BLUECAT™

Cloud Resolver

Solution Brief

Table of Contents

Summary:

Cloud adoption provides highly scalable infrastructure and platform solutions for enterprises that want to build agile hybrid cloud applications and services. For cloud adoption to succeed, network teams must keep pace with the rapid changes that DevOps and other stakeholders demand. With self-service automation IT initiatives and different flavors of cloud-native DNS, network teams are buried under the manual task of managing conditional forwarding rules that come with constant zone changes in the cloud.

Seventy-two percent of CloudOps and network teams have reported poor visibility in the cloud. And it's no wonder why query resolution consistency drops, especially as cloud networks become more complex — with multiple clouds, regions, and virtual networks. The compounding effect on manual forwarding rules becomes unmanageable and can grind business to a halt. Responding to cloud- resolution challenges requires automated discovery and resolution that is cloud-native, cloud-aware, and cloud-agnostic. This solution brief presents how BlueCat Cloud Resolver uniquely tames cloud DNS by simplifying zone discovery and conditional forwarding rule management to improve service delivery.

Application and Service Innovation Has Changed

The term "born in the cloud" was once reserved for companies far removed from the data center or not involved with legacy systems. But with shifts in work mobility, SaaS, and SD-WAN, enterprises of all types build cloud-first applications and services. When developing or moving services into the cloud, DevOps teams accelerate with IaaS (Infrastructure as a Service) and PaaS (Platform as a Service) using responsive, reliable, and provisioned DNS. As a result, network teams must deliver cloud network solutions with scalable routing options.

The Chore of DNS Routing

Despite many "as a Service" innovations for DevOps, core DNS routing has remained unchanged for network teams. Whether enterprises use a point-to-point VPN tunnel or reverse proxy for applications and services, manual routing changes and IP address allocation for clients are still required. As a result, the lack of DNS innovation has been traded for 3rd party workarounds to delegate network space and route queries for resolution.

These DNS workarounds, including cloud-native DNS, bring back old and complex DNS resolution challenges. The good news is that solving these challenges with agnostic and purpose-built enterprise DNS for hybrid cloud is easy. Yet CloudOps forces the adoption of cloud-native DNS, like Route 53 from AWS, to achieve a smooth service-driven cloud workflow for DevOps. Unfortunately, this decision removes network teams' visibility and control to ensure private endpoint resolution to and across hybrid cloud.

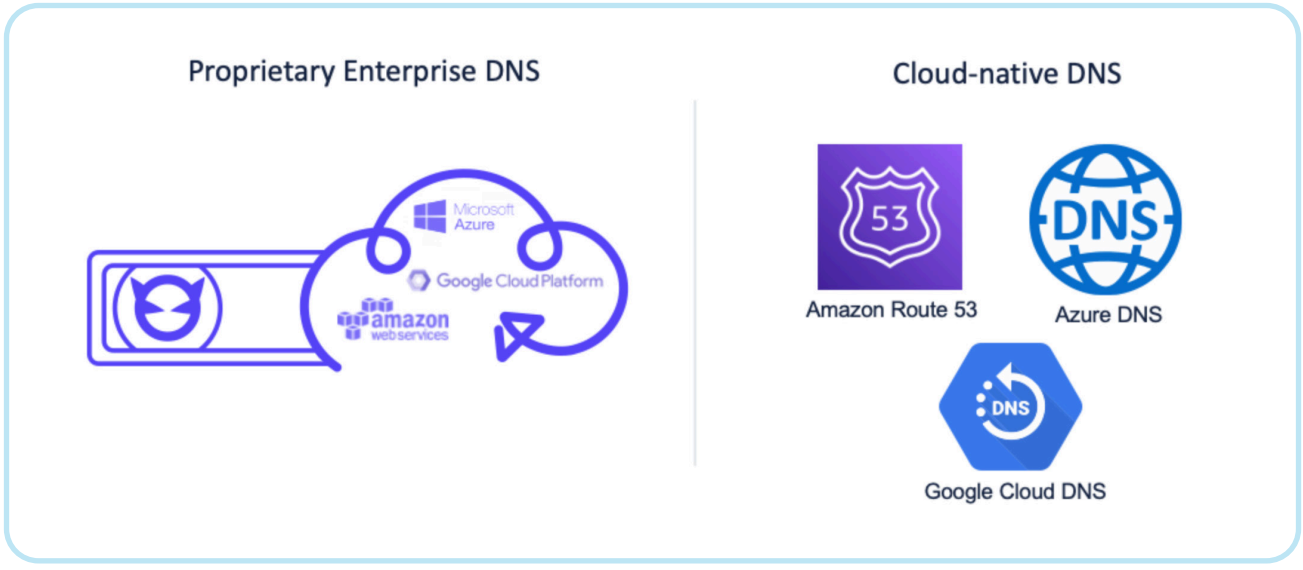

Cloud Twist On Old DNS Challenges

Before the challenges posed by the cloud, network teams had to tame the complexity of split islands of DNS with on-premises Microsoft environments. These islands posed a compounding resolution challenge as enterprises expanded legacy networks over wide and distributed data centers. Administrators would create conditional forwarding rules to confront this challenge and ensure endpoint resolution across island resources. However, they buried themselves under an avalanche of manual conditional forwarders without a unified DNS, DHCP, and IPAM environment. As a result, they could not maintain conditional forwarders with the rate of change to zones and records, which created inconsistent resolution and business disruption.

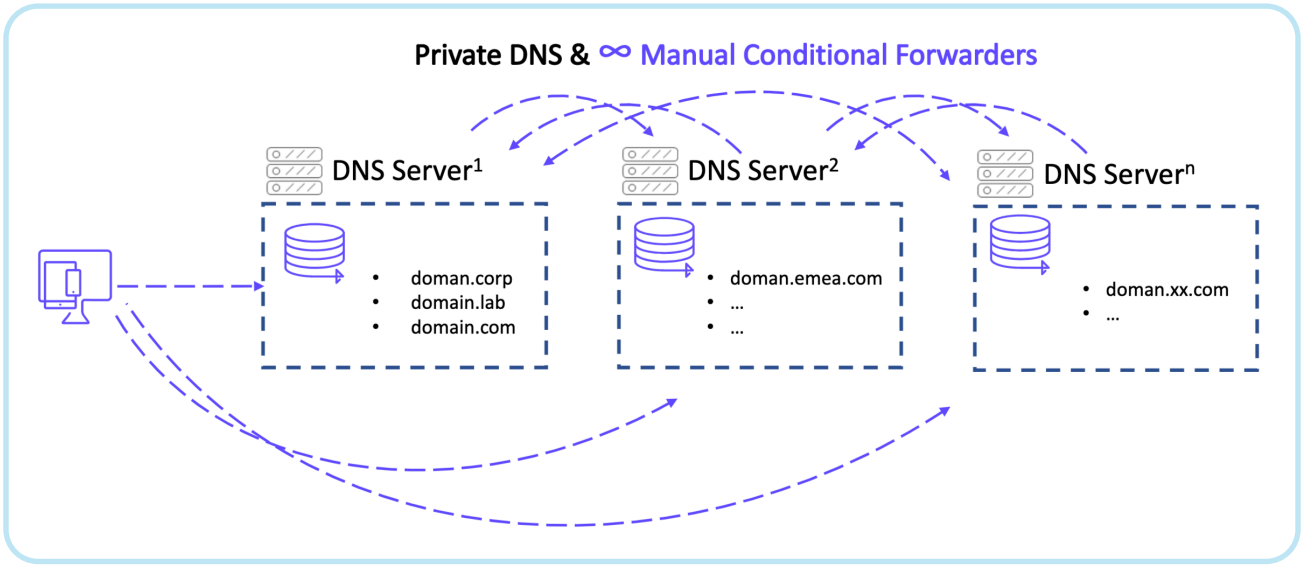

To tame the complexity of siloed DNS, network teams turned to enterprise DDI solutions, like BlueCat, that provide a single source of truth. By migrating away from legacy Microsoft environments, placing purpose-built DDI boxes in data centers, and using intelligent forwarding, they could tame complexity through improved visibility and control. Now, these teams can keep up with changes across the enterprise and confidently manage resolution to critical applications and services.

However, with the advent of the cloud and cloud-native DNS adoption, the same old on-premises challenges have resurfaced but with the following cloud twists:

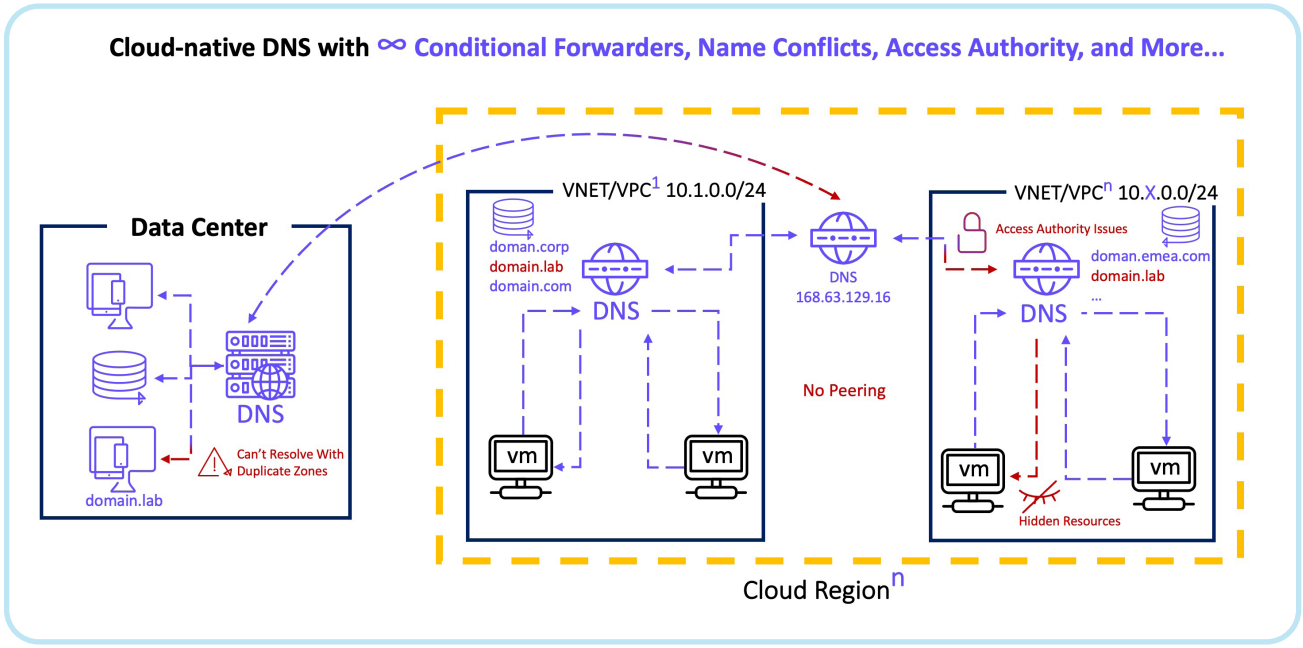

Conditional Forwarders:

With the addition of cloud, network teams must manage conditional forwarding rules at the data center and multi-cloud vendors. This compounding complexity exacerbates the manual work required to manage, track, and change forwarding rules. In addition, logging in and out of cloud tools, when enterprises need IT to move at machine speed, requires dedicated resources and quick muscle memory to navigate multi-DNS UI's, forwarding limitations, and workarounds without increasing downtime risks and SLAs.

Access Authority:

Different security schemas, like IAM for AWS, exist across cloud vendors that delegate access to users, applications, or services. Navigating access authority for private cloud endpoints can be a challenge when moving at the speed of application development. In addition, cloud mature enterprises may have multiple cloud accounts and regions that complicate resolution to different resources. Even if conditional forwarders are maintained when changes occur, query resolution will fail if endpoints don't have immediate access authority.

Zone Overlap:

Often, naming conventions are duplicated in the data center and multi-cloud environments, as some network teams prefer to have zone and resource naming consistent in hybrid environments. As a result, duplicate zone names can create conflicts when private cloud endpoints seek resolution to data center resources. In addition,

naming conflicts for zones within clouds can occur when private endpoints are trying to resolve outside or across cloud virtual networks or virtual private clouds.

Manual Reverse zones:

The lack of up-to-date reverse zones provided by cloud servers to data center authoritative servers can bring business to a griding halt. In many cases creating reverse zones would require manual bandwidth network teams don't have. Without them, security and configuration tools can't do reverse lookups needed to deliver successful and secure resolution the moment new applications and services are stood-up.

Private Cloud Endpoints:

These network interfaces have their own private IP address from a virtual private network and need to connect to services or resources widely distributed across cloud vendor regions and VPCs/VNETs. The challenge for these private cloud endpoints is knowing where new zones live or changes made to them across complex cloud architectures. Of course, the temptation is to add on more conditional forwarders that will bury networking teams with manual management they can't maintain.

The struggle is real for cloud and network teams in adopting the cloud and resolving the past challenges. Unfortunately, most organizations adopting cloud are not off to a good start, with 72% not having good visibility into changes made in cloud networks (2021 EMA House Divided Report). So how do you solve these challenges caused by cloud-native DNS without stifling application innovation for hybrid cloud? Let's uncover the answer.

Rethinking DNS Using BlueCat Cloud Resolver

Network teams using cloud-native DNS need to think differently about zone discovery, reverse zone creation, and manual conditional forwarders. Without placing proprietary virtual DNS and DHCP servers in the cloud, network teams must:

1. automate DNS zone discovery in real-time
2. provide hybrid and cross-cloud resolution with far fewer namespaces
3. work natively alongside cloud-native DNS without costly runtime.

Luckily, network teams can use BlueCat Cloud Resolver, the first cloud-native DNS resolver that provides immediate resolution to and across any private virtual network. Once placed in a region, it becomes cloud-aware, discovering all DNS zones and creating a single BlueCat Edge namespace for any endpoint in the data center or cloud to resolve queries. The solution is also cloud-agnostic, allowing network teams to embrace any combination of private cloud vendors (AWS, Azure, GCP) without getting buried in DNS conditional forwarding rules. Learn more about BlueCat Cloud Resolver in product documentation.

Use Case 1 (Private links)

Icon Legend: Conflict · Duplicate Zones · Ideal private path · Public path if private resolution fails

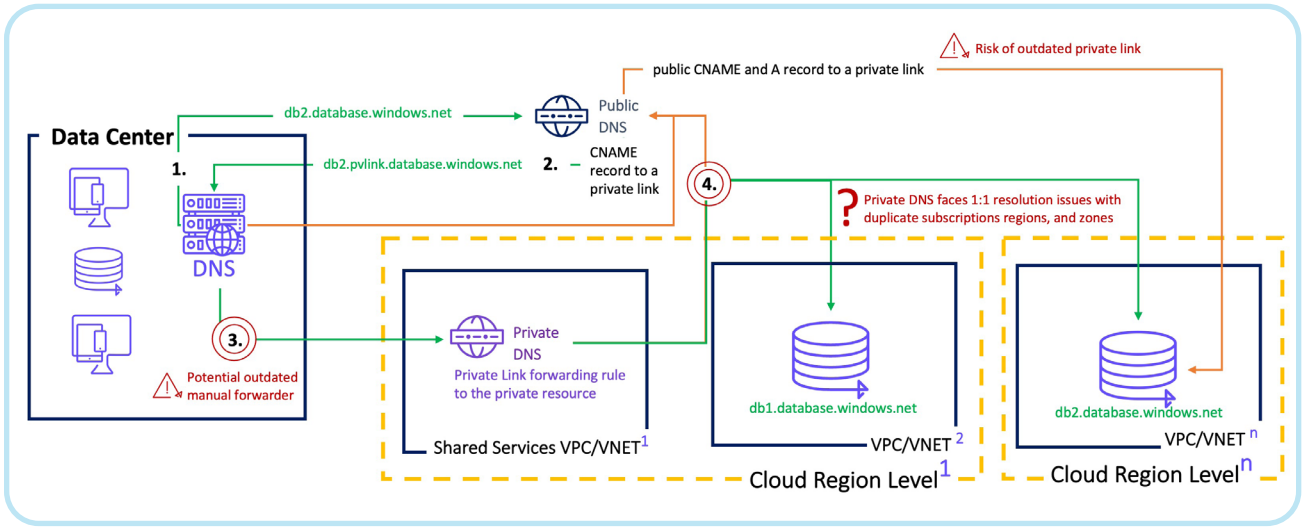

Before: *Lack of consistent resolution for private cloud endpoints and the data center*

Data center and private cloud endpoints have the same discovery and resolution challenges when dealing with private links. They both need to tie public DNS and Private DNS together. If the ideal private path fails, the public path to resolution is exclusively taken and increases network complexity. Tieing both paths is traditionally done by maintaining complex forwarding rules. Network teams also don't have control over the naming conventions of private links, increasing the risk of manual errors when configuring or updating forwarding rules. In addition, private DNS has a 1:1 resolution limitation that can't intelligently distinguish between duplicate zone names when trying to access a specific resource in a VPC/VNET.

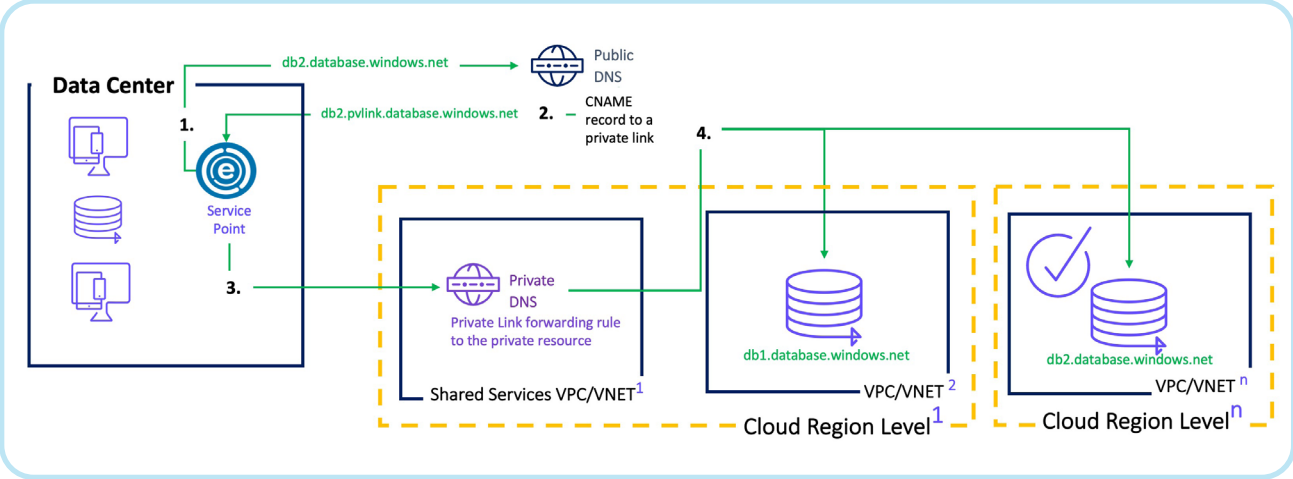

After: *Complete confidence in accessing cloud resources using private links for cloud endpoints*

BlueCat Cloud Resolver tears down discovery barriers to resolution for endpoints using private links. It can intelligently overcome routing conflicts to resources where duplicate zones exist or where VPCs/VNETs do not have peering enabled. In addition, using the cloud API, Cloud Resolver can automate resolution to any private link without falling back to a public IP addresses. That means accelerating access to critical data and services for endpoints wherever they may live.

Use case 2 (Overlapping zones and non-routable networks)

Icon Legend: ? Unknown Path Conflict

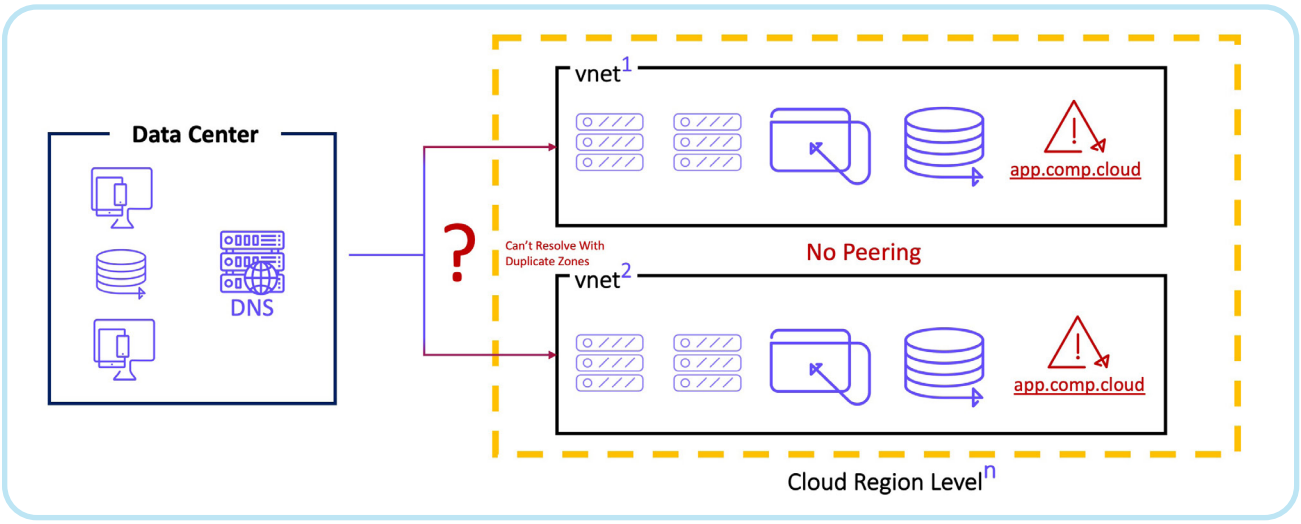

Before: *On-premises resolution fails due to a lack of visbility into cloud and zone conflicts*

As network teams take a decentralized approach to managing the cloud, as demanded by CloudOps, internal stakeholders are less likely to ensure that zones across regions and VPCs/VNETs do not have duplicate zone names. In addition, in some security circumstances, peering is not enabled for queries to be made directly to VPCs/VNETs. As a result, queries are blocked or fail to get the correct answer. Finally, native resolvers still can't get past 1:1 forwarding limitations that make it challenging to resolve overlapping zones across subscriptions, regions, and VPC/VNET architectures.

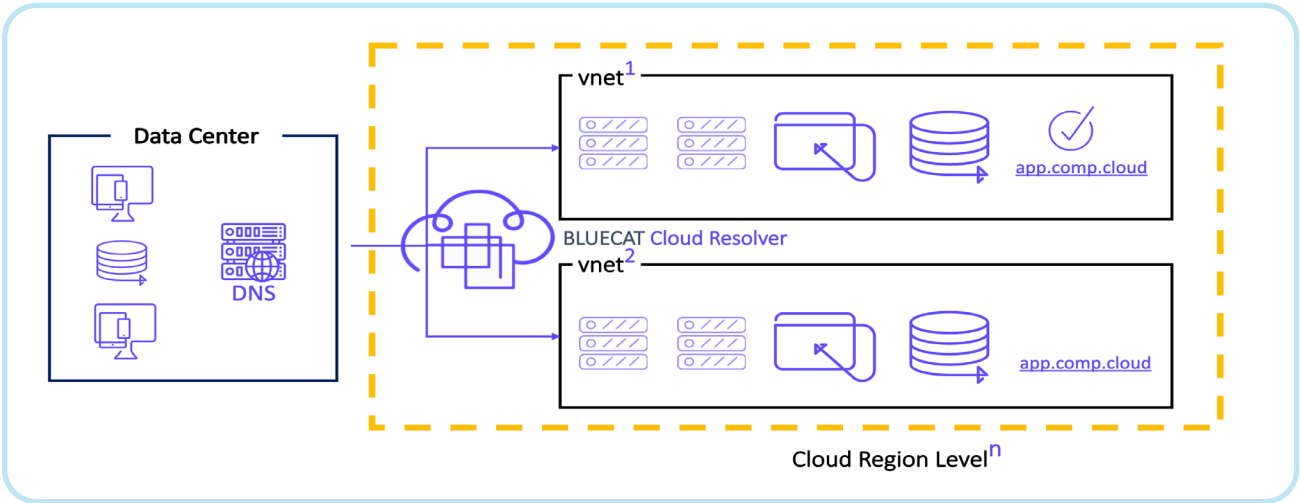

After: *Access the right zones and records regardless of overlapping names for a successful resolution*

BlueCat Cloud Resolver is not limited to 1:1 association for conditional forwarding or

disabled peering within any cloud environment. It goes beyond discovering overlapping zones across subscriptions, regions, and VPCs/VNETs to create one single resolution layer and database that can resolve to any resource in the cloud. It eliminates the need to manage conditional forwarding rules in the cloud by automating discovery using a cloud vendor API to offer the most optimal and successful resolution path.

Use case 3 (Cross zone resolution / multiple aubscriptions / regions)

Icon Legend: ? Unknown Path

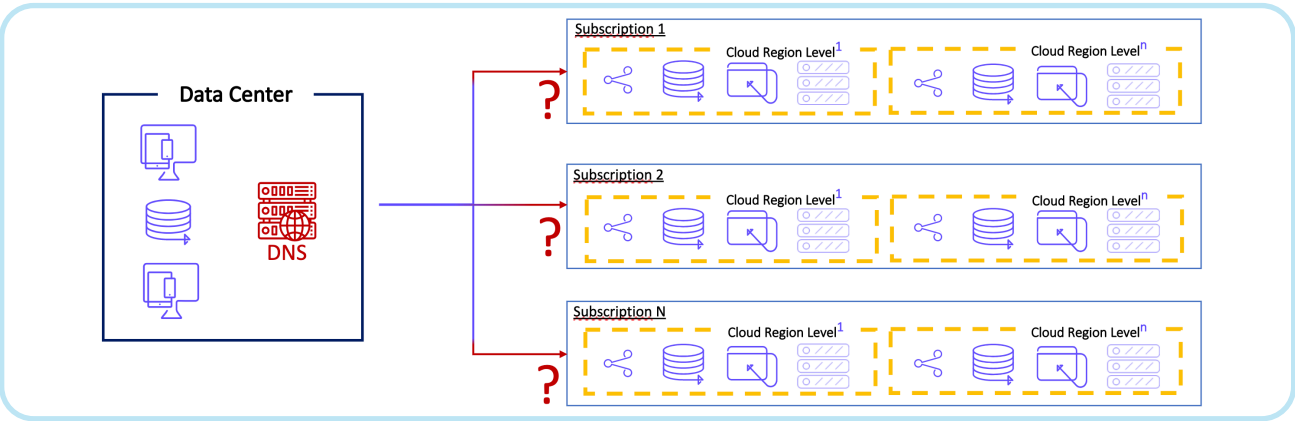

Before: *DNS resolution trips up across siloed regions and multi-cloud subscriptions*

Every enterprise manages multi-cloud or even a single-cloud vendor differently. It is common to find multiple subscriptions and tenants for different business units within a single cloud vendor or multi-cloud. This setup creates complexity at the DNS layer, where multiple conditional forwarding rules must be considered and managed. Furthermore, shared service cloud DNS still has 1:1 forwarding limitations that can cause reply errors to endpoint queries.

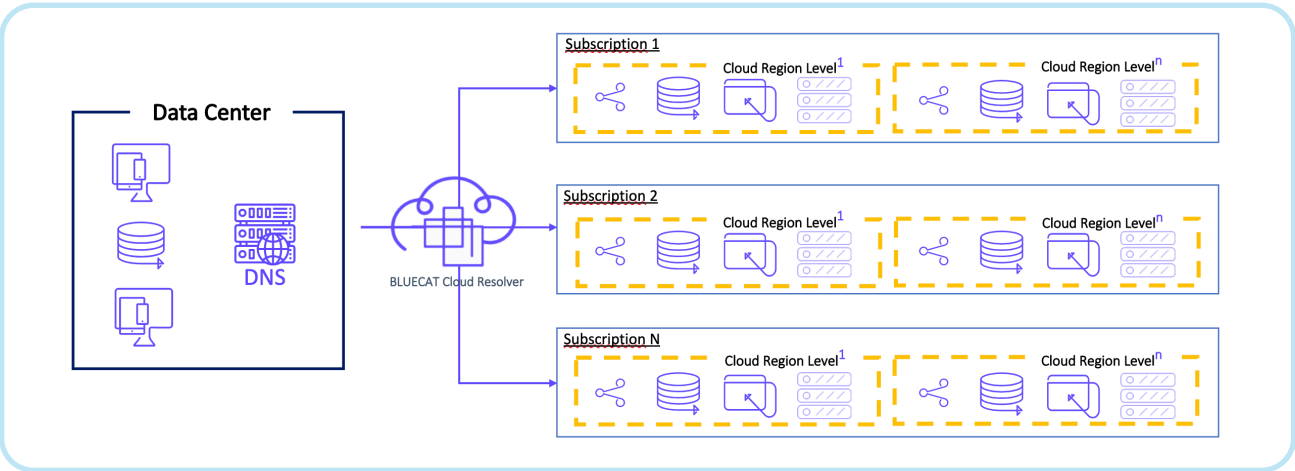

After: *Bridge the multi-cloud divide with agnostic and cloud-native DNS*

Cloud Resolver is not tied to a single subscription or region. It can resolve across multiple subscriptions, regions and cloud platforms as it's both a cloud-native and cloud-aware resolver. With one resolution layer and databse maintained near realtime through automated discovery, on premisis and cloud endpoints can get consistant answers regardless of where resources live or move.

Use case 4 (Reverse zones)

Icon Legend: 

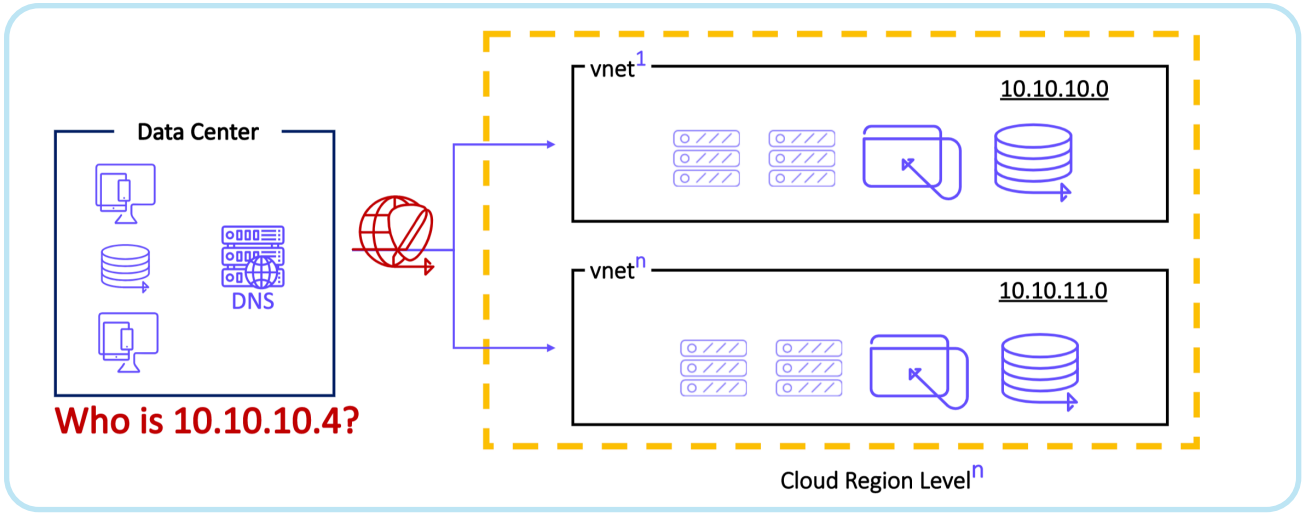

Before: *Manual creation of reverse zones slows down time to access resources and productivity*

A common issue among many cloud resolvers is answering queries from endpoints at the data center with a reverse zone. Reverse zones are critical to providing a secure connection to enterprise apps and services that require them. Unfortunately, network teams may not embrace the cloud due to internal security compliance requirements for reverse zones. To meet security requirements, admins try and create reverse zones manually, which impacts SLAs and slows down application innovation.

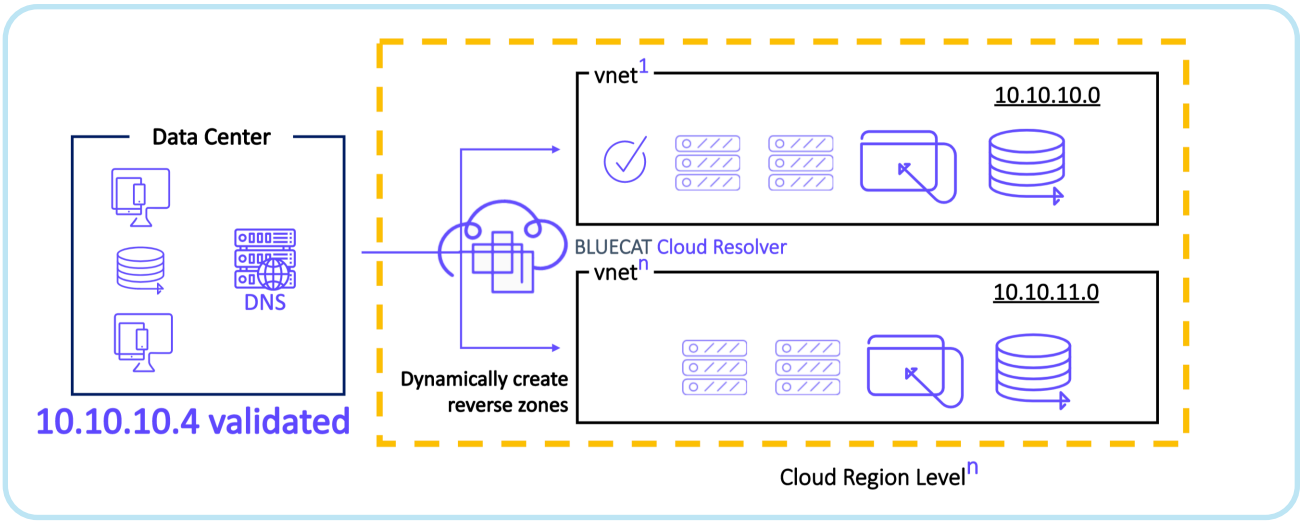

After: *Provide real-time access to access resources by eliminating manual reverse zone*

BlueCat Cloud Resolver is cloud-aware with powerful automated discovery and provides reverse zones for any query to cloud resources. That way, network teams can extend automation to reverse zone creation and accelerate application innovation without compromising security requirements.

Interacting with BlueCat Cloud Resolver

To interact with BlueCat Cloud Resolver, BlueCat provides a CI from BlueCat Edge. In addition, within BlueCat Edge, you can see all the zones and accompanying reverse zones for any cloud (AWS, Azure, GCP) and on-premises environment.

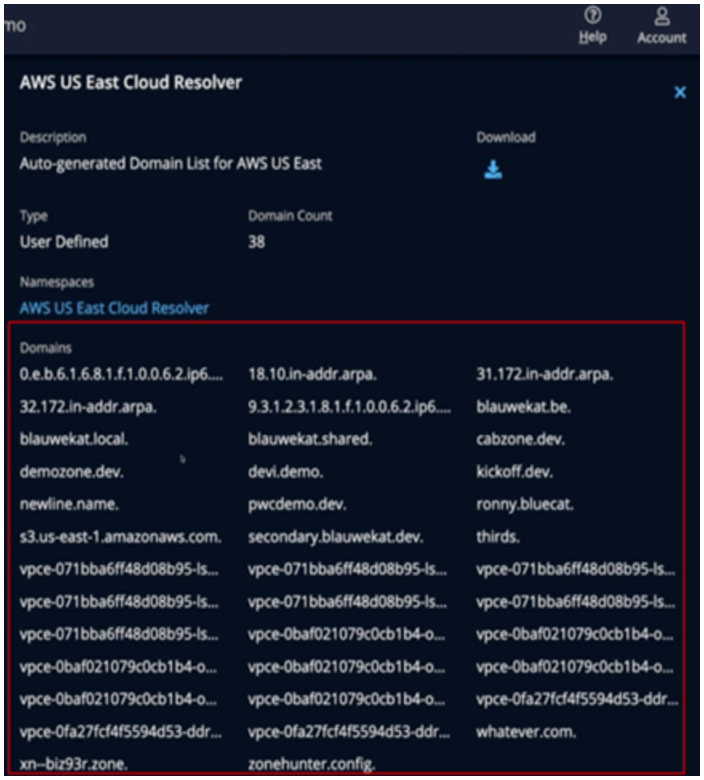

Finally, users can also see the auto-generated namespace for a single cloud region instead of managing an infinite amount of conditional forwarders for every VPC or cloud resource.

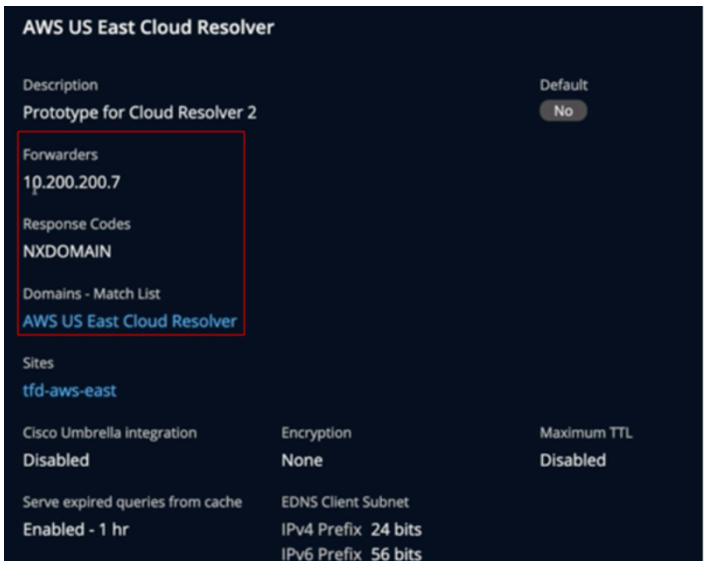

With friendly Prometheus Exposition Format generated data, network teams can easily monitor telemetry and resolution from cloud resolver.

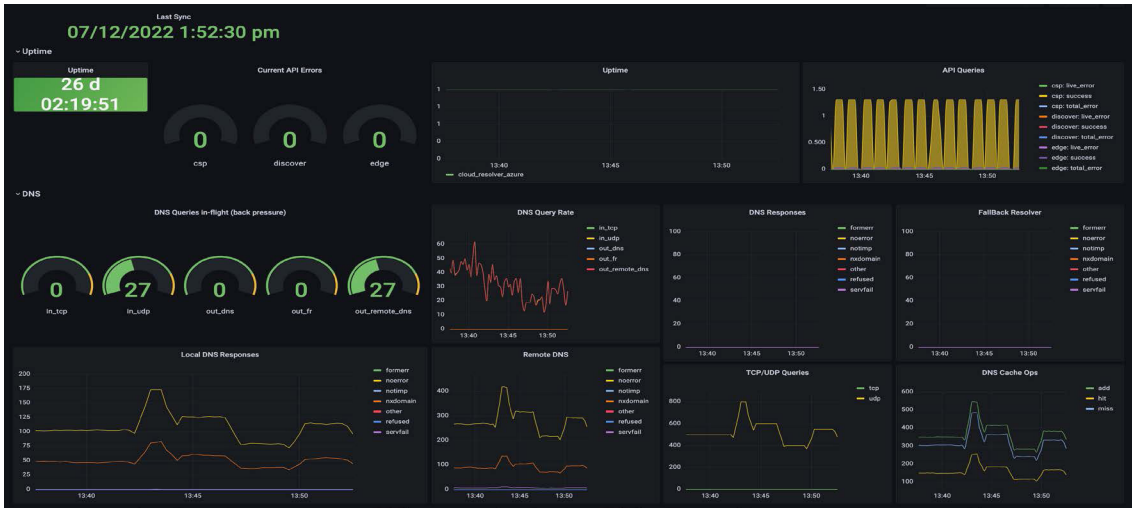

What is BlueCat Edge?

BlueCat Edge is BlueCat's intelligent DNS resolver and caching layer that leverages existing DNS infrastructure to provide unprecedented visibility and control over DNS traffic. In addition, it intelligently manages DNS forwarding rules governing the hybrid cloud to quickly deliver SaaS-based services while monitoring IP addresses to protect the enterprise against cyber attacks. For network teams adopting hybrid cloud, the art of the possible can be realized with Intelligent forwarding and is only limited by architecting creativity.

BlueCat DNS Edge also gives cybersecurity, and network teams shared visibility and control over internal and external DNS traffic. Through a single platform, enterprises can mitigate and eliminate the ways attackers can exploit DNS, detect and block cyberattacks, and investigate threats such as domain name generation algorithms and tunneling/data exfiltration incidents. In addition, using the combination of BlueCat's advanced Threat Protection feeds and policy-based networking/security management, enterprises can simplify DNS operations, tighten security, and improve network performance in ways traditional DNS solutions alone cannot achieve.

Namespaces in BlueCat Edge

To ensure both high performance and a localized experience, network teams must architect optimal and creative DNS resolution routes. The magic behind BlueCat's intelligent DNS forwarding is a namespace, a group of one or more DNS forwarders, and can optionally include rules to match the source (IP address or network range) or destination (DNS zones).

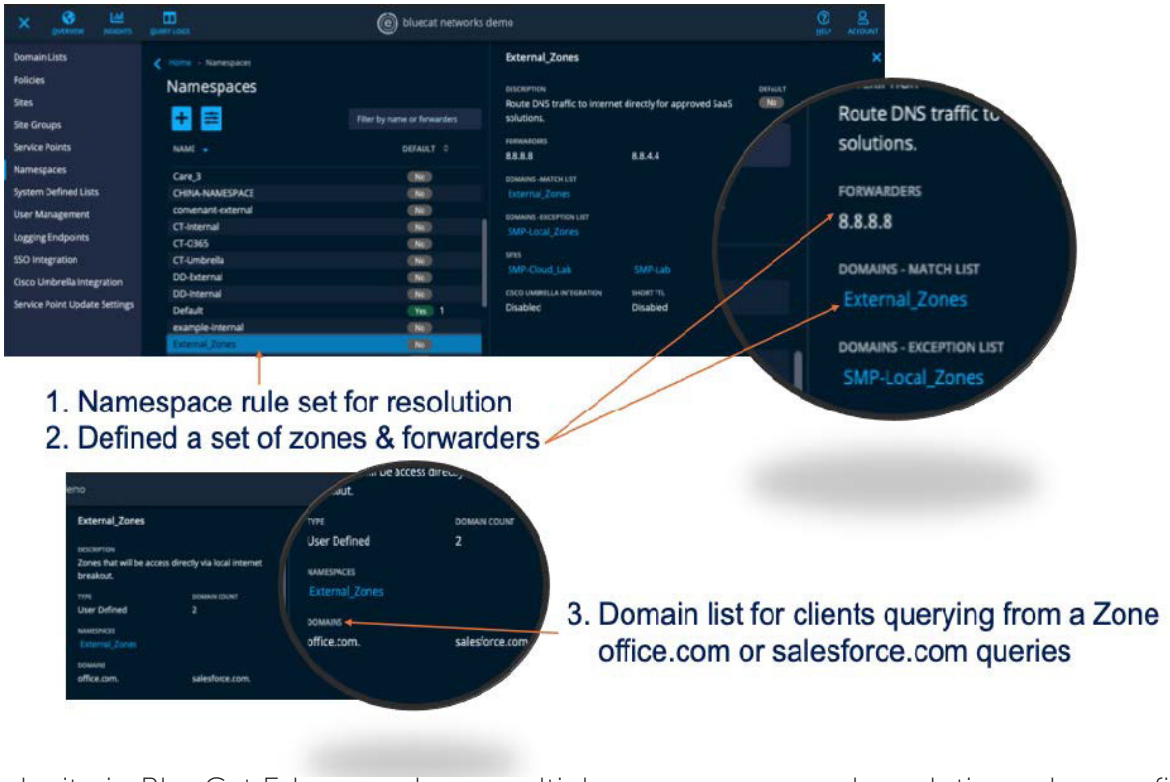

Each site in BlueCat Edge may have multiple namespaces and resolution rules configured that simplify DNS redundancies, overlapping zones, or complex forwarding rules.

Looking for better DDI tools & expertise to optimize your cloud investment?

You're in luck.

We've got what you need.

Contact us to discuss your cloud challenges today

bluecatnetworks.com