Drive rapid change to innovate faster

NetOps 2.0. SASE. Cloud-first. Cloud smart. Everyone seems to have a different buzzword for the “network of the future”.

All of these concepts essentially describe the same thing: a highly automated, fully integrated, cloud-driven network that promotes innovation throughout the enterprise.

These networks are designed for speed. Cloud architectures, SD-WAN, and intent-based networking deliver superior network performance that DevOps and cloud teams need to deliver business-critical applications and services quickly.

This is all great. In theory.

The reality: Most networks can't deliver this ideal future state.

Why is delivery of this ideal network so hard?

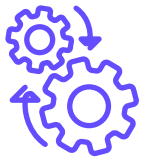

Everything's manual Automation is only as fast as the slowest process in your network stack. When your network team is manually provisioning IP addresses using spreadsheets or configuring the network one server at a time, they simply can't deliver at the pace of a DevOps or cloud team, let alone the automated processes that support them.

Nothing's integrated Running applications and services across hybrid environments is a mess if they can't draw on an API-first, automation-friendly network infrastructure. The same goes for SD-WAN, intent-based networking, and virtualization tools - to operate well, they need to draw from an integrated foundation of core network services.

Shadow IT If DevOps and cloud teams don't see what they need, they will probably build it on their own. That leads to an even more fractured, decentralized enterprise where nobody on the network or security teams can even see what's going on, let alone manage day-to-day operations.

Fragile networks Custom-built core network systems simply aren't designed to scale up and down at the speed of DevOps. They can't handle the demands of higher-level strategic initiatives like cloud, virtualization, and automation.

All of these roadblocks share a common thread: **core network infrastructure.**

Core network infrastructure is usually the last thing that IT leaders think about when they start a digital transformation project. Their eye is on the end goal - that sexy, innovative network that drives business outcomes. It's only after they're far down the road of digital transformation that they realize the key role that core network infrastructure plays.

How do IT leaders come to realize that their core network infrastructure is slowing down the pace of change? Here are a few examples:

DevOps teams wait for days for IP address provisioning, causing them to put development on hold until the network team can get around to fulfilling the service ticket.

The network crashes frequently when shadow IT results in spiraling network conflicts and misconfigurations go unaddressed.

The network team can't meet their SLAs, as they fall behind the sudden increase in service ticket volume.

Thankfully, none of this is inevitable. Core network infrastructure doesn't have to be the weakest link in your digital transformation project.

Core network infrastructure can be the catalyst for higher-level innovation.

At BlueCat, we have a name for this dynamic, open, scalable, secure, automated system. It's called Adaptive DNS.

Our vision is simple. DNS, DHCP, and IPAM (DDI) should be the foundation of the fully integrated, automated network you desire. With that end goal in mind, we've built a fully integrated, automated way to manage DDI - one that provides reliable, feature-rich core network services at the edge, in the core, and everywhere in between.

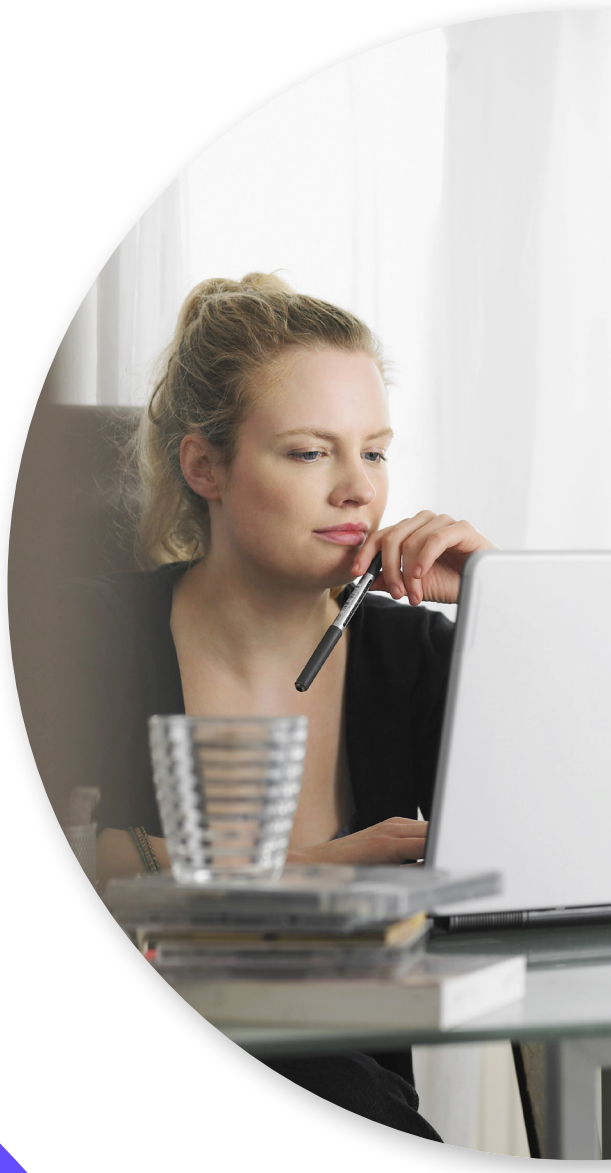

How does that vision play out at a tactical level?

Single source of truth for DDI management, enabling automated integration with cloud infrastructure, SD-WAN, virtualization engines, and your own applications.

Self-service provisioning gives DevOps and cloud teams the IP space they need, fast - preventing them from resorting to opaque shadow IT solutions.

Open APIs and zero-touch automation reduce the burden of one-off configurations and network changes.

BLUECAT™

With Adaptive DNS, everyone gets what they need to drive innovation.

Network teams eliminate the drudgery of DDI service tickets through automation. DevOps and cloud teams get the responsive, reliable core network services they need. Everyone gets an API-first, automation-friendly infrastructure that supports their vision of tomorrow's network.

Did we mention that all of this is actually attainable? You're probably wondering how it all works. That's good - we'd love to explain it all a bit more.

Learn more about BlueCat's DNS automation solutions

bluecatnetworks.com

Increase resilience of critical infrastructure

When's the last time you had a network outage?

Maybe things are so good that you can't even remember the last time. (If so, congratulations.) Or maybe your last outage happened more recently than you care to remember. Maybe that outage was only the latest in a string of incidents stretching over the past week, month, or quarter. Maybe those outages are becoming more frequent, and lasting longer.

Fragile networks are easy to spot. Here are some of the most common symptoms:

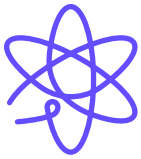

Decentralized Siloed operations are a telltale sign of a fragile network. When everyone's responsible for their own little fiefdom, they tend to charge forward without understanding the consequences. This often leads to conflicts and misconfigurations which spiral out of control.

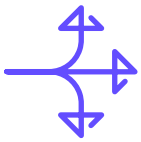

Complex On a fragile network, everything runs through a dizzying array of configurations and processes - so many, in fact, that nobody really understands how they all work together. One missed connection is all it takes to bring the whole system to a halt.

Unpredictable When nobody knows how changes will impact the network, that's a clear indicator of fragility. If you're scared to tackle strategic initiatives or roll out new systems, something's seriously wrong.

Manual If the network team is spending most of its time simply managing back-end configurations, you're probably dealing with a fragile network. When any workload change seems unmanageable, it's a sign that scaling up is out of the question.

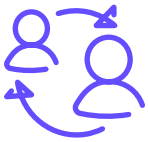

Compartmentalized Fragile networks can't operate without the knowledge of a few people (or a single person). If you're aligning change windows around vacations of certain people, that's an indicator that your network isn't resilient enough to operate without them.

The first step in dealing with network fragility is acknowledging that you have a problem.

With apologies to Hemmingway, networks become fragile gradually, then suddenly. When outages become all too frequent, it's usually the result of infrastructure problems that developed over months or years. Every IT manager reaches a point where the network they built simply can't deliver the performance everyone expects.

Maybe it's a major outage that threatens the business. Maybe it's a steady drip of incidents that suddenly become noticeable to internal users or customers.

BLUECAT™

If any of this sounds familiar, you probably know where we're going to point the finger. It's not at the network team - they're just trying to stay afloat. It's not your end-users or executives - they just want a network that delivers.

The problem is your DNS, DHCP, and IPAM (DDI).

When it functions well, you're not supposed to even notice DDI. It just blends into the background. You can take it for granted.

When your DDI is complex, your network is inherently fragile. That fragile network drives up back-end operational costs, increases risk, and prevents you from meeting basic service expectations. You've probably got compliance issues, too.

It's an even bigger problem if you're trying to move to the cloud, implementing automation, or operating a complex global enterprise. (And who isn't doing one or all of those things?)

Your DDI should be:

Dynamic DDI should provide visibility and insight into the state of users, devices and applications across hybrid environments

Open DDI systems should integrate seamlessly with hybrid cloud and other mission-critical IT systems, using open APIs and built code.

Secure DDI should act as an intelligent control plane for rapid threat detection and response on both internal and external pathways.

Scalable DDI systems should be easy to deploy whenever and wherever you want them to be.

Automated DDI should implement smart security policies, remediate threats, and analyze data at machine speed.

BLUECAT™

In a nutshell, that's BlueCat's vision for your network. We call it Adaptive DNS.

Of course, there's a lot more to it than that. We can talk your ear off about high availability, IPv6, GSS-TSIG, direct internet access, and all the other nitty-gritty technical details. But all the bits and bytes really point to one thing: DDI that you can take for granted. DDI that doesn't crash your network.

Isn't that what you want, after all?

But seriously, though, here's where we get into the weeds.

Learn more about BlueCat's DDI solutions

bluecatnetworks.com

Leverage DNS data to reduce risk

Breaches happen. It's not a question of if. It's a question of when and how.

Sure, there are ways to reduce risk upfront. Security controls, security standards, and security training can make it more difficult to penetrate your network. But nothing is foolproof.

Nobody can control how a breach will happen. What you can control is the response.

Root Cause Analysis

Most enterprises are drowning in security alerts. When a breach happens, you need to be able to pluck the relevant incident(s) out of the alert haystack and trace it back to a specific device and action.

Timeliness

On average, the mean time to remediate a breach (MTTR) is eight hours. That figure should be shocking to any security manager. Every business should be working to bring down its response time.

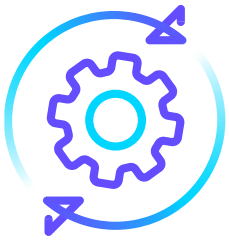

Mitigation

Once you know the cause of a threat, the cleanup process can begin. That usually means isolating the impacted area (a device, a network), removing malicious software, and testing to ensure that the breach is truly gone.

Reducing your MTTR is achievable, if you have full visibility into what's happening on your network and the ability to apply security policies consistently. **That's a huge "if".**

Here's the reality:

Most security and network teams can't see what's happening on their networks - particularly when hybrid cloud is involved.

They can track data flowing through strategic choke points and analyze it in a SIEM. They can react to alerts. But they don't have a way to trace those alerts back to "patient zero" in a timely manner, let alone test if their mitigation strategies are actually working.

What if there was a consistent source of data that flowed through every application and device on your network?

Surprise! The data does exist. It's sitting on your network right now, just waiting to be used. **It's called DNS.**

Yes, that DNS. The protocol that's been around since the dawn of the internet. The data that flows effortlessly through your network every day, working its magic quietly in the background.

BLUECAT™

Every user, device, and application on your network produces a treasure trove of DNS queries every day. Think about what you could learn about your network if all of that data was at your disposal:

Malicious activity:

91% of malware uses DNS to navigate through networks, establish command and control, and exfiltrate data. Paying attention to malicious DNS queries could stop these exploits in their tracks.

Network performance:

DNS query patterns are often the first indicators of significant errors and misconfigurations. Stopping a pattern of NXDOMAIN responses can save a lot of remediation time (and admin hours, and bandwidth costs) down the road.

BLUECAT™

If DNS data is so amazing, why don't more security teams use it?

It's hard to compile.

If you're managing DNS through a decentralized system like Microsoft DNS or BIND, capturing DNS data means manually compiling it, server by server. That's weeks or months of work for every incident.

It's hard to act on.

Compiling DNS data is difficult enough. Applying security policies to it across the enterprise? For decentralized DNS architectures, that's an even greater lift.

It belongs to the network team.

In most IT organizations, DNS is managed by the network team, not the security team. Compiling all that data and acting on it usually involves complicated interdepartmental negotiations around who has access, what's delivered, and when.

BlueCat Adaptive DNS turns your DNS data into security gold.

BlueCat Adaptive DNS brings the untapped potential of your DNS data to life. By deploying a simple VM at the "first hop" of every network query, BlueCat provides unprecedented visibility into what's happening on your network and the ability to control the entire enterprise.

Collect, filter, and analyze DNS data in real-time:

BlueCat compiles and analyzes all of your DNS data. And we mean "all" - including the internal DNS queries which make up 60% of your network traffic, the source IPs of compromised devices, and much more. With this data at your fingertips, security teams can find the source of a breach in minutes, not weeks.

Configure, deploy, and enforce DNS policies consistently:

BlueCat enforces security policies through DNS right at the network edge, ensuring a consistent approach across hybrid cloud environments, IoT devices, and internal traffic. No need for time-consuming agents.

Integrate with leading security and network tools:

BlueCat adds DNS security to the defense-in-depth strategy you're probably already using. From powerful Crowdstrike threat feeds to integration with Cisco Umbrella, BlueCat integrates seamlessly with your existing security playbook.

We're guessing that your interest is piqued. 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 DNS security solutions

bluecatnetworks.com

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.

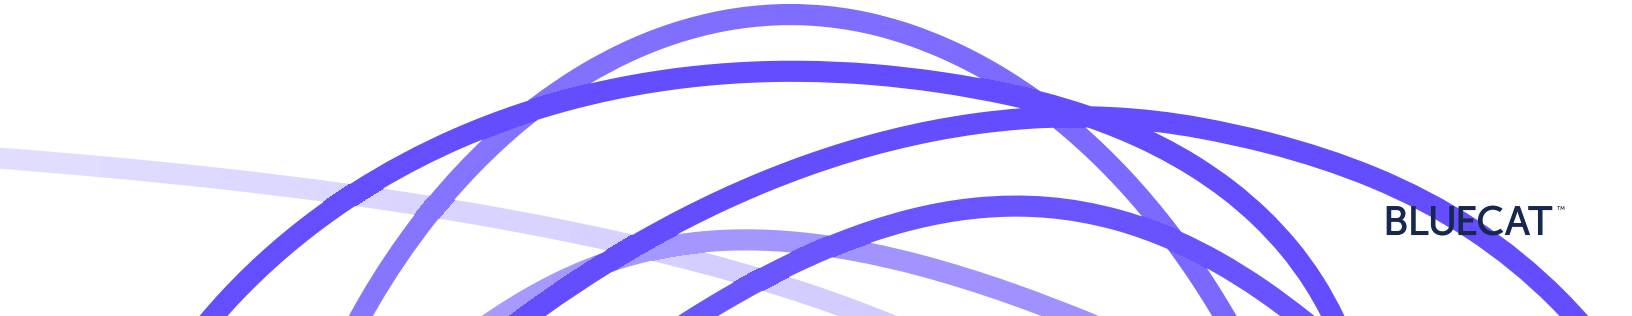

BlueCat Adaptive DNS

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.

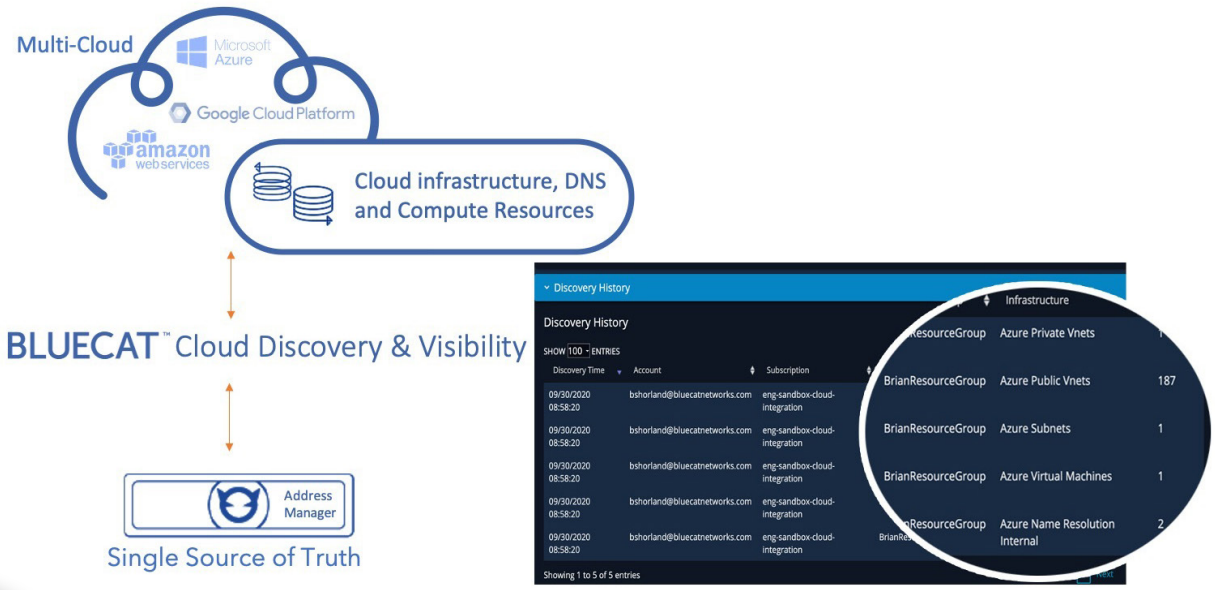

Establishing a Consistent, Centralized Platform for DDI

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.

BLUECAT™

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.

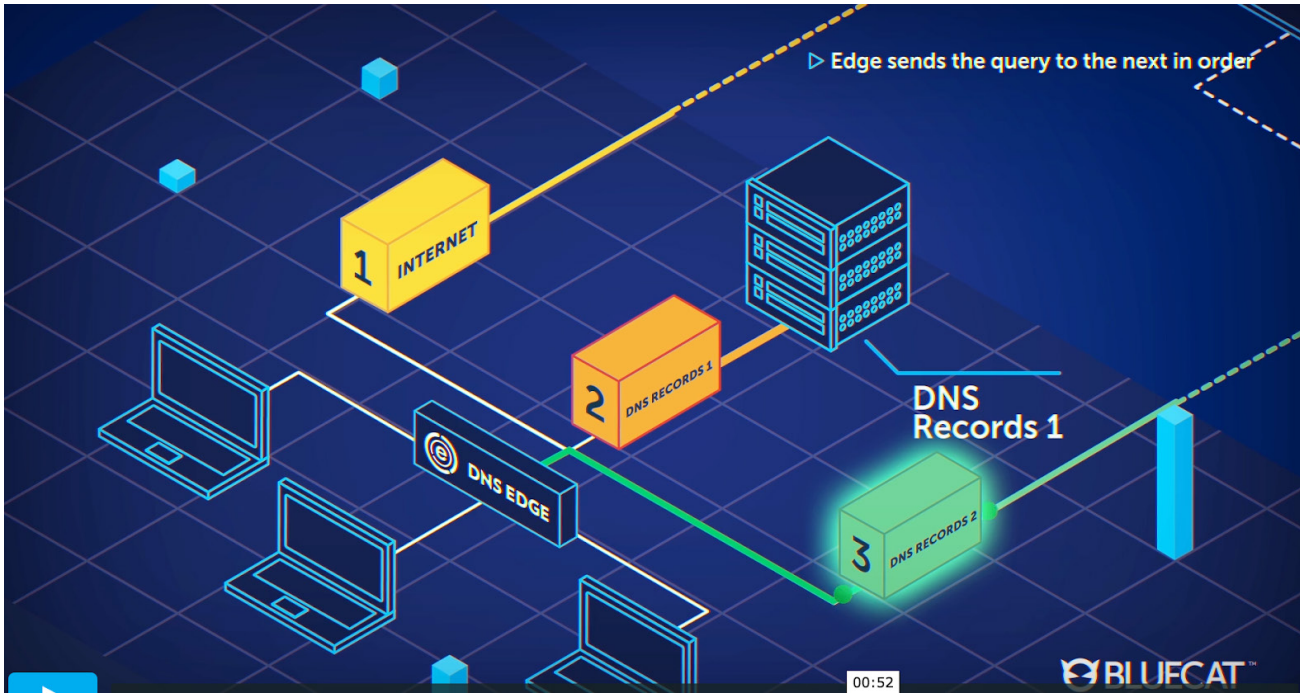

Please see our video for more information on how BlueCat's Intelligent Forwarding works.

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.

BLUECAT™

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.

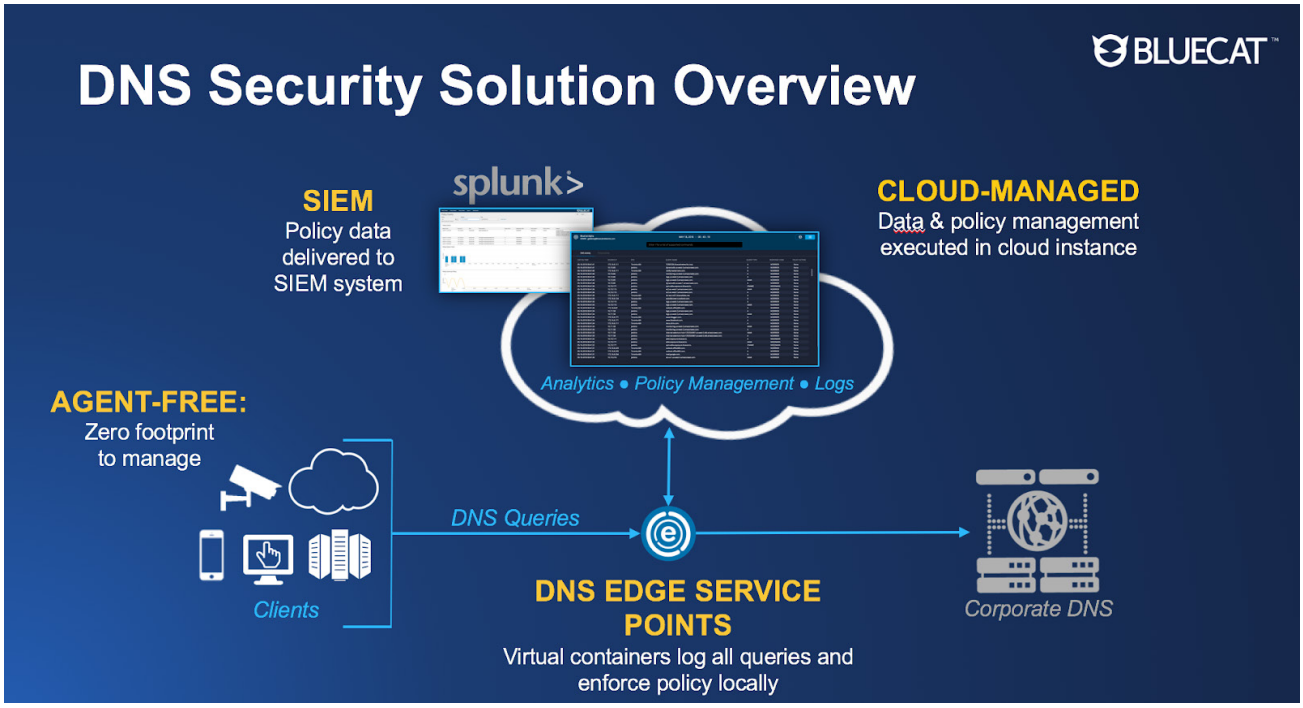

Please see our DNS Edge video for more information on how BlueCat's DNS security application works at the first hop.

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

THE COST OF FREE:

How much are you REALLY paying for Microsoft DNS?

BLUECAT™

bluecatnetworks.com

If you've played Jenga before, then you know how Microsoft's Domain Name System (DNS) plays out as networks evolve.

Everything starts out in perfect alignment. The structure is simple, but it holds together well.

Over time, the foundation starts to erode. Pieces of critical infrastructure are moved around. A new network region here, a hundred new employees there – each layer of complexity puts new strains on the system. Mergers and acquisitions create an awkward tangle of network pathways. New security layers lead to additional complications.

You know what's coming, and try to avoid it. You patch. You reallocate resources. You hire more system administrators to manage an increasingly unstable architecture. You build an entire organization to manually respond to network demands.

Then it happens – the moment of reckoning.

Administrators spend so much time fixing zones and domains that they have little time for their "real jobs". Downtime slows ordinary business functions to a crawl. Domains get stuck in a circular resolution process through overlapping regions and zones. An inflexible network architecture makes new initiatives either too costly or impossible. Poor visibility leads to compromises on security.

The costs of Microsoft DNS may start as a slow drip, but at a certain point they become a torrent that threatens network stability and constrains strategic initiatives.

The costs of Microsoft DNS may start as a slow drip, but at a certain point they become a torrent...

How can network administrators and CIOs keep their networks from reaching this problematic state?

In this eBook, we'll examine the true cost of Microsoft's "free" DNS by looking at the business implications and hard numbers associated with a dysfunctional network architecture.

The Domain Name System (DNS) lies at the core of every network. DNS acts as the "phone book" for every query, channeling traffic to its proper destination.

Most networks start with a simple, easy to administer architecture. Administrators just want to get the system running with as little expense as possible. It's not surprising, then, that so many network administrators go with Microsoft as the default service for DNS. Microsoft's DNS tools are free, and at a basic level they work.

This early in the game, few IT administrators have a long-term perspective on how Microsoft DNS will constrain business initiatives or ultimately weaken their system architecture. Enabling strategic business initiatives and managing network complexity don't even appear on the radar.

Over time, the business logic of sticking with Microsoft DNS will gradually erode for any organization. Microsoft DNS is included in the standard toolkit, but that means that it only handles standard tasks. As organizations evolve, they need a DNS management system that can handle changing requirements and increasing complexity. If administrators don't pay close enough attention to the infrastructure needs that underpin these changes, the network can quickly slide into dysfunction. In this context, the cost of remaining with Microsoft can be quite high.

DOWNTIME BY THE NUMBERS

In a 2016 study, companies surveyed reported an average of five downtime events each month, with the cost of each downtime event ranging from $1 million a year for a typical midsize company to more than $60 million for a large enterprise.[1]

Tactical Constraints

SLOW ZONE TRANSFERS

Complex, overlapping zones in Microsoft DNS often lead to latency and dropped connections. It can sometimes take several hours for IP address changes to filter through a Microsoft-based DNS schema spread across multiple regions.

STALE RECORDS

When network complexity reaches a critical point, Microsoft tools can produce a DNS database that is never fully up to date. Changes are quickly overcome by events in other zones, resulting in a continuous circle of updates that never fully resolves.

NO ERROR PREVENTION

As DNS complexity mounts, the slip of a finger can result in misdirected traffic that snowballs through chains of connected servers. Microsoft's DNS tools have no mechanism to identify or correct the source of a "fat finger" issue. Tracing the origin of a problem can result in hours or days of downtime.

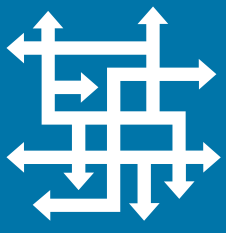

COMPLEXITY

Many standard tasks in Microsoft DNS environments are onerous, increasing the probability of human error, poor service, and outages. Without centralized and automated management capabilities, updates require hands-on support from IT support personnel.

Strategic Constraints

DEV OPS

Agile operating environments require a flexible, easily adaptable network architecture. Testing new iterations of software, creating temporary zones for a development push, and de-provisioning unneeded parts of the network, are all difficult to accomplish on the fly in an admin suite reliant on Microsoft DNS.

AUTOMATION

Automation eliminates manual processes that used to consume IT departments. Unfortunately, Microsoft DNS tools do not support automation in any form, hindering the automation of business processes like the ability to stand up and tear down domain names quickly or leverage APIs.

SECURITY

Microsoft's DNS tools were not built with security in mind, even though an estimated 91% of malware uses DNS to maneuver through target networks. When a breach or incident occurs, the patchwork nature of Microsoft DNS makes it difficult for network administrators to identify, isolate, and mitigate harmful activity.

Microsoft DNS Horror Stories:

The Nuclear Football

"We're one of the world's largest brands. We operate in 150 countries with 150,000 employees. Yet, just three network admins manage Microsoft changes using specially assigned laptops. They refer to these laptops as their 'nuclear football'. Once someone mistakenly deleted a zone that took out the intranet and Exchange for half a day. As most don't know, there is no 'Delete Undo' function in Microsoft DNS so it was lost altogether. Had they not had an off-line copy in the lab, they would not have been able to restore those critical applications."

Information Technology Director,
Multinational Manufacturer

The MacGyver Delusion

MacGyver was famous for working his way out of any jam. He could escape from a maximum security facility with little more than duct tape, a Swiss army knife, and dental floss.

Some savvy network administrators think they can MacGyver their way around the shortcomings of Microsoft DNS. They devise tools and work-arounds. They conjure up hybrid solutions that integrate BIND and other "flavors" of DNS on top of a Microsoft foundation.

Yet these tend to be short-term patches rather than long-term solutions. Custom software layered on top of Microsoft DNS may have the appearance of a well-oiled machine, but there are significant risks in pursuing this strategy.

RISK #1: PASSING THE STRESS TEST

Adapted Microsoft DNS solutions may work reasonably well during a time of normal operations, but they quickly fail during times of stress on the network. A surge in network traffic, DNS routing errors caused by human error, or integration with a new network tool – all of these can bring a patchwork solution to its knees.

RISK #2: THE COST OF ADAPTATION

Like the Microsoft DNS tools they are built on, work-around solutions lack the flexibility to adapt to an increasingly complex network. When DNS practices start to diverge on different parts of the network, or if a different variety of DNS management comes into the picture (through an acquisition, for example), custom solutions based on Microsoft DNS will consume resources and time to adapt to the new situation.

RISK #3: TURNOVER

Microsoft DNS work-arounds also have a single point of failure – the person or team responsible for creating them. If the one knowledgeable person in Microsoft DNS leaves the organization, the work-around they created suddenly becomes endangered. Any change in network architecture might require the creation of a new tool, or a costly adaptation of the existing one.

RISK #4: THE COST OF INTEGRATION

Patchwork solutions are never seamless. Building new layers on top of Microsoft DNS inevitably creates more complexity and a greater chance of something slipping through the cracks. The cost of developing, managing, and deploying these integrations over time can add up quickly.

In the end, work-arounds, patchwork solutions, and hybrids end up demonstrating the need for a comprehensive resolution to the fundamental problems of Microsoft DNS. They are not a long term solution. They merely delay the inevitable move to a more systematic, unified approach.

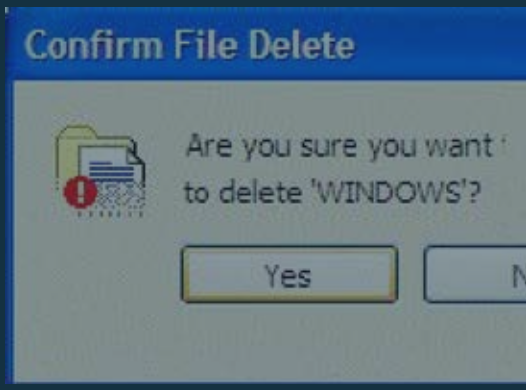

Microsoft DNS Horror Stories:

I deleted EVERYTHING

"I did a deployment and it deleted EVERYTHING... Then I had no choice but to sit there while it rebuilt everything. In the meantime, Active Directory (AD) replicated all the delete operations to the other Domain Controllers (DC) and all the DNS data magically disappeared from AD. I ended up with an outage of nearly an hour!"

Sr. Network Engineer, University

Migration Challenges and Opportunities

Despite the well-documented shortcomings of Microsoft DNS, risk-averse network administrators are often reluctant to move away from it. Migrating to any new network system comes with risks and potential costs; DNS management platforms are no different. Many Microsoft DNS customers rightly fear the ripple effects from a disruption of this core service.

At a certain point, the risks of continuing with Microsoft DNS start to outweigh any concerns with migration, but there are also ways to minimize risk. With a measured, clearly mapped out strategy to move DNS into a centralized management system, migration from Microsoft DNS can be accomplished with few hiccups.

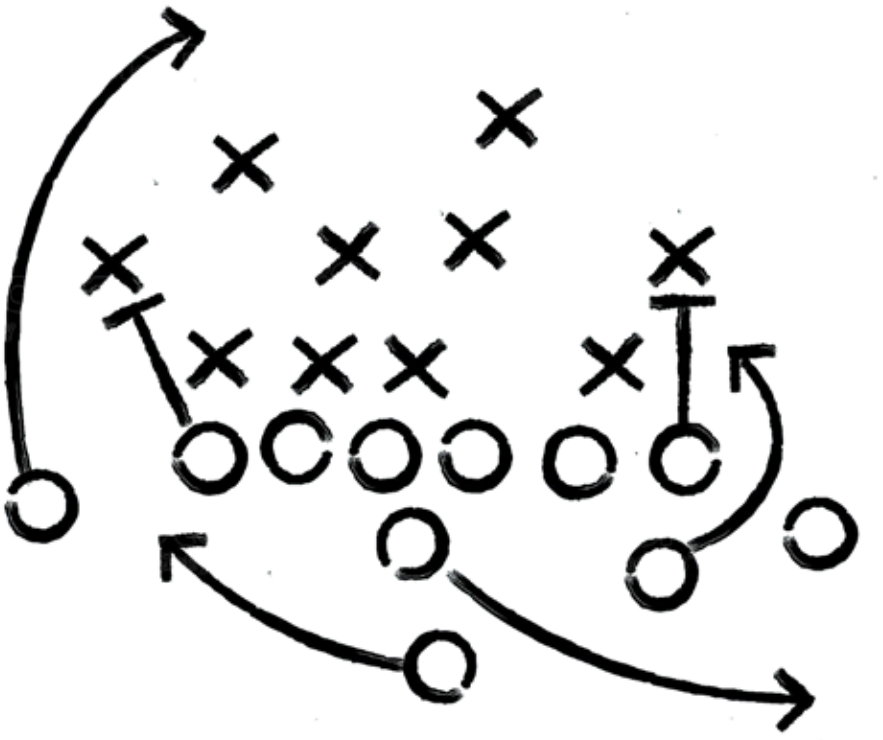

Changing from one DNS management system to another can be stressful, but it can also yield concrete dividends. Here are just a few things a network administrator can expect to have visibility into when switching from a Microsoft DNS system to a unified management platform:

ORPHANED DATA

As Microsoft DNS architectures evolve, information inevitably gets lost in the shuffle. A unified DNS system will show network administrators these "bridges to nowhere", allowing the network to once again encompass the entire universe of available data.

INACCURATE DATA

When DNS records are misconfigured, stale, or incorrectly deleted, network traffic comes to a halt. A single point of truth for DNS data can identify these broken resolves, correcting (or at least identifying) inaccurate information to keep queries humming.

OVERLAPPING DATA

Redundancy has its merits in network administration, but there is a limit to its usefulness. Active management of a DNS system allows system administrators to reduce the inefficiency of overlapping information.

NON-STANDARD DATA

To achieve maximum performance, DNS architectures need standardized rules and procedures throughout the network. Active management of DNS allows administrators to eliminate non-standard data which can derail normal operations.

INEFFICIENT DEFINITIONS

As the complexity of DNS architecture grows, queries can be routed in odd ways that negatively impact network performance. A DNS management platform provides administrators with the strategic view they need to streamline operations.

DOWNTIME BY THE NUMBERS

Service provider problems and internal human errors each make up nearly 25% of downtime.[2]

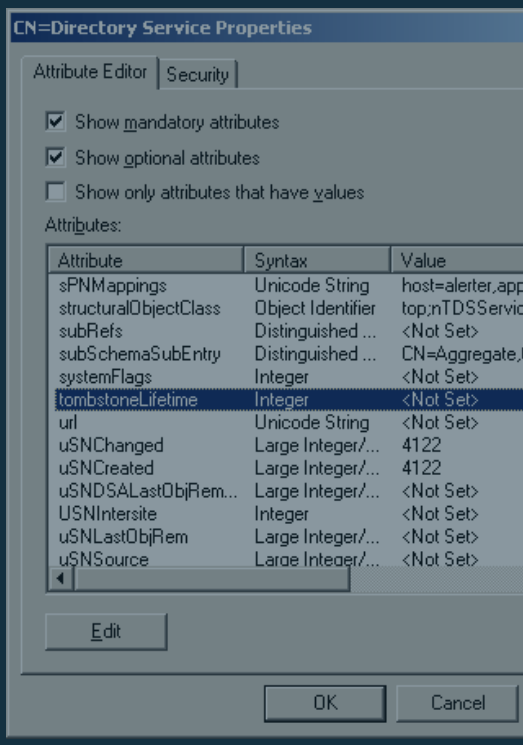

Microsoft DNS Horror Stories:

Data still lurks

"Delete operations don't actually delete the data, as you would expect. Data still lurks there until Active Directory (AD) purges the data as part of the tombstoning process. In the meantime, it adds all the records again as new AD objects, so your AD object database grows massively with every deployment."

IT Systems Administrator,
Major Retailer

Quantifying the Cost of "Free"

Microsoft DNS tools are included with the standard network package, but that doesn't make them free. As networks scale and evolve, the constraints of Microsoft DNS become significant. They also become quantifiable – measurable in terms of administrator hours, downtime, and likelihood of a security breach.

At BlueCat, we've helped hundreds of organizations move from Microsoft DNS to a flexible, automated, intuitive DNS management system that meets their needs. In the process, we've learned a lot about the true cost of Microsoft DNS, both at the tactical level and the strategic level.

What is your organization actually paying for Microsoft DNS?

Our team of DNS experts work with organizations across every industry to provide IT administrators and CIOs with hard numbers about the business cost of the status quo. With baseline knowledge about your IT operations, we can provide you with instant feedback on your cost of free, and more importantly how an Enterprise DNS solution can add immense value to your organization.

Let's Get Started!

[1,2] https://www.networkcomputing.com/networking/high-price-it-downtime/856595126