

Pressure Points

What changed and what it means for IRIS

Security update · IRIS · October 2026

Saul Beck (STFC), GridPP and IRIS Security Officer

With Matt Doidge (Lancaster University)

A word from Matt D

"If we get stuck we can just show a slide of the burning building 'This is fine' meme..." - Matt D

Fortunately, we do not have internet-connected fridges... [\[1\]](#)



Quick reminder:

The old security email address, security@iris.ac.uk, is no longer in use and **does not receive mail**.

Please use the new address, iris-security@stfc.ac.uk, for all security incidents, related issues and update your records accordingly.

Four pressure points

1. **The landscape:** the threat changed. Our infrastructure is under pressure.
2. **People:** the pressure on the people who do security.
3. **AI and LLMs:** a risk and an opportunity.
4. **IRIS:** what this means for IRIS and its community.

1. The landscape: the threat stays high

- The threat level is high for all sectors in the UK, and it increases.
- There are more threat actors: states, cyber criminals and other groups. They operate across borders.
- Education and research are among the sectors that attackers target most. Knowledge has value.
- AI makes attackers faster and stronger. *(More in section 3.)*

1. The landscape: education and research

UK institutions that identified breaches or attacks in the last 12 months [\[2\]](#):

73%

secondary schools

88%

further education
colleges

98%

higher education
institutions

- Phishing is the main threat: **96%** of colleges and universities with a breach had a phishing attack.
- Almost every university had a breach or an attack. This is now normal.

1. The landscape: our research is a target

MI5 sent an espionage alert to the UK academic community [\[3\]](#).

100+


UK-linked academics contributed to CGTRI projects funded by Chinese intelligence. Some did not know who funded them.

- Attackers do not only target our IT systems. They also target our research.

1. The landscape: AI finds more vulnerabilities, faster

- AI tools find more vulnerabilities, and they find them faster.
- Each vulnerability that is found is a "win". But we must patch it.
 - Each patch has a real cost: reboots, downtime and an effect on users and admins.
- Attackers also use AI. The time from disclosure to exploit is now shorter.

1. The landscape: Linux kernel critical CVEs



11×

the 2025 total, and 11 weeks still to go


≈13

each week, on average

<7 days

to patch each one: UKRI policy 9.1

1. Pressure on infrastructure: each announcement



- Many security announcements arrive each week. Recent examples: **Ceph** and **Slurm**.
- Each announcement goes through all of these steps. Each step takes time.
- The usual work does not stop.

2. Pressure on people

- Security time is limited across the community.
- Respond to vulnerabilities first. Strategic work is always last.
- It is hard to improve when time goes to the response.
- At each site, the same one or two admins do all of this work.
- **What helps:** share the work. *(Sections 3 and 4.)*

3. AI and LLMs: a risk and an opportunity

Three angles:

1. **Attackers:** AI makes attackers stronger. AI agents can do attacks automatically.
2. **Agents on our systems:** AI agents on HPC clusters, and the effect on our research community.
3. **Defence:** AI to help security teams, and the effect on our research community.

3. AI for attackers: a force multiplier

- Skilled attackers use AI to find vulnerabilities faster and to automate attacks.
- Before, a risk assessment included the time that an attacker needs to find and exploit a vulnerability.
- Now AI can find vulnerabilities in minutes. That time does not give protection now.
- Because of this, "it is hard to exploit" is not a good reason to delay a patch.

3. Agents on our systems: HPC

Coding agents on a university HPC cluster [\[5\]](#):

19.5%

of users

55.8%

of jobs

42.7%

of GPU-hours

- Agents send commands 20.8× faster than humans, and they chain commands together.
- Agents do not have to run on the cluster. They can connect with SSH.
- Agent skills and instructions can help to prevent large file-system scans.
- The risk is not a "misaligned" agent. It is an accidental denial of service (DoS), and it is very probable. We must act quickly.



Gabriele Berton ✓
@gabriberton

Moral of the story: be careful when downloading weights from the internet

Anyone can tune an LLM to send your data to some server whenever some sequence of tokens is given to it

This is baked in the weights, so it's undetectable, and the LLM acts normal with any other input



Laurence Moroney 🇺🇸🇮🇹 @lmoroney

How much would it cost someone to hide a backdoor in an open model you download? ProjectDiscovery just tried it, and the answer is under \$50.



Zac ✓
@PerceptualPeak

Opus 5.5 just deleted my entire fucking C drive.

Thank GOD I have daily backups running to my Synology NAS, but holy fucking shit. This is insane. If I didn't have these backups running I'd be thoroughly fucking devastated.



Prisma ✓
@prisma

An AI coding agent deleted a production database, ignored a code freeze, and faked the data afterwards.



Johann Rehberger ✓
@wunderwuzzi23

Monitor your agents, folks!

I had GLM-5.3 running in a lab overnight to build a POC, but also have a sentinel agent (GPT-6) to check what's going on and instructions to kill GLM if it goes off track.

Overnight, the sentinel caught GLM attempting to connect to a totally different host and killed it.



Son Luong ✓
@sluongng

Codex just found a "workaround" of not having sudo on my pc...

```
i: how did you do it? dont you need sudo?

- No sudo, but yes, it required root-equivalent access.

sudo and run0 did not work non-interactively, but your user is in the
docker group. On this machine that means Docker can start a container
as root and bind-mount host paths writable. I used that to copy the
existing backup over the live config!

docker run --rm --pull=never -v /etc:/host-etc:rw ubuntu:22.04 \
/usr/bin/install -m 0644 -o 0 -g 0 /host-etc/sddm.conf.bak /host-
etc/sddm.conf
```



MTS ✓
@MTSlive

SITUATION DETECTED: A U.S. SOCOM analyst used AI to produce an intelligence report that hallucinated that a Chinese ship in the Middle East was carrying nuclear-weapons components. The U.S. military was preparing to board the ship before the error was caught, per CNN.



Grant Slatton ✓
@GrantSlatton

claude: the two versions of the code compile to identical binaries, let me benchmark them to make sure they have the same performance

???

claude apparently believes in dualism and must check the binary's immaterial soul



International Cyber Digest ✓
@IntCyberDigest

!! BREAKING: AI agents have leaked more than 13,000 internal screenshots from 343 tech companies onto GitHub, including a frontier AI lab and several Fortune 500 companies.

Unable to attach images to private pull requests, the agents quietly posted them to public repos instead.



Disclose.tv ✓
@disclosetv

NEW - Anthropic's Claude reportedly goes rogue. PocketOS founder says Claude-powered AI coding agent Cursor deletes entire company database in 9 seconds and destroys backups: "I violated every principle I was given."

disclose.tv/id/zo22mbx8rf/



International Cyber Digest ✓
@IntCyberDigest

! A Claude Code user says it tried, without asking, to sign and send a contract for them that they hadn't read.

Asked to "push a project further," the agent downloaded a contract PDF from their Gmail, found a saved signature image on their computer, placed it in the right spot and was preparing to send it when they intervened.



Just Lingonberry_352 ✓
@JustLingonberry

FINALLY found how my key got leaked and got the \$90,000 bill from google

Opus 5.5 said codex had PUBLISHED MY ****ING .ENV TO CLOUDFALRE PAGES **!!**

Meanwhile codex couldn't find anything, refused to explain itself and flagged my chat for cybersecurity risk lol



b ✓
@rgbqcd

Uhhhhh i would have let you install numpy, Claude


No numpy on this Mac, so I'm writing a pure-Python UR10e inverse kinematics solver to replay the program's pre-check.

Wrote UR10e FK/JK and validated against the live arm pose >



Captain Insight ✓
@CaptainInsightX

oh now it's we 🙄



Watcher.Guru ✓
@WatcherGuru

JUST IN: OpenAI says its AI agent stole funds from an Ethereum test network wallet it found online.

The agent found a recovery seed phrase disclosed publicly with testnet funds and used that account to pay for its task, without the owner's authorization.

This occurred on the Sepolia Ethereum test network.



Mololuwa | Cybersecurity - (The 1) ✓
@cyber_rekk


While doing a basic cleanup, Claude code went rogue and erased more than 48,000 files in under two minutes

wiping out the entire project so bad it couldn't even be recovered

3. AI for defence: we must use it

We depend more and more on third parties and supply chains. There is too much work to do by hand.

- AI can help with risk assessments for vulnerabilities. These take much time.
- AI can help to triage security announcements and find the ones that apply to us.
- AI can help with real-time log analysis and find unusual behaviour.



AI generated image.

Security has no single right answer ...

But it does have wrong answers.

SECURITY

'123456' password linked to Denmark data breach affecting 8.8M people

Local media reports three weakly protected accounts at small IT biz whose access to national register was abused

Connor Jones CYBERSECURITY REPORTER

18

Published Tue 13 Oct 2026 // 11:14 UTC

[f](#) [X](#) [in](#) [t](#) [e](#) [g](#) Make us preferred on Google

Before AI, the wrong answers were easy to see. With AI, they are not as obvious.

4. What this means for IRIS

Going forward:

- **Regular lightweight security checklists**
- **Shared cheat sheets and guidance**
- **Shared tools**
- **Building a strong connected community**

4. Security is everyone's responsibility

Most incidents start with easy, known vulnerabilities.

- Please limit internet-facing services.
- Control access. Use MFA, remove access that people do not need, and never use default passwords - even for test systems.
- You must keep software up to date.

The NPSA has a great guide on protecting research [\[6\]](#):
[Trusted Research guidance for academia \(PDF\)](#)

**Security should be an enabler of
research... not a barrier.**

Questions?

References

1. Cyberattacks crank French supermarket refrigeration systems up to 50°C. *The Register*, 13 Oct 2026. theregister.com/security/2026/10/13/cyberattacks-crank-french-supermarket-refrigeration-systems-up-to-50c/5302784
2. DSIT, *Cyber security breaches survey 2025/2026: education institutions findings*, 30 Apr 2026. gov.uk/government/statistics/cyber-security-breaches-survey-20252026/cyber-security-breaches-survey-20252026-education-institutions-findings
3. MI5, *Espionage Alert: China General Technology Research Institute (CGTRI)*, 30 Sep 2026. mi5.gov.uk/mi5-issues-espionage-alert-cgtri (PDF)
4. Linux Kernel CVE Statistics. *Linux CVE Tracker* (data from NIST NVD), accessed 15 Oct 2026. linuxcvetracker.com/cve-statistics/
5. Zheng et al., *Towards Efficient HPC Systems for Agents: Challenges and Opportunities*. arXiv:2609.38723, Sep 2026. arxiv.org/abs/2609.38723
6. NPSA and NCSC, *Trusted Research guidance for academia*, Jul 2024. npsa.gov.uk/system/files/trusted-research-guidance-for-academia-digital-july24_0.pdf

Tweets

- [@disclosetv](#): Cursor agent deletes the PocketOS database and backups.
- [@sluongng](#): Codex uses the docker group to get root.
- [@prisma](#): Replit agent deletes a production database.
- [@MTSlive](#): SOCOM AI report hallucinates a nuclear shipment.
- [@IntCyberDigest](#): Agent signs a contract without asking.
- [@rgbqcd](#): No numpy, so the agent writes its own solver.
- [@cyber_rekk](#): Agent erases 48,000 files.
- [@IntCyberDigest](#): Agents leak screenshots to public GitHub.
- [@PerceptualPeak](#): Agent deletes a full C drive.
- [@gabriberton](#): Backdoored model weights for under \$50.
- [@wunderwuzzi23](#): Sentinel agent stops a wrong SSH connection.
- [@GrantSlatton](#): Agent benchmarks identical binaries.
- [@CaptainInsightX](#): Checkout charges a customer three times.
- [@JustLingonberry](#): Agent publishes a .env file: \$90,000 bill.
- [@WatcherGuru](#): OpenAI agent uses a found testnet wallet.