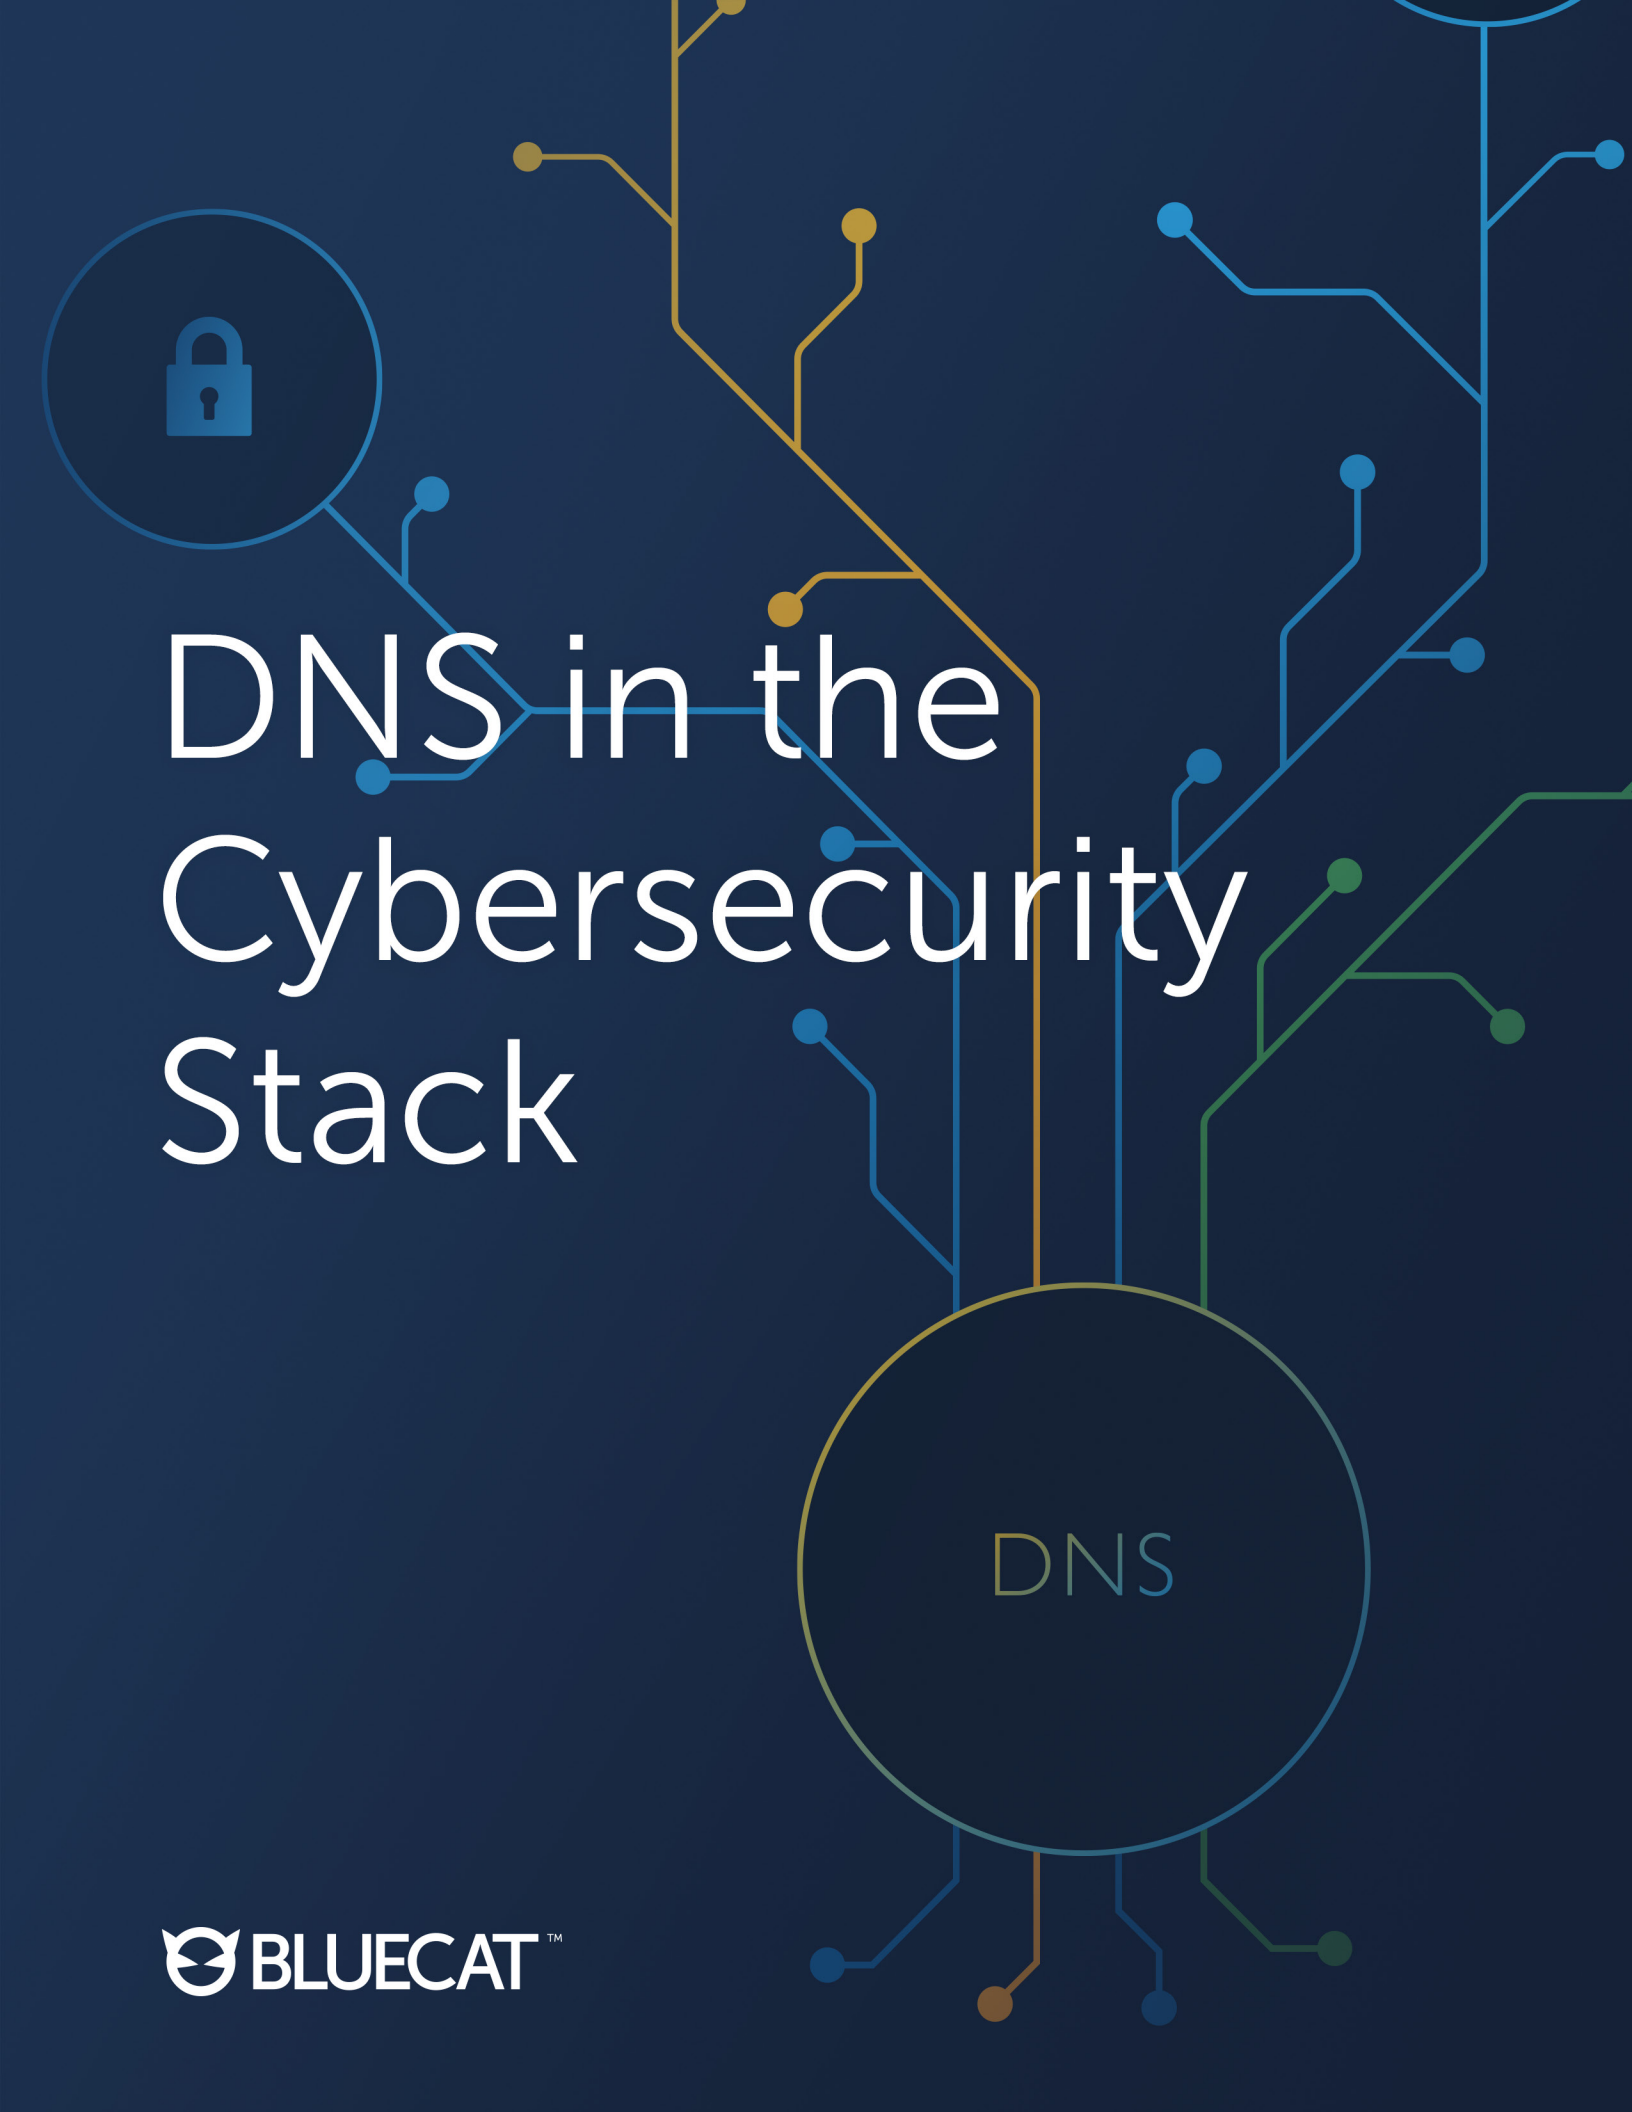

DNS in the Cybersecurity Stack

BLUECAT™

WHERE DOES DNS FIT IN THE SECURITY STACK?

When most people think about DNS, they think about core network infrastructure. That's only natural. DNS has been directing queries since the dawn of the internet, and remains the "plumbing" every network relies on.

The central position of DNS in the network is starting to get it some attention in cybersecurity circles. The use of DNS as an attack vector is a more recent phenomenon, but one which is spreading fast. Studies show that over 91% of malware now uses DNS to navigate through a network, establish command and control, and/or communicate with outside resources. With DNS playing such a critical role in the activities of malicious actors, it was only a matter of time before the world of cybersecurity started to take notice.

Even as the security implications of DNS for cybersecurity become more commonly recognized and understood, a similar consensus around where DNS belongs in a cybersecurity strategy remains elusive.

Just as DNS permeates different layers of the Open Systems Interconnection (OSI) network model, it also defies categorization in the typical "security stack" CISOs build to protect their networks. The fact that DNS also complements and supplements other solutions also makes its overarching security value difficult to fully appreciate.

In the absence of a recognized category or context, DNS will remain an underutilized and underappreciated security tool. Without a framework to evaluate the effectiveness of DNS against other forms of network security, the default will be to treat it like just another security tool that claims to be unique but offers no solid proof. That's a shame, because DNS truly has the potential to disrupt and reshape our entire understanding of how cybersecurity operates.

Clearly cybersecurity will never be a black-and-white discipline. No two networks are alike, and no two organizations have the same risk profile. Yet just as there are clear patterns of cyberattacks, common practices have emerged for cybersecurity as well. As the core of all network communication, it makes sense that DNS should play a role in security – whether it happens over a small network or a large enterprise.

In this eBook, we attempt to place DNS security tools in the context of a cybersecurity stack. By defining the discreet security value of DNS and demonstrating how it complements other cybersecurity tools, we hope to empower threat hunters, CISOs, and network administrators to better understand how DNS can play into their defense-in-depth strategy.

THE BASELINE CYBERSECURITY TOOLKIT

There is no standard conceptual framework or accepted reference model for cybersecurity, but we can talk about common categories of solutions which usually form part of a cybersecurity stack:

Network security tools protect the boundaries or critical pathways of a network. Scanning network traffic for anomalies, known malware signatures, and unauthorized access, these tools are usually designed to alert cybersecurity operators to a potential issue and provide a foundation for remediation. Next-Gen Firewalls (NGFW), Intrusion Detection Systems (IDS), Intrusion Prevention System (IPS), Network Access Control (NAC), and NetFlow monitors all fall into this broad category of products.

Endpoint security systems focus on the client devices which often serve as "ground zero" for cyberattacks. While the tactics and potential attack vectors involved in these systems vary significantly, endpoint security solutions are built around the idea that prevention at the lowest level of the network will decrease overall risk by preventing the spread of malware at the source. Antivirus software, Data Loss Prevention (DLP), Patch Management, Client Firewalls, Intrusion Detection Systems (IDS), and Intrusion Prevention Systems (IPS) all fall into this general category.

Application security systems apply scrutiny to the narrow connection between critical applications and surrounding networks. By protecting the applications and data which malicious actors are most likely to target, application security systems attempt to raise the difficulty of compromise. Web Application Firewalls, database security tools, and code scanners are just a few examples of widely used application security tools.

Data security tools protect against exfiltration, alteration, or unauthorized access to critical information. By hardening the target of malicious actors, these tools make it far more difficult and costly for a cyberattack to occur. Examples of data security systems include encryption, Identity and Access Management (IDAM), Data Loss Prevention, and Data Rights Management tools.

Cloud security systems secure connection points between assets within the network boundary and a third party cloud, providing additional visibility and protection in these critical pathways. By inspecting and intervening in data flows between on-prem and cloud assets, these solutions prevent data leakage and larger compromise. Cloud Access Security Brokers (CASB), content filtering, and traffic inspection systems often fall into this category.

Every CISO and cybersecurity administrator has to choose for themselves which combination of tools fit their specific needs. Those needs can be driven by any number of factors, including size, regulatory compliance, extent of cloud usage, presence of IoT devices, and of course the budget for all of these things.

So while most organizations will have at least one tool from each of these categories, some will have far more. The value of any cybersecurity tool is not only in what it prevents, but also in how it interacts with other tools to form a complete cybersecurity stack.

BLUECAT'S APPROACH TO DNS SECURITY

Just as DNS permeates every layer of the network, it also has a role to play in the "cyber kill chain". Cybersecurity systems tend to identify malware and other threats on the basis of signatures – DNS is a particularly compelling signature detector because it is used just about everywhere on the network.

BlueCat's approach to DNS security leverages the power of this ubiquitous data source to touch multiple use cases in the cybersecurity stack at the same time, providing a cross-cutting capability which is unique in the market.

BlueCat sits on a service point on the "first hop" recursive server, providing visibility into every DNS query which comes from a client – whether it's a computer, mobile device, or IoT sensor. BlueCat then applies security policies to that DNS traffic, allowing cybersecurity personnel to monitor, block, or redirect DNS queries rather than having them automatically resolve. BlueCat also creates a comprehensive log of all client-side DNS data which can be pushed into a SIEM or analyzed directly in our UI.

The unique features and approach of BlueCat's DNS security offering puts it in an interesting place within the cybersecurity stack. It protects critical pathways like a standard network security solution, but does so for both internal and external DNS traffic – a far cry from standard boundary-level filters and firewalls. It protects endpoints by identifying and blocking malicious behavior, but without the agents which can slow down client performance. It protects transit points to public clouds, private clouds, and on-prem resources, but also protects the pathways which lead to those transit points.

There is no silver bullet in cybersecurity, and DNS is no exception. DNS is a vital part of the "cyber kill chain", but it cannot carry the burden of cybersecurity on its own. Like all of these tools, DNS is most effective when placed in context. DNS security provides comprehensive visibility into what's happening on the network, but multiple types of data are often required to create a complete topology of cybersecurity.

HERE'S WHAT WE'LL COVER

Identity and Access Management

WHO IS ON THE NETWORK?

Network access control

Cloud access security brokers

Identity and Access Management Systems

Asset Management

WHAT IS ON THE NETWORK?

Web Application Firewalls

SIEMs

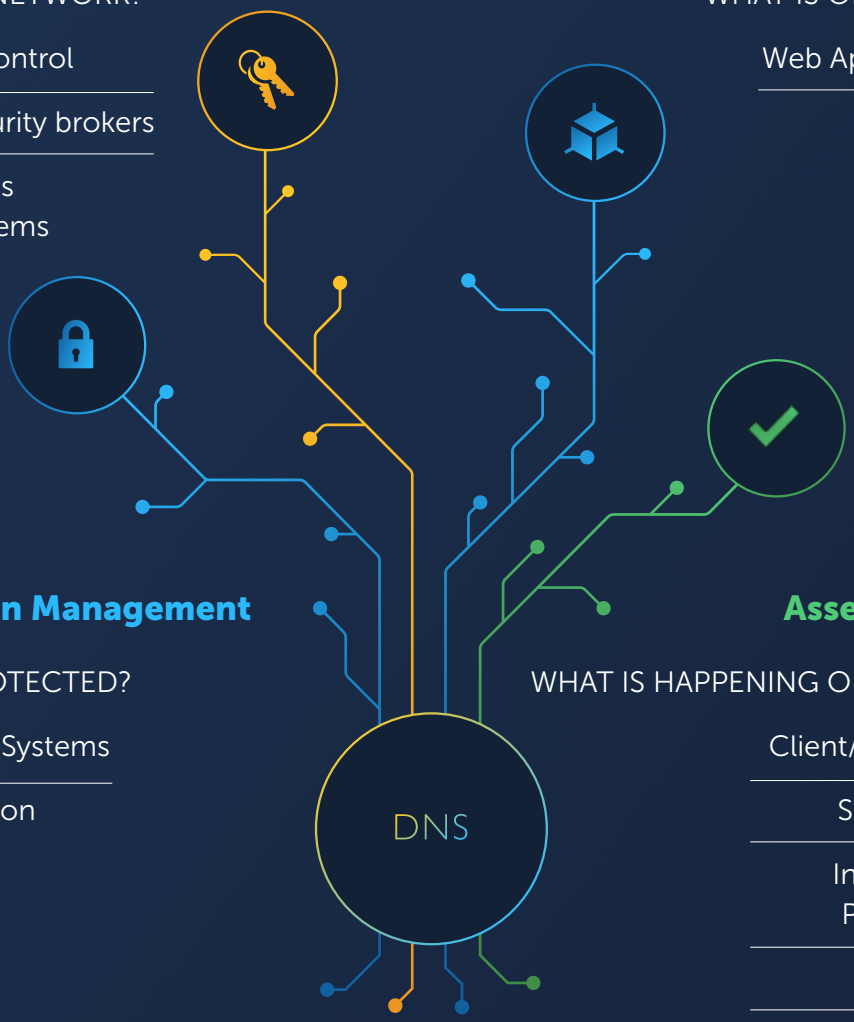

Data Protection Management

HOW IS DATA PROTECTED?

Database Security Systems

Data Loss Prevention

Asset Management

WHAT IS HAPPENING ON THE NETWORK?

Client/Endpoint Firewalls

Signature Detectors

Intrusion Detection/ Prevention Systems

Perimeter Security

Public DNS resolvers

Next Generation Firewalls

Netflow

Packet Capture/Inspection

IDENTITY AND ACCESS MANAGEMENT

Network Access Control (NAC)

Network access control systems are designed around the premise that every device (and by extension, every user) on a network should only have access to certain information. By controlling device connections and configurations through security policies, network access control systems restrict devices and users to authorized areas of the network only. This effectively reduces the attack surface for insider threats by restricting lateral movement beneath the network boundary.

Limitations of NAC systems: Most NAC systems are hardware-based, requiring on-device agents to identify and instruct devices when they connect to the network. This works fine for standard computers and network equipment, but can grow cumbersome and unwieldy when applied to IoT and mobile devices. Since these devices often rely on wireless connections and have limited capacity to house additional software elements, NAC systems often introduce performance limitations or have difficulty consistently tracking connected devices as they move around the network.

BlueCat and NAC systems: Since it applies security policies through network traffic rather than hardware-based signatures, BlueCat's security platform can perform many of the functions of NAC systems without the need for on-device agents. Using DNS as a method of identifying devices and their intent on the network, BlueCat applies security policies to network traffic rather than the device itself. This has the clear advantage of controlling IoT and mobile devices which may not be compatible with on-device agents. It also makes it easier to consistently track devices which may transit through different IP spaces over time.

Once a device is authorized through a NAC system, traffic is generally allowed through unless another system alerts the administrator to shut it down. BlueCat takes more of a zero trust approach, constantly monitoring network traffic and applying security policies to all devices, whether they were authorized in the past or not.

Cloud Access Security Brokers (CASBs)

Gartner created the term "cloud access security broker" to describe a series of products which place security policy enforcement points between cloud service consumers and cloud service providers. This ensures a consistent approach to securing assets even when those assets sit in the cloud – that is, outside of the traditional network boundary. CASBs can include multiple types of security policy enforcement such as authentication, single sign-on, authorization, credential mapping, device profiling, encryption, tokenization, logging, alerting, and malware detection/prevention.

Limitations of Cloud Access Security Brokers: CASBs sit at critical junctures between on-prem and cloud resources, creating "choke points" to isolate and block malicious activity. CASBs provide critical visibility in that narrow area, but often have little insight into the broader context of network data transfers. Without larger visibility into the source and destination of data, it can be difficult for CASBs to filter out critical vulnerabilities from noise.

BlueCat and Cloud Access Security Brokers: While CASBs reach into other areas such as credential management and authentication that BlueCat does not cover, in terms of malware detection CASBs are generally limited by their position in the network. Through the comprehensive client-side data it logs and acts upon, BlueCat offers both wider visibility and more granular information on network activity than most CASBs. Adding the broad store of client intent demonstrated by DNS data to the narrower picture of cloud access created by CASBs will ensure a layered, more broad-ranging approach to network security across hybrid environments.

Identity and Access Management (IDAM) systems

Identity and access management systems ensure that only those who have the "need to know" are able to access critical applications or data. IDAM systems also attach an identity to actions on the network, allowing administrators to monitor activity and identify any anomalies which may surface. From simple log-on passwords to biometrics and continuous identity verification, identity systems can be applied in many different forms and at various levels of sophistication. IDAM systems are often integrated with other data management applications such as Microsoft Active Directory.

Limitations of Identity and Access Management *systems:* IDAM systems are great at the "who" part of cybersecurity, but have little to no insight into the "what". Just having an identity management system, while a significant hurdle for malicious actors, is not a security panacea. Only when identity is attached to network activity logs can malicious activity be prevented in real time or quickly investigated in the event of a breach.

BlueCat and Identity and Access Management systems: BlueCat is the ideal method to identify malicious activity by a client device, but in a forensic investigation scenario administrators will want to dig deeper into who (not just which computer) attempted to compromise the network. This is where correlations between DNS and IDAM data can prove extraordinarily valuable, painting a more detailed picture of who was responsible for particular acts. Both systems are critical to painting a complete picture of network activity – each complements the other.

Database Security Systems

Database security is a specific niche, designed to protect critical data from misuse, exfiltration, or compromise. Through a combination of data encryption, access control, and network segmentation, cybersecurity administrators generally use a layered approach to protecting databases against a variety of threat vectors. Usually databases are secured through a series of standard security services which already exist for other data sources on the network, but there are also database-specific security programs out there. These are usually proprietary solutions offered as part of the database itself.

Limitations of Database Security Systems: Threats to databases and the information they hold can come in many different forms. Often exfiltration is the goal, but deletion or changes to the data itself can also constitute a threat. No single security system can adequately protect against the full range of what malicious actors can dream up. In addition, some database security measures extend beyond the digital realm, requiring physical separations and disaster recovery plans.

BlueCat and Database Security Systems: BlueCat adds a strong layer of protection to any database, and can be configured to provide protection against multiple threat vectors. Specifically, security policies in BlueCat can effectively wall off the database and apply very specific access control rules, blocking unnecessary or inappropriate connections before they make it to the database itself. By monitoring internal DNS queries, BlueCat can also detect attempts by malware to move laterally in search of valuable database assets, protecting the network around the database and reducing the overall attack surface.

Data Loss Prevention (DLP)

Data loss prevention is another broad category of technologies which generally attempt to identify and block unauthorized leaks of sensitive information from a network. Firewalls, intrusion detection and prevention systems, and network access control systems all take on the basic characteristics of a data loss prevention system from time to time, even as they fulfill other security requirements unrelated to data loss per se.

Data loss prevention systems often employ sophisticated machine learning to scan for sensitive information, putting security policies in place to protect it. As sensitive information can be found "at rest", along transit points in the network, and on multiple connected devices, any data loss prevention system must employ a layered approach to cover the wide variety of potential vulnerabilities involved.

Limitations of Data Loss Prevention systems: Since there are so many ways to steal information, data loss prevention systems are by definition always playing catch-up. Many are built to react to the latest trends in data loss rather than proactively block exfiltration. With hackers constantly inventing new types of exfiltration techniques, attempting to predict the next great data heist is a fool's game.

BlueCat and Data Loss Prevention systems: Like most data loss prevention systems, BlueCat cannot claim to be a comprehensive solution. Even so, BlueCat acts on one of the most prominent vulnerabilities in today's networks. 91% of malware uses DNS to move through the network, establish command and control, and exfiltrate data. Gaining control of the vital network pathways which malicious actors have long used to steal data offers clear advantages to cybersecurity staff.

In particular, BlueCat offers unique protection against DNS tunneling, a common data exfiltration technique. By embedding data in DNS queries themselves, malicious actors take advantage of open ports to slip sensitive data past standard filters and firewalls. BlueCat identifies and blocks unauthorized DNS tunneling at the source, preventing this common type of data exfiltration.

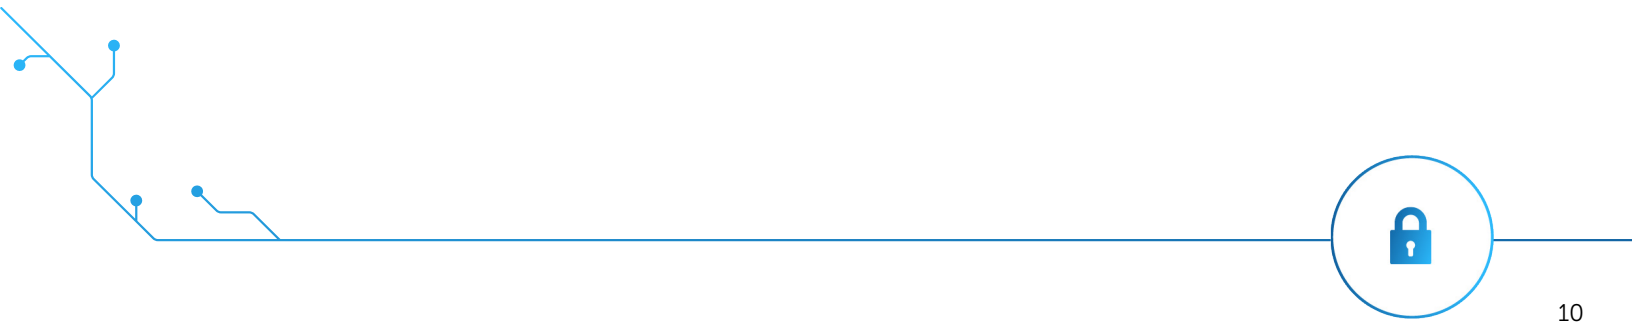

Web Application Firewalls

Web application firewalls are designed to secure http traffic by monitoring and filtering data flows between an application and a client. While they filter against a wide range of threats, web application firewalls were originally designed to counter specific types of vulnerabilities, such as cross-site forgery, cross-site-scripting (XSS), file inclusion, and SQL injection. Web application firewalls can also help to obscure the identities of internal clients and servers, making them less vulnerable to exploitation.

Limitations of Web Application Firewalls: Securing http traffic is a laudable goal, but sometimes it can be a step too late. By the time data is flowing through the http layer, DNS has already established a connection between the client and a remote server. Controlling the initial DNS query can prevent http-based threats from emerging in the first place. Since DNS essentially indicates the intent behind an http connection, it often makes sense to examine the validity of that intent first.

BlueCat and Web Application Firewalls: BlueCat's DNS security system precedes web application firewalls, examining the initial DNS connection before the http protocol takes over. It's a good bet that malicious http traffic will also flow through malicious domains, so it makes sense to check the DNS record against threat profiles before allowing any http connection to proceed. As in the case of packet inspection, having a web application firewall as an additional security layer on top of DNS protections offers protection against specific http-based threats listed above, so it would make sense to have both as part of a security stack.

SIEMs

Security Information and Event Management (SIEM) software is a standard piece of the security toolkit, even if it isn't part of the security stack per se. SIEMs are needed primarily because of what the security stack produces – a whole lot of data. Making sense of that information, organizing it into a useful form, and using it as a repository for investigations is at the heart of what security managers and threat hunters need to perform their everyday work.

Limitations of SIEMs: SIEMs are essential for managing data, but that is basically all they do. Making sense of data is one thing, but acting on it is another thing altogether. In an ideal scenario, a SIEM will help to identify an issue which security software can then correct. Unfortunately, SIEMs often act as aggregators of aggregators – culling the masses of data from deployed third-party sensors and pointing out anomalous behavior. Threat mitigation through the application of security policies, adjustment of configurations, or other actions usually aren't automatic or obvious from the dashboards and data feeds managed by a SIEM.

BlueCat and SIEMs: BlueCat is integrated with all of the major SIEMs, including Splunk, ArcSight, and QRadar. At a basic level, BlueCat interacts with these platforms like all other third party inputs, sending data from a deployed sensor into the SIEM's centralized reporting portal. Yet BlueCat differs in three key ways from the standard sensor input.

First, BlueCat filters information before it reaches the SIEM, sending only the anomalous activity on. Since many SIEMs charge by the amount of data ingested, this can save a considerable amount of money versus the prospect of ingesting all of your DNS logs.

Second, BlueCat provides the complete narrative of a DNS action on the network - including the query itself, the response, metadata, and contextual analysis. Having all of this in one place is far easier than trying to piece it together from separate sensors within a SIEM.

Third and most importantly, Edge provides the vital link between monitoring and action. Not only does Edge point out where anomalous activity is happening, it also allows for immediate action in the form of client-level security policies which monitor or block DNS activity at the source. This functionality goes beyond mere aggregation and reporting of data; it creates the essential actions which make networks more secure. Those actions can also be integrated into a Security Orchestration, Automation, and Response (SOAR) engine to accelerate response times.

Client/Endpoint Firewalls

Client or endpoint-based firewalls have been around for a long time – these are the agents which sit on connected devices and constantly scan against known threat signatures. Those signatures must be constantly updated to account for the latest vulnerabilities and threat trends, and are often backward-looking.

Limitations of Client/Endpoint Firewalls: Client and endpoint firewalls are primarily designed to protect traditional computers and servers, not mobile or IoT devices with limited memory and computing capacity. Since they deploy as device-based agents which continually churn in the background, these firewalls can consume significant memory resources, exacting a cost in system performance. There is also a human factor at work – many client/endpoint firewalls require periodic permission from users to download the latest signatures. When these downloads aren't automatic, they essentially negate the value of having an endpoint firewall in the first place and make for inconsistent coverage across the enterprise.

BlueCat and Client/Endpoint Firewalls: Device-based agents are needed to protect against common brute-force hacks, but their utility is limited to specific types of endpoint devices. BlueCat covers a broader range of client machines (IoT devices in particular) and eliminates the need for agents which can slow down network performance. Unlike client/endpoint firewalls, BlueCat does not require constant updating of threat data, reducing the human factor and ensuring that the entire enterprise has consistent protection against threats.

Signature Detectors

Many cybersecurity systems use the common activities associated with malware – beaconing, lateral movement, etc. – to alert threat hunters to anomalous activity on the network. These systems are constantly at work, sifting through various data sources in the background as they look for activity which meets a pre-programmed threat profile. Increasingly, these systems use machine learning and artificial intelligence to map networks, develop data on "normal" network activity, and generate deeper insights into potential vulnerabilities.

Limitations of signature detectors: Signature detectors look at a relatively broad spectrum of cyber threats. Rather than taking a single data source or attack vector into account, they are programmed to look for a wide variety of anomalous activity. That can be a disadvantage in the sense that these types of cyber tools can produce a lot of data – much of it false positives. If used in concert with a SIEM or other tools to filter alerts, however, the workload produced by these systems can be manageable.

Since they are constantly trolling the network and churning through various data sources, signature detectors can also introduce latency issues. It takes bandwidth to capture and process all that information, and sometimes the total cost of those solutions can be larger than just the software's initial price tag – particularly in the cloud, where additional usage can increase the price of compute.

The dynamic nature of artificial intelligence and machine learning in signature detectors is both an advantage and a potential weakness. These new tools can adapt to new types of cyber threats, but since they are constantly shifting positions on the network the view they create can be inconsistent and lacking in context.

BlueCat and signature detectors: BlueCat's DNS security platform is itself a type of signature detector. While it focuses on the DNS protocol specifically, the advantage of this approach is that almost all cyber threats use DNS at some point. The difference with BlueCat is that it not only identifies anomalous behavior, but also helps to prevent and remediate it. Signature detectors are one step better than perimeter security systems (covered below), in that they can often pinpoint the source of anomalous activity – even to the source IP level. That is usually the end of their functionality, however.

BlueCat has a distinct advantage over most signature detectors in that it can use security policies to block, monitor, or redirect malicious traffic at the source. It not only alerts threat hunters to the presence of an issue, but actively starts to do something about it. This preventive step is more than just "cutting through the noise" of constant alerts – it actively prevents alerts by dealing with threats before they metastasize.

Intrusion Detection/Prevention Systems

Intrusion detection and prevention is a broad category, but there is a common underlying workflow: searching for anomalous behavior on a network and then mitigating that behavior. Underneath that broad umbrella, there are various methods for identifying potential threats. Some rely on deviations from baseline norms, while others search for specific signatures commonly associated with malicious activity. Some look at specific network devices, while others are deployed at chokepoints to identify movement across the network as a whole.

Limitations of Intrusion Detection/Prevention Systems: Intrusion detection and prevention systems go deeper than boundary level protections, delving into core network systems to deliver a more granular defense. While the benefits are significant, there are inevitable tradeoffs with this approach.

Most intrusion detection and prevention systems are static – they sit in one place and/or monitor a specific type of traffic. This means that dynamic threats which change as they move through the network or do their work through gradual accumulation of fragmentary data are unlikely to be detected. These systems can also be expensive to deploy, as they have to sit in multiple network locations and are often hardware-based. Intrusion detection systems also create a lot of data, all of which must be purged of false positives.

It's also worth noting that DNS is a known blind spot for most intrusion detection and prevention systems.

BlueCat and Intrusion Detection/Prevention Systems: Of all the comparable pieces of the security stack, BlueCat is most like an intrusion detection and prevention system. In a nutshell, it collects and analyzes data which streams through critical points in the network, identifying anomalies or known malicious signatures and blocking them based on security policies.

The difference with BlueCat is that it uses the core protocol behind every network communication. Every other intrusion detection and prevention system on the market makes a strategic choice about what it will (and will not) pay attention to. That mix of signatures is part of the "secret sauce" which different vendors will claim are more significant or impactful. Rather than provide a menu of specific malicious activities to block or ignore, BlueCat is a smorgasbord. It looks at the full swath of network communications, including the data exfiltration strategies which use DNS itself.

Perimeter Security

Perimeter security systems such as firewalls, content filters, router-based protections, and honeypots are some of the most widely used cybersecurity protections on the market today. At least part of the appeal of these solutions is their simplicity and ease of use. Just deploy a box or a virtual machine at your network boundary, and disallow all known malicious traffic. The specific tactics and features vary slightly, but the theme is the same: keep malicious inbound traffic from entering the network and block malicious queries from getting out to the internet.

Limitations of Perimeter Security: Perimeter security is a no-brainer that every cybersecurity architect starts with. Blocking or diverting obviously malicious traffic is the least than anyone can do. Yet the position of perimeter security systems ultimately limits their effectiveness. To be sure, blocking traffic at the network boundary is necessary. Where that traffic is blocked makes a difference, however. Applying security policies at the network boundary simply isn't enough to capture the full range of threats flowing into the network and operating underneath the network perimeter.

Network layers and recursive servers effectively obscure the ability of perimeter defenses to identify the source of malicious actors within the network. Perimeter defenses also tend to block one part of the conversation – the inbound traffic – rather than dealing with the full input and output of malicious activity. This necessarily limits the scope of response for any perimeter security solution.

BlueCat and Perimeter Security: BlueCat and perimeter security systems work together like bifocals. For everyday visibility into threats that are a good distance away, perimeter security will do the trick. For close-in threats that appear underneath the network boundary, a different kind of visibility is needed. In the end you'll need both to see all relevant activity, but each has its specialty.

Even more, BlueCat provides the context that perimeter security systems need to be effective and relevant across a wider range of threats. By capturing all of the query and response information at the client level, BlueCat can see what's going on long before those queries make it to the network perimeter. For internal queries, this allows BlueCat to deal with a type of malicious activity that perimeter defenses would miss altogether. For external queries, it helps to correlate the source or destination IP with the malicious traffic identified on the network boundary, even through multiple recursive layers.

BlueCat's position on the network also acts as a pre-filter for perimeter security systems. By blocking malicious activity at the source, BlueCat lowers the amount of activity that even makes it to the network firewall, allowing the perimeter security system to concentrate its resources on the inbound queries which it was primarily deployed to filter against. In this sense, BlueCat is a security system which both reduces dependence on perimeter firewalls and supplements their functionality.

Public DNS Resolvers

As the cybersecurity community started to recognize DNS as an attack vector, various efforts were launched to make DNS more secure. In 2009, Google launched a free public DNS service at 8.8.8.8. In 2017, IBM and its partners launched Quad9 (9.9.9.9) – another DNS service offered free of charge. OpenDNS was also a free DNS service, until it was bought by Cisco and turned into a paid product (Cisco Umbrella).

While the features, functionality, and technical deployment of these solutions vary slightly, the concept is generally the same. Anyone can send their DNS through these public services and filter against known malicious domains. With support for Anycast and DNSSEC, these solutions offer a relatively easy way to implement basic security for DNS queries.

Limitations of Public DNS Resolvers: As with so many other things in this world, with public DNS resolvers you get what you pay for. These services offer a degree of baseline protection against common, known threats, but that's about it. They don't offer any analytics, customization, or the ability to ingest outside threat feeds. Since they are used for external DNS only, these services provide no insights into threats against internal DNS traffic.

BlueCat and Public DNS Resolvers: BlueCat's approach to DNS security goes far deeper than any public DNS resolution service in the same way that it adds more value than perimeter security systems (see below). Since it sits inside the network and at the client level, BlueCat can block malicious DNS queries for all traffic – both internal ("east-west") and external ("north-south").

Public DNS resolvers offer no insight into what's happening with your DNS – they simply block queries using threat feeds and a business logic which remains opaque to end-users. BlueCat puts the user in charge of security policies, allowing for customization of threat feeds and definition of malicious activity according to the unique contours of the network.

Next Generation Firewalls

Traditional firewalls are relatively simple – they examine information flowing through ports and protocols, and block anything which appears on a blacklist. Next-generation firewalls add new functionality into the mix, including the use of application-level information, deep packet inspection, and intelligence gathered elsewhere on the network. They are a more flexible and ultimately more all-encompassing type of firewall – one which looks not just at the flow of information but also what is contained within that flow.

Limitations of Next-Generation Firewalls: To be sure, a next-generation firewall is more sophisticated and ultimately more effective than a traditional firewall. At the same time, any firewall which sits on the network boundary has natural limitations on what it can detect and when.

As with any perimeter-focused security tool, next-generation firewalls focus solely on the network boundary rather than the complete picture of network traffic. Since they provide no visibility into the internal workings of a network, next-generation firewalls often leave advanced persistent threats and malicious insiders with a blank check – the only chance a boundary-level protection system has to catch them is when they attempt to communicate with the outside internet.

While they do a fine job of scrubbing outbound queries through known ports and protocols, next-generation firewalls are not equipped to protect against attacks which utilize the protocols themselves. DNS tunneling effectively skirts next-generation firewalls, and can easily be used for data exfiltration. The use of domain-generating algorithms (whose short half-life means they won't appear on any blacklist) is also a way of using DNS to get around even the broader functionality a next-generation firewall provides.

BlueCat and Next-Generation Firewalls: Boundary level firewalls (whether traditional or next-gen) are always going to be necessary, but users should plan for the situations where they offer little or no functionality. DNS happens to be one of those areas, protection of internal network assets is another.

BlueCat's DNS security solution is an ideal complement to a next-generation firewall in that it directly addresses these blind spots. BlueCat's platform is built to identify and block malicious DNS tunneling activity, protecting networks from malicious exfiltration. BlueCat provides innovative ways to identify and block domain-generating algorithms which even sophisticated next-generation firewalls cannot detect in real-time. And of course, BlueCat's position at the client level allows for visibility and action before things go wrong instead of waiting until malicious queries get to the network boundary.

NetFlow

NetFlow is a broad term used for the ability to monitor data as it moves through a network. Originally developed by Cisco in the 1990s as a part of its hardware offering, NetFlow capabilities are now provided by many vendors through deployments of both software and hardware.

At a basic level, NetFlow solutions capture information about the quantity and characteristics of information flowing through a network at key points – usually hardware like routers or network switches. This provides network administrators with information about traffic patterns which can be used to optimize network performance. It can also be used to reconstruct a data transfer, which can be useful in forensic investigation scenarios.

Limitations of NetFlow: NetFlow can build an entire traffic session end to end, and is useful for in-depth investigations when the contours of malicious activity are already known. At the same time, NetFlow cannot capture full packets, and since it must travel through enabled routers retrieval time is limited due to storage size. On its own, NetFlow cannot identify or block malicious activity – it is a passive monitor only.

BlueCat and NetFlow: NetFlow is only useful as an investigatory tool if malicious activity is already identified, preferably with specific information about the source device. Since it sits on the client-facing "first hop" of any transaction, BlueCat's DNS security system can identify the source device and pinpoint the need for further investigation through the application of security policies. Unlike the hardware-based NetFlow sensors which must be installed on every part of the network, BlueCat has a light touch that is software-based, providing full coverage of network activity on both internal ("east-west") and external ("north-south") traffic without the need to re-architect the entire network.

Packet Capture/Inspection

Packet capture and inspection technology goes beyond monitoring data patterns to provide visibility into the actual payload that flows through networks. By inspecting the content of information rather than merely its size and destination, packet capture and inspection can help to prevent data loss and ensure data integrity. In a forensic investigation, packet capture and inspection can provide a great deal of detail into what happened to information as it was transmitted through a network.

Limitations of Packet Capture/Inspection: As with NetFlow, packet capture and inspection can only happen in places on the network where sensors are set up to gather information, limiting its reach on the network. The storage requirements for packet capture information can quickly become costly, which usually leads to data being kept for a limited amount of time. While packet capture and inspection provides a great deal of information, this can also be a downside – sorting through the sheer volume of data can be time-consuming and difficult if malicious activity isn't identified in advance.

BlueCat and Packet Capture/Inspection: BlueCat's DNS security capabilities complement most packet capture/inspection tools. By providing client-facing intelligence about malicious queries, BlueCat can help to pinpoint the specific packets which are worth investigating further. BlueCat also proactively blocks queries, promoting a more active defense which lowers reliance on packet capture and inspection tools. Over time, the patterns of DNS activity (both legitimate and malicious) demonstrated by BlueCat can indicate where packet capture and inspection tools will provide the most value.

ABOUT BLUECAT

At BlueCat, we leverage the power of DNS to rationalize, secure, and optimize the performance of any network. Many of the world's largest, most critical, and most complex networks rely on BlueCat to deliver the innovative DNS capabilities they require to turn their IT systems into a competitive advantage.

Network security is complicated. There are thousands of solutions out there, each of which is unique in its own way. While no single product can offer comprehensive security, we happen to believe that DNS offers value across the enterprise in ways which make it both efficient and cost-effective. That's why we wrote this eBook - to show exactly how our approach to network security through DNS sets us apart from just about every other kind of tool out there.

We've shown you how DNS fits into your cybersecurity stack. Now all that's left is to show you how BlueCat's DNS security platform fits into your operational cadence. We'd love to show you exactly how it all works.

CONTACT US