What All Is ITS Logging/Monitoring?

What Information Is ITS Logging and Monitoring, To Keep the Campus Safe?

2017 was a watershed year for cybersecurity. In the wake of highly publicized breaches like the Equifax hack, new vulnerabilities like Spectre and Meltdown, and widespread ransomware attacks, many Carleton users have asked what all ITS is doing--in particular how we stay aware of what is happening on our network and help keep users as safe as reasonably possible.

The purpose of this page is to summarize the basics of what we are doing and give users, and internal ITS staff, a clearer sense of what we do (and don't) know about the security status of our network, and what we can and can't find out. Not all of the links provided here lead to user-visible pages. Some contain sensitive information. If you have questions, call the Helpdesk (x5999) or talk to the campus IT security officer.

Please note: In general, ITS does **not** track intimate details of what individual users are doing. Rather, we log and track normal activity in aggregate, and respond to things like exceptional activity spikes, indications of compromise, and malware. *We look in detail at individual activity **only to the extent needed to respond appropriately**.* (For example, ITS may respond to an alert that a user has logged in simultaneously from two different countries.) We may also take actions like notifying a user if it appears that their account has been compromised and, rarely, locking their account temporarily, to try to limit damage to the user's information and resources.

Asset Tracking

Centrally, ITS tracks what users and departments are assigned what hardware, and we also track what hardware needs replacement and when. This is done through asset management software (currently WebHelpDesk), and is keyed to the unique Carleton ID (that is, the CCID) of the hardware, i.e., to the number on the barcode sticky affixed to most deployed hardware. Everyone in ITS can see this data. We need it to know exactly where to go if a particular physical asset needs attention or replacement, and to tie information to that asset, and track it, if a user (or in some cases, an automated system) notifies us of a problem.

Patching

Windows machines on campus take their operating system updates from a local update server (WSUS). That update server tells us what machines have updated, when, and which specific updates were applied. We need this information in order to ensure that everyone's computer is up to date and not vulnerable to any obvious intrusions or attacks. Out of all the various defenses users can employ to combat such attacks, simply keeping devices up to date is by far the simplest and most effective. See also below on the KBox, under Software Deployment, Packaging.

Software Deployment, Packaging

The software inventory, operating system, and general configuration parameters of all computers deployed to users and computing labs are recorded in software we purchase from Dell Computing, called KBox. The KBox helps us, among other things, package up software for easy installation. Often a particular piece of software will require a special license key, or it will need to be pointed at a particular device on campus. The KBox can automate these configuration steps. It can also tell us if, for example, a piece of software needs updating and it can (usually) perform the update. This is particularly useful when a serious security issue has been discovered.

ITS staff and a few student workers can see KBox data. We need to see it in order to understand and diagnose software issues that users call us about. We also use it to get software installed to all the right places in the right ways, with the right license keys and settings. And we use it to find machines that are not updating their software correctly and are therefore exposing users to possible compromise, or opening the user and Carleton up to licensing violations.

Cloud Lock, Spirion

Files in Google Drive, and on storage internal to our network, are scanned for personally identifiable information, like social security numbers. In the case of Google Drive, CloudLock does the scanning, and these scans are fully automated. Users get direct email if a possible issue gets uncovered. In contrast to CloudLock, Spirion (formerly Identity Finder) only runs under user direction. Users, that is, install the software from the KBox and run it when desired. It also uncovers personally identifiable or otherwise sensitive data, but works on local and networked file storage, that is, things on your local computer, rather than in the cloud.

Keyserver Logging, License Management

Data is logged regarding what software is executed, where, and when, via keyserver software, Sassafras. We do this, when needed, for audit and compliance purposes. Sassafras also tracks, indirectly, when a machine is in use. General information at this level is visible to a handful of people including desktop computing and software asset management specialists, the IT security officer, and the data warehouse administrator.

Authentication

All attempts at authenticating are logged, not only by individual computers, but also by applicable enterprise applications (e.g., Colleague, OnBase, Advance, Slate, the Carleton website), domain controllers, our web-based login pages, and associated dual-factor authentication services (Duo). These systems need to record who logged in, when, and from where not only for auditing and troubleshooting purposes, but also to afford us an extra measure of assurance, when the data is pooled and analyzed, that the people logging in are who they say they are. Access to this data is limited to application managers and/or to a few core systems staff and the IT security officer. Alerts are generated when anomalous logins (rapid spikes, logins from impossibly disparate physical locations, lockouts due to dual-factor failures, etc.) are detected.

Network-Level Logging

In general, nearly all devices through which Carleton network traffic passes log that traffic. These logs typically include the source and destination IP address, MAC address, and various other relevant details. Additional detail gets logged for WiFi via PacketFence. This data, collectively, allows us to locate and fix problems and bottlenecks, and it helps us diagnose problems when we (or our users) discover them. We normally don't look at this data in detail unless there is a problem, such as a user experiencing frequent WiFi disconnects. The number of people who can see this data is very limited (a few core systems and networking staff and the IT security officer).

The Firewall

Additionally, ITS protects the campus with a central firewall, which inspects traffic as it flows through and looks for infiltration attempts, brute-force scans of our network, and other hostile activity, and logs and/or outright blocks that activity, depending on its severity. ITS does not examine logs produced by the firewall, except when performing forensics, troubleshooting, or examining alert notices. The number of people who can do this is limited to a few core systems staff and the IT security officer.

Machine-Level Logging

Individual machines (Linux servers, Windows servers, and Windows desktops) also log what they are doing, locally. There are, for one thing, antivirus and local firewall logs. There are also general system logs. On administrative machines maintained by Carleton these latter logs are written not only to local disk, but also forwarded to a central syslog server, then on to a central repository (Splunk). Windows program executions are recorded separately, and for some departments that handle a lot of sensitive data and have requested it, we actually limit what programs can be installed and run (AppLocker). ITS knows nothing about what the programs being executed are really doing. But we do use the basic information provided in logs to detect malware-related infestations, and also to perform forensics when and if hostile activity is discovered.

Administrative Windows users desiring a higher level of monitoring, or who have access to sensitive data and require such monitoring, may install an additional Splunk Universal Forwarder, which will perform extra logging. Sysmon may also be installed in addition, for users requiring the highest level of logging and auditability.

Log data forwarded to Splunk is particularly useful in detecting campus-wide events, and for looking back to see what happened, even if attackers have erased logs on the machine(s) they attacked to cover their tracks. Forwarded log data also allow us to infer what “normal” behavior is and construct alerts when something abnormal is detected, like a login from a staff member in France, and a login from that same person a few minutes later in China. The number of people who can perform this sort of analysis and set up alerts is very limited (a few core systems staff and the IT security officer), and once set up, the alerts only fire when an anomaly is detected. When that happens, ITS's analysis process only brings in what detail is needed in order for us to take appropriate action, like verifying that the anomaly is in fact due to compromise, informing users as needed, and locking applicable accounts temporarily to prevent anything worse from happening.

Richard Goerwitz
January 2018

- Security Incident Response and Breach Notification Procedure
- What All Is ITS Logging/Monitoring?
- When and why does my password expire?
- Internet Explorer Security Zones
- McAfee VirusScan Mac