BLUECAT™

Solve DNS challenges in the cloud with **BlueCat Cloud Resolver**

Contents

Introduction

Cloud adoption provides highly scalable infrastructure and platform solutions for enterprises that want to build agile applications and services. For cloud adoption to succeed, network teams must keep pace with the rapid changes that DevOps teams 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.

According to research firm Enterprise Management Associates, only 28% of both cloud and networking professionals believe that they have very good visibility into changes made in cloud networks (*A House Divided: Dysfunctional Relationships Between Network and Cloud Teams Put Cloud Strategies at Risk, April 2021*). 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 white paper discusses how BlueCat Edge customers can use BlueCat Cloud Resolver to uniquely tame cloud DNS by simplifying zone discovery and conditional forwarding rule management to improve service delivery.

Even with cloud, the chore of DNS routing remains

The term "born in the cloud" was once reserved for organizations far removed from the data center or not involved with legacy systems. But with shifts in work mobility, software as a service (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 infrastructure as a service and platform as a service using responsive, reliable, and provisioned DNS. As a result, network teams must deliver cloud network solutions with scalable routing options.

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 third-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 solutions for hybrid cloud is easy. Yet cloud teams force the adoption of cloud-native DNS, like Route 53 from Amazon Web Services (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 environments.

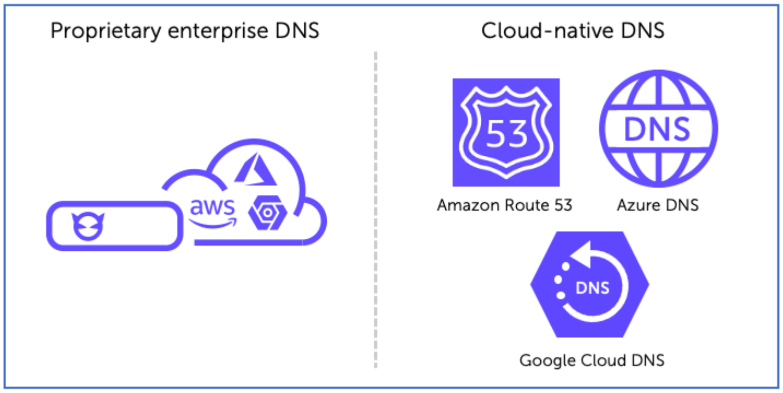

A cloud twist on old DNS challenges

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 IP address management (DDI) 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.

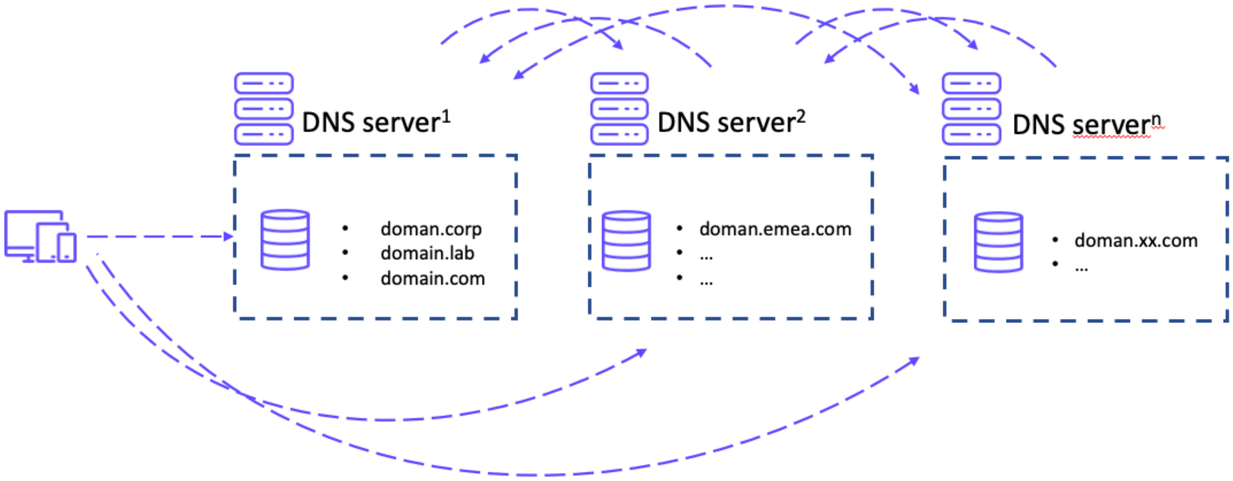

To tame the complexity of siloed DNS, network teams turned to enterprise DDI solutions like BlueCat that provide a single source of truth. By migrating away from legacy Microsoft environments, placing purpose-built DDI boxes in data centers, and using intelligent forwarding, they could tame complexity through improved visibility and control. With BlueCat's integrated DDI platform, network 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 DNS challenges have resurfaced but with a few cloud twists.

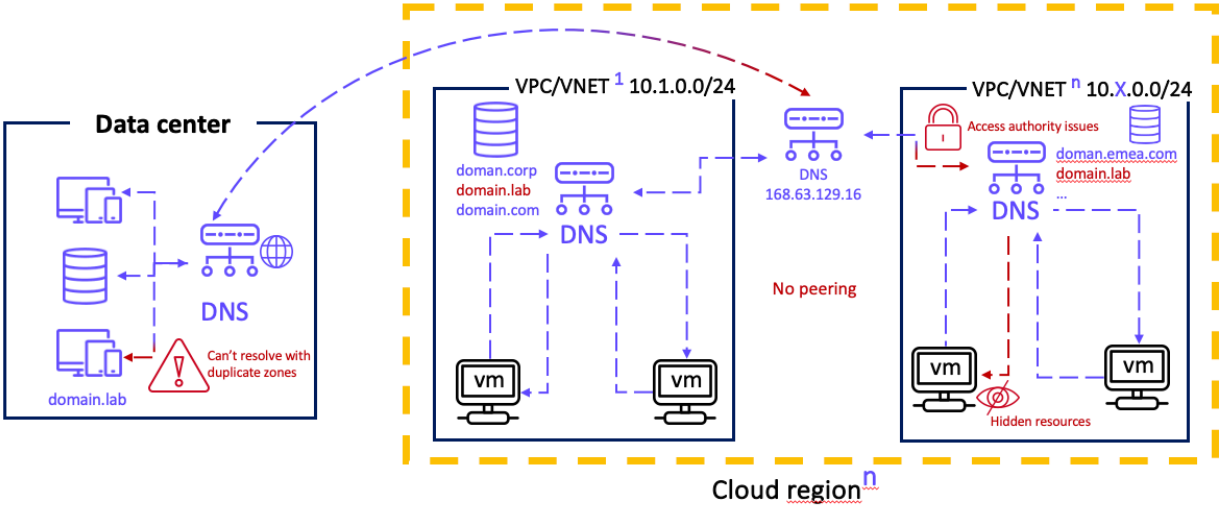

Conditional forwarders

Network teams must manage conditional forwarding rules for both the data center and multiple cloud service providers. This compounding complexity exacerbates the degree of 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 multiple DNS user interfaces, forwarding limitations, and workarounds. And all of it must be done without increasing the risk of downtime or impacting service-level agreements.

Access authority

Different security schemas, like Identity and Access Management for AWS, allow public cloud service providers to 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. Furthermore, enterprises may have multiple cloud accounts and regions that complicate resolution to different resources. And 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 cloud environments, as some network teams prefer to have zone and resource naming consistent in hybrid environments. However, 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 grinding halt. In many cases, creating reverse zones would require manual bandwidth that network teams don't have. But 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. But they need to connect to services or resources widely distributed across cloud service provider regions and AWS Virtual Private Cloud (VPC) or Microsoft Azure Virtual Network (VNet). The challenge for these private cloud endpoints is knowing where new zones live or when changes are made to them across complex cloud architectures. Of course, the temptation is to add more conditional forwarders that will bury networking teams with manual management that they can't maintain.

When it comes to cloud adoption, the struggle is real for cloud and network teams that are trying to resolve DNS challenges. Most organizations adopting cloud are not off to a good start, with 72% of both cloud and networking professionals believing that they do not have very good visibility into changes made in cloud networks.

How do you solve these challenges caused by cloud-native DNS without stifling application innovation for hybrid cloud environments?

Use cases: Rethinking DNS with 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:

- automate DNS zone discovery in real-time;
- provide hybrid and cross-cloud resolution with far fewer namespaces; and
- 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.

Note: To use Cloud Resolver, you must have a subscription to AWS, Azure, or Google Cloud, as well as a BlueCat Edge environment.

Once placed in a region, Cloud Resolver 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. Cloud Resolver is also cloud-agnostic, allowing network teams to embrace any combination of private cloud service providers (AWS, Azure, or Google Cloud) without getting buried in DNS conditional forwarding rules.

Learn more in the use cases below and in the BlueCat Cloud Resolver Administration Guide.

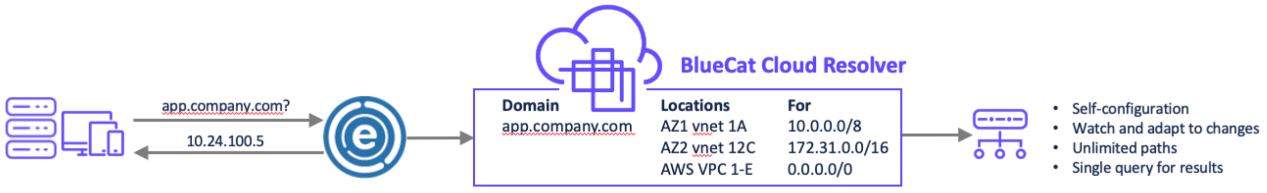

Domain	Locations	For
app.company.com	AZ1 vnet 1A	10.0.0.0/8
	AZ2 vnet 12C	172.31.0.0/16
	AWS VPC 1-E	0.0.0.0/0

What is BlueCat Edge?

BlueCat Edge is our intelligent DNS resolver and caching layer that provides unprecedented visibility and control over DNS traffic. You can quickly and easily deploy Edge in any hybrid cloud environment. As the first hop of any DNS query, Edge intelligently directs DNS traffic, tames conditional forwarding rules, blocks malicious DNS queries, and helps monitor and collect all DNS query and response data for diagnostics and investigations.

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, which 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.

Use case 1: Private links

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. Tying 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 or VNet.

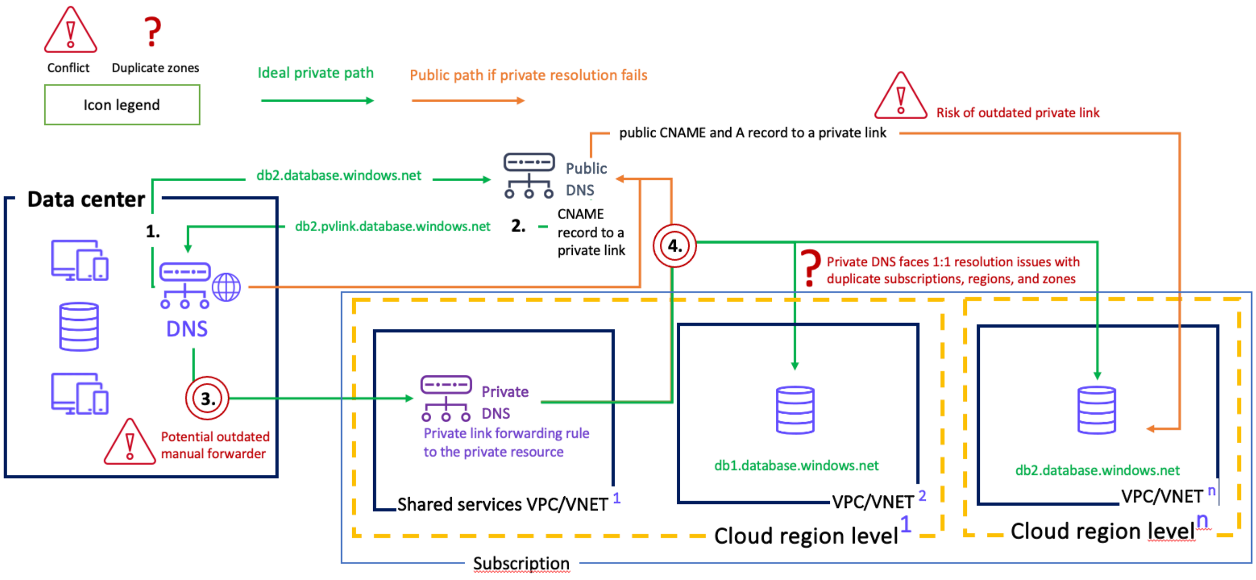
Conflict
Duplicate zones
Icon legend
Ideal private path
Public path if private resolution fails
Risk of outdated private link
public CNAME and A record to a private link
db2.database.windows.net
Public DNS
Data center
1.
db2.pvlink.database.windows.net
2.
CNAME record to a private link
4.
Private DNS faces 1:1 resolution issues with duplicate subscriptions, regions, and zones
DNS
3.
Potential outdated manual forwarder
Private DNS
Private link forwarding rule to the private resource
Shared services VPC/VNET 1
db1.database.windows.net
VPC/VNET 2
Cloud region level 1
db2.database.windows.net
VPC/VNET n
Cloud region level n
Subscription

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 or 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 address. The result is accelerated access to critical data and services for endpoints wherever they may live.

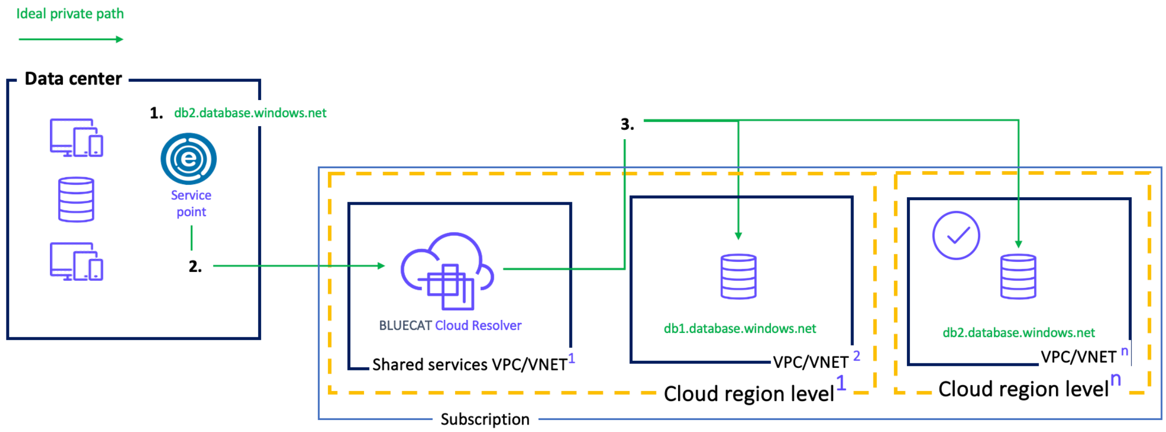

Use case 2: Overlapping zones and non-routable networks

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

As network teams take a decentralized approach to managing the cloud as demanded by cloud teams, internal stakeholders are less likely to ensure that zones across regions and VPCs or VNets do not have duplicate zone names. In addition, in some security scenarios, peering is not enabled for directly making queries to VPCs or 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 or VNet architectures.

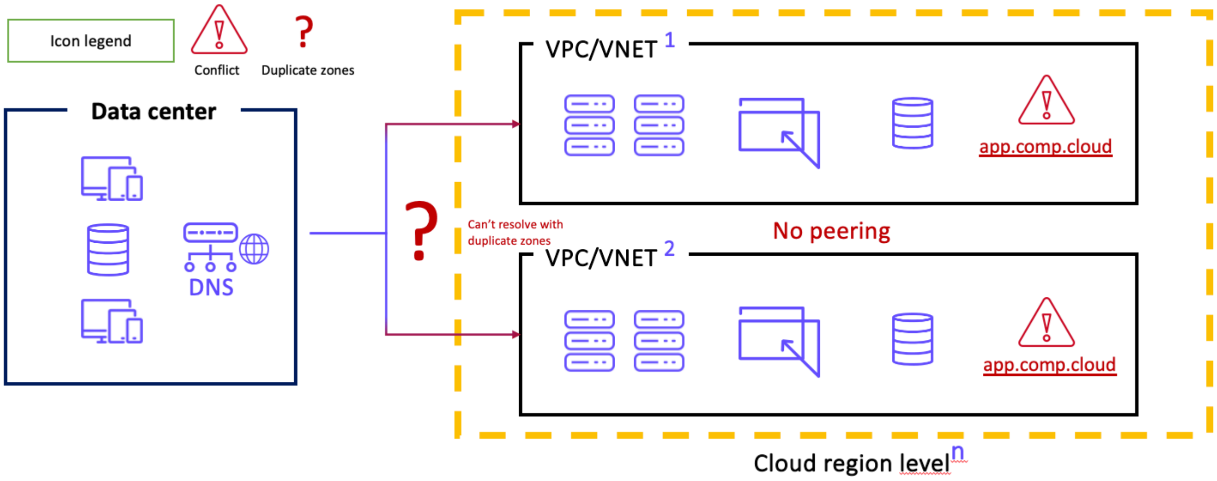

After: Access the right zones and records regardless of overlapping names for a successful resolution

BlueCat Cloud Resolver is not limited to a 1:1 association for conditional forwarding or disabled peering within any cloud environment. It goes beyond discovering overlapping zones across subscriptions, regions, and VPCs or 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.

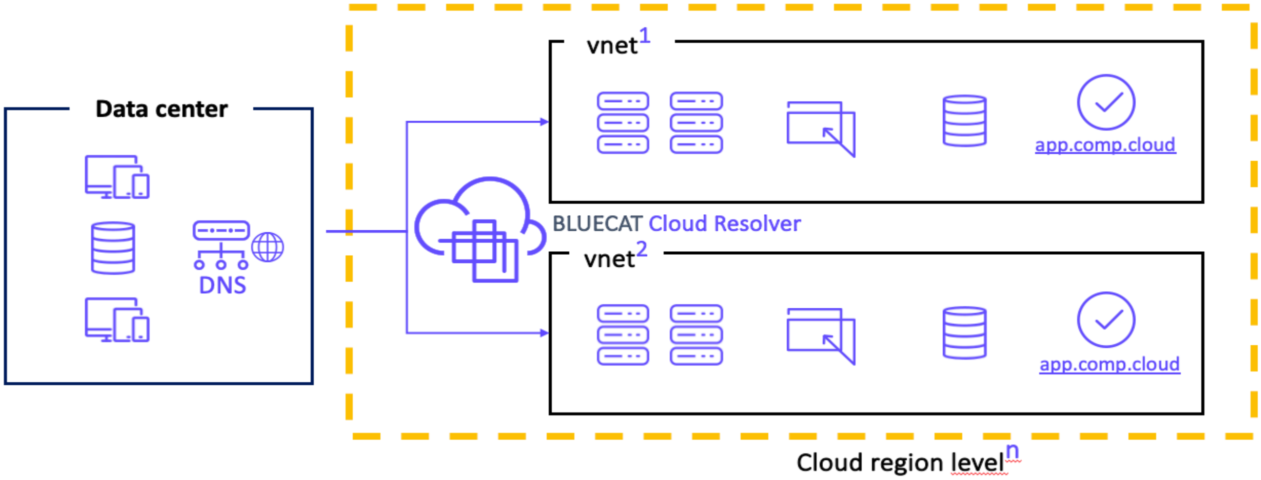

Use case 3: Cross-zone resolution for multiple subscriptions and regions

Before: DNS resolution trips up across siloed regions and multiple cloud subscriptions

Every enterprise manages multiple cloud service providers—or even a single provider—differently. It is common to find multiple subscriptions and tenants for different business units within a single or multiple cloud service provider(s). 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.

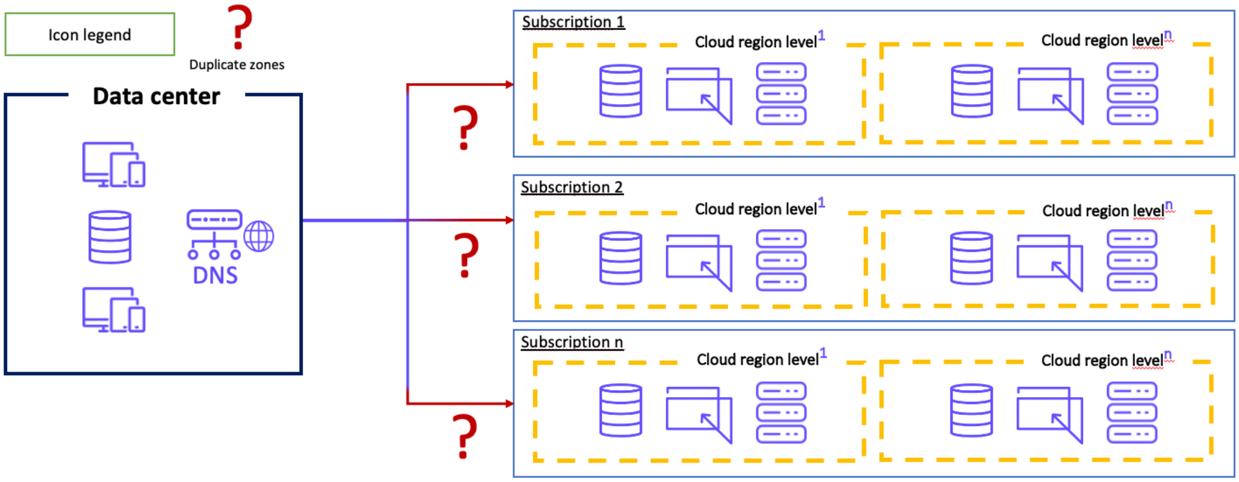

After: Bridge the multicloud divide with agnostic and cloud-native DNS

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 database maintained near real-time through automated discovery, on-premises and cloud endpoints can get consistent answers regardless of where resources live or move.

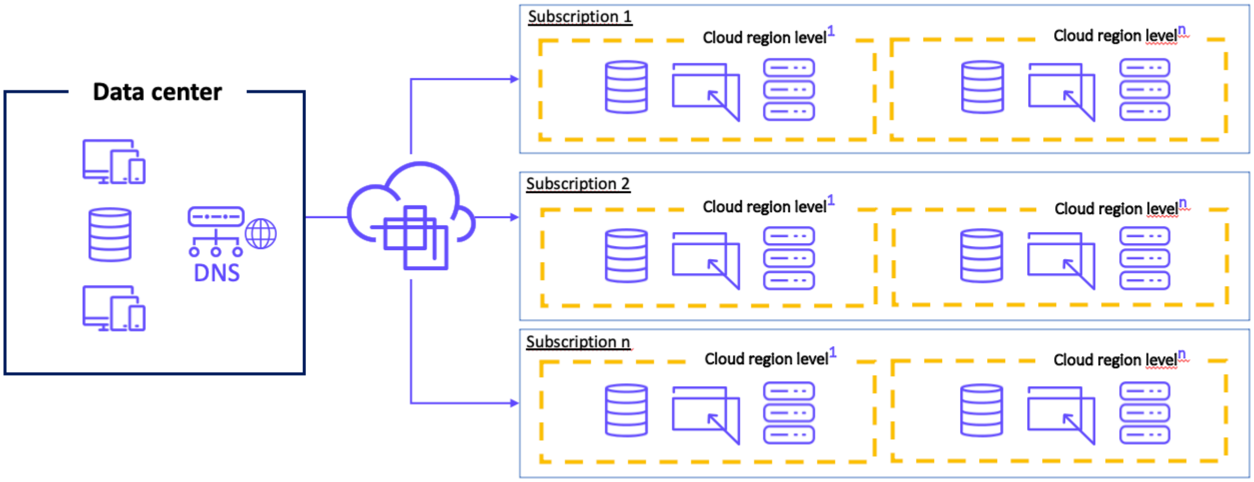

Use case 4: Reverse zones

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 applications 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 to create reverse zones manually, which impacts service-level agreements and slows down application innovation.

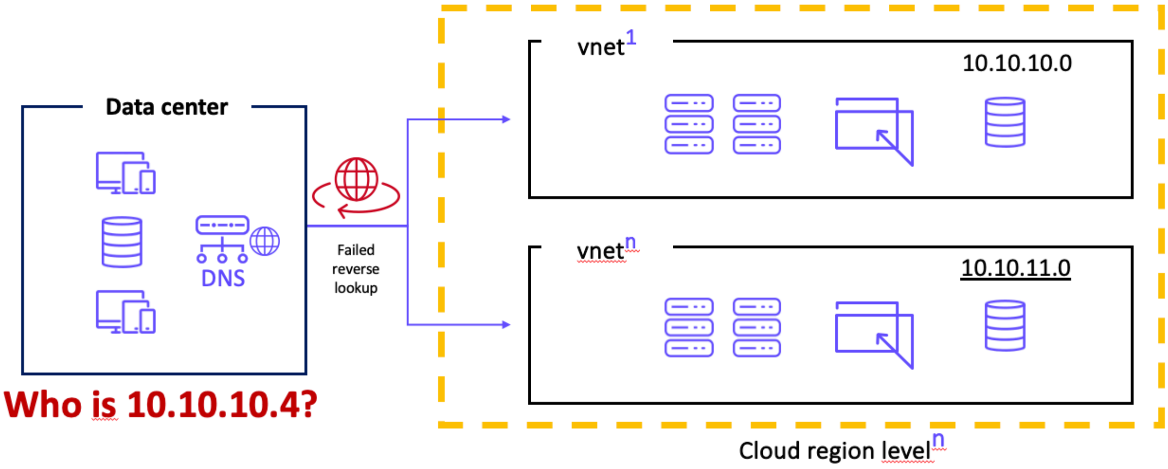

After: Provide real-time access to access resources by eliminating manual reverse zones

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

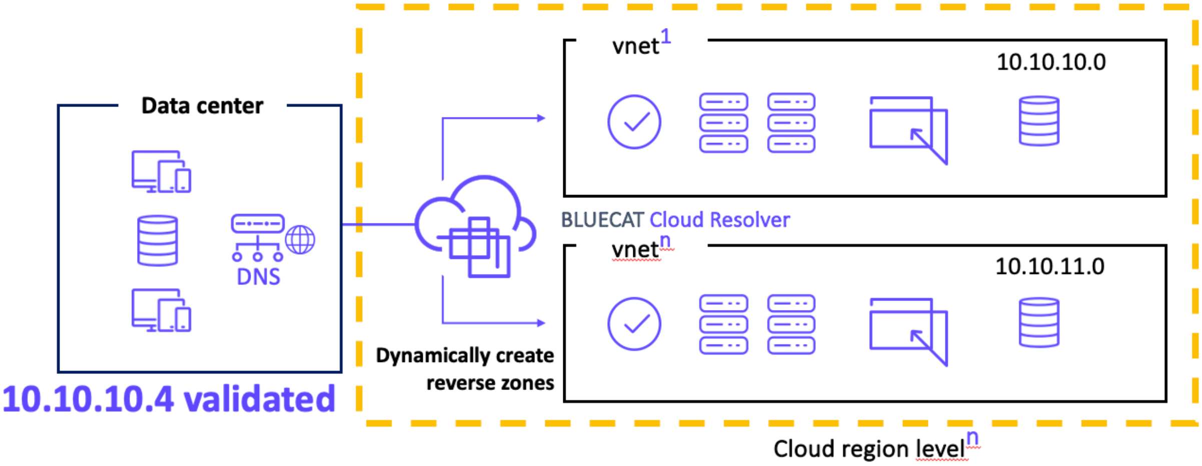

BlueCat Cloud Resolver visibility benefits

Using the BlueCat Edge user interface, Cloud Resolver allows you to view all the DNS zones and accompanying reverse zones for any cloud (AWS, Azure, or Google Cloud) or on-premises environment.

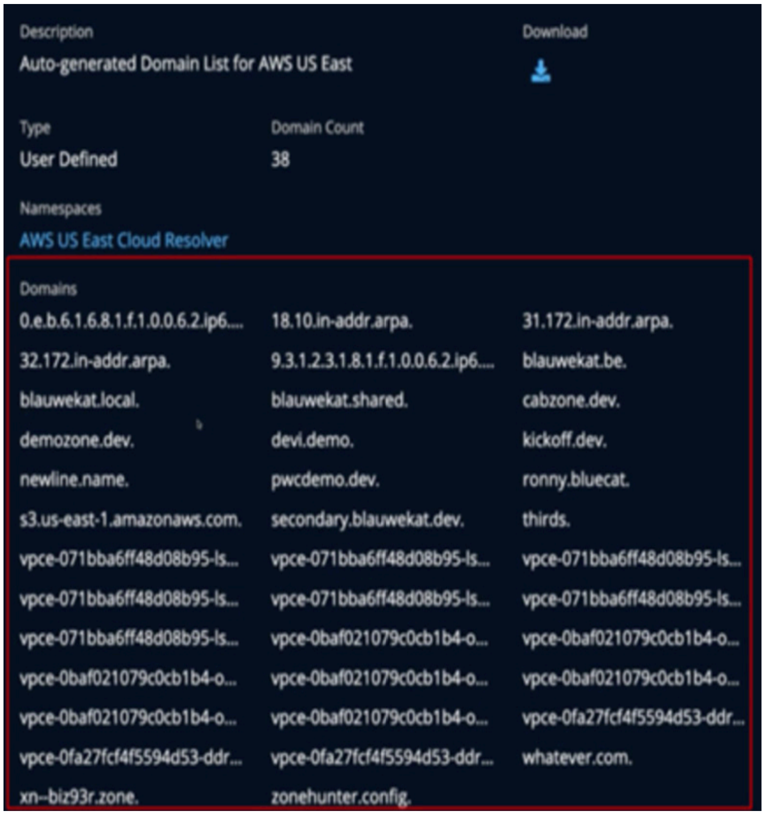

Description		Download
Auto-generated Domain List for AWS US East		

Type	Domain Count
User Defined	38

Namespaces

AWS US East Cloud Resolver

Domains

0.e.b.6.1.6.8.1.f.1.0.0.6.2.ip6....	18.10.in-addr.arpa.	31.172.in-addr.arpa.
32.172.in-addr.arpa.	9.3.1.2.3.1.8.1.f.1.0.0.6.2.ip6....	blauwekat.be.
blauwekat.local.	blauwekat.shared.	cabzone.dev.
demozone.dev.	devi.demo.	kickoff.dev.
newline.name.	pwcdemo.dev.	ronny.bluecat.
s3.us-east-1.amazonaws.com.	secondary.blauwekat.dev.	thirds.
vpce-071bba6ff48d08b95-ls...	vpce-071bba6ff48d08b95-ls...	vpce-071bba6ff48d08b95-ls...
vpce-071bba6ff48d08b95-ls...	vpce-071bba6ff48d08b95-ls...	vpce-071bba6ff48d08b95-ls...
vpce-071bba6ff48d08b95-ls...	vpce-0baf021079c0cb1b4-o...	vpce-0baf021079c0cb1b4-o...
vpce-0baf021079c0cb1b4-o...	vpce-0baf021079c0cb1b4-o...	vpce-0baf021079c0cb1b4-o...
vpce-0baf021079c0cb1b4-o...	vpce-0baf021079c0cb1b4-o...	vpce-0fa27fcf4f5594d53-ddr...
vpce-0fa27fcf4f5594d53-ddr...	vpce-0fa27fcf4f5594d53-ddr...	whatever.com.
xn--biz93r.zone.	zonehunter.config.	

You can also see the auto-generated namespace for a single cloud region instead of managing an infinite number of conditional forwarders for every virtual private cloud or cloud resource.

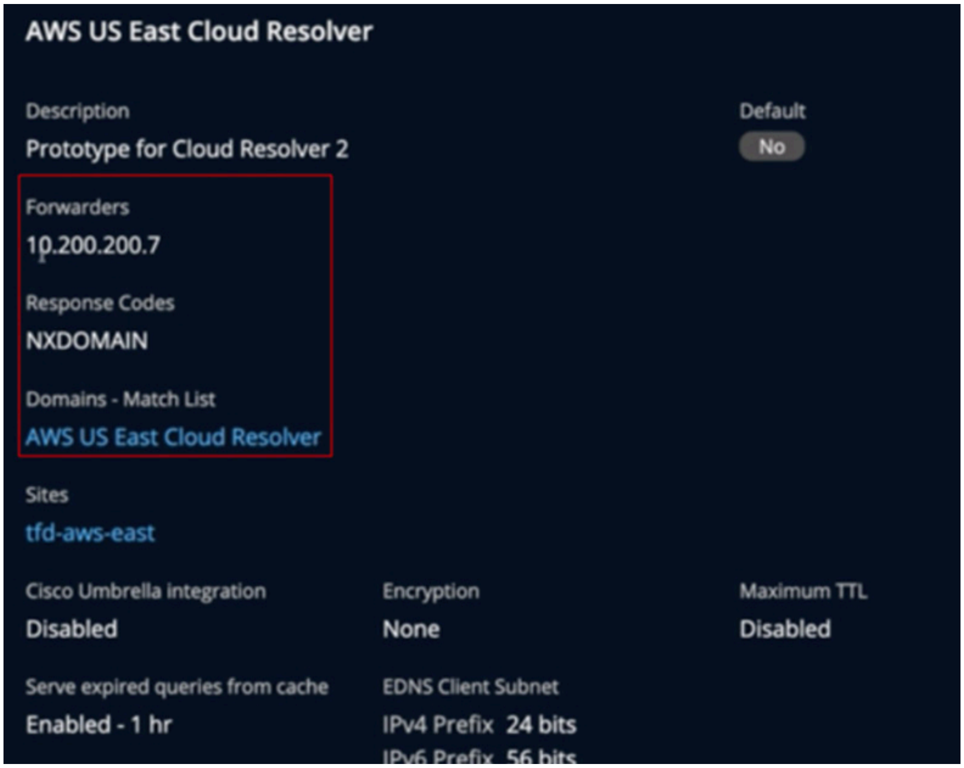

AWS US East Cloud Resolver

Description
Prototype for Cloud Resolver 2

Default
No

Forwarders
10.200.200.7

Response Codes
NXDOMAIN

Domains - Match List
AWS US East Cloud Resolver

Sites
tfd-aws-east

Cisco Umbrella integration
Disabled

Encryption
None

Maximum TTL
Disabled

Serve expired queries from cache
Enabled - 1 hr

EDNS Client Subnet
IPv4 Prefix 24 bits
IPv6 Prefix 56 bits

Looking for better DDI tools and 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