Identifier System Attack Mitigation Methodology

Introduction

This document is part of ICANN's effort to contribute to enhancing the Stability, Security, and Resiliency (SSR) of the Internet's system of unique identifiers ("Internet Identifiers") by working with the Community to identify and increase awareness of related attacks and to promote broader adoption of attack mitigation practices. This effort also addresses Recommendation #12 of the Security, Stability & Resiliency (SSR) Review Team (SSR-RT) by creating an Identifier System Attack Mitigation Methodology. Specifically, this document identifies and prioritizes types of attacks against the Identifier System, providing a stepping-off point for ICANN to coordinate with the Community to develop a series of short technical documents ("Tech Notes") on actual high-impact attacks and emerging high-risk vulnerabilities. This document will be periodically updated to reflect evolution of both the Identifier System and the cybercrime landscape, supporting on-going efforts within both ICANN and the Community to mitigate attacks that pose the greatest risk to Identifier System SSR.

Attack Mitigation Methodology

ICANN is proposing a new Identifier System Attack Mitigation Methodology to:

- Identify, prioritize, and periodically refresh a list of top Identifier System attacks;
- Develop guidance on actual high-impact attacks and emerging high-risk vulnerabilities;
- Describe corresponding attack mitigation practices that are commonly considered useful; and
- Encourage broader adoption of those practices via contracts, agreements, incentives, etc.

This document represents the first component of this methodology. Specifically, it explores, briefly defines, and prioritizes types of attacks that have had an actual impact on Identifier System SSR, along with emerging high-risk vulnerabilities. This prioritized list is intended to serve as a launch point for ICANN's Office of the CTO, SSR Department (OCTO-SSR) to develop an on-going series of attack-specific Tech Notes to grow awareness within the Community and encourage broader adoption of effective mitigation practices.

Each Tech Note will draw on Community experience and insights to define a type of attack against the Identifier System, using examples to illustrate past/potential attacks and their impacts. Tech Notes will cite available research to illustrate effective attack mitigation practices already in use by some Community members.

OCTO-SSR will develop a plan to socialize each Tech Note, soliciting feedback from and encouraging broader adoption by the ICANN Community. Draft Tech Notes can be updated to reflect Community feedback on mitigation practices and their effectiveness.

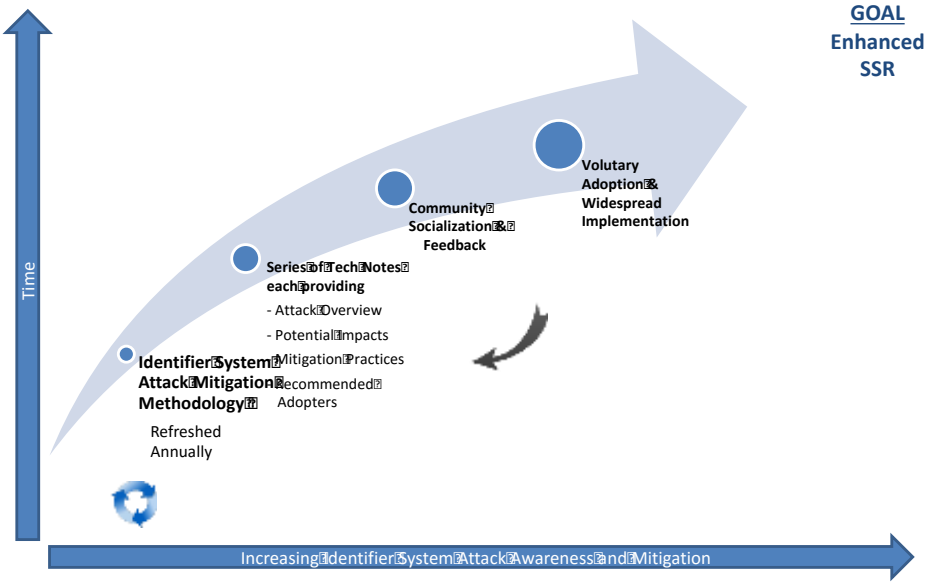

Finally, this document will be periodically revisited to re-examine and re-prioritize types of attacks and emerging vulnerabilities, reflecting evolution of the Identifier System and the Identifier System attack landscape. By implementing this methodology, ICANN hopes to promote greater attack awareness and mitigation within the Community, thereby contributing to enhanced Identifier System SSR.

Identifier System Assets, Vulnerabilities and Threats

ICANN's technical mission includes helping to coordinate, at the overall level, the allocation of the Internet's system of unique identifiers: specifically, top-level domain names, blocks of Internet Protocol (IP) addresses and autonomous system (AS) numbers allocated to the Regional Internet Registries, and protocol parameters as directed by the IETF. As a result, ICANN is responsible for helping to coordinate and collaborate with the relevant operational communities to develop ways in which the stability, security and resiliency of the Identifier Systems can be improved, as well as facilitating the policies associated with these parts of the Identifier Systems.

The role and remit of ICANN's Office of the CTO, SSR Department (OCTO-SSR) is further based on the following definitions:

- **Security** – The capacity to protect and prevent misuse of Internet unique identifiers;
- **Stability** – The capacity to ensure that the Identifier System operates as expected and that users of unique identifiers have confidence that the system operates as expected; and
- **Resiliency** – The capacity of the Identifier System to effectively withstand, tolerate and survive malicious attacks and other disruptive events without disruption or cessation of service.

Within this document, we refer to attacks, assets, vulnerabilities, threats, and risk.[1] To identify and prioritize past/potential attacks that jeopardize Identifier System SSR, we must take into account all four factors and their interactions.

- **Assets** – *What we're trying to protect*. Attackers can use an Identifier System(s) to target and attack property, information, or people, e.g., by using the DNS to implement denial of service attacks. They can also attack aspects of an Identifier System by using it in a manner that is abusive or malicious, for example, by using fraudulent registration information or hijacking names or addresses. Importantly, these forms of attacks can be executed in tandem. These assets include:
    - Authoritative domain name servers and recursive and stub resolvers, as well as domain name registration data, registries, registrars, and registrants.
    - IP addresses and autonomous system numbers (ASNs) employed by the global Internet routing system, along with associated network infrastructure components (e.g., routers, switches, address management systems) and regional Internet registries.

[1] Based on https://www.icann.org/news/blog/threats-vulnerabilities-and-exploits-oh-my

- Protocol parameters and the implementations of the associated protocols that make use of those parameters, both individually and within the context of larger systems that incorporate those protocols and protocol parameters.

Each of these assets performs a different critical function; a specific attack against one may threaten the security, stability, or resiliency of the Identifier System as a whole.

- **Vulnerabilities** – *Weaknesses or gaps that can be exploited.* Vulnerabilities may be flaws within Identifier System assets themselves, or within measures intended to protect them. These vulnerabilities include design defects, coding errors, configuration mistakes, and other gaps that weaken an asset's attack resistance, stability, or resiliency.

- **Threats** – *What we're trying to protect against.* Threats include both entities and events which may exploit an asset's vulnerabilities. Threats to Identifier System assets include attacks against domain name servers and name resolvers, or network elements such as routers or switches. The threat landscape includes attacks against domain and address registration services as well as natural disasters such as storms that trigger power and network outages which degrade Identifier System operations.

- **Risk** – *The probability that threats will exploit vulnerabilities to obtain, damage or destroy assets.* In this document, we attempt to identify high-risk types of attacks against Identifier System assets. To do so, we focus on high-impact vulnerabilities that are being actively exploited by threats, as well as new vulnerabilities at high risk of exploitation.

Top Identifier System Attacks

As a key component of Internet infrastructure, the Identifier System is a high-value attack target. Some types of attacks against the Identifier System are theoretical – for example, protocol design vulnerabilities that have been reported but not (yet) actively exploited. Other types of attacks against the Identifier System are oft-attempted but widely deflected by mitigation practices – for example, malformed Domain Name System (DNS)[2] queries and responses that are easily recognized and discarded by firewalls. In between these two extremes lie types of attacks which are either more sophisticated or more difficult to easily deflect than our examples, and these in particular pose noteworthy risk to Identifier System SSR. Several of these types of attacks are summarized below, as a starting point for developing further guidance about how to more effectively mitigate them on a broader scale.

Route Insertion Attacks

Some of the prominent Internet routing attacks are not attacks against the routing system itself, but rather use the routing system to make ("announce"" or "advertise") IP addresses usable for spam or other malicious activities known and reachable across the Internet. Such *insertion attacks* allow

[2] This document uses the acronym "DNS" as the common abbreviation for "Domain Name System." However, in the Affirmation of Commitments, "DNS" is also defined as "domain names; Internet protocol addresses and autonomous system numbers; protocol port and parameter numbers. ICANN coordinates these identifiers at the overall level, consistent with its mission."

malicious traffic to originate from these IP addresses and be delivered to targets or victims across the Internet, or to redirect legitimate traffic to malicious destinations.

Route insertion attacks commonly exploit the Border Gateway Protocol (RFC 1771), a routing protocol that is used to exchange network reachability information among autonomous systems (AS), defined as one or more IP networks that operate under a single routing policy (e.g., an ISP, a broadband access provider, or a hosting or cloud operator). In a route insertion attack, attackers use "someone else's IP addresses". They take control over blocks of IP addresses, identify an exploitable network element such as a misconfigured or compromised router and utilize this system to announce these IP blocks as reachable from this *injection* point. These actors may do so in an attempt to conduct a variety of attacks from hosts to which they've assigned addresses from these IP blocks, and then effectively "disappear" after the attacks.

The article Internet Address Hijacking, Spoofing and Squatting Attacks[3] offers a taxonomy of the many kinds of route insertion attacks, including autonomous system number (ASN) squatting attacks, ASN spoofing attacks, and IP prefix exploitation attack. Broader awareness and adoption of practices to deter route insertion attacks, such as measures to protect address registrant accounts against misuse, strong authentication for account access, and filtering/blocking of fraudulent route advertisements, could enhance the stability and resiliency of the Identifier System.

DNS Denial of Service Attacks

Most DNS queries are processed by a relatively small number of software products, installed on off-the-shelf hardware running general-purpose operating systems. As a result, zero-day or unpatched known vulnerabilities in these widely-deployed platforms can be used to attack a very large number of DNS resolvers and authoritative name servers. In particular, Denial of Service (DoS) attacks can exploit these vulnerabilities to degrade or prevent domain name resolution; if not mitigated, such attacks could significantly impact the DNS portion of the Identifier System. For example, BIND, the Berkeley Internet Name Domain, a product of the Internet Systems Consortium[4], is open source software deployed on production DNS servers that has been used as a reference implementation and the base for other DNS products. Recently, all versions of BIND version 9 were found to be vulnerable to a new DoS exploit that can use a single packet to crash those servers. Not long after this flaw was announced, real-world attacks began, triggering a worldwide rush to patch BIND version 9 servers.

There are many such DNS DoS attacks, ranging from volumetric attacks designed to overwhelm DNS platform assets or the infrastructure supporting those assets, to malformed, long, or specially-crafted packets sent to crash DNS servers. Practices to deter DoS attacks, such as proactive/unattended patching, OS/application hardening, infrastructure overprovisioning, and routine vulnerability assessment, are unevenly adopted and too often ignored. Broader adoption of these and other anti-DoS practices could enhance the stability and resiliency of the Identifier System.

---

3 "Internet Address Hijacking, Spoofing and Squatting Attacks," The Security Skeptic, June 2011.

4 https://www.isc.org/downloads/bind/

DNS DDoS Attacks

Distributed Denial of Service (DDoS) attacks against the DNS also have the potential to severely degrade or prevent domain name resolution. For example, a DNS reflection and amplification DDoS attack exploits a vulnerable recursive DNS resolver by using thousands of hosts to send a very high volume of DNS queries. Those queries get reflected through a vulnerable recursor to a victim IP address. The victim's server is overwhelmed with DNS traffic that exhausts the resources of that victim and possibly the exploited recursor itself. According to the Open DNS Resolver Project, of 32 million DNS resolvers now responding to queries in some fashion, 28 million have been found to be at risk of exploitation by this type of attack.

This is just one kind of DNS DDoS attack; other kinds of DDoS attacks could target DNS root servers, stub resolvers, authoritative domain name servers, and the network resources used to deliver DNS queries and responses. However, all of these kinds of DDoS attacks originate from a large, distributed set of source IPs, using the DNS or other stateless protocol to consume Identifier System resources to degrade or prevent domain name resolution. As such, a common set of mitigation practices could conceivably be applied more broadly to reduce the overall risk and potential impact of DDoS attacks. For example, source IP address filters, response rate limits, and anti-spoofing measures can all be used to narrow the possible magnitude of DNS DDoS attacks. Broader adoption of these and other anti-DDoS practices could enhance the stability and resiliency of the Identifier System.

Web Services Attacks impacting Identifier Resources

Registry and, in the case of DNS, Registrar operations systems represent another critical asset within the Identifier System. Most of these systems have web services components, including resource registration services and control panels for updating or transferring domain names or address blocks. These web applications are largely Internet-facing customer access portals connected to back-end database(s) where other critical assets such as resource registrations and databases are stored.

Unfortunately, cross-site scripting, SQL injection, and cross-site request forgery attacks which exploit web application vulnerabilities are extremely popular among attackers and succeed at an alarming rate. According to Verizon's 2015 Data Breach Investigations Report, over 9.4% of all confirmed data breaches resulted from web application attacks, with organized crime being the most frequent threat actor utilizing this type of attack. When vulnerabilities in web applications associated with Registry and Registrar systems are successfully exploited, attackers can gain administrative control over registration, other databases, and associated personal identifying information.

Practices such as proactive patching, application hardening, vulnerability assessment, and multi-factor authentication (see Account Compromises) have been proven to reduce the frequency and severity of data breaches resulting from web application attacks. Broader adoption of web application attack mitigation practices could raise the security posture of affected Registrar and Registry systems, enhancing the overall stability of the Identifier System.

Resource Registration Account Compromises

Compromised resource registration accounts can be used to hijack registered resources such as domain names and address blocks, register fraudulent subdomains or sub-allocations, alter records to render legitimate sites unreachable or redirect victims to malicious sites, and facilitate subsequent attacks that exploit the affected infrastructure. For example, a Registrar's Domain Name Account management system which supports only weak authentication leaves accounts vulnerable to password cracking, which (absent other mitigation measures) can lead to malicious DNS use as described in SAC 044: A Registrant's Guide to Protecting Domain Name Registration Accounts. According to Verizon's 2015 Data Breach Investigations Report, over 95% of web application attacks now involve harvesting credentials from customer devices, then logging into web applications with them.

Although vulnerabilities and the methods used to exploit them vary, this type of attack exploits Domain Name Registration accounts as a vector to harm the registrant and its employees, customers, or other users of affected domain names, reducing the stability of the DNS portion of the Identifier System. Failure to adequately protect these accounts is due in part to diversity and poor understanding of readily-available account security practices. A uniform baseline of security measures across DNS registrars could help to elevate the default security posture of DNS registration accounts. For example, multi-factor authentication and DNS change alerts are proven ways to deter and quickly detect registration account compromise. Broader adoption of these and other Domain Name Registration Account security measures could elevate the security of these critical resources, enhancing the stability of the Identifier System.

Identifier Hijacking via Social Engineering

Domain Name , ASN, or IP address block hijacking occurs when an unauthorized entity gains control over a domain name, ASN, or IP address block registered to someone else. While these identifiers can be hijacked as the result of registration account compromise (above), hijacking may also result from social engineering attacks. For example, attackers may use public information obtained from the Internet or the domain name's WHOIS record to send phishing messages to the registrant or their designated representatives, tricking those victims into ceding control over or even transferring the domain name. Following a successful phish, the attacker can then modify or add name server or resource records to the domain name's authoritative zone file, with impacts ranging from financial losses to brand damage. Such attacks have long plagued the Identifier System. For example, several costly social engineering hijack attacks were described a decade ago in the SSAC's 2005 Report on Domain Name Hijacking. Despite advances in underlying technology, these kinds of socially-engineered identifier hijacks continue to grow, ranging from isolated high-profile incidents to new hijack methods that place thousands of domain names at risk.

In all such attacks, the vulnerable asset being exploited is a person responsible for the unique identifier's registration – a designated administrative or technical contact for the identifier, or an authorized representative of the registrant. By exploiting human vulnerabilities through social engineering, attackers can gain control of the hijacked identifier, effectively bypassing any other registration account security measures. And, while identifier hijacking is certainly not a new kind of attack, exploit methods

continue to evolve along with Internet technology – for example, growth in phishing via mobile notifications and social media. Raising awareness about these evolving social engineering attacks could help to reduce the risk and potential impact of domain name, ASN, and IP address block hijacking, thereby enhancing the stability of the Identifier System.

DNS Zone File Attacks

A DNS Zone File is a critical asset within the Identifier System, containing resource records that map one or more domain names to IP addresses and other resources within that domain. However, Zone Files can play a role in many attacks.

For example, attackers may perform reconnaissance on a domain name by transferring or walking the domain's Zone File to obtain name space and topology information. Attacks that gather published Zone File data to identify attack targets may be precursors to subsequent attacks. For example, naming conventions used within a Zone File may reveal something about which hosts have valuable data (e.g., accounting.example.com, creditcardDB.example.com). It is therefore important that DNS administrators understand that Zone File data is effectively public and should be treated as such.

However, attacks that modify a domain name's authoritative Zone File have clear potential to impact Identifier System SSR by adding, deleting or changing resource records, with consequences similar to domain name hijacking. For example, recent successful attacks have started with a criminal probing a web server on a shared host that also runs an authoritative DNS server. By exploiting vulnerabilities in that web server software, the criminal can gain root privileges and then alter the Zone File.

All of these attacks focus on Zone Files as critical assets. Practices such as DNSSEC can protect the integrity of Zone Files, but are not yet universally adopted. Broader awareness of Zone File vulnerabilities and adoption of common DNS Zone File attack mitigation practices[5] could elevate the security of individual Zone Files and domain names, enhancing the stability of the DNS portion of the Identifier System.

Parallel Root Name System Risks

Geographically-diverse root name server instances, operated by the root name server operators as identified by the IANA "root hints,"[6] are almost universally used in the DNS resolution process to refer resolvers to the appropriate authoritative name server for resolution of domain names. However, some organizations operate their own root within their internal network(s), deploying privately-operated root name server instances to resolve queries for a specified set of internal top-level domains (TLDs). This relatively common technique, sometimes referred to as "split horizon" DNS, creates a "parallel" root which can peacefully co-exist with the Internet's public namespace except for situations in which there is an overlap in TLDs. For example, if an organization creates regional 2-letter character domains, e.g., "ap" for their Asia Pacific presence, "na" for North America, "as" for Asia, etc., the users of that organization's DNS may have trouble connecting to sites in Namibia (ISO-3166 code "NA") and American

[5] For example, practices described in "Top Level Domain Incident Response: A "Recovery" Checklist" developed by OCTO-SSR in cooperation with subject matter experts from Mark Monitor, Network Startup Resource Center (NSRC), and Farsight, Inc.
[6] https://www.iana.org/domains/root/files

Samoa (ISO-3166 code "AS"). However, since this parallel root is constrained to be within an internal network, the risks are relatively low and easily mitigated by the network's operator.

A more troublesome example of a parallel root is one in which the namespace is not constrained to internal networks. In cases such as this, organizations offer an alternative namespace to the public, allowing people to obtain top-level domains independently of the Internet's public namespace. For example, in addition to the Internet's root, at least one parallel root organization resolves .biz, resulting in .biz domain names at times resolving to different IP addresses, with direct impact on user experience, security, and functionality.

As illustrated in SAC009, the potential for conflicting roots poses risk to the stability of the Identifier System. However, parallel roots are privately-operated and thus procedures for adding new names to those roots vary. For example, domain names under .bit must be resolved by a parallel root operated by Namecoin. To resolve .bit domain names, users must redirect all or some DNS requests to a Namecoin-aware name server such as OpenNIC, or install a Namecoin-aware browser plug-in such as FreeSpeechMe. Such procedures introduce complexity and vulnerabilities commonly incurred when redirecting DNS through any third party. Broader understanding of these kinds of potential vulnerabilities and related mitigation practices could enhance the stability of the Identifier System.

DNS and Surveillance Attacks

Internet privacy concerns are widely addressed in part through use of encrypted communication protocols such as SSL/TLS and IPsec and encrypted messaging protocols such as S/MIME. However, even when user data is encrypted hop by hop, end to end, or at rest, concerns remain about vulnerabilities that may be exploited for purposes of Internet surveillance. For example, much can be learned by monitoring the clear text protocols used to find communication partners and establish secure connections between them. One such protocol is DNS.

DNS traffic can potentially be intercepted at various points throughout the Internet, including on the host initiating a DNS query, on the name server that resolves that query, or anywhere along the path traversed by the query and response. In addition to intercepting DNS queries and responses, man-in-the-middle attacks on encrypted channels become possible if a host can be directed via DHCP or other means to use a fraudulent name server which maps some or all names to malicious destinations.

In 2016, the IETF standardized a method for users to communicate with the DNS resolver that they use. The new protocol[7] tunnels DNS over TLS to prevent sending that traffic in clear text. Although this standard is quite new and thus thinly implemented, increased use will protect more DNS traffic. At the same time, DNS resolver operators who are concerned about minimizing the exposure of users' requests have started to use a process call "QNAME minimization."[8] Although the DNS queries from those resolvers still are clear text, queries to higher-level authoritative servers (such as to the root and to second-level domains like ".com") contain only the minimum information needed instead of the full

7 https://tools.ietf.org/wg/dprive/
8 https://tools.ietf.org/html/draft-ietf-dnsop-qname-minimisation-09

domain name. Broader understanding of new protocols and practices such as these to minimize DNS traffic surveillance attacks by untrusted parties could enhance the security of the Identifier System.

DNS Misuse as Covert Channel

As a critical component of Internet infrastructure, the DNS protocol is generally permitted to pass freely through Internet firewalls, routers, and other security systems. Although some of these systems inspect the IP traffic sent through port 53 (see Surveillance), many do not – creating a vulnerability in network defenses that can be exploited to create cover channels. For example, a host infected with BotNet software may communicate with its Command and Control server over port 53 to receive further instructions, download additional malware, or exfiltrate data stolen from the infected host. Messages sent by trojan software such as Feederbot and Morto have even been carefully crafted to look like DNS TXT lookups in order to evade detection.

DNS can be misused in other ways – for example, to tunnel user traffic out of a public hotspot that might otherwise require payment for Internet access. In all such attacks, the host exploits an egress rule that permits outbound traffic on port 53. Broader understanding of this vulnerability and practices that can be applied to firewalls, routers, and other security systems to deter exploitation could help to enhance the security of the Identifier System.

Status and Evolution

The above, non-exhaustive list of attacks against the Identifier System has been put forth for consideration within ICANN and by Identifier System security experts throughout the Community. These attacks are intended to serve as a starting point for answering the following questions:

1. **Are there actual high-impact attacks or emerging high-risk vulnerabilities that should be added?**
2. **Which of these types of attacks have already been (or should be) addressed elsewhere?**
3. **Of the remainder, which should ICANN attempt to address over the next 12-24 months?**
4. **Is there a body of mitigation knowledge (work) that can be leveraged to improve Identifier System SSR?**

These answers are to be recorded as updates to this ever-evolving Methodology document and used as the basis for OCTO-SSR, in consultation with IETF, the Regional Internet Registries and community experts, to draft and then socialize an on-going set of Tech Notes, one for each top priority type of attack. Over time, this section will serve as an index of available Tech Notes, their current status, and the target audience for each.

Annex A: Background

This document is part of ICANN's effort to contribute to enhanced Identifier System SSR by working with the Community to identify and increase awareness of related attacks and to promote broader adoption of attack mitigation practices. This effort also addresses Recommendation #12 of the Security, Stability & Resiliency (SSR) Review Team (SSR-RT), which commits ICANN to work with the Community to identify SSR-related "best practices" and support the implementation of such practices.

Recommendation #12 was made in the SSR-RT's Final Report, which considered ICANN's FY12 SSR Methodology, an annual plan detailing ICANN priorities for the next fiscal year in promoting a healthy, stable and resilient unique Identifier System. The SSR-RT also considered the ICANN Strategic Plan – July 2012-June-2015 which enumerates goals and activities related to ICANN's areas of influence.

In its Final Report, the SSR-RT supported ICANN's programs and objectives to collaborate and coordinate with TLD operators and others to enhance DNS security, stability and resiliency, stating that "*by engaging with Registries and Registrars, ICANN can play an important role in promoting the development and implementation of SSR-related best practices.*"

To reach this recommendation, the SSR-RT considered many collaborative Community work projects involving DNS technical and operational issues. In addition to activities undertaken by ICANN itself, the SSR-RT examined issues that are subject to ICANN's influence, but not its control. For these SSR-related activities, the SSR-RT found that ICANN sometimes relies on specific parties to take the lead – for example, GNSO policy development activities in connection with the Registrar Accreditation Agreement (RAA). In other cases, ICANN follows the lead of Community Working Groups responsible for developing solutions to address Internet security – for example, coordination of emergency response teams (DNS-CERT). The SSR-RT supported this collaborative approach but expressed concern over "the slow progress of some Community efforts to address registration abuse and other SSR-related issues."

Similarly, the SSR-RT found that ICANN recognized the importance of implementing effective mechanisms to prevent, identify and respond to malicious abuse of the unique Identifier System. However, the SSR-RT expressed concern that SSR-related issues had not always been addressed in a timely manner, further observing that "It can also be difficult to achieve consensus on practices that may create costs and operational burdens on contracted parties."

These concerns led the SSR-RT to formulate Recommendation #12, urging ICANN to harness its Community relationships to contribute to informed decision-making and further enhance Identifier System SSR.