Informatics System Compromise - May 2021

Stephen Quinney <squinney@inf.ed.ac.uk>

Wednesday 2nd June, 2021

Abstract

In May 2021 the security of one of the self-managed servers within the School of Informatics became compromised and it was subsequently used to send large quantities of spam email. This resulted in the central mail services for the University being added to many Realtime Blackhole List (RBL) registers, consequently much of the outgoing email coming from the University was rejected for up to 24 hours after the source of the spam was blocked.

The purpose of this report is to examine the cause of that incident and consider what lessons may be learnt so that the situation does not recur. A timeline of events and details of the investigatory process are provided. The findings from that investigation are utilised as the basis for a discussion of what worked well and which aspects need to be improved. Finally a list of proposals are presented for consideration.

1 Background

Within the School of Informatics the majority of servers are managed by the School Computing Team. This means that, even where people own the hardware, they do not have any administrative privileges. This is a secure mode of management which suits most needs but occasionally the demands of research mean that greater access rights are required. In this situation the owners take over all management responsibilities for the server and we refer to these machines as "*self-managed*". The School provides a policy [9] which states the basic requirements to which administrators of self-managed systems must adhere. Currently approximately a quarter of the School's servers fall into this self-managed category. The Computing Team retains full control over the network configuration for all machines such as DNS, DHCP and firewall settings along with access to the internet itself. Although this incident involved a self-managed server, to ensure a high quality investigation it was carried out by a member of the Informatics Computing Team.

In addition to the firewall provided for the University network by Information Services there is a secondary firewall for the School network. This firewall has a "*default deny*" policy which means that external inbound access is only granted on request after due consideration and requirements are regularly reviewed. This ensures that the attack surface for the School's computing infrastructure is as small as possible. In this particular case the system was only externally accessible via the Secure Shell (SSH) protocol. To give some context, only 13 machines within the School are accessible via SSH, 8 of which are self-managed. SSH access was permitted to this particular host to allow access for external collaborators. Ideally self-managed machines are only accessed by using either the VPN service or by going through the School SSH gateway [4] but in this case the administrative process of acquiring visitor accounts for external colleagues was disruptively slow.

In a previous incident in November 2011 the primary Informatics SSH service was compromised. As part of the response to that incident a report [1] was produced, the report into this recent incident intentionally takes a similar approach. The earlier report resulted in the

creation of Incident Management Procedures [2] for the School and guidelines on the procedures to follow when investigating compromised machines [3].

2 Impact

This incident resulted in the mail service for the University, which supports both Office 365 and staffmail, being given a bad reputation by many Realtime Blackhole List (RBL) providers, for example SpamCop and Sorbs. This had a major impact on the ability of members of the University to send email to external contacts for a period of up to 24 hours. This is a rare disruption to normal service which will have caused confusion, wasted staff time and, in the worst case, may have led to people missing important deadlines. Nearly three weeks later it was noted by Information Services that, although the University mail service mostly recovered quite quickly, some mail to external sites is still being rejected. This means that IS staff are having to actively manage the routing of that mail to minimise the impact.

It must be recognised that dealing with this type of incident is also very disruptive to normal work for the Computing Team and, in this, case also for our colleagues in Information Services. We estimate that our response to this incident within the School has required about 1 week of FTE staff time. It also required a great deal of effort from IS staff who spent a day striving to get mail flowing again. Thanks to their hard work, and some clever thinking, they were able to redirect the flow of mail in such a way that most mail was being accepted again in much less than 24 hours. There must also be a recognition of the personal impact on the staff involved. Dealing with these situations can be quite stressful as people feel under great pressure to resolve the problem quickly.

It is important to note that incidents such as this have the potential to cause embarrassment and reputational damage to the University. It is becoming increasingly common for such incidents to be reported by the mainstream media. Furthermore, there is the potential for reputational damage to the School itself. In terms of the way in which our IT infrastructure is managed Informatics is in a somewhat privileged position. For example, much of our network traffic is allowed to flow unhindered by the University firewall. This gives us greater control and flexibility over how we manage our services compared to some other schools. This is permitted because over many years we have built a reputation for having secure computing systems which can be relied upon to not cause problems. Although in this case the server was self-managed, from an external point of view all our internet facing services will be seen to be owned by Informatics and will be held to the same standard.

3 Timeline of Events

Information has been gathered from various sources to give an overview of the incident and the initial response. This includes logs from the Informatics mail service and the Computing Team chatroom. The firewall access request is recorded in the Informatics RT Service ticket 107910 [6]. For details of the response from the Information Services team please refer to UniDesk ticket I210510-0471 and the alert they issued [5].

Thursday 29th April 12:08	External SSH access through Informatics firewall requested
Friday 30th April 11:12	External SSH access permitted
Saturday 8th May 03:31	Attacks begin on SSH
Saturday 8th May 06:30	First successful compromise
Sunday 9th May 21:28	System begins sending spam email
Sunday 9th May 22:10	nagios warnings regarding SMTP service
Monday 10th May 09:10	Issue recorded in UniDesk by Information Services
Monday 10th May 09:53	IS alert regarding "University mail relays blocklisted"
Monday 10th May 10:01	Informatics staff contacted via Slack chat
Monday 10th May 10:08	Informatics mail service stopped
Monday 10th May 10:31	Attempt to close SSH firewall hole
Monday 10th May 10:57	Informatics mail service restarted
Monday 10th May 11:27	Final access of compromised account
Monday 10th May 11:39	All network access blocked
Friday 14th May	Compromised system investigated

4 Investigation

The investigation into the compromise of security for this system was begun on Friday 14th May. We would normally aim to begin an investigation sooner but that was not possible due to Covid-19 restrictions on access to the Informatics server rooms.

The system is managed by academic staff, rather than the Informatics Computing Team, and does not use either of the standard DICE computing environments (Scientific Linux 7 or Ubuntu LTS 20.04). It also does not use any of the services provided (e.g. Kerberos, LDAP, AFS). Discussions with the owner revealed that, as it contains a design sample CPU, attempts to use Ubuntu 20.04 had been unsuccessful, due to stability issues, and it was instead using Redhat Enterprise Linux (RHEL) version 8.3. No confidential data was held on the system, it is just a standard RHEL installation plus various Open Source code which had been checked out from repositories.

The system is a completely standalone server so there is no shared user database and no shared file-systems or other resources. It is installed alongside an identical machine with the aim of doing distributed experiments between them and, as such, it is expected that frequent SSH connections would be made between the two machines. The second machine was not externally accessible and there is no suspicion of a security compromise but as a precaution network access was also disabled. All users of the system were informed of the incident and asked to change passwords on other systems where necessary as a precaution.

The system contains a number of SSD disks which are combined using the Logical Volume Manager (LVM). This is split into 4 volumes, 3 of which are associated with the RHEL 8.3 installation - root, swap and home - the fourth appeared to be left over from the previous Ubuntu install, it was ignored during this investigation.

4.1 The Investigation Process

To make it easier for the Computing Team to examine the machine, whilst it remained disconnected from the network, it was booted using an Ubuntu 20.04 installer on a USB device. This approach carries the risk of losing important information about the attack if it is only in memory but means the investigator can work in a known safe environment. Furthermore, attackers will often attempt to compromise the Linux kernel to hide processes and data in the file-system. By booting into a known safe environment a complete copy of the file-systems may be made. Copies of the root and home partitions associated with the RHEL 8.3 installation were made using the low-level "dd" tool. The copies were made onto an external USB drive,

this was rather slow but avoided the need to reconnect the system to the network. The data could then be easily examined as necessary without having to work in a horribly noisy server room.

4.2 Detection of the Problem

Within the School of Informatics services are monitored using nagios. On the Sunday evening problems related to connecting via SMTP to the mail relay server smtp.inf were reported on several occasions. By chance this was noticed at the time but not followed up immediately as it was considered to be likely caused by a temporarily high load due to legitimate email which does very occasionally happen. Also, as the Computing Team does not provide any out of hours cover there was no expectation that problems should have an immediate response at the weekend. The Computing Team first became aware of the spam email problem on the Monday morning when we were contacted by the administrators for the central mail service provided by IS. The sending of large quantities of spam email resulted in the University being given a bad reputation by some external spam monitoring services. That led to mail delivery being blocked by some external sites.

4.3 Initial Response

On becoming aware of the spam email issue on Monday morning the Computing Team responded quickly by disabling the local mail service. That resolved the immediate problem of large quantities of spam being sent from the School servers but could only ever be a short term measure as it also denied legitimate users. Users were notified of the service being unavailable via the front page of the computing.help site. All remaining spam was removed from the mail queue and after ensuring that the source had been correctly identified access to the mail service from that machine was blocked using explicit sendmail access rules before the service was restarted.

At that stage suspicions were raised that the machine might have been compromised. In particular the relatively recent opening of the firewall hole for SSH into this machine seemed like it might be linked to the incident. As a response to that the SSH firewall access was blocked as a precaution. This will have disconnected the attackers from the machine but potentially left it still able to send spam email. After further discussions it was later decided to completely disconnect the machine (and its partner) from the network. That ensured the machine was no longer a risk to the wider Informatics network.

4.4 How was it able to send spam email?

As there is an expectation of the machine being accessed by external collaborators it is hosted on a restricted subnet (`129.215.218.64/26`) which is within the Informatics firewall but is outside the inner perimeter of that firewall. Machines on that subnet cannot access services on any other Informatics subnets without extra firewall rules being added. It was discovered that, to permit connections from a few mail servers managed by Information Services, overly generous firewall rules were in place on the mail relay which allowed connections from a large range of addresses within the University. Thus the compromised machine was allowed to connect to the two Informatics mail servers - mail.inf.ed.ac.uk and smtp.inf.ed.ac.uk - via the SMTP protocol. To provide more fine-grained security the mail servers are also configured to restrict access. The sendmail service on smtp.inf was configured to allow unauthenticated relaying from any host with an address in the inf.ed.ac.uk domain, that consequently permitted the attacker to send the spam emails. The sendmail service on mail.inf was better configured to only permit relaying for certain trusted subnets but due to the combination of using just part of

the wider `129.215.218.0/24` subnet and the rather limited functionality provided in sendmail the attacker was still permitted to connect. Notably only 4 connections were made to mail.inf, the attacker seems to have focused on using smtp.inf. Information provided by IS staff revealed that the attackers also sent mail directly via the central mail services. It seems likely that the attacker identified all possible routes for sending mail using a combination of the Mail Exchanger (MX) and Sender Policy Framework (SPF) records for the system. IS staff added local firewall rules to block all access to their mail servers from the compromised machine.

4.5 Exploiting the SSH Service

In response to a request made by the owner on 29th April the system had recently been made externally accessible via the Secure Shell (SSH) protocol. The intention was that this would be used by external collaborators to easily access the servers for running experiments. External access via all other protocols was denied by the Informatics firewall. There were no limitations on access to this SSH service, it could be reached from anywhere for any user with an account on the system. Inspection of the configuration (included below) shows that the SSH service was listening on the standard port (22), public key (enabled by default) and password authentication were supported and login to the root account was permitted. There was no evidence of any local active defences to mitigate against attacks (e.g. fail2ban).

```
AcceptEnv LANG LC_CTYPE LC_NUMERIC LC_TIME LC_COLLATE LC_MONETARY LC_MESSAGES
AcceptEnv LC_IDENTIFICATION LC_ALL LANGUAGE
AcceptEnv LC_PAPER LC_NAME LC_ADDRESS LC_TELEPHONE LC_MEASUREMENT
AcceptEnv XMODIFIERS
AuthorizedKeysFile .ssh/authorized_keys
ChallengeResponseAuthentication no
GSSAPIAuthentication yes
GSSAPICleanupCredentials no
HostKey /etc/ssh/ssh_host_ecdsa_key
HostKey /etc/ssh/ssh_host_ed25519_key
HostKey /etc/ssh/ssh_host_rsa_key
PasswordAuthentication yes
PermitRootLogin yes
PrintMotd no
Subsystem sftp /usr/libexec/openssh/sftp-server
SyslogFacility AUTHPRIV
UsePAM yes
X11Forwarding yes
```

The number of users who are expected to have access to the system is currently limited to just three members of the School of Informatics. To allow management of the system all of those users have access to the root account via the Linux sudo utility by virtue of being members of the wheel unix group. The configuration for sudo showed that users were required to enter their password to execute commands so it was not possible for any compromised account to immediately become the root user.

On inspection of the home volume for the system it was noted that there is a fourth user account with the username “user”. This account has the UID/GID of 1000, which is always the first ID to be used when creating user accounts on Debian/Ubuntu systems. Discussions with the external engineer revealed that this account originated from the custom install image which was prepared for the machine. A simplistic trial and error process showed that this account was easily exploitable as it has an incredibly weak, and easy to guess, password.

Examination of the `wtmp` file - which records all user logins along with information about date, time and origin - revealed that the "user" account had been accessed at least 6 times from 4 different locations on the Sunday and Monday. The details were as follows:

```
user   pts/0   185.222.57.251   Mon May 10 11:25 - 11:25 (00:00)
user   pts/0   185.222.57.251   Mon May 10 10:58 - 10:58 (00:00)
user   pts/0   41.190.2.11      Mon May 10 06:36 - 06:36 (00:00)
user   pts/0   156.96.61.117    Sun May 9 21:36 - 21:36 (00:00)
user   pts/0   156.96.62.254    Sun May 9 21:26 - 21:26 (00:00)
user   pts/0   41.190.2.97      Sun May 9 21:20 - 21:20 (00:00)
```

Through further investigations it was noted that the Linux audit daemon had been running. Although impossible to say with any great certainty, examination of the logs suggests that the account was independently compromised on something like 12 occasions. There were 43 successful SSH connections using the "user" account from 24 different hosts beginning at 06:30 on Saturday 8th May and ending on 11:27 on Monday 10th May. The difference in numbers between auditd and `wtmp` is probably due to the connections not starting a full session.

Looking at the passwd and group entries for the "user" account reveals something quite interesting.

```
user:x:1000:1000::/home/user:/sbin/nologin
user:x:1000:
```

The shell for the account had been set to `/sbin/nologin`. At first glance this seems like quite a sensible method of disabling the account to ensure it is safe and unusable. It seems quite reasonable to naively assume that this would make it impossible to login using any method, not just for SSH.

Network access (but not firewall access) was temporarily re-enabled to allow examination of the behaviour of the SSH service. Testing on the compromised system has confirmed that any attempt to login results in the connection being closed before it can be used with a helpful message of "This account is currently not available". It is important to note that, although effective in blocking shell access, there is still crucial information leakage here. It only behaves in this way for a valid account when the correct password is provided so an attacker can still use this as a method for brute-force attacks on the user/password database.

In fact, the weak point here is specifically the SSH service. Although it is impossible to start a shell or execute any command our research has revealed that it is possible to initiate a TCP port forwarding session. An article on the askubuntu website from 4 years ago gives full details [7], note that article refers to a helpful blog article which is now only available via the Wayback Machine [8].

This shows that using the `/sbin/nologin` shell is **NOT** effective as a strategy for denying access to systems for specific accounts and must never be relied upon. Access to an account may be fully disabled in the passwd/shadow file but that will render it completely unusable. Much better is to explicitly specify the allowed users, or groups of users, and deny access to all others through either of the SSH configuration or the Linux PAM stack (or even better use both).

In this particular case the situation was further weakened by both the choice of the username and the "user" account having an exceptionally weak password which is clearly one of the first that is tried by most attackers. The lack of any active defences like fail2ban also meant that when the attacks began they were never deflected away which allowed them to continue as long as necessary to achieve success.

5 Discussion

5.1 What Worked Well?

In a situation like this it is very easy to focus entirely on all the negatives and only examine what went wrong, why that happened and where the fault lies. It is genuinely useful though to take stock of the entire situation and see what was done well. If nothing else the hope is that this should be useful to other sites who may wish to avoid going through a similar "learning process".

5.1.1 Restricted Subnet

The use of a restricted subnet is a good strategy as it effectively limits the scope for extending an attack into our internal network. Not all self-managed machines which are accessible through the Informatics firewall are treated in this way, we should consider if this approach can be expanded.

5.1.2 Firewall Rules

As previously noted, the Informatics firewall policy is "*default deny*" for inbound traffic. This means that our exposure to attacks via the SSH protocol is strictly limited. Although it didn't help in this particular case it clearly prevents attackers from being able to reach large numbers of, potentially insecure, self-managed servers.

5.1.3 Mail server access

It was already understood that relaying of mail should only be supported for some groups of machines. Limits were already in place to restrict access, in this case the restrictions just weren't quite tight enough.

5.1.4 Immediate Response

Once we were aware of the issue the relaying of mail via the Informatics mail service was blocked very quickly. A prompt response is very important in such situations to minimise the amount of time the service appears in the RBLs and to avoid the potential for reputational damage.

5.1.5 The Investigation

As mentioned before, this investigation was carried out by the Computing Team. This was done by experienced staff with the relevant knowledge and time to guarantee a high quality investigation. Current policy does not explicitly mention how security incidents related to self-managed servers should be handled, the policy should be modified to make this approach our standard practice.

5.2 What Could have been Better?

In an incident with such a clear cause as this one there is a tendency to focus on the single obvious issue. Whilst it is important to review that problem and ensure there is no recurrence of the situation it is also beneficial to review the entire incident response process. Major incidents are, thankfully, rare, when they do occur they shine a spotlight on current policies and procedures. They provide a test of how well the team can respond to the challenge. This will inevitably reveal some areas which could benefit from improvement.

5.2.1 Communication with Information Services

IS recorded the incident in UniDesk (I210510-0471) at 9:10 and issued an alert at 9:53 [5].

It was only at 10:01 that IS first informally contacted a member of the Informatics Computing Team directly via the Slack chat service. Once the School was notified the response was fairly rapid but clearly almost an hour was lost because this was not immediately reported to our incident response email address. It seems likely that IS staff are not generally aware of how we would like security incidents to be reported. Furthermore, a quick inspection of our computing.help website doesn't reveal any information on how to report suspected security incidents. Clearly the communication mechanisms between IS and Informatics staff need to be improved to ensure rapid resolution of serious issues. It was noted in the Computing Team chatroom that an Informatics "Incident Response" team in UniDesk would be useful to improve the handling of these situations. We should also provide relevant information for external users who may wish to report issues.

5.2.2 Immediate Response

Our immediate response was focused on stopping the flow of spam and it was only slightly later that the potential for a compromised machine was properly considered. At the time there was some indecision amongst the team over the correct response for dealing with a self-managed system. The response was somewhat ad-hoc and not done entirely according to our policy, it would be worthwhile to refresh the Computing Team's awareness of the guidelines. Also, if a lead investigator had been selected to direct the process it is likely that a more coherent and organised approach would have been taken. Our guidance on when a machine should be considered to be compromised, and thus should be disconnected from the network, would benefit from a review.

Another issue with the response was that the attempt to remove the SSH firewall access at 10:31 was unsuccessful. It was only at 11:27 that access to the machine was properly blocked when the network port was disabled. This was caused by a known issue with our LCFG configuration management system. There are ways to workaround that problem, that information needs to be added to the incident response guidelines.

5.2.3 Communication with Users

During the incident members of the School were not adequately notified of the fact that email to external sites would be delayed. Given the issue affected email we decided not to send anything using that facility. A note was placed on the front page of the computing.help website but nothing was added to the twitter feed. The use of twitter for any significant incident should be standard practice, we may also want to consider if it is possible to send important alerts using Microsoft Teams. The "Communication With Users" section of the Incident Management Policy does not mention twitter and should be reviewed to ensure it is clear and up-to-date.

5.2.4 Weak Management

Due to stability issues with the hardware support in Ubuntu an external engineer was engaged to help and they instead installed RHEL 8.3 with patched firmware. Given the existence of the "user" account it appears that the local staff who manage the server did not thoroughly review the state of the system after installation to ensure that it was reasonable. Whenever external staff are involved in the configuration of systems there must always be a thorough review afterwards before the system is considered safe to be deployed back into service.

5.2.5 Poor Oversight

When SSH access through the Informatics firewall was allowed there was no prior review of the security of the system by the Computing Team. This lack of oversight of self-managed machines means that they cannot be guaranteed to match the high standards of security expected of the DICE systems managed by the Computing Team. From an external point of view any machine within the inf.ed.ac.uk domain is an "Informatics Service", no matter who it is managed by, the expectation will be that they are all similarly trustworthy.

The Computing Team are currently in the process of producing a course on the University Learn facility to assist staff with the basics of managing their systems. For externally accessible services where there is a greater risk of accounts being compromised, such as SSH, it may be that an additional checklist of security requirements would be helpful.

6 Proposals for Improvements

It is useful to distil the previous discussions into a set of distinct proposals. Some of these proposals are for small, very specific, tasks which need to be done immediately, others will require more effort and potentially a great deal of discussion. The existence of an item in this list of proposals does not automatically mean that it is guaranteed to be done but it will at least be given some consideration.

1. Restrict access to mail servers:

The firewall holes which allow communication between the Informatics mail servers and the IS mail service must be restricted as tightly as possible.

Unauthenticated relaying must also be limited to the smallest possible set of trusted Informatics hosts.

2. Consider rate-limiting access to the mail services:

The relaying of mail via the Informatics mail services should be rate-limited to minimise the damage caused when an account or system is compromised.

3. Consider enhanced monitoring of the mail services:

Options for automated monitoring of the mail services should be considered. The aim would be to raise alerts whenever the mail servers appear to be used to send spam. This would allow the prompt identification of compromised systems or accounts.

The sheer quantity of spam email generated meant that the IS mail service was added to various RBL registers. That is what made us aware of the security compromise. It might be useful to consider, if messages had been sent at a much slower rate, how long could the compromised system have remained undetected?

4. Update the Incident Management Policy:

The Informatics policy for Incident Management [2] was last updated in August 2011 and now needs to be reviewed and updated. In particular the section on "Communication with Users" needs to be clarified and newer facilities, such as twitter, should be mentioned.

Once updated the information should be promoted to the Computing Team to ensure everyone is aware of the guidance available.

5. Review guidance on compromised machines:

The guidance provided on how to respond to potentially compromised machines [3] was last updated in 2018. A quick check shows that some of the links to useful sites and tools are no longer functional. There is also an implicit assumption that the machine involved is a managed system, guidance should consider the different scenarios and provide distinct procedures for managed and self-managed machines where necessary. All of the details should be reviewed in the light of recent incidents.

6. Improve Incident Response Communications:

We need to create an Incident Response Team for Informatics in the UniDesk system. This would allow IS to easily pass issues directly to us in a way which we notice immediately. To assist external people we should also promote a way of contacting the team on the Informatics computing.help website.

7. Review the policy on self-managed servers:

The current policy [9] does not include a statement on how to respond to a suspected compromise of security on a self-managed system. In particular there needs to be be some detail noting that the Computing Team must be permitted access.

8. Consider other self-managed servers:

Not all self-managed servers are on restricted subnets like the compromised system. We should consider whether it would be significantly safer if all self-managed servers with firewall holes were handled in the same way.

9. Require an Alternate Port:

Although this must really be considered to be "security through obscurity" there is a good deal of anecdotal evidence that configuring an SSH service to listen on an alternate TCP port massively reduces the level of attacks. Such a requirement for self-managed servers would create little inconvenience for users but may significantly reduce the likelihood of a compromised system.

10. Penetration Testing:

If the Computing Team had carried out targeted penetration testing on the compromised system prior to the firewall access being permitted the issue would have been rapidly identified. We should consider regular network scanning and penetration testing of all self-managed servers, not just those which are externally accessible, to identify potential security issues.

11. Security Checklists for External Access:

Before opening up external access to services such as SSH we could require system managers to complete a check list of requirements. This would ensure that important security configuration is not overlooked and help gauge the potential competence of the system managers.

12. Consider if direct SSH access is appropriate:

The School must carefully consider if it is appropriate for self-managed machines which are intended for experiments to be directly accessible. As previously noted, any compromise of security has the potential for significant negative impact on the School and University. Further to the impacts already discussed there is also a significant risk of leaking sensitive data.

It is rare that direct access to services such as SSH is absolutely essential. They can instead be made, almost transparently, accessible by using the School VPN or using the managed SSH service as a "jump host". If there are external collaborators they may be issued with DICE visitor accounts which permit access to those services.

13. Higher standards for self-managed servers:

Services such as SSH which provide the facility for full shell access inevitably increase the risk of the system being compromised. If we are to retain external SSH access to self-managed machines they must be maintained to the same high standards expected of those managed by the Computing Team.

Managing such a service to the necessary standard requires a good working knowledge of the service itself along with the basics for the particular platform (e.g. how to apply updates in a timely manner) and many aspects of the security frameworks provided (e.g. account management, PAM, auditd, fail2ban). In addition to the requisite knowledge, the managers of a system need sufficient time to regularly monitor the status of the service. Running services which are exposed to the internet requires constant vigilance.

7 Conclusions

Beyond identifying the cause of the incident - an issue which is easy to resolve - the aim of our review is to identify achievable positive outcomes which will improve the security of the computing systems within the School. Our review highlights a range of issues which require some attention. These issues can be grouped into 4 broad themes: mail service configuration, incident management procedures, communication and self-managed server policy.

In light of the way in which it was easily abused by attackers access to the Informatics mail service is now being thoroughly reviewed. A separate report will be provided which covers that aspect in greater detail. It seems very likely that other schools, and the central IS service, would benefit from carrying out similar reviews of their mail systems. The days of allowing unlimited unauthenticated relaying of email for the majority of machines on a network have surely come to an end.

Our review has highlighted that our guidelines on how to respond to such events need to be regularly reviewed and institutional awareness of the various policies needs to be occasionally refreshed. These events are relatively rare and unused knowledge quickly fades so the guidance must be of sufficient breadth and depth to allow staff to rapidly respond to an event. It must also be kept up-to-date to reflect the ever changing types of threats.

The Informatics Computing Team believe strongly in being open and transparent with members of the School and colleagues within the wider University IT community wherever possible. This approach to communications has many benefits, in particular it helps build strong relationships based on mutual trust. Incidents such as this will naturally erode trust, that erosion may be repaired by executing an effective review and clearly delivering on the recommendations. Reporting on these incidents also has the potential for the findings to have

a much wider impact and improve computing security across the whole University. Our review has highlighted that some aspects of our incident response communication procedures need improvement. It is inevitable that major incidents will occur again so we must endeavour to ensure they are handled more efficiently.

The primary issue which our review has revealed is that within the School a thorough review of the policy regarding self-managed machines is required. It is clear from the high number of recent cases that Universities are now a prime target for attacks, in particular ransomware and cryptojacking. Few academic staff will have the necessary combination of skills and experience to safely run internet-facing services and also the spare time to be adequately vigilant. Given the wide range of security threats which direct internet access poses in the modern world we must reconsider what exposure is appropriate and reasonable for self-managed servers.

Acknowledgements

I would like to acknowledge all the valuable contributions made to this report by my colleagues in the Computing Team with the School of Informatics. The response to this system compromise has been very much a team effort, in particular I am grateful to Ian Durkacz for clarifying the network configuration and Neil Brown for explaining the mail service configuration.

References

[1] *A report on the root compromise of the Informatics SSH service* - `http://www.dice.inf.ed.ac.uk/publications/2011/ssh_compromise_report.pdf`

[2] *Incident Management Procedures* - `https://wiki.inf.ed.ac.uk/twiki/pub/DICE/WebHome/incident-management-policy.pdf`

[3] *Compromised Machine Investigations* - `http://www.dice.inf.ed.ac.uk/org/CompromisedMachineInvestigation.html`

[4] *Informatics SSH Service* - `https://computing.help.inf.ed.ac.uk/external-login`

[5] *IS Alert 10083* - `https://alerts.is.ed.ac.uk/index.cfm?fuseaction=view_alert&alert_id=10083`

[6] *SSH access request* - `https://rt4.inf.ed.ac.uk/Ticket/Display.html?id=107910`

[7] *Ask Ubuntu article* - `https://askubuntu.com/questions/886044/ssh-broken-in-for-user-with-no-shell`

[8] *Blog article* - `https://web.archive.org/web/20200222023633/http://commandline.ninja/2012/05/06/binfalse-sbinnologin-and-ssh/`

[9] *Self-Managed Machine Policy* - `https://computing.help.inf.ed.ac.uk/self-managed-policy`