UNIVERSITY OF CAMBRIDGE
Department of Computer Science and Technology

ITSC_MINUTES_03/10/2024

IT Strategy Committee, Department of Computer Science and Technology 3rd October 2024 at 10:30 am SW00, William Gates Building

MINUTES

Present

Richard Mortier, Chair [RM] Thomas Sauerwald, Deputy HoD [TS]

Mark Cresham, Secretary [MC] Daniel Porter, IT Support Manager [DP]

Tim Jones, UTO Rep [TJ] Sam Nallaperuma, Research Staff Rep [SN]

Nic Lane, GPU resourcing strategy lead [NL] Malcolm Scott, IT Infrastructure Specialist [MS]

*Rob Harle, Director UG Teaching [RH] **Sarah Bainsfair, Financial Analyst [SB]

*Attended remotely

**Attended as a guest

1. **Apologies for absence**

Abraham Martin Campillo, Helen Francis.

2. **Approval of the Minutes of the previous meeting**

Unconfirmed minutes of the meeting held on 20th June 2024 were approved.

3. **Matters arising**

None

4. **Actions from the previous meeting** (i) GPU Upgrades

MS - reported the purchase of 2 servers, 4 large GPUs, and 8 low-power GPUs. A lower end server will be used for development, with flexible GPU configurations.

MS - applied for EPSRC funding for a GPU pool. Offered funding covers 1 server with 4 GPUs, but it’s challenging to integrate with the current setup for development purposes.

NL - suggested a small cloud GPU trial, which could release in-house GPUs for PhD students and increase flexibility. PyTorch Lightning on AWS was proposed as a trial platform, but concerns about costs (£140/month per user) were raised. A potential trial with Part II and Part III students was discussed.

SB - highlighted financial constraints around cloud services, as these cannot be capitalised and would need to come from current budgets.

RH - identified that GPUs are a high priority for all student groups, including undergraduates, ACS, and PhD students.

The feasibility of managing accounts was discussed, with the Undergrad Teaching Office being identified as the potential candidate responsible for creating user accounts at the start of the academic year.

Tracking usage and obtaining test credits from companies such as AWS was discussed.

<u>ACTION</u>

MS and NL will discuss and decide whether to implement a trial for a full course or limit it to select projects.

(ii) Recovering of Finances

DP - discussed the visibility of PhD laptop costs, noting that SB has separated these into distinct codes for better clarity.

The HoD has requested the development of a clear policy for recovering these costs from grants.

5. Standing items

RM - noted that AM would provide an update on the door locking system at a future meeting.

6. Main business

(i) ITSC membership and process

Members were reminded to review the terms of reference and to note that some budgets would expire soon.

Discussion on whether a Tech Committee member should be included on the committee. RH and DP confirmed their involvement in the Postgraduate Forum.

(ii) Hybrid Meeting Rooms

A smaller hybrid meeting room setup is being trialled from UIS; that supports Teams, Zoom and Google Meet via Zoom

The smaller setup is being trialled in a medium sized room, against advice from UIS and IT Services.

If unsuitable, the equipment will be repurposed in a smaller room.

<u>ACTION</u>

DP will pick up the conversation with UIS regarding the hybrid meeting room trial.

(iii) State of the Nation

GN09 Space

MS - provided updates on the GN09 server room, which is reaching capacity both in terms of space and power.

Options for relocating or retiring older servers were discussed.

The possibility of developing a policy for managing and decommissioning old hardware was raised, with MS acting as the gatekeeper for GN09.

MS - noted that we have a tentative offer of space from Kingston, but it only meets some of our needs, and users are unable to make modifications.

The committee supported that requests for the addition of GN09 equipment must be authorised by MS.

The possibility of machine logging was discussed.

Legacy Services

Ongoing efforts to reduce unnecessary services were discussed.

MS - identified legacy mailboxes as a particular service that could be deprecated.

Discussions included the financial implications of running legacy services, particularly in the context of supporting more suitable services.

ACTION

DP and MS will compile a list of legacy services for review.

Network Upgrades

The committee discussed the upcoming replacement of network equipment, which will cost around £1 million.

MS - proposed a phased replacement plan, starting with firewall upgrades costing £75k, with the possibility of spending up to £100k. The committee supported the expenditure.

ACTION

MS will start planning of the firewall replacement.

Separation of Department Websites

RM - raised concerns about managing the departmental (cst) and teaching (cl) websites, which are currently running in parallel.

Concerns regarding a single point of failure for the teaching website (CL) were discussed.

The discussion of alternative teaching systems was considered. Moodle was identified as the recommended but imperfect solution for teaching resources.

The committee agreed that the teaching site (cl) should become self-contained and thus need not provide the same 'look and feel' as the main departmental (cst) website.

ACTION

RM will inform mgk25 of the plan to separate the (cl) website from (cst) website.

Extra Funding for COs

NL - suggested revising grant overhead percentages to better reflect the rising costs of running departmental services, particularly for COs.

A blanket increase was considered but noted that it might depend on the type of grant and funder.

7. Any Other Business

None

8. Date of next meeting

(i) Date: 2nd December 2024
Time: 14:00 - 15:30
Location: SW00

ITSC “State of the Nation” report - Support

October 2024

Daniel Porter

Overview

We’ve made good progress streamlining many support processes and I’m trying to instil a general “philosophy” that should guide most aspects of our work. This is to focus on delivering scalable services and to make use of standard offerings where applicable to create time to deal with bespoke needs. We are also taking the time to query whether an existing process is still the best way of proceeding.

Recent improvements have made a large impact on our general workflow and freed up more staff time across the team to allow us to get a better grip on the workload and take on more pro-active projects and process improvement.

Progress

Taken from ‘Summer Update 2024’ e-mail distributed to Department

Team name

Firstly, we would like to launch the new name for our team – **IT Services**. This was formally approved some time ago, but we are choosing to launch it now alongside some other revamped services. We feel that this name is now a more accurate reflection of the function being performed by our team in the Department. We understand that “sys admin” has been a long-standing term of reference for us and we know this is likely to continue amongst those who are familiar with it, but formally we are now known by a new name.

Ticketing system

We have made several changes behind the scenes of our ticketing system to enable us to better manage incoming user support requests. This has involved creating multiple queues for different types of work as well as a triage process to ensure that tickets end up with the relevant people in a timely manner. We are still iterating these improvements, but we are hopeful it will lead to a much better service and quicker response times.

We also have a new helpdesk address for IT enquiries: service-desk@cst.cam.ac.uk

E-mails to existing addresses will continue to function, but the new address will be publicised going forwards.

Lecture Theatre upgrade project

The aging AV equipment in LT1, LT2 & FW26 has been completely replaced with a standard University specification designed by the UIS's AV team. This brings the Department in line with teaching facilities offered around many of the University's shared spaces such as the West Hub and the Titan Teaching rooms, as well as many other Departments.

In addition to replacing the old equipment, we are also launching new AV support processes. A dedicated AV emergency phone number is provided on signage in each Lecture environment to enable presenters to reach the IT Services team in the event of AV issues that are preventing sessions from proceeding. We also have a webform available to report any other minor issues that occurred, which should help us gather the information required to investigate them.

Full information about the Lecture Theatre equipment and support processes can be found on our dedicated webpages: https://www.cst.cam.ac.uk/av

Printers

We have now completed the second phase of the Department's DS-Print roll-out, with four new devices including two in corridors/alcoves and one in the GC18 Reprographics room

Full information about the DS-Print offering in the Department, including a full list of devices locations & installation methods can be found on our dedicated webpages: https://www.cst.cam.ac.uk/dsprint

Poster Printing

We have simplified our poster printing offering and created a new webform to enable users to easily submit their requests for printing rather than e-mailing them.

Full information about the Poster Printing service and a link to the request form can be found on our webpage: https://www.cst.cam.ac.uk/local/sys/printers/posters

Device Purchasing

We are working on the streamlining equipment purchasing process, particularly where only non-specialist equipment is required. As part of this we have built some webpages to outline our standard offerings and will be providing webforms to enable people to submit their orders in an easier fashion and help us process them quicker. Some areas of this are not yet complete and will be rolled out gradually throughout the year.

The purchasing webpages can be found here: https://www.cst.cam.ac.uk/local/sys/purchasing/

Main Challenges

There is still a large amount of technical debt on legacy systems as well as some issues with non-common processes across the team.

Technical

- Database (aka dbwebserver) ○ Collection of bespoke web apps & services
    - Taking opportunity to re-work parts where possible onto more sustainable platforms
- Support Services ○ Aspects were neglected in the past (e.g. PSS Support) ○ Formalising processes for support & devices to help combat this
- User Admin processes ○ Need to fully re-think onboarding & offboarding processes
    - Should be a collaborative effort with other offices
- Standardisation ○ Lack of uniformity across many systems in different ways causes problems
    - Taking the time to address these to create an environment that's easier to understand
    - E.g. Mailing lists, Shared Mailboxes, Accounts, Device setup

Team

There has been a slight disconnect within the team with historical processes appearing to be very siloed and where people have developed their own processes/procedures.

There is a planned office reshuffle which will bring the three team members together in a single office which I hope will create a greater unity, more opportunity for collaborative working and process uniformity.

This will also provide a single physical point of presence for users looking for in-person support.

The increased size of the team has made it easier to now identify and address single points of knowledge and ensure processes can be performed by multiple members of the team.

Future Plans

Small selection of aims for the near future

- Recruitment ○ Hoping to replace ex-Chris Hadley post with new Windows SysAdmin
  ▪ Parallel post to Piete Brooks replacement
- SharePoint ○ Creating sites for [initially] PSS to host shared team data ○ Better for hybrid working ○ Removes reliance on filer and issues with permissions
- Department Website ○ Review content & platform migration ○ A good catalyst for updating at our own online documentation
- Purchasing ○ Creating a policy for equipment provision to enable expectations to be set ○ Working with Finance Analyst on budgeting process & financial planning
- Printing ○ Complete move to DS-Print (pending successful Linux testing)
- Hybrid Meeting rooms ○ To provide Teams & Zoom capable hybrid meeting facilities ○ Test pilot of one room (GC15?) to see viability
- Domain Joined laptops ○ Will enable Department accounts to be used on staff laptops ○ User files automatically backed up/synced via OneDrive
- WebApp server ○ Joint venture with Infrastructure Team ○ Aim to provide a place for the teams to serve apps to users ○ E.g. Password

collection/resetting, VPN credentials ○ May play a part in replacing components of dbwebserver

ITSC “State of the Nation” report: Infrastructure October 2024

Malcolm Scott
26 September 2024

I will take the opportunity of a change to the committee membership to give a broader overview than usual of the services provided by the Infrastructure team, their current work and future plans, as well as infrastructure services that we procure from elsewhere.

Many of the services in our portfolio do not need a discussion in the meeting, but their existence and status is noted here for reference.

Apologies for the length of this report – even so, there are some services we provide and ongoing projects in my team that I have not mentioned.

Strategy for our overall service portfolio

The Infrastructure team, and the IT team collectively, runs a lot of services. This practice dates back to our close involvement in the early development of computing when we were amongst the first to have to figure out what we could do with a networked computer, what would be needed to provide people with access to it and what services we could run on it, in the process writing our own code and implementing novel systems to achieve this. Some conventions still evident today in our systems predate UNIX. We were (or rather Piete was) amongst the first to configure a modern email service akin to what we have today. Many of these technologies of which we were early implementers or adopters have since become commonplace in some form in every organisation and are now available as a standard off-the-shelf service from a range of different suppliers.

We should not keep running these services just because we have always run them, even if we are quite attached to our implementation – we should focus our time on the things where we *do* still have a unique need due to our ongoing work at the forefront of computer science, that nobody else can do adequately for us.

In many cases there is in fact an existing solution available within the University; we are not the only department to use IT, and there are standard University solutions to many common IT requirement. Even in cases where our in-house service was the only good option very recently, some central services are catching up and overtaking ours, and we should be open to the possibility of new opportunities to make use of University-wide facilities that allow us to further focus our time on more needs specific to computer science research and teaching.

We should also be open to possible off-the-shelf solutions which although not *perfect* for us, are nevertheless *good enough*. When we were first to implement a service (email, for example), we may have made choices that turned out not to match what the rest of the world later decided to do (for example, email sub-address formats, or moderation practices for mailing lists). Systems that have grown with us for decades will inevitably be ingrained into our working practices and expectations. But we should look at what other departments and organisations are doing, and where we see a difference in our processes that means they can use a standard system whereas we have something bespoke, critically examine that difference. If we can change our practices to better align with the rest of the University or the world, that will make us more efficient, aid our working with other institutions, and again free up time that was locked up in the maintenance of the bespoke system.

Our overarching IT strategy is to avoid wasting time and resources by running services that duplicate others available to us – to let other people solve our problems when they can – so that we have time to specialise.

(Though of course we must also accept when things move in the opposite direction: services that once were commonplace can become niche. This is what led to UIS's closure of MCS Linux, for example: other departments and colleges stopped wanting it, leaving us as the last major users. In such cases we must accept that we need to bring newly-specialist services in-house, when we are the last people who still care about something.)

Staffing

Piete Brooks has semi-retired and is now part time (working the equivalent of one day per week). Recruiting for his replacement should be underway very shortly; at the time of writing we are almost ready to advertise the post. Piete's ongoing focus will be to pass knowledge of our myriad bespoke systems on to the new staff member, and the new staff member will focus on modernising and replacing some of those systems, as well as supporting me to maintain our infrastructure services.

I would also like to recruit a specialist Research Support staff member with HPC experience, who would focus on the complex needs of researchers (including ACS/Part III students working on their projects) who need either bespoke solutions or help to make use of complex systems already available to them. My hope is that this will free up some of my time to work on outstanding core infrastructure issues, and also to reduce the demand for yet more expensive local resources such as GPUs by more effectively utilising what is already available. Regardless, this will allow us to devote more time to helping researchers.

Core departmental infrastructure

Networking

Core network

The department's core network consists of two pairs of 10Gbps switches (Cisco 4500-X), one named "gatwick" in the William Gates Building (GN09) and the other named "heathrow" in the University's West Cambridge Data Centre (WCDC). Both pairs of switches have historically performed three functions in tandem (effectively providing four-way redundancy):

- Routing for every logical partition (VLAN) of our network

- Firewalling between untrusted and untrusted VLANs, and between the internet/University and our network
- Aggregation of the connections to wiring cupboard switches and data centre rack switches

The Cisco 4500-X platform reaches end of support on 31 October 2025, and after that point there is a risk that a vulnerability will be discovered that Cisco will not fix and that we cannot work around. We need to plan on procuring replacement switches. These are costly and have a long lead time (in some cases a year or more). We have not planned the replacement yet. The cost will be in the region of £100k.

We have already ended our support contract for the current hardware, because it no longer offers good value: we are unlikely to need software support as the setup has been stable for many years and there will be no more major updates release by Cisco, and we can protect ourselves against hardware failures more cheaply by maintaining a pool of spare second-hand hardware.

At the same time I am restructuring the network. The current combined functions above – and especially using switches as firewalls – seriously narrows down our options for replacements (for example, many otherwise-high-end switches do not have enough capacity for ACLs, which we are using to implement our firewall policies), and has in any case resulted in performance issues on the current hardware as we are stretching it to (or past) the limits of its intended capabilities.

The current highly-centralised routing architecture in which every VLAN must extend both to GN09 and WCDC, and a lot of traffic must pass over our inter-site links, is also not optimal either for performance or for reliability: a single fault on one switch can take down (and has in the past taken down) our entire network.

I have already begun decentralising routing to an extent. Some parts of GN09 are now routed using their local distribution switch rather than by gatwick and heathrow; this means that a fault downstream of these distribution switches would be much less likely to affect other parts of the network, and network traffic between systems in GN09 would never have to go over the inter-site link to be routed by heathrow before coming back again. There is a lot more work to do before these goals are fully realised however. And there are trade-offs, most notably that it will no longer be possible to have systems in offices on the same VLAN as servers in data centres, and there will now always be multiple routers in between offices and servers – but this is not as significant an issue as it historically was, as practically all network protocols these days will cope with intermediate routers. It will also mean that devices must change IP addresses when they move to a different part of the network; I am considering ways that the operational overhead of this can be reduced. Subdividing the network means also that our IPv4 address space becomes more fragmented, and thus slightly less efficiently used, though we are lucky to have plenty of IPv4 addresses so this is probably not going to be a serious problem. In any case I am also taking opportunities to free up IPv4 addresses assigned to legacy uses (e.g. DTG and SRG VLANs) and to switch certain systems completely to IPv6 where this will not negatively impact functionality.

Border network and firewalls

The border network (i.e. connections to the University network) and firewalls are currently the same as the core network, but as part of the aforementioned core network project I plan to separate these functions onto dedicated hardware. I am still considering options and don't yet have a firm plan, but one option I am considering is to implement new firewalls in software on general-purpose servers rather than on switch/router hardware or on a dedicated firewall appliance. Benefits might include:

- Better value for money on the purchase (firewall appliances get extremely expensive if resilience across multiple sites is needed, and also if high speed links such as ours are needed)
- Very much lower ongoing support cost from the manufacturer
- Very much shorter lead times
- Much lower time to procurement, as we would not have to evaluate several manufacturers' appliances before committing
- More local development effort required, though it better aligns with our existing skills (we have no experience with commercial firewall appliances)
- Future upgradeability: individual components such as network cards could be added or upgraded, or the software could be moved to a new server without impact, compared with the large amount of work required to move to a completely new hardware firewall appliance which might be managed in a completely different way
- Risk of unexpected bottlenecks: the performance of (for example) Linux's software firewall code on modern high-speed server hardware is not fully known

This approach will cost approximately £34k initially, though may need additional servers.

We have not yet selected a software platform for the firewalls. The plan had initially been to evaluate VyOS, an open source network operating system which we are already using for some more-minor functions, for this purpose; however I have concerns about the sustainability of the VyOS project and we may need to seek an alternative. Arista also has a version of their network operating system that can run on general-purpose hardware, which will need further investigation.

Office wired network

The office network consists of 38 Cisco 2960-X switches. Cisco has announced the end of support for this platform by 31 October 2027 and we should ideally plan on replacing them before then, although the actual risk of their end of life leading to a security incident is relatively low as security issues would most likely affect only their management interfaces, which are reasonably well protected from attack by being on a restricted VLAN. As above, we have ended the support contract already and already have a stock of spare hardware. Due to their age, these switches are starting to fail but used spares are not hard to source.

One option could be to not replace them, but to instead opt into UIS's Network Defragmentation project, which would involve UIS replacing the switches at their expense as well as managing them. That project had initially targeted "simple" networks; our requirements may be more complex. It is unlikely that an outsourced network would be suitable for the diverse needs of research servers in our data centre, but it is possible that it would be suitable for offices.

If we instead have to buy our own switches, the approximate cost would be in the region of £200k for a like-for-like replacement (1Gbps ports); however a few researchers are starting to ask about faster connections already. Due to the long lifecycle of switches, we should consider how our needs might increase over the next decade. However, providing office ports faster than 1Gbps would be more costly.

Wireless network

The wireless network in the William Gates Building is fully managed by University Information Services. They are currently undertaking a major project to replace it entirely, with new access points and a higher density to improve coverage. For the first time this will involve some access points being installed in offices (on the ceiling), rather than just in corridors. Every new access point will have new network wiring to current standards, to support future upgrades with higher speed uplinks from access points.

One of the wireless networks operated through UIS's platform, "wgb", is an open network for visitors. As unauthenticated guests should not be allowed access to the University network or the national academic network, the "wgb" wifi is routed via a commercial cloud service provider (Mythic Beasts) where we run a virtual router. There is a risk that anyone near the building could anonymously use this network for illegal or immoral activities, which would ultimately be traced back to the department. We have so far had one report of a user allegedly illegally downloading copyrighted material, which was detected by the copyright holder's agent and reported to Mythic Beasts. At the same time, UIS's offerings for ad-hoc visitors are improving; visitors can now self-register for internet access via UIS's UniOfCam-Guest network. We could consider withdrawing the "wgb" network in future to reduce the risk, though we know that it does get used by members of the department as an easy option for quickly getting online without having to register, and for Internet-of-Things devices.

Data centre networks

The GN09 server network comprises:

- 35 Cisco 2960-X rack switches for 1Gbps server connections and infrastructure monitoring and control; end of support: 31 October 2027; replacement cost: ~£170k. As above, ideally we should plan on replacing them, though the likelihood of their end of support leading to an actual security concern is relatively low.

- 5 Arista 7050 rack switches for 10-40Gbps server connections; end of support: 13 December 2025

- 2 Cisco 4500-X rack switches for 10Gbps server connections; end of support: 31 October 2025

- 2 Arista 7050 aggregation switches providing the uplinks for all of the above and interconnection to the rest of our network; end of support: 13 December 2025

The approximate replacement cost for the 9 faster switches will be in the region of £450k. We can sweat the assets for a while as the switches are functionally perfectly adequate for several more years, though there is a risk of a security-critical issue that the manufacturer does not patch. Additionally, it is likely that over time 1Gbps will be seen as inadequate for a research server's network connection, which could mean that we need more of the faster 10Gbps+ switches in future (and fewer of the slower 1Gbps switches).

Our network in WCDC has a further two Cisco 2960-X switches which also need replacement within a few years as above. (We also have two much newer Cisco 9300L switches in WCDC bought for a recent expansion and serving the Morello project, which will not need replacing for a long time.) There is currently no separate distribution switch in WCDC for higher-speed server connections or other switches; we use our core/border switches (heathrow) for this purpose – but may in future separate this function onto dedicated hardware.

Out-of-band emergency network access

We operate an out-of-band router for emergency remote access to critical infrastructure, allowing some faults to be investigated and mitigated without physically being in the building even if the whole network is nonfunctional. This uses the aforementioned router at Mythic Beasts and the O2 4G network to provide a link into a few critical pieces of equipment that does not rely on the University network or our own network.

Uninterruptible power supplies

Each wiring cupboard has a small battery-backed UPS to keep the office network running for at least a few minutes in the event of a power interruption; these are undergoing a gradual rolling replacement programme to replace obsolete hardware (and at the same time to replace worn-out batteries). Recent electrical work in the building has served as an opportunity to test these; usually each time we do electrical work we find one or two more old UPSes that were not working optimally, and those are prioritised for replacement.

Our departmental data centre GN09 has two UPSes to power critical infrastructure and a subset of our servers. The larger of these UPSes will power roughly half of the room for about 10 minutes. The smaller one only powers critical infrastructure and should last for about an hour. We have a maintenance contract for these, with the supplier regularly servicing and testing them. A full battery replacement in the large UPS will take place within the next couple of months.

Unlike most other data centre operators such as UIS, we do not have a diesel generator to power our data centre for longer periods. However, a couple of years ago we undertook a project to allow for temporary connection of a generator, and perhaps a future permanent generator. This has been used to good effect multiple times recently to keep critical services running during electrical works that required isolation of the main switch room.

We are aware that the Roger Needham Building next door has a generator for their data centre, but the data centre is largely unused and hence the generator has spare capacity. We are attempting to explore the possibility of having an electrical cable run across the room so that we can make use of the generator there (and also their separate electrical supply, as we could make use of that without starting the generator during work within the William Gates Building).

In any case, however, we do not have a way to power the cooling system (chiller, water pumps, water/air heat exchangers in GN09) from a generator. This limits the utility of a generator, as we would have to shut down most servers anyway in order to control the room temperature. However it would enable us to keep our most-critical infrastructure running essentially indefinitely.

External connectivity

For completeness: our connectivity to the rest of the University and onward to Janet and the internet is provided (for a fee) by UIS; the form in which we receive these services presents us with two 10Gbps fibre links to GN09 and two to WCDC, both via the Granta Backbone Network, with each pair terminated at the far end on two different UIS routers for resilience. Our routers negotiate with theirs via BGP to route IPv4 and IPv6 packets between our network and theirs.

We also use the Granta Backbone Network for private "dark fibre" links between the two sections of our network; the William Gates Building and WCDC. Although originally set up as two 10Gbps links using two fibre pairs, we now operate them as four 10Gbps links using each fibre bidirectionally, so that heathrow, gatwick and the GN09 data centre aggregation switches can form a resilient ring topology.

Core network services

DNS resolution (recursive)

We operate multiple DNS resolvers, as Ubuntu servers (mainly VMs) running ISC BIND. In recent years we have improved the resilience of these, by arranging for the main DNS IP addresses to transfer automatically to a known-working server (using VRRP) in the event of an outage, to avoid clients having to wait for a timeout when the first DNS server is not available. This has been working well.

DHCP

We are gradually expanding DHCP provision on our network (as some VLANs did not provide DHCP for historical reasons). This will allow us to more easily make changes to the network in future, as most clients can automatically pick up our changes within a few hours, and also helps to simplify the PC deployment workflow for the support team.

The department's main DHCP servers are rather old hardware and are in need of replacement (the CPUs have known vulnerabilities, though they are not exploitable; the age of the servers also means that they are harder to remotely manage as they have obsolete BMCs). This will not be especially costly as DHCP will run on a very low-end server.

Additionally, the DHCP server software that we use is ISC DHCPd, which although a very widely used open source product has been discontinued by ISC. For the moment it remains supported by Ubuntu. But eventually we will need to move to alternate DHCP server software, which will involve some development effort as we have a lot of automation and configuration specific to ISC DHCPd which will need rewriting.

Some parts of the network are experimentally using various other DHCP servers, generally because it was not practical to extend the main DHCP servers' network connections to particular self-contained networks; the requirement for this is reducing as we have recently implemented support for DHCP relaying and we are likely to centralise DHCP further.

Our allocation and management of IP addresses currently uses a rather archaic process. UIS is currently evaluating IP address management products for their own use, and I plan to make use of the results of this evaluation to see whether we can implement a better system ourselves.

Network boot, deployment

GN09 has a deployment system (originally built for the Model Data Centre) to allow easy remote installation of operating systems on servers. The recent consensus amongst the IT team is that using this for desktops in offices as well would be useful, which already works to an extent. We are considering how support workflows can be best adapted (and the deployment system adapted) to better make use of this.

DNS hosting (authoritative)

The department's DNS zones (cl.cam.ac.uk, cst.cam.ac.uk, etc.) are primarily hosted on two recently-upgraded Ubuntu servers running ISC BIND, aided by various other institutions providing resilience by hosting a mirror of our main zones on their own DNS servers.

The way that we maintain the content of the zones is somewhat archaic by modern standards; we have long sought a better alternative but have yet to find one that meets our needs particularly well. We remain on the lookout for a good product.

Our DNS zones are secured using DNSSEC; we were a fairly early adopter of this technology, and since then standard practices and cryptographic algorithms have evolved, and our setup is not quite up to modern security standards. The risk remains very low at present but at some point will modernise this setup.

Infrastructure monitoring

We use Observium Professional (with some local modifications) for infrastructure monitoring, logging and graphing. This has been working well, although the effort required to maintain our local modifications means that we are not always running a current version.

We use Nagios for alerting, some of our instances of which are still running an obsolete version (Nagios 3) and will need upgrading or replacing. We may take the opportunity to revisit how we configure it. It is likely that this will be given as a project to our upcoming new recruit, as it may be a good way to gain familiarity with our infrastructure.

We also use Netdisco mainly to track the history of switch ports in the office network; this is also due an upgrade but has so far been a low priority.

VPNs

We have two departmental VPNs: one run by UIS (their Managed VPN service) and one run in-house. The in-house VPN was originally set up during the first Covid lockdown to mitigate capacity issues on the UIS service caused by the unexpectedly high number of remote workers, but turned out to be beneficial as an ongoing service as it offers some features not provided by UIS, such as IPv6 and split routing. My hope is that eventually UIS will add these features to their Managed VPN service so that we can retire our own, but until then, maintaining an in-house VPN is not proving to be particularly onerous.

Linux LDAP

Our Linux LDAP service enables account information, group information and filesystem locations to be shared across managed Lab Linux systems. It has remained largely unchanged (except for routine operating system upgrades) for some time; however we recently encountered scaling issues as the number of clients surpassed a default limit. This caused some disruption on a couple of recent occasions but is now believed to be resolved.

Windows clients use a different LDAP service (Active Directory). These days, Linux can also use Active Directory's LDAP, and it is no longer really necessary to maintain two separate LDAP services. However the effort involved in migrating Linux clients to Active Directory LDAP would probably be greater than the effort required to continue to maintain a second LDAP service for the foreseeable future.

Kerberos

Managed Linux and Windows clients use Active Directory's Kerberos service. Microsoft has been making some changes during the past couple of years to improve the security of the Kerberos protocol, and initially we experienced major problems with Linux clients as a result – on our first attempt to install the relevant update to the Active Directory servers, Linux clients became unable to authenticate users at all until we reverted the update. After that, we had to put all Active Directory OS upgrades on hold until the situation was resolved.

I am pleased to report that we have since been able to get Linux clients working with up-to-date Active Directory Kerberos, and the Active Directory servers are now fully up-to-date – more on this below. Microsoft has announced that there will be further changes to the Kerberos protocol soon and we will have to monitor the effect of these changes carefully.

In the longer term, there are indications that Microsoft has started to consider Active Directory to be a legacy product, and Kerberos and LDAP to be legacy protocols. Their replacement (Entra) is rather different, and much less useful to Linux clients. We will monitor ongoing developments.

Remote access (slogin)

The slogin service is a remotely-accessible Linux server accessible by anyone in the department, for light use such as managing their data on filer or connecting onwards to other specialist systems.

We have recently made changes to the way we manage and advertise the multiple slogin servers in order to try to make it clearer which one users should currently be using, and to improve maintainability by providing a standby server that we can use as a backup or to test upgrades. (For example, currently the main service is running Ubuntu 22.04, and we have Ubuntu 24.04 on the standby server; when we're confident that that is ready for use, we will swap the standby and active servers.)

However, most users continue to use one standalone server – ely – specifically, because for several years it was the primary or only slogin server. We would like users to switch to slogin.cl.cam.ac.uk and will communicate this when there is a reason to retire ely.

Time synchronisation (NTP, PTP)

We operate several time synchronisation servers for use within the department (and in some cases as a public service), including two "stratum 1" NTP servers which use GPS satellites as a highly accurate time source. Our main "stratum 1" server (a Meinberg M300 with GPS receiver and a highly accurate local oscillator, originally quite a high-end and expensive piece of equipment) is failing due to its age. We have procured a replacement (a LeoNTP 1200), which is a lower-cost solution but should perform perfectly adequately for our needs, especially as we still have a newer Meinberg M600 mainly for servers in GN09 (which also provides a PTP service for research requiring time-sensitive distributed measurement). The LeoNTP 1200 is not yet in service due to the need to mount a new antenna on the roof to get an adequate GPS signal.

The lower-accuracy "stratum 2" NTP service (ntp1[abcd]) used by most general-purpose machines in the department has also been rehomed on various newer devices which are capable of acting as a NTP server in parallel with their main duties, and which have been observed to happen to have a stable clock. The service is currently housed on two TrueNAS storage servers, one of the GN09 aggregation switches and one VM on a particularly stable host; these are split equally between GN09 and WCDC.

SSL/TLS certificates

We make use of both Lets Encrypt (a free public certification authority) and JISC's certificate service (provided by Sectigo, and 'resold' (for free) to us by UIS). Both allow for automated certificate issuance and renewal, though we have various bespoke in-house systems to obtain and deploy certificates on behalf of the systems and services that use them. We are gradually moving to the JISC platform from Lets Encrypt as Sectigo has a simpler system for verifying certificates for private services, of which we have many that still nevertheless need to be secure, whereas Lets Encrypt is simplest for services that are open to the internet.

Virtual machine platforms

Main departmental VM pools

Our main VM pool in WCDC hosts numerous virtual servers for our departmental operations and for research and teaching. Any member of the department can request VMs from us, whether for internal development or hosting internet-accessible applications or any other reasonable use.

We bought a new set of servers in order to move this service off warranty-expired, slow and less-reliable servers last year; however the migration of VMs from the old pool to the replacement has been very slow due to other pressing work and because the VMs on the old pool are rather poorly tracked, with many in unknown conditions and with unknown owners. We aim to better track VM lifecycles on the new pool, which will require some development work. We also aim for VMs on the new pool to be entirely on the new decentralised network setup rather than one of the network-wide VLANs (see ”Core networking” above). In general we have been rehoming our own VMs onto the new pool when they happen to undergo other extensive maintenance.

We have a secondary pool in GN09 mainly hosting a second instance of various core services for resilience – for example users connecting to VPN2 may end up routed either through a VM on the main pool or another on the secondary pool. The secondary pool also hosts the management and monitoring systems for GN09 itself. The VM hosts were last replaced in 2022 (a quicker project than the main pool due to the smaller number of VMs) and the current servers are functioning well. However the VM disks are still stored on older servers purchased in 2015 alongside the previous VM hosts.

All our VM pools run XCP-ng, which is a free and open source fork of XenServer. Commercial support is available for XCP-ng from its main developer, Vates; we do not currently pay for this as we have the expertise in-house to support it ourselves, but it remains an option for the future.

UIS Infrastructure-as-a-Service

UIS has recently started to offer use of their own virtualisation platform, IaaS, to us for our own VMs. This service is still a work in progress but is in a state where we can start to host some services there. The first such VM on UIS’s platform is one of the new Active Directory servers (see below); with three such servers we can now offer even greater resilience, with Active Directory remaining available even in the event of a hypothetical major outage affecting both GN09 and our WCDC network.

The main caveat at present is that IPv6 does not work well on IaaS due to a bug under investigation. I am assured that they are working on this and hope to have it rectified soon; this matters to us because we have some IPv6-only systems.

VM management: Xen Orchestra

We use the open source product Xen Orchestra to manage VMs across all of our VM hosts and those we run for others. This is also directly accessible by users for self-service management of their own VMs, authenticated via their University accounts (with two-factor authentication). We have developed some local improvements to Xen Orchestra, and have made contributions back to the open source project.

Like XCP-ng, commercial support is available for Xen Orchestra but we do not currently need it; it remains an option for the future.

Storage platforms

Strategy for PSS data

Data associated with the various administrative roles in the department (except for SQL databases; see below) has historically been stored on Filer, as it is the most resilient system that we operate ourselves. However, since that choice was made, cloud services have become available and mature as a viable option, and are becoming the de-facto standard location for such data in the University and beyond due to the benefits they provide, in particular secure remote access to data and better facilities for collaboration.

Furthermore, we know that file permissions on the relevant parts of Filer have not been closely managed, leading to some data being more (or occasionally less) accessible than it should be.

The strategy across the IT teams is to move administrative data gradually into the University's Microsoft SharePoint tenancy. Daniel is leading that initiative; the implication for my team is that some aspects of data storage will no longer be managed by the Infrastructure team and will instead become the responsibility of UIS and Microsoft, via the Support team.

This will slightly reduce storage demands on Filer, but as that will continue to be used for many other things (research data in particular) we will continue to maintain and upgrade it.

Filer

Filer is a pair of clusters of NetApp FAS2720 storage appliances, optimised for highly resilient data storage replicated in almost real time between GN09 and WCDC for disaster recovery. These were purchased in 2019 to replace an earlier set of NetApp systems. The purchase initially had a support arrangement with Q Associates (now Logicalis) lasting a little over 5 years, so this is about to need renewal. At the same time we are planning a modest capacity increase.

Filer is optimised for resilience, not performance. It performs particularly poorly for workloads involving reading and writing a large number of small files, for example git repositories and compiling large code trees. Additionally we occasionally observe suspected bugs with Filer's NFS implementation leading to even higher latency on some operations. We know that some researchers find Filer too slow for day-to-day use, particularly if using our traditional Linux setup whereby user home directories are stored on Filer. However we think it is still useful as a repository for valuable research outcomes.

We are gradually moving towards configuring desktop PCs and research servers to store home directories on local SSDs rather than on Filer, with Filer displayed as a location on which to keep important documents, though this must be accompanied by careful communication about the implications as data stored only on a single local SSD can easily be irrevocably lost due to hardware faults or user error.

Filer is configured to keep snapshots of old file versions for a long time – initially this was configured as 5 years and was reduced to 4 recently as a stop-gap to free up some space. The Head of Department's view is that keeping old data for overly long is risky due to our data protection obligations and in particular the implications of retaining data belonging to people who have left the department, and we therefore plan to significantly reduce the length of time for which we keep old snapshots – though will still keep them for a reasonable period that we hope will be long enough for lost data to be detected and recovered from. It is likely that we will keep personal home directories' snapshots for much less time than shared research group data. The details are not yet decided and the committee's input is sought.

Archive

Archive is a less-resilient bulk storage server, originally intended for long-term archival of large research datasets though that requirement proved not to be as significant as we expected. It is now used for any data too large to be held on Filer.

The server was bought in 2015 as a Spectra Verde DPE appliance, and since then has had a colourful history befitting the rainbow LED arrays on its fascia. The Verde DPE was Spectra's experiment into the use of hard disks using Shingled Magnetic Recording for archival. They ultimately deemed this experiment to be a catastrophic failure due to a poor interaction with SMR's high write latency and very large block size with their choice of the ZFS filesystem. After myriad issues they eventually replaced all the hard drives with 2.5 Petabytes of new high-end conventional hard drives in 2019. Their software platform, however, remained inflexible and buggy and their support became ineffectual. We allowed the support contract with them to lapse.

The lack of software updates since the support contract lapsed started to become a problem due to lack of support for new protocol extensions, which current versions of Windows have started to require.

Conveniently, although sold as a "black box" (actually bright green) storage appliance with which customers are only provided with a bespoke management application with no access to the underlying operating system, Spectra built it from mostly standard hardware components (aside from a custom LED lighting controller!) and a standard ZFS filesystem, and we have been able to rebuild it around a general-purpose front-end server running the open-source TrueNAS storage operating system, connected to Spectra's disk enclosures housing the drives they provided in 2019. So far the SMB (Windows-style) volumes have been moved across to the TrueNAS front-end, with the NFS (Linux/Unix-style) volumes to follow soon. Besides providing up-to-date software, TrueNAS also enables new functionality such as Kerberos-secured NFS and multi-protocol volumes; users who had been struggling with these limitations of the old server are already benefiting from this.

Like XCP-ng and Xen Orchestra (see above), commercial support for TrueNAS is available from its developer iXsystems but we are not currently making use of this; it remains an option for the future. iXsystems has demonstrated that even in the absence of a support contract, they are willing to help us for free when we encounter bugs.

VM disk storage

For the last few years, the VM disks associated with our XCP-ng pools have also been stored on various TrueNAS servers, albeit on very different hardware to Archive's: VMs on the main pool and on the GPU VM pool are stored on high-speed NVMe flash memory on modern Dell servers. The first of these servers has just reached the end of its Dell warranty, but is still very much in use and highly useful; we may seek extended hardware support for these servers, or replace them and transplant the SSDs which represent the bulk of the cost.

The VM disk data on each of these servers is replicated to another TrueNAS system (in most cases off-site) for safety.

GPU VM data storage

VMs on the GPU VM pool store user home directories on the same TrueNAS server that hosts the VM disks, but on a separate NFS filesystem that can be shared between the VMs. We chose not to implement per-user quotas on this filesystem so as not to be a barrier to data-storage-heavy experiments. The storage requirements have grown much more quickly than anticipated and we have had to upgrade the storage in this server several times, as machine learning datasets have grown rapidly. Some of these upgrades have been paid for by the MPhil SIF funds due to the MPhil students' use of this infrastructure.

UIS services: RDS, IFS

We have gained some experience with UIS's Research Data Store for providing additional storage to users on HPC for a modest fee. A couple of ACS students this year needed more space than HPC ordinarily provides, and RDS proved to be an easy and effective solution to the problem once we understood how to use it.

UIS also offers an Institutional File Store service, on which we have a free initial allocation of 4TB via the School. This is essentially a resilient NetApp cluster like Filer, albeit less flexible as they do not give direct configuration access. We are not currently using this but may do so in future, particularly if we have storage linked to accounts in the University domain rather than the departmental domain. One possible future use may be for undergraduates' data on the student SSH servers.

Data centres

The departmental data centre, GN09, has 29 server racks and is physically approximately 90% full. It is also close to its cooling and power limits, though those are harder to quantify.

Cooling capacity is dictated by the reduction in chiller efficacy on very hot days; during June 2022 when Cambridge experienced record-breaking temperatures, the chiller remained working but had to run both of its cooling circuits continuously during the hottest few days to keep up with the load, whereas normally we expect one of its circuits to remain idle as a standby in case of a fault affecting the active circuit. Had the chiller experienced any fault during that time, we would most likely have had to immediately turn off a large proportion of our servers. The thermal load in GN09 has increased since 2022 due to GPU server purchases.

Power capacity is mainly dictated by a desire to keep the runtime of the UPS to a reasonable level; if we increase the electrical demand on the UPS much further, its runtime on battery will be reduced to just a few minutes. The runtime is already insufficient to gracefully shut down servers, especially as in the event of an unexpected power outage we would not know immediately whether we should start to shut servers down or whether the power might return within a few minutes. (The aforementioned electrical connection to the Roger Needham Building would help with this.)

We have four racks in the UIS's West Cambridge Data Centre, two of which were provided specifically for the Morello project on the understanding that the power requirement in those racks would be very low. The two higher-capacity racks are close to capacity.

UIS has tentatively offered us space in a new shared facility in the Genome Campus in Hinxton, which supposedly has a lot of spare capacity for systems with high power and cooling requirements. Obviously its location would be inconvenient for research systems that researchers would need frequent access to for experiments. It may however be suitable for our GPU servers, which could free up space in GN09 for less power-hungry uses – though these are reducing as even previously low-power research (such as systems and networking) is starting to involve GPUs.

I have received feedback from one researcher who really wants a lot of accessible dedicated space for networking experiments, separate from our departmental data centre which imposes operational restrictions such as the requirement to maintain tidy wiring. The department previously had SE18 for this purpose, prior to the need to convert it to office space, at which point the plan was for IT staff and researchers to share GN09. My impression is that this arrangement is working for most people, with just a couple of exceptions.

But I must be abundantly clear that researchers are not and will not ever be permitted to bring new hardware into any of our departmental data centre spaces without approval from the Infrastructure team (generally, me) due to the need to plan capacity. There has been one instance of this recently that was only noticed by coincidence of timing, that would have added a very large and unsustainable thermal and power load to GN09 and cannot possibly be accommodated there or in any other facility that we currently have access to – even though the researcher in question already has the hardware and nowhere else to put it.

Active Directory

Active Directory, as our main authentication and user management system (as well as our system for managing Windows machines) is at present jointly managed by the Infrastructure and Support teams.

We have successfully undertaken a significant project to replace two of the three Active Directory servers which were running an end-of-life operating system, with two new servers running the latest version of Windows. The new servers are virtual machines, one of which is on UIS IaaS (see above) and the other is a VM running on the secondary pool in GN09. These are running alongside the existing third server, a physical machine in WCDC, which was already running a newer version of Windows that is still in support for now. However the hardware is ageing and we plan to replace it to bring it in line with the others, probably within a year or two.

Email

Our primary email service is Exchange Online, via the University's tenancy, though we support the use of other mail stores.

Incoming mail processing and forwarding

Mail to addresses @cl.cam.ac.uk is routed via a UIS anti-spam and anti-malware service initially, then through departmental Mail Transfer Agents (MTAs) before being forwarded onward to the recipient's mail store (or to another system such as a mailing list service).

Mail to addresses @cst.cam.ac.uk is routed via the same UIS anti-spam and anti-malware service, then until very recently went to the UIS Managed Mail Domain service for routing to its ultimate destination, performing the same function as the departmental MTAs. We had hoped to migrate cl.cam.ac.uk to the same Managed Mail Domain service; however UIS have now withdrawn the service (their replacement service within Exchange does not meet our needs: no address can have multiple recipients, and no API is provided to programmatically update forwarding addresses). Consequently we have had to move cst.cam.ac.uk mail forwarding in-house onto the departmental MTAs, temporarily making it the same as cl.cam.ac.uk.

UIS are also withdrawing the front-end anti-spam and anti-malware service on which mail lands from the internet. I know from experience that running such a service ourselves would involve a lot of ongoing work to continually tune anti-spam rules and to keep track of the changing requirements of major email services.

We plan to outsource both the anti-spam/malware and mail forwarding functions to Forward Email. Some development work is needed to implement synchronisation of our email forwarding rules to Forward Email's API, which will be undertaken during October.

There may be some subtle changes to some of the more esoteric email addressing schemes that the departmental MTAs support, though we will try to ensure compatibility with the schemes that we observe to be seeing ongoing use. The exact behaviour of the spam filter will inevitably change as well. These changes will be communicated to the department alongside an opportunity to test Forward Email before it goes live.

Forward Email is based in the US and we will be transferring some personal data to them: a list of department members' email addresses. The details of this data transfer have been discussed with and approved by the University's Head of Data Protection and Information Compliance.

Mailing lists

Our primary mailing list service is lists.cam.ac.uk, UIS's Sympa system. We have been working to close down our legacy in-house mailing list system; at the time of writing there are still a few lists on that system and we hope to get the last ones moved or removed before Forward Email goes live.

Fastmail

The department offers an alternative mail store service for people for whom the University's Exchange Online offering is inadequate, for example users relying for research on mail not being altered in transit (which Exchange does, both for security and due to its implementation), or who need standard mail protocols such as IMAP and SMTP (Exchange does not fully adhere to Internet standards). This service is provided by Fastmail, who charges a small monthly fee for each mailbox. The department funds this where the user can justify its use and accepts that this is an advanced service for which we offer minimal support. We do not currently require a high burden of justification.

Legacy mailboxes

A few long-standing members of the department store their email messages on Filer, via a legacy route on our MTAs that will save mail to a file. After the move to Forward Email, ongoing support for this will require us to keep running MTAs just for this purpose. We plan to continue to support legacy mailboxes for the time being, but would ultimately like to decommission the MTAs entirely. We will start encouraging such users to migrate to a modern email service such as Exchange Online or Fastmail.

Web hosting

www.cst.cam.ac.uk

The main "externally-facing" department website is hosted by UIS, currently on their Drupal 7 platform. We will need to migrate to Drupal 10 soon. The involvement of the Infrastructure team in this move will be minimal as the concerns are mainly centred around management and migration of content, rather than technical issues.

www.cl.cam.ac.uk

The "old" department website is still used for various functions for which Drupal is unsuitable, for example research groups who maintain their own HTML files, or who have dynamic web pages, or need to serve data directly from Filer. www.cl.cam.ac.uk will continue to exist for the foreseeable future for these purposes and others, maintained by the Infrastructure team.

One particular ongoing use of www.cl.cam.ac.uk is our teaching web pages, which are currently maintained in such a way as to try to appear as part of www.cst.cam.ac.uk, to the extent that a mirror of large parts of the structure and content of the Drupal site is maintained on Filer so that the transition between the Teaching pages and the rest of the website appears seamless.

We don't think this is sustainable, and – if the University's central solution to teaching web pages (Moodle) is not deemed adequate – we propose to make the Teaching pages into a self-contained site that does not try to share a navigation structure or design with the public (Drupal) website.

Research websites

Research groups are increasingly asking for self-contained websites separate from the departmental website, with their own unique visual identity. Several groups have found their own solutions for implementing hosting these (generally on some free cloud service such as GitHub Pages or Google Sites). We don't consider this to be a problem, necessarily, but I am conscious that not much thought has been given to the ongoing maintenance of these sites after the person who set them up has left the department, or to the impact of this on our own brand and website (less of the information about our research is on our website than it used to be), or even to a consistent naming scheme for their URLs (some are in cl.cam.ac.uk; some are in cst.cam.ac.uk (e.g. mobicentre.cst.cam.ac.uk); some have their own domains (ocamllabs.io) or use a cloud service's domain (mlatcl.github.io).

Wikis

The department has an internal wiki service, used by a few research groups and by the Infrastructure team. This is based on MoinMoin which is obsolete software that will not even run on modern operating systems, so we will need to replace this service.

UIS also has a managed wiki service, which we also make use of, but they too have announced its closure.

UIS's GitLab service does have basic wiki functionality, though without public access; some further investigation is needed of whether this might meet our needs.

Internal IT team documentation is likely to move to SharePoint / OneNote.

We do not believe that implementing and running a new wiki service would be an especially good use of our time.

Lab Linux

"Lab Linux" is our shorthand for an installation of Linux, generally Ubuntu, which is provided, customised and managed by the IT team. Over the course of several decades of providing Linux for members of the department, we have accumulated a wide array of customisations that we routinely apply to a Linux installation.

Image maintenance

We generally provide only "Long Term Support" releases of Ubuntu, starting after its first patch release, which are released every two years; the latest (24.04.1) was just released. Building and testing a new image requires detailed knowledge of our environment, scripts and customisations as well as the bespoke imaging system, and has always been done by Piete. When this task is passed to Piete's successor, this will be an opportunity for a fresh set of eyes to look at the way this is done and it is possible that they will choose to make changes.

Desktop image deployment

As mentioned above, we plan to make more use of network booting (PXE) by the support team for deploying Lab Linux to a PC. The support team also plans to use PXE to install Windows. We plan to explore a combined solution that uses a similar simple process for both operating systems.

We have also briefly discussed using ZFS to image Linux systems.

However any decision on future deployment development effort should wait until new staff members are in place as they will be implementing, maintaining and/or using it.

Packages

As part of our customisations to Ubuntu (and previously Red Hat / Fedora / CentOS), we maintain a large number of simple 'deb' (and 'rpm') packages. These are currently built using a legacy process, initially as 'rpm' packages and then converted to 'deb'. This process is in need of modernisation; however the standard process for building ‘deb’ packages directly is high-overhead for our use case. This may be a conundrum for the new member of staff to consider.

Laptops

We do not currently provide Lab Linux on laptops, as our customisations are mostly concerned with integration with the Lab LDAP, Kerberos and Filer which will not work when the laptop leaves the building. Thus far, anyone asking for Linux on a laptop is expected to install Linux themselves, and refer to our documentation (which is admittedly rather terse) to make use of Lab services from their standalone laptop. Generally we think that anyone specifically asking Linux will already have sufficient experience to manage this. However it remains for consideration whether this is the best approach.

Business applications and processes

Database server

Most of the structured data associated with our business needs are stored in an instance of Microsoft SQL Server. This covers a wide range of types of data including staff records, visitor records, training records, inventory, network registrations, network wiring, user account information, group memberships and many other things. These databases have grown organically over an extended period of time, with complex interdependencies between tables, a lot of custom code baked into database views and stored procedures, inconsistent setups and multiple old versions of some tables and fields. The database needs a coherent redesign, separating data into self-contained tables with single sources of truth for each class of record and well-specified interfaces.

The database server itself is believed to be worth keeping as a platform for the redesigned database, albeit with some issues in need of rectification, for example the backup strategy (which currently uses replication, which is fragile).

The database server is currently a joint responsibility between the Infrastructure and Support teams, though neither team has specialist SQL Server knowledge at present; we hope to rectify this with future recruitment but are currently gaining experience of managing it ourselves.

Database web applications (dbwebserver)

dbwebserver hosts most of the user-facing interface to the aforementioned databases, implemented as Active Server Pages on Windows. Much of this has become unmaintainable, and alongside the database reimplementation, we plan to reimplement the user interface applications. We are currently evaluating options for hosting such applications in a modern and maintainable way, with a proposed new password management application as a useful proof-of-concept.

User administration

The processes and interfaces for creating user accounts form one of the applications currently provided by dbwebserver, largely the domain of the Support team and used by administrative staff; however the deletion of lapsed accounts falls to the Infrastructure team and I note it here because the system for archiving user data is not working and needs significant attention. Also, there is an ongoing problem affecting password issuance for the annual influx of PhD, MPhil and Part III students each October, which for the last couple of years has required a manual workaround. These processes and systems are in particular need of attention.

Request Tracker

Request Tracker, which manages IT support tickets and requests (as well as other ticket queues for other teams such as Building Services) is managed by the Support team, but the software installation and management of the underlying server have been taken over by the Infrastructure team since Chris Hadley's retirement. We have for some time been preparing to migrate to a new major version of Request Tracker on a new server running a current version of Ubuntu; this project has been completed earlier this year and we are making use of some of the new features. Further incremental improvements to the server setup are ongoing, and a minor upgrade alongside a new Ubuntu version is upcoming.

Teaching servers

MPhil / Part III GPU servers

We run a small GPU cluster for the ACS course, currently comprising two servers each with 4x NVIDIA P100 GPUs (of which one server is currently used as a shared development host and the other hosts the small number of VMs needed by niche projects). These GPUs are obsolete now and we cannot buy new GPUs compatible with the existing servers, so we have purchased two new servers each with 4x NVIDIA L40S GPUs to replace this pool, plus another server with 8x NVIDIA L4 GPUs (these are low-power GPUs with reduced capability, which we hope to be useful to allow more students to develop and test code in parallel using GPUs). However, time pressures have meant that these servers are not installed yet. We hope to have them ready in time for ACS projects.

It remains to be seen whether the L4 server is useful for development; this is an experiment. If it is not as useful as we hope, we may make one of the L40S servers into a development server.

The GPU demands of ACS projects are very much greater than can be satisfied by this small cluster alone; we ask them to use HPC as their primary GPU resource, with the departmental GPUs mainly for development and testing, with perhaps a small minority of students using GPU VMs, mainly interactively. This strategy mostly worked in 2023-24, though an unfortunately-timed string of HPC outages caused problems for a while and students had to wait a long time for their jobs to complete.

Undergraduate SSH servers

Undergraduate students have historically not used any departmental IT service; their needs were met by UIS, and in particular the Public Workstation Facility / Managed Cluster Service / Desktop Services which provided Windows and Linux workstations in the Intel Lab and remote access to Linux servers via SSH. All of those services have been withdrawn in recent years, replaced only by a University Managed Desktop service which only offers Windows desktop PCs, with no remote access nor Linux.

To replace UIS's MCS Linux, last year we implemented a Linux SSH server for undergraduates, to which they can log in from either the Intel Lab Windows PCs or from their own computers remotely. This seems to have been popular, or at least well-used.

Per their suggestion, we set up this server to use UIS's Active Directory authentication service (Blue AD), so that students could log in without having to be issued with a departmental account. UIS provided detailed instructions on how we should implement this so that students could log in using a password, and that password would be checked against their University account. However after the service had gone into use, a different team in UIS objected to this setup and required urgent changes to how authentication works, specifically to disable password authentication entirely (which had been the only authentication method implemented by the instructions that they originally provided).

We are currently implementing a new system to enable users of this server to set up and manage SSH keys for access, via a web application. We hope to have this ready before the start of term. Every user of this service will need to take action to maintain their access; instructions will be communicated once the system is ready. We apologise for the last-minute nature of this change before the start of term.

Specialist purchasing

Besides the purchase of bespoke research servers discussed above, the Infrastructure team is involved with various other specialist procurement processes in the department.

Nonstandard desktops

The Support team has developed a catalogue of standard desktops and laptops that meet the needs of most people. However, some researchers need specialist desktop PCs, for example if they need a high-end GPU at their desk which cannot be supplied as an upgrade to the standard PC. We will routinely get involved with the specification of such machines to try to ensure that wherever possible, the machine meets our requirements for remote management – in particular that it has a management controller (BMC).

Credit card purchases

Although not usually related to IT infrastructure, for historical reasons I am the sole de-facto port of call for researchers who need goods or services that cannot be procured through the University's proper channels and which can only be bought on a departmental credit card. Often these are miscellaneous non-IT goods needed for practical experiments, or are gift vouchers to reward participants in experiments. This takes time away from the maintenance of IT infrastructure. Although the University of course does not want to make it easy to sidestep procurement procedures, some of our research would not be possible without this.

Research support

General

I see it as an important part of the IT services in the department that we are generally able to help researchers with a wide range of technical issues, whether or not the issues involve departmental IT systems in any way. Not only does this hopefully benefit our research, but having a central view into the problems being solved by different researchers can help inform future service developments – for example, recent discussions have led to a promising avenue for a potential new service for running containers – or at least the sharing of expertise and solutions between research groups. My hope is that we can in future devote more staff time to this.

High-performance compute and GPUs

We have a central departmental GPU cluster on which any researcher can request a VM with one (real) GPU attached via PCIe passthrough. Most of these were funded from EPSRC Core Equipment funds, supplemented by some research funds generously contributed by PIs who make use of GPUs but were willing to contribute towards a shared facility rather than buying their own.

As originally purchased, the research GPU pool had 40 RTX8000 GPUs split across four large servers. The first batch (one server with 10 GPUs) worked well; the subsequent servers proved somewhat unreliable due to several hardware quality issues both with the GPUs and with the large servers. In any case these GPUs are approaching the end of their useful life and we are starting to replace them as funds become available.

We have previously replaced one server (10 GPUs) as well as a couple of particularly unreliable GPUs from other servers with four newer servers each with four NVIDIA A100 GPUs.

This year we applied again for EPSRC Core Equipment funding to replace about half of the remaining RTX8000 GPUs, but were only awarded sufficient funds to buy one server with four L40S GPUs. As four GPUs is not sufficient for an effective VM pool, we may install the L40S server as a shared development system replacing dev-gpu-1 which is currently a VM with one A100 GPU (which would return to the VM pool to marginally increase the capacity for VMs with A100 GPUs.

Due to data centre capacity constraints (cooling, power and space) we are unable to meaningfully increase the size of the GPU pool to meet the very high level of demand. Our strategy is to primarily provide our GPUs for development and testing, then direct people towards the University HPC facility to run their experiments at scale as they have a very large GPU cluster (360 NVIDIA A100s plus a new large cluster of Intel GPUs). However, that facility is overused too and researchers are unhappy with the job queueing times and lack of provision for interactive use, and some are unhappy with their choice of Intel GPUs for a recent expansion rather than NVIDIA.

We are also unable to buy any more A100 GPUs as they are no longer available. There is no obvious successor model; the newer L40S is a slightly lower-end GPU which although better than the A100 several respects, performs very substantially worse than the A100 at one class of operation (FP64) and also has much less (and slower) memory. The new high-end H100 is very capable but costs about three times as much as the A100 did.

The RTX8000 GPUs displaced from the departmental pool, as well as the displaced server, have found homes in a couple of research groups – one of which had suffered a failure of their own GPU server, and another of which has existing GPU-capable servers and wants to expand into a research area that needs GPUs. On the one hand it is good that we are continuing to benefit from this purchase beyond the conclusion of their original application, and in spite of their unreliability in that application; however we had intended to remove them from GN09 so that they did not continue to use valuable power and cooling capacity, but they have stayed there, and are quite wasteful of power compared with newer GPUs.

Model data centre

The model data centre is a shared cluster of 80 servers interconnected by a fast internal network, originally built in 2014-5 (both as my first project for the department, and as the first project in the redeveloped GN09 which went on to inform how we would manage the whole data centre). This equipment is now a decade old; the network infrastructure has long been obsolete and is no longer as interesting to research as it once was. Some of the servers are suffering from hardware faults. Many of the servers are however still seeing a reasonable level of use, albeit not for the purpose for which they were originally intended. No experiment ever ended up scaling to use all 80 servers as we originally thought; however a few users have used 20+ servers simultaneously. There have been more distributed systems projects and fewer networking projects than we originally anticipated.

It is clear from the ongoing use that a pool of physical servers is useful to some in the department, particularly for use cases that cannot be virtualised such as performance measurements and work that requires management of the virtualisation layer. However, modernising this cluster would not be cheap, and would be time-consuming (my original employment in the department was dedicated full-time to building the model data centre the last time around).

Research servers

Several PIs choose to purchase dedicated servers for their projects and research groups. Generally I work with the PI to build a specification that both meets their needs and will work well in our data centres, which is mutually beneficial as it saves time for me and my team as well as better enabling us to set up their servers quickly and keep them working reliably. When the PI asks for GPUs or other resource-intensive equipment, I make a point to discuss the data centre capacity constraints imposed on us all, but ultimately have not stood in their way if they believe that they need dedicated GPUs despite the existence of HPC and shared GPU facilities.

However, we currently have a backlog of at least 16 servers either awaiting ordering or awaiting installation, but deprioritised behind other urgent work.

Many long-standing researchers also have old servers still in use, or at least still installed in the data centre. Older servers require disproportionately more maintenance effort, especially as they start to develop problems. They also tend to use disproportionately more power and/or space and require more cooling. I would welcome a discussion on how we can fairly and considerately reduce the demands of obsolete hardware, particularly in cases where the researcher does not have funds to replace it or does not see a need to.

Bespoke infrastructure setup and maintenance

A few research groups with more complex needs that could not be handled by individual standalone servers have at various times opted for bespoke infrastructure supported by the Infrastructure team (or in some cases researchers built and maintained their own infrastructure without our involvement). Particular common building blocks for such infrastructure are the components that we use ourselves for departmental infrastructure, for example:

- Dedicated XCP-ng VM pools: we have now built several of these for a couple of groups (in one case providing GPUs, similarly to the departmental GPU VM pool). Typically there is a level of joint responsibility for the operation of the pool and the VMs, though the split varies according to the needs and wishes of the users. On one pool, the Infrastructure team sets up VMs as if on our own infrastructure, and the users jointly manage the individual VMs alongside us, as is common practice for Lab-managed research systems. On another pool, we maintain the VM infrastructure but users manage VMs themselves.

- TrueNAS storage: we currently run multiple TrueNAS servers for one group which has several servers and needed shared storage accessible from all of them, and off-site backups of the data to another dedicated off-site server. Another group has been considering TrueNAS for their needs as well.

Software licensing

We have historically run licence servers for groups that have purchased or been given software licences. This led to a proliferation of licence servers in various states of disrepair, some with obsolete licences or with unknown owners or users, and all rather fragile due to the varied and unknown software requirements of the licence server applications. We are trying to move to a model where researchers with software licences operate their own licence servers, and we simply provide them with a VM to run them on as we would for any other research use. There are however complications around certain licences which although are logically used by certain groups, strictly belong to the department as a whole with no clear owner.

Cloud

Amazon Web Services

The departmental IT team does not use AWS ourselves, but we have an AWS Organization in which researchers can request accounts, and a billing arrangement that allows usage to be charged to research grants.

We have a small number of historic AWS accounts via a reseller, Digital Spaces, for which we are paying about £10 per month in total, with no knowledge of (or ability to find out) what these accounts are being used for or by whom – or indeed whether the accounts have been abandoned without terminating all of their services.

We are attempting to move our AWS Organization to the JISC OCRE framework now available via UIS which didit not exist when we joined AWS, but which would now give us preferential pricing (and may also present the opportunity to adopt and gain control of the Digital Spaces accounts). However, for some reason and despite UIS attempting to chase our request, JISC has not acted on our request since we submitted it many months ago.

Google Compute Platform

We have a departmental project and billing arrangement set up with GCP, similar to our AWS Organization though in the case of GCP this exists within the University's tenancy rather than standalone. So far the only use of GCP has been by projects who have been awarded free credit by Google. However GCP seems to treat credit as shared by all users within our department, so this may become complicated to manage if we ever get a mix of paying and non-paying users.

UIS IaaS

UIS Infrastructure-as-a-Service is essentially another cloud service available to us, though with a narrower feature set than AWS or GCP. IaaS is free for departmental operations, but is charged for research use or profit-generating activities. The charges are generally slightly cheaper than similar commercial services.

Mythic Beasts

We have an account with a local small cloud service provider, Mythic Beasts, who have a close relationship with the University. We currently use them for one VM for departmental infrastructure use (routing of non-academic traffic to the internet; see above) as well as for domain registrations, and routinely suggest them to members of the department who need a service that we are not in a position to economically provide, such as managed dynamic web servers.

Hardware lab

The department's hardware lab formally falls within the remit of the Infrastructure team, and in practice is run solely by Tom Bytheway. He provides, and has equipment for, services such as:

- 3D modelling, design and printing
- Electronics design and assembly
- Hardware design and fabrication
- Laser cutting

His services have been popular amongst researchers with hardware needs and he has solved a wide range of practical problems for them. He is hoping to expand the facilities offered, and is currently working with the Building Services team with the aim to set up a dedicated workshop and fabrication space separate from the electronics facilities.

He reports that the laser cutter has suffered some damage recently due to a researcher operating it incorrectly, so it is temporarily out of action and some costs will be incurred to have the manufacturer’s specialist technician attend to repair and recalibrate it.

Hardware repairs

Nick Batterham provides the department with, amongst other services, the ability to repair various kinds of devices such as laptops, printers and phones which can no longer be economically repaired by their original manufacturer.

Teaching

Nick is also involved in the teaching of electronics in the tripos.

Tom is also interested in becoming more involved in the teaching of hardware-related skills, most likely initially as a research skills module teaching soldering, but perhaps more in future. We believe that teaching of such skills can lead researchers to make better use of the facilities and services that he provides, and perhaps open up avenues of research or possibilities for practical projects that they would not have otherwise thought of.

The same applies to the use of GPUs, on which I have been involved in giving another research skills lecture.

Printing

Daniel will report on the rollout of UIS’s DS-Print service, and new printers. The main consequence for the Infrastructure team is that we can start to consider the requirement of our CUPS print server, which until the advent of DS-Print had been the primary way for members of the department to print, but is now hopefully superceded for most users.

CUPS has historically been difficult to run, and was managed by Chris Hadley prior to his retirement; we hope to remove this from our service portfolio, or at least scale it back to specific use cases.

This is dependent on DS-Print working well for Linux users, which is currently not certain and will be subject to ongoing investigation. UIS has provided a Linux DS-Print client, but its installation is fiddly and may require some work to reliably automate.

Weather station

The Digital Technology Group (and prior to them, LCE and their predecessors) has long run a weather station as a public service, with sensors on the roof of the William Gates Building – though the weather station predates the building and was moved here. The Digital Technology Group has closed, but the department considers the weather station to be a useful ongoing service, so it has by default fallen to the Infrastructure team to oversee.

The weather station was abandoned with its software in the middle of a rewrite which was never quite completed. Thankfully we have a couple of new volunteers in the department to try to complete this work and maintain the software.

Some quite pressing work is needed to improve the interface between the software and the hardware; the Infrastructure team has set up new hardware and a new home for the software and public interface, and we are awaiting the new volunteers making the necessary software changes. Progress has been slow and there is a lingering security issue which will become ever more pressing to rectify.

Physical security system

The door locks in the William Gates Building are managed by a very old system, with unmaintained software running on obsolete hardware. The server will be maintained by the Infrastructure team (in collaboration with the Support team, as it runs Windows) for as long as is feasible, but it will fail sooner or later. We are at serious risk of disruption to use of the building if this fails.

UIS had begun a project to procure a University-wide physical access control system, with our building set to be the pilot so as to give us a new system as quickly as possible. However the project appears to have stalled. Discussion of next steps will take place in the Buildings and Environment Committee but as a highly risky piece of IT infrastructure I mention it here for completeness.

University IT policies

The University's Information Services Committee has imposed several new University-wide policies intended to reduce cybersecurity and reputational risks to the University. The Acceptable Use Policy, applicable to every individual who uses IT facilities in the University, came into force in April 2024 with a transition period lasting until April 2025. Amongst other things it requires available operating system updates to be installed promptly, which we know is not the case on some systems. A Systems Management Policy, applicable to those such as us who run multi-user systems, also came into force in April 2024 with a transition period lasting until April 2026. This latter policy will require some substantial changes to how we operate systems and work as IT staff. We are still considering the full implications.

Some areas of research such as cybersecurity may need to deliberately contravene these policies, and in order to be allowed to do so they will need to apply for an exemption. We have had contradictory advice saying both that exemptions need to be case-by-case, and that we may (if we justify it) be granted a blanket exemption for the department. We are not aware of any such requests so far.

There is also a new policy proposed to limit who can send and receive email as a cam.ac.uk email address, including departmental subdomains. Currently our alumni automatically keep their departmental email address for life (forwarding to a mail store elsewhere); this is not permitted under the new proposed policy, although there are mechanisms whereby we can allow alumni with a specific need (for example, their email address is published on a paper) to apply to a “Retention Service” operated by the department to keep their email forwarding active. As written, the Retention Service appears to be labour-intensive to operate due to the requirement to manually review applications and for the end user to apply regularly for renewal.