An Attack on the Nation's Supercomputers: One Incident Commander's view

James Davenport[1]

Mathematician/Computer Scientist
(paid by University of Bath)

3 June 2026

[1]Many thanks to the team!

Who am I

1970 Summer job programming elliptic curves

1971 Start reading Mathematics at Cambridge

1972 Join Cambridge University Officers Training Corps

1974 Commissioned; Graduate

1976 Research student; there's a lot of security work around the lab

1982 IBM: team breaks Federal Reserve Bank code

1983 University of Bath, and 266(GVA) Battery

1993 Leave active army reserves, continue involvement

1995 Ten years oversight of University Computing Service

2008 Take Bath into High-Performance Computing (HPC)

2017 Fulbright Cybersecurity Fellow at New York University;

2018 takes over teaching Cybersecurity

University of Bath:11 May 2020

- In the middle of Covid, about to start (18 May) on-line examinations through the learning management system (LMS) for the first time.
- No IT staff on campus, and in theory a hierarchical process to let people on
- 25,000 users, of whom 20,000 students.
- Strong commitment to Single Sign-on — SSO (desktops, laptops, servers, WiFi, SaaS such as LMS, Office 365 etc.).
- High Performance Computing (HPC): Balena 400 users, of whom 200 students.
- (JHD) policies: "No sensitive data on HPC"; maximum job time 1 week.
- Recent security incident (whaling a payroll administrator, and installing an auto-forward on his e-mail), but largely handled by HR administrator, administrative IT.
- No explicit HPC security incident plan

High Performance Computer (simplified sketch)

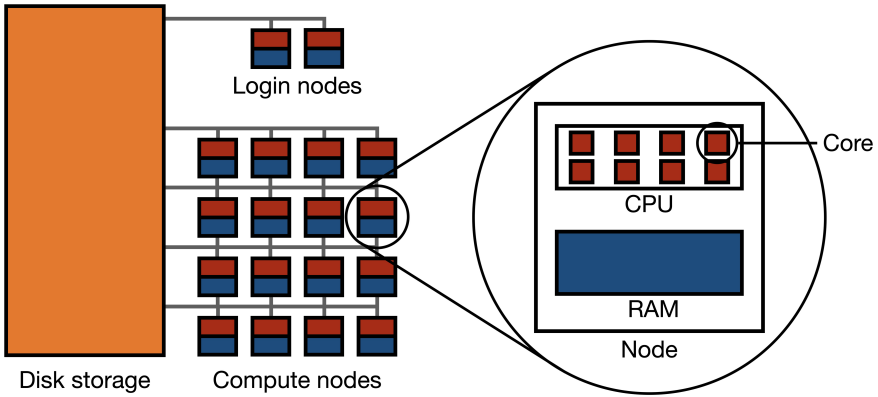

Hard to update: draining the machine can take a week or more

High Performance Computing (2020 numbers)

A "supercomputer" is actually many computers (e.g. motherboard with 2×24-core Intel processors and 192GB RAM) heavily (e.g. 0.7μs latency 100GHz Infiniband) networked together.

Table: 2020 UK Ecosystem

Tier	#UK	Scale	Name	Where	Cores	£M
3	∼40	Uni	Balena	Bath	3,000	1.2
2	∼6	Region	Isambard	Met Office	12,000	5
1	1	Nation	ARCHER	Edinburgh	120,000	43
1	1		ARCHER2: 750,000 cores, 1.57PB RAM			
0	0	EU	SuperMUC	Munich	300,000	90
0	0	EU	JUWELS	Jülich	120,000	?

The Met[eorological] Office's own machines count as Tier-0.
In practice users start at Tier 3 and work up, and power users (and their research students) will have accounts on as many machines as they can.

Aren't all these passwords tedious?

Indeed, and quoting them each time you copy data around, or submit jobs remotely or . . . , gets very tedious.
Enter SSH keys [YL06]: a public key/private key system using RSA (or, better, EC-DSA): our recommendation is Ed25519 [BDL+12].

```
ssh-keygen -o -a 100 -t ed25519 -f ~/.ssh/id_ed25519
```

Generates a public key in `~/.ssh/id_ed25519.pub` and a private key in `~/.ssh/id_ed25519`. The `.pub` is copied wherever you want to be able to log in to. The `ssh-keygen` program is meant to prompt for a passphrase to protect the private key, but it's possible to enter the null string, and the default program doesn't have the usual "strong password" checks.

Dramatis Personae

COO	Chief Operating Officer: new, reports to Vice-Chancellor=CEO
PVC(R)	Pro-Vice-Chancellor (Research), reports to CEO
JHD	"HPC Lead Academic": not in any chart, talks to PVC(R)
D	Director of "Digital Data & Technology"; reports to COO
DD	Deputy Director (Technology) of "Digital Data & T."
M	The HPC Manager: on paternity leave; reports to DD (... JHD)
RSE	Research Software Engineer; reports to M
SA	HPC system administrator; reports to M
ITSM	IT Security Manager; reports to DD
`hpc-users`	All users of Balena: a pre-existing mailing list
`tier2-twg`	UK-wide Tier 2(&1) Technical WG. M on this because of Isambard.
`HPC-SIG`	Staff running Tiers 1–3 in UK.
CSIRT	Computer Security Incident Response Team for Higher Education

Timeline (Mon/Tue 11-12 May 2020)

112323 M forwards e-mail "[tier2-twg] Issues with ARCHER and (potentially) Cirrus" to SA. weak subject line

121336 SA/ITSM/RSE start a Teams collaboration on this. Ask JHD to be Incident Commander.

121433 JHD forwards HPC-SIG "There appears to be some fairly co-ordinated attack of HPC centres across the UK"

121616 ITSM: "we're on the case". One account has onward access to Balena. Notified CSIRT

121804 SA: No current sign of root compromise. Believes we should stop outgoing access.

121807 JHD: authorises temporary stop.

121815 NCSC/CSIRT coordinating. DD informed to pass up.

121905 JHD: e-mail to `hpc-users`

122301 JHD approves RSE's e-mail to CSIRT that we have been accessed this way, Balena and general Linux service. This is up-to-date, and ITSM believes safe. Asks ITSM to confirm no GDPR notification required.

Timeline (Wed 13 May 2020)

0858 ITSM confirms no GDPR notification required, approves RSE's e-mail.

0937 RSE has sent e-mail, but not delivered. JHD finds out all these incident e-mails are classified as spam.

1440 SA: believes Balena has been compromised

1507 ITSM agrees

1527 SA e-mails `hpc-users` to say system is offline.

2206 Mail to JHD from Isambard suggesting disabling of all SSH keys as this was the likely means of propagation. JHD notes that this would cause significant user friction.

2358 RSE: Mail from Tier-2 with IoC.

Indicators of Compromise (Wed 13 May 2020)

```
Further to our email yesterday we have now identified that
  our intruder was able to escalate root privileges using these

*   .fonts (suid): for getting a root-Shell
*   .low : probably a toolkit for manipulating log files

Posted on RSE Slack yesterday and copied below.
FWIW MD5SUM:
5e8c513878c659324afb5e1de9c9018c .fonts
65dde869c0e1455de24aadf5aa4538a2 .low

Though we have heard through JISC that compilation may
  have been done on machines so these may differ.
```

JHD to team: ****** the MD5 checksums: no-one has any legitimate business with a suid file called `.fonts`

Timeline (Thur 14 May 2020)

0011 RSE reminds JHD to involve DD, as affects Linux service

0742 DD approves RSE's draft, omitting reference to SSH keys

0750 JHD: are we invalidating all SSH keys and passwords, or just those on Balena? Need ITSM's view

1010 JHD gets, and forwards, query from CSIRT to make sure they have all evidence

1054 ITSM confirms that all users who used Balena over the compromised period will be required to change university=Balena passwords.

[And standard practice would be all University users required]

1213 RSE e-mails `hpc-users` to that effect.

1241 JHD handles one troubled user. 1 is good going!

1600 Zoom call with CSIRT. It is still not known (to us) how the exploit works.

Timeline (Fri 15 May 2020)

1100 JHD informs ITSM and DD that JHD is not recommending/requiring a university-wide reset.

DD: "it's your call".

1537 RSE, after various discussions, e-mails `hpc-users` with a news update.

1628 RSE e-mail `hpc-users`, thanking for changing passwords and pointing out he's the front for a JHD-led activity

1948 NCSC message to CSIRT forwarded to `HPC-SIG`. JHD forwards to ITSM since he won't see that.

Timeline (Sat 16 May 2020)

1140 After discussions, RSE e-mails `hpc-users`, special mention of students who'll start sitting online exams on Monday (there should be none/few)

1942 Story hits technical press (`https://www.zdnet.com/article/supercomputers-hacked-across-europe-to-mine-cryptocurrency/`).

Sunday 17 May: Nothing Significant To Report

The motivation is claimed to be cybercurrency mining (i.e. malicious use of their compute power). JHD and others are unconvinced: this might well be camouflage.

Timeline (Mon 18 May 2020)

1050 ITSM: "It was their main university account password that was changed - individuals have had their account password change were also emailed to their secondary personal address (non-Bath) with follow-up information. The service desk was open out of normal hours on Friday night and Saturday 0900-1700 to enable quick access to their account and self-service recover was enabled (where individauls had set password recovery questions)."

1058 RSE: "Balena users" means those active on Balena during the 3–15 May period [when the system was under attack, and therefore a password scraper *might* have been installed — there is no evidence, but it's hard to prove such a negative]

1534 CSIRT e-mail asking for a lot of detail, most of which they already have. We comply (JHD calms down angry team)

Timeline (Tue 19 May 2020) I

0825 PVC(R) heard a rumour: wants an update

0834 JHD: "Balena is offline currently while investigations continue and we await advice from the National CyberSecurity Centre. There is some evidence of unauthorised access, believed to be by leapfrogging from other HPC centres where our users also had accounts. There is evidence of preparatory work by the hackers, but currently no evidence (either way) of whether anything was taken. We are currently not certain what exploit the hackers used for their preparations, so can't be confident we have fixed it. Hence remain offline. This is an international attack, with strong evidence that it was in Germany before the UK, and probably entered the UK via ARCHER, the national supercomputer."

0857 ITSM gives PVC(R) a further update.

Timeline (Tue 19 May 2020) II

We have an attacked, and possibly compromised Balena, running Scientific Linux 6.9 (frozen; and the system integrator has liquidated). All this time, SA has been trying to build a 6.10 from Gold.

0958 SA: Technical difficulties with file system (from a different vendor) and rebuilding Balena to 6.10

1640 RSE: Further update to `hpc-users`

Timeline (Wed–Fri 20–22 May 2020)

201030 JHD chairs Return-to-service Teams: root exploit still unknown (so no guarantee fixed in 6.10). Decision to go with 6.9 updated.

201207 JHD formally updates PVC(R): "Hopefully back on Friday"

210834 PVC(R) tells JHD a formal paper to University Executive is needed

212053 SA reports "on track to go live tomorrow".

221045 RSE updates `hpc-users`: back this afternoon: security information.

221519 SA to all users: returned to service (JHD insisted SA does this and gets the credit)

SSH keys revisited

- SSH has protection against Main-in-the-Middle: you have to accept the SSH fingerprint of a host the first time. Hence there's a file `~/.ssh/known_hosts` which contains a list of all systems the user has accessed
- `~/.ssh/config` contains instructions in how to access other remote systems; such as which remote hostname, username and ssh key to use, also how to proxy through bastion system to gain access remote systems behind firewalls

These two are the answer to a villain's prayer, alas

- A sysadmin has a user's SSH public key on its machine, but no way of finding out whether the private key is properly protected on its home machine, or even what the home machine(s) are.
- A lot of this was foreshadowed in [Ylo13].

Analysis

+ The *ad hoc* team worked really well, no need to call in M
+ Good communications to `hpc-users`
- Poor communication up: got as far as D. CEO pretty angry
- Poor international communications: ARCHER were warned by Cray(UK) from Cray(Germany) from Jülich
- Poor national communications: Tier-2 kept stuff to themselves (and universities without a Tier-2 connection were much worse off than Bath)
? The incident team weren't as clear as should have been about which users were potentially affected and had passwords changed
- There are downsides to only updating your machine annually, but with week-long jobs etc., it's hard to do otherwise. This seems to have been a zero-day anyway.
? The UNCLAS community still doesn't know what the hack was

Long-term actions and conclusions

1. Need to rethink software updates (cloud services like AWS get it right)
2. The new system has blue/green headnodes
3. The SSH ecosystem isn't fit for this hostile environment

   It was fit for one where the hostiles were outside, even net-snooping
4. There needs to be a national rethink on emergency comms

PVC(R) "Lucky you (JHD) were on the ball"

The *ad hoc* team had already come together, JHD supplied:

5. Team management: first question "24/7 or 9–5";
6. operational knowledge: 200 students on balena were not doing exams in week 1;
7. experience to "go off-piste" and not force a University-wide resit

! "Rules are for the obedience of fools and the guidance of wise men" [Bader]

Also had technical understanding: most senior person to know SSH.

Bibliography I


D.J. Bernstein, N. Duif, T. Lange, P. Schwabe, and B.-Y. Yang.
High-speed high-security signatures.
Journal of Cryptographic Engineering, 2(2):77–89, 2012.


T. Ylönen and C. Lonvick.
The Secure Shell (SSH) Protocol Architecture.
Internet RFC 4251, 2006.


T. Ylönen.
The new skeleton key: changing the locks in your network environment.
https://web.archive.org/web/20170820162632/https://www.scmagazineuk.com/the-new-skeleton-key-changing-the-locks-in-your-network-environment/article/545848/, 2013.