

1 Managing the Risks of Top-Level Domain
2 Name Collisions

3 Findings for the Name Collision Analysis Project (NCAP) Study 1

4 Karen Scarfone, Scarfone Cybersecurity

5

6 February 12, 2020

7

DRAFT

8 Table of Contents

9	1	Study Overview	1
10	2	Name Collision Primer	3
11	2.1	Domains.....	3
12	2.2	Name Collisions	4
13	3	Review of Previous Work.....	6
14	3.1	DNS Wildcard Address Records: 2003 – 2009	6
15	3.2	Collisions from Failure to Renew a Domain: 2006.....	8
16	3.3	Initial TLD Delegation Concerns: 2008 – 2013	8
17	3.3.1	Invalid TLD Queries Reaching Root Servers.....	8
18	3.3.2	Certificates for Internal Domains That May Also Become gTLDs.....	9
19	3.3.3	Verisign Labs Report on New gTLD Security and Stability.....	10
20	3.3.4	PayPal Concerns about Delegating Certain gTLDs	10
21	3.3.5	Internet-Draft on TLD Delegation Procedures.....	11
22	3.4	gTLD Risk Profiles: 2013 – 2014.....	11
23	3.4.1	ICANN Report from Interisle Consulting Group.....	11
24	3.4.2	ICANN Proposal on New gTLD Collision Risk Mitigation	12
25	3.4.3	Public Comments on ICANN Proposal.....	13
26	3.4.4	ICANN Proposal on New gTLD Collision Occurrence Management	15
27	3.4.5	DNS-OARC Workshop Session on High-Risk Strings Collisions.....	16
28	3.4.6	SSAC Advisory SAC 062 on Mitigating Name Collision Risk.....	17
29	3.4.7	SLD Blocking List Effectiveness	17
30	3.5	Research on Name Collision Causes: 2013 – 2016.....	18
31	3.5.1	Search List Processing and FQDN Usage	18
32	3.5.2	Causes of Internal Domain Leakage.....	20
33	3.5.3	Detection of Leaking Clients.....	21
34	3.6	Name Collision Occurrence Management Framework: 2014 – 2015	22
35	3.6.1	JAS Global Advisors Phase One Report Draft.....	22
36	3.6.2	Public Comments on Phase One Report Draft	24
37	3.6.3	JAS Global Advisors Final Phase One Report.....	25
38	3.6.4	SSAC Response to the Final Phase One Report.....	26
39	3.6.5	Approval of the Name Collision Occurrence Management Framework	28
40	3.6.6	Controlled Interruption for New ccTLDs.....	28
41	3.6.7	JAS Global Advisors Phase Two Report.....	29

42	3.7	Potential Changes to Existing gTLD Processes: 2016 – present.....	30
43	3.7.1	ICANN New gTLD Subsequent Procedures (SubPro) Working Group	30
44	3.7.2	Requests to Delegate corp, home, and mail.....	32
45	4	The Known Harm of Name Collisions and the Technical Impact of Controlled Interruption..	34
46	4.1	Preparation.....	34
47	4.2	Name Collision Reports.....	35
48	4.2.1	Reports to ICANN	35
49	4.2.2	Reports to Others	37
50	5	Datasets for Name Collision Studies	38
51	5.1	Datasets Used in Past Studies.....	38
52	5.2	Additional Datasets Needed for Studies 2 and 3	39
53	6	Bibliography	40
54	7	Acronyms.....	52
55			

DRAFT

56 Acknowledgments

57 Scarfone Cybersecurity thanks The Internet Corporation for Assigned Names and Numbers
58 (ICANN), the ICANN Office of the Chief Technology Officer (OCTO), and ICANN’s Name
59 Collision Analysis Project Discussion Group (NCAP DG) for their support and insights during the
60 development of this draft report.

61

DRAFT

62 1 Study Overview

63 This report presents the findings for Study 1 of the Name Collision Analysis Project (NCAP) [1].
64 The purpose and scope of Study 1 were defined in a July 2019 Request for Proposal (RFP) [2]. The
65 initial draft of this report addresses the first two goals of Study 1, as stated in the RFP:

- 66 “1. Production of a summary report on the topic of name collision that brings forth important
67 knowledge from prior work in the area. The report will be a primer for those new to the
68 subject. The report will be based on an examination of all relevant prior work on the issue
69 of name collisions.
- 70 2. Creation of a list of datasets used in past name collision studies; an identification of gaps, if
71 any; and creation of a list of additional data sets that would be required to successfully
72 complete Studies 2 and 3.” [2]

73 For the purposes of Study 1, the term *name collision* “refers to the situation where a name that is
74 defined and used in one namespace may also appear in another. Users and applications intending to
75 use a name in one namespace may attempt to use it in a different one, and unexpected behavior
76 may result where the intended use of the name is not the same in both namespaces. The
77 circumstances that lead to a name collision could be accidental or malicious.

78 Study 1 concerns name collisions in the context of top-level domains (TLDs), where the
79 conflicting namespaces are:

- 80 • the global Internet Domain Name System (DNS) namespace reflected in the root zone
81 overseen by the Internet Assigned Numbers Authority (IANA) Function; and
- 82 • any other namespace, regardless of whether that other namespace is intended for use with
83 the DNS or any other protocol.” [2]

84 Also from the RFP:

85 “Name collision refers to the situation in which a name that is used in one namespace may
86 be used in a different namespace, where users, software, or other functions in that domain
87 may misinterpret it. In the context of top level domains, the term ‘name collision’ refers to
88 the situation in which a name that is used in the global Domain Name System (DNS)
89 namespace defined in the root zone as published by the root zone management (RZM)
90 partners ICANN and VeriSign (the RZM namespace) may be used in a different namespace
91 (non-RZM), where users, software, or other functions in that domain may misinterpret it.”
92 [2]

93 The initial draft of this report contains the following sections addressing tasks from the RFP:

- 94 • Section 2 contains a name collision primer (task 2a).
- 95 • Section 3 provides a review of pertinent previous work (tasks 1 and 2c).

96 • Section 4 details evidence of harm caused by name collisions (task 2b) and discusses the
97 technical impact of name collision mitigation techniques employed to date (task 2d).

98 • Section 5 assesses datasets used in past name collision studies, identifies additional datasets
99 that would be needed for Studies 2 and 3, and discusses the availability of those additional
100 datasets (tasks 3, 4, and 5, respectively).

101 All sources referenced in this report are cited in Section 6, Bibliography.

102

DRAFT

103 2 Name Collision Primer

104 This section explains the basics of name collisions. Readers who are already well-versed on the
105 topic of name collisions should still read this section because it defines new terms for the purposes
106 of this report and establishes the scope for the report. Other concepts in this section are based on
107 material from the ICANN Acronyms and Terms tool. [3]

108 2.1 Domains

109 A *domain name* maps to a piece of data, like an IP address. For example, icann.org is a domain
110 name for the ICANN organization. You can use “icann.org” to reach ICANN’s computing
111 resources, like websites and email servers, instead of typing in icann.org’s IP address every time
112 you want to access an ICANN website.

113 Every domain name consists of one or more labels, and the labels go from most specific on the left
114 to least specific on the right. The icann.org domain name has two labels. “org” is the label for the
115 *top-level domain (TLD)*. “icann” is the label for the *second-level domain (SLD)*. The SLD is a
116 domain name that is associated with a TLD—in other words, an SLD is *registered* to a TLD. There
117 are usually many SLDs registered to a single TLD. Many domain names have three or more levels,
118 such as “www.icann.org”, but for the purposes of this explanation, we will focus on the highest
119 two levels (TLDs and SLDs) only.

120 In the past, there were a small number of *generic TLDs (gTLDs)* like com and org. [4] A few more
121 were added in 2000 and in 2004. [5] Efforts began in 2005 to consider adding many more gTLDs,
122 and in October 2013 the first of these new gTLDs was made available for usage on the Internet, a
123 process better known as *delegation*. The gTLDs are frequently referred to as “names” or “strings”,
124 so when you see a term like “delegated strings”, it just means that a new gTLD was made available
125 on the Internet. For more information on gTLDs, see ICANN’s resources [6] [7] [8].

126 In addition to gTLDs, there are also TLDs specific to country names—*country code TLDs*
127 (*ccTLDs*). The original ccTLDs were all two letters long, such as fr and us, taken from the two-
128 letter country codes in International Organization for Standardization (ISO) 3166, *Codes for the*
129 *representation of names of countries and their subdivisions*. [9] The two-letter ccTLDs have since
130 become known as *ASCII ccTLDs*. In 2009, ICANN approved an effort to delegate new
131 internationalized domain name ccTLDs (IDN ccTLDs) through what is called the IDN ccTLD Fast
132 Track Process. [10] IDN ccTLDs use non-Latin characters, such as the alphabet of the primary
133 language spoken in a particular country. For more information on ccTLDs, see ICANN’s resources
134 [11] [12] [13].

135 In this report, usage of “TLD” refers to both gTLDs and ccTLDs.

136 2.2 Name Collisions

137 There are many forms of name collisions. To understand what name collisions are and which types
138 are in scope for this report, let’s look at four examples, which are based on Section 2.3.3 of the
139 study RFP [2]. Each example maps to one or more of the situations described in the RFP.

Explanation of Name Collision Type	Mapping to RFP
<p>1. Suppose that Alice uses .EXAMPLE internally only as her top-level domain, which works without ambiguity because .EXAMPLE is not a TLD delegated on the Internet. If Alice types “www.example” in a web browser, it would take her to her own website. The next year, .EXAMPLE is delegated as a new TLD. Now when Alice tries to access “www.example”, it’s no longer clear whether she is trying to access her own website or the new public domain on the Internet. The .EXAMPLE used internally by Alice and the .EXAMPLE used publicly by someone else <i>collide</i>. This report will refer to these as <i>duplicate name collisions</i>—the collision is caused by the same TLD being used in two places at the same time.</p>	<p>A.a: User Alice intentionally uses .EXAMPLE in a non-RZM context and .EXAMPLE is now delegated in the public DNS. User Alice suffers adverse impact as a result.</p> <p>A.b: User Alice unintentionally uses .EXAMPLE in a non-RZM context (for example as the result of a software behavior) and .EXAMPLE is now delegated in the public DNS. User Alice suffers adverse impact as a result.</p>
<p>2. Suppose that Alice uses shortened forms of domain names—for example, she might type “dashboard.example” instead of “dashboard.example.com”—and there’s a list of domain suffixes like “.com” that automatically get appended to what she typed in order to find the desired domain. This is known as <i>search list processing</i>, and this works as long as there’s no TLD for .example. However, the next year, .EXAMPLE is delegated as a new TLD. When Alice wants to go to “dashboard.example.com” and types “dashboard.example”, she’ll be taken to the latter instead of the former. This report will refer to these as <i>shortened name collisions</i>—the collision is caused by someone using a shortened name that matches a TLD being used elsewhere at the same time.</p>	<p>A.c: Registrant Alice uses EXAMPLE as a label anywhere except as a non-RZM TLD, and relies on search list processing where the label EXAMPLE is the terminal label, as an intermediate step in that search list processing. (<i>e.g. User searches for dashboard.example.com by typing in dashboard.example</i>) .EXAMPLE is now registered in the public DNS and the search list processing behavior of Alice now changes.</p>

Explanation of Name Collision Type

Mapping to RFP

3. Suppose that there is a public domain EXAMPLE.COM, and Alice uses it as her domain. The next year, .EXAMPLE is delegated as a new TLD. Some external users might have search list processing that automatically appends the “.com” domain suffix to requests, so some queries for .EXAMPLE domains may mistakenly go to .EXAMPLE.COM instead. Alice will be receiving traffic that she was not intended to receive. This report will refer to these as *search list name collisions*—the collision is caused by the search list not recognizing .EXAMPLE as a new TLD and instead going through its search list to try to find the domain.

B.a: Registrant Alice uses EXAMPLE.COM (or EXAMPLE.TLD where TLD is any current TLD in the public DNS) and .EXAMPLE is now registered in the public DNS. Registrant Alice now receives multiple queries as a result of search list processing of users of domains under .EXAMPLE

4. Suppose that Alice registers a TLD or SLD and uses it for some time, then lets it expire. Subsequently someone else registers the same domain and delegates it. Now queries looking for the old domain (for Alice) will go to the new domain (for someone else). This report will refer to these as *re-registered name collisions*—the collision is caused by someone registering a domain that was previously registered by someone else.

B.b: Registrant Alice uses .EXAMPLE as a TLD in the public DNS and then lets the registration expire. Registrant Bob then registers and delegates .EXAMPLE. Traffic intended for Alice’s use of .EXAMPLE is now received by Bob’s use of .EXAMPLE

B.c: Registrant Alice uses EXAMPLE.COM and then lets the registration expire. Registrant Bob then registers and delegates EXAMPLE.COM. Traffic intended for Alice’s use of EXAMPLE.COM is now received by Bob’s use of EXAMPLE.COM

140

141 All four of these types of name collisions are in scope for Study 1. Only duplicate name collisions
142 and shortened name collisions (types A.a, A.b, and A.c from the RFP) are in scope for Section 5 of
143 this report (on data sets for Studies 2 and 3). No other types of name collisions are in scope for any
144 parts of Study 1.

145 For more information on name collisions, see ICANN’s resources [14].

146 3 Review of Previous Work

147 This section provides a review of previous work on name collisions. All reviewed work meets at
148 least one of the following criteria from the Study 1 RFP:

- 149 “i. Peer reviewed paper
- 150 ii. Report/Analysis based on data
- 151 iii. Qualitative research on name collision experience
- 152 iv. Proposed or agreed technical standards” [2]

153 The review is broken into several sections based on timeframe and topic area. It is largely
154 chronological, but some items are intentionally out of sequence—for example, it may have taken a
155 few years to finalize a standard on a particular topic, so that standard is included in the topic area’s
156 section, where it fits thematically, instead of a later section where it would fit chronologically.

157 In cases where the previous work includes correspondence on a particular document, such as
158 public comments on a draft report, the review points to the archived correspondence as a whole
159 and does not list or mention each piece of correspondence. In some cases, particular pieces of
160 correspondence are mentioned and discussed. This does not imply that only the cited
161 correspondence is relevant; often several parties made similar points, so one or a few instances are
162 cited as examples, and readers are encouraged to read the others if they desire.

163 Each subsection within this section indicates which type or types of name collisions are applicable
164 to its contents, if any.

165 3.1 DNS Wildcard Address Records: 2003 – 2009

166 *Applicability: No name collision types (background material)*

167 In September 2003, Verisign launched what they called Site Finder. Site Finder changed how
168 requests for nonexistent domain names were handled by adding a DNS wildcard address record
169 that matched to every com and net address that didn’t otherwise have a match. People and services
170 were used to the previous behavior and were unaware it was changing, so the sudden deployment
171 of DNS wildcard address records inadvertently caused a lot of problems. Then the workarounds for
172 the problems caused even more problems. [15] [16] [17] [18]

173 Note that while Site Finder’s launch was the event that brought a great deal of attention to the
174 subject of DNS wildcard address records, the possibility of domain name requests being resolved
175 in unintended ways was not a new one, with formal treatments of the subject going back to at least
176 1993. [19] What made Site Finder so noteworthy was that it affected many people and services at
177 one time. Site Finder in particular, and the use of DNS wildcard address records more broadly, did
178 not cause name collisions; however, they are relevant to this study because there are obvious
179 parallels between wildcard address records and name collisions. Both involve domain name
180 queries being resolved in unexpected ways that can disrupt Internet usage for affected parties. So

181 reviewing the recommendations for avoiding another Site Finder-like incident helps indicate
182 potential ways of avoiding negative impacts from name collisions as well.

183 The ICANN Security and Stability Advisory Committee (SSAC) conducted a review of Site Finder
184 and DNS wildcard address records, and issued their findings and recommendations in report SAC
185 006 in July 2004. [16] Most pertinent to this study is recommendation 4 from SAC 006:

186 “Changes in registry services should take place only after a substantial period of notice,
187 comment and consensus involving both the technical community and the larger user
188 community. This process must (i) consider issues of security and stability, (ii) afford ample
189 time for testing and refinement and (iii) allow for adequate notice and coordination with
190 affected and potentially affected system managers and end users. Thirty years of experience
191 show that this strategy ensures robust engineering and engenders trust in the systems and
192 the processes surrounding their maintenance and development.”

193 Also of particular relevance for this study is the Reserved Names portion of the “Problems
194 encountered in recent experiences with wildcards” section of [15], which is also duplicated by
195 [16]:

196 “This sort of wildcard usage is incompatible with any use of DNS which relies on reserving
197 names in a registry with the express intent of not adding them to the DNS zone itself. An
198 example of such a use is the JET-derived IDN approach of ‘registry restrictions’ and
199 ‘reserved names’, which depends on the existence of names that are reserved and can be
200 registered only by the holder of some related name, but which do not appear in the DNS.
201 By some readings of the current ICANN IDN policy, support for that ‘reserved name’
202 approach is required. To accomplish the goal of reduced consumer confusion, the reserved
203 names must not be resolvable at all. This reserved name approach appears to be completely
204 incompatible with this sort of wildcard usage: since the wildcard will always cause a result
205 to be returned, even for a reserved name which does not appear in the zone, one can support
206 either one or the other, but not both.”

207 In November 2006, the SSAC posted SAC 015 [17], an advisory explaining why wildcards should
208 not be used for TLDs. It explained how wildcards work and gave examples of problems that
209 resulted in the past from wildcard resource records in TLDs. SAC 015 referenced previous SSAC
210 work and also cited a report from the ICANN Registry Services Technical Evaluation Panel [20]
211 that reached the same conclusion as SAC 006 did: wildcards were too risky to be used in TLDs.

212 SAC 032 [18] was released in June 2008. It contained a broader discussion of DNS response
213 modification, with wildcards part of that discussion. SAC 032 provided preliminary
214 recommendations for addressing DNS response modification, including this: “SSAC concurs with
215 the IAB and recommends that entrusted agents should not use DNS wildcards in a zone without
216 informing the domain registrant of the risks identified in this Report and elsewhere, that entrusted
217 agents should not generate wildcards and synthesized responses without the informed consent of
218 the registrant, and that entrusted agents should provide opt-out mechanisms that allows clients to
219 receive the original DNS answers to their queries.”

220 SAC 041, published in June 2009, summarized the SSAC’s study of DNS wildcarding since 2004
221 and advised “ICANN that new TLDs, including both new gTLDs and new ccTLDs, should not use

222 DNS redirection and synthesized DNS responses. [...] The redirection and synthesizing of DNS
223 responses by TLDs poses a clear and significant danger to the security and stability of the domain
224 name system.” [21] The references to redirection and synthesizing included the use of DNS
225 wildcard address records.

226 For additional information on Site Finder, see the ICANN Archives for Verisign’s Wildcard
227 Service Deployment. [22]

228 3.2 Collisions from Failure to Renew a Domain: 2006

229 *Applicability: Re-registered name collisions*

230 Re-registered name collisions can occur when someone fails to renew a domain and someone else
231 subsequently acquires the same domain. This topic was extensively discussed in SAC 010 [23] and
232 SAC 011 [24], both from June 2006.

233 SAC 010 provided information and guidance for registrants. From SAC 010: “...registrants may
234 not appreciate that expired domain names are commonly registered to another registrant within a
235 few weeks or months of the date of expiry of the domain name registration agreement. The new
236 registrant may not use the domain name for the same purposes as a former registrant. Incidents
237 show that previously registered domain names may be exploited, at the expense of the reputation of
238 a former registrant. In this Advisory we refer to this form unexpected consequence as *reputational*
239 *harm.*”

240 SAC 011 provided more technical information on the situation, with examples of the disruptions
241 that could be caused by a domain name not being renewed and a malicious party subsequently
242 renewing it. In this situation, the attacker could receive traffic that was intended for the
243 organization that originally had the domain.

244 Both SAC 010 and SAC 011 made it clear that it is ultimately the registrant’s responsibility to
245 ensure they renew their domains in a timely fashion. The importance of keeping contact
246 information up to date was emphasized.

247 3.3 Initial TLD Delegation Concerns: 2008 – 2013

248 3.3.1 Invalid TLD Queries Reaching Root Servers

249 *Applicability: Duplicate name collisions, re-registered name collisions*

250 ICANN started work in mid-2008 to figure out processes for parties to apply for new gTLDs and
251 for ICANN to evaluate the applications. [25] In 2009, the SSAC was asked to look at the issue of
252 invalid TLD queries reaching root servers, stemming from someone registering a TLD that others
253 had already been using and the root server had been responding to. At that time, an estimated 26%
254 of all query load at root servers was invalid TLDs. One of the earliest mentions of this problem
255 was in a June 2009 blog posting. [26] The SSAC studied the issue and released their SAC 045
256 report in November 2010. [27] SAC 045 cited the possibility of someone applying for a TLD that

257 had appeared in queries before or had been issued before, and the problems with queries that had
258 been failing suddenly succeeding once the TLD was delegated.

259 SAC 045 acknowledged how difficult it would be to eliminate inadvertent instances of such
260 problems: “It is likely that many of the same conditions that cause the current set of invalid TLD
261 queries to appear at the root level of the DNS will persist despite efforts to encourage end users,
262 private networks, software and equipment manufacturers to correct configuration and
263 programming errors.”

264 SAC 045 had recommendations for reducing other instances of query ambiguity. Recommendation
265 2 said, “Prohibit the delegation of certain TLD strings. Internet Engineering Task Force (IETF)
266 Request for Comments (RFC) 2606, ‘Reserved Top Level Domain Names,’ currently prohibits a
267 list of strings, including test, example, invalid, and localhost.” Section 2.2.1.2 of ICANN’s *gTLD*
268 *Applicant Guidebook* [25] released in June 2012 specified names that could not be gTLDs; these
269 included the names from RFC 2606 [28], plus a few dozen more, in what was termed the Top-
270 Level Reserved Names List.

271 3.3.2 Certificates for Internal Domains That May Also Become gTLDs

272 *Applicability: Duplicate name collisions*

273 In November 2012, the SSAC became aware of a problem with certificate issuance that could
274 negatively affect new gTLD delegation. At that time, the application period for new gTLDs had
275 already closed but no new gTLDs had yet been delegated. In March 2013, the SSAC released SAC
276 057. [29] SAC 057 was an advisory cautioning the ICANN Board about certificate authorities
277 (CAs) issuing certificates for internal domains that are also TLDs. One type of certificate is called
278 an Internal Name certificate and is meant for an organization’s internal use only, so the CA cannot
279 resolve the name or look up the owner. Certificates can also contain Subject Alternative Names,
280 which are supposed to be additional names for the same domain. At that time, the typical practice
281 was for CAs to confirm that the internal domains were not the same as an already-delegated TLD;
282 CAs were not checking the list of applied-for new gTLDs to see if there was a match. This could
283 allow someone to get an Internal Name certificate matching a domain name that would soon be a
284 TLD. SAC 057 presented a case study showing how an SSAC member was able to get a certificate
285 issued for the www.site domain.

286 The most important findings from SAC 057 related to name collisions were the following:

- 287 • From Finding 3: “There are at least 37,000 internal name certificates used in thousands of
288 enterprises. [...] with the introduction of new gTLDs, namespace collisions and other man-
289 in-the-middle attacks (see Finding 4) will become more apparent. In addition, because
290 many of the applied for TLDs are common, generic terms the risk of collisions increases.”
- 291 • Finding 4: “The practice for issuing internal name certificates allows a person, not related
292 to an applied for TLD, to obtain a certificate for the TLD with little or no validation, and
293 launch a man-in-the-middle attack more effectively.”

294 SAC 057’s recommendations included “requesting that they [CAs] treat applied for new gTLDs as
295 if they were delegated TLDs as soon as possible....”

296 3.3.3 Verisign Labs Report on New gTLD Security and Stability

297 *Applicability: Duplicate name collisions, shortened name collisions, search list name collisions*

298 Verisign Labs issued two similar versions of a technical report, “New gTLD Security and Stability
299 Considerations” in March 2013 (version 2.1 [30] and version 2.2 [31]). The report noted the lack of
300 data and metrics on queries for TLDs, and how this meant the impact of delegating new gTLDs
301 would not be quantifiable. The report emphasized the technical and logistical complexity of new
302 gTLD adoption for registry operators, and it pushed back on ICANN’s timelines for the new gTLD
303 program, saying it did not give registry operators enough time to prepare.

304 Section V of the Verisign Labs report discussed name collisions specifically:

- 305 • Subsection A referenced the SAC 057 advisory and acknowledged the benefits of the
306 advisory’s recommendation to have CAs “treat applied for new gTLDs as if they were
307 delegated TLDs” when issuing Internal Name certificates. However, it also criticized the
308 changes to the proposed gTLD delegation processes as still providing a window of
309 opportunity for attackers.
- 310 • Subsection C cited the lack of studies to identify the complex, subtle issues of name
311 collisions. One particularly noteworthy statement was, “...the introduction of .info over a
312 decade ago highlighted just what sort of obvious and nuanced interdependencies may exist
313 as new gTLDs are delegated and made available on the Internet while applications and
314 other systems are ill-prepared.”

315 3.3.4 PayPal Concerns about Delegating Certain gTLDs

316 *Applicability: Duplicate name collisions*

317 In mid-March 2013, between the release of versions 2.1 and 2.2 of the Verisign Labs report,
318 PayPal sent a letter to ICANN [32] regarding SAC 045 [27] and RFC 6762 [33]. The PayPal letter
319 warned ICANN of issues with delegating certain gTLDs:

320 “ICANN should consider not just the potential costs and unwanted network traffic sent to
321 applicants for these names, but the substantial and severe costs imposed on the general
322 Internet community arising from delegation of names that have been common *de facto*
323 private network suffixes for nearly two decades. At minimum, the top ten observed invalid
324 TLDs plus those recommended for use by RFC 6762 should be permanently reserved for
325 private use to prevent large scale disruption and damage to the millions of users and
326 systems that rely upon them today. A more prudent approach would be to consider the
327 negative externalities for each of the applied for new gTLDs.”

328 The 13 names that PayPal recommended be permanently reserved were: invalid, wpad, home,
329 belkin, corp, lan, domain, localdomain, localhost, local, intranet, internal, and private.

330 3.3.5 Internet-Draft on TLD Delegation Procedures

331 *Applicability: Duplicate name collisions*

332 On May 2, 2013, an Internet-Draft proposing procedures for TLD delegation was released. [34] It
333 was authored by representatives of NLnet Labs, Dyn, and Google. It was updated twice in the
334 following three months, with the last draft published on August 1. [35] The Internet-Draft
335 specifically addressed situations where queries for a never-delegated TLD had already been seen
336 on the Internet, and the TLD was subsequently going to be delegated.

337 The Internet-Draft emphasized the potential consequences if commonly used internal TLDs were
338 delegated as public TLDs, citing SAC 045 and RFC 6762. These consequences included security
339 issues. Section 2 stated, “Responsible administration of the public namespace therefore requires
340 great care in permitting public delegation of any name where there is good reason to suppose it is
341 in widespread use as a private namespace....” Section 2.1 gave a hypothetical example of a name
342 collision caused by an organization using an internal subdomain, corp, that subsequently was
343 delegated as a TLD, so it was no longer clear how queries for corp names should be resolved.

344 Section 3 recommended that zone operators monitor the frequency of queries for nonexistent
345 domains. Such domains that receive the most queries should not be delegated as public TLDs.
346 Section 3.2 sketched out parts of a methodology zone operators could use to determine which
347 names are most likely to be problematic. However, the authors did not update the Internet-Draft
348 further, and it expired.

349 3.4 gTLD Risk Profiles: 2013 – 2014

350 3.4.1 ICANN Report from Interisle Consulting Group

351 *Applicability: Duplicate name collisions*

352 The next major milestone was the August 2, 2013 release of an ICANN report from Interisle
353 Consulting Group titled, “Name Collision in the DNS: A study of the likelihood and potential
354 consequences of collision between new public gTLD labels and existing private uses of the same
355 strings, version 1.5.” [36] It was studying duplicate name collisions for gTLDs only: an internally
356 used domain is subsequently delegated as a new gTLD. The Interisle study reviewed much of the
357 same body of work as this NCAP Phase 1 report (through mid-2013), but its most significant
358 contribution was to analyze data sets collected from several root servers of all DNS requests they
359 received (totaling 94 billion, from the DNS-OARC Day in the Life of the Internet data [37]), plus a
360 global DNS resolver service providing 53 billion requests it saw.

361 The findings of the Interisle study that are most noteworthy within the context of this report were:

- 362
- 363 • **“The potential for name collision with proposed new gTLDs is substantial.** Based on
364 the data analyzed for this study, strings that have been proposed as new gTLDs appeared in
365 3% of the requests received at the root servers in 2013. Among all syntactically valid TLD
labels (existing and proposed) in requests to the root in 2013, the proposed TLD string

366 home ranked 4th, and the proposed corp ranked 21st. DNS traffic to the root for these and
367 other proposed TLDs already exceeds that for well-established and heavily-used existing
368 TLDs.”

369 • **“The delegation of almost any of the applied-for strings as a new TLD label would**
370 **carry some risk of collision.** Of the 1,409 distinct applied-for strings, only 64 never
371 appear in the TLD position in the request stream captured during the 2012 ‘Day in the Life
372 of the Internet’ (DITL) measurement exercise, and only 18 never appear in any position. In
373 the 2013 DITL stream, 42 never appear in the TLD position, and 14 never appear in any
374 position.”

375 • **“The designation of any applied-for string as ‘high risk’ or ‘low risk’ with respect to**
376 **delegation as a new gTLD depends on both policy and analysis.** This study provides
377 quantitative data and analysis that demonstrate the likelihood of name collision for each of
378 the applied-for strings in the current new gTLD application round and qualitative
379 assessments of some of the potential consequences. Whether or not a particular string
380 represents a delegation risk that is ‘high’ or ‘low’ depends on policy decisions that relate
381 those data and assessments to the values and priorities of ICANN and its community; and
382 as Internet behavior and practice change over time, a string that is ‘high risk’ today may be
383 ‘low risk’ next year (or vice versa).”

384 • **“For a broad range of potential policy decisions, a cluster of proposed TLDs at either**
385 **end of the delegation risk spectrum are likely to be recognizable as ‘high risk’ and**
386 **‘low risk.’** At the high end, the cluster includes the proposed TLDs that occur with at least
387 order-of-magnitude greater frequency than any others (corp and home) and those that occur
388 most frequently in internal X.509 public key certificates (mail and exchange in addition to
389 corp). At the low end, the cluster includes all of the proposed TLDs that appear in queries
390 to the root with lower frequency than the least-frequently queried existing TLD; using 2013
391 data, that would include 1114 of the 1395 proposed TLDs.”

392 In summary, the Interisle study concluded that there was a risk of name collision with practically
393 any new gTLD, but that most gTLDs would be low risk because there were few queries already
394 being seen by the root servers for those domain names. A small number of TLDs were already
395 mistakenly requested so often that it would surely cause significant disruptions to delegate them as
396 public gTLDs.

397 3.4.2 ICANN Proposal on New gTLD Collision Risk Mitigation

398 *Applicability: Duplicate name collisions*

399 A few days after the Interisle study [36] was released, ICANN posted a proposal called “New
400 gTLD Collision Risk Mitigation.” [38] It proposed how the risks identified in the Interisle study
401 [36] could be mitigated. The proposal was based on the data sets used by Interisle and information
402 provided by CAs on the domains specified within Internal Name certificates they had issued. The
403 proposal defined three risk profiles for applied-for new gTLDs:

- 404 • **Low-risk:** there were fewer queries being received for the not-yet-delegated TLD at the
405 root servers as there were for other delegated TLDs that were “empty”. The low-risk
406 profile fit roughly 80% of the applied-for new gTLDs.
- 407 • **High-risk:** the number of queries being received for the not-yet-delegated TLD at the root
408 servers was an order of magnitude higher than other such queries. The high-risk profile
409 would fit two names, home and corp. Also, corp was the string most often seen in Internal
410 Name certificates.
- 411 • **Uncalculated-risk:** there was not enough information to determine if these were low or
412 high risk. This was roughly 20% of all applied-for new gTLDs.

413 The proposal included recommendations for mitigating the risk for each of the three risk profiles.
414 Low-risk TLDs could be delegated, with a mandatory 120-day waiting period between signing an
415 agreement and activating names. Also, for at least the first 30 days after first delegating a TLD, the
416 registry operator would not activate names under the TLD, and during that time would notify the
417 appropriate contacts for any IP address that requested a name under the TLD. High-risk TLDs were
418 not to be delegated for the time being. Uncalculated-risk TLDs were not to be delegated until
419 further study was completed, and applicants would also be expected to “provide evidence of the
420 results from the steps taken to mitigate the name collision risks to an acceptable level.”

421 ICANN also posted an announcement that gave an overview of the Interisle report, the mitigation
422 proposal, and other information ICANN was making available related to the topic. [39]

423 3.4.3 Public Comments on ICANN Proposal

424 *Applicability: Duplicate name collisions*

425 There were dozens of responses to the ICANN proposal [38] over the next few months. Some of
426 these also commented on the findings of the Interisle report. The entire archive of approximately
427 80 public comments is available online. [40] There was also a report from ICANN summarizing
428 the public comments. [41]

429 There were public comments from over 15 companies about the ICANN proposal and the Interisle
430 report overstating the risk from new gTLDs, especially those in the uncalculated-risk profile. The
431 public comment summary [41] listed several of these concerns, including the following:

- 432 • “...all applied for new TLDs other than .corp and .home represent a combined 0.016% of
433 the total query rate in the 2012 DITL data provided by Interisle. This figure and the
434 potential reasons that these queries are taking place simply do not warrant mitigation
435 through a 3-6 month delay.”
- 436 • “Merely counting the number of requests for each string is completely insufficient when
437 judging risk. The true origin of the ‘collision’ must be taken into account. The vast
438 majority of requests provided in Table 12 either posed no potential risks or risks that could
439 be handled with simple mitigations.”

- 440 • “Basing risk measurement on total query counts is fundamentally flawed, especially when
441 using data collected after the new TLD applications were posted. The Interisle report
442 makes no mention of investigating the possibility that some of the requests were issued
443 intentionally.”
- 444 • “Risks listed by Interisle or Verisign already exist and many are prevalent in existing
445 gTLDs such as .com. Future studies would gain credibility if the listed risks were compared
446 against the situation in current gTLDs.”

447 There were also numerous comments criticizing the methodology used in the Interisle study and
448 questioning the findings of that study. The comment from Donuts [42] indicated that their own
449 analysis found a lower rate of requests for applied-for gTLDs than Interisle’s analysis did, because
450 Donuts accounted for time to live (TTL) for DNS answers, and that Interisle had admitted they had
451 insufficient time to perform their analysis. Donuts downplayed the risk of name collision, which
452 included stating that, “In order to make a fair comparison of the relative risk regarding collision,
453 it’s critical to point out that Verisign, as manager of the .COM registry, experiences collision at a
454 rate of at least 2,000 names per day for the studied period in 2013, and at least 16,000 names per
455 day for the study period in 2012....” Donuts also provided an explanation for the prevalence
456 of .home requests: 92% of them were from Google Chrome querying for random SLDs within
457 the .home TLD, seeking replies that there was no such domain. Donuts provided its own set of
458 recommendations for addressing name collisions, and stated that “no applied-for TLDs need
459 mitigation, with the possible exception of a very few.”

460 There were dozens of comments questioning the uncalculated-risk profile, with most asserting that
461 only a few TLDs should be high-risk and all others should be considered low-risk and allowed to
462 proceed with applications. One of these was from DigiCert. [43] DigiCert performed its own
463 analysis of the Interisle data combined with “additional data on certificates, SLD information, and
464 total number of domains.” DigiCert looked at potential collisions at all levels (not just the top
465 level). Their conclusion was that six TLDs should be considered high-risk: corp, home, mail, ice,
466 global, and ads. All other applied-for TLDs should be considered low-risk.

467 One of the submitted comments was a study titled “Namespace Expansion” from JAS Global
468 Advisors and simMachines. [44] This study used the same data that was the basis for the Interisle
469 study, and it analyzed it to look for queries for applied-for gTLDs with a focus on the SLD names
470 in the queries and the IP addresses making these queries. The study provided statistics, not
471 conclusions. It was based on the assumption that “there is risk inherent in interacting on the
472 Internet”, so this study was trying to help differentiate unusual risks from typical risks.

473 Another comment in support of treating all but a few applied-for TLDs as low risk came from the
474 New gTLD Applicant Group (NTAG). [45] Although the NTAG agreed that the two high-risk
475 profile strings should not immediately proceed with delegation, the NTAG did not find justification
476 for delaying any others. They stated, “A Verisign analysis using data from January 2006, prior to
477 the launch of several active TLDs, found that .xxx received more queries before delegation than
478 any other new TLD. Despite having more queries than of all of the TLDs currently under
479 consideration in the ‘Uncategorized Risk’ category, .xxx was delegated in 2011. This TLD
480 launched without incident, and no public complaints or technical issues have been identified
481 since.”

482 Verisign Labs submitted a report analyzing the risk for three applied-for TLDs: website, coffee,
483 and club. [46] The website and coffee domains were initially classified as low risk, while the club
484 domain was considered uncalculated risk. Their analysis indicated greater levels of risk for all
485 three domains than originally estimated, and their report criticized the original analysis
486 methodology as being inadequate both in terms of the length of time data was collected (two days)
487 and in the importance given to the number of queries.

488 Verisign Labs also did its own interdisciplinary study on the risk that gTLD delegation could cause
489 to end users. It published this study in late August 2013 and submitted it as a comment. [47]
490 Verisign Labs proposed a methodology for measuring risk for applied-for TLDs. Based on their
491 analysis, they discovered several cases where a particular string or strings meant for internal use
492 was reaching root servers because of proxies, Internal Name certificates, and other reasons, and
493 they believed delegating applied-for gTLDs would put the users in these cases at immediate risk
494 from man-in-the-middle attacks. The study was very cautious about delegating more gTLDs, and it
495 recommended more study and more implementation of existing recommendations for mitigating
496 the risks.

497 A final example of a public comment on the ICANN proposal proposing a different risk analysis
498 methodology was Neustar's report, *A Methodology for Assessing Collision Risk and New gTLDs*.
499 [48] It said that "ICANN's mitigation strategy rests entirely on the possibility of collision, not the
500 consequences." Also, "ICANN already has all the data and research necessary to calculate the risk
501 and develop mitigation strategies that are carefully tailored to the specific risk associated with each
502 TLD." Neustar proposed its own methodology for assessing impact based on "(i) TLD query
503 volume; (ii) query source IP address volume; (iii) queried second-level domain volume; and (iv)
504 volume of SSL certificates." By far the highest scoring TLDs were corp and home, with mail in
505 third and all others far behind mail. Accordingly, they proposed having those three TLDs as high-
506 risk and all others as low-risk.

507 3.4.4 ICANN Proposal on New gTLD Collision Occurrence Management

508 *Applicability: Duplicate name collisions*

509 ICANN released a second proposal on October 4, 2013. [49] The first proposal was on collision
510 risk mitigation; this subsequent proposal was on managing collisions that occurred. The proposal
511 stated that ICANN would have a name collision occurrence management framework developed.
512 The framework would be used for each applied-for TLD to assess the likelihood of collisions and
513 their potential impact, and to help create a name collision occurrence assessment for the TLD.
514 Each assessment would include suggested mitigations for SLDs within that TLD, such as blocking
515 particular SLDs (temporarily or indefinitely). The proposal also stated ICANN would perform
516 outreach to raise public awareness of name collisions and to educate network operators and
517 software and equipment manufacturers about name collisions and how they can mitigate them.

518 Appendix I of the October proposal provided ICANN's response to the public comments on the
519 August proposal. Responses of particular interest are as follows:

- 520 • "ICANN agrees that other parameters, besides request frequency, should be considered in
521 assessing the threat, particularly the potential for harm caused by name collisions. ICANN

522 will adopt the advice regarding the use of the other proposed parameters when developing a
523 collision occurrence management framework.”

524 • “ICANN will adopt the idea by NTAG and others to block Second Level Domain names
525 (SLDs) that are being queried.”

526 • “ICANN will enable an affected party to report and request the suspension of a domain
527 name that by virtue of name collisions is causing severe harm.”

528 • “DotGreen requested that strings in the uncalculated-risk category be allowed to proceed to
529 contracting. Similarly, other commenters complained about ICANN not allowing these
530 strings to proceed to contracting when the public comment period for the proposal is still
531 open. ICANN understands the interest of applicants to see their strings move as fast as
532 possible through the new gTLD process and will remove that restriction. The adoption of
533 the blocking of SLDs makes this restriction unnecessary.”

534 On October 7, 2013, the New gTLD Program Committee (NGPC) passed a resolution to have the
535 proposal implemented. [50]

536 3.4.5 DNS-OARC Workshop Session on High-Risk Strings Collisions

537 *Applicability: Duplicate name collisions*

538 The Domain Name System Operations Analysis and Research Center (DNS-OARC) held a
539 workshop session on high-risk strings collisions on October 5, 2013. There were four presentations
540 in the session:

541 • Jim Reid from Interisle [51] spoke on the data analysis Interisle performed for their study
542 for ICANN. He explained many of the logistical issues they experienced while attempting
543 to analyze terabytes of data in a matter of weeks.

544 • Roy Hooper from Demand Media [52] discussed numerous challenges encountered when
545 attempting to perform additional analysis of the same data Interisle had analyzed.

546 • Andrew Simpson from Verisign [53] presented the results of research he and his colleagues
547 had performed on queries for applied-for gTLDs to see if there was any significance to
548 their origins (e.g., a disproportionate number coming from a particular country). This could
549 help identify countries at greater risk from a particular gTLD being delegated.

550 • Andrew Sullivan from Dyn [54] spoke about an Internet-Draft he was co-authoring with
551 Olaf Kolkman from NLnet Labs and Warren Kumari from Google. [55] The idea was that
552 “test delegations be used to enable empirical research on the extent of the possible
553 disruption prior to actual allocation and delegation of any label in the root zone.” The
554 Internet-Draft proposed a methodology for doing the test delegations and collecting the
555 necessary data. (Note that the Internet-Draft was updated twice in the following few
556 months, but the authors eventually let it expire.)

557 3.4.6 SSAC Advisory SAC 062 on Mitigating Name Collision Risk

558 *Applicability: Duplicate name collisions*

559 On November 7, 2013, the SSAC released its SAC 062 advisory on mitigating name collision risk.
560 [56] This advisory was based on the Interisle study [36], the August ICANN proposal [38], the
561 October ICANN proposal [49], and SSAC’s own analysis of the subject. SAC 062 stated that the
562 SSAC generally agreed with the October proposal, and it made a few additional recommendations:

- 563 • The first was for ICANN to work with the Internet Architecture Board (IAB), IETF, and
564 potentially others to determine which domain names should be reserved, both TLDs and
565 lower-level names.
- 566 • The second involved trial delegation. The concept was to delegate a TLD with a short time
567 to live, then collect data on queries for that TLD. The trial could cause name collisions for
568 a short time, which might be temporarily disruptive but would also allow issues to be
569 identified and addressed before permanent delegation occurred.
- 570 • The third was having policies and processes in place to roll back delegation of a TLD, if
571 the TLD was causing security or stability issues that couldn’t be immediately mitigated
572 through other means.

573 3.4.7 SLD Blocking List Effectiveness

574 *Applicability: Duplicate name collisions*

575 As mentioned in Section 3.4.4, one of the outcomes of the public comments on ICANN’s August
576 proposal [38] was ICANN’s decision to adopt SLD blocking, as announced in ICANN’s October 4
577 proposal. [49]

578 Verisign Labs published a preliminary analysis of SLD blocking list effectiveness on November 5,
579 2013. [57] This was based on a group of gTLDs that had SLD blocking lists released on October
580 29. The gTLDs were: camera, clothing, equipment, guru, holdings, lighting, singles, ventures, and
581 voyage. The initial results of the analysis indicated that the SLD blocking lists were “ineffective.”
582 It also raised questions about how the SLDs on the blocking lists were selected.

583 Another Verisign Labs report was released on November 15 on SLD blocking effectiveness. [58]
584 This report continued the analysis in the November 5 report, expanding it to include newly
585 published SLD blocking lists for 16 more gTLDs released on November 6. The report asserted that
586 “a fundamental reason that SLD blocking based on DITL datasets is ineffective is that the set of
587 SLDs in queries evolves. New SLDs appear in queries all the time.” Verisign Labs’ analysis
588 indicated that “the number of SLDs observed in the DITL data for the first time each year is on a
589 significant upward trend.”

590 On November 17, 2013, ICANN posted an announcement about SLD blocking. [59] The
591 announcement indicated that for some gTLDs, SLD blocking was not sufficiently effective:

592 “The gTLDs that were considered ineligible were those for which the growth of the number
593 of SLDs queried year over year significantly exceeded the average growth rate for all
594 applied for gTLDs in at least two of the DITL years (2006-2012), and for which one of the
595 years in which this was observed was the most recent year, 2012. The analysis of this data
596 showed that for some strings, the variance of SLDs queried varied so significantly from
597 year to year that the mechanism of blocking SLDs might not be an effective way of
598 addressing the name collision issue.”

599 The announcement then listed 25 applied-for gTLDs that would not be delegated for these reasons,
600 plus a mention that home and corp would also not be delegated.

601 On March 8-10, 2014, the Workshop and Prize on Root Causes and Mitigation of Name Collisions
602 (WPNC) was held. [60] All the talks at this workshop were related in some way to name collisions.
603 The talks listed below pertained to SLD blocking lists. Note that this report also summarizes talks
604 with other material in their subject areas, such as the creation of the Name Collision Occurrence
605 Management Framework (see Section 3.6). See RFC 8023, *Report from the Workshop and Prize
606 on Root Causes and Mitigation of Name Collisions* for a summary of the workshop. [61]

- 607 • Verisign Labs personnel gave a talk and released a paper on using block lists to prevent
608 collisions. [62] Their work was almost entirely focused on data analysis to attempt to
609 quantify the effectiveness of SLD blocking. Their conclusion was that SLD queries change
610 so often that SLD blocking would not be effective in mitigating name collisions.
- 611 • Paul Hoffman from the VPN Consortium [63] commented during his talk on the
612 uncertainty about the effectiveness of SLD blocking.
- 613 • Another came from RTFM. [64] This work was also focused on data analysis, but with the
614 purpose of determining if SLD blocking would cause harm to “naïve DNS clients” like
615 “stub resolvers and forwarding-only devices. If these query the root servers, they can
616 receive referral responses that they are unable to process and that would result in undefined
617 behavior.” The conclusion of the work was that these types of clients were unlikely to be
618 harmed by SLD blocking.

619 3.5 Research on Name Collision Causes: 2013 – 2016

620 3.5.1 Search List Processing and FQDN Usage

621 *Applicability: Duplicate name collisions, shortened name collisions, search list name collisions*

622 The Réseaux IP Européens Network Coordination Centre (RIPE NCC) posted an article on
623 November 12, 2013 about new gTLDs and search list processing [65]. Search list processing is
624 when a person specifies only a portion of a domain name instead of a fully qualified domain name
625 (FQDN), and the search list adds to that domain name and attempts to resolve it as an FQDN,
626 trying again with another addition if the first one fails, and so on. Although search list processing
627 was not a novel topic at the time—it had been discussed in numerous documents before—this was
628 when it got greater attention on its own. The RIPE NCC article argued that the risks from name
629 collisions caused by searches for internal domains leaking out to the Internet were much lower than

630 had been claimed. However, it also appealed for operating systems and web browsers to handle
631 searches for names in standard ways so that the leaking of queries for internal domains would stop.

632 ICANN released a blog posting on December 6, 2013 on managing name collision occurrences,
633 focusing on issues with search list processing [66]. The blog posting referenced version 1.0 of the
634 *Guide to Name Collision Identification and Mitigation for IT Professionals*, which ICANN had
635 released the day before the blog posting. [67] Both the blog posting and the guide had the same
636 motivation: to stop the leaking of queries for internal domains, which would prevent name
637 collisions from occurring, by encouraging organizations to migrate from internal-only, shortened
638 domain names to public FQDNs.

639 In February 2014, the SSAC released its SAC 064 advisory on DNS search list processing. [68]
640 This advisory discussed the inconsistencies in how search list processing was performed by
641 operating systems, web browsers, email clients, and other software. Although there were RFCs
642 (1123 [69], 1535 [19], and 1536 [70]) with search list guidelines, these RFCs were informational
643 and had not been widely adopted, and there were also concerns that the RFCs were not as clear and
644 specific as they needed to be. The SAC 064 advisory documented some of the differences in search
645 list processing among commonly used client operating systems. Most importantly, Section 4 of
646 SAC 064 proposed improvements to search list processing that would reduce the likelihood of
647 name collisions.

648 In March 2014, there were several talks at the Name Collision Workshop pertaining to search list
649 processing:

- 650 • Warren Kumari [71] spoke about the need to educate developers on not using shortened
651 names instead of FQDNs, and on the value of reserving .alt as a local-only domain name.
- 652 • Paul Hoffman from the VPN Consortium [63] talked about what organizations could do to
653 help mitigate name collisions. He encouraged organizations to stop using shortened names
654 and to use FQDNs instead. He also mentioned that “determining the so-called ‘potential for
655 collisions’ for a private namespace is nearly impossible.”
- 656 • Colin Strutt from Interisle [72] presented on the corp.com domain. This domain had been
657 registered in 1994 by Mikey O’Connor, but no SLDs within it were registered, so any
658 queries for this domain were likely to be queries for organizations’ internal .corp names
659 that, because of search list processing, had leaked onto the Internet. Mikey O’Connor and
660 NetChoice sponsored a small study of this query data by Interisle. The corp.com domain
661 was receiving approximately 2 million queries a day from a wide variety of IP addresses,
662 domains, and countries. The study also attempted to contact some of the organizations and
663 internet service providers where the queries were coming from in order to get them to stop
664 issuing these queries, but this met with little success.
- 665 • Casey Deccio from Verisign Labs [73] spoke on quantifying risk from name collisions by
666 creating a name collision model that takes search lists into account.
- 667 • Andrew Simpson from Verisign [74] spoke about detecting search lists. He experimented
668 with real-world systems to observe their name resolution behavior and compare this with

669 DITL data on queries for nonexistent domains. He also gave a second talk [75] on DNS
670 query analysis techniques that might be of use in name collision discovery.

671 3.5.2 Causes of Internal Domain Leakage

672 *Applicability: Duplicate name collisions*

673 Verisign Labs published a paper in November 2014 on leakage of the .onion domain, which is a
674 non-delegated TLD meant for Tor usage only. [76] The nature of Tor is such that .onion queries
675 should not leak onto the Internet, but analysis of root server data on queries indicated leakage was
676 definitely happening. The reasons for the leakage were unknown, but there were several
677 possibilities cited:

- 678 • User error
- 679 • Client software misconfiguration
- 680 • Browser prefetching
- 681 • Third-party applications or plug-ins
- 682 • Search list processing
- 683 • Web crawlers
- 684 • Malware

685 Another paper on .onion domain leakage was submitted in March 2016 and published in October
686 2017. [77] Written by university researchers, one of whom was also a co-author on the Verisign
687 Labs 2014 paper [76], this paper used data from DITL and other sources, and it provided a more
688 rigorous analysis of the data for .onion domain leakage than the 2014 paper. It also looked at some
689 potential causes of the leakage in greater detail:

- 690 • User error and misconceptions: the researchers surveyed graduate students in a computer
691 security class about the .onion domain, and only half the students who considered
692 themselves “very knowledgeable” about Tor knew the special function of the .onion
693 domain.
- 694 • Browser prefetching and web crawlers: some web browsers would try resolving the links
695 on a webpage in advance of anyone clicking on those links, so that could cause .onion links
696 to try to be resolved when Tor isn’t running. To evaluate this, the researchers did a website
697 crawl and looked for strings ending in “.onion”, and they found that 17% of the instances
698 of .onion queries seen in the DITL data corresponded to strings seen during their crawl.
- 699 • Malware: the researchers looked for .onion queries for SLDs known to host malware, but
700 there was not a clear correlation.

701 RFC 7686, *The “.onion” Special-Use Domain Name*, was published in October 2015. [78] It
702 explained the unique role of the onion domain name, and it defined how queries for onion names
703 should be resolved, which would prevent further leakage.

704 In May 2017, there was a presentation about the Operational Research Data from Internet
705 Namespace Logs (ORDINAL) dataset. [79] This presentation gave numerous examples of
706 protocols and applications that misused DNS by using it for authentication, not identification—
707 essentially, they trusted whatever result they got from DNS as being sufficient confirmation of the
708 legitimacy of the destination, instead of subsequently performing authentication with the
709 destination to verify it.

710 3.5.3 Detection of Leaking Clients

711 *Applicability: Duplicate name collisions*

712 As discussed in Section 3.6.3, the JAS Global Advisors report mentioned observations that
713 indicated some attackers were purposely choosing domains with collisions so they could take
714 advantage of those collisions. There have been works published since that timeframe regarding
715 how attackers could utilize name collisions by detecting internal queries leaking from clients onto
716 the Internet and registering the searched-for names.

717 Two of these papers focused on vulnerabilities in the Web Proxy Auto-Discovery (WPAD)
718 protocol. Both papers were published in May 2016; one was authored by Verisign Labs personnel
719 only [80] and the other was co-authored by Verisign Labs personnel and University of Michigan
720 researchers [81]. Also published in May 2016 was an alert from the National Cyber Awareness
721 System on the WPAD name collision vulnerability. [82]

722 The WPAD issue involved internal-only domain names not being found when laptops using
723 WPAD were used on external networks, so the laptops were sending DNS queries to the Internet.
724 Attackers aware of this behavior could register the domains the laptops were erroneously trying to
725 reach and perform man-in-the-middle attacks on the laptops. The problem was first found on
726 Microsoft laptops, but it was soon confirmed that Apple and Linux laptops had the same problem.
727 The recommended mitigation was to disable WPAD if not needed, otherwise to hard-code proxy
728 addresses instead of using WPAD to acquire them. The [80] and [81] papers both highlighted that
729 transient devices like laptops might encounter name collisions more frequently than other devices
730 because transient devices go from one network to another.

731 The same researchers who authored [81] plus two additional University of Michigan researchers
732 wrote a conference paper published in November 2017 on client-side name collision
733 vulnerabilities. [83] This paper covered a broader range of vulnerabilities than just WPAD. The
734 authors created a general name collision threat model for clients querying internal-only names that
735 were being leaked onto public networks. The authors then analyzed DITL query data for 2011
736 through 2016 to find evidence of internal services (WPAD and many others) being exposed. They
737 found that 115 registered services and an undetermined (but large) number of unregistered services
738 were exposed, and they chose 48 of the most commonly seen services for further analysis. The
739 researchers then looked for vulnerabilities in those services and determined that 93.8% of them
740 were vulnerable, for reasons such as lack of server authentication or accepting a different server
741 certificate than the one expected without notifying the user. Further discussion of the contents of

742 the paper is outside the scope of this document, because the service vulnerabilities were not name
743 collision related; the relevance of the paper is that leakage of internal names associated with
744 services puts those services at high risk of exploitation if an attacker registers a particular name
745 collision domain.

746 Four of the co-authors of the papers mentioned above filed a patent application on March 24, 2017.
747 [84] This patent proposed ways to detect internal names leaking, especially for the WPAD
748 vulnerability, and remediate the problems causing the leaks. As of this writing, the patent
749 application is still pending.

750 3.6 Name Collision Occurrence Management Framework: 2014 – 751 2015

752 3.6.1 JAS Global Advisors Phase One Report Draft

753 *Applicability: Duplicate name collisions, shortened name collisions, search list name collisions,*
754 *re-registered name collisions*

755 The Name Collision Occurrence Management Framework was to be developed so it could be
756 applied to any newly requested gTLD to assess risk and identify mitigations before the gTLD was
757 delegated. JAS Global Advisors was selected by ICANN in November 2013 to create the Name
758 Collision Occurrence Management Framework. The phase one report draft for their work was
759 released for public comment on February 26, 2014. [85] JAS Global Advisors also presented on
760 this work at the Name Collision Workshop in March 2014 [86] and at the March ICANN meeting
761 in Singapore [87] [88]. Among their findings at that time were the following:

- 762 • DNS name collisions happened frequently and have happened before the delegation of each
763 new TLD since at least 2007. Issues caused by collisions date back to approximately 1987.
764 There was “no evidence to suggest that the security and stability of the global Internet DNS
765 itself is at risk.”
- 766 • There were several causes of these collisions, including shortened name usage, search list
767 processing differences, misunderstandings about DNS, expired registrations, human error,
768 and intentional acquisition of colliding names.
- 769 • Other types of namespaces have had collisions, with numerous examples from phone
770 numbers and mailing addresses. These were handled through advance notification of the
771 transitions, and in having a grace period of some sort when feasible.
- 772 • The corp, home, and mail TLDs should not be delegated because of their existing
773 widespread internal use by organizations. RFC 6762 [33] “suggests that .corp and .home
774 are safe for use on internal networks.” Also, “.mail has been hardcoded into a number of
775 installations...and has a large global ‘installed base’ that is likely to have significant inertia
776 comparable to .corp and .home.”

- 777 • There should be processes in place to act if a TLD delegation presents “clear and present
778 danger to human life.”

779 The phase one report draft compared the disruption caused by a name collision to the disruption
780 caused by failing to renew a domain. “Like unintended expirations, DNS namespace collisions can
781 be viewed as a notification problem. The system administrator utilizing the colliding namespace
782 (either knowingly or unknowingly) must be notified and take action to preserve the security and
783 stability of their systems.”

784 The report discussed at length how new gTLDs could be delegated using a method called
785 “controlled interruption.” Instead of simply delegating a new gTLD and allowing traffic that
786 previously would have gone elsewhere to inadvertently reach the newly delegated gTLD instead,
787 queries for the gTLD would receive a response that directs them to a different address that
788 essentially indicates an error has occurred. The report discussed two options for this “different
789 address”—a honeypot or a loopback address—and recommended loopback addresses because this
790 “prevents traffic from leaving the requestor’s network and blocks a malicious actor’s ability to
791 intercede.”

792 JAS Global Advisors recommended that a standard loopback address should be used for all
793 controlled interruptions: 127.0.53.53 (with 53 chosen because it is the port number associated with
794 DNS). Having the same unusual IP address returned in all controlled interruption replies should
795 help system administrators to identify the problem. They recommended having a 120-day
796 controlled interruption period:

797 “Registries that have not yet been delegated to the root zone shall implement controlled
798 interruption via wildcard records; registries that have elected the ‘alternative path to
799 delegation’ shall implement controlled interruption by adding appropriate resource records
800 for the labels appearing in their respective block lists. Following the 120-day controlled
801 interruption period, registries will not be subject to further collision-related restrictions.
802 ...we believe the 120-day controlled interruption period offers a conservative buffer
803 between potential legacy usage of a TLD and the new usage.”

804 In other words, for a new TLD on “alternative path,” reply with the loopback address for SLDs on
805 the blocking list only. For all other new TLDs, reply with the loopback address for every SLD.
806 That is the equivalent of a wildcard, and the draft report recommended that wildcard records be
807 permitted for the purpose of controlled interruption for TLDs not on “alternative path,” since the
808 TLD would not have any registrant data during that period.

809 Section 2.1.2 of the report draft briefly discussed a trial JAS Global Advisors had performed of
810 controlled interruption, and “despite publishing phone numbers and email addresses via http and
811 Whois, in the event the controlled interruption caused harm, not a single call or email was
812 received.”

813 The report draft also mentioned alternatives to controlled interruption, “including several honeypot
814 approaches, use of DNAME, and various 2LD string-by-string and TLD-by-TLD approaches.
815 While we eventually concluded that controlled interruption approach offers the most value and
816 presents the least risk, discussion of alternatives is worthwhile.” See Section 3.6.3 for further

817 discussion of the alternatives to controlled interruption and other contents of the phase one report
818 draft, as presented in the final version of the report.

819 3.6.2 Public Comments on Phase One Report Draft

820 *Applicability: Duplicate name collisions, shortened name collisions, search list name collisions,*
821 *re-registered name collisions*

822 Over 25 public comments were submitted on the phase one report draft. The complete set of public
823 comments is available at [89]. An ICANN report summarizing these comments was also posted.
824 [90]

825 Most of the comments received largely agreed with the report draft and its recommendations,
826 except for the topic of 120-day controlled interruption. Many felt it should be shorter, perhaps 38,
827 45, or 60 days, although one felt 120 days was too short to allow organizations to remediate
828 problems.

829 There was some disagreement about whether using the 127.0.53.53 loopback address or a honeypot
830 would be better. The majority felt the address would be better than a honeypot. An alternate
831 solution proposed was to use a public IP address and website that could provide information on the
832 nature of the problem to end users.

833 On the topic of not delegating the corp, home, and mail TLDs, there was no consensus. Some felt
834 they should all be permanently reserved, while others thought more discussion and evaluation was
835 needed, and yet others disagreed with permanently reserving any of them.

836 Several commenters mentioned that collisions happened all the time in .com and nothing was done
837 about that, but great scrutiny was being given to a much smaller problem with collisions involving
838 new gTLDs.

839 Verisign released preliminary public comments on the phase one report draft in late February 2014
840 and updated those comments on March 31. [91] [92] Verisign released an additional set of
841 comments on the phase one report draft on April 21, 2014, in part to clarify the preliminary
842 comments. [93] [94] Topics of the Verisign comments included the following:

- 843 • Verisign did not find the expected elements of the Framework in the draft report; instead,
844 the draft report focused on using controlled interruption when delegating new gTLDs, as if
845 an undefined Framework were already being applied.
- 846 • Controlled interruption had not been tested. There were two scenarios where it might not
847 succeed in notifying organizations of an impending change in name resolution. The first
848 scenario was an “alternative path” delegation where someone queries for an SLD that isn’t
849 on the blocking list. The second scenario involved use of WPAD. Implementations of
850 WPAD vary, and a non-recommended WPAD implementation might receive a controlled
851 interruption reply, ignore it, and continue with its search list processing.
- 852 • “There is therefore a reasonable case to be made, at least for some new gTLDs and SLDs,
853 that the controlled interruption should be done more selectively.” Verisign termed this

854 “selective interruption” and said it “requires careful qualitative analysis” to determine
855 when it can be used instead of the broader controlled interruption. There were also some
856 drawbacks to the selective interruption approach. Another alternative approach from other
857 public commenters (United TLD, the NTAG, and the China Internet Network Information
858 Center [CNNIC]) was to allow all SLDs in a new gTLD to be interrupted except for those
859 SLDs that have already been delegated.

- 860 • Verisign was opposed to using external honeypots because of the likelihood of sensitive
861 data inadvertently being leaked over unsecured networks (e.g., the Internet) to the
862 honeypot.
- 863 • Expired registration name collisions (the term used in this NCAP Phase 1 report, not the
864 JAS Global Advisors report or Verisign comments) were not actually a form of name
865 collision.
- 866 • “If controlled interruption is adopted, the only way to get a better understanding of the
867 appropriate period is by qualitative analysis of the effectiveness of the mitigation measure
868 in practice.”

869 3.6.3 JAS Global Advisors Final Phase One Report

870 *Applicability: Duplicate name collisions, shortened name collisions, search list name collisions,*
871 *re-registered name collisions*

872 The final phase one report was released on June 4, 2014 [95]. Its assertions of no significant
873 problems caused by name collisions are noteworthy:

874 “We do not find that the addition of new Top Level Domains (TLDs) fundamentally or
875 significantly increases or changes the risks associated with DNS namespace collisions.”

876 “As we write this update, 275 New gTLDs have been delegated and over 835,000 second
877 level registrations have been added.” ... “Neither JAS nor ICANN is aware of even a single
878 instance of a problematic collision. Of course, this fact certainly doesn’t ‘prove the
879 negative’ but it also can’t be ignored at this point.”

880 Significant changes from the draft report included the following:

- 881 • The length of the controlled interruption period was dropped from 120 days to 90 days.
- 882 • A discussion of controlled interruption for IPv6 addresses was added. The assertion was
883 that there was not a significant problem with IPv6-only hosts having name collisions, and
884 there was not an IPv6 counterpart to 127.0.53.53 that could be used for controlled
885 interruption.
- 886 • A discussion of having staggered controlled interruption periods instead of continuous
887 controlled interruptions was added. This idea was suggested by Google in the public
888 comments. The report did not favor staggered controlled interruptions because it would
889 cause “intermittent failures, which are maddening and hard to diagnose from a system

890 administrator perspective. Moreover, we found that systems configured in a way to create
891 collision-related effects in the existing DNS namespaces routinely experience and tolerate
892 intermittent failures....”

893 Section 3.1 of the final report discussed alternatives to using controlled interruption with loopback
894 addresses:

895 • **String-by-string approaches (TLD and SLD)** (as detailed in [96]): “JAS’ assessment is,
896 with the exception of .corp, .home, and .mail, that the risk of a collision in the newly
897 applied-for TLD namespaces causing more than a highly localized disruption is low after
898 the recommended mitigation technique is applied. String-by-string and TLD-by-TLD
899 approaches add significant complexity and potential for unintended consequences while
900 adding little if any security value. Not a good tradeoff. As such, we recommend an
901 approach that addresses the root causes and does not delineate between specific strings
902 unnecessarily.”

903 • **Honeypots:** Honeypots would be more useful from a notification standpoint than loopback
904 addresses, but honeypots have several drawbacks. In addition to the potential exposure of
905 sensitive data across networks and to the honeypot itself, honeypot use would also make it
906 possible for someone to “game” things by sending queries to make it look like the new
907 gTLD’s delegation should be delayed. Also, the data on the honeypot could be subject to
908 privacy laws and regulations from numerous jurisdictions.

909 • **Use of DNAME records:** While wildcard DNAME records could point to something like
910 “you-need-to-change-your-dns-config-see-collisions-dot-icann-dot-org.”, DNAME has
911 only been well supported since around the year 2000, whereas loopback addresses have
912 been well supported since around 1989. Also, “DNAME-based approaches don’t
913 necessarily interrupt, negating the whole purpose of controlled interruption. The DNAME
914 redirect to return NXDOMAIN means folks can continue on as they’re currently doing.
915 They won’t notice anything so they won’t fix it, defeating the purpose of the interruption.”

916 The final report also added a new Section 3.3, “Collisions in Existing DNS Namespace”. To
917 measure collisions within existing TLDs, JAS Global Advisors registered some SLDs and found
918 that “these registrations immediately generated a surprising amount of traffic.” They noted that
919 they used tools meant to aid people with “domain drop catching” and “squatting”. *Domain drop*
920 *catching* is, in the parlance of the NCAP Phase 1 study, a re-registered name collision that is
921 performed immediately after an expired domain becomes available, typically for malicious
922 purposes. *Squatting*, which refers to someone registering a TLD to prevent someone else from
923 registering it, is not a form of name collision and is outside the scope of the NCAP study. The
924 volume of queries received immediately after registration indicated the tools may have had access
925 to feeds with data on queries to nonexistent domains, and took advantage of that data to
926 intentionally cause name collisions and benefit from them.

927 3.6.4 SSAC Response to the Final Phase One Report

928 *Applicability: Duplicate name collisions, shortened name collisions, search list name collisions,*
929 *re-registered name collisions*

930 The SSAC released SAC 066, their comment on the final phase one report, on June 6, 2014. [97]
931 The Recommendations from SAC 066 that differed from the recommendations in the final phase
932 one report and are applicable to the NCAP Phase 1 study are as follows (with recommendation text
933 bolded and supporting text not bolded):

- 934 • **“Instead of a single controlled interruption period, ICANN should introduce rolling**
935 **interruption periods, broken by periods of normal operation, to allow affected end-**
936 **user systems to continue to function during the 120-day test period with less risk of**
937 **catastrophic business impact.** Controlled interruption periods starting at 24 hours and
938 eventually lengthening to 30 days would be separated by periods of at least 3 days, to allow
939 users or system administrators to identify or develop and put in place solutions or
940 workarounds.”
- 941 • **“ICANN should perform an evaluation of potential notification approaches against at**
942 **least the requirements provided by the SSAC prior to implementing any notification**
943 **approach.”** This was due to SSAC’s concerns that use of the 127.0.53.53 loopback address
944 would only effectively notify some system administrators, but not typical end users. The
945 SSAC felt that there was “a wealth of operational expertise” in handling sensitive data sent
946 to honeypots and that privacy concerns should not preclude the use of honeypots instead of
947 loopback addresses.
- 948 • **“ICANN should implement a notification approach that accommodates IPv6-only**
949 **hosts as well as IPv4-only or dual-stack hosts.”**
- 950 • **“ICANN should consider not taking any actions solely based on the JAS Phase One**
951 **Report.”**
- 952 • **“ICANN should seek to provide stronger justification for extrapolating findings based**
953 **on one kind of measurement or data gathering to other situations.”** This was in
954 response to this assumption from the Phase One report: “The modalities, risks, and
955 etiologies of the inevitable DNS namespace collisions in the new TLD namespaces will
956 resemble the collisions that already occur routinely in other parts of the DNS.” The SSAC
957 questioned whether this assumption was valid.

958 Appendix A of SAC 066 discussed four alternative notification approaches:

- 959 1. “Do nothing. Users...will experience failures or misconnections and come to realize their
960 configurations are problematic only after the new gTLD and domains within that gTLD are
961 delegated and elicit operational impacts to their systems.” This approach was deemed
962 unacceptable.
- 963 2. “Perform qualitative analysis of query sources as measured at root and TLD servers and
964 provide proactive user notification.” This approach would require the root servers to have
965 measurement capabilities they did not yet possess, so it was not an option in the short term.
- 966 3. “Implement structured, short-term test periods (‘controlled interruption’), in which end
967 users utilizing a proposed gTLD will experience a failure, and then be given time (after
968 each short-term test period) for planning and effectuating remediation efforts specific to

969 their environment. This approach triggers the errors in a more controlled environment, and
970 can be used as an early warning system to notify potentially impacted parties. There are two
971 variations to notification in this approach:”

972 a. Loopback address usage (127.0.53.53)

973 b. Redirection to honeypot

974 3.6.5 Approval of the Name Collision Occurrence Management Framework

975 *Applicability: Duplicate name collisions, shortened name collisions, search list name collisions,*
976 *re-registered name collisions*

977 ICANN approved the Name Collision Occurrence Management Framework on July 30, 2014. [98]
978 ICANN stated in the introduction to the Framework that they took into consideration the JAS
979 Global Advisors final Phase 1 report [95], the public comments submitted on the draft Phase 1
980 report [90], and the SAC 062 [56] and SAC 066 [97] documents.

981 The Framework required registry operators to do continuous controlled interruption for each new
982 gTLD for a minimum of 90 days. It was stated that “there is already a mechanism in place (name
983 collision reporting) for affected parties to find temporary relief from name collision harm, if
984 needed, making the intermittent approach an unnecessary burden” for registries. There was not yet
985 an IPv6 counterpart to the IPv4 loopback address, but ICANN was to collaborate with other
986 organizations in finding a suitable mechanism for IPv6-only hosts.

987 The Framework also stated that ICANN would defer the corp, home, and mail TLDs indefinitely
988 and would work with other organizations to determine how to handle them long-term.

989 A few days after approval of the Framework, ICANN released version 1.1 of the *Guide to Name*
990 *Collision Identification and Mitigation for IT Professionals*. [99] This new version included the
991 requirements from the Framework, such as the 90-day continuous controlled interruption period. A
992 few days after that, ICANN released a Name Collision Occurrence Assessment document for
993 gTLD applicants and registry operators. [100] This included details of implementing continuous
994 controlled interruption and responding to requests from ICANN regarding name collision report
995 handling. In November 2014, ICANN released Addendum to Name Collision Occurrence
996 Assessment, which pertained to trademark claims. [101] ICANN also released a Frequently Asked
997 Questions (FAQ) on the Framework for registries [102] and a FAQ on name collisions for IT
998 professionals [103].

999 3.6.6 Controlled Interruption for New ccTLDs

1000 *Applicability: Duplicate name collisions, shortened name collisions, search list name collisions,*
1001 *re-registered name collisions*

1002 In February 2014, while the Name Collision Occurrence Management Framework [98] was under
1003 development, ICANN released a name collision briefing document that summarized ICANN’s
1004 ongoing efforts to address name collisions for new gTLDs. [104] Section 4 of the briefing

1005 document pointed out the potential relevance of the gTLD name collision work for ccTLDs: “The
1006 issue of name collision is not unique of new gTLDs and could present in new ccTLDs too, both
1007 ASCII and IDN. ICANN is requesting the ccNSO to review the name collision issue and its
1008 implications for new ccTLDs.” The briefing document also stated the following:

1009 “Until advice is received from the ccNSO, ICANN plans to send each new ccTLD manager
1010 the same kind of interim report that new gTLDs received for the alternate path to
1011 delegation. It will remain the responsibility of the local Internet Community and the ccTLD
1012 manager to either: 1) proceed to delegation while temporarily blocking the SLDs identified
1013 in the report; 2) temporarily defer delegation until receipt of their full collision occurrence
1014 assessment and implementation of the measures described; or 3) some other course of
1015 action determined by the local Internet Community and the ccTLD manager.”

1016 Shortly after the Name Collision Occurrence Management Framework [98] was approved for new
1017 gTLDs, ICANN recommended it also be used for each new ccTLD. This included having a
1018 continuous controlled interruption period of at least 90 days. The same resources ICANN had
1019 already made available for name collisions for new gTLDs (as described in Section 3.6.5) were
1020 also relevant for new ccTLDs, and people were pointed to those resources for more information.
1021 [105]

1022 3.6.7 JAS Global Advisors Phase Two Report

1023 *Applicability: Duplicate name collisions, shortened name collisions, search list name collisions,*
1024 *re-registered name collisions*

1025 Release of JAS Global Advisors’ phase two report was delayed because it would have disclosed a
1026 vendor security vulnerability. The final report wasn’t publicly released until October 2015 [106].
1027 The first few sections were duplicates of the JAS Global Advisors final phase one report [95], with
1028 a few notable changes and additions that reflected events occurring during the delay:

- 1029 • “...several vendors have...included detection and messaging around the 127.0.53.53
1030 response. For example, recent builds of Google’s Chrome browser now include the new
1031 error ‘ERR_ICANN_NAME_COLLISION’ which provides specific and richer error
1032 messaging to the user over a general connection timeout.”
- 1033 • There was a new Section 3.1.5, “Effectiveness of Controlled Interruption” that discussed
1034 gTLD delegations and real-world name collisions.
 - 1035 ○ There had been more than 650 gTLD delegations, and ICANN had “received fewer
1036 than 30 reports of disruptive collisions since the first delegation in October of 2013.
1037 None of these reports have reached the threshold of presenting a danger to human
1038 life.”
 - 1039 ○ “As expected, controlled interruption caused some instances of limited operational
1040 issues as collision circumstances were encountered with new gTLD delegations.
1041 While some system administrators expressed frustration at the difficulties, overall it
1042 appears that controlled interruption in many cases is having the hoped-for outcome.

1043 ... JAS would characterize the overall response as ‘annoyed but understanding and
1044 generally positive.’”

1045 ○ “JAS also is aware of specific examples where controlled interruption, for whatever
1046 reason, did not cause underlying DNS issues to be remedied.” In regard to one
1047 example: “JAS suspects that in this specific instance, controlled interruption was
1048 probably not disruptive enough to get the attention of operators; or if it did get the
1049 attention of operators, the issue was not viewed as important enough to cause
1050 action. Based on JAS’ knowledge of the specific circumstances surrounding this
1051 operator, it is unlikely that a longer controlled interruption period or an entirely
1052 different approach to controlled interruption would have made a difference.”

1053 • JAS Global Advisors tested HTTP honeypots in SLDs known to have high volumes of
1054 collisions; reaching a honeypot would return a web page with contact information for JAS
1055 and a request to contact JAS. They received no replies. “Reviewing our HTTP logs, less
1056 than 8% of DNS resolutions ultimately led to the retrieval of one of our HTTP honeypot
1057 pages. Reviewing the HTTP logs further, less than 12% of those 8% reported an HTTP
1058 user-agent that could be considered a user-facing application (i.e. a Browser).”

1059 The substantive new material in the phase two report started in Section 4.1, which said that JAS
1060 analysis showed many of the queries to nonexistent domains were generated by malware. It also
1061 mentioned the queries generated by Google Chrome, which queried for random SLDs within
1062 the .home TLD to try to get replies that the domains did not exist. Taken together, queries to
1063 nonexistent domains automatically generated by malware and Google Chrome represented “nearly
1064 80% of the random and pseudo-random labels we detected in DITL datasets and in excess of 41%
1065 of the total NXDOMAIN traffic described in the DITL datasets. This is consistent with the
1066 observation that the ‘Alternate Path to Delegation’ Second Level Domain (SLD) Collision Block
1067 Lists published by ICANN are comprised largely of these seemingly random, pseudo-random,
1068 machine-generated or otherwise linguistically nonsensical labels.”

1069 Section 5 of the phase two report elaborated on the material briefly discussed in Section 3.3 of the
1070 phase one report, where JAS Global Advisors had registered some SLDs in order to measure
1071 collisions within existing TLDs. The phase two report indicated that they registered over 50 SLDs.
1072 Through their research and analysis, they eventually discovered a vulnerability in Microsoft
1073 products, which was the disclosure-related issue that caused the delay in releasing the phase two
1074 report.

1075 3.7 Potential Changes to Existing gTLD Processes: 2016 – 1076 present

1077 3.7.1 ICANN New gTLD Subsequent Procedures (SubPro) Working Group

1078 *Applicability: Duplicate name collisions, shortened name collisions, search list name collisions,*
1079 *re-registered name collisions*

1080 The purpose of the New gTLD Subsequent Procedures (SubPro) Working Group is to use “the
1081 community’s collective experiences from the 2012 New gTLD Program round to determine what,
1082 if any changes may need to be made to the existing Introduction of New Generic Top-Level
1083 Domains policy recommendations from 8 August 2007.” [107]

1084 A July 2018 initial report from the SubPro Working Group [107] indicated that their Work Track 4
1085 would address name collisions. Section 2.7.8, “Name Collisions,” of the initial report (pages 156-
1086 164) discussed the changes that had occurred regarding name collisions since 2012. The SSAC had
1087 previously provided input to the SubPro Working Group, as documented in SAC 094 (May 22,
1088 2017). [108]

1089 The SubPro Working Group’s initial report included a set of preliminary recommendations for
1090 name collisions:

- 1091 • “2.7.8.c.1: Include a mechanism to evaluate the risk of name collisions in the TLD
1092 evaluation process as well during the transition to delegation phase.
- 1093 • 2.7.8.c.2: Use data-driven methodologies using trusted research-accessible data sources
1094 like *Day in the Life of the Internet* (DITL) and *Operational Research Data from Internet*
1095 *Namespace Logs* (ORDINAL).
- 1096 • 2.7.8.c.3: Efforts should be undertaken to create a ‘Do Not Apply’ list of TLD strings that
1097 pose a substantial name collision risk whereby application for such strings would not be
1098 allowed to be submitted.
- 1099 • 2.7.8.c.4: In addition, a second list of TLDs should be created (if possible) of strings that
1100 may not pose as high of a name collision risk as the ‘Do Not Apply’ list, but for which
1101 there would be a strong presumption that a specific mitigation framework would be
1102 required.
- 1103 • 2.7.8.c.5: Allow every application, other than those on the ‘do not apply’ list, to file a name
1104 collision mitigation framework with their application.
- 1105 • 2.7.8.c.6: During the evaluation period, a test should be developed to evaluate the name
1106 collision risk for every applied-for string, putting them into 3 baskets: high risk, aggravated
1107 risk, and low risk. Provide clear guidance to applicants in advance for what constitutes high
1108 risk, aggravated risk, and low risk.
- 1109 • 2.7.8.c.7: High risk strings would not be allowed to proceed and would be eligible for some
1110 form of a refund.
- 1111 • 2.7.8.c.8: Aggravated risk strings would require a non-standard mitigation framework to
1112 move forward in the process; the proposed framework would be evaluated by an RSTEP
1113 panel.
- 1114 • 2.7.8.c.9: Low risk strings would start controlled interruption as soon as such finding is
1115 reached, recommended to be done by ICANN org for a minimum period of 90 days (but
1116 likely more considering the typical timeline for evaluation, contracting and delegation).

1117 • 2.7.8.c.10: If controlled interruption (CI) for a specific label is found to cause disruption,
1118 ICANN org could decide to disable CI for that label while the disruption is fixed, provided
1119 that the minimum CI period still applied to that string.”

1120 The SSAC provided feedback on the SubPro Working Group’s initial report in SAC 103, posted
1121 October 3, 2018. [109]

1122 As of this writing, the final report from the SubPro Working Group is not yet available.

1123 3.7.2 Requests to Delegate corp, home, and mail

1124 *Applicability: Duplicate name collisions, shortened name collisions, search list name collisions*

1125 **Background on corp, home, and mail being reserved**

1126 As previously discussed in Section 3.3.1, the ICANN SSAC’s SAC 045 report in 2010 [27]
1127 recommended prohibiting the delegation of certain domain names as TLDs. ICANN’s *gTLD*
1128 *Applicant Guidebook* [25] released in 2012 specified prohibited names that included the reserved
1129 TLD names from RFC 2606 [28]—test, example, invalid, and localhost—plus a few dozen more,
1130 in what was termed the Top-Level Reserved Names List. Most of the additional names were
1131 specific to internet infrastructure, like apnic, iab, iana, icann, ietf, and ssac, while a few were more
1132 general, such as local.

1133 In February 2013, RFC 6761, *Special-Use Domain Names* [110] defined how the RFC 2606 names
1134 should be treated, with RFC 6762, *Multicast DNS* [33] providing additional guidance on handling
1135 usage of the reserved names from RFC 6761.

1136 Section 3.4 of this report discussed in detail the evaluation of various TLD names, and Section 3.6
1137 covered the recommendations from the JAS Global Advisors Phase One final report [95] and the
1138 ICANN approval of the Name Collision Occurrence Management Framework [98], which
1139 prevented delegation of the corp, home, and mail TLDs for the time being.

1140 Since at least 2013, perhaps earlier, parties have been asking for the corp, home, and mail TLDs to
1141 be delegated. In August 2016, a group of applicants for those three TLDs sent a letter to ICANN
1142 asking for the names to be released because the risks that were present some years ago have been
1143 mitigated. [111] The letter included the following:

1144 “.HOME, .CORP, and .MAIL were originally put on the high-risk list due to an anticipated
1145 combined effect of conflict with internal name certificate authority use and the number of
1146 queries to the root where no name existed (sometimes referred to as ‘name collisions’). The
1147 unreliability of self-assigned certificates, however, was mitigated last year with the
1148 reassignment of certificates to internal names and private IP addresses (i.e., for internal
1149 networks). This effective mitigation, coupled with the completion of controlled interruption
1150 of new gTLDs without incident, presents evidence that risks anticipated by the JAS report
1151 were grossly overstated.

1152 These results, at a minimum, call for a new examination to determine whether the basis for
1153 the Board’s earlier decision to stymie .HOME, .CORP, and .MAIL remains valid, and

1154 whether the original assumptions and recommendations continue to hold, given current
1155 experience. Just as the name collision issues were mitigated in all other gTLDs, the same
1156 likely is true for these three gTLDs.”

1157 RFC 8244, *Special-Use Domain Names Problem Statement* was released in October 2017. [112]
1158 Among the challenges it discussed were those involving reserving additional domain names so
1159 they are not publicly delegated as TLDs. RFC 8244 referenced an Internet-Draft from 2015,
1160 *Additional Reserved Top Level Domains*, that was not finalized and expired. [113] That Internet-
1161 Draft proposed classifying the corp, home, and mail domain names as reserved in compliance with
1162 RFC 6761. [110] RFC 8244 also referenced RFC 7788, *Home Networking Control Protocol*, [114]
1163 which specified in Section 8 the use of “.home” as the default “network-wide zone” for name
1164 resolution on a home network.

1165 In response to the August 2016 letter, the ICANN Board approved resolutions on November 2,
1166 2017 regarding the corp, home, and mail strings. [115] The resolutions indicated that “the effect of
1167 name collisions on interoperability, resilience, security and/or stability of the DNS is not fully
1168 understood” and “the Board has made no determination as to the efficacy or feasibility of potential
1169 mitigation mechanisms for Name Collision, and remains focused on minimizing or avoiding risk to
1170 the security and stability of the DNS.”

1171 Consequently, the Board asked the ICANN SSAC “to conduct a study... to present data, analysis
1172 and points of view, and provide advice to the Board regarding the risks posed to users and end
1173 systems if .CORP, .HOME, .MAIL strings were to be delegated in the root, as well as possible
1174 courses of action that might mitigate the identified risks,” as well as a study on several questions
1175 related to name collisions in general. That was the driver for this NCAP Phase 1 study and report.

1176 Note that while there has been continued interest in delegating corp, home, and mail, there has also
1177 been continued interest in not delegating them and in reserving additional names. For example,
1178 there was a 2017 Internet-Draft proposing reservation of “.internal” as a TLD. [116] There is
1179 another Internet-Draft, started in 2014 and still in progress as of this writing, proposing “.alt” as a
1180 reserved domain name not to be used for DNS. [117]

1181

1182 4 The Known Harm of Name Collisions and the 1183 Technical Impact of Controlled Interruption

1184 The study RFP [2] specified that this report must include the following:

- 1185 • Study task 2b: “summarizes the known (evidenced) harm of name collisions”
- 1186 • Study task 2d: “documents any mitigations/actions taken so far, specifically including
1187 controlled interruption, and the technical impact of those mitigations only (no examination
1188 to be undertaken of the non-technical impacts such as resourcing or costs)” (note: the first
1189 part of this was already documented in Section 3.6)

1190 Much of the publicly available information on the known harm of name collisions is not relevant
1191 for evaluating current and future risks because it is outdated. Generally, what has happened is there
1192 has been increased awareness of a particular cause of name collisions, so that cause was addressed
1193 and future harm was avoided. An example is re-registered name collisions, as discussed in Section
1194 3.2. There were definitely organizations harmed by their domains expiring and subsequently being
1195 registered and misused by others, but this has been a known issue for many years, and
1196 organizations have full control over and responsibility for preventing this form of name collision.
1197 Similarly, the duplicate name collision risks from Internal Name certificates (see Section 3.3.2)
1198 were addressed by CAs changing their processes.

1199 Accordingly, this section of the report summarizes the known harm of name collisions for TLDs
1200 since controlled interruption for new TLD delegation began. Most of the harm or potential for
1201 harm should have occurred during the 90-day controlled interruption periods, which became
1202 mandatory for new gTLD delegation starting in August 2014 and were recommended for new
1203 ccTLD delegation in October 2014 [118], so this section of the report focuses on known harm of
1204 name collisions that occurred after controlled interruption was mandated.

1205 This section of the report also describes, documents, and analyzes the technical impact of
1206 controlled interruption. Controlled interruption is intended to reduce harm—for example, by
1207 preventing an organization’s network traffic from inadvertently leaking to another organization—
1208 but it can still cause harm, such as by causing that network traffic to be routed to the special
1209 loopback address. Any discussion of harm from name collisions will be closely tied with a
1210 discussion of the technical impact of controlled interruption, so both topics are discussed jointly in
1211 this section.

1212 4.1 Preparation

1213 As described in Section 3.6.1, controlled interruption was proposed for use to help mitigate name
1214 collision risks for new gTLDs. Section 3.6.5 explained that the Name Collision Occurrence
1215 Management Framework [98] was approved on July 30, 2014, and it required registry operators to
1216 do continuous controlled interruption for each new gTLD for a minimum of 90 days. The same
1217 controlled interruption measures were recommended for each new ccTLD on October 2, 2014.
1218 [118]

1219 Attempts to query a new TLD during the controlled interruption period for an “A” record (an IP
1220 address) would result in a reply utilizing the loopback address 127.0.53.53. The idea was that this
1221 address would be unexpected and unusual, with the repeated “53” values implying the relationship
1222 to DNS. DNS queries looking for text records (“TXT”) would return the following: “Your DNS
1223 configuration needs immediate attention see <https://icann.org/namecollision>”. Other types of DNS
1224 queries would return an answer containing the string “your-dns-needs-immediate-attention.” as
1225 part of the domain name. Doing a subsequent query for that domain name would return the
1226 127.0.53.53 address. [99]

1227 ICANN also increased awareness of controlled interruption through other means. This ranged from
1228 creating online technical resources like webpages [14] and the *Guide to Name Collision*
1229 *Identification and Mitigation for IT Professionals* [99] to having social media [119], articles [120],
1230 and even Google ads [121] that referenced the 127.0.53.53 address, controlled interruption, and
1231 ICANN’s name collision resources website [14].

1232 Finally, ICANN provided a webform so any parties adversely affected by a name collision
1233 (including a controlled interruption) could report it. [122] The page currently says, in part, “If you
1234 believe your name collision meets the criteria above (i.e. your system is suffering demonstrably
1235 severe harm as a consequence of name collision or you have a reasonable belief that the name
1236 collision presents a clear and present danger to human life), please use the form below to submit
1237 your report to ICANN.”


1238 4.2 Name Collision Reports

1239 There is no way to quantify the number of name collisions encountered during controlled
1240 interruption periods, let alone the nature of the collisions, such as severity, length of time, or name
1241 collision cause. Any study of actual name collisions during controlled interruption will be largely
1242 anecdotal. To get a somewhat broader view of name collisions, this section looks at reports made
1243 both to ICANN and to others (e.g., system administrator sites, user forums, bug tracking systems).

1244 Note that, as mentioned in Section 3.6.7 of this report, JAS Global Advisors stated in their phase
1245 two report [106] that there were no significant problems from the delegation of new gTLDs (as of
1246 that writing, approximately 650 gTLDs).

1247 4.2.1 Reports to ICANN

1248 The lower (blue) line in Figure 1 shows the number of name collision reports ICANN received by
1249 half-year. [123] The upper (orange) line shows the number of new gTLDs delegated during the
1250 same half-year periods. [8] Note that as of this writing, there have been a total of 57 IDN ccTLDs
1251 delegated through the IDN Fast Track Program since the first new IDN ccTLD requests were
1252 submitted in 2009. [124] Because there have been so few new ccTLDs compared to new gTLDs,
1253 and even compared to new collision reports, a line on the graph for the new ccTLDs would not be
1254 distinguishable from zero values, so it has been omitted.



1255

1256

Figure 1: Name Collision Reports to ICANN by Half Year

1257

The following statements are based on the data used for Figure 1:

1258

- The vast majority of new TLDs delegated since July 2014 have not been the subject of any name collision reports to ICANN.

1259

1260

- For every one report in the second half of 2014, there were approximately eight TLDs delegated. During 2015, the ratio was roughly one report to 26 TLDs, and in 2016 it was one report to 57 TLDs.

1261

1262

1263

- During the three-year period from 2017 through 2019, there was only one report to ICANN.

1264

1265 Additional analysis was performed on the name collision reports ICANN received. A few of those
 1266 reports were incomplete, so the following statements are based on analysis of the complete reports
 1267 only:

1268

- Each report specified how many days after the new TLD's delegation the problem began. As a reminder, controlled interruption was to last at least 90 days after initial delegation.

1269

1270

- The range was from 1 day to 991 days (roughly 2.7 years).

1271

- The median value was 23 days.

1272

- About one-fourth of the reported problems were detected within seven days of delegation. Just over half the problems were detected within 30 days of delegation.

1273

1274 ○ Nearly 30% of the problems were not detected until after 90 days of delegation.
1275 However, 80% of the problems not detected at 90 days were still not detected after
1276 180 days, and half of those were not yet detected even a year after delegation.

1277 • Around 60% of the reporting organizations said their corporate network was affected. Just
1278 over 25% said individual computers were affected, and over 10% cited applications or
1279 application development.

1280 Of all the reports to ICANN, only one led to action by a registry. In that case, a large organization
1281 had reported disruption of its services on the first day after new TLD delegation. The registry
1282 operator for the new TLD voluntarily chose to temporarily stop controlled interruption for that
1283 TLD. After the affected organization updated its systems to correct the problem, the registry
1284 operator was able to resume controlled interruption for the TLD.

1285 4.2.2 Reports to Others

1286 For the purposes of this study, a member of the ICANN NCAP Discussion Group created and
1287 provided a list of URLs for 50 publicly known instances of name collisions identified through
1288 controlled interruption. Each instance was reviewed, and the 33 instances where the nature of the
1289 problem could be determined based on the available information were further evaluated. Note that
1290 these reports are strictly anecdotal, so while some insights can be gleaned from analyzing them, the
1291 accuracy of each report cannot readily be verified, and thus drawing specific conclusions from
1292 individual reports is unwise.

1293 Most of the 33 evaluated instances involved duplicate name collisions, where there was internal-
1294 only use of a domain that was subsequently publicly delegated. In nearly half of those cases, dev
1295 was the TLD in question, with the prod, bar, and box TLDs each also cited in multiple cases, and
1296 several other TLDs cited once. The rest of the 33 instances involved shortened name collisions.

1297 Several of the evaluated instances affected an individual, typically someone using a domain on a
1298 personally owned computer or home network until public delegation of that domain caused the
1299 home configuration to stop working.

1300 To look for additional publicly known instances of name collisions besides those on the list of 50,
1301 searches were conducted using terms such as “127.0.53.53”, “name collision”, “controlled
1302 interruption”, and “outage” to identify news articles, blog postings, forum discussions, and other
1303 accounts of the technical impact of name collisions and controlled interruption. No significant new
1304 information was found other than additional instances of name collisions found through controlled
1305 interruption, similar to those on the list of 50. It was noted that the volume of new accounts of
1306 name collision-related problems has dropped sharply over the past few years, with only a handful
1307 of such postings made during all of 2019.

1308

1309 5 Datasets for Name Collision Studies

1310 The study RFP [2] specified that this report must include the following:

- 1311 • Study task 3: “Identify datasets used in past studies and determine if those datasets are still
1312 available and any constraints there may be regarding access.”
- 1313 • Study task 4: “Identify gaps in the datasets used by previous studies, resulting in a list of
1314 additional datasets or data providers that would be necessary to successfully complete
1315 Studies 2 and 3.”
- 1316 • Study task 5: “Assess the potential availability of these additional datasets.”

1317 Section 5.1 discusses the first item (datasets from past studies), and Section **Error! Reference**
1318 **source not found.** covers the other two (identify gaps in datasets from past studies, list what is
1319 needed to fill those gaps, and assess the availability of items on the list).

1320 5.1 Datasets Used in Past Studies

1321 Most past studies of name collisions have used data from DNS-OARC Day in the Life of the
1322 Internet (DITL) [37]. Authors of work cited in this report that used data from DITL include
1323 Demand Media [52], DigiCert [43], Donuts [42], ICANN [38], Interisle [36] and [51], JAS Global
1324 Advisors [95] and [106], JAS Global Advisors and simMachines [44], Verisign [74], Verisign Labs
1325 [58], and Verisign Labs and University of Michigan [77] and [83].

1326 According to [37], DITL data is currently available for every year from 2006 through 2018, and
1327 “access to this data requires a current OARC paying membership, or in lieu of payment...a
1328 mutually beneficial form of in-kind membership.” The lowest-priced paid membership as of this
1329 writing is \$1100 per year, which allows two people to participate. [125] Note that [37] states that
1330 OARC members have access to OARC analysis machines, and that OARC requires “that the data
1331 may not be copied off OARC servers to any other host or network beyond OARC’s access and
1332 control.”

1333 Another dataset mentioned by a previous name collision report is the Operational Research Data
1334 from Internet Namespace Logs (ORDINAL) dataset [79]. ORDINAL is housed by the Information
1335 Marketplace for Policy and Analysis of Cyber-Risk & Trust (IMPACT). It appears that access to
1336 IMPACT is free, and researchers in the United States and several other countries approved by the
1337 US Department of Homeland Security (DHS) are eligible for IMPACT access. [126] ORDINAL
1338 data is being provided on an ongoing basis by JAS Global Advisors, and ORDINAL “contains
1339 robust DNS protocol layer data, select application layer data, standard activity logs, received select
1340 transmissions, and packet captures of associated activity originally intended to study the impact of
1341 DNS namespace collisions. The dataset is generated via Internet activity to sensor nodes which are
1342 linked to high activity Domain Names.” [127]

1343 In addition to the DITL and ORDINAL datasets, there are also ICANN name collision reports with
1344 pertinent information on actual name collisions, their characteristics, and their outcomes. Sanitized
1345 summaries of all name collision reports received to date were provided for the purposes of this

1346 report, and it is assumed up-to-date summaries could be provided for the authors of Studies 2 and 3
1347 as needed and appropriate. [123]

1348 Finally, there was a dataset particular to the corp.com domain, as mentioned in 2014 in [72] and
1349 discussed in Section 3.5.1 of this report. The corp.com domain was receiving many queries that
1350 were believed to be leaking from internal .corp domains. The current availability of data for the
1351 corp.com domain and any constraints on its access are unknown and would need to be assessed
1352 early during the performance of Study 2. See [128] for more information on the status of corp.com
1353 as of this writing. A recent update indicated that new corp.com data is no longer being provided to
1354 ORDINAL. [129]

1355 5.2 Additional Datasets Needed for Studies 2 and 3

1356 The plans for Studies 2 and 3 are outlined in Section 3.3 of the *SSAC Proposal for the Name*
1357 *Collision Analysis Project* from February 2019. [1] Study 2, “Name Collision Root Cause and
1358 Impact Analysis, and Data Repository” would involve gathering datasets in a data repository,
1359 conducting an analysis of that data to understand the root cause of most name collisions, and
1360 testing mitigation strategies. Study 3, “Analysis of Mitigation Options,” would be a deeper
1361 examination of mitigation strategies, with specific guidance to be produced on the potential
1362 delegation of the corp, home, and mail TLDs, as well as other TLDs likely to cause name
1363 collisions.

1364 It is not obvious that additional datasets would be needed for Studies 2 and 3. Information on
1365 previous and recent leakage of corp, home, and mail should already be captured in the DITL and
1366 ORDINAL datasets. A current dataset for corp.com could be valuable for comparing current
1367 leakage of the corp domain to 2014-era leakage. Similar datasets for the home and mail
1368 counterparts to corp.com (e.g., home.com and mail.com) might also be valuable, although much of
1369 the same information might be available through the DITL and ORDINAL datasets.

1370 As for identifying causes of name collisions, they have already been established in some cases,
1371 usually by individuals researching a particular leaked TLD to find its origin. There is unlikely to be
1372 any dataset that would contain root causes; identifying root causes is generally going to require
1373 research on a case-by-case basis. Based on previous research and studies, such as [76], it seems
1374 quite likely that there is not a single root cause for most name collisions, but rather several types of
1375 root causes.

1376

6 Bibliography

- [1] ICANN SSAC, "SSAC Proposal for the Name Collision Analysis Project," February 2019. [Online]. Available: <https://community.icann.org/download/attachments/79437474/NCAP%20Proposal%20for%20Board%20%28revised%20by%20OCTO%20based%20on%20V2.5BTClean%29%20REDACTED.pdf?api=v2>.
- [2] ICANN, "Project Overview for the Name Collision Analysis Project (NCAP) Study 1: Request for Proposal," 9 July 2019. [Online]. Available: <https://www.icann.org/en/system/files/files/rfp-ncap-study-1-09jul19-en.pdf>.
- [3] ICANN, "ICANN Acronyms and Terms," [Online]. Available: <https://www.icann.org/icann-acronyms-and-terms/icann-acronyms-and-terms/en/nav/A>.
- [4] J. Postel, "RFC 1591, Domain Name System Structure and Delegation," March 1994. [Online]. Available: <https://www.ietf.org/rfc/rfc1591.txt>.
- [5] ICANN, "New gTLD Program," October 2009. [Online]. Available: <https://archive.icann.org/en/topics/new-gtlds/factsheet-new-gtld-program-oct09-en.pdf>.
- [6] ICANN, "About the Program," [Online]. Available: <https://newgtlds.icann.org/en/about/program>.
- [7] ICANN, "New Generic Top-Level Domains," [Online]. Available: <https://newgtlds.icann.org/en/>.
- [8] ICANN, "Delegated Strings," [Online]. Available: <https://newgtlds.icann.org/en/program-status/delegated-strings>.
- [9] International Organization for Standardization (ISO), "ISO 3166 Country Codes," [Online]. Available: <https://www.iso.org/iso-3166-country-codes.html>.
- [10] ICANN, "IDN ccTLD Fast Track Process," [Online]. Available: <https://www.icann.org/resources/pages/fast-track-2012-02-25-en>.
- [11] ICANN, "Resources for Country Code Managers," 25 February 2012. [Online]. Available: <https://www.icann.org/resources/pages/cctlds-21-2012-02-25-en>.
- [12] ICANN, "IDN ccTLD Fast Track String Evaluation Completion," 19 February 2014. [Online]. Available: <https://www.icann.org/resources/pages/string-evaluation-completion-2014-02-19-en>.

- [13] ICANN, "Frequently Asked Questions: IDN ccTLDs by Country," 25 February 2012. [Online]. Available: <https://www.icann.org/resources/pages/faqs-5b-2012-02-25-en>.
- [14] ICANN, "Name Collision Resources & Information," [Online]. Available: <https://www.icann.org/resources/pages/name-collision-2013-12-06-en>.
- [15] Internet Architecture Board, "IAB Commentary: Architectural Concerns on the Use of DNS Wildcards, September 2003," 19 September 2003. [Online]. Available: <https://www.iab.org/documents/correspondence-reports-%20documents/docs2003/2003-09-20-dns-wildcards/>.
- [16] ICANN SSAC, "Redirection in the com and net Domains," 9 July 2004. [Online]. Available: <http://www.icann.org/committees/security/ssac-report-09jul04.pdf>.
- [17] ICANN SSAC, "SAC 015: Why Top Level Domains Should Not Use Wildcard Resource Records," 10 November 2006. [Online]. Available: <https://www.icann.org/groups/ssac/documents/sac-015-en>.
- [18] ICANN SSAC, "SAC 032: Preliminary Report on DNS Response Modification," June 2008. [Online]. Available: <https://www.icann.org/en/system/files/files/sac-032-en.pdf>.
- [19] E. Gavron, "RFC 1535, A Security Problem and Proposed Correction With Widely Deployed DNS Software," October 1993. [Online]. Available: <https://tools.ietf.org/rfc/rfc1535.txt>.
- [20] ICANN Registry Services Technical Evaluation Panel (RSTEP), "Report on Internet Security and Stability Implications of the Tralliance Corporation search.travel Wildcard Proposal," 2 November 2006. [Online]. Available: <https://www.icann.org/en/system/files/files/tralliance-report-09nov06-en.pdf>.
- [21] ICANN SSAC, "SAC 041: Recommendation to prohibit use of redirection and synthesized responses by new TLDs," 10 June 2009. [Online]. Available: <https://www.icann.org/en/system/files/files/sac-041-en.pdf>.
- [22] ICANN, "ICANN Archives, Verisign's Wildcard Service Deployment," [Online]. Available: <https://archive.icann.org/en/topics/wildcard-history.html>.
- [23] ICANN SSAC, "SAC 010: Renewal Considerations for Domain Name Registrants," June 2006. [Online]. Available: <https://www.icann.org/en/system/files/files/renewal-advisory-29jun06-en.pdf>.
- [24] ICANN SSAC, "SAC 011: Problems caused by the non-renewal of a domain name associated with a DNS Name Server," June 2006. [Online]. Available: <https://www.icann.org/en/system/files/files/renewal-nameserver-07jul06-en.pdf>.

- [25] ICANN, "New gTLD Application Guidebook," 4 June 2014. [Online]. Available: <https://newgtlds.icann.org/en/applicants/agb/guidebook-full-04jun12-en.pdf>.
- [26] G. Kirikos, "Most Popular Invalid TLDs Should Be Reserved," 18 June 2009. [Online]. Available: http://www.circleid.com/posts/20090618_most_popular_invalid_tlds_should_be_reserved/.
- [27] ICANN SSAC, "SAC 045: Invalid Top Level Domain Queries at the Root Level of the Domain System," 15 November 2010. [Online]. Available: <https://www.icann.org/en/system/files/files/sac-045-en.pdf>.
- [28] D. Eastlake and A. Panitz, "RFC 2606, Reserved Top Level DNS Names," June 1999. [Online]. Available: <https://tools.ietf.org/html/rfc2606>.
- [29] ICANN SSAC, "SAC 057: SSAC Advisory on Internal Name Certificates," 15 March 2013. [Online]. Available: <https://www.icann.org/en/system/files/files/sac-057-en.pdf>.
- [30] Verisign Labs, "Verisign Labs Technical Report #1130007 version 2.1: New gTLD Security and Stability Considerations," March 2013. [Online]. Available: <https://forum.icann.org/lists/comments-name-collision-05aug13/pdfY5loOoWatX.pdf>.
- [31] Verisign Labs, "Verisign Labs Technical Report #1130007 version 2.2: New gTLD Security and Stability Considerations," 28 March 2013. [Online]. Available: <https://www.verisign.com/assets/gtld-ssr-v2.1-final.pdf>.
- [32] B. Hill and B. Smith, "Re: Proposed delegation of invalid names from SAC 045 and RFC 6762," 15 March 2013. [Online]. Available: <https://www.icann.org/en/system/files/correspondence/hill-smith-to-chehade-crocker-15mar13-en.pdf>.
- [33] S. Cheshire and M. Krochmal, "RFC 6762, Multicast DNS," February 2013. [Online]. Available: <https://tools.ietf.org/html/rfc6762>.
- [34] O. Kolkman, A. Sullivan and W. Kumari, "Internet-Draft draft-kolkman-cautious-delegation-00, A Procedure for Cautious Delegation of a DNS Name," 2 May 2013. [Online]. Available: <https://www.ietf.org/archive/id/draft-kolkman-cautious-delegation-00.txt>.
- [35] O. Kolkman, A. Sullivan and W. Kumari, "Internet-Draft draft-kolkman-cautious-delegation-02, A Procedure for Cautious Delegation of a DNS Name," 1 August 2013. [Online]. Available: <https://www.ietf.org/archive/id/draft-kolkman-cautious-delegation-02.txt>.
- [36] Interisle Consulting Group, "Name Collision in the DNS: A study of the likelihood and potential consequences of collision between new public gTLD labels and existing private

- uses of the same strings, version 1.5," 2 August 2013. [Online]. Available: <https://www.icann.org/en/system/files/files/name-collision-02aug13-en.pdf>.
- [37] DNS-OARC, "Day In The Life of the Internet (DITL)," [Online]. Available: <https://www.dns-oarc.net/oarc/data/ditl>.
- [38] ICANN, "New gTLD Collision Risk Mitigation: Proposals to mitigate the collision risks between new gTLDs and existing private uses of the same strings," 5 August 2013. [Online]. Available: <https://www.icann.org/en/system/files/files/new-gtld-collision-mitigation-05aug13-en.pdf>.
- [39] ICANN, "Addressing the Consequences of Name Collisions," 5 August 2013. [Online]. Available: <https://www.icann.org/news/announcement-3-2013-08-05-en>.
- [40] ICANN, "[comments-name-collision-05aug13] Chronological Index," [Online]. Available: <https://forum.icann.org/lists/comments-name-collision-05aug13/index.html>.
- [41] ICANN, "Report of Public Comments: Proposal to Mitigate Name Collision Risks," 5 August 2013. [Online]. Available: <https://forum.icann.org/lists/comments-name-collision-05aug13/pdf3wmJxwMJoR.pdf>.
- [42] Donuts, "Donuts' Comments Regarding Proposal to Mitigate Name Collision Risks," 5 August 2013. [Online]. Available: <https://forum.icann.org/lists/comments-name-collision-05aug13/pdfuanEVzPqbD.pdf>.
- [43] DigiCert, Inc., "Letter from DigiCert to the ICANN Board," 27 August 2013. [Online]. Available: <https://forum.icann.org/lists/comments-name-collision-05aug13/pdfUNz0liz2VL.pdf>.
- [44] J. Schmidt, K. White, D. Conrad and A. Muller-Molina, "Namespace Expansion," 17 September 2013. [Online]. Available: <https://forum.icann.org/lists/comments-name-collision-05aug13/pdf0r8YJwS4iG.pdf>.
- [45] ICANN New gTLD Applicant Group (NTAG), "NTAG Comments on ICANN Name Collision Report," 14 August 2013. [Online]. Available: <https://forum.icann.org/lists/comments-name-collision-05aug13/msg00001.html>.
- [46] E. Osterweil, "Illustrating the Need to Undertake Qualitative Impact Assessments for Applied-For Strings: .WEBSITE, .COFFEE, and .CLUB," 17 September 2013. [Online]. Available: <https://forum.icann.org/lists/comments-name-collision-05aug13/pdf5H5Sqf0igA.pdf>.
- [47] Verisign Labs, "Verisign Labs Technical Report #1130008 Version 1.1: New gTLD Security, Stability Resiliency Update: Exploratory Consumer Impact Analysis," 22 August

2013. [Online]. Available: <https://forum.icann.org/lists/comments-name-collision-05aug13/pdfu6z5kKHEV5.pdf>.
- [48] Neustar, "A Methodology for Assessing Collision Risk and New gTLDs," 17 September 2013. [Online]. Available: <https://www.home.neustar/resources/whitepapers/new-tlds-dns-collision>.
- [49] ICANN, "New gTLD Collision Occurrence Management: Proposal to manage the collision occurrences between new gTLDs and existing private uses of the same strings," 4 October 2013. [Online]. Available: <https://www.icann.org/en/system/files/files/resolutions-new-gtld-annex-1-07oct13-en.pdf>.
- [50] ICANN New gTLD Program Committee, "Approved Resolutions | Meeting of the New gTLD Program Committee," 7 October 2013. [Online]. Available: <https://www.icann.org/resources/board-material/resolutions-new-gtld-2013-10-07-en>.
- [51] J. Reid, "DITL Crunching for gTLD Name Collision Study," 5 October 2013. [Online]. Available: <https://indico.dns-oarc.net/event/1/contributions/46/attachments/38/166/Reid-Crunching.pdf>.
- [52] R. Hooper, "Abusing Resources to Process 7TB of PCAP Data...Or how not to fork-bomb yourself," 5 October 2013. [Online]. Available: <https://indico.dns-oarc.net/event/1/contributions/49/attachments/41/169/DNS-OARC-Abusing-Resources.pdf>.
- [53] A. Simpson, D. McPherson, E. Osterweil, M. Thomas and D. Wessels, "Regional Affinity for Applied for gTLD Strings," 5 October 2013. [Online]. Available: https://indico.dns-oarc.net/event/1/contributions/37/attachments/44/174/gTLD_Regional_Affinity.pdf.
- [54] A. Sullivan, "Using Test Delegations from the Root Prior to Full Allocation and Delegation," 5 October 2013. [Online]. Available: https://indico.dns-oarc.net/event/1/contributions/42/attachments/49/180/Sullivan-Test_delegations.pdf.
- [55] O. Kolkman, A. Sullivan and W. Kumari, "Internet-Draft draft-kolkman-root-test-delegation-00, Using Test Delegations from the Root Prior to Full Allocation and Delegation," 20 September 2013. [Online]. Available: <https://tools.ietf.org/html/draft-kolkman-root-test-delegation-00>.
- [56] ICANN SSAC, "SAC 062: SSAC Advisory Concerning the Mitigation of Name Collision Risk," 7 November 2013. [Online]. Available: <https://www.icann.org/en/system/files/files/sac-062-en.pdf>.
- [57] Verisign Labs, "Preliminary Analysis of SLD Blocking Effectiveness," 5 November 2013. [Online]. Available: <https://www.icann.org/en/system/files/correspondence/kaliski-to-atallah-crain-05nov13-en.pdf>.

- [58] Verisign Labs, "Continued Analysis of SLD Blocking Effectiveness," 15 November 2013. [Online]. Available: <https://www.icann.org/en/system/files/correspondence/kaliski-to-atallah-crain-15nov13-en.pdf>.
- [59] ICANN, "Reports for Alternate Path to Delegation Published," 17 November 2013. [Online]. Available: <https://newgtlds.icann.org/en/announcements-and-media/announcement-2-17nov13-en>.
- [60] "Workshop and Prize on Root Causes and Mitigation of Name Collisions (WPNC)," [Online]. Available: <http://namecollisions.net/program/index.html>.
- [61] M. Thomas, A. Mankin and L. Zhang, "RFC 8023, Report from the Workshop and Prize on Root Causes and Mitigation of Name Collisions," November 2016. [Online]. Available: <https://www.rfc-editor.org/in-notes/rfc8023.html>.
- [62] M. Thomas, Y. Labrou and A. Simpson, "The Effectiveness of Block Lists in Preventing Collisions," 9 March 2014. [Online]. Available: http://namecollisions.net/downloads/wpnc2014_paper_effectiveness_block_lists.pdf.
- [63] P. Hoffman, "Name Collision Mitigation for Enterprise Networks," 10 March 2014. [Online]. Available: http://namecollisions.net/downloads/wpnc14_slides_hoffman_name_collision_mitigation.pdf.
- [64] J. Reid, "Analysing the use of the RA and RD Bits in Queries to Root Servers," 9 March 2014. [Online]. Available: http://namecollisions.net/downloads/wpnc2014_paper_reid.pdf.
- [65] G. Huston, "New gTLD Concerns: Dotless Names and Name Collisions," 12 November 2013. [Online]. Available: <https://labs.ripe.net/Members/gih/dotless-names>.
- [66] D. Piscitello, "Managing Name Collision Occurrences," 6 December 2013. [Online]. Available: <https://www.icann.org/news/blog/managing-name-collision-occurrences>.
- [67] ICANN, "Guide to Name Collision Identification and Mitigation for IT Professionals, Version 1.0," 5 December 2013. [Online]. Available: <https://www.icann.org/en/system/files/files/name-collision-mitigation-05dec13-en.pdf>.
- [68] ICANN SSAC, "SAC 064: SSAC Advisory on DNS 'Search List' Processing," 13 February 2014. [Online]. Available: <https://www.icann.org/en/system/files/files/sac-064-en.pdf>.
- [69] R. Braden, "RFC 1123, Requirements for Internet Hosts -- Application and Support," October 1989. [Online]. Available: <https://tools.ietf.org/html/rfc1123>.

- [70] A. Kumar, J. Postel, C. Neuman, P. Danzig and S. Miller, "RFC 1536, Common DNS Implementation Errors and Suggested Fixes," October 1993. [Online]. Available: <https://tools.ietf.org/html/rfc1536>.
- [71] W. Kumari, "ALT Special Use TLD," 9 March 2014. [Online]. Available: http://namecollisions.net/downloads/wpnc14_slides_internet_engineeringpanel_kumari.pdf.
- [72] C. Strutt, "Looking at corp.com as a Proxy for .corp," 9 March 2014. [Online]. Available: http://namecollisions.net/downloads/wpnc14_slides_strutt_looking_at_corpcom.pdf.
- [73] C. Deccio and D. Wessels, "What's in a Name (Collision): Modeling and Quantifying Collision Potential," 10 March 2014. [Online]. Available: http://namecollisions.net/downloads/wpnc2014_paper_deccio.pdf.
- [74] A. Simpson, "Detecting Search Lists in Authoritative DNS," 10 March 2014. [Online]. Available: http://namecollisions.net/downloads/wpnc2014_paper_simpson.pdf.
- [75] M. Thomas and A. Simpson, "Analysis Techniques for Determining Cause and Ownership of DNS Queries," 9 March 2014. [Online]. Available: http://namecollisions.net/downloads/wpnc14_paper_simpson_thomas.pdf.
- [76] M. Thomas and A. Mohaisen, "Measuring the Leakage of Onion at the Root: A measurement of Tor's .onion pseudo-top-level domain in the global domain name system," November 2014. [Online]. Available: <https://www.verisign.com/assets/labs/Measuring-the-Leakage-of-Onion-at-the-Root.pdf>.
- [77] A. Mohaisen and K. Ren, "Leakage of .onion at the DNS Root: Measurements, Causes, and Countermeasures," October 2017. [Online]. Available: <http://seal.cs.ucf.edu/doc/17-tnet.pdf>.
- [78] J. Appelbaum and A. Muffett, "RFC 7686, The ".onion" Special-Use Domain Name," October 2015. [Online]. Available: <https://tools.ietf.org/html/rfc7686>.
- [79] "Introducing: The ORDINAL Dataset," 13 May 2017. [Online]. Available: <https://www.icann.org/en/system/files/files/presentation-ordinal-datasets-colliding-domains-13may17-en.pdf>.
- [80] Verisign, "Enterprise Remediation for WPAD Name Collision Vulnerability," 23 May 2016. [Online]. Available: https://www.verisign.com/assets/Enterprise_Remediation_for_WPAD_Name_Collision_Vulnerability.pdf.

- [81] Q. Chen, E. Osterweil, M. Thomas and Z. Mao, "MitM Attack by Name Collision: Cause Analysis and Vulnerability Assessment in the New gTLD Era," May 2016. [Online]. Available: <https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=7546529>.
- [82] Cybersecurity and Infrastructure Security Agency (CISA), "Alert (TA16-144A), WPAD Name Collision Vulnerability," 23 May 2016. [Online]. Available: <https://www.us-cert.gov/ncas/alerts/TA16-144A>.
- [83] Q. Chen, M. Thomas, E. Osterweil, Y. Cao, J. You and Z. Mao, "Client-Side Name Collision Vulnerability in the New gTLD Era: A Systematic Study," November 2017. [Online]. Available: https://www.ics.uci.edu/~alfchen/alfred_ccs17.pdf.
- [84] E. Osterweil, D. McPherson, M. Thomas and Q. Chen, "Detecting and Remediating Highly Vulnerable Domain Names Using Passive DNS Measurements". United States of America Patent US20170279846A1, 24 March 2017.
- [85] JAS Global Advisors, "Mitigating the Risk of DNS Namespace Collisions: A Study on Namespace Collisions in the Global Internet DNS Namespace and a Framework for Risk Mitigation, Phase One Report," 24 February 2014. [Online]. Available: <https://www.icann.org/en/system/files/files/name-collision-mitigation-26feb14-en.pdf>.
- [86] J. Schmidt, "Mitigating the Risk of DNS Namespace Collisions," 10 March 2014. [Online]. Available: http://namecollisions.net/downloads/wpnc14_slides_jas_framework_session.pdf.
- [87] J. Schmidt, "Name Collision Mitigation Update," 24 March 2014. [Online]. Available: <https://archive.icann.org/meetings/singapore2014/en/schedule/mon-name-collision/presentation-name-collision-24mar14-en.pdf>.
- [88] "Transcript of ICANN Name Collision Mitigation meeting in Singapore," 24 March 2014. [Online]. Available: <https://archive.icann.org/meetings/singapore2014/en/schedule/mon-name-collision/transcript-name-collision-24mar14-en.pdf>.
- [89] ICANN, "[comments-name-collision-26feb14] Chronological Index," [Online]. Available: <https://forum.icann.org/lists/comments-name-collision-26feb14/index.html>.
- [90] ICANN, "Report of Public Comments: Mitigating the Risk of DNS Namespace Collisions," 10 June 2014. [Online]. Available: <https://www.icann.org/en/system/files/files/report-comments-name-collision-10jun14-en.pdf>.
- [91] Verisign, "Preliminary Comments on 'Mitigating the Risk of DNS Namespace Collisions' Phase One Report," 24 February 2014. [Online]. Available: <https://forum.icann.org/lists/comments-name-collision-26feb14/pdfNPWfDHk1pu.pdf>.

- [92] B. Kaliski, "Name Collisions in the Domain Name System," 17 April 2014. [Online]. Available: <http://www.verisign.com/assets/Verisign-Kaliski-Collisions-US-Telecom-04162014.pptx>.
- [93] Verisign, "Additional Comments on "Mitigating the Risk of DNS Namespace Collisions" Phase One Report," 21 April 2014. [Online]. Available: <https://forum.icann.org/lists/comments-name-collision-26feb14/pdfTWUAZM3gBN.pdf>.
- [94] B. Kaliski, "The Real Uneven Playing Field of Name Collisions," 16 May 2014. [Online]. Available: <https://blog.verisign.com/security/the-real-uneven-playing-field-of-name-collisions/>.
- [95] JAS Global Advisors, "Mitigating the Risk of DNS Namespace Collisions: A Study on Namespace Collisions in the Global Internet DNS Namespace and a Framework for Risk Mitigation, Phase One Report," 4 June 2014. [Online]. Available: <https://www.icann.org/en/system/files/files/name-collision-mitigation-study-06jun14-en.pdf>.
- [96] Verisign, "Re: ICANN's Proposal to Mitigate Name Collision Risks - .CBA Case Study," 15 September 2013. [Online]. Available: <https://forum.icann.org/lists/comments-name-collision-05aug13/pdfJUJTT9vS7d.pdf>.
- [97] ICANN SSAC, "SAC 066: SSAC Comment Concerning JAS Phase One Report on Mitigating the Risk of DNS Namespace Collisions," 6 June 2014. [Online]. Available: <https://www.icann.org/en/system/files/files/sac-066-en.pdf>.
- [98] ICANN, "Name Collision Occurrence Management Framework," 30 July 2014. [Online]. Available: <https://www.icann.org/en/system/files/files/name-collision-framework-30jul14-en.pdf>.
- [99] ICANN, "Guide to Name Collision Identification and Mitigation for IT Professionals, Version 1.1," 1 August 2014. [Online]. Available: <https://www.icann.org/en/system/files/files/name-collision-mitigation-01aug14-en.pdf>.
- [100] ICANN Global Domains Division, "Name Collision Occurrence Assessment," 4 August 2014. [Online]. Available: <https://newgtlds.icann.org/sites/default/files/agreements/name-collision-assessment-04aug14-en.htm>.
- [101] ICANN Global Domains Division, "Addendum to Name Collision Occurrence Assessment," 14 November 2014. [Online]. Available: <https://newgtlds.icann.org/sites/default/files/agreements/name-collision-assessment-addendum-14nov14-en.htm>.

- [102] ICANN, "Frequently Asked Questions: Name Collision Occurrence Management Framework for Registries," [Online]. Available: <https://www.icann.org/resources/pages/name-collision-ro-faqs-2014-08-01-en>.
- [103] ICANN, "Frequently Asked Questions: Name Collision for IT Professionals," [Online]. Available: <https://www.icann.org/resources/pages/name-collision-it-pros-faqs-2014-08-01-en>.
- [104] ICANN, "Briefing on Name Collision Risk for New TLDs," 18 February 2014. [Online]. Available: <https://www.icann.org/en/system/files/files/name-collision-risk-18feb14-en.pdf>.
- [105] ICANN, "Name Collision Occurrence Mitigation for New ccTLDs," 2 October 2014. [Online]. Available: <https://www.icann.org/resources/pages/cctld-mitigation-2014-10-02-en>.
- [106] JAS Global Advisors, "Mitigating the Risk of DNS Namespace Collisions: A Study on Namespace Collisions in the Global Internet DNS Namespace and a Framework for Risk Mitigation, Final Report," 28 October 2015. [Online]. Available: <https://www.icann.org/en/system/files/files/name-collision-mitigation-final-28oct15-en.pdf>.
- [107] ICANN GNSO New gTLD Subsequent Procedures Working Group, "Initial Report on the new gTLD Subsequent Procedures Policy Development Process (Overarching Issues & Work Tracks 1-4)," 3 July 2018. [Online]. Available: <https://gns0.icann.org/sites/default/files/file/field-file-attach/subsequent-procedures-initial-overarching-issues-work-tracks-1-4-03jul18-en.pdf>.
- [108] ICANN SSAC, "SAC 094: SSAC Response to the New gTLD Subsequent Procedures Policy Development Process (PDP) Working Group Community Comment 2," 22 May 2017. [Online]. Available: <https://www.icann.org/en/system/files/files/sac-094-en.pdf>.
- [109] ICANN SSAC, "SAC 103: SSAC Response to the new gTLD Subsequent Procedures Policy Development Process Working Group Initial Report," 3 October 2018. [Online]. Available: <https://www.icann.org/en/system/files/files/sac-103-en.pdf>.
- [110] S. Cheshire and M. Krochmal, "RFC 6761, Special-Use Domain Names," February 2013. [Online]. Available: <https://tools.ietf.org/html/rfc6761>.
- [111] .Home Registry Inc. et al, "Letter to Members of the ICANN Board," 24 August 2016. [Online]. Available: <https://www.icann.org/en/system/files/correspondence/home-registry-inc-et-al-to-icann-board-24aug16-en.pdf>.
- [112] T. Lemon, R. Droms and W. Kumari, "RFC 8244, Special-Use Domain Names Problem Statement," October 2017. [Online]. Available: <https://tools.ietf.org/html/rfc8244>.

- [113] L. Chapin and M. McFadden, "Internet-Draft draft-chapin-additional-reserved-tlds-02, Additional Reserved Top Level Domains," 2 March 2015. [Online]. Available: <https://www.ietf.org/archive/id/draft-chapin-additional-reserved-tlds-02.txt>.
- [114] M. Stenberg, S. Barth and P. Pfister, "RFC 7788, Home Networking Control Protocol," April 2016. [Online]. Available: <https://tools.ietf.org/html/rfc7788>.
- [115] ICANN, "Approved Board Resolutions for 02 Nov 2017, section 2a, Consideration of .CORP, .HOME, and .MAIL and other Collision Strings, Rationale for Resolutions 2017.11.02.29-2017.11.02.31," 2 November 2017. [Online]. Available: <https://www.icann.org/resources/board-material/resolutions-2017-11-02-en#2.a>.
- [116] W. Kumari, "Internet-Draft draft-wkumari-dnsop-internal-00, The .internal TLD," 2 July 2017. [Online]. Available: <https://www.ietf.org/archive/id/draft-wkumari-dnsop-internal-00.txt>.
- [117] W. Kumari and A. Sullivan, "Internet-Draft draft-ietf-dnsop-alt-tld-12, The ALT Special Use Top Level Domain," 23 August 2019. [Online]. Available: <https://datatracker.ietf.org/doc/draft-ietf-dnsop-alt-tld/>.
- [118] ICANN, "Name Collision Occurrence Mitigation for New ccTLDs," 2 October 2014. [Online]. Available: <https://www.icann.org/resources/pages/cctld-mitigation-2014-10-02-en>.
- [119] ICANN, 27 September 2014. [Online]. Available: <https://twitter.com/icann/status/515976677734629377>.
- [120] S. Vaughan-Nichols, "ICANN offers fix for domain name collisions," 15 August 2014. [Online]. Available: <https://www.zdnet.com/article/icann-offers-fix-for-domain-name-collisions/>.
- [121] djchuang, "What does 127.0.53.53 mean? It's a system alert notification thing.," 10 September 2014. [Online]. Available: <https://djchuang.com/127-0-53-53-mean-system-alert-notification-thing/>.
- [122] ICANN, "Report a Name Collision," [Online]. Available: <https://forms.icann.org/en/help/name-collision/report-problems>.
- [123] ICANN, *Summary of sanitized name collision reports*, January, 2020.
- [124] ICANN, "ICANN Announces Successful String Evaluation for the European Commission and Laos IDN ccTLDs," 5 June 2019. [Online]. Available: <https://www.icann.org/news/announcement-2019-06-05-en>.

- [125] DNS-OARC, "DNS-OARC Participation Agreement," January 2020. [Online]. Available: <https://www.dns-oarc.net/files/agreements/oarc-participation.pdf>.
- [126] IMPACT, "Join Impact," [Online]. Available: <https://www.impactcybertrust.org/joinus>.
- [127] IMPACT, "Dataset Details for DS-0794," [Online]. Available: https://impactcybertrust.org/dataset_view?idDataset=794.
- [128] M. Larson, "[NCAP-Discuss] [Ext] Revised draft of NCAP Study 1 report," 5 February 2020. [Online]. Available: <https://mm.icann.org/pipermail/ncap-discuss/2020-February/000202.html>.
- [129] J. Schmidt, "[NCAP-Discuss] Latest draft of Study 1," 12 February 2020. [Online]. Available: <https://mm.icann.org/pipermail/ncap-discuss/2020-February/000233.html>.

1379

1380

1381

DRAFT

7 Acronyms

Acronym	Definition
APNIC	Asia-Pacific Network Information Centre
CA	Certificate Authority
ccNSO	Country Code Names Supporting Organization
ccTLD	Country Code Top-Level Domain
CI	Controlled Interruption
CISA	Cybersecurity and Infrastructure Security Agency
CNNIC	China Internet Network Information Center
DHS	Department of Homeland Security
DITL	Day in the Life of the Internet
DNS	Domain Name System
DNS-OARC	Domain Name System Operations Analysis and Research Center
FAQ	Frequently Asked Questions
FQDN	Fully Qualified Domain Name
GNSO	Generic Names Supporting Organization
gTLD	Generic Top-Level Domain
HTTP	Hypertext Transfer Protocol

Acronym	Definition
IAB	Internet Architecture Board
IANA	Internet Assigned Numbers Authority
ICANN	Internet Corporation for Assigned Names and Numbers
IDN	Internationalized Domain Name
IETF	Internet Engineering Task Force
IMPACT	Information Marketplace for Policy and Analysis of Cyber-Risk & Trust
IP	Internet Protocol
ISO	International Organization for Standardization
IT	Information Technology
MITM	Man in the Middle
NCAP	Name Collision Analysis Project
NCAP DG	Name Collision Analysis Project Discussion Group
NGPC	New gTLD Program Committee
NTAG	New gTLD Applicant Group
OCTO	Office of the Chief Technology Officer
ORDINAL	Operational Research Data from Internet Namespace Logs
PDP	Policy Development Process

Acronym	Definition
RFC	Request for Comments
RFP	Request for Proposal
RIPE NCC	Réseaux IP Européens Network Coordination Centre
RSTEP	Registry Services Technical Evaluation Panel
RZM	Root Zone Management
SLD	Second-Level Domain
SSAC	Security and Stability Advisory Committee
SSL	Secure Sockets Layer
SubPro	Subsequent Procedures (Working Group)
TLD	Top-Level Domain
TTL	Time to Live
URL	Uniform Resource Locator
WPAD	Web Proxy Auto-Discovery
WPNC	Workshop and Prize on Root Causes and Mitigation of Name Collisions
2LD	Second-Level Domain

1383

1384