Report Prepared for ICANN

Review of The Domain Abuse Activity Reporting system (DAAR) and methodology

By: Marcus J. Ranum, consultant

Executive Summary

I have performed a rigorous review of the DAAR system and its methodology, with particular attention paid to DAAR's methodology for combining existing data from generic Top-Level Domain (gTLD) registries and ICANN accredited registrars with 3rd party reputation data. A primary concern is that the DAAR system accurately and credibly measures security threats in the composite data and can be viewed with confidence. Since DAAR is a historical roll-up of data sources – DNS metadata plus 3rd party classification data. Its accuracy will be as good as those data sources; however, those sources are accepted with high confidence in the industry already and they are what make the internet DNS function properly. The composite data in DAAR is a useful publication: quality historical reference for the community, researchers, and registrars. The DAAR collection and compilation methodology is well-explained and documented.

A second question is whether the DAAR studies a meaningful set of data that are worth studying: whether spam domains are a security threat for the internet user community. They are. In addition to placing a massive unnecessary, valueless load on the DNS infrastructure, spam domains have important roles in spam infrastructure deployments which are popular vehicles for attacks. DAAR documentation explains this but misses the opportunity to explain that the massive streams of unimportant spam-based attacks provide camouflage for more targeted and significant attacks. ICANN's Security and Stability Advisory Committee (SSAC) is correct to identify spam as an important part of the registrar ecosystem that is worth studying. The body of this report will contain extended discussion of the spam problem.

To conclude: the DAAR system is a straightforward implementation of a good idea. It will assist policy-makers and researchers who want to study Internet and domain abuse from a daily, periodic, or historical perspective.

The remainder of this report will look at sub-topics in more detail and will conclude with a few relatively minor technical suggestions and the reasoning behind them.

Scope and Purpose of This Review

I have been tasked with validating the DAAR operating principles, system design, documentation, and outputs, to assess whether they meet industry/community norms for reliability. My review process has been to examine the system through its operator's interface, existing documentation, FAQs, and industry and research whitepapers about the spam ecosystem. Since the operator's interface is not intended to be made publicly accessible, I did not examine its implementation for software security issues; my analysis was strictly on the data, data sources, compilation methods, and description of the system and its operation.

I am a computer security practitioner with over 25 years consulting experience in system design and implementation; I am holder of an ISSA lifetime achievement award and am an ISSA fellow. The process I used for this review is typical for this sort of project; I have performed similar implementation and design reviews for a number of significant security products, using this methodology.

Method

The method of DAAR is sound: the sources, collection method, combination and counting methods are well thought-out. Combining multiple data-sources generally brings about a question of normalization – how to calibrate the increments on the multiple sources – but DAAR avoids that problem by not attempting to resolve it. The block lists that DAAR is matching against may potentially have subtle differences in their scoring but because DAAR is reporting data at a large scale, any subtle differences are literally going to be lost in the noise. If someone were to decide to resolve a single entry in DAAR they would only be able to learn how the various RBL providers have scored it. Any complaints about the RBL scoring are not ICANN's problem; they are the RBL providers', or the registrars. From a methodological standpoint, DAAR relies on the RBL's methodologies – that's good from the perspective of both knowledge-base management and scoping the problem domain: if there are any complaints, they will be about the RBLs, not DAAR. The RBL maintainers have strong incentives to ensure their classifications are as accurate as they can be and have a credible system in place for correction/redress; there will be occasional mis-classifications, but they will be corrected and the corrections will be transparent to the DAAR system. Likewise, the zone information and TLD data collected from the registrars will be as accurate as it can be; it is meta-data about how the DNS functions and if there are any problems with it the specific registrars will fix them transparently to the DAAR system.

We are satisfied that there are incentives to ensure accuracy in the data upon which DAAR is built. Because the data are drawn from dynamic systems, there will be changes and corrections over time, but the data is currently accurate enough to make the DNS and anti-spam commercial products function effectively; it will continue to be at least that accurate in the future. The DAAR design would allow for additional RBLs to be added should new ones become available or dropped should one be determined to be inaccurate or out of maintenance.

The DAAR project FAQs collection does a good job of explaining the relationship between DAAR's data sources and its method in constructing the data. From the description of the DAAR system, an experienced system builder could independently build their own identically functioning version if they wanted to. It is critically important to frame any questions about DAAR correctly, since it is inevitable that some organizations will examine it critically: are they saying nasty things about us? Will people draw conclusions from this data that we don't like? The way DAAR is described is good; it is very neutral, informative, and non-threatening.

Data Correlation Threats

Whenever data is presented based on combining it with other data, there is a potential threat that someone will "subtract" the additional data and be able to extract the original data set. DAAR avoids this problem by design: none of the data that is used is secret, and the way it is compiled can be performed by anyone. Therefore, there is no secret data within DAAR that risks being exposed.

There is always a potential when data is exposed, for that data to be combined with some other data in a way that is problematic. The recent experience with Strava making their heat-maps available, resulting in classified US Government facilities being disclosed, is one example of that sort of incipient public relations disaster. In the case of the DAAR data, there does not appear to be any such problem – mostly because the data is publicly collectible (with more effort) already.

Quality of Correlational Feeds

The RBL feeds that DAAR correlates against are the best that are available. There is an epistemological challenge that can be raised against the RBL feeds, but not the results of combining the feeds with registry or registrar data – if there are charges of inaccuracy, they are deflected over to the maintainers and producers of the black-lists. The DAAR FAQ document and DAAR whitepaper provide a good explanation of the relationship between the RBL feeds. Whois and DNS zone data; there is no significant threat other than "there may be some complaints" from organizations that have not succeeded in complaining to the RBL maintainers.

The RBL maintainers constantly receive complaints that they are inaccurate. Sometimes, they are – for example my personal e-mail server was on the RBL for a while because I inherited an IP address that had been used for spam-sending. I had to use a redress system to get my address removed from the list. Some might consider that a complaint that the list was wrong, but I consider it "how reputation lists work." The consumer of an RBL's data does so out of a desire to protect their systems by blocking spam – even in a case like my domain, the RBL is working correctly by providing the customer's spam filtering system advisory information about the history and reputation of any server that contacts them. Inside the commercial spam-blocking products that use RBLs, the RBLs are only part of an overall scoring process that is used to decide if any given message gets through. The people who complain about reputation lists are generally the people who consistently are being scored as abusers – because, they're abusers. It is normal, in other words, for there to be some complaining about RBLs. If spammers aren't complaining about them, they aren't working properly. None of this is ICANN's problem – it's a dialogue between the spammers and the RBL maintainers.

The developers of DAAR were wise to avoid taking on the problem of vetting or weighting the RBLs, which would amount to making a judgement that one RBL was more or less accurate or had a better redress process. Avoiding the "redress problem" the way that DAAR does – by leaving it a problem for the RBL maintainers – is also a good strategy. Since the RBLs are widely used operationally and in products, they are accurate enough. Therefore, demographic data correlated with the RBL data is also accurate enough. Since the data is changing day over day, "accurate" means "correctly collected and compiled."

Does DAAR "Name and Shame"?

Spamming domains are named explicitly in the RBL datasets; that is how the RBLs work. The DNS functions to map names to addresses and vice-versa. Therefore, it will always be possible to map an RBL entry to a registrar or an address – DAAR is not "naming and shaming" anyone when the system makes the names of registrars or domains available, that is information that is already in the RBL, DAAR is just gathering it into one place and presenting it as a unified dataset.

Spam Is A Security Threat

The DAAR FAQ and whitepaper both make the claim that spam is a security threat. This is apparent in the reactions of the academic computer security community, the security products market, and in the end customers' practices. For example, I first taught day-long classes on spam-blocking for the USENIX conference in Atlanta in 2004. Before then, and ever since, there has been an academic interest in spam blocking techniques, including some really great fundamental academic research in statistical methods for spam classification. Paul Graham's 2002 paper "*A Plan for Spam*" is one example – Graham broke open the entire field of using machine-learning techniques for spam classification and blocking. Meanwhile, on the commercial front, leading spam-blocker technology provider Barracuda Networks was acquired for $1.6 billion in 2017 and Google acquired spam-blocking mail service Postini for $625 million in 2007. Researchers and businesses are placing that level of value on spam-blocking: it is a *problem*.

Someone claiming that "spam is not a security threat" would be making a suspiciously self-serving argument given that a great deal of threat activity is predicated on exactly that principle. For example, Cisco's Talos "threat detection center" specifically treats spam as a *threat* not an annoyance, as does Google's Gmail service.

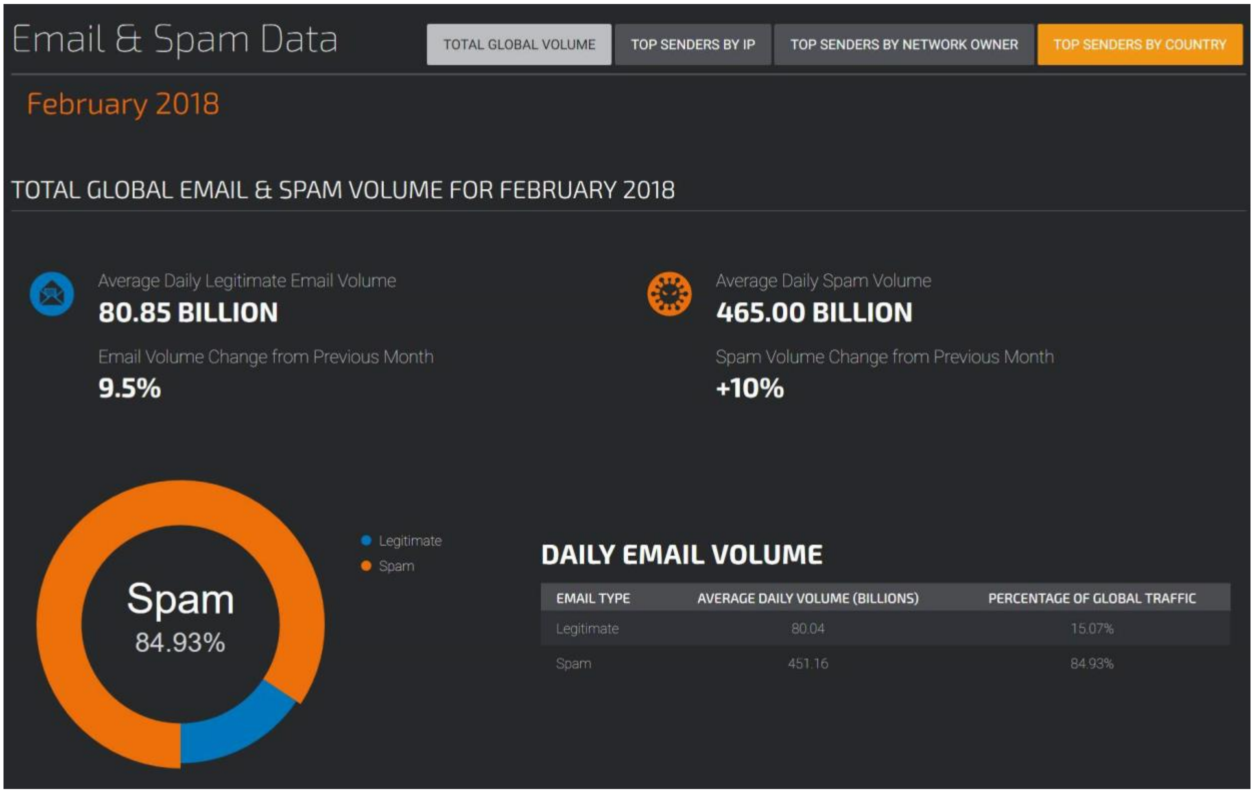

Email & Spam Data

February 2018

TOTAL GLOBAL EMAIL & SPAM VOLUME FOR FEBRUARY 2018

Average Daily Legitimate Email Volume
80.85 BILLION

Email Volume Change from Previous Month
9.5%

Average Daily Spam Volume
465.00 BILLION

Spam Volume Change from Previous Month
+10%

DAILY EMAIL VOLUME

EMAIL TYPE	AVERAGE DAILY VOLUME (BILLIONS)	PERCENTAGE OF GLOBAL TRAFFIC
Legitimate	80.04	15.07%
Spam	451.16	84.93%

(Cisco Talos)

In the 465 billion spam messages that Cisco is measuring right now, there is a percentage of spam messages that carry PDF attachments with exploit/attack code in them. In fact, the Russian hacking attacks during the 2016 US election relied on both targeted email attacks and spam email attacks – over 100,000 pieces of spam carrying malware-laden document attachments, in a single campaign of attacks.

Securelist reports:

- 40% of spam emails were less than 2 KB in size.
- The most common malware family found in mail traffic was **Trojan-Downloader.JS.Sload**
- The Anti-Phishing system was triggered 246,231,645 times.

Malware downloaders such as Trojan-Downloader.JS.Sload are pieces of javascript code, embedded in spam emails, which direct the reader's mail client to retrieve attack code from a staging site; the attack code then takes over the readers' computer and installs a backdoor. The 246 million times Securelist detected Trojan-Downloader.JS.Sload are just one of thousands of spam-based attack vectors.

In the security community, we track spam-based malware in order to defeat it. But we also keep historical statistics about the amount of spam-based malware we deal with, in order to understand and communicate about the magnitude and importance of the problem. DAAR is another historical data-view into the roles that domain names have in spam attacks and email abuse; it will be valuable to the security community.

There are extreme "free speech" arguments that have been made in favor of spam (if one is arguing *against* spam-blocking, one is arguing *for* spam) and currently those arguments are working their way through the court system in Georgia, Maryland, Washington, and Virginia. The question is not "is spam annoying?" but rather "does spam constitute a security threat?" To that, the security community gives an unequivocal "yes" on two grounds: 1) spam is used as a vehicle for mass attacks and 2) massive amounts of spam provide concealing noise for targeted attacks.

The first generation of spam-based security threats depended on software bugs in email clients. For example, Microsoft Outlook used to automatically fetch and display images that were linked into a message. An attacker could craft an image that exploited a software bug in a specific version of Outlook and would then launch a spam campaign delivering that message to millions of email users. A significant percentage of those email users would be running vulnerable versions of Outlook and would have their system taken over. The only defenses against this sort of attack are constantly updating one’s browser and email client and using a spam-blocking system at the edge of the network or in front of the email client.

The current state of the art for a spam campaign is to combine advertisement (of real or nonexistent products) *and* an attack into a single message. The recipient is presented with an advertisement for some junk, with a link to “buy it now!” and a link to “unsubscribe from these emails.” If they click on “buy it now” they get vectored to a shopping cart somewhere; if they click on “unsubscribe” their browser is directed to a drive-by malware dropper.

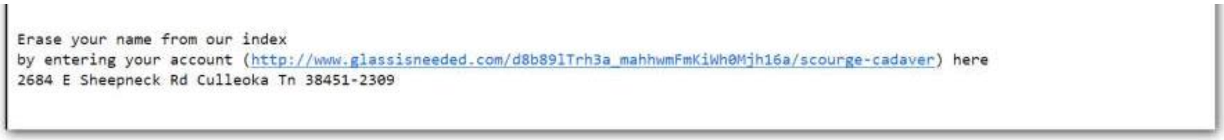

```
Erase your name from our index
by entering your account (http://www.glassisneeded.com/d8b891Trh3a_mahhwmFmKiWh0Mjh16a/scourge-cadaver) here
2684 E Sheepneck Rd Culleoka Tn 38451-2309
```

The constructed URL in this spam is intended to defeat URL filtering; “scourge-cadaver” seems to be an inauspicious randomly-generated name. In the malware ecosystem randomly-generated names and URLs are a common attempt to defeat network security systems that attempt to blacklist domains or URLs that are known to host malware. A vulnerable browser attempting to retrieve that URL gets back a load of JavaScript that attempts to generate a deceptive pop-up message for a “Windows anti-malware cleaner.” Drive-by malware droppers are particularly pernicious: they identify the user’s browser, map the version against a list of vulnerabilities and exploits, launch the exploit, and install persistent backdoors in the user’s system. A new trend in drive-by malware is to direct a user to a URL with code that performs bitcoin mining; it’s not very efficient but the spammer may make some money, and they don’t care because they’re not paying for the electricity. At the point where someone is attempting to direct a user to a URL that runs attack code, spam is a security threat.

Many corporations deploy spam-blockers specifically to prevent their users from encountering dangerous URLs and drive-by malware. Spam domains are often used to host such sites, as well as fake copies of real sites that are used to harvest users’ logins and passwords – there is an attack technique called “typosquatting” in which a hacker registers an abuse domain name based on a common typo, e.g.: “microsft.com” or they spam out messages from “0racle.com” where the ‘O’ is a zero. Some couple the spam-blockers to URL filters in firewalls or web proxies; those systems also use RBLs along with other detection techniques. It’s the same data that DAAR uses, and it’s matched and evaluated the same way; the difference is that the spam blocker takes a policy-based action on the RBL match, whereas DAAR keeps a historical record that a particular address was on a particular list at a particular time.

The fact that the spammers are automatically generating attack URLs and registering many abuse domains in response to the blocking technique is proof that they know they are trying to bypass the user’s endpoint security; they know that they are attacking systems.

A subtle point regarding spam as a security problem is that the great cascade of harmless spam serves as cover for directed e-mail attacks. Because organizations filter spam, the users that receive targeted e-mail (phishing) attacks are *more* likely to open a message that is crafted to not look like spam. Even though targeted e-mail attacks are smaller than a full-blown spam campaign, they are often 10,000-100,000 messages at a time. For example, a typical phishing attack might look like a message from a bank’s customer service department, directing users to log in and update their account; the destination domain name might be an abuse domain like *mybigbank-support-desk.com* where the real domain is *mybigbank.com*. Not all phishing attacks employ abuse domains, but the more successful ones do.

Phishing attacks are undeniably part of the spam/anti-spam ecosystem and end-user and corporate responses to them bear that out.

In this discussion, I have been focusing exclusively on email spam. There are other forms of spam and they are all, also, security threats. Forum comment-spam can be used to jump viewers to drive-by malware downloads. Social media spam can also be used to direct viewers to malware-dropper sites – this was part of the arsenal of techniques used by the Russian hackers to influence the 2016 election – users would be directed to fake sites, their accounts taken over, and then used to vote up selected articles. Text message/SMS spam is also used to direct users to download apps that contain backdoors, or to websites that contain malware droppers. All of the varieties of spam have been weaponized by hackers and spammers and responding against these attacks has been a tremendous expense to the software and security industry: every email client, every browser, every discussion forum – all must be carefully coded and constantly updated to prevent them from being used as attack vectors by spammers. The maintenance cost of keeping software resistant against data-driven attacks runs into the hundreds of millions of dollars, annually. If one adds in the cost to the end user, of re-imaging compromised systems, and constantly updating their software, spam-delivered attacks represent a massive global "tax" on software that runs into the billions of dollars, annually.

Spam is a security problem. Measuring the employment of abuse domains in spam attacks is worthwhile research.

The Value of Historical Data

There are many aspects of the growth of the internet that were not measured and recorded when it started. In some cases, great effort has been made to re-derive that information (e.g: rates of emails sent, public documents posted, domains registered, website demographics). Historical views are useful in many ways: they can be used to assess long-term trends, support academic research on internet usage and growth, or to identify the growth of problem areas. Such historical views are of particular interest with regard to Internet hacking, spamming, denial of service, and other forms of high-impact abuse including domain name abuse. Historical data is the only way we can answer the question “is this particular problem getting better or worse?” DAAR will be valuable as it offers a view into abuse activity over time.

Collection, Attribution and Mapping

DAAR’s approach for collecting its data and associating top-level domains with registrars is based on how the DNS itself works; it’s as good as it *can* be.

There is no other source of meta-data for the top-level domains that can be cross-checked against, and the DNS would not function properly if the meta-data was not correct enough to use. There is a potential that some entries might be change between polling intervals, but that is acceptable because the DAAR system is intended to provide a summary view, not real-time query results – if someone wants real-time query results about a particular domain, they would use the DNS’ normal operations instead of DAAR.

The DAAR FAQ and whitepaper clearly position DAAR as a source for research, historical, and exploratory summary data; it doesn’t attempt to supercede existing operational tools so there is no likelihood that anyone will attempt to use it in that way.

The RBLs

The list of RBLs DAAR uses represents the combined state of the art in reputation scoring; there is no need to search for additional ones (though, if a newer, better, RBL emerges, there is no reason why it could not be added). Many commercial products successfully depend on one or more of the chosen RBLs, so it doesn’t seem likely anyone will challenge their accuracy *as they are reflected in DAAR*.

Someone may disagree with any particular RBL's scoring for a particular domain but that is a problem between them and the RBL maintainer.

DAAR is wise to have avoided trying to apply any additional weighting or mapping on top of the RBLs; that would open the system to accusations of favoritism.

ICANN may wish to publish a contact for new RBLs that may wish to be mapped into DAAR. That contact would probably occur through back-channels anyway - but having an official contact-point removes the possibility of complaint. The Open Data Initiative covers this point sufficiently; any organization that wants to provide information or access information has a pathway for engaging in that discussion.

Access To The Dashboard

The DAAR system is presented as a data collection and aggregation, but it will come to be seen by its users in terms of its user interface. Usually with datasets like DAAR, you find a mix of users that want to graphically explore, and users that will want to download a subset and analyze a specific item that they discovered during the exploration process. To build that explore->analyze->explore cycle, users will want the pretty user interface (and DAAR's UI *is* pretty!) because it's easier, and that will bring in issues of system query-load, whether users will just sit there mashing the button on the same query over and over, data-scraping, etc. Therefore, we recommend two things:

Continue with the plan to limit access to the DAAR administrative interface.

Since the DAAR data is updated daily, produce a static dashboard of graphical outputs to basic queries – cache the results then auto-generate an index to them, with thumbnails, etc., using a static naming-scheme in case someone wishes to deep-link to a given chart for any given time. Since the updates are daily, and the images would be moderate-sized, it would not result in a large accumulation of data over time: 365 entries/year times number of charts produced. Giving people access to "cooked" data in the form of an image would greatly simplify the question of whether anyone needs credentials to use the administrative interface; giving an outsider access to the administrative interface brings in software security questions such as whether it might be possible to perform an injection attack against the interface, etc. The best way to solve that problem is to avoid it entirely.

Edge Cases

DAAR will be used to perform comparisons, so edge-cases are going to be a problem; let us suppose that there is a registrar that only registers a single domain and the RBLs mark that domain as not abusive – the registrar appears to be "100% abuse-free" when, in fact, it is irrelevant. In statistics, edge-cases are usually dealt with using cut-offs: you simply do not consider small values that will cause large percentage movements. DAAR should not report on registrars that have fewer than 1,000 domains; if that cut-off turns out to be too low it can be adjusted later. Cut-offs should be documented as: "(not reported: registrars supporting fewer than 1,000 domains)" That would not prevent someone from registering domains to get above the cut-off; it would just cost them some money and there isn't much point in attempting to game the system in that manner. It's not worth the effort. Suppose a registrar decided to start a marketing campaign to "come use us, we're 100% abuse free!"; that would be a self-correcting problem if the marketing campaign attracted a large number of spammers. In other words, the incentive structures do not appear to encourage attempts to game the system in this manner.

Someone might decide to troll the system – trying to manipulate a percentage simply for the sake of being able to say, "I manipulated the percentage!" but, so what? The DAAR FAQ makes it very clear where the value of the data is, and isn't, and no system can control against what amount to pointless claims.

Summary of Findings

DAAR is an accurate, useful system that will offer a valuable historic view into a significant security problem that affects the internet. It uses valid methods to predictably compile data from sources that are accurate within its daily cycle. The data compiled by DAAR is maintained by third parties and ICANN does not have any responsibility for redressing any complaints about the underlying data. There are no sensitive data items in its compilation that risk being disclosed. The documentation (white paper, FAQ) is thorough and provides a detailed and accurate explanation of the system's rationale and method, and the system as implemented works as described in the documentation. Spam is a significant research topic and is an area of strong commercial activity because spam is a security threat; it is not a mere annoyance.

Appendix A: Specific Recommendations

These are details that don't fit anywhere else, or don't merit deeper discussion.

- **Cached images** – consider producing automatically-generated cacheable images of top-level query results.
- **Set Scale on X-axis** – When charts are generated, the X-axis defaults to auto-scaling. That is reasonable, except that it results in deceptive appearances if someone is visually comparing the images, i.e: if you take two screenshots of the statistics of two different registrars, it could be that one is scaled 1-50, and the other is scaled 1-5,000,000 but the histogram plot looks the same. Auto-scaling charts amounts to noise amplification. There is nothing wrong with it, because sometimes (if you're looking at a single registrar) you want the data to be auto-scaled to range, so the chart is not a little line at the bottom. If you're cutting screenshots of two registrars and comparing them visually, however, you actually need them to be at the same scale. Consider adding a fixed scale option – just let the end-user input the X axis top value.
- **Log and assess queries in use** – build a log of the queries that are generally made through the user interface; that is a good candidate-set for "most common queries" and can be used to build static images.
- **"Curation" vs "Compilation"** In the DAAR paper, the term "curation" is used to describe the process of combining RBL scores with ICANN's registrar data. In the FAQ the term "compile" is used. Choosing the correct term for the process is important, since it's the main thing people are likely to complain about regarding DAAR. Since "curation" is an aesthetic process in which an art-work is maintained and presented for a specific effect, it seems to be better avoided, as terminology. "Compile" (an automated process) implies a predictable combination of two things and is more technically accurate, less likely to raise questions, and easier to understand.

Appendix B: Brief Bio of Marcus Ranum

Marcus Ranum has been active in the UNIX networking and security community since 1989, when he took over responsibility for operating one of Digital Equipment Corporation's three internet gateways and turned it into the first successful Internet firewall product, the DEC SEAL.

Since then he has held every job position in computer security start-ups: coder, project manager, product manager, marketing strategist, sales support, customer support, VP of engineering, Chief Technology Officer, Founder, CEO, and member of the Board of Directors. In 1997 and 1998 he was on the Initial Public Offering team of V-ONE corporation and was on the board of directors of a publicly traded company, Network One. Interspersed between all of that, Marcus has written books, technical papers, blogs, taught, and consulted. He is a popular conference speaker and has keynoted most computer security conferences at one time or another. As a consultant, Marcus has worked for national governments, FORTUNE 10 firms, and the grandmother up the street whose computer got malware.

His consulting projects have ranged from helping design national-level security policies and framing legislation to writing and debugging software device drivers. He is equally comfortable working at every level in a corporation's organization. He has done a considerable amount of work in the security audit arena, as well as performing discrete services for incident response, including on several major incidents that received national-level attention.

Most recently, Marcus has been speaking and consulting for IANS, as well as consulting on security/technology strategy for a large media company in Los Angeles. Prior to that, he was a non-testifying expert in a firewall technology patent litigation and was important in his client's winning a settlement of over $100 million. Prior to that he served as Chief Security Officer for Tenable Network Security, Inc., where he was responsible for internal security operations, product lifecycle design, and end-customer deployment doctrines – as well as designing the end-user training program that is still used to teach how Nessus and Security Center are to be used.

Marcus is an ISSA Fellow and holds an ISSA Lifetime Achievement Award. He writes a bi-monthly column for SearchSecurity.