Introduction

- **This represents my views as a University postmaster.**
- **I'd be happy to sit down with the Panel to answer questions.**

I am David McBride, and I have served as a postmaster for the University since January 2017, alongside my senior colleague, David Carter. We have joint responsibility for PPSW, Hermes, the managed mailing-list system, and related services and functions.

This document is intended to inform your determination of the University's email strategy. In it, I try to point out some non-obvious risks to consider, and assert a number of constraints that our future strategy should satisfy.

In the event the Panel would like to solicit further evidence from me to help inform their deliberations, I'd be happy to do so in person or in writing.

Conclusions

- **We must continue to operate the central mail hub, PPSW, and support multiple back-end email providers, both local and remote.**
- **We must continue to provide a standards-compliant user-facing email service, whether procured from an external vendor or by developing the in-house Hermes service.**
- **We must not standardize on using Exchange Online.**
- **We should stop issuing Exchange Online accounts to new users by default.**

PPSW: The University will need to continue to operate the central email hub, PPSW. It performs a number of essential functions—such as server smarthost services, mail filtering, mail routing, and buffering in the face of service failure—which would be difficult to outsource entirely. Further, it provides the University an important point of oversight and control over email, that ensures that we retain the flexibility and technical expertise to meet current and future needs.

Multiple providers: Any one email service is unlikely to meet all University users' needs. Over and above Hermes and Exchange Online, numerous departments and colleges also operate email services, hosted locally or contracted from external providers, and rely on PPSW to link things together. We should continue to support this model, as this will help ensure that the University retains the ability to switch out different email service providers as both needs and available services evolve.

Standards-compliant services: It is important that the email services that the University procures and supplies are compliant with IETF standards—SMTP, IMAP, Sieve, etc.—because this ensures that services will remain accessible for all users with a wide range of systems and devices, rather than just a subset. It is also an important factor in ensuring that migration between different email service providers and implementations remains feasible, thus avoiding vendor lock-in and the costs this incurs.

Exchange Online: We cannot standardize on Exchange Online services for University email provision as it does not implement IETF standards—at least, not competently. While compatibility with various standards is claimed, its IMAP service implementation is sufficiently poor as to lose data in practice, it's mail handling has been seen to violate RFC MUST assertions that causes cryptographic email signatures to be rendered invalid, and it has even sometimes generated outbound email that contains invalid MIME data.

Defaulting to using Exchange Online for University mail functions was a change made without justification, and has already caused significant interoperability problems for users, operational problems for the postmasters, and has started to take us down a path towards significant vendor lock-in. We should reverse course.

Background

- **Email services are important: email is the communications system of record.**
- **Email services are complicated: they will require local technical expertise to support and operate, even if services are outsourced.**
- **Email services are sensitive and security-critical: the contents of emails can be extremely sensitive, and extraordinary access to email accounts must be carefully controlled.**
- **Standards-based services have been starved of development resources for years as a result of reorganisations, staff departures, and to support the deployment of Exchange Online.**
- **Hermes's availability compares favourably to other services.**

Important: Email services continue to be critical to the operation of the University and its members. As much as some might argue that "email is dead" based on users' pattern of using various instant messaging services for communications with friends or peers, the use of such services is not new, and email remains the formal communications systems of record for members of the University and the wider populace, both when communicating within the University and with external organisations. This seems unlikely to change in the near future.

Complicated: Email services are complicated, and can easily go wrong in subtle ways without careful service design, operation, and management of the components themselves, as well as the other systems they interact with like the DNS. This is an area where complexity tends to increase, with additional requirements and standards being introduced, and with tightening constraints to try to guard against attackers sending spam, phishing lures, and malware. Even (especially!) if services are outsourced, retaining local technical expertise and controls will be important for managing the relevant namespaces, services, and reputations in sustainable ways.

Sensitive: Email is often used to communicate privileged information, and can be extremely personally sensitive—for example, being used to communicate personal medical information, discussing workplace concerns with union representatives, documenting in-person interactions between staff and managers, and so forth.

Over and above the sensitive nature of many emails, access to an email account is frequently used as a check by other service providers to protect access to their systems. If someone forgets their Facebook password, for example, they can typically regain access to their account by soliciting a password reset link that is sent to their registered email address.

However, sometimes extraordinary access will be requested to a person's email account. The motivations vary, but can include: to investigate misuse, to attempt to locate a missing person, to set an autoresponder message on the account of a person who has died or fallen seriously unwell, or per a request from law enforcement.

Managing extraordinary access to email accounts must be done with exaggerated care, and I'm pleased to say that we have a history of being careful, principled, lawful, and effective at handling these requests as and when they occur.

It is essential that ensure that these high standards are maintained for any and all email systems operated on behalf of University members.

Under-investment: The reorganisation of computing services during the establishment of the UIS had the side-effect of reducing the number of staff providing University postmaster services from two to one for an extended period of time—leaving my senior colleague as a single point of failure for an extended period. Further, development efforts were re-directed from enhancing core email services towards instead supporting new Exchange Online facilities, and migration tools for same. While I have now been added as a second postmaster to help provide operational cover, development capacity has not increased as I still retain all of my pre-existing duties.

Availability: Though recent investment has been light, PPSW, lists.cam.ac.uk, and Hermes have an excellent availability record that compares favourably to both GMail and particularly Exchange Online, which have both suffered conspicuous failures over the past few years.

Hermes has suffered a single two-hour outage in recent years, caused by the unplanned outage of the machine-room location where Hermes was operating in early 2018. This resulted in the loss of availability of email, but not the loss of data, nor of messages in flight.

Lists.cam.ac.uk has not suffered any outage, planned or unplanned, in recent years.

PPSW has not suffered any gross outages, though the anti-malware filters have, on occasion, incorrectly rejected emails as hostile because of an bad signature. These are usually fixed within a few hours.

Finally, I would caution that the historical alerting data from the ITSS (status.uis.cam.ac.uk) should not be taken as a reliable measure of the availability or otherwise of UIS services. Reports generated by the system do not appear to include all recorded outages, such as the outage of Hermes in early 2018 caused by an unexpected machine-room shutdown. Further, not all outages result in status updates, sometimes because outages were only discovered after they have already been resolved, or because of human factors. For example, Exchange Online suffered a serious outage across Europe and the UK in January of this year that, among other issues, prevented the delivery of any email to the service for up to 9 hours. The reported health of Exchange Online was not changed to RED on the ITSS during this incident.

User needs

- **The UIS did not undertake a user needs exercise before commissioning Exchange Online services, nor stated any justification for provisioning new users on Exchange Online by default.**
- **Multiple departments will likely commission their own Hermes-like standards-compliant service if the University ceases providing one centrally.**
- **Hermes does not meet all users' needs, particularly around quotas and meeting scheduling.**
- **Exchange Online does not meet all users' needs, particularly around interoperability.**
- **Case study: the previous UIS Director mandated that UIS staff migrate from Hermes to Exchange Online; while many staff appeared to do so, they actually routed their email back to Hermes to avoid disruption.**
- **Email, in general, is hard to reliably use securely, and we should try to improve this. Deploying SPF, DKIM, and DMARC support throughout the University would help.**

Exchange Online selection process: The UIS has never showed any credible reasoning as to why it had selected Exchange Online as the ideal email service of the future. While there were some internal presentations, and even some project documentation, that supported this conclusion, those project documents were withdrawn and their findings, costings, and reported risks all repudiated—apart from the conclusion. It is clear to me that this has been an unfortunate case of policy-based evidence-making, where a decision was made first, and evidence subsequently sought to justify it, apparently heedless of the merits of any objections or concerns raised. We are continuing to live with the consequences.

Last summer, the UIS, after undertaking a significant amount of work, and again over the objections and concerns raised by some sectors of the University, took the step of changing standing procedures to cause new users to be issued with Exchange Online, rather than Hermes, email accounts as standard. There has never been a stated justification for this change.

When I asked senior staff why we were persisting with deploying this change, despite the lack of any justification for doing so, the answer I received was, paraphrased, “We've started, and so we'll finish”, with a promise that the future direction of email services will be assessed once the transition of Exchange Online by default was completed.

While Exchange Online is likely to effectively meet the needs of some users, it has not been effectively determined to what extent it meets, or fails to meet, the needs of different users and stakeholders, nor whether it is a sound strategic choice to rely on this third-party service for this critical function.

Duplication of effort: I understand that staff from a number of departments have long been expressing concern and disquiet over the possibly-limited future of the Hermes email service.

I expect that, should the Hermes service be discontinued without a suitable standards-compliant service being provided in its stead, multiple departments will likely commission their own local replacements. This will result in a set of services delivered less efficiently, with greater variability, and at greater cost than if a suitable service was provisioned centrally.

I understand that the SRCF, anticipating the decommissioning of Hermes, have already commissioned and deployed their own replacement service, called Hades.

Hermes limitations: A common complaint made about Hermes is that the storage quotas provided by default, and the upper range of quotas currently supported, are too limited. This is a fair complaint. Many users rely on being able to store a long email history to serve as a journal and record of past events, and while initially attractive, underinvestment in Hermes has meant that the service's storage capabilities have not kept pace with demand over the past decade.

Separately, users who are familiar with tightly-integrated Microsoft Exchange environments and the Outlook mail client find using Outlook with Hermes to be a limiting experience, because Hermes does not implement all of the functions that a tightly-integrated Microsoft environment provides—in particular, the easy scheduling of meetings with peers, where those peers record their availability.

Outlook, which is a popular email client on corporate Windows systems, has a history of not interoperating well with standards-compliant email services like Hermes. This results in a reduction of not just expected functionality, but also performance and reliability.

Exchange Online limitations: Outlook does not function as an IMAP client well, and Exchange and Exchange Online likewise do not function well as an IMAP server. The UIS determined that the IMAP interface provided by Exchange Online is even worse than that provided by Exchange on-premise, to the extent that it is not supportable, as it would lose user emails during normal use. This is a significant barrier to adoption by many users, particularly those who do not use Windows-centric tooling.

This lack of IMAP support particularly effects users of mobile devices, as they can no-longer use standard mail tools, and must instead either use webmail (which does not support disconnected operation), or must use mobile apps that support Exchange-specific mail interfaces, where available. I understand that, to be used, these mobile apps demand excessive levels of authority over the user's personal device—granting the University the capability to disable mobile device features, such as cameras; remotely locking the screen; and remotely erasing the device entirely. We should not have these capabilities.

I understand that Exchange Online does not smoothly integrate with third-party calendar providers, such as those hosted by Google. Users frequently operate their own personal online calendaring facilities and find it valuable to have a unified view over separate personal and professional diaries. Exchange Online's lack of interoperability in this area thus seriously limits its usefulness.

Finally, while currently only of concern to a subset of users, I have also seen cases where cryptographically-signed messages sent to Exchange Online users were modified by the service contrary to the requirements of the relevant IETF standards for email delivery, such that those messages could no-longer be verified as authentic and were rendered with red warnings, showing (correctly) that the message had been tampered with in flight. This lack of standards compliance may compromise our ability to deploy secure email technologies in future.

Migration case-study: In February 2017, all UIS staff were directed to migrate their existing Hermes email accounts to the then-new UIS Exchange Online service, barring an exemption from a Deputy Director. I was pre-exempted from this requirement, on the basis that I served as a Hermes postmaster and so should be required "eat my own dogfood"[1]. Many other staff found Exchange Online unworkable in practice, and so worked around this mandate by triggering a migration of the contents of their personal Hermes account to the new Exchange Online service—and then setting up forwarding rules in Exchange Online to route their email back to Hermes. This enabled them to continue to use their preferred service while appearing to satisfy the mandate.

Ease of secure use: Email in general is difficult for users to reliably use securely, and this has been the case for many years—certainly as far back as the ILOVEYOU replicating email-borne malware outbreak in 2000.

The attacks vary, from users receiving extortion emails that threaten to reveal damaging information unless a ransom is paid, to attackers sending phishing messages to try to influence users' behaviour, to users sending sensitive information to the wrong recipient(s) in error.

The consequences can be severe, such as the revealing of sensitive information, the download and execution of hostile malware, the unauthorised transfer of funds, or the exposure of username and password credentials to attackers, who can then perform privileged operations on behalf of their victim.

It is important that we do not blame humans for being deceived by these attackers, but, as in rail and aviation, look at systems as a whole to identify what improvements will make it harder for errors to occur. Ideally, these security features should be invisible to users.

We already use a variety of tools, such as malware-scanning, spam heuristic scoring, and real-time IP reputation services (made available to us by JISC) to try to prevent the bulk of unwanted or hostile email from landing in users' mail spools, but these mechanisms are not perfect.

[1] I prefer, "drink my own champagne".

One useful improvement would be the addition of easy-to-use message authentication, so that a user can readily perceive whether any given message is from a known person or service, or whether it is instead unauthenticated and suspect.

One approach could be to try to enroll systems and users in an OpenPGP or S/MIME-based message-signing (and, optionally, encryption) scheme, but I am concerned about the practicality of this strategy for all email users. I suspect that, if such systems are considered, they should be used for targeted use-cases only.

Rather, I think mandating the use of authenticated email transmission using TLS within and, where feasible, outside the University, and expanding the use of DKIM (server-added cryptographic email signatures) and DMARC (server-side authentication support) to University email domains is the most effective way to improve our email security posture.

Requirements for external providers

- **We must insist that any email service providers contracted by the University support Internet standards for interoperability, and resist proprietary extensions and integrations.**
- **We must insist that email providers satisfy Sender Policy Framework (SPF) requirements, preferably by routing outbound email traffic via @cam.ac.uk addresses using PPSW.**
- **We must require email providers to provide effective email sending rate-limiting controls, as these are essential for security and reputation management.**
- **We should require that extraordinary access to user data and metadata only be provided by external providers with appropriate University oversight.**

Standards-compliance: Email services that do not competently support IETF standards, such as SMTP, IMAP, MIME, and the like, present significant interoperability challenges and will risk leaving us locked in to using a single vendor. Further, we should be cautious about adopting non-standard extensions and integrations to these core-functions, as these can have the same effect.

Please note that Exchange Online does not meet what I would consider to be this minimum standard.

Outbound routing via PPSW: SPF rules places an upper bound on the number of systems that can claim to send email purporting to be from any domain, including cam.ac.uk. The systems currently authorised to send email on behalf of cam.ac.uk are currently PPSW and Exchange Online. It is difficult to expand this set without risking exceeding the limits placed on SPF rules and preventing successful email delivery for cam.ac.uk email addresses generally.

Ideally, Exchange Online would be configured to route all outbound email via PPSW, so that PPSW can remain the single authoritative source of routing and control of email for the cam.ac.uk domain. I understand that Exchange Online supports this mode of operation, but it is not currently configured.

It must be a hard requirement that external mail service providers satisfy all of the constraints imposed by SPF rules, and we should expect this to be met for those services sending email from addresses in the cam.ac.uk domain by the external provider routing outbound their email traffic via PPSW for onwards delivery.

If this expectation cannot be met, the external provider must not be used for sending email claiming to be from the root cam.ac.uk domain, and should be allocated its own sub-domain instead.

Sending rate-limits: Email sending rate-limits in Hermes and PPSW have proven to be an essential safety-feature. When appropriately configured, they will automatically limit accounts and email systems that are generating large quantities of unwanted messages, whether due to a compromise, misconfiguration, or software malfunction.

The effect of these rate-limits is to significantly limit the damage that misconfigured systems or compromised accounts can do, and avoid harming the reputation of the University core email systems in the eyes of third-party reputation systems. This ensures that the University can continue to send legitimate messages to other email service providers, and helps us to rapidly identify and contain misuse and compromised user accounts.

On occasion, a department or college will wish to send a bulk mailshot that exceeds their normal sending rate-limit; this is easily allowed for by a manual zeroing of the relevant counters by a postmaster. Colleges and departments are now in the habit of emailing us in advance of such mailshots to ensure swift and timely delivery, and this is a system that seems to work well.

One of the operational problems that Exchange Online has presented in the past year is that it does not provide effective rate-limit controls. Unlike on-premise Exchange services, where rate-limits can be customised, Exchange Online sets a high, fixed sending limit for all users of 10,000 recipients/day—which is more than the current allowance for most departments! Note also that, for the purposes of these limits, a centrally defined distribution list on Exchange Online is counted as a single recipient, regardless of the number of addresses it expands to.

At the start of October 2018, after the first cohort of undergraduate users were enrolled in Exchange Online for the first time, a small number of compromised accounts were used to distribute disproportionately large quantities of hostile messages within Exchange Online and to other users in the University. Exchange Online's own filtering was ineffective, because it considered the message flows to be intra-organisation, and so trusted not to be hostile.

As a near-term fix, my senior colleague enhanced PPSW to implement per-sender rate-limits of all email traffic arriving from the public Internet, including from Exchange Online, to try to mitigate the impact of compromised Exchange Online accounts on the rest of the University. However, this cannot prevent large numbers of hostile emails from flowing between Exchange Online accounts.

I anticipate that these kinds of outbreaks are likely to become more severe, and more frequent, as the number of Exchange Online users increases.

Oversight of extraordinary access: At present, if extraordinary access or modification to a user's Hermes account is requested by another member of the University, or, unusually, as a result of some legal demand, the postmasters liaise with the Head of User Administration in the UIS, who in turn consult with the head of the requesting University institution and/or Legal Services as necessary, to authorise the relevant actions by the Postmaster.

However, if the University contracts a third party to provide these University mail services, that third party will also be able to facilitate extraordinary access to that user's account and related metadata.

Given the immense sensitivity of email accounts, the University must insist on appropriately strict limits for the provision of such access to third parties, and ensure that some appropriate oversight mechanism exists to ensure that these are enforced.

Costs

- **Exchange Online is not free—licenses and infrastructure are paid for as part of the Office365 component of EES, for a total of approximately £35/user/year, plus staff costs. These costs are likely to substantially increase soon as a result of EES license term changes and cost increases.**
- **Hermes/PPSW/lists services are not free—costs were previously calculated to be about £10/user/year, including staff costs.**
- **The costs incurred for providing local services has been more consistent and predictable than those incurred by using external suppliers.**
- **Sourcing email services from vendors will limit accountability and control, and leave us beholden to their plans and timetables. They change their offerings at will and without consultation, and this can incur further expense.**

Internal presentations from the then UIS Deputy Director for Architecture gave a headline cost of £35/user/year for the University's Exchange Online service, though I understand this figure may also include other Office365 components. This does not include staff costs.

At around the same time, similar figures were quoted to me for the Hermes, PPSW, and the Lists services costing on the order of £10/user/year, *including* staff costs.

I understand that, as a result of changes to the terms of future EES licensing arrangements, as well as charging increases, the costs of our EES agreement—which was originally only signed to rationalise the University's Windows licensing arrangements—are due to increase significantly—perhaps by multiple tens of percent?—in the next few years.

By contrast, the costs for Hermes, PPSW, and the lists system are primarily incurred by hardware procurement and staff time, both of which are stable for a given order of magnitude of number of users.

External providers will also develop their products and deploy new features according to their own plans and timetables, and we are unlikely to be able to materially affect these. We have already had operational issues result from Exchange Online enabling, without apparent warning, the use of SRS (Sender Rewriting Scheme) when sending some mails claiming to be from the cam.ac.uk domain, which required urgent changes to the live configuration of PPSW to prevent a potentially serious email outage.

Appendix: Hermes development roadmap

- **Switch to using scale-out storage — enabling better storage efficiency, much bigger quotas, seamless high-availability, shared folders, and incremental expandability.**
- **Enable self-service email recovery.**
- **Enhance mail-filtering capabilities by expanding Sieve support and implementing the MANAGESIEVE IMAP extension.**
- **Investigate adding CardDAV (Contacts) and CalDAV (Calendars) services.**
- **Investigate replacing RoundCube webmail with SOGo, providing a refreshed interface and ActiveSync (native Outlook) support.**

It seems likely that I will be asked to take the lead on the operation and development of Hermes if the panel determines that this service should continue. As a consequence, I have been speculatively planning the design and implementation of a next generation of the Hermes service, that both preserves the lessons from the past two decades of development and takes advantage of modern software and tooling.

Storage: The current Hermes email storage architecture employs a sharding model, whereby user email accounts are distributed across several independent backend email servers. These are made to appear to be a single large mail server by a front-end proxy, that redirects incoming user connections to the specific backend server that hosts that user's data.

Each storage server also uses application-level, asynchronous replication to maintain two other copies of the email accounts it hosts, resulting in a total of three copies of every email account being stored in two (originally three) distinct locations. This provides excellent storage resilience: in the event that a single locations becomes unavailable, or is even physically destroyed, it is possible, with some manual action, to continue to provide service with little, if any, loss of data.

However, because the data replication between locations is asynchronous, data consistency between locations may not be perfect. As a consequence, we have deliberately not implemented automatic fail-over functionality between sites, as this can cause emails to be lost.

This design also prevents us from offering a shared folder facility, where two or more arbitrary accounts can access the same mail folder simultaneously. This is because each individual backend server can only access data its holds itself; if two users are on different servers, they cannot be made to see each others data.

Both of these limitations could be addressed by using modern scale-out storage techniques and tools—such as the Ceph software-defined storage system already used to good effect in the UIS, academia, and industry. While it, too, replicates data between sites, it does so synchronously, and would allow all of our servers to have access to all users' data concurrently, making shared folders possible.

Scale-out storage systems can also allow for the more efficient use of storage hardware: with the use of RAID10, triple replication, and separately stored disaster-recovery backups, any given email stored in a user's account may be recorded 10-12 times with the current design. Scale-out systems, by contrast, can supports the use of replication, automatic tiering, and erasure-coding, meaning that we can likely halve the number of stored copies without impacting resiliency. We can further be able to improve storage efficiency significantly through the use of compression and sufficiently clever single-store email server software.

As a consequence, by building a new Hermes service on this more modern storage platform, we should be able to provide much bigger quotas, for the same cost, gain automated multi-site high-availability, and add new shared mailbox functionality.

Self-service: Currently, most user requests made to the Postmaster are to seek our assistance in recovering email that was expunged in error. Currently, Hermes is designed to retain deleted emails for 28 days before destroying them permanently, meaning there is a fairly generous window for recovering from such mistakes.

However, this process is manual, and not self-service—it requires a postmaster's manual intervention. It should be relatively straightforward to configure the Hermes server software to provide users with direct, read-only access to their own expunged messages, making it straightforward for users to recover those messages at will.

Filtering: The current Sieve email filtering options are robust and quite powerful, but fairly conservative in terms of supported operations. Updating the IMAP server software that we use to a more modern revision would enable the addition of more powerful primitives, such as regular expression support. Further, adding support for the MANAGESIEVE IMAP extension would allow users with sufficiently clever email clients to manage their filtering rules directly from that client, without needing to use the out-of-band control interface provided in Hermes webmail.

Addressbook and Calendaring: The most common criticism of Hermes, after our limited default quotas, is that it does not provide integrated addressbook and calendering interfaces. This may be possible to provide in the form of CardDAV (Contacts) and CalDAV (Calendars) standards-compliant services. Various software options exist in this area; if this functionality were desirable, it seems likely we could find a viable option to deploy.

Webmail: The current Hermes webmail system, based on Roundcube but including a number of important locally-developed enhancements, hasn't been changed substantially in several years, and it lacks some useful features, notably mobile device support. The SOGo webmail system is an attractive potential replacement, in that it provides a modern interface that supports mobile devices well, is open-source (and thus extensible and customisable), and has been demonstrated to be viable by other mail service providers, such as GANDI.NET.

Further, SOGo can reportedly re-export standards-compliant mail, calendaring, and addressbook services using Outlook’s native ActiveSync protocol, meaning that it may allow us to provide Outlook users of Hermes with the reach set of capabilities they could previously only obtain from a native Exchange service.