Overall Purpose Name: **Criminal Investigation or DNS Abuse Mitigation**

Definition: The broad category of criminal investigation or DNS abuse mitigation covers all use of an RDS to support criminal and other investigations, abuse prevention, security incident response, and other activities to protect people, systems, and networks from detrimental activities. These activities range from criminal activities like extortion, phishing, and provision of child abuse materials to abusive activities including denial-of-service attacks, spam, and harassment.

Users: The primary actors in these scenarios include law enforcement, cybersecurity professionals, IT administrators, and automated protection systems.

Tasks: Using information from the RDS, these actors will, depending upon the circumstances: contact domain owners and/or the entities that provide services for an affected domain to mitigate problems, gather evidence, or notify them of compromises; expand investigations and associations to fully understand the scope of an abuse issue; identify Internet infrastructure involved with detrimental activities, inform protection systems to take protective actions; and request suspension of domain names.

Data Elements used generally for criminal investigation or DNS Abuse Mitigation

Domain WHOIS record

- Registrant (Name, Address, email address). Use - identification, information and intelligence gathering etc
- Creation date, renewal date, last updated date, expiry date. Use - is it recently registered (maybe a DGA etc) ; Is it a long time registered / historic domain - if so perform a WHOIS history check on it to look at identifying the registrant...before they changed over to a privacy/proxy registrar to hide their details
- Registrar. Use - further enquiries with an disclosure authority/court order.
- NS records (Nameserver - used to direct the traffic of your website to a specific web server at a web host.) Use - what other domains point to this NS - this could provide you with a whole host of intelligence on other domains controlled by the same person/organisation.

Network WHOIS record

Abuse contact (for further enquiries - disclosure authorities)

CIDR space of network provider (use - if they own for example a /24 - try some passive DNS to see what other domains point to these IPv4 addresses - may give you more intelligence on malicious domains associated to a rogue server etc)

DNS records

MX record. Use - which network provider provides mail for the domain ?

Bad WHOIS data of value

A false domain name, registrant, address, email

Uses - bad/false/stolen/incomplete domain whois data may give an investigation a new lead in terms of intel gathering, linked accounts showing the same false data through a registrant search of the WHOIS record for similarly registered domains.

Background: This category encompasses a broad set of use cases for querying different data elements associated with one or more domain names contained in the RDS. The data queried will depend upon the nature of the detrimental activity in question, the goals of the person or entity making the queries, and the stage of an investigation or incident at the time. For some tasks a deep set of data may be needed for a particular domain or small set of domains, while for others, a very small amount of data may be needed per domain, but for a very large number of domains. Given this wide variety of use cases, data access, and contact needs, this document will present several example use cases grouped into logical categories. At the end of the document you will find a general overview of the data elements used.

The first category distinction is between individual investigators or small teams looking into discrete incidents making ad-hoc data requests for single or small sets of domains, and automated processes that may query for information about thousands to millions of domains in a very short time period.

The second axis of differentiation of use cases differentiates between the various stages of an investigation/mitigation/protection effort. First, the use of RDS data to determine the likely involvement of a domain name as one registered and controlled exclusively to perform the detrimental activity or one that has been compromised and used against the wishes of the domain registrant. Second, the use of RDS data in order to facilitate contacting of the appropriate parties as part of an investigation, a mitigation of abusive activities, or notification of potential problems. Finally, a set of use cases for using RDS data to understand the scale and scope of domains and Internet infrastructure being used in conjunction with a particular attack or campaign.

A separate category of uses of RDS data in criminal investigation and DNS abuse mitigation encompasses those cases where the domain name itself isn't necessarily the focus of the investigation or abuse concern. Domain names can be tangentially involved in other cases ranging from online abuses to real-world crimes. Access to information in the RDS may further such investigations when it is determined for example that a potential miscreant may have registered domain names for his or her personal use or a domain name may have been associated with evidentiary e-mails. In such cases, understanding who may have registered or been involved with supporting a domain may lead to further evidence leads.

Note:

This table is largely based on current practices and currently available data unless otherwise noted.

One capability discussed in this document that exists outside of the current whois system (with some exceptions) is the concept of "reverse whois". Such services exist and provide high value

information to inform many use cases/purposes in this category. Some form of reverse whois has been proposed for a future RDS and/or the accommodation of such services within the RDS framework. This work explains how those services are used and useful today without commenting upon their appropriateness now or in the future.

Using new data elements like "social media contact" or other proposed future RDS capabilities is not explored here. Where such data elements were to be collected in the future, these use cases would need to be updated to reflect their applicable use. For example, a preferred contact method that is a unique identifier is a good candidate for pivoting on investigations to expand their scope, and of course, if a registrar prefers to receive an SMS message to report abuse, processes that involve registrar contacts would incorporate that data element.

Table of purposes and associated use cases

Section 1: Individual Investigations

Subsection 1A: Manual determination of domain status (malicious/compromised)

1A-1 Purpose Name: Manually determine if the domain of a website used for an attack is compromised or registered maliciously

Definition: Determine if domain of website used for an attack (e.g. phishing, exploit, scam, etc.) is compromised, being abused, or registered maliciously. Websites used for online abuse fall into one of three categories: compromised - hacked or exploited where unauthorized content is added to the site, abused - some hosting service is misused by bad actor, or registered maliciously by the miscreant directly. Determining this status is critical for informing the next steps of an investigation or mitigation.

Tasks:

1) Obtain a potentially abusive domain name from a report of some sort - typically an abuse report.

2) Verify abusive activity is occurring

3) Query RDS data for information about the domain including age, registrar, registrant/admin/tech/abuse contacts

4) Use known techniques and infrastructure of both "good" and "bad" actors to determine likelihood of a malicious registration. Prime factors include age of domain, nameservers of domain, registrar of domain, reseller of domain, privacy service employed (particularly for phishing), known registrant (good/bad), other known contacts. Note that for a malicious domain, the data for registrant will be false, but if it matches other known "bogus" data, this is

a positive attribution factor. One data element that will be constant between malicious registrations is the registrant e-mail address which provides control over a domain in many circumstances. A domain "handle" is also useful for such matches.

Users: Security researcher, LE researcher, automated tools used by researcher

Data: Creation date, nameservers, registrar, reseller, full available contact information for registrant and any other contacts (e-mail and contact handle most useful)

Subsection 1B: Contacting appropriate parties/taking action

1B-1 Purpose Name: Notify parties responsible for a domain name that has had its website compromised

Definition: Internet security personnel, law enforcement and other investigators working on criminal or abuse issues need to inform those parties responsible for a domain name of malicious activities and potential exposure of PII or other information related to a compromised website. These notices will lead to mitigation of the compromise and gathering of evidence related to the malicious activities related to the compromised website.

Tasks:

1)   Query RDS for information about contacts for the domain name that has had its website compromised. These contacts would typically include the technical contact (often the web host), registrant (owner), and admin contact (up-to-date responsible party) for the domain.

2)   Evaluate the information returned to determine if actual contact data is included, or if it is bogus or privacy protected to prioritize contacts towards actual people or well-defined roles.

3)   Send e-mail or make phone calls to attempt to contact responsible parties in real time.

4)   Obtain information from contactable contacts to reach actors who can mitigate issues and/or provide evidence/information

5)   Work with actors who can take action to mitigate issues and deliver information/evidence.

Users: Law enforcement personnel, security researchers, CERT teams, first responders

Data: Contact information for technical, registrant, and admin contacts including name, phone number, and e-mail address to facilitate notifications and communications.

1B-2 Purpose Name: Notify parties responsible for a domain name that has had its domain management account compromised

Definition: Internet security personnel, law enforcement and other investigators working on criminal or abuse issues need to inform those parties responsible for a domain name of malicious activities and potential exposure of PII or other information related to a take-over of a domain name management account. These notices will lead to mitigation of the account compromise and gathering of evidence related to the malicious activities related to the compromised domain. In these circumstances, currently listed contact records may be false due to miscreant capability to modify these entries.

Tasks:

1)   Query RDS for information about contacts for the domain name that has been taken over via the domain management account. These contacts would typically include the technical contact (often the web host), registrant (owner), and admin contact (up-to-date responsible party) for the domain. However, these may not be reliable since the miscreant may have changed them. Registrar abuse contact becomes primary contact to use if this is likely.

2)   Evaluate the information returned to determine if actual contact data is included, or if it is bogus or privacy protected to prioritize contacts towards actual people or well-defined roles.

3)   Determine if contact information is still reliable. Historical or certified contact information of some sort would be useful in this scenario, if it existed.

4)   Send e-mail or make phone calls to attempt to contact responsible parties in real time.

a.   At a minimum, make sure registrar is aware of the compromised account and takes action to ensure miscreant cannot re-compromise.

5)   Obtain information from contactable contacts to reach actors who can mitigate issues and/or provide evidence/information

6)   Work with actors who can take action to mitigate issues and deliver information/evidence.

Users: Law enforcement personnel, security researchers, CERT teams, first responders

Data: Registrar abuse contact primary contact point. Others, if not changed as part of incident, include contact information for technical, registrant, and admin contacts including name, phone number, and e-mail address to facilitate notifications and communications.

1B-3 Purpose Name: Notify registrar and/or reseller of malicious domain name registration for mitigation and/or evidence gathering

Definition: After an investigator has positively identified a malicious domain registration, they may take further action depending upon their goals. In most cases the goal will be to get the

domain suspended or removed from the DNS and prevented from being re-activated by the miscreant. In some cases, the investigator will be looking to obtain information from the registrar, or reseller if applicable, about how the domain was registered including items like payment details, IP address of any online order, or data behind a proxy registration.

Tasks:

1)   Determine that a domain name is malicious via an investigation.

2)   Access RDS data to obtain official abuse contact information for a registrar or information on involved reseller. Access other resources like a registrar website to get e-mail/phone for abuse desk, customer support or other relevant departments.

3)   Evaluate the information returned to determine if actual contact data is included, or if it is bogus or privacy protected to prioritize contacts towards actual people or well-defined roles.

4)   Use phone, e-mail or whatever listed contact information is available to establish communication with registrar and/or reseller if applicable.

5)   Request actions including suspension, deletion or transfer of malicious domain name, and/or further information about the actor who registered the malicious domain. In particular, evidence/information sought out will be about how the domain was registered including items like payment details, IP address of any online order, or data behind a proxy registration.

6)   Registrar or reseller takes some sort of action to the request.

7)   Escalate to registry if registrar unresponsive or refuses to take action.

Users: Law enforcement personnel, security researchers, CERT teams, first responders

Data: Registrar abuse contact primary contact point. If not available, whatever contact information is available for the registrar. If reseller involved, contact information for the reseller. E-mail address and/or phone number for these contacts is necessary. If escalation to registry is required, abuse contact information for the registry.

Subsection 1C: Scoping infrastructure involved in issue

1C-1 Purpose Name: Expand knowledge from one known malicious domain to other domains potentially part of the same issue

Definition: Investigate key attributes of a known malicious domain to find others that may be part of the same or related incidents. Since criminals/abusers often re-use common elements for registering malicious domains, once a domain has been identified as being malicious,

researchers can take key unique elements from that domain and search for other domains sharing those elements. Such unique elements often include unique nameservers, unique contact data – particularly registrant and/or admin contact e-mail or to a lesser extent, phone number. This purpose requires the existence of some sort of "reverse whois" capability where an RDS, cached database, or third party collection of RDS data can be queried on an attribute and return a list of all domains sharing that attribute. Such domains tend to cluster on less unique elements such as creation date, registrar, and reseller, but using these data elements requires other meta data for correlation. Lists of suspect domains may then be probed to see if they exhibit the same illegal/abusive behavior.

Tasks:

1)   Obtain a positively identified malicious domain from prior investigation, trusted data feed, or other high-confidence source.

2)   Query RDS for key attributes that allow for "pivoting" to other potentially related domains. Such information will include nameservers, full contact data for registrant, admin, and in some cases tech contacts (particularly unique elements like contact handle, e-mail address and phone number), and other more loosely associable elements like registrar, reseller, and creation date.

3)   Determine veracity of supplied information as an informative element. Accurate data is not necessary in this step since repeated bogus data is a strong indicator of associated abuse.

4)   Build list of domains based on reverse whois lookups on unique elements.

5)   Examine list of domains for the same abusive behavior or indicators that they may have been or will be used in a similar matter

6)   Use the gathered data to again pivot on unique elements found within the newly discovered domains.

7)   Use an investigatory tool like a relationship visualization system to "cluster" domains that have paths of relationships to look for patterns, key elements, and potential clues as to how the miscreant may create new domains in the future.

8)   Use this information to inform other processes like mitigation or criminal investigations.

Users: Law enforcement personnel, security researchers, CERT teams, first responders

Data: Full available contact information for registrant and any other contacts (e-mail, phone number, and contact handle most useful), nameservers, registrar, reseller, creation date

1C-2 Purpose Name: Examine all domains sharing one or more key elements tied to abuse to determine if a larger issue exists

Definition: Individual data elements that appear in RDS data are often identified in investigations as being associated with abuse. For example, these could include a phone number used by a serial scammer, an e-mail address used in prior security incidents or malicious registrations, a PO box of a known criminal operation. Using such elements and reverse whois queries, an investigator can find domain names likely to be associated with malicious activities and examine them for abusive behavior and/or monitor them for future activities.

1)   Obtain a positively identified malicious or suspicious data point that represents a unique attribute for a domain registration from an investigation, high-confidence data feed, or direct observation.

2)   Access RDS or other system to build list of domains based on reverse whois lookups on unique elements.

3)   Examine list of domains for the same abusive behavior or indicators that they may have been or will be used in a similar matter

4)   Use the gathered data to again pivot on unique elements found within the newly discovered domains.

5)   Use an investigatory tool like a relationship visualization system to "cluster" domains that have paths of relationships to look for patterns, key elements, and potential clues as to how the miscreant may create new domains in the future.

Use this information to inform other processes like mitigation or criminal investigations.Users: Law enforcement personnel, security researchers, CERT teams, first responders

Data: Full available contact information for registrant and any other contacts (e-mail, phone number, and contact handle most useful), nameservers, registrar, reseller, creation date

1C-3 Purpose Name: Find potentially compromised domains related to an existing hijacking or domain shadowing incident

Definition: When miscreants take over domain names in hijacking or domain shadowing attacks, they often will take over entire groups of domain names due to vulnerabilities in registrar systems, systemic use of weak or compromised passwords, or getting ahold of a domain portfolio. When one such domain is identified, understanding the scale of the compromise and potential compromise is important to determine as soon as possible to

mitigate a larger issue. Miscreants will often change elements of domains involved in such attacks to common infrastructure such as nameservers or update admin or registrant contacts to include new information to allow the miscreant control of the affected domains. Understanding the new and original information for such domains allows investigators and first responders the opportunity to mitigate all domains and not just the reported one(s). This purpose also is greatly enhanced with the existence of "who was" type services that provide historical information on how a domain name was previously listed in an RDS. Such systems are currently run by third parties, and they have been proposed for a future RDS or to be allowed under the framework of a future RDS but those ideas have yet to be explored by the working group.

Tasks:

1)   Identify a hijacked or shadowed domain name.

2)   Access the RDS to determine current attributes for the affected domain including nameservers, admin and registrant contact details, registrar, registrar abuse contact, and modification date. Admin contact information is vital in case of a hijacking since transfers are usually controlled via the admin e-mail address.

3)   Access historical records for the affected domain to obtain the same information. Particularly important is prior registrar in case of a hijacking that involved a domain name transfer.

4)   Build list of domains based on reverse whois lookups on key unique elements that have been modified for the domain.

5)   Examine list of domains for the same abusive behavior or indicators that they may have been or will be affected in a similar matter

6)   Enter notification/mitigation phase with the affected registrar(s) and legitimate registrant.

Users: Security researchers, CERT teams, first responders, registrar abuse teams

Data: nameservers, full admin and registrant contact details, registrar, registrar abuse contact, and modification date.

Section 2: Automated Processes

Subsection 2A: Automated determination of domain status (malicious/compromised)

2A-1 Purpose Name: Automatically determine if a domain used for an attack is registered maliciously

Definition: Determine if domain used for an attack (e.g. phishing, exploit, spam, scam, etc.) is compromised, being abused, or registered maliciously. Domains used for online abuse fall into one of three categories: compromised - hacked or exploited where unauthorized content is added to the site, abused - some hosting service is misused by bad actor, or registered maliciously by the miscreant directly. Determining this status is critical for informing the next steps in preventing a risky connection.

Notes: These activities are not well served by the current asynchronous and slow response that the current whois system provides. Thus while some networks incorporate such queries, most use pre-positioned reputation data aggregated by third-party specialists to make the described decisions.

Tasks:

1) Obtain a potentially abusive domain name from a live stream of data – e.g. e-mail server connection requests, outbound network requests at the DNS resolver, pre-fetching activities of browsers on a corporate network, requests to a WAF (Web Application Firewall).

2) Query RDS data for information about the domain including age, nameservers, and registrar. Other data such as contact data would be desirable, but response time is usually too slow to reliably use.

4) Use known techniques and infrastructure of both "good" and "bad" actors to determine likelihood of a malicious registration. Prime factors include age of domain, nameservers of domain, and registrar of domain.

Users: Automated security processes/systems including but not limited to e-mail servers, firewalls, DNS resolvers, and WAF's.

Data: Creation date, nameservers, registrar

2A-2 Purpose Name: Automatically create reputation score for domain names

Definition: Calculate a reputation score for a domain name that represents a scalar value or set of values for the relative risk for engaging in different communications for a domain. Publish these scores for subscribers who will make decisions on operations like e-mail delivery, network connections, web browsing requests, and other data exchanges. This includes the use of fixed algorithms on subscriber networks as well as updating scoring metrics by the reputation

provider. Scoring processes usually take into account abuse reports and white lists to better classify domains or domain elements like nameservers or registrar.

Tasks:

1)   Deploy a domain reputation scoring algorithm based on prior work and investigations into various forms of malicious and benign domain names. Modern systems use machine-learning for a majority of these tasks.

2)   Receive a new domain name to score or a previously seen one to update from one of many processes. Inputs in this step include parsing new domains out of daily zone files, observations in passive DNS sensor networks, subscriber requests based on observed connection attempts.

3)   Access RDS to obtain key elements required by the scoring algorithm. Data needed will typically be those attributes that tend to cluster for abusive domain names including nameservers, registrar, creation date, registrant contact info (particularly e-mail, phone, and name), other contact information.

4)   Obtain other meta data from other sources to improve scoring including abuse reports, other reputation lists, lists of known DGA (domain generation algorithm) domains, white lists, and known benign infrastructure.

5)   Score domain name and publish score in file, feed, or as a query response.

6)   Update algorithm, algorithm parameters, and/or domain element knowledgebase using new information obtained in processing recent domain scores.

Users: Automated processes and researchers working for organizations that provide reputation scores for domain names

Data: Creation date, expiration date, nameservers, registrar, contact information for technical, registrant, and admin contacts including name, phone number, and e-mail address

Subsection 2B: Automated notifications of abuse

Definition: Some actors who process large volumes of abuse (e.g. spam, botnets) provide automated reporting to affected entities. One of those processes is providing automated reports to registrars and registries of maliciously registered domain names. Their goal is usually to get the domains they have identified suspended or removed from the DNS and prevented from being re-activated by the miscreants who registered them.

Tasks:

1)	Determine that a domain name is malicious via a standardized investigatory process and automation..

2)	Access RDS data to obtain official abuse contact information for a registrar or information on involved reseller. Access other resources like a registrar website or abuse reporting API.

3)	Access RDS to obtain relevant information for domains being reported to include with report so registrar/registry/reseller can locate other domains with the same attributes and potentially take action.

4)	Use e-mail, abuse reporting API or whatever listed contact information is available to establish communication with registrar and/or reseller if applicable.

5)	Report relevant RDS information that may indicate miscreant activity and request actions including suspension, deletion or transfer of malicious domain name.

6)	Registrar or reseller takes some sort of action on the request.

7)	Escalate to registry if registrar unresponsive or refuses to take action.

Users: Security researchers, CERT teams, first responders

Data: Registrar abuse contact primary contact point. If not available, whatever contact information is available for the registrar like an abuse reporting form or API. If reseller involved, contact information for the reseller. E-mail address and/or phone number for these contacts is necessary. If escalation to registry is required, similar abuse contact information for the registry. Reported information will typically include registrant name, e-mail, admin e-mail, phone numbers for registrant and admin contacts, and nameservers.

Subsection 2C: Automatically scoping infrastructure involved in issue

2C-1 Purpose Name: Automatically expand knowledge from one or more known malicious domains to other domains potentially part of the same issue

Definition: Automatically investigate key attributes of a known malicious domain to find others that may be part of the same or related incidents. This is the same purpose as 1C-1 except at scale, so the data elements involved will typically be more narrowly constrained (typically nameservers and key e-mail addresses) to allow for fully automated processing. Since criminals/abusers often re-use common elements for registering malicious domains, once a domain has been identified as being malicious, researchers can configure automated processes to take key unique elements from that domain and search for other domains sharing those elements. Such unique elements often include unique nameservers, unique contact data –

particularly registrant and/or admin contact e-mail or to a lesser extent, phone number. This purpose requires the existence of some sort of "reverse whois" capability where an RDS, cached database, or third party collection of RDS data can be queried on an attribute and return a list of all domains sharing that attribute. Such domains tend to cluster on less unique elements such as creation date, registrar, and reseller, but using these data elements requires other meta data for correlation. Lists of suspect domains may then be automatically probed to see if they exhibit the same illegal/abusive behavior.

Tasks:

1)   Obtain a positively identified malicious domain from prior investigation, trusted data feed, or other high-confidence source.

2)   Query RDS for key attributes that allow for "pivoting" to other potentially related domains. Such information will include nameservers, full contact data for registrant, admin, and in some cases tech contacts (particularly unique elements like contact handle, e-mail address and phone number), and other more loosely associable elements like registrar, reseller, and creation date.

3)   Determine veracity of supplied information as an informative element. Accurate data is not necessary in this step since repeated bogus data is a strong indicator of associated abuse.

4)   Build list of domains based on reverse whois lookups on unique elements.

5)   Examine list of domains for the same abusive behavior or indicators that they may have been or will be used in a similar matter

6)   Use the gathered data to again pivot on unique elements found within the newly discovered domains.

7)   Use an investigatory tool like a relationship visualization system to "cluster" domains that have paths of relationships to look for patterns, key elements, and potential clues as to how the miscreant may create new domains in the future.

8)   Use this information to inform other processes like mitigation or criminal investigations.

Users: Automated processes configured by security researchers and incident response teams.

Data: Full available contact information for registrant and any other contacts (e-mail, phone number, and contact handle most useful), nameservers, registrar, reseller, creation date

Section 3: Supporting investigations into crimes/abuse not directly involving a domain name

Subsection 3A: Determining domain ownership or involvement with operating a domain name

3A-1 Purpose Name: Determining domain ownership or involvement with operating a domain name tied to real-world criminal/abuse activities

Definition: Domain names can be tangentially involved in other cases ranging from online abuses to real-world crimes. Access to information in the RDS may further such investigations by providing ownership or operational connections to a domain name that has come up as evidence or a potential lead in a case focused on behavior not primarily tied to that domain. For example, e-mails may indicate that a miscreant used a domain name to commit fraud or some other act, or an e-mail address tied to a threat like a botnet was used to register one or more domain names.

1)   Determine that a domain name is potentially indirectly involved with in a crime or incident via an investigation.

2)   Access RDS data to obtain full registrant and potentially admin contact data.

3)   Evaluate the information returned to determine if actual contact data is included, or if it is bogus or privacy protected to determine if the domain registration is providing actual data.

4)   Use real data to significantly supplement tangential investigation. Add bogus or suspect data to investigatory file.

5)   If applicable, pivot off data found in this use case to expand to other potentially related domains.

Users: Law enforcement personnel, security researchers, CERT teams, first responders

Data: Full available contact information for registrant and any other contacts (e-mail, phone number, and contact handle most useful), nameservers, registrar, reseller, creation date