Crowd-Sourcing the Detection of Compromised User Accounts

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

Thursday, 20th March 2014

1 Background

I work as a Computing Officer for the University of Edinburgh School of Informatics. My primary role involves managing the base OS for all the Linux machines within the School and maintaining the configuration management system (LCFG) which is used to do that efficiently. The School has approximately 1170 managed Linux machines, of which 320 are student lab desktops, 450 are research desktops and the other 400 are an assortment of servers.

One of my secondary interests is the security of our systems at the user access level (e.g. shell or web applications) and that will be the focus of this discussion. I'm not directly involved with security at the firewall or networking level.

With an infrastructure of this scale it is imperative that the system security is managed efficiently. To assist in achieving this we heavily restrict external access using a default-deny firewall. External access to services (e.g. web or SSH) is opened up only where necessary. For example, login shell access to the infrastructure is channelled through a small number of gateway hosts. This makes it much easier to harden the configuration and monitor login activity without significantly inconveniencing users.

In 2011 we discovered that the Informatics network had become seriously compromised. This led to a full-scale review of our approach to system security. This proved to be a highly valuable process which I have discussed in previous talks to FLOSS UK. The review produced a set of proposals for improvement. One of the most important proposals was to enhance the logging of system events, in particular to send all system logs to a central log server. Previously system logs had only been stored on the disk of the local host which is a high risk strategy. Without a remote log server, in the event of a system being compromised the logs might be tampered with or removed entirely which leaves them unavailable or, at least, untrustworthy for forensic analysis.

The following discussions will examine ways in which we have greatly enhanced our system compromise detection processes by analysing the wealth of data gathered through this central log server.

2 My castle is bigger than yours

Until recently most efforts to improve system security have been almost entirely focused on defensive hardening. The basic tenet being that if we build a big enough castle with thick enough walls then no attacker will ever be able to break through. Coupled with this strategy is the idea that attackers will focus their efforts on the weakest sites so we only have to build defences which are stronger than average.

Much of what is considered to be the current "state of the art" is based on automatically hardening defences in response to specific threats. Typically this is done by running a tool, such as fail2ban, which monitors authentication logs and applies a complete access ban (typically this is done with a firewall or tcpwrappers) for any host which causes too many authentication failures.

This defensive strategy has one weak point which is that for any system to be genuinely useful it must permit access to users (i.e. some times we have no choice but to lower the drawbridge and raise the portcullis). This means, for instance, in the case of the automated attack response we have to allow for a certain number of failures within a certain time frame to handle the case where the user has a typo in their password or the caps lock key is turned on.

Having a strong digital fortress is essential and until that has been achieved it is not worth expending a great deal of effort on monitoring the activity of users.

3 The Human Perimeter

Once we have our nice solid concentric rings of defences it is tempting to sit back and believe that we are completely safe. This is a risky strategy though as it does not consider the true perimeter of the network. As noted previously, a network is only useful if it is accessible. If a network is accessible to authorised users then it is also accessible to attackers via those users. Your users provide a means of access to your network and thus they form the real outer perimeter of your network. They form the front line for most attacks and they will be the route through which most compromises of a system occur.

4 Account Compromise is Inevitable

Users are only human and thus they will occasionally make poor decisions and genuine mistakes will occur. They will visit web sites which install key loggers, they will fall prey to social engineering style scams, they will write down their passwords and lose the pieces of paper. Additionally, a highly-skilled well-motivated attacker is always going to find a way into your network. Consequently we have to accept that, irrespective of the strictness of any security protocol or the amount of care which is taken, the compromise of user accounts is inevitable.

Once we accept that the compromise of user accounts is an unavoidable aspect of running a useful service we are obliged to go beyond a simplistic digital fortress model. In particular, it becomes necessary to augment our approach by monitoring the activity of our users so that we can attempt to identify unusual behaviour patterns. When an account is compromised we need rapid detection and a fast response so that we can plug the breach and minimise the damage.

At the current time many sites are monitoring authentication failures but few are actively monitoring authentication successes. The remainder of this talk will discuss how Informatics moved beyond solely monitoring login failures. We will show that it is possible to enhance security without intruding deeply into the privacy of our users.

5 Why do we care about compromised user accounts?

Before considering how system logs should be analysed it is worth reviewing why we might be particularly concerned about the compromise of individual user accounts. On a Linux system a typical user does not have a great deal of power (they aren't root, for instance) so why should we worry?

Probably the most important reason to be concerned about the compromise of user accounts is that there are a multitude of local root exploits in the code we are running. Some of these might already be known (if you have not yet had the opportunity to update to the latest kernel, for instance) but many will not be in the public domain. We can mitigate against certain classes of these problems by using security frameworks such as SELinux but can never entirely rule out all privilege escalation routes. Once an account has been compromised an attacker may quietly keep checking your systems until they do become vulnerable. Are you going to spot the once-a-month login from some peculiar location which goes on to scan your systems for privilege escalation opportunities?

Another issue is that the attacker may have no interest in gaining root-level access to your system. Your systems may be of interest because they have access to greater than normal resources (many of our servers have gigabit connections through to JANET, for example). This could be abused to launch denial-of-service (DOS) attacks against other sites. Once launched, these types of attacks are normally spotted very quickly but it would clearly be better if they were prevented in the first place.

Finally, there are rare occasions when a particular user is specifically targeted. This might be because of the work in which they are involved, the data to which they have access or their profile within the wider community. There could be situations in which users are working on sensitive commercial projects where they have signed non-disclosure agreements and they are a target for industrial espionage. You might have users who work on projects which are of a contentious nature (e.g. they have animal rights implications). You might have users who are high profile (maybe a Nobel laureate or involved in TV programmes).

6 What is Monitoring?

As discussed previously, we have accepted that the compromise of user accounts is inevitable and we believe that this does pose a genuine risk to the integrity of our systems. We are consequently obliged to consider how we can monitor authentication events to identify any suspicious activity.

It is worth considering what constitutes the authentication monitoring process. Monitoring of authentication activity can broadly be split down into three separate phases:

Collection : The log files are gathered into a central location and, filtered to collect the entries related to authentication. These events are processed to get the data of interest (e.g. time-stamp, username, source).

Analysis : The authentication data is examined to find any successful logins which appear to be suspicious.

Response : Where suspicious behaviour is identified, access to the account is blocked. The authorities are notified. Forensic analysis is carried out to identify the intrusion method.

Typically sites will be collecting authentication data even if it is not in an immediately accessible form (e.g. syslog files on individual hosts).

Traditionally, where any monitoring is done, the analysis and response phases are not well separated and usually are carried out by a security team using some fairly basic techniques.

7 Data Collection Strategy

Within the School of Informatics we store our log files centrally and they are processed automatically to find events of interest. We extract various information associated with each login attempt. This is broken down into the following parameters:

- date/time
- source host
- destination host
- username
- protocol
- authentication method
- success/fail

We keep this information for 120 days in case we discover an account has been compromised and need to carry out forensic analysis of historical records. After that the data is anonymised and kept for 6 months to provide a longer-term base line with which we can compare current activity levels.

This data is imported into a PostgreSQL database from our raw log files using a locally developed system named BuzzSaw. We also keep the raw log files so that we have access to the entire data set if necessary.

8 BIG DATA

The Informatics log database currently holds approximately 12 million entries that are directly related to authentication. In the world of modern computing this quantity of information-rich data immediately provokes thoughts of data mining. We should be able to look for trends and attempt to identify unusual activity and suspicious behaviour. There are now several projects which aim to use artificial intelligence to carry out this type of data mining in a fully automated fashion. The aim being to automatically deny access to accounts whenever suspicious activity has been identified. This might seem like a very good thing but it can easily lead to a situation which we believe is highly undesirable. The recent furore over the dubious data collection and analysis tactics of government agencies (e.g. GCHQ and the NSA) has shown that people do not want to live in a surveillance state where their lives are tracked and monitored in fine detail. We should never lose sight of the fact that the data we have collected relates to individual people who have a right to their privacy. The collection of authentication records may seem reasonable innocuous but from that data it is possible to infer a great deal about the daily lives of our users. We

could be able to deduce when they start work, when they finish work, if they are on holiday or where they like to hang out when not in the office.

We want to demonstrate here that a desire to avoid intruding into the privacy of our users does not mean that the authentication data we have collected cannot be analysed but a careful, cautious approach must be taken.

9 Data Protection: The rights of the individual

It is worth reviewing the guidance laid down in the Data Protection Act which covers the collection of this type of data. This Act is regularly cited in discussions regarding personal privacy but how many people actually know the precise details?

The Data Protection Act gives rights to individuals in respect of the personal data that organisations hold about them. The Act says that: "*Personal data shall be processed in accordance with the rights of data subjects under this Act*". This is the sixth data protection principle, and the rights of individuals that it refers to are:

1. a right of access to a copy of the information comprised in their personal data
2. a right to object to processing that is likely to cause or is causing damage or distress
3. a right to prevent processing for direct marketing
4. a right to object to decisions being taken by automated means
5. a right in certain circumstances to have inaccurate personal data rectified, blocked, erased or destroyed
6. a right to claim compensation for damages caused by a breach of the Act

10 Profiling is hard

Clearly there are situations in which automatically analysing and profiling user behaviour is a sensible and necessary approach. I would be concerned if my bank did not flag up and block my credit card if it were to suddenly be used in China when it appears that I am still at home in Edinburgh.

Most people have probably been on the wrong end of this type of system at some point in their lives. I've had great difficulties ordering laptops online late at night. My father-in-law loves travelling in Europe, over the last few years he's had trouble paying in restaurants on the first night of each trip even though he has informed his bank prior to departure. Everyone has this sort of anecdotal evidence that reliable profiling of user behaviour is non-trivial. So, automated behaviour profiling systems should be used with great caution and only where truly necessary. In the case of my bank I am willing to accept the occasional inconvenience as I really do want to have some confidence that a genuine bad guy would also be blocked.

The solution here is to understand the risk profile of the assets you need to protect and to be aware of what collection and analysis strategies your users are willing to accept. In discussions with our users we have found that they are typically happy with the idea of their authentication data being gathered in case it is required as part of an investigation. They are much more likely to object to the concept of automated analysis, particularly if it is coupled to automated response systems (e.g. blocking their account when a login comes from an unusual location). It is fairly normal for users in the academic community to travel widely and collaborate with other academics across the world. If their work were to suddenly be interrupted by automated processes they would, justifiably, get quite grumpy. Always remember that your users have a right to object to decisions based on their personal data being taken by automated means.

We are left with the problem of how to monitor user login activity without relying entirely on an automated decision making process. Clearly we do not want to have a system administrator manually monitoring all activity and making decisions based on their judgements. This is unlikely to scale well and having another person doing the monitoring is likely to be seen as a intrusion into their rights of privacy and thus even less acceptable to users than the automated approach.

11 Involve the Users

It is well known that people are bad at evaluating risk and will often make poor decisions when it comes to the security of their computing accounts. Most often this comes down to a simple desire to "get stuff done". We can imagine a scenario where someone is on holiday but needs to check their email in case something really important occurs. They may be in the situation where the only access is that dodgy looking internet cafe which they would not normally consider going anywhere near. They want to be a good employee and keep abreast of developments in some critical project so they decide to ignore sensible security precautions "just this once".

Clearly what is needed here is a way to educate users about the risks they are taking. Merely laying down the commandment "thou shalt not use internet cafes" has been demonstrated to be a poor strategy. There needs to be a way of highlighting, and regularly reinforcing, the idea that some access methods are safer than others.

Additionally, rather than viewing users as a general menace and a risk to the security of your infrastructure they must be considered to be capable of providing part of the solution. Each individual will have much better knowledge of their recent account activity than you can ever infer from the authentication which has been collected.

12 Help the Users

Users have to be trusted to be responsible for their own accounts. Asking them to monitor the activity on their own accounts encourages and empowers them to take ownership and solves the problems of scale whilst involving a human in the process. To take true ownership of their computing resources users require assistance. The identification of activity which is indicative of an account compromise can only be achieved if the authentication data is presented in an easily accessible format. Any human interface must draw attention to entries which are less trusted and minimise the effort required to spot an event which is out of the ordinary.

The web interface to the authentication logs which we provide for our users is intentionally quite simple. There is a per-month view which gives a list of logins for each day (see figure 1). Each login event entry shows the time, the authentication method, the source host and the destination host. Particular attention is given to the source location. If it is outwith the Informatics network then it is highlighted in bright yellow to draw the attention of the user. We also interpret the address in the best way possible - if a host name is available we show that, otherwise we use whois information. We also add a flag icon to denote the country of origin and where city information is available that is displayed.

The location information is found using the free MaxMind GeoIP database (`http://www.maxmind.com/en/geolocation_landing`), this data has to be treated with a certain amount of caution as it is not perfectly accurate. It can also occasionally be confusing, and potentially alarming, for users as they may not understand that their ISP can route traffic through many different physical locations within the UK. For example, the IP address I am dynamically allocated by my ISP occasionally shows up as being from Newcastle instead of Edinburgh.

13 Good Corporate Environment

We feel it is essential that a corporate environment is created in which users feel they can report any potential security breaches without fear of recriminations or embarrassment. Typically people will realise they have been scammed quite soon after the event. Most people have probably experienced the "I shouldn't have done that" sensation at some point or another. The sooner a user reports this to the computing staff the sooner we can help the user fix the situation.

We feel that this approach of encouraging users to be involved in monitoring the security of their account also encourages user to feel free to report any other security problems they have experienced.

14 Create Good Habits

It is not sufficient to just provide users with the ability to monitor their own account activity. Given human nature there would be an initial surge of interest followed by a slow decline of involvement. People

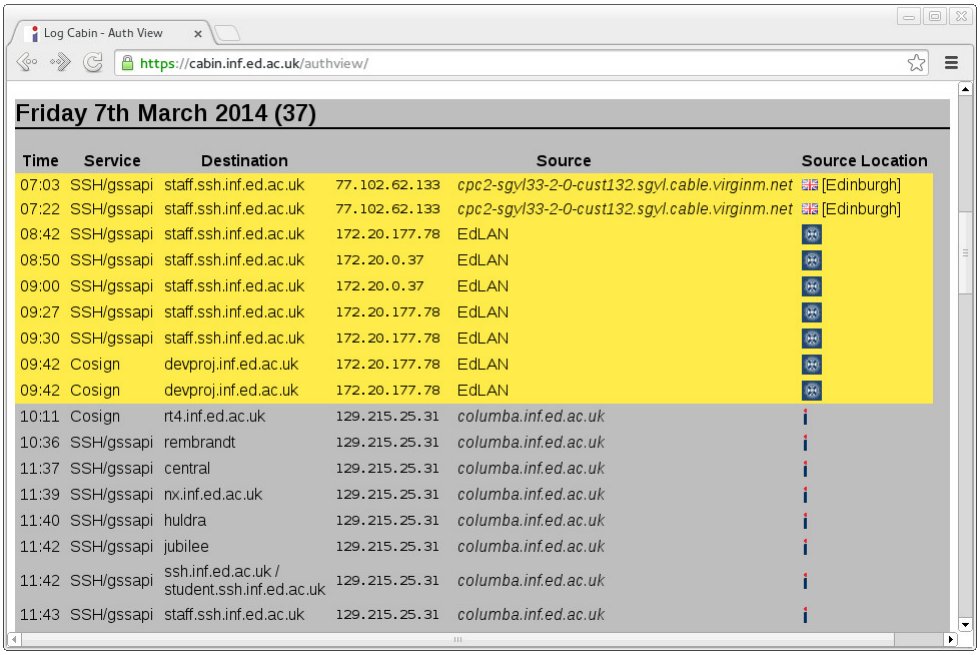

Figure 1: Excerpt from my authentication log

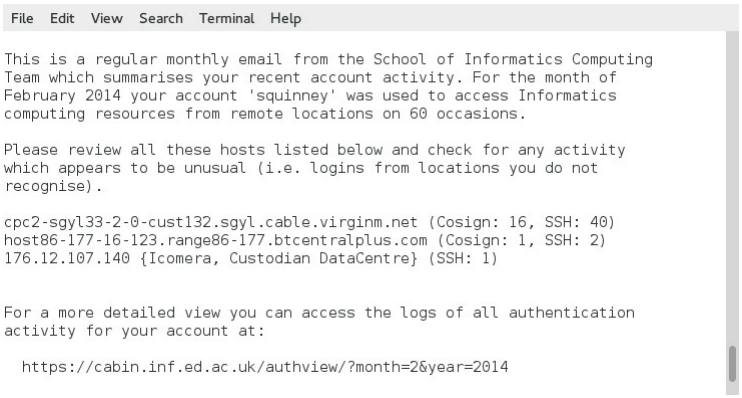

Figure 2: Excerpt from my monthly summary

need to be actively encouraged to keep checking on a regular basis. The aim is to reinforce the process until it becomes a good habit. I suspect that most people check their bank statements whenever they arrive to ensure there have not been any fraudulent transactions. People tend to adopt good practices when they see that there is a lot at stake (e.g. their monthly pay or life savings). They will also adopt them more willingly when the process is simple. A paper bank statement arriving through the post greatly simplifies the process of checking, would as many people check their statements if they were online only?

A monthly "check your login activity" reminder email is likely to very rapidly come to be considered as just another piece of spam. Instead we decided to achieve our objective by sending monthly summary emails (the equivalent of the bank statement) to all users whose accounts have had some login activity in the previous month. The information contained within the monthly summary (see figure 2 for an example) is a heavily reduced down version of the logs. This information-rich report is designed to help users quickly spot anything out of the ordinary and trigger interest in viewing the full logs.

This system has been in service for just over one year and we believe it is producing good results. We have not yet discovered any genuine unauthorised access but there have been many events which have provoked the interest of our users. Typically a user will spot a login attempt from an unusual location or at an unusual time and will contact our support team for help verifying the event. It might be that they just require a small amount of additional information (e.g. the name of the ISP which provides internet

access for the IP from which they connected). If the login was via the web we can usually provide extra information from apache log files such as which pages were visited, the client operating system version and the web browser name and version. This information is usually enough to help them remember why they were checking the Informatics website from a random night club at 2am using their friend's smartphone!

15 Future Developments

The current interface is intentionally quite simple but there are clearly a few ways in which we could improve the user experience. We had quite a lot of useful feedback from our users when we launched the service and that will direct our future developments.

White lists : One particular issue that was suggested by several people is the ability to white list particular source hosts. For example, where a user has a static IP address for their home internet access. In this case there is no need to highlight the address as being untrusted. Additionally we could allow users to supply their own "tag" which would be used in place of the country/city location indicator. This could help to substantially increase the signal to noise ratio.

Alternate Views : There was also considerable interest in the provision of alternative views. The current monthly/daily view is very straightforward but if we want users to more easily spot an unusual source location it might be better to group the logs by source location and show counts for each. This would provide a very rapid way to spot, for example, a single odd login from another country which might have otherwise been lost in a very long full listing.

Report Issue Button : We could make it even easier to report suspicious login events by providing a button for single-click reporting.

Automated Analysis : Having said that we don't want to do automated analysis there are some ways in which this could be used without intruding into the privacy of our users. We could have a process which monitors for unusual logins and then contacts the user directly by a trusted mechanism such as SMS or email. This still avoids the issues with automated response systems and retains the idea that only the user themselves is involved in the process of monitoring their account.

Monthly Prizes : We could run a monthly random prize draw for all users who have viewed and verified the login activity for their accounts. This is a fun strategy which is probably quite likely to keep more people interested in the longer term.

16 Summary

We have discussed how the building of a strong digital fortress is good practice but cannot be considered a completely sufficient system security solution. The threats posed to our infrastructure in the modern computing world mean that we have to treat the compromise of user accounts as an inevitable aspect of running a useful service. Given the serious consequences of any account compromise we feel we are obliged to monitor authentication activity so that unauthorised access is prompty identified. We have described a strategy which we feel provides a solution to this problem without intruding into the privacy of our users. As an alternative to traditional monitoring methods, the new approach we have introduced treats users as part of the security solution rather than just considering them to be the root cause of the problem. This new partnership model has empowered our users to take responsibility for the security of their own digital resources. We believe that by providing our users with the necessary tools and trusting them to regularly verify their authentication activity we can reinforce the idea that their digital resources have significant value. This new approach also has further benefits, it is highly scalable and is much less prone to the false positives which are associated with many automated monitoring systems. We have also identified a number of ways in which we can further enhance our monitoring systems as we continue to work on the Sisyphean task of keeping our systems secure.