ICANN Registry Services Technical Evaluation Panel

Report on Internet Security and Stability Implications

of the

Tralliance Corporation

search.travel Wildcard Proposal

November 2, 2006

Preface

This report presents the findings of a technical evaluation of the proposal[1] by Tralliance Corporation to introduce a new service called "search.travel" into the operation of the .travel top-level domain.

On 8 November 2005 ICANN adopted[2] a consensus policy developed by its Generic Names Supporting Organization (GNSO) concerning the review and approval of requests by gTLD registry operators for new registry services.[3] This policy was implemented on 25 July 2006[4] as the Registry Services Evaluation Policy.[5] The policy provides for the evaluation of a proposed registry service by a team of experts selected from a standing Registry Service Technical Evaluation Panel (RSTEP)[6] when ICANN determines that the service could raise significant security or stability issues.

The process begins with a preliminary determination by ICANN that an RSTEP review is or is not required for a particular proposed registry service.[7] If ICANN determines that a review is required, an RSTEP review team investigates and evaluates the proposed service with respect to its potential impact on security or stability, as defined by the consensus policy:

> **Security**—An effect on security by the proposed Registry Service shall mean (a) the unauthorized disclosure, alteration, insertion, or destruction of Registry Data, or (b) the unauthorized access to or disclosure of information or resources on the Internet by systems operating in accordance with all applicable standards.
>
> **Stability**—An effect on stability shall mean that the proposed Registry Service (a) is not compliant with applicable relevant standards that are authoritative and published by a well-established, recognized, and authoritative standards body, such as relevant Standards-Track or Best Current Practice RFCs sponsored by the IETF, or (b) creates a condition that adversely affects the throughput, response time, consistency, or coherence of responses to Internet servers or end systems operating in accordance with applicable relevant standards that are authoritative and published by a well-established, recognized, and authoritative standards body, such as relevant Standards-Track or Best Current Practice RFCs and relying on Registry Operator's delegation information or provisioning services.

[1] http://www.icann.org/registries/rsep/tralliance_request.pdf
[2] http://www.icann.org/minutes/resolutions-08nov05.htm
[3] The ICANN Board resolution adopting the GNSO consensus policy (see footnote 2) specifies that implementation of the policy in contractual terms should be guided by the provisions of the .NET registry agreement (http://www.icann.org/tlds/agreements/net/net-agreement-new.html), which includes a precise definition of "Registry Services."
[4] http://www.icann.org/announcements/rsep-advisory-25jul06.htm
[5] http://www.icann.org/registries/rsep/rsep.html
[6] http://www.icann.org/registries/rsep/rstep.html
[7] The consensus policy also provides for the separate review of potential competition issues, which lie outside the scope of the RSTEP review.

The review team completes its evaluation within 45 days, and prepares a written report of its findings, containing:

> (a) a detailed statement description of the technical issue(s) raised by the proposed registry service, and the assumptions, information,[8] analysis, reasons, and information reasoning upon which the panel review team's evaluation is based;
>
> (b) the team's expert assessment of the potential impact of the proposed registry service on security or stability; and
>
> (c) a response to any specific questions from ICANN that were included in the referral from ICANN staff in its request for the RSTEP review.

The review team's report is delivered to the ICANN Board as input to the Board's consideration of the proposed registry service and action on the registry operator's request to deploy the service within the context of its contract with ICANN.

It is important to recognize that the RSTEP review is a technical evaluation of a proposed registry service with respect to the likelihood and materiality of effects on security and stability, including whether the proposed registry service creates a reasonable risk of a meaningful adverse effect on security or stability. Because many other questions and issues may be relevant to the overall assessment of a proposed registry service, it is not a recommendation to the ICANN Board concerning whether or not the Board should approve or reject the registry operator's proposal.

[8] RSTEP review teams are expected to gather information from as many sources as necessary in order to conduct a thorough and comprehensive evaluation, including, but not limited to, information provided by the registry operator, by ICANN, and by contributors to the ICANN public comment forum that is associated with each registry service request.

Contents of the Report

1. Introduction

1.1 The Tralliance Proposal

Tralliance Corporation proposes to introduce a wildcard[9] at the apex of the .travel sTLD (that is, *.travel) in order to redirect queries for name strings that are not found in the .travel zone to a web form that:

- advertises the availability of the name string for registration as a .travel domain name; and,
- provides a search box pre-loaded with the name string which, if used, would return .travel results with higher rankings than results from other TLDs.

Tralliance compares this service to the wildcard currently in operation at the apex of the .museum TLD.

The proposal has two components:[10]

(1) The insertion of a wildcard resource record into the .travel zone of the DNS:

"To effect the wild card redirection, we will insert a wild card into the apex of the .travel zone. The use of wild cards is well documented in various RFCs. Upon inserting the wild card, all DNS queries for any domain not found within the zone will receive a response containing the IP address of the search.travel web site."

(2) An HTML-based form that serves as a tool for further searching of travel and tourism sites:

"Under the proposed service, when a travel consumer types in a travel word or term as a DNS look-up they will land at a page that indicates that the name is not registered and is available for registration to eligible entities. Further, the entered name will be parsed into the search box on that page, without having the user type it in again. This search will return results that meet the profile description of the term, with .travel TLD results returned with higher rankings."

Under this proposal, any DNS query for a name string for which no exactly matching resource record can be found would receive, instead of a "does not exist" (*nxdomain*) response, a synthesized response containing the IP address of the *search.travel* server. All queries for domain names

[9] DNS wildcards are defined in RFC 1034; an update on "The Role of Wildcards in the Domain Name System" is provided by RFC 4592.
[10] The descriptions of the two components are quoted from the *Tralliance Application to ICANN for New Registry Service*.

not found in .travel would receive the same response (modulo load balancing of the web servers supporting *search.travel*).

Tralliance proposes a phased implementation of the .travel wildcard. In the initial phase, Tralliance would insert the wildcard for a period of one hour and then remove it, with the goal of gathering operational and performance data and later analyzing it for any impact. In their application, Tralliance describes a test plan for the insertion of the wildcard that includes workload tests, tests on the target web server, and query stream testing against .travel DNS servers. After a series of these one-hour tests, Tralliance would then try longer deployments of the wildcard with tests of 2, 4, 8, 12, and finally 24 hours. Following these tests, and assuming that the test results show no significant adverse impacts, Tralliance would put the wildcard in place permanently. During this testing period, Tralliance would mark the web page supporting the *search.travel* service as "beta," removing this designation once the testing was complete.

The proposal would have no impact on the *whois* service. Those names currently registered in .travel would continue to appear in the whois database. Those names not registered would not appear in *whois*, although DNS queries for those names would still be answered with the *search.travel* IP address.

The Tralliance proposal would also capture standard DNS query log files, and the Web server supporting *search.travel* would maintain standard HTTP log files.

Tralliance intends to address issues related to SMTP mail being sent erroneously to the *search.travel* wildcard address by not opening ports other than TCP port 80 (the default port for HTTP) at *search.travel*. The intended result is that, because port 25 (the default SMTP port) is not open at *search.travel*, the SMTP HELO message would time out.[11] To enable the developers of spam filters (for example) to distinguish between domain names under .travel that actually exist and those that are synthesized by the wildcard, Tralliance proposes to publish the IP address(es) of *search.travel* on its web site, with the assumption that applications could thereby recognize query responses that were generated by the wildcard mechanism rather than by the existence of a matching resource record, and that deliberate queries for the domain name *search .travel* itself could be handled as a special case.

[11] Because the landing site does not respond to SMTP connections, it appears as if a server is in place that does not have a mail server running. When port 25 is not open, the attempt to open the underlying TCP connection will be rejected with the ICMP error "connection refused," and the exchange will never get as far as the SMTP HELO message. The situation is complicated, however, by the possible interposition of firewalls and other devices that may filter ICMP packets, blocking the "connection refused" message; in this case the client's TCP SYN will eventually time out. ICMP "connection refused" and a TCP SYN timeout are typically handled very differently by SMTP implementations.

1.2 RSTEP Process Summary

1.2.1 Activities

RSTEP evaluated the Tralliance proposal with respect to its potential impact on the security and stability of the Internet. In order to inform its work, the panel took advantage of previous analyses of wildcards in the apex of TLD zones, consulted with outside experts, and engaged Tralliance in clarifying discussion.

During the period of the panel's work (starting with the referral from ICANN to the Chair of the Registry Services Technical Evaluation Panel on September 18, 2006), the panel took the following actions:

- Participated in ten conference calls attended by the panel and the Chair of the Registry Services Technical Evaluation Panel;
- Convened a clarifying conference call with Tralliance and representatives of Neustar on October 13, 2006;
- Reviewed the feedback of the open public comment process initiated by ICANN on 19 September 2006;
- Consulted with external experts in registry services related to security, stability, and wildcard implementation;
- Reviewed the feedback of ICANN's SSAC on wildcard deployments; and
- Reviewed the feedback of the Internet Architecture Board on wildcard deployments.

1.2.2 Public Comments

ICANN opened a public comment forum for the Tralliance *search.travel* proposal on September 19, 2006. The comment period closed on October 18, 2006.

A total of 14 substantive comments were made in the forum. Thirteen were made by individual members of the community. One comment was the consensus position of the ICANN At Large Advisory Committee. No other supporting organization or constituency within ICANN commented on the proposal. Abstracts of the public comments can be found in the References section of this report.

1.2.3 Gathering of Supporting Material and Data

In the early part of the panel's work, a great deal of supporting material related to wildcards in the apex of a TLD zone was gathered and reviewed.

In addition, external experts were consulted and, where possible, empirical data were gathered about the operation of TLD wildcards.

The available supporting material was collected in a reference library available to panel members during the 45-day review period. The publicly available material is widely available and the References section of this document provides abstracts and, where available, URLs for the source documents that the panel reviewed.

1.2.4 Discussions with Tralliance

The panel arranged for a conference call with representatives of Tralliance and those companies supporting the operation of the registry. The goal of this discussion was to clarify the panel's understanding of specific aspects of the proposal. This conference call took place on 13 October 2006.

1.3 Key Definitions

1.3.1 Security

An effect on security by the proposed Registry Service shall mean (A) the unauthorized disclosure, alteration, insertion or destruction of Registry Data, or (B) the unauthorized access to or disclosure of information or resources on the Internet by systems operating in accordance with all applicable standards. (Definition comes from GNSO Recommendation, located at http://gnso.icann.org/issues/registry-services/final-rpt-registry-approval-10july05.htm#5.)

1.3.2 Stability

An effect on stability shall mean that the proposed Registry Service (A) is not compliant with applicable relevant standards that are authoritative and published by a well-established, recognized and authoritative standards body, such as relevant Standards-Track or Best Current Practice RFCs sponsored by the IETF or (B) creates a condition that adversely affects the throughput, response time, consistency or coherence of responses to Internet servers or end systems, operating in accordance with applicable relevant standards that are authoritative and published by a well-established, recognized and authoritative standards body, such as relevant Standards-Track or Best Current Practice RFCs and relying on Registry Operator's delegation information or provisioning services. (Definition comes from GNSO Recommendation, located at http://gnso.icann.org/issues/registry-services/final-rpt-registry-approval-10july05.htm#5.)

1.4 Members of the RSTEP Panel for this Proposal

The five members of the RSTEP Panel for the Tralliance *search.travel* proposal are:

- Patrik Fältström (Cisco; Sweden)
- Lars-Johan Liman (Autonomica; Sweden)
- Cricket Liu (Infoblox; USA)
- Mark McFadden (internet policy advisors, llc; USA)
- Paul Mockapetris (Nominum; USA)

The members of the panel were assisted in their work by the Chair of the Registry Services Technical Evaluation Panel:

- Lyman Chapin (Interisle Consulting Group; USA)

Staff support was provided to the panel by ICANN:

- Patrick Jones - Registry Liaison Manager

2 Summary of Findings

Tralliance Corporation proposes to introduce a wildcard at the apex of the .travel sTLD (that is, *.travel) in order to redirect queries for name strings that are not found in the .travel zone to a web form that:

(a) advertises the availability of the name string for registration as a .travel domain name; and,
(b) provides a search box pre-loaded with the name string which, if used, would return .travel results with higher rankings than results from other TLDs.

The effect of such a wildcard would be to replace the "does not exist" response (*nxdomain*) that is normally returned by a name server when it finds no match for a query with a response that gives the address of the search.travel web site as the successful result of the query.

For example, a query for Deutschland.travel, which (presumably) matches a name string in an existing .travel zone resource record, would return the registered address associated with that name; a query for Duetchland.travel (*sic*), presumably a misspelling, would return the address of the search.travel web site rather than "does not exist." Tralliance compares this service to the wildcard currently in operation at the apex of the .museum TLD.

With respect to technical feasibility, we believe that Tralliance could implement the service that they have proposed. The test plan and prior art seem adequate to ensure this.

Our technical evaluation of this proposed registry service with respect to the likelihood and materiality of effects on security and stability concludes that it does create a reasonable risk of a meaningful adverse effect on security and stability. This report presents a detailed description of the technical issues raised by the proposed service, and the assumptions, information, and reasoning upon which our evaluation is based.

The principal findings that lead us to this conclusion may be summarized as follows:

The fundamental difficulty presented by the proposed .travel wildcard is that redirection would affect all current and future applications and protocols that rely on the DNS. The effects of redirection could not, given the current state of Internet standards and practice, be restricted to simple HTTP web traffic (the context in which the benefits of the service are

intended to be realized). The wildcard would change the definition of a host address and disable the technique that many applications use to detect (and potentially correct) erroneous or misleading input. For example:

- Misaddressed mail could be delayed, perhaps by days.
- Spam filters would become less effective.
- The locally-optimized algorithms that some applications, including web browsers and proxies, use to decide what to do with a name string that does not correspond to a properly registered domain name would no longer operate correctly.
- Some resolver search list algorithms would interact with the wildcard to inadvertently match non-existent domain names and produce erroneous or surprising results.
- Easily detected errors in configuration files and clickable links would become difficult to find as the wildcard transforms hard errors into fake soft errors.

Our investigations discovered many such applications that obey all applicable Internet standards—that is, the difficulty is not limited to misconfigured or legacy applications. Because the proposed wildcard changes the expected behavior of the DNS in such a fundamental way, it is impossible to anticipate all of its side effects without testing each and every mail server and agent, every instant message application and agent, every VOIP server, proxy, and user agent, every parental control system—basically every application on the Internet. We are not persuaded that either the "small size" of the current .travel domain or the publication by the .travel registry operator of the IP address(es) of the search.travel server(s) adequately mitigates these concerns.

Beyond the mere existence of these problems, the empirical evidence shows that the problems and their associated costs are often felt by the user, the user's organization, or ISP rather than the .travel registry, registrars, or registrants. It is evident that although a .travel wildcard will not destroy the Internet, its introduction will impair the Internet's existing portfolio of applications, and complicate the development and implementation of new services such as DNSSEC, Internationalized Domain Names, and protocols not yet imagined.

The potential added value of the proposal is to redirect web users away from difficult-to-understand error messages toward easy-to-use search tools placed in a common "landing page." We observe that there are alternatives to embedding search in landing sites that do not result in the stability issues that are associated with wildcards. Automatic searching on name errors, when implemented in a web browser, eliminates all of the problems caused by the wildcard's effect on protocols other than HTTP,

and gives the user a consistent experience regardless of which top-level zone contains the mistyped domain name. We see no reason to believe that the .travel landing site has any technical advantage over non-wildcard approaches to improving user experience.

If it were deployed, the landing site's operation would have to be constrained in order to protect the user's privacy and security. In the event of deployment of this proposal, the user is particularly vulnerable to phishing. We do not see technology as a solution to this problem. Instead, the solution lies within the contractual relationship between ICANN and Tralliance—a relationship which is beyond the scope of this panel. If a user typing our example "Duetchland.travel" were redirected to a page that promoted travel to Fiji, it would be fair to view the coherence of the system as compromised.

Finally, some believe that the small size of the .travel domain reduces the scope of any potential negative effects of adding a wildcard. If reciprocal fairness results in other registries being able to implement wildcards as a result of a decision about .travel, the relative size of the .travel zone does not change its impact on the stability of the Internet. We believe that the problems with wildcards will grow as .travel grows. As other top-level zones deploy wildcards, and registrants in .travel use the zone for normal traffic, the problems will worsen.

As an example, one of the suggested "fixes" to the problems the wildcard would cause to email is to publish a list of landing sites. Once this list was available, all mail servers would need to be reconfigured to equate mail addressed to these sites with mail addressed to a non-existent address. Not only does this unfairly burden administrators on the Internet, but the use of wildcards by more top-level zones would expand this list, offering opportunities for DOS attacks by manipulating the list, and increasing the ongoing workload and the chance of error.

In summary, while we believe that Tralliance could implement the service that they have proposed, we also conclude that the proposal does create a reasonable risk of a meaningful adverse effect on security and stability to the public Internet.

3 Security and Stability—Issues

3.1 Architectural and Theoretical Implication of Wildcards

The original purpose of Internet technology was to allow multiple users and applications to operate over multiple different types of networks. To avoid the so-called "N-squared" problem, in which each user and application would have to learn the details of each and every network, a common set of intermediate standards was created including IP and TCP. Applications and users above the intermediate layers only need to learn to interface to the intermediate standards in order to use the networks hidden below IP and TCP. The same holds for the network types that wish to make themselves available for use by applications.

This design is also known as the hourglass model, which builds up from a lower layer which has a large number of network technologies (e.g., Ethernet, ATM, point-to-point circuits) through a layer which has only IP, through a transport layer that has a few choices (i.e., TCP and UDP) and on to a wide diversity of middleware and applications.

The central theme is that by standardizing the middle layers where there are few choices, we get to interconnect a very large number of higher-level technologies to a large number of lower-level technologies without an explosion of the effort to do so. We even get to add new choices above and below, so long as we preserve the standardization of the intermediate layer.

DNS's role in this is to mirror this structure and provide a database for configuration information that can also be layered; for example mail routing is handled with the mail exchange (MX) type, which in turn uses the address (A) datatype. The DNS keeps track of its own distribution using the Name Server (NS) type, which also relies on the address type.

The key point here is that changing the semantics of any of these data types in the DNS potentially affects all users of that datatype. The whole point of publishing the standard is to allow others to reuse the definitions for other applications, some of which become known via registered port numbers, and others that are unknown. So, for example, HTTP could use A records even though HTTP didn't exist at the time the A record was standardized.

Of course, the Internet and DNS do see continued evolution; the original fixed formats for IN-ADDR.ARPA and addresses in the A RR have evolved to include private address space and classless addressing; host

names now regularly include leading digits (e.g., 7up.com, 3com.com, etc.) and also include encoded formats for international characters. All of these changes in DNS led to consequences for the users of the system. In some cases the changes led to sequences of disruption when changes in one part of the system led to changes or countermeasures elsewhere that had further unintended consequences.

The original impetus for wildcards was to provide a mechanism for routing mail between systems in the US and UK, *inter alia*, which used completely different protocols and required translation in special gateways. In RFC 882, "wildcards":

> ***In certain cases, an administrator may wish to associate default resource information for all or part of a domain. For example, the CSNET domain administrator may wish to establish IN class mail forwarding for all hosts in the CSNET domain without IN capability.***

The purpose of the address RR is to associate an Internet class address to a host name.

Thus, when a .travel wildcard creates an A RR in response to a query for a name that doesn't have explicit data, it creates two kinds of potential problems:

- The first is that algorithms that wish to test for the existence of a particular host no longer work. For example, anti-spam algorithms often do this to detect forged addresses. In essence, we have broken a tool that others depend upon.

- The second is that algorithms that assume that the A RR is genuine may be misled. For example, mail to a mistyped email address will be directed to a host that is not a mail server and delivery retried until the sender gives up.

Experience and common sense suggest that the effects are not cataclysmic; otherwise, rogue hackers might have crashed the Internet by now. However, assessing the costs entails examining all known protocols that use the A RR type, as well as types that depend on A, such as NS, MX, etc. In some cases we can rule out any effects, in others we can estimate the effects, but the success at standardization means that we can't know all of the effects without knowing all of the different uses by every Internet user. We can know that frequently the problems will not be seen by .travel domain operators, or their clients, but by uninvolved third parties. The A RR type is the type that has the most direct and indirect uses, so changing its definition should be done with caution.

Some defend (perhaps any) use of wildcards if they syntactically follow the form set by the RFCs. This is a necessary, but not sufficient, proof. Hackers use syntactically correct DNS packets in many DDOS and identity theft attacks, which does not legitimize them. Even the original use of wildcards with mail gateways took place in an Internet where the protocol set was limited to essentially TELNET (remote login), FTP, and mail, and didn't really create any side effect with any of them.

The proper approach to the TLD wildcard is to understand that it:

- impairs the general utility of the DNS database to applications unknown and yet to be invented;
- has empirically observed effects which are discussed in the following sections;
- affects an unknown number of systems; and,
- will create costs to third parties.

and balance these against the perceived benefits.

The benefits of the search site for .travel should also be compared to the similar search functions that could be implemented in the web browser, by an ISP, or by a third party DNS service provider.

3.2 Security Issues Related to the Proposal

3.2.1 Impact of the proposal on privacy of .travel users

When an application communicates over the Internet, it usually starts by looking up the IP address of the endpoint it wishes to connect to. Given the IP address, the application can create a connection between the originator and the termination of the connection, or flow. Communication is initiated over that flow, but exactly how that works depends on how the underlying protocol works. In some protocols the originator begins by sending some data, in other protocols the termination point starts.

If two entities want to communicate, and these entities are in the same country, then all the traffic between may stay within the same jurisdiction. In those cases, the rule and laws of that jurisdiction would apply to those flows. An example of this would be the wiretapping of a conversation and the retention of the identities of those conversing—in Europe this would be covered in the EU's Data Retention Directive. A misspelling would cause the communication to fail, not to be redirected. All of the data that

makes up the (failed) connection stays in the country the two entities are in.

But if a wildcard existed in the zone, a misspelling would not cause a failed connection, but instead would redirect the traffic to a site chosen by the registry. If the server and the client are in the same country, and the flow between then remains within that country's boundaries, that nation's regulatory framework would certainly apply to the traffic. However, the location of this site (and service) might be in a different country than the communicating parties. In this case, regional or international (for instance, European) regulatory rules might apply.

Information about the flow of traffic—for example, which endpoints are communicating and which protocol they're speaking—might be protected under the privacy laws of the country in which the endpoints reside. However, this information might not be protected in the country in which the redirection site is hosted, and in which this traffic is exposed.

It is also worth noting that redirection may create unexpected privacy requirements for the operator of the zone containing the wildcard. Further, the flow will be redirected (in the cases where the wildcard applies) so that it might cross different juridictions than the originator expected. This is obviously true for HTTP traffic but just as true for other applications. The jurisdiction of the landing site may impose privacy requirements that would be different from the requirements on traffic unaffected by the wildcard.

3.3 Stability Issues Related to the Proposal

3.3.1 Impact on the stability of the DNS

3.3.1.1 Impact of the Proposal on the Ability to Deploy a Secure DNS

The data and protocol extensions to add security to the DNS (DNSSEC) are defined in RFCs 4034 and 4035, respectively. Wildcards have been given considerable attention in these documents, and there is a clear strategy for handling them in DNSSEC.

In DNSSEC each "set of DNS resource records" (which is a well defined term) is accompanied by a matching "resource record signature" (RRSIG) record. RRSIG records are normally pre-computed from the textual format of the resource records signed. This isn't possible in the case of wildcards, however, since it's impossible to predict the domain name the client is going to look up, which will match the wildcard and synthesize a resource record.

DNSSEC's solution is to add a field, called "label count," to the RRSIG record. The label count signals how many "labels" (parts of the domain name) the signature covers. A signature covering the wildcard domain name can then be calculated. When comparing signatures, a name server ignores the fact that the wildcard can change into any possible name.

Example:

> *If the wildcard IP address record "*.example.se" is signed with an RRSIG record, the signature will only cover "example.se". If someone queries for "stupidly-long-test-string.example.se," which happens to match the wildcard, the returned data would be the domain name "stupidly-long-test-string.example.se," the IP address in the wildcard record, and the RRSIG of the wildcard record with a label count of "2". (The actual answer will contain more information, but that information is irrelevant to the issue at hand.) This last number means that only the "example.se" part of the name is covered by the signature. The answer will also contain another record that proves that the exact domain name ("stupidly-long-test-string.example.se") doesn't exist.*

This example applies to any zone, including top-level domains. (In the TLD case, the label count in the RRSIG will of course be "1".) Also, the principles are exactly the same regardless of the data type in the wildcard record (address (A) record, mail exchange (MX) record, or any other type).

The result of this strategy is that DNSSEC gives the client far better ability to recognize situations where a wildcard is being used when compared with classic DNS.

Finally, it should be noted that DNSSEC has not been deployed on a large scale in the public Internet. As DNSSEC is deployed in larger zones and in a broader set of circumstances, new and unforeseen problems may be discovered.

Conclusion

The standards-based use of wildcards in any zone, including TLD zones, theoretically has no negative impact on the ability to deploy DNSSEC. The effect of deploying DNSSEC is to give DNS clients an enhanced ability to determine when a wildcard is used. However, the limited deployment of DNSSEC gives the Internet community little empirical information about possible deployment problems that may not have been anticipated during the development of the DNSSEC standards.

3.3.1.2 Impact of the Interception of DNS NXDOMAIN Responses

There are four major ways a NXDOMAIN response from an authoritative DNS server can be turned into a "search" or "approximate match"—something that some people claim is important to improve the end-user experience.

(1) The application (e.g., a web browser) that performs a DNS query could fall back to a default URL instead of just reporting an error such as "domain does not exist" when the DNS response returns NXDOMAIN. This behavior would depend on the application, allowing the end user to choose his desired behavior by using a different implementation or by changing configuration settings in the application.

(2) The resolver at the client could pass a fallback response back to the end user. The application would never see the NXDOMAIN response, which would render the solution described in (1) ineffective. This is implemented by some ISPs with an interest in providing a value-added service such as a keyword system (in which the ISPs usually take for granted that the only application the end user uses is "the Web"). In this case, the end user sees a different result depending on which resolver he uses. The user may have a limited ability to select his desired behavior by choosing to use a particular resolver.

(3) The registry could add a wildcard to the zone to catch "misspellings." This, like solution (2), implies that the registry running the service can guess which service is used (normally "the Web"). If this solution is employed, neither solution (1) nor (2) will work. The result will be different depending on which zone the misspelling is made in.

(4) If the IP address of the search page is known the local ISP can establish a web server to take its place. Then, as future traffic is supposed to go to the search page, it is routed to the locally established web server. A similar re-routing of traffic can be implemented if the hostname is known, and queries for that hostname are caught in the recursive resolver used by the end user. This form of hijacking is similar to, and has similar effects to, catching the NXDOMAIN in the recursive resolver described in (2).

If we look at the options from the consumer's perspective, it is clear that the end user has more control over behavior in the case of NXDOMAIN

responses in solution (1) than (2), and more control in solution (2) than (3). Solution (1) is in control of the end user, while solution (3) is in control of the registry, while (2) is somewhere in between. We can also see that in solutions (1) and (2), there is the opportunity for providers of the services to "compete" by either selling competing software (such as web browsers), give instructions to end users how to configure the software, or (as in solution (2)) help users configure their resolvers.

The biggest difference between solutions (1), (2) and (3) is that in (3), the registry is a *de facto* monopoly and therefore the only organization that can employ this particular solution. Moreover, since solution (3) prevents solutions (1) and (2) from working, the registry thereby forecloses competing solutions. This has a detrimental impact on innovation and competition for such fallback services—including the option for the end user *not* to use such services.

3.3.2 Impact on the stability of the Internet's Applications

3.3.2.1 Impact on HTTP and Web User Experience and Deployment

The whole purpose of the .travel wildcard is to assist web surfers who have attempted to access a domain name that is not registered to a .travel-specific web site. Such a name may be reserved (as in some country codes), not registered yet, or simply not a name that will ever be legitimate, such as "NewYYork.travel" (sic), which is probably the result of a typo or misspelling.

In this section we consider the implications for web users and those that develop and maintain web sites and applications. We do not consider the possible indirect benefits from potential revenue generation that might subsidize additional services for .travel users.

Users of the search site

The central benefit of the .travel wildcard to the user is that it automatically directs the web surfer to a web page that will help the user find the correct site or offer to sell the name to the user, and display some information and advertisements. The thesis is that .travel is better organized (e.g., as opposed to germany.com) into a travel-specific hierarchy, its registrants are vetted, and hence in addition to being automatic, it can provide better search results.

The automatic aspect of this benefit is not particularly unique. Users of Internet Explorer, Firefox, or other browsers have the option of having failed queries redirected to a search engine via a plugin or the browser itself. There may be some small benefit to knowing that the user is interested in travel, so that "Washington" would be understood to be a

reference to a place and not a person. The result of typing "Washington travel" (or its misspellings) into a search engine and accessing the .travel search engine with "Washington" would not be that different. Users who knew what they were searching for, as opposed to inputting a known domain name (say from a business card) probably started in a search engine anyway, or perhaps a travel-oriented portal such as Orbitz, Travelocity, Travel.com, or .travel's search site.

The organization and supervision of the .travel domain could potentially be beneficial by providing better "editorial control" of the data in the .travel domain. For example, web surfers that travel often suffer from knowing the name of a particular hotel and not being able to find its web site amid the clutter of tens or hundreds of web sites that seek to represent it or divert you to their choices in subtle and not-so-subtle ways. However, there is no guarantee that .travel can achieve such a goal in the face of commercial and scaling pressure.

The very importance of the travel industry that led to the formation of the .travel domain has also led to an abundance of travel information in other domains that compete for user's attention. The issue of whether the .travel search site will provide better information as more of a battle for mindshare and branding than one where .travel can have a significant technical edge.

While the .travel wildcard adds a new search facility for users, it also detracts from the consistency of the web surfing experience.

Alternatives / Competition

As we have seen, the user has many competing offerings for services that will deal with non-existent domain names, with distinct properties:

- As previously mentioned, web browsers, as an example application, often transform non-existent name results into a search query using a search engine supplied as a default or configured by a user.
    - Advantages: The option is user-controlled and selected, can have access to user-specific information to refine the search, and doesn't affect any other applications.
- Some ISPs configure their name servers to watch for address queries for domain names used in the searches generated in the previous case, and redirect those domain names to their own search site.
    - Advantage: This only affects selected searches, and not other applications.

    - Disadvantage: Some see this as unethical, illegal, or both. Search providers will clearly take countermeasures. This option is also incompatible with DNSSEC.
- Some ISPs have caching servers that do the same for queries that are upstream from the user but before the authoritative server.
    - Advantages: The caching servers may have geolocation information, and be able to optimize the results returned.
    - Disadvantages: This option either affects all applications or requires careful filtering to avoid these problems. Some problems may be unavoidable.
- Caching name servers run by non-ISPs (e.g., OpenDNS) that provide DNS service to any user in return for transforming non-existent domain results into revenue-generating traffic.
    - Advantages: Users can choose whether to use the service.
    - Disadvantages: The option adds to network latency, and the caching servers have no context to optimize the search.
- Authoritative servers, such as that proposed for .travel, can redirect users via wildcards.
    - Advantages: The authoritative name server can use domain-specific information to provide marginally better search results
    - Disadvantages: either affects all applications or requires careful filtering to avoid these problems.

From a purely architectural point of view, the browser-based search implementation provides the user with the greatest choice and fewest side effects.

In the real world, the promise of revenue generation will motivate implementation of several of these, and there are technical implications to their interaction, which may grow into competition:

- The ISP is supposed to use the list of IP addresses of the .travel web servers to avoid delayed mail transactions, but can also use them to effectively disable the .travel wildcard by transforming any result containing these addresses into an NXDOMAIN response.
- The ISP can use filtering to force users to use the ISP's name servers and transformations.
- .travel could change its addresses to defeat ISP filtering, but that would also mean that mail would be delayed.

The long term interests of the user are best served when DNS is able to provide a reliable and secure channel between the domain registrants and the users of the information.

Other scenarios will involve users who see inconsistent results when they connect to the Internet in different locations or via different service providers. Tampering by intermediaries degrades the user's security and coherence.

Web site extensions and concerns

While the addition of a wildcard address record is the focus of the .travel issue, there are some other concerns with regard to the operation of the search web site itself.

It is easy to check if a wildcard is in use. However, there is no comparable method to determine if users are being tracked. Tools to track users are often part of the technology used on the Web—and outsourced search services complicate the issue.

While there are some technical measures that could be used (For example the detection of web bugs, á la Sitefinder) this is probably a matter where ICANN and .travel need to consider limits and compliance measurement.

Since there is no way to technically measure whether users' privacy is being degraded by tracking in the web site, this issue needs to be addressed in regulations or contractual relationships.

The collection of a user's activity, when it can be matched to an individual user, may be unethical—or in some jurisdictions, illegal.

The jurisdiction in which the search web site resides may add conditions on the processing and use of data collected while tracking individual sessions or flows.

One of the results observed in Sitefinder was that web content filtering tools in Tennessee schools were defeated [SSAC, page 18] in that it was possible to access objectionable sites via Sitefinder's web page which were not otherwise accessible.

3.3.2.2 Side effects in web browsing

AJAX and other web applications

More and more, the web is a portal to applications such as email, spreadsheets, and other services that are effectively outsourced applications. All of these services will see the same set of issues as applications running directly on hosts. Having an AJAX transaction access the suggested search page instead of the HTTP point of access that was

planned might lead to very unpredictable results in the JavaScript engine as well as in the user interface.

Links vs. typed names

Much web navigation is achieved via clicking on links as opposed to using explicitly typed URLs. In cases where these are set up incorrectly, wildcard redirection can delay recognition of the problem and create the need for additional debugging.

Impacts on Local Customization of Browsers

A common feature of modern browsers is to allow for the use of local languages in the display of menus, toolbars and error messages. Two crucial customizations are possible: a version of the browser whose application tools (menus and dialog boxes, etc.) have been adapted to a local language; and a version of the browser that adapts content in the document window to a local language.

When a browser adapts content based on a user's language preferences it can display text in an alternative character set—possibly rendering the content more usable. This clearly optimizes the browsing experience for the user. This customization is also usually extended to displaying traditional HTTP error messages in the language of the user's choice.

In the presence of a wildcard, the local customization of the browser is disregarded. Instead of getting an error that indicates that the document request was not found (in the preferred language of the user), the target of the wildcard redirection is displayed instead. This is very likely not in the language preferred by the user.

It is possible to attempt to provide translated copies of the target of the wildcard redirection and display a specific version of the web page based on external information (information in the HTTP header, information about the IP address). However, these attempts to provide language customization based on external information often fail. It is difficult for any entity to maintain multiple copies of content in all languages that might be requested. Web pages displayed in a particular language based on the origin IP address fail in the presence of roaming users, IP tunnels and anonymization strategies. As an example, an American browser traveling to Stockholm and requesting the home page of Google gets the Google home page translated into Swedish—despite the local language needs of the specific user.

3.3.2.3 HTTP vs. HTTPS

With a wildcard in place, a Web user's experience of redirected, secure HTTP sessions would change. Either the trust model involving the certificates would change, or the client would experience timeouts related to not receiving responses to the start of the secure sessions. In either case, an attempt to use a URL with an "https:" schema and a redirected DNS query will impact the user.

SSL, and its replacement TLS, provides three services for HTTP traffic: authentication of the endpoints of HTTP sessions, confidentiality of the traffic exchanged between a browser and a client, and message integrity. The existence of a wildcard during the setup of a secure web session has a single side effect.

When the user of a browser types in a URL such as:

`https://notreally.there.travel/index.html`

the browser must first resolve the name "notreally.there.travel." In the case where wildcards are present, the resolution does not, as we have seen, inform the browser that the domain name does not exist. Using the IP address provided by the resolver, the browser begins the process of setting up the secure session.

In the most common use of SSL and TLS, the browser and client exchange keys once the server has been authenticated. The first step in this process is the responsibility of the browser. It must send a "hello" message to the server to indicate that it wishes to negotiate the start of a SSL/TLS session. The rules for the negotiation are enumerated in the TLS Handshake Protocol, part of the overall TLS standard.

Crucial to the TLS Handshake Protocol is the opening sequence of messages. The client sends a "client hello" message to which the server must respond with a "server hello" message, or else a fatal error will occur and the connection will fail. When a web browser is directed to a web site such as *search.travel* by a wildcard address record and attempts to set up a TLS connection, the registry operator has the choice of whether or not to respond to the TLS "client hello" message.

If nothing is listening at port 443, the SSL/TLS handshake will never be attempted due to a "connection refused" at the transport layer. If firewalls intervene and filter ICMP packets, the result may be a TCP SYN timeout instead.

If the registry operator simply chooses to not respond to the "client hello"—while still listening on port 443 — a fatal error occurs and the

SSL/TLS session is not set up. Note that, compared to a situation where there are no wildcards and the client receives the NXDOMAIN response from the DNS, the client must perform an extra step to discover that the expected SSL/TLS session will fail. The "client hello" exposes more than simply the identity of the client. It also tells the server information about the security capabilities of the browser, including the security protocols supported by the browser, what cryptographic suites are available at the browser and what data compression methods are possible.

If the registry chooses to respond to the browser's "client hello," it sends a "server hello" and a server certificate for inspection by the client. In the "server hello" message, the server nominates the version of SSL and the ciphers and key lengths to be used in future messages between browser and server, chosen from the selection offered in the client hello.

More importantly, the server sends its digital certificate to the client for inspection. Almost all modern browsers automatically check the certificate (depending on configuration) and warn the user if it's not valid. A browser may consider a certificate invalid if it is out of date or does not point to a certification authority that is explicitly trusted. Another important case where the browser may consider a certificate invalid is if the IP address/domain name pair of the certificate does not match the requested pair. In the presence of wildcards, almost all modern browsers would warn the redirected user that the digital certificate presented was not valid.

In either case, whether the registry decides to ignore SSL/TLS "client hello" messages or respond to them, the underlying behavior of the secure Web session changes. The scope of the behavior change is small when individual sessions are considered, but has the potential to be large as the scale of the zone involved, and the number of secure Web sessions, increases. It is worth considering whether the exposure of the capabilities of the browser, even when the redirected server has no intention of setting up a secure session with the browser, represents a potential security risk for the user of the browser.

3.3.2.4 Impact on SMTP

Internet mail service depends heavily on DNS for routing messages. The bulk of Internet mail is transferred using the (Extended) Simple Mail Transfer Protocol ((E)SMTP) according to RFC 2821 and its predecessors, primarily RFCs 821 and 974 (full standard and historic, respectively).

SMTP uses DNS and its mail routing capabilities for the following purposes:

I) To reach a mail server closer to the destination

The sending mail server is obliged by the standards to first look for the MX (mail exchange) record for the domain name in the email address to which it is sending a message. If one or more MX records exist, the server is obliged to use them. If no MX record exists, the mail server is expected to look for an A (internet address) record, and where available, use it as a replacement for the wanted MX record. The standards clearly state that A records may only be used if no MX record exists.

This fallback to the use of A records is for backward compatibility: in the days before MX records, the right hand side of a mail address would be interpreted as the name of the specific computer the recipient user was using, rather than the mail domain he/she belonged to. The A record would then refer to the appropriate computer, which, of course, was expected to operate a mail service.

If a mail server is instructed to send a message to a non-existent mail domain in a zone that contains no wildcard records, the normal outcome is that DNS returns NXDOMAIN back to the mail server. The mail server can then immediately issue an error report to the sending user with a clear message letting him know that the domain name doesn't exist. The user is promptly informed that the message hasn't gone through, as well as why it didn't, and can react promptly to the problem.

1) If a wildcard A record exists, but no corresponding wildcard MX record is entered into a zone, the consequence will be the following in the SMTP case.

The mail server will issue the MX query for the domain name, but since the wildcard record is specific to the data type (i.e., A record in this case), no MX record exists. However, it is no longer obvious to the DNS server that the *domain name* doesn't exist, so the response is no longer NXDOMAIN. Instead, it is reported as an empty answer, which is interpreted as “the domain name exists, but not connected with the type of data you're looking for.” Since the domain name seems to exist, the mail server will then fall back to looking up the A record for the non-existent domain. This record does not exist *per se*, but due to the wildcard, one will be synthesized “on the fly.” An A record will be returned to the mail server, and the mail domain will appear to exist. The mail server will now attempt to deliver the message to the IP address in the A record.

Here two things can happen, depending on whether a mail service is operated on that IP address or not.

a) In the case that a mail server runs at the IP address, the sending mail server will initiate a transaction and, depending on its configuration, the receiving mail server will either accept or reject the message. If it

rejects the message, the error code reported back to the sending mail server can be used to inform the sending user of the situation, and depending on which message the receiving mail server issues, the user may or may not understand what the problem is. If the message is accepted by the receiving mail server, the obvious conclusion to the sending mail server is that the message has been transferred to the right party, and there is no need to inform the sending user at all. This has the two major drawbacks: First, no one will know what happened to the message, and it may take some time before anyone notices that it has disappeared, if anyone notices at all. Second, the recipient party suddenly has access to a mail message that was in no way intended for him or her, which is quite harmful from an integrity perspective.

b) In the case that a mail service is *not* operated at the IP address, the sending mail server will attempt to send the message to the (non-existent) mail service at the IP address reported by the DNS. This contact attempt will fail immediately, but the standard then instructs the sending mail server to repeatedly try to send the message in the hope that this is a temporary failure (maybe a software problem or a network outage). This is not deemed to be an error situation, and the mail server will try for an extended period of time (often 3 to 5 days) before finally reporting back to the sending user that the transaction failed. The error message, however, will not be that the mail domain doesn't exist, but that its mail server doesn't respond.

The user is implicitly led to believe that the domain name exists and that there is some other technical problem, which may be resolved in the future. This misleading message will reach the sending user several days after the submission of the message, at which point in time the content of the message is quite likely to be outdated. Some mail servers will issue a warning message after a couple of hours saying that the initial attempt at delivery failed, and that delivery will be retried, but the message returned to the sending user is still not that the address is wrong. Also, in this scenario, system resources are wasted on the sending mail server to keep track of the message and its status, to issue repeated DNS queries, to make repeated attempts to deliver it, etc. This may not seem important, until you multiply it by the number of messages that large email operators handle, which often is in the order of millions of messages per day.

In the case of .travel the proposal is to *not* operate a mail server on the host with the address of the wildcard A record.

2) Now, consider the case where the wildcard A record is accompanied by a wildcard MX record.

The major difference from case 1) is that the MX lookup will succeed and yield one (or several) potential recipient mail servers. This gives the zone administrator the ability to give the sending mail server a *different* Internet address to try to deliver the message to. The failure modes are the same as in cases 1) and 2) above, except that the zone administrator may choose to use a totally bogus network address, such as 127.0.0.1 (often referred to as "localhost") which means "your own machine." This may or may not be detected by the sending mail server. If it isn't detected, the server will attempt to send the message to itself, which will at best produce a very strange error message back to the sending user. If it *is* detected, it will produce an equally strange, but different error message to the user. None of these will inform the user that the mail domain does not exist. There is also the possibility that the message will be queued in the attempts to send it to the listed host (whether that host address is bogus or not), and therefore waste resources unnecessarily.

II) Verify sender domain

There is another case in which the mail server makes intelligent use of DNS. The SMTP standard has the concept of an "envelope sender." This corresponds to the address on the reverse side of the paper envelope, to which the postal service will return the paper message if it is undeliverable. The function is exactly the same in the electronic version: this is the address to which an error message should be sent, if there is a need to send one.

The envelope sender is the first useful thing a sending mail server tells its recipient counterpart. The reason is that among the first things the recipient server wants to do, is to make sure it can return an error message to the sender, if there is the need to send one. Hence, when the recipient mail server hears this envelope sender address, it will immediately (before continuing the transaction) look up the mail domain of this address, to make sure it can reach the sender. This is done in the same fashion as described above (it "pretends" to send an error report, and performs all the corresponding lookups). If the mail domain doesn't exist, it will reject the incoming message on the basis that it will be unable to send an error report back to the sending user, if need be, and leaves it up to the sending mail server (which has obviously accepted the message, and hence *has* a way to report back to the user) to send an error report back.

If the message is sent from a non-existent mail domain in a zone in which there is a wildcard, the sender's mail domain will *seem* to exist (either a wildcard MX or A record gives this impression). If it seems to exist, the recipient mail server will accept the message, and if it later is unable to forward the message appropriately, it will be unable to send an error report back (quite possibly wasting additional resources trying to do so).

This method of verifying the return path also has a limiting effect on unsolicited commercial email (UCE, or “spam”), which is often produced to seem to come from mail domains that do not exist. If the recipient mail server is unable to verify the sender's mail domain, it will reject the message. A wildcard in the zone the fake domain name falls within will make the fake domain name seem to exist, and hence cause the spam message be accepted, i.e., the wildcard will be to the benefit of those who profit from sending spam.

III) Find submission mail server.

DNS is also used in another part of mail handling, when a mail user agent (MUA, the program that is used for reading and writing mail messages) needs to find its mail server(s). The MUA will typically be configured with the name of a mail server to which it will submit outgoing mail messages. It will look up the address of this mail server in DNS. If the domain name of the mail server is mistyped, it will normally be noticed immediately when the MUA tries to send a message, but with a wildcard in the parent domain, the mail server will seem to exist, and the MUA will attempt to submit its message to it. This may either succeed or not, again depending on whether a mail server is operated on this host. The error message to the user will again be unclear and misleading. Note, however, that in this particular case, only the A record is consulted, not the MX record(s).

Conclusions

Adding a wildcard record to a zone will affect SMTP (email) service to or from domain names ending in .travel. It will have a negative impact on the clarity and promptness of error reports returned to sending users, it is likely to waste resources at mail operators, and it will to some extent impact the ability of mail servers to reject mail from illegitimate mail addresses.

3.3.2.5 Search Lists and the DNS

The search list is a feature of most modern DNS resolvers that allows users to specify partial domain names at the command line, in graphical user interfaces, or in configuration files. The resolver attempts to intelligently “complete” these domain names by appending the domain names in the search list and looking up the result. The importance of the search list feature can be seen in its widespread adoption in commercial software—search lists are implemented in all Microsoft and UNIX systems.

The search list was originally designed to ease the transition between the host table, which used simple, single-label host names, and DNS, which uses multiple-label domain names. Users with the domain name of their local zone in their computer's search list could use just the first label of the domain name of a host to access it. For example, a user with a computer in the foo.example zone would likely have foo.example in his resolver's search list, and could type just “ssh host1” to reach host1.foo.example.

Most resolvers offer two ways to configure the search list: an explicit method and an implicit one. A user configuring the search list explicitly specifies the domain names in the search list in the order he'd like them appended. For example, this BIND resolver directive sets the search list to include the domain names foo.example and bar.example, in that order:

```
search foo.example bar.example
```

The search list may be set implicitly simply by setting a resolver's local domain name. The search list usually includes the local domain name and additional domain names derived by removing successive leading labels of the local domain name. For example, a resolver with the local domain name foo.bar.example might derive a search list which included the domain names foo.bar.example and bar.example.

How this derivation is done may vary, from one resolver to another, based on configuration, or even between versions of a particular resolver. Given the local domain name foo.bar.example, one resolver might also include the domain name example in the search list while another wouldn't. Still other resolvers, particularly those included in Windows operating systems, may have multiple “local” domain names, and so may derive a search list that includes domain names “devolved,” in Microsoft's terminology, from those multiple local domain names.

Another area in which resolvers may implement the search list differently is in when it's applied; that is, when the elements of the search list are appended to possibly incomplete domain names. Some resolvers examine the domain name the user typed for a trailing dot, which is taken as a cue indicating that the domain name is absolute, or written relative to the root. Some resolvers count the number of dots in the domain name to determine whether it's likely absolute; if the number exceeds some threshold, possibly configurable, the domain name will be looked up as-is before applying the search list.

3.3.2.6 Impact on services other than HTTP and SMTP

The addition of a wildcard A record to the .travel zone would have several negative effects on services other than HTTP and SMTP. In fact, the wildcard record has the potential to adversely affect almost any TCP/IP-

based service because of its interaction with both features of and common implementation flaws in stub resolvers.

The first issue is mentioned in RFC 1535, which notes the following security issue with wildcards:

> "After registering the EDU.COM domain, it was discovered that an unliberal application of one wildcard CNAME record would cause *all* connects from any .COM site to any .EDU site to terminate at one target machine in the private edu.com sub-domain.
>
> Further, discussion reveals that specific hostnames registered in this private subdomain, or any similarly named subdomain may be used to spoof a host.
>
> Example: **`harvard.edu.com. CNAME targethost`**
>
> Thus all connects to Harvard.edu from all .com sites would end up at targethost, a machine which could provide a Harvard.edu login banner.
>
> This is clearly unacceptable. Further, it could only be made worse with domains like COM.EDU, MIL.GOV, GOV.COM, etc."

This and similar issues could afflict anyone with a search list which includes the domain name "travel" or a nonexistent domain name such as "mispelled.travel." This search list might be configured explicitly, term by term, to include "travel," or might be implicitly derived from a domain name that simply ended in "travel," such as "corp.carlson.travel."

As applied to the .travel wildcard, the issue is this: Any computer with the domain name "travel" in its search list may, depending on the nuances of its resolver's implementation, append the string "travel" to domain name arguments typed on the command line or domain names entered in configuration files before looking up the domain name literatim. With the wildcard A record in place, these lookups ending in "travel" will never return the NXDOMAIN response necessary to cause the resolver to try looking up the argument exactly as it was typed. If the query is for an A record, the lookup will return the A record of *search.travel*.

This, in turn, will cause one of several problems. For example, imagine the user correctly types the domain name of his intended destination at the command line:

```
% ssh host.foo.example
```

Instead of a login prompt, the user will see something like:

ssh: connect to host host.foo.example.travel port 22: Connection refused

A user reading the error carefully may notice that the domain name of the host ssh tried to connect to is not the domain name he intended. However, many users wouldn't notice that subtlety. Moreover, many programs don't produce output this clear.

If the user inadvertently mistypes the domain name of his intended destination at the command line, the resolver will still return the address of *search.travel* to the program and likely attempt to connect to that host. Depending on the nature of the program, this may expose sensitive data to eavesdropping, either by the maintainers of *search.travel* or by someone with access to an intermediate network.

Worse, incorrect domain names entered in configuration files, where their use may not result in immediate feedback to the user (or any feedback, for that matter), may go unnoticed for much longer than if the domain name lookup returned a simple NXDOMAIN response. At the very least, unexpected responses (the address of *search.travel* or a NODATA response) and consequent misleading error messages will make troubleshooting more difficult.

A common bug in applications and resolver libraries will cause similar results. Some programs don't initially recognize IP addresses typed as arguments or in configuration files as IP addresses. Because IP addresses are syntactically legal domain names, these programs look up IP addresses as domain names first. Upon receiving a response indicating that these “domain names” don't exist, the programs attempt to use them as IP addresses.

An analysis of the queries received by a root name server done by CAIDA shows that this is a common occurrence: Of the “bogus” queries received by a replica of f.root-servers.net over a 24-hour period, between 12% and 18% were queries for the A records of domain names that were already IP addresses.[12]

When coupled with a wildcard in .travel, this phenomenon will cause surprising, undesirable and possibly dangerous results on computers with search lists which include "travel." On such computers, IP addresses used as arguments or in configuration files may be misinterpreted as domain

[12] “DNS Measurements at a Root Server,” Nevil Brownlee, kc Claffy and Evi Nemeth, http://www.caida.org/publications/papers/2001/DNSMeasRoot/dmr.pdf

names and have the string ".travel" appended to them by the resolver. These domain names would then match the wildcard A record in .travel and be mapped to the address of *search.travel*. Thus any IP address, used in the wrong context on a computer configured with such a search list, would be mapped to the address of *search.travel*.

The net effect is much the same as in the earlier scenario, where arbitrary domain names are mapped to the address of *search.travel*. In this case, however, the result is even more unexpected, as many users and administrators use IP addresses rather than domain names specifically for situations in which they want to bypass DNS entirely and ensure the use of a particular IP address.

3.3.2.7 Impact of Wildcards on the Deployment of IDNs

The fundamental principle of Internationalized Domain Names (IDNs) is the ability to convert a non-ASCII domain name into an ASCII label that is compatible with the existing Domain Name System. The conversion from a non-ASCII domain name is done through a two-step process of normalization and conversion.

The <<ToASCII>> algorithm does not change labels containing all ASCII characters, but if a label has at least one non-ASCII character, it applies the two-step normalization and conversion process. Normalization is done through a process called <<nameprep>>. The remaining conversion is important for our consideration of the current proposal. The IDNA conversion process translates a label that includes non-ASCII characters into an ASCII-Compatible Encoding (ACE) using an algorithm called Punycode, then prepends the 4-character string "xn—" to that translation. The 4-character string is called the ASCII Compatible Encoding prefix.

The Punycode conversion does not encode the language, script or glyph during translation. In simple terms, one cannot look at a Punycode ASCII label and determine the language (or languages) of the original label.

What is the impact of looking up a domain name that begins with an ACE label in zone that contains a wildcard?

- If the zone supports IDNs, then the label may exist in the zone. If the label exists, then DNS functions normally: returning the resource records requested in the query.
- If the zone supports IDNs, but the label does not exist in the zone, the wildcard returns the address pointed to in the wildcard.
- If the zone does not support IDNs, in all cases the wildcard returns the address pointed to in the wildcard.

Suppose a query is made for an address record attached to ACE-encoded domain name in .travel once the wildcard is in place. Assuming that .travel makes no special provision for the support of IDNs, the result is that the client receives the address of the web server provided by the wildcard. The resulting web page may or may not be in the language expected by the client. The panel notes that, while there are other mechanisms for attempting to provide content localization, the .travel domain will not be able to accomplish this through the use of IDNs with a wildcard in place.

The panel also notes that, given the long history of IDNs, there have been many alternative, non-IETF, approaches to solving the problem of non-ASCII character sets in domain names. In the cases where client-side applications intercept and redirect DNS queries, the .travel wildcard will not have its intended effect.

To request a web page, as an example, using an IDN and to get a different page in a different language is an unexpected result. TLD operators could avoid this problem by filtering incoming queries that have an ACE prefix and immediately returning an NXDOMAIN result to the client. The panel notes that Tralliance makes no statement in its application regarding support for IDNs and any interaction between the proposed wildcard and IDNs.

4 References

4.1 Introduction

The members of the RSTEP panel that carried out the Security and Stability implications analysis benefited from a rich and substantial library of reference material that helped in their analysis of the Tralliance proposal. The combination of the history of wildcards in the DNS with the process for introducing new registry services created a large body of publicly available information related to the Tralliance proposal.

It is not the intent of the panel to duplicate or restate any of the reports, analysis or guidelines provided by previous parties.

However, as a service to the ICANN Board as it considers the Tralliance proposal and to members of the Internet community considering wildcards in the DNS, the panel has collected and annotated the references that it found useful during its work in October and November of 2006.

These references include the material that is specific to Tralliance's application for implementation of a new registry service. They also include references that are germane to the discussion of security and stability implications of inserting wildcards in the apex of TLD zones.

The panel has provided brief annotations on the material provided in this reference so that the reader may understand how the source material was used during the panel's work.

4.2 Material Specific to this Application

4.2.1 Tralliance Application to ICANN for New Registry Service

Document metadata:

- Dated: 2006-08-23
- Author: Tralliance Corporation
- Length: 19 pages (printed from Web page)
- URL: http://www.icann.org/registries/rsep/tralliance_request.pdf

This is the document that Tralliance sent to ICANN to request that a new registry service be approved. It contains a description of the service, how it will be implemented, the benefits proposed and a discussion of the contractual implications of the proposal. This document is in a standard format supplied by ICANN. Tralliance added two attachments: a technical description of the service and a description of how whois would be affected.

4.2.2 ICANN Letter to SSAC

Document metadata:

- Dated: 2006-09-01
- Author: Kurt Pritz
- Length: 1 page
- URL: http://www.icann.org/registries/rsep/icann-to-ssac-01sep06.pdf

The ICANN Registry Services Evaluation Policy allows ICANN to avail itself of expert advice during the preliminary examination of any new registry service. In particular, ICANN is allowed to ask if new service applications should be referred to the Standing Technical Review Panel. This letter formally requests advice from ICANN's Security and Stability Advisory Committee on whether or not the Tralliance proposal should be referred to RSTEP.

4.2.3 SSAC Response to ICANN

Document metadata:

- Dated: 2006-09-06
- Author: Steve Crocker
- Length: 2 pages (email)
- URL: http://www.icann.org/registries/rsep/ssac-to-icann-06sep06.pdf

This is the SSAC's response to ICANN's 2006-09-01 letter. It quotes the earlier SSAC report from 2004 and advises against Tralliance's proposal.

4.2.4 ICANN Notice of Referral to Tralliance

Document metadata:

- Dated: 2006-09-13
- Author: Patrick Jones
- Length: 2 pages
- URL:
http://www.icann.org/registries/rsep/icann-to-tralliance-13sep06.pdf

ICANN is required, in its Registry Services Evaluation Policy, to notify any applicant for new registry services that the application is to be referred to RSTEP. This process gives the applicant a chance to confirm that they wish to proceed with the process and discusses the process by which the review team will be selected.

4.2.5 Tralliance Response to ICANN

Document metadata:

- Dated: 2006-09-14
- Author: Cherian Mathai
- Length: 2 pages
- URL:
http://www.icann.org/registries/rsep/tralliance-to-icann-14sep06.pdf

This letter states Tralliance's intent to continue with the New Registry Service process and indicates Tralliance's dissatisfaction with the SSAC response of 2006-09-06.

4.2.6 Referral of Tralliance Request from ICANN to RSTEP

Document metadata:

- Dated: 2006-09-18
- Author: Patrick Jones
- Length: 2 pages
- URL:

http://www.icann.org/registries/rsep/icann-to-rstep-18sep06.pdf

This letter to the chair of the RSTEP provides notification that ICANN is going to use the 45-day process of technical evaluation for the Tralliance proposal. It outlines the timetable for the RSTEP process and provides the starting impetus for the technical evaluation.

4.2.7 ICANN Public Comments on Tralliance Proposal

Document metadata:

- Dated: September 18 through October 18, 2006
- Author: Various postings from 13 separate authors
- Length: 14 separate forum postings
- URL:

http://forum.icann.org/lists/tralliance-comments

ICANN opened a public forum for comment on the Tralliance proposal on September 18, 2006. The public forum was open for a month and saw 14 separate postings. Notable postings included a statement from the At-Large Advisory Committee, former and present members of the Internet Architecture Board and representatives of Tralliance.

The members of the panel used the following abstract to help guide their reading of the public comments.

[#1] Danny Younger (http://forum.icann.org/lists/tralliance-comments/msg00000.html) wonders who is more authoritative: the IAB recommendation which provides a set of criteria for instituting a wildcard in a TLD and ICANN's own SSAC which is quoted as saying that wildcards ought to be phased out.

[#2] Frank Schilling (http://forum.icann.org/lists/tralliance-comments/msg00001.html) worries about precedent setting for larger registries when making decisions about smaller registries. In particular, he believes that a precedent set for Tralliance will be later used by a larger registry to establish a similar service. He suggests that the proposal be denied and that contractual conditions be set so that this kind of service is not open to all namespaces.

[#3] W Lipiner (http://forum.icann.org/lists/tralliance-comments/msg00003.html) is opposed to the wildcard approach because it will "channel the public into their coffers."

[#4] John Levine (http://forum.icann.org/lists/tralliance-comments/msg00004.html) submitted the consensus position of the At Large Advisory Committee. It is a very long set of comments and it strongly urges ICANN to reject the proposal. Notably it argues that the testbed TLDs should not be taken as precedent for new services in other TLDs. The ALAC response also discusses the false "success" aspect of queries answered by a wildcard. The ALAC response also spends a considerable amount of time on the impact of wildcards on non-HTTP based applications.

[#5] George Kirikos (http://forum.icann.org/lists/tralliance-comments/msg00002.html) cites SiteFinder extensively as a reason to deny the request. He suggests that .travel should be re-evaluated to see if it really belongs in the root. He is concerned that Tralliance is monetizing search in the .travel TLD.

[#6] Briger Backman (http://forum.icann.org/lists/tralliance-comments/msg00005.html) writing on behalf of The Travel Partnership Corporation, which is the non-profit organization formed to sponsor the .travel TLD supports the introduction of the wildcard. The response concentrates on the lack of a search service specific to .travel domains and suggests that landing on an error page (404 error?) gives a negative impression to visitors. Of particular note is a quote: "It would be extremely disappointing if, after all of the years of work to bring the .travel TLD to the Internet, ICANN would consider denying this important service especially since the Security and Stability Advisory Committee (SSAC) has clearly stated that such a service is already in existence for small and well-defined top level domains." TTPC also believes that Tralliance has met the standard set by the IAB in their 2003 paper.

[#7] Noel Perkins (http://forum.icann.org/lists/tralliance-comments/msg00006.html) simply pasted in Ron Andruff's CircleID response to Brett Fausett's article.

[#8] John Klensin (http://forum.icann.org/lists/tralliance-comments/msg00007.html) provides a very long posting with substantial technical comment. Perhaps a good summary comes from the first paragraph: "no, it wasn't a good idea before, it isn't a good idea now, and it is not going to turn into a good idea in the future." He says that wildcards, in general, are a bad idea. He also says, "I think

that one might be able to accept such a wildcard if it did not lead to any unpredictable behavior that would not be expected from all relevant hosts in the domain."

[#9] Olaf Kolkman (http://forum.icann.org/lists/tralliance-comments/msg00008.html) writes on behalf of the Internet Architecture Board and suggests that the conclusions of the 2003 IAB statement still hold true for the .travel proposal.

[#10] Ron Andruff (http://forum.icann.org/lists/tralliance-comments/msg00009.html) wrote a response to the ALAC paper. This was done in an issue-response format for each one of the issues raised in the paper posed by John Levine.

[#11] Ron Andruff (http://forum.icann.org/lists/tralliance-comments/msg00010.html) posted a second copy of the response above.

[#12] Ken Fockler (http://forum.icann.org/lists/tralliance-comments/msg00011.html) posted individual comments although we are aware that he advises .travel on ICANN related matters. An interesting quote from his posting gives a flavor of his response: "Now seems an opportunity for some small movement, with caution and perhaps even conditions to address possible issues, for a defined space to offer some added value, if desired, for users."

[#13] Edward Hasbrouck (http://forum.icann.org/lists/tralliance-comments/msg00012.html) suggests that "ICANN should not consider or act on this proposal while a request for independent review of the decision to approve .travel is pending." The author of the posting has only concerns about the process of any decision related to .travel and no substantive technical concerns or remarks. He does make the comment that approving a further service in the .travel sTLD would make it more difficult to reconsider the original approval of .travel.

[#14] Thomas Barrett (http://forum.icann.org/lists/tralliance-comments/msg00013.html) is from EnCirca, an ICANN registrar. EnCirca supports the proposal by Tralliance saying that it is a "low risk to internet stability and should be approved." The comments of this registry suggest that whatever impacts might be felt, the implication is small because the registry is small. The posting also says: "No evidence exists that use of the wildcard feature causes any stability or security issues in the DNS." Furthermore, ".museum has shown for five years that the DNS wildcard feature in small TLD namespaces does not introduce stability issues to the global internet."

4.2.8 Current .travel TLD Registry Agreement

On 5 May 2005, ICANN and Tralliance Corporation entered into a Sponsored TLD Registry Agreement under which Tralliance Corporation sponsors the .travel top-level domain. The agreement and the appendices were examined by the panel during its evaluation of the proposed Registry Service. The registry agreement can be seen at the following location: http://www.icann.org/tlds/agreements/travel/

4.3 Supporting Material and Reports

In addition to the documents directly related to Tralliance's proposal, the RSTEP panel made extensive use of other existing reference materials and reports. Since wildcard implementations have been controversial in the past, there was a substantial body of material available for the panel to consider in tandem with the request materials.

What follows is a summary of the public materials that this RSTEP panel used during its consideration of Tralliance's proposal. This list is not exhaustive since the RSTEP panel was able to take advantage of materials not in the public domain and materials subject to non-disclosure.

4.3.1 SSAC Report on Redirection in the .COM and .NET Domains

Document metadata:

- Dated: 2004-07-09
- Author: ICANN Security and Stability Advisory Committee
- Length: 85 pages
- URL:
  http://www.icann.org/committees/security/ssac-report-09jul04.pdf

This is a report by the Security and Stability Advisory Committee (SSAC) in the wake of the initial SiteFinder implementation. SSAC held a pair of public meetings in 2003 and then analyzed the input from both face-to-face and online participation. The result is a document with eight "findings" and four "recommendations" including those that were passed from the SSAC to ICANN at the time the Tralliance proposal was considered.

4.3.2 IAB Commentary on the use of DNS Wildcards

Document metadata:

- Dated: 2003-09-19
- Author: Internet Architecture Board
- Length: 8 printed pages (from web page)
- URL:
  http://www.iab.org/documents/docs/2003-09-20-dns-wildcards.html

The IAB produced a paper called "Architectural Concerns on the use of DNS Wildcards" in response to the SiteFinder implementation. This paper includes a primer on DNS wildcards and a short description of the problems associated with using wildcard records. The IAB paper identifies problems related to web browsing, email, spam filters, interactions with other protocols and automated tools. The paper ends with a set of principles, conclusions and recommendations.

4.3.3 VeriSign's Description of SiteFinder Implementation

Document metadata:

- Dated: 2003-08-27
- Author: Verisign Naming and Directory Services
- Length: 8 pages
- URL: http://www.verisign.com/static/002702.pdf

This document describes VeriSign's implementation of SiteFinder. Included in the discussion is information about how TTL values are set for the wildcarded A records, the responses made to other protocol requests and the filtering that took place for specific ports. The document finishes with a description of the monitoring done on the redirected traffic and how VeriSign intended to communicate with other operators of name servers and networks.

4.3.4 Crocker Presentation on WildCard Issues

Document metadata:

- Dated: 2004-07-21
- Author: Steve Crocker
- Length: 19 presentation slides
- URL:
  http://icann.org/presentations/crocker-ccnso-kl-21jul04.pdf

This presentation acted as an introduction to the SSAC paper on the introduction of wildcards in .COM and .NET. It recaps the findings and recommendations of the SSAC report and identifies the constituencies that participated in the development of the report.

4.3.5 Klensin Presentation on Technical Issues

Document metadata:

- Dated: 2003-10-26
- Author: John Klensin
- Length: 14 presentation slides
- URL:
  http://www.icann.org/presentations/klensin-wildcard-carthage-27oct03.pdf

This presentation addresses the technical issues that surround innovation in the Internet. While the topic is a superset of the wildcard issue, wildcards are used as an example of the problems and responsibilities related to innovation in Internet technology.

4.3.6 MuseDoma Statement on Wildcard Records

Document metadata:

- Dated: 2003-10-06
- Author: Museum Domain Management Association
- Length: 5 printed pages (web page)
- URL: http://musedoma.museum/policy/wildcard/

MuseDoma crafted a statement about its implementation of wildcard A records under the .museum TLD. This was originally intended as input for the SSAC public meeting process (see above). MuseDoma identified the key differences between the .museum wildcard that the SiteFinder experience. It also provided a history of the process used to establish and implement the .museum wildcard.

4.3.7 VeriSign Response to IAB Commentary

Document metadata:

- Dated: 2003-10-06
- Author: Russell Lewis
- Length: 9 pages
- URL: http://www.icann.org/correspondence/verisign-response-iab-06oct03.pdf

VeriSign answers the IAB commentary (see above) by noting that the SiteFinder service was based on IETF standards. The remainder of the document addresses, in a point-by-point style, the objections raised in the IAB commentary on wildcards in the apex of a TLD. The report also addresses the two key IAB recommendations on “understanding the risks of wildcard introduction,” and “informed consent.”

4.3.8 VeriSign Response to SSAC Report

Document metadata:

- Dated: 2003-10-06
- Author: James Ulam
- Length: 3 pages
- URL: http://forum.icann.org/wildcard-comments/msg00204.html

This letter is not exactly a response to the SSAC report, but a request to restructure the public meetings that took place during the development of the SSAC report. In the letter, VeriSign makes note of the public meetings and its concern that they were being set up to "gather data to support the conclusions and recommendations already issued in its report." The report being responded to is not the final SSAC report but the "Security and Stability Advisory Committee Recommendations Regarding Verisign's Wildcard Service" which can be found at: http://www.icann.org/correspondence/secsac-to-board-22sep03.htm.

4.3.9 Extract from Signposts in Cyberspace

Document metadata:

- Dated: 2005-03-31
- Author: Committee on Internet Navigation and the DNS
- Length: 392 pages
- URL: http://www.nap.edu/books/0309096405/html/

The panel used an extract of a section of this book that detailed the specific problems found with application software during the implementation of SiteFinder.