Comparison of Proposed Alerting and Data Collection Techniques

A critical component of the NCAP work involves the introduction of passive and active collision assessment as methods to collect name collision data. This introduction is likely to bring up a variety of questions regarding how these differ from each other and from controlled interruption. This short document offers a comparison of the different methods from the perspectives of alerting effectiveness, operational continuity, security and privacy, user experience, root cause identification, public response, and telemetry.

Overview

We first present a high-level overview of each technique. Subsequent sections describe the mechanisms in more detail as part of their comparison. Throughout the text, Phase 1 refers to the resolution of colliding names in the public DNS, while Phase 2 refers to the transport- and application-level communications that are dependent on and follow that resolution.

Passive Collision Assessment

With passive collision assessment, an application attempts to resolve a domain name (Phase 1), which results in a DNS query to the client's DNS resolver and—depending on configuration—to one or more public authoritative servers. The ultimate response from public authoritative servers is a negative response (e.g., name error or NXDOMAIN).

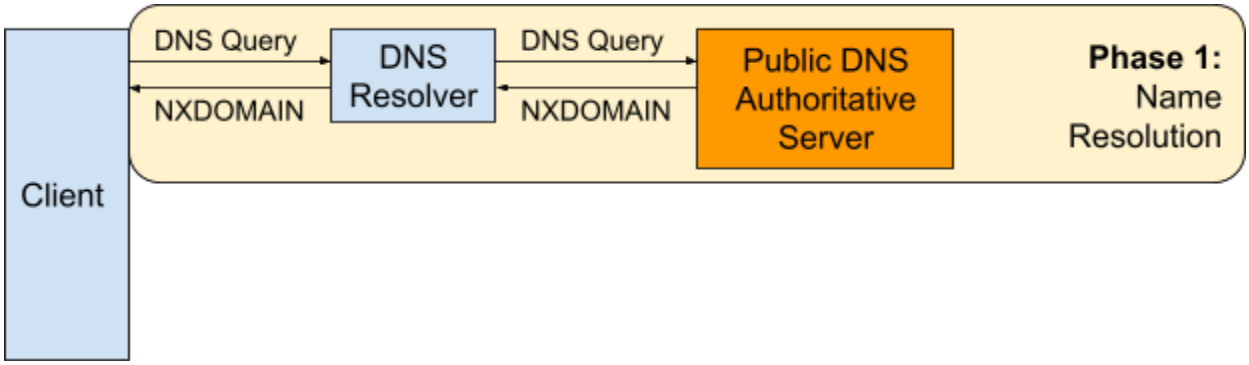

Controlled Interruption

With controlled interruption, if an application's attempt to resolve a domain name (Phase 1) results in a DNS query to public authoritative servers, the ultimate response from the authoritative servers is 127.0.53.53 for queries of type A (IPv4 address); queries of type AAAA (IPv6 address) result in a negative response. Any attempts by the application to communicate with that destination (Phase 2) will stay local to the client.

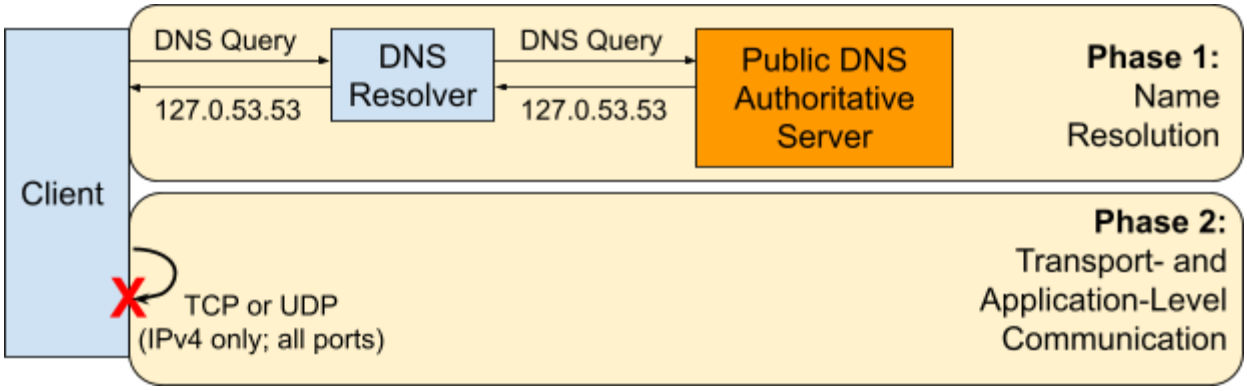

Active Collision Assessment

With active collision assessment, if an application's attempt to resolve a domain name (Phase 1) results in a DNS query to public authoritative servers, the ultimate response from the authoritative servers is the IPv4 or IPv6 address (for A and AAAA queries, respectively) of the server designated to handle application-layer requests from clients on select ports. Any attempts by the application to communicate with that destination (Phase 2) will leave the client.

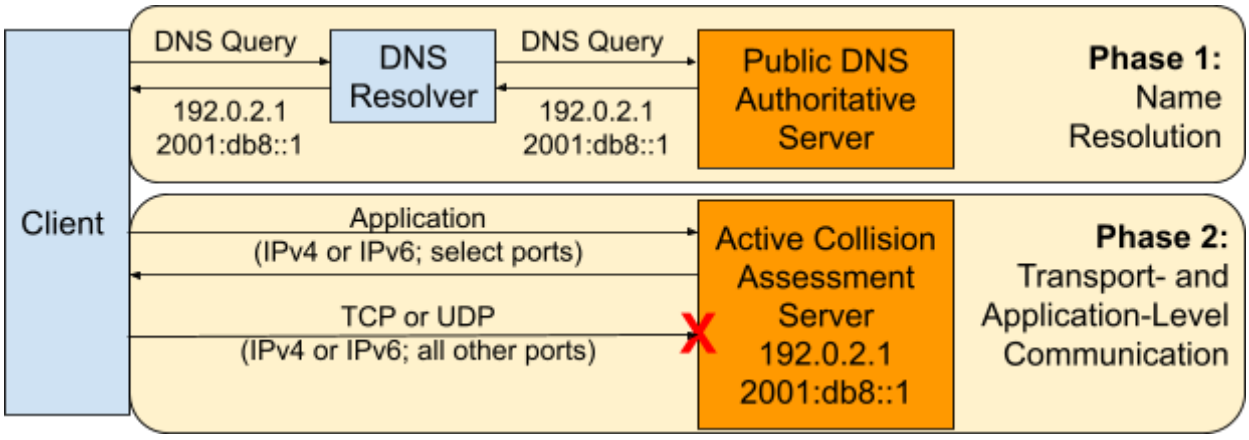

Alerting Effectiveness and Coverage

What population of potentially affected users, systems, and applications are expected to be reached by the alerting mechanism?

With passive collision assessment, there is no intent to alert end-user systems and applications. Because only negative responses are returned from authoritative servers, applications are not expected to behave differently. Thus, this section mostly applies to controlled interruption and active collision assessment.

DNS Resolution of Queried Names. One question involves the resolution of colliding names—whether or not they resolve to addresses in the public DNS (Phase 1). With passive

collision assessment, there are no positive answers for queried names, so the names do not resolve. However, with both controlled interruption and active collision assessment, *all* colliding queries reaching public authoritative servers are answered with the address appropriate for the mechanism. Thus, alert success is based on the likelihood of queries reaching the public authoritative servers. Two extreme cases are the following:

- No *queries reach public authoritative servers.* This might be the case, for example, if systems are on a corporate network for which the recursive DNS resolvers are configured to answer authoritatively for the colliding namespace *and* those systems never leave the corporate network.
- All *queries reach public authoritative servers.* This might be the case where the offending query is one of the intermediate queries issued in the course of search list processing and, prior to delegation, results in negative response.

The former case constitutes name collision *potential*, which would only alert only if one of the configuration requisites changes while the mechanism (controlled interruption or active collision assessment) is still deployed. The latter would affect *all* end systems using the colliding namespace. Other configurations, including variants of these, might result in alerting affecting some subset of systems some fraction of the time.

	DNS Resolution of Queried Names
Controlled Interruption	Resolution of queried names depends on DNS configuration and system mobility
Active Collision Assessment	Resolution of queried names depends on DNS configuration and system mobility
Passive Collision Assessment	Queries names do not resolve

Application Coverage. Resolution of a domain name is not only dependent on the query reaching the public authoritative server; it is also dependent on the capabilities of the application that initiated the resolution and the network connectivity of the system on which the application runs. Applications that support IPv4, when run on a system with IPv4 connectivity, will attempt to resolve the domain name to an IPv4 address; similarly, applications that support IPv6, when run on a system with IPv6 connectivity, will attempt to resolve the domain name to an IPv6 address.

Active collision assessment allows applications to resolve affected domain names to both IPv4 and IPv6 addresses. However, with controlled interruption, domain names can only resolve to IPv4 addresses. Because of this, only applications with IPv4 connectivity will be affected by controlled interruption. Applications and systems that are IPv6-only will neither resolve colliding domain names to IP addresses (Phase 1) nor attempt application-level communication (Phase 2) with controlled interruption.

The absence of IPv6 in controlled interruption is discussed in section of 3.1.3 of the JAS report[1]. That discussion concentrates on two questions: the need for IPv6; and the IPv6 address that would be used.

With regard to the need for IPv6, the question is the prevalence of IPv6-*only* systems, i.e., where the IPv4 mechanism would not even be a possibility. The JAS report discusses the practicalities, expectations, and even measurements associated with IPv6-only systems. From that standpoint, they conclude that controlled interruption responses for IPv6 addresses are unnecessary. Unfortunately, the data from the report is not helpful in supporting this claim. The measurement of DNS *resolvers* that "appear to be" IPv6-only simply cannot be used to approximate end-user systems, which are the primary candidates for name collisions and controlled interruption; they are simply two very different things. Additionally, the report's reference to the number of systems accessing Google over IPv6 is not supportive, as that does not imply *IPv6-only* activity (even so, the percentage of IPv6-capable clients has risen from 3% to 40% since the JAS report was written[2]). Despite the lack of self-presented evidence supporting the claims from the JAS report, there is data that suggests that it is impractical to expect a measurable presence of IPv6-only systems. Among that supporting data, only 26% of the top 500 Web sites have AAAA records published[3]. NAT64 and related systems are a special consideration. In a NAT64 environment, a client is effectively IPv6-only, relying on special infrastructure to make IPv4-only resources available to IPv6-only clients. A study of how applications on NAT64 systems behave when confronted with name collisions and controlled interruption is desirable but beyond the scope of this work.

With regard to the IPv6 address that might be used for controlled interruption, the JAS report authors considered ::1 (loopback), ::53 (public), addresses within fd00::/8 (site-local), and addresses within fe80::/10 (link-local). However, because of the "the potential for unintended consequences" and the relative immaturity of IPv6 implementations compared to IPv4 implementations, the JAS report recommended avoiding the risk associated with "experimenting in the 'fringes' of v6 for what is very likely a small benefit."[4] More experimentation and analysis can be done to test controlled interruption using addresses from these and possibly other ranges, but it is beyond the scope of this work.

In summary, there is currently no IPv6 address for controlled interruption. Such would require additional considerations and, perhaps most importantly, testing. The need for IPv6 in controlled interruption seems low, but remains unclear without thorough studies, including in environments like NAT64.

	Application Coverage
Controlled Interruption	Only applications using IPv4 are affected

[1] JAS Report (https://www.icann.org/en/system/files/files/name-collision-mitigation-study-06jun14-en.pdf)
[2] https://www.google.com/intl/en/ipv6/statistics.html Last visited Aug 26, 2022.
[3] http://www.delong.com/ipv6_alexa500.html Last visited Aug 25, 2022.
[4] JAS Report.

Active Collision Assessment	Applications using either IPv4 or IPv6 are affected
Passive Collision Assessment	No applications are affected

Operational Continuity, Security, and Privacy

How might users or systems be negatively impacted by interruption to service or subjected to exploit or privacy violations?

There are five levels of impact to consider with regard to operational continuity, security, and privacy.

DNS Query Surveillance. Colliding DNS queries (Phase 1) are observed by public authoritative servers, as well as Internet service providers, and any other operators of infrastructure on the path between recursive and authoritative servers. In the case of passive collision assessment, only a *fraction* of both name collision queries and query names are leaked to public authoritative servers because of negative caching at the DNS resolver. With controlled interruption and active collision assessment, all *query names* will be observed. The *rate* of queries observed for any given query name is a lower bound of the rate observed by the recursive resolver, again because of caching.

Communication Interruption. Application communications (Phase 2) are interrupted in connection with the intent to alert. *All* cases of controlled interruption and *select* cases of active collision assessment result in this level of impact. In the case of active collision assessment, communications to ports designated to receive incoming communications result in communication interception (discussed hereafter); all others result in communication interruption and application inference (discussed next). Examples of interruption caused by controlled interruption are well documented in the reports submitted to ICANN via their Web submission form and in the Root Cause Analysis document.

Application Inference. Application-layer protocols and sometimes the applications themselves can be inferred by observation of destination (TCP or UDP) port in application-layer communication attempts (Phase 2)—whether or not those attempts are intercepted. This level of impact applies to all cases of active collision assessment.

Communication Interception. Application communications (Phase 2) are intercepted. This will *never* be the case with controlled interruption because, by design, communications destined for 127.0.53.53 will never leave the local system, where they might be intercepted by foreign systems. With active collision assessment, select ports are designated to accept incoming communications, and application-layer data is exchanged between client and server.

Data Exfiltration. Potentially sensitive application-layer data is sent to the server that intercepted application-layer communications (Phase 2). Thus, clients subjected to active collision assessment are subject to data exfiltration, and only in the case where ports have been configured to accept incoming communications. In some cases the data sent might be innocuous. For example, in an SMTP transaction, the initial communication from the client (after the greeting from the server) is merely a HELO or EHLO message, used to identify the client to the server. HTTP, on the other hand, is a request-response protocol, in which the client sends application-layer data first (the request) before the server sends application-layer data (the response). An HTTP request contains information that might be considered sensitive, including the path being requested, the values of any query string name-value pairs, cookie values, and request data associated with POST request. Thus, with HTTP, by the time the server responds, the potentially sensitive information has already been transmitted.

	Operational Continuity, Security, and Privacy
Controlled Interruption	**DNS Query Surveillance:** all qnames **Communication Interruption:** all **Application Inference:** none **Communication Interception:** none **Data Exfiltration:** none
Active Collision Assessment	**DNS Query Surveillance:** all qnames **Communication Interruption:** all **Application Inference:** all **Communication Interception:** select **Data Exfiltration:** select
Passive Collision Assessment	**DNS Query Surveillance:** all SLDs, fraction of qnames **Communication Interruption:** none **Application Inference:** none **Communication Interception:** none **Data Exfiltration:** none

User Experience

What is the experience of the end user, in terms of application behavior, path to resolution, etc?

With passive collision assessment, only negative responses are returned from authoritative servers (Phase 1). Therefore, applications are not expected to behave differently for the user. One possible exception to this is the case where applications querying for the TLD itself expect an NXDOMAIN (rather than a NODATA) and behave differently because of it. We do not know of any such applications, but we note the possibility for completeness.

That being said, this section mostly applies to controlled interruption and active collision assessment, which involve transport- and application-level communication (Phase 2). The user

experience is dependent in part on what is experienced by the application. The application experience, in turn, is dependent on what type of name collision configuration is in place. We consider communication interruption and communication interception as two of those configurations.

Communication Interruption

Applications communicating with a system whose intention is to interrupt might result in one of many behaviors, depending on factors such as the transport-layer protocol used (TCP or UDP), kernel-specific routing or access policy, firewall and endpoint protection software, and the application itself. One consideration is the timing of an interruption error, for which we give two primary outcomes:

- *Quick-Response Error.* The application detects the communication error relatively quickly and possibly notifies the user. Such errors are consistent with TCP RST packets, which come from the kernel of the "server" to which the application attempted to connect, as well as ICMP port unreachable packets, which are typically sent by a kernel in response to UDP messages destined to a port that is not listening.
- *Timeout.* A potentially lengthy period of time passes before the application detects the error and possibly notifies the user. This is because neither a TCP RST nor an ICMP port unreachable message are received, so the application must wait for communications to time out.

Quick-response errors are expected almost exclusively when controlled interruption is in use. This is because: 1) hosts subject to controlled interruption only communicate with the loopback interface of the host itself, never leaving the host or the network; 2) firewalls on the loopback interface typically do not make sense, allowing queries to be received by the kernel; and 3) the kernel is the responder, doing so in one the two ways described previously.

For active collision assessment, the timing of the error response experienced by the application depends on the configuration of the network path between the client and the server and the destination port. Here are several scenarios:

- *Stealth Firewall*. If an intervening firewall, anywhere on the path, including the server itself, drops packets associated with communication to a given port, then the application will experience a timeout.
- *Active Firewall*. If an intervening firewall responds with a TCP RST or an ICMP error message, then the application will experience a quick-response error.
- *Server Rejection*. If no firewall intervenes, and the server is not listening on a given port, then the kernel responds with a TCP RST, and the application will likely experience a quick-response error.
- *Communication Interception.* If no firewall intervenes, and the server is listening on a given port, then this constitutes communication interception and is covered in the next section.

While we have attempted to enumerate the errors that might be experienced by users, the list should not be taken as definitive for two reasons. First, while the application will almost certainly experience the timeout or quick-response error associated with transport-layer communication issues, how the application handles that response varies, and what the user

sees might be different. Second, in some cases the user experience does not come directly from the application that experiences the communication interruption; rather, their experience is with an application that depends on the application experiencing the interruption. For example, one name collision report submitted to ICANN described clients "freezing" when they encountered controlled interruption. Without additional qualitative data regarding the experiences of users, we can only speculate.

	Error Response - Application Experience
Controlled Interruption	Quick-Response Error
Active Collision Assessment	Quick-Response Error or Timeout, depending on network configuration and application port
Passive Collision Assessment	No Error

	Error Response - User Experience
Controlled Interruption	Application Dependent
Active Collision Assessment	Application Dependent
Passive Collision Assessment	No Error

Communication Interception

As mentioned previously, communication interception applies only to active collision assessment. We describe the user experience from the perspective of several different clients, protocols, and ports.

- *Web Browser / HTTP.* If the server listens on the standard Web port (80), and the browser requests Web content over HTTP (no TLS), then the server will return whatever content it wishes in response, i.e., about the name collision. The user will almost certainly see different content than they were expecting.
- *Web Browser / HTTPS.* If the server listens on the standard HTTPS port (443), it is expected that browsers will initiate a TLS handshake. At the point in which the browser receives the server's certificate, most browsers will attempt to validate it, using the domain name in the Uniform Resource Locator (URL). In the case that the certificate is successfully validated by the client, two things can be said: 1) the server can return whatever content it wishes in response, i.e., about the name collision; and 2) the browser and the user will have a greater trust in the server because of the validation. However, in the case that the certificate does *not* validate, the browser will prompt the user with a warning about the certificate, discouraging them from continuing with a button like "I accept the risks and wish to proceed anyway."

While the details of a proposed active collision assessment implementation are not included in this study, there are some practical limitations that must be considered. First, while there is precedent for a TLS certificate having many domain names, including wildcard domains, having a wildcard domain with the asterisk (*) immediately below a TLD is rare if even possible. Additional research would need to be done to conduct the feasibility of such a setup. Even so, the asterisk (*) in any wildcard domain can only be substituted for a single label. Thus, even with a wildcard, it is infeasible for the server to have a certificate that includes all domain names that might resolve to the server's address with active collision assessment.

	User Experience - HTTP / HTTPS Browsers
Controlled Interruption	Not applicable
Active Collision Assessment	**HTTP:** unexpected content received **HTTPS:** TLS certificate errors
Passive Collision Assessment	Not applicable

- *Other clients and protocols.* For applications other than Web browsers and protocols other than HTTP, user experience depends on the application. For example:
    - A non-browser client attempting to access content from (or upload content to) a server over HTTP might not get the desired HTTP response—whether it is status code, header value, or response body—causing it to fail. The user experience depends on how automated the process is and how notifications are configured.
    - A non-browser client attempting to access content from (or upload content to) a server over HTTPS might abort on validation of the server's TLS certificate. Again, the user experience depends on the extent of process automation and notification configuration.
    - An SSH user attempting to connect to a server whose identity is already known will be met with a warning message that "IT IS POSSIBLE THAT SOMEONE IS DOING SOMETHING NASTY!" If the user pushes through anyway, they would fail to login without a legitimate account, but only after having saved the intercepting (false) server's public key.
    - A mail user agent attempting to connect to a mail server over IMAP, Outlook Web Access, POP3, or the like, would encounter account issues, due to either a bad TLS validation or a login failure.

	User Experience - Other Clients and Protocols
Controlled Interruption	Not applicable
Active Collision Assessment	**Non-browser HTTP:** unexpected content received, other unknown errors **Applications that use TLS:** TLS certificate errors

	SSH: man-in-the-middle attack errors
Passive Collision Assessment	Not applicable

Local firewall alerts
Some firewall software raises alerts about anomalous connection attempts and/or prompts the user for permission to allow communications to proceed, even if that communication originates from the local system. This behavior would only be manifest with controlled interruption. This behavior is expected to be rare, but is known to happen. One example was found in the Web search results for 127.0.53.53, documented in the Root Cause Analysis document.

	User Experience - Local Firewall Alerts
Controlled Interruption	Rare but possible
Active Collision Assessment	Not applicable
Passive Collision Assessment	Not applicable

Root Cause Identification

How useful is the technique in leading users towards the root cause and a possible resolution?

Because root cause identification implies that users or systems have experienced some sort of alert by way of interruption or interception (Phase 2), this really only applies to controlled interruption and active collision assessment.

Controlled interruption and active collision assessment have distinct strategies for helping users—and ultimately system administrators—identify the root cause of the problems they have experienced because of name collisions. Active collision assessment presents the application—and, where possible, the user—with an explicit message (i.e., actual text) that they have encountered a name collision, returned as application-layer content in the server response. The most obvious example is a Web browser attempting to retrieve content over HTTP or HTTPS, but the server returns the custom message about name collisions instead. However, as mentioned previously, retrieving this content is fraught with challenges related to the user experience. In particular, most anything other than a Web browser over HTTP suffers from warnings that are deterring at best and possibly impassible. Controlled interruption makes no attempt to intercept client-server communications but instead leaves a “hint” to the user or system administrator by way of a specific IP address (127.0.53.53), which, if the subject of a Web search, *should* result in information about ICANN, name collisions, and controlled interruption.

Both approaches have their advantages and disadvantages. Finding the hints left by controlled interruption requires technical expertise that is well beyond that of a typical user; users in that

scenario might simply need to get support for a computer, application, or network that "doesn't work." In the case of active collision assessment, even if a user is able to retrieve and read the content prepared by the active collision assessment server server—with or without technical challenges—there remain questions about how effectively that content will be processed and acted on by the user. For example, with the message presented, would the user know which steps to take next, whom to contact, etc? In both scenarios, the problem (hopefully) finds its way to system or network administrators. In the best case, those administrators are knowledgeable and are able to track down the relevant information and resolve the problem—not just for the short term, but also for the long term. In the worst case, the root cause is not properly identified, and the problem not appropriately fixed. This might be because the message isn't sufficiently understood, or because the problem is difficult to reproduce. The survey results in the Root Cause Analysis document showed that 50% of those that observed 127.0.53.53 successfully associated them with ICANN. The analysis of the Web search results for "127.0.53.53", also contained in the Root Cause Analysis document, showed a success rate of 76%. However, for active collision assessment, only a properly designed and executed user study could properly measure the effectiveness of both the user experience and the messaging.

While we have no data specifically related to active collection assessment, we do have three collections of data that provide insights into the relative effort and intuitiveness associated with identifying the cause of a name collision with controlled interruption, all from the Root Cause Analysis document: 1) the name collision reports submitted to ICANN via their Web form[5]; 2) the results from a Web search for "127.0.53.53"; and 3) responses to the survey issued to network operators as part of the root cause analysis. With the name collision reports submitted to ICANN, the submitters obviously found the Web page for the submission form. However, it is unclear whether finding that page came from observing the controlled interruption IP address (127.0.53.53) or from other troubleshooting measures. We also cannot say anything about cases in which name collisions were experienced but the form was not discovered or cases in which the users found the page but were deterred from making a submission because of the statement on the form indicating that only those experiencing "a clear and present danger to human life" should submit a report. Nevertheless, we can learn several things from the reports with regard to root cause identification.

First, relatively few of those that reported name collisions observed the controlled interruption IP address. Of the 34 reports, only eight (34%) mentioned "controlled interruption" or referred to the IP address "127.0.53.53." We cannot definitively conclude that reports that did not include such references did not observe the controlled interruption IP address. However, responses to the survey support these low numbers: only 28% of those that experienced name collision problems observed the controlled interruption IP address, and only half of those found it useful in identifying the root cause. The Web search results showed a slightly higher figure of those that made the association between the controlled interruption IP address and ICANN: 76% of those that observed 127.0.53.53 also mentioned ICANN.

[5] https://www.icann.org/en/forms/report-name-collision

Second, an analysis of the reports includes very little evidence to suggest that Web browsers are a common way for name collisions to be manifest. The root cause analysis summarizes the experiences of individuals and organizations that submitted name collision reports and rates the level of impact described in those reports. Seventeen (50%) of the 34 reports were rated as either having “severe” or “significant” impact. Of those 17, 14 (82%) did not mention specific applications but described more general problems. For example: “Network down”; “causing … laptops to crash”; “cannot resolve DNS”; and “Users cant loggon to local domain.” The remaining three (18%) mentioned applications more specifically: mail and network shares. Only two reports—from the remaining 10 reports categorized as having “small-scale” impact—possibly referred to Web browsers: “Internet browsing issues from LAN”; “can't access to some servers.” We are limited to the data that was submitted in the reports and have had no avenue to pursue follow-up information. Nevertheless, it is backed up by the analysis of the Web search results for “127.0.53.53,” which shows that only 20% of results for which an application was identified were related to Web browsers. Based on this data, we expect that alerts that would be noticed in browsers are in the minority, diminishing the effectiveness of root cause identification with active collision assessment.

	Root Cause Identification
Controlled Interruption	**Low** - hint often not observed (34%) or not understood (24% - 50%)
Active Collision Assessment	**Low** - name collisions experienced in Web browsers are few (12 - 20%)
Passive Collision Assessment	Not applicable

Public Response

In what ways might the techniques be received in the public, with ICANN and others being accountable for complaints and fallout associated with design and execution of the mechanism?

This section is intended to provide a comparison of the possible public reception of the techniques, based on the experience of end users, system administrators, or other parties with the deployment of the techniques themselves or similar mechanisms.

As previously mentioned, we do not expect passive collision assessment to affect the user experience, and we expect any impact on security and privacy to be negligible at best. Therefore, we do not expect expressions of sentiment of any kind associated with passive collision assessment; in fact, we hardly expect passive collision assessment to be noticed. Our description hereafter applies to controlled interruption and active collision assessment.

Because we have eight years of deployment experience with controlled interruption, we have some insights into the public reception associated with its deployment. The results of the Web

search for “127.0.53.53” were analyzed for sentiment, as documented in the Root Cause Analysis document. In 94% of cases, the public comments surrounding a name collision were neutral in nature—neither positive nor negative. Only one result (6%) conveyed a very negative sentiment, towards both ICANN and the registry of the affected TLD. Based on this data, we suspect that controlled interruption is generally low risk with regard to public reception.

On the other hand, we have no deployment experience with active collision assessment. Therefore, we can only refer to experiences associated with the deployment of similar proposals and deployments from the past, while being careful to distinguish key differences. VeriSign’s Site Finder effectively employed a technique almost identical to that being proposed[6]. The major difference is that Site Finder introduced a wildcard A record into the `com` and `net` zones, while active collision assessment proposes introducing a wildcard into a TLD not previously delegated. Thus, active collision assessment is expected to affect a much smaller and more targeted population of users—namely those with name collision issues. The public outcry associated with Site Finder was extensive. The same day that it was rolled out, there were calls to submit complaints to VeriSign, ICANN, and the United States Department of Commerce, patches to DNS resolver code to bypass or blacklist the Site Finder mechanism, and even calls inflict a denial-of-service attack on VeriSign using their own framework[7][8]. Thus, the general sentiment of Site Finder was negative.

We emphasize that active collision assessment has a different motivation, comes at a different time, and affects a smaller and more targeted population of systems and users than Site Finder. Nonetheless, it provides useful insights in evaluating potential public reception because of its similarity in technique to active collision assessment.

	Public Response
Controlled Interruption	**Neutral (94%)**, based on actual deployment experience
Active Collision Assessment	**Unknown, Possibly negative**, based on experience with Site Finder
Passive Collision Assessment	**No reactions anticipated**

Telemetry

How much data is available to investigative parties, and what type of effort will it take to collect and analyze it?

[6] https://web.archive.org/web/20041109202247/http://www.verisign.com/static/002702.pdf
[7] https://slashdot.org/story/03/09/16/0034210/resolving-everything-verisign-adds-wildcards
[8] https://mailman.nanog.org/pipermail/nanog/2003-September/166467.html

Earlier sections describe the reach of each data collection technique as well as their impact on security and privacy. These same attributes can be used to assess their effectiveness in terms of data collection.

DNS Queries. With all three techniques, DNS queries (Phase 1) can be collected by the DNS servers that are configured to answer authoritatively for the given TLD namespace. In the case of passive collision assessment, a fraction of DNS queries and query names will be collected, based on the average query rate and the negative cache time-to-live value. In the case of controlled interruption and active collision assessment, *every* name queried will be observed at the authoritative server, even though the rate of queries for given query names by end systems is masked by the caching behavior of DNS resolvers.

Because collection of query data at root servers is already common practice, the value of passive collision assessment might be called into question. In the context of telemetry, we now address some of the advantages of passive collision assessment over simply collecting DNS queries at the root DNS servers.

- *Real-time Availability.* The root servers are operated by 12 different organizations, and each has their own policies with regard to DNS query logging. This includes the specific query components that are recorded, how long data is maintained, whether or not it is anonymized, and how it is shared. Thus, root server data is not generally available to third parties. An exception to this is the annual collection of DNS queries carried out by many root server operators and known as the Day in the Life (DITL). DITL data is made available to members of the DNS Operations, Analysis, and Research Center (DNS-OARC). In contrast, authoritative DNS servers operated in connection with passive collision assessment would facilitate continuous and near real-time analysis—not simply data with year granularity.
- *Consolidated Control.* Considering that 12 distinct operators operate the root servers, getting relatively comprehensive query information—outside of DITL—would require cooperation, coordination, and consent from each operator. However, by delegating a TLD to specific authoritative servers designated for passive collision assessment, no coordination or consent is needed.
- *Increased Query and qname Volume.* The DNS servers associated with passive collision assessment are configured to answer authoritatively for a more specific namespace than the root zone. The impact can be explained as follows. A resolver that learns that a *TLD* does not exist can infer without further queries to the root that *subdomains* of that TLD do not exist. Therefore, the resolver doesn't return to the root server with a query under that TLD until the negative cache value for the root zone expires. However, when the TLD *does* exist (i.e., with passive collision assessment), the resolver must query the servers authoritative for the TLD for every *second-level domain (SLD)* that it doesn't know about. Once the resolver learns that the SLD does not exist, it does not return to the TLD authoritative servers again for query names under that SLD until the negative cache value for the TLD expires. Thus, while only a fraction of query names are observed in either case, the root might only observe one query per negative cache value for a given TLD and resolver, while passive collision might observe one query for a given

SLD and resolver, in the same time period. Additionally, the negative cache value of the TLD is under control of the TLD operator. Thus, it could be reduced to increase the query rate.

- *Reduced QNAME Minimization Effects.* Many DNS resolvers have implemented a feature known as QNAME minimization, in which only the minimum labels are issued in a query to an authoritative server, reducing the authoritative server's visibility into full query names being issued. At the root servers, only the TLD would be seen in the strictest cases of QNAME minimization. At the TLD authoritative server, however, the SLD would be observed in such cases.
- *Reduced Aggressive Negative Caching Effects.* Many DNS resolvers have implemented aggressive negative caching for DNS zones signed with DNSSEC. With aggressive negative caching, a resolver can *infer* domain names that do not exist with hints provided gratuitously by authoritative DNS servers in connection with queries for *other* domain names. The result is that queries that might otherwise be asked of authoritative servers can be answered by the resolver without querying the authoritative servers. This yields fewer queries and fewer query names from resolvers that support aggressive negative caching. However, this only applies to DNS zones that are DNSSEC-signed. Thus, it applies to the root zone and root servers, but as long as TLDs to which passive collision assessment is applied are not DNSSEC-signed, aggressive negative caching does not apply.

IPv4/IPv6. With active collision assessment, data can be collected with regard to which IP address family is used to attempt application-layer communication (Phase 2). This is not possible with either controlled interruption or passive collision assessment.

Transport-Layer Protocol and Ports. With active collision assessment, the transport-layer protocol and destination TCP or UDP port can be collected to infer the application-layer protocol with which communication is being attempted (Phase 2). This is not possible with either controlled interruption or passive collision assessment.

Application-Layer Data. With active collision assessment, not only can the destination port be observed, but also some amount of application-layer data, depending on the protocol and the logging configured (Phase 2). For example, for an HTTP/HTTPS request from a Web browser, the whole HTTP request could be logged, including request method, path, user-agent, and more. This is not possible with either controlled interruption or passive collision assessment. Additionally, for TCP-based communications, once a connection has been established, there is reasonable assurance that the client IP address was not spoofed by an off-path entity.

	Telemetry
Controlled Interruption	**DNS queries:** all qnames; end-system query volume masked by caching **Application:** no telemetry

Active Collision Assessment	**DNS queries:** all qnames; end-system query volume masked by caching **Application:** IPv4 and IPv6; TCP/UDP usage and destination ports; application-layer data
Passive Collision Assessment	**DNS queries:** all SLDs, fraction of qnames, end-system query volume masked by caching **Application:** no telemetry

Generated Measurements of Collision Potential

Deployment of the proposed alerting and data collection techniques will result in telemetry data consistent with their respective capabilities. As colliding namespaces are used by end systems and users, and the queries reach public authoritative DNS servers, the activity is captured in DNS query and application logs. However, as noted in the section on “Alerting Effectiveness and Coverage”, there are network environments in which DNS queries *would* collide—should they be allowed to reach public authoritative DNS servers—but the network configuration of these systems prohibits those queries from reaching public authoritative DNS servers. The proposed data collection techniques currently have no way to measure this name collision potential.

To address this possible gap, two additional measurement techniques have been proposed, which will test for the private use of namespace within a network. Both techniques use the same configuration as passive collision assessment. That is, the TLD is delegated, but any queries associated with the TLD result in NXDOMAIN. Thus, these measurement techniques can be considered in the same light as passive collision assessment.

Overview

We now discuss specifics of each measurement technique and assess how each would contribute to name collision telemetry.

RIPE Atlas probes. The RIPE Atlas platform has thousands of measurement probes embedded in “host” networks around the world. These probes can be used by researchers to run measurements, which include issuing queries against the DNS resolvers designated for that network—that is, the resolvers used by other systems on the host network. The proposed measurement involves issuing queries under a given TLD to the DNS resolvers used by a RIPE Atlas probe and observing the authoritative DNS server logs for those queries. Networks for which queries are not observed at the authoritative servers can be inferred as using those TLDs as a private namespace.

Ad-based measurements embedded in Web pages. Special ads, embedded in Web pages, can cause the browser that renders them to fetch a given resource (e.g., an image) over HTTP/HTTPS. Fetching the resource requires a DNS lookup using the DNS resolvers used by

the browser. The ads themselves show up in browsers of users world-wide, based on the algorithms of the ad company. In addition to their primary purpose of promotional marketing, ads can be used to run Internet-related measurements. The proposed measurement involves placing an ad that requires the browser to fetch a resource hosted at a domain name under a given TLD and observing the authoritative DNS server logs for queries associated with requests for that resource. Networks for which queries are not observed at the authoritative servers can be inferred as using those TLDs as a private namespace.

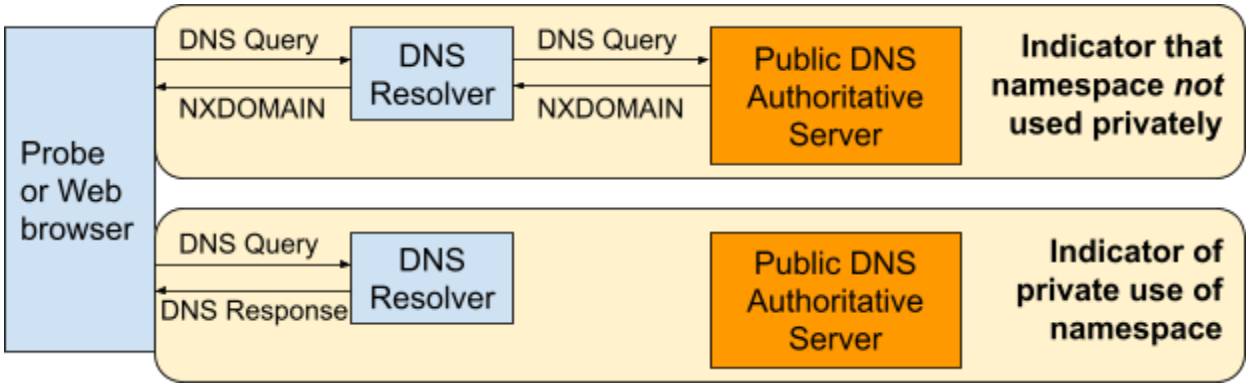

Limitations

The proposed measurement techniques promise to enhance the telemetry associated with name collisions, in particular where name collision potential could not otherwise be identified and quantified. Nevertheless, there are limitations with the techniques, which we now discuss.

- While the measurement techniques would identify *networks* for which TLDs are being answered internally—rather than communicating with the public DNS—the resulting measurement data does not necessarily reflect actual activity by *end users and systems*.
- Although several different configurations and usage models result in name collisions, these measurement techniques only address a subset of those—in particular those that involve private use of namespace.
- Queries observed at authoritative DNS servers—both TLD and root servers—will include queries from both actual end systems *and* the active measurements herein proposed. Without further filtering and processing, the queries from the active measurements will affect the data and metrics associated with “normal” behavior. At the very least, the two types of queries should be made distinguishable from one another to make accurate and meaningful assessments of the data. This is possible at the TLD authoritative servers by using query names whose second label is distinguishable. However, for measurements at root servers, this might not be possible due to a growing percentage of resolvers that use qname minimization and for which the second label will not be visible.
- For the RIPE Atlas measurements, not all probes will point at DNS resolvers that are used by end users and systems.
- For the RIPE Atlas measurements, data will only be gathered for networks that host a RIPE Atlas probe.
- For the ad-based measurements, not all browsers are configured to use the DNS resolvers associated with the network to which they are connected, so their experience may or may not be typical of the network.

- qname minimization has been deployed in many DNS resolvers across the Internet. As such, observing the full query name at the public authoritative DNS servers can only be expected a fraction of the time. Thus, any identifiers associated with query names must be embedded in the second label.