

Domain Shadowing: A Stealthy Use of DNS Compromise for Cybercrime

Executive Summary

Cybercriminals compromise domain names to attack the owners or users of the domains directly, or use them for various nefarious endeavors, including phishing, malware distribution, and command and control (C2) operations. A special case of DNS hijacking is called domain shadowing, where attackers stealthily create malicious subdomains under compromised domain names. Shadowed domains do not affect the normal operation of the compromised domains, making it hard for victims to detect them. The inconspicuousness of these subdomains often allows perpetrators to take advantage of the compromised domain's benign reputation for a long time.

Current threat research-based detection approaches are labor-intensive and slow as they rely on the discovery of malicious campaigns that use shadowed domains before they can look for related domains in various data sets. To address these issues, we designed and implemented an automated pipeline that can detect shadowed domains faster on a large scale for campaigns that are not yet known. Our system processes terabytes of passive DNS logs every day to extract features about candidate shadowed domains. Building on these features, it uses a high-precision machine learning model to identify shadowed domain names. Our model finds hundreds of shadowed domains created daily under dozens of compromised domain names.

Emphasizing the difficulty of discovering shadowed domains, we found that only 200 domains were marked as malicious by vendors on VirusTotal out of 12,197 shadowed domains automatically detected by us between April 25 and June 27, 2022. As an example, we give a detailed account of a phishing campaign leveraging 649 shadowed subdomains under 16 compromised domains such as `bancobpmmavfhxcc.barwonbluff.com[.]au` and `carriernhoousvz.brisbanegateway[.]com`. The perpetrators leveraged the benign reputation of these domains to spread fake login pages harvesting credentials. VT vendor performance is much better for this specific campaign, marking as malicious 151 out of the 649 shadowed domains

– but still less than one quarter of all the domains.

Palo Alto Networks provides protection against shadowed domains leveraging our automated classifier in multiple [Palo Alto Networks Next-Generation Firewall cloud-delivered security services](#), including [DNS Security](#) and [Advanced URL Filtering](#). Additionally, customers can leverage [Cortex XDR](#) to alert on and respond to domain shadowing when used for C2 communications.

Related	
Unit 42	DNS Security, Credential Harvesting
Topics	

How Domain Shadowing Works

Cybercriminals use domain names for various nefarious purposes, including communication with C2 servers, malware distribution, scams and phishing. To help perpetrate these activities, crooks can either purchase domain names (malicious registration) or compromise existing ones (DNS hijacking/compromise). Avenues for criminals to compromise a domain name include stealing the login credential of the domain owner at the registrar or DNS service provider, compromising the registrar or DNS service provider, compromising the DNS server itself, or abusing [dangling domains](#).

Domain shadowing is a subcategory of DNS hijacking, where attackers attempt to stay unnoticed. First, cybercriminals stealthily insert subdomains under the compromised domain name. Second, they keep existing records to allow the normal operation of services such as websites, email servers and any other services using the compromised domain. By ensuring the undisturbed operation of existing services, the criminals make the compromise inconspicuous to the domain owners and the cleanup of malicious entries unlikely. As a result, domain shadowing provides attackers access to virtually unlimited subdomains inheriting the compromised domain's benign reputation.

When attackers change the DNS records of existing domain names, they aim to target the owners or users of these domain names. However, criminals often use shadowed domains as part of their infrastructure to support endeavors such as generic phishing campaigns or botnet operations. In the

case of phishing, crooks can use shadowed domains as the initial domain in a phishing email, as an intermediate node in a malicious redirection (e.g., in a [malicious traffic distribution system](#)), or as a landing page hosting the phishing website. In the case of botnet operations, a shadowed domain can be used, for example, as a proxy domain to conceal C2 communication.

In Table 1, we collect example shadowed domains used as part of a recent phishing campaign automatically discovered by our detector. The attackers compromised several domain names that have existed for many years and thus built up a good reputation. We can observe that the IP addresses of these domains (and IPs of their benign subdomains) are located in either Australia (AU) or the United States (US). Suspiciously, all the shadowed domains have IP addresses located in Russia (RU) – a different country and [autonomous system](#) from the parent domains. Furthermore, all shadowed domains in this campaign use an IP address from the same /24 [IP subnet](#) (the first three numbers are the same in the IP address). An additional indicator of malice we noticed is that all the malicious subdomains shown were activated around the same time and were operational for a relatively short period.

FQDN	IP Addres s	C C	First Seen	Last Seen	Time Active*
halont.edu[.]au	103.152 .248[.]1 48	A U	2020- 11-23	2022- 06-28	~ 9 years
training.halont.edu[.]au	103.152 .248[.]1 48	A U	2020- 12-08	2021- 05-02	~ 7 years

training.halont.edu[.]au*	62.204.41[.]218	R U	2022-04-17	2022-05-06	< 1 month
ocwdvmjjj78krus.halont.edu[.]au	62.204.41[.]218	R U	2022-04-04	2022-04-04	< 1 day
baqrxmgfr39mfpp.halont.edu[.]au	62.204.41[.]218	R U	2022-04-01	2022-04-01	< 1 day
barwonbluff.com[.]au	27.131.74[.]5	A U	2018-12-13	2022-06-28	~ 19 years
bancobpmmavfhxcc.barwonbluff.com[.]au	62.204.41[.]247	R U	2022-03-07	2022-06-06	~ 3 months
tomsvprfudhd.barwonbluff.com[.]au	62.204.41[.]77	R U	2022-03-07	2022-03-07	< 1 day
brisbanegateway[.]com	101.0.12[.]230	A U	2015-04-23	2022-06-24	~ 12 years
carriernhoousvz.brisbanegateway[.]com	62.204.41[.]218	R U	2022-03-07	2022-03-08	~ 2 days

vembanadhouse[.]com	162.215 253[.]1 10	U S	2019- 09-04	2022- 06-28	~ 17 years
wiguhllnz43wxvq vembanadhouse[.]com	62.204. 41[.]21 8	R U	2022- 03-25	2022- 03-25	< 1 day

Table 1. Example of compromised domains and their shadowed subdomains. *Time active column is based on the time first seen in pDNS, Whois, or archive.org. **It seems that the subdomain *training.halont.edu[.]jau* was deactivated, and later the attacker accidentally hijacked it via [DNS wildcarding](#). FQDN stands for Fully Qualified Domain Name and CC stands for the country-code of the IP address.

How to Detect Domain Shadowing

To address issues with threat hunting-based approaches to detect shadowed domains – such as lack of coverage, delay in detection and the need for human labor – we designed a detection pipeline leveraging passive DNS traffic logs (pDNS) based on work by [Liu et al.](#) Building on observations similar to the ones discussed in Table 1, we extracted over 300 features that could signal potential shadowed domains. Using these features, we trained a machine learning classifier that is the core of our detection pipeline.

Design Approach for the Machine Learning Classifier

We can arrange the features into three groups – those specific to the candidate shadowed domain itself, those related to the candidate shadowed domain’s root domain and those related to the IP addresses of the candidate shadowed domain.

The first group is specific to the candidate shadowed domain itself. Examples of these FQDN-level features include:

- Deviation of the IP address from the root domain's IP (and its country/autonomous system).
- Difference in the first seen date compared to the root domain's first seen date.
- Whether the subdomain is popular.

The second feature group describes the candidate shadowed domain's root domain. Examples are:

- The ratio of popular to all subdomains of the root.
- The average IP deviation of subdomains.
- The average number of days subdomains are active.

The third group of features is about the IP addresses of the candidate shadowed domain, for example:

- The apex domain to FQDN ratio on the IP.
- The average IP country deviation of subdomains using that IP.

As we generate over 300 features – where many of them are highly correlated – we perform feature selection in order to use only the features that will contribute most to the machine learning classifier's performance. We use the [Chi-squared test](#) to find the best features individually and mutual [Pearson correlation](#) to decrease the weight of highly correlated features.

We can select classifiers with different performance and complexity tradeoffs depending on the desired use case. Using a random forest classifier, we can achieve 99.99% accuracy, 99.92% precision and 99.87% recall using only the 64 best features and allowing each of 200 trees in the random forest to use at most eight features and to have a maximum depth of four. A simpler classifier – using only the top 32 features where each tree can only use at most four features and have a depth of two – can achieve 99.78% accuracy, 99.87% precision and 92.58% recall.

During a two-month period, our classifier found 12,197 shadowed domains averaging a couple hundred detections every day. Looking at these domains in VirusTotal, we find that only 200 were marked as malicious by at least one vendor. We conclude from these results that domain shadowing is an active threat to the enterprise, and it is hard to detect without leveraging automated machine learning algorithms that can analyze large amounts of DNS logs.

A Phishing Campaign Using Shadowed Domains

Next, we dive deeper into the phishing campaign we used as an example in Table 1. Clustering – based on IP address and root domains – the results from our detector, we found 649 shadowed domains created under 16 compromised domain names for this campaign. Figure 1 is a screenshot of `barwonbluff.com[.]au`, one of the compromised domains. Even though it seems to operate normally, attackers have created many subdomains under it that they can use in phishing links such as `hxxps[:]/snaitchbumxzzwt.barwonbluff[.]com.au/bumxzzwt/xxx.yyy@target.it`.

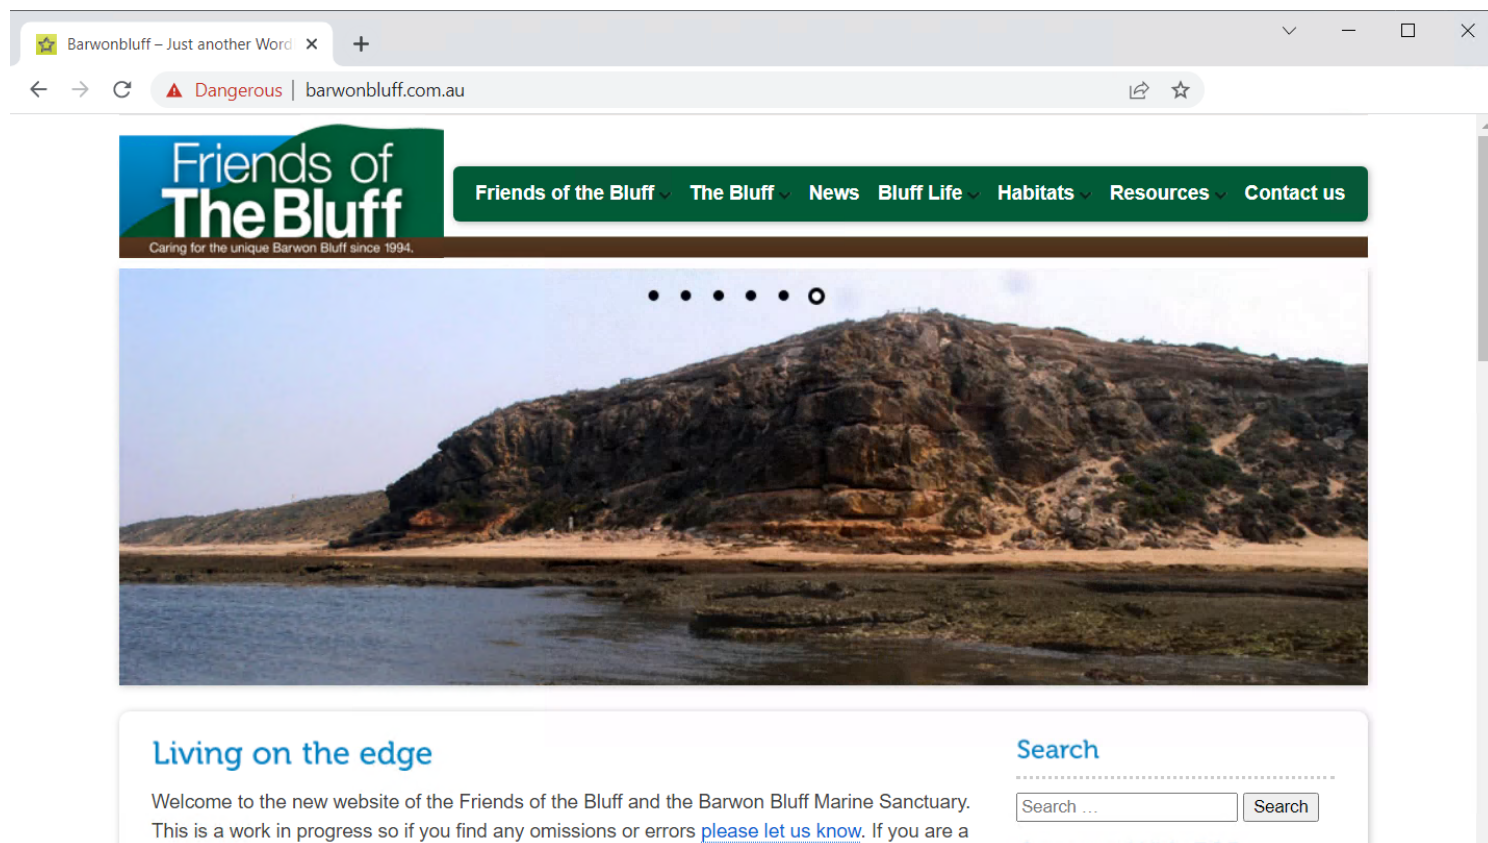


Figure 1. Screenshot of `barwonbluff.com[.]au` – an originally benign domain.

When users click on the above phishing URL, they are redirected to a landing page, as shown in Figure 2. The phishing page on `login.elitepackagingblog[.]com` wants to steal Microsoft user credentials. To avoid falling for similar phishing attacks, users need to check the domain name of the website they are visiting and the lock icon next to the URL bar before entering their credentials.

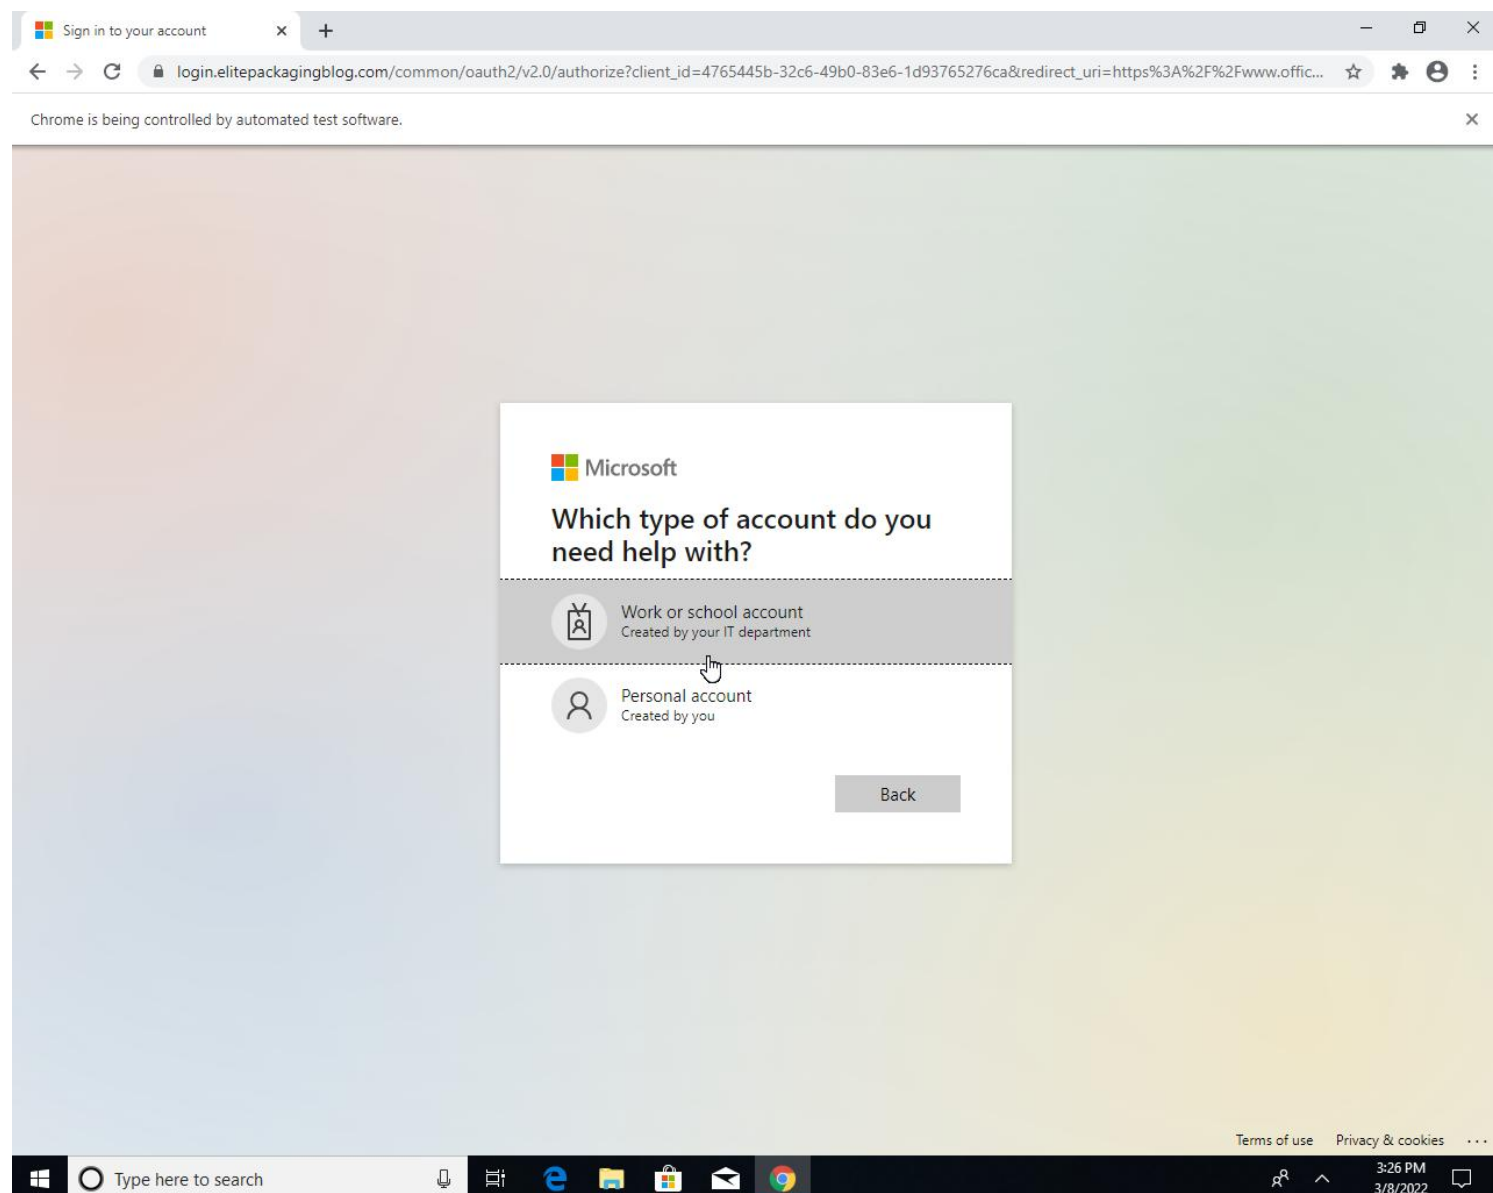


Figure 2. Screenshot of the phishing landing page on elitepackagingblog[.]com, where victims are redirected from the snaitechbumxzzwt.barwonbluff[.]com.au shadowed domain. Source: Joe Sandbox.

Figure 3 is a screenshot of halont.edu[.]au after the website owners found out that their domain name was compromised. Unfortunately, we observed many shadowed domains created under this domain name before the owners realized it was hacked. These cases further emphasize the necessity to automatically detect these domains because it is hard for domain owners to discover that they are compromised.

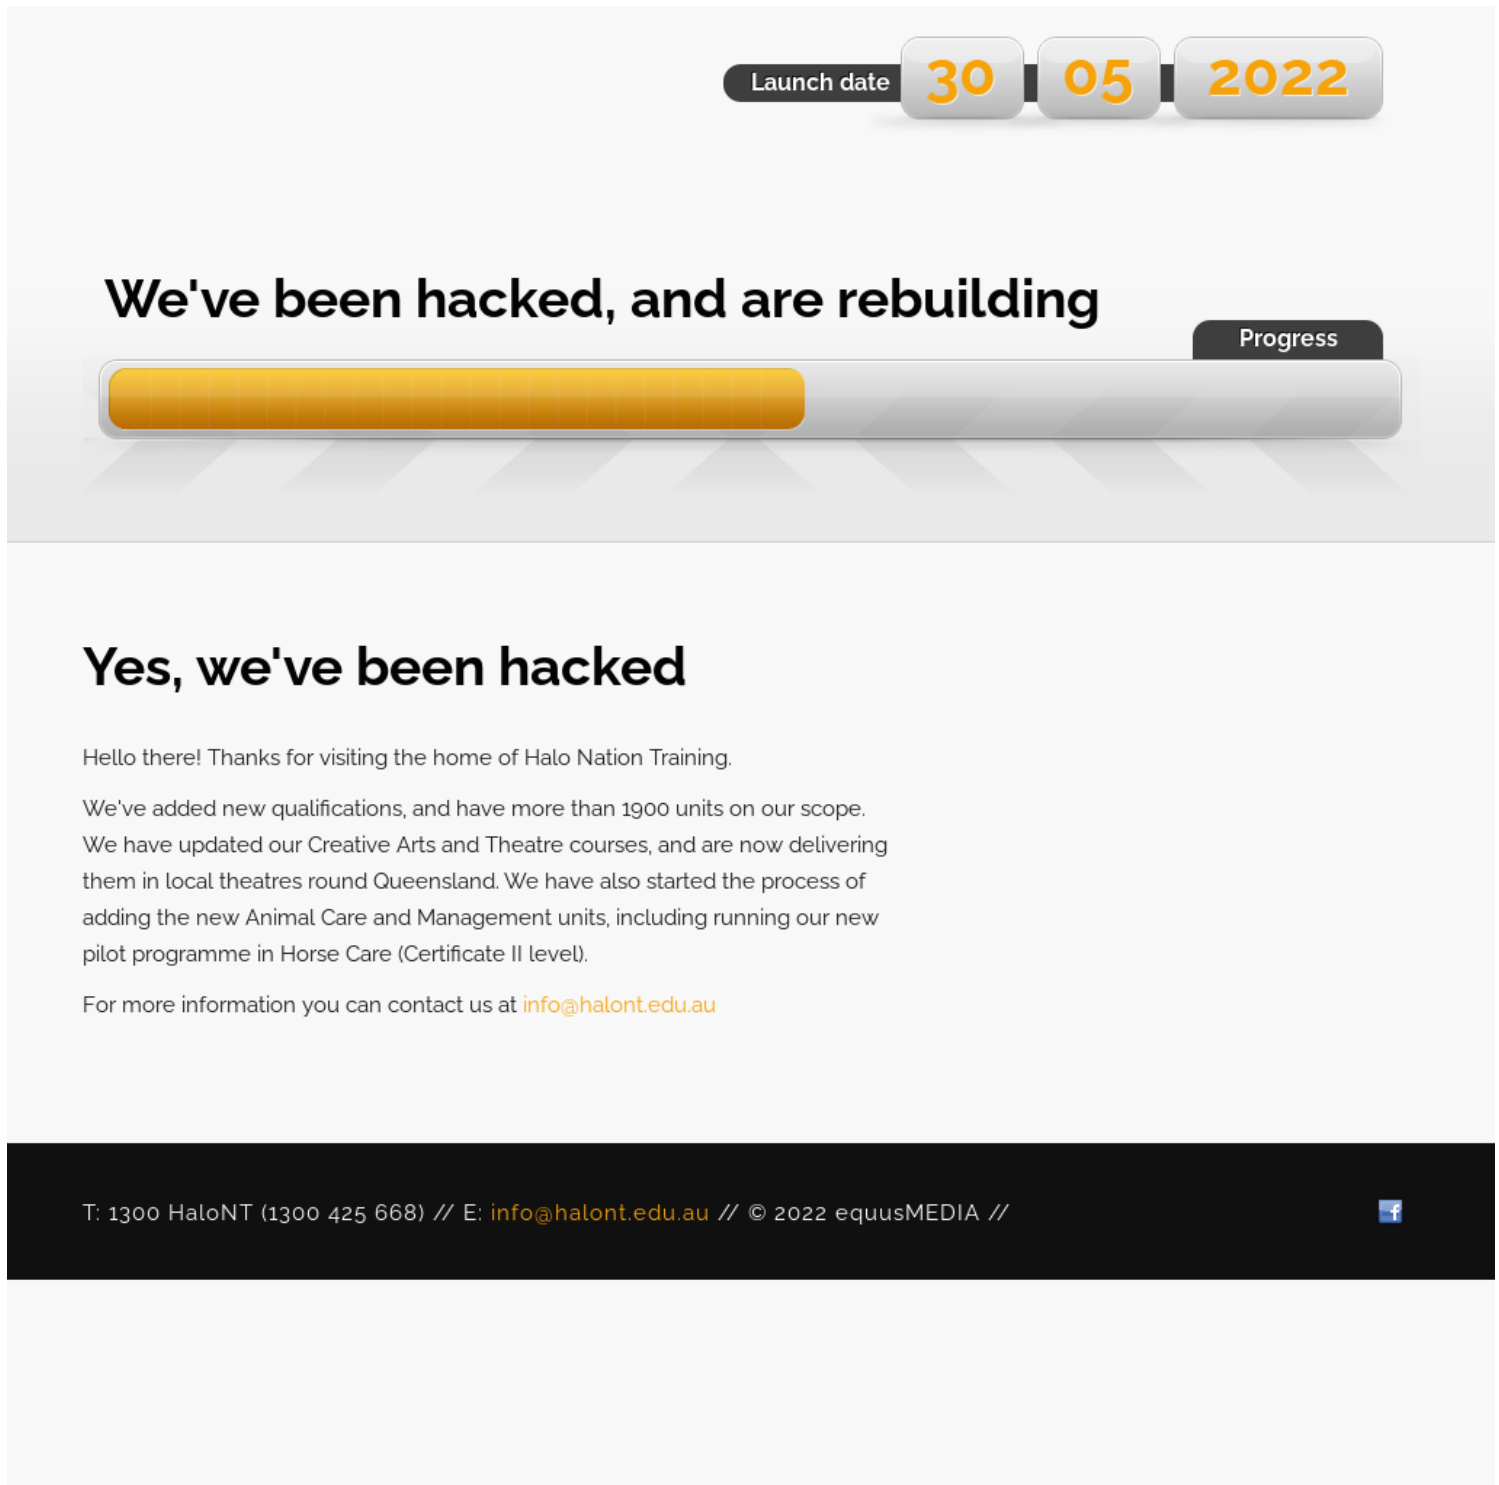


Figure 3. Screenshot of halont.edu[.]au, an originally benign domain that is being rebuilt after compromise.

Conclusion

Cybercriminals use shadowed domains for various illicit ventures, including phishing and botnet operations. We observe that it is challenging to detect shadowed domains as vendors on VirusTotal cover less than 2% of these domains. As traditional approaches based on threat research are too slow and fail to uncover the majority of shadowed domains, we turn to an automated detection system based on pDNS data. Our high-precision machine learning-based detector processes terabytes of DNS logs and discovers hundreds of shadowed domains daily. Palo Alto Networks offers multiple security subscriptions – including [DNS Security](#) and [Advanced URL Filtering](#) – that leverage our detector to protect against shadowed domains. Additionally, customers can leverage [Cortex XDR](#) to alert on and respond to domain shadowing when used for command and control communications.

Acknowledgements

We want to thank Wei Wang and Erica Naone for their invaluable input on this blog post.

Indicators of Compromise

halont.edu[.]au
training.halont.edu[.]au
ocwdvmjjj78krus.halont.edu[.]au
baqrxmgfr39mfpp.halont.edu[.]au
barwonbluff.com[.]au
bancobpmmavfhxcc.barwonbluff.com[.]au
snaitechbumxzzwt.barwonbluff[.]com.au
snaitechbumxzzwt.barwonbluff.com[.]au/bumxzzwt/xxx.yyy@target.it
tomsvprfudhd.barwonbluff.com[.]au
brisbanegateway[.]com
carriernhoousvz.brisbanegateway[.]com
vembanadhouse[.]com
wiguhllnz43wxvq.vembanadhouse[.]com
login.elitepackagingblog[.]com
login.elitepackagingblog[.]com/common/oauth2/v2.0/authorize?client_id=4765445b-32c6-49b0-83e6-

1d93765276ca&redirect_uri=https%3A%2F%2Fwww.office.com%2Flandingv2&response_type=code%20id_token&scope=openid%20profile%20https%3A%2F%2Fwww.office.com%2Fv2%2FOfficeHome.All&response_mode=form_post&nonce=637823463352371687.MDY0MjMzYjMtOWNlZC00ODA5LWE1YWQtOWMyMTIwYTZiOTIwODZiNTMyN2MtZWQ3ZC00Mzg4LWJjMzktNGQxYjQlMDFkNmNi&ui_locales=en-US&mkt=en-US&state=q81i2V5Z572r5P2TuEfGYg0HZLgy9vMW3HMxjfeMMm60rJIIPgKe4SKR8D86gIjkNlgD6cd8jK754mEWDiHZtRQ1pzeGpqaVJOckSmAUGOWUcOxbKCr2sPnoBds6H7fZCJdLqcotpA2NF3vvVbRDSSWk3xhQuxnXOoJoN2pj0RhiR97YEUKUwqEEsCoboffTLGgVrjaDy_ASgmhE_7mkvYE6YsXicgxoeZDqhrjxB_vFcTt_u7olrrAYcWIv-0vZ4vPVToJ7Nwqlf6BHPz7zPQ&x-client-SKU=ID_NETSTANDARD2_0&x-client-ver=6.12.1.0&sso_reload=true#ODQuMTccGFvbGEucGVsbGVnYXRhQHNUYWl0ZWNoLml0=

Additional Resources

[Don't Let One Rotten Apple Spoil the Whole Barrel: Towards Automated Detection of Shadowed Domains](#)