



EARLY WARNING SERVICE FOR CYBERSECURITY INCIDENTS

A PRACTICAL APPROACH FOR SIGNIFICANTLY REDUCING
TIME-TO-DISCOVERY AND IMPROVING SECURITY POSTURE

A CASE STUDY IN HIGHER EDUCATION



TABLE OF CONTENTS

Executive summary	3
Main takeaways	3
Evaluating the usefulness of external security notifications	4
Project timeline	4
Project outcomes	5
Time from first observation to fix	5
Insights gained from continuous external monitoring	7
Asset definitions for monitoring	7
Notification types	8
Compromises	8
Vulnerable services	8
Building Early Warning Service	9
Future work	10
Conclusions	10



Executive Summary

Many IT teams and management at colleges and universities of all sizes are struggling with increasing cybersecurity incidents, the impact of COVID-19 and the expanded e-learning environment. Lack of resources and tight budgets are preventing them from taking action. At the same time, both the likelihood and costs of a breach are increasing.

There is a wealth of valuable external security information actionable for IT teams but not easily accessible. It is often too expensive for individual enterprises to acquire the data and too cumbersome to filter out what is relevant. Arctic Security offers a passive early warning service solution that provides clients with information on active threats and vulnerabilities from publicly available data on the Internet. We believe that democratizing access to these resources is demonstrably beneficial to colleges.

This white paper presents the results of a six-month project that evaluated the usefulness and efficacy of that external information for cybersecurity notifications. Arctic Security conducted the study as a collaborative project with the Independent College Enterprise, a consortium of 10 private colleges in the USA's Appalachian region.

The colleges participating in the study were impressed and surprised by the amount of useful information available from these external sources. They appreciated how convenient it was to start consuming the data, which revealed easily fixable gaps in their security posture and allowed them to track new kinds of malicious activity in their networks.

Four main takeaways

The average time-to-fix dropped from 81 days at the beginning of the project to just 1.3 days for the participants.

With a reporting accuracy of 98%, there were very few false positives in the six months of the project.

The number of weekly security issue observations for the participants was halved compared to the control group.

The security posture of the participants significantly improved without further encumbering IT teams.



It was a surprise that we actually had things that needed to be cleaned up. I thought we had fairly good security practices. This exercise showed us where we could improve.

– CIO of a consortium college





Evaluating the usefulness of external security notifications

Arctic Security is a Finnish cybersecurity company that has previously focused on helping national cybersecurity authorities by providing an incident notification platform. Having seen the effects of such systems in critical infrastructure space, Arctic Security is now focused on bringing the capability to a broader audience of MSSPs, enterprises, and universities.

Independent College Enterprise is a consortium of 10 private colleges in West Virginia, Virginia, North Carolina, Tennessee, and Alabama. Arctic Security proposed the joint project to the consortium to evaluate the usefulness and efficacy of external cybersecurity notifications.

Five of the consortium member colleges chose to participate actively in the project. Neither the participants nor the control group received any remuneration from Arctic Security, as the project's goal was to evaluate the notifications without bias.

Notifications were sent over a period of 21 weeks and delivered via email when they occurred. A notification email would contain one or more events. We use the terms 'incident' and 'event' mostly interchangeably in the text, but a 'compromise notification' is always considered to be an incident. The colleges committed to responding to the reported incidents within a few days of the notification and documenting and describing the case and where it happened. Arctic Security advised in responding to a few of the incidents.

Arctic Security selected a control group with similar characteristics to validate the project's findings. A set of private colleges in the same region as the participants made them as comparable as possible. There were two selection criteria for the control group members. First, their internet-facing assets needed to be discoverable without requiring the project team to solicit information directly from the colleges. Second, they would need to have at least one notification in the full observation period to ensure that the networks attributed to them are actually in use and that each one could have an event.

All five participating colleges had incidents and received notifications during the project, although there was significant variation in the types of reported issues with individual colleges. In the set of 31 potential regional colleges in which we could find their asset information online, 18 had

at least one incident during the entire 42-week observation period and were included in the control group.

Project timeline

The project's observation period started from the beginning of 2020, while the first notifications for the participants began in the second week of May. The observation period included in the project analysis was divided into two halves for data analysis, spanning 21 weeks before and after notifications were sent to the colleges. The past observations had already been collected and stored in a global historical database maintained by Arctic Security. For analysis, the data was retrieved and then correlated with both participants and the control group.




Fig. 1: Daily observations of unique IP addresses with likely compromised systems in networks belonging to US universities and colleges in 2020. A cluster of anomalies of up to 6,000 unique daily IPs were observed March-May, likely artifacts of other external malicious activity.

The project spans a relatively long period that includes significant academic calendar events, including the summer break. The summer break period is marked with a decreased number of events that would merit a notification. There were also significant changes in the overall operating environment due to COVID-19-related issues, starting from around week 14.

As shown in Fig. 1, we could observe a substantial change in the incident patterns in the background of all academic institutions in the USA beginning in the middle of March, where the overall trend was an increase in security issues that became visible from the internet. In the months leading to the project, the number of reported security events attributed to the participating colleges was higher than the control group's average.



We observed a significant change in the rate of events at the onset of the pandemic for the monitored colleges, visible in Fig. 2. However, the participants and the control group exhibited opposite trends. This shift may be due to how these colleges approached the switch to remote work and teaching. For the control group, events spiked to a level roughly three times higher than the average earlier in the year, while the participants, in turn, recorded fewer events.

By the time the project's notifications started on the second week in May, the pandemic's effects on security incidents were normalizing to a new stable level, as can be seen in Fig. 1. The onset of the summer break led to an overall decrease in events for all monitored colleges, reflecting the degree of activity in those months.

Project outcomes

In the first weeks, the participants received a significant number of notifications, which were resolved and documented. Warnings were sent to the participants on a daily schedule at 09:00. After that initial period, as new events came in, the colleges resolved them quickly, and typically a notification would typically only occur for a few days before the recipient handled it.

The observed trends in the amounts of notifications support the notion that you can't fix what you can't see. Many of the issues reported to the participants were topics that do not spontaneously raise attention. If they weren't already taken into account already when IT deployed its services, the staff




Fig. 2: Number of weekly notifications per college during the project, and the resulting trend lines for participants and control group.

“It was surprising how much you were able to inform us. I have 2,500 devices here, and you were able to say, “you’ve got one that’s really bad here. It’s talking to a really bad server”

– IT Director at a consortium college

might never have noticed them. All of the participants had regular updates and patch schedules.

Still, sometimes things fall through the cracks, and a particular patch doesn't get applied, a vital patch gets delayed, or some systems are misconfigured and allow unexpected behavior. Continuous external monitoring of assets catches these events and lets the IT team know what needs to be fixed.

Time from first observation to fix

Another way to look at the project results is how many days an observed incident is visible to the world before it gets resolved.

Industry estimates for the time that it takes for organizations to detect and fix compromises and vulnerabilities range from 50 to 300 days, depending on the type.

Reducing this time-to-fix delay is essential for preventing damages, data leaks, and extensive costs associated with breaches. One way to achieve this is by informing organizations about these problems directly.

An issue that stops being visible could be the result of different outcomes since fixing the problem, whether on purpose or by accident, looks the same from an external perspective. A changed IP address or a retired server looks the same as well; the host stops being reported as being vulnerable or compromised. For the participants, we know the reason – the problem was fixed – but for the control group that assertion can't be made.



Arctic Security analyzed the project results to see what kind of impact the direct notifications had on time-to-fix. You can see the results in Fig. 3, where a general upward trend is apparent.




Fig. 3: Average number of days elapsed for an observed incident or vulnerability before it is fixed.

The upward trend in the early months of the year is an artifact of the data collection method since the collection starts from January 1, which doesn't allow us to see how long the issues – first caught and fixed on that date – had been present. Over time, the time-to-fix settles on a level that is typical for the measured organizations.

Fig. 3 also shows a spike of resolved cases for the project participants in May. It resulted from colleges dealing with the reported security issues that had been present in their networks before the project started.

Those issues had persisted for 81 days on average at that moment. After that initial phase, the time-to-fix begins to approach one day for the following months.

The reduction is the result of the colleges generally fixing any newly detected issues after the first notification. The average time-to-fix for the participants starts to rise again in September, as seen in Fig. 3.

One of the participating colleges decided not to address some of the reported low priority security issues in July-August as those servers were soon due to be upgraded or decommissioned in the autumn. When the college eventually resolved those issues later in September and October, it raised the average time-to-fix time for the participant group for those months. The decision to not fix a security problem doesn't remove the risk; even though, after triage, it may not always make sense to dedicate limited resources to eliminate the risk for a very short-term benefit.

We did not remove any known factors from the data in order to remain as accurate to the situation as possible. Arctic Security continued to notify participants of the issues until they fixed them, which shows up in the analysis. For the control group, the steadily climbing time-to-fix rate approaches industry averages of more than 100 days, and it is likely to continue at that level unless they find a way to detect and fix the issues.

The time elapsed from the time of the incident to its detection is a crucial metric for any organization to evaluate the security measures' effectiveness. In this project's context, it had additional significance, as quick reaction times also help to document findings before the tracks of an incident become too old. Swift action increases the likelihood that the case will be appropriately investigated and resolved.

It took more time for the participants to react to the notifications in the first weeks of the project, but this improved quickly. After three weeks, participants usually achieved the target response time of 2-3 days. Arctic Security debriefed the college project liaison person once they had resolved the case to document it.

Average times shown in Fig. 3 only tell us half of the story about how long incidents and vulnerabilities can remain undetected. Table 1 lists by month the number of days for how long the most persistent incident or vulnerability had been visible in the participants or controls network. To be counted for that month, the issue had to be fixed (participants) or disappear (control). These are the cases that are the worst offenders, in contrast to the average values presented in Fig. 3.

Month	Participants	Control
Jan	0	0
Feb	0	1
Mar	61	9
Apr	0	96
May	122	56
Jun	53	165
Jul	2	186
Aug	1	154
Sep	96	138
Oct	71	226

Table 1: Largest number of days that an incident or a vulnerability was observed to have been present in the network.



Insights gained through continuous external monitoring

As the project advanced, it became apparent that it is crucial to have the ability to locate the systems in the internal network based on the external IP address, detection time, and communications port that was used for the reported activity. For the colleges that already had that capability in place, reacting to these notifications only required them to learn to use their existing systems to pinpoint the local system in their network. This investigation was not an everyday activity before the project, so the first notifications triggered a learning activity to investigate the case.

In some cases, locating the system required installing third-party log collection systems as the existing security systems did not provide adequate access or precise enough logs of the internet traffic. The college's IT staff then used these log collectors and SIEM systems to locate the offending internal system after a new notification.

During the project, identifying this need led to one of the colleges highlighting the use case while acquiring and provisioning a new firewall system. The old system did not allow proper searching of the logs. After the deployment of the new system, they were able to find the culprits in their networks with ease.

During an incident investigation, one of the colleges learned that the way that their internal network was built prevented the discovery of the offending internal system. They achieved partial visibility into parts of their networks in their investigation. Still, full visibility would require changes in the network architecture, which is difficult to accomplish until the next round of general upgrades.

Working with incident response cases also reveals security-related gaps in the architecture. It's essential to be aware and test for these kinds of issues while designing and implementing the network architecture, but it is not a well-known best practice yet.

Investigating and fixing these low-risk problems is an important way to advance the IT teams' knowledge about IT security matters and help the teams to understand their network better.

When there is a severe incident later down the line, the preliminary work has been done. The team can respond and mitigate the problem without first overcoming obstacles unrelated to the issue.

One of the persistent issues for a participating college turned out to be a student engaging in internet scanning activity from their campus. The IT team was surprised that this kind of information was available to them in the notifications. It isn't easy to know whether the student's motives are black or white hat activity and whether they also targeted the college network. External monitoring reveals scanning when it is also targeting the internet at large but can't tell you about scanning activities that are internal to the network.

While it is not strictly illegal in the USA to scan other networks in search of weaknesses and exploitable services, it is usually against the campus network's acceptable use policy. Indeed, it merits a serious discussion with the student. Being aware of this kind of activity helps the college prevent damage to their services and mitigate legal risk. It may expose them to civil legal liability if the victim of the scanning and exploitation decides to pursue legal damages.

“ *“It could make it easier to sleep at night. There are things that you know and things that you don't know. One of the things that you don't know is what people are doing on your network in the middle of the night.”*

– CIO of a consortium college

”

Asset definitions for the monitoring

For external monitoring to work, there needs to be as accurate an understanding as possible of the college's externally visible assets. The precision of the asset definition differs between the project participants and the control group. More accurate coverage was possible since the project participants shared their understanding of their assets with Arctic Security.



The project's monitoring scope included all network ranges used by the employees, enterprise services, and the students. Some participants also shared information about their cloud infrastructure assets, which were monitored in the same way. Common for all the colleges, Arctic Security also monitored their top-level .edu domain if the domain appeared in any known malicious URLs or phishing campaigns.

For the control group, we based their asset definition on publicly available information extracted from online databases. This data acquisition method introduces some uncertainty to the results, as the online databases do not necessarily cover all the assets that are in use by the college. Often some networks that are used by a college are labeled to belong to their ISP, which makes external discovery difficult. For example, online databases would not have revealed complete asset information for two of the consortium colleges participating in the project, attributing their networks to the regional service provider.

The lack of knowledge means that the results for the control group are not likely to be complete. While captured and archived by Arctic Security, some of their security issues have not been properly attributed to them and omitted from the data analysis. That does not significantly impact the usefulness of this project's results, as this systematic error is likely to make the control group appear to be doing better and be more secure than the actual situation. However, it is relevant for interpreting the results when comparing the participants with the control group.

Notification types

The project focused on informing the participants about security events in two main categories, compromised systems, and vulnerable services.

Compromises

First, there were notifications about cases indicating a compromised system in the college network. While Arctic Security collects data from various public and private sources, Arctic Security mostly sourced the data for the notifications during the project from a commercial partner. We observed several types of incidents that fell into this category: botnet drones observed in the network, vulnerability exploitation activity originating from the network, and vulnerability scanning (scanner) activity

originating from the network. When a system is reporting this kind of action, it is likely already compromised.

Each of these incident types typically results from a system under the control of a malicious third party, which is the primary reason to notify the network owner. Sometimes scanning and exploitation activity can also be intentional, as was the case with the student.

Compromise type	Participants	Control
Botnet Drone	8	49
Exploitation	3	0
Scanner	39	168
Total	50	217

Table 2: Reported compromise types over the observation period.

Examples of findings in this category included a video camera system provided by a vendor in the campus fraternity network that was compromised. Neither the college nor the vendor was aware of it. It was a typical IoT compromise, where insecure hardware gets exposed to the internet and then promptly abused. It was also unnecessarily accessible from the internet, which was remediated after the notification. Another case was a malware-infected file server, which was likely being used for crypto-mining at the time of discovery. That case could have later escalated into a data leak or ransomware.

Vulnerable services

The second category of notifications is about current and potential vulnerable services. Typically, the source material for these kinds of notifications is mined from services such as Shodan.io, among others. These issues range from recently disclosed vulnerabilities in applications, such as weaknesses in Microsoft's SMB3 identified in 2020, or services vulnerable to the SSL Freak attack. Freak is now a five-year-old vulnerability that allows man-in-the-middle attacks.

Remediation can be as simple as whitelisting the service for added security, as was the case with a legitimate SFTP service. Patching the service to the latest secure version or a configuration change preventing the use of an insecure authentication options are also typical remedies.



Vulnerability type	Participants	Control
Expired x509 cert	24	41
Freak	23	11
IIS WebDAV	0	5
Moxa NPort	0	6
Open ARD	0	1
Open BACnet	0	3
Open IPMI	24	0
Open mDNS	0	27
Open MySQL	0	14
Open Telnet	0	1
Open VNC	0	24
Palo Alto SAML	0	25
Recursive DNS	41	39
SSLv2	21	91
Vulnerable RDP	0	69
Vulnerable SMB3	1	39
Total	184	614

Table 3: Reported open and vulnerable service types over the observation period.

Notifications for these vulnerabilities were sent to participants and fixed. Vulnerable services such as open databases, exposed hardware management interfaces such as IPMI, expired x509 certificates, and services that still use weak crypto libraries also fall in this category.

Since the control group was larger than the participants, they also had a wider variety of issues shown in Table 3 that the EWS service could have notified them about, such as exploitable Palo Alto routers and open MySQL databases.

After analysis with the participants, we found that fewer than two percent of the overall notifications sent to them were false positives.

Reports considered to be false positives were related to the same type of malware infections collected from the same source. The colleges tracked both incidents to iPhones using the college Wi-Fi, which were unlikely to be compromised

since the notification source is focused on reporting compromised PC based systems. They were most likely triggered by an unfortunate visit to a website that contained an URL pointed to a known bad domain.

Building Early Warning Service

Collecting and processing the information used for the notifications would require significant personnel and time investment for an enterprise or college seeking to replicate the service. The purpose of the Early Warning Service (EWS) is to remove this obstacle and make the information conveniently available.

In the study, Early Warning Service provided participating colleges information based on Arctic Security's long experience with supporting national cybersecurity authorities to create similar services. Those services have been very successful in their niche environments. Still, governments suffer from inherent limitations that make it difficult for them to scale the service for a wider audience, beyond critical infrastructure, especially beyond their national borders.

Choosing and validating the data sources used for the notifications depends on Arctic Security's extensive experience in this field. Individual data sources that provide high-quality data often are quite expensive, with annual subscriptions ranging from the €10,000 mark to the €100,000 mark. These kinds of data providers don't usually sell access to small enough subsets of the data to make it either affordable or easy to process, instead preferring to sell as-is access to the entire data set.

Arctic Security's EWS handles the data processing on behalf of the customers, which helps with a common problem with the bulk data providers who are offering massive amounts of data that makes it difficult to find the information that is relevant for the recipient.

Filtering the small slice of relevant data from the mass of the whole is cumbersome. Globally, there are tens of millions of daily events to parse and match their respective affected network owners.

EWS works at this scale, breaking data into small subsets to make it easy to consume, automating the process so that recipients can focus more of their time on fixing rather than searching for problems.



By merging data from many commercial data vendors and providing affordable access to individual colleges and enterprises, it also makes the access affordable. EWS provides a 24/7/365 monitoring service, and the notifications are sent on a schedule that is convenient for the subscriber.

“

“This is a 24/7 job, and you feel better if there is somebody watching it. I would love to have a full-time network security person who did nothing but security, but for financial reasons that is not available. This service is invaluable to us, and that’s personal as well as business. It lets me rest much better at night.”

– CIO of a consortium college

”

Future work

This project focused on smaller private universities, but the discussed problem set is not limited to them. Larger universities have to deal with the same issues, and while they have more options to acquire or build services in-house, their budgets are perennially tight. The same lack of cybersecurity professionals affects organizations of any size.

Fig. 4 presents a set of large, well-known universities who had the most affected systems recorded in 2020. Typical observations for them include about 200 unique IP addresses reported per month, with some monthly peaks. The same




Fig. 4: Unique affected IPs for the academic institutions (anonymized) with most affected systems in 2020. Most of these are very large universities.

caveats as with the control group apply here; the accuracy and completeness of the observed events are only approximate as they have been correlated with publicly available information about these institutions' network assets.

Many of these issues have been present and visible in their networks for periods comparable to our control group. Unlike the participating colleges in this study, these organizations likely have sizable IT staff that includes many security experts. Yet, there are still things that fall through the cracks. It would be illuminating to run the same kind of project with a large university or university system to observe the impact of an EWS-type service on their daily routine.

Conclusions

At the end of the study, all participants found the notifications useful and concluded that the service had improved their security posture with only a small amount of administrative overhead. Notifications were accurate and receiving them induced a new daily routine of checking for security-related issues using both the notifications and the internal monitoring capabilities.

There was a significant improvement on the time-to-fix, from 81 to 1.3 days, measured from the initial visibility of the event to when it was fixed.

The industry standard response times measured in more than a hundred days could be observed for the control group.

Many cybersecurity solutions and network monitoring systems suffer from alert fatigue, as there are too many false positives in relation to actual incidents. We were able to show that by relying on high-quality data sources, we can mostly avoid the problem. The number of notifications that were false positives was so low (2%) that some of the participants experienced zero false positives in the span of the project.

Security improvements were accomplished without further affecting the college's networks with any kind of hardware installation. EWS is a completely external solution that proved to be easy to implement for the participants.



“

“You did this with no performance degradation. We're not screening traffic, we don't filter here. With every other device I've ever seen that does this kind of work, there are performance issues. This, however, was done without any intrusion whatsoever to the clients.”

– IT Director of a consortium college

”

The looming risk of severe compromises that can lead to reputation damage and leak of private information is a fact that occupies the minds of most IT organizations. Still, without budgets and personnel, that has not been easy to mitigate. While external monitoring will never catch all the security breaches, it does help to capture and notify you of many of the problems that are already visible to the outside world. Several participants highlighted the benefit of the peace of mind they got from having coverage from an additional security monitoring service.

In conclusion, the study shows how affordable security notifications can significantly improve the IT staff's ability to fix and remediate cybersecurity vulnerabilities and improve their ability to investigate cases. It brings up issues that would otherwise go unnoticed and has a small impact on the IT workload.

Early Warning Service

Arctic Security and our partners offer Early Warning Service to customers globally.

Subscribe, register your internet-facing asset information with Arctic Security, and start to receive prompt notifications of your organization's cybersecurity issues.

These issues are already visible to the internet, you should know about them too.

For more information, reach out to us at contact@arcticsecurity.com