Using BlueCat Adaptive DNS in the Cloud

Executive Summary

This document describes the challenges that hybrid, multi-cloud architectures present for DDI teams, specifically related to DNS, DHCP, and IP address management (DDI). These include:

1. Lack of visibility into cloud DNS leaves network teams guessing at how networks and IP space are allocated, leading to data errors, conflicts and outages.
2. Cloud and on-premises DDI are separate entities preventing centralized control and management of enterprise DDI.
3. Complex DNS forwarding rules govern resolution across cloud(s) and data center, consuming management resources and threatening misconfigurations that lead to outages.
4. The inability to deliver SaaS-based services in an optimal way impacts end user experience and increases costs.
5. The inability to apply consistent security policies across cloud and data center environments, and the lack of full visibility into the activity in each one.

This paper discusses BlueCat's approach to solving these challenges through the application of its Adaptive DNS solution for hybrid cloud environments.

The Cloud Challenge

Being in the cloud means moving fast. Cloud and DevOps teams are constantly standing up new compute, tearing it down, and moving workloads. For developers, this is a pretty exciting situation to be in. The entire cloud environment is built to give them what they want, when they want it.

When all of this innovation is happening in the cloud, the consequences for core network infrastructure are usually an afterthought. Cloud and DevOps teams use cloud-native DNS services or free, stand-alone DNS resources such as BIND, which are spun up on the fly. They often don't know what that means for the rest of the network, and they probably don't care. They just want to keep moving.

But the dynamic nature of cloud development introduces risks that need to be managed with consistent security policies across the network. Network administrators care deeply about how cloud development and adoption impacts legacy network infrastructure. They're the ones who have to deal with the back-end chaos that results from everything the cloud and DevOps teams do. Cloud and DevOps teams expect all this stuff to "just work". Network teams have to make it work.

Below are a few of the challenges network teams face when dealing with the consequences of cloud DDI infrastructure, and how BlueCat solves some of these challenges.

Visibility

Infrastructure teams at large organizations have a mandate to understand holistically how the business is allocating IP space in order to optimize network performance and deliver critical services. But all too often, they have precious little insight into what's going on in the cloud. This creates numerous challenges.

Data conflicts, errors and outages: Whenever cloud/DevOps teams stand up networking components in the cloud, they assign IP space to those components. In the absence of a single source of truth for assigning IP space across environments, those IP ranges may already be assigned to another area of the on-prem network. These conflicts cause network outages that reduce productivity (and profitability) to zero for as long as they persist.

Unnecessary cloud expense: Despite the promises of cloud vendors, the cloud isn't cheap. While it is easy and natural for DevOps teams to spin up resources on demand, it is often less natural to remove those expensive resources when they are no longer needed. Without any visibility into what services have been created, it's hard to keep track of cloud usage. The bills keep coming in even though no value is being delivered. DNS records are a common and efficient way to track application consumption across hybrid environments. But only if DDI has been deployed enterprise-wide as a homogenous, single source of truth.

Control

Maintaining centralized visibility and control over core infrastructure resources is critical to error-free, rapid delivery of network services across the enterprise. When the cloud creates autonomous areas of the network with their own DDI resources, that centralized system begins to erode, and with it the ability to deliver services efficiently and effectively.

At global enterprises, few cloud-native applications and services can operate unfettered from legacy data systems that contain customer, financial or product information. That means even cloud-first solutions must go back to the data center to complete transactions. That necessarily involves traversing fractured DNS management tools.

The inevitable result is service delivery delays as staff work to integrate disparate DDI systems. And when developers used to moving at cloud speed must slow down, they get frustrated and start to look for alternatives. This serves to further fragment the IT landscape, increasing complexity and adding more DNS silos.

Because cloud is often its own island of compute, networking, and DNS, it can be difficult to connect everything together with on-premises infrastructure. In addition to service delivery delays, this makes it difficult to seamlessly move application workloads around to meet user needs. In the end, the promise of the cloud becomes difficult to achieve because it isn't fully integrated into the rest of the environment.

Imposing change controls on application developers doesn't work. Even asking these teams to simply document their rapid changes and feed these back to teams responsible for network infrastructure is viewed as archaic and bureaucratic.

The result? IT teams are left holding the bag. Their plans to properly manage IP space across single clouds or hybrid cloud environments become impossible to implement. When the lack of a centralized authority for DNS resolution results in data conflicts that bring down the network, the cloud and DevOps teams somehow avoid the blame. Instead, network teams are faulted for "slowing things down." It's a classic problem of network admins having all the responsibility but only some of the authority over network infrastructure.

Complexity

In order to overcome the challenges associated with a patchwork of hybrid cloud infrastructure, the network team often needs to build and manage conditional forwarding rules to bridge the gaps between different environments. Best practices provided by cloud vendors push organizations to spin up private clouds in many regions with separate networks to isolate critical functions and provide security controls. Add that to the scale and speed of cloud adoption at many large enterprises and it is easy to end up with thousands of conditional forwarding rules to patch everything together. As workloads move and new applications are developed, these rules will need to be constantly updated.

This creates unmanageable complexity at many organizations. It falls to a single person or a small group of experts to maintain this rat's nest of routing and forwarding rules. Only they actually understand how the system of rules was built and how to maintain it. Their full-time job becomes the maintenance of these rules. Resources are pulled away from more important things, like ensuring that critical new services are brought to end users and customers quickly.

Ordinarily, administrators want to apply some kind of automation to sort out this mess. The challenge is that none of the cloud-native DDI services support automation outside of their own environment. Maintaining dozens of "islands of automation" for each of the cloud instances is just more trouble than it's worth. So organizations turn to complex overlay solutions that again require constant maintenance and development to keep up with business requirements.

Optimization

Of course, not all cloud services are going to be provided from an organization's own hybrid cloud environments using their own DNS services. Companies are increasingly likely to first consume services directly from public cloud SaaS-based service providers. This presents its own challenge: How to effectively connect users to those services without having to route all of their DNS and application traffic back to a centralized location. The cost of MPLS-based WAN services makes that a costly proposition. So organizations have begun projects to access SaaS services using local internet links to ensure higher performance, localized end-user experience, and reduced operating costs. Intelligent routing of DNS traffic to services that may exist in the data center, a company-controlled hybrid cloud, or out on the public internet is a massive challenge for network administrators.

A second challenge is providing an appropriate, localized end-user experience for remote workers. If DNS traffic is routing to centralized data centers, not only do WAN costs increase, but all end users receive pointers to resources that are local to that data center, not to where they are actually located. So a user in Germany trying to access a SaaS-based solution may end up accessing that service from servers halfway around the globe, and in the wrong language to boot! This severely degrades both the performance and usability of the solution being accessed.

Security

Network security is a hard-enough task when all the infrastructure is on-premises. Moving to the cloud introduces even more complications.

Suddenly administrators are securing information in someone else's data centers, triangulating against someone else's infrastructure, and dealing with someone else's software running through the network. On top of that, there's that whole class of cloud-specific malware, which takes advantage of the unique architecture of the cloud to exploit new security vulnerabilities.

The shared responsibility model used by most public cloud providers offers cold comfort for network security teams. On one hand, the sheer scale of resources cloud providers devote to physical and data security is beyond what most companies or even governments could deliver on their own. On the other hand, cloud customers are on the hook to secure everything outside of the cloud provider's infrastructure – not an easy task by any means.

In an ideal world, customers should be able to simply extend the security architecture created for on-prem environments into the cloud. Everything would be consistent, and the security controls would simply scale into a new environment. In reality, most security teams don't even have visibility into what's happening in the cloud. Actual control over events seems like a pipe dream.

DNS is the common denominator that can bridge the security gaps inherent in hybrid cloud environments. That's because every query on the network – whether on-prem or in the cloud, legitimate or malicious – uses DNS.

When customers have visibility into what's happening in DNS, they can create consistent security controls across the enterprise. More specifically, if customers have visibility into internal DNS records and data – DNS at the level of devices, VMs, and containers – they can apply security policies regardless of where individual assets sit.

And even better, in the event of a security incident, that visibility becomes critical to aid in the investigation of the cause, scope, and impact. Faster identification of the issue and the ability to quickly apply DNS-based policy controls to limit the damage has a huge positive impact on the incident's risk and cost.

BlueCat Adaptive DNS

BlueCat's mission is to reduce the complexity caused by inefficient, disconnected network services in the cloud through an approach we call Adaptive DNS. This approach gives network administrators the power to thrive in a complex, hybrid cloud world. They won't get buried in an avalanche of conditional forwarders, disparate DNS services, and conflicting sources of information.

When it comes to implementing DDI in the cloud, there's no single architecture that works for everyone. Every network is different, and the goals supported by every network vary widely. A more flexible architectural approach is needed – one that doesn't rely on boxes in the data center. This approach needs to offer easy integration with cloud-native tools already being leveraged by cloud teams, as well as traditional DDI solutions. It needs to be a "deploy anywhere" solution – in the cloud, in a data center, or at the network edge – with a pricing model that doesn't lock buyers into rigid deployment choices.

BlueCat helps organizations find the cloud visibility and control they need without disrupting the pace of innovation by:

- **Automating discovery of cloud DNS data, networks, and other resources:** BlueCat's powerful, cloud-agnostic discovery platform lets administrators quickly identify, map, and import existing cloud data and infrastructure into a centralized management platform. Once there, it can be viewed and managed alongside data from existing networks and data centers. Because discovery happens regularly across all cloud platforms, administrators are assured that they will see any changes made by other organizations as they manage cloud-based services and infrastructure.

- **Establishing a consistent, centralized platform for DDI:** BlueCat's unified platform acts as a single source of truth of IP, namespace, and DNS records, regardless of how or where those records are assigned on the network. With BlueCat, administrators can extend core DDI infrastructure into the cloud or integrate it with cloud-native DDI services. This isn't necessarily an either/or decision, and there are several strategies that offer different paths to the same goal. BlueCat provides the flexibility to tackle this issue in the most appropriate way.

- **Deploying network automation and orchestration tools:** Once a single source of truth for DDI is in place, cloud and DevOps teams can operate quickly at scale by calling on those DDI resources with BlueCat's network automation tools. Or, even better, they can use the major cloud orchestration platforms that organizations are already embracing, such as Terraform. Self-service provisioning connected to an automated DDI infrastructure is a prime example of the value that fully integrated infrastructure can provide to cloud and DevOps teams. It gives them the power to get the resources they need quickly without creating more problems for their network infrastructure colleagues.

- **Providing advanced network services designed for the cloud:** BlueCat's network services lower costs, enhance user experience, and increase efficiency by optimizing cloud operations. With direct internet access and traffic steering, BlueCat cuts through the complexity of the cloud, easing the management burden on network teams.

- **Enhancing security through visibility and control over network activity:** Using BlueCat's core network services, network and security teams get complete visibility into cloud operations, as well as the ability to apply strong security policies uniformly across the enterprise.

Enabling Universal Discovery and Visibility in Hybrid Environments

From a DNS perspective, the common challenge organizations face in supporting a hybrid cloud architecture is the ability to enable bi-directional resolution and maintain complete visibility across all platforms in a central location. That is the promise of DDI after all – to be able to manage an organization's IP, namespace and DNS records across all environments, whether on-premises to cloud, cloud to on-premises, cloud to cloud, within tenant, across tenants, or out to the internet. But that becomes difficult in hybrid and multi-cloud architectures where the cloud vendors control IP distribution.

To solve this challenge, BlueCat provides IT teams with automated workload discovery, IP addressing, and DNS deployment within existing cloud deployments. This allows for real-time visibility of dynamic workload changes within multi-cloud environments. By enabling a consistent approach to DNS and IP visibility across the entire network, BlueCat's Adaptive DNS reduces provisioning errors and DNS namespace conflicts.

Dynamic visibility is more than just cloud DNS record assignment. BlueCat provides visibility into the creation of entire address blocks/networks, private network blocks/subnets in Azure VNet/AWS VPC, workload instances, and related IP addresses and DNS names.

To do this, BlueCat starts by polling the cloud provider to fully document its IP infrastructure. This discovery process first occurs at the region level. A unique configuration is then created within BlueCat for each discovered region. All public address space, network blocks, and subnets within the cloud provider region are then dynamically created. This occurs independently of whether IP space is actually being utilized, allowing for dynamic allocation/reallocation of internet-facing public IP addresses that may be utilized on compute workloads.

Any private address space contained with any discovered private networks (Azure VNet/AWS VPC) is then dynamically added to BlueCat's IP address management (IPAM) solution and represented as network blocks and subnets.

In BlueCat's approach to DNS in the cloud, IT teams do not have to manually create network blocks and subnets. Instead, automated documentation of IP address space in the cloud and on premises is dynamically created, providing networking teams with a single pane of glass for managing all IP space. This includes visibility into cloud address space that may have been running for extended periods without being formally documented.

The second phase of the discovery process focuses on workloads within private networks. Any compute discovered, whether started or not, is added dynamically to the BlueCat IPAM solution as a device instance. These instances hold additional metadata including machine size, owner, and whether the instance is started or stopped. When the device instance is added to BlueCat, any IP addresses, both public and private, are added to the infrastructure discovered in the first phase.

BlueCat's Cloud Discovery and Visibility solution is unique in how it dynamically represents cloud compute that is currently running and disassociated from internal corporate DNS domains. BlueCat documents DNS records that are automatically allocated by cloud solutions upon device instance creation as metadata. Even more importantly, BlueCat domains are associated with corporate naming policy, allowing for easier service discovery by internal resources.

Dependent on the cloud provider, BlueCat also:

- Creates unique DNS views and zones in its address management solution for any public or private hosted zones with cloud DNS services such as Amazon's Route 53 or Azure DNS.
- Documents any IP-based load-balancing devices utilizing cloud-native capabilities as special device instances.

BlueCat's approach to phased discovery will allow for future discovery enhancements such as documenting actions initialized in the cloud, like AWS CloudFormation or Azure templates.

BlueCat ensures that any change to network infrastructure done in the cloud – from cloud assignment of a single IP address to creation of entire networks via orchestration tools – is reflected in BlueCat's address management system. This allows application and cloud teams to operate unfettered in hybrid environments, while ensuring that infrastructure teams can see, and get out ahead of, potential IP conflicts. These conflicts may cause errors that pose serious risks to business continuity.

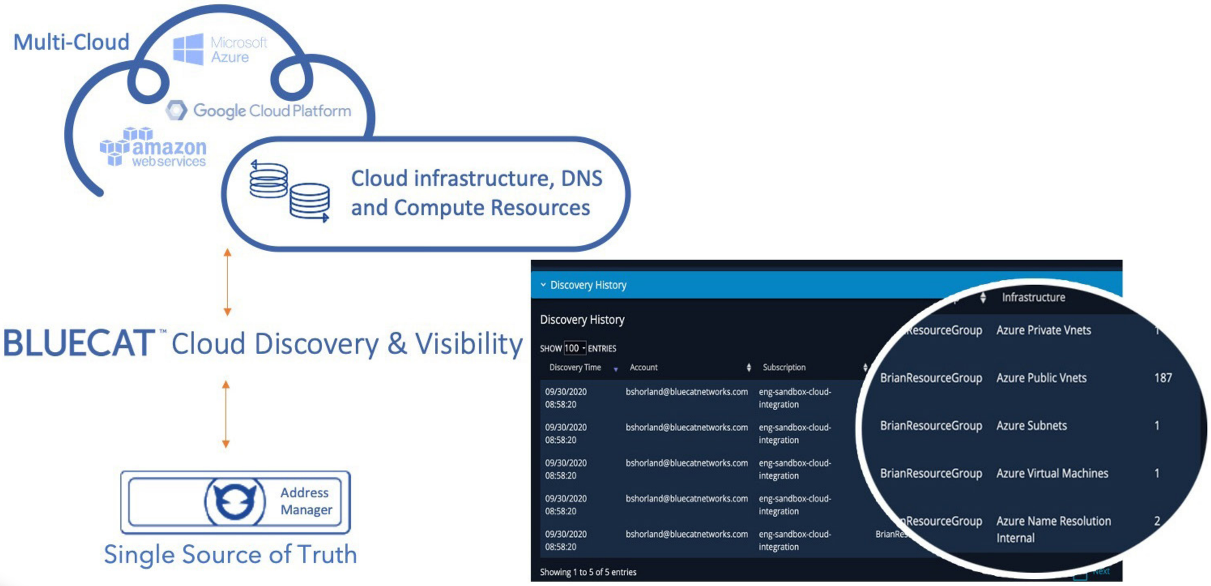

Discovery History

Discovery History

SHOW 100 ENTRIES

Discovery Time	Account	Subscription		Infrastructure	
09/30/2020 08:58:20	bshorland@bluecatnetworks.com	eng-sandbox-cloud-integration	esourceGroup	Azure Private Vnets	1
09/30/2020 08:58:20	bshorland@bluecatnetworks.com	eng-sandbox-cloud-integration	BrianResourceGroup	Azure Public Vnets	187
09/30/2020 08:58:20	bshorland@bluecatnetworks.com	eng-sandbox-cloud-integration	BrianResourceGroup	Azure Subnets	1
09/30/2020 08:58:20	bshorland@bluecatnetworks.com	eng-sandbox-cloud-integration	BrianResourceGroup	Azure Virtual Machines	1
09/30/2020 08:58:20	bshorland@bluecatnetworks.com	eng-sandbox-cloud-integration	BrianRe nResourceGroup	Azure Name Resolution Internal	2

Showing 1 to 5 of 5 entries

Establishing a Consistent, Centralized Platform for DDI

While gaining visibility into and reconciliation of cloud-provisioned IP and DNS records is certainly a big step forward for network teams struggling to keep up with application development and deployment in the cloud, it is only the first step. Why settle for visibility when IT staff can gain control and management of DNS services in the cloud equal to what is done on-premises? And if they can do so without sacrificing the speed and agility that DevOps teams crave?

Extending on-premises DDI management capabilities to cloud environments allows administrators to provide consistent, localized, secure services to those locations, resulting in several key benefits for cloud teams.

- **Improved DNS performance:** By providing local DNS services from a centrally managed platform, DNS administrators can ensure that cloud applications and services have local access to the DNS data that they require to operate. Instead of sending recursive DNS queries to the data center to find the authoritative information required to process a user request, the data is local and retrieved instantly to service the need.
- **Consistent automation:** Cloud experts demand that the everyday tasks required to build and maintain services are as highly automated as possible so that they can focus on delivering value to customers. Extending DDI to cloud environments is a critical step in automating DNS tasks, since many automation requirements must extend beyond a cloud-native DNS platform in order to be fully effective. Providing a local automation endpoint that can span multi-cloud and on-premises environments lets automation be built once and applied globally. Integrations with cloud orchestration solutions such as Terraform allow cloud teams to work in tools that they are familiar with while ensuring back-end consistency and visibility into their changes.

- **Centralized control:** Cloud teams routinely utilize multiple cloud platforms, multiple instances, and hundreds or thousands of individual networks. Managing all of the various DNS and IPAM capabilities that those environments may require can slow down the real work of the cloud team. Centralized control on a common platform is managed from a single point but with the flexibility to delegate management of cloud-facing data to the right consumers. This allows for speed and flexibility while maintaining control and consistency across all locations.

Taming Complexity and Embracing SaaS with Intelligent Forwarding

BlueCat gives network teams the control they need over pathways of data and compute flowing through hybrid cloud environments. Once the foundation of a standardized DDI infrastructure is in place, BlueCat uses automation to solve the problem of conditional forwarders.

The concept is simple. Instead of managing a complex and changing set of single-option DNS resolution paths, network teams can provision multiple resolution possibilities. If the first DNS query comes back with an NXDOMAIN response, the query will automatically re-route to the next priority location. The solution continues to attempt multiple pathways until it finds the right answer.

Managing multiple resolution pathways across a hybrid cloud environment is much easier when they are all represented in a single IPAM interface. That enables network administrators to have the enterprise-level visibility and control needed to operate hybrid cloud environments at scale, taking full advantage of the nimble DevOps and cloud development tools.

Leverage DNS for Direct Internet Access

Of course, overcoming DNS routing challenges doesn't just apply to internal and private cloud environments. End users need to be able to consume appropriate, localized, and authorized SaaS-based services as well. For these use cases, BlueCat's Intelligent Forwarding capabilities allow administrators to leverage local ISP links from in-country DNS providers to ensure the best user experience. And because the rules of DNS resolution are consistently managed across the enterprise, it is easy to allow some SaaS services to be consumed directly while others might be restricted to specific authorized locations or networks. Whatever the requirement, the usage of these services is captured, giving administrators the visibility required to ensure they are utilized appropriately and securely.

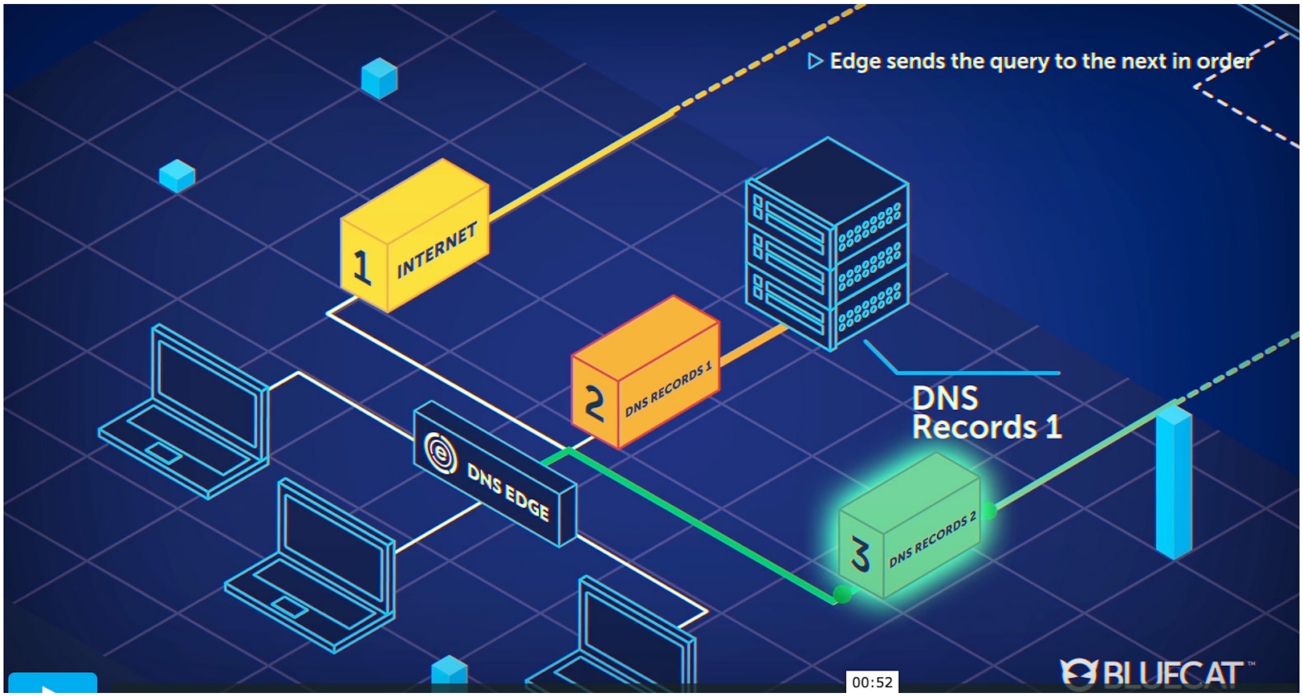

Please see our video for more information on how BlueCat's Intelligent Forwarding works.

Enabling Cloud Security Through DNS Query Logging and Policy Management

BlueCat creates a consistent security posture across the enterprise by leveraging the information flowing through DNS infrastructure. It does this by managing DNS right at the client level, logging and applying security policies to DNS queries at the "first hop". This provides the baseline visibility that security and network teams need to implement needed controls in the cloud and on-prem.

BlueCat's DNS security policies reduce attack surface by blocking malicious or inappropriate queries at the source. BlueCat also uses DNS security policies to prevent lateral movement between clouds, underneath the external filters and firewalls that many advanced persistent threats and malicious insiders seek to avoid. The policies applied to DNS can vary according to the threat – security teams can monitor, redirect, or block queries based on how the threat should be treated.

BlueCat allows security teams to triangulate threat data against a source IP to quickly identify the origin of cloud-based threats. BlueCat provides the detailed DNS logs and query data security teams need to identify patterns and anomalies that are the first indicators of compromise. For example, BlueCat can identify DNS tunneling, which could be hiding data exfiltration. In addition, these DNS logs can be easily passed to leading SIEM solutions and data analysis tools for further analysis and remediation.

To extend the scope of security and control, BlueCat offers a powerful integration with Cisco Umbrella. The integration gives customers answers to the most critical questions in network security – who, what, when and why – with a non-intrusive deployment framework.

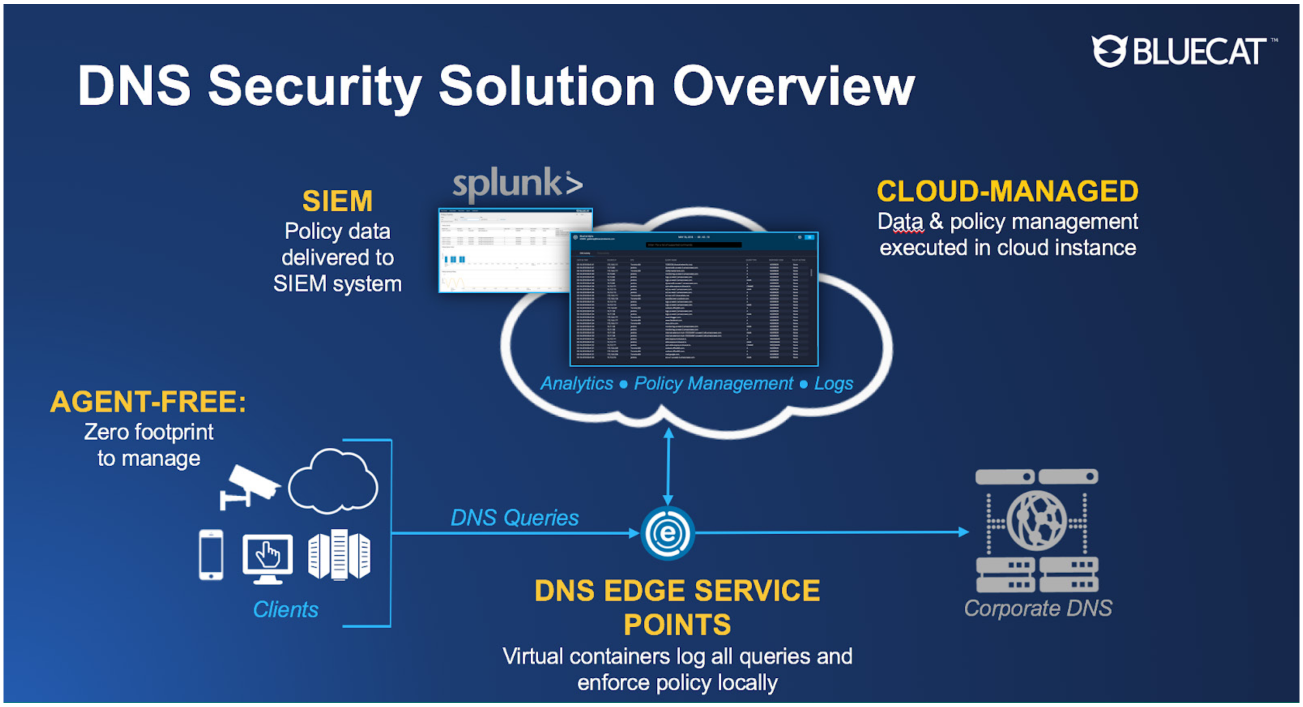

DNS Security Solution Overview

BLUECAT™

SIEM
Policy data
delivered to
SIEM system

CLOUD-MANAGED
Data & policy management
executed in cloud instance

AGENT-FREE:
Zero footprint
to manage

DNS EDGE SERVICE POINTS
Virtual containers log all queries and
enforce policy locally

Please see our DNS Edge video for more information on how BlueCat's DNS security application works at the first hop.

Conclusion

The adoption of hybrid and multi-cloud architectures allows delivery of mission-critical services to internal stakeholders and customers with unprecedented speed. At the same time, it also introduces compounding complexity for networking teams struggling to keep up with rapidly changing environments. Too often, the result is network and data conflicts, errors and costly outages, or, at the very least, a poor user experience when customers cannot access the services or applications they need.

BlueCat's Adaptive DNS platform helps overcome these challenges. It allows networking, cloud and application delivery teams to manage complexity and take advantage of the obvious benefits of hybrid cloud environments. It does this by establishing a single source of truth for namespace, IP address, and DNS record information in a centralized DDI platform. BlueCat Adaptive DNS deployed in the cloud ensures network connectivity, business continuity and data security, no matter where workloads and compute reside.

You've probably got a lot of questions about how it all works.

You're in luck.

We've got all the technical detail you need right here.

Learn more about BlueCat's solutions

bluecatnetworks.com