The Technology Leader's Guide to DNS Enterprise

Paul Vixie Webinar

October 31st, 2019

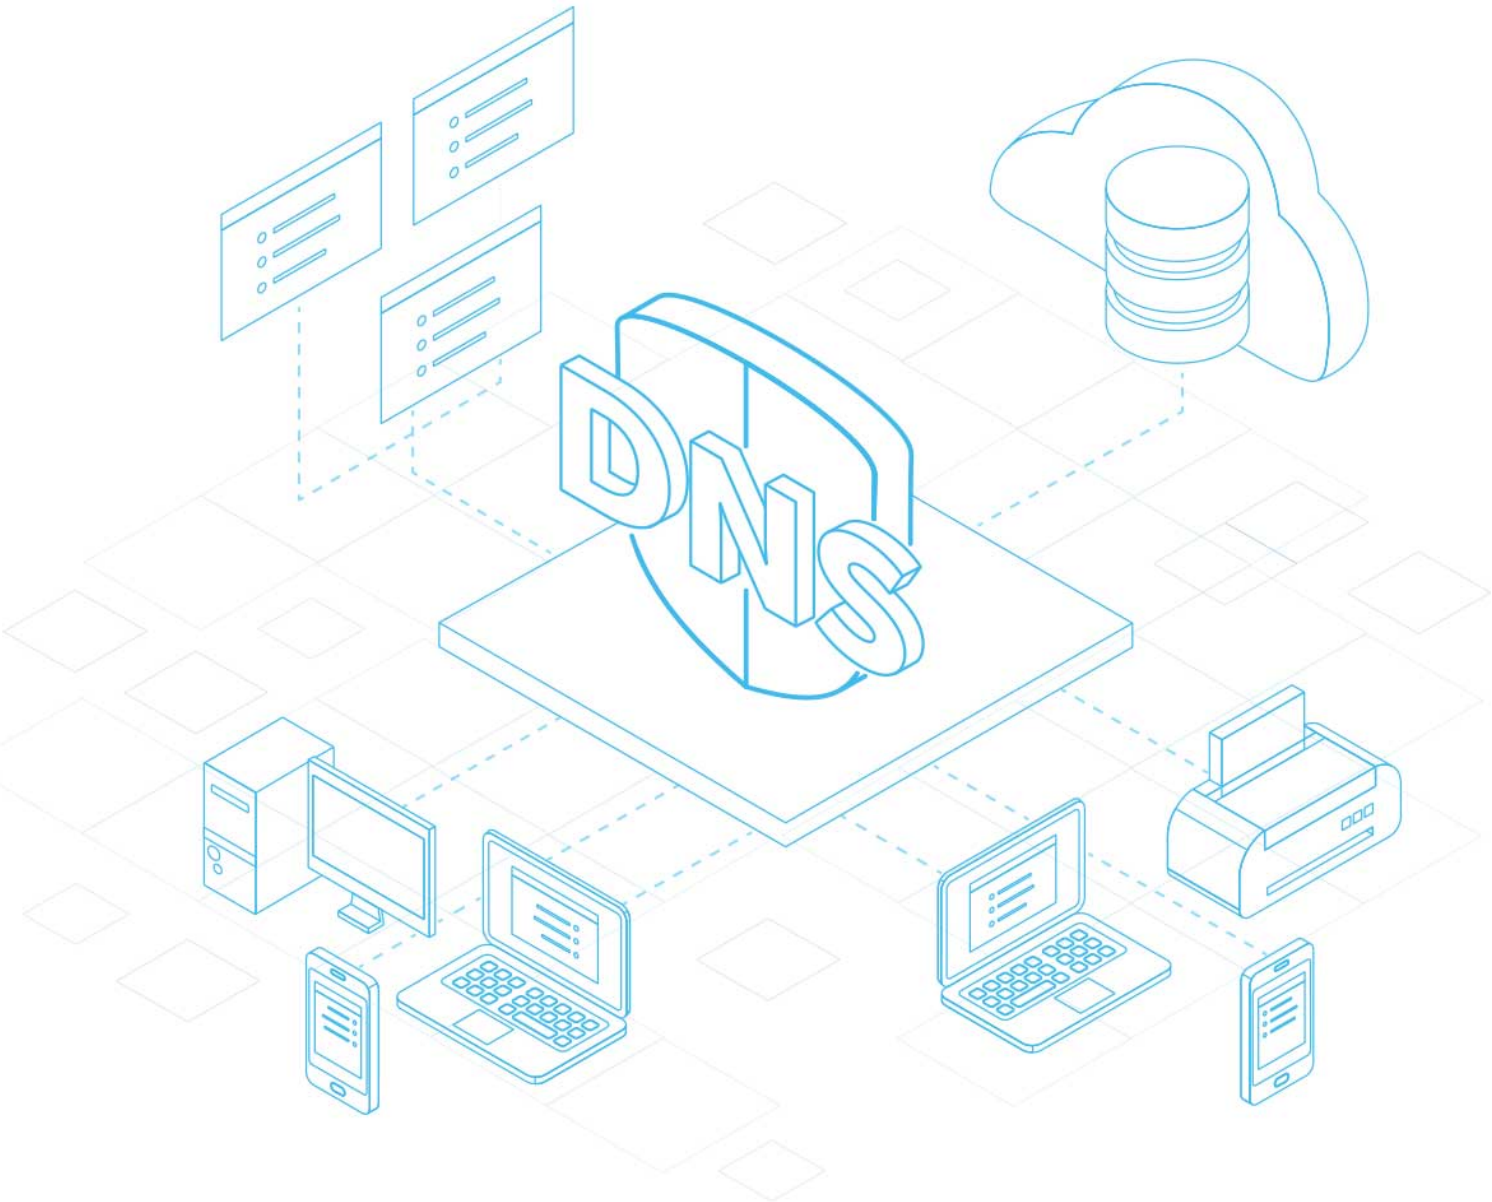

Erin: All right, so why don't we get started while some people still get logged in. Hi everyone. Thank you so much for joining us today and happy Halloween and I actually see that we have some key competitors on the call today, so big welcome to you guys. My name is [Erin 00:02:52] and I'm the marketing manager here at BlueCat and with us today we have our esteemed guest, Paul Vixie, who is the CEO of Farsight Security and of course Mr. Andrew Wertkin BlueCat's Chief Strategy Officer. We're aiming for today's conversation to be in the 20 to 30 minute range plus some time for questions of course.

Erin: If at any time you do have a question, by all means ask it in the Q&A feature, we'll do our best to answer questions on the fly, but if we don't get to it during the conversation, we'll make sure that we get to all your questions at the end and of course we are recording the call and we'll share it with everyone afterwards. If you miss something or you have to step away or drop, we've got you covered and I think that's it. I'm going to pass it over to Mr. Wertkin to kick us off with the discussion.

Andrew Wertkin: Thank you Erin and welcome everybody. Thanks for joining and we're thrilled to have a Paul Vixie on with us for this webinar. We want to talk about DNS in the context of the enterprise both from a network operations and security operations standpoint. This is an area where we've been doing a great deal of work at BlueCat and really coming from our customer's requirements and the things that they're trying to drive in their enterprise and the just absolute wonderful usability of DNS both from a visibility perspective to understand what's going on on the network and then also from a control point standpoint, whether that's on the network side or the security side.

Andrew Wertkin: DNS is a highly efficient protocol that allows for some pretty interesting things. On the network side we see more and more requirements around traffic steering. How do we make sure that the right answer's given to every single client and that right answer being what is the closest healthiest endpoint for whatever service I'm trying to resolve as the... frankly, the inside the network becomes more and more like the Internet as our customers are investing in building their own internal cloud native applications that run on many, many endpoints throughout the enterprise, whether they're hosted on public cloud or built into their private clouds, that the amount of change that goes into planning for and deploying and traffic steering goes up pretty dramatically as well.

Andrew Wertkin: Egress management is another area where customers are seeing a lot of churn and we'll get into that in a bit as there's more and more Direct Internet Access for services like Office 365 that are being consumed locally as opposed to everything back hauling to the data center and out of [inaudible 00:05:39] link there. Performance management is another key area that we'll talk about as well. On the security side, DNS becomes an amazing source for visibility of what's being accessed and certainly a great control point to prohibit access for instance, and then for forensics if something does happen frankly from the network side or the security side of something, if there is some indicator that

suggests that there might be for instance, commanded control running on a specific endpoint, DNS becomes highly usable from a forensic standpoint, so I understand what was the context here. What other things did this device do? Sure, I can block what I noticed to be bad but what did it look up internally? What else did it go externally and who else might have gone to these endpoints in the past?

Andrew Wertkin: It's an area that we think about a lot. We build technology around, we speak to our customers and prospects a great deal about because we think it's just under harvested set of data that can be used for a great deal of things. Next slide. I'm sure many, many people on this webinar know... in general a lot of conversations about DNS just start with why it matters inside the enterprise and we talk a lot about the complexities of managing DNS inside of enterprise and obviously that's how BlueCat built its business and it's really around this private namespace. The segmented DNS, the DNS that runs inside a corporation that has nothing to do with public DNS except at certain endpoints and how that can be managed.

Andrew Wertkin: This diagram on the right is a visualization that some of our crew did during a hackathon a couple of years ago, but it's a visualization of the DNS inside customers networking and we pulled this out of our control system, and so each of those little roundels are a DNS server in the enterprise in this specific case, and so there's plenty of them and all the little pie slices are different DNS roles. They might be primary zones. They could be secondary zones, they might be forwarding rules, they could be stubs zones, all the different roles that any of those servers might be deployed and all the lines connecting them are either resolution paths or dependencies for things like zone transfers and so what you end up building is a fairly complex looking set of servers and roles that are all communicating with each other based on the specific use case.

Andrew Wertkin: All done to provide basically a single requirement, which is we need to make sure that every endpoint is getting the right answer to every query as quickly as possible, so there might be more servers for survivability. There could be more servers because we want to make sure that there are the zone data local to every geography and that it's across a couple of different data centers so that we can plan for things like faults. All of that work and that configuration, which is driven through our applications and other applications are done to drive that requirement and given this complexity and we always joke around and when there's a problem on the network it's always DNS and oftentimes it is DNS and it becomes very difficult to understand where that problem emerges from without having a sense of what's actually flowing through the system.

Andrew Wertkin: I can query successfully from this endpoint, but I can't query it from this other endpoint. Why is that the case? Did a zone expire? Is there issues with zone transfers? Is there a problem with some service that's in the resolution path? Somebody who's getting a stale record. Why are they getting a stale record? Is there a caching issue? All of these things require some level of visibility.

Andrew Wertkin: On the internal side some of our customers and some companies in general also override the Internet. They deploy an internal route zone so that the endpoints can't resolve public queries if they wanted to. There's no resolution path. They override the entire Internet and the resolution is actually done by proxies from a web perspective and that's changing a bit and customers are opening up more egress points for things like Direct Internet Access where all of a sudden you can't put proxies in 250 different egress points or they're adopting cloud security type solutions, but so that causes some complexity if I've overwritten the Internet, especially if I've reused internal and external zones not in a split view sort of way, but in almost willy-nilly way because these things have never had to meet or marry in the past.

Andrew Wertkin: You see, differences in the utilization of DNS in this private enterprise space versus the public space as well like on average, about 50% of the queries we see in our customer base are internal queries versus things we're cursing out to the Internet. In an average enterprise only about 65% of those private DNS queries resolve to, at the end of the day A/AAAA records where on the public side the stuff that's going to the internet over 90% of them are resolved to A/AAAA records.

Andrew Wertkin: We're also exposing a lot of other use of DNS inside the enterprise and those other queries are going to be everything from SRV records to find the local active directory controller or Kerberos related or just a variety of different records that are used internaly to control how I onboard my machine onto the network and so you just see a different utilization pattern all together on the internal side versus the external side and if there's an issue with DNS, obviously all of these other systems that depend on it, when it's hard to think of one that doesn't, but for instance, like active directory or anything else then has a very severe problem as well.

Andrew Wertkin: This stuff needs to be up and it needs to be managed well and that requires understanding what's going on with the system. Next slide. There's a bunch of key challenges and by the way, this diagram on the right here is the same as the last one but it's a completely different customer and you can see the difference between the two. They manage their DNS very, very different. There's a style to how some corporations think about managing DNS. Is it distributing primary zones to local sites, having big servers with all the secondaries. There's... the complexity leads to many different combinations of how this is done.

Andrew Wertkin: There are some key challenges. One is just the acceleration of change and this really comes from the dev ops types use cases. The more things change internally, the more compute that's deployed and how rapidly compute can be deployed leads to an acceleration of change in DNS that requires pretty responsive capabilities to deliver these changes out to the enterprise. Some of our customers used to have a single change window every day where they would deploy any changes to DNS and now are hitting our system with thousands of API calls an hour in some case, driving in changes to DNS so that

they can support the business who is building and building and building applications and new endpoints.

Andrew Wertkin: That acceleration of change becomes difficult to manage but then with all that change, it's also more likely that something will be done incorrectly or something might not be working and so it's just more of a reason why visibility to this data is critical. We're seeing all sorts of different types of implementations now as customers think about DNS in terms of network segmentation. How do I deploy DNS in a way where the only resource records available for certain networks are the ones that they should be looking up for instance.

Andrew Wertkin: We see a lot of different architectures now with fault zones. How do we make sure that if a data center goes down, we don't take down the entire business. We reduce what the blast radius is of any server going down as much as possible and so we're seeing some very interesting new deployment architectures as our customers drive a service deployed in a way where reduce what can happen if any node goes down, period.

Andrew Wertkin: Again, I talked about Direct Internet Access and this becomes a big challenge especially depending on how the DNS was managed in the past but we're seeing lots of issues with zone management on the inside and the outside. We're seeing customers who want only to allow certain types of traffic out from the branches and stores and anywhere but the data center but the other stuff, they want still to go back haul and go out through their proxy and all of the security architecture that have been placed. In some cases it's opening up to the Internet in more and more sites, but only opening up a small part of the Internet from a DNS perspective.

Andrew Wertkin: Securing the servers is a constant thing that we're concerned about and by securing the servers, I mean, making sure it's up and available and that's everything from vulnerability management, to penetration testing, to ensuring that we don't DOS an internal server. How do we potentially rate limit or start refusing traffic that might harm the servers and then also utilizing DNS to improve the security posture of the company in general, where given this visibility and given the control, given our ability to very efficiently block something that we don't want people to go to because we know it might be bad or it might be risky or it's not even business relevant, and Paul, will talk much more about this if we know that already, then one of the most efficient ways to stop people from going there is simply blocking that resource record.

Andrew Wertkin: Inside an enterprise and depending on the type of enterprise it is there is only certain types of services that you would expect people to use and given how regulated that industry is, or that enterprise is, you'll see a narrow set of things that they're going to access where they would expect people to access on the Internet and so it's not just about blocking things that we know to be bad, but it's always the gray list stuff. It's the stuff we don't know if it's bad or good, but

for sure it's probably not business relevant because it's on... coming from a dynamic DNS provider for instance, or it's on a TLD that is used only for abuse, so just sort of a low rent TLD that is never business relevant.

Andrew Wertkin: There's lots of contexts other than whether we know something's good or bad and DNS becomes a very important part of utilizing that context to efficiently control the traffic. Next slide. This is just one example. In this specific case, this is a real example, it's security agent. It's the internal DNS traffic the security agent is calling, resource records. It's looking up inside the network to do its job and over the short period of time, this agent queried its own internal records 15 billion times across 50 or 60,000 different endpoints, a ton of queries.

Andrew Wertkin: In fact, sometimes we've mapped this thing doing about 5,000 queries per second and this is... not going to mention which vendor it is but it seemingly, well, I know it's not actually using the stub resolver. It's not using the operating system's DNS configuration. It's literally... has its own application code to query DNS and it clearly doesn't cache anything. It's constantly querying the same two internal endpoints that have a TTL of eight days or something. These records don't change. Their internal servers don't move around. They're static compute on static IP addresses and it calls it like an insane number of time and in fact, this blip at the bottom where you see the amount of traffic growing dramatically. In fact, it goes up to 13,000 queries per second. It occurred because a patch was installed and this thing couldn't communicate anymore, so there's a problem with the agent and it just started queering like crazy.

Andrew Wertkin: In a matter of hours the amount of DNS traffic inside the network over doubled and my point there is, the mechanism by which this customer found out there was an issue, was analyzing what was happening in DNS, and so from a performance, from a visibility, from a debugging, what's going on. I bring this use case in just to show sort of the power of being able to see, collect and monitor and analyze this stuff. Next slide. With that, I want to shift over now to Paul. Paul, why don't you take control and go on from here.

Paul Vixie: Thank you Andrew. First I'm going to I say this is my first webinar with BlueCat. I've been working with BlueCat since sometime in the early 2000s, because they were at the time a BIND bundler. You guys were using the BIND as your protocol, the engine inside of your appliance and I was the BIND guy and then I was just sort of running the company that was to BIND company, and so I am glad that we have finally progressed to the point where we're doing joint webinars. I'm glad that we have some competitors of BlueCat on the line here.

Paul Vixie: I have always been pretty open to working for what's best for the Internet, a rising tide that lifts all boats and so I am excited for anybody who thinks that running their own DNS instead of outsourcing it to some cloud provider, might be on this call and like be ready to hear some tips and techniques. DNS as a control point for cybersecurity has been, I don't know, a special hill of mine that I've been climbing since the early 2000s. Some of you know that the DNS

response rate limiting capability that first appeared in BIND and is now generally available in other servers was done by my team at ISC because we could see that DNS was being used as a reflecting amplifier and that was bad.

Paul Vixie: Some of you may remember that in BIND 8 we had a very similar config file to BIND 9, but there was nothing like views but it turned out that a lot of enterprises needed to have a different view of DNS inside their network than what people outside the network with them see and so that also was a control point for cybersecurity, and I think pretty much all name servers now have the view capability that was first pioneered by the BIND 19. We also developed the first really standardized DNS firewalls called RPZ, response policy zones that allows your average recursive name server operator to subscribe in real time to... we continuously updated corpus of policy for how responses will be perturbed deliberately in an enterprise context and that also is now present in several name servers and at some point we get the RFC published.

Paul Vixie: Finally, my team at Farsight has created the dnstap system. T-A-P that is an open source and open data, open protocol for collecting telemetry from name servers and that sort of was meant to give name server operators the capabilities that Andrew just described as being present in the BlueCat system and we want that to be the norm. We want everybody everywhere to be able to get telemetry out of their operating server and put it into a set of tools that might be multi-vendor in nature. You might have more than one type of appliance you'd like to get the same telemetry protocol from all of them.

Paul Vixie: I really have been pushing on DNS as a control point for cybersecurity for a very long time and I've also been pushing for the idea of running your name servers locally and you might not be able to do that for the authority servers that publish your content because you might really have a reason why you need a content delivery network to do that for you or you might just want to have worldwide resilience on that publishing capability that you can't afford to sort of go build by renting servers in Iraq somewhere in 20 different time zones.

Paul Vixie: It may be that on the authority side there are some pretty good reasons why you would sort of pay someone else to be your secondary name server and we do that at Farsight, but as far as the recursive name server, that intermediary that you sort of point all of your DHCP clients at and get it to sort of give you answers if it knows them, maybe perturb those answers with local policy, be the subject of telemetry analysis so that you can find out exactly who you're talking to, look for anomalies. All of that. I have never seen a reason to push outside of the boundary of the enterprise network and we've gradually seen, first we commercialized the Internet and started moving that recursive name server function into our ISPs and then a number of different companies decided, hey, we could do this Internet wide with any cast, and so you've got sort of 8.8 and 1.1 and 9.9.

Paul Vixie: You've got all these different people doing that and really I don't see a need. I think that there are plenty of ways for an enterprise to either get an appliance like the BlueCat or one of their competitors or build it themselves out of open source. There are too many ways to run this inside and keep control and have the power of your own fate in your own hands, and so I really am... these are all kinds of hobby horses of mine and I never miss an opportunity to tell people, run your own recursive name server. I think if I had gotten run over by a truck, that's probably what my family would have chiseled onto my gravestone, to run your own recursive name server.

Paul Vixie: The things that will happen if you go outside your local server or outside your local perimeter in order to get these recursive answers for your applications and your users include both sort of monitoring other third parties, possibly your ISP or somebody who's monitoring your ISP or maybe competitors that are further away than your ISP, might be able to see your traffic and it also sort of opens you up to various types of cache poisoning attacks or it means that the policy for DNS sec, which is a security protocol we added to DNS might not be set exactly the way you'd like it to be, and so I just want to really recommend that anyone who is outsourcing recursive DNS please reconsider and anyone who is running their own recursive DNS who doesn't know about how to do policy, local policy enforcement so that maybe the infected nodes in your network can no longer reach their command and control servers because you don't want that to happen, you know, that's in your power.

Paul Vixie: I really am speaking specifically of managed private networks. I'm not talking about ISPs. I know there are ISPs who behave this way. I don't happen to subscribe to one and I'm glad that I have choices. ISPs should not be doing this kind of thing, but unless it's a parental controls feature that somebody is paying for it, that is, it just should not be the default. Let's go to the next slide. I mentioned that some of the problems of sort of what happens when you put your DNS server too far from its clients but I want to go deeper into that because the problem is much worse than most people realize.

Paul Vixie: If some intruder breaks into your network and that might be a human intruder or it could be a hardware intruder, it might be an IOT box that's doing things you didn't know about, or maybe it's a poisoned supply chain, so that the IOT device is doing something that not even its maker knows that it's doing or poisoned software supply chain like that BIOS update that added a whole second operating system to your motherboard. There are all sorts of intruders and they aren't all people and they don't wear hoods even when they are people, but one of the things that they will want to do after they have broken in is probably steal your stuff.

Paul Vixie: They're going to find some treasure trove of source code or business records or something that is of value to them, they can sell out on the dark web or maybe perhaps they're doing industrial espionage, they're planning on selling it to your competitors. They will often use DNS to do this because it's possible to build a

tunnel on top of DNS and if you permit your DNS to be completely open between sort of all of the end user clients, virtual servers, laptops, whatever, inside your network, if they can make direct queries through the DNS that are neither monitored nor filtered then there are plenty of software packages out there that they can deploy that will basically build a VPN over DNS itself and I've seen this work.

Paul Vixie: I've seen a full motion video over Webex work on top of a tunnel that was built on DNS. DNS is faster than you think apparently, and this is the kind of thing that I think most CSOs would say, I hate that. I don't want that to be possible, but it will be possible unless you run your own recursive name server inside your perimeter and you monitor it and you put in filters. There are a number of sources of filtering information. We can get a list of sort of all the popular DNS VPN so that you can block them but you don't have that option if you're using a name server provided by some cloud provider.

Paul Vixie: I've already mentioned distributed denial-of-service. Back in the old days we only had trusted people using the Internet. You had to have a government contract but then it got commercialized and anyone who could pay for an Internet connection now has one and that really changed everything because the protocols we built were meant for trusted scientists to go exchange information with each other and trusted scientists were never going to spoof an IP source address in order to then cause maybe some DNS server to answer a question that was never asked by the apparent source of that question.

Paul Vixie: That kind of DDoS vulnerability is kind of built into the DNA of the Internet itself and would be very difficult to extract or modify at this stage, and so what we have to do instead is make sure that none of our name servers are able to be reached by people who shouldn't or are willing to do things like DDoS that they should not be doing because we are all responsible for whatever flows outward from our networks.

Paul Vixie: Now, there are plenty of other examples but I'm going to race through some of these because I'd like to get to the Q&A. I'm very interested in what people have to ask. I just want to say you can often detect brand infringement by looking at third-party DNS traffic and that is a type of traffic that my company makes available. We have a very large network of passive DNS sensors and it's possible and actually have some companies who do this to purchase a feed from us that you can then take a look at, the real time traffic that is going on between other parties is not coming through your name service, is not coming through your firewall. You're not going to see it unless you have access to a network like ours but once you can see it, then you can start looking for variations, whether it's an IDN variation using internationalized domain names or just swapping the [IS 00:32:07] for 1s and Os for zeroes.

Paul Vixie: There's variations on your brand name that will tell you that you're being attacked even though it isn't touching any of your servers, any of your firewalls,

that is the kind of thing that DNS makes possible as long as we continue to be able to observe it. I'm going to skip the rest of this. It's obvious that people even spammers need good DNS and the ability to stop them from getting good DNS is a good way to keep them from spamming.

Andrew Wertkin: Hey Paul.

Paul Vixie: Yes sir.

Andrew Wertkin: One interesting aspect and something that we concentrated on, a great deal at BlueCat is given that the egress of these queries to the Internet, the recursion to the Internet is... basically any server on the public side that sees that query is going to be attributed to the egress IP address of the corporation as opposed to some IP address of a local client on RFC 1918 IP address. It's netted in other words. One of the things that we concentrate on heavily is trying to solve that problem for our customer because yes, I know I've blocked something, it's bad, but I don't necessarily know which endpoint on my internal side actually executed that query unless I start trying to correlate logs across multiple name servers, especially given the way things like caching work as well.

Andrew Wertkin: We've been doing a lot of work now to take that initial attributable source IP address and make sure that we can drive the intelligence back to that. I think that becomes a very critical part of this process.

Paul Vixie: Well, I think you're right and I think you're highlighting one of the big differences between rolling your own a recursive name servers out of open source software and there's some very good open source software. I want to shout out to my friends that do Unbound, the Knot server that comes from CZ.NIC Labs and also PowerDNS which comes from a company in the Netherlands and BIND, which I have a hand in myself. There are at least four very good name servers that people can roll on their own. However, BlueCat and various commercial competitors have often built something that is much harder than just install a bunch of open source software and hope for the best and one of the values that you can get is what you described, which is centralized log management and also centralized configuration management so that you can treat a brace of similarly configured named servers as if they were a single object and I think that that is one of the values that the open source community is never going to get around to fulfilling.

Paul Vixie: If you want something that works that way, you're going to have to become a commercial customer of somebody who wants to fill that niche because the open source community's goal is more about the protocol and more about making the Internet grow and less about helping commercial entities than manage it easily or secure it easily. There is, I think, a very valuable service that is performed by all of the companies who provide DNS appliances including BlueCat, which as I say, I've been working with for almost 20 years. Yes, what

you say is so. On the other hand you're calling into question sort of the net problem in general.

Paul Vixie: When we were fighting the Conficker worm back in 2008, 2009, I used to run one of the sinkholes and I used to then determine, hey, there's a net exit gate way for some large American bank that is trying to connect to the Conficker sinkhole a lot of times per day, and from that we can estimate that they have a population of, I don't know, let's say it was 3000 different infected computers in that bank and I would then contact that bank and say, hey, you got a problem and if they could worry about this at all, which often they couldn't because they had bigger problems that day than Conficker, but if they could worry about this at all, they'd say, okay. Send us the logs and I send them the logs, all I've got is the IP address of their net exit gateway and the port number and the timestamp and the question that was asked.

Paul Vixie: I never ran into a bank who could take 3000 of those and quickly turn it around into a list of, well, who were the internal clients that were using that port number at that time because just logging of this kind was not prevalent and it really has to become so, not just for DNS stuff but also for web fetches and everything else. It is extremely important that every enterprise operator be able to correlate the externally visible part of some event, which just has the net gateway address and a port number and a timestamp into something that is internally meaningful so that the security team knows where to go and what to do, and that's an investment that I'm sad to say, not every mid to large company has yet made but it's vital that they do so for the reasons that Andrew just stated.

Paul Vixie: One more mention of response policy zones. My company Farsight Security does not really do reputation. Some of you know that I started the first anti-spam company back in the mid '90s, then I got sued by a lot of people who did not like me calling them spammers. We're not in the business of calling people spammers or anything else. What we are in the business doing here is saying this is what we've observed. This is what we know. This is our confidence interval in what we're telling you. This is the latency, this is everything we can tell you so that you as our customer can make your own informed choice in light of your own local policy as to what to do about that indicator, so that's a big shift that has taken place in my mind in the last 22 years is to not take on the role of sort of policy maker.

Paul Vixie: I really want to feed other people's policies. I don't want to make them for... make those policies myself and to that end, we actually have a commercially available RPZ feed. That's the response policy zone protocol that's a free. It's in your name server most likely where we can actually feed a real time and continuously updated policy corpus that just has in it the domain names that we have first observed recently and what that's meant to do is to enable the people who are running their own recursive name servers as I advise you to do to then maybe treat newly observed domains badly.

Paul Vixie: There was a time that it took three days to get a new .com name. Now it takes 30 seconds because of competition, but I don't think there are a whole lot of domain names that get created within 30 seconds that are actually valid cooperative domain names that aren't created for malicious purposes. In other words, the primary market for a 30 second turnaround is probably somebody who means you ill and so at my house, I'm a customer of my own company, of course. We eat our own dog food as they say. It's not possible to resolve a domain name that was first observed by my company Farsight in the last 30 seconds... no, 10 minutes. I'm using the 10 minute feed.

Paul Vixie: There is a 30 second feed. There is a 24 hour feed, whatever you think is new, we've got a feed for you, and that's kind of cool because if you aren't willing to resolve something because it's less than 10 minutes old and by the time you are willing to resolve it, it's probably been taken down or it's in the Spamhaus feed or it's in the Abusix feed or it's in some other thing because it has behaved badly and now it has a known reputation but until it has a known reputation, I'm willing to just shun it and so far, no complaints. The stuff works pretty well even though it might sound very simple.

Andrew Wertkin: I think as I was talking about before in terms of... especially from an enterprise where there's an expectation of the sorts of services that may be used, certainly a domain that's been live for just 10 minutes is... that sort of timeout, a business relevant service cropping up in that short period of time most likely not and especially if there's a white list of services that we do expect to potentially launch a new domain or new endpoint.

Paul Vixie: Next slide please. Okay. This is going to be dual use, Andrew and I both have some comments along these lines. In spite of our best efforts, those of us in the field to remind people that they could just run their own recursive name server inside their perimeter where they can defend against both perturbation of replies and observation of requests, a lot of people have done it and they had moved their name server or they moved the sort of client to server path to be this really long path, maybe it's to their ISP, maybe it's all the way across the Internet to some anycast provider and when you do that, you dramatically increase your attack surface because there are a lot more intermediate network providers who are able to therefore see what you're doing and possibly send their own answer faster than the real answer might otherwise reach you, and so it's created this apparent need to now protect this overly long path.

Paul Vixie: Now, I want to remind everybody, again, the right way to protect this path is do it in house so that you don't have that problem, but if you're going to use an externally advertised name server that belongs to your ISP or belongs to some anycast provider, then it falls to you to find some way to encrypt your questions and let them encrypt the answers so that nobody on this overly long, unnecessarily long data path can then see what you're doing or change what you hear, and so the IETF took this as a project after the Edward Snowden disclosures of 2013 and agreed that everything ought to be encrypted if it's

going across the open Internet and two different teams inside the IETF decided that they would take on the project of encrypting DNS and then the first team to produce results made something called DNS over TLS.

Paul Vixie: TLS is the encryption protocol used by HTTPS and a whole bunch of other encrypted protocols, so it was not an invention for DNS. It was just, gee, we've got TLS, let's use it and let's make DNS live over TLS and this is gradually rolling out right now. For example, the Android 9Pie operating system offers an option that they call private DNS and it is DoT and it is possible to deploy DoT in your name servers. We're working now on getting it to work on the content servers because we want to secure that part of the path too, but the first use is going to be from the client to the intermediate recursive server. I want to give some kudos to this team. They did a fine job.

Paul Vixie: This was a fairly well-focused and uncontroversial working group within the IETF who stuck to their business and got something done that is a legitimate improvement not just in encryption, but it's a legitimate improvement over the protocols DNS was using before, which of course are UDP port 53 and TCP port 53. If you combine this, which is TCP port 853 with another thing called TCP Fast Open, then it's my view that you could replace UDP completely. You could just stop using UDP for DNS once you have DoT and TCP Fast Open.

Paul Vixie: That would be kind of cool because UDP/53 has got a lot of other problems like source address spoofing and DDoS and whatnot, but the thing that is best of all about their design is that they put it on its own port number and that means that if you have a policy, which is... would be common for an enterprise or a for any other managed private networks such as home network, it would be common to say that in my firewall I have a rule that only my name server can reach external name servers, and you would do that by restricting based on the board number you'd say, hey, if it's headed to some distant outside server on TCP port 853, then it has to be coming from one of my local name servers and if it isn't, then the firewall doesn't let it happen and this is perfect because that's how a lot of us have been controlling the old DNS protocols on UDP/53 and TCP/53.

Paul Vixie: They basically kept to the status quo. They said, we don't want to change the security perimeter. We don't want to move anything around. We want every control that used to be possible to still be possible, but we want it to be secure and we want it to be better engineered and so hats off to those guys. However, in parallel with that, a different group of people that are driven more by the web and less by DNS also part of the IETF, but different group of people, although that little bit of overlap decided that's great what you did there, and it's going to make some things better, but it's not going to solve our problem and our problem is that a lot of ISPs, including LTE providers, mobile IP providers are assuming the rights of a managed private network. They're actually doing the type of surveillance that we are afraid of and they are doing the type of answer substitution that we're afraid of.

Paul Vixie: Often this is for ad insertion or whatever and given that we've got these ISPs that are trying to behave more like managed private networks and then we've got ISPs that are in authoritarian regimes that behave very much like private networks, we need something that DoT does not provide, which is that we have to hide the DNS transactions inside of HTTPS where they are indistinguishable from other HTTPS traffic and are thus able to get out whether or not the network operator intends to block them and this is a huge controversy, right?

Paul Vixie: I don't like that there are these ISPs that are assuming the rights of managed private network operators and they're looking at my DNS traffic or they're perturbing it, they're sending me to advertising servers. They're doing things that I don't like and that that should stop and I do want to thank the Europeans for GDPR because in countries where GDPR is the rule of law, that is the law of the land none of that is legal and so that's not happening. If you're in GDPR, you do not need to worry about your ISP behaving badly for DNS, but in the rest of the world apparently it's fair game. Anyway, those were their circumstances and so they came up with their own protocol that puts DNS on top of the web and they did it deliberately to prevent on path interference in DNS operations, and I myself am not very comfortable with that.

Paul Vixie: I don't think that IETF ought to be telling me as a network operator what traffic I have to carry. That should be up to me to do and they especially ought not do it on purpose as they've done in this case, because this does disintermediate any CSO who might like to use DNS filtering to protect their network, but also disintermediate parents who are using DNS for parental controls and the reason is the DoH cannot tell whether you are a person who should be behaving this way or you are an ISP that should not be behaving this way, so they paint us all with the same brush and so I'm pretty uncomfortable about this and I just want to mention that if you're going to encrypt your DNS you probably have to support both of these protocols because some of your endpoints will only speak one or the other and you certainly should speak both of these protocols so that anyone who wants to encrypt their DNS traffic to you has the ability to do it.

Paul Vixie: However, this DoH thing is a new exfiltration opportunity for intruders and is going to call for some pretty significant remedies in your outbound proxies and in your outbound treatment of HTTPS.

Andrew Wertkin: To be fair to some of... or maybe even to highlight another potential issue, part of what the ISPs are doing is ensuring that the answer that comes back is the one that cost the least and by that I mean you're connecting to a backend service like Netflix or something and they have optimized their network to ensure that from a performance standpoint, efficiency standpoint or even cost to them, but usually that means performance standpoint, that you're connecting to the right Netflix and one of my issues with DNS over HTTPS is, now it's passing through that network operator who you're actually counting on to provide you that service. It's being answered by a server that might be far away. It might not be in the same geo location. It might be... you don't necessarily know where it is

and therefore you might get the wrong answer and by wrong answer again, I mean the answer that doesn't provide you the user with the most efficient path and best experience with that service.

Paul Vixie: Well, I think it's a very valid observation, but along those lines I want to point out that there's another group within the Internet Research Task Force, the IRTF is working on something that does what you just said, but much more deliberately and that is, it's the resolver list DNS project. The plan is that you won't have to ask DNS questions because all of the DNS bindings you might need to fetch all the images and JavaScript files and style sheets and so forth that are referenced by a webpage, all of the DNS bindings you need to fetch those objects will be embedded in the parent object and there will be no DNS sec information there by which a browser might know whether this is a defacement and this is fake information or whatever.

Paul Vixie: There's a lot of interest in speeding this up and it is not all compatible with the will of existing DNS operators and part of that is the problem that Andrew just described.

Andrew Wertkin: Yeah. We get a lot of questions and there are certainly some questions that have been asked during this webinar that are along the lines of, okay what do I do as an enterprise to try now to stop, block, remove access to DNS over HTTPS and I think part of the way it's being deployed, and there's two very well known entry points for DNS over HTTPs. Now, with the browsers and in one case there's a canary domain to shut it off. Inside an enterprise, if this specific domain resolves, then the browser won't use it and what I like more of though is the style of checking to see if there's a DNS over HTTPS service on the configured stub resolver and connecting to that or providing some level of these ones we know, we being a big corporation, but these ones that we know are... kind of have specific logging policies or are known to provide a good service.

Andrew Wertkin: It's one of the reasons frankly that BlueCat were going to allow and deploy in developing a DoH and DoT endpoints on our servers so that we can be that first hop and then basically be the termination of encryption so that we can still give the appropriate data to NetOps and SecOps because again, inside an enterprise there's not that assumption of privacy. You're at work. That's one way we're controlling it, but just in general, Paul when you get this question, because I'm sure you get it a lot. What do you tell people in terms of how to control the utilization of DoH inside their corporation or home?

Paul Vixie: Well. First to your comment about browser behavior, I agree very much that an application who wants to do DNS has to respect the system configuration. No browser should be selecting a DNS server for you no matter what motive they have or how much trust they might have in whoever they're choosing for you. They just should not do that. This is something that if you're going to try to use the DoH protocol to work with a name server that's already been chosen by the user, then that's a simple upgrade. That's just saying, I'm going to use a privacy

protected protocol to reach a destination that the user's already decided they like and I've got no problem with that, but no application whether a browser or anything else ought to be making that choice on behalf of the user because there's no such thing really as informed consent here.

Paul Vixie: Even if they were to use opt in instead of opt out, there's no way that you're... let's say my parents would be able to successfully navigate a choice like that. It is really something that the operating system ought to be in control of but as to the question, sort of what do we do about DoH? It begs the question, do we only have to worry about browsers, right? Because if we could argue with the browser vendors and other applications who want to do DNS, we could argue with our OS vendors as to how they're going to use DoH and DoT, then that's one set of concerns and one set of arguments and questions.

Paul Vixie: However, that is not the thing upper most in my mind. What's upper most in my mind is that every bot net from now on is going to be coded to use DOH, because if you're writing malicious software that you intend to infect people with, then of course you're not going to respect the operating systems choice, that would cause you to be monitored by the SOC team and who wants that. What we've done here is to create a new class of exfiltration risk that we can expect every intruder whether hardware, software or [meetware 00:56:41] is going to be using and so it's going to be broadly necessary for managed private networks to upgrade their outbound HTTPS treatment to say it's all got to use a proxy.

Paul Vixie: Everything's got to go through a proxy because I have to find out if somebody is trying to reach some distant name server by DoH and I have to stop that happening. Now, this is especially complicated by the recent release of TLS version 1.3 and the ability to now encrypt the SNI. SNI is kind of like the host keyword in the old HTTP system, and it's how you let a server know which of the many virtual websites that it hosts that you actually want to talk to. Which certificate do you want?

Paul Vixie: As soon as they're encrypting that then the traditional method of kind of a lightweight proxy for HTTPS where you sort of get in the middle of a TCP connection that wasn't strictly speaking headed for you, but you're the firewall, so you can answer it and then you just say, okay, I am the other end. Tell me what you want and then that gave this next gen firewall the opportunity to make some policy decisions possibly to enforce compliance regulations for the industry they're in.

Paul Vixie: That's not going to be possible once the world is using TLS 1.3 with encrypted SNIs and that means it's going to have to be an explicit proxy, probably socks or it could be an HTTPS proxy and you're going to have to force all of your outbound HTTPS traffic through an explicit proxy, basically strip searching everybody as they try and leave the building in digital terms just to make sure that you are in compliance with corporate policy, and I kind of hate that. That's

the wrong cultural direction for the Internet to go in but the choices that have been made by people who have the power to force those choices upon us lead us to no other future.

Andrew Wertkin: Yeah. It's sort of ironic that in the enterprise context, by using that explicit proxy, by decrypting everything, you're actually creating less privacy for those users than simply using a control point to stop them from going to a certain site for instance, or to execute your corporate policy like no web mail or not allowed to post anything over certain... whatever the case might be, so it sort of works against it that way. We've got a couple minutes left.

Andrew Wertkin: I just wanted to get to a couple more of the questions and these came in through our customer select channel, which is a great way that we're collaborating with customers these days. There's been a couple of questions on DSTs, you know, what about encrypting between the authoritative servers for things like zone transfers and anything else which is very different thing but for sure, Paul, I'm going to ask you your opinion, but it... where traffic can be encrypted, it should be encrypted and zone transfers is a perfect example of that.

Paul Vixie: Yes, zone transfers are a perfect example, but the cache miss traffic that flows between the intermediate recursive server and the distant content server also has to be encrypted and DoT was designed to be that protocol and we haven't seen broad adoption on that data path for this protocol but I think we will, and I do want to say that [Andre Sere 01:00:22] of Internet Systems Consortium posted to the DNS operations mailing list a few days ago saying we'd really like to do this but nobody wants to fund it, could somebody please offer a grant of some kind so that BIND 9 could have the DoT on the server to server transaction.

Paul Vixie: Anybody who might be thinking of trying to push us into that future should pick somebody who whether it's ISC or NLnet Labs or CZ.NIC Labs or the PowerDNS will pick somebody and help them build it because what you want is not just something you can use, you everybody to be using this because you will not be safer than them.

Andrew Wertkin: For sure. Let me see. I think most of the other questions I think you've either addressed during your talk track or along the same lines. Given that we're at the top of the hour, I think I'll wrap and thank you, Paul for your time. It was great to hear your thoughts and certainly get an update on what Farsight is doing and I encourage the conversation to continue with BlueCat and for any of the participants in any of these areas, whether it's via Slack or any other mechanism and let's continue talking about this because I think there's a lot changing now. There's things we're doing. There's things that you can do with or without us but it's a very... there's a lot of good conversation comes out of this, so thank you again, Paul.

Paul Vixie: Well, thank you again for inviting me. This has been terrific and I hope we do it again, so make it a regular thing.

Andrew Wertkin: Wonderful. Absolutely.

Erin: All right. Thank you both so much. I know we're at the top of the hour, so everyone's got a job. I just want to remind you all we will be sharing out the recording, so look for that in your inbox tomorrow and again, a big thank you to Andrew and to Paul for taking the time and having a lovely chat this hour. Thank you.