To the ICANN Board,

We, the members of the New gTLD Applicant Group, are greatly concerned about the substance and tone of recent staff recommendations concerning security and stability issues with the New gTLD program. These recommendations were made in response to a report ordered by ICANN and delivered by Interisle Consulting Group.

The NTAG believes that the quantitative results Interisle published do not support the recommendations of ICANN staff.

As you well know, the security and stability of the Internet has been a primary focus of the policy development process that created the New gTLD program. Numerous investigations into gTLD safety have happened and continue to happen in the GNSO, SSAC, RSSAC, and other relevant bodies, and as conscientious citizens of the ICANN community we participate in and applaud these efforts.

With such a long history of productive discussions, it is dismaying to see a great deal of uncertainty about the New gTLD program being introduced at such a late point in the process.

We applicants believe that some valid issues have been raised and we are fully willing to engage with such parties individually or as a group. We are particularly aware that the delegation of gTLDs that correspond to widely used internal namespaces can lead to the leakage of sensitive data and other unpredictable effects. We fully support efforts to prevent these problems, as long as the mitigations are developed based on legitimate evidence of risk and are fairly applied, as compared with past TLD delegations.

The NTAG, including several applicants for the affected strings, agrees that the two strings listed as a "high risk"[1] should be delayed while further studies are conducted. We believe that none of Interisle's findings give cause to delay the new TLD program and that none of the 20% of strings classified as "unknown risk" pose any danger to the DNS or the Internet community and should therefore proceed unhindered.

This letter represents our initial feedback on this potential name collision issue, and is intended to provide guidance that can be implemented quickly without placing the DNS or Internet users at risk. Unfortunately the process ICANN started with Interisle did not include any facility for applicants to study the same data, but members of the NTAG are actively analyzing the DITL captures and other data sets and we will continue to comment with more quantitative detail.

[1] .corp and .home

Our Concerns with the Study

Our initial concerns with the report and staff response follow:

The Problem has been Overstated - Overall a Very Small Number of Requests
Figure 1 of the Interisle report represents an astonishing fact that should influence all discussions on this topic: 45% of requests to the TLD DNS servers are for non-existing TLDs. However, **only 3%** of the total requests conflict with strings that are actually being considered under the new TLD program. Even this 3% figure may be overstated due to the difference in TTL treatment and the behavior of caching resolvers.

Little Focus on the Real Risks
In Section 5 of the report, "Name Collision Etiology", Interisle attempts to explain the origin of the spurious requests for non-existing domains. We believe that merely counting the number of requests for each string is completely insufficient when judging risk, and that any reasonable conclusions made from the data must take into account the true origin of the "collision".

When we looked at the breakdown of requests provided in Table 12, we found that the vast majority of requests either posed no potential risks or risks that could be handled with simple mitigations. It is also important to realize that inadvertent requests of the same pattern in existing TLDs pose the same or greater risk, and that no data was provided to put these results in that context.

In the appendix to this initial letter, we present our own analysis of the risks and possible mitigations for each of the eleven categories proposed by Interisle.

Previous Expansions Caused No Known Issues
A Verisign analysis using data from January 2006[2], prior to the launch of several active TLDs, found that .xxx received more queries before delegation than any other new TLD[3] . Despite having more queries than of all of the TLDs currently under consideration in the "Uncategorized Risk" category, .xxx was delegated in 2011. This TLD launched without incident, and no public complaints or technical issues have been identified since.

In addition, most of the other TLDs listed in Table 1 of the Verisign report, including .asia, .kp, .ax, .um and .cw, also demonstrated much higher numbers of NXDOMAIN responses than all 279 of the "uncategorized" strings, and again all were delegated with no noticeable impact. In fact, the least "dangerous" current gTLD on the chart, .sx, had 331 queries per million in 2006. This is a

[2] "New gTLD Security, Stability, Resiliency Update: Exploratory Consumer Impact Analysis", Verisign Labs Technical Report #1130008 Version 1.0, August 5, 2013, http://techreports.verisignlabs.com/docs/tr-1130008-1.pdf
[3] Ibid, Page 2, Tables 1 and 2

higher density of NXDOMAIN queries than all but five[4] proposed new TLDs. Again, .sx was launched successfully in 2012 with **none** of the problems predicted in these reports.

These successful delegations alone demonstrate that there is no need to delay any more than the two most risky strings.

These Risks Already Exist

As we discuss in the Appendix, many of the risks listed by Interisle or Verisign already exist and many are prevalent in existing gTLDs such as .com. Just as the NTAG would not ask ICANN to halt .com registrations while a twelve month study is performed on these problems, we believe there is no reason to introduce a delay in diversifying the Internet's namespace due to these concerns. Future studies would gain credibility if the listed risks were compared against the situation in current gTLDs.

Risk Measurement is Easily Tampered With

Any model of risk measurement that is based on total query counts is fundamentally flawed. This is especially true when using data collected after the new TLD applications were posted, such as the 2013 DITL corpus.

Query counts are very easily gamed by any Internet connected system, allowing for malicious actors to create the appearance of risk for any string that they may object to in the future. It would be very easy to create the impression of a wide-spread string collision problem with a home Internet connection and the abuse of the thousands of available open resolvers.

It is also possible that publicity around the new gTLD program has caused confused consumers and businesses to try out proposed strings and generate millions of invalid requests.

The Interisle report makes no mention of investigating the possibility that some of these requests were issued intentionally. ICANN should be wary of the precedent that such an easily manipulated metric will be used to make multi-million dollar business decisions.

[4] Ibid, the five potential TLDs with more requests than .sx are home, corp, ice, global and med.

Recommended Next Steps for the Board

There is no reason for ICANN to delay the 279 “uncategorized” names any further, and reasonable protections can be put in place while the existing new TLD calendar is executed. None of these strings pose any more risk than .xxx, .asia and other currently operating TLDs. We believe that the Board can take the following four steps to accelerate the mitigation process without bringing this important expansion of the Internet’s namespace to a halt.

1) Proceed with IDNs

IDNs were not listed in Interisle report, confirming no name collision issues were seen in either certificates[5] or DNS queries[6]. We suggest ICANN to continue to proceed to delegation for IDNs without requiring 120-day waiting period or 30-day mitigation process while staff, applicants and the Board work on deciding risk assessment and mitigation for ASCII TLDs.

2) Proceed with the “Unknown Risk” Strings using the “Low Risk” Mitigations

We believe strongly that the expansion of the namespace will improve the safety, stability and performance of the Internet. We recognize that a small number of applied for names may possibly pose a risk to current operations, but we believe very strongly that there is no quantitative basis for holding back strings that pose less measurable threat than almost all existing TLDs today. This is why we urge the board to proceed with the applications classified as “Unknown Risk” using the mitigations recommended by staff for “Low Risk” strings. We believe the 80% of strings classified as “Low Risk” should proceed immediately with no additional mitigations.

3) Accelerate Handling of the Certificate Collision Issue

NTAG members have discussed the handling of the CA collision issue with prominent members of the Certificate Authority industry and believe that a much more efficient solution exists than the current agreement with the CA/Browser Forum. We believe that the Board can write to the CA/B Forum today and inform them that all but a handful[7] of new TLDs are very likely to be delegated in the next two years and, for the benefit of their customers, the 120 day revocation process should begin today.

---

[5] Not a single IDN gTLD is listed in Appendix C

[6] The IDN gTLD receiving the most queries in 2013 ranked only 364, well into the “low risk” range; in the more reliable 2012 data, the IDN gTLD with the most traffic ranked 565 amongst applied-for gTLDs.

[7] We suggest the “likely to be delegated” list be comprised of strings that have at least one applicant that has passed IE, minus the two “high risk” strings (.corp and .home). Our counterparts in the CA industry have indicated that the largest stumbling block for them was revoking .corp certificates, and that a request for mass revocation without .corp would receive consideration.

Based upon conversations with members of the CA/B Forum, we believe that such a request would be well received by the CA industry.

4) Encourage Investigation by Applicants

ICANN and Interisle should post more detailed breakdowns[8] for each applied-for string in an easily analyzed format so the community can perform more meaningful analysis. As detailed packet capture data cannot be published due to DNS OARC policies, NTAG asks for ICANN to publish raw packet capture data from the L-root after appropriate anonymization[9] so that applicants can perform their own analysis and prepare their own responses.

Please see the appendix for some more technical mitigations supported by the NTAG.

We appreciate your diligent attention to this matter and your continued support of the ICANN community's shared mission to expand the namespace, encourage innovation and diversify the Internet's critical infrastructure.

Sincerely,
The NTAG

[8] Useful metrics for each applied-for string include the counts for the top 1000 SLDs causing NXDOMAINs under each TLD, a list of the top requesting Autonomous Systems, and a breakdown into the 12 query types indentified by Interisle for each TLD.

[9] Converting the requesting IP to AS Number will preserve privacy while allowing the applicant to notify affected parties.

Appendix: Risk Breakdown per Category of Collision

Let's look at the breakdown of the categories Interisle created, their frequency and our assessment of the potential risks and mitigations.

SLD is an existing TLD

2012: 19%, 2013: 15%

Interisle's Description:

Common examples of domain names in this category are:

- *<something>.com.home.*
- *<something>.<CC>.home.*

in which <CC> is a two-letter country code (for example .uk. or .sg.).

These appear to be examples of a valid domain name that has had a TLD appended (incorrectly)—for example, by a commercial off-the-shelf (COTS) router or cable modem used at home or in a small office.

Security Risk: Very low. Interisle implies that the vast majority of these collisions are caused by a misconfiguration by one network and mostly affect one name, .home.

Potential Mitigations: ICANN could restrict the selling of existing gTLDs in the second level while this issue is addressed with the small number of offending ISPs. ccTLDs are already restricted.

SLD is also a proposed TLD

2012: 6%, 2013: 5%

Interisle's Description:

Common examples of domain names in this category are:

- *<something>.<company>.corp.*

in which <company> is the name of a company that matches a proposed TLD.

These appear to be examples of an internal domain name being used (incorrectly) outside the administrative boundary of the company in which it is defined.

Security Risk: Serious for a small number of TLDs. The majority of these requests are likely caused by poorly configured corporate systems being used outside of their home networks. This creates a security issue that can be exploited by any attacker on the local segment, and no changes ICANN makes to the root will increase or decrease this risk.

Potential Mitigations: Interisle does not provide a breakdown of occurrences of this issue, but .corp almost certainly accounts for a plurality if not a majority of this issue. Holding .corp for further study should address the majority of this issue, and for other TLDs the applicant can contact each <company> that issues the largest number of requests so they can address this issue.

SLD is a random 10-alphabetic-character string

2012: 23%, 2013:46%

Interisle's Description:

Often spotted near each other in the traces, examples of domain names in this category are

- *`lfbviakqaw.home.`*
- *`mdqrerrefm.home.`*
- *`uprxbvqnxa.home.`*

These domain names comprise ten apparently random alphabetic characters for the SLD and a proposed TLD (often, but not only, .home.).

Apparently the Google Chrome browser, as a defense against domain name hijacking, generates the pattern of three "random" 10-character alphabetic host names.29 30 When combined with the possibly incorrect addition of a TLD (such as .home.), for example by a home or branch office COTS router, this yields the pattern observed.

While we are sure that not every use of this pattern (10-alphabetic-character SLD and proposed TLD) is caused by this Chrome behavior, we believe from our sampling that it is likely to be the dominant cause.

Security Risk: None. These are random domains specifically generated to detect DNS hijacking, and any incidental collision with real SLDs will have no security effect.

Potential Mitigations: None needed.

Name is www.<proposedTLD>

2012: 0%, 2013:0%

Interisle's Description:

Common examples in this category include:

- *www.youtube.*
- *www.google.*
- *www.yahoo.*
- *www.amazon.*

Possible explanations for these might include typographical errors, in which the TLD (for example, .com.) was omitted unintentionally. A pattern of use introduced by browsers in the 1990s in which .com was appended by default if no TLD was provided is likely to be responsible for this behavior becoming a user habit.

Security Risk: None. All of the strings listed were applied for by the company the user wishes to reach, and in fact enabling this kind of behavior is one of the reasons for .brand TLDs.

Potential Mitigations: No additional mitigations necessary. The Legal Rights Objection mechanism in the new TLD program already sufficiently handles the unlikely possibility of a malicious actor hijacking a .brand.

Name includes _ldap or _kerberos at the lowest level

2012: 3%, 2013:3%

Interisle's Description:

Name includes _ldap or _kerberos at the lowest level Patterns observed show many requests of one of the forms:

- *_ldap._tcp.dc._msdcs.<etc.>*
- *_ldap._tcp.pdc._msdcs.<etc.>*
- *_ldap._tcp.<etc.>._sites.dc._msdcs.<etc.>*
- *_ldap._tcp.<etc.>._sites.gc._msdcs.<etc.>*
- *_ldap._tcp.<etc.>._sites.<etc.>*
- *_ldap._tcp.<etc.>*
- *_kerberos._tcp.dc._msdcs.<etc.>*
- *_kerberos._tcp.<etc.>._sites.dc._msdcs.<etc.>*
- *_kerberos._tcp.<etc.>._sites.<etc.>*
- *_kerberos._tcp.<etc.>*

Typically these are queries for SRV records, although sometimes they are requests for SOA records. They appear to be related to Microsoft Active Directory services.

Security Risk: No additional risk from new TLDs. There are serious attacks possible for local segment attackers who hijack Active Directory DNS requests, but there is absolutely no additional risk from new TLDs. Names with an underscore are legal in DNS but are not allowed as SLDs, so as long as this practice holds there is no possibility of a conflict that creates a security risk.

Potential Mitigations: None needed. ICANN should consider adding queries for these and other disallowed names into the SLA monitoring process to insure compliance by new and existing registries.

Name includes _dns-sd at one level, often the 3rd or 4th level of the name

2012: 1%, 2013:1%

Interisle's Description:

Patterns observed show many requests of the form:

- *`<something>._dns-sd._udp.<proposedTLD>`*

Occasionally two levels occur after `._udp.`

Typically these are lookups for PTR records, although sometimes they are queries for TXT records. They appear to be related to Apple's service discovery service (Bonjour or multicast DNS).

Security Risk: No additional risk from new TLDs. There are serious attacks possible for local segment attackers who Boujour requests, but there is absolutely no additional risk from new TLDs. Names with an underscore are legal in DNS but are not allowed as SLDs, so as long as this practice holds there is no possibility of a conflict that creates a security risk.

Potential Mitigations: None needed.

Name starts with "File moved-http://"

2012: 18%, 2013:0%

Interisle's Description:

The predominant form of DNS name that has been observed in this category is of the form:

- *`File moved-http://www.whatismyip.<M>.<N>.Home.`*

in which <M> and <N> are one- to three-digit numbers (presumably two quads of an IPv4 address).

These all appear to be queries for A records.

Security Risk: None. This seems to be a misconfiguration related to whatismyip.com that was fixed in 2012. No queries were seen in 2013.

Potential Mitigations: None needed.

Name includes _sip, _sipinternal, _sipinternaltls, _sipfederationtls, or _sips at the lowest level

2012: 0%, 2013:0%

Interisle's Description:
Patterns observed show many requests of the form:

- *_sip._tcp.<etc.>*
- *_sip._udp.<etc.>*
- *_sip._tls.<etc.>*
- *_sipinternal._tcp.<etc.>*
- *_sipinternaltls._tcp.<etc.>*
- *_sipfederationtls._tcp.<etc.>*

These are mostly SRV record lookups, which are typically associated with communication services such as Voice over IP (VoIP) telephony and video messaging.

Security Risk: None. Names with an underscore are legal in DNS but are not allowed as SLDs, so as long as this practice holds there is no possibility of a conflict that creates a security risk.

Potential Mitigations: None needed.

Name includes _xmpp-client, _xmpp-server, or _xmpconnect at the lowest level

2012: 0%, 2013:0%

Interisle's Description:
Patterns observed show many requests of the form:

- *_xmpp-client.<etc.>*
- *_xmpp-server.<etc.>*
- *_xmppconnect.<etc.>*

Typically the first two are queries for SRV records and the latter are queries for TXT records. These names are likely associated with attempts to discover XMPP32 (a.k.a. Jabber) messaging services.

Security Risk: None. Names with an underscore are legal in DNS but are not allowed as SLDs, so as long as this practice holds there is no possibility of a conflict that creates a security risk.

Potential Mitigations: None needed.

Name includes mail at the lowest level, and/or as the SLD

2012: 0%, 2013:0%

Interisle's Description:

While many DNS names include `mail` *at one level, many appear to fall under one of the categories described above as "SLD is an Existing TLD" or "SLD is Also a Proposed TLD." These are predominantly lookups for A or AAAA records, but requests for other resource record types such as MX, TXT, and SRV have also been observed.*

Security Risk: Moderate. Some of these requests will be unqualified requests in mail agents and browsers for "mail". If so, then there is already a large risk of compromise when faced with a local segment attacker. The risk is elevated slightly by the delegation of strings that currently see a large number of `mail.<tld>` requests, but this risk is appropriately dealt with with SLD controls.

Potential Mitigations: More information is needed on which proposed TLDs are most affected, but it is likely that eliminating .corp from the list will mostly mitigate this risk. Disallowing registration of the mail SLD on other TLDs with a high number of `mail.<tld>` requests would also be appropriate but should not slow delegation.

Name comprises just the proposed TLD

2012: 0%, 2013:0%

Interisle's Description:

Examples of these are where the proposed TLD appears as the only part of the DNS name.

Commonly appearing proposed TLD strings are:

- *.home.*
- *.cisco.*
- *.honda.*
- *.toshiba.*
- *.ericsson.*

Security Risk: None. All of the strings listed were applied for by the company the user wishes to reach, and in fact enabling this kind of behavior is one of the reasons for .brand TLDs, although these names will not resolve unless apex A records are allowed.

Potential Mitigations: No additional mitigations necessary. The Legal Rights Objection mechanism in the new TLD program already sufficiently handles the unlikely possibility of a malicious actor hijacking a .brand, and the process for requesting permission to use an apex record make exploiting this risk almost impossible.