~~SECRET//ORCON,NOFORN~~

UNITED STATES

FOREIGN INTELLIGENCE SURVEILLANCE COURT

WASHINGTON, D.C.

[redacted]

NOTICE OF FILING OF GOVERNMENT'S SUPPLEMENT TO ITS SUBMISSIONS OF JUNE 1st AND JUNE 28TH, 2011

THE UNITED STATES OF AMERICA, through the undersigned Department of Justice attorney, respectfully submits the attached supplement in further support of the

~~SECRET//ORCON,NOFORN~~

Classified by: Tashina Gauhar, Deputy Assistant Attorney General, NSD, DOJ
Reason: 1.4(c)
Declassify on: ~~16~~ August 2036

arguments set forth in submissions of June 1st and June 28th, 2011, concerning the above-referenced matters. This supplement explains the methodology behind and sets forth the results of a manual review by the National Security Agency (NSA) of a statistically representative sample of the nature and scope of the Internet communications acquired through NSA's FISA Amendments Act Section 702 upstream collection during a six-month period. The Government respectfully submits that the data provided herein supplements and supports the Government's Responses to the Court's Briefing Order of May 9th, 2011, and supplemental questions of June 17, 2011, and will further assist the Court in concluding that the certifications and procedures submitted in the above-referenced matters satisfy the requirements of the Act and are consistent with the Fourth Amendment to the Constitution of the United States. ~~(S//OC,NF)~~

Given the complex nature of the information provided in this supplement, the United States is prepared to provide any additional information the Court believes would aid it in reviewing these matters. The Government may also seek to supplement and/or clarify the information provided herein as appropriate during any hearing that the Court may hold in the above-captioned matters. ~~(S//OC,NF)~~

Respectfully submitted,

[redacted]

National Security Division
United States Department of Justice

(U//FOUO) NSA Characterization of Upstream Data: Process and Results

I. (U) Introduction

~~(TS//SI//NF)~~ This report explains the methodology behind and provides the results of a manual review of a statistically representative sample of Internet communications acquired through NSA's FISA Amendments Act (hereinafter "FAA") section 702 upstream collection during a six-month period.[1] The purpose of this review was to assemble data to assist the Court in understanding the nature and scope of the communications acquired through NSA's upstream collection. The data assembled consisted of:

- The volume of transactions containing single, discrete communications to, from, or about a selector used by a person targeted in accordance with NSA's section 702 targeting procedures (hereinafter "tasked selector") versus transactions containing multiple communications (hereinafter "Multi-communication Transactions" or "MCT") not all of which may be to, from, or about a tasked selector;[2]
- The types of discrete communications contained within MCTs [redacted]; and

[1] ~~(TS//SI//NF)~~ Additionally, as described on pages 8-9 of the Government's June 1, 2011 Response to the Court's Briefing Order of May 9, 2011, NSA conducted two tests of FAA 702 upstream collection in May 2011 using information from NSA's technical databases in an attempt to determine the likelihood of collecting an Internet transaction between a user in the United States and [redacted] NSA also attempted to further determine the extent to which those tests might be statistically representative of NSA's 702 upstream collection and repeated these tests in July 2011 using alternative data sets. Because of the technical limitations for automatically identifying transactions containing multiple communications, NSA assesses that the results of these tests are not comparable to each other or with the results of the separate manual analysis discussed herein. Furthermore, for the same reason of technical limitation, the results do not express as high a degree of granularity and accuracy as the manual analysis discussed herein, which took more than one month of careful review by experienced analysts to complete. None of the results discussed herein and in the Government's June 1 Response, however, are inconsistent.

[2] ~~(TS//SI//NF)~~ As described on pages 27-28 of the Government's June 1, 2011 Response to the Court's Briefing Order of May 9, 2011, NSA's inability to separate out individual pieces of information from Internet communications acquired by NSA's upstream collection systems does not extend to all forms of transactions. NSA has developed the capability to [redacted] identify transactions which [redacted] and, in certain other limited instances, transactions where an "active user" (as described more fully below) is a tasked selector. Based on a test of this capability from July 16th-29th 2011, NSA estimates that approximately only [redacted] of NSA's current upstream collection under FAA section 702 could be identified through [redacted] processes as communications to, from or about NSA's tasked selector. As reflected by the results of this manual review, this figure is significantly under-representative of the total proportion of NSA's upstream collection assessed to be communications to, from or about a tasked selector.

- The volume of MCTs that NSA assesses contain a wholly domestic communication not to, from, or about a tasked selector.[3]

II. (U) How the Statistically Representative Sample Was Assembled

~~(TS//SI//NF)~~ NSA assembled the sample of communications acquired through its upstream collection by first identifying all Internet communications acquired under section 702 – i.e., both from NSA upstream collection and collection from Internet service providers either by or with the assistance of the Federal Bureau of Investigation (hereinafter "PRISM collection") -- during a six-month period from January 1st through June 30th, 2011, and present within [redacted] as of July 14, 2011. As of that date, 140,974,921 Internet communications were present within [redacted] Of these, 127,718,854 (or approximately 91%) were acquired from PRISM collection, and 13,256,067 (or approximately 9%) were acquired through NSA's upstream collection.[6]

~~(TS//SI//NF)~~ The approximately 13.25 million Internet communications acquired through NSA's upstream collection (hereinafter "transactions") were then "shuffled" by NSA statisticians to ensure a random sample (i.e., any sample drawn would be statistically representative of the total 13.25 million transactions). NSA statisticians estimated that a manual review of a sample of approximately 50,000 of these randomized transactions would enable characterization of all 13.25 million transactions with a statistically high level of confidence and precision.[7]

III. (U) How the Manual Review Was Conducted and the Results of the Review

~~(TS//SI//NF)~~ Under the leadership of NSA's Deputy Director, an experienced interdisciplinary team consisting of experienced intelligence analysts, attorneys from NSA's Office of General Counsel, representatives from NSA's Office of the Director of Compliance, NSA statisticians, representatives from NSA's Network Analysis Center, and representatives from NSA's Office of Oversight and Compliance was assembled to conduct the review described herein and compile this report. A team of experienced NSA

---

[3] ~~(TS//SI//NF)~~ This aspect of the review required analysts to perform intensive analysis on discrete communications which did not contain the target's selector within MCTs, to determine if the sender and all intended recipients of those discrete communications were located in the United States. Such in-depth analysis is not typically conducted by analysts in their daily foreign intelligence analysis. Instead, an analyst would tend to focus his or her attention on those discrete communications within the MCT that are to, from, or about their assigned target, and would only perform a deeper inspection of those communications to confirm they were not wholly domestic if they were in-fact pertinent to the analyst's evaluation of foreign intelligence information and therefore worth further analysis for potential use.

[4] ~~(TS//SI//NF)~~ [redacted]

[5] ~~(TS//SI//NF)~~ This figure does not include Internet communications that were acquired during this six-month period but were purged prior to July 14, 2011.

[6] ~~(TS//SI//NF)~~ *See* Figure A of Appendix A, attached hereto.

[7] ~~(TS//SI//NF)~~ Details for the basis for NSA's statistical assertions are set forth in Appendix B, attached hereto.

intelligence analysts was assigned to conduct a manual review of the transactions. Ultimately, that team of NSA intelligence analysts collectively reviewed a total of 50,440 individual transactions.

~~(TS//SI//NF)~~ In order to ensure consistency among the analysts in their review, before beginning the manual review, the team members were trained to recognize MCTs and how to characterize the discrete communications contained within them. The team members were given training materials created specifically for this effort, which included screenshots depicting typical examples of the types of transactions acquired through NSA's upstream collection. NSA's Office of General Counsel, Office of Oversight and Compliance, and Office of the Director for Compliance reviewed all training materials and provided guidance throughout the manual review.

~~(TS//SI//NF)~~ For quality assurance, some transactions (approximately 10 out of every 5,000) underwent independent reviews by more than one analyst. In addition, the team lead performed spot reviews of transactions that had already undergone review (approximately 1 out of every 100). The team lead also personally reviewed any transaction that team members were unable to immediately characterize as clearly being a discrete communication or an MCT, as well as any MCT identified as potentially concerning a person located in the United States. Both the quality assurance overlap and the reviews performed by the team lead revealed no discrepancies among how analysts characterized any of the transactions subjected to these overlapping reviews.

~~(TS//SI//NF)~~ In conducting the manual review, NSA analysts took the following steps and made the following findings:

1. **Determined if the transaction was a single, discrete communication or an MCT.**[8] If the transaction was determined to be a single, discrete communication, no further analysis was done. Transactions determined to be MCTs were further analyzed, as described below.

   - Of the 50,440 transactions reviewed, 45,359 (approximately 90%) were determined to be single, discrete communications. The remaining 5,081 transactions (approximately 10%) were determined to be MCTs.[9]

2. **Characterized the discrete communications within the 5,081 MCTs as being** [redacted]

   - Of the 5,081 MCTs reviewed, [redacted]

[8] ~~(TS//SI//NF)~~ For any objects that the initial reviewer was uncertain about how to characterize (e.g., if the transaction contained data requiring further processing to render it intelligible to the analyst), the team lead performed a second review. As a result, each of 50,440 transactions reviewed were able to be characterized as being either a single, discrete communication or an MCT.

[9] ~~(TS//SI//NF)~~ *See* Figure B of Appendix A.

[10] ~~(TS//SI//NF)~~ [redacted]

[redacted][12]

3. **Determined whether the 5,081 MCTs contained any discrete communications as to which the sender and all intended recipients were located in the United States.** As discussed in more detail below, in many cases NSA analysts were able to make these determinations based on the location of the "active user" of the MCT.[13] In other cases, NSA had to rely on content analysis because the MCT did not contain technical information sufficient to identify the active user or to determine the active user's location. There were, however, instances where the MCT did not contain sufficient technical information or content for NSA to assess whether the MCT contained any wholly domestic communications.

   - Of the 5,081 MCTs, 713 (approximately 14%) had a tasked selector as the active user [redacted]. No further analysis of these MCTs was done to determine whether they contained wholly domestic communications. That is because the user of the tasked selector, who by operation of the NSA targeting procedures is a person reasonably believed to be located outside the United States, would be either the sender or an intended recipient of each of the discrete communications contained within the MCT.[14] Accordingly, all of the discrete communications within those MCTs would have at least one communicant reasonably believed to be located outside the United States (i.e., the target) and thus would not be wholly domestic.

   - Of the 5,081 MCTs, 2,668 (approximately 52%) had an active user that was not a tasked selector but was nonetheless an electronic communications account/address/identifier

[11] ~~(TS//SI//NF)~~ *See* Figure C of Appendix A.

[12] ~~(TS//SI//NF)~~ [redacted]

[13] ~~(TS//SI//NF)~~ When NSA acquires an Internet transaction between an individual using an electronic communications account/address/identifier and his/her service provider, that individual is the "active user" for that transaction. Such transactions can have, at most, one "active user."

[14] ~~(TS//SI//NF)~~ In this context, a communication to or from the target includes communications to or from the tasked selector itself (e.g., an e-mail sent to a tasked e-mail account), as well as communications where the tasked selector appears in other communications attributable to the target [redacted] *See In re DNI/AG Certification* [redacted] Docket No. 702(i)-08-01, Mem. Op. at 17 n.14 (USFISC Sept. 4, 2008).

reasonably believed to be used by a person located outside the United States.[15] No further analysis of these MCTs was done to determine whether they contained wholly domestic communications. That is because the foreign-based active user would be either a sender or intended recipient of each of the discrete communications within the transaction. Accordingly, all of the discrete communications within those MCTs would have at least one communicant reasonably believed to be located outside the United States (i.e., the foreign-based active user) and thus would not be wholly domestic.

- Of the 5,081 MCTs, 8 (approximately 0.16%) contained an electronic communication account/address/identifier of a non-targeted active user who appeared to be located in the United States, but none of the discrete communications within the MCT were determined to be wholly domestic because at least one of the communicants to each discrete communication was reasonably believed to be located outside the United States. Specifically, the 8 MCTs were determined to concern six non-targeted active users (i.e., two of the MCTs were duplicates):
  - Four MCTs (including both duplicates) [redacted] contained at least one e-mail message from a tasked selector as well as other e-mail messages from accounts/addresses/identifiers reasonably believed to be used by a person located outside the United States.[16] [redacted]
  - Three MCTs [redacted] with the users of accounts/addresses/identifiers who were reasonably believed to be located outside the United States.[17]
  - One MCT [redacted] where further technical analysis revealed that the active user was reasonably believed to be located outside the United States.
- Of the 5,081 MCTs, 10 (approximately 0.2%) contained an electronic communication account/address/identifier of a non-targeted active user who was located in the United States, and the MCTs contained at least one discrete communication that was wholly

[15] ~~(TS//SI//NF)~~ To determine the location of the non-targeted active user, NSA performed the same sort of [redacted] analysis it would perform before tasking an electronic communications account/address/identifier in accordance with its FAA Section 702 targeting procedures.

[16] ~~(TS//SI//NF)~~ To determine the location of the senders of each of these discrete e-mail messages, NSA performed the same sort of [redacted] analysis it would perform before tasking an electronic communications account/address/identifier in accordance with its FAA Section 702 targeting procedures.

[17] ~~(TS//SI//NF)~~ To determine the location of [redacted] NSA performed the same sort of [redacted] analysis it would perform before tasking an electronic communications account/address/identifier in accordance with its FAA Section 702 targeting procedures.

domestic. Specifically, all 10 of these MCTs were [redacted] and all 10 involved U.S.-based persons using [redacted].[18] For all 10 of these MCTs, only [redacted] was present. The [redacted] did not include [redacted].

- 9 of the 10 [redacted] were attributed to a single U.S.-based user. Each of these 9 [redacted] 10 total e-mail messages. The 9 [redacted] were not completely duplicative, but many of the 10 e-mail messages [redacted] were duplicative.
  - Two of the messages [redacted] in each of the 9 [redacted] contained a tasked selector and thus were not assessed to be wholly domestic.
  - Three of the messages [redacted] in each of the 9 [redacted] were [redacted] which is located in the United States) and thus were assessed to be wholly domestic.
  - The remaining e-mail messages [redacted] were between the U.S.-based user and persons reasonably believed to be located outside the United States (and thus not assessed to be wholly domestic) or whose location was unknown.[19]
- The other [redacted] was attributed to a different U.S.-based user. This [redacted] 15 total e-mail messages:
  - One of the [redacted] e-mail messages was from a tasked selector and thus was not assessed to be wholly domestic.
  - One of the [redacted] e-mail messages appeared to be a message that the U.S.-based user sent to himself [redacted] and thus was assessed to be wholly domestic.
  - One of the [redacted] e-mail messages appeared to be a message sent by an associate [redacted] account and thus was assessed to be wholly domestic.
  - The remaining e-mail messages [redacted] were between the U.S.-based user and persons reasonably believed to be

[18] ~~(TS//SI//NF)~~ [redacted]

[19] ~~(TS//SI//NF)~~ To determine the location of the other communicants, NSA performed the same sort of [redacted] analysis it would perform before tasking an electronic communications account/address/identifier in accordance with its FAA Section 702 targeting procedures.

located outside the United States and thus were not assessed to be wholly domestic.[20]

- Of the 5,081 MCTs, 1,682 (approximately 33%) required further, in-depth [redacted] analysis because they lacked information sufficient for NSA to readily identify the active user or determine the active user's location. In most of these cases, the transactions did not contain enough information for NSA to readily determine which electronic communication account/address/identifier appearing in the transaction was that of the active user. In other cases, NSA was able to determine which electronic communication account/address/identifier appearing in the transaction was that of the "active user," but NSA was unable to determine the active user's location. NSA's further [redacted] analysis of these 1,682 MCTs revealed:
  - For 1,220 of these 1,682 MCTs, NSA analysis of [redacted] data indicated that they were characteristic of a foreign use [redacted]
  - For 152 of these 1,682 MCTs, NSA analysis of [redacted] data indicated that they were [redacted]
  - For 86 of these 1,682 MCTs, NSA analysis of a combination of technical data and content revealed that they appeared to contain communications of persons located outside the United States (e.g., through further content analysis, NSA analysts were able to identify the active users of some MCTs and information indicative of those users' locations).
- Of the 5,081 MCTs, NSA cannot determine whether 224 MCTs contained wholly domestic communications, because these MCTs lack information sufficient for NSA to identify the active user or determine the active user's location. Nevertheless, NSA has no basis to believe any of these MCTs contain wholly domestic communications.
  - For 182 of these 224 MCTs, NSA technical analysis indicates that they were characteristic of [redacted]
  - For 1 of these 224 MCTs, NSA initially determined that it contained an electronic communication account/address/identifier of a non-targeted active user who appeared to be located in the United States, but whose location could not be determined upon further technical analysis. Specifically, [redacted]

[20] ~~(TS//SI//NF)~~ To determine the location of the other communicants, NSA performed the same sort of [redacted] analysis it would perform before tasking an electronic communications account/address/identifier in accordance with its FAA Section 702 targeting procedures.

[redacted].

- 23 of these 224 MCTs were not further analyzed because, although they were present in [redacted] as of the date the sample was assembled, they were subsequently purged and/or placed on NSA's Master Purge List.
- 18 of these 224 MCTs could not be further characterized by NSA analysts.

IV. (U) Conclusions Drawn from the Random Sample

~~(TS//SI//NF)~~ Based on a random sample of the approximately 13.25 million total Internet communications acquired by NSA through "upstream" techniques pursuant to FAA section 702 for the six-month period discussed, NSA assesses that the volume of transactions containing multiple communications not all of which may be to, from, or about a tasked selector is approximately between 1.29 and 1.39 million (9.70%-10.45%).[21] With respect to the types of discrete communications contained within multi-communication transactions manually reviewed by NSA analysts, [redacted]

~~(TS//SI//NF)~~ As described in Appendix B, which details NSA's Statistical Methodology for this review, the data compiled during the above-discussed manual review of a random sample of Internet communications acquired during a six-month period can be used to characterize with a statistically high degree of confidence (i.e., a simultaneous confidence level of 95% for these intervals collectively) the nature and scope of the entirety of the approximately 13.25 million Internet communications from

[21] ~~(TS//SI//NF)~~ As calculated in the attached Appendix detailing NSA's Statistical Methodology for this review, these figures are based on the 45,359 of the 50,440 transactions (89.93%) manually reviewed by NSA analysts as containing single, discrete communications and the 5,081 transactions (10.07%) manually reviewed by NSA analysts as containing multiple communications. *See also* Step 1, *supra* page 3.

[22] ~~(TS//SI//NF)~~ [redacted]

[23] ~~(TS//SI//NF)~~ [redacted]

[24] ~~(TS//SI//NF)~~ [redacted]

which the random sample was drawn. Specifically, NSA assesses that of these approximately 13.25 million Internet communications acquired through NSA upstream collection:

- between approximately 11.87 and 11.97 million (89.55%-90.30%) are transactions that contain only single, discrete communications to, from, or about a tasked selector;
- between 168,853 and 206,922 (1.27%-1.56%)[25] are transactions that contain multiple communications, all of which are either to or from a tasked selector;
- between 1,042,838 and 1,113,947 (7.87%-8.53%)[26] are transactions that contain multiple communications, at least one of which is to, from, or about NSA's tasked selector, but all of which are believed to either be to or from non-targeted persons reasonably believed to be located outside the United States;
- between 48,609 and 70,168 (0.37%-0.53%)[27] are transactions that contain multiple communications, at least one of which is to, from, or about NSA's tasked selector, and at least one of which is a communication between non-targeted persons (i.e., not to, from or about a tasked selector) that lacks sufficient information for NSA to identify the location of the sender and all intended recipients of that communication; and
- between 996 and 4,965 (0.0075%-0.0375%) contain a wholly domestic communication not to, from, or about a tasked selector.

~~(TS//SI//NF)~~ In sum, while there was insufficient information present for 224 multi-communication transactions for NSA analysts to characterize the likelihood that they may contain wholly domestic communications (the majority of which were attributable to [redacted], for the reasons explained in detail

[25] ~~(TS//SI//NF)~~ As calculated in the attached Appendix, these figures are based on 713 of the 5,081 MCTs (14.03%) and 50,440 total transactions (1.41%) reviewed by NSA analysts as containing a tasked selector as the active user [redacted]. See also Step 3, *supra* page 4.

[26] ~~(TS//SI//NF)~~ As calculated in the attached Appendix, these figures are based on 4,134 of the 5,081 MCTs (81.36%) and 50,440 total transactions (8.19%) reviewed by NSA analysts as containing discrete communications believed to be to or from non-targeted persons located outside the United States. More specifically, this total includes the following MCTs manually reviewed by NSA analysts: 2,668 that had an active user reasonably believed to be a person located outside the United States; 8 that included at least one communicant reasonably believed to be located outside the United States for each communication therein; 1,220 that are characteristic o[redacted] 152 that are indicative of [redacted] and 86 that all communications contained therein were to or from persons located outside the United States. *See* Step 3, *supra* pages 4-6.

[27] ~~(TS//SI//NF)~~ As calculated in the attached Appendix, these figures are based on 224 of the 5,081 MCTs (4.41%) and 50,440 total transactions (0.44%%) reviewed by NSA analysts that lacked sufficient information to identify the active user or the active user's location. *See* Step 3, *supra* page 6.

above, NSA has no basis to believe any of the remaining Internet communications reviewed in the 50,440 sample are wholly domestic beyond those 10 discussed above.[28] Moreover, each of those 10 Internet communications has been placed on NSA's Master Purge List.

----- *The remainder of this page intentionally left blank.* -----

[28] ~~(TS//SI//NF)~~ *See* Figure D of Appendix A.

(U) VERIFICATION

(U) I declare under penalty of perjury that the facts set forth in the foregoing "NSA Characterization of Upstream Data: Process and Results" are true and correct based upon my best information, knowledge and belief. Executed pursuant to Title 28, United States Code, § 1746, on this 16th day of August, 2011.

[redacted]

Signals Intelligence Directorate Compliance Architect
National Security Agency

Appendix A

Fig. A Total FAA 702

140,974,921 Internet Communications

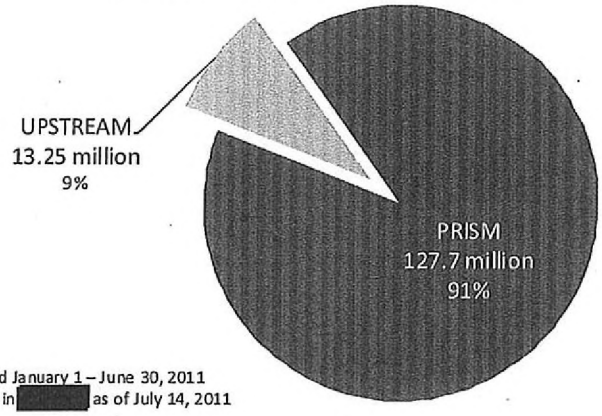

Acquired January 1 – June 30, 2011
Present in [redacted] as of July 14, 2011

Fig. B Total Upstream Sample

50,440 objects manually reviewed

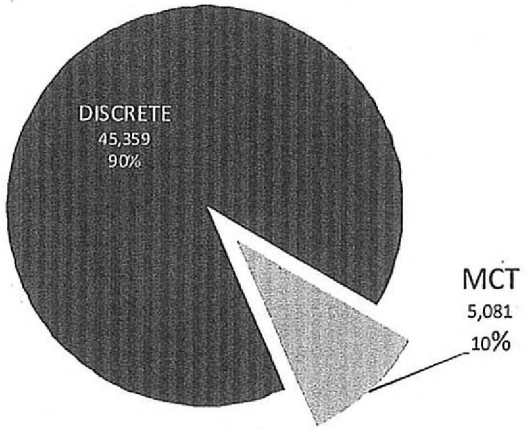

Fig. C MCT Type

5,081 objects

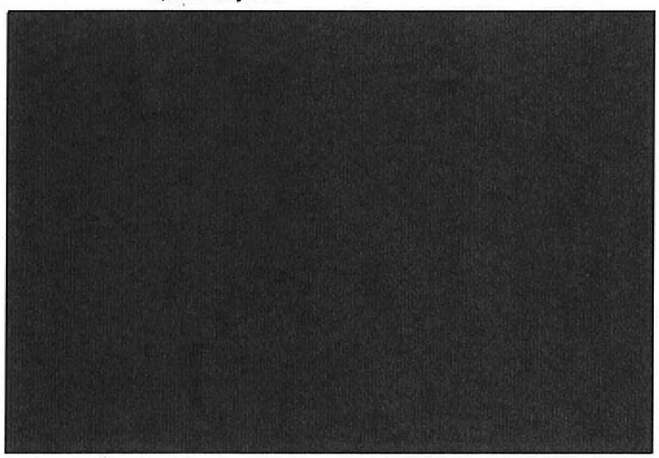

Fig. D Summary

50,440 objects

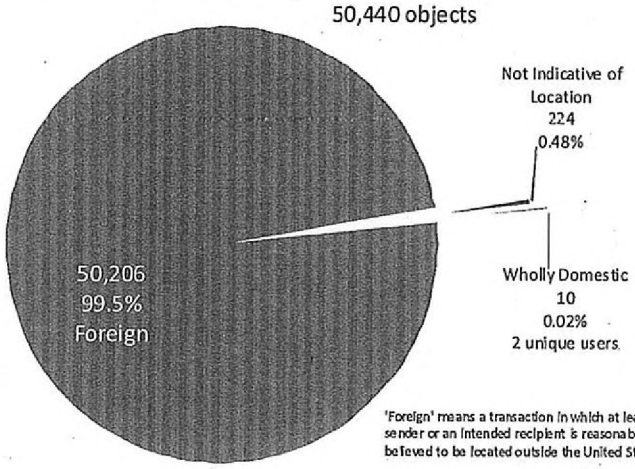

'Foreign' means a transaction in which at least the sender or an intended recipient is reasonably believed to be located outside the United States.

Derived From: NSA/CSSM 1-52
Dated: 20070108
Declassify On: 20360801

Appendix B: Statistical Methodology – FAA Section 702 Upstream Manual Review

~~(TS//SI//NF)~~ Using statistical analysis NSA determined the proportions of transactions satisfying certain criteria (e.g., proportion of FAA Section 702 upstream Internet transactions that are Multi-communication Transactions (MCT) versus transactions containing single, discrete communications). As further described below, transactions were categorized in various ways. The categorization process can be complex; to minimize categorization error, NSA used a statistical approach involving actual examination of an appropriate sample of transactions by experienced intelligence analysts. (The use of only a sample is a concession to the large volume of transactions and the labor-intensive nature of the categorization process.) That is, NSA traded "categorization error" for "statistical error"; the latter refers to the fact that by considering only a randomly sampled portion of the universe of transactions, NSA estimated the true proportions (as they exist in the universe) -- with error bounds and levels of confidence that can be stated justifiably.

~~(TS//SI//NF)~~ **THE SAMPLE.** As discussed more fully in the "NSA Characterization of Upstream Data: Process and Results," NSA identified 13,256,067 transactions acquired through NSA's FAA 702 upstream collection during a six-month period from January 1st through June 30th, 2011. Of those approximately 13.25 million transactions, a team of experienced intelligence analysts carefully examined 50,440 over a nearly one-month time period. The transactions were presented to the analysts in a randomized order, ensuring that a simple random sample would serve as the basis for conclusions – supported by statistical theory – about the true proportions of the 13.25 million-transaction universe.

~~(TS//SI//NF)~~ **ESTIMATES AND CONFIDENCE INTERVALS.** The proportions formed from the sampled transactions serve as unbiased estimates of the corresponding proportions of the 13,256,067-transaction universe. Further, for (six) selected proportions, NSA states a confidence interval for each. Collectively, these intervals have a simultaneous confidence level of 95%. This means that the intervals were produced by a procedure calibrated to produce, for at least 95% of the sample sets NSA could have drawn, intervals which all cover the corresponding true (i.e., universal) proportions. Individually, each interval has a higher level of confidence associated with it; component confidence levels are quoted below.

~~(TS//SI//NF)~~ For each of the six categories, NSA also states a confidence interval for the actual number of that category's transactions within the 13,256,067-transaction (January-June, 2011 upstream) universe. Such an interval is simply an equivalent representation of the corresponding proportion-interval (it is obtained by multiplying the endpoints of the proportion-interval by 13,256,067), and so the inclusion of such intervals does not affect the (95%) level of simultaneous confidence.

~~(TS//SI//NF)~~ Specifically: By sampling a subset of the universe (or *population*) of upstream transactions, NSA estimated the following six proportions. (Hereinafter, N denotes 13,256,067 – the size of that universe; M denotes the (unknown) actual number of MCTs in that universe).

- M / N: the proportion of the population comprising MCTs;

Derived From: NSA/~~CSSM~~ 1-52
Dated: ~~20070108~~
Declassify On: 20360801

- 1-(M / N): the proportion of the population comprising discrete transactions;

- the proportion of the population comprising MCTs in which all communications are either to or from NSA's tasked selector (hereinafter labeled "Target" MCTs);

- the proportion of the population comprising MCTs in which all communications are believed to either be to or from non-targeted persons located outside the United States (hereinafter labeled "Foreign" MCTs);

- the proportion of the population comprising MCTs in which the nature of one or more communications between non-targeted persons lacked sufficient information for NSA analysts to identify the location of the sender and all intended recipients (hereinafter labeled "Unknownable" MCTs);

- the proportion of the population comprising MCTs that NSA analysts assessed contain a wholly domestic not to, from, or about a tasked selector (hereinafter labeled "Confirmed Wholly Domestic").

~~(TS//SI//NF)~~ (The first of these proportions equals the total of the last four.) In the following, lower-case letters denote transaction counts as realized in the sample, in categories corresponding to their upper-case counterparts. That is, n is the number of transactions sampled (this turned out to be 50,440), and m is the number of MCTs in the sample.

~~(TS//SI//NF)~~ **OUTLINE OF PROCEDURE.** NSA designed a procedure that accepts a size-n *simple random sample*[1] of the population, and produced from it estimates and confidence intervals for the six "true"[2] proportions NSA sought. The estimates NSA produced are simply the corresponding proportions as found in the sample – e.g., the sample proportion m/n was NSA's estimate of the population proportion M/N; such a sample proportion is unbiased[3] for its population counterpart, meaning that were a sample proportion to be computed for each of the possible size-n samples that could be drawn, the average of these sample proportions would equal the "true" (population) proportion.

---

[1] ~~(TS//SI//NF)~~ A simple random sample is one that is drawn in a way that ensures that all possible size-n subsets of the (size-N) population have an equal chance of being selected; this sampling technique enables statistically justifiable claims by avoiding potential (known or unknown) sources of bias in the population (e.g., a periodic trend in the population over time).

[2] ~~(TS//SI//NF)~~ "True" refers to proportions that relate to the entire population, which cannot be determined for certain, as n is smaller than N.

[3] ~~(TS//SI//NF)~~ Unbiasedness means that the estimate is aiming for the right "target"; however, it indicates nothing about the precision of the estimate. An estimation procedure can be unbiased whether it is based on a small or large sample size n.

~~(TS//SI//NF)~~ To express precision appropriately, NSA designed its procedure to produce confidence intervals – one for each of the (six) population proportions of interest – having a simultaneous confidence level of 95%. This means that:

- Based on a sample, the procedure will produce a collection of intervals, each asserted to contain the true (population) proportion it targets.

- Because the procedure operates on a random sample, the interval endpoints are *random variables*; the particular collection of intervals a particular sample yields may fail to cover one or more of the population proportions it targets. But the procedure is designed so that this failure probability – *whatever* the true proportions are – is no more than 5%; that is, for at least 95% of the (size-n) simple random samples it might process, the procedure will produce intervals which *all* cover their targeted population proportions.

- In order to achieve this level of confidence about a collection of intervals simultaneously, the procedure is designed so that the respective failure probabilities associated with the component intervals total no more than 5%. In particular, this 5% was allocated as follows:

    - 2.5% to the proportion of "Confirmed Wholly Domestic";
    - 0.67% to each of the "Target," "Foreign," "Unknown" proportions;
    - 0.5% to the proportion of MCT (i.e., M / N). As the proportions of discrete and MCT transactions are complementary (i.e., they total 1), the confidence interval for the proportion of discrete transactions is obtained by subtracting each of the endpoints for the MCT-interval from 1 – and it is the case that one of these intervals will cover its population target if and only if the other does. Therefore, there is no need to separately allocate "failure probability" to the proportion-of-discrete.

~~(TS//SI//NF)~~ The probability of drawing a sample resulting in one or more "failing" intervals is no more than the sum of the failure probabilities of the respective component intervals, hence the claim of 95% confidence for the procedure outlined here. The "no more" qualification makes this technique conservative: relationships (complicated and left unanalyzed) between the random variables involved may make the practical confidence level higher; 95% represents a worst-case claim. To achieve simultaneous 95% confidence, the 5% failure probability could have been allocated in any way. (Broadly: the lower the confidence level (i.e., the higher the failure probability), the narrower the intervals the procedure will produce. An extreme example: a procedure for 100% confidence intervals would produce uselessly wide intervals, as it would have to be able to claim that its intervals cover truth for *every* possible size-n sample it could have received.) This procedure for simultaneous intervals is conservative in a further way: Just as the sum of the discrete and MCT proportions equals 1, so does the sum of the discrete, "Target," "Foreign," "Unknown," and "Confirmed Wholly Domestic" proportions. It is difficult to exploit this latter constraint properly; NSA utilized the conservative method described here to ensure that its assertions about the procedure's performance are valid.

~~(TS//SI//NF)~~ **CONFIDENCE-INTERVAL PROCEDURE FOR A SINGLE PROPORTION.** As outlined above, the procedure for (95%) simultaneous confidence intervals was achieved by producing component confidence intervals based on (individually higher) levels of confidence (e.g., 99.5% for M/N). The construction of component confidence intervals can be understood via the following example, using the M/N target. For the sample of size n to be observed, m represents the (random) number of MCTs to be realized in the sample. Formally, m has a *hypergeometric* distribution (arising from sampling transactions "without replacement"); to make the mathematical computations tractable, NSA approximated this distribution by a *binomial* distribution corresponding to sampling *with* replacement (in which each sampled transaction would be replaced after it is drawn, and hence would be eligible to be drawn multiple times). This approximation is uniformly conservative; i.e., it will result in wider intervals. The proportion to be estimated, M/N, appears as the (unknown) parameter (now denoted p) of this binomial distribution. Treating m as a binomial random variable based on n *trials*, NSA used an accepted method (the *Clopper-Pearson* method) as the basis to devise its confidence-interval procedure for p. (Below, the notation $B(n,q)$ refers to an n-trial binomial random variable having parameter q.) Upon observing m, NSA:

- Determines, for each of various proportions x between 0 and 0.5%, parameters q and r such that
  - x is the probability that a $B(n,q)$ random variable takes a value of at least m (but if m=0, take q to be 0);
  - (0.5% - x) is the probability that a $B(n,r)$ random variable takes a value no larger than m (but if m=n, take r to be 1).

  r exceeds q; the pair $[q,r]$ determines an interval.

- Determines the narrowest of all such intervals $[q,r]$ and reports it as the (99.5%) confidence interval for $p = M/N$.

~~(TS//SI//NF)~~ Practically, the q's and r's can be computed using *inverse Beta functions*, and computer software can find the narrowest interval efficiently.

Remainder of this page intentionally left blank.

RESULTS:

	# of transactions in sample	Sample proportion (of 702 upstream)	Confidence interval for corresponding universal proportion	Confidence interval for the actual number (of the 13.25 million)
Discrete	45,359	0.8993	0.8955 – 0.9030	11,870,284 – 11,970,275
MCT	5,081	0.1007	0.0970 – 0.1045	1,285,792 – 1,385,783

	# of transactions in sample	Sample proportion (of MCT)	Confidence interval for corresponding universal (MCT) proportion	Confidence interval for the actual number (of the 13.25 million)
TARGET	713	0.01414	0.01274 – 0.01561	168,853 – 206,922
FOREIGN	4,134	0.08196	0.07867 – 0.08532	1,042,838 – 1,130,947
UNKNOWABLE	224	0.004441	0.003667 – 0.005293	48,609 – 70,168
CONFIRMED WHOLLY DOMESTIC	10	0.0001983	0.00007508 – 0.0003746	996 – 4,965

[redacted]

Remainder of this page intentionally left blank.

VERIFICATION

I declare under penalty of perjury that the facts set forth in this Appendix are true and correct based upon my best information, knowledge and belief. Executed pursuant to Title 28, United States Code, Section 1746, on this 11th day of August, 2011.

[redacted]
[Statistician]
National Security Agency