

Quantifying the financial savings Protective DNS (PDNS) brings to the UK public sector

Finding a data-based model of losses prevented through the deployment of PDNS

Key findings

- » Over 60% of organisations caught at least two unique threats each week, with 2 out of every 10 organisations seeing anywhere from 4 to over 20 different threat varieties every week.
- » On average, a typical strain of malware is seen in attempted network connections with client systems for 59 days, with individual C&C families present for just two weeks.
- » The types of incidents prevented by PDNS include: **general incidents** (mainly consuming internal staff time), higher-impact **ransomware** events (less frequent but larger impacts), and rare but potentially devastating **major incidents** receiving public attention.
- » PDNS prevents an average of **4005 general incidents** and **1400 ransomware incidents** each year based on the sample.
- » Across the three incident types, based on the sample data provided, PDNS typically provides **yearly savings of at least £59M**.
- » Threats and attacks vary, but almost always (9 out of 10 years) PDNS **prevents losses of £48M – £223M**.
- » For rarer but potentially catastrophic major events, PDNS safeguards the UK public sector against a **1-in-20 year loss of over £223M**.

Setting the stage

The Domain Name System (DNS) is the directory of the internet, translating human readable domain names, such as www.gov.uk, into the numeric IP addresses, such as 151.101.0.144, that computers require for network communication. As the central lookup service of the internet, DNS is a key opportunity for visibility where network devices attempt to visit. Nominet, on behalf of the National Cyber Security Centre (NCSC), offers **Protective DNS** (PDNS) to organisations in the UK public sector to prevent access to known malicious locations. Regardless of the platform – Linux, MacOS, and Windows systems, and even Internet of Things devices – PDNS offers visibility and protection that is centrally provisioned and cost-free at point of use.

PDNS works by analysing DNS requests, inspecting them against a library of constantly updated ‘known bad’ locations, and intervening against malicious or harmful sites. These actions to block or redirect contact with malicious resources allow PDNS to safely, quickly, and transparently stop internet traffic from reaching malicious locations; it prevents new infections and keeps infected systems (which rely on being able to ‘phone home’ to get commands/exfiltrate data) from propagating or worsening an attack. PDNS is a highly practical implementation of

active cyber defence that does not require active action or oversight on the part of the organisation to deliver value. When a domain is classified as ‘known malicious’, a set of rules prevents DNS resolution and entities (such as servers, workstations, Linux, MacOS, Windows, etc.) from connecting, preventing many common threats such as malware/ransomware, viruses, spyware, etc., from causing harm. Beyond blocking, organisations may investigate and review their usage data, providing more opportunities for improving protection.

Nominet commissioned the [Cyentia Institute](#) to perform an independent review on the value of PDNS. Nominet provided an anonymised, representative, 12-month sample of activity observed by covered entities. While this data is reliable for the nature and relative volumes of both allowed and filtered traffic seen across all protected parts of the UK government, readers should bear in mind that this is not the entirety of the UK public sector nor the entire PDNS customer base – see [Limitations](#).

Customer distribution across geography

There is a wide breadth of public sector organisations protected by PDNS in our sample¹. Our analysis covers 185 entities in England, 47 in Scotland, 24 in Wales, and 11 in Northern Ireland. PDNS customers vary in size from very large central government agencies to smaller local councils. Additionally, not all of these UK government agencies are the same – some agencies collect, manage, and maintain very sensitive (sometimes classified) information and data, which gives them a greater threat profile and requires higher-priority protection, while other agencies deal with lower-level sensitivity and have a smaller attack surface. Nonetheless, there are commonalities between organisations which support some important conclusions about the overall UK public sector.

¹This sample of PDNS data does not include the growth of the service since September 2020 and is a sanitised subset of data. After this time the number of organisations has significantly increased. In 2021 it protected 900+ organisations in addition to the 1,000+ organisations that sit under the Health & Social Care Network (HSCN).

Query volumes: the good, the bad, and the prevalent

The overall volume of DNS traffic observed between September 2019 and August 2020 is relatively stable (see Figure 1). We can see a slow upward trend through the months (as would be expected with normal adoption and use over time) and can see when a significant amount of the sample started working from home in March 2020. For a picture of how much of this traffic is blocked, see Figure 2.




Figure 1: Total Traffic Volume Across PDNS Customers

to the total traffic volumes seen in Figure 2. Blocked queries are measured in millions per month², as opposed to the billions of total queries per month. Optimistically, we can conclude that most of the internet traffic from the UK public sector organisations to legitimate, uninfected resources is good and free of malice.

But simply counting up the number of blocks isn't a good measure of the amount of harm being reduced. Even a low number of bad traffic can cause significant harm to organisations.




Figure 2: Blocked Query Volume per Month

²There is one customer that experienced an anomalous traffic burst in August 2020 that, in consultation with Nominet, we have filtered from this view.

³Calculated as a geometric mean of the monthly totals.

⁴This does not suggest that every employee comes in contact with bad traffic every month. Some employees will undoubtedly not come into contact while others may come into contact multiple times a month. These averages serve only to illustrate the scale of the problem.

Commonalities among organisations




Figure 3: Percent of Organisations with Blocked Queries

Figure 3 shows the percentages of organisations that had interventions by PDNS, i.e. malicious IPs or domains to block. At first glance this may seem uninteresting because it is so flat, however, this chart demonstrates a remarkable consistency in the number of organisations with PDNS-blockable traffic.

Over 96% of organisations made suspicious DNS requests – either an infected internal system making an outbound request or a healthy system requesting access to a known malicious location – that needed blocking.

Those remaining 4 out of 100 organisations consist principally of lab environments and other networks not fully configured. At 96% of organisations experiencing interventions by PDNS, there’s opportunities for nearly every organisation, large or small, to benefit from PDNS.

It’s one thing to say that organisations will come into contact with bad things on the internet, but we want to help answer the question: what is the risk to the organisation? We explore this using a graphical aid that appears a couple of times throughout this report – Figure 4.

This is a loss exceedance curve, a common chart from the field of insurance, where communicating both the uncertainty and variability of risky events is an everyday problem. Each point on the line shows the chance (percentages marked on the y-axis) of an organisation having at least the number of blocks indicated on the x-axis. Our goal with this chart is to help our readers answer the question: how severe is something likely to be?

Putting this into practice, we see that our first highlighted point indicates a 95% chance that an organisation will have at least three events

(unique instances of blocked queries). Nine out of ten organisations deal with as few as three and as many as 14,000 blocked queries every month. That is a massive spread and reflects the wide variety of organisations covered by PDNS. Half of the organisations experience 137 events or more, implying that while everyone deals with blocked queries, the number ramps up for some customers very quickly. With the anonymised data we have, we can’t tell if the number of blocks relates to differing threat profiles. Regardless, all organisations face hostile traffic and need defending.




Figure 4: Likelihood of Blocked Queries Seen by an Organisation in a Typical Month

What does normal look like?

Figure 5 shows the average block proportion of PDNS for the UK public sector organisations within our sample. PDNS blocks, on average, 1 out of every 62 thousand requests⁵. However, there is a remarkable variation in that statistic, with some organisations blocking 1 out of every 100 requests and others blocking just one in every 10 million requests.

Referring to the numbers previously mentioned in this report, 13 billion constitutes a lot of traffic to be analysed, creating a massive haystack in which to search for needles, i.e., the malicious IPs and domains. Fortunately, in addition to benefits such as increased visibility and analytics, PDNS consistently and continuously protects against known threat varieties that persist on the internet and are most likely to be seen in organisations' traffic.




Figure 5: Proportion of Blocked Queries

⁵A 0.04% overall block ratio was discussed in the [Active Cyber Defence: The Fourth Year](#) report. In this analysis, we are calculating the ratio of blocked traffic on a per organisation basis and looking at the median rate seen across organisations. When looking at all traffic in aggregate, we continue to see the same magnitude of numbers as reported in the Active Cyber Defence report.

What threats are being identified?

We now turn to examining what we have identified in those organisations we are protecting. There is a vast assortment of strains of different threat varieties. As these strains evolve and mutate, they require that defenders develop different signatures and techniques.

To help understand the scope of the problem facing defenders, the loss exceedance curve can be useful.

We return to this form below in Figure 6.

VOLUME OF UNIQUE THREATS SEEN AT ORGANISATIONS

In using the loss exceedance curve to understand the number of unique threats – distinct variants of malware and other forms of malicious code – we can better understand the landscape. Over 60% of organisations caught at least two unique threats each week, with 2 out of every 10 organisations seeing anywhere from 4 to over 20 different threat varieties every week. That presents a challenge for triage, if there is no immediate form of proactive defence in place.




Figure 6: Unique Threats Seen at each Organisation

Looking across the UK public sector




Figure 7: Unique Threats Seen Across PDNS Customer Base

In Figure 7 we zoom out from a single organisation to look at the entirety of our sample of PDNS end customers in the UK public sector, facing a weekly average⁶ of 243 unique threats detectable via DNS. This gives us a median of 235 unique threats per week, with the most common range of 205 – 311, accounting for week-to-week fluctuations.

This is a representation of the variety of threats that are coming into contact with the UK public sector. Without proactive prevention at an external layer (i.e., the internet, external to the organisation), all of these varieties must be handled once that traffic hits the internal network.

⁶A geometric mean.

These attack volumes are a crucial problem for enterprises worldwide, straining already under-staffed cyber security teams. To handle the constant barrage, organisations have turned to a variety of technologies. PDNS is a first line of defence against these threats, saving significant time and effort on triage and remediation, and freeing up analyst/operator time for higher-level analysis. In other words, PDNS can eliminate what security teams already know will be problematic and gives them the freedom to focus on those threat varieties most demanding of human intelligence.

Proportion of different threat types seen by an organisation

What kinds of threats does a typical organisation encounter in its internet traffic over a given week?

Assigning categories can be a difficult area to draw consensus as some varieties of threats can exhibit characteristics from multiple categories. For simplicity we will use the five categories assigned by PDNS as it tracks threats. Each of these categories can be ranked in terms of the likelihood of harm they present to users. These categories, in the order of most likely to result in loss to least likely are:



C&C

Command and Control (C2). This category is populated by intelligence sources that identify domains being used as C2 points for the orchestration or distribution of malware. Such domains may issue instructions that affect malware behaviour, may host a collection point for exfiltrated data, or service other similar use cases for an attacker to take advantage. Disrupting C2 traffic will degrade a botnet's function, rendering the endpoints inaccessible to C2 servers and effectively neutralising the infection in some cases.



SPYWARE

This category is populated by intelligence sources that identify domains that are used in spyware campaigns. This includes malware on endpoints that collects information such as usernames and passwords.



MALWARE

This category is populated by intelligence sources that can identify domains being used as part of a malware campaign. A domain tagged with malware is known to host malicious file downloads, malicious content, or act as a rendezvous point for malware command and control.



COMPROMISED

A legitimate good domain name that has been compromised by a malicious actor. Such a compromised domain may be hosting malicious files, may act as a C2, or be used in attacks on targets. An attacker may have gained access to a website using stolen credentials or a vulnerability in the website infrastructure, for example, a content management system (CMS). It may not be immediately obvious to casual users if a website is compromised as the original content may be left intact.



UNKNOWN

This category is populated by intelligence sources that identify unclassifiable malicious domain names. Each domain is confirmed as malicious but is of an indeterminate variety. Unknown threats are often tagged as such to protect sensitive sources or avoid giving misleading information where clear attribution is not yet available.

Threat type distribution

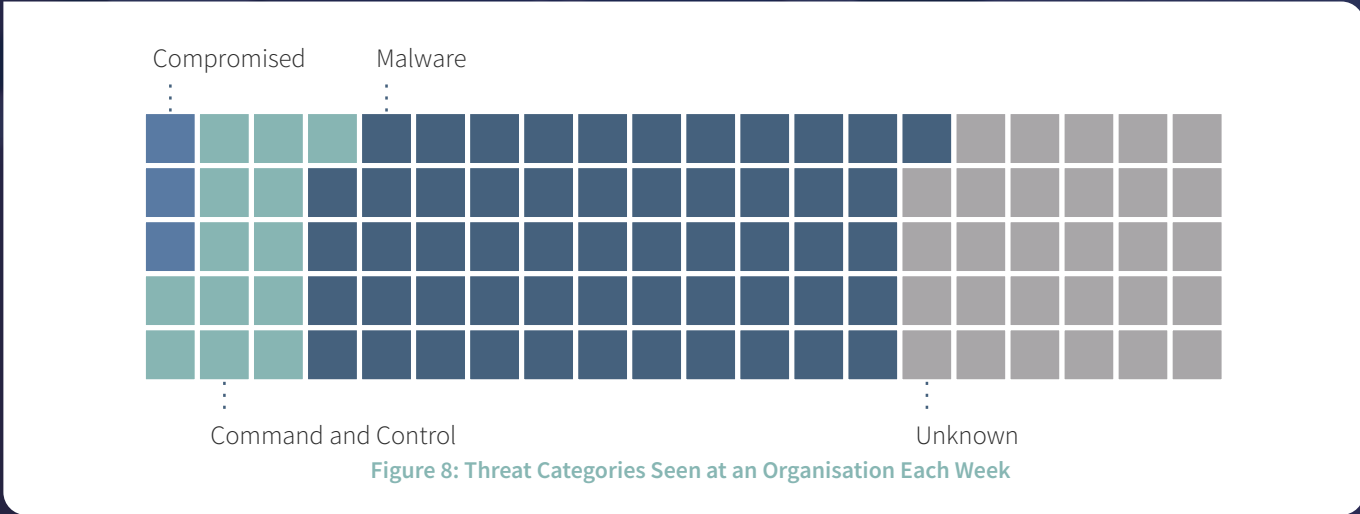
Figure 8 provides a graphical illustration of the average⁷ distribution of threat types observed at a typical organisation in any given week. Each square represents 1% of the threats faced in a particular week. If an organisation were to see 100 different threats, all things being equal, this is what that traffic would have looked like. There is one noticeable absence in this diagram: spyware. This category occurs relatively rarely, together representing less than a tenth of a single square.

The command-and-control category, denoted by the dark blue squares, needs special attention. Appearing in 13 out of every 100 events, this category is most closely related to instances of successful compromise within organisations. Command and control activity is difficult to identify and more difficult to eradicate; once an attacker compromises a machine (user device, server, host, etc.), the infected machine communicates with the attacker's server and carries out malicious commands. When C&C is found, it often means

that a threat actor has achieved persistence, thus making it a concerning category.

While the C&C category represents a substantial danger, organisations should not ignore the other concerns. The malware category includes items that could, if successfully deployed at a client, result in eventual C&C activity. These categories can be considered different phases of an attack's lifecycle⁸, from first contact and infection, to the communication back to the threat actor. During all of these phases, controls such as PDNS can, via blocking of communications, intervene to block progress in the lifecycle of compromise and exploitation.

Threat varieties rapidly change, evolving and mutating as they attempt to achieve their ends. In Figure 9 we see the results of analysis which looks across these threat categories and measures the average duration from a threat variant's first appearance to the last time that it is seen.



⁷A geometric mean.

⁸Our use of the lifecycle is informal in this report. For a rigorous evaluation of the lifecycle of attacks, reference the MITRE ATT&CK framework.

Duration of named threats

On average, a typical strain of malware is seen in attempted network connections with client systems for 59 days, with individual C&C families present for just two weeks. This sizable difference reflects that, against a background of relatively long-lived threats, the families of threats which present the potential for greatest harm to organisations (C&C activity) change frequently. Organisations seeking protection against the malware causing the most harm need to constantly update their defences.

If you consider the high volume of traffic (Figure 1) versus the duration of the threat (Figure 9), it is clear that identification and prevention leave an organisation significantly less vulnerable for less time than detection, response and recovery. Tying this data together with the number of unique threats (a weekly average of 244 unique threats due to malicious DNS) leaves a lot of unique threats for defenders to deal with, and some of them – like the more pernicious command and control – don't stay around very long. By the time a defender tracks down these unique threats, versus warding them off at the pass with PDNS, the attacker could already have the data.




Figure 9: Typical Duration of Threat Categories as noted among PDNS Customers

What this might mean if undefended

Now that we understand the scope of what is being screened via PDNS, let us envision a world where we do not have this control. What are the forms and sizes of harm that we are avoiding by having PDNS? For a data-grounded way of answering this question, we'll combine observed real world PDNS data with independently gathered and verifiable breach information. By applying data and modelling, we can get an understanding of the benefit PDNS has on its customers.

THREE CATEGORIES OF INCIDENTS

Just as threats have a wide diversity, not all cyber loss incidents are equal. The losses experienced by firms as a result of cyber security incidents range from large and well-publicised breaches that appear in the headlines to the less talked about, but far more commonplace events that organisations deal with during regular operations. We will divide these breaches into three broad categories: general incidents, ransomware incidents, and major breach incidents.

USING PDNS DATA TO IDENTIFY GENERAL INCIDENTS

We begin with the general incidents. As referenced earlier in this paper, threats in the C&C category typically represent malicious code actively trying to reach out after a successful infection to receive commands and exfiltrate data. Blocks in this category are a strong signal of threats that have evaded other controls and managed to establish a presence in a client's network.

Within this population of events, we have to deal with the relationship of blocked threats possibly representing multiple systems. DNS uses caching at multiple levels of the network stack, meaning a single blocked query may represent one, dozens, or possibly even hundreds of systems in a given customer's network. We don't have a reliable way to estimate this effect, so keep in mind that these numbers therefore reflect a conservative lower threshold. In reality, the numbers could be even higher.

Now that we can infer how many of these most severe threats each organisation is likely to have seen over the course of the sample period, we can perform statistical sampling⁹ to create simulations of a hypothetical year based on the sample of users. By running through multiple simulations (see the Methodology section for the detailed process) we create a distribution of good years (when organisations avoid many incidents through other means), bad years (when threat actors overcome more client defences), and typical years. Together this gives us a distribution, illustrated in Figure 10, showing how many loss events were prevented. We can estimate 4,005 events (on average) in this category being prevented from disrupting clients' operations every year.

These general incidents are just part of the picture. Following is the second of the three incident types in our model: ransomware incidents.




Figure 10: General Incidents Prevented by PDNS

⁹Bootstrap resampling.

FOCUSING IN ON RANSOMWARE INCIDENTS

Discussions on the threat landscape have recently been dominated with concerns about the effects of ransomware. These high profile incidents present a concern to many organisations, with data being destroyed, locked away, or exfiltrated by threat actors demanding payment.

The PDNS data gives an indication of many threat varieties that are related to known ransomware strains; we'll use them, in combination with the general incident data from the preceding section, to estimate the number of these high impact events.

Ransomware incidents are quite common in our observed traffic. When we see an average of over 1,400 incidents being stopped every year, as displayed in Figure 11, we are deriving this based upon observed traffic patterns in the UK public sector.




Figure 11: Ransomware Incidents Prevented by PDNS

ACCOUNTING FOR MAJOR INCIDENTS

In extreme cases, incidents can expand from their initial compromise in both scope and impact, reaching a level where they enter common discourse and public discovery. These are the events that surface via news articles, court cases, press releases, etc. In its recent work from the **IRIS series**, Cyentia looked at what we can learn from this population of publicly verifiable cyber losses. In the IRIS series, we observed that an organisation in the public sector has a 5.4% chance of experiencing one or more of these major events in a given year¹⁰. We also know from prior work leveraging the **Verizon DBIR** that approximately a third of key incident patterns have DNS in their kill chain. Armed with the rate of occurrence and number of events that are relevant to DNS controls, we now have the parameters to build our third model.

This model (see Figure 12) represents how many of the major incidents could have been avoided in the UK across PDNS's user base.

¹⁰Figure 6 of IRIS 20/20 https://www.cyentia.com/wp-content/uploads/IRIS2020_cyentia.pdf

Examples of incidents that occurred in organisations eligible for PDNS include **Hackney Borough Council**, estimated costs to be around £10 million and the **Redcar and Cleveland Borough Council**, financial impact assessment put costs at £8.7 million.

The **Irish Health Service Executive** projected that the recovery costs from a ransomware attack in May 2021 would be \$600 million (£442 million).

Famously, the WannaCry attack that affected the NHS in May of 2017 affected a large number of NHS trusts, disrupting operations and resulting in costs that have exceeded, as of a **Department of Health** report, over £92 million.

While these are extreme cases, similar attacks continue to be attempted through the present day. PDNS should be seen as a component in a layered approach to security. The Wannacrys of tomorrow are represented in the general and ransomware events in our model, with the potential of growing into the major incident class discussed here.

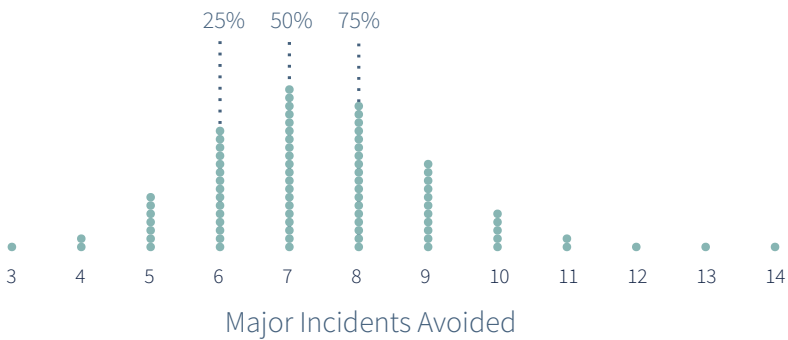


Figure 12: Major Incidents Prevented by PDNS

Primary savings to the UK public sector

With the frequency of the three different types of incidents (general incidents, ransomware incidents, and major incidents) which we are modelling as prevented by PDNS, we can now estimate how much savings this represents for the UK economy. Just as each class of event requires a different frequency of occurrence models, they also require different models of costs. Note that our modelling is limited to verifiable public information on the costs for incidents. Secondary costs such as lost productivity and non-monetary costs are not included.

Non-monetary costs are encapsulated in other reports such as the NCSC Annual Review 2021, where it points to food supply shortages, local fuel price increases, citizens being denied access to public services, and instances where at-risk children's details have been lost. Specifically with regards to the attacks on the local councils mentioned previously, non-monetary costs for Redcar and Cleveland Borough Council included staff needing to use pens and paper, as appointment bookings, planning documents, social care advice, and council housing complaints systems went offline. In the case of Hackney Borough Council, land searches and

planning applications were disrupted and the online portal for paying rent, service charges, and checking balances was temporarily unavailable. Normal systems used to process reports around noise nuisance, anti-social behaviour and missed waste collections were impacted, causing responses and investigations to be slower.

In the Department of Health report with regards to the WannaCry attack, it refers to 19,000 appointments that were cancelled as a result of service disruption in one-third of hospital trusts and around 8% of GP practices. The Scottish Environment Protection Agency (SEPA) was hit by a devastating ransomware attack on Christmas Eve. With no access to emails, files and with data stolen, the agency's core purpose – to protect the environment – was curtailed. 1,200 staff were prevented from doing their work, historical data has been permanently lost, and 1.2GB of data, more than 4,000 files, was exposed by the attackers because SEPA refused to pay the ransom.

Instead of incorporating these difficult to measure non-monetary costs, we have chosen to be conservative in our estimation and remind the reader that there are often significant additional secondary and non-monetary costs with incidents that are hard to quantify or generalise and are consciously excluded from the scope of our analysis.

For general incidents—those that usually stay at a lower level of impact (though with the possibility of expanding into larger events)—the major driver of cost is time to respond and clean up from infections that do not result in data compromise. We have calibrated estimates¹¹ of the time involved in events of this nature, ranging from half an hour in the simplest case, up to two hours for a larger response. While not all governmental organisations have similar staffing models, this is a conservative estimate to avoid over attributing costs without justification. These events can also develop into larger events that have bigger costs, which is captured in our ransomware and major incident categories.

¹¹Derived based upon practitioner expertise and in consultation with Nominet.

The cost of ransomware

For ransomware events, hard data on the amounts of currency involved in ransomware demands (and the amounts actually paid) is difficult to come by. After surveying multiple vendor reports, academic publications, and many news sources¹², we have a distribution¹³ of £15K at the low end, £36.5K as the most likely, and £55K at the upper end¹⁴. These costs include both paid ransoms (when such payments take place) and response costs. There are certainly loss events that are larger– and it is for these larger events (both ransomware events and other loss occurrences), where we delve into the major incident distribution.

¹²Including the US FBI Internet Crime Report and a convenience sample or reports indexed at the Cyentia Cybersecurity Reference Library.

¹³Most of the ransomware statistics we found are denominated in USD. The log transformed beta PERT distribution in USD is \$20K at the low end, \$50K for most likely, and \$75K for the upper end.

¹⁴See [Limitations](#)

For these publicly discoverable events, we have the advantage of numbers we can reference from the IRIS 20/20 loss distribution for the public sector. For details on how these estimates are derived, take a look at the IRIS at cyentia.com/iris.

Bringing this all together, we come to our conclusion in Figure 13, below. This reprise of our loss exceedance curve chart combines the three types of prevented incidents into the aggregate loss curve for all the PDNS customers. Like the other curves we've seen in this report, there is a long tail of extreme values, reflecting those years where costs could have been much larger without the protection of PDNS.

The range of losses which is more representative of the typical year demonstrates that in nine out of ten years PDNS safeguards between £48 million and £223 million. A typical (50% chance) year sees savings of at least £59 million. By having the PDNS service, there are also those large 1-in-20 year savings of more than £223 million.




Figure 13: Aggregate Loss Exceedance Curve of PDNS Savings

Revisiting what we've done

Our objective to find a data-based model of losses prevented through the deployment of PDNS has involved a number of steps that are summarised below.

RATE OF EVENTS:

As the strongest indicator of probable 'bad things' happening within an organisation, we used the number of unique C&C threats blocked in our sample of PDNS data to establish the rate of events most likely to result in some form of loss, had they not been blocked. We used this rate to run simulations of the number of incidents across the PDNS customer sample in a hypothetical year.

IDENTIFYING CATEGORIES OF INCIDENTS:

Diving deeper into the PDNS data, we identified the probable individual threat families seen which are ransomware related, giving us a relative ratio of these higher impact events versus the general (but still significant) population of events.

ACCOUNTING FOR MAJOR INCIDENTS:

We used the IRIS 20/20 dataset to determine the rate of publicly discoverable events – with their frequently large losses – occurring in the public sector. This rate is combined with prior work together with the GCA, using the Verizon DBIR to identify the number of common attack patterns with DNS in the kill chain. Together, these data points allow us to create a distribution for the likely number of major incidents which are DNS related in the PDNS customer base.

CALCULATING AVOIDED LOSSES:

With these three incident types, we used IRIS data and calibrated estimates to determine how much each instance of these incidents would have cost firms, had they not been blocked. Together, these culminate in the loss exceedance curve in Figure 13, demonstrating the losses against which PDNS is helping its customers protect themselves.

Conclusion

The data in this report demonstrates the overall value of PDNS to the UK public sector. It has shown the visibility of threats PDNS provides to those using the service, with 60% of organisations catching at least two unique threats each week. Most significantly, it has enabled us to put a monetary value on the service for the first time, by identifying what the UK public sector is saving through avoided losses - which equates to tens of millions of pounds every year.

As this research shows, public sector organisations face a wide variety of threats. DNS traffic sees a vast assortment of strains of different threat varieties, each with the capability to evolve. While many threats identified by PDNS are not significant, the ones that are have a big impact. PDNS intelligence has been shown to identify and help mitigate high-impact ransomware events and devastating major incidents that would carry a high cost for public sector organisations.

At the same time, PDNS helps organisations cope with the most common and persistent general threats. By acting as a first line of defence against these threats, PDNS saves organisations significant resources on triage and remediation, and frees up analyst time to focus on the threats that most require human intervention.

While no individual security control is a cure-all, PDNS's simplicity of use and broad applicability means that it plays an economically significant role in safeguarding the UK public sector at scale, by helping to prevent avoidable losses through rapid threat intelligence and proactive defence.

Appendix: Methodology

Our goal was to estimate the distribution of data breach losses in the UK public sector which are prevented through the presence of PDNS. To do this, we constructed a Monte-Carlo simulation derived from the OpenFAIR quantified risk management framework. This approach allowed us to quantify the range of possible total losses prevented by PDNS and their likelihood.

Briefly, the model is constructed as follows: using PDNS data provided by Nominet and external data on losses, we simulate the number of adverse, DNS related, threat events likely seen by public sector organisations in the UK. For each hypothetical event, we sample from a distribution of global (dominated by the US and believed to be appropriate for the UK market) losses appropriate for the type of event to determine the losses that would have occurred had the event materialised. We simulate losses across three categories of events: publicly discoverable events, ransomware events, and general incidents. These three categories of events are simulated based upon observed PDNS data and independent and publicly verifiable IRIS 20/20 data.

Leveraging prior work (Measuring the Economic Value of DNS, GCA 2019), commonly seen attack patterns are identified in the large scale data collection seen in the Verizon DBIR. This work established that approximately one-third of all the measured incidents in the Verizon DBIR involved DNS in the kill chain and are therefore subject to prevention by DNS-based controls such as PDNS. We use this rate to determine how many of the IRIS 20/20-derived public discoverable events are likely to have involved DNS as part of the prevention process. This gives us the probability that an organisation encounters a DNS-preventable, publicly discoverable, loss event over a given twelve month period. For each simulated period, we drew from a binomial distribution using this probability of success (where success equals compromise) to determine the number of these loss events seen across the PDNS customer base for a given period.

With PDNS data, blocked queries are identified as belonging to one of several broad categories of threat families (i.e, C&C, Spyware, etc.). PDNS is not the sole

defensive measure in organisations. To extrapolate from observed blocked events to loss events that may have occurred had PDNS not been present, we limit our study to the C&C category. C&C is the PDNS categorisation for traffic trying to reach command and control sites. These C&C sites are not the first step in infection and generally are only contacted when a successful breach of an organisation's controls occurs. Blocking C&C events is a critical event in the kill chain of many forms of malware. By blocking either the ability of malware to exfiltrate data out of an organisation or the malware's capability to receive commands, organisations may stop various forms of harm.

For each unique threat strain seen by PDNS in the C&C category, Cyentia manually investigated the characteristics of the threat using open source intelligence. We manually tagged each threat strain that exhibited ransomware or other extortion like activities as one of its known behaviours. This tagging was used to distinguish ransomware events from more general malware infestations. We performed this to distinguish ransomware events, which tend to have a larger distribution of financial impacts, from the general population of threats whose financial impacts are dominated by the time spent remediating infections.

For both the ransomware and general incident types, losses were drawn from two separate beta PERT distributions of potential losses. We assume that larger losses, being more newsworthy, are represented in our collection of IRIS 20/20 data derived, in part, from news reports, forensic reports, and public disclosures.

These three categories of loss events (public discoverable, ransomware, and general incidents) are combined into an overall distribution of events for the UK public sector. It is this combined distribution of annualised losses that we use for our loss exceedance curve and overall reporting.

This Monte Carlo process was performed using defined seeds for repeatability within each category of loss. 10,000 simulations were performed.

Limitations

For the class of major, publicly discoverable, events, we did not attempt to correct for multiple events occurring at a single organisation over the course of a sample period. It is possible for an organisation to experience multiple major events over the course of a given year. As we are looking at a particular slice of these major loss incidents, we chose this course as a conservative estimate rather than attempt to model repeat events. This affects our results by making them lower than what real world organisations may exhibit. We are assuming that our IRIS 20/20 data, focused on the US public sector is similar to the threat profile faced by the UK public sector.

The data provided to Cyentia is a sample of the overall PDNS traffic. Certain organisations were removed by Nominet due to having unique security requirements or other characteristics not typical of the general population of protected clients. Nominet believes the subset of data provided is representative of the typical organisation in the UK public sector. We also do not account for potential false positives in the PDNS sampled data. As the majority of loss valuations are from the publicly discoverable events, we believe the effect of these limitations are minimal.

Month on month, there is variability of threat feeds, and different things are being blocked. It is the opinion of Nominet that this April – August period is representative of the volumes and nature of blocked queries over a given 12 month period. For the purposes of deriving distributions of particular event type rates, we calculate the per organisation rate of events seen over the four month period and scale that up to a 12 month period to obtain a yearly rate. Again, these distributions are outweighed by the effects of the public discoverable effects, minimising the risk of this materially affecting results.

When measuring DNS queries, the function of caching in the DNS protocol can obscure the total volume of queries made by potentially infected hosts. We addressed these issues by looking at distinct instances of named threat families within the C&C category over a month. This does have the potential to undercount the number of events, but is a conservative estimate to avoid the dangers of inflating data.

Finding accurate figures of the cost of ransomware to public sector organisations can be difficult. The report uses a conservative estimate based on independent research.