \
  'https://vulnerability.circl.lu/api/vlai/severity-classification' \
  -H 'Content-Type: application/json' \
  -d '{
    "description": "An authentication bypass in the API component of
      Ivanti Endpoint Manager Mobile 12.5.0.0 and prior allows attackers
      to access protected resources without proper credentials."
  }'

{"severity": "High", "confidence": 0.8225}
```

From Vulnerability Data to Intelligence

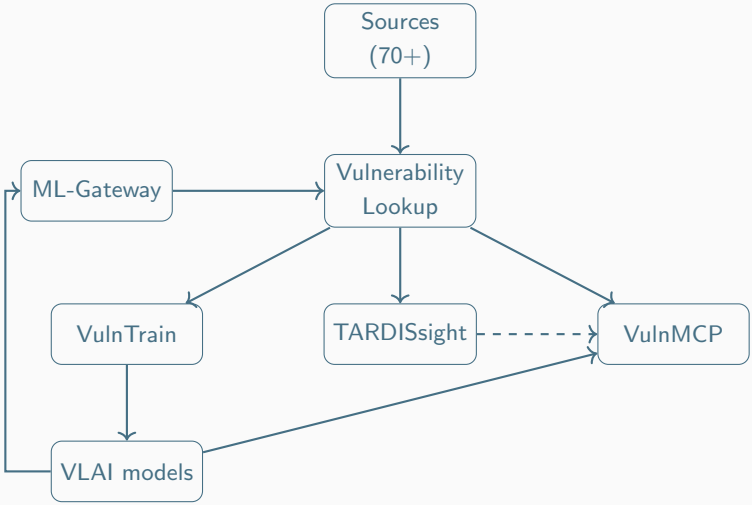

Closing

What's next

- Deeper analysis of **sightings** and exploitation context.
- Insights generation from **patches** (commit + diff).
- **CPE guessing**
- **Assisted vulnerability description redacting**
- Continued participation in collaborative research (e.g. EU AIPITCH project).

Open data, open models, open code — feedback and contributions are welcome.

References

https://www.vulnerability-lookup.org

https://vulnerability.circl.lu

https://github.com/vulnerability-lookup

https://huggingface.co/CIRCL

https://arxiv.org/abs/2507.03607

https://arxiv.org/abs/2604.16038

GPU Efficiency in VLAI Model Training

https://social.circl.lu/@circl

Thank you

- Questions, ideas, new sources:
  `https://github.com/vulnerability-lookup/vulnerability-lookup`
- Contact: info@circl.lu

CIRCL
Computer Incident
Response Center
Luxembourg