UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

NETSKOPE, INC.,

Petitioner,

v.

FORTINET, INC.,

Patent Owner.

PTAB Case No. IPR2023-00459

Patent No. 10,084,825

DECLARATION OF MICHAEL FRANZ IN SUPPORT OF PETITION FOR *INTER PARTES* REVIEW OF U.S. PATENT NO. 10,084,825

TABLE OF CONTENTS

I. INTRODUCTION 1

A. Qualifications 2

1. Education 2

2. Work Experience 2

3. Publications 6

4. Curriculum Vitae 7

B. Materials Reviewed 7

C. Level of Ordinary Skill in the Art 8

D. Summary of Opinions 9

II. OVERVIEW OF THE TECHNOLOGY 10

A. Priority Date of the Claims 10

B. Overview of Relevant Technology When the '825 Patent Was Filed 11

1. Coordinated Work Existed Long Before Computers 11

2. Networked Computers Have Coordinated Their Work for Decades 13

3. Computing "Appliances" 22

C. The '825 Patent 23

D. The Challenged Claims 28

E. Claim Construction 28

III. UNPATENTABILITY OF THE '825 PATENT CLAIMS 28

A. Standards for Invalidity 28

1. Obviousness 28

B. Ground I: Chandra renders claims 1-5, 12, 15, 16, 21, 22, and 27 obvious. 30

1. Claim 1 33

2. Claim 2 48

3. Claim 3 50

Netskope Exhibit 1002

4. Claim 4 51
5. Claim 5 53
6. Claim 12 53
7. Claim 15 54
8. Claim 16 55
9. Claim 21 55
10. Claim 22 56
11. Claim 27 56
C. Ground 2: Chandra and Buruganahalli render claims 6, 7, 17, 18, 23, 24, 28, and 29 obvious. 57
1. Motivation to Combine 60
2. Claim 6, 17, 23, and 28 61
3. Claims 7, 18, 24, and 29 68
D. Ground 3: Chandra and Kumar render claims 10 and 11 obvious. 70
1. Motivation to Combine 72
2. Claim 10 72
3. Claim 11 75
E. Ground 4: Chandra and Chandra-2 render claims 13, 14, 19, 20, 25, 26, 30, and 31 obvious. 76
1. Claims 13, 19, 25, 30 76
2. Claims 14, 20, 26, 31 77
F. Ground 5: Keohane renders claims 1-5, 10-12, 15, 16, 21, 22, and 27 obvious 78
1. Claim 1 82
2. Claim 2 94
3. Claim 3 97
4. Claim 4 99
5. Claim 5 101
6. Claim 10 102

Netskope Exhibit 1002

7. Claim 11 104
8. Claim 12 104
9. Claim 15 105
10. Claim 16 106
11. Claim 21 107
12. Claim 22 107
13. Claim 27 108
G. Ground 6: Keohane and Buruganahalli render claims 6, 7, 17, 18, 23, 24, 28, and 29 obvious. 108
1. Motivations to combine 108
2. Claims 6, 17, 23, and 28 109
3. Claims 7, 18, 24, and 29 112
H. Ground 7: Keohane and Kumar render claims 10 obvious. 115
1. Motivation to Combine 115
2. Claim 10 115
I. Ground 8: Keohane and Chandra-2 render claims 13, 14, 19, 20, 25, 26, 30, and 31 obvious. 117
1. Claims 13, 19, 25, 30 117
2. Claims 14, 20, 26, 31 118
IV. CONCLUSION 119

Netskope Exhibit 1002

LIST OF APPENDICES

Appendix A	*Curriculum Vitae* of Michael Franz, Ph.D.
Appendix B	National Security Agency/Central Security Service Information Assurance Directorate - Mobile Access Capability Package (June 19, 2015)
Appendix C	National Security Agency/Central Security Service Information Assurance Directorate - Multi-site Connectivity Capability Package (May 4, 2016)
Appendix D	Litigation Summary of Michael Franz, Ph.D.

Netskope Exhibit 1002

I. INTRODUCTION

1. I, Michael Franz, have been retained by Petitioner Netskope, Inc. ("Petitioner") to investigate and opine on certain issues relating to United States Patent No. 10,084,825 ("the '825 patent") (Ex. 1001) in their Petition for *Inter Partes* Review of that patent. The Petition requests that the Patent Trial and Appeal Board ("PTAB" or "Board") review and cancel claims 1-7 and 10-31 of the '825 patent. I have analyzed the claims of the '825 patent in numerical order.

2. The opinions set forth in this declaration are based on my personal knowledge, my professional judgment, and my analysis of the materials and information referenced in this declaration and its exhibits.

3. I am being compensated for consulting services, including time spent testifying at any hearing that may be held, at my standard consulting rate of $750 per hour. I am also reimbursed for reasonable and customary expenses associated with my work in this case. I receive no other forms of compensation related to this case. My compensation does not depend on the outcome of this *inter partes* review or the co-pending district court litigation and *inter partes* reviews, and I have no other financial interest in this *inter partes* review.

4. I understand that the '825 patent has been assigned to Fortinet, Inc.

5. This declaration is based on the information currently available to me. To the extent that additional information becomes available, I reserve the right to

Netskope Exhibit 1002

continue my investigation and study, which may include a review of documents and information that may be produced, as well as testimony from depositions that have not yet been taken.

A. Qualifications

1. Education

6. I completed my undergraduate studies with a Diplomingenieur from the Swiss Federal Institute of Technology in Zurich ("ETH Zurich") in 1989. In 1994, I obtained my Doctorate of Technical Sciences from ETH Zurich. My dissertation was entitled "Code-Generation On-the-Fly: A Key to Portable Software."

2. Work Experience

7. I am a tenured Full Professor of Computer Science in the Donald Bren School of Information and Computer Sciences at the University of California, Irvine ("UCI"). I am also, by courtesy, a Full Professor of Electrical Engineering and Computer Science in the Henry Samueli School of Engineering at UCI. In 2016, the University awarded me the title of distinction of "Chancellor's Professor" and in 2022, I was further elevated to "Distinguished Professor," the highest campus-level title of distinction at my university.

8. I have served as a visiting professor at ETH Zurich, the University of Klagenfurt in Austria, the Technical University of Brunswick, and the University of Ulm in Germany.

Netskope Exhibit 1002

9. I have been elevated to Fellow of the Institute of Electrical and Electronics Engineers (IEEE), the global engineering society. Fellow is the highest of three grades of membership that are awarded based on merit. In every year, IEEE limits the number of new Fellows to one tenth of one percent of the membership, which currently stands at about 400,000 members.

10. I have also been elevated to Fellow of the Association for Computing Machinery (ACM), the global professional society for computer scientists. Fellow is the highest of ACM's four grades of membership. ACM's rules for Fellows are even more restrictive than IEEE's, limiting the total number of Fellows in absolute terms to 1% of the membership, which currently stands at about 100,000 members. In recent years, ACM has typically elevated no more than 50 individuals to Fellow status in a single year; in 2015, the year I was advanced, there were 42 new Fellows.

11. I have also been elevated to Fellow of the American Association for the Advancement of Science (AAAS), the world's largest general scientific society, with over 120,000 members.

12. Lastly, I am a member of the inaugural cohort of Fellows of the International Federation for Information Processing (IFIP), the global organization for researchers and professionals working in the field of computing. Established in 1960 under the auspices of UNESCO, IFIP is recognized by the United Nations and links some 50 national and international societies and academies of science with a

Netskope Exhibit 1002

total membership of over half a million professionals. IFIP is an international, non-governmental organization that operates on a non-profit basis.

13. In April of 2021, the Association for Computing Machinery (ACM) named me the recipient of the ACM Charles P. "Chuck" Thacker Breakthrough in Computing Award. This is the third highest award across all of computer science given by a professional society, behind only the Turing Award ("the Nobel Prize of Computing") and the ACM Prize in Computing (for early to mid-career contributions).

14. Furthermore, I am a recipient of the IEEE Computer Society's Technical Achievement Award. At most 5 of these awards are given annually by the IEEE Computer Society, the largest of the IEEE's technical societies with a current membership of more than 60,000 members. I am also a recipient of the Humboldt Research Award, also known as the "Humboldt Prize." The award, given by the Alexander von Humboldt Foundation of Germany and funded by the German federal government, recognizes renowned researchers outside of Germany whose "fundamental discoveries, new theories or insights have had a significant impact on their own discipline and who are expected to continue producing cutting-edge achievements in the future." It is the highest award given by the Foundation to researchers based outside of Germany.

Netskope Exhibit 1002

15. I have led pioneering research on both the security and the performance aspects of downloadable code in client-server settings such as what we today call "Web 2.0." Two of my six awarded U.S. Patents concern themselves with security aspects of such programs downloaded over a network, one concerns itself with performance aspects of such code, while two further patents concern themselves with detecting and preventing malware attacks in general.

16. My research has had a real and lasting impact on a great many people. I am the co-inventor (with one of my former Ph.D. students) of the "Trace Tree" compilation technique, for which the United States Patent and Trademark Office has awarded U.S. Patent No. 8,769,511. I collaborated with the non-profit Mozilla Foundation to incorporate this technique into the Firefox web browser, where it became the basis of the "TraceMonkey" JavaScript engine, eventually used by several hundred million people every day.

17. Over the course of my career so far, I have been the Principal Investigator on several high-profile research projects with a total budget of well over $20M. My expertise in software systems with distinct emphases on the security and performance of client-server and mobile computing has been sought out repeatedly by the Federal Government, and I have been participating in many high-level invitation-only meetings on Critical Infrastructure Protection and on Cyber Security organized by the National Intelligence Community, the Department of Defense, the

Netskope Exhibit 1002

Department of Homeland Security, and the Department of Energy. The overwhelming majority of the federally-funded projects that I have worked on over the course of my career have been related to computer security.

18. Before I reached the maximum allowed number of consecutive terms on the editorial board, I was an Associate Editor of one the flagship journals of the IEEE, the IEEE Transactions on Dependable and Secure Computing (TDSC). I have also served on the editorial boards of two further peer-reviewed scholarly journals focusing on software engineering, Software – Practice and Experience (SPE) and Computer Science – Research and Development (CSRD). I have served on the program committees of most major academic conferences that are related to the various themes of my research. I have served as the primary advisor to 35 completed Ph.Ds. and currently serve as the primary advisor on several further dissertations in progress.

3. Publications

19. In addition to my dissertation, I co-authored "Automated Software Diversity," released in 2015.

20. I have published 36 reviewed journal and magazine articles since 1993, and well over 100 conference and workshop papers.

21. I am an inventor on six issued U.S. patents.

Netskope Exhibit 1002

4. Curriculum Vitae

22. A copy of my curriculum vitae is attached as Appendix A to this declaration.

B. Materials Reviewed

23. My opinions expressed in this declaration are based on documents and materials identified in this declaration, including the ’825 patent and its prosecution history, the prior art references and background materials discussed in this declaration, and the other references specifically identified in this declaration. I have considered these materials in their entirety, even if only portions are discussed here. The following is an exemplary list of materials on which I based my opinions.

Ex. 1001	U.S. Patent No. 10,084,825 to Xu (“the ’825 patent”)
Ex. 1003	File History of U.S. Patent No. 10,084,825
Ex. 1004	U.S. Patent Application Publication 2007/0204018 A1 (“Chandra”)
Ex. 1005	U.S. Patent Application Publication 2014/0282843 A1 (“Buruganahalli”)
Ex. 1006	U.S. Patent Application Publication 2009/0144818 A1 (“Kumar”)
Ex. 1007	U.S. Patent Application Publication 2017/0093917 A1 (“Chandra-2”)
Ex. 1008	U.S. Patent Application Publication 2008/0134332 A1 (“Keohane”)
Ex. 1009	Microsoft Computer Dictionary Fifth Edition (2002)
Ex. 1010	U.S. Patent Publication 2006/0156401 A1 (“Newstadt”)
Ex. 1011	U.S. Patent Publication 2007/0237147 A1 (“Quinn”)

Netskope Exhibit 1002

Appendix B	National Security Agency/Central Security Service Information Assurance Directorate - Mobile Access Capability Package (June 19, 2015)
Appendix C	National Security Agency/Central Security Service Information Assurance Directorate - Multi-site Connectivity Capability Package (May 4, 2016)

24. I have also relied on my own experience and expertise in computer systems and computer systems security.

C. Level of Ordinary Skill in the Art

25. I am not an attorney and offer no legal opinions. I have been informed about certain aspects of the law for purposes of my analyses and opinions.

26. I understand that in analyzing questions of invalidity and infringement, the perspective of a person having ordinary skill in the art ("POSITA") is often implicated, and the Court may need assistance in determining that level of skill.

27. I understand that the claims and written description of a patent must be understood from the perspective of a POSITA. I have been informed that the following factors may affect the level of skill of a POSITA: (1) the educational level of the inventor; (2) the type of problems encountered in the art; (3) the prior-art solutions to those problems; (4) the pace of innovation; (5) the sophistication of the technology; and (6) the educational level of active workers in the field. A person of ordinary skill in the art is also a person of ordinary creativity in the art.

Netskope Exhibit 1002

28. Based on my experience in network security system design and network workflow coordination, as well as my reading of the '825 patent, it is my opinion that a person of ordinary skill with respect to the subject matter of the '825 patent at the time of the alleged priority date of the '825 patent in May 2017 would have had at least a B.S. degree in computer science, computer engineering, or electrical engineering (or equivalent experience) and would have had at least two years or more of practical experience with computer networking and computer security. This definition is flexible and additional educational experience in computer science could make up for less work experience and vice versa.

29. I am at least a person of ordinary skill in the art and was so on the date to which the '825 patent claims priority. As shown by my qualifications and my curriculum vitae attached as Appendix A, I am aware of the knowledge and skill possessed by a person of ordinary skill in the art at the time of the priority date of the '825 Patent. In performing my analysis, I have applied the standard set forth above.

D. Summary of Opinions

30. I have reviewed and analyzed the '825 Patent.

31. Based on my review and analysis, it is my opinion that claims 1-7 and 10-31 of the '825 Patent are invalid based on the following grounds:

Ground	Basis	Reference(s)	Challenged Claims
1	§ 103	Chandra	1-5, 12, 15, 16, 21, 22, 27

Netskope Exhibit 1002

2	§ 103	Chandra + Buruganahalli	6, 7, 17, 18, 23, 24, 28, 29
3	§ 103	Chandra + Kumar	10, 11
4	§ 103	Chandra + Chandra-2	13, 14, 19, 20, 25, 26, 30, 31
5	§ 103	Keohane	1-5, 10-12, 15, 16, 21, 22, 27
6	§ 103	Keohane + Buruganahalli	6, 7, 17, 18, 23, 24, 28, 29
7	§ 103	Keohane + Kumar	10
8	§ 103	Keohane + Chandra-2	13, 14, 19, 20, 25, 26, 30, 31

II. OVERVIEW OF THE TECHNOLOGY

A. Priority Date of the Claims

32. I have been informed that a U.S. patent application may claim the benefit of the filing date of an earlier patent application if the earlier patent application disclosed each limitation of the invention claimed in the later-filed U.S. patent application. I have also been informed that priority is determined on a claim-by-claim basis so that certain claims of a patent may be entitled to the priority date of an earlier-filed patent application even if other claims of the same patent are not entitled to that priority date.

33. I have also been informed that a patented claim is invalid if the claimed invention was patented, described in a printed publication, or in public use, on sale, or otherwise available to the public before the effective filing date of the claimed invention, or the claimed invention was described in an issued patent or a published patent application that was effectively filed before the effective filing date of the claimed invention.

34. I understand that the ’825 patent claims a priority date of May 8, 2017.

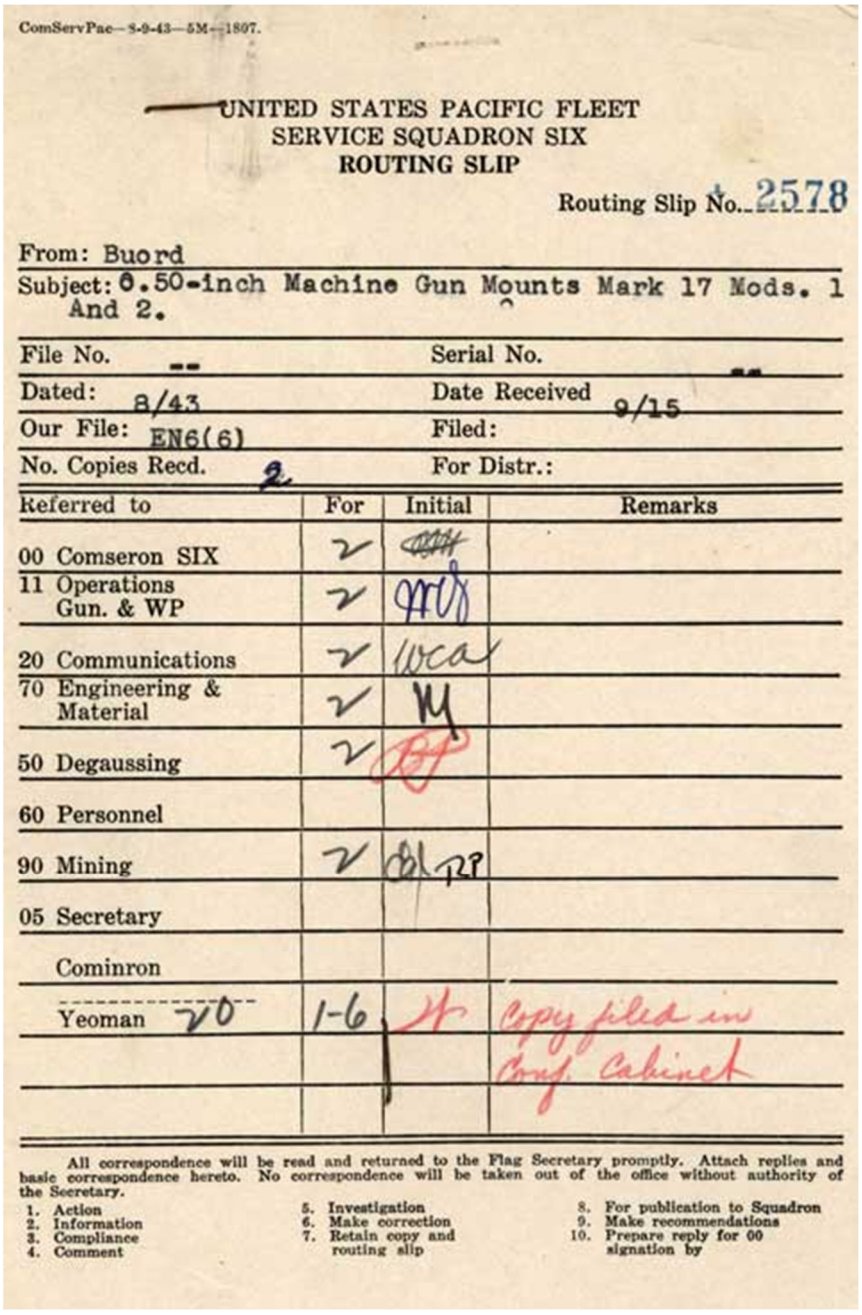
Netskope Exhibit 1002

B. Overview of Relevant Technology When the '825 Patent Was Filed

1. Coordinated Work Existed Long Before Computers

35. The concept of coordinating workflow via "routing slips" on which different parties "signed off" their task items as the routing slip was passed from one such party to the next is very old. As an example, the picture below (from the National Archives, dated August 1943) shows a routing slip used for distributing an ordnance pamphlet to members of the United States Pacific Fleet.

Netskope Exhibit 1002



ComServPac—8-9-43—5M—1807.

UNITED STATES PACIFIC FLEET
SERVICE SQUADRON SIX
ROUTING SLIP

Routing Slip No. 2578

From: Buord

Subject: 0.50-inch Machine Gun Mounts Mark 17 Mods. 1 And 2.

File No. --	Serial No. --
Dated: 8/43	Date Received 9/15
Our File: EN6(6)	Filed:
No. Copies Recd. 2	For Distr.:

Referred to	For	Initial	Remarks
00 Comseron SIX	2		
11 Operations Gun. & WP	2		
20 Communications	2		
70 Engineering & Material	2		
50 Degaussing	2		
60 Personnel			
90 Mining	2		
05 Secretary			
Cominron			
Yeoman 20	1-6		copy filed in Corr. Cabinet

All correspondence will be read and returned to the Flag Secretary promptly. Attach replies and basic correspondence hereto. No correspondence will be taken out of the office without authority of the Secretary.

1. Action
2. Information
3. Compliance
4. Comment
5. Investigation
6. Make correction
7. Retain copy and routing slip
8. For publication to Squadron
9. Make recommendations
10. Prepare reply for 00 signation by

36. As can be seen in the example, there were different parties that each needed to perform a specific task on the payload that was being routed. The specific tasks that different parties performed were not explicit from the routing slip (above), but such tasks would have been separately specified in advance in the form of, e.g., workplace manuals, for many organizations. Each party of the organization would

Netskope Exhibit 1002

have a clear understanding of not just its own tasks, but also the tasks that other parties were performing.

37. After completing their respective tasks, parties attached their unique signature certifying their fulfillment of their specific tasks and passed the complete packet (routing slip and payload) to the next party. Besides the previously specified standardized tasks, the routing slip also provided for the signed attestations of additional tasks performed (e.g., "copy filed in conf. cabinet"). At any time, a recipient of the routing slip could verify (by checking the signature) which other recipient(s) had already performed their tasks and would know what the respective tasks assigned to other recipient consisted of. Based on the assumption that each previous party had actually performed all of the tasks assigned to them, a recipient could then elect to skip the tasks already performed, or could possibly repeat them.

2. Networked Computers Have Coordinated Their Work for Decades

38. The '825 patent describes mechanisms that replicate these well-known concepts in the context of computer networks.

39. By the early 2000's, most businesses and medium to large organizations were deploying computer networks and the general public had become aware of the threat of computer viruses, worms, and other malware. An outright computer security industry had emerged, offering a wide array of computer security products, both hardware and software. A company or organization would often

Netskope Exhibit 1002

purchase an array of different solutions from different vendors to build and secure their network; for example, a company might purchase firewalls and routers in addition to virus/malware detection and mitigation products. Consequently, such networks might contain a hodgepodge of products potentially performing partially overlapping tasks.[1]

40. To address this problem, a slew of developments were made to improve device coordination. For example, below are four examples of patent literature dating long before the '825 patent that show how POSITAs were attuned to the problem of uncoordinated devices and to solving that problem by having each device in a sequence indicate (by adding additional signaling information, which the different prior solutions variously refer to as "tagging the network traffic" and "adding signature headers") what action(s) it had performed so that downstream devices could determine which action(s) they still needed to perform.

A. U.S. Patent Publication 2006/0156401 A1 ("Newstadt"), filed January 12, 2005

[1] *See* U.S. Patent No. 8,000,329, 1:26-60.

Netskope Exhibit 1002

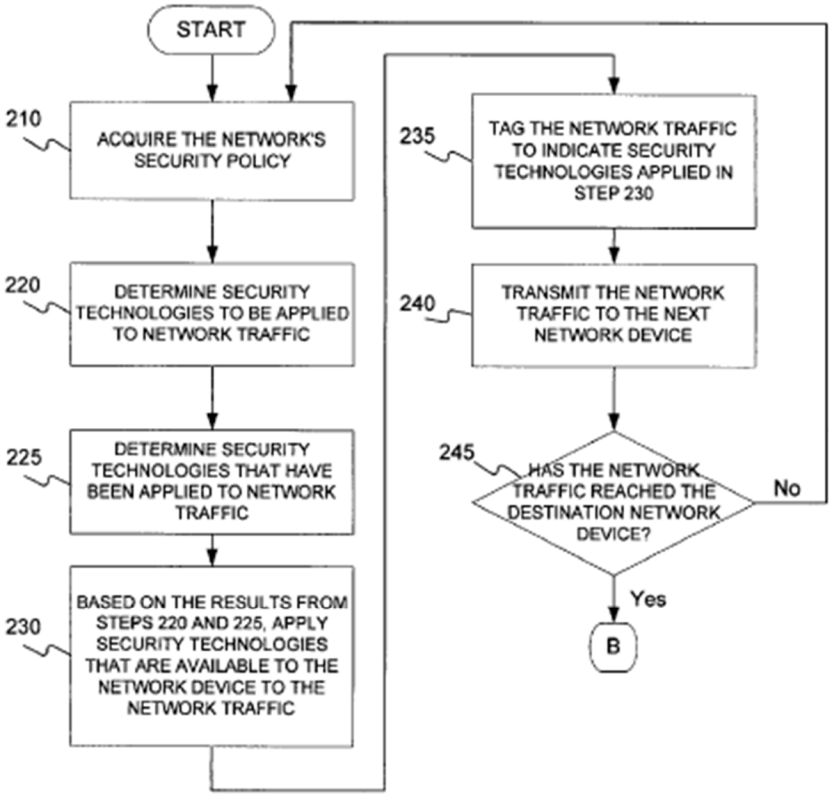


Fig. 2

Netskope Exhibit 1002

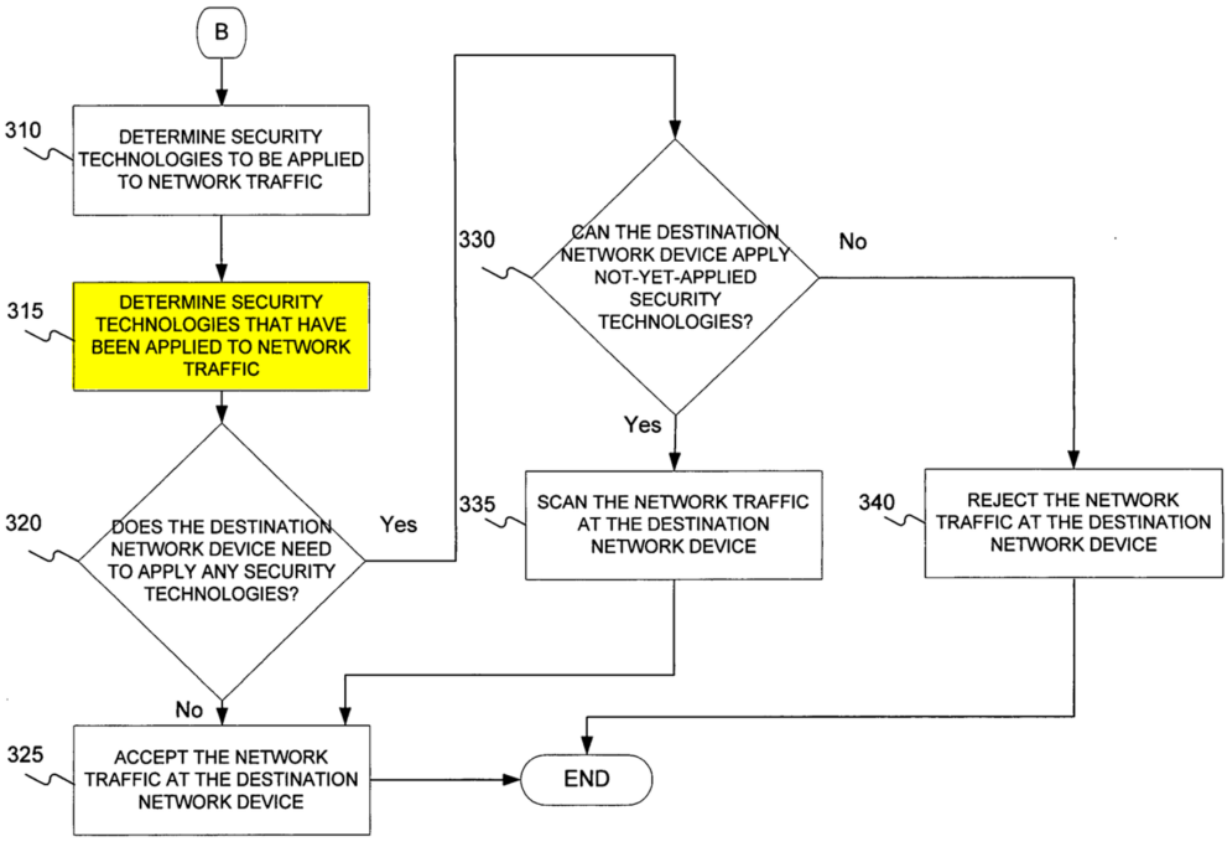


Fig. 3 (Annotated)

B. U.S. Patent Publication No. 2008/134332 A1 ("Keohane"), filed December 4, 2006

Netskope Exhibit 1002

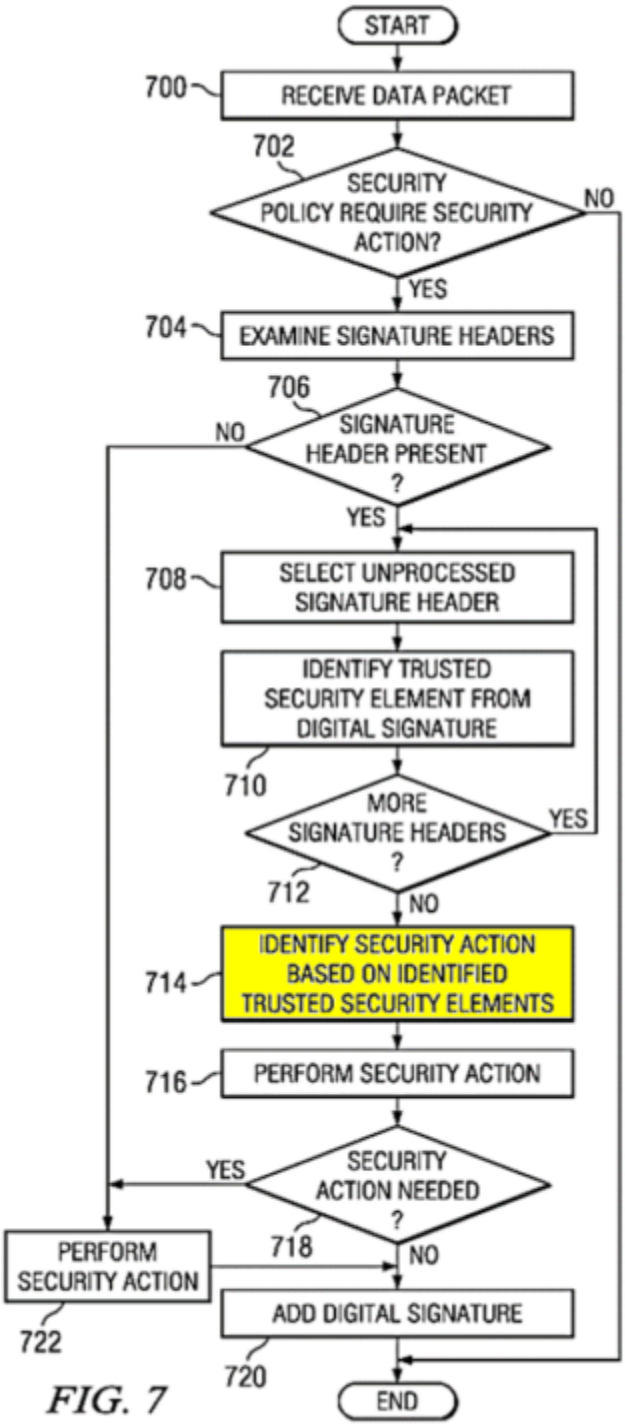
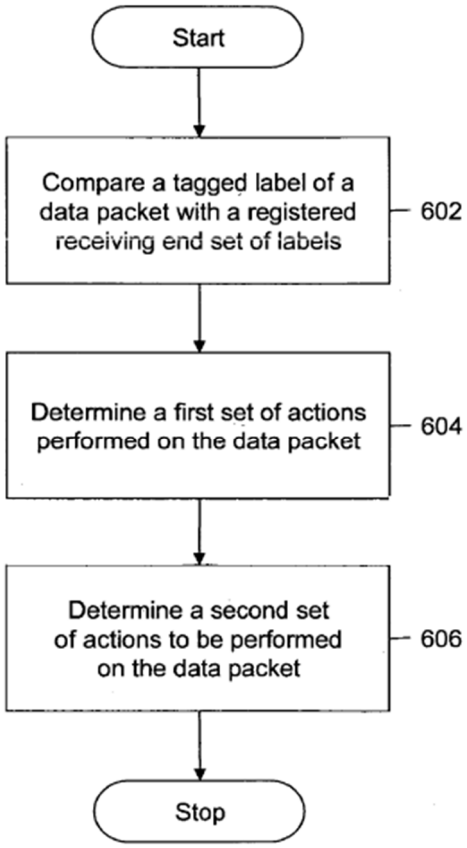


FIG. 7
(Annotated)

C. U.S. Patent Publication 2007/0204018 A1 ("Chandra"), filed February 24, 2006

Netskope Exhibit 1002



D. U.S. Patent Publication 2007/0237147 A1 ("Quinn"), filed April 7, 2006

Netskope Exhibit 1002

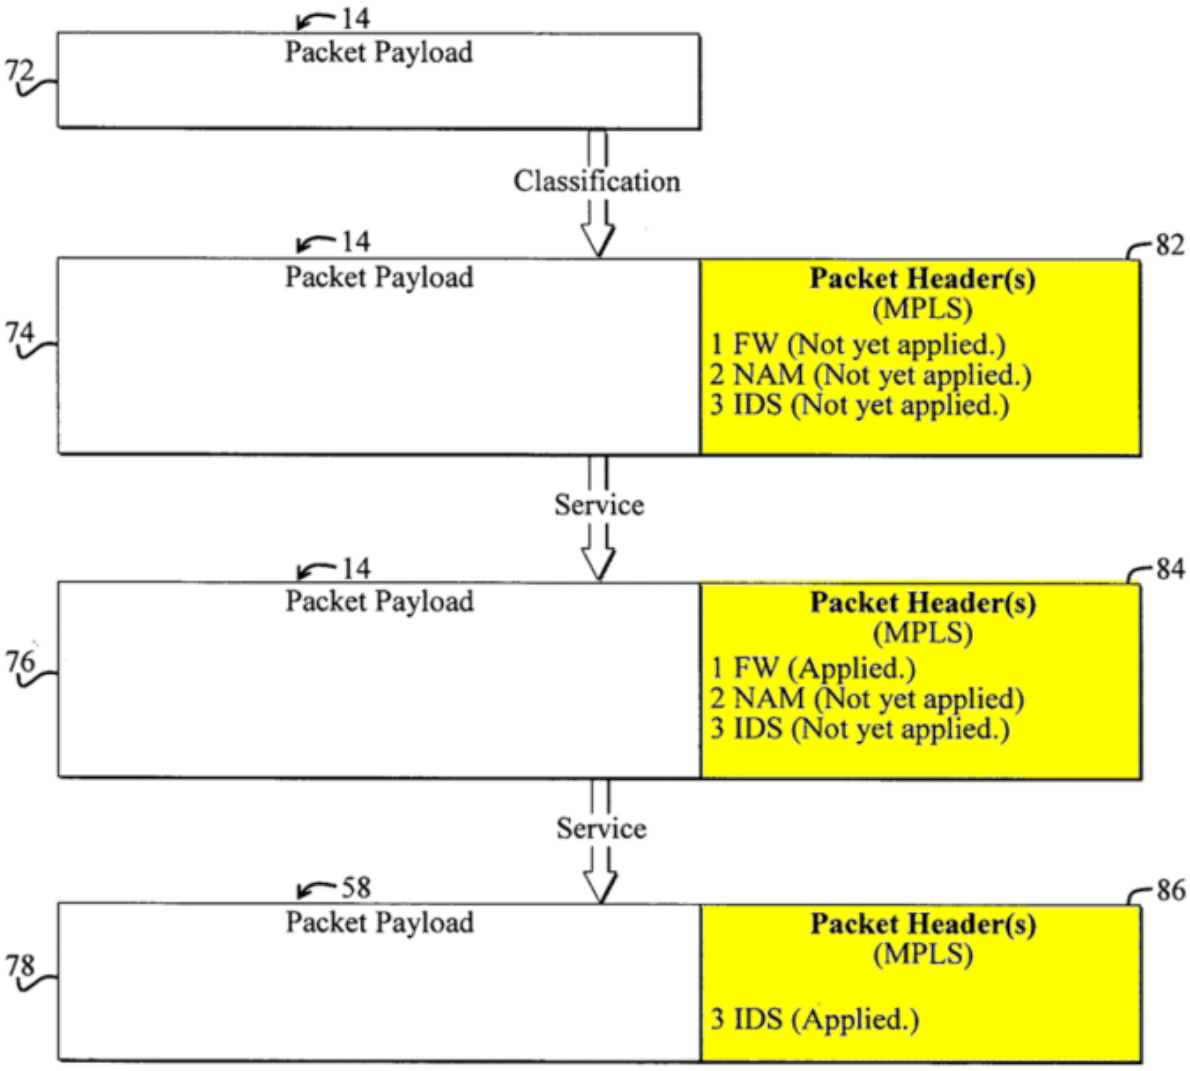


Fig. 2 (Annotated)

41. One known method for coordinating actions was to use information in packets. In particular, as people of ordinary skill in the art (and even many laypeople) would know, computer network packets typically are composed of a "header" and a "payload". A header contains the network routing information as well as space for

Netskope Exhibit 1002

additional information (such as "flags" and other fields). A payload contains the data content of the network packet.

42. The widely used TCP header format provides for a variable number of "option" bytes that can be added to a TCP packet header. These "option" bytes provide a simple mechanism for future extensibility of the format without losing backward compatibility with devices that are not aware of the extension. For example, such an extension might add extra data (such as additional flags and other fields) to TCP packets for a specific purpose. Downstream network devices that are unaware of the semantics of such option bytes are able to simply skip over them and hence remain backward compatible with the previous standard, while suitably programmed downstream network devices can decode the content of the option bytes and act on the extra information. Hence, option bytes provide a convenient channel to communicate additional information in a network to downstream devices without having to modify the payload and without losing backward-compatibility with devices that are unaware of these extensions.

43. The ability to use headers for coordinating actions was well-known over a decade before the '825 patent. For example, Newstadt, filed in 2005, disclosed placing a security marker in a packet header. The security marker could include a list of security technologies that had been applied to the network traffic. Alternatively, it could include a list of security technologies that had not been

Netskope Exhibit 1002

applied to the network traffic. The security marker could be authenticated using a public key signature scheme. The information in the header could be used by downstream devices to determine which security technologies had not yet been applied to the network traffic and which operations they still needed to perform.[2]

44. As another example, Quinn, filed in 2006, disclosed using packet headers to specify the actions to be performed on a packet as it progressed through network appliances and updating the packet header as it progressed through the different modules.[3]

45. In the context of these well-known techniques, the '825 patent describes just another application of using this conventional mechanism for encoding "flags" in the packet header:

> The flag may be preferably added to the Internet Protocol (IP) options field of the original network packet, although other fields, such as type of service (TOS) field, may also be used. The member node may extract the flag that is set by the other member node from these fields of the incoming packet.

('825 patent, 6:59-64.) "Flag" and "member node," in turn, correspond directly to information that were expressed on a traditional paper routing slip as in the 1943 example (above), in which "member node" corresponds to a party on the paper

[2] Newstadt, ¶¶ 29-34.
[3] Quinn, ¶¶ 39-50.

Netskope Exhibit 1002

routing slip and "flag" corresponds to the task and human signature associated with that party.

46. For example, the '825 patent states:

> At block 304, the member node may determine what network security operations have been performed on the incoming packet by the member nodes at previous hops. In one example, the member node may assume that all security operations that the member nodes at previous hops are capable of executing have been performed when a flag associated with a member node is present within the incoming packet. In another example, the flag is added to each packet to indicate that one or more security operations have been executed by the flagging member node and other member nodes along the transmission path do not need to execute the same operation again.

('825 patent, 7:33-45.) This mechanism is no different at all from the workflow coordination typically accomplished through routing slips by paper-based organizations eighty years ago. It describes exactly the division of labor that forms the core management principle of any competent organization. In my opinion, there is no innovative concept here at all.

3. Computing "Appliances"

47. The early 2000's also saw the emergence into widespread use of the term "appliance computing" and the related concept of an "appliance" (such as "computing appliance," "network appliance," "security appliance," "network storage appliance") to describe a device connected to a network that provides specific services, typically with a reduced administration and maintenance

Netskope Exhibit 1002

complexity than a general-purpose server would require. Many different types of network-attached computing devices that perform one or more network security function have been advertised as "(network) security appliances" over the course of the past decades; even the humble "firewall" is now frequently being marketed as a "network firewall appliance" (at the time of this writing, there are many such products available for purchase on Amazon).

C. The '825 Patent

48. The '825 patent concerns systems and methods for "coordinating security operations among members of a cooperative security fabric (CSF)…" ('825 patent, 2:3-5.)

49. Figure 1 (below) depicts an exemplary architecture according to an embodiment of the '825 patent. The architecture includes a private network (120) that has network appliances (e.g., network security appliances (121, 130a and 130b, 140a-140d)) and subnets (141a-141d). The exemplary architecture also includes subnet (151) and subnet network security appliance (150), which connect to private network (120) through an Internet-based virtual private network (VPN). (*Id.*, 4:58-67.) The '825 patent explains that "subnet 151 and subnet network security appliance 150 may be seen as a part of private network 110 [*recte* 120]." (*Id.*, 5:1-2.)

Netskope Exhibit 1002

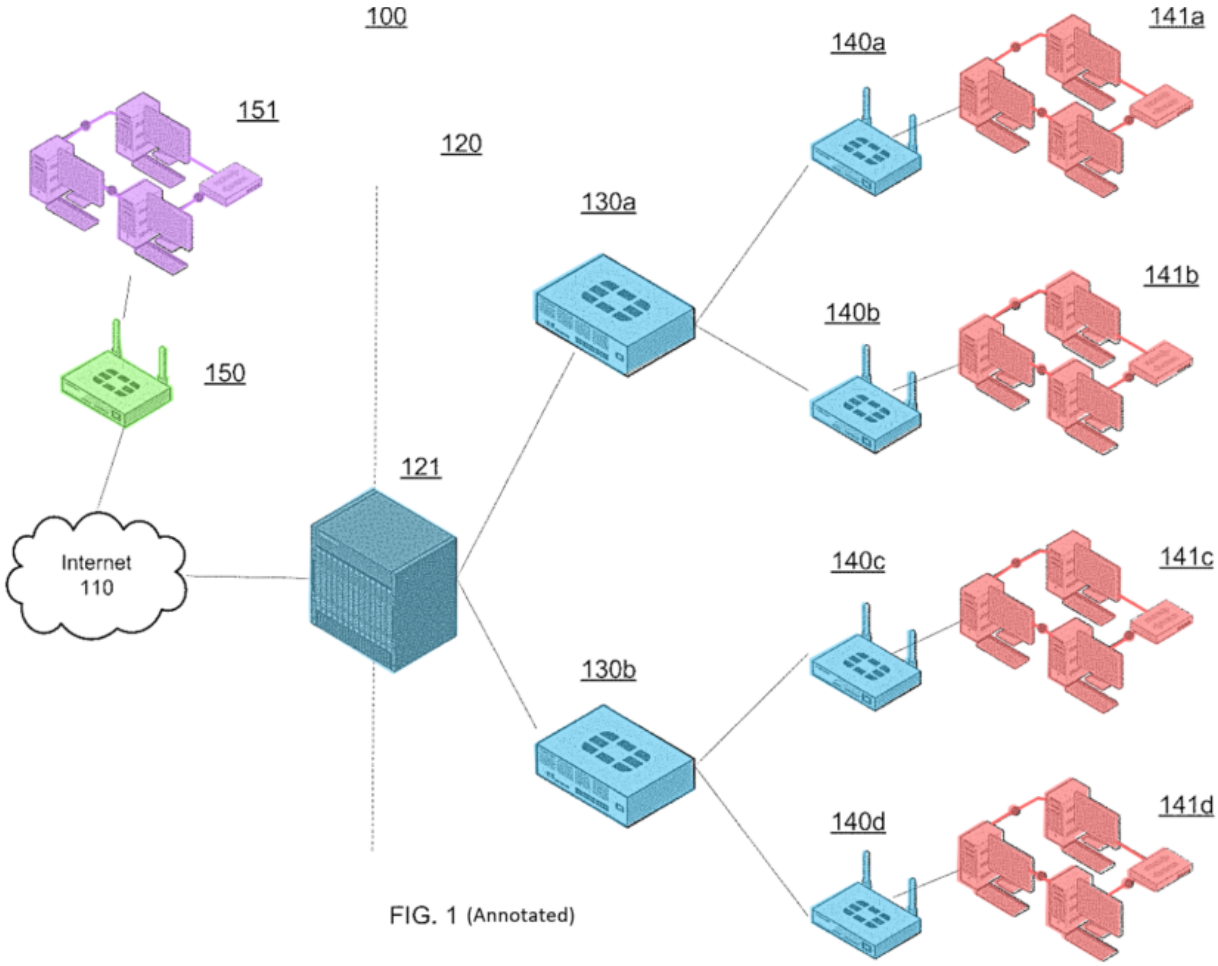


FIG. 1 (Annotated)

50. Figure 2 (below) depicts a topology of the network appliances shown in Figure 1. The '825 patent describes the network appliances' position in the architecture using a basic tree analogy: root node (220) corresponds to network security appliance (121), branch nodes (230a, 230b) correspond to network security appliances (130a, 130b) and leaf nodes (240a-240d) correspond to network security appliances (140a-140d). (*Id.*, 5:52-53, 60-32, 6:1-3.) The '825 patent explains that a POSITA would appreciate that the topology could include other combinations of root, branch, and leaf appliances. (*Id.*, 6:14-16.) For example, the topology could

Netskope Exhibit 1002

involve as few as two network security appliances, "such as the network path in which a leaf node 250 is connected to root node 220 directly." (*Id.*, 6:16-20.) Together, the network security appliances "form a cooperative security fabric that coordinates operations among members and avoids redundant operations." (*Id.*, 5:45-47.) The '825 patent does not define the term "CSF" or give it a reference numeral. Rather, the '825 patent describes the CSF simply as the result of appliances coordinating security operations performed on the network traffic passing through them.

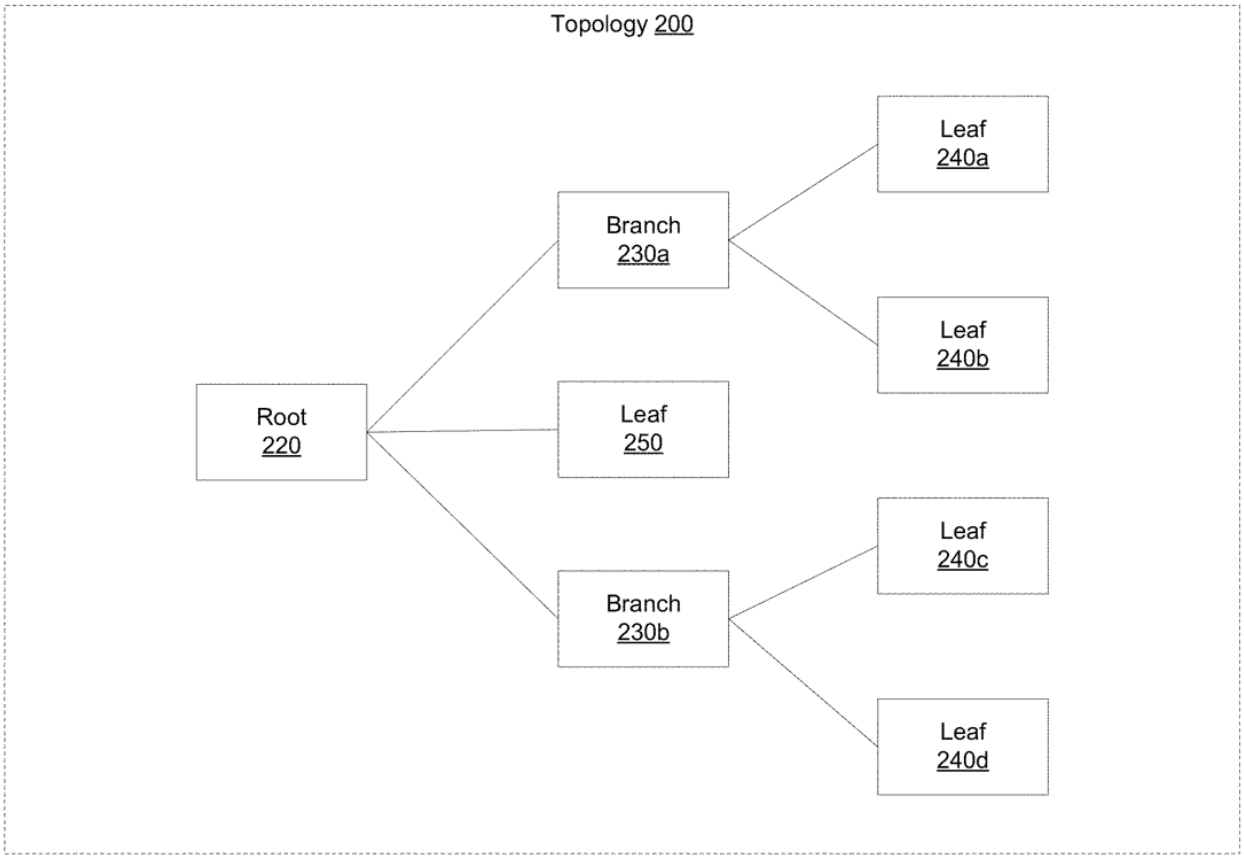


FIG. 2

Netskope Exhibit 1002

51. In operation, a network appliance "may add one or more flags (e.g., a single bit flag or a collection of bit fields) to a network packet after the network packet is inspected by the member node" and transmit the network packet with the flag "to another member node at the next hop." (*Id.*, 6:28-34.) The flag may be added to the header field of the network packet (similar to what Quinn, Newstadt, and Keohane did, as discussed above). (*Id.*, 9:29-32.) Once that next-hop member node receives the network packet, it "may check if the flag is carried by the network packet" and, if so, it "may determine to skip the operations that have been reported via the flag(s) as having been executed by the last member node." (*Id.*, 6:34-39.) This workflow coordination is the same mechanism used in the 1943 routing slip example and the Quinn, Newstadt, Keohane, and Chandra inventions discussed above. Figure 3 (below) illustrates this basic process flow. (*Id.*, 6:49-8:41.)

Netskope Exhibit 1002

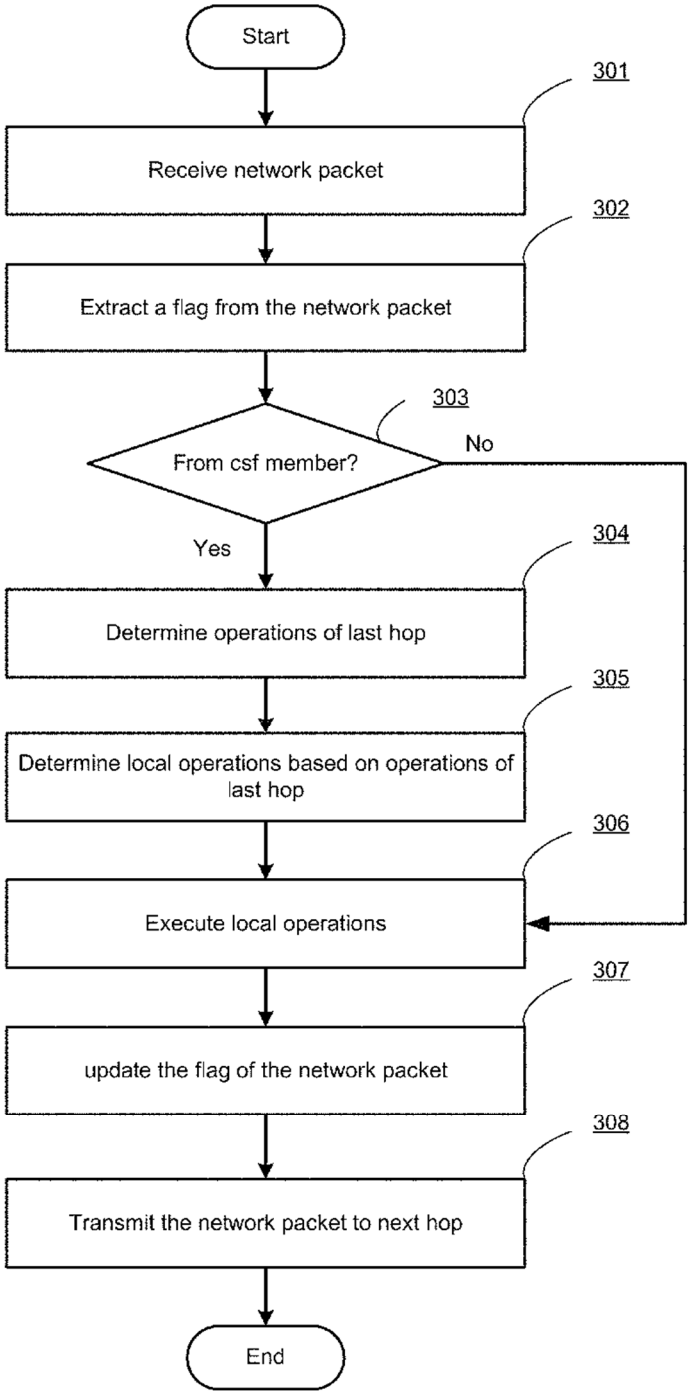


FIG. 3

Netskope Exhibit 1002

D. The Challenged Claims

52. I understand that Petitioner is challenging claims 1-7 and 10-31 of the ’825 patent.

E. Claim Construction

53. I understand that claim terms generally are construed in accordance with the ordinary and customary meaning they would have to a POSITA at the time of the invention in light of the claim language, the specification, and the prosecution history. I understand that dictionaries and other extrinsic evidence may be considered as well, though such evidence is typically regarded as less significant than the intrinsic record in determining the meaning of the claim language.

54. For all terms of the challenged claims of the ’825 patent, I have afforded them their ordinary and customary meaning in light of the intrinsic evidence at the time of the invention in May 2017.

III. UNPATENTABILITY OF THE ’825 PATENT CLAIMS

A. Standards for Invalidity

1. Obviousness

55. I am informed and understand that a patent cannot be properly granted for subject matter that would have been obvious to a person of ordinary skill in the art at the time of the alleged invention, and that a patent claim directed to such obvious subject matter is invalid under 35 U.S.C. § 103. It is also my understanding that in assessing the obviousness of claimed subject matter, one should evaluate

Netskope Exhibit 1002

obviousness in light of the prior art from the perspective of a person having ordinary skill in the art at the time the alleged invention was made (and not from the perspective of either a layman or a genius in that art). It is my further understanding that the question of obviousness is to be determined based on:

- The scope and content of the prior art;

- The difference or differences between the subject matter of the claim and the prior art (whereby in assessing the possibility of obviousness one should consider the manner in which a patentee and/or a Court has construed the scope of a claim);

- The level of ordinary skill in the art at the time of the alleged invention of the subject matter of the claim; and

- Any relevant objective factors (the "secondary indicia") indicating nonobviousness, including evidence of any of the following: commercial success of the products or methods covered by the patent claims; a long-felt need for the alleged invention; failed attempts by others to make the alleged invention; copying of the alleged invention by others in the field; unexpected results achieved by the alleged invention; praise of the alleged invention by the alleged infringer or others in the field; the taking of licenses under the patent by others and the nature of those licenses; expressions of surprise by experts and those skilled in the art at the subject matter of the claim; and whether the patentee proceeded contrary to accepted wisdom of the prior art.

Netskope Exhibit 1002

- Any relevant objective factors (the "secondary indicia") indicating obviousness: independent invention of the claimed invention by others before or at about the same time as the named inventor thought of it; and other evidence tending to show obviousness.

B. Ground I: Chandra renders claims 1-5, 12, 15, 16, 21, 22, and 27 obvious.

56. ***Chandra***: Over a decade before the '825 patent and much like it, Chandra recognized the problem of network appliances not coordinating operations and, consequently, causing redundant actions and wasted resources. (Chandra, ¶¶ 4-7.) Also much like the '825 patent, Chandra disclosed an invention addressing that problem in which network appliances coordinated operations using labels that passed between appliances. (Chandra, Abstract, ¶ 30.)

57. Chandra's Figure 2 (below) depicted an embodiment of its invention for obviating redundant actions. (*Id.*, ¶ 24.) The embodiment consisted of a transmitting end (202) and a receiving end (204), each of which included a network device (206, 212), a Customer Edge (CE) (208, 214), and a Provider Edge (PE) (210, 216). (*Id.*, ¶¶ 24-26.)

Netskope Exhibit 1002

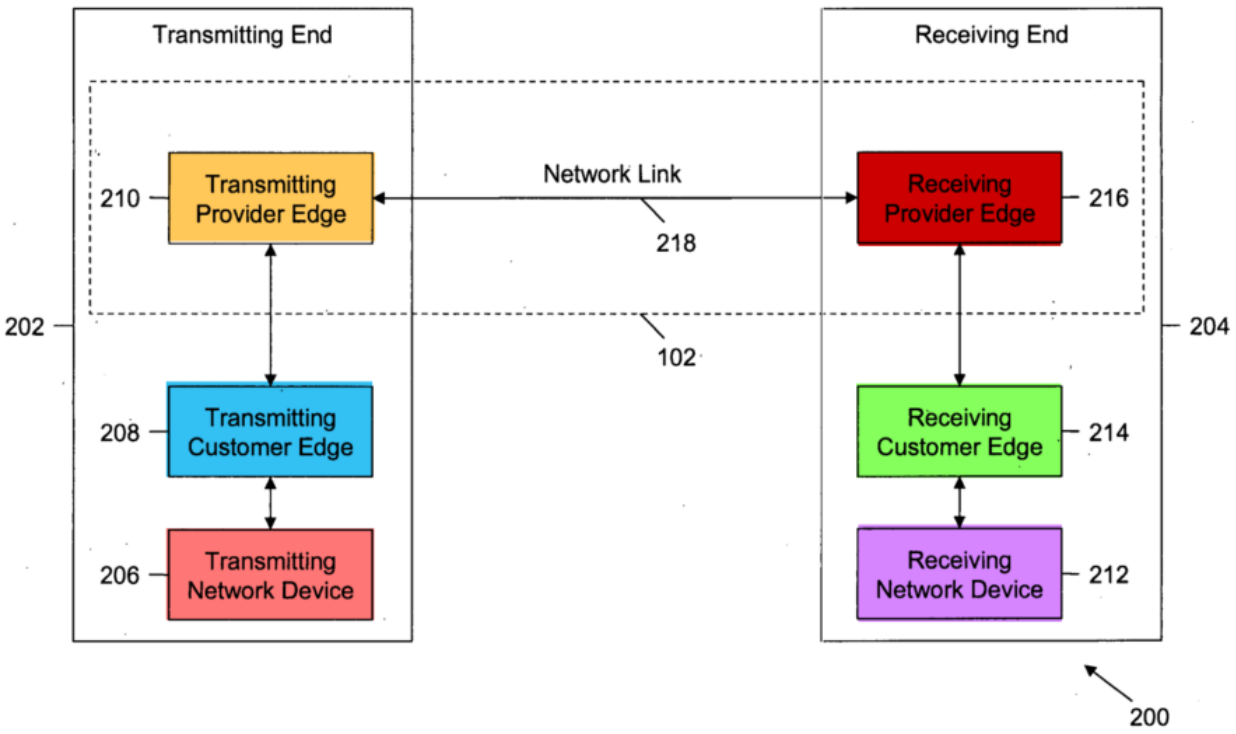


FIG. 2 (Annotated)

Transmitting CE (208) and receiving CE (214) were security-based devices, like "Intrusion Detection Systems (IDSs), fire walls, etc.," that performed security actions on a data packet. (*Id.*, ¶ 30.)

58. At setup, transmitting CE (208) and receiving CE (214) authenticated each other and exchanged an identical set of "CE labels to be used for tagging [] data packets." (*Id.*, ¶¶ 33, 34.) Transmitting CE (208) and receiving CE (214) then each registered and stored the labels in a table. (*Id.*, ¶ 52.) The tables below depicted exemplary labels used by transmitting CE (208) and receiving CE (214) to coordinate actions performed on a data packet. The receiving CE labels were associated with the address of transmitting CE (208).

Netskope Exhibit 1002

Transmitting CE labels

Label	Rules	Receiving CE
1256	Rules A and B	Address of receiving CE 214
1257	Rules C and D	Address of receiving CE 214
<untagged>	Rule E	Address of receiving CE 214
. . .		

Receiving CE labels

Label	Rules	Transmitting CE
1256	Rules A and B	Address of transmitting CE 208
1257	Rules C and D	Address of transmitting CE 208
<untagged>	Rule E	Address of transmitting CE 208
. . .		

59. (*Id.*, ¶¶ 51-52.) The transmitting CE used the labels to indicate a first set of actions it performed on network traffic, while the receiving CE used them to determine the first set of actions performed by the transmitting CE and determine the actions it must perform. (*Id.*)

60. In operation, transmitting network device (206) transmitted a packet to transmitting CE (208). Then, transmitting CE 208 "perform[ed] a first set of actions on the data packet, and tag[ged] the data packet with a CE label indicating the first set of actions that ha[d] been performed." (*Id.*, ¶ 31.) Transmitting CE (208) transmitted the tagged data packet to receiving CE (214) through transmitting PE (210) and receiving PE (216). (*Id.*, ¶¶ 31, 32.) Receiving CE (214) analyzed the label to identify the first set of actions performed by transmitting end (202) and determine a second set of actions to perform on the data packet. (*Id.*) The second

Netskope Exhibit 1002

set of actions could be redundant actions, non-redundant actions, or no action. (*Id.*, ¶ 18.)

61. For example, using the exemplary CE labels above, if transmitting CE (208) performed security actions corresponding to rules A and B on the data packet, it would tag the data packet with a label 1256. "When this tagged data packet reache[d] receiving CE [], receiving CE [] refer[red] to the receiving CE table, and determine[d] that actions corresponding to rules A and B ha[d] been performed on the data packet by transmitting CE []. Therefore, receiving CE [] [could] obviate th[o]se redundant actions." (*Id.*, ¶ 53.)

1. Claim 1

a. 1[pre]: "A method comprising:"

62. Chandra disclosed a "method that enable[d] the obviation of [] redundant actions by tagging a data packet with a label indicating a first set of actions performed on the data packet at a transmitting end." (*Id.*, ¶ 18.)

b. 1[a]: "receiving, by a first network security appliance of a private network that is participating in a cooperative security fabric (CSF), network traffic that is to be inspected by the first network security appliance;"

63. Chandra disclosed a receiving CE (first network security appliance) of a private network (e.g., a local area network or Virtual ***Private Network***) that coordinated security actions with other appliances and received and inspected data packets (network traffic).

Netskope Exhibit 1002

64. ***"first network security appliance of a private network"***: Chandra disclosed a first network security appliance in the form of a receiving CE.

65. Figure 2 (below) depicted a receiving CE (214) that could be a "security-based device." (Chandra, ¶ 30.) For example, receiving CE (214) could be "Intrusion Detection Systems (IDSs), fire walls, etc." (*Id.*, ¶ 30.) A POSITA would have found it obvious that IDSs and firewalls were network security appliances because they were well-known devices that performed network security operations. The '825 patent, for example, broadly discloses examples of network security appliances, including the Fortigate firewall device. ('825 patent, 4:16-18.)

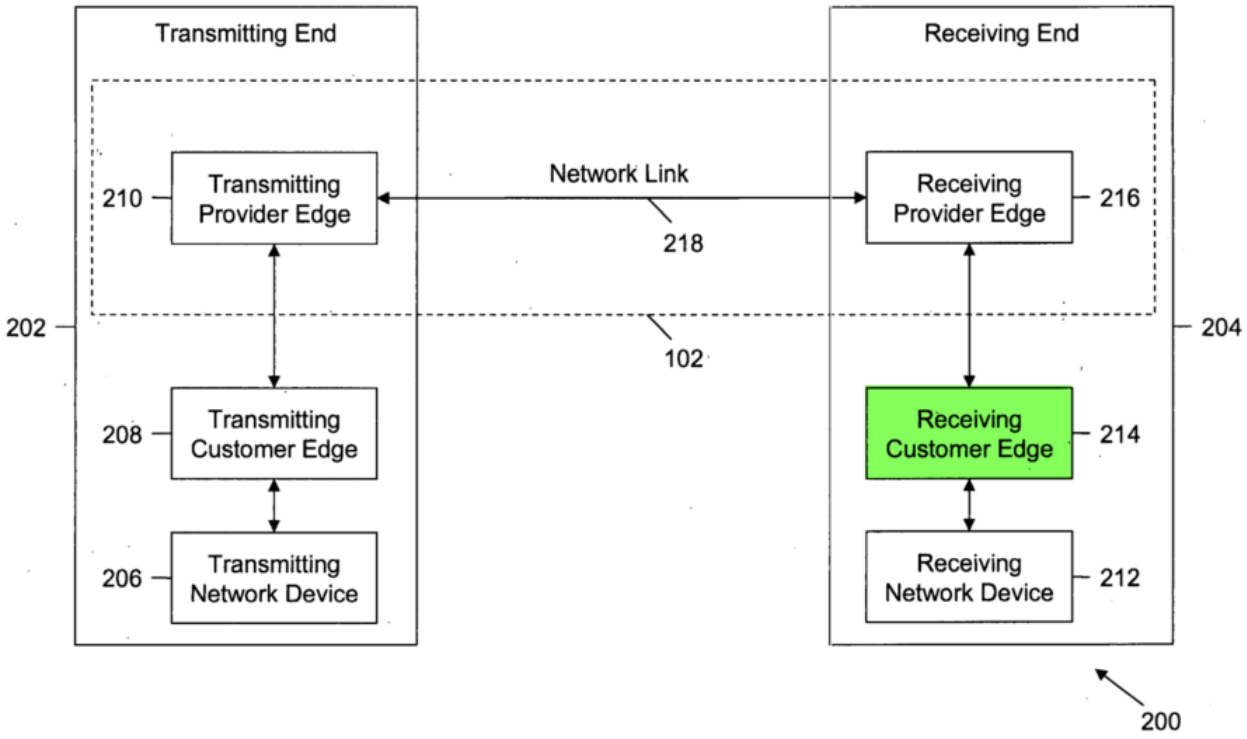


FIG. 2 (Annotated)

Netskope Exhibit 1002

66. Chandra also disclosed that the receiving CE could be implemented in a private network—for example, a local network or virtual private network. Chandra explained that network (102) (shown above) "provide[d] a communication medium to network devices [] for data to be transmitted from one network device to another network device" and could be, for example, a Local Area Network (LAN). (Chandra, ¶ 21.) A POSITA would have found it obvious that a LAN could be a private network because, long before the '825 patent, LANs (e.g., an enterprise network) were commonly setup as private networks using private network addresses. Moreover, while Figure 2 (above) did not expressly show the receiving CE within network (102), Chandra explained that "transmitting end 202 and receiving end 204 [could] be any identifiable part or region of network 102," which obviously includes receiving CE. (*Id.*, ¶ 29.)

67. Additionally, Chandra explained that its invention—including the receiving CE—could be implemented in a Virtual ***Private Network***. (*Id.*, ¶¶ 4-7, 58-68, Fig. 8.) Chandra explained that a VPN included "a security device, such as a firewall, to control access." (*Id.*, ¶ 5.) As an example, in Figure 8 (below), Chandra illustrated various elements of a system for obviating redundant actions in a Virtual Private Network (VPN), in accordance with an embodiment of its invention. Here, the transmitting and receiving CEs were parts of a private network that used an encrypted, private tunnel connecting two ends of communication.

Netskope Exhibit 1002

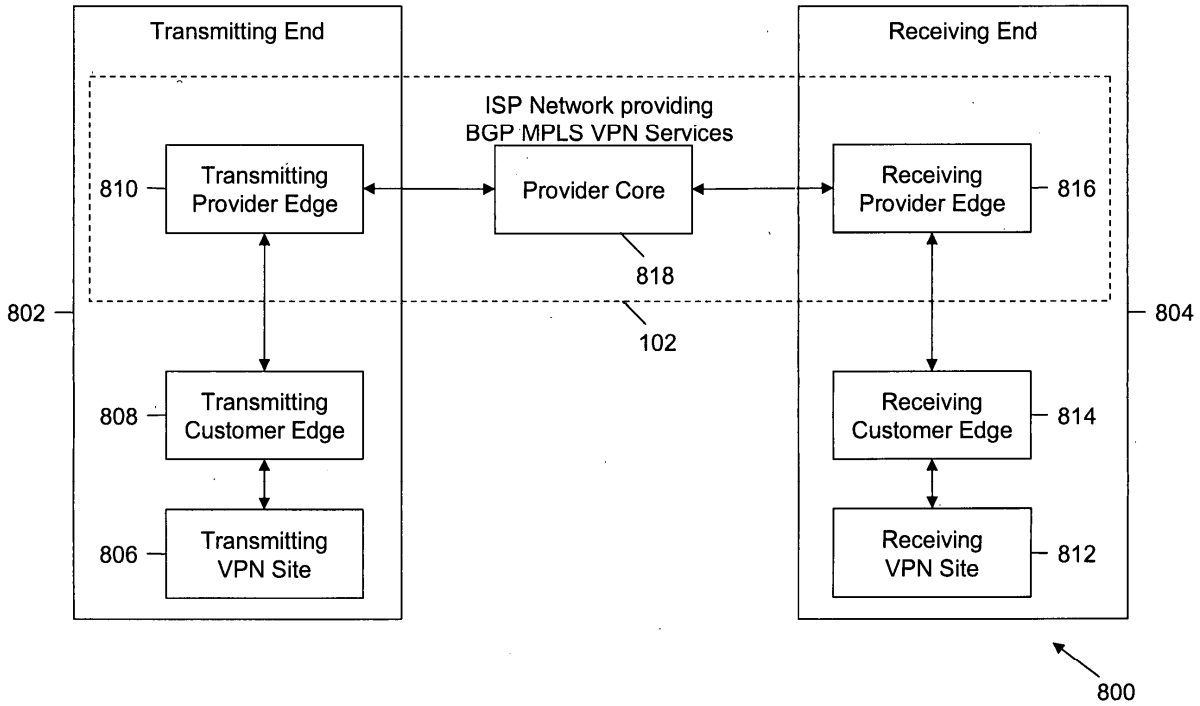


FIG. 8

68. The '825 patent similarly discloses that a VPN can be a private network. '825 Figure 1 (below) depicted "a subnet 151 and a subnet network security appliance 150 [] connected to private network 120 through the Internet 110 via VPN." ('825 patent, 4:65-67.) Subnet (151) and subnet network security appliance (150), according to the '825 patent, "may be seen as part of private network 110." (*Id.*, 5:1-2.) This is consistent with Chandra's disclosure. (Chandra, ¶ 4 ("The VPN is a network that uses a public network, like the Internet, to connect remote sites or users.").) For example, the transmitting VPN site could be a private enterprise network at the main corporate office and the receiving VPN site could be

Netskope Exhibit 1002

a remote employee's computer connecting back to the enterprise network or the subnet of a branch office connecting back to the main enterprise network, and vice versa. The receiving CE, like appliance 150 in '825 Figure 1, would be part of the private network. Thus, a POSITA would have found it obvious that a device (e.g., a receiving CE) operating in a VPN would be on a private network.

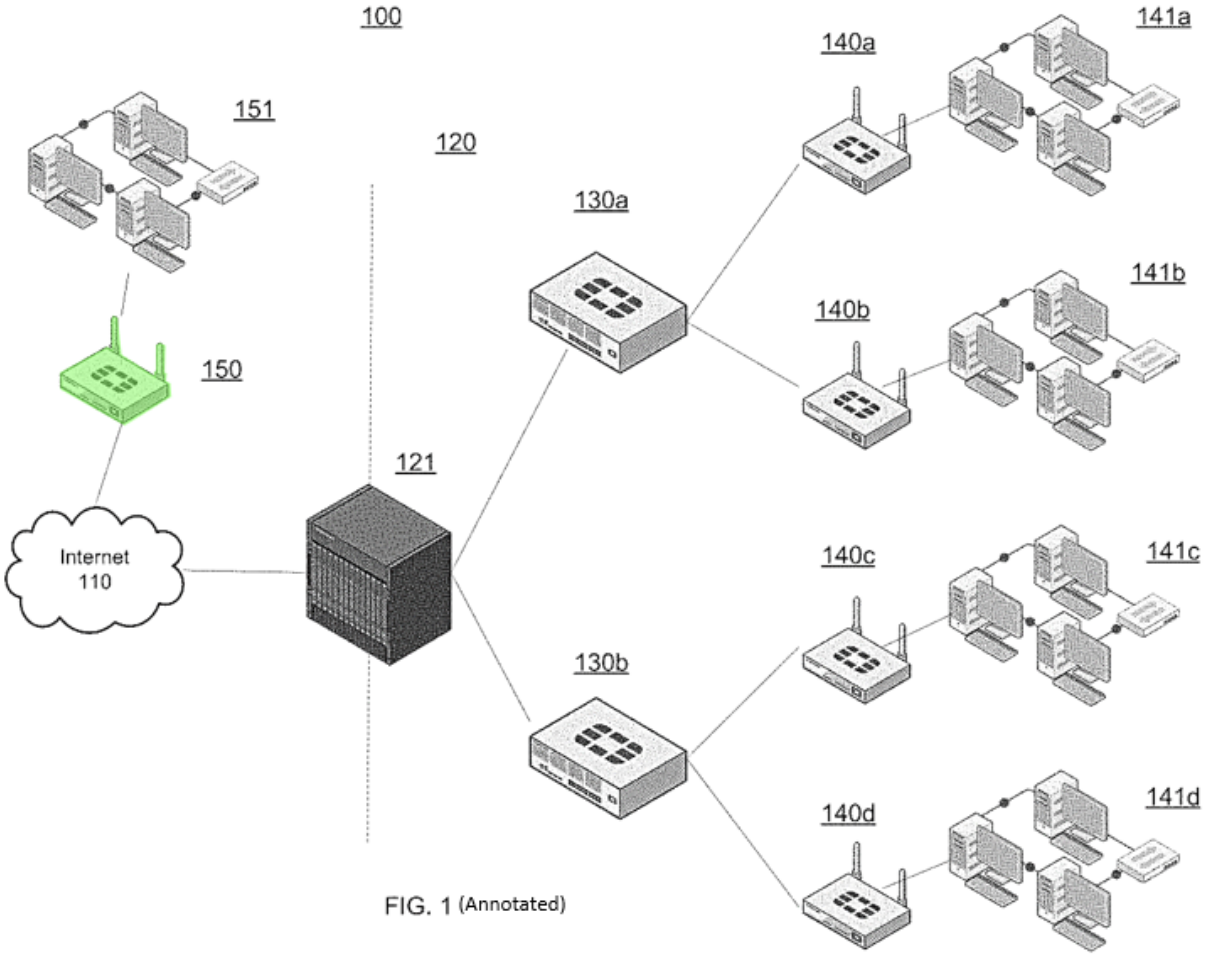


FIG. 1 (Annotated)

69. ***"participating in a cooperative security fabric (CSF)"***: Chandra disclosed that receiving CE (214) cooperated with transmitting CE (208) to coordinate security actions performed on network traffic. (*Id.*, Abstract, ¶¶ 18, 24-

Netskope Exhibit 1002

26, 30.) At setup, receiving CE (214) and transmitting CE (208) authenticated each other and exchanged labels that could be used to coordinate security actions. (*Id.*, ¶¶ 19-20, 33-34.) This exchange of information enabled coordinated enforcement of rules, which helped obviate redundant operations. (*Id.*, ¶¶ 18-20.)

70. These disclosures in Chandra line-up with the '825 patent's discussion about the CSF. The '825 describes the CSF primarily by what it achieves—it "coordinates operations among members and avoids redundant operations." ('825 patent, 5:45-47, 1:64-68.) As discussed, that is precisely what the cooperating transmitting CE and receiving CE achieve. Moreover, the '825 patent does not set a minimum number of appliances needed to form a CSF and, in fact, explains that as few as two are needed. (*See, e.g.*, *id.*, 6:18-20.) Therefore, Chandra disclosed that the receiving CE participated in a CSF.

71. ***"receiving…network traffic that is to be inspected by the first network security appliance"***: Chandra disclosed that the receiving CE received network traffic (data packets) before *inspecting* it. (Chandra, ¶¶ 32 ("Next, receiving PE 216 ***transmits the tagged data packet to receiving CE 214***. Receiving CE 214 pops off the tagged CE label. Based on this CE label, receiving CE 214 identifies the first set of actions. According to the identified first set of actions, receiving CE 214 determines a second set of actions that has to be performed on the data packet.") (emphasis added), 37, 46 ("***On receiving the data packet***…receiving CE 214

Netskope Exhibit 1002

compare[d] the CE label tagged on the data packet with the registered receiving end set of labels.") (emphasis added).)

c. 1[b]: "determining, by the first network security appliance, whether the network traffic has been transmitted from a second network security appliance that is participating in the CSF based on a flag carried by one or more packets of the network traffic;"

72. Chandra disclosed that the receiving CE (first network security appliance) determined whether network traffic had been transmitted from a transmitting CE (second network security appliance) using a label (flag) carried by the packet.

73. Before addressing how Chandra maps to this limitation, it is helpful to understand what happened to a packet before it reached the receiving CE. As explained above, at setup, transmitting and receiving CEs authenticated each other and exchanged labels for coordinating security operations. (Chandra, ¶¶ 31-35.) The CE labels were exchanged and registered because "receiving CE 214 need[ed] to recognize the CE label." (*Id.*, ¶ 47.)

74. Once transmitting CE (208) received network traffic from a transmitting network device (206), it performed a first set of actions on the traffic "and tag[ged] the data packet with a CE label indicating the first set of actions that ha[d] been performed." (*Id.*, ¶ 31.) The network traffic then travelled from

Netskope Exhibit 1002

transmitting CE (208) to transmitting PE (210) to receiving PE (216) and, finally, to receiving CE (214). (*Id.*, ¶¶ 31-32.)

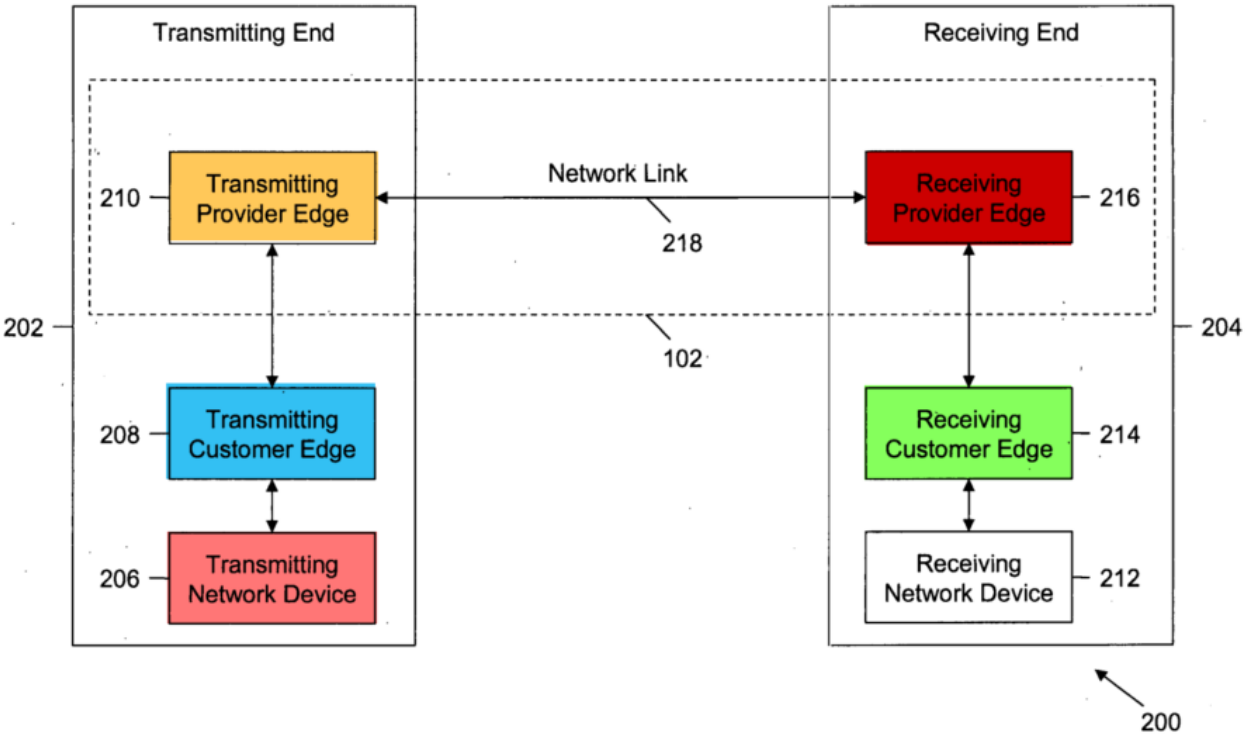


FIG. 2 (Annotated)

75. Returning to how Chandra disclosed this limitation. Once the receiving CE received network traffic, it "checked ***if the data packet ha[d] been tagged with a CE label*** present in the registered receiving end set of labels." (*Id.*, ¶ 48 (emphasis added); *see also, e.g.*, ¶¶ 46, FIGs. 6-7.) This was depicted in step 602, below.

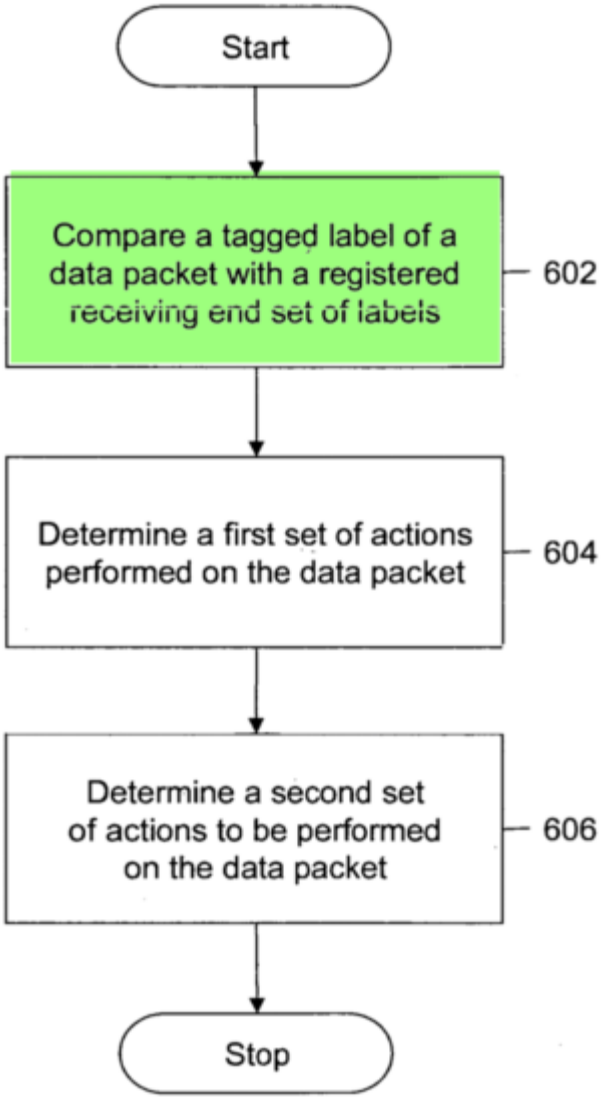


FIG. 6

76. In other words, the receiving CE determined if the packet had a CE label that corresponded to a label it had previously exchanged, after authentication, with the transmitting CE (a second security appliance participating in the CSF). (Id.,

Netskope Exhibit 1002

¶ 47-48.) This was done because network traffic could have "not been tagged with a CE label or the tagged CE label [was] not recognizable by receiving CE 214," in which case a separate set of actions was done. (*Id.*, ¶¶ 47-49; *see also, e.g.*, *id.*, ¶ 37 ("At step 306, receiving CE 214 identifie[d] the CE label. This CE label [was] present in the registered receiving end set of labels, and therefore, [could] be recognized by receiving CE 214.").) Thus, because each CE label used by transmitting CE (208) to tag the data packet was registered at receiving CE (214), a POSITA would find it obvious that when receiving CE (214) checked the data packet for a label (flag), the presence of a registered CE label in the data packet indicated that the data packet had been transmitted from transmitting CE (208). This is consistent with the '825 patent's disclosure that a flag's presence on network traffic serves as "an indicator to subsequent member nodes along the transmission path that…the incoming packet was transmitted from the member node." ('825 patent, 6:54-59, 6:65-67.) Receiving CE (214) could also look up the CE label, which was associated with the address of transmitting CE (208), in the registered label table (below) to further confirm that the data packet was transmitted from transmitting CE (208).

77. Additionally, the CE label was a number such as "1256" or "1257." (*Id.*, ¶ 52.) A POSITA would find it obvious that the CE label was a flag because

Netskope Exhibit 1002

the ’825 patent described the recited “flags” broadly, e.g., “a single bit flag or a collection of bit fields.” (’825 patent, 6:28-34.)

d. 1[c]: “when the network traffic is determined to have been transmitted from the second network security appliance: determining, by the first network security appliance, one or more network security operations that are executed by the second network security appliance;”

78. Chandra disclosed that, once the receiving CE (first network security appliance) determined that the data packet (network traffic) had been transmitted from the transmitting CE (second network security appliance), it determined the security operations performed by the transmitting CE. Specifically, for step 604 in Figure 6 (below), Chandra disclosed that “the ***first set of actions performed on the data packet at transmitting CE 208 [was] determined***.” (*Id.*, ¶ 46 (emphasis added); *see also, e.g.*, *id.*, ¶ 32 (“Based on this CE label, receiving CE 214 identifie[d] the first set of actions.”).)

Netskope Exhibit 1002

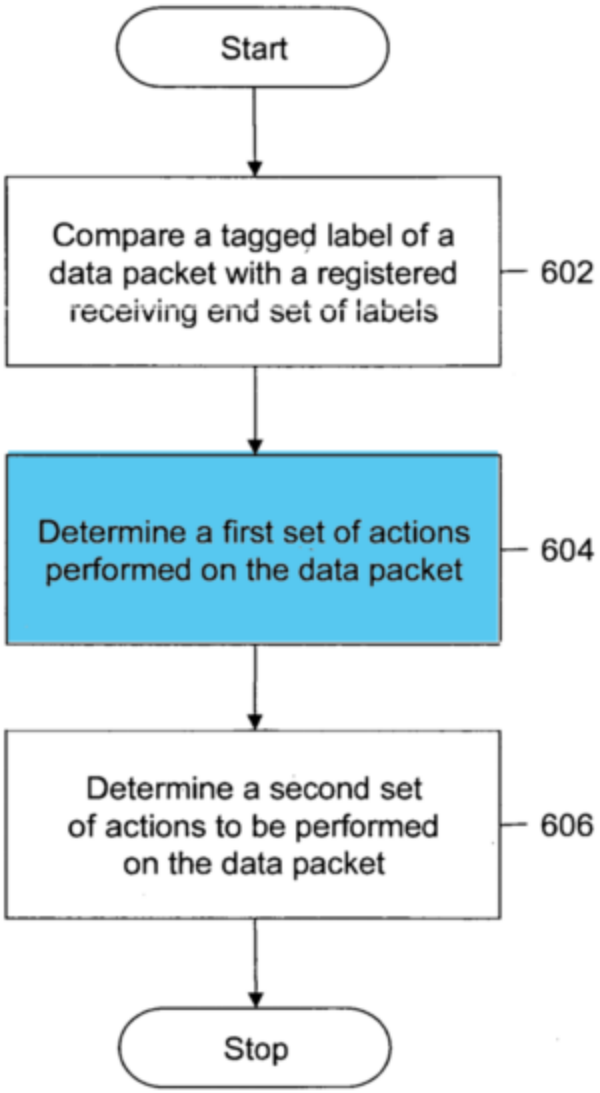


FIG. 6 (Annotated)

79. Chandra also disclosed that if the data packet was not tagged with a CE label or was tagged with an unknown label, a third set of actions was performed on the data packet. (Chandra, ¶ 48.)

Netskope Exhibit 1002

e. 1[d][i]: "determining, by the first network security appliance, one or more local network security operations that are to be executed by the first network security appliance based on the one or more network security operations that are executed by the second network security appliance,"

80. Chandra disclosed that the receiving CE (first network security appliance) determined the actions it would perform based on the actions performed by the transmitting CE (second network security appliance).

81. At step 606 in Figure 6 (below), receiving CE (214) determined the actions it would perform on the data packet based on the first set of actions performed by transmitting CE (208). (*Id.*, ¶ 46 ("At step 606, the second set of actions that ha[d] to be performed on the data packet at receiving end 204 [was] determined."); *see also, e.g.*, *id.*, ¶¶ 32 ("***According to the identified first set of actions***, ***receiving CE 214 determine[d] a second set of actions that ha[d] to be performed on the data packet***.") (emphasis added), 18.)

82. A POSITA would have found it obvious that because the second set of actions were locally determined and performed at and by receiving CE (214). they were "local" network security operations.

Netskope Exhibit 1002

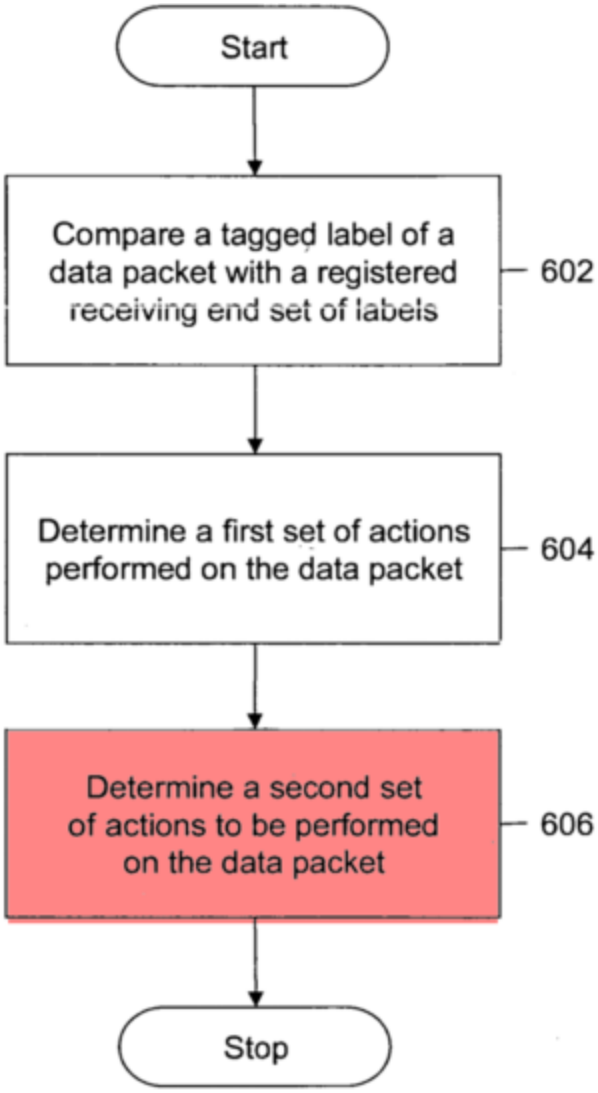


FIG. 6 (Annotated)

f. 1[d][ii]: "including skipping, by the first network security appliance, the one or more network security operations that are executed by the second network security appliance;"

83. Chandra disclosed the receiving CE (first network security appliance) not performing (skipping) the one or more network security operations that were executed by the transmitting CE (second network security appliance). (*Id.*, ¶ 37 ("The second set of actions [to be performed by receiving CE] ***includes non-redundant actions that are not present in the first set of actions*** [already performed by transmitting CE]. Therefore, ***redundant actions are obviated and not performed at receiving CE []***.") (emphasis added).)

84. For example, based on the exemplary receiving CE labels shown below, if receiving CE (214) received a data packet tagged with the CE label "1256," it would determine "that actions corresponding to rules A and B ha[d] been performed on the data packet by transmitting CE 208. Therefore, ***receiving CE 214 [could] obviate th[o]se redundant actions***." (*Id.*, ¶ 53 (emphasis added).)

Label	Rules	Transmitting CE
1256	Rules A and B	Address of transmitting CE 208
1257	Rules C and D	Address of transmitting CE 208
<untagged>	Rule E	Address of transmitting CE 208
. . .		

g. 1[e]: "and executing, by the first network security appliance, the one or more local operations on the network traffic."

85. Chandra disclosed that the receiving CE (first network security appliance) executed the second set of actions (one or more local operations) that it had determined earlier on the data packet (network traffic). (*Id.*, ¶ 46 ("Once

Netskope Exhibit 1002

determined, the ***second set of actions [was] performed on the data packet at receiving CE*** [].") (emphasis added).). As before, the term "local" refers to the operations executed at and by a specific device, i.e., in this case, the receiving CE (first network security appliance).

2. Claim 2

a. 2[a]: "The method of claim 1, wherein said determining, by the first network security appliance, one or more network security operations of the first network security appliance to be executed by the first network security appliance based on the one or more network security operations that are executed by the second network security appliance"

86. *See* element 1[d][i].

b. 2[b]: "further comprises determining, by the first network security appliance, one or more network security operations that are to be executed by the first network security appliance except the one or more network security operations that are executed by the second network security appliance based on local network polices."

87. Chandra disclosed that the receiving CE (first network security appliance) determined one or more actions to perform, obviating the one or more actions that are executed by the transmitting CE (second network security appliance) based on local network polices.

88. Chandra disclosed that the "[t]he second set of actions [performed by receiving CE] ***include[d] non-redundant actions*** that [were] not present in the first

Netskope Exhibit 1002

set of actions [performed by transmitting CE]." (Chandra, ¶ 37 (emphasis added); *see also*, *e.g.*, ¶¶ 18, 35.)

89. Returning to the example in limitation 1[d][ii], receiving CE (214) determined that actions corresponding to rules A and B had been performed by transmitting CE (208) and skip those actions. (*Id.*, ¶ 53.) A POSITA would have found it obvious from this disclosure that receiving CE (214) could then perform actions that had not been performed by transmitting CE (214), such as actions corresponding to rule C and/or rule D. (*See, e.g.*, *id.*, ¶ 35 ("Since the ***second set of actions include[d] non-redundant actions***, redundant actions [were] obviated.") (emphasis added).)

Label	Rules	Transmitting CE
1256	Rules A and B	Address of transmitting CE 208
1257	Rules C and D	Address of transmitting CE 208
<untagged>	Rule E	Address of transmitting CE 208
. . .		

90. ***"local network polices"***: It is unclear from the claim language whether the "based on local network policies" limitation applies to the first network security appliance's operations, or the second network security appliance's operations, or both. Chandra disclosed this limitation according to any interpretation. A POSITA would find it obvious that in Chandra, there were two sets of local policies: one set

Netskope Exhibit 1002

local to the receiving CE and a second set local to transmitting CE. Each set of policies were stored at and executed at and by the respective CE.

3. Claim 3

a. 3[a]

91. *See* element 2[a].

b. 3[b]: "further comprises determining, by the first network security appliance, that the one or more network security operations that are executed by the second network security appliance are to be executed by the first network security appliance."

92. Chandra disclosed that the receiving CE (first network security appliance) determined one or more actions to perform that are executed by the transmitting CE (second network security appliance).

93. Chandra disclosed that "the second set of actions to be performed at the receiving end could include "***nonredundant actions that ha[d] not been performed on the data packet at the transmitting end***" and "***[could] include redundant actions that ha[d] been performed*** on the data packet at the transmitting end." (*Id.*, ¶ 18 (emphasis added).) While skipping some actions already performed reduced redundancy, in some circumstances, repeating some actions already performed, such as running another virus scan of the same type, might be appropriate or simply more convenient. For example, a downstream device might perform certain actions that are so inexpensive to execute that for these actions, it may not be worth performing the redundancy analysis at all as it may be easier to just always execute this cheap

Netskope Exhibit 1002

action regardless. As another example, a broad scan performed by an upstream security appliance might contain elements that have a partial overlap with a more focused scan performed downstream. Being able to simply repeat any redundant action without having to carefully design a downstream scan that explicitly omits the partial work already performed upstream makes the overall configuration much simpler. Consequently, a POSITA would have found it obvious that the actions performed by receiving CE could include both redundant and nonredundant actions. In addition, in general, a POSITA would have found it entirely obvious that at any point one could repeat redundant actions. This is just like the earlier paper-based example from 1943 cited above in which any recipient of a package with a routing slip can elect to perform any action "locally", including any action that has been flagged as having already been performed by a previous recipient along the route.

4. Claim 4

a. "The method of claim 1, further comprising when the network traffic is not transmitted from the second network security appliance, executing, by the first network security appliance, one or more local network security operations specified by one or more local policies."

94. Chandra disclosed—or at least suggested—that if the receiving CE (first network security appliance) received a data packet that was not transmitted from the transmitting CE (second network security appliance), the receiving CE performed at least one local action specified by a local policy.

Netskope Exhibit 1002

95. Chandra disclosed that when "the data packet ha[d] not been tagged with a CE label or the tagged CE label [was] not recognizable by receiving CE 214…a third set of actions c[ould] be performed on the data packet." (Chandra, ¶ 47). It would be obvious to a POSITA that the absence of specific labels associated with the transmitting CE might be indicative of traffic not being transmitted from the transmitting CE, and it would also be obvious that the third set of actions executed by the receiving CE would be one or more local network security operations specified by one or more local policies.

96. In Figure 7 (below), Chandra disclosed that the receiving CE determined if a received data packet had a label that matched the registered set of labels. (*Id.*, ¶¶ 47-48, 53.) If the data packet did not have a recognized label (e.g., the traffic originated from the Internet), the receiving CE identified the packet as non-secure and performed a third set of actions on the packet. (*Id.*) Because each CE label used by the transmitting CE to tag the data packet was registered at the receiving CE, Chandra taught or suggested that when the receiving CE checked the data packet at step 702, the presence of a registered CE label in the data packet indicated that the data packet was transmitted from the transmitting CE, and an unregistered CE label or an absence of CE label indicated that the data packet was not transmitted from the transmitting CE.

Netskope Exhibit 1002

5. Claim 5

a. "The method of claim 1, wherein the flag indicates the network security operations executed by the second network security appliance."

97. Chandra disclosed that the CE label (flag) indicated the actions executed by the transmitting CE (second network security appliance). Specifically, Chandra disclosed that, once transmitting CE (208) performed a first set of actions on a packet, "it tag[ged] the data packet with a CE label ***indicating the first set of actions that ha[d] been performed***." (Chandra, ¶ 31 (emphasis added).)

98. A POSITA would have found it obvious that Chandra's "actions" corresponded to "network security operations."

6. Claim 12

a. "The method of claim 1, wherein the one or more packets in which the flag is carried comprises a first packet of a session."

99. While Chandra did not expressly disclose that the packets in which the flag was carried comprised a first packet of a session, a POSITA would have found that limitation obvious based on Chandra's disclosures and the knowledge of a POSITA.

100. The most commonly used internet protocol, TCP/IP is "connection-oriented": a connection needs to be explicitly established between a client and a server before data can be sent. There is also a mechanism to explicitly close connections and thereby free up resources. The sequence of interactions between

Netskope Exhibit 1002

any two communicating parties, starting with the opening of a connection, encompassing any subsequent data exchange between these parties, and ending with the closing of the connection, is commonly also called a "session".

101. It would have been obvious to a POSITA that for the specific objective of communicating a flag or set of flags between a connection initiator and the connection end point of a TCP/IP connection, including the flag(s) in the first packet would be sufficient since all of the packets in a session travel back and forth between these two end points. They do not, however, necessarily travel along the same route between these end points. A POSITA would have known that if there is just a single route between the two end points of a connection, then all subsequent packets in the connection will travel along the same path as the first packet of the session, whereas if there are multiple alternative routes, the path may be different but the end points will be the same. Hence, it would have been obvious to a POSITA that for connections without alternative routes, the flags would actually be redundant in packets in a session after the initial packet.

7. Claim 15

a. 15[pre]-15[d][i]

102. *See* limitations 1[pre]-1[d][i].

Netskope Exhibit 1002

b. 15[d][ii]: "including determining one or more network security operations that are to be executed by the first network security appliance except the one or more network security operations that are executed by the second network security appliance based on local network polices;"

103. *See* limitation 2[b]. For example, Chandra disclosed that the "[t]he second set of actions [performed by receiving CE] ***include[d] non-redundant actions*** that [were] not present in the first set of actions [performed by transmitting CE]." (Chandra, ¶ 37 (emphasis added); *see also*, *e.g.*, ¶¶ 18, 35.)

c. 15[e]

104. *See* limitation 1[e].

8. Claim 16

a. "The method of claim 15, wherein the flag indicates the network security operations executed by the second network security appliance."

105. *See* claim 5.

9. Claim 21

a. 21[pre]-21[d][i]

106. *See* limitations 1[pre]-1[d][i].

b. 21[d][ii]: "including determining, by the first network security appliance, that the one or more network security operations that are executed by the second network security appliance are to be executed by the first network security appliance;"

107. *See* limitation 3[b]. For example, Chandra disclosed that "the second set of actions [to be performed at the receiving end] ***[could] include redundant***

Netskope Exhibit 1002

actions that ha[d] been performed on the data packet at the transmitting end." (*Id.*, ¶ 18 (emphasis added).) As stated before, a POSITA would have found it entirely obvious that at any point one could repeat redundant actions. This is just like the earlier paper-based example from 1943 cited above in which any recipient of a package with a routing slip can elect to perform any action "locally", including repeating any action that has been flagged as having already been performed by a previous recipient along the route.

c. 21[e]

108. *See* limitation 1[e].

10. Claim 22

a. "The method of claim 21, wherein the flag indicates the network security operations executed by the second network security appliance."

109. *See* claim 5.

11. Claim 27

a. 27[pre]-27[b][i]

110. *See* limitations 1[pre]-1[b].

b. 27[b][ii]: "wherein the flag indicates the network security operations executed by the second network security appliance;"[4]

111. *See* claim 5.

---

[4] "[T]he network security operations" in this claim lacks the proper antecedent basis and is assumed to refer to the network security operations that the second network security appliance

Netskope Exhibit 1002

c. 27[c]-27[e]

112. *See* limitations 1[c], 1[d][i], and 1[e].

C. Ground 2: Chandra and Buruganahalli render claims 6, 7, 17, 18, 23, 24, 28, and 29 obvious.

113. ***Buruganahalli***: Much like the '825 patent, Buruganahalli disclosed a method for network security devices to coordinate security actions. Specifically, Buruganahalli disclosed that, once a network device performed actions on traffic, it used a metadata tag carried by network traffic to indicate the actions performed by the device, thereby eliminating the need to repeat those actions. (*See, e.g.*, Buruganahalli, ¶¶ 26-30.)

114. Specifically, Buruganahalli's Figure 2 (below) depicted network (18) that included various network devices, each of which "may be configured to perform complex processing…on network data [] per packet basis looking for malicious content." (Buruganahalli, ¶ 26.) The network devices could be network security devices "such as, for example, an intrusion protection system, a data loss prevention system, a fire wall, a load balancer," etc. (*Id.*, ¶ 30.)

executes on the data packet before transmitting the data packet to the first network security appliance.

Netskope Exhibit 1002

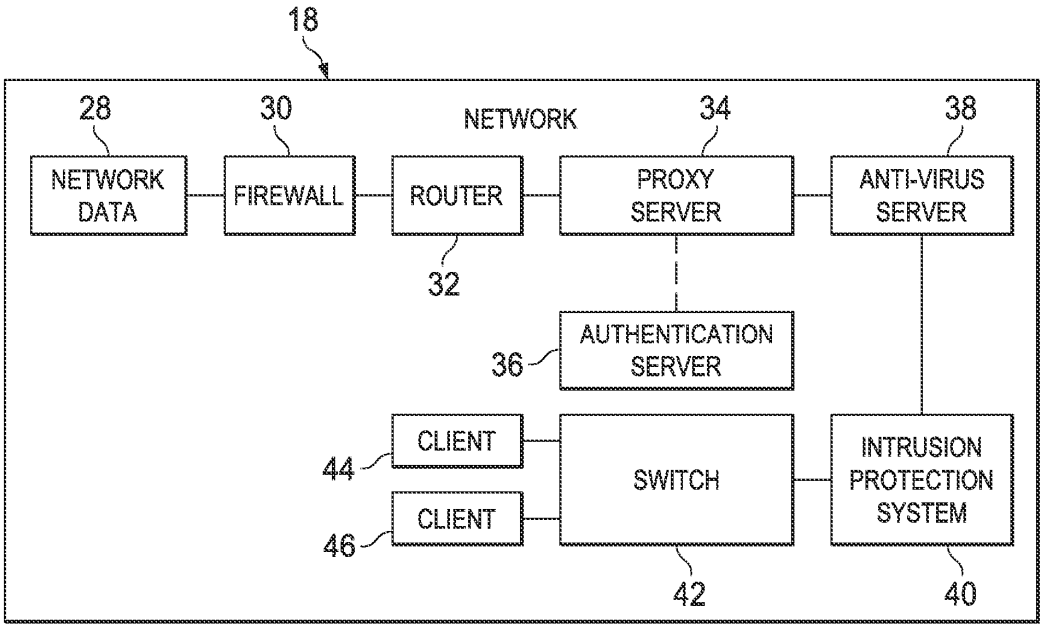
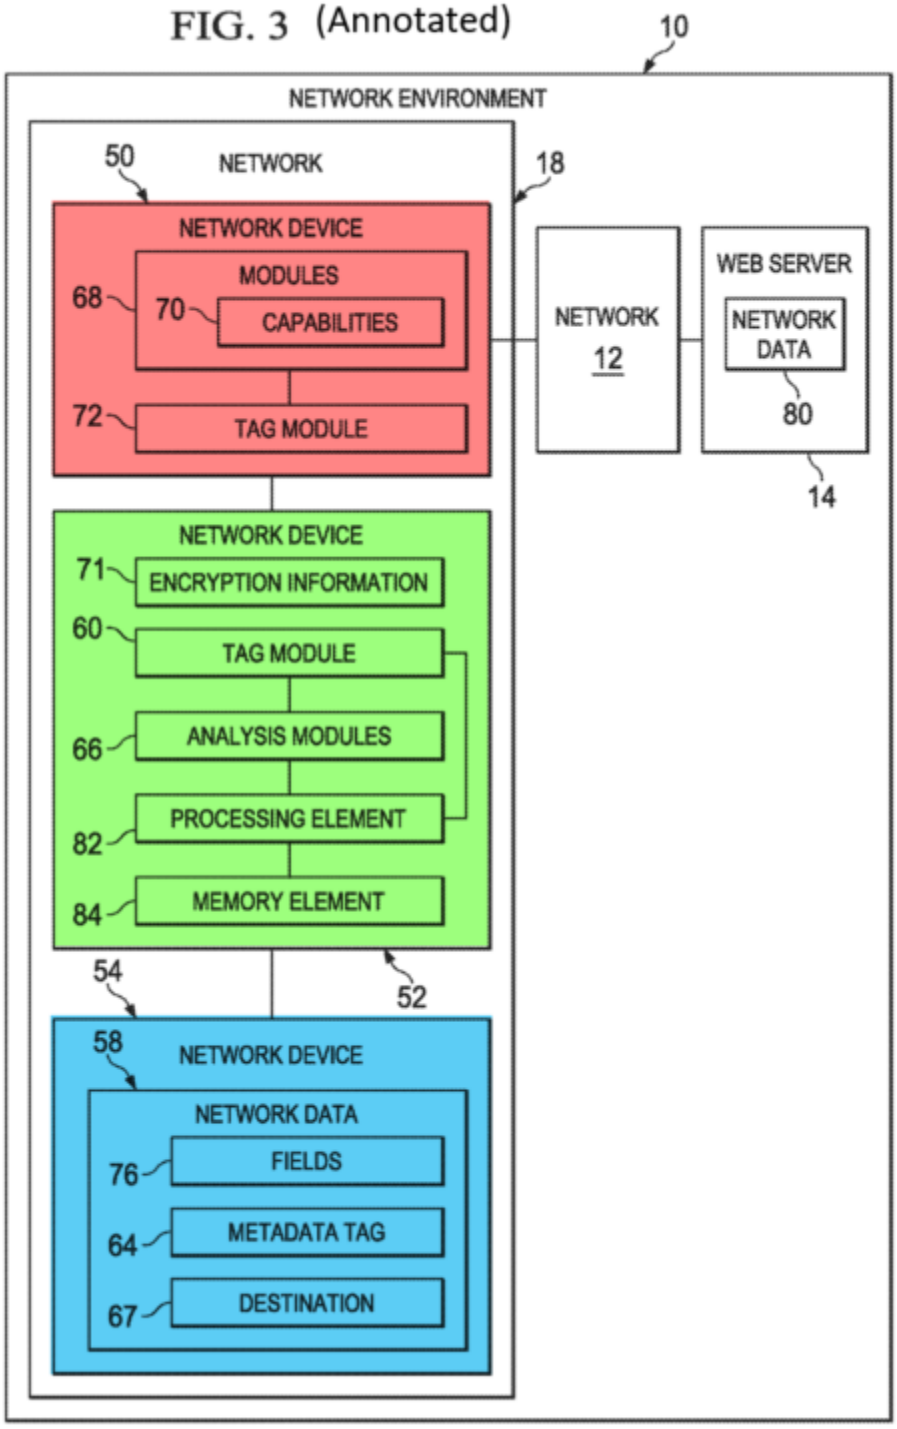


FIG. 2

In operation, the network data carried a metatag tag that included information regarding "processing performed by a previous network device through which network data [] ha[d] already traveled." (*Id.*, ¶ 27.) The next network device receiving the network data analyzed the metadata tag to determine what additional processing to perform "without having to repeat actions already performed by the previous network device." (*Id.*, ¶ 27; *see also, e.g.*, *id.*, ¶ 28.) Buruganahalli's Figure 3 (below) depicted an embodiment that followed this process.

Netskope Exhibit 1002

FIG. 3 (Annotated)



Network device (50), network device (52), and network device (54) were connected within network (18). (*Id.*, ¶ 32.) When network data (58) traveled from network device (54) to network device (52), it carried metadata tag (64) that included

"processing of network data 58 by network device 54." (*Id.*, ¶ 36.) For example, the information could include a virus scan result or an intrusion protection scan result. (*Id.*) If the network data included a metadata tag, analysis module 66 then performed a "number of actions based on…metadata tag 64." (*Id.*, ¶ 38.) Network device 52 could then "add or update…metadata tag 64" to indicate the processing it performed and send network data (58) to network device (50). (*Id.*, ¶ 40, 46.)

1. Motivation to Combine

115. A POSITA would have been motivated to combine Chandra and Buruganahalli with a reasonable expectation of success for several reasons.

116. First, Chandra and Buruganahalli both concerned coordinating network security actions to reduce redundant actions. For example, Chandra disclosed network security devices transmitting CE and receiving CE coordinating operations using labels. (Chandra, Abstract, ¶ 30.) Likewise, Buruganahalli disclosed network devices in a data path coordinating operations using a metadata tag carried by data packets. (Buruganahalli, Abstract, ¶¶ 27, 28.) Thus, a POSITA looking to improve Chandra would have looked to Buruganahalli's complementary disclosures.

117. Second, combining Chandra and Buruganahalli would have provided a method to reduce redundant security operations with improved flexibility, security, and convenience. Specifically, as discussed in detail below, a POSITA would be

Netskope Exhibit 1002

motivated to add Buruganahalli's features of updating and removing a metadata tag to Chandra.

2. Claim 6, 17, 23, and 28

118. These claims recite similar limitations and are, thus, analyzed together.

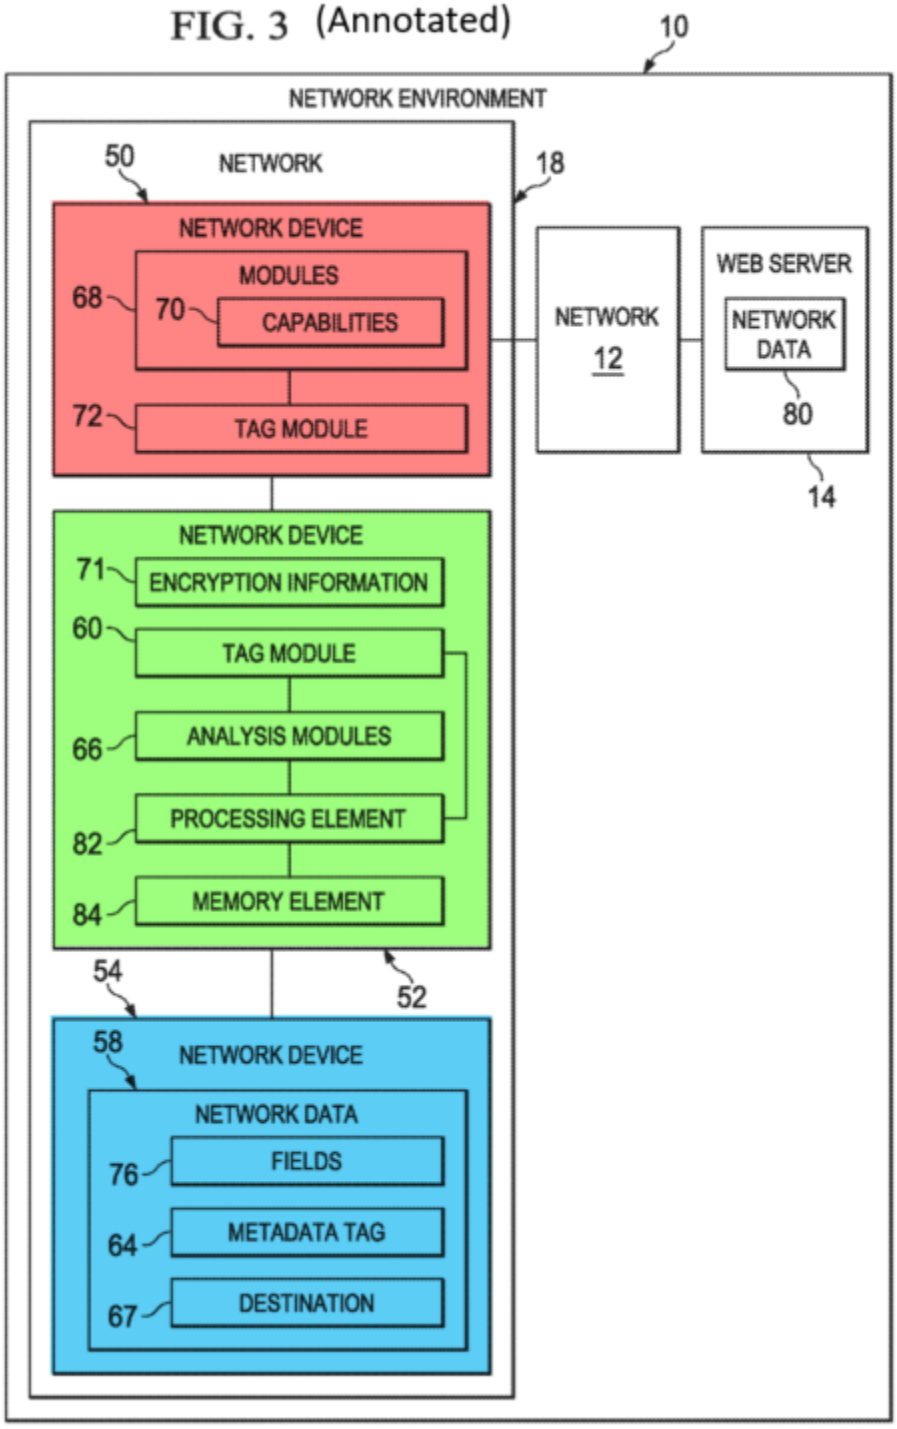
a. 6[a]/17[a]/23[a]/28[a]: "The method of claim [1/15/21/27], further comprising: updating, by the first network security appliance, the flag to indicate the network security operations executed by the first network security appliance;"

119. Buruganahalli disclosed that a network device updated a metadata tag (flag) to indicate the operations performed by the network device.

120. Buruganahalli's Figure 3 (below) depicted "[n]etwork device 50, network device 52, and network device 54 [] connected within network 18." (Buruganahalli, ¶ 32.)

Netskope Exhibit 1002

FIG. 3 (Annotated)



121. First, network device (52) received network data (58) from network device (54). Network data (58) included metadata tag (64). After receiving the data, analysis module (66) performed a "number of actions based on…metadata tag 64"

Netskope Exhibit 1002

on network data (58). (*Id.*, ¶ 38.) Afterwards, tag module (60) "***add[ed] or update[d]*** fields 76 to metadata tag 64 to form metadata tag 64" and "[f]ields 76 may separate the different types of analysis performed on network data 58." (*Id.*, ¶ 40.) A POSITA would have found it obvious that these disclosures in Buruganahalli disclosed updating the fields in the metadata tag to indicate the analysis performed on network data by the device that had received the data. Updating the metadata tag by subsequent network devices would create a cumulative record of the operations performed on the network data so far.

122. A POSITA would have been motivated to add Buruganahalli's feature of updating the metadata tag to Chandra. If each device added its own label indicating the actions it performed to the data packet, that might increase the time it took for the next device to read and process the actions performed by the prior devices, increasing latency and wasting resources. In addition, the space requirement for storing the labels would keep increasing along the data path. To prevent these issues, a POSITA would be guided by Buruganahalli's teaching regarding updating the metadata tag to use in Chandra, so that each CE along the data path updated the CE label before passing the data packet to the next CE instead of adding a new label.

Netskope Exhibit 1002

b. 6[b]/17[b]/23[b]/28[b]: "and transmitting, by the first network security network appliance, the network traffic with the updated flag to a next network appliance participating in the CSF."

123. Buruganahalli disclosed network device (52) transmitted the network traffic with the updated flag to network device (50), which was another device participating in the CSF.

124. Buruganahalli disclosed that "network device 52 [could] be configured to send network data 58 to network device 50," which obviously included metadata tag 64. (*Id.*, ¶ 46; *see also, e.g.*, *id.*, ¶¶ 44, 48.) Network device (50) participated in a CSF because it was coordinating actions with at least network device (52) to reduce redundancies. (*See, e.g.*, *id.*, ¶ 27.)

125. In the Chandra+Buruganahalli combination, a POSITA would incorporate Buruganahalli's disclosure of updating a tag into receiving CE updating a label and transmitting the updated label to a next network appliance participating in the CSF. There could be a variety of well-known reasons to implement Buruganahalli's disclosure. For example, Chandra disclosed that "a plurality of such CEs [could] be associated with a single PE." (Chandra, ¶ 28.) In this situation, a POSITA would find it obvious that in the multi-CE data path, a receiving CE would benefit from the updated label from the previous CE indicating what actions had already been performed. For another example, receiving CE (214) could transmit the data packet with updated label to receiving network device (212). This would

Netskope Exhibit 1002

inform receiving network device (212) of the actions performed, allowing receiving network device (212) to know the packet came from an appliance participating in the CSF and/or the operations those appliances performed. Receiving network device (212) participated in the CSF because it could understand the tag and/or was an intended recipient of the packet.

126. In yet another implementation, the receiving CE could transmit the data packet with updated label to a device in an extended VPN network. Specifically, Chandra disclosed a system with a transmitting end and a receiving end, shown below in both a general configuration (Figure 2) and a VPN configuration (Figure 8). (Chandra, ¶ 24, 58.)

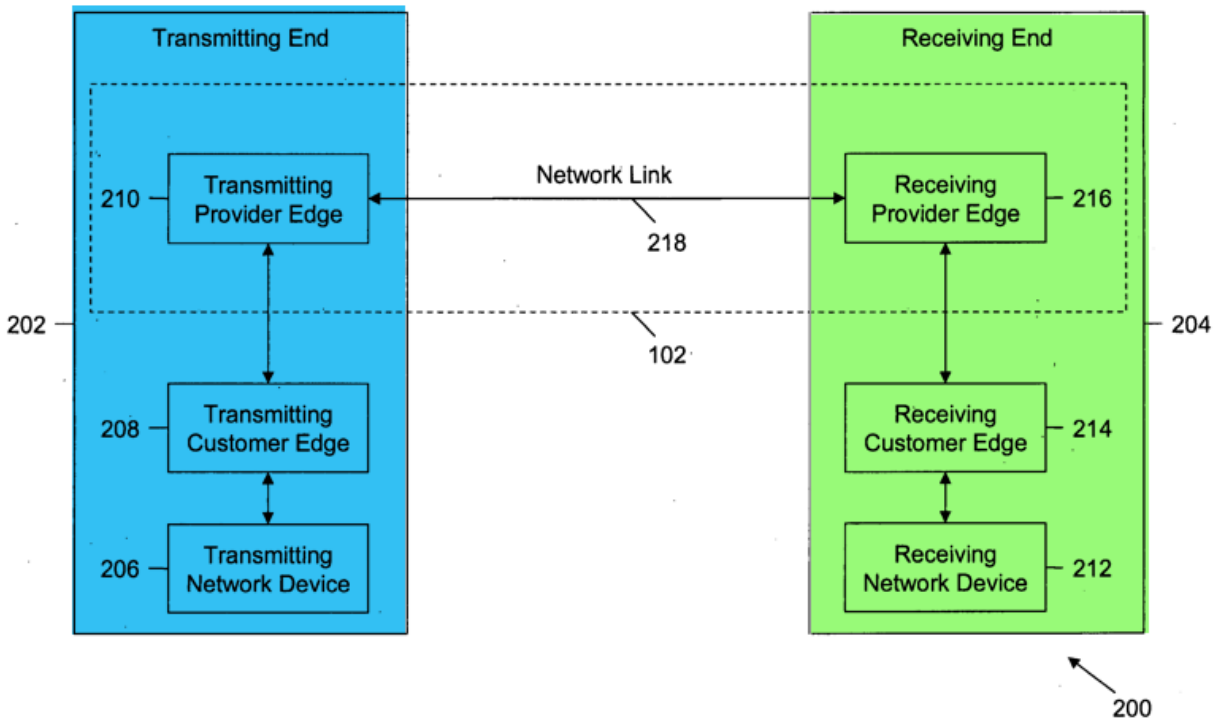


FIG. 2 (Annotated)

Netskope Exhibit 1002

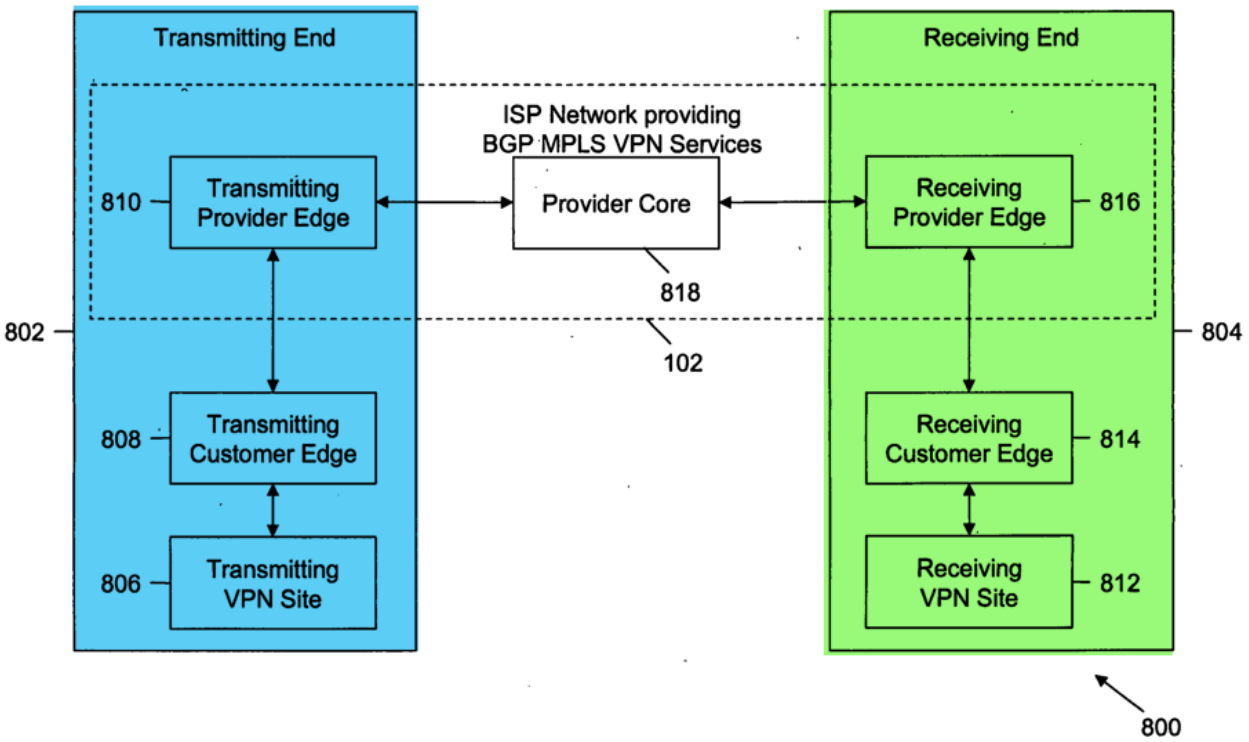


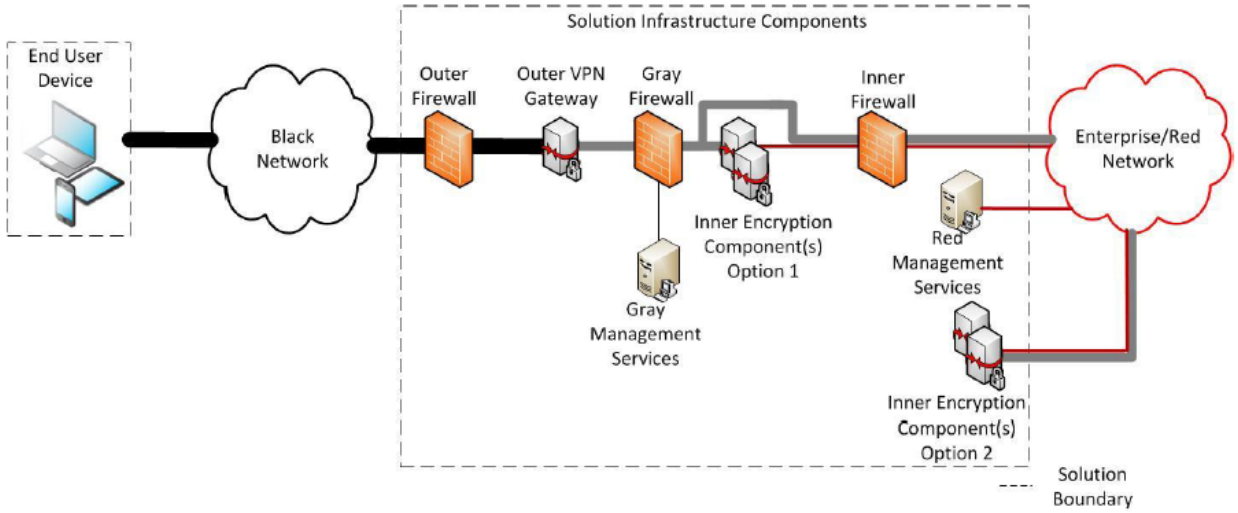
FIG. 8 (Annotated)

127. A POSITA interested in current developments in network security and computer security would have been aware of the National Security Agency's public (unclassified) outreach program to industry in which they discussed recommended reference architectures to achieve "commercial solutions for classified". These were published in a series of documents called "capability packages." (Appendix B, C.) In particular, the NSA's public outreach literature detailed how to protect classified data in transit across an untrusted network using multiple encrypted tunnels. This would have provided significant motivation for a POSITA to aggregate multiple embodiments disclosed in Chandra in sequence in the same system and combine this with Buruganahalli.

Netskope Exhibit 1002

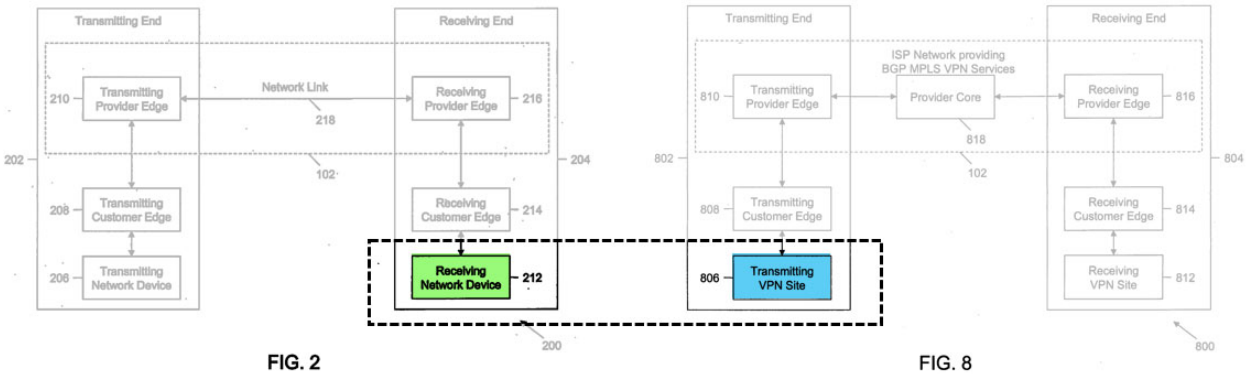
128. Long before the ’825 patent, it was well known how to connect one private network (e.g., a remote site of a company) with another private network (e.g., an enterprise network) using a VPN tunnel. This way, a receiving network device of the private network becomes a gateway to the VPN. In fact, when Chandra discloses that “The VPN is a network that uses a public network, like the Internet, to connect ***remote sites*** or users,” a POSITA would understand that a “remote site” might constitute an entire subnetwork such as a remote branch office of an organization that is connected to the main office’s private network using a VPN link.

129. A POSITA studying the NSA’s reference articles would find that the recommended structure for handling classified information consisted of separating protected data on a private “red network”. This red network can be connected to an outside “black network through a series of network appliances, including firewalls and VPN gateways. (Appendix B, 10, 21.)

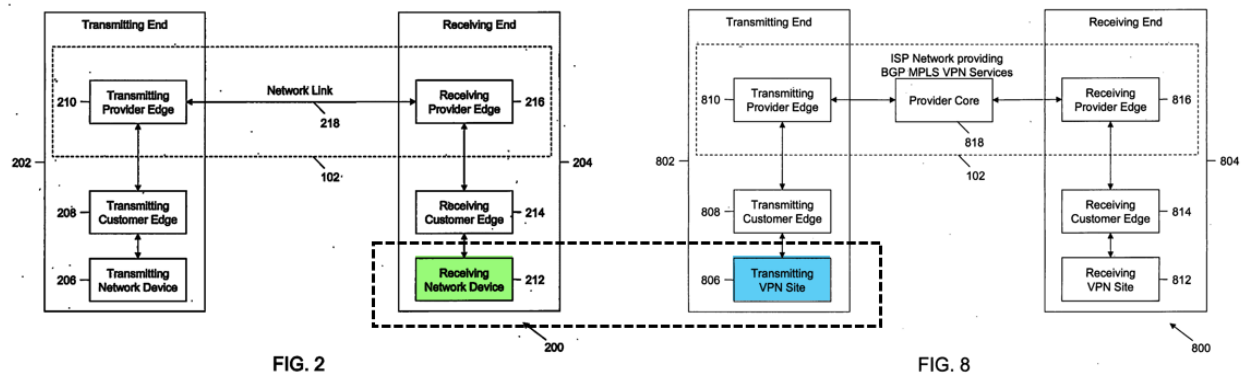


Netskope Exhibit 1002

130. Especially with this additional guidance and motivation to combine in mind, a POSITA would have found it obvious that the VPN in Chandra's Figure 8 could be attached as an extension to the private network in Chandra's Figure 2. Receiving network device (212) in Figure 2 could be coupled with transmitting VPN site (806) and become a gateway that connected the networks in Figures 2 and 8. (*See, e.g.*, Chandra, ¶ 56.) This extended private network is shown in the combined version of Chandra's Figures 2 and 8 (below).



(Redrawn and Annotated)



(Redrawn and Annotated)

3. Claims 7, 18, 24, and 29

131. These claims recite similar limitations and are, thus, analyzed together.

Netskope Exhibit 1002

a. "The method of claim [1/15/21/27], further comprising: removing, by the first network security appliance, the flag from the network traffic; and transmitting, by the first network security network appliance, the network traffic without the flag to a next network appliance."

132. Buruganahalli disclosed a network device removing the metadata tag carried by the network data before sending the data to the next device.

133. Specifically, Buruganahalli disclosed that "last network device to receive the network data prior to sending the data outside of the group of network devices ***may remove the metadata tag***. In other embodiments, any network device ***may remove the metadata tag***." (Buruganahalli, ¶ 29 (emphasis added); *see also, e.g.*, *id.*, ¶ 47.) In operation, Buruganahalli disclosed that the last network device in the group of devices may remove the metadata tag before transmitting the data to an "undesired network device." (*Id.*, ¶ 29.)

134. Moreover, a POSITA would have been motivated to incorporate into Chandra Buruganahalli's teaching to remove a tag from a tagged data packet before passing the data packet to another device for several reasons. First, removing the tag before passing the data packet would keep the network's labeling mechanism private, preventing malicious devices from intercepting the data packet and exploiting the labeling mechanism. Indeed, Buruganahalli recognized this benefit itself, remarking that the tag could be removed before the packet went to "undesired network devices." (*Id.*, ¶ 29; *see also, e.g.*, *id.*, ¶ 47.) Second, removing the tag

Netskope Exhibit 1002

would prevent compatibility issues in subsequence devices that might not be compatible with the tags, which would likely be proprietary. For example, receiving CE (814) could remove the CE label from the data packet before transmitting the untagged packet to receiving VPN site (812), when it determined that receiving VPN site (812) did not need to receive the CE label. (*See* Chandra, ¶ 68). This is supported by Buruganahalli, which disclosed removing the metadata tag when the next destination was an undesired network device. (*Id.*, ¶¶ 46-47.)

D. Ground 3: Chandra and Kumar render claims 10 and 11 obvious.

135. ***Kumar***: Kumar disclosed a method for a network node that received a network packet to authenticate the network node that sent the network packet based on a security tag containing the sending node's signature in the network packet. (Kumar, Abstract.)

136. In Figure 1A (below), Kumar disclosed a method in which sending node (20) and receiving node (60) established an authenticated session. (*Id.* ¶¶ 31, 32.) Then, receiving node (60) sent a unique client ID and session key to sending node (20). (*Id.* ¶ 33, 36.) Sending node (20) created a digital signature using the client ID and session key. (*Id.* ¶ 36.) Sending node (20) then used the digital signature to create a security tag and inserted the security tag in the outgoing packet. (*Id.* ¶¶ 45, 46.) The security tag may be inserted "in the IP header field." (*Id.* ¶ 38.) Sending node (20) then sent the packet with the security tag to receiving node (60).

Netskope Exhibit 1002

(*Id.* ¶ 49.) Receiving node (60) authenticated the security tag by recreating the digital signature using the same method as sending node (20) and then comparing the result against the digital signature contained in the security tag. (*Id.* ¶ 50.) An authentic security tag entailed that the packet was transmitted from an authenticated source (sending node (20)), so receiving node (60) accepted the packet and passed it on. (*See id.* ¶ 55.) Kumar explained that its system could broadly be implemented in the network or firewall layers and (*id.*, ¶ 64) and could use "any device on [a] network or sub-network that interact[ed] with the stream of packets being secured" (*id.*, ¶ 22).

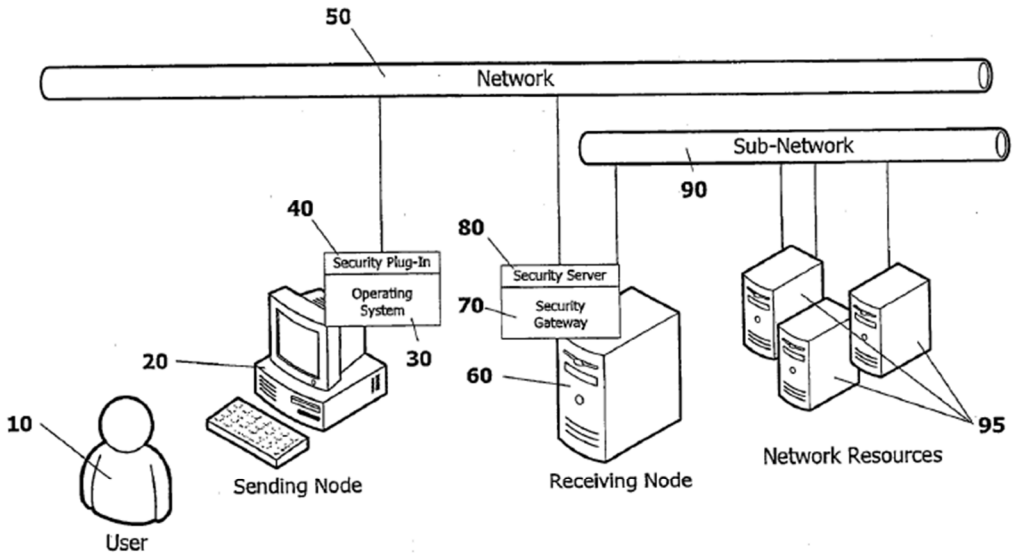


Fig. 1A

Netskope Exhibit 1002

1. **Motivation** to Combine

137. Kumar illustrated the combination of two well-known techniques: (1) the use of digital signatures to authenticate another party in a communication and (2) inserting additional information in a data stream that can be decoded by "downstream" devices. Kumar also discusses the traditional technique of inserting such additional information to be communicated "downstream" into packet headers, which has the advantage that no payloads need to be altered. A POSITA would have been motivated to incorporate these two common techniques into Chandra because, as discussed below, they were well-known techniques for securing data packets transmitted between network devices, and were often used predictably and reliably.

138. Additionally, combining Chandra and these techniques would have provided a method to coordinate security operations with improved convenience. The motivation for such improvements is described in detail below.

2. Claim 10

a. **"The method of claim 1, wherein the flag further comprises or is associated with a signature of the second network security appliance and wherein the method further comprises authenticating, by the first network security appliance, the second network security appliance based on the signature."**

139. Although Chandra did not expressly disclose a signature, a POSITA guided by Kumar would find it obvious that Chandra's method could include a CE label (flag) comprising or being associated with a signature of the transmitting CE

Netskope Exhibit 1002

(second network security appliance), and the receiving CE (first network security appliance) authenticating the transmitting CE (second network security appliance) based on the signature.

140. Kumar illustrated the common knowledge that a transmitting device could use information such as client ID, session key, public key, etc., to create a digital signature. (*See* Kumar, ¶ 36.) The transmitting device could then use the digital signature to create a security tag and insert the security tag in the data packet. (*See* Kumar, ¶¶ 45, 46.) When the receiving device received the data packet, it could authenticate the security tag using various methods. (*Id.*, ¶ 49.) It could recreate the digital signature using the same method as the transmitting device and then comparing the result against the digital signature contained in the security tag. (*See* Kumar, ¶ 50.)

141. A POSITA would have been motivated to incorporate the common technique of using a digital signature to authenticate a source into Chandra. Chandra disclosed a system of network security devices performing security actions on data packets. (Chandra, Abstract, ¶ 30.) In such a system, it was important for a network security device to authenticate the source of the data packet. This ensured that the received data packet was transmitted from a trusted security device and that the packet was not tampered with by unknown devices when it traveled between devices. A network security device would treat data packets transmitted from unknown

Netskope Exhibit 1002

devices differently from those transmitted from recognized and trusted security devices, as Chandra itself disclosed. (*See, e.g.*, Chandra, ¶¶ 47-48, 53.) Digital signature source authentication was the most popular, practical, and reliable authentication method. Thus, a POSITA looking to improve Chandra would have looked to incorporate digital signature authentication.

142. In practice, Chandra discloses that "before transmitting CE 208 and receiving CE 214 c[ould] exchange the CE labels, they ha[d] to authenticate each other." (Chandra, ¶ 32.) They could use "an authentication mechanism such as a Message-Digest Algorithm 5 (MD5) based authentication." (*Id.*, ¶ 33).

143. A POSITA would find it obvious to take this feature further. Rather than authenticating just at the beginning when the labels are exchanged, the authentication could be applied continuously on every packet exchange. In addition to the address of transmitting CE (208), a POSITA would find it obvious for the CE label to include more information on transmitting CE 208, including a crypto-based digital signature of transmitting CE (208). Transmitting CE (208) would tag the data packet with the CE label and transmit it to receiving CE (214), and receiving CE (214) would decrypt the signature in the CE label and authenticate transmitting CE (208), confirming that the data packet was transmitted from transmitting CE (208).

Netskope Exhibit 1002

3. Claim 11

a. "The method of claim 1, wherein the flag is carried in an Internet Protocol (IP) header field of the one or more packets."

144. Kumar disclosed the well-known idea that a security tag could be inserted "in the IP header field" of a data packet. (Kumar, ¶ 38)

145. A POSITA would have been motivated to incorporate this well-known technique of inserting the tag in the IP header field in Chandra's method. Around the priority date of the '825 patent, IP was the most standard and popular protocol for network communication. The majority of networks transmitted data packets among network devices using IP; given that Chandra's method could be implemented in a VPN environment, which operated over the public internet, a POSITA would have understood that Chandra's solution was at the very least able to interoperate with IP. An IP data packet was composed of payload and header. A POSITA would have found it obvious that if the IP packet was to carry a label, the header field of the IP packet would be the most natural place to record such a label; a POSITA would likely not want to modify the payload. Moreover, by allowing for the variable length "options" field in the header, the IP specification itself provides for a natural way to insert such information into the header without jeopardizing compatibility with devices that are unaware of the extension. Therefore, in

Netskope Exhibit 1002

Chandra's method, a POSITA would have found it obvious to record the CE label in the IP header of the data packets.

E. Ground 4: Chandra and Chandra-2 render claims 13, 14, 19, 20, 25, 26, 30, and 31 obvious.

146. The claims challenged under this ground merely require that the network security devices perform specific, well-known operations. As shown below, Chandra-2 proved that those operations were common knowledge.

147. A POSITA would have found it obvious that Chandra's security devices could perform these operations described in Chandra-2 because each operation was well-known and commonly implemented and performed in computer networks. Indeed, incorporating them into Chandra would have been implementing very common elements to yield predictable results. Thus, a POSITA looking to improve Chandra would have implemented and incorporated these common operations, giving Chandra's receiving and transmitting CEs a broader range of operations.

1. Claims 13, 19, 25, 30

148. These claims recite similar limitations and are, thus, analyzed together.

a. "The method of claim [1/15/21/27], wherein the one or more network security operations include network traffic logging."

149. Chandra-2 described that a "***network security device*** c[ould] be a device providing…***logging***." (Chandra-2, ¶ 23 (emphasis added).) A POSITA

Netskope Exhibit 1002

would have known that network traffic logging is a well-known and widely used network security operation that is especially important for forensic analysis after a vulnerability or breach has been discovered. A POSITA would have found it obvious that "logging" provided by a network security device might include "network traffic logging." A POSITA would also have found it obvious that the CE in Chandra, as a network security device, could perform network security logging.

2. Claims 14, 20, 26, 31

150. These claims recite similar limitations and are, thus, analyzed together.

a. "The method of claim [1/15/21/27], wherein the one or more network security operations include one or more of antivirus scanning, intrusion prevention (IPS), content filtering, data leak prevention, antispam, antispyware, reputation-based protections, event correlation, network access control, vulnerability management, load balancing, application control and traffic shaping."

151. Chandra-2 described that a "***network security device*** c[ould] be a device providing one or more of the following features: ***antivirus, intrusion prevention (IPS), content filtering, data leak prevention, antispam, antispyware…reputation-based protections, event correlation, network access control, vulnerability management, load balancing and traffic shaping***." (Chandra-2, ¶ 23 (emphasis added).) In addition, Chandra-2 described that "functionality that may be provided by a ***network appliance*** include…***content filtering***." (*Id.*, ¶ 24 (emphasis added).) Indeed, the '825 patent recognizes that

Netskope Exhibit 1002

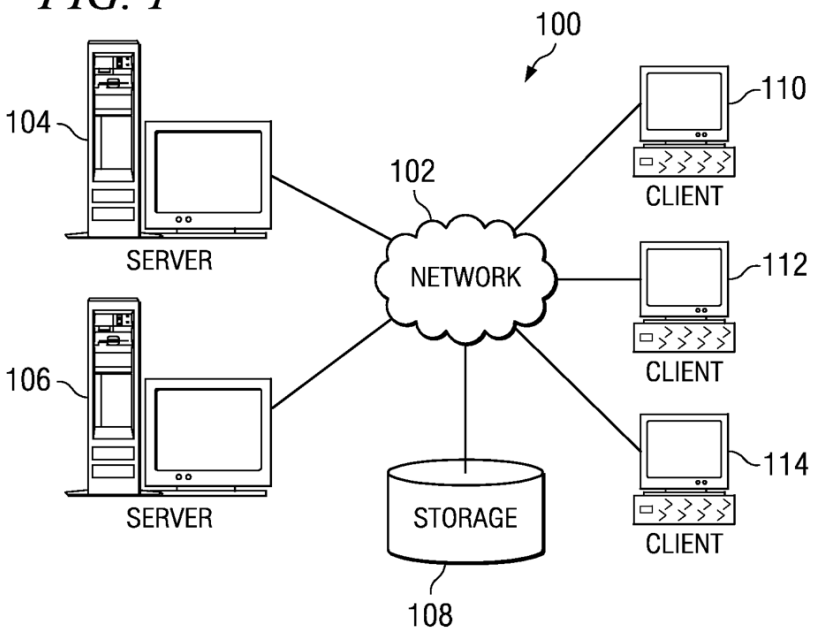
DLP and IPS were well-known. ('825 patent, 9:23-26 ("operations, such as antivirus, DLP, IPS…are *well known* to those skilled in the art and further description thereof will be omitted") (emphasis added).) A POSITA would have found it obvious that the CE in Chandra could perform one or all of these well-known operations.

F. Ground 5: Keohane renders claims 1-5, 10-12, 15, 16, 21, 22, and 27 obvious.

152. ***Keohane***: Similar to Chandra, Keohane was filed over ten years before the challenged patent and disclosed coordinating operations among security devices to avoid redundant security actions. (Keohane, Abstract, ¶¶ 9, 10.)

153. Keohane's Figure 1 (below) disclosed an embodiment of network data processing system (100) that included server (104) and server (106) connected to clients (110, 112, and 114) through network (102). (*Id.*, ¶ 22, 23.)

FIG. 1



The clients and servers included "software and/or hardware used to perform security

Netskope Exhibit 1002

actions." (*Id.*, ¶ 23.) Moreover, network (102) included devices to perform security actions, like "routers, gateways, and firewalls." (*Id.*)

154. Figure 3 (below) depicted an exemplary multi-level security system according to Keohane. The exemplary system included a data packet (306), firewall (300), mail gateway (302), client (304), and a security database (308). (*Id.*, Fig. 3.)

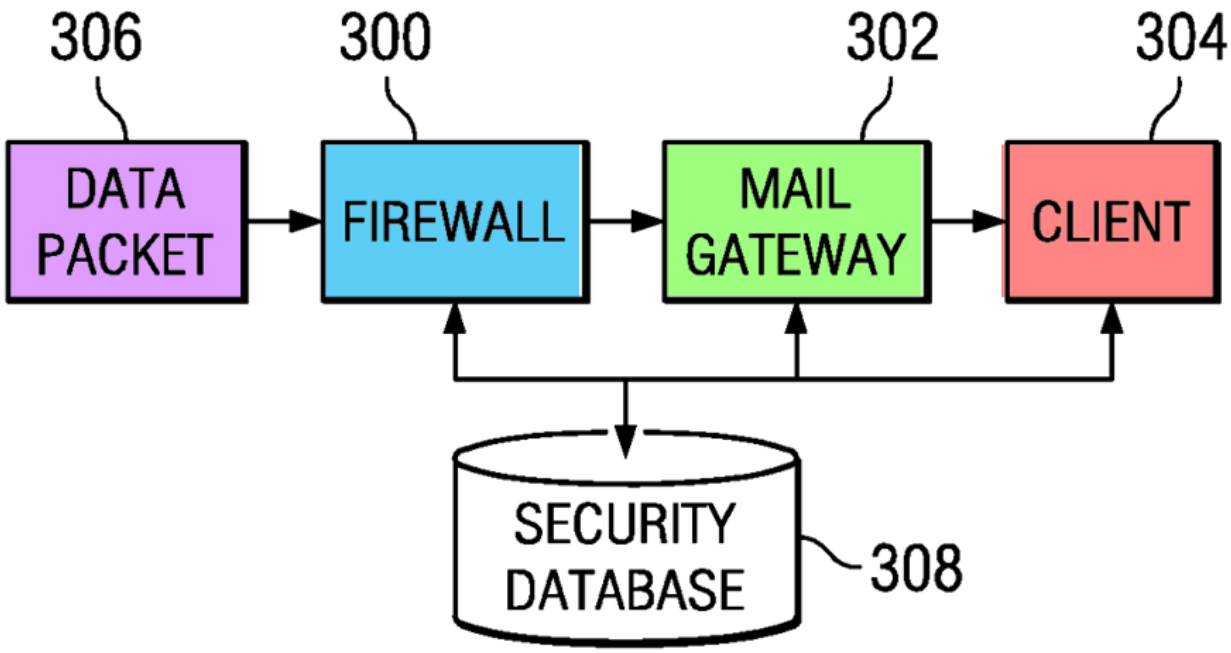


FIG. 3
(Annotated)

The members of the system cooperated and coordinated with each other to efficiently perform security actions on data packets. In Figure 3's example, "firewall 300, mail gateway 302, and client 304 form different levels of network security in a network." (*Id.*, ¶ 33.) To start, firewall (300) received data packet (306), then performed one or more security actions (e.g., determining if packet was from a trusted source) on

Netskope Exhibit 1002

it. (*Id.*) After firewall (300) performed the security actions, it added a digital signature to data packet (306). (*Id.*) The signature was added to the data packet's header and "indicate[d] that a security action ha[d] been performed on data packet 306." (*Id.*, ¶ 34.) In this example, the signature indicated that firewall (300) determined that data packet (306) was from a trusted source. (*Id.*) Firewall (300) then passed data packet (306) to mail gateway (302). (*Id.*, ¶ 35.)

155. Once mail gateway (302) received the packet, it examined the data packet's header to "determine whether a signature from a trusted security element [was] present in data packet 306." (*Id.*, ¶ 35.) In this example, that signature was present, so mail gateway (302) knew that firewall (300) had performed that security operation, but firewall (300) had not scanned for malware. (*See id.*) Consequently, "mail gateway 302 perform[ed] another security action to scan for viruses and, thereafter, "add[ed] a signature to data packet 306." (*Id.*, ¶ 36.) Mail gateway (302) passed the packet to client (304), where the process of determining a security operation from a signature, performing that operation, and adding a signature to the data packet, was repeated. (*Id.*, ¶ 37.)

156. Keohane illustrated its method of performing security actions on data packets in Figure 7's flowchart (below). (*Id.*, ¶ 52.)

Netskope Exhibit 1002

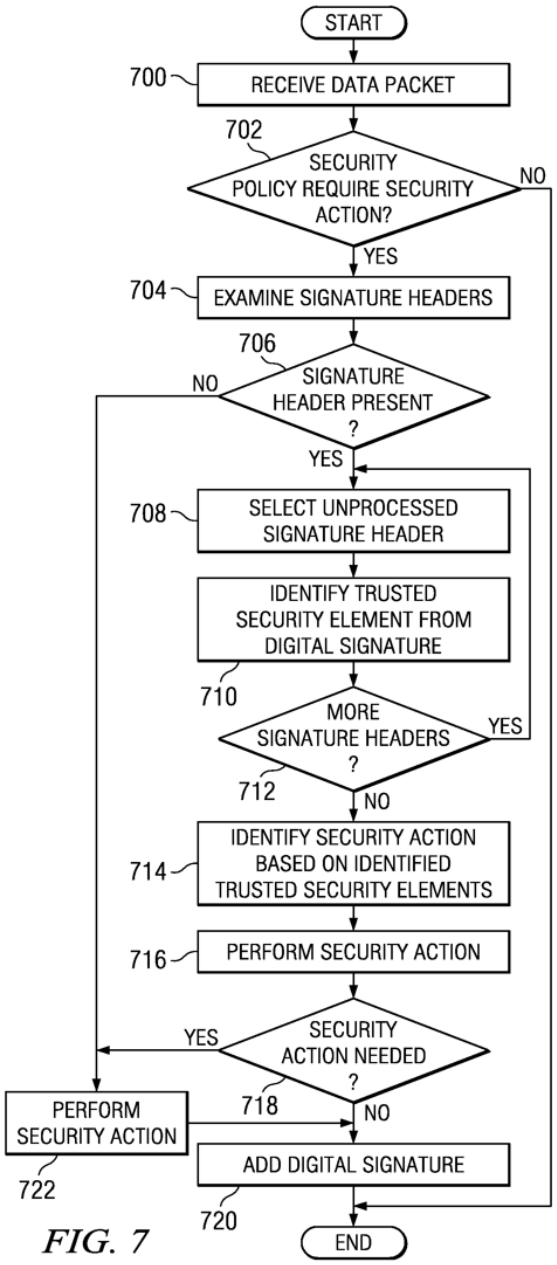


FIG. 7

157. Keohane further disclosed that the security action performed on the data packet could be indicated by inserting a code that corresponded to the action into the signature. For example, a code of "01" or "policy 1" could indicate that a trusted source determination had been performed, while a code of "02" or "policy 2" could indicate that a virus scan had been performed. (*Id.*, ¶ 38.)

Netskope Exhibit 1002

1. Claim 1

a. 1[pre]: "A method comprising:"

158. Keohane disclosed a method for screening data packets. (*Id.*, Abstract ("A computer implemented method…for screening data packets.").)

b. 1[a]: "receiving, by a first network security appliance of a private network that is participating in a cooperative security fabric (CSF), network traffic that is to be inspected by the first network security appliance;"

159. Keohane disclosed a firewall (first network security appliance) of a private network (e.g., intranet, LAN, VPN) coordinating security actions with other devices and receiving and inspecting data packets (network traffic).

160. ***"first network security appliance of a private network"***: Keohane disclosed a first network security appliance in, for example, the form of mail gateway (302) (shown below).

Netskope Exhibit 1002

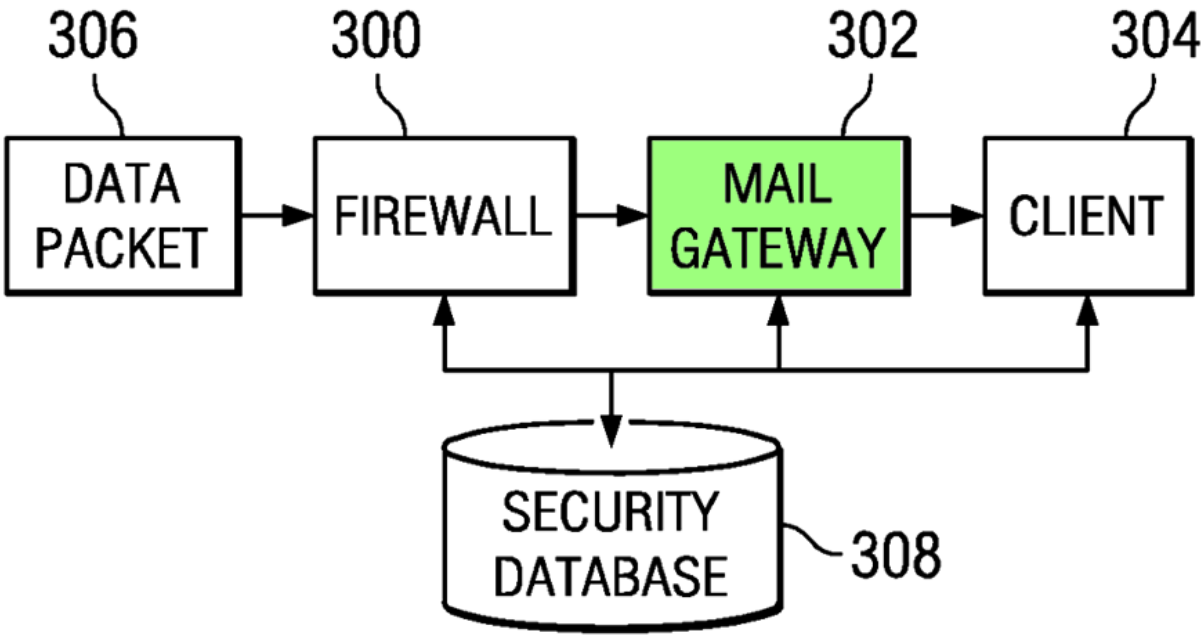
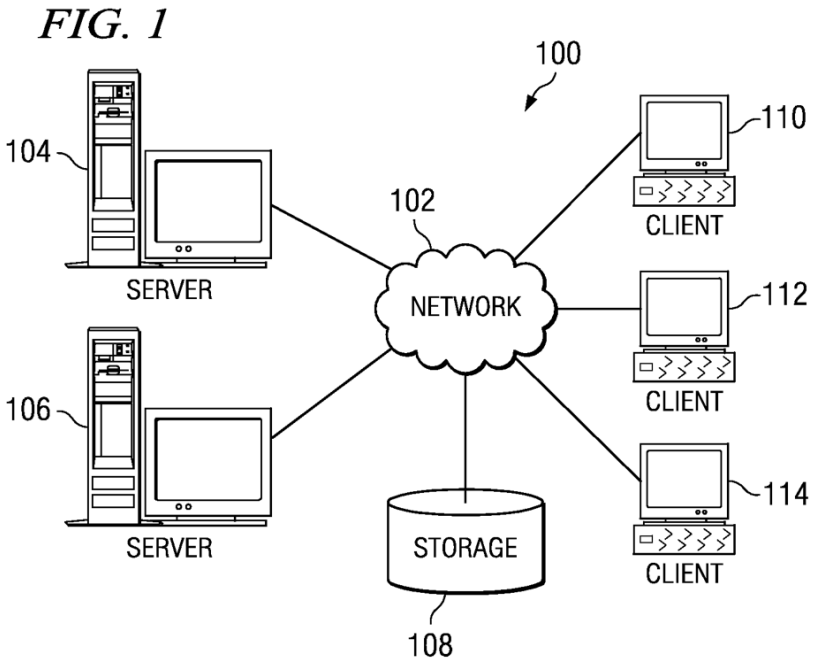


FIG. 3
(Annotated)

(*Id.*, ¶ 23, FIG. 3.) While Keohane talks of "security devices" and doesn't use the term "appliance," a POSITA would have understood that several of the devices described as "security element" by Keohane would also fall under the common nomenclature of a "security appliance." Among these are the firewall and the mail gateway. (*Id.*, ¶¶ 23 ("network 102…includes devices such as…***gateways***…that may…***perform security actions with respect to malware***.") (emphasis added), 36.)

161. Additionally, Keohane disclosed that mail gateway (302) could be on a private network, e.g., an intranet, LAN, or VPN. (*Id.*, ¶¶ 24, 32.) Specifically, Keohane disclosed that network (102) (below) provided communications links between various devices in the data processing system and "include[d] devices such

as routers, ***gateways***, and firewalls that may…perform security actions with respect to malware." (*Id.*,¶ 22-23 (emphasis added).) This disclosure clearly included mail gateway (302). Keohane also disclosed that network (102) could be an intranet, LAN, or VPN. A POSITA would understand that an intranet was a private network. (*See, e.g.*, Ex. 1009 ("intranet - A private network based on Internet protocols such as TCP/IP but designed for information management within a company or organization.").) Additionally, as discussed above for Chandra (*see supra,* §§ A1b), a POSITA would have found it obvious that a LAN and VPN were private networks.



162. In addition to mail gateway (302) being a first network security appliance, the first network security appliance could also be any network appliance in the appliance chain (e.g., client (304)) that received network traffic because Figure 3 was not limiting, and a POSITA would understand that the order and number of

Netskope Exhibit 1002

appliances in Keohane's system could be adjusted according to implementation. But for the sake of simplicity, Petitioner focuses its mapping of the first network security appliance on mail gateway (302).

163. ***"participating in a cooperative security fabric (CSF)"***: Keohane disclosed that mail gateway (302) coordinated security actions with other appliances. To coordinate actions, each participating appliance inserted a signature into a packet to indicate the security operations performed by it. (*Id.*, ¶ 11.) Downstream devices determined whether the signature originated from a trusted security element. (*Id.*) If so, the receiving appliance determined the security operations performed by the prior devices and, based on those operations, determined which operations it would perform. (*Id.*) While not expressly disclosed in Keohane, a POSITA would have understood that, prior to operation, the appliances participating in the system would have been configured so that each knew which code corresponded to which operation. (*See, e.g.*, ¶¶ 37 ("These ***signatures indicate[d] that the data packet [was] from a trusted source and that no viruses were found*** in data packet 306."), 38 ("[D]etermination of what security actions have been taken may be ***identified by a code in the signature***.") (emphasis added).) Coordinating security actions allowed the system to avoid redundant actions. (Keohane, ¶¶ 11, 37 (disclosing that "the amount of time and resources used in scanning data packet 306 [was] reduced" and "[u]nneeded redundancy [did] not occur, saving time in processing data.").)

Netskope Exhibit 1002

164. Members of Keohane's security system performed the same function as the '825 patent's CSF, which "coordinates operations among members and avoids redundant operations." ('825 patent, 5:45-47.) Therefore, a POSITA would have found it obvious that mail gateway (302) participated in a CSF.

165. ***"receiving…network traffic that is to be inspected by the first network security appliance"***: Keohane disclosed that mail gateway (302) received data packet (306) from firewall (300) to inspect it. (*Id.*, ¶ 35 ("When data packet 306 [was] received by mail gateway 302, mail gateway 302 examine[d] data packet 306.").)

166. Figure 7 (below), for example, depicted a process flow for performing a security action in Keohane's security elements (e.g., mail gateway). (Keohane, ¶¶ 35 ("[A] trusted security element [was] a network element, such as a firewall, mail gateway, router, or some other security intermediary that perform[ed] a security action on a data packet."), 52.) At step 700, the security element (e.g., mail gateway) "receiv[ed] a data packet." (*Id.*, ¶ 53.)

Netskope Exhibit 1002

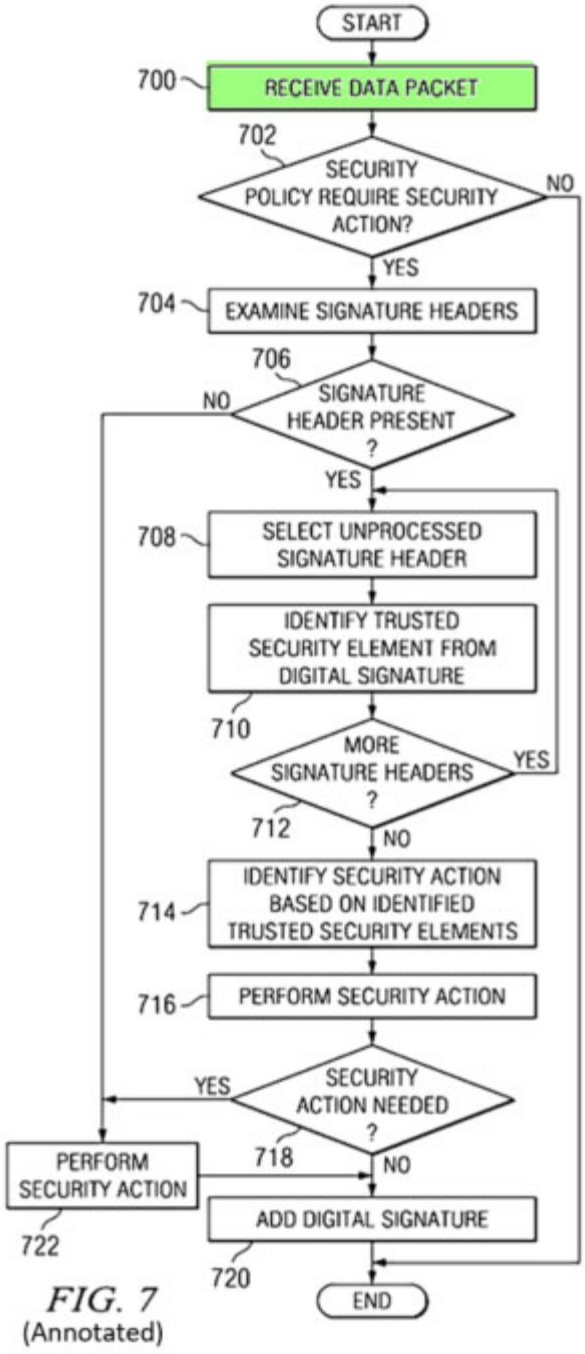


FIG. 7
(Annotated)

167. Additionally, client (304) received data packet (306) from mail gateway (302) to inspect it. (*Id.*, ¶ 37 ("Then, data packet 306 is passed to client 304. Client 304 examines data packet 306.").)

c. 1[b]: "determining, by the first network security appliance, whether the network traffic has been transmitted from a second network security appliance that is participating in the CSF based on a flag carried by one or more packets of the network traffic;"

168. Keohane disclosed that mail gateway (302) (first network security appliance) determined whether received packet (network traffic) had been transmitted from another network security element (e.g., firewall (300)) (second network security appliance) based on the packet's signature (flag).

169. Keohane disclosed that, after mail gateway (302) received data packet (306) traffic, it "examine[d] data packet 306 to ***determine whether a signature from a trusted security element [was] present***…." (Keohane, ¶ 35 (emphasis added); *see also, e.g.*, ¶ 32 ("When a data packet [was] received, a determination [was] made as to whether a signature of a trusted security element [was] present in the header of the data packet.").) In Step 704 of Figure 7 (below), for example, Keohane disclosed that a device analyzed the signature in a packet's header to determine if the packet originated from a trusted security element (e.g., firewall (300)). (*Id.*, ¶ 54.)

Netskope Exhibit 1002

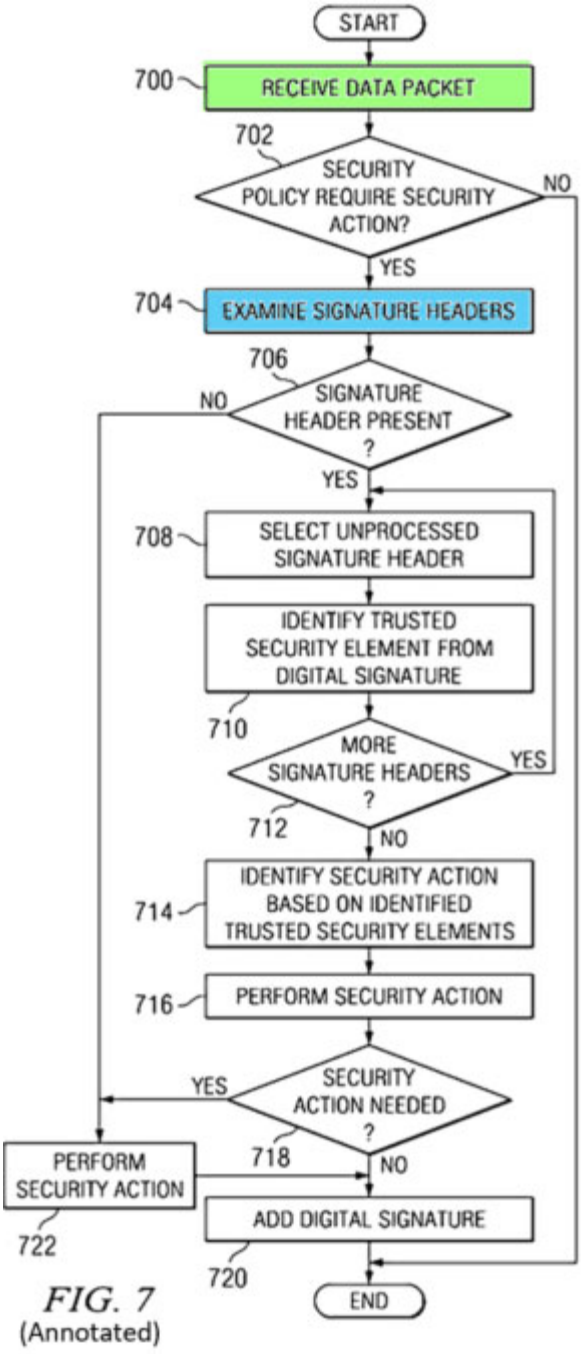


FIG. 7
(Annotated)

170. Keohane disclosed that the signature in a packet's header was used to determine if the packet had been transmitted from another device in the system. Keohane explained that "the signature [was] a digital signature added to the header" of a data packet. (*Id.*, ¶ 34.) The signature could contain a code (e.g., "01" or "02")

Netskope Exhibit 1002

to indicate what actions had been taken on a packet. For example, a "code of '02' or 'policy 2' may indicate that a virus scan has been performed by the trusted element." (*Id.*, ¶ 38.) A POSITA would have found it obvious that the digital signature was a flag as described by the '825 patent because the '825 patent broadly described a flag. ('825 patent, 6:28-34 ("flag (e.g., a single bit flag or a collection of bit fields.").)

171. Further, a POSITA would have found it obvious that the presence of the signature indicated that the data packet was transmitted from a trusted security element participating in the CSF. This is consistent with the '825 patent's disclosure that a flag's presence in network traffic "indicat[ed] that the network packet [was] transmitted from a member node." ('825 patent, 6:54-59, 6:65-67.) In addition, a POSITA would have found it obvious that the signature disclosed in Keohane was a well-known authentication method that would be used by the receiving security element to determine if the packet had been transmitted from another trusted security element. Put simply, if the packet header had a signature that the receiving device could recognize as being from a trusted security element, then the traffic came from another device (second network security appliance) in the CSF.

d. 1[c]: "when the network traffic is determined to have been transmitted from the second network security appliance: determining, by the first network security appliance, one or more network security operations that are executed by the second network security appliance;"

Netskope Exhibit 1002

172. If data packet had been transmitted from firewall (300), Keohane disclosed that mail gateway (302) analyzed the signature carried by the data packet to determine the actions performed on the data packet by firewall (300).

173. Keohane disclosed that mail gateway (302) checked the signature carried by data packet (306) from firewall (300) and determined that firewall (300) had performed the security action of determining that data packet (306) was from a trusted source. (*See* Keohane, ¶¶ 35-36.) Similarly, client 304 checked the signature carried by data packet 306 from mail gateway 302 and determined that mail gateway 302 had performed the security action of virus scan. (*See id.*, ¶ 37 ("a virus scan…has already been performed by mail gateway 302").)

174. As discussed in element 1[b], Keohane disclosed that the codes within the signature could indicate the actions performed by a prior device. (Keohane, ¶ 38.) Accordingly, Keohane disclosed that "recipients of data packet 306 may identify what security actions ha[d] been performed on data packet 306 to determine what additional security actions [were] to be performed." (*Id.*, ¶ 40.)

e. 1[d][i]: "determining, by the first network security appliance, one or more local network security operations that are to be executed by the first network security appliance based on the one or more network security operations that are executed by the second network security appliance,"

175. Keohane disclosed that mail gateway (302) determined the actions it would perform on a data packet based on the actions performed by firewall (300).

Netskope Exhibit 1002

176. Keohane disclosed that a security element (e.g., mail gateway) performed a security action "selected based on the identified previous security action performed on the data packet [by the previous security element]". (*Id.*, ¶ 12; *see also, e.g.*, ¶ 55 ("[A] security action [was] identified based on the identified security elements from the digital signatures (step 714)."), ¶ 58 ("In response to the signature of the trusted security element being present, a security action [was] then performed in the data packet based on the previous security action."), ¶ 32.)

177. For example, in connection with Figure 3, Keohane disclosed that mail gateway (302) performed a virus scan on data packet (306) because the identified security action performed by firewall (300) did not include a virus scan. (*See id.*, ¶ 36 ("Although, firewall 300 allowed data packet 306 to pass, this security element did not perform any scanning for malware. ***As a result, mail gateway 302 perform[d] another security action to scan for viruses.***") (emphasis added).) Similarly, client (204) performed a spyware scan on data packet (306) because the identified security action performed by mail gateway (302) did not include a spyware scan. (*See id.*, ¶ 37 ("Client 304 scan[ned] data packet 306 for spyware. This security action [was] performed because the other scans did not include scanning data packet 306 for this type of malware. Client 304, however, d[id] not need to perform a virus scan in this example because this type of security action ha[d] already been performed by mail gateway 302.").) Additionally, Keohane disclosed that "recipients of data packet

Netskope Exhibit 1002

306 may identify what security actions have been performed on data packet 306 to determine what additional security actions are to be performed." (*Id.*, ¶ 40.)

f. 1[d][ii]: "including skipping, by the first network security appliance, the one or more network security operations that are executed by the second network security appliance;"

178. Keohane disclosed a security element (e.g., mail gateway) skipping one or more security actions that had already been performed by a previous security element.

179. For example, in connection with Figure 3, Keohane disclosed that "[c]lient 304…does not need to perform a virus scan…because this type of security action has already been performed by mail gateway 302." (*Id.*, ¶ 37.) Keohane also explained that, if mail gateway (302) had also performed a virus and anti-spyware scan, then "client 304 would taken [sic] no action…." (*Id*; *see also, e.g.*, *id.*, ¶ 12 ("The security action may be an absence of an additional action to check the data packet.").) Finally, Keohane disclosed that a security element could simply skip all actions depending on the prior actions performed on the data packet. (Keohane, ¶ 33 ("The security action also may ***include taking no action depending on the prior actions taken on a particular data packet***.") (emphasis added).)

g. 1[e]: "and executing, by the first network security appliance, the one or more local operations on the network traffic."

Netskope Exhibit 1002

180. Keohane disclosed that the security element executed the one or more local operations on the data packet (network traffic). (*Id.*, ¶ 11 ("A ***security action [was] performed on the data packet*** based on the previous security action performed on the data packet in response to the signature of the trusted security element being present."), ¶ 36 ("As a result, mail gateway 302 ***perform[ed] another security action to scan for viruses***.") (emphasis added), ¶ 37 ("Client 304 then ***perform[ed] an additional security action*** based on the security actions already performed.").)

2. Claim 2

a. 2[a]: "The method of claim 1, wherein said determining, by the first network security appliance, one or more network security operations of the first network security appliance to be executed by the first network security appliance based on the one or more network security operations that are executed by the second network security appliance"

181. *See* element 1[d][i].

b. 2[b]: "further comprises determining, by the first network security appliance, one or more network security operations that are to be executed by the first network security appliance except the one or more network security operations that are executed by the second network security appliance based on local network polices."

182. Keohane disclosed that mail gateway (302) determined to perform a security operation that had not been executed by firewall (300) based on local policies.

Netskope Exhibit 1002

183. Keohane disclosed that mail gateway (302) determined to perform a virus scan on a data packet but did not determine to perform a trusted source action because firewall (300) had already performed that action. (Keohane, ¶¶ 34-36.) Specifically, mail gateway determined to perform a virus scan because, "[a]lthough firewall [] allowed data packet [] to pass, [it] did not perform any scanning for malware." (*Id.*, ¶ 36.) Consequently, mail gateway "perform[ed] another security action to scan for viruses." (*Id.*) Thus, mail gateway skipped performing a trusted source action and, instead, performed a malware scan that had not been performed by the firewall.

184. As another example, Keohane disclosed that client (304) determined the security actions performed by firewall (trusted source analysis) and mail gateway (virus scan). (*Id.*, ¶ 37 ("These signatures indicate that the data packet [was] from a trusted source and that no viruses were found in data packet 306.").) Keohane then explained that "[c]lient 304 [] ***perform[ed] an additional security action based on the security actions already performed***." (*Id.* (emphasis added).) That action could be a spyware scan because other devices had not performed that action. (*Id.*; *see also, e.g.*, *id.*, ¶ 40 ("In this manner, other recipients of data packet 306 may identify what security actions ha[d] been performed on data packet 306 to determine what additional security actions [were] to be performed.").)

Netskope Exhibit 1002

185. ***"local network polices"***: It is unclear from the claim language whether the "based on local network policies" limitation applies to the first network security appliance's operations, or the second network security appliance's operations, or both. Keohane disclosed this limitation according to any interpretation. A POSITA would find it obvious that in Keohane, there were two sets of local policies: one set local to the receiving security element and a second set local to the transmitting security element. Each set of policies were stored at and executed at and by the respective security element.

186. Keohane disclosed that each element in the system determined the local action(s) it would perform using local network policies. For example, a code of "policy 1" indicated that a determination of whether the source of the data packet was a trusted source had been performed and a code of "policy 2" indicated that a virus scan had been performed. (*Id.*, ¶ 38.) The codes could be stored locally on each device. (*Id.*)

187. Similarly, Keohane disclosed that, once a device received traffic, "[a] determination [was] made as to ***whether a security policy requires a security action***." (*Id.*, ¶ 53 (emphasis added).) The security policy was "a set of rules used to determine what security action to take." (*Id.*) Keohane explained that "[t]he ***security policy [could] be local*** to the component executing the process." (*Id.* (emphasis added).) A POSITA would find it obvious that these locally stored codes and

Netskope Exhibit 1002

security policies were local network policies instructing the member elements on what security actions to perform or not perform.

3. Claim 3

a. 3[a]:

188. *See* element 2[a].

b. 3[b]: "further comprises determining, by the first network security appliance, that the one or more network security operations that are executed by the second network security appliance are to be executed by the first network security appliance."

189. Keohane disclosed—or at least suggested—that mail gateway (302) determined to perform at least one security operation that had been executed by firewall (300) and skipping the remaining operations.

190. Keohane disclosed that the security operation performed by the downstream security appliance (e.g., mail gateway (302)) "include[d] ***performing another scan on the data packet,*** a different type of scan, or no action at all." (Keohane, ¶ 32 (emphasis added).) Thus, Keohane at least suggested that the network security element could simply perform another scan on the packet rather than a different type of scan or no action at all.

191. To the extent Keohane did not expressly disclose this limitation, a POSITA would have found it obvious. While Keohane disclosed avoiding redundant actions, in some circumstances, running another scan of the same type might be appropriate or simply more convenient. For example, a downstream device

Netskope Exhibit 1002

might perform certain actions that are so inexpensive to execute that for these actions, it may not be worth performing the redundancy analysis at all as it may be easier to just always execute this cheap action regardless. As another example, a broad scan performed by an upstream security appliance might contain elements that have a partial overlap with a more focused scan performed downstream. Being able to simply repeat any redundant action without having to carefully design a downstream scan that explicitly omits the partial work already performed upstream makes the overall configuration much simpler.

192. On the other hand, Keohane disclosed that the security operation performed by the security element "include[d] ***performing another scan*** on the data packet [or] ***a different type of scan***." (Keohane, ¶ 32.) This suggested that (1) one scan of the same type as a scan that had already been performed on the data packet by the previous security element could be performed, while (2) other scans that had already been performed were skipped and a different type of scan was performed instead. In general, a POSITA would have found it entirely obvious that at any point one could repeat redundant actions. This is just like the routing slip example from 1943 cited above in which any recipient of a package with a routing slip can elect to perform any action "locally", including repeating any action that has been flagged as having already been performed by a previous recipient along the route.

Netskope Exhibit 1002

4. Claim 4

a. "The method of claim 1, further comprising when the network traffic is not transmitted from the second network security appliance, executing, by the first network security appliance, one or more local network security operations specified by one or more local policies."

193. Keohane disclosed—or at least suggested—that if a security element received a data packet that lacked a signature header, certain security actions were performed on the packet.

194. In Step 706 (below), Keohane disclosed that a determination was made as to whether a signature header was present in a packet. (Keohane, ¶ 56.) If not, the process proceeded to Step 722, which was to perform a security action.

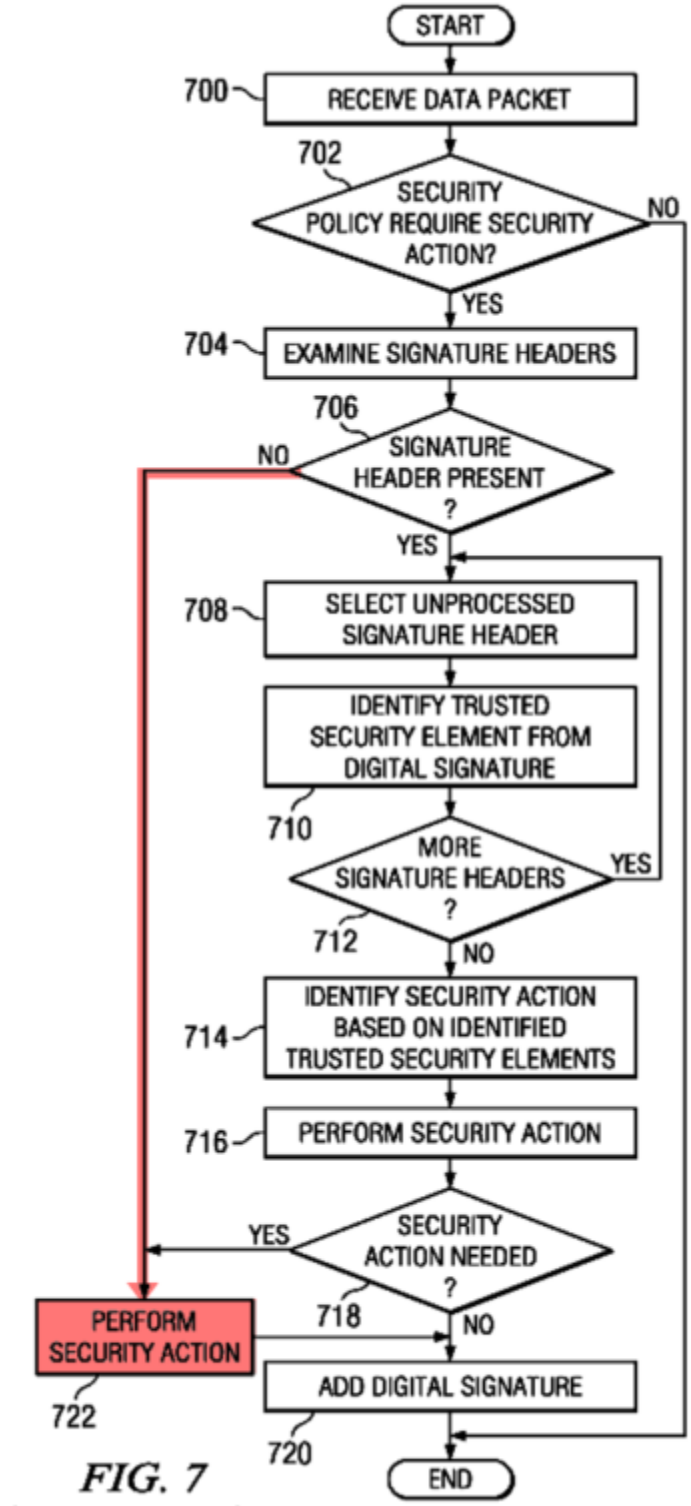


FIG. 7
(Annotated)

Netskope Exhibit 1002

195. A POSITA would have found it obvious that if the packet did not have a signature, it had not been transmitted by an upstream device (second network security appliance) participating in the cooperative security network because Keohane disclosed that each device added a signature to the packet after processing it. (*See, e.g.*, *id.*, ¶ 32.) Thus, a packet without a signature could not have come from an upstream device participating in Keohane's mechanism.

196. Additionally, the policies specifying the security action would be local policies because Keohane disclosed that a "security policy may be local to the component executing the process." (*Id.*, ¶ 53.) Moreover, the actions performed by the security element would be a local operation because it was being performed by the local device that received the traffic and would be executing the actions specified in the local policy. While Keohane did not expressly identify what the local actions could be in this scenario, a POSITA would have found it obvious that the actions could include dropping the packet or performing one or more security actions on it.

5. Claim 5

a. "The method of claim 1, wherein the flag indicates the network security operations executed by the second network security appliance."

197. Keohane disclosed that the "signature indicate[d] that a previous security action ha[d] been performed [by a previous security element] on the data packet." (*Id.*, ¶ 32.) For example, Keohane disclosed that firewall (300) added a

Netskope Exhibit 1002

signature to data packet 306 that "indicate[d] that a security action ha[d] been performed on data packet 306 [by firewall 300]." (*Id.*, ¶ 34; *see also, e.g.*, ¶¶ 36-37 (explaining that mail gateway (302) added a signature to the data packet indicating that a security action had been performed on the data packet).) Keohane explained that the security actions could be identified by a code in the signature. For example, a code of "01" could indicate that source of the data packet was a trusted source and a code of "02" could indicate that a virus scan had been performed. (*Id.*, ¶ 38.)

6. Claim 10

a. "The method of claim 1, wherein the flag further comprises or is associated with a signature of the second network security appliance and wherein the method further comprises authenticating, by the first network security appliance, the second network security appliance based on the signature."

198. Keohane disclosed that a flag in a packet included a signature of a security element (e.g., firewall) and that later security elements that received the packet used the signature to know if the packet came from a trusted source.

199. Keohane disclosed that, after a security element performed an action, the security element added a signature to the data packet. (Keohane, ¶ 34 ("Before passing data packet 306 to mail gateway 302, firewall 300 add[ed] a signature to data packet 306.").) Keohane explained that the signature was a digital signature added to the packet's header and it indicated that a security action had been performed on the packet. (*Id.*) When a downstream device received a packet, it

Netskope Exhibit 1002

analyzed the signature header to identify what actions had been performed. (*Id.*, ¶ 54.)

200. Additionally, Keohane disclosed that, after a security element received a data packet, it "examine[d] data packet [] to ***determine whether a signature from a trusted security element [was] present***…." (Keohane, ¶ 35 (emphasis added); *see also, e.g.*, ¶ 32 ("When a data packet [was] received, a determination [was] made as to whether a signature of a trusted security element [was] present in the header of the data packet.").) A POSITA would find it obvious that determining whether the signature was from a trusted security element included authenticating the signature and its source. A POSITA would also find it obvious that the signature disclosed in Keohane was a well-known authentication method that would be used by the receiving security element to determine if the packet had been transmitted from another trusted security element. Put simply, if the packet header had a signature that the receiving device could recognize as being from a trusted security element, then the traffic came from another device (second network security appliance) in the CSF. Therefore, it would have been obvious to a POSITA that Keohane's downstream security element (first network security appliance) would be able to authenticate that a signature was from an upstream security element (second network security appliance).

Netskope Exhibit 1002

7. Claim 11

a. "The method of claim 1, wherein the flag is carried in an Internet Protocol (IP) header field of the one or more packets."

201. Keohane disclosed the signature (flag) being carried in an IP header field of the one or more packets. Keohane taught that the disclosed method "may be implemented in a data packet using Internet protocol version 6 (IPv6)…In this type of protocol, data packets include[d] a multi-header chain followed by the payload." (*Id.*, ¶ 40.) Keohane further disclosed that "the signature [was] located in a header of the data packet." (*Id.*, ¶ 12.)

8. Claim 12

a. "The method of claim 1, wherein the one or more packets in which the flag is carried comprises a first packet of a session."

202. While Keohane did not expressly disclose that the packets in which the flag was carried comprised a first packet of a session, a POSITA would have found that limitation obvious based on Keohane's disclosures and the knowledge of a POSITA.

203. As a POSITA would have known, the most commonly used internet protocol, TCP/IP is "connection-oriented": a connection needs to be explicitly established between a client and a server before data can be sent. There is also a mechanism to explicitly close connections and thereby free up resources. The sequence of interactions between any two communicating parties, starting with the

Netskope Exhibit 1002

opening of a connection, encompassing any subsequent data exchange between these parties, and ending with the closing of the connection, is commonly also called a "session".

204. It would have been obvious to a POSITA that for the specific objective of communicating a flag or set of flags between a connection initiator and the connection end point of a TCP/IP connection, including the flag(s) in the first packet would be sufficient since all of the packets in a session travel back and forth between these two end points. They do not, however, necessarily travel along the same route between these end points. A POSITA would have known that if there is just a single route between the two end points of a connection, then all subsequent packets in the connection will travel along the same path as the first packet of the session, whereas if there are multiple alternative routes, the path may be different but the end points will be the same. Hence, it would have been obvious to a POSITA that for connections without alternative routes, the flags would actually be redundant in packets in a session after the initial packet.

9. Claim 15

a. 15[pre]-15[d][i]

205. *See* limitations 1[pre]-1[d][i].

Netskope Exhibit 1002

b. 15[d][ii]: "including determining one or more network security operations that are to be executed by the first network security appliance except the one or more network security operations that are executed by the second network security appliance based on local network polices;"

206. *See* limitation 2[b]. For example, Keohane disclosed that mail gateway (302) determined to perform a virus scan on a data packet but did not determine to perform a trusted source action because firewall (300) had already performed that action. (Keohane, ¶¶ 34-36.) As another example, Keohane disclosed that client (304) determined the security actions performed by firewall (trusted source analysis) and mail gateway (virus scan). (*Id.*, ¶ 37.) Keohane then explained that "[c]lient 304 [] ***perform[ed] an additional security action based on the security actions already performed***." (*Id.* (emphasis added).) That action could be a spyware scan because other devices had not performed that action. (*Id.*; *see also, e.g.*, *id.*, ¶ 40.)

c. 15[e]

207. *See* limitation 1[e].

10. Claim 16

a. "The method of claim 15, wherein the flag indicates the network security operations executed by the second network security appliance."

208. *See* claim 5.

Netskope Exhibit 1002

11. Claim 21

a. 21[pre]-21[d][i]

209. *See* limitations 1[pre]-1[d][i].

b. 21[d][ii]: "including determining, by the first network security appliance, that the one or more network security operations that are executed by the second network security appliance are to be executed by the first network security appliance;"

210. As discussed in limitation 3[b], Keohane disclosed—or at least suggested—that mail gateway (302) determined to perform at least one security operation that had been executed by firewall (300) and skipping the remaining operations. A POSITA would have found it entirely obvious that at any point one could repeat redundant actions. This is just like the 1943 routing slip example cited above in which any recipient of a package with a routing slip can elect to perform any action "locally", including any action that has been flagged as having already been performed by a previous recipient along the route.

c. 21[e]

211. *See* limitation 1[e].

12. Claim 22

a. "The method of claim 21, wherein the flag indicates the network security operations executed by the second network security appliance."

212. *See* claim 5.

Netskope Exhibit 1002

13. Claim 27

a. 27[pre]-27[b][i]

213. *See* limitations 1[pre]-1[b].

b. 27[b][ii]: "wherein the flag indicates the network security operations executed by the second network security appliance;"[5]

214. *See* claim 5.

c. 27[c]-27[e]

215. *See* limitations 1[c], 1[d][i], and 1[e].

G. Ground 6: Keohane and Buruganahalli render claims 6, 7, 17, 18, 23, 24, 28, and 29 obvious.

1. Motivations to combine

216. A POSITA would have been motivated to combine Keohane and Buruganahalli with a reasonable expectation of success for several reasons.

217. First, Keohane and Buruganahalli both concerned coordinating network security actions to reduce redundant actions. For example, Keohane disclosed coordinating operations among security elements along a data path using signatures to avoid redundant security actions. (Keohane, Abstract, ¶¶ 9, 10.) Likewise, Buruganahalli disclosed network devices in a data path coordinating operations

[5] "[T]he network security operations" in this claim lacks the proper antecedent basis and is assumed to refer to the network security operations that the second network security appliance executes on the data packet before transmitting the data packet to the first network security appliance.

Netskope Exhibit 1002

using a metadata tag carried by data packets. (Buruganahalli, Abstract, ¶¶ 27, 28.) Thus, a POSITA looking to improve Keohane would have looked to Buruganahalli's complementary disclosures.

218. Second, combining Keohane and Buruganahalli would have provided a method to reduce redundant security operations with improved flexibility, security, and convenience. Specifically, as described in detail below, a POSITA would be motivated to add Buruganahalli's features of updating and removing a metadata tag to Keohane.

2. Claims 6, 17, 23, and 28

219. These claims recite similar limitations and are, thus, analyzed together.

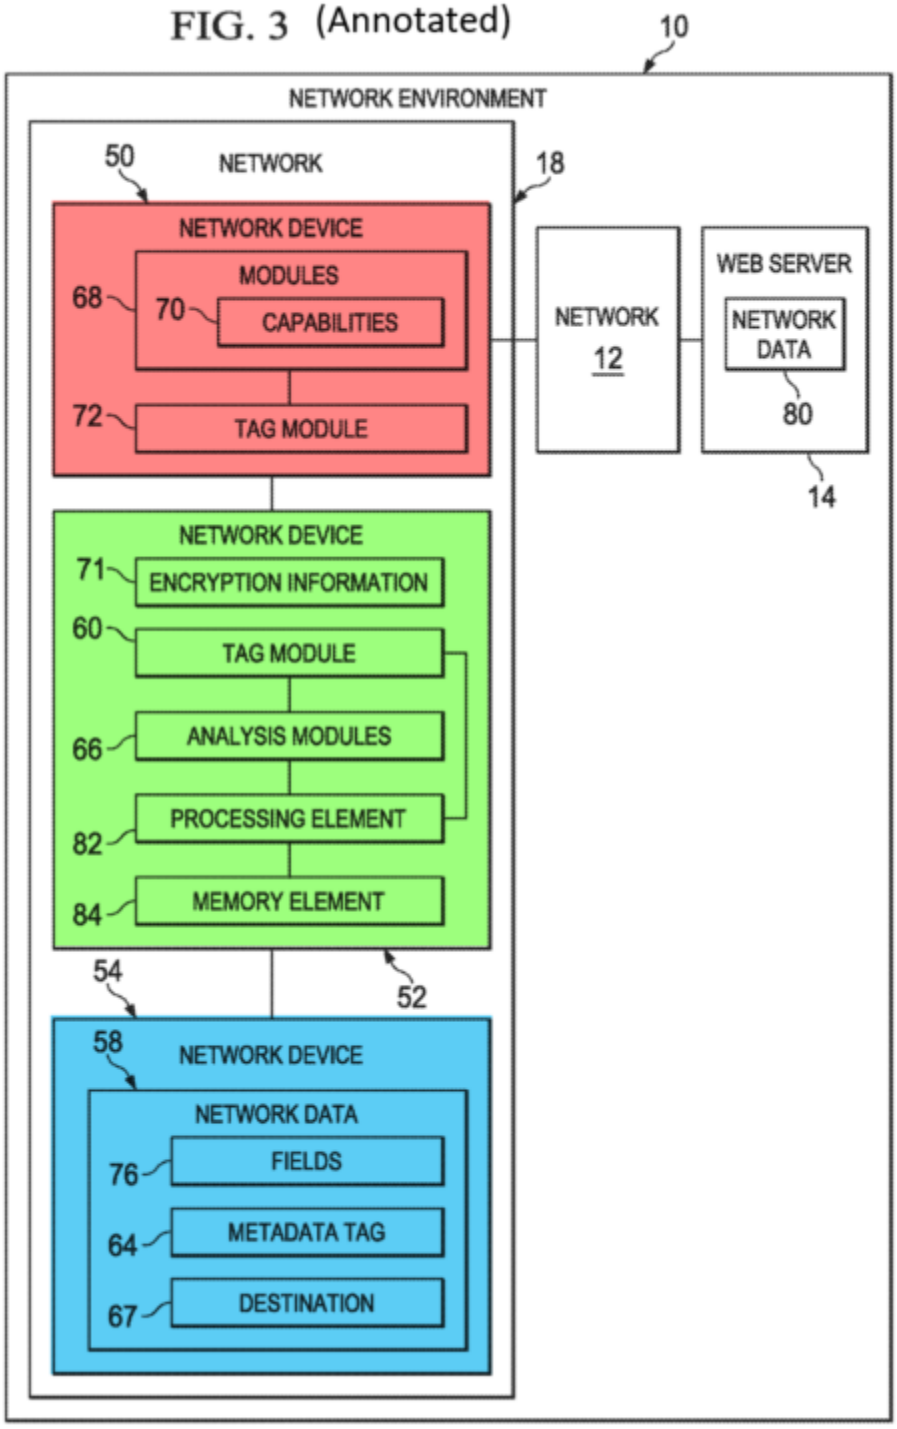
a. 6[a]/17[a]/23[a]/28[a]: "The method of claim [1/15/21/27], further comprising: updating, by the first network security appliance, the flag to indicate the network security operations executed by the first network security appliance;"

220. Buruganahalli disclosed that a network device updated a metadata tag (flag) to indicate the operations performed by the network device.

221. Buruganahalli's Figure 3 (below) depicted "[n]etwork device 50, network device 52, and network device 54 [] connected within network 18." (Buruganahalli, ¶ 32.)

Netskope Exhibit 1002

FIG. 3 (Annotated)



222. First, network device (52) received network data (58) from network device (54). Network data (58) included metadata tag (64). After receiving the data, analysis module (66) performed a "number of actions based on…metadata tag 64"

Netskope Exhibit 1002

on network data (58). (*Id.*, ¶ 38.) Afterwards, tag module (60) "***add[ed] or update[d]*** fields 76 to metadata tag 64 to form metadata tag 64" and "[f]ields 76 may separate the different types of analysis performed on network data 58." (*Id.*, ¶ 40.) A POSITA would have found it obvious that these disclosures in Buruganahalli disclosed updating the fields in the metadata tag to indicate the analysis performed on network data by the device that had received the data.

223. A POSITA would have been motivated to add Buruganahalli's feature of updating the metadata tag to Keohane. If each security element added its own signature indicating the actions it performed to the data packet, that might increase the time it took for the next element to read and process the actions performed by the prior elements, increasing latency and wasting resources. In addition, the size requirement for storing all the signatures would keep increasing along the data path. To prevent these issues, a POSITA would be guided by Buruganahalli's teaching regarding updating the metadata tag to use in Keohane, so that each security element along the data path updated the signature before passing the data packet to the next element instead of adding a new signature.

Netskope Exhibit 1002

b. 6[b]/17[b]/23[b]/28[b]: "and transmitting, by the first network security network appliance, the network traffic with the updated flag to a next network appliance participating in the CSF."

224. Buruganahalli disclosed network device (52) transmitted the network traffic with the updated flag to network device (50), which was another device participating in the CSF.

225. Buruganahalli disclosed that "network device 52 [could] be configured to send network data 58 to network device 50," which obviously included metadata tag 64. (*Id.*, ¶ 46; *see also, e.g.*, *id.*, ¶¶ 44, 48.) Network device (50) participated in a CSF because it was coordinating actions with at least network device (52) to reduce redundancies. (*See, e.g.*, *id.*, ¶ 27.)

226. In the Keohane+Buruganahalli combination, a POSITA would incorporate Buruganahalli's disclosure of updating a tag into Keohane's security element updating the signature and transmitting the updated signature to the next security element along the data path.

3. Claims 7, 18, 24, and 29

227. These claims recite similar limitations and are, thus, analyzed together.

Netskope Exhibit 1002

a. "The method of claim [1/15/21/27], further comprising: removing, by the first network security appliance, the flag from the network traffic; and transmitting, by the first network security network appliance, the network traffic without the flag to a next network appliance."

228. Buruganahalli disclosed a network device removing the metadata tag carried by the network data before sending the data to the next device.

229. Specifically, Buruganahalli disclosed that "last network device to receive the network data prior to sending the data outside of the group of network devices ***may remove the metadata tag***. In other embodiments, any network device ***may remove the metadata tag***." (Buruganahalli, ¶ 29 (emphasis added); *see also, e.g.*, *id.*, ¶ 47.) In operation, Buruganahalli disclosed that the last network device in the group of devices may remove the metadata tag before transmitting the data to an "undesired network device." (*Id.*, ¶ 29.)

230. Moreover, a POSITA would have been motivated to incorporate Buruganahalli's teaching to remove a tag from a tagged data packet before passing the data packet to another device into Keohane for several reasons. First, removing the tag before passing the data packet would keep the network's labeling mechanism private, preventing malicious devices from intercepting the data packet and exploiting the labeling mechanism. Indeed, Buruganahalli recognized this benefit itself, remarking that the tag could be removed before the packet went to "undesired network devices." (*Id.*, ¶ 29; *see also, e.g.*, *id.*, ¶ 47.)

Netskope Exhibit 1002

231. Second, as suggested by Keohane, removing the tag would prevent compatibility issues in subsequence devices that might not be compatible with the tags, which would likely be proprietary. Keohane disclosed that in some embodiments, the packet header in which the signature was stored was "a new type of extension header." (Keohane, ¶ 48.) Keohane's mail gateway (302) could determine that such a header was not compatible with client (304) that only used standard header format. Therefore, mail gateway (302) could remove the signature from the data packet before transmitting the untagged packet to client (304). This provided a motivation to incorporate Buruganahalli, which disclosed removing the metadata tag when the next destination was not within the network's predefined set of destinations, such as when the next destination was not one of the network devices that could read and utilize the tag to cooperate network operations. (*See id.*, ¶¶ 46-47.)

232. For these reasons, a POSITA would be guided by Buruganahalli's teaching regarding removing the metadata tag to use in Keohane, so that a security element on the network data path could remove the signature before passing the data packet to the next device.

Netskope Exhibit 1002

H. Ground 7: Keohane and Kumar render claims 10 obvious.

1. Motivation to Combine

233. Kumar illustrated a well-known technique regarding transmitting data packets between network security appliances: source authentication using a digital signature. In fact, a principal application of digital signatures in general has always been to authenticate a sender. A POSITA would have been motivated to incorporate this common technique into Keohane because, as discussed below, it is a well-known technique for authenticating data packets transmitted between network devices, and was often used predictably and reliably.

234. Additionally, combining Keohane and this technique would have provided a method to coordinate security operations with improved convenience. The motivation for such improvements is described in detail below.

2. Claim 10

a. "The method of claim 1, wherein the flag further comprises or is associated with a signature of the second network security appliance and wherein the method further comprises authenticating, by the first network security appliance, the second network security appliance based on the signature."

235. A POSITA guided by Kumar would find it obvious that Keohane disclosed a flag comprising or being associated with a signature of a transmitting security element, and the receiving security element authenticating the transmitting security element based on the signature.

Netskope Exhibit 1002

236. Kumar illustrated the common knowledge that a transmitting device could use information, such as client ID, session key, public key, etc., to create a digital signature. (*See* Kumar, ¶ 36.) The transmitting device could then use the digital signature to create a security tag and insert the security tag in the data packet. (*See* Kumar, ¶¶ 45, 46) When the receiving device received the data packet, it could authenticate the security tag using various methods. (*Id.*, ¶ 49.) It could recreate the digital signature using the same method as the transmitting device and then comparing the result against the digital signature contained in the security tag. (*See* Kumar, ¶ 50.)

237. A POSITA would have been motivated to incorporate the common technique of source authentication in Keohane's method. Keohane disclosed that after the receiving security element found a signature header in the data packet, it identified a trusted transmitting security element from the signature. (Keohane, ¶ 54.) Therefore, a POSITA would find it obvious that the receiving security element was able to authenticate the signature as being uniquely associated with the trusted transmitting security element. Although Keohane was not explicit about how the receiving security element was able to authenticate the signature and the transmitting security element, a POSITA would find it obvious that Keohane's system utilized a method of digital signature source authentication, which could be the one illustrated in Kumar.

Netskope Exhibit 1002

I. Ground 8: Keohane and Chandra-2 render claims 13, 14, 19, 20, 25, 26, 30, and 31 obvious.

238. The claims challenged under this ground merely require that the network security devices perform specific, well-known operations. As shown below, Chandra-2 proved that those operations were common knowledge.

239. A POSITA would have found it obvious that the Keohane's security elements could perform these operations because each operation was well-known and commonly implemented in computer networks. Indeed, incorporating them into Keohane would have been implementing very common elements to yield predictable results. Thus, a POSITA looking to improve Keohane would have implemented incorporated these common operations, giving Keohane's security elements more range of operations.

1. Claims 13, 19, 25, 30

240. These claims recite similar limitations and are, thus, analyzed together.

a. "The method of claim [1/15/21/27], wherein the one or more network security operations include network traffic logging."

241. Chandra-2 explained that a "***network security device*** c[ould] be a device providing…***logging***." (Chandra-2, ¶ 23 (emphasis added).) A POSITA would have known that network traffic logging is a well-known and widely used network security operation that is especially important for forensic analysis after a vulnerability or breach has been discovered. A POSITA would have found it

Netskope Exhibit 1002

obvious that "logging" provided by a network security device might include "network traffic logging." A POSITA would also have found it obvious that the security element in Keohane, as a network security device, could perform network security logging.

2. Claims 14, 20, 26, 31

242. These claims recite similar limitations and are, thus, analyzed together.

a. "The method of claim [1/15/21/27], wherein the one or more network security operations include one or more of antivirus scanning, intrusion prevention (IPS), content filtering, data leak prevention, antispam, antispyware, reputation-based protections, event correlation, network access control, vulnerability management, load balancing, application control and traffic shaping."

243. Chandra-2 described that a "***network security device*** c[ould] be a device providing one or more of the following features: ***antivirus, intrusion prevention (IPS), content filtering, data leak prevention, antispam, antispyware…reputation-based protections, event correlation, network access control, vulnerability management, load balancing and traffic shaping***." (Chandra-2, ¶ 23 (emphasis added).) In addition, Chandra-2 described that "functionality that may be provided by a ***network appliance*** include…***content filtering***." (*Id.*, ¶ 24 (emphasis added).) Indeed, the '825 patent recognizes that DLP and IPS were well-known. ('825 patent, 9:23-26 ("operations, such as antivirus, DLP, IPS…are *well known* to those skilled in the art and further description thereof

Netskope Exhibit 1002

will be omitted") (emphasis added).) A POSITA would have found it obvious that the security elements in Keohane could perform one or all of these well-known operations.

IV. CONCLUSION

244. Based on the foregoing analysis, it is my opinion that claims 1-7 and 10-31 of the '825 patent are obvious over the prior art.

Netskope Exhibit 1002

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; that these statements were made with knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under 18 U.S.C. § 1001; and further that such willful false statements may jeopardize the validity of the application or any patent issued thereon. I declare under penalty of perjury under the laws of the United States of America that the foregoing is true and correct.

Executed on 1/31/23 in Irvine, CA.

DR. MICHAEL FRANZ

Netskope Exhibit 1002

Michael Franz

444 Information & Computer Science Building
University of California, Irvine
Irvine, CA 92697-3435
franz@uci.edu

Major Research Emphases

- Secure and Trustworthy Computing, Critical Cyber-Infrastructure Protection. Mobile code security; secure and efficient mobile program representations; code verification. Language-based security. Information flow; system-level end-to-end security properties. Moving target defenses, automatically generated software diversity, n-variant systems.
- Software Execution Environments. Compilers, virtual machines, and machine code generation and optimization. On-the-fly, feedback-directed and continuous compilation and optimization; binary translation; trace-based compilation. Code generation for embedded systems, heterogeneous architectures, and mobile computing; compiling for low power consumption. Automatic parallelization. Memory management.
- Software Engineering. Software architectures for secure systems; minimizing the trusted code base. Component-oriented programming languages and their implementation. Software reliability and robustness.

Education

Doctor of Technical Sciences, ETH Zürich, Switzerland; February 1994
Dissertation Title: "Code-Generation On-the-Fly: A Key to Portable Software"
Advisor: Niklaus Wirth

Diplomingenieur, ETH Zürich; May 1989

Academic Appointments

2022 – present	***Distinguished Professor***
2016 – 2022	***Chancellor's Professor***
2006 – present	*Professor of Computer Science* (with tenure)
2001 – 2006	*Associate Professor* (with tenure)
1996 – 2001	*Assistant Professor* Department of Computer Science (since January 2003) Department of Information and Computer Science (until January 2003) **The Donald Bren School of Information & Computer Sciences** University of California, Irvine
2007 – present	*Professor of Electrical Engineering & Computer Science* (by courtesy) Department of Electrical Engineering & Computer Science **The Henry Samueli School of Engineering** University of California, Irvine
1994 – 1995	*Senior Research Associate ("Oberassistent") and Lecturer* Institut für Computersysteme ETH Zürich, Switzerland

Visiting Appointments

June 2019 – June 2022	*Guest Professor* (partially on Sabbatical from UC Irvine) Technical University of Braunschweig, Germany (Host: Prof. Dr. Ina Schäfer)
August 2010 – September 2011	*Visiting Professor* (on Sabbatical from UC Irvine) ETH Zurich, Switzerland (Host: Prof. Dr. Thomas Gross)
January – September 2002	*Visiting Researcher* (on Sabbatical from UC Irvine) University of California, Berkeley (Host: Prof. Dr. George Necula)
Summer Semester 2000	*Visiting Professor* University of Klagenfurt, Austria (Host: Prof. Dr. Laszlo Böszörményi)
Summer Semester 1998	*Visiting Professor* University of Ulm, Germany (Host: Prof. Dr. Peter Schulthess)

Major Professional Honors

- *ACM Charles P. "Chuck" Thacker Breakthrough in Computing Award, Association for Computing Machinery (ACM),* "for the development of just-in-time compilation techniques that enabled fast and feature-rich web services on the internet." This personal award of $100,000 recognizes "individuals who have made a surprising or disruptive leapfrog in computing ideas or technologies;" 2020.
- *Fellow, American Association for the Advancement of Science (AAAS),* "for distinguished contributions to computer science, particularly to the areas of just-in-time compilation and optimization and techniques for computer security;" 2019.
- *Inaugural Fellow, International Federation for Information Processing (IFIP).* The Fellow Award recognizes individuals of the highest professional standing and expertise in one of IFIP's constituent societies (which includes ACM) who have also contributed directly to IFIP. Established in 1960 under the auspices of UNESCO, IFIP is the global organisation for researchers and professionals working in the field of information and communication technologies. IFIP is recognised by the United Nations and links some 50 national and international societies and academies of science with a total membership of over half a million professionals. 2019.
- *Humboldt Research Award*, Alexander von Humboldt Foundation. This personal award of €60,000 is granted in recognition of a researcher's entire achievements to date to academics whose fundamental discoveries, new theories, or insights have had a significant impact on their own discipline and who are expected to continue producing cutting-edge achievements in the future; 2018.
- *Fellow, Association for Computing Machinery (ACM),* "for contributions to just-in-time compilation and optimization and to compiler techniques for computer security;" 2015.
- *Fellow, The Institute of Electrical and Electronics Engineers (IEEE),* "for contributions to just-in-time compilation and to computer security through compiler-generated software diversity;" 2015.
- *IEEE Computer Society Technical Achievement Award,* 2012, "for pioneering contributions to just-in-time compilation and optimization and significantly advancing Web application technology."
- *University of California, Irvine, Distinguished Mid-Career Faculty Award for Research*, 2010. This is the Academic Senate's highest honor for research. One such award at most is given yearly to an Assistant Professor, one to an Associate or Full Professor Step I-IV (the "Mid-Career Award"), and one to a Professor Step V or higher.
- *National Science Foundation CAREER Award*, 1997.
- *Fulbright Scholarship*, 1989.

Teaching Honors

- *Dean's Award for Graduate Student Mentoring*, Donald Bren School of Information and Computer Sciences, UC Irvine, awarded three times in 2007, 2016, and 2022.
- *Outstanding Professor of the Year Award*, Graduating Class of 2007, UC Irvine.

Institutional Affiliations

- Director, *Secure Systems and Software Laboratory,* Donald Bren School of Information and Computer Sciences, UC Irvine; since September 2007.
- Charter Faculty Member, *The California Institute for Telecommunications and Information Technology (Cal-(IT)^2),* one of four California Institutes for Science and Technology.
- Charter Faculty Member, *Security Computing and Networking Center (SCoNCe)* (previously named *Center for Cyber-Security and Privacy*), Donald Bren School of Information and Computer Sciences, UC Irvine.

Noteworthy Contributions With Wide Impact

I am the co-inventor (with my former Ph.D. student Andreas Gal) of the "Trace Tree" compilation technique, which has been transitioned successfully from academic research into one of the most widely distributed open-source projects. From version 3.5 (June 2009) onwards, the JavaScript engine in *Mozilla's Firefox* browser has been based directly on my academic research (see publication C.58).

Furthermore, since version 4.0 (March 2011), the *Firefox* browser additionally contains the "Compartmental Memory Manager" developed in collaboration between my lab and Mozilla (see publication C.69). No fewer than four of my former students with completed Ph.D.s are now employed full-time at Mozilla.

Funding

Current Grants and Awards

- *Office of Naval Research,* N00014-22-1-2232, "Attack Surface Reduction for Binary Programs BinRec Transition;" 1st March 2022 – 28th February 2023, $289,039 (**sole PI**. PM Dr. Samuel M. Weber.
- *DARPA I2O Small Business Technology Transfer (STTR) Program,* W31P4Q-20-C-0052, "IA2: Intent-Capturing Annotations for Isolation and Assurance—Phase II," 1st August 2021 – 31st July 2023, $306,075 (**sole PI** on sub-award from prime contractor Immunant, Inc.). PM Dr. Sergey Bratus.
- *Office of Naval Research,* N00014-21-1-2409, "HONEY-MON: Combining Survivability with Cyber Deception;" 1st May 2021 – 30th April 2024, $1,654,816 (**sole PI)**. PM Dr. Daniel Koller.

Past Grants and Awards

- *DARPA, Assured Micropatching (AMP) Program,* N66001-20-C-4027, "Patches Assured up to Trace Equivalence (PATE)," 15th July 2020 – 14th January 2022, $897,378 (**sole PI** on sub-award from prime contractor Galois, Inc. / Tristan Ravitch). PM Dr. Sergey Bratus.
- *Office of Naval Research,* N00014-17-1-2782, "Attack Surface Reduction for Binary Programs;" 30th September 2017 – 30th September 2021, $3,157,799 (**lead PI**. This is a collaborative award with Herbert Bos of Vrije Universiteit Amsterdam, Netherlands. My share of the award is $2,337,935). PM Dr. Sukarno Mertoguno.
- *DARPA I2O Small Business Technology Transfer (STTR) Program,* W31P4Q-20-C-0052, "IA2: Intent-Capturing Annotations for Isolation and Assurance—Phase I," 1st September 2020 – 31st March 2021, $67,500 (**sole PI** on sub-award from prime contractor Immunant, Inc.). PM Dr. Sergey Bratus.
- *National Science Foundation, Secure & Trustworthy Cyberspace (SaTC) Program,* CNS-1619211, "TWC:Small: Hydra—Hybrid Defenses for Resilient Applications: Practical Approaches Towards Defense In Depth;" 1st July 2016 – 30th June 2020, $499,981 (**sole PI**). PM Dr. Sol J. Greenspan.
- *United States Air Force & Air Force Research Laboratory,* FA8750-16-C-0260, "Thunderlane Phase II;" 1st September 2018 – 12th March 2020, $457,672 (**sole PI** on sub-award for $457,672 from prime contractor Assured Information Security, Inc. / Adam Hovak).
- *DARPA, Cyber Fault-tolerant Attack Recovery (CFAR) Program,* FA8750-15-C-0124, "Robust, Assured Diversity for Software Security (RADSS)." In August of 2017, **award was increased by $217,597** and the duration extended to the end of March 2019. The modified award now runs 13th May 2015 – 31st March 2019, $2,199,227 (**sole PI** on sub-award for $2,199,227 from prime contractor Galois, Inc. / Stephen Magill). PMs Dr. John Everett and Dr. Jacob Torrey.
- *United States Air Force & Air Force Research Laboratory,* FA8750-16-C-0260, "Thunderlane;" 19th September 2016 – 24th May 2017, $45,000 (**sole PI** on sub-award for $45,000 from prime contractor Assured Information Security, Inc. / Philip White).
- *National Science Foundation, Secure & Trustworthy Cyberspace (SaTC) Program,* CNS-1513837, "ENCORE—ENhanced program protection through COmpiler-REwriter cooperation;" 1st July 2015 – 30th June 2018, $1,199,953 (**lead PI**. This was a collaborative award with Matthias Payer of Purdue University and Kevin Hamlen of The University of Texas at Dallas. My share of the award was $619,267.) PM Dr. Sol J. Greenspan.
- *DARPA, Cyber Fault-tolerant Attack Recovery (CFAR) Program,* FA8750-15-C-0124, "Robust, Assured Diversity for Software Security (RADSS)," 13th May 2015 – 1st November 2018, $1,975,630 (**sole PI** on sub-award for $1,975,630 from prime contractor Galois, Inc. / Stephen Magill). PM Dr. John Everett.
- *DARPA, Cyber Fault-tolerant Attack Recovery (CFAR) Program,* FA8750-15-C-0085, "RAVEN," 5th May 2015 – 31st March 2019, $702,271 (**sole PI** on sub-award for $702,271 from prime contractor Apogee Research, LLC / Tiffany Frazier). PM Dr. John Everett.

- *DARPA, I2O Vetting Commodity IT Software and Firmware (VET) Program,* N66001-13-C-4057, "Heterogeneous Compilations for Detection of Malice in Embedded Systems," 1st February 2015 – 30th June 2015, $64,999 (**sole PI** on sub-award for $64,999 from prime contractor Apogee Research, LLC / Tiffany Frazier). PM Dr. Timothy Fraser.

- *National Science Foundation, Computing and Communications Foundations Program,* IIP-1439439, "I-Corps: Hardening Programs Against Cyber Attacks," 1st June 2014 – 30th November 2015, $50,000 (**sole PI**). PM Dr. Rathinda Dasgupta.

- *DARPA, I2O Clean-Slate Design of Resilient, Secure Hosts (CRASH) Program & Transformative Apps Program,* D11PC20024, "Defending Mobile Apps Through Automated Software Diversity." In May of 2014, award was increased by $247,830 and duration extended to 30th September 2015. The modified award now runs 4th February 2011 – 30th September 2015, $2,095,432 (**sole PI**). PMs Dr. Howard Shrobe and Dr. Robert Laddaga.

- *DARPA, I2O Mission-Oriented Resilient Clouds (MRC) Program,* N66001124014, "Meta-Circular Software Diversity for Intrusion Tolerant Clouds," 1st July 2012 – 31st October 2015, $456,809 (**sole PI** on this sub-award for $456,809, which is part of a larger project led by Yair Amir awarded to Johns Hopkins University). PMs Dr. Howard Shrobe and Dr. Robert Laddaga.

- *DARPA, I2O Clean-Slate Design of Resilient, Secure Hosts (CRASH) Program & Transformative Apps Program,* D11PC20024, "Defending Mobile Apps Through Automated Software Diversity." In June of 2012, award was increased by $467,442 and duration extended by an additional year. Modified award now runs 4th February 2011 – 3rd February 2015, $1,847,602 (**sole PI**). PMs Dr. Howard Shrobe, Dr. Robert Laddaga, and Dr. Mari Maeda.

- *National Science Foundation, Computing and Communications Foundations Program,* CCF-1117162, "SHF: CSR: Small: Fine-Grained Modularity and Reuse of VM Components," 1st August 2011 – 31st July 2014, $499,867 (**sole PI**). PM Dr. Bill Pugh.

- *DARPA, Clean-Slate Design of Resilient, Secure Hosts (CRASH) Program & Transformative Apps Program,* D11PC20024, "Defending Mobile Apps Through Automated Software Diversity," 4th February 2011 – 3rd February 2014, $1,380,162 (**sole PI**). PMs Dr. Howard Shrobe and Dr. Mari Maeda.

- *Samsung Telecommunications America,* Richardson, Texas, Agreement No. 51070, "Fine-Grained Modularity and Reuse of Virtual-Machine Components," 1st January 2011 – 31st December 2011, $349,965 (**sole PI**). PM Venky Raju.

- *National Science Foundation, Trusted Computing Program,* CNS-0905684, "Next-Generation Infrastructure for Trustworthy Web Applications," 1st September 2009 – 31st August 2012, $600,000 (**lead PI**, award is split evenly with co-PI C. Flanagan of UC Santa Cruz). PM Dr. Karl Levitt.

- *California MICRO Program* and industrial sponsor *Sun Microsystems, Inc.,* Project No. 07-127, "Trace Compilation for a Server Java Virtual Machine," 24th August 2007 – 30th June 2009, $81,500 ($50,000 gift from sponsor, $31,500 matching cash contribution from MICRO, waiver of overhead charges applies to the total grant amount; **sole PI**).

- *National Intelligence Community, Enterprise Cyber Assurance Program (NICECAP),* FA8750-07-2-0085, "Leveraging Parallel Hardware to Detect, Quarantine, and Repair Malicious Code Injection," 17th May 2007 – 17th August 2009, $1,020,375 (**sole PI**). PM Dr. Carl Landwehr. (This solicitation drew 265 responses, of which 11, including this one, were funded. Among the 11 funded projects, 4 were from M.I.T. and one each from Carnegie-Mellon, Columbia, Cornell, Stanford, and UT Austin. UC Irvine was the only university in the competition to receive a grant awarded to a sole Principal Investigator.)

- *National Science Foundation, Trusted Computing Program,* CNS-0627747, "MLS-VM: Design and Implementation of a Next-Generation Information-Centric Target Platform for Trusted Internet Computing," 1st September 2006 – 31st August 2010, $400,000 (**sole PI**). PM Dr. Helen Gill.

- *National Science Foundation, Embedded and Hybrid Systems Program,* CNS-0615443, "Virtual-Machine Techniques for Resource-Constrained Devices: Reconciling Reliability With Reusability and Low Development Costs in the Embedded Systems Space," 1st July 2006 – 30th June 2010, $300,000 (**sole PI**). PM Dr. Helen Gill.

- *United States Homeland Security Advanced Research Projects Agency (HSARPA),* FA8750-05-2-0216, "Adding Mandatory Access Control to Virtual Machines", 2nd May 2005 – 1st November 2007, $312,483 (**sole PI**). PM Dr. Douglas Maughan. (My proposal was the only one of 80 submissions in the category "Vulnerability Prevention" that got funded by DHS. Overall, the Homeland Security solicitation drew 583 responses, of which 17, including this one, were funded.)

- *California MICRO Program* and industrial sponsor *Microsoft Research,* Project No. 04-032, "Executing Legacy Machine Code on a Safe Virtual Machine," 11th August 2004 – 30th June 2005, $46,881 (waiver of overhead charges applies to the total grant amount; **sole PI**).

- *Deutsche Forschungsgemeinschaft (DFG)* [German National Science Foundation], AM-150/1-3, "SafeTSA: Entwicklung syntaxorientierter Verfahren zur sicheren und effizienten Ausführung von mobilem Code," 1st March 2004 – 28th February 2006, Euro 140,000 (equal co-PI with W. Amme and W. Rossak of the University of Jena, Germany). (This is a new grant that provides continuing support for an earlier DFG-funded research project listed below.)

- *National Science Foundation, Information Technology Research (ITR),* CCR-0205712, "Virtual Power for a Wireless Campus: Orchestrated Modeling, Analysis, Composition and Compilation Strategies for Distributed Embedded Systems," 1st September 2002 – 31st August 2005, $2,000,796 (**lead PI** with C. Krintz and R. Wolski of UC Santa Barbara). PM Dr. Helen Gill. (Award is split $500,000 to Franz, Krintz and Wolski each, with a further $500,000 going to an internal sub-contract at UC Irvine with Senior Personnel P. Chou, N. Dutt, and T. Givargis.)

- *National Science Foundation, Trusted Computing Program,* CCR-TC-0209163, "Practical Language-Based Security, From the Ground Up," 1st August 2002 – 31st July 2005, $300,000 (**sole PI**). PM Dr. Carl Landwehr.

- *DARPA Information Systems Office,* F30602-99-1-0536, "New Approaches to Mobile Code: Reconciling Execution Efficiency With Provable Security," follow-on effort, 22nd June 2002 – 30th September 2003, additional $207,632 (**sole PI**). PM Dr. Jaynarayan H. Lala.

- *Deutsche Forschungsgemeinschaft (DFG)* [German National Science Foundation], AM-150/1-1, "SafeTSA: Entwicklung syntaxorientierter Verfahren zur sicheren und effizienten Ausführung von mobilem Code," 23rd August 2001 – 1st February 2004, Euro 135,000 [corresponding to 270,000 Deutsche Marks] (equal Co-PI with W. Amme and W. Rossak of the University of Jena, Germany).

- *National Science Foundation, Operating Systems and Compilers Program,* CCR-0105710, "Design and Implementation of Component-Oriented Programming Languages," 1st July 2001 – 30th June 2004, $240,000 (**sole PI**). PM Dr. Xiaodong Zhang.

- *Department of Defense, Critical Infrastructure Protection and High Confidence, Adaptable Software (CIP/SW) Research Program of the University Research Initiative,* N00014-01-1-0854, "A Comprehensive Context for Mobile-Code Deployment," 1st May 2001 – 30th September 2004, $981,121, (**lead PI** with B. Fleisch of UC Riverside). PMs Frank Deckelman and Dr. Ralph Wachter. (Award is split $793,201 to Franz and $187,920 to Fleisch. According to the ONR website, "the competition drew 115 white papers, from which 74 proposals were received. After a thorough evaluation by technical expert teams, 20 of these proposals were selected for funding.")

- *National Science Foundation, Next Generation Software Program,* EIA-9975053, "TMO Based Modeling and Design of Reliable Next-Generation Complex Software," 15th August 1999 – 14th August 2002, $550,000 (with K. Kim, Principal Investigator, and P. C.-Y. Sheu, Department of Electrical and Computer Engineering, UC Irvine). PM Dr. Frederica Darema. ($117,000 of the total allocated to co-PI Franz.)

- *National Science Foundation, Operating Systems and Compilers Program,* CCR-9901689, "Graph-Based Mobile-Code Representations for High-Performance Portable Software," 1st September 1999 – 31st August 2002, $180,000 (**sole PI**). PM Dr. Mukesh Singhal.

- *California MICRO Program* and industrial sponsor *Microsoft Research,* Project No. 99-039, "An Infrastructure for Dynamic Optimization at Run-Time," 2nd August 1999 – 30th June 2000, $38,000 (waiver of overhead charges applies to the total grant amount; **sole PI**).
- *DARPA Information Systems Office,* F30602-99-1-0536, "New Approaches to Mobile Code: Reconciling Execution Efficiency With Provable Security," 22nd June 1999 – 21st June 2002, $720,741 (**sole PI**). PM Dr. Jaynarayan H. Lala.
- *National Science Foundation CAREER Award*, CCR-9701400, "Dynamic Optimization of Software Component Systems," 1st March 1997 – 28th February 2001, $205,000 (**sole PI**).

Supplementary Awards

- National Science Foundation, Research Experiences for Undergraduates (REU) Award Supplement for Grant CNS-0905684, Summer 2011, $16,000.
- National Science Foundation, Research Experiences for Undergraduates (REU) Award Supplement for Grant CNS-0905684, Summer 2010, $8,000.
- National Science Foundation, Research Experiences for Undergraduates (REU) Award Supplement for Grant CNS-0627747, Summer 2007, $6,000.
- National Science Foundation, Research Experiences for Undergraduates (REU) Award Supplement for Grant CNS-0615443, Summer 2007, $6,000.
- National Science Foundation, Research Experiences for Undergraduates (REU) Award Supplement for Grant CCR-0205712, Summer 2004, $6,000.
- National Science Foundation, Research Experiences for Undergraduates (REU) Award Supplement for Grant CCR-0205712, Summer 2003, $10,000.
- National Science Foundation, Research Experiences for Undergraduates (REU) Award Supplement for Grant CCR-9701400, Summer 1998, $5,000.

Unrestricted Gifts

- Oracle Corporation, $100,000; May 2016.
- Qualcomm Corporation, $40,000; May 2015.
- Oracle Corporation, $140,000; August 2014.
- Mozilla Corporation, $83,000; August 2014.
- Oracle Corporation, $33,000; September 2013.
- Adobe Corporation, $25,000; August 2011.
- Google Corporation, $61,000; June 2011.
- Adobe Corporation, $35,000; August 2010.
- Adobe Corporation, $40,000; March 2010.
- Mozilla Corporation, $85,000; December 2009.
- Sun Microsystems, $80,000; May 2009.
- Google Corporation, $50,000; January 2008.
- Mozilla Corporation, $85,000; May 2007.

- Intel Corporation, $30,000; April 2006.
- Intel Corporation, $30,000; June 2005.
- Sun Microsystems Laboratories, $56,031; September 2004
- Intel Corporation, $30,000; July 2004.
- Microsoft Research, $33,183, April 2004.

Other Gifts

- Amazon Corporation., $18,000 in Amazon Web Services credit; September 2012.

Publications

Awarded Patents

P.6 P. Larsen (lead), S. Brunthaler, and M. Franz; *Error Report Normalization*; United States Patent No. 10,430,265; filed October 2015, issued October 2019.

P.5 M. Franz (lead), A. Homescu, S. Brunthaler, and P. Larsen; *Code Randomization for Just-In-Time Compilers*; United States Patent No. 9,250,937; filed November 2014, issued February 2016.

P.4 A. Gal (lead) and M. Franz; *Dynamic Incremental Compiler and Method*; United States Patent No. 8,769,511; filed February 2007, issued July 2014.

P.3 M. Franz (lead), W. Amme, and J. von Ronne; *Safe Computer Code Formats And Methods For Generating Safe Computer Code*; United States Patent No. 8,392,897; filed August 2006, issued March 2013.

P.2 M. Franz (lead), A. Gal, and B. Salamat; *Multi-Variant Parallel Program Execution to Detect Malicious Code Injection*; United States Patent No. 8,239,8367 B1; filed March 2008, issued August 2012.

P.1 M. Franz (lead), W. Amme, and J. von Ronne; *Safe Computer Code Formats And Methods For Generating Safe Computer Code*; United States Patent No. 7,117,488; filed October 2001, issued October 2006.

Books

B.2 P. Larsen, S. Brunthaler, L. Davi, A.-R. Sadeghi, and M. Franz; *Automated Software Diversity;* Morgan & Claypool, San Rafael, California, ISBN 978-1-6270-5734-9 (paperback), ISBN 978-1-6270-5755-4 (ebook); December 2015. doi:10.2200/S00686ED1V01Y201512SPT014

B.1 M. Franz; *Code-Generation On-the-Fly: A Key to Portable Software,* Doctoral Dissertation No. 10497, ETH Zürich; published in book form by Verlag der Fachvereine, Zürich, ISBN 3-7281-2115-0; March 1994.

Edited Volumes

E.1 M. Franz and P. Papadimitratos (Eds.); *Trust and Trustworthy Computing (Proceedings of the 9th International Conference, TRUST 2016 Vienna, Austria, August 29–30, 2016);* Springer, Heidelberg, ISBN 978-3-319-45571-6 (paperback), ISBN 978-3-319-45572-3 (ebook); August 2016. doi:10.1007/978-3-319-45572-3

Peer-Reviewed Book Chapters

BC.10 S. Crane, A. Homescu, P. Larsen, H. Okhravi, and M. Franz; "Diversity and Information Leaks;" in P. Larsen and A.-R. Sadeghi (Eds.), *The Continuing Arms Race: Code-Reuse Attacks and Defenses,* ACM Books, Vol. 18, Morgan & Claypool Publishers, ISBN 978-1-97000-183-9, pp. 61–81; 2018. doi:10.1145/3129743.3129747

BC.9 T. Jackson, A. Homescu, S. Crane, P. Larsen, S. Brunthaler, and M. Franz; "Diversifying the Software Stack Using Randomized NOP Insertion;" in S. Jajodia, A. K. Ghosh, V. S. Subrahmanian, V. Swarup, C. Wang, X. S. Wang (Eds.), *Moving Target Defense II: Application of Game Theory and Adversarial Modeling,* Springer Advances in Information Security, Vol. 100, ISBN 978-1-4614-5415-1, pp. 151–174; 2013. doi:DOI 10.1007/978-1-4614-5416-8_8

BC.8 T. Jackson, B. Salamat, A. Homescu, K. Manivannan, G. Wagner, A. Gal, S. Brunthaler, Ch. Wimmer, and M. Franz; "Compiler-Generated Software Diversity;" in S. Jajodia, A.K. Ghosh, V. Swarup, C. Wang, and X.S. Wang (Eds.), *Moving Target Defense: Creating Asymmetric Uncertainty for Cyber Threats*; Springer, ISBN 978-1-4614-0976-2, pp. 77–98; September 2011. doi:10.1007/978-1-4614-0977-9_4

BC.7 M. Franz, W. Amme, M. Beers, N. Dalton, P.H. Fröhlich, V. Haldar, A. Hartmann, P. S. Housel, F. Reig, J. von Ronne, Ch.H. Stork, and S. Zhenochin; "Making Mobile Code Both Safe And Efficient;" in J. Lala (Ed.), *Foundations of Intrusion Tolerant Systems*; IEEE Computer Society Press, ISBN 0-7695-2057-X, pp. 337–356; December 2003. doi:10.1109/FITS.2003.1264941 (Expanded version of conference paper C.15)

BC.6 M. Franz; "A Fresh Look At Low-Power Mobile Computing;" in L. Benini, M. Kandemir, J. Ramanujam (Eds.), *Compilers and Operating Systems for Low Power*; Kluwer Academic Publishers, Boston, ISBN 1-4020-7573-1, pp. 209–220; September 2003. doi:10.1007/978-1-4419-9292-5_12 (Expanded version of conference paper C.18)

BC.5 M. Franz; "Safe Code: It's Not Just For Applets Anymore;" in L. Böszörményi and Peter Schojer (Eds.), *Modular Programming Languages: Proceedings of the Sixth Joint Modular Languages Conference (JMLC 2003),* Klagenfurt, Austria; Springer Lecture Notes in Computer Science, No. 2789, ISBN 3-540-40796-0; pp. 12–22; August 2003. (Full Text of Invited Keynote Address)

BC.4 J. von Ronne, A. Hartmann, W. Amme, and M. Franz; "Efficient Online Optimization by Utilizing Offline Analysis and the SafeTSA Representation;" in J. Powers and J. T. Waldron (Eds.), *Recent Advances in Java Technology: Theory, Application, Implementation*; Computer Science Press, Trinity College Dublin, Dublin, Ireland, ISBN 0-9544145-0-0, pp. 233–241; November 2002. (Expanded version of conference paper C.22)

BC.3 M. Franz; "Oberon: The Overlooked Jewel;" in L. Böszörményi, J. Gutknecht, G. Pomberger (Eds.), *The School of Niklaus Wirth: The Art of Simplicity*; Morgan Kaufmann, San Francisco; ISBN 1-55860-723-4, pp. 41–53; September 2000.

BC.2 J. Gutknecht and M. Franz; "Oberon with Gadgets: A Simple Component Framework;" in M. Fayad, D. Schmidt, R. Johnson (Eds.), *Implementing Application Frameworks: Object-Oriented Frameworks at Work*; Wiley, ISBN 0-4712-5201-8, pp. 323–338; September 1999.

BC.1 M. Franz; "Adaptive Compression of Syntax Trees and Iterative Dynamic Code Optimization: Two Basic Technologies for Mobile-Object Systems;" in J. Vitek and Ch. Tschudin (Eds.), *Mobile Object Systems: Towards the Programmable Internet*; Springer Lecture Notes in Computer Science, No. 1222, ISBN 3-540-62852-5, pp. 263–276; February 1997. doi:10.1007/3-540-62852-5_19

Strongly Reviewed Journal & Magazine Articles

Note: Several conference proceedings have appeared as "special issues" of journals. My contributions to such journal special issues that contain regular conference proceedings are ***not*** *included in this section but are listed under "conference papers" below.*

J.36 B. Belleville, W. Shen, S. Volckaert, A.M. Azab, and M. Franz; "KALD: Detecting Direct Pointer Disclosure Vulnerabilities;" in *IEEE Transactions on Dependable and Secure Computing (TDSC), Vol 18, No. 3, pp. 1369–1377; May/June 2021 (first online publication May 2019).* doi:10.1109/TDSC.2019.2915829

J.35 M. Franz; "Making Multivariant Programming Practical and Inexpensive;" in *IEEE Security and Privacy,* Vol. 16, No. 3, pp. 90–94; May 2018. doi:10.1109/MSP.2018.2701161

J.34 N. Burow, S.C. Carr, J. Nash, P. Larsen, M. Franz, S. Brunthaler, and M. Payer; "Control-Flow Integrity P^3: Protection, Precision, and Performance," in *ACM Computing Surveys (CSUR),* Vol. 50, No. 1, Article No. 16; April 2017. doi:10.1145/3054924

J.33 A. Homescu, T. Jackson, S. Crane, S. Brunthaler, P. Larsen, and M. Franz; "Large-scale Automated Software Diversity—Program Evolution Redux;" in *IEEE Transactions on Dependable and Secure Computing (TDSC),* Vol. 14, No. 2, March/April 2017. doi:10.1109/TDSC.2015.2433252

J.32 G. Wagner, P. Larsen, S. Brunthaler, and M. Franz; "Thinking Inside the Box: Compartmentalized Garbage Collection;" in *ACM Transactions on Programming Languages and Systems (TOPLAS),* Vol. 38, No. 3, Article No. 9; May 2016. doi:10.1145/2866576

J.31 P. Larsen, A. Homescu, S. Brunthaler, and M. Franz; "Automatic Software Diversity;" in *IEEE Security and Privacy,* Vol. 13, No. 2, pp. 30-37; March 2015. doi:10.1109/MSP.2015.23

J.30 G. Savrun-Yeniceri, W. Zhang, H. Zhang, E. Seckler, C. Li, S. Brunthaler, P. Larsen, and M. Franz; "Efficient Hosted Interpreters on the JVM;" in *ACM Transactions on Architecture and Code Optimization (TACO),* Vol. 11, No. 1, Article No. 9; February 2014. doi:10.1145/2532642

J.29 P. Larsen, S. Brunthaler, and M. Franz; "Security through Diversity: Are We There Yet?;" in *IEEE Security and Privacy,* Vol. 12, No. 2, pp. 28–35; March 2014. doi:10.1109/MSP.2013.129

J.28 Ch. Kerschbaumer, E. Hennigan, P. Larsen, S. Brunthaler, and M. Franz; "Information Flow Tracking meets Just-In-Time Compilation;" in *ACM Transactions on Architecture and Code Optimization (TACO),* Vol. 10, No. 4, Article No. 38; December 2013. doi:10.1145/2541228.2555295

J.27 G. Wagner, A. Gal, and M. Franz; "Slimming a Java Virtual Machine by way of Cold Code Removal and Optimistic Partial Program Loading;" in *Science of Computer Programming,* Vol. 76, No. 11, pp. 1037–1053; November 2011. doi:10.1016/j.scico.2010.04.008 (Expanded version of conference paper C.53)

J.26 B. Salamat, T. Jackson, G. Wagner, Ch. Wimmer, and M. Franz; "Run-Time Defense Against Code Injection Attacks Using Replicated Execution;" in *IEEE Transactions on Dependable and Secure Computing (TDSC),* Vol. 8, No. 4; July 2011. doi:10.1109/TDSC.2011.18

J.25 W. Amme, J. von Ronne, Ph. Adler, and M. Franz; "The Effectiveness of Producer-Side Machine-Independent Optimizations for Mobile Code;" in *Software—Practice and Experience)*, Vol. 39, No. 10, pp. 923–946; July 2009. doi:10.1002/spe.v39:10 (Expanded version of conference paper C.40)

J.24 E. Yardimci and M. Franz; "Mostly Static Program Partitioning of Binary Executables;" in *ACM Transactions on Programming Languages and Systems (TOPLAS)*, Vol. 31, No. 5, Article No. 17; June 2009. doi:10.1145/1538917.1538918

J.23 A. Gal, Ch.W. Probst, and M. Franz; "Java Bytecode Verification via Static Single Assignment Form;" *ACM Transactions on Programming Languages and Systems (TOPLAS),* Vol. 30, No. 4, Article No. 21, pp. 1–21; July 2008. doi:10.1145/1377492.1377496

J.22 E. Yardimci and M. Franz; "Dynamic Parallelization of Binary Executables on Hierarchical Platforms;" *The Journal of Instruction-Level Parallelism,* Vol. 10, Paper 6, ISSN 1942-9525, pp. 1–24; June 2008. http://www.jilp.org/vol10/v10paper6.pdf (Expanded version of conference paper C.41)

J.21 M. Franz; "Containing the Ultimate Trojan Horse;" *IEEE Security and Privacy,* Vol. 5, No. 4, pp. 52–56; July 2007. doi:10.1109/MSP.2007.77

J.20 W. Amme, J. von Ronne, and M. Franz; "SSA-Based Mobile Code: Implementation and Empirical Evaluation;" *ACM Transactions on Architecture and Code Optimization (TACO),* Vol. 4, No. 2, Article No. 13; June 2007. doi:10.1145/1250727.1250733

J.19 V. Venkatachalam, M. Franz, and Ch.W. Probst; "A New Way Of Estimating Compute Boundedness And Its Application To Dynamic Voltage Scaling;" *International Journal of Embedded Systems (IJES)*, Vol. 3, No. 1/2, pp. 17–30; 2007. doi:10.1504/IJES.2007.016030

J.18 V. Venkatachalam and M. Franz; "Power Reduction Techniques For Microprocessor Systems;" *ACM Computing Surveys (CSUR),* Vol. 37, No. 3, pp. 195–237; September 2005. doi:10.1145/1108956.1108957

J.17 M. Franz, D. Chandra, A. Gal, V. Haldar, Ch.W. Probst, F. Reig, and N. Wang; "A Portable Virtual Machine Target For Proof-Carrying Code;" *Science of Computer Programming,* (Special Issue on Interpreters, Virtual Machines, and Emulators), Vol. 57, No. 3, pp. 275–294; September 2005. doi:10.1016/j.scico.2004.09.001 (Expanded version of conference paper C.28)

J.16 M. Franz, P.H. Fröhlich, and A. Gal; "Supporting Software Composition at the Programming-Language Level;" *Science of Computer Programming,* (Special Issue on New Software Composition Concepts), Vol. 56, Nos. 1–2, pp. 41–57; April 2005. doi:10.1016/j.scico.2004.11.004

J.15 W. Amme and M. Franz; "Effiziente Codegenerierung für mobilen Code;" *Informatik-Spektrum,* Vol. 26, No. 4, pp. 237–246; August 2003. doi:10.1007/s00287-003-0317-1

J.14 T. Kistler and M. Franz; "Continuous Program Optimization: A Case Study;" *ACM Transactions on Programming Languages and Systems (TOPLAS),* Vol. 25, No. 4, pp. 500–548; July 2003. doi:10.1145/778559.778562

J.13 T. Kistler and M. Franz; "Continuous Program Optimization: Design and Evaluation;" *IEEE Transactions on Computers,* Vol. 50, No. 6, pp. 549–566; June 2001. doi:10.1109/12.931893

J.12 T. Kistler and M. Franz; "Automated Data-Member Layout of Heap Objects to Improve Memory-Hierarchy Performance;" *ACM Transactions on Programming Languages and Systems (TOPLAS),* Vol. 22, No. 3, pp. 490–505; May 2000. doi:10.1145/353926.353937

J.11 T. Kistler and M. Franz; "A Tree-Based Alternative to Java Byte-Codes;" *International Journal of Parallel Programming,* Vol. 27, No. 1, pp. 21–34; February 1999. doi:10.1023/A:1018740018601 (Expanded version of conference paper C.05)

J.10 M. Franz; "The Java Virtual Machine: A Passing Fad?;" *IEEE Software,* Vol. 15, No. 6, pp. 26–29; November 1998. doi:10.1109/52.730834

J.09 M. Franz; "Open Standards Beyond Java: On the Future of Mobile Code for the Internet;" *Journal of Universal Computer Science (j.ucs),* Vol. 4, No. 5, pp. 521–532; May 1998. doi:10.3217/jucs-004-05-0522 (Expanded version of conference paper C.08)

J.08 M. Franz; "Java: Anmerkungen eines Wirth-Schülers" (in German); *Informatik-Spektrum,* Vol. 21, No. 1, pp. 23–26; February 1998. doi:10.1007/s002870050086

J.07 M. Franz and T. Kistler; "Slim Binaries;" *Communications of the ACM,* Vol. 40, No. 12, pp. 87–94; December 1997. doi:10.1145/265563.265576

J.06 M. Franz; "The Programming Language Lagoona: A Fresh Look at Object-Orientation;" *Software-Concepts and Tools,* Vol. 18, No. 1, pp. 14–26; March 1997.

J.05 M. Franz; "Dynamic Linking of Software Components;" *IEEE Computer,* Vol. 30, No. 3, pp. 74–81; March 1997. doi:10.1109/2.573670

J.04 M. Brandis, R. Crelier, M. Franz, and J. Templ; "The Oberon System Family;" *Software—Practice and Experience,* Vol. 25, No. 12, pp. 1331–1366; December 1995.

J.03 M. Franz; "Protocol Extension: A Technique for Structuring Large Extensible Software-Systems;" *Software—Concepts and Tools,* Vol. 16, No. 2, pp. 86–94; July 1995.

J.02 M. Franz; "The Case for Universal Symbol Files;" *Structured Programming,* Vol. 14, No. 3, pp. 136–147; October 1993.

J.01 M. Franz; "Emulating an Operating System on Top of Another;" *Software—Practice and Experience,* Vol. 23, No. 6, pp. 677–692; June 1993.

Strongly Reviewed Conference and Workshop Papers

Note: Several conference proceedings have appeared as "special issues" of journals. They are included in this section rather than under "journal articles" above, and for faster identification have been marked with an asterisk. Talks given at conferences are annotated in this section and are not listed again under "presentations" below.

C.118 G. Kim, M. Franz, and J. Kim; "The Ticket Price Matters in Sharding Blockchain;" in *6th International Workshop on Cryptocurrencies and Blockchain Technology (CBT 2022),* Copenhagen, Denmark, Springer Lecture Notes in Computer Science, to appear; September 2022.

C.117 G. Kim, S. Hong, M. Franz, and D.K. Song; "Improving Cross-Platform Binary Analysis using Representation Learning via Graph Alignment;" in *ACM SIGSOFT International Symposium on Software Testing and Analysis (ISSTA 2022),* Seoul, South Korea, pp. 151–163; July 2022. doi:10.1145/3533767.3534383 61 papers accepted out of 250 submissions = 25.6%

C.116 P. Kirth, M. Dickerson, S. Crane, P. Larsen, A. Dabrowski, D. Gens, Y. Na, S. Volckaert, and M. Franz; "PKRU-Safe: Automatically Locking Down the Heap Between Safe and Unsafe Languages;" in *EuroSys 2022*, Rennes, France, ACM Press, pp. 132–148; April 2022. doi:10.1145/3492321.3519582 **Best Paper Award** 45 papers accepted out of 161 submissions = 27.6%

C.115 C. Deshpande, D. Gens, and M. Franz; "StackBERT: Machine Learning Assisted Static Stack Frame Size Recovery On Stripped and Optimized Binaries;" in *14th ACM Workshop on Artificial Intelligence and Security (AISec 2021),* Seoul, South Korea, pp. 85–95; November 2021. doi:10.1145/3474369.3486865

C.114 A. Dabrowski, K. Pfeffer, M. Reichel, A. Mai, E. Weippl, and M. Franz; "Better Keep Cash in Your Boots—Hardware Wallets Are the New Single Point of Failure;" in *2021 ACM Workshop on Decentralized Finance and Security (DeFi21),* Seoul, South Korea; November 2021. doi:10.1145/3464967.3488588

C.113 K. Pfeffer, A. Mai, A. Dabrowski, M. Gusenbauer, P. Schindler, E. Weippl, M. Franz, and K. Krombholz; "On the Usability of Authenticity Checks for Hardware Security Tokens;" in *USENIX Security 2021,* Vancouver, British Columbia, USENIX Association, ISBN 978-1-939133-24-3, pp. 37–54; August 2021. 248 papers accepted out of 1,319 submissions = 18.8%

C.112 A. Voulimeneas, D. Song, P. Larsen, M. Franz, and S. Volckaert, "dMVX: Secure and Efficient Multi-Variant Execution in a Distributed Setting;" in *14th European Workshop on Systems Security (EuroSec 2021),* Edinburgh, Scotland, ISBN 978-1-4503-8337-0, pp.41–47; April 2021. doi:10.1145/3447852.3458714

C.111 P. Larsen and M. Franz; "Adoption Challenges of Code Randomization," in *7th ACM Workshop on Moving Target Defense (MTD 2020);* November 2020. doi:10.1145/3411496.3421226.

C.110 D. Song, F. Hetzelt, J. Kim, B. Kang, J. Seifert, and M. Franz; "Agamotto: Accelerating Kernel Driver Fuzzing with Lightweight Virtual Machine Checkpoints;" in *USENIX Security 2020,* Boston, Massachusetts, ISBN 978-1-939133-17-5, pp. 2541–2557; August 2020. 157 papers accepted out of 977 submissions = 16%

C.109 Z. Kenjar, T. Frassetto, D. Gens, M. Franz, and A. Sadeghi; "V0LTpwn: Attacking x86 Processor Integrity from Software;" in *USENIX Security 2020,* Boston, Massachusetts, ISBN 978-1-939133-17-5, pp. 1445–1461; August 2020. 157 papers accepted out of 977 submissions = 16%

C.108 P. Rajasekaran, S. Crane, D. Gens, Y. Na, S. Volckaert, and M. Franz; "CoDaRR : Continuous Data Space Randomization against Data-Only Attacks;" in *15th ACM ASIA Conference on Computer and Communications Security (ACM ASIACCS 2020),* Taipei, Taiwan; June 2020 (moved to October 2020). 67 papers accepted out of 308 submissions = 22%

C.107 A. Voulimeneas, D. Song, F. Parzefall, Y. Na, P. Larsen, M. Franz, and S. Volckaert; "Distributed Heterogeneous N-Variant Execution;" in C. Maurice, L. Bilge, G. Stringhini, and N. Neves (Eds)., *Detection of Intrusions and Malware, and Vulnerability Assessment (DIMVA 2020), 17th International Conference,* Lisbon, Portugal, Springer Lecture Notes in Computer Science Vol. 12223, ISBN 978-3-030-52682-5, pp. 326–348; June 2020. doi:10.1007/978-3-030-52683-2_11 13 papers accepted out of 45 submissions = 28.9%

C.106 A. Altinay, J. Nash, T. Kroes, P. Rajasekaran, D. Zhou, A. Dabrowski, D. Gens, Y. Na, S. Volckaert, C. Giuffrida, H. Bos, and M. Franz; "BinRec: Dynamic Binary Lifting and Recompilation;" in *EuroSys 2020,* Heraklion, Greece, ACM Press, ISBN 978-1-4503-6882-7; April 2020. doi:10.1145/3342195.3387550 43 papers accepted out of 234 submissions = 18%

C.105 T. Park, K. Dhondt, D. Gens, Y. Na, S. Volckaert, and M. Franz; "NoJITsu: Locking Down JavaScript Engines;" in *2020 Network and Distributed Systems Security Symposium (NDSS 2020),* San Diego, California, Internet Society, ISBN 1-891562-61-4; February 2020. doi:10.14722/ndss.2020.24262 88 papers accepted out of 506 submissions = 17%

C.104 D.K. Song, J. Lettner, P. Rajasekaran, Y. Na, S. Volckaert, P. Larsen, and M. Franz; "SoK: Sanitizing for Security;" in *40th IEEE Symposium on Security and Privacy,* San Francisco, California, ISBN 978-1-5386-6660-9, pp. 187–207; May 2019. doi:10.1109/SP.2019.00010 84 papers accepted out of 673 submissions plus 10 revised papers from the previous year = 12.5%

C.103 D.K. Song, F. Hetzelt, D. Das, Ch. Spensky, Y. Na, S. Volckaert, G. Vigna, Ch. Kruegel, J.-P. Seifert, and M. Franz; "PeriScope: An Effective Probing and Fuzzing Framework for the Hardware-OS Boundary;" in *2019 Network and Distributed Systems Security Symposium (NDSS 2019),* Internet Society, ISBN 1-891562-55-X, San Diego, California; February 2019. doi:10.14722/ndss.2019.23176 89 papers accepted out of 521 submissions = 17%

C.102 T. Kroes, A. Altinay, J. Nash, Y. Na, S. Volckaert, H. Bos, M. Franz, and Ch. Giuffrida; "BinRec: Attack Surface Reduction Through Dynamic Binary Recovery;" in Y. Shoshitaishvili, M. Naik (Eds.), *FEAST '18: Proceedings of the 2018 Workshop on Forming an Ecosystem Around Software Transformation,* Toronto, Canada, ACM Press, ISBN 978-1-4503-5997-9, pp. 8–13; October 2018. doi:10.1145/3273045.3273050

C.101 B. Belleville, H. Moon, J. Shin, D. Hwang, J.M. Nash, S. Jung, Y. Na, S. Volckaert, P. Larsen, Y. Paek, and M. Franz; "Hardware Assisted Randomization of Data;" in M. Bailey, Th. Holz, M. Stamatogiannakis, and S. Ioannidis (Eds.), *21st International Symposium on Research in Attacks, Intrusions, and Defenses (RAID 2018),* Heraklion, Crete, Greece, Springer Lecture Notes in Computer Science Vol. 11050, ISBN 978-3-030-00469-9, pp. 337–358; September 2018. doi:10.1007/978-3-030-00470-5_16 33 papers accepted out of 145 submissions = 23%

C.100 J. Lettner, D.K. Song, T. Park, S. Volckaert, P. Larsen, and M. Franz; "PartiSan: Fast and Flexible Sanitization via Run-time Partitioning;" in M. Bailey, Th. Holz, M. Stamatogiannakis, and S. Ioannidis (Eds.), *21st International Symposium on Research in Attacks, Intrusions, and Defenses (RAID 2018),* Heraklion, Crete, Greece, Springer Lecture Notes in Computer Science Vol. 11050, ISBN 978-3-030-00469-9, pp. 403–422; September 2018. doi:10.1007/978-3-030-00470-5_19 33 papers accepted out of 145 submissions = 23%

C.99 M. Qunaibit, S. Brunthaler, Y. Na, S. Volckaert and M. Franz; "Accelerating Dynamically-Typed Languages on Heterogeneous Platforms Using Guards Optimization;" in T. Millstein (Ed.), *2018 European Conference on Object-Oriented Programming (ECOOP 2018);* Amsterdam, Netherlands, LIPIcs–Leibniz International Proceedings in Informatics, Vol. 109, ISBN 978-3-95977-079-8, pp. 16:1–16:29; July 2018. doi:10.4230/LIPIcs.ECOOP.2018.16 26 papers accepted out of 66 submissions = 39%

C.98 T. Park, J. Lettner, Y. Na, S. Volckaert and M. Franz; "Bytecode Corruption Attacks Are Real—And How To Defend Against Them;" in C. Giuffrida, S. Bardin, and G. Blanc (Eds)., *Detection of Intrusions and Malware, and Vulnerability Assessment (DIMVA 2018), 15th International Conference,* Saclay, France, Springer Lecture Notes in Computer Science Vol. 10885, ISBN 978-3-319-93410-5, pp. 326–348; June 2018. doi:10.1007/978-3-319-93411-2_15 18 papers accepted out of 59 submissions = 30%

C.97 P. Biswas, A. Di Federico, S.A. Carr, P. Rajasekaran, S. Volckaert, Y. Na, M. Franz, and M. Payer; "Venerable Variadic Vulnerabilities Vanquished;" in *USENIX Security 2017,* Vancouver, British Columbia, ISBN 978-1-931971-40-9, pp. 186–198; August 2017. 85 papers accepted out of 522 submissions = 16%

C.96 S. Volckaert, B. Coppens, B. De Sutter, K. De Bosschere, P. Larsen, and M. Franz; "Taming Parallelism in a Multi-Variant Execution Environment;" in *EuroSys 2017,* Belgrade, Serbia, ISBN 978-1-4503-4938-3, pp. 270–285; April 2017. doi:10.1145/3064176.3064178 41 papers accepted out of 182 valid submissions = 22%

C.95 R. Rudd, R. Skowyra, D. Bigelow, V. Dedhia, Th. Hobson, S. Crane, Ch. Liebchen, P. Larsen, L. Davi, M. Franz, A.-R. Sadeghi, and H. Okhravi; "Address Oblivious Code Reuse: On the Effectiveness of Leakage Resilient Diversity;" in *2017 Network and Distributed System Security Symposium (NDSS 2017),* Internet Society, ISBN 1-891562-46-0, San Diego, California; February 2017. doi:10.14722/ndss.2017.23477 68 papers accepted out of 423 submissions = 16%

C.94 S. Volckaert, B. Coppens, A. Voulimeneas, A. Homescu, P. Larsen, B. De Sutter, and M. Franz; "Secure and Efficient Application Monitoring and Replication;" in *2016 USENIX Annual Technical Conference (ATC 2016),* Denver, Colorado, ISBN 978-1-931971-30-0, pp. 167–179; June 2016. 47 papers accepted out of 266 submissions = 17.6%

C.93 J. Lettner, B. Kollenda, A. Homescu, P. Larsen, F. Schuster, L. Davi, A.-R. Sadeghi, T. Holz, and M. Franz; "Subversive-C: Abusing and Protecting Dynamic Message Dispatch;" in *2016 USENIX Annual Technical Conference (ATC 2016),* Denver, Colorado, ISBN 978-1-931971-30-0, pp. 209–221; June 2016. 47 papers accepted out of 266 submissions = 17.6%

C.92 K. Braden, S. Crane, L. Davi, M. Franz, P. Larsen, Ch. Liebchen, and A.-R. Sadeghi; "Leakage-Resilient Layout Randomization for Mobile Devices;" in *2016 Network and Distributed System Security Symposium (NDSS 2016),* Internet Society, ISBN 1-891562-41-X, San Diego, California; February 2016. doi:10.14722/ndss.2016.23364 60 papers accepted out of 389 submissions = 15.4%

C.91 S. Crane, S. Volckaert, F. Schuster, Ch. Liebchen, P. Larsen, L. Davi, A.-R. Sadeghi, T. Holz, B. De Sutter, and M. Franz; "It's a TRAP: Table Randomization and Protection against Function Reuse Attacks;" in *22nd ACM Conference on Computer and Communications Security (CCS 2015),* Denver, Colorado, ACM Press, ISBN 978-1-4503-3832-5, pp. 243–255; October 2015. doi:10.1145/2810103.2813682 128 papers accepted out of 646 submissions = 19.4%

C.90 M. Conti, S. Crane, L. Davi, M. Franz, P. Larsen, Ch. Liebchen, M. Negro, M. Qunaibit, and A.-R. Sadeghi; "Losing Control: On the Effectiveness of Control-Flow Integrity under Stack Attacks;" in *22nd ACM Conference on Computer and Communications Security (CCS 2015),* Denver, Colorado, ACM Press, ISBN 978-1-4503-3832-5, pp. 952-963; October 2015. doi:10.1145/2810103.2813671 128 papers accepted out of 646 submissions = 19.4%

C.89 G. Savrun-Yeniceri, M. L. Van de Vanter, P. Larsen, S. Brunthaler, and M. Franz; "Efficient and Generic Event-based Profiler Framework for Dynamic Languages;" in *2015 International Conference on Principles and Practices of Programming on the Java platform: Virtual machines, Languages, and Tools (PPPJ'15),* Melbourne, Florida, ACM Press, ISBN 978-1-4503-3712-0, pp. 102–112; September 2015. doi:10.1145/2807426.2807435

C.88 C. Stancu, Ch. Wimmer, S. Brunthaler, P. Larsen, and M. Franz; "Safe and Efficient Hybrid Memory Management for Java;" in *International Symposium on Memory Management 2015 (ISMM'15),* Portland, Oregon, ACM Press, ISBN 978-1-4503-3589-8, pp. 81-92; June 2015. doi:10.1145/2754169.2754185

C.87 S. Crane, Ch. Liebchen, A. Homescu, L. Davi, P. Larsen, A.-R. Sadeghi, S. Brunthaler, and M. Franz; "Readactor: Practical Code Randomization Resilient to Memory Disclosure;" in *36th IEEE Symposium on Security and Privacy,* San Jose, California; May 2015. doi:10.1109/SP.2015.52 55 papers accepted out of 407 submissions = 13.5%

C.86 S. Crane, A. Homescu, S. Brunthaler, P. Larsen, and M. Franz; "Thwarting Cache Side-Channel Attacks Through Dynamic Software Diversity;" in *2015 Network and Distributed System Security Symposium (NDSS 2015),* San Diego, California; February 2015. doi:10.14722/ndss.2015.23264 51 papers accepted out of 302 submissions = 16.9%

C.85 V. Mohan, P. Larsen, S. Brunthaler, K. Hamlen, and M. Franz; "Opaque Control-Flow Integrity;" in *2015 Network and Distributed System Security Symposium (NDSS 2015),* San Diego, California; February 2015. doi:10.14722/ndss.2015.23271 51 papers accepted out of 302 submissions = 16.9%

C.84 M. Murphy, P. Larsen, S. Brunthaler, and M. Franz; "Software Profiling Options and Their Effects on Security Based Code Diversification;" in *First ACM Workshop on Moving Target Defense (MTD 2014),* Scottsdale, Arizona, ACM Press, ISBN 978-1-4503-3150-0, pp. 87–96; November 2014. doi:10.1145/2663474.2663485

C.83 W. Zhang, P. Larsen, S. Brunthaler, and M. Franz; "Accelerating Iterators in Optimizing AST Interpreters;" in *2014 ACM International Conference on Object Oriented Programming Systems Languages & Applications (OOPSLA 2014),* Portland, Oregon, ACM Press, ISBN 978-1-4503-2585-1, pp. 727–743; October 2014. doi:10.1145/2660193.2660223 52 papers accepted out of 186 submissions = 28%

C.82 C. Stancu, Ch. Wimmer, S. Brunthaler, P. Larsen, and M. Franz; "Comparing Points-to Static Analysis with Runtime Recorded Profiling Data;" in *2014 International Conference on Principles and Practices of Programming on the Java platform: Virtual machines, Languages, and Tools (PPPJ 2014),* Cracow, Poland, ACM Press, ISBN 978-1-4503-2926-2, pp. 157–168; September 2014. doi:10.1145/2647508.2647524

C.81 P. Larsen, A. Homescu, S. Brunthaler, and M. Franz; "SoK: Automated Software Diversity;" in *35th IEEE Symposium on Security and Privacy,* San Jose, California, IEEE, ISBN 978-1-4799-4686-0, pp. 276-291; May 2014. doi:10.1109/SP.2014.25 44 papers accepted out of 334 submissions = 13%

C.80 Ch. Kerschbaumer, E. Hennigan, P. Larsen, S. Brunthaler, and M. Franz; "Information Flow Tracking meets Just-In-Time Compilation;" in *High Performance and Embedded Architecture and Compilation Conference (HiPEAC 2014),* Vienna, Austria; January 2014. doi:10.1145/2541228.2555295

C.79 Ch. Kerschbaumer, E. Hennigan, P. Larsen, S. Brunthaler, and M. Franz; "CrowdFlow: Efficient Information Flow Security;" *16th Information Security Conference (ISC 2013),* Dallas, Texas); November 2013. Springer Lecture Notes in Computer Science, Vol. 7807, ISBN 978-3-319-27658-8, pp. 321–340; December 2015. doi:10.1007/978-3-319-27659-5

C.78 A. Homescu, P. Larsen, S. Brunthaler, and M. Franz; "librando: Transparent Code Randomization for Just-in-Time Compilers;" in *20th ACM Conference on Computer and Communications Security (CCS 2013),* Berlin, Germany, ACM Press, ISBN 978-1-4503-2477-9, pp. 993–1004; November 2013. doi:10.1145/2508859.2516675
105 papers accepted out of 530 submissions = 19.8%

C.77 G. Savrun-Yeniceri, W. Zhang, H. Zhang, C. Li, S. Brunthaler, P. Larsen, and M. Franz; "Efficient Interpreter Optimizations for the JVM;" in *2014 International Conference on Principles and Practices of Programming on the Java platform: Virtual machines, Languages, and Tools (PPPJ'13),* Stuttgart, Germany, ACM Press, ISBN 978-1-4503-2111-2, pp. 113–123; September 2013. doi:10.1145/2500828.2500839

C.76 S. Crane, P. Larsen, S. Brunthaler, and M. Franz; "Booby Trapping Software;" in *2013 New Security Paradigms Workshop (NSPW 2013),* Banff, Canada, ACM Press, ISBN 978-1-4503-2582-0, pp. 95–106; September 2013. doi:10.1145/2535813.2535824

C.75 E. Hennigan, Ch. Kerschbaumer, P. Larsen, S. Brunthaler, and M. Franz; "First-Class Labels: Using Information Flow to Debug Security Holes;" in M. Huth, N. Asokan, S. Capkun, I. Flechais, and L. Coles-Kemp (Eds.), *Trust and Trustworthy Computing, 6th International Conference (TRUST 2013),* London, United Kingdom, Springer Lecture Notes in Computer Science, Vol. 7904, ISBN 978-3-642-38907-8, pp. 151–168; June 2013. doi:10.1007/978-3-642-38908-5_12

C.74 Ch. Kerschbaumer, E. Hennigan, P. Larsen, S. Brunthaler, and M. Franz; "Towards Precise and Efficient Information Flow Control in Web Browsers;" in M. Huth, N. Asokan, S. Capkun, I. Flechais, and L. Coles-Kemp (Eds.), *Trust and Trustworthy Computing, 6th International Conference (TRUST 2013),* London, United Kingdom, Springer Lecture Notes in Computer Science, Vol. 7904, ISBN 978-3-642-38907-8, pp. 187–195; June 2013. doi:10.1007/978-3-642-38908-5_14

C.73 A. Homescu, S. Neisius, P. Larsen, S. Brunthaler, and M. Franz; "Profile-guided Automated Software Diversity," in *2013 IEEE/ACM International Symposium on Code Generation and Optimization (CGO 2013),* Shenzhen, China, IEEE, ISBN 978-1-4673-5524-7, pp. 204–214; February 2013. doi:10.1109/CGO.2013.6494997
33 papers accepted out of 117 submissions = 28%

C.72 A. Homescu, M. Stewart, P. Larsen, S. Brunthaler, and M. Franz; "Microgadgets: Size Does Matter In Turing-complete Return-oriented Programming," in *6th USENIX Workshop on Offensive Technologies (WOOT '12),* Bellevue, Washington; August 2012. https://www.usenix.org/system/files/conference/woot12/woot12-final9.pdf

C.71 Ch. Wimmer, S. Brunthaler, P. Larsen, and M. Franz; "Fine-Grained Modularity and Reuse of Virtual Machine Components;" in *11th Annual International Conference on Aspect-Oriented Software Development (AOSD '12),* Potsdam, Germany, ACM Press, ISBN 978-1-4503-1092-5, pp. 203–214; March 2012. doi:10.1145/2162049.2162073

C.70 M. Chang, B. Mathiske, E. Smith, A. Chaudhuri, M. Bebenita, A. Gal, Ch. Wimmer, and M. Franz; "The Impact of Optional Type Information on JIT Compilation Of Dynamically Typed Languages," in *7th Dynamic Languages Symposium (DLS 2011),* Portland, Oregon, ACM Press, ISBN 978-1-4503-0939-4, pp. 13–24; October 2011. doi:10.1145/2047849.2047853

C.69 G. Wagner, A. Gal, Ch. Wimmer, B. Eich, and M. Franz; "Compartmental Memory Management in a Modern Web Browser," in *International Symposium on Memory Management 2011 (ISMM'11),* San Jose, California, ACM Press, ISBN 978-1-4503-0263-0; June 2011. doi:10.1145/1993478.1993496

C.68 T. Jackson, B. Salamat, G. Wagner, Ch. Wimmer, and M. Franz; "On the Effectiveness of Multi-Variant Program Execution for Vulnerability Detection and Prevention;" in *6th International Workshop on Security Measurements and Metrics (MetriSec'10),* Bolzano-Bozen, Italy, ACM Press, ISBN 978-1-4503-0340-8, Article No. 7; September 2010. doi:10.1145/1853919.1853929

C.67 M. Franz; "E unibus pluram: Massive-Scale Software Diversity as a Defense Mechanism;" in *2010 Workshop on New Security Paradigms (NSPW'10),* Concord, Massachusetts, ACM Press, ISBN 978-1-4503-0415-3, pp. 7-16; September 2010. doi:10.1145/1900546.1900550

C.66 M. Bebenita, M. Chang, K. Manivannan, G. Wagner, M. Cintra, B. Mathiske, A. Gal, Ch. Wimmer, and M. Franz; "Trace-Based Compilation in Execution Environments without Interpreters;" in A. Krall, H. Mössenböck (Eds.), *8th International Conference on the Principles and Practice of Programming in Java 2010 (PPPJ'10),* Vienna, Austria, ACM Press, ISBN 978-1-4503-0269-2, pp. 59–68; September 2010. doi:10.1145/1852761.1852771

C.65 K. Manivannan, Ch. Wimmer, and M. Franz; "Decentralized Information Flow Control on a Bare-Metal JVM;" in *Sixth Annual Workshop on Cyber Security and Information Intelligence Research (CSIIRW'10),* Oak Ridge, Tennessee, ACM Press, ISBN 978-1-4503-0017-9; April 2010. doi:10.1145/1852666.1852738

C.64 T. Jackson, Ch. Wimmer, and M. Franz; "Multi-Variant Program Execution for Vulnerability Detection and Analysis;" in *Sixth Annual Workshop on Cyber Security and Information Intelligence Research (CSIIRW'10),* Oak Ridge, Tennessee, ACM Press, ISBN 978-1-4503-0017-9; April 2010. doi:10.1145/1852666.1852708

C.63 Ch. Wimmer and M. Franz; "Linear Scan Register Allocation on SSA Form;" in *The Eighth International Symposium on Code Generation and Optimization (CGO 2010),* Toronto, Canada, ACM Press, ISBN 978-1-60558-635-9, pp. 170–179; April 2010. doi:10.1145/1772954.1772979

C.62 A. Yermolovich, Ch. Wimmer, and M. Franz; "Optimization of Dynamic Languages Using Hierarchical Layering of Virtual Machines;" in *5th Symposium on Dynamic Languages (DLS 2009)*, Orlando, Florida, ACM Press, ISBN 978-1-60558-769-1, pp. 79–88; October 2009. doi:10.1145/1640134.1640147 Awarded **"Most Notable Paper of DLS 2009"** (Test of Time Award) at DLS 2019.

C.61 Ch. Wimmer, M. Cintra, M. Bebenita, M. Chang, A. Gal, and M. Franz; "Phase Detection using Trace Compilation;" in *7th International Conference on the Principles and Practice of Programming in Java 2009 (PPPJ 2009),* Calgary, Alberta, ACM Press, ISBN 978-1-60558-598-7, pp. 172–181; August 2009. doi:10.1145/1596655.1596683

C.60 Ch. Kerschbaumer, G. Wagner, Ch. Wimmer, A. Gal, Ch. Steger, and M. Franz; "SlimVM: A Small Footprint Java Virtual Machine for Connected Embedded Systems;" in *7th International Conference on the Principles and Practice of Programming in Java 2009 (PPPJ 2009),* Calgary, Alberta, ACM Press, ISBN 978-1-60558-598-7, pp. 133–142; August 2009. doi:10.1145/1596655.1596678

C.59 M. Bebenita, M. Chang, A. Gal, and M. Franz; "Stream-Based Dynamic Compilation for Object-Oriented Languages;" in M. Oriol and B. Meyer (Eds.), *Objects, Components, Models and Patterns,* 47th International Conference (TOOLS-EUROPE 2009), Zurich, Switzerland, Springer Lecture Notes in Business Information Processing (LNBIP), Vol. 33, ISBN 978-3-642-02570-9, pp. 77–95; June 2009. doi:10.1007/978-3-642-02571-6_6

C.58 A. Gal, B. Eich, M. Shaver, D. Anderson, B. Kaplan, G. Hoare, D. Mandelin, B. Zbarsky, J. Orendorff, J. Ruderman, E. Smith, R. Reitmaier, M. R. Haghighat, M. Bebenita, M. Chang, and M. Franz; "Trace-based Just-in-Time Type Specialization for Dynamic Languages;" in *Programming Language Design and Implementation (PLDI 2009),* Dublin, Ireland, ACM Press, ISBN 978-1-60558-392-1, pp. 465–478; June 2009. doi:10.1145/1542476.1542528

C.57 B. Salamat, T. Jackson, A. Gal, and M. Franz; "Intrusion Detection Using Parallel Execution and Monitoring of Program Variants in User-Space;" in *EuroSys'09,* Nuremberg, Germany, ACM Press, ISBN 978-1-60558-482-9, pp. 33–46; April 2009. doi:10.1145/1519065.1519071

C.56 M. Franz; "Information-Flow Aware Virtual Machines: Foundations For Trustworthy Computing;" in *Cybersecurity Applications and Technologies Conference for Homeland Security (CATCH 2009),* Washington, D.C., IEEE Computer Society Publications, ISBN 978-0-7695-3568-5, pp. 91–96; March 2009. doi:10.1109/CATCH.2009.45

C.55 M. Chang, E. Smith, R. Reitmaier, A. Gal, M. Bebenita, Ch. Wimmer, B. Eich, and M. Franz; "Tracing for Web 3.0: Trace Compilation for the Next Generation Web Applications;" in *2009 ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments (VEE 2009),* Washington, D.C., ACM Press, ISBN 978-1-60558-375-4, pp. 71–80; March 2009. doi:10.1145/1508293.1508304

C.54 L. Wang and M. Franz; "Automatic Partitioning of Object-Oriented Programs for Resource-Constrained Mobile Devices with Multiple Distribution Objectives;" in *The 14th IEEE International Conference on Parallel and Distributed Systems (ICPADS'08),* Melbourne, Victoria, Australia, December 2008. doi:10.1109/ICPADS.2008.84

C.53 G. Wagner, A. Gal, and M. Franz; "SlimVM: Optimistic Partial Program Loading for Connected Embedded Java Virtual Machines;" in L. Veiga, V. Amaral, N. Horspool, and G. Cabri (Eds.), *Principles and Practice of Programming in Java 2008 (PPPJ 2008),* Proceedings of the 6th International Conference, Modena, Italy, ACM Press, ISBN 978-1-60558-223-8, pp. 117–126; September 2008. doi:10.1145/1411732.1411749 **(Best Paper Award)**

C.52 A. Yermolovich, A. Gal, and M. Franz; "Portable Execution of Legacy Binaries on the Java Virtual Machine;" in L. Veiga, V. Amaral, N. Horspool, and G. Cabri (Eds.), *Principles and Practice of Programming in Java 2008 (PPPJ 2008),* Proceedings of the 6th International Conference, Modena, Italy, ACM Press, ISBN 978-1-60558-223-8, pp. 63–72; September 2008. doi:10.1145/1411732.1411742

C.51 A. Noll, A. Gal, and M. Franz; "CellVM: A Homogeneous Virtual Machine Runtime System for a Heterogeneous Single-Chip Multiprocessor;" in *2008 ISCA Workshop on Cell Systems and Applications,* Beijing, China; June 2008.

C.50 B. Salamat, A. Gal, and M. Franz; "Reverse Stack Execution in a Multi-Variant Execution Environment;" in *2008 DSN Workshop on Compiler and Architectural Techniques for Application Reliability and Security (CATARS'08),* Anchorage, Alaska; June 2008.

C.49 B. Salamat, A. Gal, T. Jackson, K. Manivannan, G. Wagner, and M. Franz; "Multi-Variant Program Execution: Using Multi-Core Systems to Defuse Buffer-Overflow Vulnerabilities;" in *2008 International Conference on Complex, Intelligent and Software Intensive Systems (CISIS 2008),* Barcelona, Spain, IEEE Computer Society Press, ISBN 978-0-7695-3109-0, pp. 843–848; March 2008. doi:10.1109/CISIS.2008.136

C.48 M. Franz; "Eliminating Trust From Application Programs By Way Of Software Architecture;" in *Software Engineering 2008 (SE 2008),* Munich, Germany, Lecture Notes in Informatics (LNI) No. 121, GI-Edition, Gesellschaft für Informatik, Bonn, ISBN 978-3-88579-215-4, pp. 112–126; February 2008.

C.47 M. Franz; "Understanding and Countering Insider Threats In Software Development;" in P. Kropf, M. Benyoucef, and H. Mili (Eds.), *2008 International Montreal Conference on e-Technologies (MCETECH 2008),* Montreal, Canada, IEEE Computer Society Publications, ISBN 978-0-7695-3082-6, pp. 81–90; January 2008. doi:10.1109/MCETECH.2008.32

C.46 D. Chandra and M. Franz; "Fine-Grained Information Flow Analysis and Enforcement in a Java Virtual Machine;" in *23rd Annual Computer Security Applications Conference (ACSAC 2007),* Miami Beach, Florida, IEEE Computer Society Publications, ISBN 0-7695-3060-5, pp. 463–474; December 2007. doi:10.1109/ACSAC.2007.37

C.45 M. Bebenita, A. Gal, and M. Franz; "Implementing Fast JVM Interpreters In Java Itself;" in V. Amaral, L. Veiga, L. Marcelino, and H. C. Cunningham (Eds.), *Principles and Practices of Programming in Java, Proceedings of the 5th International Conference (PPPJ 2007),* Lisbon, Portugal, ACM Press, ISBN 978-1-59593-672-1, pp. 145–154; September 2007. doi:10.1145/1294325.1294345

C.44 A. Gal, M. Bebenita, and M. Franz; "One Method At A Time Is Quite a Waste of Time;" in *Implementation, Compilation, Optimization of Object-Oriented Languages, Programs and Systems (ICOOOLPS'2007),* Berlin, Germany, Report No. 2007-5, Technische Universität Berlin, ISSN 1436-9915, pp. 11–16; July 2007.

C.43 M. Franz, A. Gal, and Ch.W. Probst; "Automatic Generation of Machine Emulators: Efficient Synthesis of Robust Virtual Machines for Legacy Software Migration;" in W.-G. Bleek, J. Raasch, H. Züllighoven (Eds.), *Software Engineering 2007 (SE 2007),* Hamburg, Germany, Lecture Notes in Informatics (LNI) No. 105, GI-Edition, Gesellschaft für Informatik, Bonn, ISBN 978-3-88579-199-7, pp. 83–94; March 2007.

C.42 A. Gal, Ch.W. Probst, and M. Franz; "HotpathVM: An Effective JIT Compiler for Resource-Constrained Devices;" in *Second International Conference on Virtual Execution Environments (VEE 2006),* Ottawa, Canada, ACM Press, ISBN 1-59593-332-6, pp. 144–153; June 2006. doi:10.1145/1134760.1134780

C.41 E. Yardimci and M. Franz; "Dynamic Parallelization of Binary Executables on Hierarchical Platforms;" in *Computing Frontiers 2006,* Ischia, Italy, ACM Press, ISBN 1-59593-302-6, pp. 127–138; May 2006. doi:10.1145/1128022.1128040

C.40 Ph. Adler, W. Amme, M. Franz, and J. von Ronne; "Producer-Side Platform-Independent Optimizations and Their Effects on Mobile-Code Performance;" in *The 10th IEEE Annual Workshop on Interaction between Compilers and Computer Architectures (INTERACT-10),* Austin, Texas; February 2006. 14 submitted, 8 accepted

C.39 V. Haldar, D. Chandra, and M. Franz; "Dynamic Taint Propagation for Java;" in *Twenty-First Annual Computer Security Applications Conference (ACSAC 2005),* Tucson, Arizona, IEEE Computer Society Publications, ISBN 0-7695-2461-3, pp. 274–282; December 2005. doi:10.1109/CSAC.2005.21

C.38 V. Haldar, D. Chandra, and M. Franz; "Practical, Dynamic Information-Flow for Virtual Machines;" in *2nd International Workshop on Programming Language Interference and Dependence (PLID'05),* London, England; September 2005.

C.37 A. Gal, Ch.W. Probst, and M. Franz; "Average Case vs. Worst Case Margins of Safety in System Design;" in Ch. F. Hempelmann, V. Raskin (Eds.), *New Security Paradigms Workshop 2005 (NSPW 2005),* Lake Arrowhead, California, ACM Press, ISBN 1-59593-317-4, pp. 25–32; September 2005. doi:10.1145/1146269.1146279

C*.36 A. Gal, Ch.W. Probst, and M. Franz; "Structural Encoding of Static Single Assignment Form;" in *4th International Workshop on Compiler Optimization Meets Compiler Verification (COCV'05),* Edinburgh, Scotland; April 2005. Revised post-conference version published as *Electronic Notes in Theoretical Computer Science (ENTCS),* Vol. 141, No. 2, pp. 85–102; November 2005. doi:10.1016/j.entcs.2005.02.045

C*.35 W. Amme, J. von Ronne, and M. Franz; "Quantifying the Benefits of SSA-Based Mobile Code;" in *4th International Workshop on Compiler Optimization Meets Compiler Verification (COCV'05),* Edinburgh, Scotland; April 2005. Revised post-conference version published as *Electronic Notes in Theoretical Computer Science (ENTCS),* Vol. 141, No. 2, pp. 103–119; November 2005. doi:10.1016/j.entcs.2005.02.046

C*.34 A. Gal, Ch.W. Probst, and M. Franz; "Integrated Java Bytecode Verification;" in *First International Workshop on Abstract Interpretation of Object-Oriented Programming Languages (AIOOL'05),* Paris, France; January 2005. Also published as *Electronic Notes in Theoretical Computer Science (ENTCS),* Vol. 131, pp. 27–38; May 2005. doi:10.1016/j.entcs.2005.01.020

C.33 V. Haldar and M. Franz; "Symmetric Behavior-Based Trust: A New Paradigm for Internet Computing;" in Carla Marceau, Simon Foley (Eds.), *New Security Paradigms Workshop 2004 (NSPW 2004),* White Point, Nova Scotia, ACM Press, ISBN 1-59593-076-0, pp. 79–84; September 2004. doi:10.1145/1065907.1066039 (This paper was one of 4 papers selected for the "Highlights of NSPW 2004" session at *ACSAC 2004*.)

C.32 J. von Ronne, N. Wang, and M. Franz; 'Interpreting Programs in Static Single Assignment Form;" in *ACM SIGPLAN 2004 Workshop on Interpreters, Virtual Machines and Emulators (IVME'04),* Washington, D.C., pp. 23–30; June 2004. doi:10.1145/1059579.1059585

C.31 M. Beers, Ch.H. Stork, and M. Franz; "Efficiently Verifiable Escape Analysis;" in M. Odersky (Ed.), *18th European Conference on Object-Oriented Programming (ECOOP 2004),* Oslo, Norway, Springer Lecture Notes in Computer Science, Vol. 3086, ISBN 3-540-22159-X, pp. 75–95; June 2004. doi:10.1007/b98195

C.30 V. Haldar, D. Chandra, and M. Franz; "Semantic Remote Attestation: A Virtual Machine Directed Approach to Trusted Computing;" in *3rd USENIX Virtual Machine Research & Technology Symposium (VM'04),* San Jose, California, ISBN 1-931971-20-X, pp. 29–41; May 2004. **(Best Paper Award)**

C.29 Ch.W. Probst, A. Gal, and M. Franz; "Code Generating Routers: A Network-Centric Approach to Mobile Code;" in *2003 IEEE 18th Annual Workshop on Computer Communications (CCW'2003),* Dana Point, California, IEEE Press, ISBN 0-7803-8239-0, pp. 179–186; October 2003.

C.28 M. Franz, D. Chandra, A. Gal, V. Haldar, F. Reig, and N. Wang; "A Portable Virtual Machine Target For Proof-Carrying Code;" in *ACM SIGPLAN 2003 Workshop on Interpreters, Virtual Machines and Emulators (IVME'03),* San Diego, California, pp. 24–31; June 2003. doi:10.1145/858570.858573

C.27 J. von Ronne, A. Hartmann, W. Amme, and M. Franz; "Code Annotation for Safe and Efficient Dynamic Object Resolution;" in *2003 Workshop on Compiler Optimization meets Compiler Verification (COCV 2003),* Warsaw, Poland, April 2003. doi:10.1016/S1571-0661(05)82597-6

C.26 A. Gal, M. Franz, and D. Beuche, "Learning from Components: Fitting AOP for System Software;" in *Second AOSD 2003 Workshop on Aspects, Components, and Patterns for Infrastructure Software (ACP4IS'2003),* Boston, Massachusetts; March 2003.

C.25 V. Haldar, Ch.H. Stork, and M. Franz; "The Source Is The Proof;" in C. Serban, S. Saydjari (Eds.), in *ACM SIGSAC 2002 Workshop on New Security Paradigms (NSPW 2002),* Virginia Beach, Virginia, ACM Press, ISBN 1-58113-598-X, pp. 69–73; September 2002. doi:10.1145/844102.844114 (This paper was one of 4 papers selected for the "Best of NSPW 2002" session at ACSAC 2002.)

C.24 V. Haldar and M. Franz; "Towards Trusted Systems, From The Ground Up;" in *Tenth ACM SIGOPS European Workshop: Can We Really Depend On An OS? (EW 2002),* Saint-Emilion, France, ACM Press, pp. 251–254; September 2002. doi:10.1145/1133373.1133426

C.23 M. Franz; "Enhancing Class Files: A Migration Path to Better Mobile-Code Representations;" in D. Bakken (Ed.), *DSN Fast Abstracts, International Conference on Dependable Systems and Networks (DSN 2002),* Washington, D.C., June 2002.

C.22 J. von Ronne, A. Hartmann, W. Amme, and M. Franz; "Efficient Online Optimization by Utilizing Offline Analysis and the SafeTSA Representation;" in *Proceedings of the 2nd Workshop on Intermediate Representation Engineering for Virtual Machines (IRE 2002),* Dublin, Ireland, June 2002.

C.21 D. Chandra, Ch. Fensch, W.-K. Hong, L. Wang, E. Yardimci, and M. Franz; "Code Generation at the Proxy: An Infrastructure-Based Approach to Ubiquitous Mobile Code;" in *Fifth ECOOP Workshop on Object-Orientation and Operating Systems (ECOOP-OOOSWS 2002),* Málaga, Spain, June 2002.

C.20 A. Gal, P.H. Fröhlich, and M. Franz; "An Efficient Execution Model for Dynamically Reconfigurable Component Software;" in *7th International Workshop on Component-Oriented Programming (WCOP 2002),* Málaga, Spain, June 2002.

C.19 P.H. Fröhlich and M. Franz; "On Certain Basic Properties of Component-Oriented Programming Languages;" in *First OOPSLA Workshop on Language Mechanisms for Programming Software Components,* Tampa Bay, Florida; October 2001.

C.18 M. Franz; "A Fresh Look At Low-Power Mobile Computing;" in *Compilers and Operating Systems for Low Power 2001 (COLP 01),* Barcelona, Spain, pp. 15.1–15.6; September 2001.

C*.17 Ch.H. Stork, P. S. Housel, V. Haldar, N. Dalton, and M. Franz; "Towards Language Agnostic Mobile Code;" in N. Benton and A. Kennedy (Eds.), *First Workshop on Multi-Language Infrastructure and Interoperability (BABEL'01),* Florence, Italy; September 2001. Also published as *Electronic Notes in Theoretical Computer Science (ENTCS),* Vol. 59, No. 1, pp. 142–157; November 2001. doi:10.1016/S1571-0661(05)80458-X

C.16 W. Amme, N. Dalton, J. von Ronne, and M. Franz; "SafeTSA: A Type Safe and Referentially Secure Mobile-Code Representation Based on Static Single Assignment Form;" in *ACM Sigplan Conference on Programming Language Design and Implementation (PLDI 2001),* Snowbird, Utah, pp. 137–147; June 2001. doi:10.1145/378795.378825

C.15 W. Amme, N. Dalton, P.H. Fröhlich, V. Haldar, P. S. Housel, J. von Ronne, Ch.H. Stork, S. Zhenochin, and M. Franz; "Project transPROse: Reconciling Mobile-Code Security With Execution Efficiency;" in *The Second DARPA Information Survivability Conference and Exhibition (DISCEX II),* Anaheim, California; IEEE Computer Society Press, ISBN 0-7695-1212-7, pp. II.196–II.210; June 2001. doi:10.1109/DISCEX.2001.932172

C.14 W. Amme, N. Dalton, M. Franz, and J. von Ronne; "A Type-Safe Mobile Code Representation Aimed At Supporting Dynamic Optimization At The Target Site;" in *Third ACM Workshop on Feedback-Directed and Dynamic Optimization (FDDO-3),* Monterey, California, December 2000. (**Best Paper Award**; additionally and independently, the paper's presentation was one of three simultaneous winners of the **Best Presentation Award**.)

C.13 P.H. Fröhlich and M. Franz; "Stand-Alone Messages: A Step Towards Component-Oriented Programming Languages;" in J. Gutknecht and W. Weck (Eds.), *Modular Programming Languages: Proceedings of the Fifth Joint Modular Languages Conference (JMLC 2000),* Zurich, Switzerland; Springer Lecture Notes in Computer Science, No. 1891, ISBN 3-540-67958-8, pp. 90–103; September 2000. doi:10.1007/10722581_9

C.12 M. Franz, P.H. Fröhlich, and T. Kistler; "Towards Language Support for Component-Oriented Real-Time Programming;" in *The Fifth IEEE International Workshop on Object-Oriented Real-Time Dependable Systems (WORDS'99F),* Monterey, California, November 1999; IEEE Computer Society Press, ISBN 0-7695-0616-X; April 2000. doi:10.1109/WORDSF.1999.842343

C.11 T. Kistler and M. Franz; "Computing the Similarity of Profiling Data: Heuristics for Guiding Adaptive Optimizations;" in *Proceedings of the Workshop on Profile and Feedback-Directed Optimization,* Paris, France, October 1998.

C.10 M. Franz; "On the Architecture of Software Component Systems;" in R.N. Horspool (Ed.), *Systems Implementation 2000, (Proceedings of the IFIP TC2 WG2.4 Working Conference on Systems Implementation 2000: Languages, Methods and Tools, Berlin, Germany),* Chapman & Hall, ISBN 0-412-83530-4, pp. 207–220; February 1998.

C.09 M. Franz and T. Kistler; "Does Java Have Alternatives?;" in D.J. Richardson and D. Wile (Eds.), *Proceedings of the Third California Software Symposium (CSS'97),* Irvine, California, pp. 5–10; November 1997.

C.08 M. Franz; "Beyond Java: An Infrastructure for High-Performance Mobile Code on the World Wide Web;" in S. Lobodzinski and I. Tomek (Eds.), *Proceedings of WebNet 97,* World Conference of the WWW, Internet, and Intranet, Association for the Advancement of Computing in Education; ISBN 1-880094-27-4, pp. 33–38; October 1997. (**Best Paper Award**)

C.07 M. Franz; "Run-Time Code Generation as a Central System Service;" in The Sixth Workshop on Hot Topics in Operating Systems (HotOS VI), IEEE Computer Society Press, ISBN 0-8186-7834-8, pp. 112–117; May 1997. doi:10.1109/HOTOS.1997.595192

C.06 M. Franz; "Toward an Execution Model for Component Software;" in Proceedings of the First International Workshop on Component-Oriented Programming (WCOP 1996), subsequently published as M. Mühlhäuser (Ed.), *Special Issues in Object-Oriented Programming: Workshop Reader of the 10th European Conference on Object-Oriented Programming (ECOOP'96),* dpunkt Verlag, Heidelberg, ISBN 3-920993-67-5, pp. 144–149; March 1997.

C.05 T. Kistler and M. Franz; "A Tree-Based Alternative to Java Byte-Codes;" in *Proceedings of the International Workshop on Security and Efficiency Aspects of Java,* Eilat, Israel; January 1997.

C.04 M. Franz; "Compiler Optimizations Should Pay for Themselves;" in P. Schulthess (Ed.), *Advances in Modular Languages: Proceedings of the Joint Modular Languages Conference,* Universitätsverlag Ulm, ISBN 3-89559-220-X, pp. 111–121; September 1994.

C.03 M. Franz; "Technological Steps toward a Software Component Industry;" in J. Gutknecht (Ed.), *Programming Languages and System Architectures: Proceedings of the International Conference,* Zurich, Switzerland, Springer Lecture Notes in Computer Science, No. 782, pp. 259–281; March 1994. doi:10.1007/3-540-57840-4_36

C.02 M. Franz; "Immediate Object-Level Software Reuse on Different Target Architectures using Fast On-The-Fly Code Generation;" in *Position Paper Collection of the Second International Workshop on Software Reusability,* Lucca, Italy; March 1993.

C.01 M. Franz and S. Ludwig; "Portability Redefined;" in *Proceedings of the Second International Modula-2 Conference,* Loughborough, England, pp. 216–224; September 1991.

Selected Further Conferences, Workshops, and Other Publications

*Note: Technical reports that have subsequently been published as book chapters, conference papers, or journal articles are **not** listed again here.*

CPC04 E. Yardimci, N. Dalton, Ch. Fensch, and M. Franz; "Azure: A Virtual Machine for Improving Execution of Sequential Programs on Throughput-Oriented Explicitly-Parallel Processors;" in *Proceedings of the 11th International Workshop on Compilers for Parallel Computers (CPC 2004),* Seeon, Germany, Shaker Verlag, pp. 61–174; July 2004.

PLOS04 A. Gal, Ch.W. Probst, and M. Franz; "Executing Legacy Applications on a Java Operating System;" in *Proceedings of the ECOOP Workshop on Programming Languages and Operating Systems 2004 (ECOOP-PLOS 2004),* Oslo, Norway; June 2004.

TR.04-09 A. Gal, Ch.W. Probst, and M. Franz; *Complexity-Based Denial of Service Attacks on Mobile-Code Systems*; Technical Report No. 04-09, School of Information and Computer Science, University of California, Irvine; April 2004.

CPC03 N. Dalton, Ch. Fensch, E. Yardimci, and M. Franz; "A Virtual Machine for Improving Native-Code Execution on Explicitly Parallel Processors;" in *Proceedings of the 10th International Workshop on Compilers for Parallel Computers (CPC 2003),* Amsterdam, The Netherlands, pp. 261–270; January 2003.

CPC01 J. von Ronne, M. Franz, N. Dalton, and W. Amme; "Compile Time Elimination of Null- and Bounds-Checks;" in *Ninth International Workshop on Compilers for Parallel Computers (CPC 2001),* Edinburgh, Scotland, pp. 325–334; June 2001.

TR.98-34 M. Franz and T. Kistler; *Splitting Data Objects to Increase Cache Latency*; Technical Report No. 98-34, Department of Information and Computer Science, University of California, Irvine; October 1998.

TR.90-142 M. Franz; *MacOberon Reference Manual*; Technical Report No. 142, Departement Informatik, ETH Zürich; November 1990.

TR.90-141 M. Franz; *The Implementation of MacOberon*; Technical Report No. 141, Departement Informatik, ETH Zürich; October 1990.

Professional Activities

Major Honors and Awards

- *ACM Charles P. "Chuck" Thacker Breakthrough in Computing Award, Association for Computing Machinery (ACM),* "for the development of just-in-time compilation techniques that enabled fast and feature-rich web services on the internet." This award of $100,000 recognizes "individuals who have made a surprising or disruptive leapfrog in computing ideas or technologies;" 2020.
- *Fellow, American Association for the Advancement of Science (AAAS),* 'for distinguished contributions to computer science, particularly to the areas of just-in-time compilation and optimization and techniques for computer security;" 2019.
- *Inaugural Fellow, International Federation for Information Processing (IFIP).* Among the inaugural cohort of only 18 inductees when the Fellow grade was first introduced. The Fellow Award recognizes individuals of the highest professional standing and expertise in one of IFIP's constituent societies who have also contributed directly to IFIP. Established in 1960 under the auspices of UNESCO, IFIP is the global organisation for researchers and professionals working in the field of information and communication technologies. IFIP is recognised by the United Nations and links some 50 national and international societies and academies of science with a total membership of over half a million professionals. 2019.
- *Humbold Research Award*, Alexander von Humboldt Foundation. This award is granted in recognition of a researcher's entire achievements to date to academics whose fundamental discoveries, new theories, or insights have had a significant impact on their own discipline and who are expected to continue producing cutting-edge achievements in the future. Award of €60,000; 2018.
- *Innovator of the Year Award, UCI Applied Innovation & The Beall Family Foundation),* Award of $10,000; 2018.
- *Fellow, Association for Computing Machinery (ACM),* "for contributions to just-in-time compilation and optimization and to compiler techniques for computer security;" 2015.
- *Fellow, The Institute of Electrical and Electronics Engineers (IEEE),* "for contributions to just-in-time compilation and to computer security through compiler-generated software diversity;" 2015.
- *Dean's Award for Research*, Donald Bren School of Information and Computer Sciences, UC Irvine, 2015.
- *IEEE Computer Society Technical Achievement Award*, 2012, "for pioneering contributions to just-in-time compilation and optimization and significantly advancing Web application technology."
- *IEEE Orange County Chapter Outstanding Engineer Award*, 2012.
- *University of California, Irvine, Distinguished Mid-Career Faculty Award for Research*, 2010. This is the Academic Senate's highest honor for research. One such award at most is given yearly to an Assistant Professor, one to an Associate or Full Professor Step I-IV (the "Mid-Career Award"), and one to a Professor Step V or higher.
- *Distinguished Scientist, Association for Computing Machinery (ACM),* "Created early mobile code system. Leads key research group on Virtual Machines and Mobile-Code Security. Co-Founder of the ACM Sigplan VEE Conference;" 2006.
- *Senior Member, The Institute of Electrical and Electronics Engineers (IEEE),* 2006.
- *National Science Foundation CAREER Award,* 1997.
- I was awarded a *Fulbright Scholarship* (for graduate study in the United States) in 1989, but subsequently declined this award in order to join the research group of Prof. Niklaus Wirth at ETH Zürich.

Service to the Professional Community: Ongoing

- Program Committee Member, *2023 IEEE Symposium on Security and Privacy ("Oakland"),* San Francisco, California; May 2023.
- Member, *IEEE Computer Society Fellows Evaluation Committee*; 2023.
- Member, *IFIP Working Group 11.10 ("Critical Infrastructure Protection"),* 2018 – present.
- Journal Editorial Board Member, *Software—Practice and Experience (SPE);* since July 2010.
- Emeritus Member, *IFIP Working Group 2.4 ("Software Implementation Technology"),* since July 2018 (previously, I was a Full Voting Member from 2002 – 2018, and before that, an Observer from 1998 – 2002).
- Charter Faculty Member, *Security Computing and Networking Center (SCoNCe)* (previously named *Center for Cyber-Security and Privacy*), Donald Bren School of Information and Computer Sciences, UC Irvine, May 2005 – present.
- Charter Member, *The California Institute for Telecommunications and Information Technology (Cal-(IT)^2),* one of four California Institutes for Science and Technology, December 2000 – present.

Service to the Professional Community: Past

Editorial Board Member

- Journal Editorial Board Member, *Software-Intensive Cyber-Physical Systems (SICS);* October 2009–December 2020. Prior to 2017, the journal was published under the name *Computer Science—Research and Development (CSRD).*.
- Journal Editorial Board Member, *IEEE Transactions on Dependable and Secure Computing (TDSC);* March 2015–December 2019. (IEEE Computer Society policy limits the service of Associate Editors to two consecutive 2-year terms.)

Program Committee Member

- *2022 USENIX Security Conference,* Boston, Massachussetts; August 2022.
- *2022 Network and Distributed System Security Symposium (NDSS),* San Diego, California, April 2022.
- *24th International Symposium on Research in Attacks, Intrusions and Defenses (RAID 2021),* Donostia / San Sebastian, Spain; October 2021.
- *26th European Symposium on Research in Computer Security (ESORICS 2021),* Darmstadt, Germany; October 2021.
- *2021 USENIX Security Conference,* Vancouver, British Columbia; August 2021.
- *2021 IEEE Symposium on Security and Privacy ("Oakland"),* San Francisco, California; May 2021.
- *18th International Conference on Cryptology and Network Security (CANS 2020),* Vienna, Austria; December 2020.
- *27th ACM Conference on Computer and Communications Security (ACM CCS 2020),* Orlando, Florida; November 2020.
- *5th Workshop on Forming an Ecosystem Around Software Transformation (FEAST 2020),* Orlando, Florida; November 2020.
- *4th Workshop on Attacks and Solutions in Hardware Security (ASHES 2020),* Orlando, Florida; November 2020.
- *7th ACM Workshop on Moving Target Defense (MTD 2020),* Orlando, Florida; November 2020.
- *2020 IEEE Symposium on Security and Privacy ("Oakland"),* San Francisco, California; May 2020.
- *GI SICHERHEIT 2020),* Göttingen, Germany; March 2020.
- *6th ACM Workshop on Moving Target Defense (MTD 2019),* London, England; November 2019.

- *3rd International Workshop on Software Protection (SPRO-2019),* London, England; November 2019.
- *26th ACM Conference on Computer and Communications Security (ACM CCS 2019),* London, England; November 2019.
- *2019 IEEE Symposium on Security and Privacy ("Oakland"),* San Francisco, California; May 2019.
- *2018 Dynamic Languages Symposium (DLS18),* Boston, Massachusetts; November 2018.
- *25th ACM Conference on Computer and Communications Security (ACM CCS 2018),* Toronto, Ontario, Canada; October 2018.
- *5th ACM Workshop on Moving Target Defense (MTD 2018),* Toronto, Ontario, Canada; October 2018.
- *2018 Secure Development Conference (SecDev 2018),* Cambridge, Massachussetts, September/October 2018.
- *19th World Conference on Information Security Applications (WISA 2018),* Jeju Island, South Korea; August 2018.
- *First Workshop on Software Debloating and Delayering (SALAD '18),* Amsterdam, Netherlands; July 2018.
- *38th IEEE International Conference on Distributed Computing Systems (ICDCS 2018),* Vienna, Austria; June 2018.
- *GI SICHERHEIT 2018),* Constance, Germany; April 2018.
- *2017 ACM/IFIP/USENIX International Middleware Conference (Middleware 2017),* Las Vegas, Nevada; December 2017.
- *4th ACM Workshop on Moving Target Defense (MTD 2017),* Dallas, Texas; October 2017.
- *2017 Secure Development Conference (SecDev 2017),* Cambridge, Massachussetts, September 2017.
- *International Symposium on Engineering Secure Software and Systems (ESSoS'17),* Bonn, Germany; July 2017.
- *15th International Conference on Applied Cryptography and Network Security (ACNS 2017),* Kanazawa, Japan; July 2017.
- *37th IEEE International Conference on Distributed Computing Systems (ICDCS 2017),* Atlanta, Georgia; June 2017.
- *2017 ACM Asia Conference on Computer and Communications Security (ASIACCS 2017),* Abu Dhabi, UAE; April 2017.
- *8th IEEE International Workshop on Information Forensics and Security (WIFS 2016),* Abu Dhabi, UAE; December 2016.
- *3rd ACM Workshop on Moving Target Defense (MTD 2016),* Vienna, Austria; October 2016.
- *2nd International Workshop on Software Protection (SPRO-2016),* Vienna, Austria; October 2016.
- *23rd ACM Conference on Computer and Communications Security (ACM CCS 2016),* Vienna, Austria; October 2016.
- *2016 International Conference on Principles and Practices of Programming in Java (PPPJ'2016),* Lugano, Switzerland; September 2016.
- **Program co-Chair**, *9th International Conference on Trust and Trustworthy Computing (TRUST 2016),* Vienna, Austra; August 2016.
- *14th International Conference on Applied Cryptography and Network Security (ACNS 2016),* London, United Kingdom; June 2016.
- *International Symposium on Engineering Secure Software and Systems (ESSoS'16),* Egham, United Kingdom; March 2016.
- *First IEEE European Symposium on Security and Privacy 2016 (EuroS&P2016),* Saarbrücken, Germany; March 2016.
- *2nd ACM Workshop on Moving Target Defense (MTD 2015),* Denver, Colorado; October 2015.
- *22nd ACM Conference on Computer and Communications Security (CCS 2015),* Denver, Colorado; October 2015.

- *2015 International Conference on Principles and Practices of Programming in Java (PPPJ'2015),* Melbourne, Florida; September 2015.
- *8th International Conference on Trust and Trustworthy Computing (TRUST 2015),* Heraklion, Greece; August 2015.
- *1st International Workshop on Software Protection (SPRO-2015),* Florence, Italy; May 2015.
- *IEEE Workshop on Web 2.0 Security and Privacy 2015 (W2SP'2015),* San Jose, California; May 2015.
- *International Symposium on Engineering Secure Software and Systems (ESSoS'15),* Milan, Italy; March 2015.
- *First ACM Workshop on Moving Target Defense (MTD 2014),* Scottsdale, Arizona; November 2014.
- *2014 New Security Paradigms Workshop (NSPW 2014),* Victoria, British Columbia, Canada; September 2014.
- *2014 International Conference on Principles and Practices of Programming in Java (PPPJ'2014),* Krakow, Poland; September 2014.
- *IEEE Workshop on Web 2.0 Security and Privacy 2014 (W2SP'2014),* San Francisco, California; May 2014.
- *ACM International Conference on Computing Frontiers 2014 (CF 14),* Cagliari, Italy; May 2014.
- *The Next Generation Malware Attacks and Defense Workshop (NGMAD),* New Orleans, Louisiana; December 2013.
- *2013 International Conference on Principles and Practices of Programming in Java (PPPJ'2013),* Stuttgart, Germany; September 2013.
- *2013 New Security Paradigms Workshop (NSPW 2013),* Banff, Alberta, Canada; September 2013.
- *6th International Conference on Trust and Trustworthy Computing (TRUST 2013),* London, United Kingdom; June 2013.
- *28th Annual Computer Security Applications Conference (ACSAC 2012),* Orlando, Florida; December 2012.
- *2012 New Security Paradigms Workshop (NSPW 2012),* Bertinoro, Italy; September 2012.
- *11th International Conference on Generative Programming and Component Engineering (GPCE 2012),* Dresden, Germany; September 2012.
- *2012 IEEE International Conference on Privacy, Security, Risk and Trust (PASSAT 2012),* Amsterdam, The Netherlands; September 2012.
- *5th International Conference on Trust and Trustworthy Computing (TRUST 2012),* Vienna, Austria; June 2012.
- *Eighth ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments (VEE 2012),* London, United Kingdom; March 2012.
- *27th Annual Computer Security Applications Conference (ACSAC 2011),* Orlando, Florida; December 2011.
- **Program co-Chair**, *International Workshop on Programming Language And Systems Technologies for Internet Clients (PLASTIC 2011),* Portland, Oregon; October 2011.
- *6th Workshop on Programming Languages and Operating Systems (PLOS 2011),* Cascais, Portugal; October 2011.
- *Third IEEE International Conference on Privacy, Security, Risk and Trust (PASSAT2011),* Boston, Massachusetts; October 2011.
- *2011 New Security Paradigms Workshop (NSPW 2011),* Sonoma, California; September 2011.
- *4th International Conference on Trust and Trustworthy Computing (TRUST 2011),* Pittsburgh, Pennsylvania; June 2011.
- *ACM Sigplan Conference on Programming Language Design and Implementation (PLDI 2011),* San Diego, California; June 2011.
- *5th International Multidisciplinary Conference on e-Technologies (MCETECH 2011),* Les Diablerets, Switzerland; January 2011.

- **Program Chair**, *26th Annual Computer Security Applications Conference (ACSAC 2010),* Austin, Texas; December 2010. 237 submitted papers, 39 accepted.
- *19th ACM/IEEE/IFIP International Conference on Parallel Architectures and Compilation Techniques (PACT 2010),* Vienna, Austria; September 2010.
- *2010 IEEE International Conference on Privacy, Security, Risk and Trust (PASSAT-10),* Minneapolis, Minnesota; August 2010.
- *ACM SIGPLAN/SIGBED Conference on Languages, Compilers and Tools for Embedded Systems (LCTES 2010),* Stockholm, Sweden; April 2010.
- *ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments (VEE 2010),* Pittsburgh, Pennsylvania; March 2010.
- *International Symposium on Engineering Secure Software and Systems (ESSoS 2010),* Pisa, Italy; February 2010.
- **Program co-Chair**, *25th Annual Computer Security Applications Conference (ACSAC 2009),* Honolulu, Hawaii; December 2009.
- *12th Information Security Conference (ISC 2009),* Pisa, Italy; September 2009.
- *2009 New Security Paradigms Workshop (NSPW 2009),* Oxford, United Kingdom; September 2009.
- *2009 IEEE International Conference on Privacy, Security, Risk and Trust (PASSAT-09),* Vancouver, British Columbia, Canada; August 2009.
- *2009 International Conference on Principles and Practices of Programming in Java (PPPJ'2009),* Calgary, Alberta, Canada; August 2009.
- *47th International Conference on Objects, Models, Components, and Patterns (TOOLS-EUROPE 2009),* Zurich, Switzerland, June/July 2009.
- *4th Montreal Conference on eTechnologies (MCETECH),* Ottawa, Canada; May 2009.
- *Compiler Construction 2009 (CC 2009),* York, United Kingdom; March 2009.
- *2008 IEEE Symposium on Security and Privacy,* Oakland, California; May 2008.
- *2008 Annual IEEE Computer Society/ACM International Symposium on Code Generation and Optimization (CGO 2008),* Boston, Massachusetts; March 2008.
- *23rd Annual Computer Security Applications Conference (ACSAC 2007),* Miami Beach, Florida; December 2007.
- *2007 International Conference on Principles and Practices of Programming in Java (PPPJ'2007),* Monte de Caparica/Lisbon, Portugal; September 2007.
- *New Security Paradigms Workshop (NSPW 2007),* Washington Valley, New Hampshire; September 2007.
- *Workshop on Linguistic Support for Modern Operating Systems (PLOS 2006),* October 2006.
- *The Second Workshop on Advances in Trusted Computing (WATC'06 Fall),* Tokyo, Japan, November-December 2006.
- *Seventh Joint Modular Languages Conference (JMLC 2006),* Oxford, United Kingdom, September 2006.
- *2006 International Conference on Principles and Practices of Programming in Java (PPPJ'2006),* Mannheim, Germany, September 2006.
- *New Security Paradigms Workshop (NSPW 2006),* Dagstuhl, Germany, September 2006.
- *New Security Paradigms Workshop (NSPW 2005),* Lake Arrowhead, California, September 2005.
- *ECOOP Workshop on Programming Languages and Operating Systems (ECOOP-PLOS 2005),* June 2005.
- *Third International Workshop on Compiler Optimization Meets Compiler Verification (COCV 2005),* Edinburgh, Scotland, April 2005.
- *3. Arbeitstagung Programmiersprachen (ATPS 2004) of the German Computer Society (GI),* Ulm, Germany, September 2004.

- *ECOOP Workshop on Programming Languages and Operating Systems (ECOOP-PLOS 2004),* Oslo, Norway, June 2004.
- *ACM SIGPLAN 2004 Conference on Languages, Compilers, and Tools for Embedded Systems (LCTES'04),* Washington, D.C., June 2004.
- *Second Annual IEEE/ACM International Symposium on Code Generation and Optimization (CGO 2004),* San Jose, California, March 2004.
- *Third International Workshop on Compiler Optimization Meets Compiler Verification (COCV 2004),* Barcelona, Spain, March/April 2004.
- *ACM SIGSAC New Security Paradigms Workshop 2003 (NSPW-2003),* Ascona, Switzerland, September 2003.
- *Sixth Joint Modular Languages Conference (JMLC 2003),* Klagenfurt, Austria, August 2003.
- *ACM SIGPLAN 2003 Workshop on Interpreters, Virtual Machines and Emulators (IVME'03),* San Diego, California, June 2003.
- *Second International Workshop on Compiler Optimization Meets Compiler Verification (COCV 2003),* Warsaw, Poland, April 2003.
- *4th Annual Workshop on Binary Translation (WBT-2002),* Charlottesville, Virginia, September 2002.
- *Fifth ECOOP Workshop on Object-Orientation and Operating Systems (ECOOP-OOOSWS 2002),* Málaga, Spain, June 2002.
- *Fifth IEEE International Symposium on Object-Oriented Real-Time Distributed Computing (ISORC 2002),* Washington, D.C., April–May 2002.
- *11th International Conference on Compiler Construction (CC'2002),* Grenoble, France, March 2002.
- *Fourth IEEE International Symposium on Object-Oriented Real-Time Distributed Computing (ISORC 2001),* Magdeburg, Germany, May 2001.
- *Fifth Joint Modular Languages Conference (JMLC 2000),* Zurich, Switzerland, September 2000.
- *Third Workshop on Distributed Communities on the Web (DCW 2000),* Quebec City, Canada, June 2000.
- *Third IEEE International Symposium on Object-Oriented Real-Time Distributed Computing (ISORC 2000),* Newport Beach, California, March 2000.
- *European Symposium on Programming (ESOP 2000),* Berlin, Germany, March/April 2000.
- *Workshop on Binary Translation* (in conjunction with the International Conference on Parallel Architectures and Compilation Techniques, PACT '99), Newport Beach, California, October 1999.
- *ACM Sigplan 1999 Workshop on Compiler Support for System Software (WCSSS'99),* Atlanta, Georgia, May 1999.
- *Fourth California Software Symposium (CSS'98),* Irvine, California, October 1998.
- *Workshop on Principles of Abstract Machines* (in conjunction with the joint international symposia SAS'98 and PLILP/ALP'98), Pisa, Italy, September 1998.
- *ACM Sigplan Conference on Programming Language Design and Implementation (PLDI'98),* Montreal, Canada, June 1998.
- *Fourth Joint Modular Languages Conference (JMLC'97),* Linz, Austria, March 1997.

Session Chair

- *2022 USENIX Security Conference,* Boston, Massachussetts; August 2022.
- *2021 USENIX Security Conference,* Vancouver, British Columbia; August 2021.
- *2021 IEEE Symposium on Security and Privacy ("Oakland"),* San Francisco, California; May 2021.
- *27th ACM Conference on Computer and Communications Security (ACM CCS 2020),* Orlando, Florida; November 2020.

- *2020 IEEE Symposium on Security and Privacy ("Oakland"),* San Francisco, California; May 2020.
- *26th ACM Conference on Computer and Communications Security (CCS 2019),* London, United Kingdom; November 2019.
- *2019 IEEE Symposium on Security and Privacy ("Oakland"),* San Francisco, California; May 2019.
- *25th ACM Conference on Computer and Communications Security (CCS 2018),* Toronto, Canada; October 2018.
- *38th IEEE International Conference on Distributed Computing Systems (ICDCS 2018),* Vienna, Austria; June 2018.
- *2018 ACM Asia Conference on Computer and Communications Security (ASIACCS 2018),* Incheon, South Korea; June 2018.
- *Third IEEE European Symposium on Security and Privacy 2018 (EuroS&P2018),* London, United Kingdom; April 2018.
- *Usenix Security 2017,* Vancouver, British Columbia; August 2017.
- *23rd ACM Conference on Computer and Communications Security (CCS 2016),* Vienna, Austria; October 2016.
- *First IEEE European Symposium on Security and Privacy 2016 (EuroS&P2016),* Saarbrücken, Germany; March 2016.
- *22nd ACM Conference on Computer and Communications Security (CCS 2015),* Denver, Colorado; October 2015.
- *10th Conference on High Performance and Embedded Architecture and Compilation (HiPEAC 2015),* Amsterdam, Netherlands; January 2015.
- *6th International Conference on Trust and Trustworthy Computing (TRUST 2013),* London, United Kingdom; June 2013.
- *22nd International Conference on Compiler Construction (CC 2013),* Rome, Italy; March 2013.
- *28th Annual Computer Security Applications Conference (ACSAC 2012),* Orlando, Florida; December 2012.
- *11th International Conference on Generative Programming and Component Engineering (GPCE 2012),* Dresden, Germany; September 2012.
- *27th Annual Computer Security Applications Conference (ACSAC 2011),* Orlando, Florida; December 2011.
- *ACM Sigplan Conference on Programming Language Design and Implementation (PLDI 2011),* San Diego, California; June 2011.
- *26th Annual Computer Security Applications Conference (ACSAC 2010),* Austin, Texas; December 2010.
- *19th ACM/IEEE/IFIP International Conference on Parallel Architectures and Compilation Techniques (PACT 2010),* Vienna, Austria; September 2010.
- *ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments (VEE 2010),* Pittsburgh, Pennsylvania; March 2010.
- *23rd Annual Computer Security Applications Conference (ACSAC 2007),* Miami Beach, Florida; December 2007.
- *Seventh Joint Modular Languages Conference (JMLC 2006),* Oxford, United Kingdom, September 2006.
- *Invitational Workshop on the Future of Virtual Execution Environments,* Armonk, New York; September 2004
- *New Security Paradigms Workshop (NSPW 2004),* White Point, Nova Scotia, September 2004.
- *The Fourth IEEE International Conference on Peer-to-Peer Computing (P2P 2004),* Zurich, Switzerland, August 2004.
- *Southern California Parallel Processing and Computer Architecture Workshop,* Los Angeles, California, May 2004.
- *Sixth Joint Modular Languages Conference (JMLC 2003),* Klagenfurt, Austria, August 2003.
- *Ninth International Workshop on Compilers for Parallel Computers (CPC 2001),* Edinburgh, Scotland, June 2001.
- *Fourth IEEE International Symposium on Object-Oriented Real-Time Distributed Computing (ISORC 2001),* Magdeburg, Germany, May 2001.

- *Workshop on Binary Translation* (in conjunction with the International Conference on Parallel Architectures and Compilation Techniques, PACT '99), Newport Beach, California, October 1999.
- *ACM Sigplan Conference on Programming Language Design and Implementation (PLDI'98),* Montreal, Canada, June 1998.
- *Fourth Joint Modular Languages Conference (JMLC'97),* Linz, Austria, March 1997.

Other Service

- Member, *IEEE Computer Society Fellows Evaluation Committee*; 2022.
- Member, *Fraunhofer CyberStar Award for Graduate Students in Israel Selection Committee*; 2020–2021.
- Member, *IEEE Computer Society Fellows Evaluation Committee*; 2020.
- Member, *IEEE Computer Society Publication Board Best Paper Award (BPA) Committee for IEEE Transactions on Dependable and Secure Computing (TDSC);* 2019.
- Member, *Fraunhofer CyberStar Award for Graduate Students in Israel Selection Committee*; 2019–2020.
- Member, *IEEE Computer Society Fellows Evaluation Committee*; 2019.
- Full Voting Member, *IFIP Working Group 2.4 ("Software Implementation Technology"),* 2002 – 2018 (elevated to Emeritus Member in July 2018).
- Member, *IFIP Working Group 11.3 ("Data and Application Security and Privacy"),* 2008 – 2017.
- Shadow PC Member, *2017 ACM Asia Conference on Computer and Communications Security (ASIACCS 2017),* Abu Dhabi, UAE; April 2017.
- Member, *IEEE Computer Society Fellows Evaluation Committee*; 2016.
- Nomination Committee Member, MacArthur Fellows Program, *John D. and Catherine T. MacArthur Foundation*; 2016.
- External Review Committee Member, *28th European Conference on Object-Oriented Programming (ECOOP'2014),* Uppsala, Sweden; July/August 2014.
- External Review Committee Member, *ACM Research Conference on Object-Oriented Programming (OOPSLA 2013),* Indianapolis, Indiana; October 2013.
- External Review Committee (ERC) Member, *ACM Sigplan Conference on Programming Language Design and Implementation (PLDI 2013),* Seattle, Washington; June 2013.
- Paper Shepherd, *2011 New Security Paradigms Workshop (NSPW 2011),* Sonoma, California; September 2011.
- Organization Committee Member and Sponsorship co-chair, *EuroSys 2011*, Salzburg, Austria; March 2011.
- Organizing Committee Member (Student Travel Chair), *Fourteenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '09),* Washington, D.C.; March 2009.
- Panels Chair and Conference Committee Member, *24th Annual Computer Security Applications Conference (ACSAC 2008),* Anaheim, California; December 2008.
- Steering Committee Member, *ACM SIGPLAN/SIGOPS/USENIX International Conference Series on Virtual Execution Environments (VEE),* 2004 – 2008.
- Local Arrangements Chair, *IFIP WG2.4 Working Meeting,* Arrowhead, California, May 2007.
- Local Arrangements Co-Chair, *New Security Paradigms Workshop (NSPW 2006),* Dagstuhl, Germany, September 2006.
- **Founding Steering Committee Co-Chair** (with Sam Midkiff of Purdue University), *ACM SIGPLAN/SIGOPS/USENIX International Conference Series on Virtual Execution Environments (VEE),* September 2004 – June 2005.
- **General Chair**, *ACM SIGPLAN 2004 Workshop on Interpreters, Virtual Machines and Emulators (IVME 2004),* Washington, D.C., June 2004.

- Observer, *IFIP Working Group 2.4,* February 1998 – November 2002 (elected to full membership on November 14th).
- Tutorials Chair, *ACM Sigplan Conference on Programming Language Design and Implementation (PLDI 2000),* Vancouver, Canada, June 2000.
- Local Arrangements Co-Chair, *Third IEEE International Symposium on Object-Oriented Real-Time Distributed Computing (ISORC 2000),* Newport Beach, California, March 2000.
- Executive Committee Member, *The Institute for Software Research at UC Irvine,* July 1999–January 2002.
- Charter Member, *The Institute for Software Research at UC Irvine,* July 1999.
- Session Organizer and Host, *Bay Area Round Table (BART),* Palo Alto, California, February 1999.
- Executive Committee Member, *Irvine Research Unit in Software (IRUS)* [precursor to The Institute for Software Research], January 1996 – June 1999.
- Swiss Delegate to *IFIP Technical Committee No. 2,* "Software: Theory and Practice," 1995–1996 term.
- Program Committee Chair, *Oberon Track at the First Joint Annual Conference of the Gesellschaft für Informatik and the Schweizer Informatiker Gesellschaft,* Zürich, September 1995.
- Executive Committee Member, *Special Interest Group on Oberon of the Schweizer Informatiker Gesellschaft,* 1994–1996.
- Organizing Committee Member, *Conference on Programming Languages and System Architectures,* Zürich, March 1994.

Grant Application Review Panel Member

- *National Science Foundation, Program on Software and Trusted Computing (SaTC),* Arlington, Virginia, January 2016.
- *National Science Foundation, Program on Software and Trusted Computing (SaTC),* Arlington, Virginia, October 2012.
- *National Science Foundation, Program on Software and Trusted Computing (SaTC),* Arlington, Virginia, May 2012.
- *National Science Foundation, Program on Computer and Network Systems,* Arlington, Virginia, April 2009.
- *National Science Foundation, Program on Foundations of Computing Processes and Artifacts,* Arlington, Virginia, February 2007.
- *National Science Foundation, CAREER Program in CyberTrust,* Arlington, Virginia, November 2005.
- *National Science Foundation, CAREER Program in Networking and Security,* Arlington, Virginia, November 2003.
- *National Science Foundation, Program in Embedded & Hybrid Systems,* Arlington, Virginia, June 2002.

Invited Keynotes, Presentations and Panels at Conferences

- M. Franz; "Fast and Furious: How the Web got Turbo Charged Just In Time…," ACM Breakthrough Lecture, *ASPLOS 2022,* Lausanne, Switzerland, March 2022.
- M. Franz; "Cyber Attacks And Defenses: Trends, Challenges, and Outlook," *CyberSecurity@KAIST Workshop,* Daejeon, South Korea, June 2018.
- M. Franz; "From Fine Grained Code Diversity to Execute-No-Read: The Cat and Mouse Game Between Attackers and Defenders Continues," *2nd ACM Workshop on Moving Target Defense (MTD 2015),* Denver, Colorado; October 2014.

- M. Franz; "Biologically Inspired Software Defenses," *Fifteenth High Confidence Software and Systems Conference (HCSS 2015),* Annapolis, Maryland; May 2015.
- M. Franz; "Code Diversity and Biologically Inspired Computer Defenses" (Invited Keynote), *TTI/Vanguard Reprogramming Programming),* Arlington, Virginia; September/October 2014.
- M. Franz; "Software Diversity as a Cyber Defense" (Invited Keynote), *The Next Generation Malware Attacks and Defense Workshop (NGMAD),* New Orleans, Louisiana; December 2013.
- M. Franz; "Eliminating the Insider Threat in Software Development by Combining Parallelism, Randomization and Checkpointing" (Invited Keynote Address); *Fourth Annual Cyber Security and Information Intelligence Research Workshop (CSIIRW'08),* Oak Ridge National Laboratory, Oak Ridge, Tennessee; May 2008.
- M. Franz; "Security and Privacy in Service Oriented Architectures" (Panelist); *21st Annual IFIP WG 11.3 Working Conference on Data and Applications Security (DBSEC'07),* Redondo Beach, California; July 2007.
- M. Franz; "Erinnerungen und Ausblicke: Was haben wir gelernt? Und was soll die nächste Generation lernen?" (Invited Panelist); *Tag der Informatik,* ETH Zurich, Switzerland, October 2006.
- M. Franz; "A New Approach to Embedded Java" (Invited Keynote Address); *Mobile Information & Communication Systems, Scientific Conference,* Zurich, Switzerland, October 2006.
- M. Franz; "Pervasive Security" (Panelist); *Software Security Panel, National Science Foundation, Trusted Computing Program, PI Meeting,* Pittsburgh, Pennsylvania; August 2004.
- M. Franz; "Safe Code: It's Not Just For Applets Anymore" (Invited Keynote Address); *Sixth Joint Modular Languages Conference (JMLC 2003),* Klagenfurt, Austria, August 2003.
- M. Franz; "Pervasive Security" (Panelist); *Trusted Computing Panel, National Science Foundation, Trusted Computing Program, PI Meeting,* Baltimore, Maryland; August 2003.
- M. Franz; "The Source is The Proof" (Panelist); *NSPW Panel, 18th Annual Computer Security Applications Conference (ACSAC-18),* Las Vegas, Nevada; December 2002.
- M. Franz; "Extensible Programming: Ein neues Paradigma für die Softwareentwicklung" (Invited Keynote Address, in German); *Moderne Programmierparadigmen,* conference sponsored by Gesellschaft für Informatik, FH Braunschweig-Wolfenbüttel, Germany; October 1994.

Meeting Participation By Invitation († = I gave a presentation, ‡ = my student gave a presentation, ⋆ = I presented a poster)

Note: Presentations at conferences with proceedings are documented under "Publications" above and are not listed again here.

- ‡*Black Hat Briefings 2020,* Las Vegas, Nevada, August 2020.
- †*Third AmSEC Workshop on Systems Security,* Amsterdam, Netherlands, October 2019.
- ‡*Black Hat Briefings 2019,* Las Vegas, Nevada, August 2019.
- ‡*Qualcomm Product Security Summit,* San Diego, California, May 2019.
- *DARPA Cyber Fault-tolerant Attack Recovery (CFAR) PI Meeting,* Chantilly, Virginia, January 2018.
- †*IFIP WG2.4 Working Meeting,* Essex, Vermont, October 2017.
- *DARPA Cyber Fault-tolerant Attack Recovery (CFAR) PI Meeting,* Chantilly, Virginia, May/June 2017.
- †*IFIP WG2.4 Working Meeting,* Dresden, Germany, December 2016.
- *DARPA Cyber Fault-tolerant Attack Recovery (CFAR) PI Meeting,* Chantilly, Virginia, November 2016.

- *DARPA Cyber Fault-tolerant Attack Recovery (CFAR) PI Meeting,* Arlington, Virginia, April 2016.
- *DARPA Cyber Fault-tolerant Attack Recovery (CFAR) PI Meeting,* Arlington, Virginia, January 2016.
- †*M.I.T. Invitational Think-Shop on Multi-Spectrum Metrics for Cyber Defense,* sponsored by the National Science Foundation, Arlington, Virginia; December 2015.
- *DARPA Cyber Fault-tolerant Attack Recovery (CFAR) PI Meeting,* Arlington, Virginia, November 2015.
- †*DARPA Cyber Fault-tolerant Attack Recovery (CFAR) PI Meeting,* Arlington, Virginia, August 2015.
- †*DARPA Cyber Fault-tolerant Attack Recovery (CFAR) Kick-Off PI Meeting,* Arlington, Virginia, May 2015.
- †*DARPA Joint Clean-Slate Design of Resilient, Secure Hosts (CRASH) & Mission-Oriented Resilient Clouds (MRC) PI Meeting,* Jacksonville, Florida, September 2014.
- †*IFIP WG2.4 Working Meeting,* Asilomar, Pacific Grove, California, February 2014.
- †*DARPA Joint Clean-Slate Design of Resilient, Secure Hosts (CRASH) & Mission-Oriented Resilient Clouds (MRC) PI Meeting,* San Diego, California, January 2014.
- †*M.I.T. Invitational Think-Shop on Multi-Spectrum Metrics for Cyber Defense,* sponsored by the National Science Foundation, Cambridge, Massachusetts; October 2013.
- *Facebook Faculty Summit,* Menlo Park, California, August 2013.
- †*DARPA Joint Clean-Slate Design of Resilient, Secure Hosts (CRASH) & Mission-Oriented Resilient Clouds (MRC) PI Meeting,* Park Ridge, New Jersey, May 2013.
- ⋆*National Security Agency, First Annual Science of Security (SoS) Community Meeting,* National Harbor, Maryland, November 2012.
- †*DARPA Joint Clean-Slate Design of Resilient, Secure Hosts (CRASH) & Mission-Oriented Resilient Clouds (MRC) PI Meeting,* San Diego, California, November 2012.
- *DARPA Mission-oriented Resilient Clouds (MRC) Program, PI Meeting,* San Diego, California, October 2012.
- †*DARPA Clean-Slate Design of Resilient, Secure Hosts (CRASH) Program, PI Meeting,* Boston, Massachusetts, May 2012.
- *DARPA Colloquium on Future Directions in Cyber Security,* Arlington, Virginia, November 2011.
- †*DARPA Clean-Slate Design of Resilient, Secure Hosts (CRASH) Program, PI Meeting,* Arlington, Virginia, November 2011.
- †*2nd Army Research Office (ARO) Workshop on Moving Target Defense,* Fairfax, Virginia, October 2011.
- †*DARPA Clean-Slate Design of Resilient, Secure Hosts (CRASH) Program, PI Meeting,* San Jose, California, May 2011.
- *Microsoft Research Faculty Summit,* Redmond, Washington, July 2010.
- †*IFIP WG2.4 Working Meeting,* Berg en Terblijt, Netherlands, January 2010.
- *Networking and Information Technology Research and Development (NITRD) Program, National Cyber Leap Year Summit,* Arlington, Virginia, August 2009.
- †*National Intelligence Community, Enterprise Cyber Assurance Program (NICECAP), PI Meeting,* Washington, D.C., September 2008.
- *Google Faculty Summit,* Mountain View, California, July 2008.

- †*National Intelligence Community, Enterprise Cyber Assurance Program (NICECAP), Reverse Site Visit,* Jessup, Maryland, January 2008.
- †NCDI Workshop on Game-changing Solutions for Cyber Security (jointly sponsored by NSF, DHS, IARPA, NSA, ONR, and OSD), College Park, Maryland, November 2007.
- †*National Intelligence Community, Enterprise Cyber Assurance Program (NICECAP), PI Meeting,* Boston, Massachusetts, September 2007.
- *U.S. Department of Energy Workshop on Cyber Security Research Needs for Open Science,* Bethesda, Maryland, July 2007.
- †*IFIP WG2.4 Working Meeting,* Arrowhead, California, May 2007.
- †*National Intelligence Community, Enterprise Cyber Assurance Program (NICECAP), Program Kick-Off Meeting,* Chantilly, Virginia, March 2007.
- †*U.S. Department of Homeland Security, S&T CyberSecurity R&D PI Meeting,* Menlo Park, California, February 2007.
- †*National Science Foundation Safe Computing Workshop,* Albuquerque, New Mexico, November/December 2006.
- †*U.S. Department of Homeland Security, S&T CyberSecurity R&D PI Meeting,* Arlington, Virginia, August 2006.
- †*IFIP WG2.4 Working Meeting,* Glasgow, Scotland, July 2006.
- *The First Workshop on Advances in Trusted Computing,* Tokyo, Japan, March 2006.
- †*U.S. Department of Homeland Security, S&T CyberSecurity R&D PI Meeting,* Menlo Park, California, January 2006.
- *National Science Foundation, Trusted Computing Program, PI Meeting,* Newport Beach, California, September 2005.
- †*U.S. Department of Homeland Security, BAA 04-17, Program Kick-Off Meeting,* Arlington, Virginia, July 2005.
- *Microsoft Academic Days in Silicon Valley,* Mountain View, California, October 2004.
- *Microsoft Research 2004 Faculty Summit,* Redmond, Washington, August 2004.
- †*Southern California Parallel Processing and Computer Architecture Workshop,* Los Angeles, California, May 2004.
- †*ONR Critical Infrastructure Protection, Mobile Code Program, Final Review,* Annapolis, Maryland, May 2004.
- †*IFIP WG2.4 Working Meeting,* Brisbane, Australia, March 2004.
- †*IFIP WG2.4 Working Meeting,* Santa Cruz, California, August 2003.
- †*ONR Critical Infrastructure Protection, Mobile Code Program, PI Meeting,* Ithaca, New York, July 2003.
- †*ONR Critical Infrastructure Protection, Mobile Code Program, Review Meeting,* Arlington, Virginia, June 2003.
- *DARPA Organically Assured and Survivable Information Systems (OASIS) Program, PI Meeting,* Fort Lauderdale, Florida, January 2003.
- †*ONR Critical Infrastructure Protection, Mobile Code Program, PI Meeting,* Irvine, California, January 2003.
- †*IFIP WG2.4 Working Meeting,* Dagstuhl, Germany, November 2002.
- †*DARPA Organically Assured and Survivable Information Systems (OASIS) Program, PI Meeting,* Santa Rosa, California, August 2002.

- †*ONR Critical Infrastructure Protection, Mobile Code Program, PI Meeting,* State College, Pennsylvania, July 2002.
- †*IFIP WG2.4 Working Meeting,* Simon's Town, South Africa, March 2002.
- †*Southern California Parallel Processing and Computer Architecture Workshop,* Irvine, California, February 2002.
- †*ONR Critical Infrastructure Protection, Mobile Code Program, PI Meeting,* Melbourne, Florida; January 2002.
- †*DARPA Organically Assured and Survivable Information Systems (OASIS) Program, PI Meeting,* Santa Fe, New Mexico; July 2001.
- †*ONR Critical Infrastructure Protection, Mobile Code Program, PI Meeting,* Arlington, Virginia; July 2001.
- †*Symposium on Research in Mobile Computing Systems,* Zurich, Switzerland; May 2001.
- †*DARPA Organically Assured and Survivable Information Systems (OASIS) Program, PI Meeting,* Norfolk, Virginia; February 2001.
- *University of Washington and Microsoft Research Summer Institute 2000, "Accelerating the Pace of Software Tools Research: Sharing Infrastructure",* hosted by C. Chambers, D. Notkin, A. Srivastava, and B. Zorn; Seattle, Washington; August 2000.
- †*DARPA Intrusion Tolerant Systems (ITS) Program, PI Meeting,* Honolulu, Hawaii; July 2000.
- ‡*17th Gesellschaft für Informatik (GI) Workshop on Programming Languages and Computing Concepts (with Special Emphasis on Software Components),* Bad Honnef, Germany; May 2000.
- †*DARPA Intrusion Tolerant Systems (ITS) Program, PI Meeting,* Aspen, Colorado; February 2000.
- †*DARPA Intrusion Tolerant Systems (ITS) Program, PI Meeting,* Phoenix, Arizona; August 1999.
- *National Science Foundation CAREER Program, PI Meeting,* Washington, D.C.; January 1999.
- †*Southern California Parallel Processing and Computer Architecture Workshop,* Irvine, California; March 1998.
- †*International Workshop on Component-Oriented Programming,* Linz, Austria; July 1996.
- *Third International Workshop on Workstation Operating Systems,* Key Biscayne, Florida; April 1992.

Administrative Service

- University of California, Inaugural Member, *UCI Cybersecurity Oversight Committee,* 2022-2023.
- Department of Computer Science, *Faculty Search Committee, Positions in Systems,* 2017–18, 2018–19, 2019–20, 2020–21, 2022–23.
- Donald Bren School of Information and Computer Science, Chair, *Computing and Network Policy Committee,* 2016–17, 2017–18, 2018–19, 2019–20, 2020–21. 2021-22.
- University of California, Vice Chair, *Irvine Campus Council on Planning and Budget,* 2015–2016, 2016–2017.
- University of California, Member, *Irvine Campus Council on Planning and Budget,* 2014–2015.
- Donald Bren School of Information and Computer Science, *Software Engineering Steering Committee,* 2013–2014, 2014–2015, 2015–2016, 2016–2017, 2017–2018.
- Department of Computer Science, *CS Graduate Admissions Committee,* 2016–2017.
- Department of Computer Science, *CS Admission and Graduate Student Planning Committee,* 2013–2014, 2014–2015, 2015–2016.

- University of California, Irvine, *5-year Organized Research Unit Review Committee for the Center for Embedded Computer Systems (CECS),* 2012.
- Donald Bren School of Information and Computer Science, *Executive Committee,* 2007–2008.
- Donald Bren School of Information and Computer Science, Chair, *Computing and Network Policy Committee,* 2005–2006, 2006–2007, 2008–2009.
- Donald Bren School of Information and Computer Science, *Marketing and Outreach Committee,* 2004–2005.
- University of California, *Irvine Campus Council on Undergraduate Admissions and Relations with Schools and Colleges,* 2000–2004.
- Donald Bren School of Information and Computer Science, Chair, *Faculty Search Committee, Position in Security and Cryptography,* 2002–2003.
- Donald Bren School of Information and Computer Science, *Committee on Graduate Policy,* 2002–2003.
- ICS Department, Chair, *Committee on Space Policy,* 2001–2002.
- ICS Department, *Faculty Search Committee, Position in Cryptography and Security,* 2000–2001.
- ICS Department, *Committee on Graduate Policy,* 2000–2001.
- ICS Department, *Ad-Hoc Faculty Search Committee, "Systems" Position,* 1999–2000.
- ICS Department, *Committee on Educational Policy,* 1999–2000.
- ICS Department, *Executive Committee,* 1998–1999.
- ICS Department, *Committee on Undergraduate Policy,* 1998–1999.
- ICS Department, *Faculty Search Committee, Multiple Positions in Interdisciplinary Applications of Computer Science,* 1998–1999. (Committee reviewed 170 applications = 4 linear feet of files and filled three open faculty positions.)
- University of California, *Irvine Campus Committee on Undergraduate Admissions and Relations with Schools and Colleges,* 1997–2000.
- ICS Department, *Committee on Graduate Policy,* 1997–1998.
- ICS Department, *Committee on Graduate Admissions,* 1997–1998.
- ICS Department, *Faculty Search Committee, Position in "Informatics,"* 1997–1998.
- University of California, *Irvine Campus Representative Assembly,* 1996–1997.
- ICS Department, *Committee on Personnel,* 1996–1997.
- ICS Department, *Faculty Search Committee, Position in Software Engineering,* 1995–1996.

Teaching Activities

Teaching Awards

- *Dean's Award for Graduate Student Mentoring*, Donald Bren School of Information and Computer Sciences, UC Irvine, 2022.
- *Dean's Award for Graduate Student Mentoring*, Donald Bren School of Information and Computer Sciences, UC Irvine, 2016, "For his outstanding mentoring of doctoral students over the last decade."
- *Dean's Award for Graduate Student Mentoring*, Donald Bren School of Information and Computer Sciences, UC Irvine, 2007.
- *Outstanding Professor of the Year Award*, Graduating Class of 2007, UC Irvine.

Post-Doctoral Habilitation Theses Supervised

- Dr. Christian Herrman, Universität Ulm, Germany; thesis: "Verbesserte prozedurale Programmiersprachen" (Improved Procedural Programming Languages); March 2007.

Post-Doctoral Fellows Supervised

1. Dr. Wolfram Amme
   (January–December 2000; first subsequent position: Privatdozent at the *University of Jena,* Germany).
2. Dr. Won-Kee Hong
   (October 2001–October 2002; first subsequent position: Assistant Professor at *Daegu University,* South Korea).
3. Dr. Fermin Reig
   (October 2001–July 2003; first subsequent position: Postdoc at *University of Nottingham,* United Kingdom).
4. Dr. Roxana Diaconescu
   (January 2003–September 2004; first subsequent position: PostDoc at *California Institute of Technology (Caltech),* Pasadena, California).
5. Dr. Christian Probst
   (January 2003–May 2005; first subsequent position: Assistant Professor at the *Technical University of Denmark (DTU)*, Lyngby, Denmark).
6. Dr. Andreas Gal
   (January 2007–February 2010, first subsequent position: Researcher at *Mozilla,* Mountain View, California).
7. Dr. Christian Stork
   (March 2007–September 2008).
8. Dr. Christian Wimmer
   (July 2008–April 2011, first subsequent position: Principal Member of Technical Staff, *Oracle Sun Labs,* Redwood Shores, California).
9. Dr. Stefan Brunthaler
   (April 2011–June 2015, first subsequent position: Key Researcher at *SBA Research,* Vienna, Austria).
10. Dr. Per Larsen
    (September 2011–June 2015, first subsequent position: Chief Executive Officer of *Immunant,* Irvine, California).
11. Dr. Stijn Volckaert
    (December 2015–July 2018, first subsequent position: Assistant Professor at *KU Leuven,* Belgium).

12. Dr. Yeoul Na
(July 2016–March 2020, first subsequent position: Software Engineer at *Apple,* Cupertino, California).

13. Dr. David Gens
(March 2019–June 2022; first subsequent position: Software Engineer at *Cerebras Systems,* Sunnyvale, California).

14. Dr. Adrian Dabrowski
(May 2019–April 2022; first subsequent position: Researcher at *CISPA Helmholtz Center for Information Security,* Saabrücken, Gemany).

15. Dr. Dokyung Song
(January–March 2021, first subsequent position: Assistant Professor at *Yonsei University,* Seoul, South Korea).

16. Dr. Felicitas Hetzelt (since August 2022).

Graduated Ph.D. Students (Principal Advisor and Dissertation Committee Chair)

1. Thomas Kistler
(affiliated in April 1995, candidacy: February 1998, final defense: November 1999; thesis: "Continuous Program Optimization;" first employment after graduation: *Transmeta, Inc.,* Santa Clara, California.

2. Peter H. Fröhlich
(affiliated in September 1998; advanced to candidacy in May 2001; final defense in March 2003; thesis: "The Structure of Component-Oriented Programming Languages;" first employment after graduation: *University of California, Riverside,* California).

3. Jeffery von Ronne
(affiliated in September 1999; advanced to candidacy in February 2003; final defense in July 2005; thesis: "A Safe and Efficient Machine-Independent Code Transportation Format Based on Static Single Assignment Form and Applied to Just-In-Time Compilation;" first employment after graduation: *University of Texas at San Antonio*).

4. Vivek Haldar
(affiliated in August 2000; advanced to candidacy: November 2002; final defense: February 2006; thesis: "Semantic Remote Attestation;" first employment after graduation: *Google,* Santa Monica, California).

5. Efe Yardimci
(affiliated in August 2001; advanced to candidacy: November 2003; final defense: March 2006; thesis: "Exploiting Parallelism to Improve the Performance of Sequential Binary Executables;" first employment after graduation: *Advanced Micro Devices (AMD),* Santa Clara, California).

6. Christian H. Stork
(affiliated in September 1998; advanced to candidacy: May 2001; final defense: August 2006; thesis: "WELL: A Language-Agnostic Foundation for Compact and Provably Safe Mobile Code;" first employment after graduation: Postdoc at *University of California, Irvine*).

7. Deepak Chandra
(affiliated in August 2001; advanced to candidacy: March 2004; final defense: September 2006; thesis: "Information Flow Analysis and Enforcement in Java Bytecode;" first employment after graduation: *Google,* Irvine, California).

8. Andreas Gal
(affiliated in January 2002; advanced to candidacy: December 2003; final defense: November 2006; thesis: "Efficient Bytecode Compilation and Verification in a Virtual Machine;" first employment after graduation: Postdoc at *University of California, Irvine*).

9. Matthew Beers
(affiliated in September 1999; advanced to candidacy: July 2002; final defense: March 2007; thesis: "Shifting the Burden of Code Optimization to the Code Producer," first employment after graduation: *Ocean Tomo* Intellectual Capital Equity, San Francisco, California).

10. Ning Wang
(affiliated in September 2001; advanced to candidacy: September 2004; final defense: May 2007; thesis: "From Assumptions to Assertions: A Sound and Precise Points-to Analysis for the C Language," first employment after graduation: *Fortify Software*, Palo Alto, California).

11. Vasanth Venkatachalam
(affiliated in September 2002; advanced to candidacy: September 2003; final defense: May 2007; thesis: "Self-Calibrating Processor Speed: A New Feedback Loop For Dynamic Voltage Scaling Control;" first employment after graduation: *Advanced Micro Devices (AMD),* Austin, Texas).

12. Lei Wang
(affiliated in June 2001; advanced to candidacy: September 2004; final defense: June 2009; thesis: "Automatic Program Partitioning to Alleviate Resource Constraints of Object-Oriented Applications;" first employment after graduation: *Microsoft,* Redmond, Washington).

13. Babak Salamat
(affiliated in January 2007; advanced to candidacy: May 2007; final defense: June 2009; thesis: "Multi-Variant Execution: Run-Time Defense Against Malicious Code Injection Attacks;" first employment after graduation: *Yahoo,* Sunnyvale, California).

14. Michael Bebenita
(affiliated in January 2007; advanced to candidacy in May 2009; final defense: October 2011; thesis: "Trace-Based Compilation and Optimization in Meta-Circular Virtual Execution Environments;" first employment after graduation: *Mozilla,* Mountain View, California).

15. Gregor Wagner
(affiliated in September 2007; advanced to candidacy in May 2009; final defense: October 2011; thesis: "Domain Specific Memory Management in a Modern Web Browser;" first employment after graduation: *Mozilla,* Mountain View, California).

16. Mason Liu Chang
(affiliated in June 2007; advanced to candidacy in May 2009; final defense: February 2012; thesis: "Efficient Analysis and Optimization of Dynamically Typed Languages;" first employment after graduation: *Mozilla,* Mountain View, California).

17. Todd Morris Jackson
(affiliated in September 2007; advancement to candidacy in June 2009; final defense: May 2012; thesis: "On the Design, Implications, and Effects of Implementing Software Diversity for Security;" first employment after graduation: *Google,* Mountain View, California).

18. Christoph Kerschbaumer
(affiliated in Summer 2010; advancement to candidacy in November 2011; final defense: March 2014; thesis: "Probabilistic Information Flow Control in Modern Web Browsers;" first employment after graduation: *Mozilla,* Mountain View, California).

19. Eric Hennigan
(affiliated in July 2008; advancement to candidacy in April 2011; final defense: December 2014; thesis: "From FlowCore to JitFlow: Improving the Speed of Information Flow in JavaScript;" first employment after graduation: *Google,* Mountain View, California).

20. Marcelo Cintra
(affiliated in December 2007; advancement to candidacy in November 2009, final defense: April 2015; thesis:

"Just-in-Time Compilation Techniques for Hardware/Software Co-Designed Processors;" first employment after graduation: *Intel,* Santa Clara, California).

21. Andrei Homescu
(affiliated in Fall 2010; advancement to candidacy in March 2012, final defense: April 2015; thesis: "Securing Statically and Dynamically Compiled Programs using Software Diversity;" first employment after graduation: *Immunant,* Irvine, California).

22. Codrut Stancu
(affiliated in Summer 2012; advancement to candidacy in May 2013, final defense: May 2015; thesis: "Safe and Efficient Hybrid Memory Management for Java;" first employment after graduation: *Oracle,* Redwood Shores, California).

23. Wei Zhang
(affiliated in Spring 2011; advancement to candidacy in November 2011, final defense: June 2015; thesis: "Efficient Hosted Interpreters for Dynamic Languages;" first employment after graduation: *Twitter,* San Francisco, California).

24. Stephen Crane
(affiliated in Fall 2011; advancement to candidacy in August 2013, final defense: June 2015; thesis: "Enhancing and Extending Software Diversity;" first employment after graduation: *Immunant,* Irvine, California).

25. Gulfem Savrun Yeniceri
(affiliated in Fall 2010; advancement to candidacy in January 2013, final defense: November 2015; thesis: "Efficient Interpreters and Profilers for Hosted Dynamic Languages;" first employment after graduation: *Intel,* Santa Clara, California).

26. Julian Lettner
(affiliated in Fall 2013; advancement to candidacy in March 2016, final defense: August 2018; thesis: "Finding and Mitigating Memory Corruption Errors in Systems Software;" first employment after graduation: *Apple,* Cupertino, California).

27. Brian Belleville
(affiliated in Fall 2013; advancement to candidacy in June 2016, final defense: August 2018; thesis: "Security Applications of Static Program Analysis;" first employment after graduation: *Google,* Mountain View, California).

28. Mohaned Qunaibit
(affiliated in Summer 2014; advancement to candidacy in March 2016, final defense: March 2019; thesis: "Accelerating Dynamically-Typed Language on Heterogeneous Platforms;" first employment after graduation: *Oracle,* Redwood Shores, California).

29. Joseph Nash
(affiliated in Summer 2015; advancement to candidacy in May 2018, final defense: May 2020; thesis: "Binary Recompilation via Dynamic Analysis and the Protection of Control and Data-Flows Therein;" first employment after graduation: *Advanced Micro Devices (AMD),* Boxborough, Massachussetts).

30. Alexios Voulimeneas
(affiliated in Summer 2015; advancement to candidacy in June 2018, final defense: May 2020; thesis: "Building the Next Generation of Security Focused NVX Systems: Overcoming Limitations of N-Variant Execution;" first employment after graduation: Post-Doctoral Researcher at *K U Leuven,* Belgium).

31. Prabhu Karthikeyan Rajasekaran
(affiliated in Spring 2015; advancement to candidacy in May 2018, final defense: May 2020; thesis: "Practical Run-Time Mitigations Against Data-Oriented Attacks;" first employment after graduation: *Google,* Mountain View, California).

32. Anil Altinay
(affiliated in Summer 2015; advancement to candidacy in June 2018, final defense: May 2020; thesis: "Dynamic Binary Lifting and Recompilation;" first employment after graduation: *Google,* Mountain View, California).

33. Taemin Park
(affiliated in Summer 2015; advancement to candidacy in May 2018, final defense: May 2020; thesis: "Comprehensive Protection for Dynamically-Typed Languages: Avoiding the Pitfalls of Language-Level Sandboxing;" first employment after graduation: *Intel,* Santa Clara, California)).

34. Dokyung Song
(affiliated in Fall 2016; advancement to candidacy in May 2019, final defense: December 2020; thesis: "Precise and Efficient Dynamic Analysis of Systems Software;" after a brief period as a Post-Doctoral Researcher at UCI started as an Assistant Professor at *Yonsei University,* Seoul, South Korea)).

35. Paul Kirth
(affiliated in Fall 2016; advancement to candidacy in March 2019, final defense: November 2021; thesis: "Practical Methods for Automatic Intra-Process Compartmentalization with MPK;" first employment after graduation: *Google,* Mountain View, California).

Graduate Students Supervised as Principal Academic Advisor and Committee Chair

Advanced to Ph.D. Candidacy (in order of advancement date)

1. Min-Yi Hsu (since Fall 2018; advancement to candidacy in March 2021)
2. Fabian Parzefall (from Summer 2018); advancement to candidacy in November 2021)
3. Mitchel Dickerson (from Summer 2018; advancement to candidacy in May 2022)
4. Chinmay Diwakar Deshpande (from Fall 2019; advancement to candidacy in September 2022)

Not Yet Advanced to Candidacy (in order of affiliation date)

5. Tommaso Monaco (from Fall 2020)
6. André Rösti (from Fall 2022)

Graduated Ph.D. Students (Co-Advisor and "Opponent" During Final Dissertation Defense)

1. Christian Wimmer, University of Linz, Austria
(final defense: March 2008; thesis: "Automatic Object Inlining in a Java Virtual Machine").
2. Stefan Brunthaler, Technical University of Vienna, Austria
(final defense: February 2011; thesis: "Purely Interpretative Optimizations").
3. Thomas Würthinger, Johannes-Kepler University of Linz, Austria
(final defense: April 2011; thesis: "Dynamic Code Evolution for Java").
4. Christian Häubl, Johannes-Kepler University of Linz, Austria
(final defense: February 2015; thesis: "Generalized Trace Compilation for Java").
5. Stijn Volckaert, University of Ghent, Belgium
(final defense: October 2015; thesis: "Advanced Techniques for Multi-Variant Execution").
6. Victor van der Veen, Vrije Universiteit Amsterdam, Netherlands
(final defense: October 2019; thesis: "When Memory Serves Not So Well—Memory Errors 30 Years Later").
7. Robert E. Bühren, Technical University of Berlin, Germany
(final Defense: February 2022; thesis: "Resource Control Attacks against Encrypted Virtual Machines).

8. Dominik Maier, Technical University of Berlin, Germany
(final Defense: June 2022; thesis: "Automated Security Testing of Unexplored Targets through Feedback-Guided Fuzzing").

9. Felicitas Hetzelt, Technical University of Berlin, Germany
(final Defense: June 2022; thesis: "Security Analysis of Hardware-OS Interfaces in Linux").

Other Ph.D. Students

Dissertation Committee Member

- Byron Hawkins, UC Irvine
(final defense: August 2017; committee chair: Brian Demsky; thesis: "Introspective Intrusion Detection").

- Andreas Gerstlauer, UC Irvine
(final defense: April 2004, committee chair: Daniel D. Gajski; thesis: "Modeling Flow for Automated System Design and Exploration").

- Ana Lucia Velloso Azevedo, UC Irvine
(final defense: October 2002; committee chair: Alexandru Nicolau; thesis: "Annotation-based Compiler Technology").

- Chang Liu, UC Irvine
(final defense: August 2002, committee chair: Debra J. Richardson; thesis: "Redundant Arrays of Independent Components").

- Martin Burtscher, University of Colorado at Boulder
(final defense: April 2000, committee chair: Benjamin Zorn; thesis: "Improving Context-Based Load Value Prediction").

- Jianwen Zhu, UC Irvine
(final defense: September 1999; committee chair: Daniel D. Gajski; thesis: "Behavioral Synthesis from an Extensible Object Oriented Language").

Candidacy Committee Member

- Semen Pyankov, UC Irvine
(candidacy: February 2022; committee chair: Alexander Veidenbaum).

- Claudio Parra, UC Irvine
(candidacy: December 2020; committee chair: Isaac Scherson).

- Zhihao "Zephy" Yao, UC Irvine
(candidacy: May 2020; committee chair: Ardalan Amiri Sani).

- Wail Alkowaileet, UC Irvine
(candidacy: November 2019; committee chair: Michael J. Carey).

- Ivan Oliveira Nunes, UC Irvine
(candidacy: December 2018; committee chair: Gene Tsudik).

- Tyler Kaczmarek, UC Irvine
(candidacy: December 2015; committee chair: Gene Tsudik).

- Lu Fang, UC Irvine
(candidacy: May 2014; committee chair: Guoquing Xu).

- Nicolae Savoiu, UC Irvine
(candidacy: September 1999; committee chair: Alexandru Nicolau).

M.Sc. Students Graduated from UC Irvine with Thesis Option

- Wail Alkowaileet, M.S. thesis committee member, graduated October 2013 (thesis: "NUMA-aware multicore Matrix Multiplication;" committee chair: Isaac Scherson).
- Alexander Yermolovich, primary M.S. advisor / committee chair, completed M.S. degree in May 2009 (thesis: "Efficient Execution of Binary and Guest Virtual Machines on Platform Independent Host Virtual Machines").
- Mason Liu Chang, primary M.S. advisor / committee chair, completed M.S. degree in May 2009 (thesis: "Tracing for Web 3.0 – Trace Compilation for the Next Generation Web Applications").
- Songmei Han, primary M.S. advisor, graduated with a M.S. in Computer Science in June 2003 (she also received a Ph.D. in Cognitive Science, for which Barbara Dosher was the advisor); subsequently a tenure-track Assistant Professor of Cognitive Science and Computer Science at SUNY Oswego and now Usability Engineer at Apollo Group.
- Anjum Gupta, M.S. thesis committee member, graduated June 2003 (thesis: Design and Implementation of an Adaptive Cache on a Configurable Processor; committee chair: Rajesh Gupta).

Other Graduate Advising

- Jason Alaya, primary M.S. advisor, Fall 2020 – Spring 2022, graduated June 2022.
- Dixin Zhou, primary M.S. advisor, Spring 2018 – Spring 2019, graduated June 2019.
- Faraz Zaerpoor, primary M.S. advisor, Fall 2016 – Spring 2018.
- Anton Vasick, primary M.S. advisor, Summer 2015 – Spring 2018.
- Mark Murphy, research advisor, Fall 2010 – Fall 2015.
- Divya Varshini Agavalam Padmanabhan, primary M.S. advisor, graduated June 2016.
- Nikhil Gupta, primary M.S. advisor, graduated June 2016.
- Roeland Singer-Heinze, primary M.S. advisor, graduated June 2016.
- Stephen Neisius, primary M.S. advisor, graduated Summer 2014.
- Karthikeyan Manivannan, primary advisor, 2007– 2011.
- Sergiy Zhenochin, primary M.S. advisor / committee chair, graduated Fall 2001.
- Prashant Saraswat, primary M.S. advisor / committee chair, graduated Fall 2001.
- Hans-Christian Stadler, primary M.S. advisor / committee chair, graduated June 1998.

Undergraduate Honors Students Graduated from UC Irvine

- Eric Thomas Parsons (summer research advisor, Summer 2018).
- Muneeb Baig (honors research advisor); graduated Magna Cum Laude and Phi Beta Kappa in 2007; honors thesis: "Optimizing Array Bound Checking During Trace-Based Compilation".
- Michael Masukawa (honors research advisor); graduated Summa Cum Laude and Phi Beta Kappa in 2007; honors thesis: "Dynamic Taint Propagation in Java Web Applications".
- Jesse Morrow (honors research advisor); graduated Magna Cum Laude and Phi Beta Kappa in 2005.
- Matthew Chu (honors research advisor); graduated Phi Beta Kappa in 2004.
- Zachary Mouri (honors research advisor); graduated Phi Beta Kappa in 2004.
- Ronald Harvest (honors research advisor); graduated Summa Cum Laude and Phi Beta Kappa in 1999.
- Calvin Shen (honors research advisor); graduated Cum Laude in 1999.

Other Undergraduate Advising

- Rasmus Tjalk-Boggild, visiting from DTU Lynby, Denmark, faculty research advisor; Summer 2016.
- Thomas Bourgenolle, visiting from ENSTA ParisTech, Paris, France, faculty research advisor; Summer 2015.
- Martin Imre, visiting from Technical University of Vienna, Austria, faculty research advisor; Summer 2015.
- Dominik Infuehr, visiting from Technical University of Vienna, Austria, faculty research advisor; Summer 2015.
- David Poetzsch-Heffter, visiting from University of Kaiserslautern, Germany, faculty research advisor; Summer 2015.
- Aditiya Verma, visiting from IIT (BHU) Varanasi, India, faculty research advisor; Summer 2015.
- William Lee, Troy Tech Senior Internship, supervisor; Summer 2015.
- Christos Ioannidis, visiting from University of Thessaly, Greece, faculty research advisor; Summer 2014.
- Mohit Mishra, visiting from Indian Institute of Technology, Varanasi, faculty research advisor; Summer 2014.
- Michalis Papamichail, visiting from Aristotle University of Thessaloniki, Greece, faculty research advisor; Summer 2014.
- Martin Schleiss, visiting from Technical University of Vienna, Austria, faculty research advisor; Summer 2014.
- Henry Elias Hernandez, faculty research advisor, Summer 2014.
- Daniel Nima Salehi, faculty research advisor, Summer 2014.
- Michael Stewart, faculty research advisor, Summer/Fall 2011.
- Jeffrey Bosboom, NSF Research Experiences for Undergraduates Summer Scholar, faculty advisor, 2011.
- Shawn Merrill, NSF Research Experiences for Undergraduates Summer Scholar, faculty advisor, 2011.
- Chris Austin, NSF Research Experiences for Undergraduates Summer Scholar, faculty advisor, 2010.
- Daniel A. Ehrenberg, Carleton University, NSF Research Experiences for Undergraduates Summer Scholar, faculty advisor, 2010.
- Sean Kocol, honors research advisor, 2009.
- Jonathan Mood, honors research advisor, 2009.
- Adrian Tran, honors research advisor, 2009.
- Yaoxiang Zhou, honors research advisor, 2008.
- Raymond Yu, honors research advisor, 2007/08.
- Stephen C. Reed, California Alliance for Minority Participation in Science (CAMP) Summer Scholar, faculty advisor, 2004.

Visiting Diploma Students Supervised at UC Irvine

- Urs Fässler, ETH Zürich, Switzerland, co-supervised with Th. Gross; March–September 2012.
- Alen Stojanov, École polytechnique fédérale de Lausanne (EPFL), Switzerland, co-supervised with M. Odersky; September 2011–March 2012.
- Dominik Lichtenauer, Johannes-Kepler University of Linz, Austria, co-supervised with H. Mössenböck, September 2011–March 2012.
- Stefan Rath, Technische Universität Graz, Austria, co-supervised with Ch. Steger; July–December 2010.
- Franz Maier, Technische Universität Graz, Austria, co-supervised with Ch. Steger; March–September 2010.
- Christoph Kerschbaumer, Technische Universität Graz, Austria, co-supervised with Ch. Steger; March–September 2008.
- Giacomo Amorosa, ETH Zürich, Switzerland, co-supervised with J. Gutknecht; February–August 2008.
- Katharina Seke, Technische Universität Graz, Austria, co-supervised with Ch. Steger; July 2006–October 2006.
- Gregor Wagner, Technische Universität Graz, Austria, co-supervised with Ch. Steger; March 2006–October 2006.
- Albert Noll, Technische Universität Graz, Austria, co-supervised with Ch. Steger; March 2006–September 2006.
- Michael Rauch, Technische Universität Graz, Austria, co-supervised with Ch. Steger; March 2006–September 2006.
- Isabella Thomm, Universität Erlangen-Nürnberg, Germany, co-supervised with W. Schroeder-Preikschat; November 2005–February 2006.
- Michael Stilkerich, Universität Erlangen-Nürnberg, Germany, co-supervised with W. Schroeder-Preikschat; December 2004–February 2005 and November 2005–February 2006.
- Nicolas Marochow, Fachhochschule Braunschweig-Wolfenbüttel, Germany, co-supervised with R. Rüdiger; September 2004–January 2005.
- Jan Peterson, Universität Jena, Germany, co-supervised with W. Amme; May–September 2004.
- Tobias Körner, Fachhochschule Braunschweig-Wolfenbüttel, Germany, co-supervised with R. Rüdiger; March–August 2003.
- Alexander Apel, Universität Jena, Germany, co-supervised with W. Amme; September–November 2003.
- Christian Rattei, Fachhochschule München, Germany, co-supervised with K. Köhler; April–November 2000.
- Joachim Büchse, ETH Zürich, Switzerland, co-supervised with J. Gutknecht,' 1998.
- M. Burtscher, ETH Zürich, 1996.
  (work conducted at Irvine but thesis submitted in Zurich while Franz still had a formal association with ETH)
- M. Dätwyler, ETH Zürich, 1996.
  (work conducted at Irvine but thesis submitted in Zurich while Franz still had a formal association with ETH)

Diploma Students Supervised at ETH Zurich († = co-supervised with N. Wirth)

- E. Brandenberger, Oberon Module Interchange auf Intel-Prozessoren, 1996.
- D. Posva, Dynamische Reoptimierung auf einem RISC, 1996.
- M. Sperisen, Executable Content in WWW-Dokumenten: Java, 1996.
- †H. Buchser, Portable Objektfiles und codegenerierender Lader, 1995.
- †H. Domjan, Metaprogrammierung, 1995.

- †O. Dreer, “Slim Binaries” auf Macintosh, 1995.
- †Th. Kistler, Smartest Recompilation, 1995.
- †Ch. Denzler, A Message Mechanism for Oberon, 1993.
- †E. Oertli, Oberon-2 für Macintosh, 1993.
- †I. Posva, Elimination redundanter Tests durch Programmanalyse, 1993.
- †Th. Bühlmann, Call Optimization for the MacOberon Compiler, 1992.
- †S. Ludwig, A Portable Object and Symbol File Format for Oberon, 1991.
- †S. Meier, Zeichenerkennung mittels Strukturanalyse, 1990.

i

National Security
Agency/Central Security Service

INFORMATION ASSURANCE DIRECTORATE

MOBILE ACCESS CAPABILITY PACKAGE

This Commercial Solutions for Classified (CSfC) Capability Package (CP) describes how to protect classified data (including Voice and Video) in Mobile Access Solutions transiting Wired Networks, Domestic Cellular Networks, and Trusted Wireless Networks to include Government Private Cellular Networks and Government Private Wi-Fi networks.

Version 1.1
June 19, 2015

Netskope Exhibit 1002 - Appendix B

CHANGE HISTORY

Title	Version	Date	Change Summary
Commercial Solutions for Classified (CSfC) Mobile Access Capability Package (CP) release for Public Comment	0.8	November 3, 2014	• Initial release of CSfC Mobile Access guidance for public comment. • Incorporates End User Device (EUD) Solution Designs from VPN version 3.0 CP. • Incorporates content from Mobile Security Guide version 2.3.
Commercial Solutions for Classified (CSfC) Mobile Access CP release for Public Comment	1.0	April 2, 2015	• Removed "Non-MDF Validated" EUD type • Removed EUD design utilizing two VPN Gateways • Removed option to utilize separate computing platform with VPN Client installed to provide Outer layer of encryption • Changed restrictions on control plane traffic • Added Tactical Solution Implementation Appendix • Added requirements for End User Device • Added requirements for Retransmission Device • Corrected language in requirement MA-CR-9 and made consistent with the MACP Compliance Matrix

TABLE OF CONTENTS

1 Introduction 8
2 Purpose of This Document 8
3 Use of This Document 9
4 Description of the Mobile Access Solution 9
4.1 Networks 12
4.1.1 Enterprise/Red Network 12
4.1.2 Gray Network 13
4.1.3 Black Network 14
4.1.4 Data, Management, and Control Plane Traffic 16
4.2 High-Level Design 17
4.2.1 End User Devices 18
4.2.2 Independent Site 21
4.2.3 Multiple Sites 22
4.3 Rationale for Layered Encryption 22
4.4 Authentication 24
4.5 Other Protocols 24
4.6 Availability 25
5 Infrastructure Components 25
5.1 Outer Firewall 26
5.2 Outer VPN Gateway 26
5.3 Gray Firewall 27
5.4 Gray Management Services 27
5.4.1 Outer Certificate Authority (Located on Gray Network) 27
5.4.2 Gray Administration Workstation 28
5.4.3 Gray Certificate Revocation Services 28
5.4.4 Gray Security Information and Event Management (SIEM) 29
5.5 Inner Encryption Components 29
5.5.1 Inner VPN Gateway 30
5.5.2 Inner TLS-Protected Server 30

5.5.3 Inner SRTP Endpoint 31
5.5.4 Inner Firewall 31
5.6 Red Management Services 32
5.6.1 Certificate Authorities (Located on Red Network) 32
5.6.2 Red Administration Workstations 32
5.6.3 Red Security Information and Event Management (SIEM) 32
5.6.4 Red Certificate Revocation Services 33
6 End User Device Components 33
6.1 Outer VPN Component 33
6.1.1 Outer VPN Gateway 34
6.1.2 Outer VPN Client 34
6.2 Virtual private Network (VPN) End User Device (EUD) 34
6.2.1 Inner VPN Client on Computing Device 35
6.3 Transport Layer Security (TLS) End User Device (EUD) 35
6.3.1 TLS Client 36
6.3.2 SRTP Client 36
7 Mobile Access Configuration and Management 37
7.1 Solution Infrastructure Component Provisioning 37
7.2 EUD Provisioning 37
7.3 Administration of Mobile Access Components 38
7.4 EUDs for Differing Classification Domains 39
8 Continuous Monitoring 39
8.1 Monitoring Network Traffic 39
8.2 Monitoring Log Data 41
9 Key Management 42
9.1 Distribution Of Certificate Revocation Lists 45
10 Threats 47
10.1 Passive Threats 47
10.2 External (Active) Threats 48
10.2.1 Rogue Traffic 48

10.2.2 Malware and Untrusted Updates 49
10.2.3 Denial of Service 49
10.2.4 Social Engineering 50
10.3 Insider Threats 50
10.4 Supply Chain Threats 50
10.5 Integrator Threats 52
11 Requirements Overview 52
11.1 Capabilities 52
11.2 Threshold and Objective Requirements 53
11.3 Requirements Designators 54
12 Requirements for Selecting Components 56
13 Configuration Requirements 59
13.1 Overall Solution Requirements 60
13.2 Configuration Requirements for All VPN Components 62
13.3 Configuration Requirements For Inner and Outer VPN Components 65
13.4 Inner VPN Components 66
13.5 Outer VPN Components 67
13.6 TLS-Protected Server & SRTP Endpoint Requirements 68
13.7 Retransmission Device requirements 69
13.8 End User Devices Requirements 71
13.9 Port Filtering Requirements for Solution Components 75
13.10 Configuration Change Detection Requirements 77
13.11 Device Management Requirements 77
13.12 Continuous Monitoring Requirements 79
13.13 Auditing Requirements 81
13.14 Key Management Requirements 84
13.14.1 General Requirements 84
13.14.2 Certificate Issuance Requirements 85
13.14.3 Certificate Renewal and Rekey Requirements 87
13.14.4 Certificate Revocation and CDP Requirements 87

14 Requirements for Solution Operation, Maintenance, and Handling 89
14.1 Requirements for the Use and Handling of Solutions 89
14.2 Requirements for Incident Reporting 92
15 Role-Based Personnel Requirements 94
16 Information to Support AO 97
16.1 Solution Testing 98
16.2 Risk Assessment 99
16.3 Registration of Solutions 99
17 Testing Requirements 99
17.1 Product Selection 100
17.2 Physical Layout of Solution 102
17.3 TLS-Protected Server Configurations 103
17.4 End User Device Configurations 103
17.5 Retransmission Device Configuration 107
17.6 Inner And Outer VPN Component Configurations 108
17.7 Key Management 111
17.8 Solution Filtering configurations 118
17.9 Configuration Change Detection 119
17.10 Continuous Monitoring 120
17.11 Audit 121
17.12 EUD With Multiple Connections 124
17.13 Incident Reporting Guidance 125
17.14 Implementation of Guidance 126
17.15 Solution Functionality 127
Appendix A. Glossary of Terms 128
Appendix B. Acronyms 132
Appendix C. References 135
Appendix D. End User Device Implementation Notes 138
Appendix E. Tactical Solution implementations 143

TABLE OF FIGURES

Figure 1. Two Layers of Encryption Protect Data across an Untrusted Network 10
Figure 2. Acceptable Black Transport Networks 15
Figure 3. EUD Solution Designs 19
Figure 4. EUDs Connected to Independent Site 21
Figure 5. Multiple Mobile Access Solution Infrastructures supporting EUDs 22
Figure 6. MA Solution Continuous Monitoring Points 40

LIST OF TABLES

Table 1. Overview of Mobile Access CP Terminology 10
Table 2. Acceptable Black Transport Networks 15
Table 4. Certificate Authority Deployment Options 44
Table 5. Capability Designators 53
Table 6. Requirement Digraphs 54
Table 7. Product Selection Requirements 56
Table 8. Overall Solution Requirements 60
Table 9. Approved Suite B Algorithms (IPSec) 62
Table 10. Approved Algorithms (TLS) 63
Table 11. Approved Algorithms (Secure Real-Time Protocol) 64
Table 13. Inner VPN Components Requirements 66
Table 14. Outer VPN Components Requirements 67
Table 15. TLS-Protected Server & SRTP Endpoint Requirements 68
Table 16. Requirements for Retransmission Device 69
Table 17. Requirements for End User Devices 71
Table 18. Port Filtering Requirements for Solution Components 75
Table 19. Configuration Change Detection Requirements 77
Table 20. Requirements for Device Management 77
Table 21. Continuous Monitoring Requirements 79
Table 22. Auditing Requirements 81
Table 23. PKI General Requirements 84
Table 24. Certificate Issuance Requirements 85

Table 25. Certificate Renewal and Rekey Requirements 87
Table 26. Requirements for Certificate Revocation and CDPs 87
Table 27. Requirements for the Use and Handling of Solutions 89
Table 28. Incident Reporting Requirements 92
Table 29. Role-Based Personnel Requirements 96
Table 30. Test Requirements 98
Table 31. Tactical Implementation Requirements Overlay 144

1 INTRODUCTION

The Commercial Solutions for Classified (CSfC) Program within the National Security Agency (NSA) Information Assurance Directorate (IAD) publishes Capability Packages (CPs) to provide configurations that empower IAD customers to implement secure solutions using independent, layered Commercial Off-the-Shelf (COTS) products. The CPs are product-neutral and describe system-level solution frameworks documenting security and configuration requirements for customers and/or Solution Integrators.

IAD delivers this CSfC Mobile Access (MA) CP to meet the demand for mobile data in transit solutions (including Voice and Video) using approved cryptographic algorithms and National Information Assurance Partnership (NIAP) evaluated components. These algorithms, known as Suite B algorithms, are used to protect classified data using layers of COTS products. MA CP Version 1.0 is the initial release to provide customer requirements for domestic and international voice, video, and data capabilities from a mobile End User Device. This CP builds on lessons learned from a CSfC Initial Solution that implemented secure voice and data capabilities using a set of Secure Sharing Suite (S3) algorithms, modes of operation, standards, and protocols. The CSfC Initial Solution leveraged lessons learned from customer pilot solutions utilized to protect classified data in mobile environments. The Initial Solution and Pilots included a layered use of COTS products for the protection of classified information.

The CSfC MA CP Version 1.0 builds on the End User Device (EUD) designs of the Virtual Private Network (VPN) CP Version 3.0 as well as the Mobility Security Guide Version 2.3. VPN Remote and Local EUD Designs which complied with the VPN CP Version 3.0, dated July 22, 2014, are expected to need only minor changes to comply with the VPN EUD Design of this CP.

2 PURPOSE OF THIS DOCUMENT

This CP provides high-level reference designs and corresponding configuration requirements that allow customers to select COTS products from the CSfC Components List available on the CSfC web page (http://www.nsa.gov/ia/programs/csfc_program), for their MA solution and then to properly configure those products to achieve a level of assurance sufficient for protecting classified data while in transit. As described in Section 12, customers must ensure that the components selected from the CSfC Components List will permit the necessary functionality for the selected capabilities. To successfully implement a solution based on this CP, all Threshold requirements, or the corresponding Objective requirements applicable to the selected capabilities, must be implemented, as described in Sections 11-13.

Customers who want to use this CP must register their solution with NSA. Additional information about the CSfC process is available on the CSfC web page (www.nsa.gov/ia/programs/csfc_program).

3 USE OF THIS DOCUMENT

This document, the CSfC Mobile Access CP Version 1.0, dated February 13, 2015, has been approved by the IAD Director and will be reviewed twice a year to ensure that the defined capabilities and other instructions still provide the security services and robustness required. Solutions designed according to this CP must be registered with NSA/IAD. Once registered, a signed IAD Approval Letter will be sent validating that the Mobile Access solution is registered as a CSfC solution validated to meet the requirements of the latest Mobile Access CP and is approved to protect classified information. Any solution designed according to this CP may be used for one year and must then be revalidated against the most recently published version of this CP.

Please provide comments on usability, applicability, and/or shortcomings to your NSA/IAD Client Advocate and the MA CP maintenance team at Mobile_Access@nsa.gov. MA CP solutions shall also comply with Committee on National Security System (CNSS) policies and instructions. Any conflicts identified between this CP and NSS or local policy should be provided to the MA CP Maintenance team.

The following Legal Disclaimer relates to the use of this CP:

This CP is provided "as is." Any express or implied warranties, including but not limited to, the implied warranties of merchantability and fitness for a particular purpose are disclaimed. In no event shall the United States Government be liable for any direct, indirect, incidental, special, exemplary or consequential damages (including, but not limited to, procurement of substitute goods or services, loss of use, data, or profits, or business interruption) however caused and on any theory of liability, whether in contract, strict liability, or tort (including negligence or otherwise) arising in any way out of the use of this CP, even if advised of the possibility of such damage.

The user of this CP agrees to hold harmless and indemnify the United States Government, its agents and employees from every claim or liability (whether in tort or in contract), including attorney's fees, court costs, and expenses, arising in direct consequence of Recipient's use of the item, including, but not limited to, claims or liabilities made for injury to or death of personnel of User or third parties, damage to or destruction of property of User or third parties, and infringement or other violations of intellectual property or technical data rights.

Nothing in this CP is intended to constitute an endorsement, explicit or implied, by the U.S. Government of any particular manufacturer's product or service.

4 DESCRIPTION OF THE MOBILE ACCESS SOLUTION

This CP describes a general MA solution to protect classified information as it travels across either an untrusted network or a network of a different classification level. The solution supports connecting EUDs to a classified network via two layers of encryption terminated on the EUD, provided that the EUD and the network operate at the same security level. The MA solution uses two nested, independent

tunnels to protect the confidentiality and integrity of data (including Voice and Video) as it transits the untrusted network. The MA Solution utilizes Internet Protocol Security (IPsec) as the outer tunnel and, depending on the Solution Design, IPsec or Transport Layer Security (TLS) as the Inner layer of protection.

Throughout this CP, the term "Inner Encryption Component" is used to refer generically to the component (device or software application) that terminates the Inner layer of encryption. An Inner Encryption Component can be a VPN Component or a TLS Component that is in the infrastructure or part of an EUD. The term "VPN Component" refers generically to both VPN Gateways and VPN Clients in situations where the differences between the two are unimportant. The Term "TLS Component" is used to denote a component that implements TLS between the infrastructure (TLS-Protected Server or Secure Real-time Transport Protocol (SRTP) Endpoint) and EUDs (TLS Client or SRTP Client) in accordance with this CP (see Sections 5.5 and 6 respectively). Finally, there are two EUD solution designs: VPN EUD and TLS EUD. The term "EUD" is used to refer generically to both designs where the differences between them are unimportant.

Table 1. Overview of Mobile Access CP Terminology

	VPN EUD	TLS EUD
Inner Encryption Component	IPsec provided by VPN Client	TLS or SRTP provided by TLS-Protected Server, SRTP Endpoint, TLS Client, **OR** SRTP Client
Outer Encryption Component	IPsec provided by VPN GW **OR** VPN Client	IPsec provided by VPN GW **OR** VPN Client

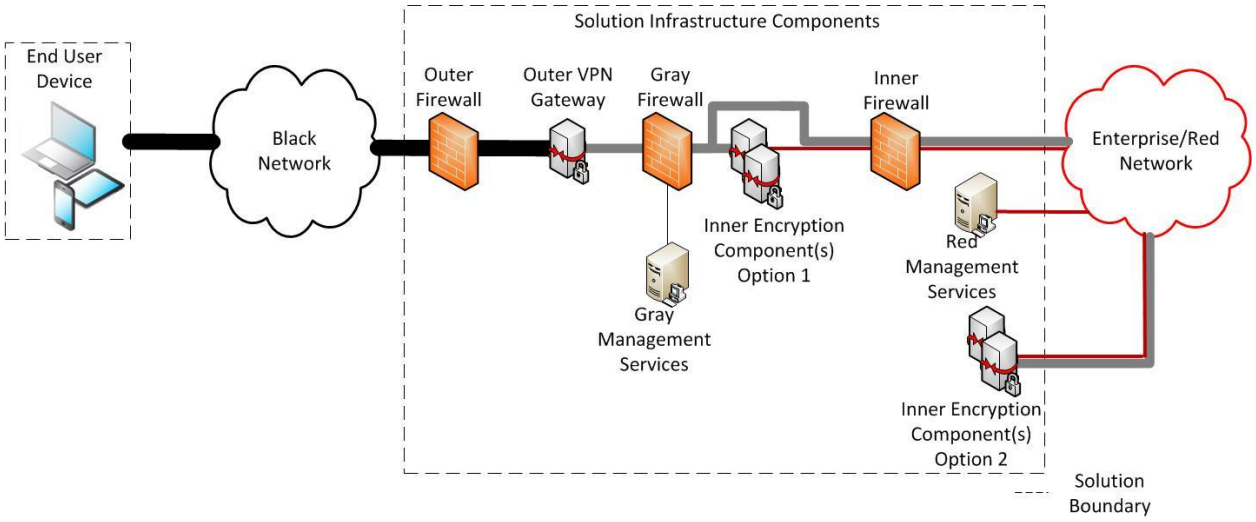


Figure 1. Two Layers of Encryption Protect Data across an Untrusted Network

As shown in Figure 1, before being sent across the untrusted network, classified data is encrypted twice: first by an Inner Encryption Component, and then by an Outer VPN Component. At the other end of the data flow, the received packet is correspondingly decrypted twice: first by an Outer VPN Component, and then by an Inner Encryption Component.

All Inner Encryption components, including those with only a connection to the Enterprise/Red network, are within the CSfC Solution Boundary. The MA CP allows the use of existing Classified Enterprise Network Encryption Components to provide the Inner layer of protection (e.g. Web Servers and Voice over Internet Protocol (VoIP) Desktop Phones), provided that the components are selected, configured, and tested in accordance with the requirements of this CP. When existing Classified Enterprise Network services are utilized as the Inner layer of encryption, they are within the MA Solution Boundary and reside on the internal side of the Inner firewall. There is no limit to the number of Inner Encryption Components permitted to terminate the Inner layer of encryption.

MA Solution components are managed using Red Management Services for Inner Encryption Components, and Gray Management Services for Outer Encryption Components. The Gray Management Services include an Administration Workstation to manage the Outer VPN Gateway, Gray firewall, a Security Information and Event Monitoring (SIEM) Component, and any additional components located between the Outer VPN Gateway and Inner Encryption Components. Gray Management Services may also manage a locally run Outer Certificate Authority (CA), Certificate Revocation List (CRL) Distribution Point (CDP), and/or Intrusion Detection System (IDS)/Intrusion Protection System (IPS). The Red Management Services include an Administration Workstation to manage the Inner Encryption Components, Inner firewall and other components within the Enterprise/Red Network. The Red Management Services may also manage a locally run Inner Certificate Authority (CA), and optionally, a Locally-run Outer CA. In addition, the MA CP allows customers to leverage an existing Enterprise Public Key Infrastructure (PKI) to issue Certificates to Outer VPN Components and Inner Encryption Components. To utilize an existing enterprise Root CA at least two separate subordinate CAs must be utilized: One to issue Certificates for Outer VPN Components and the other to issue Certificates for Inner Encryption Components.

The EUDs of the MA CP are form-factor agnostic. Typical MA CP EUDs include smart phones, tablets, and laptops. An MA CP EUD can be composed of multiple physical devices (for example a VPN Gateway and a Computing Device) all collectively referred to as the EUD. Although the CP allows flexibility in the selection of the EUD, the customer and Solution Integrator must ensure that EUDs meet all applicable requirements for the planned solution design. Section 4.2.1 describes in detail the differences between the VPN EUD and TLS EUD solution design options.

The MA CP instantiations are built using products from the CSfC Components List (see Section 12). Customers who are concerned that their desired products are not yet on the CSfC Components List are encouraged to contact the vendors to encourage them to sign a Memorandum of Agreement (MoA) with NSA and start the National Information Assurance Partnership (NIAP) evaluation process, which will

enable them to be listed on the CSfC Components List. Products listed on the CSfC Components List are not guaranteed to be interoperable with all other products on the Components List. Customers and Integrators should perform interoperability testing to ensure the components selected for their MA Solution are interoperable. If you need assistance obtaining vendor POC information, please email csfc_components@nsa.gov.

While CSfC encourages industry innovation, trustworthiness of the components is paramount. Customers and their integrators are advised that modifying a NIAP-evaluated component in a CSfC solution may invalidate its certification and trigger a revalidation process. To avoid delays, customers or integrators who feel it is necessary to modify a component should engage the component vendor and consult NIAP through their Assurance Continuity Process (see http://www.niap-ccves.org/Documents_and_Guidance/ccves/scheme-pub-6.pdf) to determine whether such a modification will affect the component's certification. In case of a modification to a component, NSA's CSfC Program Management Office requires a statement from NIAP that the modification does not alter the certification, or the security, of the component. Modifications which trigger the revalidation process include, but are not limited to the following: modifying the original equipment manufacturers' code (to include digitally signing the code) or not leveraging the baseline NIAP-evaluated configuration.

4.1 NETWORKS

This CP uses the following terminology to describe the various networks that comprise a MA solution and the types of traffic present on each. The terms Red, Gray, and Black refer to the level of protection applied to the data as described below.

4.1.1 ENTERPRISE/RED NETWORK

Red data consists of unencrypted classified data while Gray data consists of singly encrypted classified data. A Red network contains Red data and can contain Gray data. The Red network begins at the internal interface(s) of Inner Encryption Components located between the Gray firewall and Inner firewall. When a solution is implemented without an Inner Encryption Component between the Gray firewall and Inner firewall, then the Red network begins on the internal interface of the Inner firewall. EUDs access the Red network through the two layers of nested encryption described in this CP. For example, an Inner VPN Gateway, located between the Gray firewall and Inner firewall, terminates the Inner layer of IPsec encryption from a VPN EUD. Once a successful IPsec connection is established, the EUD is given access to classified services such as web, email, Virtual Desktop Infrastructure (VDI), voice, etc.

In some instances when the MA Infrastructure is designed to support TLS EUDs, the TLS-Protected Server or SRTP Endpoint, which terminates the Inner layer of encryption, will reside in the Enterprise/Red network on the internal side of the Inner firewall. In this design, the Inner firewall implements an Access Control List (ACL) which only permits traffic to reach TLS-Protected Servers and SRTP Endpoints which are within the CSfC Solution Boundary and terminate the Inner layer of

encryption. For example, an organization can choose to include Enterprise/Red network Web Servers within the CSfC Solution Boundary. The Web Servers would meet all applicable requirements of the MA CP and be configured to terminate an Inner layer of TLS traffic (Hypertext Transfer Protocol Secure (HTTPS)) originating from an EUD TLS Client (Web Browser). An alternative to this approach is to implement a TLS-Protected Server that includes both Gray and Red network interfaces located between the Gray firewall and Inner firewall. This TLS-Protected Server terminates the TLS connection from the EUD and proxies web data to the Enterprise/Red Web Servers which in this case would not be in the CSfC Solution Boundary. A similar situation exists for SRTP, by using a VoIP Gateway/Border Controller to terminate the SRTP traffic for an EUD, and relaying the data to the Enterprise/Red network. When a VoIP Gateway/Border Controller terminates the Inner layer of SRTP, Desktop Phones in the Enterprise/Red network are not included in the Solution Boundary.

The Enterprise/Red networks are under the control of the solution owner or a trusted third party. Enterprise/Red networks may only communicate with an EUD through the MA solution if both operate at the same security level.

4.1.2 GRAY NETWORK

Gray data is classified data that has been encrypted once. Gray networks are composed of Gray Data. Gray networks are under the physical and logical control of the solution owner or a trusted third party. Gray Data may extend into the Enterprise/Red network when enterprise TLS-Protected Servers (i.e. Web server) or SRTP Endpoints (i.e. VoIP Desktop Phone) terminate the Inner layer of encryption as described in Section 4.1.1. The Gray network is physically treated as a classified network even though all classified data is singly encrypted. If a solution owner's classification authority determines that data on a Gray Network is classified, perhaps by determining the Internet Protocol (IP) addresses are classified at some level; then, the MA solution described in this CP cannot be implemented, as it is not designed to provide two layers of protection for any classified information on the Gray Network.

Gray network components consist of the Solution Infrastructure Outer VPN Gateway, Gray firewall, and Gray Management Services. All Gray network components are physically protected at the same level as the Enterprise/Red network components of the Solution Infrastructure. Gray Management Services are physically connected to the Gray firewall and includes, at a minimum, an Administration Workstation. The MA CP requires the management of Gray Network components through the Gray Administration Workstation. As a result, neither Enterprise/Red nor Black administration workstations are permitted to manage the Outer VPN Gateway, Gray firewall, or Gray Management Services. Additionally, the Gray Administration Workstation is prohibited from managing Inner Encryption Components. These Inner Encryption Components must be managed from an Enterprise/Red Administration Workstation.

4.1.3 BLACK NETWORK

A Black network contains classified data that has been encrypted twice. The network connecting the Outer VPN Components together is a Black network. Black networks are not necessarily (and often will not be) under the control of the solution owner, and may be operated by an untrusted third party. The MA CP allows EUDs to operate over any Black network when used in conjunction with a Government-owned Retransmission Device (RD) or a physically separate VPN Gateway to establish the Outer IPsec Tunnel. A RD provides a connection to the MA solution infrastructure via any Black network and interfaces to the EUD using Wi-Fi or an Ethernet cable however, the CP does not permit the use of Ethernet over USB. Black networks include non-domestic cellular carrier networks, public Wi-Fi networks, wired connections (to the Internet for example), and any other wireless or wired networks.

The Government-owned RD is a category which includes Wi-Fi Hotspots and Mobile Routers. On the external side, the RD can be connected to any type of medium (e.g. Cellular, Wi-Fi, SATCOM, Ethernet) to gain access to a Wide Area Network. On the internal side the RD is connected to EUDs either through an Ethernet cable or Wi-Fi. When the RD is a Wi-Fi access point connected to the EUD (or multiple EUDs), the Wi-Fi network shall implement Wi-Fi Protected Access II (WPA2) with either Pre-Shared Key (PSK) or WPA2 Enterprise. The EUD shall be configured to only permit connections to authorized RDs. RDs are only permitted to establish connectivity to the Black network, and may not be placed between an Outer Encryption Component and Inner Encryption Component.

The CP also allows connectivity without the use of a RD or physically separate VPN Gateway if any of the following Transport Networks are utilized: Domestic Cellular Providers, Government Private Cellular Networks, or Government Private Wireless Networks. Domestic Cellular Providers enable connectivity through cellular Base Stations geographically located within the United States of America. Government Private Cellular Networks are defined as cellular Base Stations which are owned and operated exclusively by the United States Government (such as in Tactical Environments). Finally, Government Private Wireless Networks denotes Wi-Fi connectivity by a Wireless Local Area Network (WLAN) accredited by a Government Authorizing Official (AO). These Wi-Fi networks must comply with applicable organization policies. Within the Department of Defense (DOD) the applicable policy is DOD Instruction (DODI) 8420.01. At a minimum these Wi-Fi networks must implement WPA2 with PSK; however, WPA2 with certificate-based authentication is preferred. When Government Private Wireless Networks utilize certificate-based authentication, they cannot share the Outer Tunnel CA or Inner Tunnel CA certificate management services. WPA2 protects the Black transport network, but does not count as one of the layers of CSfC Data-in-Transit encryption.

Finally, the CP allows connectivity without a RD or physically separate VPN Gateway when implementing the VPN EUD Solution Design connecting through a Wired Internet connection (e.g. Residential Internet connection using Ethernet between the EUD and a Router).

Table 2. Acceptable Black Transport Networks

	VPN EUD	TLS EUD
Any Black Transport Network	Government RD **OR** VPN Gateway	Government RD **OR** VPN Gateway
Domestic Cellular, Government Private Cellular, or Government Private Wireless	No additional requirements	No additional Requirements
Wired Internet Connection	No additional requirements	No additional Requirements

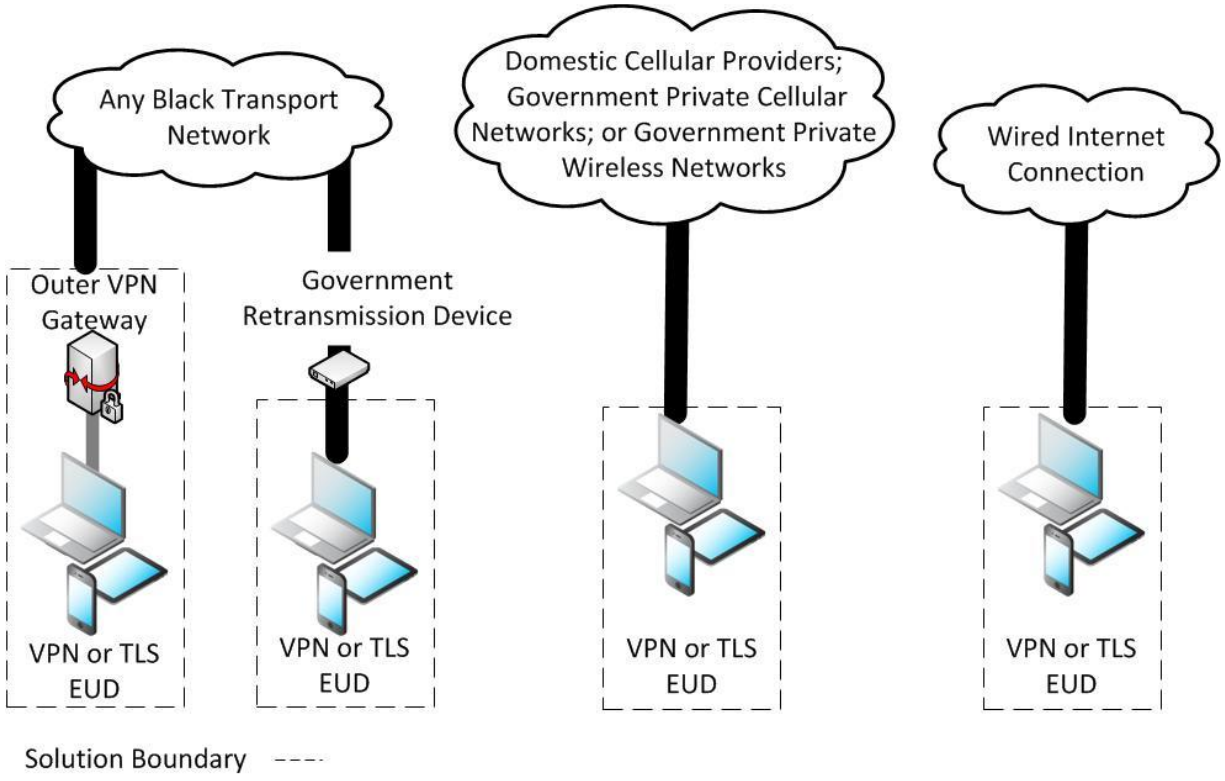


Figure 2. Acceptable Black Transport Networks

As shown in Figure 2, both EUD designs can connect to the Mobile Access solution over Domestic Cellular, Government Private Cellular, or Government Private Wireless Networks without the need for a separate standalone piece of hardware. Additionally, when a VPN EUD uses a Wired Internet connection there is no need for a separate, standalone piece of hardware. When connecting over any other black transport network, EUDs must use a standalone VPN Gateway or a Government RD to connect to the Mobile Access Solution. When an EUD includes a standalone VPN Gateway, that

Gateway is utilized to establish the Outer Layer of IPsec to the government infrastructure and is included within the CSfC Solution Boundary. Additionally, the VPN Gateway must be wired to the computing platform which terminates the Inner layer of encryption. Although only required as described above, a standalone VPN Gateway can be utilized for any of the EUD Solution designs to connect to any transport network. Similarly, an EUD can utilize a Government RD to connect to any transport network. The Government RD is outside the CSfC Solution Boundary, but acts as an intermediary between the desired transport network and the EUD.

4.1.4 DATA, MANAGEMENT, AND CONTROL PLANE TRAFFIC

Data plane traffic is classified information, encrypted or not, that is being passed through the MA solution. The MA solution exists to encrypt and decrypt data plane traffic. All data plane traffic within the Black network is encapsulated within the Encapsulating Security Payload (ESP) protocol. All data plane traffic within the Gray network is encapsulated within ESP, TLS, or SRTP.

Management plane traffic is used to configure and monitor solution components. It includes the communications between a system administrator and a component, as well as the logs and other status information forwarded from a solution component to a SIEM or similar repository. Management plane traffic on Red and Gray networks is encapsulated within the Secure Shell version 2 (SSHv2), ESP, or TLS protocol.

Control plane traffic consists of standard protocols necessary for the network to function. Control plane traffic is typically not initiated directly on behalf of a user (unlike data traffic) or a system administrator (unlike management traffic). Many, but not all, control plane protocols operate at Layer 2 or Layer 3 of the Open Systems Interconnection (OSI) model. Examples of control plane traffic include, but are not limited to, the following:

- Network address configuration (e.g. Dynamic Host Configuration Protocol (DHCP), Neighbor Discovery Protocol (NDP), etc.)
- Address resolution (e.g. Address Resolution Protocol (ARP), NDP, etc.)
- Name resolution (e.g. Domain Name System (DNS), etc.)
- Time synchronization (e.g. Network Time Protocol (NTP), Precision Time Protocol (PTP), etc.)
- Route advertisement (e.g. Routing Information Protocol (RIP), Open Shortest Path First (OSPF), Intermediate System to Intermediate System (IS-IS), Border Gateway Protocol (BGP), etc.)
- Certificate status distribution (e.g. Online Certificate Status Protocol (OCSP), HTTP download of CRLs, etc.)

The MA CP explicitly prohibits the use of most control plane traffic for EUDs consisting of a single computing device to provide both the Inner and Outer layer of encryption (see Appendix D). The MA CP does not allow route advertisement or certificate status distribution to ingress/egress from the Black Transport Network for these EUDs. As a result, the implementing organization will require procedures to be in place to handle a situation in which the certificate of an Outer VPN Gateway is revoked. EUDs are configured to enable all IP traffic, with the exception of network Internet Key Exchange (IKE), address configuration, time synchronization, and name resolution traffic, to flow through the Outer IPsec VPN Client. EUDs selected from the CSfC Components List can ensure that IP traffic flows through the Outer IPsec VPN Client by placing the device in the NIAP evaluated configuration. Upon establishing the Outer VPN tunnel, the CP does not impose detailed requirements restricting control plane traffic in the Gray and Red networks.

Restrictions are also placed on control plane traffic for the Outer VPN Gateway. The Outer VPN Gateway is prohibited from implementing routing protocols on external and internal interfaces. The Outer VPN Gateway can rely on the Outer firewall to perform routing functionality.

Except as otherwise specified in this CP, the usage of specific control plane protocols is left to the solution owner to approve, but any control plane protocols not approved by the solution owner must be disabled or blocked.

Data plane and management plane traffic are generally required to be separated from one another by using physical or cryptographic separation. Use of a Virtual Local Area Network (VLAN) alone is not sufficient to separate data plane and management plane traffic. As a result, a solution may, for example, have a Gray Data network and a Gray Management network which are separate from one another, where the components on the Gray Management network are used to manage the components on the Gray Data network. The Gray Management network is separated from the Gray Data network via the Gray firewall. The Gray firewall utilizes an ACL to ensure that only appropriate Gray Management Services (i.e. Administration Workstation, SIEM or Network Time Server) can communicate with the Outer VPN Gateway and EUDs that have established an Outer VPN tunnel. Given that some control plane traffic is necessary for a network to function, there is no general requirement that control plane traffic be similarly separated, unless otherwise specified.

4.2 High-Level Design

The MA solution is adaptable to support multiple capabilities, depending on the needs of the customer implementing the solution. The supported EUD capabilities are mutually exclusive, therefore if a customer chooses to implement an EUD using two layers of IPsec, then the Inner TLS Client would not be included as part of that EUD implementation. Similarly, if a customer only needs a Secure Voice Capability, then the Inner IPsec Component would not be included as part of that EUD implementation. Although the EUD Solution Designs are mutually exclusive, the infrastructure may be configured to support both EUD solution designs (see Appendix D). This enables implementation of both types of

EUDs based on use cases and device features. Any implementation of the MA solution must satisfy all of the applicable requirements specified in this CP, as explained in sections 12 and 13.

4.2.1 End User Devices

This CP uses the concept of an EUD, which includes a computing device such as a smart phone, laptop, or tablet. The EUDs provide two layers of protection for data in transit to tunnel through the black network and access classified data on the Enterprise/Red network. In some instances, an EUD encompasses more than one piece of hardware (e.g. Computing Device and VPN Gateway) each of which perform a layer of encryption. Where more than one piece of hardware is used, each component is included as part of the EUD and are within the CSfC Solution Boundary. There are two EUD designs which can be implemented as part of a MA solution. Each of the EUD designs share many requirements in common, but also have unique requirements specific to that design:

1) **IPsec-IPsec (VPN EUD):** Utilizes two IPsec tunnels to connect to the Enterprise/Red network. Such an EUD includes both an Inner VPN Client and Outer VPN Component to provide the two layers of IPsec. Throughout the document this EUD design is referred to as the "VPN EUD". VPN EUDs can be implemented utilizing combinations of IPsec VPN Clients and IPsec Gateways (see Appendix D). For example, a VPN EUD can be implemented on a computing device with two VPN Clients running on separate IP stacks. Similarly, the MA CP allows a VPN EUD to utilize a VPN Gateway to provide the Outer layer of IPsec encryption and a VPN Client installed on a computing device to provide the Inner layer of encryption.

2) **IPsec-TLS (TLS EUD):** Utilizes an Outer layer of IPsec encryption and an Inner layer of TLS encryption to access the Enterprise/Red network. Throughout the document this EUD design is referred to as "TLS EUD". The Outer layer of encryption can be provided by either an IPsec VPN Client or a standalone IPsec VPN Gateway. The Inner layer of encryption is then provided by a TLS Client. The EUD TLS Client includes a number of different options which can be selected, in accordance with the CP requirements, to meet the operational needs of the customer. The EUD TLS Clients include, but are not limited to, Web Browsers, Email Clients, and VoIP Applications. Traffic between the TLS EUD Client and the TLS-Protected Server is encrypted with TLS or in some instances SRTP. When SRTP is utilized, session keys are first exchanged using Session Initiation Protocol (SIP) over TLS.

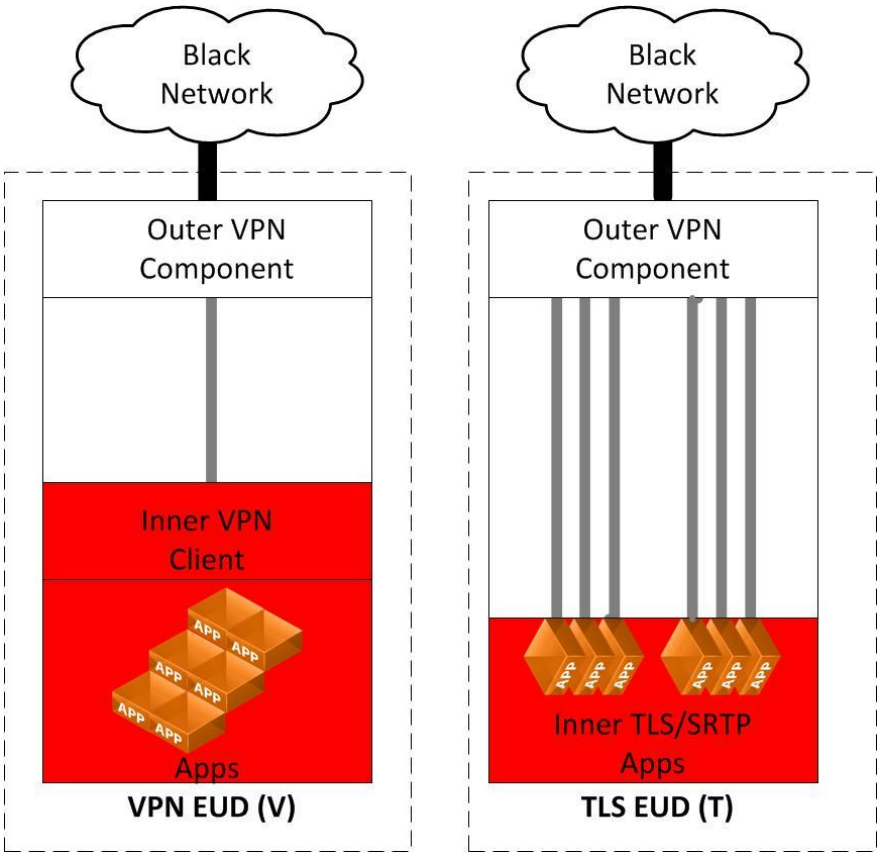


Figure 3. EUD Solution Designs

Figure 3 depicts the two EUD Solution Designs available as part of the MA CP. In each design the Outer VPN Component is utilized to establish an IPsec tunnel to the Outer VPN Gateway of the Mobile Access Solution infrastructure. In either EUD design this Outer VPN Component could be either a VPN Client or a VPN Gateway selected from the CSfC Components List. If a standalone VPN Gateway is utilized to provide the Outer IPsec tunnel, then it must be wired directly to the EUD (i.e. Ethernet cable).

The Inner layer of encryption for VPN EUDs is provided by a VPN Client. The Inner VPN Client must be selected from the CSfC Components List (see Section 12). If VPN Clients are used for both the Inner and Outer layers of encryption then they must utilize a different IP stack, and generally are implemented using Virtualization.

The Inner layer of encryption for TLS EUDs is provided by either TLS or SRTP. Every application which performs TLS or SRTP must be selected from the CSfC Components List.

The Mobile Access CP allows three different deployment options pertaining to the use and handling of an EUD while powered off:

1. **EUD with DAR:** To implement Data-at-Rest on an EUD, the DAR solution shall be approved by NSA – either as a tailored solution, or compliant with NSA's Data at Rest Capability Package (DAR CP) for the protection of information classified at the level of the Enterprise/Red network connected to the EUD. Specification of such a DAR solution is outside the scope of this CP, but can be found in the Data at Rest CP. Positive control of the EUD must be maintained at all times. The EUD is considered lost if out of positive control for twenty (20) minutes or more where positive control is defined by the AO.

2. **Thin EUD:** The EUD can be designed to prevent any classified information from being saved to any persistent storage media on the EUD. Possible techniques for implementing this include, but are not limited to: using VDI configured not to allow data from the Enterprise/Red network to be saved on the EUD, restricting the user to a non-persistent virtual machine on the EUD, and/or configuring the EUD's operating system to prevent the user from saving data locally. Since the EUD does not provide secure local storage for classified data, its user is also prohibited by policy from saving classified data to it. The EUD in this case must enable the native platform DAR protection to protect the private keys stored on it from disclosure and to increase the difficulty of tampering with the software and configuration. This option is not permitted if any of the private keys or certificates stored on the EUD are considered classified by the AO. Positive control of the EUD must be maintained at all times. The EUD is considered lost if out of positive control for twenty (20) minutes or more where positive control is defined by the AO.

3. **Classified EUD:** The EUD can be used exclusively with physical security measures approved by the AO. EUDs are not subject to special physical handling restrictions beyond those applicable for classified devices, since they can rely on the environment they are in for physical protection. If this design option is selected, then the EUDs must be treated as classified devices at all times. The EUD in this case must enable the native platform DAR protection to protect the private keys stored on it from disclosure and to increase the difficulty of tampering with the software and configuration. Positive control of the EUD must be maintained at all times. The EUD is considered lost if out of positive control for twenty (20) minutes or more where positive control is defined by the AO.

While powered on, an EUD is classified at the same level of the Enterprise/Red network that it communicates with through the MA solution, since classified data may be present in volatile memory and/or displayed on screen. To mitigate the risk of accidental disclosure of classified information to unauthorized personnel while the EUD is in use, the customer must define and implement an EUD user agreement that specifies the rules of use for the system. The customer must only grant users access to

an EUD after they agree to the user agreement and receive training on how to use and protect their EUD. There is no limit to the number of EUDs that may be included in a MA solution.

4.2.2 INDEPENDENT SITE

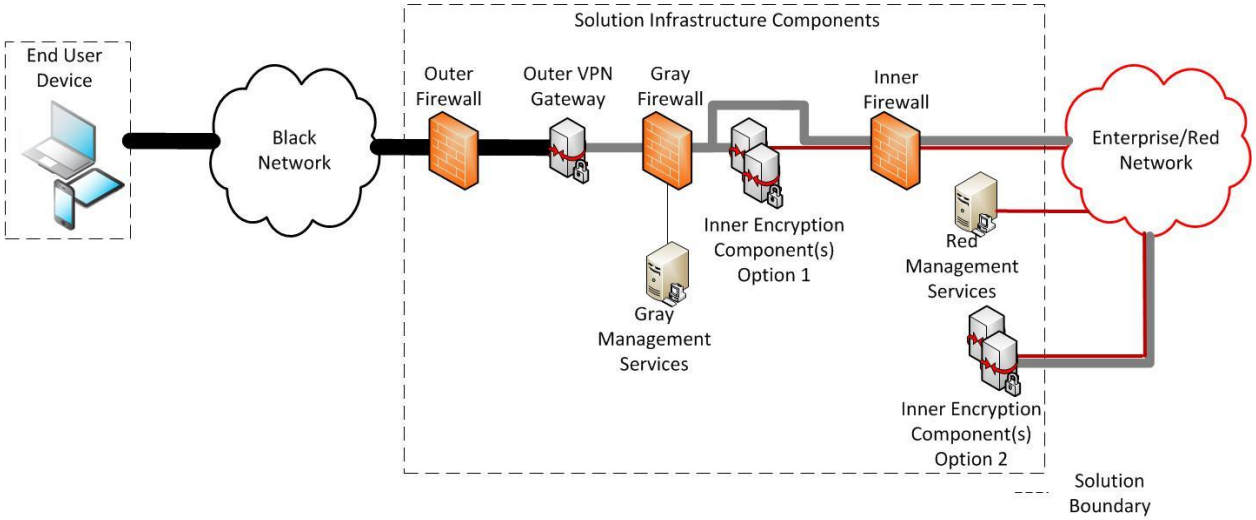


Figure 4. EUDs Connected to Independent Site

Figure 4 depicts a single Red network connected to EUDs that operate at the same security level, through the MA solution. Here, the Enterprise/Red network has at least two Encryption Components associated with it: One or more Inner Encryption Components connected to the Red network, and an Outer VPN Gateway between the Inner Encryption Components and the Black network. The Inner Encryption Component(s) can terminate either in the Grey network as depicted in Option 1 or within the Enterprise/Red network as depicted in Option 2. There are two layers of encryption between any EUD communicating with the Enterprise/Red network: one IPsec tunnel between their Outer VPN Components, and a second IPsec, TLS or SRTP tunnel depending on the selected EUD Design(s).

For independent sites, administration is performed at that site for all components within the Solution Boundary including the Outer VPN Gateway, Gray Management Services, Inner Encryption Components, Red Management Services, firewalls, and EUDs. Independent Sites are not interconnected with other Infrastructure Sites through the MA Solution. Therefore management, data plane, and control plane traffic between Solution Infrastructure sites are outside the scope of the MA CP. If two or more sites must be interconnected, customers may also register the MA solution against the VPN CP or utilize an NSA-Certified encryptor.

Note that while Figure 4 depicts only a single EUD, this solution does not preclude the use of a large number of EUDs being implemented.

4.2.3 MULTIPLE SITES

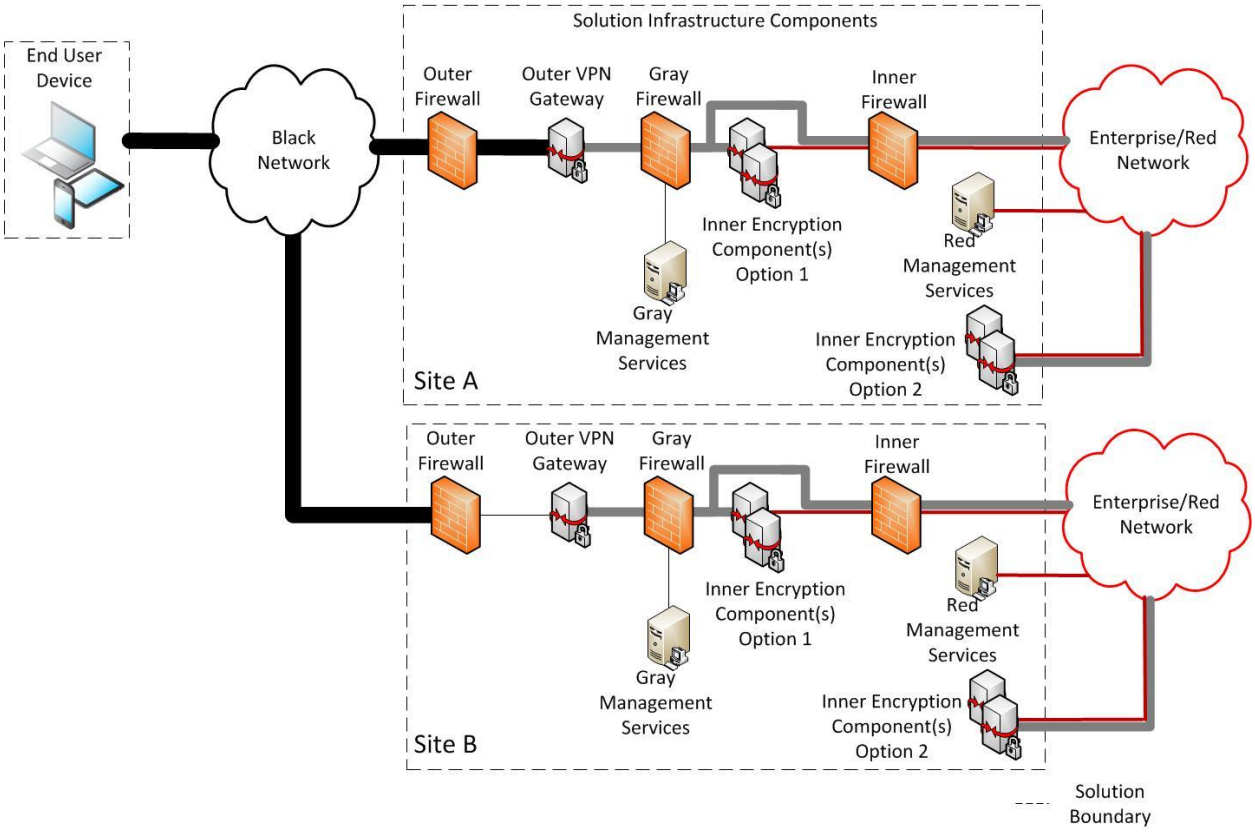


Figure 5. Multiple Mobile Access Solution Infrastructures supporting EUDs

Figure 5 depicts two MA Solution Infrastructures which an EUD can connect to in order to access Enterprise/Red network services. Customers may want to implement multiple Solution Infrastructures to support Continuity of Operations or provide better performance based on geographic location of EUDs or Enterprise/Red services. The Multiple Solution Infrastructures may be interconnected using a NSA-approved solution such as the Virtual Private Network (VPN) CP or a NSA-Certified encryptor; however, connectivity of Solution Infrastructure Components is outside the scope of the MA CP.

Note that, while Figure 5 depicts only two sites, this solution can scale to include numerous sites, with each additional site having the same design as those in the figure.

4.3 RATIONALE FOR LAYERED ENCRYPTION

A single layer of Suite B encryption, properly implemented, is sufficient to protect classified data in transit across an untrusted network. The MA solution uses two layers of Suite B encryption not because

of a deficiency in the cryptographic algorithms themselves, but rather to mitigate the risk that a failure in one of the Components, whether by accidental misconfiguration, operator error, or malicious exploitation of an implementation vulnerability results in exposure of classified information. The use of multiple layers of protection reduces the likelihood of any one vulnerability can be used to exploit the full solution, particularly if the layers exhibit suitable independence.

If an Outer VPN Component is compromised or fails in some way, the Inner Encryption Component can still provide sufficient encryption to prevent the immediate exposure of classified data to a Black network. In addition, the Gray firewall can indicate that a failure of the Outer VPN Gateway has occurred, since the filtering rules applied to its external network interface will drop and log the receipt of any packets not associated with an Inner Encryption Component. Such log messages indicate that the Outer VPN Gateway has been breached or misconfigured to permit traffic to pass through to the Inner Encryption Component that is not allowed.

Conversely, if instead the Inner Encryption Component is compromised or fails in some way, the Outer VPN Gateway can likewise provide sufficient encryption to prevent the immediate exposure of classified data to a Black network. As in the previous case, the Gray firewall filtering rules applied to its internal network interfaces will drop and log the receipt of any packets not associated with an Inner Encryption Component. Such log messages indicate that the Inner Gateway has been breached or misconfigured to permit traffic to pass through to the Outer VPN Gateway that is not allowed.

If both the Outer and Inner Gateways are both compromised or fail simultaneously, then it may be possible for classified data from the Enterprise/Red network to be sent to a Black network without an adequate level of encryption. The security of the MA solution depends on preventing this failure mode by promptly remediating any compromises or failures in one Encryption Component before the other also fails or is compromised.

Diversity of implementation is needed between the components in each layer of the solution in order to reduce the likelihood that both layers share a common vulnerability. The CSfC Program recognizes two ways to achieve this diversity. The first is to implement each layer using components produced by different manufacturers. The second is to use components from the same manufacturer, where that manufacturer has provided NSA with sufficient evidence that the implementations of the two components are independent of one another. The CSfC web page (http://www.nsa.gov/ia/programs/csfc_program) contains details for how a manufacturer can submit this evidence to NSA and what documentation must be provided. Customers wishing to implement a solution in accordance with this CP that uses products from the same manufacturer in both layers should contact their NSA/IAD Client Advocate to confirm that NSA has accepted the manufacturer's claims before implementing their solution.

4.4 AUTHENTICATION

The MA solution provides mutual device authentication between Outer VPN components and Inner Encryption components via public key certificates. This CP requires all authentication certificates issued to Outer VPN components and Inner Encryption components be Non-Person Entity (NPE) certificates, except in the case when TLS EUDs are implemented. Following the two layers of device authentication, VPN EUDs require the user to authenticate to the network before gaining access to any classified data. TLS EUDs typically utilize a user certificate to authenticate to TLS-Protected Server(s). In these instances, the user certificate authenticates the Inner layer of TLS encryption as well as authenticating the user for access to the requested classified data. Although not as common, TLS EUDs can be implemented using mutual device authentication between Outer and Inner Encryption components, similar to VPN EUDs. Following the two layers of device authentication, the user must then authenticate to the network before gaining access to any classified data.

In addition to authentication for the Outer and Inner layer of Encryption, the MA CP requires user-to-device authentication. This authentication occurs between the user and the Computing Device (which processes Red data) of an EUD. In some instances the computing device may be physically separate from the component of the EUD which provides the Outer layer of encryption (for example, a VPN Gateway provides the Outer Layer of encryption). The MA CP requires EUD components to utilize a minimum of a four-character, case–sensitive, alpha-numeric password to authenticate to the device. This password can be used both for decrypting the Platform Encryption as well as for unlocking the screen. EUD components, which are selected from the Mobile Platform section of the CSfC Components, are able to utilize a relatively short authentication factor since they are backed by a hardware root of trust which is evaluated during the NIAP certification.

The MA CP defines requirements for an Outer CA to issue NPE (i.e. device) authentication certificates to Outer VPN Components. The CP also defines requirements for an Inner CA to issue NPE authentication certificates to VPN EUDs and Inner VPN Gateways. TLS EUDs and TLS-Protected Servers primarily utilize User Authentication Certificates issued by an Inner CA. These User certificates are generally issued by the CA that is responsible for issuing enterprise server certificates, which may be different than the Inner CA described within this MA CP.

4.5 OTHER PROTOCOLS

Throughout this document, when IP traffic is discussed, it can refer to either IPv4 or IPv6 traffic, unless otherwise specified. In addition, Red, Gray and Black networks can run either IPv4 or IPv6 and each network is independent from the others in making that decision. In the remainder of the document, if no protocols or standards are specified then any appropriate protocols may be used to achieve the objective.

Public standards conformant Layer 2 control protocols such as ARP are allowed as necessary to ensure the operational usability of the network. This CP is agnostic with respect to Layer 2; specifically, it does not require Ethernet. Public standards conformant Layer 3 control protocols such as Internet Control Message Protocol (ICMP) may be allowed based on local AO policy, but the default configuration of this solution is for all Layer 3 control protocols to be disabled. Red and Gray network multicast messages and Internet Group Management Protocol (IGMP) or Multicast Listener Discovery (MLD) may also be allowed depending on local AO policy. Multicast messages received on external interfaces of the Outer VPN component shall be dropped.

It is expected that the MA solution can be implemented in such a way as to take advantage of standards-based routing protocols that are already being used in the network. For example, networks that currently use Generic Routing Encapsulation (GRE) or OSPF protocols can continue to use these in conjunction with this solution to provide routing as long as the AO approves their use.

4.6 AVAILABILITY

The high-level designs described in Section 4.2 are not designed with the intent of automatically providing high availability. Supporting solution implementations for which high availability is important is not a goal of this version of the CP. However, this CP does not prohibit adding redundant components in parallel to allow for component failover or to increase the throughput of the MA solution, as long as each redundant component adheres to the requirements of this CP. The CP does not limit the number of Outer VPN Gateways or Inner Encryption components which can be implemented for high availability in a MA Solution.

5 INFRASTRUCTURE COMPONENTS

In the high-level designs discussed in the previous section, all communications flowing across a Black network are protected by at least two layers of encryption, implemented using an outer IPsec VPN tunnel and an Inner layer of IPsec, TLS, or SRTP encryption. Mandatory aspects of the solution infrastructure also include Administration Workstations, IDS or IPS, Security Information Event Management (SIEM), firewalls, and CAs for key management using Public Key Infrastructure (PKI).

Each infrastructure component is described in more detail below. The descriptions include information about the security provided by the components as evidence for why they are deemed necessary for the solution. Components are selected from the CSfC Component List in accordance with the Product Selection requirements of this CP (see Section 12).

This section also provides details on additional components that can be added to the solution to help reduce the overall risk. However, where indicated in the text, these are not considered mandatory components for the security of the solution; therefore, this CP does not place configuration requirements on those optional components.

5.1 OUTER FIREWALL

Configuration and enforcement of network packet handling rules is fundamental to the security provided by the MA Solution Firewalls. The Outer firewall is located at the edge of the Mobile Access Solution Infrastructure and is connected to the Black Transport Network. The External Interface of the Outer firewall only permits IPsec IKE and ESP traffic with a destination address of the Outer VPN Gateway and AO-approved control plane traffic. The Internal Interface of the Outer firewall only permits IPsec traffic with a source address of the Outer VPN Gateway, AO-approved control plane traffic, and AO-approved Management Traffic. The minimum filtering required by the Outer firewall is primarily based on source IP addresses, source ports, destination IP addresses, and destination ports (see Section 13.9).

Although the Outer firewall is located on the perimeter of the network and thus more exposed to external attacks, the Outer firewall rules prohibit unauthorized data flows, which helps mitigate Denial of Service (DoS) attacks and resource exhaustion. This solution does not require that a single Outer firewall be utilized; however, all firewalls implemented as part of the Mobile Access Solution shall conform to the port filtering requirements in Section 13.9.

In addition to performing the functions described in this CP, the Outer firewall may also use AO-approved routing protocols on the Black network with the exception of prohibited control plane traffic as described in this CP. The Outer Firewall, selected from the CSFC Components List, must be physically separate from the Outer VPN Gateway as depicted in Figure 4.

5.2 OUTER VPN GATEWAY

Authentication of peer VPN Components, cryptographic protection of data in transit, and configuration and enforcement of network packet handling rules are all aspects fundamental to the security provided by VPN Gateways.

The external interface of the Outer VPN Gateway is connected to the internal interface of the Outer firewall. The VPN Gateway establishes an IPsec tunnel with peer Outer VPN Components, which provides device authentication, confidentiality, and integrity of information traversing Black networks. VPNs offer a decreased risk of exposure of information in transit since any information that traverses a Black network is placed in a secure tunnel that provides an authenticated and encrypted path between the site and an EUD. The Outer VPN Gateway is implemented identically for all the high-level designs, including when implemented as the Outer VPN Component for EUDs (see Section 6.1.1). Similar to the Outer firewall, the external interface of the Outer VPN Gateway only permits IPsec traffic and AO-approved control plane traffic. The internal interface of the Outer VPN Gateway is configured to only permit traffic with an IP address and port associated with Inner Encryption Components, Gray Management Services (i.e. SIEM and Administration Workstation), or Control Plane Component (i.e. DNS and NTP Servers in the Gray).

The Outer VPN Gateway is prohibited from implementing routing protocols on external and internal interfaces. The Outer VPN Gateway can rely on the Outer firewall and/or Gray firewall to perform routing functionality. The Outer VPN Gateway, selected from the CSFC Components List, must be physically separate from the Outer firewall and Gray firewall as depicted in Figure 4.

The Outer VPN Gateway cannot route packets between Gray and Black networks; any packets received on a Gray network interface and sent out a Black network interface must be transmitted within an IPsec VPN tunnel configured according to this CP.

For load balance or other performance reasons, multiple Outer VPN Gateways that comply with the requirements of the CP are acceptable.

5.3 GRAY FIREWALL

Gray firewalls are responsible for filtering traffic to only allow the proper traffic to flow to/from the Inner Encryption Components, Outer VPN Gateway, and Gray Management Services. Since this filtering is primarily based on the source and destination ports and addresses in the packet, the Gray networks themselves must use an addressing scheme that supports the necessary filtering (such as using separate address ranges for the Gray interfaces of Inner Encryption Components supporting different TLS-Protected Server devices). The Gray firewalls are configured to only allow the necessary type of traffic based on the type(s) of EUD Design implemented. For example, the Gray firewall will only allow valid IP addresses or address ranges to perform DNS queries to a Gray DNS server.

For load balance or other performance reasons, multiple Gray firewalls that comply with the requirements of the CP are acceptable. The Gray firewall, selected from the CSFC Components List, must be physically separate from the Outer VPN Gateway and Inner Encryption Components as depicted in Figure 4.

5.4 GRAY MANAGEMENT SERVICES

Secure administration of components in the Gray network and continuous monitoring of the Gray network are essential roles provided by the Gray Management Services. Gray Management Services are composed of a number of components which each can play a distinct role in the overall security of the solution. The MA CP allows flexibility in the placement of some Gray Management Services as described below. All components within the Gray Management Services are either connected directly to the Gray firewall or indirectly connected to the Gray firewall (i.e. multiple Gray Management Services connected to a switch which is connected to the Gray firewall). Finally, the Gray Management Services are physically protected as classified devices.

5.4.1 OUTER CERTIFICATE AUTHORITY (LOCATED ON GRAY NETWORK)

An Outer CA located on the Gray network issues digital certificates for the Outer VPN Components in the solution. These certificates are used for authentication in establishing the Outer IPsec tunnels between

pairs of VPN Components. If an Outer CA is located in the Gray network then the CP also requires a physically separate Inner CA located in the Enterprise/Red network. The Inner CA issues certificates to the Inner Encryption Components. This separation provides key management separation between the two independent layers of encryption.

To improve integration with existing Enterprise PKI and enable remote renewal of EUD Certificates, the MA CP allows for flexibility in the placement of Certificate Authorities. As an alternative to implementing the Outer CA as a Gray Management Service, customers can choose to implement the Outer CA on the Enterprise/Red network (See Section 9). When the Outer CA is located in the Enterprise/Red network it is not managed by the Gray Management Services.

5.4.2 Gray Administration Workstation

The Gray Administration Workstations is responsible for maintaining, monitoring, and controlling all security functionality for the Outer VPN Gateway, Gray firewall, and all Gray Management Service components. The Gray Administrative Workstations are not permitted to maintain, monitor, or control Inner Encryption Components or Red Management Services. All Mobile Access Solutions will have at least one Gray Administrative Workstation. Section 7 provides more detail on management of Mobile Access Solution components.

5.4.3 Gray Certificate Revocation Services

CRL Distribution Points (CDP) and Online Certificate Status Protocol (OCSP) Responders are servers other than a CA that makes revocation information available to components. Outer CDPs and OCSP Responders are deployed on the internal side of the Outer VPN Gateway for which revocation information is being made available. Collectively Outer CDPs and OCSP Responders are referred to as Gray Certificate Revocation Services. The Gray Certificate Revocation Services ensure the Outer VPN Gateway can verify the status of the certificates used by the Outer VPN Component of EUDs. The Gray Certificate Revocation Services also provide certificate revocation information to Gray Management Services.

CDPs and OCSP Responders are not required components of the Mobile Access CP, but if not utilized the organization must implement other means, such as whitelists, to ensure that once a certificate is revoked it cannot successfully establish an Outer IPsec Tunnel with the Solution Infrastructure. The Mobile Access CP also allows for an Inner CDP to be stood up in the Gray Management Services. Placing an Inner CDP in the Gray Management services allows EUDs to check the certificate status of the Inner Encryption component prior to establishing a tunnel. To utilize an Inner CDP in the Gray Management Services, an AO must determine that CRLs generated by the Inner CA are unclassified. Additionally, these CRLs must be moved from the Enterprise/Red network to the Gray Management Services using an AO approved method (e.g. Cross Domain Solution).

The use of CDPs in a Mobile Access Solution is discussed in detail in Section 9.1.

5.4.4 Gray Security Information and Event Management (SIEM)

The Gray SIEM collects and analyzes log data from the Outer VPN Gateway, Gray firewall, and other Gray Management Service components. Log data may be encrypted between the originating component and the Gray SIEM with SSHv2, TLS, or IPsec to maintain confidentiality and integrity of the log data. At a minimum, an auditor reviews the Gray SIEM on a weekly basis. The SIEM is configured to provide alerts for specific events including if the Outer VPN Gateway or Gray firewall receive and drop any unexpected traffic which could indicate a compromise of the Outer firewall or Outer VPN Gateway respectively. These functions can also be performed on an Enterprise/Red SIEM if AO-approved one-way taps are utilized as described in this CP (see Section 8.2).

5.5 Inner Encryption Components

The MA CP allows for the use of up to three different types of Inner Encryption Components: Inner VPN Gateway, Inner TLS-Protected Server, or Inner SRTP Endpoint. Inner VPN Gateways are always located between the Gray firewall and Inner firewall. An Inner VPN Gateway will always have at least two interfaces. One external interface connected to the Gray firewall and one internal interface connected to the Inner firewall.

Inner TLS-Protected Servers and Inner SRTP Endpoints are permitted to use a single interface or multiple interfaces. Similar to the Inner VPN Gateway, Inner TLS-Protected Servers and SRTP Endpoints with multiple interfaces can have one external interface connect to the Gray firewall and one internal interface connected to the Inner firewall. If implemented with a single interface, then that interface establishes the Inner layer of encryption and provides the classified data to the TLS EUD. An example of a TLS-Protected Server with a single interface is a web server on the Enterprise/Red network which terminates the Inner layer of encryption with HTTPS and directly returns the content to the TLS EUD. An example of a SRTP Endpoint with a single interface is a VoIP Desktop Phone in the Enterprise Network which terminates the SRTP encryption and sends the voice traffic from the VoIP Desktop Phone. The TLS-Protected Servers and SRTP Endpoints can be placed either between the Gray firewall and Inner firewall or within the Enterprise/Red network.

A MA Solution Infrastructure may support both TLS EUDs and VPN EUDs. When supporting both TLS EUDs and VPN EUDs the solution infrastructure will always include an Inner VPN Gateway between the Gray firewall and Inner firewall. This Inner VPN Gateway will terminate the Inner layer of IPsec traffic for all VPN EUDs. Additionally, the solution infrastructure will include one or more TLS-Protected Servers. The TLS-Protected Servers are either placed between the Gray firewall and Inner firewall or on the internal side of the Inner firewall. When placed between the Gray firewall and Inner firewall the TLS-Protected Server(s) must be placed in parallel with the Inner VPN Gateway such that the TLS-Protected Server is not dependent on the Inner-VPN Gateway to reach the Gray firewall or Inner firewall. When the solution includes a TLS-Protected Server located in the Enterprise/Red network the Gray firewall includes a physical path directly to the Inner firewall. Proper configuration of the Gray

firewall and Inner firewall ACL is critically important when supporting TLS-Protected Servers in the Enterprise. The Gray firewall must ensure that only packets destined for the IP address of a TLS-Protected Server within the Solution Boundary are permitted to flow to the Inner firewall. Similarly, the Inner firewall must ensure that only packets destined for the IP address of a TLS-Protected Server within the Solution Boundary are permitted to the Enterprise/Red network (see Appendix D).

For load balance or other performance reasons, multiple Inner Encryption Components that comply with the requirements of the CP are acceptable.

5.5.1 INNER VPN GATEWAY

Similar to the Outer VPN Gateway, the Inner VPN Gateway provides authentication of peer VPN Components, cryptographic protection of data in transit, and configuration and enforcement of network packet handling rules. The Inner VPN Gateway is located between the Gray firewall and the Inner firewall. The Inner VPN Gateway is required to be implemented if supporting VPN EUDs.

The External Interface of the Inner VPN Gateway is connected to the internal interface of the Gray firewall. The VPN Gateway establishes an IPsec tunnel with peer Inner VPN Components, which provides a second layer of device authentication, confidentiality, and integrity of information traversing Black networks. Similar to the Outer VPN Gateway, the external interface of the Inner VPN Gateway only permits the egress of IPsec traffic and AO-approved control plane traffic. The internal interface of the Inner VPN Gateway is configured to only permit traffic with an IP address and port associated with Red network services.

The Inner VPN Gateway cannot route packets between Red and Gray networks; any packets received on a Red network interface and sent to a Gray network interface must be transmitted within an IPsec VPN tunnel configured according to this CP. The Inner VPN Gateway, selected from the CSFC Components List, must be physically separate from the Gray firewall and Inner firewall as depicted in Figure 4.

5.5.2 INNER TLS-PROTECTED SERVER

Inner TLS-Protected Server(s) utilizes TLS to provide confidentiality, integrity, and mutual authentication between an EUD and TLS-Protected Server. The TLS-Protected Server is either located between the Gray firewall and the Inner firewall or within an Enterprise/Red network. Inner TLS-Protected Servers are required to be implemented if supporting TLS EUDs.

This CP allows the TLS-Protected Server to utilize any protocol that is encapsulated in TLS. TLS-Protected Servers are selected from the CSfC Components List. All TLS-Protected Servers which terminate the Inner layer of encryption originating from a TLS EUD reside within the CSfC Solution Boundary, and therefore must meet all applicable requirements of the MA CP. This is true both of TLS-Protected Servers which are implemented specifically for the CSfC Solution as well as Enterprise/Red TLS-Protected Servers which terminate the inner tunnel originating from EUDs.

Examples of TLS-Protected Servers include, but are not limited to, Web Servers, SIP Servers, and Mobile Device Management (MDM) Servers. Web Servers implemented as part of the MA CP terminate the Inner layer of encryption utilizing HTTPS. SIP Servers utilize SIP over TLS for registration of EUDs and SRTP Endpoints, session setup, and session termination. When SIP Servers are included Session Description Protocol Security Descriptions (SDES) is used over the SIP TLS session for key exchange between TLS EUDs or between a TLS EUD and a SRTP Endpoint. The Inner TLS Protected-Server, selected from the CSFC Components List, must be physically separate from the Gray firewall and Inner firewall as depicted in Figure 4.

5.5.3 INNER SRTP ENDPOINT

Inner SRTP Endpoints provides cryptographic protection of data in transit. Within the MA Solution Infrastructure, SRTP Endpoints are either located between the Gray firewall and the Inner firewall or within an Enterprise Network. The Inner Layer of SRTP Encryption can also be terminated between two EUDs (see Section 6.3). Registration, session setup (including authentication and Key Exchange), and session termination for the SRTP Endpoints is performed utilizing SIP over TLS. Inner SRTP Endpoints are required to be implemented if supporting TLS EUDs that utilize SRTP.

All SRTP Endpoints which terminate the Inner layer of encryption originating from a TLS EUD reside within the CSfC Solution Boundary, and therefore must meet all applicable requirements of the MA CP. This is true for both of the SRTP Endpoints, which are implemented specifically for the CSfC Solution as well as Enterprise/Red SRTP Endpoints which terminate the inner tunnel.

Examples of Infrastructure SRTP Endpoints include VoIP Desktop Phones and VoIP Gateway/Border Controller (also known as Session Border Controllers). VoIP Desktop Phones generally reside within existing Enterprise Networks on the internal side of the Inner firewall. These Desktop Phones can provide end-to-end encryption of voice and/or video to TLS EUDs. VoIP Gateway/Border Controller terminates SRTP Traffic from a TLS EUD, and relay the data to the Enterprise/Red network. When a VoIP Gateway/Border Controller terminates the Inner layer of SRTP, Desktop Phones are not included in the Solution Boundary. Inclusion of a VoIP Gateway/Border Controller allows integration with existing enterprise voice systems.

The Inner SRTP Endpoint, selected from the CSFC Components List, must be physically separate from the Gray firewall and Inner firewall as depicted in Figure 4.

5.5.4 INNER FIREWALL

The Inner firewall is responsible for filtering to only allow the proper traffic to flow to/from the Enterprise/Red network. Since this filtering is primarily based on source and destination ports and addresses, the Inner Encryption components should use an addressing scheme that supports the necessary filtering. When the IInner layer of encryption is terminated by a TLS-Protected Server or SRTP Endpoint in the Enterprise/Red network, the Inner firewall must implement an ACL. The Inner firewall ACL is critical to ensure that only TLS-Protected Servers and SRTP Endpoints, which are configured in

accordance with this CP can terminate the Inner layer of encryption. A similar requirement is not placed on the Inner firewall when the the Inner layer of encryption terminates on an Inner VPN Gateway or TLS-Protected Server located between the Gray Firewall and Inner Firewall.

For load balance or other performance reasons, multiple Inner firewalls that comply with the requirements of the CP are acceptable.

5.6 RED MANAGEMENT SERVICES

Secure Administration of Inner Encryption Components and continuous monitoring of the Red network are essential roles provided by the Red Management Services. Red Management Services are composed of a number of components which can each play a distinct role in the overall security of the solution. The Mobile Access CP allows flexibility in the placement of some Red Management Services as described below.

5.6.1 CERTIFICATE AUTHORITIES (LOCATED ON RED NETWORK)

Mobile Access CP Solutions will always have at least one CA located in the Red network. At a minimum an Inner CA is included in the Red network to issue digital certificates for the Inner Encryption Components in the solution. These certificates are used for authentication in establishing the inner tunnel of encryption between pairs of Inner Encryption Components. The MA CP also allows Outer CAs to be included in the Red network to issue digital certificates for the Outer VPN Components. These certificates are used for device authentication in establishing the outer tunnel of encryption between pairs of VPN Components. When an Outer Tunnel CA is placed in the Red network it is critical to have AO-approved mechanisms in place to transfer revocation information to the Gray network in order to ensure it is accessible to the Outer VPN Gateway and Gray Management Services. When an Enterprise PKI capability is utilized, it is managed with existing processes and capabilities. Enterprise CAs then provide certificate management services for the MA solution over the Enterprise/Red network.

5.6.2 RED ADMINISTRATION WORKSTATIONS

The Red Administration Workstation is responsible for maintaining, monitoring, and controlling all security functionality for the Inner Encryption Components, Inner firewall, and all Red Management Service components. The Red Administrative Workstations are not permitted to maintain, monitor, or control Outer Encryption Components or Gray Management Services. All MA Solutions will have at least one Red Administrative Workstation. Section 7 provides more detail on management of Mobile Access Solution components.

5.6.3 RED SECURITY INFORMATION AND EVENT MANAGEMENT (SIEM)

Red SIEMs collect and analyze log data from the Inner Encryption Components, the Inner firewall, and other Red Management Service components. Log data may be encrypted between the originating component and the Red SIEM with SSHv2, TLS, or IPsec to ensure confidentiality and integrity. At a minimum an auditor reviews the Red SIEM on a weekly basis. The SIEM is configured to provide alerts

for specific events including if the Inner Encryption Components or Inner firewall receive and drop any unexpected traffic which could indicate a compromise of the Gray firewall, Inner Encryption Components, or Inner firewall. While Red SIEMs are not a mandatory component of the MA Solution, customers are encouraged to leverage existing Enterprise SIEM capabilities to monitor log data from Inner Encryption Components, the Inner firewall, and Red Management Services. Although a Red SIEM is not required, it is still required to analyze logs from all Inner Encryption components on at least a weekly basis. A Red SIEM may also be utilized to analyze log data from Gray network components when utilized in conjunction with one-way taps as described in this CP (see Section 8.2).

5.6.4 RED CERTIFICATE REVOCATION SERVICES

CRL Distribution Points (CDP) and Online Certificate Status Protocol (OCSP) Responders are servers other than a CA that make revocation information available to components. Red CDPs and OCSP Responders are deployed either between the Inner Encryption Component and Inner firewall or on the internal side of the Inner firewall. Red CDPs and OCSP Responders make revocation information available to Inner Encryption Components of the Solution Infrastructure. Collectively Red CDPs and OCSP Responders are referred to as Red Certificate Revocation Services. The Red Certificate Revocation Services ensure the Solution Infrastructure Inner Encryption Components can verify the status of the certificates used by the Inner Encryption Components of EUDs.

CDPs and OCSP Responders are not required components of the Mobile Access CP, but if not utilized the organization must implement other means, such as whitelists, to ensure that, once a certificate is revoked, it cannot successfully establish an Inner Tunnel with the Solution Infrastructure.

6 END USER DEVICE COMPONENTS

The MA CP supports the use of two EUDs: VPN EUDs and TLS EUDs. The MA Solution Infrastructure can support both types of EUDs; however, the EUD must be dedicated as either a VPN EUD or TLS EUD. VPN and TLS EUDs are composed of a computing device and optionally include a physically separate VPN Gateway to provide the Outer Layer of IPsec encryption. When VPN Gateways are included as part of the EUD, they must be physically connected (i.e. Ethernet cable) to the computing device (see Figure 2).

A RD to connect to the Black network is also a required component, except for the solution designs and use cases specified in section 4.1.3. and 6.1.1.

6.1 OUTER VPN COMPONENT

The allowable Outer VPN Components for both the VPN and TLS EUD are identical. Authentication of peer VPN Components and cryptographic protection of data in transit are fundamental aspects to the security provided by the EUD Outer VPN Component Gateways.

The Outer VPN Component establishes an IPsec tunnel with the Solution Infrastructure Outer VPN Gateway, which provides device authentication, confidentiality and integrity of information traversing

Black networks. VPNs offer a decreased risk of exposure of information in transit since any information that traverses a Black network is placed in a secure tunnel that provides an authenticated and encrypted path between the EUD and the Mobile Access Solution Infrastructure. The Mobile Access CP allows the use of VPN Gateways or VPN Clients to be utilized as the Outer VPN Component of EUDs.

The private keys and certificates utilized for the authentication of the Outer VPN Component are considered Controlled Unclassified Information (CUI) and must be protected with at least FIPS 140-2-validated encryption. Customers deploying Mobile Access Solutions in high-threat environments may also choose to implement controls to mitigate against tampering attacks.

6.1.1 Outer VPN Gateway

An Outer VPN Gateway can be utilized as the Outer VPN Component for EUDs. Utilizing a physically separate VPN Gateway as part of the EUD improves security by providing physical separation between the computing device and the Outer layer of encryption. When an Outer VPN Gateway is used as part of an EUD, there is no requirement to utilize a Government RD (see Appendix D). The VPN Gateway included as part of the EUD must be physically connected to the component which provides the Inner layer of encryption. VPN Gateways included as part of a EUD are selected from the *IPsec VPN Gateway* section of the CSfC Components List.

When an Outer VPN Gateway is included as part of an EUD, it provides configuration and enforcement of network packet handling rules for the Outer Layer of Encryption. The configuration settings of the Outer VPN Gateway may need to be updated when entering new environments (e.g. updating the Default Gateway).

6.1.2 Outer VPN Client

An Outer VPN Client can be utilized as the Outer VPN Component for MA EUDs. The purpose of the Outer VPN Client is to establish an IPsec tunnel to the Outer VPN Gateway of the MA Solution Infrastructure. The tunnel can be configured to automatically be established as part of the EUD's power-on process. A combination of the VPN Client and the Operating System on which it is installed is responsible for providing configuration and enforcement of network packet handling rules for the Outer layer of encryption. The Outer VPN Client is selected from the *IPSec VPN Client* section of the CSfC Components list. The VPN Client is installed on the computing device selected from the *Mobile Platform* section of the CSfC Components List.

6.2 Virtual private Network (VPN) End User Device (EUD)

VPN EUDs utilize IPsec using a VPN Client to provide the Inner layer of encryption. The purpose of the Inner VPN Client is to establish an IPsec tunnel to the Inner VPN Gateway of the MA Solution Infrastructure. The tunnel can be configured to automatically be established as part of the EUD's power-on process, following establishment of the Outer VPN tunnel. Once the Inner VPN Client establishes the

inner IPsec tunnel, any application installed on the computing device can send and receive classified data with the Enterprise/Red network.

The private keys and certificates utilized for the authentication of the Inner VPN Component are considered CUI and must be at a minimum be protected by enabling the native platform DAR protection. Customers deploying Mobile Access Solutions in high-threat environments may also choose to implement controls to mitigate against tampering attacks.

Appendix D provides more detail on the allowable configuration of VPN EUDs.

6.2.1 Inner VPN Client on Computing Device

A VPN Client can be utilized as the Inner VPN Component for VPN EUDs. The purpose of the Inner VPN Client is to establish an IPsec tunnel to the Inner VPN Gateway of the MA Solution Infrastructure. The tunnel can be configured to automatically be established as part of the EUD's power-on process. A combination of the VPN Client and the Operating System on which it is installed is responsible for providing configuration and enforcement of network packet handling rules for the Inner layer of encryption. The Inner VPN Client is selected from the *IPSec VPN Client* section of the CSfC Components list. The VPN Client is installed on the computing device selected from the *Mobile Platform* section of the CSfC Components List.

Virtualization must be utilized when an Outer VPN Client and Inner VPN Client both reside on the same computing device. The virtualization ensures that two separate IP stacks are utilized. The MA CP allows for Type 1 or Type 2 Hypervisors to be utilized to provide logically separated operating systems, each with their own IP stack.

6.3 Transport Layer Security (TLS) End User Device (EUD)

TLS EUDs utilize TLS clients or SRTP clients to provide the Inner layer of encryption. The Inner layer of TLS or SRTP is implemented by TLS clients and SRTP clients provided by individual applications installed on the computing device. Each application which sends and receives data to the Enterprise/Red network must be selected and configured in accordance with the requirements of the CP. Each application then terminates the Inner layer of encryption to TLS-Protected Servers and SRTP Endpoints within the MA Solution Infrastructure.

The private keys and certificates utilized for user authentication of the Inner TLS and SRTP clients are determined by the AO. If the private keys and certificates are considered CUI then the EUD component must, at a minimum, implement the native platform encryption. If the private keys and certificates are considered to be classified, then the EUD must be treated as classified at all times or implement a NSA-Approved DAR Solution (see Section 4.2.1). Customers deploying MA Solutions in high threat environments may also choose to implement controls to mitigate against tampering attacks.

TLS EUDs must use either a Government RD or Outer VPN Gateway to connect to the Black network, except for the use cases defined in Section 4.1.3. Appendix D provides more detail on the allowable configuration of TLS EUDs.

6.3.1 TLS CLIENT

Applications with a TLS client can be installed on the computing device and utilized for the Inner layer of TLS encryption. On TLS EUDs every application which sends or receives data through the Outer VPN Component must be configured in accordance with the requirements of this CP. For example, if a Voice Application, Web Browser, MDM Agent, and Email Client are installed on the computing device each application is configured to establish a TLS session to the TLS-Protected Server in the MA Solution Infrastructure. In some instances an application may perform both TLS and SRTP encryption. Those applications must be configured to meet requirements for both TLS clients and SRTP clients.

TLS clients utilize certificates for mutual authenticate with the TLS-Protected Server. In most cases, the TLS-client utilizes a user certificate for authentication to the TLS-Protected Server. User certificates are issued by the same PKI that issues certificates to TLS Protected Servers (e.g. customer enterprise PKI), which may be different than the Inner CA. Alternatively, the TLS client can utilize a device certificate for authentication followed by user authentication (e.g. username and password, token, smartcard, etc). A combination of the TLS Client and Computing Device Operating System is responsible for providing configuration and enforcement of network packet handling rules for the Inner layer of encryption.

6.3.2 SRTP CLIENT

Applications with a SRTP client can be installed on the computing device and utilized for the inner layer of SRTP encryption. If multiple SRTP clients are installed on the TLS EUD, then each must be configured in accordance with the requirements of this CP. SRTP Clients are generally used to encrypt real time traffic, such as voice or video. In some instances, an application may perform both TLS and SRTP encryption. Those applications must be configured to meet requirements for both TLS clients and SRTP clients.

SRTP clients utilize certificates for mutual authentication. In most cases, the SRTP-client utilizes a user certificate for authentication. User certificates are issued by the same PKI that issues certificates to TLS Protected Servers (e.g. customer enterprise PKI), which may be different than the Inner CA. Alternatively, the SRTP client can utilize a device certificate for authentication followed by user authentication (e.g. username and password, token, smartcard, etc.). A combination of the SRTP Client and Computing Device Operating System is responsible for providing configuration and enforcement of network packet handling rules for the Inner layer of encryption.

The Inner layer of SRTP encryption can terminate either at a SRTP Endpoint within the MASolution Infrastructure or on another EUD. In either case, all SRTP encrypted traffic must go through the Outer VPN Component prior to being transmitted on the Black network.

7 MOBILE ACCESS CONFIGURATION AND MANAGEMENT

The Mobile Access CP includes design details for the provisioning and management of Solution Components. The following sections describe the design in detail and Section 13 articulates the specific requirements which must be met to comply with the MA CP.

7.1 SOLUTION INFRASTRUCTURE COMPONENT PROVISIONING

Provisioning is an out-of-band process performed in a physically secured area (e.g. the enterprise within the Red network), through which Mobile Access Solution Infrastructure Components are configured and initialized before their first use. During the provisioning process, the Security Administrator configures the Outer VPN Gateway, Gray Management Services, Inner Encryption Components, and Red Management Services in accordance with the requirements of this CP. During provisioning, the Solution Infrastructure Outer VPN Gateways and Inner Encryption components generate a public/private key pair and output the public key in a Certificate Signing Request (CSR). The Security Administrator delivers the Outer VPN Gateways CSR to the Outer CA and the Inner Encryption Components CSR to the Inner CA. The appropriate CA processes the CSR for each encryption component and returns a signed X.509v3 certificate for each encryption component. The Security Administrator then installs the unique signed certificate and the certificate of the Trust Anchor.

7.2 EUD PROVISIONING

Provisioning is an out-of-band process performed in a physically secured area (e.g. the enterprise within the Red network), through which MA EUDs are initialized before their first use. During the provisioning process, the Security Administrator loads and configures the required software for the EUD. The Security Administrator also instructs the EUD to generate the requisite public/private key pairs for the EUD's Outer VPN Component and Inner Encryption Component, and to output the public keys in a specified CSR format for delivery to the Outer CA and Inner CA, respectively. For TLS EUDs that require an enterprise user certificate, the CSR is delivered to the CA in the customer's organization that has the authority to issue enterprise user certificates. This CA may not be the same as the Inner CA. If the EUD cannot generate its own key pairs or CSRs, then a dedicated management workstation is required to generate the key pairs for the EUD and construct the CSRs for delivery to the Outer CA and Inner CA. The CAs process the CSRs and return signed certificates to the Security Administrator, who installs the certificates onto the EUD. If required, the Security Administrator also installs the private keys onto the EUD. The Security Administrator then finalizes the security configuration of the EUD before it used for the first time.

If the MA solution owner is unable to remotely manage EUDs, the EUDs must be periodically re-provisioned in order to receive software and configuration updates. Re-provisioning consists of revoking the EUD's existing certificates and provisioning the EUD using a trusted baseline configuration that does not make use of any retained data originally stored on the EUD (e.g. factory reset and provision as a new device). This CP does not impose a particular frequency with which re-provisioning must take place,

although without remote management of EUDs, re-provisioning is the only means of applying security-critical patches to EUDs in accordance with requirement MA-GD-26.

Due to the time and effort needed to re-provision EUDs, it is preferable to remotely manage them when possible. With remote management capabilities, updated software and configuration data can be provided from a central management site through the MA solution to the EUD, after the EUD establishes the two MA solution tunnels (see Section 13.11).

7.3 Administration of Mobile Access Components

Each component in the solution has one or more Administration Workstations that are responsible for maintaining, monitoring, and controlling all security functions for that component. Throughout this document, these workstations are referred to as Administration Workstations. It should be understood that all of the required administrative functionality does not need to be present in each individual Workstation, but the entire set of Administration Workstations must collectively meet administrative functionality requirements.

The Administration Workstation is used for configuration review and management. Implementations will employ an SIEM in the Gray Management Services for log management of Gray Infrastructure Components except where AOs utilize one-way taps to move Gray network log data to a Red SIEM. Given the architecture of the solution, each layer has its own distinct administration LAN or VLAN: the Inner Encryption Components are managed from the Red Management Services, and the Outer VPN Gateway and supporting components are managed from the Gray Management Services. The Gray Administration Workstation, along with all Gray Management Services, is physically connected to the Gray firewall. In this manner, the Gray firewall maintains separate ACLs to permit Management Traffic to/from Gray Management Services, but prohibits such traffic from all other components. This architecture provides the separation necessary for two independent layers of protection.

Administration Workstations will be dedicated for the purposes given in the CP. For example, Administration Workstations are not to be used as the registration authority for the CA, a SIEM, or as a general user workstation for performing any functions besides management of the solution. Administration Workstations cannot be used as an enrollment workstation or provisioning workstation.

Management of all Mobile Access Solution components is always encrypted to protect confidentially and integrity, except in the case where components are locally managed through a direct physical connection (e.g. serial cable from Gray Administration Workstation to Outer VPN Gateway). Management traffic must be encrypted with SSHv2, TLS, or IPsec. Additionally, when managing components over the Black Network the SSHv2, TLS, or IPsec protocols must be configured utilizing the appropriate Suite B Algorithms. A similar requirement is not imposed if the Mobile Access Solution Infrastructure components are being managed from the same LAN or VLAN. For example, a Gray Administration Workstation residing in the Gray Management Services at the same site as the Outer VPN Gateway need not use Suite B algorithms since this traffic does not traverse an untrusted network.

In most cases, Mobile Platforms are managed over the Black network utilizing the Outer layer of IPsec and a MDM Server selected from the CSfC Components List. When an MDM Server is used to manage TLS EUDs, the MDM Server is considered a TLS-Protected Server and the MDM agent is considered a TLS Client. As a result, the MDM Server must be placed either at Inner Encryption Component Option 1 or Inner Encryption Component Option 2 as denoted in Figure 4Figure 3. As a TLS-Protected Server, the MDM Server must be configured to establish a session with the MDM Agent in accordance with the requirements in Table 14. Although not mandatory, the use of an MDM enables organizations to dynamically change policies enforced on the phone, allowing more flexibility then requiring phones to be re-provisioned. Additionally, there are several security advantages by using an MDM including the ability to perform a remote wipe of the EUD.

7.4 EUDs for Differing Classification Domains

As specified in this CP, an EUD is only authorized to communicate to Enterprise/Red networks operating at the same classification level. However, it does not preclude the possibility that an approved Cross Domain Solution (CDS) can be used within an infrastructure to provide cross domain transfer of data between EUDs operating at differing classification levels. Neither does it preclude the use of an EUD as an access CDS for multiple enclaves operating at differing classification levels if approved through the appropriate CDS approval process.

The requirements for a CDS capable of providing separation between enclaves of two or more classification levels are outside the scope of this CP. If developing a Mobile Access solution with a CDS capability, the solution must register against this CP and utilize the appropriate CDS approval processes.

8 CONTINUOUS MONITORING

The Mobile Access CP allows customers to utilize EUDs from physical environments which do not necessarily reside within a government secure facility. With the increase in accessibility also comes a need to continuously monitor network traffic and system log data within the solution infrastructure. This monitoring allows customers to detect, react, and report to any attacks which occur on their solution. This continuous monitoring also enables the detection of any configuration errors to Solution Infrastructure Components.

At a minimum, the CP requires an Auditor to review alerts, events, and logs on a weekly basis. This minimum review period allows customers in Tactical Environments to implement solutions where it may not be feasible to perform real-time monitoring. Operational and Strategic implementations of the Mobile Access CP should review alerts, events, and logs on a much more frequent period and in many cases may leverage Operations Centers to perform 24/7 monitoring of the solution.

8.1 Monitoring Network Traffic

The MA CP requires monitoring network traffic in at least one of three areas within the Solution Infrastructure. Network traffic can be monitored using an IDS; however, it is preferable to utilize an IPS

to enable real-time responses. While it is only required to monitor one of the three locations, customers monitoring all three points have the best visibility enabling detection of malicious activity or misconfiguration of components.

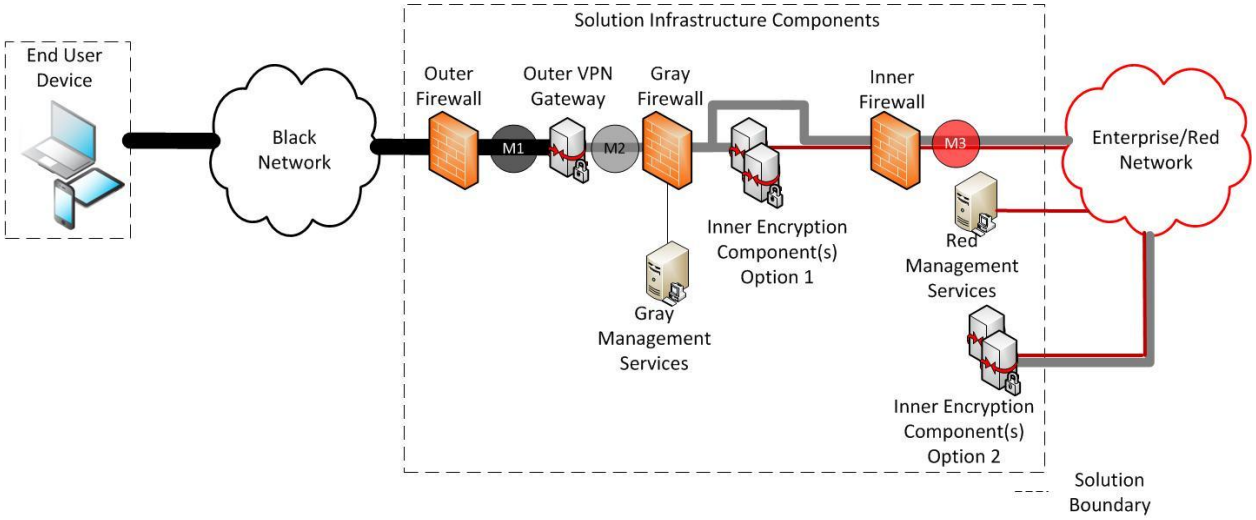


Figure 6. MA Solution Continuous Monitoring Points

Figure 6 depicts the three locations that customers can select to implement network monitoring capabilities. There are several alternatives for deploying the IDS at one or all of the Monitoring Points (M1, M2, and/or M3). Intrusion Detection Systems can ingest traffic from network taps, span ports, or in line with the solution. Similarly, an IPS can be placed either inline or in promiscuous mode. When operating in promiscuous mode, the IPS analyzes traffic at M1, M2, or M3 and issues commands to the Outer firewall, Gray firewall, or Inner firewall to block traffic flows. The following paragraphs define each of the three Monitoring Points. These descriptions detail the analysis and alerts which would be generated by the IDS. If a customer decides to implement an IPS, then it should be configured to block that traffic flow and also send an alert.

Monitoring Point 1 (M1): Located between the Outer firewall and Outer VPN Gateway. At a minimum, a M1 IDS is configured to send alert upon detection of any traffic which should have been blocked by the Outer firewall or Outer VPN Gateway. These alerts indicate a failure of the filtering or encryption functions by one of the two components and are either indicative of an improper configuration or a potential compromise. Normal traffic at M1 is very well-defined (IPsec and limited number of AO approved control plane protocols) and, as a result, will generate very few false positives. Since nearly all traffic traversing M1 is encrypted with IPsec, the IDS is limited to analyzing IP addresses, ports, protocols, and net flow data. Management of a M1 IDS occurs from a Black network.

Monitoring Point 2 (M2): Located between the Outer VPN Gateway and Gray firewall. At a minimum, a M2 IDS is configured to send an alert upon detection of any traffic which should have been blocked by the Outer VPN Gateway or Gray firewall. These alerts indicate a failure of the filtering functions by one

of the two components and are either indicative of an improper configuration or a potential compromise. Normal traffic at M2 is not as well defined as that of M1, but includes IPsec, data plane traffic encrypted with TLS or SRTP, control plane traffic, and management traffic. It is important that customers understand what ports and protocols are utilized by each of the Inner Encryption Components to ensure that firewall filters are properly configured and Intrusion Detection Systems are well tuned. Nearly all traffic traversing M2 is encrypted either with IPsec, TLS, SRTP, or SSHv2 which prevents the ability to perform deep packet inspection. Management of a M2 IDS occurs from the Gray Management Services.

Monitoring Point 3 (M3): Located between the Inner firewall and Enterprise/Red network. At a minimum, a M3 IDS is configured to send an alert upon detection of any traffic which should have been blocked by the Inner firewall. These alerts indicate a failure of the filtering function of the Inner firewall. Of the three monitoring points, M3 is the most difficult to define a normal baseline, but in many implementations, monitoring at M3 allows for deep packet inspection since traffic is not necessarily encrypted. If the Mobile Access Solution Infrastructure supports TLS EUDs that terminate within the Enterprise/Red network, then data traversing M3 will still have a single layer of encryption. Management of the M3 IDS occurs from the Red Management Services.

Monitoring Multiple Points: Although the MA CP only requires monitoring of one out of the three points, customers are encouraged to monitor all three locations. To ensure that no bypass of the Outer or Inner Encryption Components occurs, customers can implement three IDS(s)/IPS(s) located in the Black, Gray, and Red networks. Implementation of three separate components to monitor each point ensures that malicious traffic cannot inadvertently be transferred to the Enterprise/Red network, but it also increases cost and complexity for managing the solution. This approach also prevents correlation of data as it traverses throughout the solution.

Alternatively, in order to allow correlation of data at multiple points the MA CP allows a one-way tap located at M2 to feed the IDS located in an enclave of the Red Network that is isolated from the Enterprise/Red network. The one-way tap utilized must physically only allow data flow from M2 to the Enterprise/Red network. Additionally, the selected one-way tap must be approved by the AO. Similarly, the CP also permits an AO-approved one-way tap located at M1 to feed the IDS located in the Enterprise/Red network. Movement of network traffic from M3 to the Gray or Black network is explicitly prohibited. Additionally, movement of Network Traffic from M2 to the Black network is explicitly prohibited. Prior to implementing one-way taps as part of the Mobile Access Solution, AOs need to be aware of the Residual Risks associated with transferring data from Black and Gray networks to a Red network of a higher classification.

8.2 Monitoring Log Data

The Mobile Access CP requires the implementation of a SIEM component within the Gray Management Services (except in instances where an AO utilizes one-way taps to feed Gray Log data to a Red SIEM).

The Gray SIEM collects, aggregates, correlates, and analyzes log data from Gray Management Components. The SIEM also provides alerts to auditors when anomalous behavior is detected.

The Gray SIEM collects logs from the Outer VPN Gateway, Gray firewall, and any components located within the Gray Management Services. The Gray SIEM is not permitted to collect logs from the Outer firewall, Inner Encryption Components, or Inner firewall unless utilized in conjunction with AO-approved one-way taps. To protect the integrity of the data, all logs sent to the SIEM should be encrypted with TLS, SSHv2, or IPsec. At a minimum the SIEM is configured to send alerts upon receiving a log entry for blocked packets from the Outer VPN Gateway or Gray firewall. Packets blocked by the Outer VPN Gateway either indicate a failure of the Outer firewall or a failure of the Gray firewall. Similarly, packets blocked by the Gray firewall either indicate a failure of the Outer VPN Gateway, Inner Encryption Components, or Inner firewall.

In order to allow correlation of data from both Gray and Red components the MA CP allows a one-way tap located in the Gray Network to feed Gray Log Data to a Red SIEM located in an enclave of the Red Network that is isolated from the Enterprise/Red network. The one-way tap utilized must physically only allow log data to flow from Gray Components to the Enterprise/Red network. Additionally, the selected one-way tap must be approved by the AO. Similarly, the CP permits an AO-approved one way tap located in the Black Network to feed Black Component Log Data (i.e. Outer firewall) to an enclave of the Red Network that is isolated from Enterprise/Red network.

9 KEY MANAGEMENT

MA solutions utilize asymmetric algorithms (as defined in table 9 - 11) and X.509 v3 certificates for Component authentication to establish the outer and inner encryption tunnels. Each MA solution Component contains a private authentication key and a corresponding public certificate issued by an authorized CA. In addition, a Trusted CA certificate is installed as well as any other CA signing certificates that chain to the Trusted CA, so that a trusted certificate chain is established between the Component certificate and the Trusted CA certificate. Each MA Solution Infrastructure component should also contain the required CRLs to support revocation status checking of Component certificates. If CRLs are not used, other mechanisms can be implemented (e.g. whitelists) in MA solution infrastructure components.

It is preferable for the authentication keys (public/private key pair) to be generated on the security Component, where the private keys are never exported out of the Component. If the Component cannot generate its own key pair, a dedicated management workstation is required to generate the key pair for the Component. The public keys are sent in certificate requests to the Outer and Inner CAs that create and sign authentication certificates containing the public keys. If the request is for a user certificate, the request is delivered to the CA in the customer's organization that has the authority to issue enterprise user certificates. This CA may not be the same as the Inner CA. The authentication certificates are delivered to, and installed on the security Components during provisioning, along with

the private keys if they were not generated on the Component. The CAs also issue signed CRLs to provide revocation status information for the certificates issued by the CAs. CRLs are transferred to CDPs or OCSP Responders as discussed in Section 9.1, where the CRLs or Certificate Statuses are made available to MA Solution Infrastructure Components.

To provide confidentiality services within MA solutions, the Components utilize key agreement protocols (such as Elliptive Curve Diffie-Hellman(ECDH)) to generate ephemeral encryption keys. The use of ephemeral encryption keys is not part of key management discussed in this section, as CAs are not required in issuing and managing these keys.

The CAs that issue authentication certificates to MA solution Components operate either as Enterprise CAs (e.g. DoD PKI, KMI, and Agency PKI) or locally run CAs. Existing Enterprise CAs should be used whenever possible, as the advantages for using these CAs outweigh those associated with locally run CAs. Enterprise CAs have established operations, as well as Certificate Policies and Certification Practices Statements (CPSs) that customer organizations can leverage for their MA solution. These Enterprise CAs operate at Federal Department (e.g. DoD PKI, KMI) and Agency levels, and offer wide-scale interoperability across MA solutions (i.e., the certificate policies and their registered policy Object Identifiers (OIDs) are widely accepted across the Federal Department or Agency). When an Enterprise Root CA is utilized, the MA CP requires that at least two existing Subordinate CA's are used to issue certificates. One Subordinate CA issues certificates to Outer Encryption Components (known as the Outer CA) and the other CA is utilized to issue certificates to Inner Encryption Components (known as the Inner CA). To ensure that the same certificate cannot be used for authenticating both the outer and inner tunnels, the Outer CA and Inner CA are used as trust anchors to validate the outer tunnel and inner tunnel authentication certificates, respectively.

For MA solutions requiring interoperability across a Federal Department, the Department-level Enterprise CAs should be leveraged. Examples of Department-level Enterprise CAs include the DoD PKI; the NSA Key Management Infrastructure (KMI); the National Security Systems (NSS) PKI; the Intelligence Community (IC) PKI; the Department of Homeland Security (DHS) PKI; and the Department of Energy (DoE) PKI. Enterprises like this leverage Department-level Trusted CAs which reside under the same Root CA. Trusted CAs like this can be used as trust anchors in multiple MA solutions throughout a Federal Department, thereby providing certificate trust interoperability across those MA solutions. In addition, certificates issued by Department-level Enterprise CAs may assert registered policy OIDs that are acceptable for use through the Federal Department. A user with a MA EUD provisioned with certificates from a Department-level Enterprise CA could possibly use their EUD in many different MA solutions deployed throughout a Federal Department.

Similarly, MA solutions requiring interoperability across a Federal Agency should leverage Agency-level Enterprise CAs. Agency-level Enterprise CAs issue certificates only to Agency personnel and Non-Person Entities (NPEs). Enterprises like this leverage Agency-level Trusted CAs which resides under the same Root CA. Trusted CAs like this can be used as trust anchors in multiple MA solutions throughout that

Agency. Furthermore, certificates issued by Agency-level Enterprise CAs may assert registered policy OIDs that are acceptable for use through the Federal Agency. A user with a MA EUD provisioned with certificates from an Agency-level Enterprise CA could possibly use their EUD in different MA solutions deployed throughout that Federal Agency.

For both types of Enterprise CAs described above, an MA solution owner could deploy and operate independent Subordinate CAs that are issued certificates by a higher-level Enterprise CA. The benefit of this configuration is that it allows tailoring of the Subordinate CA operations to the local environment without losing the interoperability benefits gained by leveraging Enterprise CAs. However, the MA solution owner is responsible for defining and implementing CPSs for the Subordinate CAs that are approved by the Enterprise CA policy authorities.

Finally, MA solutions requiring minimal or no interoperability can deploy and operate their own locally run CAs that are independent of any Enterprise CAs. In this configuration, certificate policy and interoperability is constrained to the specific MA solution. Furthermore, the MA solution owner is required to develop and maintain CPSs that detail the operational procedures for the locally run CAs. In addition, the customer may need to develop and maintain a higher-level Certificate Policy if one does not already exist.[1] Table 4 summarizes the differences between Enterprise and ocally run CAs.

Table 3. Certificate Authority Deployment Options

CA Type	Certificate Policy	Interoperability	Operations
Department-level Enterprise	Owned and managed at the Department level (e.g. DoD PKI, NSA KMI, NSS PKI, IC PKI, DHS PKI, DoE PKI)	Department-wide	Performed by the enterprise
Agency-level Enterprise	Owned and managed at the Agency level	Agency-wide	Performed by the enterprise
Subordinate CA (Enterprise)	Owned and managed at the Department or Agency level	Department-wide or Agency-wide	Performed by the enterprise and the MA solution owner
Locally run (Non-Enterprise)	Owned and managed at the MA solution level	Constrained to the MA solution	Performed by the MA solution owner

In all CA configurations identified above, Outer CAs issue and manage authentication certificates for Outer VPN Components and Gray Management Service Components; Inner CAs, and optionally existing CAs that support enterprise services, issue and manage authentication certificate for Inner Encryption Components and Red Management Service Components. Outer CAs can be included as either part of

[1] CNSSP 25 is the governing policy for PKI solutions in support of Secret MA solutions. For MA solutions that are higher than Secret, the MA solution owner is required to develop a Certificate Policy that is approved by the local Approving Official (AO).

the Gray network or Red network. Inner CAs, including existing enterprise CAs, can only be located in the Red network.

To assist the CAs in their operations, the CAs may communicate with management services (e.g., Device Managers (DMs)) deployed in the corresponding network to support enrollment and life-cycle certificate management for MA Solution Components. Outer and Inner CAs in the Red network are limited to directly communicating with Red Management Services. Outer CAs in the Gray network are limited to directly communicating with Gray Management Services. When the CA is not located in the same network as the required management services, an AO-approved Cross Domain Solution may be utilized allowing indirect communication (for example Certificate Enrollment). The Red and Gray Management Services enable the certificate request/response process between a MA Solution Component and a CA. This CP recommends provisioning of MA Solution Components in the Red network, and that all enrollment and life-cycle certificate management be performed in accordance with the applicable Certification Practices Statements (CPSs).

This solution utilizes device authentication certificates and, in some instances, user authentication certificates. Device certificates and private keys used in the solution are considered CUI (unless determined to be higher by the AO) because they are only used for mutual authentication, not for traffic encryption or granting access to classified data. User private keys are classified to the level determined by the AO. Often user private keys are treated as classified to the level of the Enterprise/Red network. While the Capability Package enables AOs to define the classification level of User and Device private keys, the allowable options for use and handling of EUDs is dependent on that classification. If any of the private keys stored on an EUD are considered classified, then the solution must be treated as classified at all times or implement a NSA approved DAR Solution. Conversely, if the private keys stored on the EUD are determined to be CUI then the EUD can also be utilized as a Thin EUD (see Section 4.2.1).

The MA solution described here requires certificates to establish the secure tunnels between Components. Without certificates, the network cannot function. Thus, an out-of-band method must be used to issue the initial certificates to the Components. Subsequent rekeying, however, should take place over the network through this solution prior to the current key's expiration. The key validity period for certificates issued by locally run CAs cannot exceed 14 months, while the key validity period for certificates issued by an Enterprise CA are inherited from the Enterprise CA certificate policy. Updates to CRLs are distributed to Gateways within 24 hours of CRL issuance.

9.1 Distribution Of Certificate Revocation Lists

Certificate-based mutual authentication is required between MA solution components in order to establish the outer and inner encryption tunnels, as well as the encryption tunnels (i.e. device management tunnels) that are used to securely administer and manage the configurations of the outer and inner security infrastructure boundary devices. Part of the mutual authentication process includes checking the revocation status of the Component certificates to ensure they are not revoked prior to

establishing the outer and inner encryption tunnels. Certificate Revocation Lists (CRLs) are used by CAs to convey the revocation status of certificates issued by those CAs, and those CRLs need to be made available to the MA solution components.

A CDP is a web server whose sole function is to provide external distribution of, and access to CRLs issued by CAs. CDPs do not serve any other content, and, in particular, do not host any dynamically generated content. CDPs also do not provide any other services other than the distribution of CRLs. CDPs are optional in an MA solution, and they can exist in the Gray or Red networks. The Outer VPN Gateway in the Solution Infrastructure accesses an Outer CDP, located in the Gray network, to obtain CRLs and check revocation status of EUD Outer VPN Components prior to establishing the outer encryption tunnel. Furthermore, a CDP operating in the Gray network can be accessed by Gray Management Service Components to obtain CRLs and check the revocation status of the Outer VPN Gateway's certificate prior to establishing a device management tunnel with the Outer VPN Gateway. Additionally, the Mobile Access CP also allows for an Inner CDP to be stood up in the Gray Management Services. Placing an Inner CDP in the Gray Management services allows EUDs to check the certificate status of the Inner Encryption component prior to establishing a tunnel. To utilize an Inner CDP in the Gray Management Services, an AO must determine that CRLs generated by the Inner CA are unclassified. These CRLs must also be moved from the Enterprise/Red network to the Gray Management Services using an AO approved method (e.g. Cross Domain Solution).

Solution Infrastructure Inner Encryption Components access an Inner CDP, located in the Enterprise/Red network, to obtain CRLs and check revocation status of EUD Inner Encryption Components prior to establishing the inner encryption tunnel. Likewise, a CDP operating in the Red network can be accessed by Red Management Service Components to obtain CRLs and check the revocation status of the Inner Encryption Component's certificate prior to establishing a device management tunnel with the Inner Encryption Component.

An Outer CDP and an Outer CA may reside on the same or different networks. For example, the Outer CA may be operated in the Red network, while the Outer CDP operates in the Gray network. If they reside on different networks, a one-way transfer mechanism is required to periodically distribute the current CRL from the CA to the CDP. The details and procedures of the one-way transfer mechanism are left to a solution's Authorizing Official (AO).

As CRLs are digitally signed objects that contain minimally identifying information about MA solution components, there are few concerns with the confidentiality of CRLs. Therefore, CRLs can be downloaded by MA solution components over unencrypted Hypertext Transfer Protocol (HTTP). Furthermore, a CRL's integrity is protected by the digital signature of the CA that issued it, and additional integrity protection during CRL download is not required. Additionally, placement of CDPs on the Gray network for the Outer VPN Gateway and Enterprise/Red network for Inner Encryption Components reduces the exposure to external threat actors.

Use of HTTPS for CRL downloading is discouraged, as it introduces a circular dependency between the CDP and the MA solution Component attempting to download the CRL. The MA solution Component would need a CRL to determine whether the CDP's certificate is revoked before establishing an HTTPS connection to the CDP. However, the CDP cannot deliver the CRL to the MA solution Component until the MA solution Component authenticates the CDP by validating its certificate. (Note: Distributing CRLs via HTTP follows the recommendation in Internet Engineering Task Force (IETF) Request for Comments (RFC) 5280 not to use HTTPS to distribute CRLs.)

To provide redundancy and ensure that current CRLs are always made available to MA solution Components, multiple Outer and Inner CDPs may be deployed. The use of multiple CDPs is left to the discretion of the MA solution owner. Furthermore, CDPs may host delta CRLs in addition to complete CRLs. In large MA solutions, the use of delta CRLs can reduce the amount of network traffic needed to distribute updates to CRLs. A CA's Certificate Policy will define whether the use of delta CRLs is permissible.

An OCSP Responder or white lists can be utilized in lieu of a CDP Server. An OCSP Responder located in the Gray network can provide Certificate Status information to the Outer VPN Gateway. Additionally, a OCSP Responder in the Enterprise/Red network can provide Certificate Status information to Inner Encryption Components.

10 THREATS

This section details how the required components work together to provide overall security in the solution. Section 4.2 shows the boundary of the MA solution for each high-level design covered by this CP.

An assessment of security was conducted on each of the high-level designs described in this CP while making no assumptions regarding use of specific products for any of the defined components. There are several different threats to consider when evaluating the risk of transporting data over secure or unsecure networks. By examining these threats, the organization can have a better understanding of the risks they are accepting by implementing the solution and how these risks affect the Confidentiality, Integrity, and Availability of the network, systems, and data. To obtain the classified risk assessment associated with this CP, please contact the NSA via your Client Advocate.

10.1 PASSIVE THREATS

This threat refers to internal or external actors attempting to gain information from the network without changing the state of the system. Threat actions include collecting or monitoring traffic (e.g. traffic analysis or sniffing the network) passing through a network in order to gain useful information through data analysis.

The security against a passive attack targeting the data in transit across the Black network is provided by the layered IPsec tunnels. To mitigate passive attacks, two layers of Suite B encryption, Advanced Encryption Standard (AES), are employed to provide confidentiality for the solution. Use of AES is approved to protect classified information, meeting IAD and CNSSP-15 guidance for adequate confidentiality. The two Encryption Components that are used to set up the tunnels must be independent in a number of ways (see Section 12). Due to this independence, the adversary should not be able to exploit a single cryptographic implementation to compromise both tunnels.

The use of EUDs to access classified information outside of a secure physical environment opens the possibility that an attacker with physical access to the EUD's immediate environment could use surveillance techniques to obtain classified information without the user's knowledge while the EUD is in use. The organization-defined user agreement tells users of EUDs what measures they must follow when using and storing the EUD to mitigate this threat.

10.2 External (Active) Threats

This threat refers to outsiders gaining unauthorized access to a system or network, exfiltration of sensitive Red network data, or degradation of availability of the system or network. Threat actions include introducing viruses, malware, or worms with the intention to compromise the network or exfiltrate data, or to analyze the design of the network or system for future attacks. Adversaries could gain access to a VPN Gateway or EUD, and then exploit or compromise other devices on the network. DoS or Distributed DoS (DDoS) attacks compromise availability of the system, degrading/disrupting secure communication across a Black network. Further external threat actions would include social engineering attacks to assist attackers with gaining additional access to a network for the purpose of compromising a system or network, traffic injection or modification attacks, or replay attacks.

10.2.1 Rogue Traffic

One method for detecting rogue traffic from an external attack as it attempts to pass through one or both Encryption Components is by having the port filtering native to each VPN Gateway enabled and configured to audit and log any traffic that is not one of the formats described in the configuration (see Section 13.9). It is required that the port filtering be set up to block the following: 1) any traffic not coming from or going to an IP address on the network at the other site, 2) traffic not contained in IP packets other than control plane protocols needed for network operation and approved by AO policy, and 3) traffic going to unexpected ports. This will allow the Auditor(s) and/or the Security Administrator(s) to detect whether the Outer VPN Gateway has been breached, thus providing an early warning of a potential intrusion. It will also provide detection of misconfigured Outer VPN Gateways.

Another method for detecting a potential intrusion into the solution is requiring automated configuration change detection on Red and Gray Management networks to ensure that VPN Gateway configurations are not changed without the knowledge of Auditors and Security Administrators.

Auditors also ensure through the audit logs that all configuration changes are valid. This will counter attacks that take advantage of VPN Gateway misconfigurations.

End User Devices (EUDs) are protected from rogue traffic through the use of traffic filtering rules configured on their interfaces connected to Black networks to drop any traffic not necessary for connecting to the necessary Outer VPN Gateway.

CRL Distribution Points (CDPs) are protected from rogue traffic by implementing port filtering on the server. Rogue traffic to CDPs can be further mitigated by implementing a firewall or other packet filtering device between the CDP and the rest of the network.

10.2.2 Malware and Untrusted Updates

Administration Workstations and locally run CAs for Inner Tunnel Components shall be distinct and physically separate from the Administration Workstations and locally run CAs for Outer Tunnel Components. This separation minimizes the potential for malware on a single device to impact components supporting both Inner and Outer MA solution tunnels.

Each individual component of this solution has the capability to perform trusted updates through verification of a signature or hash to ensure that the update is from a reliable source, such as signed by the vendor. This mitigates threats of malicious users trying to push updates or code patches that affect the security of the component (and therefore system). The source of all updates and patches should be verified before installation occurs.

10.2.3 Denial of Service

Denial of Service (DoS) attack risks cannot be completely mitigated. MA solutions in compliance with this CP requires that the Outer firewall drop all packets that are not Internet Key Exchange (IKE), Encapsulating Security Payload (ESP), or other approved protocols on the appropriate interfaces, which significantly reduces the potential for flooding attacks. For customers that require more protection against these attacks, one option is the use of a multiple Outer firewalls between the Outer VPN Gateway and Black networks to filter and load balance traffic before it reaches the Outer VPN Gateway. Another option for customers requiring more protection is to add additional filtering based on specifics like known network IP addresses to filter traffic from devices not included in this solution, although the feasibility of doing so for EUDs is limited. Other mitigations are acceptable and up to the AO to approve their use.

A single encryption component failure is likely to result in a DoS condition. One assumption underlying this solution is that high assurance of availability is not required. If availability is critical for the customer both network redundancy and instituting DoS response procedures when loss of availability is detected can provide further protection against DoS attacks.

10.2.4 Social Engineering

It is the responsibility of the customer to define the appropriate policies and training necessary to protect against Social Engineering attacks. In addition, these types of attacks generally take advantage of other attacks detailed in this section and are already discussed.

10.3 Insider Threats

This threat refers to an authorized or cleared person or group of people with physical or logical access to the network or system who may act maliciously or negligently, resulting in risk exposure for the organization. This threat could include poorly trained employees, curious employees, disgruntled employees, escorted personnel who gain access to the equipment, dishonest employees, or those that have the means and desire to gain escalated privileges on the network.

Threat actions include insertion or omission of data entries that result in a loss of data integrity, unintentional access to an unauthorized system or network, willingly changing the configuration of an EUD, unwillingly or unknowingly executing a virus or malware, intentionally exposing the network and systems to viruses or malware, cross-contaminating a system or network with data from a higher classification to a lower classification (e.g. Secret data to an Unclassified network or system), or malicious or unintentional exfiltration of classified data. Typically, the threat from insiders has the potential to cause the greatest harm to an organization, and insider attacks are also the hardest to monitor and track.

To mitigate insider threats, separation of roles within the solution is required (see Section 15). In addition, logging and auditing of security critical functionality (see Section 13.13) is required. Also, strong authentication of the Security Administrator and Auditor are required for access to ensure accountability of these individuals. Finally, outbound filters on Encryption Components, firewalls, and EUDs are configured to block traffic leaving the internal network that does not go through the Encryption tunnels. An IDS is also deployed on at a minimum the Black, Gray, or Red network to help identify unusual or suspicious traffic that could result from a failure, misconfiguration, or attack on Inner or Outer Encryption Components.

Additionally, organizations concerned about users misbehaving when connected remotely may wish to restrict the use of EUDs to those deemed sufficiently trustworthy.

10.4 Supply Chain Threats

A critical aspect of the U.S. Government's effectiveness is the dependability, trustworthiness, and availability of the Information and Communication Technology (ICT) components embedded in the systems and networks upon which the ability to perform U.S. Government missions rely. The supply chain for those ICT components are the underpinnings of those systems and networks and supply chain attacks are attempts to proactively compromise those underpinnings.

Unfortunately, the supplier cannot always provide guarantees of a safe delivery of a component; they are only able to provide assurances based on their reliance of established procedures and processes they have developed. In a single change of hands, the component may be introduced to potential threats and compromises on many levels.

The supply chain threat refers to an adversary gaining access to a vendor, retailer, reseller, or shipper and then attempting to insert or install a modification or a counterfeit piece of hardware into a component that is destined for a U.S. Government customer in an effort to gain information or cause operational issues. This threat also includes the installation of malicious software on components of the solution. This threat is difficult to identify, and is increasingly more difficult to prevent or protect against since vendors build products containing components manufactured by subcontractors. It is often difficult to determine where different pieces of components are built and installed within the supply chain.

Threat actions include manufacturing faulty or counterfeit parts of components that can be used to disrupt system or network performance, leaving open back doors in hardware that allow attackers easy ways to attack and evade monitoring, as well as easy ways to steal data or tamper with the integrity of existing/new data. Supply Chain attacks may occur during development and production, updates, distribution, shipping, at a warehouse, in storage, during operations, or disposal. For this reason, it is imperative that all components selected for use in CSfC solutions are subject to the applicable Supply Chain Risk Management (SCRM) process to reduce the risk of acquiring compromised components.

Each component that is selected from the CSfC Components List shall go through a Product Supply Chain Threat Assessment to determine the appropriate mitigations for the intended application of the component per the organization's AO-approved Product Supply Chain Threat Assessment process (see CNSSD 505 Supply Chain Risk Management (SCRM) for additional guidance). Even after selecting components from the CSfC Components List and utilizing a rigorous acquisition process, an AO must perform due diligence when integrating commercial components for mission operations.

Doctrinal requirements are placed on Product Selection, Implementers, and System Integrators of these solutions to minimize the threat of supply chain attacks (see sections 12, 14, and 15). To further mitigate Supply Chain Threats implementing organizations should utilize the following guidance.

- Establish an ICT SCRM program which conforms to applicable policy based on external and organizational requirements and constraints. The ICT SCRM program should be integrated into the organizational business and mission processes.
- Assess all aspects of the performance of potential vendors, not only the product quality, cost, and performance, but also supply chain risk factors of vendor selection. These risk factors include political ties to foreign governments, citizenship of employees, partner affiliations, employee clearance levels, and location of suppliers and sub-suppliers.

- Ensure that each component selected from the CSfC Components List go through a Product Supply Chain Threat Assessment to determine the appropriate mitigations for the intended application of the component (see CNSSD 505 Supply Chain Risk Management and Intelligence Community Directive (ICD) 731 Supply Chain Risk Management).

- Conduct a Criticality Analysis by which mission-critical functions and components are identified and prioritized with respect to improving acquirer practices (see Defense Acquisition Guidebook, Chapter 13).

Supply chain risk management is a critical consideration in acquiring commercial products. Even after selecting components from the CSfC Components List and utilizing a rigorous acquisition process an AO must do their due diligence as composed commercial products are integrated into mission operations.

10.5 Integrator Threats

This threat refers to an integrator who has unrestricted access to all components within the solution prior to the customer purchasing and implementing the solution within their system. This is different from a Supply Chain threat in that these integrators have access to all components to be used in the solution, rather than only those being procured from a particular vendor.

Threat actions could include installing or configuring components in a manner that places the organization at risk for attack or open to an unknown vulnerability that may not be detected through normal tests, scans, and security counter-measures.

In order to mitigate this threat, integrators are required to be cleared to the highest level of data protected by the MA solution. To further reduce the integrator threat, a customer may wish to use multiple integrators, such that no one integrator has access to all components of the solution. More information on the NSA's list of trusted integrators can be found on the NSA CSfC Website in the "Criteria For CSfC Integrators" section at this link: https://www.nsa.gov/ia/programs/csfc/index.shtml.

11 REQUIREMENTS OVERVIEW

The following five sections (Sections 12 through 16) specify requirements for implementations of MA solutions compliant with this CP. However, not all requirements in the following sections will apply to each compliant solution. Sections 11.1 and 11.2 describe how to determine which set of requirements applies to a particular solution.

11.1 Capabilities

This CP provides the flexibility needed to implement a variety of designs for the implementation of the MA solution. Although most requirements are applicable to all solutions, some requirements are only applicable to implementations whose high-level designs implement certain features. For example,

requirements dealing with TLS EUDs do not include requirements for an Inner VPN Client. Table 4 lists the capabilities covered by this CP and the designators used in the requirements tables to refer to each.

Table 4. Capability Designators

Capability	Designator	Description
TLS Solution	T	Requirement that applies to the MA Solution that connects to the Red network using IPsec as the Outer layer and TLS or SRTP as the inner layer, as described in Section 6.3
VPN Solution	V	Requirement that applies to the MA solution that connects to the Red network using two IPsec tunnels, as described in Section 6.2
TLS Infrastructure	TI	Requirement that applies specifically to the infrastructure associated with the TLS solution
VPN Infrastructure	VI	Requirement that applies specifically to the infrastructure associated with the VPN solution
TLS EUD	TE	Requirement that applies specifically to the EUD associated with the TLS solution
VPN EUD	VE	Requirement that applies specifically to the EUD associated with the VPN solution
All Solution Components	All	Requirement that applies to the EUD and to the Infrastructure, regardless if it is a VPN solution or a TLS solution
CDPs	C	Requirement that applies to the MA Solution that includes CDPs, as described in section 13.14.4

Any solution that follows this CP must implement each applicable capability for their solution (i.e. all VPN EUD (V), VPN Infrastructure (VI), and VPN Solution (V) requirements for a solution supporting only VPN EUDs), and may implement multiple capabilities. The "Capabilities" column in the requirements tables in Sections 12 through 16 identifies which capabilities the requirement applies to. A requirement is only applicable to a solution if the "Capabilities" column for that requirement lists one or more of the capabilities being implemented by the solution.

11.2 Threshold and Objective Requirements

In some cases, multiple versions of a requirement may exist in this CP. Such alternative versions of a requirement are designated as being either a Threshold requirement or an Objective requirement:

- A Threshold (T) requirement specifies a feature or function that provides the minimal acceptable capability for the security of the solution.
- An Objective (O) requirement specifies a feature or function that provides the preferred capability for the security of the solution.

In general, when separate Threshold and Objective versions of a requirement exist, the Objective requirement provides a higher degree of security for the solution than the corresponding Threshold

requirement. However, in these cases meeting the Objective requirement may not be feasible in some environments or may require components to implement features that are not yet widely available. Solution owners are encouraged to implement the Objective version of a requirement, but in cases where this is not feasible solution owners may implement the Threshold version of the requirement instead. These Threshold and Objective versions are mapped to each other in the "Alternatives" column. Objective requirements that have no related Threshold requirement are marked as "Optional' in the "Alternatives" column.

In most cases, there is no distinction between the Threshold and Objective versions of a requirement. In these cases, the "Threshold / Objective" column indicates that the Threshold equals the Objective (T=O). Such requirements must be implemented in order to comply with this CP, as long as the requirement is applicable per Section 11.1.

Requirements that are listed as Objective in this CP may become Threshold requirements in a future version of this CP. Solution owners are encouraged to implement Objective requirements where possible in order to facilitate compliance with future versions of this CP.

11.3 Requirements Designators

Each requirement defined in this CP has a unique identifier consisting of the prefix "MA," a digraph that groups related requirements together (e.g. KM), and a sequence number (11). Table 5 lists the digraphs used to group together related requirements and identifies the sections in which those requirement groups can be found.

Table 5. Requirement Digraphs

Digraph	Description	Section	Table
PS	Product Selection Requirements	Section 12	Table 6
SR	Overall Solution Requirements	Section 13.1	Table 7
CR	Configuration Requirements for Inner and Outer VPN Components	Section 13.3	Table 11
IR	Inner VPN Component Requirements	Section 13.4	Table 12
OR	Outer VPN Component Requirements	Section 13.5	Table 13
TE	TLS-Protected Server & SRTP Endpoint Requirements	Section 13.6	Table 14
RD	Retransmission Device Requirements	Section 13.7	Table 15
EU	End User Device Requirements	Section 13.8	Table 16
PF	Port Filtering Requirements for Solution Components	Section 13.9	Table 17
CM	Configuration Change Detection Requirements	Section 13.10	Table 18
DM	Device Management Requirements	Section 13.11	Table 19
MR	Continuous Monitoring Requirements	Section 13.12	Table 20
AU	Auditing Requirements	Section 13.13	Table 21

Digraph	Description	Section	Table
KM	Key Management Requirements	Section 13.14	Table 22, Table 23, Table 24, Table 25
GD	Requirements for the Use and Handling of Solutions	Section 14.1	Table 26
	Role-Based Personnel Requirements	Section 15	Table 28
RP	Incident Reporting Requirements	Section 14.2	Table 27
TR	Test Requirements	Section 16.1	Table 29

12 REQUIREMENTS FOR SELECTING COMPONENTS

In this section, a series of requirements are given for maximizing the independence between the components within the solution. This will increase the level of effort required to compromise this solution.

Table 6. Product Selection Requirements

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-PS-1	The products used for the Inner VPN Gateway shall be chosen from the list of IPsec VPN Gateways on the CSfC Components List.	VI	T=O	
MA-PS-2	The products used for any Outer VPN Gateway shall be chosen from the list of IPsec VPN Gateways on the CSfC Components List.	All	T=O	
MA-PS-3	The products used for any Inner VPN Client shall be chosen from the list of IPsec VPN Clients on the CSfC Components List.	VE	T=O	
MA-PS-4	The products used for any Outer VPN Client shall be chosen from the list of IPsec VPN Clients on the CSfC Components List.	TE, VE	T=O	
MA-PS-5	The products used for the Inner and Outer CAs shall either be chosen from the list of CAs on the CSfC Components List or the CAs shall be pre-existing Enterprise CAs (e.g. DoD PKI, IC PKI)..	VI, TI	T=O	
MA-PS-6	Products used for Mobile Platform EUDs shall be chosen from the list of Mobile Platforms on the CSfC Components List.	VE, TE	T=O	
MA-PS-7	Intrusion Prevention Systems (IPS) shall be chosen from the list of IPS on the CSfC Components List.	VI, TI	O	Optional
MA-PS-8	Products used for the TLS Client shall be chosen from the TLS Client sections (e.g. VoIP Applications, Email Clients, Web Browsers, etc.) of the CSfC Components List.	TE	T=O	

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-PS-9	Products used for the SRTP Client shall be chosen from the list of VoIP Applications on the CSfC Components List.	TE	T=O	
MA-PS-10	If the solution is using a TLS-Protected Server, it shall be chosen from the list of TLS-Protected Servers on the CSfC Components List.	TI	T=O	
MA-PS-11	If the solution is using an SIP Server, it shall be chosen from the list of SIP Servers on the CSfC Components List.	TI	T=O	
MA-PS-12	If the solution is using an SRTP Endpoint, it shall be chosen from the list of SRTP Endpoints on the CSfC Components List.	TI	T=O	
MA-PS-13	Products used for the Outer firewall, Gray firewall, and Inner firewall shall be chosen from the list of Stateful Traffic Filtering Firewalls (TFFW) on the CSfC Components List.	VI, TI	T=O	
MA-PS-14	If the solution is using a MDM, it shall be chosen from the list of MDMs on the CSfC Components List.	VI, TI	T=O	
MA-PS-15	Withdrawn			
MA-PS-16	The Outer VPN Gateway and Inner Encryption Endpoints shall either: • come from different manufacturers, where neither manufacturer is a subsidiary of the other; or • be different products from the same manufacturer, where NSA has determined that the products meet the CSfC criteria for implementation independence.	All	T=O	
MA-PS-17	The Outer firewall, Outer VPN Gateway, Gray Firewall, Inner Encryption Component, and Inner firewall shall use physically separate components, such that no component is used for more than one function (see Figure 1).	VI, TI	T=O	

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-PS-18	The Outer VPN Gateway and the Inner Encryption Endpoints shall not use the same Operating System (OS). Differences between Service Packs (SP) and version numbers for a particular vendor's OS do not provide adequate diversity.	VI, TI	T=O	
MA-PS-19	The Inner and the Outer CAs shall either: • come from different manufacturers, where neither manufacturer is a subsidiary of the other; or • be different products from the same manufacturer, where NSA has determined that the products meet the CSfC criteria for implementation independence.	VI, TI	O	Optional
MA-PS-20	The Gray Network Firewall and the Inner Encryption Endpoints shall either: • come from different manufacturers, where neither manufacturer is a subsidiary of the other; or • be different products from the same manufacturer, where NSA has determined that the products meet the CSfC criteria for implementation independence.	VI, TI	T=O	
MA-PS-21	The EUD's Outer VPN Component and Inner Encryption Components shall either: • come from different manufacturers, where neither manufacturer is a subsidiary of the other; or • be different products from the same manufacturer, where NSA has determined that the products meet the CSfC criteria for implementation independence.	VE, TE	T=O	

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-PS-22	The cryptographic libraries used by the Inner CA and Outer CA shall either: • come from different manufacturers, where neither manufacturer is a subsidiary of the other; or • be different libraries from the same manufacturer, where NSA has determined that the libraries meet the CSfC Program's criteria for implementation independence.	VI, TI	O	Optional
MA-PS-23	The cryptographic libraries used by the Outer VPN Component and the Inner Encryption Components shall either: • come from different manufacturers, where neither manufacturer is a subsidiary of the other; or • be different libraries from the same manufacturer, where NSA has determined that the libraries meet the CSfC criteria for implementation independence.	VE, TE	O	Optional
MA-PS-24	Each component that is selected out of the CSfC Components List shall go through a Product Supply Chain Threat Assessment to determine the appropriate mitigations for the intended application of the component per the organization's AO-approved Product Supply Chain Threat Assessment process (see CNSSD 505 SCRM for additional guidance).	All	T=O	
MA-PS-25	Components shall be configured to use the NIAP-certified evaluated configuration.	All	T	

13 CONFIGURATION REQUIREMENTS

Once the products for the solution are selected, the next step is setting up the components and configuring them in a secure manner. This section consists of generic guidance for how to configure the components of the MA solution.

13.1 Overall Solution Requirements

Table 7. Overall Solution Requirements

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-SR-1	Network services provided by control plane protocols (such as DNS and NTP) shall be located on the inside network (i.e., Gray network for the Outer VPN Gateway and Red network for the Inner Encryption Endpoints).	VI, TI	T=O	
MA-SR-2	The time of day on Inner Encryption Endpoints, Inner firewall, and Red Management Services shall be synchronized to a time source located in the Enterprise/Red network.	VI, TI	T=O	
MA-SR-3	The time of day on the Outer VPN Gateway, Gray firewall, and Gray Management Services shall be synchronized to a time source located in the Gray Management network.	VI, TI	T=O	
MA-SR-4	Default accounts, passwords, community strings, and other default access control mechanisms for all components shall be changed or removed.	All	T=O	
MA-SR-5	All components shall be properly configured in accordance with local policy and applicable U.S. Government guidance. In the event of conflict between the requirements in this CP and local policy, this CP takes precedence.	All	T=O	
MA-SR-6	Solution Components shall receive virus signature updates as required by the local agency policy and the AO.	All	T=O	
MA-SR-7	The only approved physical paths leaving the Red network shall be through a MA solution in accordance with this CP or via an AO-approved solution for protecting data in transit.[2]	All	T=O	

[2] In some cases, the customer will need to communicate with other sites that have NSA-certified Government off-the-Shelf (GOTS) product. In particular, it is acceptable for a given site to have both an egress path via an NSA-certified product and an egress path via a CSfC Solution conforming to a CP.

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-SR-8	When multiple Inner Encryption components are placed between the Gray firewall and Inner firewall, they shall be placed in parallel.	VI, TI	T=O	
MA-SR-9	Inner Encryption components shall not perform switching or routing for other Encryption Components.	VI, TI	T=O	

13.2 CONFIGURATION REQUIREMENTS FOR ALL VPN COMPONENTS

Table 8. Approved Suite B Algorithms (IPSec)

Security Service	Algorithm Suite 2	Specifications
Overall Level of Security	192 bits	
Confidentiality (Encryption)	AES-256	FIPS PUB 197 IETF RFC 6239 IETF RFC 6379 IETF RFC 6380 IETF RFC 6460
Authentication (Digital Signature) (Threshold)	RSA 2048	FIPS PUB 186-4
Authentication (Digital Signature) (Objective)	RSA 3072 or ECDSA over the curve P-384 with SHA-384	FIPS PUB 186-4 FIPS PUB 186-4 IETF RFC 6239 IETF RFC 6380 IETF RFC 6460
Key Exchange/ Establishment	ECDH over the curve P-384 (DH Group 20)	NIST SP 800-56A IETF RFC 6239 IETF RFC 6379 IETF RFC 6380 IETF RFC 6460 NIST SP 800-56A
Integrity (Hashing)	SHA-384	FIPS PUB 180-4 IETF RFC 6239 IETF RFC 6379 IETF RFC 6380 IETF RFC 6460
Can protect	Up to Top Secret	

Table 9. Approved Algorithms (TLS)

	TLS Cipher Suites	Specifications
Overall Level of Confidentiality	192 bits	FIPS PUB 180-4 FIPS PUB 186-3 FIPS PUB 197 FIPS 800-56A IETF RFC 6460 IETF RFC 5246 IETF RFC 4492
TLS Cipher Suite (Threshold)	TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384	
TLS Cipher Suite (Objective)	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384	
Authentication (Digital Signature) (Threshold)	RSA 2048	
Authentication (Digital Signature) (Objective)	RSA 3072 or ECDSA over the curve P-384 with SHA-384	
Key Exchange	ECDHE over the curve P-384 (DH Group 20)	
Can protect	Up to Top Secret	

Table 10. Approved Algorithms (Secure Real-Time Protocol)

Security Service	Algorithm Suite 1	Algorithm Suite 2	Specifications
Overall Level of Security	128 bits	192 bits	IETF RFC 3711 IETF RFC 6188
Confidentiality (Encryption)	AES-128 in Counter Mode (CM)	AES-256 in Counter Mode (CM)	IETF RFC 3711 IETF RFC 2675
Integrity	HMAC-SHA1	HMAC-SHA1	IETF RFC 3711 IETF RFC 2104
Key Exchange (using SIP Over TLS)	TLS-SDES or DTLS	TLS-SDES or DTLS	IETF RFC 4568 IETF RFC 6347
Can protect	Up to Secret	Up to Top Secret	

13.3 CONFIGURATION REQUIREMENTS FOR INNER AND OUTER VPN COMPONENTS

Table 11 Configuration Requirements for Inner and Outer VPN Components

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-CR-1	The proposals offered by the Outer and Inner VPN Components in the course of establishing the IKE Security Association (SA) and the ESP SA for Inner and Outer Tunnels shall be configured to only offer algorithm suite(s) containing the Suite B algorithms listed in Table 8.	All	T=O	
MA-CR-2	Default, self-signed, or proprietary device certificates, which are frequently preinstalled by the vendor, for any Outer and Inner VPN Component, shall not be used for establishing SAs.	All	T	MA-CR-3
MA-CR-3	Default, self-signed, or proprietary device certificates, which are frequently preinstalled by the vendor, for any Outer and Inner VPN Component, shall be removed.	All	O	MA-CR-2
MA-CR-4	Unique device certificates shall be loaded onto the Outer and Inner VPN Gateway along with the corresponding Trust Anchor (signing) certificates.	VI, TI	T=O	
MA-CR-5	A device certificate shall be used for each Outer and Inner VPN Component authentication during IKE.	All	T=O	
MA-CR-6	Authentication performed by Outer and Inner VPN Gateways shall include a check that device certificates are authorized. This check may use a CRL, OCSP, or a whitelist.	VI, TI	T=O	
MA-CR-7	Outer and Inner VPN Component authentication with device certificates shall include a check that certificates are not expired.	All	T=O	
MA-CR-8	Withdrawn			
MA-CR-9	All IPsec connections shall use IETF standards IKE implementations (RFC 5996 or RFC 2409).	All	T=O	
MA-CR-10	All Outer and Inner VPN Components shall use Cipher Block Chaining for IKE encryption.	All	T=O	
MA-CR-11	All Outer and Inner VPN Components shall use Cipher Block Chaining for ESP encryption with an HMAC for integrity.	All	T	MA-CR-12

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-CR-12	All Outer and Inner VPN Components shall use Galois Counter Mode for ESP encryption.	All	O	MA-CR-11
MA-CR-13	All Outer and Inner VPN Components shall set the IKE SA lifetime to at most 24 hours.	All	T=O	
MA-CR-14	All Outer and Inner VPN Components shall set the ESP SA lifetime to at most 8 hours.	All	T=O	
MA-CR-15	All VPN Components shall re-authenticate the identity of the VPN Component at the other end of the established tunnel before rekeying the IKE SA.	All	T=O	

13.4 INNER VPN COMPONENTS

Table 12. Inner VPN Components Requirements

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-IR-1	The Inner VPN Component shall use Tunnel mode IPsec or Transport mode IPsec using an associated IP tunneling protocol (e.g. Transport Mode IPsec with GRE).	VI	T=O	
MA-IR-2	The packet size for packets leaving the external interface of the Inner VPN Component shall be configured to reduce packet fragmentation and impacting performance. This requires proper configuration of the Maximum Transmission Unit (MTU) (for IPv4) or Path MTU (PMTU) (for IPv6) and should consider Black network and Outer VPN component MTU/PMTU values to achieve this.	VI	O	Optional
MA-IR-3	The Inner VPN Gateway shall not allow any packets received on an interface connected to a Red network to bypass encryption and be forwarded out through an interface connected to a Gray network.	V	T=O	

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-IR-4	The Inner VPN Client of EUDs shall encrypt all traffic, with the exception of traffic necessary for the EUD to connect to the physical network (e.g. DHCP) and locate the Inner VPN Gateway (i.e. DNS lookup of the VPN Component's IP address), in accordance with this CP.	VE	T=O	
MA-IR-5	The Inner VPN Component shall not allow any packets received on an interface connected to a Gray network to bypass decryption and be forwarded out through an interface connected to a Red network.	V	T=O	

13.5 Outer VPN Components

Table 13. Outer VPN Components Requirements

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-OR-1	Outer VPN Components shall use Tunnel mode IPsec.	All	T=O	
MA-OR-2	Outer VPN Components shall not permit split-tunneling.	All	T=O	
MA-OR-3	The Outer VPN Component shall not allow any packets received on an interface connected to a Gray network to bypass encryption and be forwarded out through an interface connected to a Black network.	All	T=O	
MA-OR-4	All traffic received by the Outer VPN Component on an interface connected to a Gray network, with the exception of Control Plane traffic not prohibited in the CP, shall have already been encrypted once.	All	T=O	
MA-OR-5	The Outer VPN Client of EUDs shall encrypt all traffic, with the exception of traffic necessary for the EUD to connect to the physical network (e.g. DHCP) in accordance with this CP (see Section 4.1.4).	VE, TE	T=O	
MA-OR-6	If one or more virtual machines are used to separate Outer and Inner VPN Clients on an EUD, then the Outer VPN Client shall not run on the host operating system.	VE, TE	T=O	

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-OR-7	Outer VPN Component shall not allow any packets received on an interface connected to a Black network to bypass decryption.	All	T=O	
MA-OR-8	Withdrawn			
MA-OR-9	Outer VPN Gateways shall not perform routing.	VI, TI	T=O	

13.6 TLS-Protected Server & SRTP Endpoint Requirements

Table 14. TLS-Protected Server & SRTP Endpoint Requirements

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-TE-1	TLS Components shall utilize TLS 1.2 or later.	T	T=O	
MA-TE-2	TLS Solution Infrastructure components shall terminate the Inner layer of encryption originating from TLS EUDs.	TI	T=O	
MA-TE-3	TLS Solution Infrastructure components shall use X.509 device certificates for mutual authentication with TLS EUDs.	TI	T=O	
MA-TE-4	Default, self-signed, or proprietary certificates, which are frequently preinstalled by the vendor, for the TLS Component shall be disabled.	T	T	MA-TE-5
MA-TE-5	Default, self-signed, or proprietary certificates, which are frequently preinstalled by the vendor, for the TLS Component shall be removed.	T	O	MA-TE-4
MA-TE-6	Unique device certificates shall be loaded onto TLS Components along with the corresponding Trust Anchor (signing) certificates.	T	T=O	
MA-TE-7	TLS Components shall only use ciphers suites selected from the “TLS Cipher Suite (Threshold)” row of Table 9.	T	T	MA-TE-8
MA-TE-8	TLS Components shall only use cipher suites selected from the “TLS Cipher Suite (Objective)” row of Table 9.	T	O	MA-TE-7

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-TE-9	SRTP Components shall only use algorithms selected from Table 10 that are approved to protect the highest classification level of the Red network data.	T	T=O	

13.7 RETRANSMISSION DEVICE REQUIREMENTS

Table 15. Requirements for Retransmission Device

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-RD-1	An EUD shall only connect to Retransmission Devices (RDs) authorized by a Government AO.	VE, TE	T=O	
MA-RD-2	An RD shall provide EUDs with connectivity to the Mobile Access Solution infrastructure via any Black Network using Wi-Fi or an Ethernet cable.	VE, TE	T=O	
MA-RD-3	If the RD is configured to be a Wi-Fi access point, the Wi-Fi network shall implement WPA2 AES-CCMP PSK or WPA2-Enterprise with a TLS-based EAP method.	VE, TE	T=O	
MA-RD-4	A RD shall not be utilized to protect Gray data between an Outer VPN Gateway and EUD.	VE, TE	T=O	
MA-RD-5	If the RD is configured to be a Wi-Fi access point using PSK, then the PSK shall use a length of at least 32 hexadecimal characters (or its equivalent).	VE, TE	T	MA-EU-26
MA-RD-6	RD shall only permit connections to devices on a Media Access Control (MAC) white list.	VE, TE	O	Optional
MA-RD-7	If the RD is configured as a Wi-Fi access point, then the PSK shall not be displayed on the RD.	VE, TE	T=O	
MA-RD-8	If the RD is configured as a Wi-Fi access point, then the Service Set Identification (SSID) shall not be displayed on the RD.	VE, TE	T=O	
MA-RD-9	If the RD is configured as a Wi-Fi access point, then the MAC address of connected devices shall not be displayed on the RD.	VE, TE	T=O	

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-RD-10	The Administrator password shall not be displayed on the RD.	VE, TE	T=O	
MA-RD-11	The RD shall display the number of currently connected devices.	VE, TE	O	Optional
MA-RD-12	If the RD is configured to be a Wi-Fi access point, then Wi-Fi Protected Setup (WPS) shall be disabled.	VE, TE	T=O	
MA-RD-13	The RD shall be administered using HTTPS.	VE, TE	T=O	
MA-RD-14	The RD shall require authentication with Administrator credentials to make changes to RD settings.	VE, TE	T=O	
MA-RD-15	The RD default Administrator credentials shall be changed during provisioning.	VE, TE	T=O	
MA-RD-16	The RD shall be configured to limit the number of connected devices to the maximum required for the mission.	VE, TE	T=O	
MA-RD-17	If the RD is configured as a Wi-Fi access point, then traffic of multiple EUDs sharing the RD shall be separated (commonly referred to as Wi-Fi Privacy Separation or AP Isolation).	VE, TE	T=O	
MA-RD-18	If the RD is configured as a Wi-Fi access point, then the RD shall disable broadcasting of the SSID.	VE, TE	O	Optional
MA-RD-19	The RD shall only permit charging on USB ports and interfaces.	VE, TE	O	Optional
MA-RD-20	The RD shall not permit connected EUDs to access files stored on the RD.	VE, TE	T=O	
MA-RD-21	The RD shall require Administrator authentication prior to downloading logs or configuration files.	VE, TE	T=O	
MA-RD-22	The RD shall only allow firmware updates signed by the RD manufacturer.	VE, TE	O	Optional
MA-RD-23	The RD shall prevent the ability to boot into recovery mode.	VE, TE	O	Optional
MA-RD-24	The RD shall require user or administrator authentication prior to updating firmware.	VE, TE	O	Optional
MA-RD-25	If the RD is configured to be a Wi-Fi access point, the PSK shall use a length of at least 64 hexadecimal characters (or its equivalent).	VE, TE	O	MA-RD-5

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-RD-26	If the RD is configured to be a Wi-Fi access point using WPA Enterprise, the certificate used for authentication shall be different from the certificates utilized to authenticate the outer and inner tunnels.	VE, TE	T=O	

13.8 End User Devices Requirements

Table 16. Requirements for End User Devices

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-EU-1	EUDs which do not implement an NSA-approved DAR solution and allow a user to store classified information on the EUD shall be treated as classified at all times. (See Section 4.2.1).	TE, VE	T=O	
MA-EU-2	EUDs which implement an NSA-approved DAR solution (i.e. Data at Rest CP) shall comply with the handling requirements specified for the DAR solution.	VE, TE	T=O	
MA-EU-3	Thin EUDs which prohibit a user from storing classified information shall be treated as unclassified, or a higher classification level as determined by the AO, when powered down.	VE, TE	T=O	
MA-EU-4	The Outer VPN Client private key store shall be separate from the private key store for the Inner VPN Client.	VE TE	T=O	
MA-EU-5	The Inner and Outer VPN Clients on the EUD shall be implemented on separate IP stacks. Implementations of IPv4 and IPv6 on the same operating system are considered to be part of the same IP stack.	VE	T=O	
MA-EU-6	If the EUD is not remotely administered, then it shall only be updated and rekeyed through re-provisioning.	VE, TE	T=O	
MA-EU-7	The EUD shall not allow split-tunneling.	VE, TE	T=O	
MA-EU-8	Rekeying of an EUD's certificates and associated private keys shall be done through re-provisioning prior to expiration of keys.	VE, TE	T	MA-EU-9

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-EU-9	Rekeying of an EUD's certificates and associated private keys shall be done over the MA solution network prior to expiration of keys.	VE, TE	O	MA-EU-8
MA-EU-10	An EUD shall be deauthorized from the network and submitted for Forensic Analysis if suspected of being compromised.	VE, TE	T=O	
MA-EU-11	An EUD should be destroyed only if it has been determined to be compromised through Forensic Analysis.	VE, TE	T=O	
MA-EU-12	Users of EUDs shall successfully authenticate themselves to the services they access on the Red network using an AO-approved method.	VE, TE	T=O	
MA-EU-13	Red network services shall not transmit any classified data to EUDs until user authentication succeeds.	VE, TE	T=O	
MA-EU-14	Withdrawn			
MA-EU-15	All EUD Users shall sign an organization-defined user agreement before being authorized to use an EUD.	VE, TE	T=O	
MA-EU-16	All EUD Users shall receive an organization-developed training course for operating an EUD prior to use.	VE, TE	T=O	
MA-EU-17	At a minimum, the organization-defined user agreement shall include each of the following: • Consent to monitoring • Operations Security (OPSEC) guidance • Required physical protections to employ when operating and storing the EUD • Restrictions for when, where, and under what conditions the EUD may be used • Responsibility for reporting security incidents • Verification of IA Training • Verification of appropriate clearance • Justification for Access • Requester information and organization • Account Expiration Date • User Responsibilities	VE, TE	T=O	

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-EU-18	EUDs shall be dedicated for use solely in the Mobile Access solution, and not used to access any resources on networks other than the Red network it communicates with through the two layers of encryption.	VE, TE	T=O	
MA-EU-19	EUDs shall be remotely administered.	VE, TE	O	Optional
MA-EU-20	The EUD shall disable all transmitted Global Positioning System (GPS) and location services except Enhanced 9-1-1 (E911) or those authorized by the AO.	VE, TE	T=O	
MA-EU-21	The EUD shall disable Firmware-Over-the-Air (FOTA) updates from the cellular carrier.	VE, TE	T=O	
MA-EU-22	The EUD shall disable all wireless interfaces (e.g. Bluetooth, NFC, Cellular, 802.11) that do not pass through the VPN client.	VE, TE	T=O	
MA-EU-23	The EUD shall disable processing of incoming cellular services including voice messaging services that do not pass through the VPN client.	VE, TE	O	Optional
MA-EU-24	All EUDs shall have their certificates revoked and resident image removed prior to disposal.	VE, TE	T=O	
MA-EU-25	Passwords for user to device (EUD selected from Mobile Platform section of CSfC Components List) authentication shall be a minimum of 4 alpha-numeric case sensitive characters.	VE, TE	T=O	
MA-EU-26	Withdrawn			
MA-EU-27	For a VPN EUD that has a VPN Gateway physically attached to it, the VPN Gateway shall be the Outer layer of encryption and the VPN client on the EUD will be the Inner Layer of encryption.	VE	T=O	
MA-EU-28	Withdrawn			
MA-EU-29	If the EUD is using a physically attached VPN Gateway as the Outer layer of encryption, the communication between the EUD and the VPN Gateway shall be through a wired connection (i.e. Ethernet) or Wi-Fi using WPA2.	VE, TE	T=O	
MA-EU-30	Withdrawn			

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-EU-31	If a TLS EUD is using a standalone VPN Gateway to connect over the black transport network, the VPN Gateway shall be used to establish the Outer Layer of encryption.	TE	T=O	
MA-EU-32	If an NSA-Approved DAR Solution is not implemented on EUDs, the native platform DAR protection shall be enabled.	VE, TE	T=O	
MA-EU-33	EUDs shall use a unique X.509 v3 device certificate, signed by the Outer CA, for mutual authentication with Outer VPN Gateways.	VE, TE	T=O	
MA-EU-34	TLS EUDs shall either use a unique X.509 v3 device certificate (signed by the Inner CA) or a unique X.509 v3 user certificate (signed by an authorized enterprise services CA) for mutual authentication with TLS-Protected Servers.	TE	T =O	
MA-EU-35	VPN EUDs shall use a unique X.509 v3 device certificate, signed by the Inner CA, for mutual authentication with Inner VPN Gateways.	VE	T=O	
MA-EU-36	EUDs shall use an Access Point Name (APN) provided by a Domestic Cellular Carrier Private Network when utilizing Domestic Cellular Service as a Black Transport Network.	VE, TE	O	Optional
MA-EU-37	EUDs shall be configured for all IP traffic, with the exception of IKE, network address configuration, time synchronization, and name resolution traffic required to establish the IPsec tunnel, to flow through the IPsec VPN Client.	VE, TE	T	MA-EU-38
MA-EU-38	EUDs shall be configured for all IP traffic, with the exception of IKE, to flow through the IPsec VPN Client.	VE, TE	O	MA-EU-37
MA-EU-39	The EUD maximum password lifetime shall be less than 181 days.	VE, TE	T=O	
MA-EU-40	The EUD screen shall lock after three minutes or less of inactivity.	VE, TE	T=O	
MA-EU-41	The EUD shall perform a wipe of all protected data after 10 or less authentication failures.	VE, TE	T=O	
MA-EU-42	VPN Protection shall be enabled across the EUD.	VE, TE	T=O	
MA-EU-43	A security policy shall be configured on the EUD specific to each permitted Retransmission Device and/or Government Private Wireless network.	VE, TE	T=O	

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-EU-44	During provisioning, all unnecessary keys shall be destroyed from the EUD secure key storage.	VE, TE	T=O	
MA-EU-45	During provisioning, all unnecessary X.509 certificates shall be removed from the EUD Trust Anchor Database.	VE, TE	T=O	
MA-EU-46	All display notifications shall be disabled while in a locked state.	VE, TE	O	Optional
MA-EU-47	USB mass storage mode shall be disabled on the EUDs.	VE, TE	O	Optional
MA-EU-48	USB data transfer shall be disabled on the EUDs.	VE, TE	O	Optional
MA-EU-49	Prior to updating the Application Processor system software, the system software digital signature shall be verified.	VE, TE	T=O	
MA-EU-50	Prior to installing new applications, the application digital signature shall be verified.	VE, TE	T=O	

13.9 Port Filtering Requirements for Solution Components

Table 17. Port Filtering Requirements for Solution Components

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-PF-1	All Components within the Solution shall have all network interfaces restricted to the smallest address ranges, ports, and protocols possible.	All	T=O	
MA-PF-2	All Components within the Solution shall have all unused network interfaces disabled.	All	T=O	
MA-PF-3	CDPs shall only allow inbound HTTP traffic.	C	T=O	
MA-PF-4	For the Outer VPN Gateway interface connected to a Black network, traffic filtering rules shall be applied to both inbound and outbound traffic, such that only IKE, ESP, and control plane protocols (as defined in this CP) approved by organization-defined policy are allowed.	All	T=O	

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-PF-5	For the Inner VPN Gateway interface connected to a Gray network, traffic filtering rules shall be applied to both inbound and outbound traffic, such that only IKE, ESP, and management and control plane protocols (as defined in this CP) approved by organization-defined policy are allowed.	V	T=O	
MA-PF-6	The Inner firewall shall implement an ACL which only permits ingress/egress traffic to Inner Encryption Endpoints.	All	T=O	
MA-PF-7	Any service or feature that allows an Outer VPN Gateway or an EUD to contact a third party server (such as one maintained by the manufacturer) shall be blocked.	All	T	MA-PF-8
MA-PF-8	Any service or feature that allows an Outer VPN Gateway or an EUD to contact a third party server (such as one maintained by the manufacturer) shall be disabled.	All	O	MA-PF-7
MA-PF-9	Multicast messages received on external interfaces of Outer VPN Gateway shall be dropped.	All	T=O	
MA-PF-10	For solutions using IPv4, the Outer VPN Gateway shall drop all packets that use IP options.	All	O	Optional
MA-PF-11	For solutions using IPv4, the Outer VPN Gateway shall only accept packets with Transmission Control Protocol (TCP), User Data Protocol (UDP), Encapsulating Security Payload (ESP), or ICMP in the IPv4 Protocol field and drop all other packets.	All	T=O	
MA-PF-12	For solutions using IPv6, the Outer VPN Gateway shall only accept packets with ESP, TCP, UDP, or ICMPv6 in the IPv6 Next Header field and drop all other packets.	All	T=O	
MA-PF-13	For all Outer firewall interfaces, traffic filtering rules shall be applied to both inbound and outbound traffic, such that only IKE, ESP, and control plane protocols (as defined in this CP) approved by organization-defined policy are allowed.	VI, TI	T=O	

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-PF-14	EUDs consisting of a single computing platform shall prohibit ingress and egress of Certificate Revocation traffic (e.g. OCSP queries, HTTP GET to CDPs) on the Black interface.	VE, TE	T=O	
MA-PF-15	EUDs consisting of a single computing platform shall prohibit ingress and egress of Name Resolution traffic (e.g. DNS query/response) on the Black Interface.	VE, TE	O	Optional
MA-PF-16	EUDs consisting of a single computing platform shall prohibit ingress and egress of Network Time Protocol (NTP) traffic on the Black Interface.	VE, TE	O	Optional

13.10 Configuration Change Detection Requirements

Table 18. Configuration Change Detection Requirements

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-CM-1	A baseline configuration for all components shall be maintained by the Security Administrator and be available to the Auditor.	All	T=O	
MA-CM-2	An automated process shall ensure that configuration changes are logged.	All	T=O	
MA-CM-3	All Solution components shall be configured with a monitoring service that detects all changes to configuration.	All	O	Optional

13.11 Device Management Requirements

Only authorized Security Administrators will be allowed to administer the Components. The MA solution will be used as transport for the Secure Shell (SSH)v2, IPsec, or TLS data from the Administration Workstation to the Component.

Table 19. Requirements for Device Management

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-DM-1	Administration Workstations shall be dedicated for the purposes given in the CP and shall be physically separated from workstations used to manage non-CSfC solutions.	VI, TI	T=O	
MA-DM-2	The Inner Encryption Endpoints shall be managed from the Enterprise/Red network and the Outer VPN Gateway and Gray firewall shall be managed from the Gray network.	VI, TI	T=O	
MA-DM-3	A separate LAN or VLAN on the Enterprise/Red network shall be used exclusively for all management of Inner Encryption Endpoints and solution components within the Red network.	VI, TI	T=O	
MA-DM-4	A separate LAN or VLAN on the Gray network shall be used exclusively for all management of the Outer VPN Gateway, Gray firewall, and solution components within the Gray network.	VI, TI	T=O	
MA-DM-5	The Gray Management network shall not be directly connected to Non-secure Internet Protocol Router Network (NIPRNet) or any other Unclassified network not dedicated to the administration of CSfC solutions.	VI, TI	T=O	
MA-DM-6	All administration of solution components shall be performed from an Administration Workstation • remotely using one of SSHv2, IPsec, or TLS 1.2 or later version; or • by managing the solution components locally.	VI, TI	T=O	
MA-DM-7	Security Administrators shall authenticate to solution components before performing administrative functions.	All	T	MA-DM-8
MA-DM-8	Security Administrators shall authenticate to solution components with Suite B-compliant certificates before performing administrative functions remotely.	All	O	MA-DM-7
MA-DM-9	Security Administrators shall establish a security policy for EUDs per the implementing organization's local policy.	VE, TE	T=O	

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-DM-10	EUDs shall generate logs and send to a central SIEM in the Red network.	VE, TE	O	Optional
MA-DM-11	Security Administrators shall initiate certificate signing requests for solution components as part of their initial keying within the solution.	All	T=O	
MA-DM-12	Devices shall use Enrollment over Secure Transport (EST) as detailed in IETF RFC 7030 for certificate management.	All	O	Optional
MA-DM-13	The same Administration Workstation shall not be used to manage Inner Encryption Components and the Outer VPN Gateway.	VI, TI	T=O	
MA-DM-14	The Outer VPN Gateway and solution components within the Gray network shall forward log entries to a SIEM on the Gray Management network (or SIEM in the Enterprise/Red Network if using an AO approved one-way tap) within 10 minutes.	VI, TI	T=O	
MA-DM-15	Inner Encryption Components and solution components within the Red network shall forward log entries to a SIEM on the Red Management network within 10 minutes.	VI, TI	O	Optional
MA-DM-16	All logs forwarded to a SIEM on the Gray Management network shall be encrypted using SSHv2, IPsec, or TLS 1.1 or later.	All	O	Optional
MA-DM-17	All logs forwarded to a SIEM on a Red Management network shall be encrypted using SSHv2, IPsec, or TLS 1.1 or later.	All	O	Optional
MA-DM-18	When managing Solution components over the Black network, the management traffic shall be encrypted with Suite B algorithms IAW Table 8 and Table 9.	All	T=O	

13.12 CONTINUOUS MONITORING REQUIREMENTS

Table 20. Continuous Monitoring Requirements

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-MR-1	Traffic from the Black, Gray, or Red networks shall be monitored from an Intrusion Detection System (IDS).	VI, TI	T	MA-MR-2

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-MR-2	Traffic from the Black, Gray, or Red networks shall be monitored from an Intrusion Prevention System (IPS).	VI, TI	O	MA-MR-1
MA-MR-3	An IDS shall be deployed between the Outer firewall and Outer VPN Gateway (M1), or between the Outer VPN Gateway and the Gray firewall (M2), or on the internal side of the Inner firewall (M3).	VI, TI	T	MA-MR-4 MA-MR-5 MA-MR-6
MA-MR-4	An IDS shall be deployed between the Outer firewall and Outer VPN Gateway (M1), and between the Outer VPN Gateway and the Gray firewall (M2), and on the internal side of the Inner firewall (M3).	VI, TI	O	MA-MR-3 MA-MR-5 MA-MR-6
MA-MR-5	An IPS shall be deployed between the Outer firewall and Outer VPN Gateway (M1), or between the Outer VPN Gateway and the Gray firewall (M2), or on the internal side of the Inner firewall (M3).	VI, TI	O	MA-MR-3 MA-MR-4 MA-MR-6
MA-MR-6	An IPS shall be deployed between the Outer firewall and Outer VPN Gateway (M1), and between the Outer VPN Gateway and the Gray firewall (M2), and on the internal side of the Inner firewall.	VI, TI	O	MA-MR-3 MA-MR-4 MA-MR-5
MA-MR-7	Each IDS in the solution shall be configured to send alerts to the Security Administrator.	VI, TI	T	MA-MR-8
MA-MR-8	Each IPS in the solution shall be configured to block malicious traffic flows and alert the Security Administrator.	VI, TI	O	MA-MR-7
MA-MR-9	Each IDS in the solution shall be configured with rules that generate alerts upon detection of any unauthorized destination IP addresses.	VI, TI	T	MA-MR-10
MA-MR-10	Each IPS in the solution shall be configured with rules that block and generate alerts upon detection of any unauthorized destination IP addresses.	VI, TI	O	MA-MR-9
MA-MR-11	Each IDS in the solution shall be configured with rules that generate alerts upon detection of any unauthorized source IP addresses.	VI, TI	T	MA-MR-12
MA-MR-12	Each IPS in the solution shall be configured with rules that block and generate alerts upon detection of any unauthorized source IP addresses.	VI, TI	O	MA-MR-11

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-MR-13	A SIEM component shall be placed within the Gray network unless devices are configured to push events to an Enterprise/Red network SIEM through an AO-approved one-way tap.	VI, TI	T=O	
MA-MR-14	The SIEM shall be configured to send alerts to the Auditor when anomalous behavior is detected (i.e. blocked packets from the Outer VPN Gateway or Gray firewall).	VI, TI	T=O	
MA-MR-15	The Gray SIEM shall collect logs from the Outer VPN Gateway, Gray firewall, and any components located within the Gray Management Services.	VI, TI	T=O	
MA-MR-16	Logs sent to the Gray SIEM shall be encrypted with TLS, SSHv2, or IPsec.	VI, TI	T=O	
MA-MR-17	One-way taps deployed as part of the solution shall be approved for use by the AO.	VI, TI	T=O	
MA-MR-18	One-way taps deployed as part of the solution shall only allow monitoring data to flow from Monitoring Point 2 (M2) and/or Monitoring Point 1 (M1) to an enclave at the Red level that is isolated from the Red/Enterprise Network.	VI, TI	T=O	

13.13 AUDITING REQUIREMENTS

Table 21. Auditing Requirements

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-AU-1	VPN Gateways shall log establishment of a VPN tunnel.	T, V	T=O	
MA-AU-2	TLS-Protected Servers shall log establishment of a TLS connection.	TI	T=O	
MA-AU-3	VPN Gateways shall log termination of a VPN tunnel.	T, V	T=O	
MA-AU-4	TLS-Protected Servers shall log termination of a TLS connection.	TI	T=O	
MA-AU-5	VPN Clients shall log establishment of a VPN tunnel.	VE, TE	O	Optional
MA-AU-6	TLS Clients shall log establishment of a TLS Tunnel	TE	O	Optional
MA-AU-7	VPN Clients shall log termination of a VPN tunnel.	VE, TE	O	Optional

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-AU-8	TLS Client shall log termination of a TLS Tunnel.	TE	O	Optional
MA-AU-9	Solution components shall log all actions performed on the audit log (off-loading, deletion, etc.).	VI, TI	T=O	
MA-AU-10	Solution components shall log all actions involving identification and authentication.	VI, TI	T=O	
MA-AU-11	Solution components shall log attempts to perform an unauthorized action (read, write, execute, delete, etc.) on an object.	TI,VI	T=O	
MA-AU-12	Solution components shall log all actions performed by a user with super-user or administrator privileges.	VI, TI	T=O	
MA-AU-13	Solution components shall log escalation of user privileges.	VI, TI	T=O	
MA-AU-14	Solution components shall log generation, loading, and revocation of certificates.	All	T=O	
MA-AU-15	Solution components shall log changes to time.	VI, TI	T=O	
MA-AU-16	Each log entry shall record the date and time of the event.	All	T=O	
MA-AU-17	Each log entry shall include the identifier of the event.	All	T=O	
MA-AU-18	Each log entry shall record the type of event.	All	T=O	
MA-AU-19	Each log entry shall record the success or failure of the event to include failure code, when available.	All	T=O	
MA-AU-20	Each log entry shall record the subject identity.	All	T=O	
MA-AU-21	Each log entry shall record the source address for network-based events.	All	T=O	
MA-AU-22	Each log entry shall record the user and, for role-based events, role identity, where applicable.	All	T=O	
MA-AU-23	Auditors shall detect when two or more simultaneous VPN connections from different IP addresses are established using the same EUD device certificate.	V	O	Optional
MA-AU-24	Auditors shall detect when two or more simultaneous TLS connections from different IP addresses are established using the same EUD device certificate.	T	O	Optional

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-AU-25	Upon notification of two or more simultaneous VPN connections from different IP addresses using the same EUD device certificate, the Certificate Authority Administrator shall revoke the device certificate and provide an updated CRL to the Security Administrator.	V	O	Optional
MA-AU-26	Upon notification of two or more simultaneous TLS connections from different IP addresses using the same EUD device certificate, the Certificate Authority Administrator shall revoke the device certificate and provide an updated CRL to the Security Administrator.	T	O	Optional
MA-AU-27	The Security Administrator shall immediately drop the session upon notification of two or more simultaneous VPN connections from different IP addresses using the same EUD device certificate.	V	O	Optional
MA-AU-28	The Security Administrator shall immediately drop the session upon notification of two or more simultaneous TLS connections from different IP addresses using the same EUD device certificate.	T	O	Optional
MA-AU-29	VPN Gateways shall log the failure to download a CRL from a CDP.	VI	T=O	
MA-AU-30	TLS-Protected Servers shall log the failure to download a CRL from a CDP.	TI	T=O	
MA-AU-31	VPN Gateways shall log if the version of the CRL downloaded from a CDP is older than the current cached CRL.	VI	T=O	
MA-AU-32	TLS-Protected Servers shall log if the version of the CRL downloaded from a CDP is older than the current cached CRL.	TI	T=O	
MA-AU-33	VPN Gateways shall log if signature validation of the CRL downloaded from a CDP fails.	VI	T=O	
MA-AU-34	TLS-Protected Servers shall log if signature validation of the CRL downloaded from a CDP fails.	TI	T=O	
MA-AU-35	Auditors shall compare and analyze collected network flow data against the established baseline on at least a weekly basis.	VI, TI	O	Optional

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-AU-36	Locally-run CAs shall comply with the audit and archival requirements defined in IETF RFC 3647 Sections 4.5.4 and 4.5.5, respectively.	VI, TI	T=O	
MA-AU-37	Locally-run CAs shall comply with periodic audit and assessment requirements defined in IETF RFC 3647 Section 4.8.	VI, TI	T=O	
MA-AU-38	Audits and assessments for Outer and Inner CAs shall be performed by personnel who are knowledgeable in the CAs' operations, as well as the CAs' CP and CPS requirements and processes, respectively.	VI, TI	T=O	

13.14 Key Management Requirements

13.14.1 General Requirements

Table 22. PKI General Requirements

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-KM-1	User certificates and user private keys shall be classified to the level determined by the AO.	TE, VE	T=O	
MA-KM-2	Outer CAs shall provide services through either the Gray or Red network.	VI, TI	T = O	
MA-KM-3	Inner CAs shall provide services through the Enterprise/Red network.	VI, TI	T=O	
MA-KM-4	Locally run Inner Tunnel CAs shall be physically separate from locally-run Outer Tunnel CAs.	VI, TI	T=O	
MA-KM-5	All certificates issued by the Outer and Inner CAs for the MA Solution shall be Non-Person Entity (NPE) certificates, except in the one case when a TLS EUD requires a user certificate for the inner TLS tunnel.	VI, TI	T=O	
MA-KM-6	All certificates issued by the Outer and Inner CAs for the MA Solution shall be used for authentication only.	VI, TI	T=O	
MA-KM-7	Authentication certificates issued by the Outer and Inner CAs for the MA Solution shall be X.509 v3 certificates as defined in ITU-T Recommendation X.509.	VI, TI	T=O	

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-KM-8	Authentication certificate profiles for the Outer and Inner CAs for the MA Solution shall comply with IETF RFC 5280.	VI, TI	T=O	
MA-KM-9	All device certificates issued by the Outer and Inner CAs, and their corresponding private keys, shall be treated as CUI (or higher as determined by the AO).	All	T=O	
MA-KM-10	The key sizes and algorithms for CA certificates and authentication certificates issued to Outer VPN Components, Inner Encryption Components, and Administrative Device Components shall be as specified in CNSSP 15 (See Table 8-.	All	T=O	
MA-KM-11	Outer and Inner CAs shall not have access to private keys used in the MA Solution Components.	All	T=O	
MA-KM-12	Private keys associated with on-line, locally run Outer and Inner CAs shall be protected using Hardware Security Modules (HSMs) validated to at least FIPS 140-2 Level 2. "On-line" means the CA is always powered on and network-accessible.	VI, TI	T=O	
MA-KM-13	Outer and Inner CAs shall operate in compliance with a Certificate Policy and Certification Practices Statement that is formatted in accordance with Internet Engineering Task Force (IETF) Request for Comments (RFC) 3647.	VI, TI	T=O	

13.14.2 CERTIFICATE ISSUANCE REQUIREMENTS

Table 23. Certificate Issuance Requirements

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-KM-14	Outer VPN Components, Inner Encryption Components, and Gray and Red Management Services Components shall be initially keyed and loaded with certificates within a physical environment certified to protect the highest classification level of the MA solution network.	All	T=O	

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-KM-15	Private keys for Outer VPN Components, Inner Encryption Components and Gray and Red Management Services Components shall never be escrowed.	All	T=O	
MA-KM-16	Outer and Inner CAs shall use Public Key Cryptographic Standard (PKCS)#10 and PKCS#7 to issue authentication certificates to Outer VPN Components, Inner Encryption Components, and Gray and Red Management Services Components.	VI, TI	T	MA-KM-19
MA-KM-17	Red and Gray Management Services shall use PKCS#12 for installing certificates/keys to EUDs.	All	T	MA-KM-18
MA-KM-18	Red and Gray Management Services shall use PKCS#7 for installing certificates to EUDs.	All	O	MA-KM-17
MA-KM-19	Outer and Inner CAs shall use IETF RFC 7030 Enrollment over Secure Transport (EST) to issue authentication certificates to Outer VPN Components, Inner Encryption Components, and Gray and Red Management Services Components.	All	O	MA-KM-16
MA-KM-20	Certificate requests for Outer VPN Components, Inner Encryption Components and Gray and Red Management Services Components shall be submitted to the CA in accordance with the CA's Certificate Policy (CP) and Certification Practices Statement (CPS).	All	T=O	
MA-KM-21	Outer and Inner CAs shall issue certificates in accordance with their Certificate Policies and CPSs.	All	T=O	
MA-KM-22	Certificate Policies and CPSs for non-Enterprise, locally run CAs shall ensure the CAs issue certificates within a defined and limited name space and assert: • Unique Distinguished Names (DNs) • Appropriate key usages • A registered policy Object Identifier (OID)	All	T=O	

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-KM-23	Inner and Outer CAs shall assert at least one CRL Distribution Point (CDP) Uniform Resource Locater (URL) in certificates issued to Solution Infrastructure Outer VPN Gateways, Inner Encryption Components, and Gray and Red Management Services Components. The CDP URL specifies the location of the CAs' CRLs.	All	T=O	
MA-KM-24	The key validity period for certificates issued by non-Enterprise, locally run CAs to MA End User Devices shall not exceed 14 months.	All	T=O	
MA-KM-25	The key validity period for certificates issued by non-Enterprise, locally run CAs to MA Solution Infrastructure Components shall not exceed 36 months.	All	T=O	
MA-KM-26	Inner CAs shall only issue certificates to Inner Encryption Components and Red Network Components of MA Solutions.	All	T=O	
MA-KM-27	Outer CAs shall only issue certificates to Outer VPN Components and Gray Network Components of MA Solutions.	All	T=O	
MA-KM-28	Withdrawn			
MA-KM-29	Withdrawn			

13.14.3 CERTIFICATE RENEWAL AND REKEY REQUIREMENTS

Table 24. Certificate Renewal and Rekey Requirements

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-KM-30	Certificate renewal or rekey shall occur prior to a certificate expiring. If renewal/rekey occurs after a certificate expires, then the initial certificate issuance process is used to renew/rekey the certificate.	All	T=O	
MA-KM-31	Certificate renewal or rekey shall be performed in accordance with the CA's Certificate Policy and CPS.	All	T=O	
MA-KM-32	Inner and Outer CAs shall issue renewed/ rekeyed authentication certificates to Solution Components using PKCS#10 and PKCS#7.	All	T	MA-KM-35
MA-KM-33	Withdrawn			
MA-KM-34	Withdrawn			

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-KM-35	Inner and Outer CAs shall issue renewed/rekeyed authentication certificates to Solution Components using EST (RFC 7030).	All	O	MA-KM-32

13.14.4 CERTIFICATE REVOCATION AND CDP REQUIREMENTS

Table 25. Requirements for Certificate Revocation and CDPs

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-KM-36	Inner and Outer CAs shall revoke a certificate issued to Solution Components when the binding between the subject information and public key within the certificate issued is no longer considered valid.	VI, TI	T=O	
MA-KM-37	Inner and Outer CAs shall make certificate revocation information available in the form of CRLs signed by the CAs.	VI, TI	T=O	
MA-KM-38	CRLs shall be X.509 v2 CRLs as defined in ITU-T Recommendation X.509.	C	T=O	
MA-KM-39	CRL profiles shall comply with IETF RFC 5280.	C	T=O	
MA-KM-40	Procedures for requesting certificate revocation shall comply with the CA's Certificate Policy and Certification Practices Statement.	All	T=O	
MA-KM-41	Certificate Policies and CPSs for non-Enterprise, locally run CAs shall ensure revocation procedures address the following: • Response for a lost, stolen or compromised MA EUD • Removal of a revoked infrastructure device (i.e., VPN Gateway) from the network • Re-establishment of an MA Solution Component whose certificate was revoked • Revocation of certificates due to compromise of an MA EUD • Revocation of an authentication certificate if simultaneous use of the certificate is detected from different IP addresses	All	T=O	

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-KM-42	Inner and Outer CAs shall make CRLs available to authorized CRL Distribution Points (CDPs), so that the CRLs can be accessed by Solution Components.	C	T=O	
MA-KM-43	Enterprise CAs shall create and publish CRLs in accordance with the Enterprise CAs' Certificate Policies and CPSs.	VI, TI	T=O	
MA-KM-44	Non-enterprise, locally run CAs shall publish new CRLs at least once every 28 days.	VI, TI	T=O	
MA-KM-45	Non-enterprise, locally run CAs shall create a new CRL within one hour of a certificate being revoked.	VI, TI	T=O	
MA-KM-46	Solution Infrastructure Components shall have access to new certificate revocation information within 24 hours of the CA creating a new CRL.	VI, TI	T=O	
MA-KM-47	Non-enterprise, locally run CAs shall ensure that newly created CRLs are published at least 7 days prior to the expiration of the current CRLs.	VI, TI	T=O	
MA-KM-48	The Solution shall provide certificate revocation status information via an Online Certificate Status Protocol (OCSP) Server on the Red and Gray network that is compliant with IETF RFC 6960.	VI, TI	O	Optional
MA-KM-49	Certificate revocation status messages delivered by an OCSP server shall be digitally signed and compliant with IETF RFC 6960.	VI, TI	O	Optional

14 REQUIREMENTS FOR SOLUTION OPERATION, MAINTENANCE, AND HANDLING

14.1 REQUIREMENTS FOR THE USE AND HANDLING OF SOLUTIONS

The following requirements shall be followed regarding the use and handling of the solution.

Table 26. Requirements for the Use and Handling of Solutions

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-GD-1	All Solution Infrastructure components shall be physically protected as classified devices, classified at the level of the Enterprise/Red network.	VI, TI	T=O	
MA-GD-2	Only authorized and appropriately cleared (or escorted) administrators and security personnel shall have physical access to the solution Infrastructure components.	VI, TI	T=O	
MA-GD-3	Only authorized and appropriately cleared users, administrators, and security personnel shall have physical access to EUDs when in a classified state.	VE, TE	T=O	
MA-GD-4	All components of the solution shall be disposed of as classified devices, unless declassified using AO-approved procedures.	All	T=O	
MA-GD-5	EUDs using an NSA-approved DAR solution shall be disposed of in accordance with the disposal requirements for the DAR solution.	VE, TE	T=O	
MA-GD-6	All EUDs shall have their certificates revoked prior to disposal.	VE, TE	T=O	
MA-GD-7	Users shall periodically inspect the physical attributes of EUDs for signs of tampering or other unauthorized changes.	VE, TE	T=O	
MA-GD-8	Acquisition and procurement documentation shall not include information about how the equipment will be used, to include that it will be used to protect classified information.	All	T=O	
MA-GD-9	The solution owner shall allow, and fully cooperate with, NSA or its authorized agent to perform an IA compliance audit (including, but not limited to, inspection, testing, observation, interviewing) of the solution implementation to ensure it meets the latest version of the CP.	All	T=O	
MA-GD-10	The AO will ensure that a compliance audit shall be conducted every year against the latest version of the MA CP as part annual solution re-registration process.	All	T=O	
MA-GD-11	Results of the compliance audit shall be provided to and reviewed by the AO.	All	T=O	

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-GD-12	Customers interested in registering their solution against the MA CP shall register with NSA and receive approval prior to AO authorization to operate.	All	T=O	
MA-GD-13	The implementing organization shall complete and submit an MA CP requirements compliance matrix to their respective AO.	All	T=O	
MA-GD-14	Registration and re-registration against the MA CP shall include submission of MA CP registration forms and compliance matrix to NSA.	All	T=O	
MA-GD-15	When a new approved version of the MA CP is published by NSA, the AO shall ensure compliance against this new CP within 6 months.	All	T=O	
MA-GD-16	Solution implementation information, which was provided to NSA during solution registration, shall be updated annually (in accordance with Section 16.3) as part annual solution re-registration process.	All	T=O	
MA-GD-17	Audit log data shall be maintained for a minimum of 1 year.	All	T=O	
MA-GD-18	The amount of storage remaining for audit events shall be assessed quarterly in order to ensure that adequate memory space is available to continue recording new audit events.	All	T=O	
MA-GD-19	Audit data shall be frequently off-loaded to a backup storage medium.	All	T=O	
MA-GD-20	A set of procedures shall be developed by the implementing organization to provide guidance for identifying and reporting security incidents associated with the audit events to the proper authorities and to the data owners.	All	T=O	
MA-GD-21	The implementing organization shall develop a continuity of operations plan for auditing capability which includes a mechanism or method for determining when the audit log is reaching its maximum storage capacity.	All	T=O	

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-GD-22	The implementing organization shall develop a continuity of operations plan for auditing capability which includes a mechanism or method for off-loading audit log data for long-term storage.	All	T=O	
MA-GD-23	The implementing organization shall develop a continuity of operations plan for auditing capability which includes a mechanism or method for responding to an overflow of audit log data within a product.	All	T=O	
MA-GD-24	The implementing organization shall develop a continuity of operations plan for auditing capability which includes a mechanism or method for ensuring that the audit log can be maintained during power events.	All	T=O	
MA-GD-25	Strong passwords shall be used that comply with the requirements of the AO.	VI, TI	T=O	
MA-GD-26	Security critical patches shall be tested and subsequently applied to all components in the solution in accordance with local policy and this CP.	All	T=O	
MA-GD-27	Local policy shall dictate how the Security Administrator will install patches to solution components.	All	T=O	
MA-GD-28	Solution components shall comply with local TEMPEST policy.	All	T=O	
MA-GD-29	Software, settings, keys, and all other configuration data persistently stored on EUDs shall be handled as controlled unclassified information or higher classification.	All	T=O	
MA-GD-30	All hardware components shall be tracked through an AO-approved inventory management process that identifies each component as part of a CSfC Solution.	All	T=O	

Additional MA-GD requirements can be found in Section 15.

14.2 Requirements for Incident Reporting

Table 27 lists requirements for reporting security incidents to NSA to be followed in the event that a solution owner identifies a security incident which affects the solution. These reporting requirements are intended to augment, not replace, any incident reporting procedures already in use within the solution owner's organization. It is critical that Security Administrators, Certificate Authority

Administrators (CAAs), and Auditors are familiar with maintaining the solution in accordance with this CP. Based on familiarity with the known-good configuration of the solution, personnel responsible for the operations and maintenance of the solution will be better equipped to identify reportable incidents.

For the purposes of incident reporting, "malicious" activity includes not only events that have been attributed to activity by an adversary but also any events that are unexplained. In other words, an activity is assumed to be malicious unless it has been determined to be the result of known non-malicious activity.

Table 27 only provides requirements directly related to the incident reporting process. See Section 13.12 for requirements supporting the detection of events that may reveal that a reportable incident has occurred.

Table 27. Incident Reporting Requirements

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-RP-1	Solution owners shall report confirmed incidents meeting the criteria in MA-RP-3 through MA-RP-16 within 24 hours of detection via Joint Incident Management System (JIMS) or contacting NSA as specified in the CSfC Registration Letter issued for the solution.	All	T=O	
MA-RP-2	At a minimum, the organization shall provide the following information when reporting security incidents: • CSfC Registration Number • Point of Contact (POC) name, phone, email • Alternate POC name, phone, email • Classification level of affected solution • Name of affected Network(s) • Affected component(s) manufacturer/vendor • Affected component(s) model number • Affected component(s) version number • Date and time of incident • Description of incident • Description of remediation activities Is Technical Support from NSA requested? (Yes/No)	All	T=O	
MA-RP-3	Solution owners shall report a security failure in any of the CSfC solution components.	All	T=O	

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-RP-4	Solution owners shall report any evidence of a compromise or spillage of classified data caused by a failure of the CSfC Solution.	All	T=O	
MA-RP-5	For all Gray network interfaces, solution owners shall report any malicious inbound and outbound traffic.	All	T=O	
MA-RP-6	Solution owners shall report any evidence of an unauthorized device/user gaining access to the classified network via the solution.	All	T=O	
MA-RP-7	Solution owners shall report if a solution component sends traffic with an unauthorized destination address.	All	T=O	
MA-RP-8	Solution owners shall report any malicious configuration changes to the components.	All	T=O	
MA-RP-9	Solution owners shall report any unauthorized escalation of privileges to any of the CSfC solution components.	All	T=O	
MA-RP-10	Solution owners shall report if two or more simultaneous VPN connections from different IP addresses are established using the same EUD device certificate.	All	T=O	
MA-RP-11	Solution owners shall report any evidence of malicious physical tampering with solution components.	All	T=O	
MA-RP-12	Solution owners shall report any evidence that one or both of the layers of the solution failed to protect the data.	All	T=O	
MA-RP-13	Solution owners shall report any significant degradation of services provided by the solution.	All	T=O	
MA-RP-14	Solution owners shall report malicious discrepancies in the number of VPN connections established by Outer VPN Gateways.	All	T=O	
MA-RP-15	Solution owners shall report malicious discrepancies in the number of VPN connections established by the Inner VPN Gateway	All	T=O	
MA-RP-16	Solution owners shall report malicious discrepancies in the number of TLS connections established by the TLS-Protected Server	T	T=O	

15 ROLE-BASED PERSONNEL REQUIREMENTS

The roles required to administer and maintain the solution are defined below, along with doctrinal requirements for these roles.

Security Administrator – The Security Administrator shall be responsible for maintaining, monitoring, and controlling all security functions for the entire suite of products composing the MA solution. Security Administrator duties include but are not limited to the following:

1) Ensuring that the latest security-critical software patches and updates (such as Information Assurance Vulnerability Alerts (IAVAs)) are applied to each product.

2) Documenting and reporting security-related incidents to the appropriate authorities.

3) Coordinating and supporting product logistic support activities including integration and maintenance. Some logistic support activities may require that the Security Administrator escort uncleared personnel.

4) Employing adequate defenses of auxiliary network devices to enable proper and secure functionality of the MA solution.

5) Ensuring that the implemented MA solution remains compliant with the latest version of this CP.

6) Provisioning and maintaining EUDs in accordance with this CP for implementations that include them.

Certificate Authority Administrator (CAA) – The CAA shall be responsible for maintaining, monitoring, and controlling all security functions for the CA products. CAA duties include but are not limited to the following:

1) Administering the CA, including authentication of all components requesting certificates.

2) Maintaining and updating the CRL.

3) Provisioning and maintaining EUD certificates in accordance with this CP for implementations that include them.

Auditor – The Auditor shall be responsible for reviewing the actions performed by the Security Administrator and CAA and events recorded in the audit logs to ensure that no action or event represents a compromise to the security of the MA solution. Auditor duties include but are not limited to the following:

1) Reviewing, managing, controlling, and maintaining security audit log data.

2) Documenting and reporting security-related incidents to the appropriate authorities.

3) The Auditor will only be authorized access to Outer and Inner administrative components.

Solution Integrator – In certain cases, an external integrator may be hired to implement an MA solution based on this CP. Solution Integrator duties may include but are not limited to:

1) Acquiring the products that compose the solution.

2) Configuring the MA solution in accordance with this CP.

3) Documenting, testing, and maintaining the solution.

4) Responding to incidents affecting the solution.

End User –An End User may operate an EUD from physical locations not owned, operated, or controlled by the government. The End User shall be responsible for operating the EUD in accordance with this CP and an organization-defined user agreement. Remote User duties include, but are not limited to the following:

1) Ensuring the EUD is only operated in physical spaces which comply with the end user agreement.

2) Alerting the Security Administrator immediately upon a EUD being lost, stolen, or suspected of being tampered with.

Additional policies related to the personnel that perform these roles in an MA Solution are as follows:

Table 28. Role-Based Personnel Requirements

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-GD-31	The Security Administrator, CAAs, Auditor, EUD User, and Solution Integrators shall be cleared to the highest level of data protected by the Solution. When an Enterprise CA is used in the solution, the CAA already in place may also support this solution, provided they meet this requirement.	All	T=O	
MA-GD-32	The Security Administrator, CAA, and Auditor roles shall be performed by different people.	All	T=O	
MA-GD-33	All Security Administrators, CAAs, EUD Users, and Auditors shall meet local Information Assurance (IA) training requirements.	All	T=O	
MA-GD-34	The CAA(s) for the Inner tunnel shall be different individuals from the CAA(s) for the Outer tunnel.	All	O	Optional

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-GD-35	Upon discovering an EUD is lost or stolen, an EUD User shall immediately report the incident to their Security Administrator and Certificate Authority Administrator.	VE, TE	T=O	
MA-GD-36	Upon notification of a lost or stolen EUD, the Certificate Authority Administrators shall revoke that EUD's certificates.	VE, TE	T=O	
MA-GD-37	The Security Administrator(s) for the Inner Encryption Endpoints and supporting components on Enterprise/Red networks shall be different individuals from the Security Administrator(s) for the Outer VPN Gateway and supporting components on Gray networks.	All	T=O	
MA-GD-38	Administrators shall periodically inspect the physical attributes of infrastructure hardware for signs of tampering or other unauthorized changes.	All	O	Optional
MA-GD-39	The Auditor shall review all logs specified in this CP at least once a week.	All	T=O	
MA-GD-40	Security Administrators shall initiate the certificate revocation process prior to disposal of any solution component.	All	T=O	
MA-GD-41	Auditing of the Outer and Inner CA operations shall be performed by individuals who were not involved in the development of the CP and CPS, or integration the MA solution.	All	T=O	

16 INFORMATION TO SUPPORT AO

This section details items that likely will be necessary for the customer to obtain approval from the system AO. The customer and AO have obligations to perform the following:

- The customer, possibly with support from a System Integrator, instantiates a solution implementation that follows the NSA-approved CP.
- The customer has a testing team develop a test plan and perform testing of the MA solution, see Section 16.1.
- The customer has system certification and accreditation performed using the risk assessment information referenced in Section 16.2.

- The customer provides the results from testing and system certification and accreditation to the AO for use in making an approval decision. The AO is ultimately responsible for ensuring that all requirements from the CP have been properly implemented in accordance with the CP.
- The customer registers the solution with NSA and re-registers yearly to validate its continued use as detailed in Section 16.3.
- Customers who want to use a variant of the solution detailed in this CP will contact their NSA/IAD Client Advocate to determine ways to obtain NSA approval.
- The AO will ensure that a compliance audit shall be conducted every year against the latest version of the MA CP, and the results shall be provided to the AO.
- The AO will ensure that certificate revocation information is updated on all the Solution Components in the solution in the case of a compromise.
- The AO will ensure that any Layer 2 or Layer 3 control plane protocols that are used in the solution are necessary for the operation of the network and that local policy supports their use.
- The AO will report incidents affecting the solution in accordance with Section 14.2.

The system AO maintains configuration control of the approved solution implementation over the lifecycle of the solution. Additionally, the AO shall ensure that the solution remains properly configured with all required security updates implemented.

16.1 Solution Testing

This section provides a framework for a Test and Evaluation (T&E) plan and procedures to validate the implementation of an MA solution. This T&E will be a critical part of the approval process for the AO, providing a robust body of evidence that shows compliance with this CP.

The security features and operational capabilities associated with the use of the solution shall be tested. The following is a general high-level methodology for developing the test plan and procedures and for the execution of those procedures to validate the implementation and functionality of the MA solution. The entire solution, to include each component described in Section 5 and 6, is addressed by this test plan including the following:

1) Set up the baseline network and configure all components.
2) Document the baseline network configuration. Include product model and serial numbers, and software version numbers at a minimum.

3) Develop a test plan for the specific implementation using the test requirements from Section 17. Any additional requirements imposed by the local AO should also be tested, and the test plan shall include tests to ensure that these requirements do not interfere with the security of this solution as described in this CP.

4) Perform testing using the test plan derived in Step 3. Network testing will consist of both Black box testing and Gray box testing. A two-person testing approach should be used to administer the tests. During test execution, security and non-security related discrepancies with the solution shall be documented.

5) Compile findings, to include comments and vulnerability details as well as possible countermeasure information, into a Final Test Report to be delivered to the AO for approval of the solution.

The following testing requirement has been developed to ensure that the MA solution functions properly and meets the configuration requirements from Section 13. Testing of these requirements should be used as a minimum framework for the development of the detailed test plan and procedures.

Table 29. Test Requirements

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-TR-1	The organization implementing the CP shall perform all tests listed in Section 17.	All	T=O	

16.2 Risk Assessment

The risk assessment of the MA solution presented in this CP focuses on the types of attacks that are feasible against this solution and the mitigations that can be employed. Customers should contact their NSA/IAD Client Advocate to request this document, or visit the Secret Internet Protocol Router Network (SIPRNet) CSfC site for information. The process for obtaining the risk assessment is available on the SIPRNet CSfC website. The AO shall be provided a copy of the NSA risk assessment for their consideration in approving the use of the solution.

16.3 Registration of Solutions

All customers using CSfC solutions to protect information on National Security Systems shall register their solution with NSA prior to operational use. This registration will allow NSA to track where MA CP solutions are instantiated and to provide the AOs at those sites with appropriate information, including any significant vulnerabilities that may be discovered in components or high-level designs approved for these solutions. The CSfC solution registration process is available at http://www.nsa.gov/ia/programs/csfc_program.

Solution registrations are valid for one year from the date the solution registration is approved, at which time customers are required to re-register their solution in order to continue using it. Approved CPs will be reviewed twice a year, or as events warrant. Registered users of this CP will be notified when an updated version is published. When a new version of this CP that has been approved by the IAD Director is published, customers will have six months to bring their solutions into compliance with the new version of the CP and re-register their solution (see requirement MA-GD-15). Customers are also required to update their registrations whenever the information provided on the registration form changes.

17 TESTING REQUIREMENTS

This section contains the specific tests that allow the Security Administrator or System Integrator to ensure that they have properly configured the solution. As defined in Section 11, in order to comply with this CP, a solution must at minimum implement all Threshold requirements associated with each of the capabilities it supports, and should implement the Objective requirements associated with those capabilities where feasible. These tests may also be used to provide evidence to the AO regarding compliance of the solution with this CP. Note that the details of the procedures are the responsibility of the final developer of the solution test plan in accordance with AO-approved network procedures. The AO is ultimately responsible for ensuring that all requirements from the CP have been properly implemented.

17.1 PRODUCT SELECTION

This section contains procedures to verify that the components in this CP were selected to ensure independence in several important features.

Requirements being tested: MA-PS-1 through MA-PS-25

Procedure Description:

1) For the Inner and Outer VPN Components, perform the following:

   a) Verify that the Inner and Outer VPN Gateways are on the list of IPsec VPN Gateways on the CSfC Components List. (MA-PS-1 and MA-PS-2)

   b) Verify that the Inner and Outer VPN Clients are on the list of IPsec VPN Clients on the CSfC Components List. (MA-PS-3 and MA-PS-4)

2) For the Outer VPN Gateway and Inner Encryption Components, perform the following:

a) Verify that the Outer VPN Gateway and the Inner Encryption Endpoints either come from different independent manufacturers or that NSA has determined that sufficient implementation independence exists. (MA-PS-16)

b) Verify that the Outer VPN Gateway and Inner Encryption Components are logically separated using an NSA-approved mechanism. (MA-PS-18)

c) Verify that the Outer VPN Gateway and the Inner Encryption Endpoints are not using the same Operating System. (MA-PS-20)

d) Verify that the cryptographic libraries used by the Outer VPN Gateway and the Inner Encryption Components either come from different independent manufacturers or that NSA has determined that sufficient implementation independence exists. (MA-PS-25)

3) For each TLS-Protected Server, perform the following:

a) Verify that each TLS-Protected Server is on the CSfC Components list. (MA-PS-10)

4) For each SRTP Endpoint, perform the following:

a) Verify that each SRTP Endpoint is on the CSfC Components list. (MA-PS-12)

5) For each CA, perform the following:

a) Verify the Inner and Outer CAs came from the list of CAs on the CSfC Components List or are Enterprise CAs. (MA-PS-5)

b) Verify that the Inner and Outer CAs either come from different independent manufacturers or that NSA has determined that sufficient implementation independence exists. (MA-PS-21)

c) Verify that the cryptographic libraries used by the Inner and Outer CAs either come from different independent manufacturers or that NSA has determined that sufficient implementation independence exists. (MA-PS-24)

6) For the IPS, perform the following:

a) Verify that the IPS is from the list of IPS's on the CSfC Components Lists. (MA-PS-7)

7) For the SIP Server, perform the following:

a) Verify that the SIP Server is from the list of SIP Servers on the CSfC Components List. (MA-PS-11)

8) For the Outer, Gray, and Inner firewall, perform the following:

a) Verify that the firewalls are on the list of firewalls on the CSfC Components List. (MA-PS-13)

9) For each EUD, perform the following:

   a) Verify that the Mobile Platform EUD platform is on the list of Mobile Platforms on the CSfC Components List. (MA-PS-6)

   b) Verify that the TLS Client is chosen from the VoIP, Email, or Web Browser Applications from the CSfC Components List. (MA-PS-8)

   c) If using an SRTP Client, verify that the SRTP Client is chosen from the VoIP Applications from the CSfC Components List. (MA-PS-9)

   d) Verify that the EUD's Outer VPN Component and the Inner Encryption Components either come from different independent manufacturers or that NSA has determined that sufficient implementation independence exists. (MA-PS-23)

10) For the MDM, perform the following:

   a) Verify that the MDM is from the list of MDMs on the CSfC Components List. (MA-PS-14)

11) Verify that the Outer VPN Gateway and the Outer firewall are logically separated using an NSA-approved mechanism. (MA-PS-17)

12) Verify that the Inner Encryption Endpoints and the Gray firewall are logically separated using an NSA-approved mechanism. (MA-PS-19)

13) Verify that the Gray network firewall and the Inner Encryption Endpoints either come from different independent manufacturers or that NSA has determined that sufficient implementation independence exists. (MA-PS-22)

14) For all components used, review the mitigations in the Product Supply Chain Threat Assessment. Ensure that mitigations identified in the assessment are implemented according to the implementing organization's AO-approved Product Supply Chain Threat Assessment process. (MA-PS-26)

Expected Result:

The results of the inspection should reveal that the MA Solution components conform to the MA CP.

17.2 Physical Layout of Solution

This section contains a procedure to create an accurate record of the physical components composing the MA solution (including workstations, VPN Gateways, Inner Encryption Endpoints, CAs, and wiring). The test will also ensure that the physical implementation of the MA solution matches one of the high-level designs given in the MA CP.

Requirements being tested: MA-SR-1, MA-SR-7, MA-SR-8

Procedure Description:

1) Ensure that there are no wireless or physical connection to the solution that are not included in this CP, which may allow for traffic to leave a Red or Gray network in a manner that does not go through the MA solution (or an NSA-certified encryptor). (MA-SR-7)

2) Verify that the physical location of any network services for the Outer VPN Gateway is located on the appropriate Gray Management network. Similarly, these components for the Inner Encryption Endpoints will be located on the appropriate Red network. (MA-SR-1)

3) Verify that when using multiple Inner Encryption Components that those components are placed parallel between the Gray firewall and Inner firewall. (MA-SR-8)

Expected Result:

For Step 1, there should be no extraneous wireless or physical connections allowing data to leave Red or Gray networks besides through the MA solution (or an NSA-certified encryptor). For Step 2, Gray Management network traffic should be separate from Gray network traffic. For Step 3, Inner Encryption Components are accurately placed.

17.3 TLS-Protected Server Configurations

This section contains a procedure to ensure that the configurations for TLS-Protected Server in the MA solution follow the requirements given in this CP.

Requirements being tested: MA-TE-1 through MA-TE-9

Procedure Description:

1) For each TLS-Protected Server verify in the policy the following:

   a) Obtain the current configuration for the TLS-Protected Server.

   b) Verify that X.509 device certificates are loaded. (MA-TE-3)

   c) Verify that a unique device certificate is loaded with the corresponding CA signing certificate. (MA-TE-6)

   d) Verify the requirements MA-TE-1 through MA-TE-2, MA-TE-4, MA-TE-5, and MA-TE-7 through MA-TE-9 are configured properly.

Expected Result:

For Steps 1a- 1d, all TLS-Protected Servers shall be configured properly according to the requirements found in this CP.

17.4 End User Device Configurations

This section contains a procedure to ensure that the configurations for all the EUDs in the MA solution follow the requirements given in this CP.

Requirements being tested: MA-EU-1 through MA-EU-50, MA-RD-1 through MA-RD-4, MA-DM-9

Procedure Description:

1) For the EUD, verify in the policy the following:

   a) Ensure that, if the implementing organization's policy allows local storage of user data of classified information on the EUD, the EUDwill be treated as classified if it does not have a NSA-approved DAR. (MA-EU-1)

   b) Inspect the implementing organization's policy that EUDs which implement an NSA-approved DAR solution comply with handling requirements specified for the DAR solution. (MA-EU-2)

   c) Verify through policy that for EUDs which do not allow for local user storage that the EUD is treated as unclassified when powered down. (MA-EU-3)

   d) Verify the EUD does not allow split tunneling. (MA-EU-7)

   e) Ensure that the implementing organization's policy states that all Remote Access users must sign an organization-defined user agreement prior to using an EUD. (MA-EU-15)

   f) Verify that the implementing organization has a training program in place for Remote Access users operating an EUD. (MA-EU-16)

   g) Verify that the implementing organization has a user agreement for EUD users and also policies for each element within their user agreement document. (MA-EU-17)

   h) Ensure that the implementing organization's policy states that the EUD is dedicated for use within the MA Solution. (MA-EU-18)

   i) Ensure that the implementing organization's policy states that the EUD is to be remotely administered. (MA-EU-19)

   j) Ensure that the implementing organization's policy states that the EUDs will have their certificates revoked and resident image removed prior to disposal. (MA-EU-24)

k) Verify that the Security Administrator has an organizational security policy for the EUDs. (MA-DM-9)

2) For the EUD, perform the following:

a) Inspect the EUD configuration setting to verify that Firmware-Over-the-Air updates are disabled from the cellular carrier. (MA-EU-21)

b) Inspect the EUD configuration setting has incoming cellular services are disabled. (MA-EU-23)

c) Inspect the EUD configuration setting to verify that the wireless interfaces not passing through the VPN Client are disabled. (MA-EU-22)

d) Verify that Red network services do not transmit any classified data to the EUD until user authentication succeeds. (MA-EU-12, MA-EU-13)

e) Inspect the EUD's configuration to ensure that Global Positioning System and location services are disabled except for those authorized by the AO. (MA-EU-20)

3) For each EUD that directly connects to a Black network, perform the following:

a) Inspect the Outer VPN Client and Inner Encryption Clients on the EUD and verify that separate private key stores are used. (MA-EU-4)

b) Verify that the Inner and Outer VPN Clients on the EUD are implemented on separate IP stacks, and that the two IP stacks are not the IPv4 and IPv6 implementations on the same operating system. (MA-EU-5)

4) If the EUD is not remotely administrated, verify that the procedure given in MA-EU-6 is followed and/or is currently in place.

5) Verify that the procedures given in MA-EU-8, MA-EU-9, MA-EU-10, MA-EU-11, MA-EU-27, MA-EU-29, MA-EU-31 and MA-EU-32 are followed and/or currently in place.

6) Verify that the password length for Mobile Platform complies with MA-EU-25.

7) For Solutions that are using a Retransmission Device, perform the following:

a) Attempt to connect an unauthorized RD to the EUD to verify that the EUD will not connect to the RD. (MA-RD-1).

b) Ensure that the only connection between the EUD and the Mobile Access Solution is only via Wi-Fi or Ethernet. (MA-RD-2)

c) If the RD is using Wi-Fi, verify through configuration file that the Wi-Fi network is using WPA2 PSK or certificate-based authentication. (MA-RD-3)

d) Verify that the placement of the RD between the Outer VPN Gateway and the Inner Encryption Component. (MA-RD-4)

8) Verify the EUDs use a unique X.509 v3 device certificate, signed by the Outer CA for mutual authentication with Outer VPN Gateways. (MA-EU-33)

9) Verify TLS EUDs use a unique X.509 v3 device certificate, signed by either the Inner CA, or a unique X.509 v3 user certificate signed by an authorized enterprises service CA for mutual authentication with TLS-Protected Servers. (MA-EU-34)

10) Verify VPN EUDs use a unique X.509 v3 device certificate, signed by the Inner CA for mutual authentication with Inner VPN Gateways. (MA-EU-35)

11) Verify that, during provisioning:

a) All unnecessary keys are destroyed from the EUD secure key storage. (MA-EU-44)

b) All unnecessary X.509 certificates are removed from the EUD Trust Anchor Database. (MA-EU-45)

12) If applicable,

a) Confirm use of Domestic Cellular Service as a Black Transport Network.

b) Ensure the EUDs use an Access Point Name (APN) provided by a domestic cellular carrier private network. (MA-EU-36)

13) Verify the EUDs are configured for all IP traffic, with the exception of IKE, network address configuration, time synchronization, and name resolution traffic required to establish the IPsec tunnel, to flow through the IPsec VPN client. (MA-EU-37)

14) If applicable, verify the EUDs are configured for all IP traffic, with the exception of IKE, to flow through the IPsec VPN Client.

15) Verify the EUDs are configured as follows:

a) Maximum password life time shall be less than 181 days. (MA-EU-39)

b) Screen shall lock after three minutes of inactivity. (MA-EU-40)

c) Shall perform a wipe of all protected data after no more than 10 authentication failures. (MA-EU-41)

d) If feasible, display notifications shall be disabled while in a locked state. (MA-EU-46)

e) If feasible, USB mass storage shall be disabled. (MA-EU-47)

f) USB data transfer shall be disabled, if feasible. (MA-EU-48)

16) Ensure VPN protection is enabled across the EUD. (MA-EU-42).

17) Ensure a security policy is in place on EUDs specific to each RD and/or Government Private Wireless network to which the EUD is to be connected. (MA-EU-43)

18) Where feasible, ensure the EUDs use an Access Point Name (APN) provided by a domestic Cellular carrier Private Network when using Domestic Cellular Service as the Black Transport Network. (MA-EU-36)

19) Ensure the system software digital signatures are verified prior to the Application Processor system software is updated. (MA-EU-49)

20) Ensure application digital signatures are verified prior to installing new applications. (MA-EU-50).

Expected Result:

For steps 1-3, all EUDs shall be configured properly. For steps 4, a remotely administrated EUD shall only be rekeyed over the MA solution network prior to the expiration of keys. If this cannot be accomplished, the EUD must be re-provisioned. For Step 5, the implementing organization should have policies in place in order to address the requirements identified. For Step 6, Mobile Platform passwords shall be verified. For Step 7, an organization using a RD will ensure that the RD is configured in accordance with this CP. For steps 8-10, unique X.509 v3 certificates are used for each layer. For step 11, only keys and certificates required for operation are on the EUD. Step 12, all traffic flows through the IPsec VPN client, unless it is required to establish the IPsec tunnel. Step 13 - 20, the EUDs are configured in accordance with this CP.

17.5 RETRANSMISSION DEVICE CONFIGURATION

This section contains a procedure to ensure that the configurations for the RD in the MA solution follow the requirements given in this CP.

Requirements being tested: MA-RD-5 through MA-RD-26

Procedure Description:

1) If the RD is configured to be a Wi-Fi access point, then ensure that the following configuration is in place:

a) If the RD is using PSK, the PSK is using a length of at least 32 hexadecimal characters (or equivalent). (MA-RD-5)

b) The RD shall only permit connections to devices on a Media Access Control (MAC) white list. (MA-RD-6)

c) If the RD is using PSK, the PSK is not displayed on the RD. (MA-RD-7)

d) The SSID and MAC addresses of connected devices are not displayed on the RD. (MA-RD-8, MA-RD-9)

e) Broadcast of the SSID is disabled. (MA-RD-18)

f) Wi-Fi Protected Setup (WPS) is disabled. (MA-RD-12)

g) Traffic of multiple EUDs sharing the RD is separated. (MA-RD-17)

h) If using WPA Enterprise, then certificate used for authentication shall be different from the certificates used to authenticate the outer and inner tunnels. (MA-RD-26)

2) Verify that the Administrator password is not displayed on the RD. (MA-RD-10)

3) Verify that the RD can display the number of currently connected devices. (MA-RD-11)

4) Verify that the RD can only be administered using HTTPS. (MA-RD-13)

5) Verify that the RD is configured to limit the number of connected devices to the maximum required for the mission. (MA-RD-16).

6) Verify that the RD only permits charging on the USB ports and interfaces. (MA-RD-19)

7) Verify that connected EUDs are unable to access files stored on the RD. (MA-RD-20)

8) Verify that the default RD Administrator credentials have been changed. (MA-RD-15)

9) Verify that a user is required to authenticate with Administrator credentials prior to being able to perform the following:

a) Perform changes to the RD. (MA-RD-14)

b) Download logs or configuration files. (MA-RD-21)

c) Update firmware. (MA-RD-24)

10) Verify the RD will only allow firmware updates signed by the RD manufacturer. (MA-RD-22)

11) Verity the RD cannot boot into recovery mode. (MA-RD-23)

Expected Result: The RD is configured in accordance with this CP.

17.6 Inner And Outer VPN Component Configurations

This section contains a procedure to ensure that the configurations for the Inner and Outer VPN Components in the MA solution follow the requirements given in this CP.

Requirements being tested: MA-SR-2, MA-SR-3, MA-SR-4, MA-SR-5, MA-SR-9, MA-CR-1 through MA-CR-15, MA-IR-1 through MA-IR-5, MA-OR-1 through MA-OR-7, MA-DM-1 through MA-DM-8, MA-DM-11 through MA-DM-13, MA-DM-18

Procedure Description:

1) For the Inner and Outer VPN component in the solution, perform the following:

   a) Obtain the current configuration for the VPN Component.

   b) Verify that a unique device certificate is loaded with the corresponding CA signing certificate. (MA-CR-4)

   c) Verify that a device certificate from a CA included in the MA solution is listed in the configuration for authentication. (MA-CR-5)

   d) Ensure that the corresponding CA signing certificate and certificate revocation information are on the VPN Component. (MA-CR-6)

   e) Ensure that default accounts, passwords, community strings, and other default access control mechanisms for all components are changed or removed. (MA-SR-4)

   f) Ensure that all components are configured in accordance with local policy and applicable U.S. Government guidance, or, in the event of conflict between this CP and local policy, this CP takes precedence (MA-SR-5)

   g) Ensure that Inner Encryption components are not performing switching or routing for other Encryption Components (MA-SR-9)

   h) Verify that the requirements MA-CR-1, MA-CR-2, MA-CR-3, MA-CR-7, MA-CR-9, MA-CR-10 through MA-CR-14, are configured properly.

   i) Ensure that the time of day on the Inner Encryption Endpoints, Inner Firewall, and Red Management Services matches the current time. This should be within a small margin of error, to be determined by the AO. (MA-SR-2)

j) Ensure that the time of day on the Outer Gateway, Gray Firewall and Gray Management Services matches the current time. This should be within a small margin of error, to be determined by the AO. (MA-SR-3)

k) Verify that the VPN Components are configured to re-authenticate the identity of the VPN Components at the other end before rekeying the IKE SA. (MA-CR-15)

l) Verify that any Outer VPN Gateway in the solution is not configured to perform routing. (MA-OR-9)

2) For the Inner VPN component in the solution, use the configuration from 1a and perform the following:

a) Log into the Inner VPN component and verify that it is configured to use Tunnel or Transport mode IPsec with an associated IP Protocol (i.e. GRE). (MA-IR-1)

b) Log into the Inner VPN component and verify that the MTU (for IPv4) or the PMTU (for IPv6) has been configured to an appropriate size. (MA-IR-2)

c) Using a packet analyzer tool on the Inner VPN Gateway, verify that traffic leaving the external interface going to the Outer VPN Gateway is encrypted. (MA-IR-3)

d) Using a packet analyzer tool on the EUD, ensure that all traffic leaving the Inner VPN client is encrypted except for traffic otherwise identified in this CP as permissible to send unencrypted. (MA-IR-4)

e) Using a packet analyzer tool on the Inner VPN Gateway, verify that traffic coming through the external interface of the Inner VPN Gateway is decrypted. (MA-IR-5)

f) Verify that a separate LAN or VLAN is established on the Enterprise/Red network and using a packet sniffer, inspect traffic within the Enterprise/Red Network to ensure it is being used exclusively for all management of Inner Encryption Endpoints and solution components within the Red network. (MA-DM-3)

3) For each Outer VPN component in the solution, use the configuration from 1a and perform the following:

a) Log into the Outer VPN component and verify that it is configured to use Tunnel mode IPsec. (MA-OR-1)

b) Verify the requirements MA-OR-2, MA-OR-3, and MA-OR-4, have been properly configured.

c) Verify that a separate LAN or VLAN is established on the Enterprise/Red network and using a packet sniffer, inspect traffic within the Gray network to ensure it is being used exclusively for all

management of Outer VPN Gateway, Gray firewall and solution components within the Gray network. (MA-DM-4)

d) Using a packet analyzer tool on the EUD, ensure that all traffic leaving the Outer VPN client is encrypted except for traffic otherwise identified in this CP as permissible to send unencrypted. (MA-OR-5)

e) Using a packet analyzer tool on the Outer VPN Gateway, verify that traffic coming through the external interface of the Outer VPN Gateway is decrypted. (MA-OR-7)

4) For an EUD that uses VMs to separate Inner and Outer VPN clients, verify that the Outer VPN client is not installed on the host operating system. (MA-OR-6)

5) For all device administration, verify that requirements MA-DM-6, MA-DM-1, MA-DM-2, and MA-SR-8, MA-DM-18 are configured properly.

6) For each administration workstation, ensure the Security Administrator is required to authenticate to the component before being granted access. (MA-DM-7)

7) For each administration workstation, ensure the Security Administrator is required to authenticate to solution components using Suite B compliant certificates. (MA-DM-8)

8) Ensure that certificate signing requests are initiated by the Security Administrator as part of their initial keying within the solution. (MA-DM-11)

9) Ensure that devices use EST as detailed in RFC 7030 for certificate management. (MA-DM-12)

10) Ensure that the same Administration Workstation is not used to manage both Inner and Outer VPN Gateways. (MA-DM-13)

11) Ensure that requirement MA-DM-5 has been configured properly.

Expected Result:

For Steps 1-7, the Inner and Outer VPN Components should be configured properly according to the requirements found in this CP. For Steps 8-10, all VPN Component administration devices should be configured properly based upon the requirements of this CP. For Steps 11-12, all of the procedures have been followed or are in place.

17.7 Key Management

This section contains a procedure to ensure that key management capabilities for the MA solution follow the requirements given in this CP.

Requirements being tested: MA-KM-1 through MA-KM-49, MA-AU-36 through MA-AU-38, MA-PF-3

Procedure Description:

1) Perform the following to validate the correct deployment of CAs:

    a) ()For CAs, verify the configuration of the MA Solution to ensure that Outer CAs deliver services through either the Gray or Red networks; and, Inner CAs only deliver services through the Red network. (MA-KM-2, MA-KM-3)

    b) For CAs, verify that the Outer and Inner CAs are physically separate from one another. (MA-KM-4)

    c) For Locally-run CAs that operate on-line, verify the CAs utilize FIPS 140-2 Level 2 or higher Hardware Security Modules (HSMs) to protect the CAs' private signing keys. (MA-KM-12)

    d) For all CAs, verify that the CA does not have access to any MA Solution Component private keys. (MA-KM-11)

Expected Results: CAs are correctly deployed and in compliance with requirements MA-KM-1 through MA-KM-4, MA-KM-11 and MA-KM-12.

2) Perform the following to validate the structure of certificates issued by CAs:

    a) Obtain a sample set of test certificates issued to MA solution components.

    b) Verify the Distinguished Name in each certificate identifies a Non-Person Entity (NPE), unless that certificate is used in a TLS EUD where a user certificate is required. In this case, the certificate DN identifies a human user. (MA-KM-5)

    c) Verify the Key Usage extension in the each certificate only asserts "digitalSignature". (MA-KM-6)

    d) Verify the certificates are compliant with the data standard for Version 3 certificates defined in ITU-T Recommendation X.509. (MA-KM-7)

    e) Verify the certificates are compliant with IETF RFC 5280 profile requirements. (MA-KM-8)

    f) Verify the certificates are compliant with the key sizes and algorithms specified in Tables 9-11. (MA-KM-10)

Expected Results: The Certificate structure and content are compliant with requirements MA-KM-5 through KM-8 and MA-KM-10.

3) Perform the following to validate correct policy implementation as it relates to CAs and certificates issued by the CAs:

   a) Verify the CA has a Certificate Policy and a Certification Practices Statement (CPS) in place that is compliant with IETF RFC 3647, and that the CA operates in accordance with the policy and CPS. (MA-KM-13, MA-AU-36, MA-AU-37, MA-AU-38)

   b) Verify that the MA Solution policy states device authentication certificates issued by the CA, along with corresponding private keys, are considered Controlled Unclassified Information (CUI), and user private keys are classified to the level determined by the AO. (MA-KM-1, MA-KM-9)

Expected Results: CA policy exists and complies with requirements MA-KM-9 and MA-KM-13.

4) Perform the following steps to validate the certificate issuance capability of the MA solution:

   a) Verify that a physical environment is identified to initially load keys and certificates onto MA Solution Components, where the environment is certified to protect information at the highest classification level of the MA Solution red network. (MA-KM-14)

   b) Verify that the key and certificate provisioning process for MA Solution Components ensures private keys are never escrowed. (MA-KM-15)

   c) Generate a public/private key pair for the Outer VPN Component that complies with the key size and algorithm requirements in Tables 9-11. If the Component is capable of generating its own key pair, the key pair is to be generated on the Component. Else, the key pair is generated by a dedicated management workstation.

   d) Generate a certificate request for the Outer VPN Component, and ensure the request complies with PKCS#10.

   e) Submit the certificate request to the Outer CA, and verify that the CA returns a signed certificate using PKCS#7. (MA-KM-16)

   f) If the key pair was generated by a management workstation, install the certificate and private key using PKCS#12. If the key pair was generated by the VPN Component, install the signed certificate using PKCS#7. (MA-KM-17, MA-KM-18)

   g) Repeat steps 4c through 4f for each Inner Encryption Component of the MA Solution.

   h) If the MA Solution supports IETF RFC 7030 (Enrollment over Secure Transport (EST)), verify that the certificate request, response and installation process complies for MA Solution Components complies with EST. (MA-KM-19)

i) Verify that the certificate request and issuance processes comply with the Outer and Inner CAs' Certificate Policies and CPSs. (MA-KM-20, MA-KM-21)

j) For locally run CAs, verify the Certificate Policies and CPSs to ensure certificate issued by the CAs: 1) enforce unique Distinguished Names (DNs); 2) assert key usages as defined by MA-KM-5; and 3) assert a registered policy Object Identifier (OID).

k) For locally run CAs, examine the contents of a sample set of issued certificates to ensure that the certificates assert: 1) unique Distinguished Names (DNs); 2) key usages as defined by MA-KM-5; and 3) a registered policy OID. (MA-KM-22)

l) For all CAs, examine the contents of a sample set of issued certificates and ensure that at least one valid CRL Distribution Point (CDP) is asserted in the CDP extension of the certificates. (MA-KM-23)

m) For locally run CAs, ensure the validity periods asserted in certificates issued by the CAs do not exceed 14 months for MA End User Devices (EUDs). (MA-KM-24)

n) For locally run CAs, ensure the validity periods asserted in certificates issued by the CAs do not exceed 36 months for MA Solution Infrastructure Components. (MA-KM-25)

o) Verify that the Inner CAs can only issue certificates to Inner Encryption Components and Red Network Components. (MA-KM-26)

p) Verify that the Outer CAs can only issue certificates to Outer VPN Components and Gray Network Components. (MA-KM-27)()()

Expected Results:

For Step 4a, a physical environment exists to initially load keys and certificates onto MA Solution Components, where the environment is certified to protect information at the highest classification level of the MA Solution Red network.

For Step 4b, private keys for MA Solution Components cannot be escrowed.

For Steps 4c-4e, the certificate request/response process between the MA Solution Component and the CAs correctly implements PKCS#10 and PKCS#7.

For Step 4f, the key and certificate installation process correctly implements PKCS#12 or PKCS#7.

For Step 4g, same results as for steps 4c through 4f.

For Step 4h, the certificate request/response process between the MA Solution Component and the CAs correctly implements EST.

For Step 4i, the certificate request/response process complies with the CAs' Certificate Policies and CPSs.

For Step 4j-4n, the contents of certificates issued by the CAs comply with requirements MA-KM-22 through MA-KM-25.

For Step 4o, Inner CAs can only issue certificates to Inner Encryption Components and Red Network Components.

For Step 4p, Outer CAs can only issue certificates to Outer VPN Components and Gray Network Components.

For Step 4q, Enterprise Root CA shall issue certificates the Subordinate CAs for the Red and Gray networks respectively.

For Step 4r, the Subordinate CAs shall issue certificates for the inner and outer tunnels.

5) Perform the following steps to validate the certificate renewal and rekey capability of the MA solution:

   a) Verify that the certificate renew and rekey processes comply with the Outer and Inner CAs' Certificate Policies and CPSs. (MA-KM-31)

   b) Verify that the Outer and Inner CAs' Certificate Policies and CPSs require certificate renew and rekey be performed prior to a certificate expiring. Verify the Outer and Inner CAs' Certificate Policies and CPSs require an MA Solution Component go through the initial certificate issuance process if the certificate is expired. (MA-KM-30)

   c) Generate a new public/private key pair for the Outer VPN Component that complies with the key size and algorithm requirements in Tables 9-11. If the Component is capable of generating its own key pair, the key pair is to be generated on the Component. Else, the key pair is generated by a dedicated management workstation.

   d) Generate a certificate renew and rekey request for the Outer VPN Component, and ensure the request complies with PKCS#10.

   e) Submit the certificate request to the Outer CA, and verify that the CA returns a signed certificate using PKCS#7. (MA-KM-32)

   f) Repeat steps 5c through 5e for each Inner Encryption Component of the MA Solution.

g) If the MA Solution supports IETF RFC 7030 (Enrollment over Secure Transport (EST)), verify the certificate renew and rekey request, response and installation process complies for MA Solution Components complies with EST. (MA-KM-35)

Expected Results:

For Step 5a, the certificate renew and rekey request/response process complies with the CAs' Certificate Policies and CPSs.

For Step 5b, the CAs' Certificate Policies and CPSs require certificate renew and rekey to be performed prior to the certificate expiring. If the certificate is expired, the Certificate Policies and CPSs require the MA Solution Component go through the initial certificate issuance process.

For Steps 5c-5e, the certificate renew and rekey request/response process between the MA Solution Component and the CAs correctly implements PKCS#10 and PKCS#7.

For Step 5f, the key and certificate installation process for certificate renew and rekey correctly implements PKCS#12 or PKCS#7.

For Step 5g, same results as for steps 5c through 5f.

For Step 5h, the certificate renew and rekey request/response process between the MA Solution Component and the CAs correctly implements EST.

6) Perform the following steps to validate the certificate revocation and CDP capabilities of the MA solution:

a) Verify that the Outer and Inner CAs' Certificate Policies and CPSs define requirements and procedures for revoking MA Solution Component certificates, where certificate revocation is required when the binding between the subject information and public key within the certificate is no longer considered valid. (MA-KM-36)

b) Verify that the Outer and Inner CAs' Certificate Policies and CPSs define requirements and procedures for requesting the revocation of MA Solution Component certificates. (MA-KM-40)

c) For locally run CAs, verify the Outer and Inner CAs' Certificate Policies and CPSs define certificate revocation requirements and procedures for MA Solution Components that address: 1) response to a lost, stolen or compromised MA EUD; 2) removal of a revoked MA infrastructure device from the MA Solution network; 3) re-establishment of a MA Solution Component after certificate revocation is performed; 4) revocation of certificates when a MA EUD is considered compromised; and 5) revocation of

certificates if simultaneous use of the certificate is detected from different IP addresses. (MA-KM-41)

d) Verify that the Outer and Inner CAs has the capability to generate CRLs after certificate revocation functions are performed. (MA-KM-37)

e) Obtain CRLs from the Outer and Inner CAs and ensure their structures are compliant with the data standard for Version 2 CRLs defined in ITU-T Recommendation X.509, and with the CRL profile standard defined by IETF RFC 5280. (MA-KM-38, MA-KM-39)

f) Obtain CRLs from the Outer and Inner CAs and upload them onto the CDPs defined for the MA Solution. Depending on the MA Solution configuration, Outer CA CRLs can be uploaded onto Black and/or Outer CDPs; Inner CA CRLs can be uploaded onto Gray and/or Red CDPs.

g) Verify that MA Solution Components can access the CDPs and download the CRLs issued by the Outer and Inner CAs via HTTP. Outer VPN Components and Gray Management Service Components are able to access and download the CRL issued by the Outer CA; Inner Encryption Components and Red Management Service Components are able to access and download the CRL issued by the Inner CA. (MA-KM-42, MA-PF-3)

h) For Enterprise CAs, verify that the Certificate Policies and CPSs define requirements and procedures for publishing CRLs. (MA-KM-43)

i) For locally run CAs, verify that the Certificate Policies and CPSs define requirements and procedures for 1) publishing new CRLs at least once every 28 days; 2) creating a new CRL within one hour of a certificate being revoked; and 3) publishing a newly created CRL at least 7 days before the expiration of the current CRL. (MA-KM-44, MA-KM-45, MA-KM-47)

j) Verify the MA Solution has procedures defined to transfer new CRLs to MA Solution CDPs within 24 hours of the CRLs being created. (MA-KM-46)

k) For MA Solutions that support the On-Line Certificate Status Protocol (OCSP) to provide certificate revocation status information, verify the OCSP Servers are deployed on the Gray and Red networks to deliver OCSP responses in accordance with IETF RFC 6960. (MA-KM-48)

l) Generate an OCSP request from the Outer VPN Gateway and send the request to the OCSP Server operating in the Gray network.

m) Generate an OCSP response from the OCSP Server in the Gray network and deliver it to the Outer VPN Gateway.

n) Examine the OCSP Response, and verify that it is digitally signed and compliant with IETF RFC 6960. (MA-KM-49)

o) Repeat steps 6l-6n for all MA Solution Inner Encryption Components using an OCSP Server operating in the Red network. (MA-KM-49)

Expected Results:

For steps 6a-6c, the Certificate Policies and CPSs have requirements and procedures defined to satisfy requirements MA-KM-36, MA-KM-40 and MA-KM-41.

For steps 6d and 6e, the CAs are able to generate Version 2 CRLs that are compliant with ITU-T Recommendation X.509 and IETF RFC 5280.

For steps 6f and 6g, MA Solution Components successfully access and download CRLs from CDPs deployed in the Gray and Red networks of the MA Solution.

For steps 6h and 6i, the Certificate Policies and CPSs have requirements and procedures defined to satisfy requirements MA-KM-43 through MA-KM-45, and MA-KM-47.

For step 6j, procedures exist to transfer new CRLs to MA Solution CDPs within 24 hours of the CRLs being created. (MA-KM-46)

For steps 6k through 6o, OCSP Servers are correctly deployed in the Gray and Red networks and issue digitally signed OCSP responses compliant with IETF RFC 6960 to MA Solution Components. (MA-KM-48 and MA-KM-49)

17.8 Solution Filtering Configurations

This section contains a procedure to ensure that the filtering configurations for all the MA solution follow the requirements given in this CP.

Requirements being tested: MA-PF-1 through MA-PF-2, MA-PF- 4 through MA-PF-16

Procedure Description:

1) Perform the following steps on the each of the Solution Components:

a) Log into the component.

b) Verify through the configuration file that network interfaces are restricted to the smallest address range, ports, and protocols (MA-PF-1).

c) Verify through the configuration file that all unused network interfaces are disabled (MA-PF-2).

2) For the Outer VPN Gateway perform the following:

   a) Obtain the current configuration for the Outer VPN Gateway.

   b) Verify that the requirements MA-PF-4, MA-PF-7 or MA-PF-8, , MA-PF-9 through MA-PF-12 are met.

3) For the Inner VPN Gateway perform the following:

   a) Obtain the current configuration for the Inner VPN Gateway.

   b) Verify the requirements MA-PF-5.

4) Establish a connection from the EUD to the Red network.

   a) Using a protocol analyzer on the Outer firewall, observe inbound and outbound traffic on the Outer firewall.

   b) Verify that the only protocols that are allowed through are IKE, ESP, and control plan protocols as specified in MA-PF-13.

5) For the Inner firewall perform the following:

   a) Log into the Inner firewall.

   b) Obtain the configuration file.

   c) Verify that the Inner firewall has a whitelist for all Inner Encryption Endpoints (MA-PF-6).

6) For the EUDs consisting of a single computing platform

   a) Verify there is no ingress or egress of Certificate Revocation traffic (OCSP queries, HTTP GET to CDPs) on the Black interface. (MA-PF-14)

   b) Verify there is no ingress or egress of Name Resolution traffic (e.g. DNS query/response) on the Black interface. (MA-PF-15)

   c) Verify there is no ingress and egress of NTP traffic on the Black interface. (MA-PF-16)

Expected Result:

For Step 1, the Solution components network interfaces are configured properly. For Step 2- 4, the Outer, Inner VPN Gateway, and Outer firewall are only allowing the necessary protocols. For Step 5, verify that the Inner firewall is configured correctly. For Step 6, verify there is no

Certificate Revocation, name resolution, or NTP traffic between the EUDs and the Black network.

17.9 Configuration Change Detection

This section contains a procedure to ensure that changes made to any of the MA Solution configurations are detected by the Configuration Change Detection tool.

Requirements being tested: MA-CM-1 through MA-CM-3

Procedure Description:

1) The following shall be performed for each of the Solution components within this CP.

   a) Log into the Solution components (Outer firewall, Outer VPN Gateway, Gray firewall, Gray Management Services, Inner Encryption Endpoints, Red Management Services, EUD, and Retransmission Device (if applicable)).

   b) Compare the current version of the Solution Component's configuration with the stored baseline and ensure the current version matches the stored configuration. (MA-CM-1)

   c) Make a change to the configuration, preferably something that is not fundamental to the security of the MA solution.

   d) Look in the audit log to determine if a log entry has been generated about the configuration change and that the changes from 1c are recorded. (MA-CM-2).

   e) Inspect the monitoring service to verify that the service has detected a change in configuration. (MA-CM-3)

Expected Result:

The Auditor will validate the baseline configuration was stored in Step 1b. In Step 1d, there should be a log entry created for the configuration change in the audit log including the actual configuration change. Lastly if there was a configuration change, a monitoring service will detect a change in the configuration.

17.10 Continuous Monitoring

This section contains procedures for ensuring traffic is monitored for and alerts generated for potential unauthorized/malicious traffic. It also contains procedures for ensuring a SIEM is in place to collect logs and that it is configured correctly.

Requirements being tested: MA-MR-1 through MA-MR-18

Procedure Description:

1) Ensure that an IDS/IPS is deployed to monitor traffic in at least one of three locations (MA-MR-1 through MA-MR-6):
    a) Between the Outer firewall and Outer VPN Gateway (M1)
    b) Between the Outer VPN Gateway and the Gray firewall (M2)
    c) On the internal side of the Inner firewall (M3)
2) Ensure that each IDS/IPS in the solution is configured to send alerts to the Security Administrator, and, where possible, block malicious traffic. (MA-MR-7, MA-MR-8)
3) Ensure that each IDS/IPS in the solution is configured with rules that will generate alerts and, where possible, block traffic for any unauthorized source and destination IP addresses (MA-MR-9 through MA-MR-12).
4) Ensure that an SIEM is implemented (MA-MR-13 – MA-MR-16).
    a) Ensure the SIEM is implemented in the Gray network.
    b) Otherwise, if the SIEM is implemented within the Enterprise/Red network, ensure devices are configured to push events to an Enterprise/Red SIEM and through an AO-approved one-way tap.
    c) Send packets expected to be blocked by the Outer VPN Gateway or Gray firewall. Ensure the SIEM sends alerts to the Auditor when anomalous behavior such as this is detected.
    d) Ensure that logs from the Outer VPN Gateway Gray firewall and any other components located within the Gray Management Services are collected on the Gray SIEM.
    e) Ensure these logs are encrypted with TLS, SSHv2, or IPSEC.
5) Ensure that any one-way taps are deployed as per MA-MR-17 – MA-MR-18.
    a) Ensure that any one-way taps deployed as part of the solution are approved for use by the AO.
    b) Ensure that the SIEM implemented at the Red level that collects black and/or gray monitoring data sent through any one-way tap is deployed in an enclave isolated from the Red/Enterprise Network.
    c) Ensure that monitoring data flowing from M2 and/or M1 can transit to the SIEM if implemented at the Red level.
    d) Attempt to send other data through the one-way taps to determine if this data is blocked.

Expected Result:

For Steps 1 – 3, an IDS or IPS is in place to monitor, block where possible, and send alerts as appropriate. For Steps 4 and 5, an SIEM shall be implemented either in the Gray network or the Red/Enterprise network via one-way taps, as approved by the AO and only monitoring data will be able to transit through these taps.

17.11 AUDIT

This section contains procedures for ensuring audit events are detected, the proper information is logged for each event, and there is a procedure detailed in the CPS documentation for auditing each CA device.

Requirements being tested: MA-DM-14, MA-DM-15, MA-AU-1 through MA-AU-22, MA-AU-29 through MA-AU-35, MA-DM-10, MA-DM-16, MA-DM-17

Procedure Description:

1) Examples for testing the ability of each MA Component to audit and log audit events specified in the CP are given below. Verify that for each event logged, the applicable data regarding the event is recorded for the log entry in accordance with Section 13.13.

    a) All actions performed by a user with super-user privileges (auditor, administrator, etc.) and any escalation of user privileges. (MA-AU-12, MA-AU-13)\

        i) Log in as an administrator to the Solution Infrastructure Components or a EUD.

        ii) Perform a variety of administrator actions on the Solution Infrastructure Components or EUD.

        iii) Verify that a log entry was created for each action taken in Step ii that required super-user privileges and also states the escalation of privileges.

        iv) If performing an action on anEUD, verify that the EUD is generating logs and sends the logs to the central SIEM. (MA-DM-10)

        v) Revert back to the baseline configuration, eliminating the changes made in Step ii.

        vi) Repeat the above with the Auditor role.

    b) Changes to time. (MA-AU-15)

        i) Log in as a Security Administrator to the Solution Infrastructure Components.

        ii) Modify the system time on the Solution Infrastructure Components by at least 1 hour.

iii) Verify that a log entry was created due to the change in system time and by whom.

iv) Revert the system time back to the accurate time of day.

c) Log into and out of the MA Solution as a normal user and send traffic to the Red Network. Then log into the central SIEM as an Auditor, and inspect the audit entry for the following: MA-DM-14, MA-DM-15.

i) Verify that the log on as a normal user is logged and has an identifiable code for the type of event. (MA-AU-10, MA-AU-18)

ii) Verify that the log entry identifies the subject accessing the solution. (MA-AU-20)

iii) Verify that the log entry identifies the event. (MA-AU-17)

iv) Verify that the log entry includes the time, date, and the time zone offset. (MA-AU-16)

d) Establish and terminate a VPN tunnel. Verify in the logs, that these two events were logged. (MA-AU-1, MA-AU-3, MA-AU-5, MA-AU-7)

i) Establish and terminate a TLS connection. Verify in the logs, that these two events were logged. (MA-AU-2, MA-AU-4, MA-AU-6, MA-AU-8)

e) Log into a Solution Infrastructure Components as a Security Administrator and delete previously recorded audit log. Verify the log recorded this deletion. (MA-AU-9)

f) As the Certificate Authority Administrator, log into the audit log and attempt to delete a log entry. Verify this action is recorded with a failure code. (MA-AU-11, MA-AU-19)

g) Verify a log entry was created for the attempted unauthorized action.

2) Verify the source address, user, and for role-based events, role identity for all audit log entries is recorded. (MA-AU-21, MA-AU-22)

3) Verify that all logs forwarded to a SIEM on a Gray Management network are configured to be encrypted while in transit using SSHv2, IPSEC, or TLS with the appropriate Suite B algorithm supported by the solution. (MA-DM-16)

4) Verify that all logs forwarded to a SIEM on a Red Management network are configured to be encrypted while in transit using SSHv2, IPSEC, or TLS with the appropriate Suite B algorithm supported by the solution. (MA-DM-17)

5) Verify that the procedure MA-AU-35 is currently in place by the implementing organization and are followed.

6) Verify that the Outer VPN Gateway and Inner Encryption Components log the failure to pull the CRL from the Inner or Outer CDP. (MA-AU-29, MA-AU-30)

   a) CDP Servers shall remove all CRLs.

   b) Outer VPN Gateways and Inner Encryption Components shall attempt to pull the CRL from their respective CDPs**.**

   c) Review the Outer VPN Gateway and Inner Encryption Components audit logs to verify that a log report is generated from failure to pull the CRL.

7) Verify that the Outer VPN Gateway and Inner Encryption Components log if the version of the CRL on the Inner or Outer CDP is older than the current cached CRL. (MA-AU-31, MA-AU-32)

   a) Load the CDPs with CRLs that are older than the current cached CRLs on the Outer VPN Gateway and Inner Encryption Components.

   b) Have the Outer VPN Gateway and Inner Encryption Components attempt to pull the CRLs.

   c) Review the Outer VPN Gateway and Inner Encryption Components audit logs to verify that a log report is generated.

8) Verify that the Outer VPN Gateway and Inner Encryption Components log if signature validation of the CRL on the Inner or Outer CDP fails. (MA-AU-33, MA-AU-34)

   a) Load the CDPs with CRLs that contain an invalid signature.

   b) Have the Outer VPN Gateway and Inner Encryption Components pull the CRLs.

   c) Review the Outer VPN Gateway and Inner Encryption Components audit logs to verify that a log report is generated due to an invalid CRL signature.

9) For all Solution components, install approved certificates, generated by the approved CA, and configure the solution so that components use the certificates for authentication.

   a) Verify an entry to the Audit log has been created due to certificate loading and generation. (MA-AU-14)

   b) Initiate a revocation of certificates for the Solution components.

   c) Verify that an entry in the audit log has been created due to certificate revocation. (MA-AU-14)

Expected Result:

For Step 1, all occurrences of auditable events given should generate an entry in the audit log. For Step 2, the source address should be the MA Components' loopback address. For Steps 3-4, all logs forwarded on Red Management and Gray Management networks should be encrypted with the appropriate protocols. For Step 5, the procedure is followed and is in place. For Step 6-8, there should be an audit log entry created for each requirement. For Step 9, a log should be generated for generation and revocation of certificates.

17.12 EUD With Multiple Connections

This section contains a procedure to ensure that only one IPsec or TLS connection is allowed for the inner layer of encryption per EUD and that no other connections are permitted.

Requirements being tested: MA-AU-23 through MA-AU-28

Procedure Description:

1) For a IPsec connection, ensure that the EUD performs the following steps:
   a) The administrator will install the same device certificates on two EUDs.
   b) The administrator will authenticate to a Red network. At the same time, the Auditor will be reviewing the logs and detect that the same device certificate is coming from two different devices. (MA-AU-23)
   c) The Auditor will alert the Certificate Authority Administrator to revoke the certificates and provide an updated Certification Revocation List to the Security Administrator. (MA-AU-25)
   d) Once the Security Administrator receives notification from the Certificate Authority Administrator, the Security Administrator will drop both sessions. (MA-AU-27)
2) For TLS connection, ensure that the EUD performs the following steps:
   a) The administrator will install the same device certificates on two EUDs.
   b) The administrator will authenticate to a Red network. At the same time, the Auditor will be reviewing the logs and detect that the same device certificate is coming from two different devices. (MA-AU-24)
   c) The Auditor will alert the Certificate Authority Administrator to revoke the certificates and provide an updated Certification Revocation List to the Security Administrator. (MA-AU-26)
   d) Once the Security Administrator receives notification from the Certificate Authority Administrator, the Security Administrator will drop both sessions. (MA-AU-28)

Expected Result:

The same device certificate cannot be used for two devices. All results are expected to be pass/fail.

17.13 Incident Reporting Guidance

This section ensures that procedures are followed regarding incident reporting to NSA in the event a solution owner identifies a security incident which affects the solution.

Requirements being tested: MA-RP-1 through MA-RP-16

Procedure Description:

1) Verify that the procedures given in MA-RP-1 through MA-RP-16 were/are followed and are currently in place.

Expected Results:

For Step 1, all of these procedures have been followed or are in place.

17.14 Implementation of Guidance

This section ensures that there are procedures in place and/or that procedures were followed regarding the procurement of products and use of the MA solution. It also ensures the personnel are in place to manage and administer this solution following the guidelines given in the CP.

Requirements being tested: MA-GD-1 through MA-GD-41, MA-SR-6

Procedure Description:

1) Verify the procedures for obtaining virus signature updates as required by local agency policy and the AO were/are followed and/or are in place. (MA-SR-9)
2) Verify the procedures given in MA-GD-1 through MA-GD-8, and MA-GD-17 through MA-GD-30 were/are followed and/or are currently in place.
3) Verify that the solution owner understands that he/she shall allow and fully cooperate with an NSA-ordered IA compliance audit of this solution implementation. (MA-GD-9)
4) Verify that the AO are aware that a compliance audit will be conducted every year. (MA-GD-10)
5) Verify that the AO is aware that they shall receive the results of the compliance audit and are responsible for reviewing the completed audit. (MA-GD-11)

6) Verify that the customer is aware that when they are interested in registering their solution against this CP that NSA must grant them an approval prior to the AO authorizing the solution for operation. (MA-GD-12)

7) Verify that the customer completes and submit the compliance matrix to their AO. (MA-GD-13)

8) Verify that the customer is aware that registration and re-registration against this CP includes a submission of this CP registration forms and compliance matrix to NSA. (MA-GD-14)

9) Verify that the customer is aware that when a new MA CP is published by the NSA, the AO will comply against this new CP within 6 months. (MA-GD-15)

10) Verify that the solution owner and AO are aware that they shall provide updated solution information to NSA on a yearly basis. (MA-GD-16)

11) Verify that the personnel requirements given in MA-GD-31 through MA-GD-41 are met by the personnel supporting this implementation of the MA solution.

Expected Result:

For steps 1-10, all of these procedures have been followed or are in place.

17.15 Solution Functionality

This section contains a procedure for ensuring the implementing organization complies with the testing requirements.

Requirements being tested: MA-TR-1

Procedure Description:

1) The implementing organization's AO will inspect the test report in order to ensure all testing requirements have been met. (MA-TR-1)

Expected Result:

The report will ensure that the implementing organization complies with the MA Solution.

APPENDIX A. GLOSSARY OF TERMS

Authorization (To Operate) – The official management decision given by a senior organizational official to authorize operation of an information system and to explicitly accept the risk to organizational operations (including mission, functions, image, or reputation), organizational assets, individuals, other organizations, and the Nation based on the implementation of an agreed-upon set of security controls. (NIST SP 800-37)

Authorization Boundary – All components of an information system to be authorized for operation by an authorizing official and excludes separately authorized systems, to which the information system is connected.

Authorizing Official – A senior (Federal) official or executive with the authority to formally assume responsibility for operating an information system at an acceptable level of risk to organizational operations (including mission, functions, image, or reputation), organizational assets, individuals, other organizations, and the Nation.

Authorizing Official Designated Representative – An organizational official acting on behalf of an authorizing official in carrying out and coordinating the required activities associated with security authorization.

Authorization Package – A security package of documents consisting of the security control assessment that provides the authorizing official with essential information needed to make a risk-based decision on whether to authorize operation of an information system or a designated set of common controls.

Assurance – Measure of confidence that the security features, practices, procedures, and architecture of an information system accurately mediates and enforces the security policy. (CNSSI 4009)

Audit – The activity of monitoring the operation of a product from within the product. It includes monitoring of a product for a set of pre-determined events. Each audit event may indicate rogue behavior, or a condition that is detrimental to security, or provide necessary forensics to identify the source of rogue behavior.

Audit Log – A chronological record of the audit events that have been deemed critical to security. The audit log can be used to identify potentially malicious activity that may further identify the source of an attack, as well as potential vulnerabilities where additional countermeasures or corrective actions are required.

Availability – Ensuring timely and reliable access to and use of information. (NIST SP 800-37).

Black Box Testing – Testing the functionality of a component of the solution, such that testing is limited to the subset of functionality that is available from the external interfaces of the box during its normal

operational configuration without any additional privileges (such as given to the Security Administrator or Auditor).

Black Network – A network that contains classified data that has been encrypted twice. (See Section 4.1.3)

CP – The set of guidance provided by NSA that describes recommended approaches to composing COTS components to protect classified information for a particular class of security problem. CP instantiations are built using products selected from the CSfC Components List.

Central Management Site – A site within a MA solution that is responsible for remotely managing the solution components located at other sites (see Section 4.2.3.2).

Certificate Authority (CA) – An authority trusted by one or more users to create and assign certificates. (ISO9594-8)

Certificate Policy (CP) – A named set of rules that indicate the applicability of a certificate to a particular community and/or class of application with common security requirements. For example, a particular CP might indicate applicability of a type of certificate to the authentication of parties engaging in business-to-business transactions for the trading of goods or services within a given price range. (IETF RFC 3647)

Committee on National Security Systems Policy No. 15 (CNSSP-15) – Policy specifies which public standards may be used for cryptographic protocol and algorithm interoperability to protect National Security Systems (NSS).

Confidentiality – Assurance that the data stored in, processed by, or transmitted by the system are protected against unauthorized disclosure, and confidence that only the appropriate set of individuals or organizations would be provided the information.

Control Plane Protocol – A routing, signaling, or similar protocol whose endpoints are network infrastructure devices such as VPN Gateways or routers. Control plane protocols carry neither user data nor management traffic.

CRL Distribution Point (CDP) – A web server that hosts a copy of a CRL issued by a CA for VPN Components to download (see Section 9.1).

Cross Domain Solution (CDS) – A form of controlled interface that provides the ability to manually and/or automatically access and/or transfer information between different security domains. (CNSSI 4009)

Data Plane Protocol – A protocol that carries the data being transferred through the solution.

End User Device (EUD) – A form-factor agnostic component of the Mobile Access solution that can include a mobile phone, tablet, or laptop computer. EUDs can be composed of multiple comments to provide physical separation between layers of encryption (see Section 4.2.1 for explanation of detailed differences between VPN EUD and TLS EUD solution design options).

Enterprise/Red Network – A network that contains unencrypted classified data and can contain singly encrypted gray data (see Section 4.1.1).

External Interface – The interface of the Outer VPN Gateway that connects to the internal interface of the Outer firewall.

Federal Information Processing Standards (FIPS) – A set of standards that describe the handling and processing of information within governmental agencies.

Gray Box Testing – The ability to test functionality within a component of the solution, such that full management privileges are granted (i.e. knowing passwords for security administrator and Auditor and access to the capabilities associated with those privileges). In addition, the use of any and all testing equipment and/or testing software used inside and outside the developed solution is available.

Gray Network – A network that contains classified data that has been encrypted once (see Section 4.1.2).

Gray Firewall – A stateful traffic filtering firewall placed on the Gray network to provide filtering of ports, protocols, and IP addresses to ensure traffic reaches the correct Inner Encryption Endpoint or is dropped.

Internal Interface – The interface on a VPN Gateway or Inner Encryption Component that connects to the inner network (i.e., the Gray network on the Outer VPN Gateway or the Red network on the Inner Encryption Component).

Locally Managed Device – A device that is being managed by the direct connection of the Administration Workstation to the device in a hardwired fashion (such as a console cable).

Malicious – Any unauthorized events that are either unexplained or in any way indicate adversary activity.

Management Plane Protocol – A protocol that carries either traffic between a system administrator and a component being managed, or log messages from a solution component to a SIEM or similar repository.

Protection Profile – A document used as part of the certification process according to the Common Criteria. As the generic form of a security target, it is typically created by a user or user community and provides an implementation independent specification of information assurance security requirements.

Public Key Infrastructure (PKI) – Framework established to issue, maintain, and revoke public key certificates.

Remotely Managed Device – A device that is being managed by any other method besides that given in the definition of a Locally Managed Device.

Security Level – The combination of classification level, list of compartments, dissemination controls, and other controls applied to the information within a network.

Split-tunneling – Allows network traffic to egress through a path other then the established VPN tunnel (either on the same interface or another network interface). Split tunneling is explicitly prohibited in MA CP compliant configurations (see MA-OR-2 and MA-EU-7).

SRTP Client – A component on the EUD that facilitates encryption for voice communications.

TLS Client – A component on a TLS EUD that can provides the Inner layer of DIT encryption.

TLS Component – Refers to both TLS Clients and TLS-Protected Servers.

VPN Client – A VPN application installed on an EUD.

VPN Component – The term used to refer to VPN Gateways and VPN Clients.

VPN Gateway – A VPN device physically located within the VPN infrastructure.

VPN Infrastructure – Physically protected in a secure facility and includes Inner and Outer VPN Gateways, Certificate Authorities, and Administration Workstations, but does not include EUDs.

APPENDIX B. ACRONYMS

Acronym	Definition
AES	Advanced Encryption Standard
AO	Authorizing Official
APN	Access Point Name
ARP	Address Resolution Protocol
BIOS	Basic Input/Output System
BGP	Border Gateway Protocol
CA	Certificate Authority
CAA	Certificate Authority Administrator
CDP	CRL Distribution Point
CDS	Cross Domain Solution
CNSS	Committee on National Security Systems
CNSSI	Committee on National Security Systems Instruction
CNSSP	Committee on National Security Systems Policy
COTS	Commercial Off-the-Shelf
CP	Certificate Policy
CP	Capability Package
CPS	Certification Practice Statement
CRL	Certificate Revocation List
CSfC	Commercial Solutions for Classified
CUI	Controlled Unclassified Information
DAR	Data-At-Rest
DDoS	Distributed Denial of Service
DH	Diffie-Hellman
DHCP	Dynamic Host Configuration Protocol
DHS	Department of Homeland Security
DM	Device Management
DN	Domain Name
DNS	Domain Name System
DOD	Department of Defense
DoE	Department of Energy
DoS	Denial of Service
DSA	Digital Signature Algorithm
ECDH	Elliptic Curve Diffie-Hellman
ECDSA	Elliptic Curve Digital Signature Algorithm
ESP	Encapsulating Security Payload
EST	Enrollment Over Secure Transport
EUD	End User Device
FDE	Full Disk Encryption
FIPS	Federal Information Processing Standards

Acronym	Definition
FOTA	Firmware Over The Air
GOTS	Government Off-the-Shelf
GRE	Generic Routing Encapsulation
HSM	Hardware Security Module
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure
IAD	Information Assurance Directorate
IAVA	Information Assurance Vulnerability Alert
ICMP	Internet Control Message Protocol
ICT	Information Communication Technology
IDS	Intrusion Detection System
IETF	Internet Engineering Task Force
IGMP	Internet Group Management Protocol
IKE	Internet Key Exchange
IP	Internet Protocol
IPS	Intrusion Prevention System
IPsec	Internet Protocol Security
IPv4	Internet Protocol Version 4
IPv6	Internet Protocol Version 6
IS-IS	Intermediate System to Intermediate System
KM	Key Management
MA	Mobile Access
MDF	Mobile Device Fundamentals
MDM	Mobile Device Manager
MOA	Memorandum of Agreement
MLD	Multicast Listener Discovery
MTU	Maximum Transmission Unit
NDP	Neighbor Discovery Protocol
NIAP	National Information Assurance Partnership
NIST	National Institute of Standards and Technology
NPE	Non Person Entity
NSA	National Security Agency
NSS	National Security Systems
NTP	Network Time Protocol
O	Objective
OCSP	Online Certificate Status Protocol
OID	Object Identifier
OPSEC	Operational Security
OS	Operating System
OSI	Open System Interconnection
OSPF	Open Shortest Path First
PKCS	Public Key Cryptographic Standard

Acronym	Definition
PKI	Public Key Infrastructure
PMTU	Path Maximum Transmission Unit
POC	Point of Contact
PTP	Precision Time Protocol
RD	Retransmission Device
RFC	Request for Comment
RIP	Routing Information Protocol
RSA	Rivest Shamir Adelman algorithm
S3	Secure Sharing Suite
SA	Security Association
SCRM	Supply Chain Risk Management
SDES	Session Description Protocol Security Descriptions
SHA	Secure Hash Algorithm
SIEM	Security Information and Event Manager
SIP	Session Initiation Protocol
SIPRNet	Secret Internet Protocol Router Network
SP	Service Packs
SRTP	Secure Real-Time Protocol
SSH	Secure Shell
SSHv2	Secure Shell Version 2
T	Threshold
T&E	Test and Evaluation
TCP	Transmission Control Protocol
TFFW	Traffic Filtering Firewall
TLS	Transport Layer Security
UDP	User Datagram Protocol
URL	Uniform Resource Locator
VDI	Virtual Desktop Infrastructure
VoIP	Voice over Internet Protocol
VM	Virtual Machine
VPN	Virtual Private Network
WLAN	Wireless Local Area Network
WPA2	Wi-Fi Protected Access II

APPENDIX C. REFERENCES

CNSSI 1300	*CNSSI 1300, National Security Systems Public Key Infrastructure X.509 Certificate Policy*	October 2009
CNSSI 4009	*CNSSI 4009, National Information Assurance (IA) Glossary Committee for National Security Systems.* http://www.cnss.gov/Assets/pdf/cnssi_4009.pdf	April 2010
CNSSP 15	*CNSS Policy (CNSSP) Number 15, National Information Assurance Policy on the Use of Public Standards for the Secure Sharing of Information Among National Security Systems Committee for National Security Systems*	March 2010
CNSSD 505	*CNSS Directive (CNSSD) Number 505, Supply Chain Risk Management (SCRM)*	March 2012
FIPS 140	*Federal Information Processing Standard 140, Security Requirements For Cryptographic Modules National Institute for Standards and Technology FIPS Publication* http://csrc.nist.gov/publications/fips/fips140-2/fips1402.pdf	May 2001
FIPS 180	*Federal Information Processing Standard 180-4, Secure Hash Standard (SHS)*	March 2012
FIPS 186	*Federal Information Processing Standard 186-4, Digital Signature Standard (DSS)*	July 2013
FIPS 197	*Federal Information Processing Standard 197, Advanced Encryption Standard (AES)*	November 2001
FIPS 201	*Federal Information Processing Standard 201, Personal Identity Verification (PIV) of Federal Employees and Contractors National Institute for Standards and Technology FIPS Publication* http://csrc.nist.gov/publications/fips/fips201-1/FIPS-201-1-chng1.pdf	March 2006
IPsec VPN Client PP	*Protection Profile for IPsec Virtual Private Network (VPN) Clients*. http://www.niap-ccevs.org/pp	January 2012
NSA Suite B	*NSA Guidance on Suite B Cryptography (including the Secure Sharing Suite (S3))*. http://www.nsa.gov/ia/programs/ suiteb_cryptography/index.shtml	November 2010
RFC 2409	*IETF RFC 2409 The Internet Key Exchange (IKE)*. D. Harkins and D. Carrel.	November 1998
RFC 3647	*IETF RFC 3647 Internet X.509 Public Key Infrastructure Certificate Policy and Certification Practices Framework* Internet Engineering Task Force	November 2003
RFC 3711	*IETF RFC 3711 The Secure Real-Time Transport Protocol (SRTP).* M. Baugher and D. McGrew.	March 2004
RFC 4252	*IETF RFC 4252 The Secure Shell (SSH) Authentication Protocol*. T. Ylonen and C. Lonvick.	January 2006

RFC 4253	*IETF RFC 4253 The Secure Shell (SSH) Transport Layer Protocol*. T. Ylonen and C. Lonvick.	January 2006
RFC 4254	*IETF RFC 4254 The Secure Shell (SSH) Connection Protocol*. T. Ylonen and C. Lonvick.	January 2006
RFC 4256	*IETF RFC 4256 Generic Message Exchange Authentication for the Secure Shell Protocol (SSH)*. F. Cusack and M. Forssen.	January 2006
RFC 4302	*IETF RFC 4302 IP Authentication Header*. S. Kent	December 2005
RFC 4303	*IETF RFC 4303 IP Encapsulating Security Payload*. S. Kent	December 2005
RFC 4307	*IETF RFC 4307 Cryptographic Algorithms for Use in the Internet Key Exchange Version 2 (IKEv2)*. J. Schiller	December 2005
RFC 4308	*IETF RFC 4308 Cryptographic Suites for IPsec*. P. Hoffman	December 2005
RFC 4492	*IETF RFC 4492 Elliptic Curve Cryptography (ECC) Cipher Suites for Transport Layer Security (TLS).* S. Blake-Wilson, N. Bolyard, V. Gupta, C. Hawk Corriente, B. Moeller, and Ruhr-Uni Bochum.	May 2006
RFC 4754	*IETF RFC 4754 IKE and IKEv2 Authentication Using the Elliptic Curve Digital Signature Algorithm (ECDSA)*. D. Fu and J. Solinas.	January 2007
RFC 5246	*IETF RFC 5246 The Transport Layer Security (TLS) Protocol Version 1.2*. T. Dierks and E. Rescorla.	August 2008
RFC 5280	*IETF RFC 5280 Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*. D. Cooper, et. al.	May 2008
RFC 5759	*IETF RFC 5759 Suite B Certificate and Certificate Revocation List (CRL) Profile*. J. Solinas and L. Zieglar.	January 2010
RFC 5996	*IETF RFC 5996 Internet Key Exchange Protocol Version 2 (IKEv2)*. C. Kaufman, et. al.	September 2010
RFC 6188	*IETF RFC 6188 The Use of AES 192 and AES 256 in Secure RTP.* D. McGrew.	March 2011
RFC 6239	*IETF RFC 6239 Suite B Cryptographic Suites for Secure Shell (SSH)*. K. Igoe.	May 2011
RFC 6379	*IETF RFC 6379 Suite B Cryptographic Suites for IPsec*. L. Law and J. Solinas.	October 2011
RFC 6380	*IETF RFC 6380 Suite B Profile for Internet Protocol Security (IPsec)*. K. Burgin and M. Peck.	October 2011
RFC 6460	*IETF RFC 6460 Suite B Profile for Transport Layer Security (TLS)*. M. Salter and R. Housley.	January 2012
RFC 6818	*IETF RFC 6818 Updates to the Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*. P. Yee	January 2013
RFC 7030	*IETF RFC 7030 Enrollment over Secure Transport*. M. Pritikin, P. Yee, and D. Harkins.	October 2013

SP 800-53	*NIST Special Publication 800-53 Rev. 4, Security and Privacy Controls for Federal Information Systems and Organizations*. Joint Task Force Transformation Initiative.	April 2013
SP 800-56A	*NIST Special Publication 800-56A Rev. 2, Recommendation for Pair-Wise Key Establishment Schemes Using Discrete Logarithm Cryptography*. E. Barker, et. al.	May 2013
SP 800-56B	*NIST Special Publication 800-56B, Recommendation for Pair-Wise Key Establishment Schemes Using Integer Factorization Cryptography*. E. Barker, et. al.	August 2009
SP 800-56C	*NIST Special Publication 800-56C, Recommendation for Key Derivation through Extraction-then-Expansion*. L. Chen.	November 2011
SP 800-131A	*NIST Special Publication 800-131A, Recommendation for Transitioning of Cryptographic Algorithms and Key Lengths*. E. Barker.	January 2011
SP 800-147	*NIST Special Publication 800-147, BIOS Protection Guidelines*. D. Cooper, et. al.	April 2011

APPENDIX D. END USER DEVICE IMPLEMENTATION NOTES

Virtual Private Network (VPN) End User Devices (EUDs):

The VPN EUD can be set up using a Computing Device with the user's applications, an Inner VPN Component, and an Outer VPN Component. The Inner VPN Component is a VPN client residing on the same computing device as the user applications. The Outer VPN Component can be a VPN Gateway from the VPN Gateway section of the CSfC Components List (as shown in Figure D-1) or be a VPN Client on the same Computing Device as the user applications (as shown in Figure D-2). If all components are on the same device, virtual machines will be required to provide separate IP stacks for the Inner and Outer VPN Clients as noted in Figure D-2. A retransmission device will also be required in this case, unless, as noted in Section 4.1.3, the connection is to a Wired Network, Domestic Cellular Network, Government Private Wireless Network or a Government Private Cellular network. See Figure D-3.

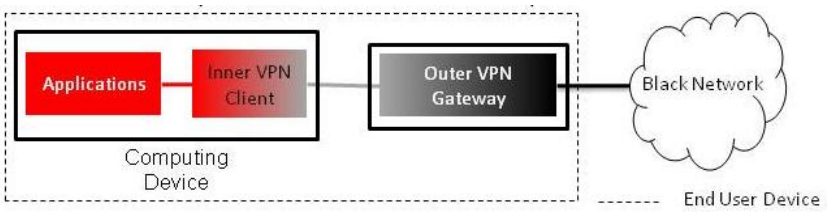


Figure D-1. VPN EUD with Inner VPN Client and Separate Outer VPN Gateway

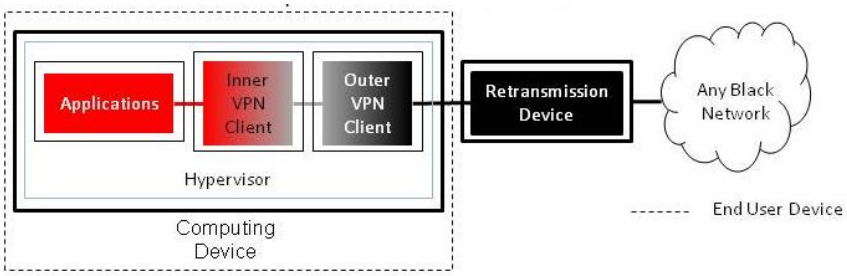


Figure D-2. VPN EUD with Inner and Outer VPN Clients in Separate Virtual Machines with Retransmission Device

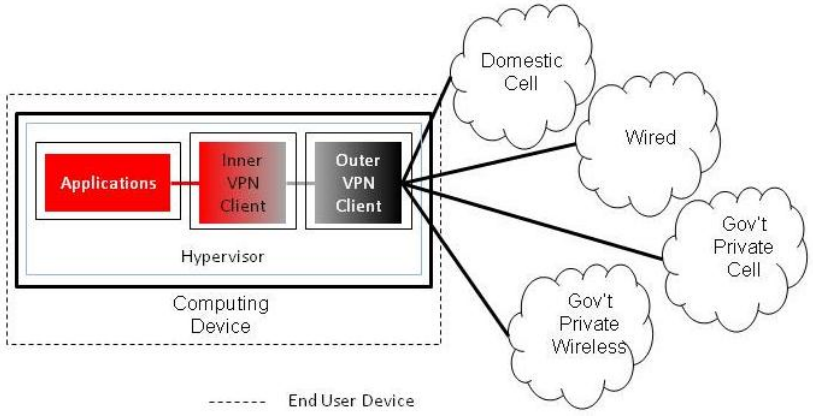


Figure D-3. VPN EUD with Inner and Outer VPN Clients in Separate Virtual Machines without Retransmission Device

Transport Layer Security (TLS) End User Devices:

The TLS EUDs can be set up using up to two separate components. These components consist of the Computing Device and the VPN Component. The Computing Device sends and receives classified data. The Outer VPN Component is either a VPN Gateway or a VPN Client. Outer VPN Gateways are always physically separate from the Computing Device and are selected from the IPsec VPN Gateway section of the CSfC Components List. VPN Clients are selected from the IPSec VPN Client section of the CSfC Components List. The Inner layer of encryption is always provided by an Application on the Computing Device which terminates either TLS and/or SRTP. Each application installed on the Computing Device must be selected from the CSfC Components List. The CSfC Components List provides several sections for which customers can select the TLS Application including Web Browser, Email Client, and VoIP Application. Physical separation between encryption components provides a number of security advantages, but also is more difficult to implement due to the required hardware users require.

For TLS EUDs, each application installed on the Computing Device is responsible for terminating the Inner Layer of Encryption. When a TLS EUD is composed of more than a single computing device, the Outer VPN Component is provided by an Outer VPN Gateway as shown in Figure D-4.

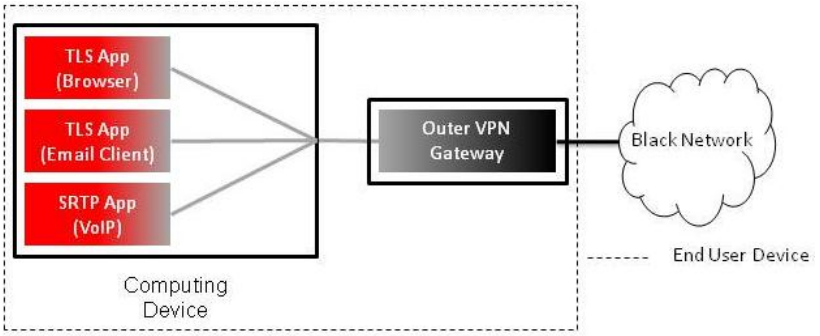


Figure D-4. TLS EUD with Separate Outer VPN Gateway

An Outer VPN client can be installed within the same computing device as the TLS Applications which provide the Inner Layer of Encryption as shown in Figure D-5. A retransmission device will also be required in this case, unless, as noted in Section 4.1.3, the connection is to a Wired Network, Domestic Cellular Network, Government Private Wireless Network or a Government Private Cellular network. See Figure D-6.

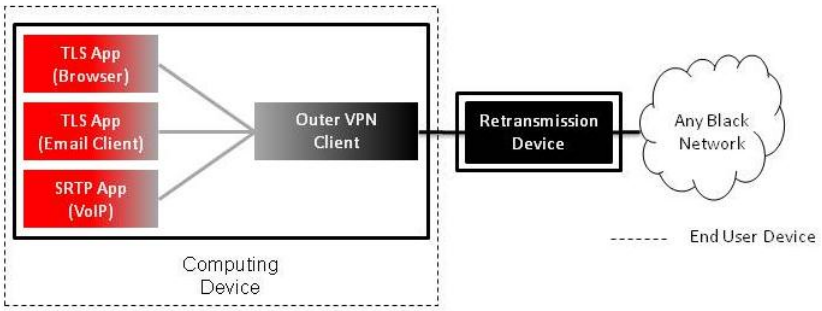


Figure D-5. TLS EUD with Integrated Outer VPN Client with Retransmission Device

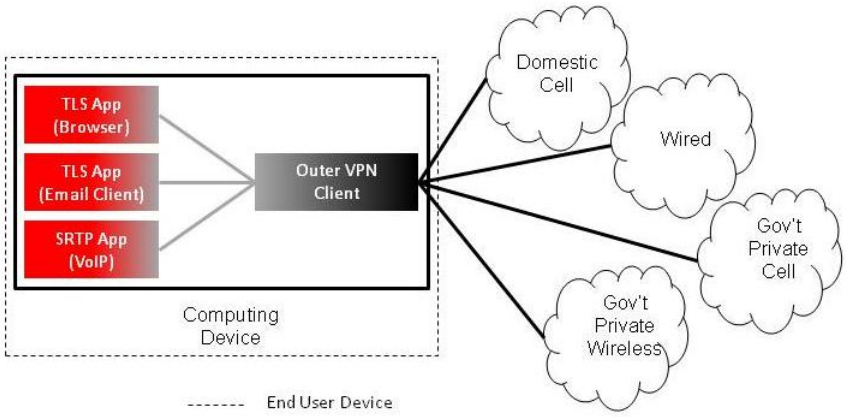


Figure D-6. TLS EUD with Integrated Outer VPN Client without Retransmission Device

Retransmission Devices:

A Government-owned Retransmission Device (RD) includes Wi-Fi Hotspots and Mobile Routers. On the external side, the RD can be connected to any type of medium (e.g. Cellular, Wi-Fi, SATCOM, Ethernet)

to gain access to the Wide Area Network. On the internal side the RD is connected to EUDs either through an Ethernet cable or Wi-Fi. See Figure D-7.

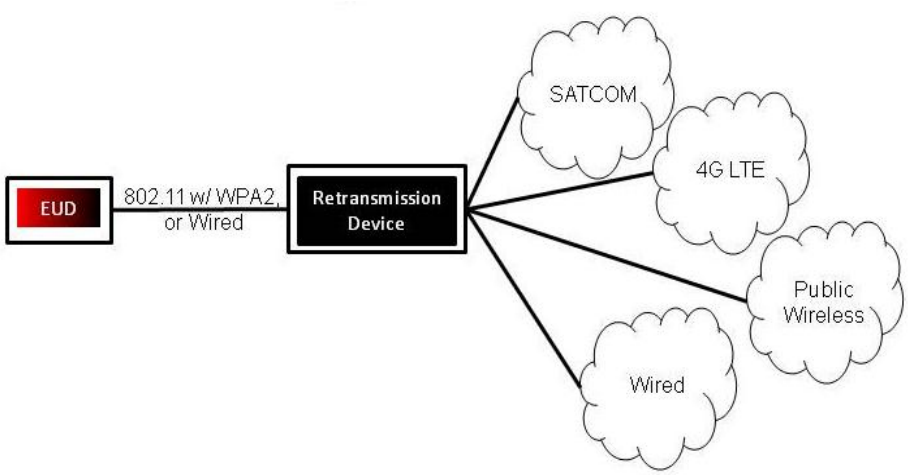


Figure D-7. Retransmission Device Connectivity

Solution Infrastructure supporting VPN and TLS EUDs

When supporting both VPN EUDs and TLS EUDs, the solution infrastructure will always include an Inner VPN Gateway between the Gray Firewall and Inner Firewall (data flow 1 in Figure D-8). Additionally, the solution infrastructure will include one or more TLS-Protected Servers. The TLS-Protected Servers are either placed between the Gray Firewall and Inner Firewall (data flow 2 in Figure D-8) or on the internal side of the Inner Firewall (data flow 3 in Figure D-8). Solutions including TLS-Protected Server(s) in the Enterprise/Red Network, TLS-Protected Server(s) between the Gray Firewall and Inner Firewall, and an Inner VPN Gateway will have at least three distinct data flows. Each Inner Encryption Components is independent and parallel to other Inner Encryption Components.

Figure X below, depicts a Mobile Access Solution which supports both TLS EUDs and VPN EUDs. While not required, the example below shows an example of implementing the TLS-Protected Servers both between the Gray Firewall and Inner Firewall and in the Enterprise/Red network.

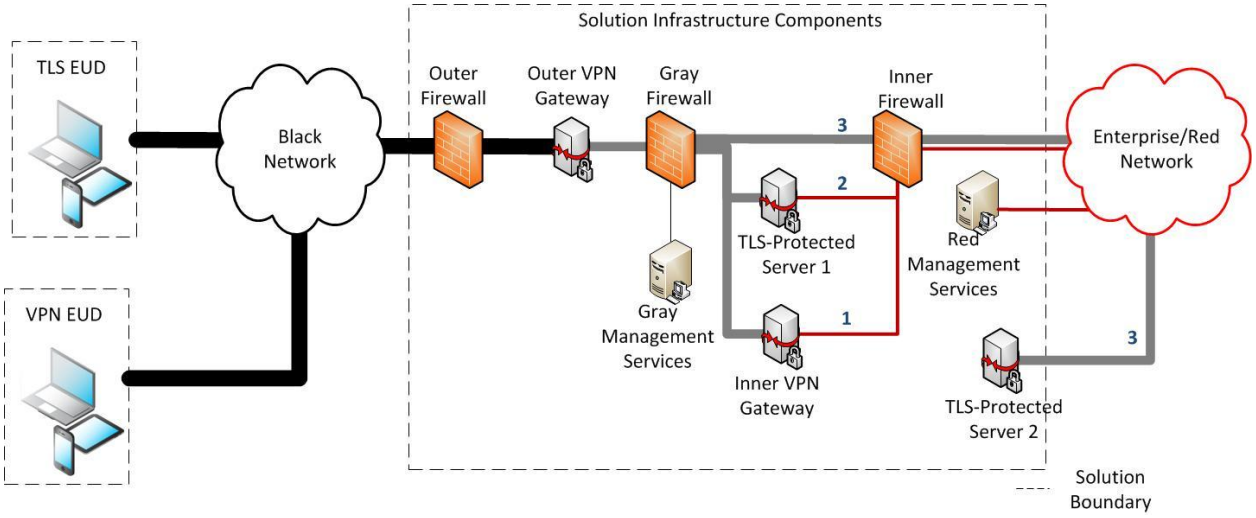


Figure D-8. Mobile Access Solution Infrastructure Supporting VPN and TLS EUDs

The following text describes each of the data flows depicted above.

1. The VPN Gateway terminates the Inner layer of IPSec traffic for all VPN EUDs, and authenticates the EUD VPN client based on device certificates. There is a physical connection between the Gray Firewall and the VPN Gateway and between the VPN Gateway and the Inner Firewall.
2. The TLS-Protected Server is placed between Gray Firewall and Inner Firewall. The TLS-Protected Server terminates the Inner layer of TLS traffic for one or more of the services available to TLS EUDs. The TLS-Protected Server could also be a Session Border Controller which terminates SRTP traffic and relays it to the appropriate destination in the Enterprise/Red network. The TLS-Protected Server authenticates the EUD's TLS client based on user or device certificates. There is a physical connection between the Gray Firewall and the TLS-Protected Server and between the TLS-Protected Server and the Inner Firewall. This connection is in parallel with the VPN Gateway such that the TLS-Protected server is not dependent on the Inner-VPN Gateway to reach the Gray Firewall or the Inner Firewall.
3. When a TLS-Protected Server is placed within the Enterprise/Red network, it is still considered part of the Mobile Access Solution. The TLS-Protected Server terminates the Inner layer of TLS traffic for one or more of the services available to TLS EUDs. The TLS-Protected Server depicted could also be a SRTP Endpoint which terminates SRTP traffic for real-time voice and video services. When the TLS-Protected Server is located in the Enterprise/Red network, there is a direct physical connection between the Gray Firewall and Inner Firewalls. This connection is in parallel with the VPN Gateway such that traffic destined for the TLS-Protected Server is not dependent on the Inner VPN Gateway or any other TLS-Protected Server to reach the Gray Firewall or the Inner Firewall. The Gray Firewall and Inner Firewalls shall be configured with an ACL to only permit traffic to TLS-Protected Servers in the Enterprise/Red network with appropriate source/destination IP addresses and ports.

APPENDIX E. TACTICAL SOLUTION IMPLEMENTATIONS

Although the majority of customers instantiating solutions based on the Mobile Access CP will be utilized for Strategic or Operational Environments, some organizations may deploy the MA CP in Tactical environments. These Tactical Environments include a specific set of Size, Weight, and Power (SWaP) constraints not found in traditional environments.

Organizations intending to deploy an MA CP Solution for Tactical Environments may utilize this Appendix which accommodates the SWaP constraints unique to their environment. This Appendix may only be utilized to protect Tactical Data classified as SECRET or below. The CP follows CNSSI 4009 which defines Tactical Data as "Information that requires protection from disclosure and modification for a limited duration as determined by the originator or information owner." In addition to protecting Tactical Data, organizations which register their solution using this Appendix must be deployed at the Tactical Edge. The CP also follows CNSSI 4009 which defines the Tactical Edge as "The platforms, sites, and personnel (U. S. military, allied, coalition partners, first responders) operating at lethal risk in a battle space or crisis environment characterized by 1) a dependence on information systems and connectivity for survival and mission success, 2) high threats to the operational readiness of both information systems and connectivity, and 3) users are fully engaged, highly stressed, and dependent on the availability, integrity, and transparency of their information systems."

If an organization's planned solution meets the two above criteria then their solution may be registered utilizing the requirement accommodations in this Appendix. The MA CP Registration form must explicitly state that the solution is being utilized in Tactical Environments and provide justification on how the above criteria are met. In general, customers registering with this Appendix will be deployed in support of Battalion and below (or equivalent) unit structure. Typically, these Tactical Environments are located in austere environments where communication infrastructure is generally limited. Due to the lack of existing communication infrastructure, the Tactical Environments are also generally characterized by the use of Government owned Black Infrastructure (Government Private Wireless Networks and/or Government Private Cellular Networks).

The below table may be utilized by customers meeting the above criteria when configuring, testing, registering, and operating their Mobile Access Solution. Any questions on the use of this Appendix should be directed to mobile_access@nsa.gov and csfc@nsa.gov.

Table 31. Tactical Implementation Requirements Overlay

Req #	Requirement Description	Capabilities	Threshold / Objective	Alternative
MA-PS-17	The Outer firewall, Outer VPN Gateway, Gray firewall, Inner Encryption Component, and Inner firewall shall use physically separate components, such that no component is used for more than one function (see Figure 1).	VI, TI	O	MA-TO-1
MA-TO-1	The Outer VPN Gateway shall be physically separate from the Inner Encryption Components	VI, TI	T	MA-PS-17
MA-EU-12	Users of EUDs shall successfully authenticate themselves to the services they access on the Red network using an AO approved method.	VE, TE	O	
MA-EU-13	Red network services shall not transmit any classified data to EUDs until user authentication succeeds.	VE, TE	O	
MA-MR-5	Each IDS in the solution shall be configured to send alerts to the Security Administrator.	All	O	
MA-MR-7	The organization shall create IDS rules that generate alerts upon detection of any unauthorized destination IP addresses.	All	O	
MA-DM-14	The Outer VPN Gateway and solution components within the Gray network shall forward log entries to a SIEM on the Gray Management network (or SIEM in the Enterprise/Red Network if using an AO approved one-way tap) within 10 minutes.	VI, TI	O	

National Security
Agency/Central Security Service

INFORMATION
ASSURANCE
DIRECTORATE

MULTI-SITE CONNECTIVITY CAPABILITY PACKAGE

This Commercial Solutions for Classified (CSfC) Capability Package describes how to protect classified data in transit across an untrusted network using multiple encrypted tunnels implemented with Internet Protocol Security (IPsec), Media Access Control Security (MACsec), or both encryption protocols.

Version 0.8
May 4, 2016

Netskope Exhibit 1002 - Appendix C

CHANGE HISTORY

Title	Version	Date	Change Summary
Commercial Solutions for Classified (CSfC) Multi-Site Connectivity (MSC) Capability Package	0.8	May 4, 2016	Initial release of CSfC Multi-Site Connectivity (MSC) guidance.

TABLE OF CONTENTS

1 Introduction 9
2 Purpose of This Document 9
3 Use of This Document 9
4 Description of the MSC Solution 11
4.1 Networks 11
4.1.1 Red Network 12
4.1.2 Gray Network 12
4.1.3 Black Network 12
4.2 Data, Management, and Control Plane Traffic 13
4.3 High-Level Design 14
4.3.1 Multiple Sites 14
4.3.1.1 Independently Managed Sites 15
4.3.1.2 Centrally Managed Sites 15
4.3.2 Multiple Security Levels 16
4.3.2.1 Networks Operating at the Same Classification Level 17
4.3.2.2 Networks Operating at Different Classification Levels 18
4.3.3 Layering Options 21
4.3.4 Authentication 23
4.4 Other Protocols 23
4.5 Availability 24
5 Solution Components 25
5.1 Outer Encryption Components 25
5.2 Gray Network Firewalls 26
5.3 Gray Management Services 26
5.3.1 Gray Administration Workstation 27
5.3.2 Outer Certificate Authority (Located on Gray Network) 27
5.3.3 Gray Certificate Revocation Services 27
5.3.4 Outer Key Generation Component 28
5.3.5 Gray Security Information and Event Management (SIEM) 28

5.4 Inner Encryption Components 28
5.5 Red Management Services 29
5.5.1 Red Administration Workstation 29
5.5.2 Inner Certificate Authority (Located on Red Network) 29
5.5.3 Red Certificate Revocation Services 30
5.5.4 Inner Key Generation Component (Located on Red Network) 30
5.5.5 Red Security Information and Event Management (SIEM) 30
5.6 Other Controls 31
6 Continuous Monitoring 31
6.1 Monitoring Network Traffic 31
6.2 Monitoring Log Data 34
7 Key Management 34
7.1 Certificates 34
7.1.1 Certificate Issuance, Renewal and Rekey 34
7.1.2 External Distribution of Certificate Revocation Lists 35
7.2 Connectivity Association Keys 37
7.2.1 Connectivity Association Key Issuance, Renewal and Rekey 38
7.2.2 Connectivity Association Key Compromise Recovery 39
8 Threats 39
8.1 Passive Threats 39
8.2 External (Active) Threats 40
8.2.1 Rogue Traffic 40
8.2.2 Malware and Untrusted Updates 41
8.2.3 Denial of Service 41
8.2.4 Social Engineering 42
8.3 Insider Threats 42
8.4 Supply Chain Threats 42
8.5 Integrator Threats 44
9 Requirements Overview 44
9.1 Threshold and Objective Requirements 44

9.2 Requirements Designators 45
10 Requirements for Selecting Components 46
11 Configuration Requirements 48
11.1 Overall Solution Requirements 49
11.2 VPN Gateway Requirements 50
11.3 MACsec Device Requirements 52
11.4 Additional Requirements for Inner Encryption Components 54
11.5 Additional Requirements for Outer Encryption Components 54
11.6 Port Filtering Requirements for Solution Components 55
11.7 Configuration Change Detection Requirements 57
11.8 Device Management Requirements 58
11.9 Continuous Monitoring Requirements 60
11.10 Auditing Requirements 61
11.11 Key Management Requirements 63
11.11.1 General Requirements 63
11.11.2 Certificate Issuance Requirements 65
11.11.3 Certificate Renewal and Rekey Requirements 67
11.11.4 Certificate Revocation Requirements 67
11.11.5 CAK Generation and Distribution Requirements 69
11.11.6 CAK Usage Requirements 70
11.11.7 CAK Update Requirements 71
11.11.8 CAK Compromise Recovery Requirements 71
11.12 Gray Network Firewall Requirements 72
12 Requirements for Solution Operation, Maintenance, and Handling 73
12.1 Requirements for the Use and Handling of Solutions 73
12.2 Requirements for Incident Reporting 75
13 Role-Based Personnel Requirements 77
14 Information to Support AO 79
14.1 Solution Testing 80
14.2 Risk Assessment 81

14.3 Registration of Solutions 81
15 Testing Requirements 82
15.1 Product Selection 82
15.2 Overall Solution 84
15.3 VPN Gateway Configurations 85
15.4 MACsec Device Configurations 86
15.5 Inner and Outer Encryption Component Configurations 87
15.6 Port Filtering 89
15.7 Configuration Change Detection 90
15.8 Continuous Monitoring 90
15.9 Auditing 92
15.10 Key Management 95
15.10.1 Certificate Authorities and Certificates 95
15.10.2 Key Generation Components and Connectivity Association Keys 99
15.11 Gray Network Firewall 101
15.11.1 Gray Network Firewall Filtering Rules 101
15.11.2 Gray Network Firewall HTTP Filtering Rules 102
15.11.3 Gray Network Firewall Management 103
15.11.4 Gray Network Firewall Address Spoofing 104
15.11.5 Gray Network Firewall HTTP Deep Packet Inspection 105
15.12 Incident Reporting Guidance 107
15.13 Implementation of Guidance 107
15.14 Solution Functionality 108
Appendix A. Glossary of Terms 109
Appendix B. Acronyms 113
Appendix C. References 116

TABLE OF FIGURES

Figure 1. Two Encryption Tunnels Protect Data across an Untrusted Network 11
Figure 2. MSC Solution Connecting Two Independently Managed Sites 14
Figure 3. MSC Solution Connecting a Central Management Site and a Remote Site 16
Figure 4. MSC Solution for Two Networks at the Same Classification Level 17
Figure 5. Using Outer Encryption Components for Gray Network Filtering 19
Figure 6. Using Standalone Gray Network Firewalls for Gray Network Filtering 20
Figure 7. Encapsulating MACsec on an Internal Interface 22
Figure 8. Encapsulating MACsec with a Separate Device 22
Figure 9. MSC Solution with Redundant Outer Encryption Components 24
Figure 10. MSC Solution Continuous Monitoring 32
Figure 11. MSC Solution using CRL Distribution Points and OCSP Responders 36

TABLE OF TABLES

Table 1. Layering Options 21
Table 2. Requirement Digraphs 45
Table 3. Product Selection (PS) Requirements 46
Table 4. Overall Solution Requirements (SR) 49
Table 5. IPsec Encryption (Approved Algorithms for Classified) 50
Table 6. VPN Gateway (VG) Requirements 51
Table 7. MACsec Encryption (Approved Algorithms for Classified) 52
Table 8. MACsec Device (MD) Requirements 52
Table 9. Additional Requirements for Inner Encryption Components (IR) 54
Table 10. Additional Requirements for Outer Encryption Components (OR) 54
Table 11. Port Filtering (PF) Requirements for Solution Components 55
Table 12. Configuration Change Detection (CM) Requirements 57
Table 13. Device Management (DM) Requirements 58
Table 14. Requirements for Continuous Monitoring (MR) 60
Table 15. Auditing (AU) Requirements 61
Table 16. General Key Management (KM) Requirements 63
Table 17. Certificate Issuance Requirements 65

Table 18. Certificate Renewal and Rekey Requirements 67
Table 19. Certificate Revocation Requirements 67
Table 20. CAK Generation and Distribution Requirements 69
Table 21. CAK Usage Requirements 70
Table 22. CAK Update Requirements 71
Table 23. CAK Compromise Recovery Requirements 71
Table 24. Gray Network Firewall (FW) Requirements 72
Table 25. Requirements for the Use and Handling of Solutions 73
Table 26. Incident Reporting Requirements (RP) 76
Table 27. Role-Based Personnel Requirements 78
Table 28. Test (TR) Requirements 81

1 INTRODUCTION

The Commercial Solutions for Classified (CSfC) program within the National Security Agency (NSA) Information Assurance Directorate (IAD) uses a series of Capability Packages (CPs) to provide configurations that will allow customers to independently implement secure solutions using layered Commercial Off-the-Shelf (COTS) products. The CPs are vendor-agnostic and provide high-level security and configuration guidance for customers and/or Solution Integrators.

IAD is delivering a generic CSfC Multi-Site Connectivity (MSC) CP to meet the demand for data in transit solutions using approved cryptographic algorithms and National Information Assurance Partnership (NIAP) evaluated components. These algorithms, known as the Commercial National Security Algorithm (CNSA) Suite, are used to protect classified data using layers of COTS products. MSC CP Version 0.8 enables customers to implement layered encryption between two or more sites. This CP takes lessons learned from multiple proof-of-concept demonstrations. These demonstrations included a layered use of COTS products for the protection of classified information.

2 PURPOSE OF THIS DOCUMENT

This CP provides high-level reference designs and corresponding configuration information that allow customers to select COTS products from the CSfC Components Lists for their MSC Solution and then to properly configure those products to achieve a level of assurance sufficient for protecting classified data while in transit. As described in Section 10, customers must ensure that the components selected from the CSfC Components Lists will permit the necessary functionality for the selected capabilities. To successfully implement a solution based on this CP, all Threshold requirements, or the corresponding Objective requirements applicable to the selected capabilities, must be implemented, as described in Section 9.

Customers who want to use this CP must register their solution with NSA. Additional information about the CSfC process is available on the CSfC web page (https://www.nsa.gov/resources/everyone/csfc).

3 USE OF THIS DOCUMENT

This document may not be used for a CSfC solution without formally obtaining support from NSA for the effort prior to presenting a solution to the implementing organization's Authorizing Official (AO). United States Government entities interested in presenting solutions to their AOs in accordance with this guidance must first obtain NSA support by submitting a request for CP application support to their NSA/IAD Client Advocate. In the future, however, customers and their solution providers will be able to use a later version of this guidance to implement solutions without such NSA/IAD involvement. Until that time, customers and solution providers may still register solutions designed according to Version 3.2 of the Virtual Private Network (VPN) CP, dated 20 August 2015; see Section 3 of that document for details.

Please provide comments on usability, applicability, and/or shortcomings to your NSA/IAD Client Advocate and the MSC CP Maintenance Team at msc_cp@nsa.gov.

Committee on National Security Systems (CNSS) Policy No. 15, *National Information Assurance Policy on the Use of Public Standards for the Secure Sharing of Information Among National Security Systems (NSS),* is in the process of being updated to reflect the recently published CNSS Advisory Memorandum (AM) Information Assurance (IA) 02-15. CNSS AM IA 02-15 expands on the guidance contained in CNSS Policy No. 15 and identifies additional public algorithms to protect information within NSS. Specifically, the following algorithms will be required to protect all NSS up to Top Secret:

- AES 256 (confidentiality)
- RSA 3072 or ECDSA P-384 (digital signature and authentication)
- RSA 3072, DH 3072 or ECDH P-384 (key exchange)
- SHA-384 (hashing and integrity)

MSC Solutions shall comply with CNSS policies and instructions. Any conflicts identified between this CP and NSS or local policy should be provided to the MSC CP Maintenance Team.

The following Legal Disclaimer relates to the use of this CP:

This CP is provided “as is.” Any express or implied warranties, including but not limited to, the implied warranties of merchantability and fitness for a particular purpose are disclaimed. In no event shall the United States Government be liable for any direct, indirect, incidental, special, exemplary or consequential damages (including, but not limited to, procurement of substitute goods or services, loss of use, data, or profits, or business interruption) however caused and on any theory of liability, whether in contract, strict liability, or tort (including negligence or otherwise) arising in any way out of the use of this CP, even if advised of the possibility of such damage.

The User of this CP agrees to hold harmless and indemnify the United States (U.S.) Government, its agents and employees from every claim or liability (whether in tort or in contract), including attorney’s fees, court costs, and expenses, arising in direct consequence of Recipient’s use of the item, including, but not limited to, claims or liabilities made for injury to or death of personnel of User or third parties, damage to or destruction of property of User or third parties, and infringement or other violations of intellectual property or technical data rights.

Nothing in this CP is intended to constitute an endorsement, explicit or implied, by the U.S. Government of any particular manufacturer’s product or service.

4 DESCRIPTION OF THE MSC SOLUTION

This CP describes a general MSC Solution to protect classified information as it travels across either an untrusted network or a network of a different classification level. The solution supports interconnecting two or more networks operating at the same security level via encryption tunnels, where the security level encompasses the classification level, list of compartments, dissemination controls, and other such controls over information. The solution provides sufficient flexibility to be applicable to many use cases of MSC implementations.

The MSC Solution uses two nested, independent encryption tunnels to protect the confidentiality and integrity of data as it transits the untrusted network. The two encryption tunnels protecting a data flow can use either Internet Protocol Security (IPsec) generated by a VPN Gateway or Media Access Control Security (MACsec) generated by a MACsec Device. VPN Gateways and MACsec Devices are implemented as part of the network infrastructure.

Throughout this CP, the term "Encryption Component" refers generically to either a VPN Gateway or a MACsec Device. "Inner Encryption Component" refers to the component that terminates the Inner layer of encryption and "Outer Encryption Component" refers to the component that terminates the Outer layer of encryption.

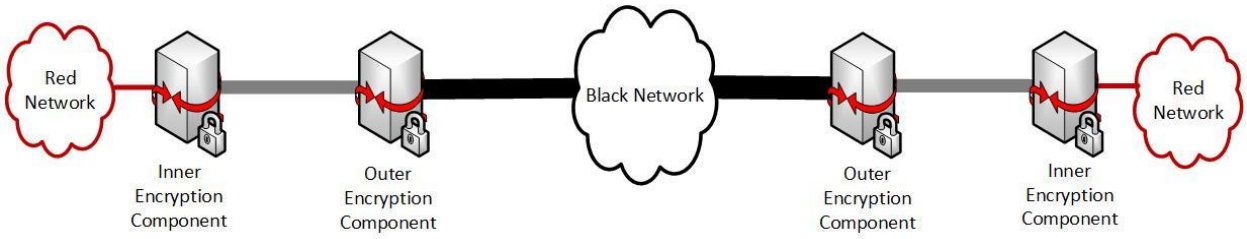


Figure 1. Two Encryption Tunnels Protect Data across an Untrusted Network

As shown in Figure 1, before being sent across the untrusted network, each packet or frame of classified data is encrypted twice: first by an Inner Encryption Component, and then by an Outer Encryption Component. At the other end of the data flow, the received packet is correspondingly decrypted twice: first by an Outer Encryption Component, and then by an Inner Encryption Component.

Products listed on a CSfC Components Lists are not guaranteed to be interoperable with all other products on the Components List. Customers and integrators should perform interoperability testing to ensure the components selected for their MSC Solution are interoperable.

4.1 NETWORKS

This CP uses the following terminology to describe the various networks that comprise a MSC Solution and the types of traffic present on each. The terms Red, Gray, and Black refer to the level of protection applied to the data as described below.

4.1.1 RED NETWORK

Red data consists of unencrypted classified data while Gray data consists of singly encrypted classified data. A Red network contains Red data and can contain Gray data. The Red network is logically located behind an Inner Encryption Component. The networks connected to one another through the MSC Solution are Red networks. Red networks are under the control of the solution owner or a trusted third party. Red networks may only communicate with one another through the MSC Solution if the networks operate at the same security level.

4.1.2 GRAY NETWORK

A Gray network contains classified data that has been encrypted once. The network between an Inner Encryption Component and an Outer Encryption Component is a Gray network. The Gray network is physically and logically under the control of the solution owner or a trusted third party. A MSC Solution compliant with this CP treats a Gray network as a classified network even though all classified data is singly encrypted. If a solution owner's classification authority determines that the data on a Gray network is classified, perhaps by determining the Internet Protocol (IP) addresses used on the Gray network interfaces are classified at some level, then the MSC Solution described in this CP cannot be implemented, as it is not designed to ensure that such information will be afforded two layers of protection.

Gray networks are either physically or cryptographically divided into two sub-networks, as follows:

- Gray Management network – The part of a Gray network that contains the management functions to run components supporting the Outer Encryption Component, including the Outer tunnel Certificate Authority (CA) and the Gray admin and audit server functions. Note, the Inner and Outer CAs can both reside within the Red network.
- Gray Data network – The part of a Gray network that carries data between Inner Encryption Components and Outer Encryption Components.

4.1.3 BLACK NETWORK

A Black network contains classified data that has been encrypted twice. The network connecting the Outer Encryption Components together is a Black network. Black networks are not necessarily (and often will not be) under the control of the solution owner, and may be operated by an untrusted third party.

4.2 DATA, MANAGEMENT, AND CONTROL PLANE TRAFFIC

Data plane traffic is classified information, encrypted or not, that is being passed through the MSC Solution. The MSC Solution exists to encrypt and decrypt data plane traffic. All data plane traffic within the Gray and Black networks is encapsulated within the IPsec and MACsec protocols.

Management plane traffic is used to configure and monitor Solution Components. It includes the communications between a system administrator and a component, as well as the logs and other status information forwarded from a Solution Component to a log server, Security Information and Event Manager (SIEM), or similar repository. Management plane traffic on Red and Gray networks is encapsulated within the Secure Shell version 2 (SSHv2), IPsec, MACsec, or Transport Layer Security (TLS).

Control plane traffic consists of other protocols necessary for the network to function that carry neither data nor management traffic. Control plane traffic is typically not initiated directly on behalf of a user (unlike data traffic) or a system administrator (unlike management traffic). Many, but not all, control plane protocols operate at Layer 2 or Layer 3 of the Open Systems Interconnection (OSI) model. Examples of control plane traffic include, but are not limited to, the following:

- Network address configuration (e.g., Dynamic Host Configuration Protocol (DHCP), Neighbor Discovery Protocol (NDP))
- Address resolution (e.g., Address Resolution Protocol (ARP), NDP)
- Name resolution (e.g., Domain Name System (DNS))
- Time synchronization (e.g., Network Time Protocol (NTP), Precision Time Protocol (PTP)
- Route advertisement (e.g., Routing Information Protocol (RIP), Open Shortest Path First (OSPF), Intermediate System to Intermediate System (IS-IS), Border Gateway Protocol (BGP)
- Certificate status distribution (e.g., Online Certificate Status Protocol (OCSP), Hypertext Transfer Protocol (HTTP) download of Certificate Revocation Lists (CRLs))

In general, this CP does not impose detailed requirements on control plane traffic, although control plane protocols may be used to implement certain requirements. For example, requirements MSC-SR-4 and MSC-SR-5 (see Section 11.1) require that time synchronization be performed, but do not require the use of any particular time synchronization protocol or technique. Notable exceptions are for IPsec session establishment and for certain certificate status distribution scenarios (see Section 7.1.2) where, given their impact on the security of the solution, this CP does provide detailed requirements. Unless otherwise specified in this CP, the usage of specific control plane protocols is left to the solution owner to approve, but any control plane protocols not approved by the solution owner should be disabled.

Data plane and management plane traffic are generally required to be separated from one another by using physical or cryptographic separation. Use of a Virtual Local Area Network (VLAN) alone is not sufficient to separate data plane and management plane traffic. As a result, a solution may, for example, have a Gray Data network and a Gray Management network which are separate from one another, where the components on the Gray Management network are used to manage the components on the Gray Data network. Given that some control plane traffic is necessary for a network to function, there is no general requirement that control plane traffic be similarly separated, unless otherwise specified.

4.3 HIGH-LEVEL DESIGN

The MSC Solution is adaptable to support capabilities for multiple sites and/or multiple security levels, depending on the needs of the customer implementing the solution. If a customer does not have a need for supporting multiple sites or multiple security levels, then those elements need not be included as part of the implementation. However, any implementation of the MSC Solution must satisfy all of the applicable requirements specified in this CP, as explained in Section 9.

4.3.1 MULTIPLE SITES

Figure 2 depicts two Red networks at different sites that operate at the same security level, connected to one another through the MSC Solution. Here, each Red network has two Encryption Components associated with it: an Inner Encryption Component connected to the Red network, and an Outer Encryption Component between the Inner Encryption Component and the Black network.

There are two layers of encryption tunnels between any pair of sites communicating directly with one another: one encryption tunnel between their Outer Encryption Components, and a second encryption tunnel between their Inner Encryption Components. Each set of Inner or Outer Encryption Components can provide encryption using either IPsec or MACsec.

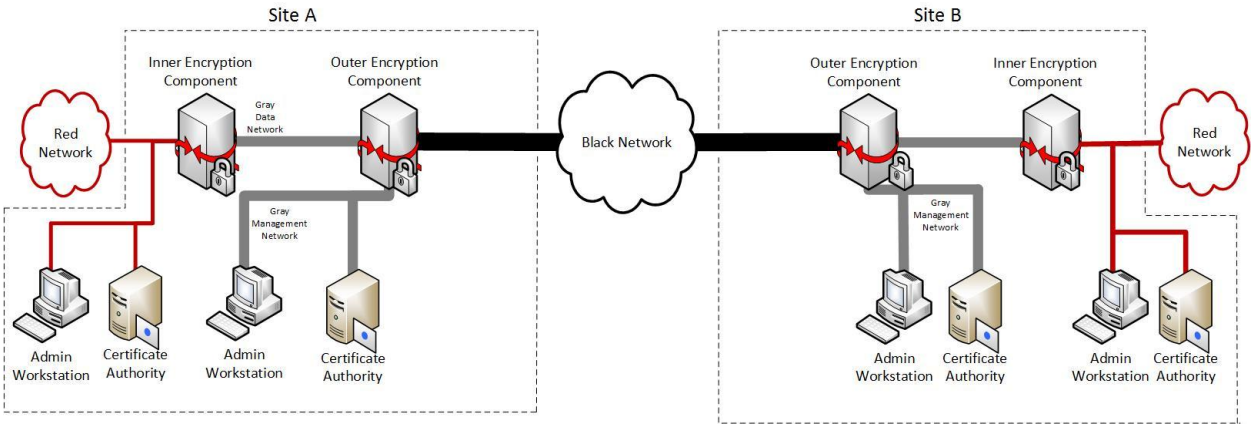


Figure 2. MSC Solution Connecting Two Independently Managed Sites

There is no limit to the number of sites that may be incorporated into a single MSC Solution.

Sites in the solution may be managed independently of one another, or may be remotely managed from a central site.

4.3.1.1 Independently Managed Sites

For independently managed sites, each site performs the administration of its own Encryption Components and has the option of using either locally-run Certificate Authorities (CAs) that they manage and control (see Figure 2) or, where available, enterprise CAs which are not necessarily managed by the solution owner. Each site needs to ensure that the Encryption Components selected interoperate with those at the other sites.

Since there is no remote management, no management traffic will cross the Black network, encrypted or otherwise. Any VPN Gateways at each site need to have the signing certificates and revocation information for the corresponding CAs used by the other sites in the MSC Solution. This high-level design requires cooperation between the various sites in the solution to ensure that all CAs used by each site are trusted at all the other sites. Similarly, any MACsec Devices at each site need to have the same Connectivity Association Key (CAK) used by other sites in the MSC Solution.

This model has the advantage of allowing communication between larger organizations that have a need to share information while maintaining independence.

Note that while Figure 2 depicts only two sites, this solution can scale to include numerous sites, with each additional site having the same design as those in the figure.

4.3.1.2 Centrally Managed Sites

If remote management is used, personnel at a single geographic site administer and perform keying for all the various sites included in the solution, as shown in Figure 3. In this case, because the administration is done by one group of Security Administrators and CA Administrators (see Section 13), they can ensure the interoperability of each site as new sites are added. A maximum of two CAs are needed: one on the Red network for all the Inner VPN Gateways and one on the Gray Management network for all the Outer VPN Gateways. If available, enterprise CAs should be used. If MACsec Devices are being used on either or both layers, CAs are not required since these devices are using CAKs.

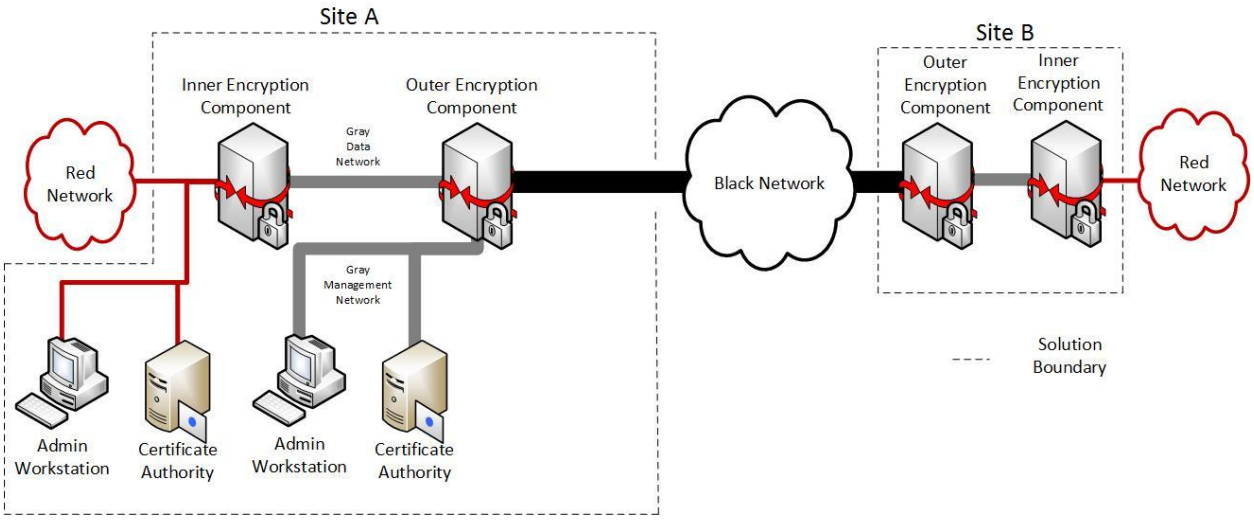


Figure 3. MSC Solution Connecting a Central Management Site and a Remote Site

Because the central management site manages the Encryption Components at the other sites over the network, encryption is used to logically separate data and management traffic as it passes between sites. Gray management traffic is encrypted using SSHv2, TLS, IPsec, or MACsec before being routed through the Outer Encryption Component to another site. The SSHv2, TLS, IPsec or MACsec serves as the inner layer of encryption for Gray management traffic, and the encryption tunnel provided by the Outer Encryption Component serves as the outer layer of encryption. Red management traffic is similarly encrypted before being routed through the Inner and Outer Encryption Components to another site. As a result, all management traffic between sites is encrypted at least twice before traversing the Black network.

This model makes it easier to add sites because of the centralized administration.

Note that while Figure 3 depicts only two sites, this solution can scale to include numerous sites, with each additional site having the same high-level design as the remotely managed site in the figure.

4.3.2 MULTIPLE SECURITY LEVELS

A single implementation of the MSC Solution may support Red networks of different security levels. The MSC Solution provides secure connectivity between the Red networks within each security level while preventing Red networks of differing security levels from communicating with one another. This enables a customer to use the same physical infrastructure to carry traffic from multiple networks. Although each Red network will still require its own Inner Encryption Component, a site may use a single Outer Encryption Component to encrypt and transport traffic that had been encrypted by Inner Encryption Components of varying security levels.

There is no limit to the number of different security levels that a MSC Solution may support. Unclassified networks can also be included behind the Outer Encryption Component.

MSC Solutions supporting multiple security levels may include independently managed sites (see Section 4.3.1.1) or centrally managed sites (see Section 4.3.1.2). Given both cases, separate CAs, CAKs, and management devices are needed to manage the Inner Encryption Components at each security level. For example, Figure 3 depicts a Central Management Site and a Remote Site, but Network 1 and Network 2 each have their own CA, CAKs, and management devices, which prevent their Inner Encryption Components from being able to authenticate with one another.

4.3.2.1 Networks Operating at the Same Classification Level

When Red networks that operate at the same classification level but at different security levels, the cryptographic separation provided by the Inner Encryption Components is sufficient to protect against unintended data flows between security levels. Two Inner Encryption Components for networks of different classification levels will be unable to mutually authenticate with each other because they trust different CAs which do not have a trust relationship with one another or they use different CAKs which will not provide authentication. This prevents the establishment of an encryption tunnel between the two components.

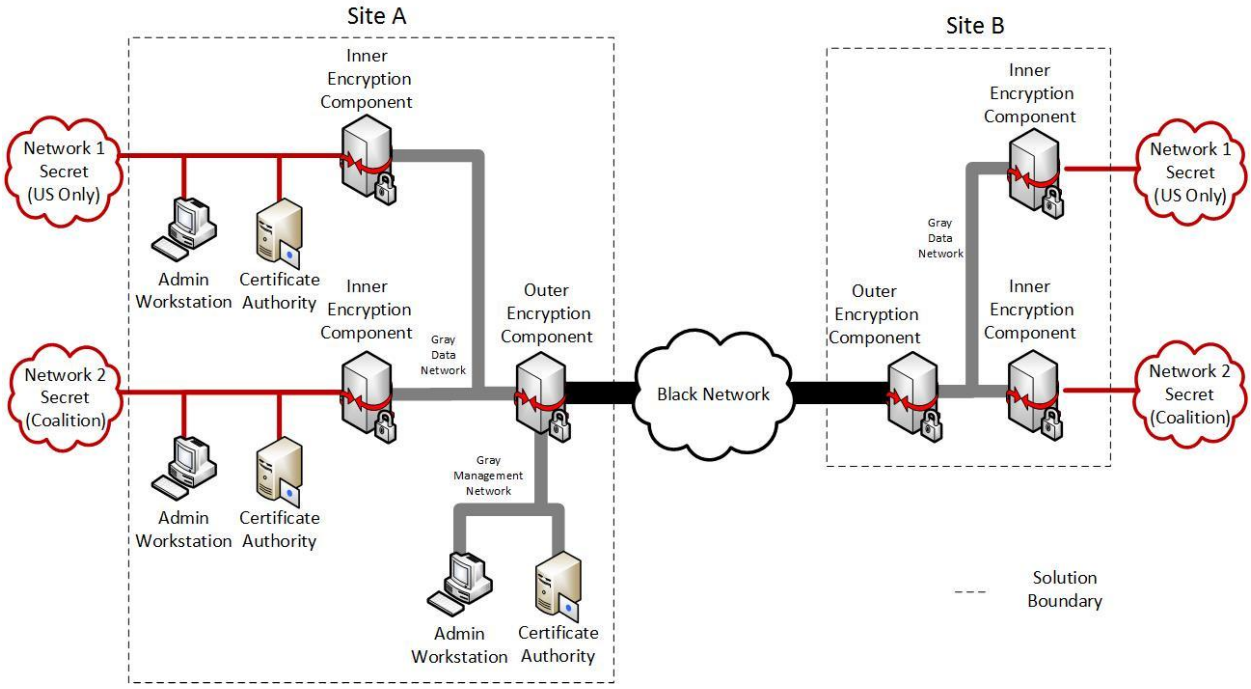


Figure 4. MSC Solution for Two Networks at the Same Classification Level

Figure 4 illustrates a MSC Solution between two sites that carries traffic between two Red networks: a Secret U.S.-only network (Network 1) and a Secret Coalition network (Network 2). Because Network 1 and Network 2 both operate at the Secret classification level, their singly-encrypted traffic can be carried over the Gray network without any additional security controls in place.

Although not required by this CP, a solution owner may choose to implement the additional security described in Section 4.3.2.2 to provide additional protection against unintended data flows between Red networks at the same classification level.

4.3.2.2 Networks Operating at Different Classification Levels

A single implementation of the MSC Solution may support Red networks of different security levels, to include unclassified networks. The MSC Solution provides secure connectivity between the Red networks within each security level while preventing Red networks of differing security levels from communicating with one another. This enables a customer to use the same infrastructure to carry traffic from multiple networks.

For Red networks of different classification levels, the cryptographic separation of their traffic on a Gray network, as described in Section 4.3.2.1, is still present. However, because the consequences of an unintended data flow between different classification levels are more severe than of one with a single classification level, an additional mechanism is necessary to further guard against such a flow from occurring.

This CP uses packet filtering within Gray networks as an additional mechanism to prevent data flows between networks of different classification levels. Any physical path through a Gray network between multiple Inner Encryption Components supporting Red networks of different classification levels must include at least one filtering component. This filtering component restricts the traffic flowing through it based primarily on the Gray network source and destination addresses, only allowing a packet through if the source and destination components are intended to communicate with one another and dropping the packet if they are not.

When multiple classification levels are being used, it is critical to enforce proper IP address assignment and firewall rulesets. The IP address assigned must be unique to that classification level such that each network's Encryption Component is only able to send and receive traffic to its respective Encryption Component at the other site.

Additionally, filtering components are included between the components used for management of the Gray networks themselves (namely, Administration Workstations and locally-run CAs) and Inner Encryption Components that support Red networks of a lower classification level than the highest-classification Red network supported by the solution. In other words, Administration Workstations and locally-run CAs on Gray networks are treated as and grouped with the Inner Encryption Component with the highest-classification Red network.

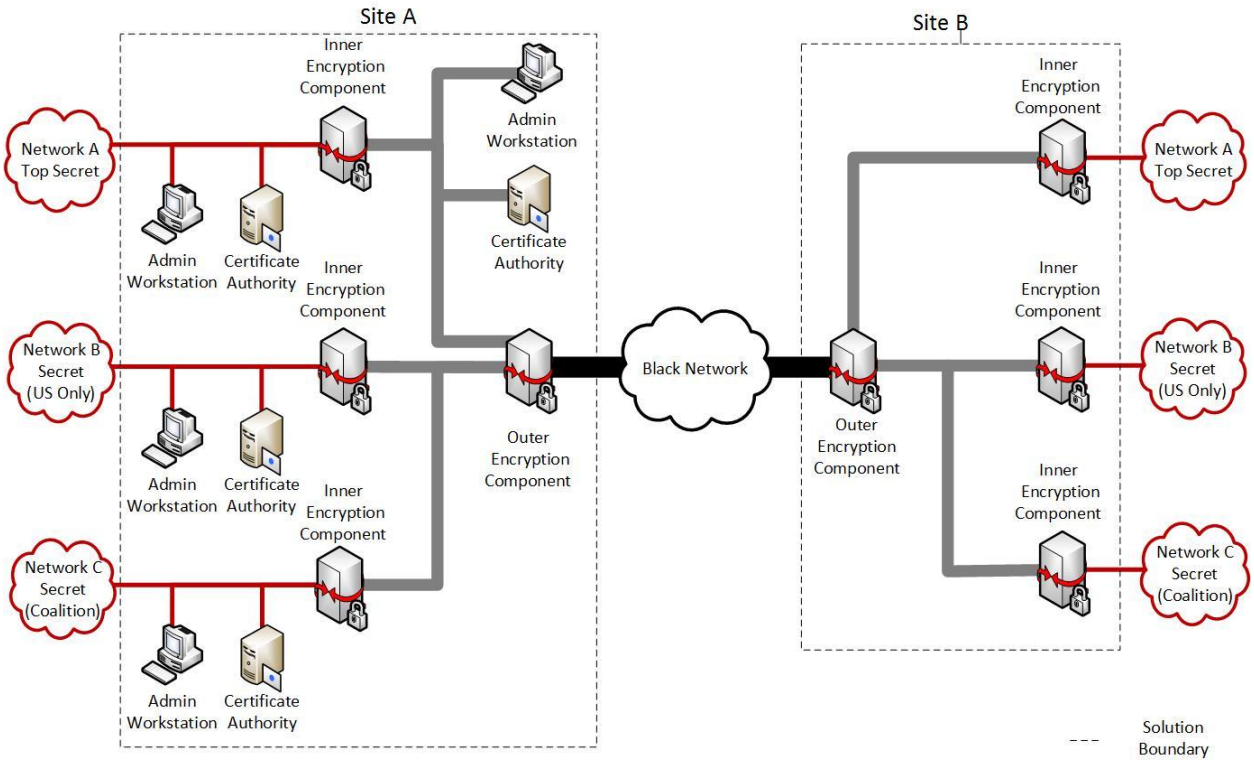


Figure 5. Using Outer Encryption Components for Gray Network Filtering

This CP provides some flexibility in where this filtering takes place within the Gray network. The simplest option is to implement the filtering on the Outer Encryption Components, as shown in Figure 6. Here, the MSC Solution is supporting three Red networks: one Top Secret network (Network 1) and two Secret networks (Networks 2 and 3). Any path through the Gray network between an Inner Encryption Component for a Secret network and an Inner Encryption Component for the Top Secret network (or for the Gray Network's Administration Workstation or CA) includes at least one Outer Encryption Component, which performs the filtering function. Effectively, the Gray network at each site is split into two separate networks which only meet at the Outer Encryption Component: one supporting the Secret networks, and one supporting the Top Secret network (and Gray network management).

Note that there is no filtering component on the path between the Inner Encryption Components for Network 2 and Network 3 within each site, but this is acceptable because Networks 2 and 3 operate at the same classification level; see Section 4.3.2.1 for more details.

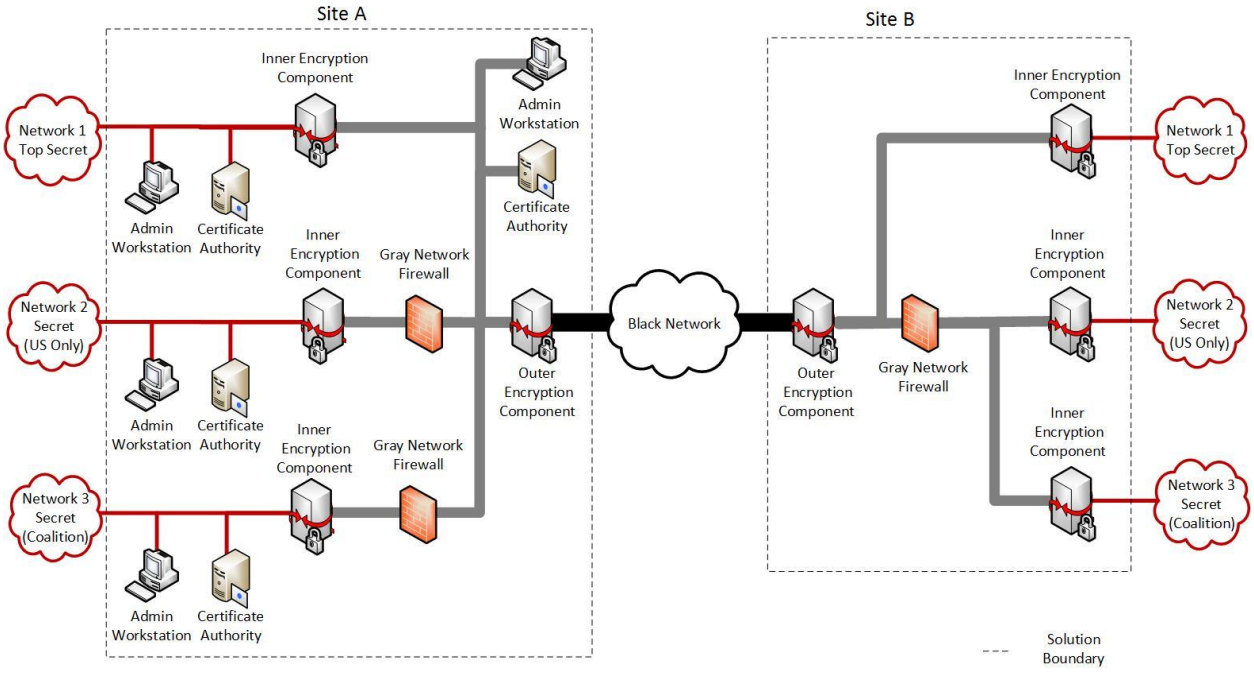


Figure 6. Using Standalone Gray Network Firewalls for Gray Network Filtering

However, some solutions may be subject to physical constraints that prevent relying exclusively on the Outer Encryption Components to provide the filtering function. To accommodate these situations one or more Gray Network Firewalls can be included in a Gray network to perform the filtering in addition to the Outer Encryption Components, as shown in Figure 6. Here, the Gray network is laid out in such a way that the paths between any two Inner Encryption Components within the same site do not pass through an Outer Encryption Component. Instead, standalone Gray Network Firewalls have been placed at each site between the Inner Encryption Components for the Secret networks and the rest of the Gray network; these Gray Network Firewalls are responsible for dropping any packets between an Inner Encryption Component for Network 1 and an Inner Encryption Component for Network 2 or 3.

Figure 6 also illustrates that there is flexibility in the specific placement of Gray Network Firewalls, as long as their placement satisfies the requirement that any path between Inner Encryption Components for networks of different classification levels is met. Site A and Site B in the figure demonstrate two possible placements of Gray Network Firewalls that would satisfy this requirement, although other acceptable placements are also possible.

Including one or more standalone Gray Network Firewalls in a solution does not remove the requirement to perform the filtering on the Outer Encryption Component as well. Outer Encryption Components are uniquely positioned to block traffic between Inner Encryption Components supporting Red networks of different classification levels when one of those Inner Encryption Components is located at a different site.

4.3.3 LAYERING OPTIONS

Each layer of the MSC Solution can use either an IPsec tunnel or MACsec tunnel. An IPsec tunnel is established between VPN Gateways. A MACsec tunnel is established between MACsec Devices. Table 1 identifies four different layering options provided by this CP.

Table 1. Layering Options

Configuration	Inner Tunnel	Outer Tunnel
1	IPsec	IPsec
2	IPsec	MACsec
3	MACsec	IPsec
4	MACsec	MACsec

NOTE: The fourth configuration (MACsec on both layers) may not be available in the final release of this CP. Customers desiring to use a configuration with MACsec on both layers should contact their Client Advocate.

MACsec was designed to provide hop-to-hop security within a Local Area Network (LAN). As MACsec-encrypted traffic arrives at an interface, it is typically decrypted, examined, and re-encrypted after determining its destination.

The MACsec-encrypted traffic needs to be encapsulated if the MACsec Device is the first layer of encryption in the MSC Solution or if the MACsec-encrypted traffic needs to traverse an IP-based network. The reason for encapsulation is to ensure the MACsec-encrypted traffic is not decrypted prior to reaching its destination and to ensure the second layer of encryption can be applied.

In some commercial MACsec Devices, encapsulation can be applied on the internal interface by creating a pseudowire (see Figure 7). If this feature is not supported, a standalone device is needed to encapsulate the MACsec-encrypted data (see Figure 8). If using a standalone device, the internal interface will be connected to the Internal MACsec Device and the external interface will be connected to the Outer Encryption Component. Since this router or switch resides in the Gray network, all requirements for Solution Components must be implemented for it.

This CP does not mandate the use of a specific protocol for encapsulation. Options include, but are not limited to, Layer 2 Tunneling Protocol version 3 (L2TPv3) and Ethernet over Multiprotocol Label Switching (EoMPLS).

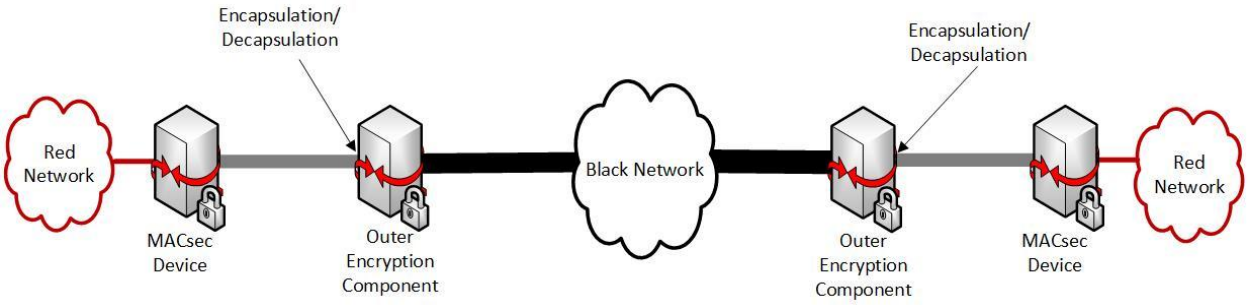


Figure 7. Encapsulating MACsec on an Internal Interface

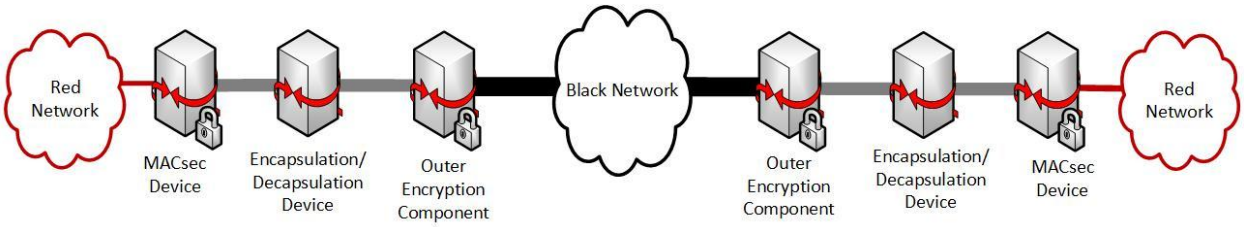


Figure 8. Encapsulating MACsec with a Separate Device

There are some scenarios when a MACsec Device provides the outer tunnel of encryption and the MACsec-encrypted traffic needs to be encapsulated prior to handing it off to the Black network. In these scenarios, this additional step falls outside the boundary of the MSC Solution. However, applying the general requirements for Solution Components is highly recommended.

In the current MACsec standard, the entire frame is encrypted with the exception of the source and destination addresses. Draft amendment IEEE 802.1AEcg provides the option of moving the Virtual LAN (VLAN) identification (ID) tag out of the encrypted payload and into the clear in the header. The benefits of moving the VLAN ID tag into the clear include service multiplexing (i.e., multiple point-to-point or multipoint services existing on a single physical interface) and providing quality of service (QoS) across a Service Provider's network. This CP allows VLAN ID tags to be used in the clear, if supported in the MACsec Device.

At high speeds, some MACsec Devices may be configured to use an eXtended Packet Number (XPN), as described in IEEE 802.1AEbw2013. Without XPN, the key space for the Secure Association Key (SAK) may be exhausted quickly at high speeds and re-keying at high speeds may interrupt traffic flow. This CP allows the XPN feature to be used, if supported in the MACsec Device.

4.3.4 AUTHENTICATION

The MSC Solution provides mutual device authentication between Outer Encryption Components and Inner Encryption Components. The method of authentication is different for VPN Gateways and MACsec Devices.

VPN Gateways authenticate via public key certificates. This CP requires all authentication certificates issued to VPN Gateways to be Non-Person Entity (NPE) certificates. This CP also requires an Inner CA when the Inner Encryption Component is a VPN Gateway and an Outer CA when the Outer Encryption Component is a VPN Gateway.

MACsec Devices authenticate using a Pre-Shared Key (PSK) called a CAK. This CP requires all CAKs and their associated Connectivity Key Names (CKNs) to be generated using an NSA-approved Key Generation Component (KGC). For each MACsec tunnel, a Key Server is identified. The Key Server authenticates the other MACsec Device and provides a SAK to provide confidentiality and integrity for the MACsec tunnel.

4.4 OTHER PROTOCOLS

Throughout this document, when IP traffic is discussed, it can refer to either IPv4 or IPv6 traffic, unless otherwise specified. In addition, Red, Gray and Black networks can run either IPv4 or IPv6, and each network is independent from the others in making that decision. In the remainder of the document, if no protocols or standards are specified then any appropriate protocols may be used to achieve the objective.

Public standards conformant Layer 2 control protocols such as ARP are allowed as necessary to ensure the operational usability of the network. Public standards conformant Layer 3 control protocols such as Internet Control Message Protocol (ICMP) may be allowed based on local AO policy, but the default configuration of this solution is for all Layer 3 control protocols to be disabled. Red and Gray network multicast messages and Internet Group Management Protocol (IGMP) or Multicast Listener Discovery (MLD) may also be allowed depending on local AO policy. Multicast messages received on external interfaces of the Outer Encryption Component shall be dropped.

It is expected that the MSC Solution can be implemented in such a way as to take advantage of standards based routing protocols that are already being used in the network. For example, networks that currently use Generic Routing Encapsulation (GRE) or OSPF protocols can continue to use these in conjunction with this solution to provide routing as long as the AO approves their use.

4.5 AVAILABILITY

The high-level designs described in Section 4.3 are not designed with the intent of automatically providing high availability, and supporting solution implementations for which high availability is important is not a goal. However, this CP does not prohibit adding redundant components in parallel to allow for component failover or to increase the throughput of the MSC Solution, as long as each redundant component adheres to the requirements of this CP.

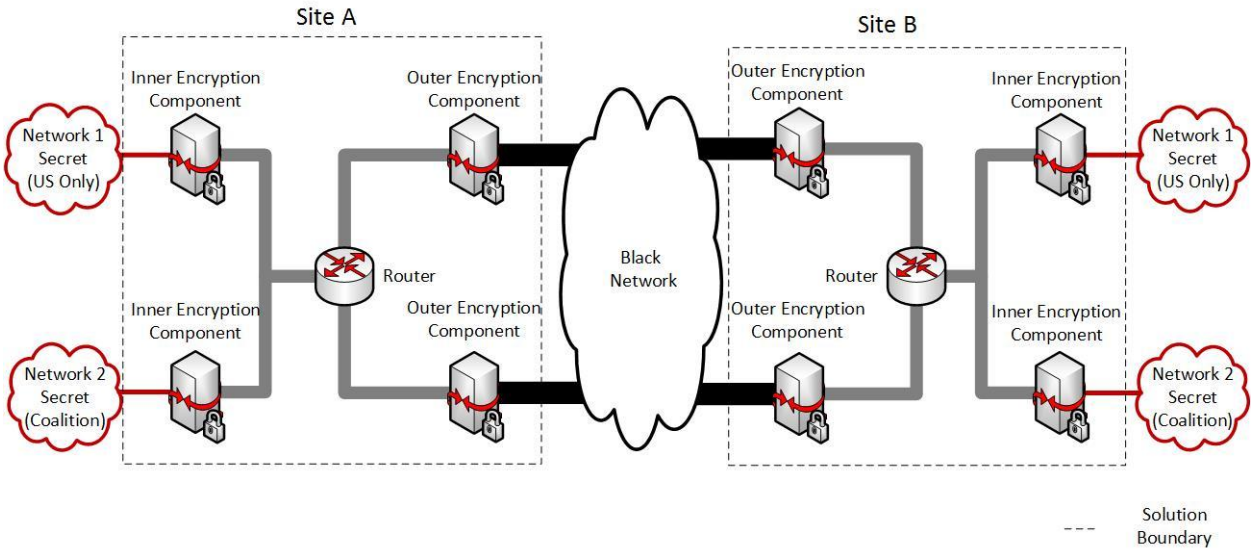


Figure 9. MSC Solution with Redundant Outer Encryption Components

For example, Figure 9 illustrates a MSC Solution between two sites where each site has a redundant Outer Encryption Component. Management components are omitted from the figure for clarity. There are two outer encryption tunnels that transit the Black network: one between the upper pair of Outer Encryption Components, and one between the lower pair of Outer Encryption Components. Each site's Gray network contains an ordinary router between the Inner and Outer Encryption Components which selects which Outer Encryption Component to route outbound packets to. This router is part of the solution only in the sense that it is part of the network infrastructure of the Gray network; this CP does not levy any security requirements on the router/switch. The MSC Solution can maintain connectivity between the two sites even if one of the Outer Encryption Components fails, because traffic will be routed through the tunnel that has not failed.

The above is only a simple example of how redundancy could be added if needed for a MSC Solution. Implementing standby or failover Encryption Components, performing load balancing between Encryption Components, or other techniques to improve the availability or throughput of the solution are outside the scope of this CP and are not discussed further.

5 SOLUTION COMPONENTS

In the high-level designs discussed in the previous section, all communications flowing across a Black network are protected by at least two layers of encryption, implemented using IPsec tunnels generated by VPN Gateways or MACsec tunnels generated by MACsec Devices. Mandatory aspects of the solution also include Administration Workstations, CAs for key management using Public Key Infrastructure (PKI), Key Generation Components (KGCs) for generating CAKs, and Gray Network Firewalls for when networks of different classification levels share the same Outer Encryption Component.

Each Solution component is described in more detail below. The descriptions include information about the security provided by the components as evidence for why they are deemed necessary for the solution. Components are selected from the CSfC Components Lists in accordance with the Product Selection requirements of this CP (see Section 10).

Additional components, discussed in Section 5.6, can be added to the solution to help reduce the overall risk. However, these are not considered mandatory components for the security of the solution; therefore, this CP does not place configuration or security requirements on the components.

5.1 OUTER ENCRYPTION COMPONENTS

The Outer Encryption Component is located at the edge of the private network and can be either a VPN Gateway or a MACsec Device. The Outer Encryption Component establishes an encrypted tunnel using IPsec or MACsec with peer Outer Encryption Components, which provides device authentication, confidentiality, and integrity of information traversing Black networks.

Although the Outer Encryption Component is a perimeter device and thus more exposed to external attacks, the Outer Encryption Component is also capable of protecting the network from unauthenticated traffic through use of an internal filtering capability. This allows specification of rules that prohibit unauthorized data flows, which helps mitigate Denial of Service (DoS) attacks and resource exhaustion. This CP does not require that the Outer Encryption Component terminate all tunnels on a single physical interface; however, all such external interfaces shall conform to the port filtering requirements in Section 11.6. The Outer Encryption Component is implemented identically for all the high-level designs covered in this CP.

Outer Encryption Components are also responsible for filtering traffic on its Gray network interfaces to prevent Inner Encryption Components for networks of different classification levels from being able to send packets to one another. Since this filtering is primarily based on the source and destination addresses in the packet on a Gray network, the Gray networks themselves must use an addressing scheme that supports the necessary filtering (such as using separate address ranges for the Gray interfaces of Inner Encryption Components supporting each Red network). This filtering on Gray network traffic is performed even for solutions that only support Red networks of a single classification level, as in that situation the actual filtering needed to comply with this CP would be simple.

In addition to performing the functions described in this CP, an Outer Encryption Component may also use AO-approved routing protocols on the Gray network it is connected to. The Outer Encryption Component cannot route packets between Gray and Black networks; any packets received on a Gray network interface and sent out a Black network interface must be transmitted within an IPsec or MACsec tunnel configured according to this CP. There is some data that will originate from the Outer Encryption Component (such as control traffic (e.g., Bidirectional Forwarding Detection (BFD)), logging and audit data, which will potentially be sent to a Gray Management network at another site) that will only go through the outer encryption tunnel. This is the only exception to having two layers of encryption for data going over a Black network and is considered acceptable given the limited intelligence value of that information and the fact that it does not contain classified data. However, management traffic on a Gray network, which originates from the Administration Workstation, must include two layers of encryption as described in this CP (see Section 11.8).

For load balance or other performance reasons, multiple Outer Encryption Components that comply with the requirements of this CP are acceptable.

5.2 Gray Network Firewalls

A MSC Solution that supports multiple Red networks may include one or more Gray Network Firewalls, as described in Section 4.3.2.2. The primary purpose of a Gray Network Firewall is to block any packets sent between Inner Encryption Components for Red networks of different classification levels. A Gray Network Firewall also blocks any packets sent between management components on the Gray network and Inner Encryption Components for Red networks that operate at a classification level other than the highest classification level of data protected by the solution. Gray Network Firewalls are physically protected as classified devices.

A standalone Gray Network Firewall, selected from the CSFC Components List, would typically only be used in solutions where the physical design of the Gray network includes paths between Inner Encryption Components for Red networks of different classification levels that do not pass through the Outer Encryption Components. Effectively, each Gray Network Firewall is another instance of the Gray network filtering performed by the Outer Encryption Component. For load balance or other performance reasons, multiple Gray Network Firewalls that comply with the requirements of this CP are acceptable.

5.3 Gray Management Services

Secure administration of components in the Gray network and continuous monitoring of the Gray network are essential roles provided by the Gray Management Services. Gray Management Services are composed of a number of components which each can play a distinct role in the overall security of the solution. This CP allows flexibility in the placement of some Gray Management Services as described below. The Gray Management Services are physically protected as classified devices.

5.3.1 GRAY ADMINISTRATION WORKSTATION

The Gray Administration Workstation is responsible for maintaining, monitoring, and controlling all security functionality for the Outer Encryption Component, Gray Network Firewall, and all Gray Management Service components. The Gray Administration Workstations are not permitted to maintain, monitor, or control Inner Encryption Components or Red Management Services. All MSC Solutions will have at least one Gray Administration Workstation.

Administration Workstations shall be dedicated for the purposes given in this CP, and shall not be used to manage any non-CSfC solutions. As such, a dedicated virtual machine on an administration device used for a non-CSfC solution cannot be used to manage CSfC solutions.

For MSC Solutions with multiple Red networks with different security levels using a single Outer Encryption Component, a separate Administration Workstation is needed to manage the Gray Management Services of each security level.

5.3.2 OUTER CERTIFICATE AUTHORITY (LOCATED ON GRAY NETWORK)

An Outer CA located on the Gray network issues digital certificates for the Outer VPN Gateways in the solution. These certificates are used for authentication in establishing an Outer IPsec tunnel between pairs of VPN Gateways.

If an Outer CA is located in the Gray network and the Inner tunnel is also provided by VPN Gateways, this CP requires a physically separate Inner CA located in the Enterprise/Red network. The Inner CA issues certificates to Inner VPN Gateways. This separation provides key management separation between the two independent layers of encryption.

To improve integration with existing Enterprise Public Key Infrastructure (PKI), this CP allows for flexibility in the placement of CAs. As an alternative to implementing the Outer CA as a Gray Management Service, customers can choose to implement the Outer CA on the Enterprise/Red network (see Section 7.1.1). When the Outer CA is located in the Enterprise/Red network it is not managed by the Gray Management Services.

5.3.3 GRAY CERTIFICATE REVOCATION SERVICES

Certificate Revocation List (CRL) Distribution Points (CDP) and Online Certificate Status Protocol (OCSP) Responders are servers other than a CA that make revocation information available to components. Outer CDPs and OCSP Responders are deployed on the external side of the Outer VPN Gateway for which revocation information is being made available. Collectively Outer CDPs and OCSP Responders are referred to as Gray Certificate Revocation Services. The Gray Certificate Revocation Services ensure the Outer VPN Gateway can verify the status of the certificates used by the Outer VPN Gateway to which it is attempting to establish a connection.

CDPs and OCSP Responders are not required components of this CP, but if not used the organization must implement other means, such as whitelists, to ensure that once a certificate is revoked it cannot successfully establish an Outer IPsec tunnel with the solution infrastructure.

The use of CDPs and OCSP Responders in MSC Solutions is discussed in detail in Section 7.1.2.

5.3.4 Outer Key Generation Component

A locally-operated Outer KGC generates CAKs only for Outer MACsec Devices. An Outer KGC is not a mandatory component of the MSC Solution. CAKs for Outer MACsec Devices can be generated on an Inner KGC.

The use of KGCs in MSC Solutions is discussed in detail in Section 7.2.

5.3.5 Gray Security Information and Event Management (SIEM)

The Gray SIEM collects and analyzes log data from the Outer Encryption Component, Gray Network Firewall, and other Gray Management Service components. Log data may be encrypted between the originating component and the Gray SIEM with SSHv2, TLS, IPsec, or MACsec to maintain confidentiality and integrity of the log data. At a minimum, an auditor reviews the Gray SIEM on a weekly basis. The SIEM is configured to provide alerts for specific events including if the Outer Encryption Component or Gray Network Firewall receives and drops any unexpected traffic which could indicate a compromise. These functions can also be performed on an Enterprise/Red SIEM if AO-approved one-way taps are used as described in this CP (see Section 6.2).

A Gray SIEM is not a mandatory component of the MSC Solution.

5.4 Inner Encryption Components

Inner Encryption Components can be either VPN Gateways or MACsec Devices. For load balance or other performance reasons, multiple Inner Encryption Components that comply with the requirements of this CP are acceptable.

Similar to an Outer Encryption Component, an Inner Encryption Component provides authentication of peer VPN Gateways or MACsec Devices, cryptographic protection of data in transit, and configuration and enforcement of network packet handling rules.

In addition to performing the functions described in this CP, an Inner Encryption Component may also use AO-approved routing protocols on the Red network it is connected to. An Inner Encryption Component shall not route packets between Red and Gray networks; any packets received on a Red network interface and sent to a Gray network interface must be transmitted within an IPsec or MACsec tunnel configured according to this CP.

When an Inner MACsec Device is used, the MACsec traffic will need to be encapsulated prior to being processed by the Outer Encryption Component, regardless of whether it's a VPN Gateway or a MACsec

Device. Some VPN Gateways and MACsec Devices allow this encapsulation to occur on the incoming interface, prior to encrypting traffic for the Outer tunnel. If the selected VPN Gateway or MACsec Device does not have this feature, a separate standalone router or switch is necessary to provide encapsulation and all requirements for Solution Components in this CP shall apply to it. Any AO-approved encapsulation protocol may be used.

5.5 RED MANAGEMENT SERVICES

Secure Administration of Inner Encryption Components is provided by the Red Management Services. Red Management Services are composed of a number of components which can each play a distinct role in the overall security of the solution. This CP allows flexibility in the placement of some Red Management Services as described below.

5.5.1 RED ADMINISTRATION WORKSTATION

The Red Administration Workstation is responsible for maintaining, monitoring, and controlling all security functionality for the Inner Encryption Components and all Red Management Service components. Red Administration Workstations are not permitted to maintain, monitor, or control Outer Encryption Components or Gray Management Services. All MSC Solutions will have at least one Red Administration Workstation.

Administration Workstations shall be dedicated for the purposes given in this CP, and shall not be used to manage any non-CSfC solutions. As such, a dedicated virtual machine on an administration device used for a non-CSfC solution cannot be used to manage CSfC solutions.

5.5.2 INNER CERTIFICATE AUTHORITY (LOCATED ON RED NETWORK)

MSC Solutions will always have at least one CA located in the Red network. At a minimum an Inner CA is included in the Red network to issue digital certificates for the Inner VPN Gateways in the solution. These certificates are used for authentication in establishing the inner IPsec tunnel between pairs of Inner VPN Gateways.

This CP also allows Outer CAs to be included in the Red network to issue digital certificates for the Outer VPN Gateways. These certificates are used for device authentication in establishing the outer IPsec tunnel between pairs of VPN Gateways. When an Outer CA is placed in the Red network it is critical to have AO-approved mechanisms in place to transfer revocation information to the Black network to ensure it is accessible to peer Outer VPN Gateways. When an Enterprise PKI capability is used, it is managed with existing processes and capabilities. Enterprise CAs provide certificate management services for the MSC Solution over the Enterprise/Red network.

If the solution is supporting Red networks of different security levels, then a separate CA is needed for the Inner VPN Gateways of each security level.

Each CA used in the solution shall have an approved Certificate Policy/ Certificate Practice Statement (CP/CPS) that addresses certificate generation, handling, distribution, storage, destruction, and key recovery and compromise recovery. Refer to NIST SP 800-57 for guidance.

If the organization has existing enterprise CAs that satisfy the requirements of this CP, those CAs should be used as part of the MSC Solution rather than setting up new CAs dedicated to this solution. Otherwise a locally managed CA will need to be deployed requiring that a CA product be selected from the CSfC Components List.

5.5.3 RED CERTIFICATE REVOCATION SERVICES

CDPs and OCSP Responders are servers other than a CA that make revocation information available to components. Red CDPs and OCSP Responders make revocation information available to Inner VPN Gateways. Collectively Red CDPs and OCSP Responders are referred to as Red Certificate Revocation Services. The Red Certificate Revocation Services ensure the Inner VPN Gateway can verify the status of the certificates used by the Inner VPN Gateway to which it is connecting.

CDPs and OCSP Responders are not required components of this CP, but if not used the organization must implement other means, such as whitelists, to ensure that, once a certificate is revoked, it cannot successfully establish an Inner Tunnel with a VPN Gateway.

The use of CDPs and OCSP Responders in MSC Solutions is discussed in detail in Section 7.1.2.

5.5.4 INNER KEY GENERATION COMPONENT (LOCATED ON RED NETWORK)

MSC Solutions will always have at least one KGC located in the Red network. At a minimum an Inner KGC is included in the Red network to issue CAKs for the Inner MACsec Devices in the solution. These CAKs are used for authentication in establishing the MACsec tunnel between a pair of Inner MACsec Devices.

This CP also allows the Inner KGC to issue CAKs for the Outer MACsec Devices. These CAKs are used for device authentication in establishing the MACsec tunnel between a pair of MACsec Devices. When the Inner KGC issues CAKs for Outer MACsec Devices, it is critical to have AO-approved mechanisms in place to transfer the CAKs to the Outer MACsec Devices.

The use of KGCs in MSC Solutions is discussed in detail in Section 7.2.

5.5.5 RED SECURITY INFORMATION AND EVENT MANAGEMENT (SIEM)

Red SIEMs collect and analyze log data from the Inner Encryption Components and other Red Management Service components. Log data may be encrypted between the originating component and the Red SIEM with SSHv2, TLS, IPsec, or MACsec to ensure confidentiality and integrity. At a minimum an auditor reviews the Red SIEM on a weekly basis. The SIEM is configured to provide alerts for specific events including if the Inner Encryption Components receive and drop any unexpected traffic which could indicate a compromise.

While Red SIEMs are not a mandatory component of the MSC Solution, customers are encouraged to leverage existing Enterprise SIEM capabilities to monitor log data from Inner Encryption Components and Red Management Services. Although a Red SIEM is not required, it is still required to analyze logs from all Inner Encryption Components on at least a weekly basis. A Red SIEM may also be used to analyze log data from Gray network components when used in conjunction with one-way taps as described in this CP (see Section 6.2).

5.6 OTHER CONTROLS

There are additional controls that could be used within this solution to potentially reduce the overall risk. A screening router can be used to filter packets from Black networks before they arrive at Outer Encryption Components and Outer CDPs or OCSP Responders. The screening router could be part of an existing Black network (e.g., Customer Edge Router), or could be added between Outer Encryption Components and existing Black network components. However, since the screening router would become part of a Black network, it is not considered to be part of the MSC Solution itself.

Additionally, if an integrator is used for implementation of this solution, the customer can require separation of roles between individuals working on Red and Gray components. The separation of roles ensures that during the development of the solution no single individual can compromise Red and Gray components simultaneously.

6 CONTINUOUS MONITORING

Continuous monitoring allows customers to detect, react to, and report any attacks which occur on their solution. This continuous monitoring also enables the detection of any configuration errors to Solution Components.

At a minimum, this CP requires an Auditor to review alerts, events, and logs on a weekly basis. This minimum review period allows customers in Tactical Environments to implement solutions where it may not be feasible to perform real-time monitoring. Operational and Strategic implementations of the MSC Solution should review alerts, events, and logs on a much more frequent period and in many cases may leverage Operations Centers to perform 24/7 monitoring of the solution.

6.1 MONITORING NETWORK TRAFFIC

This CP recommends monitoring network traffic in at least one of three areas within the solution infrastructure. Network traffic can be monitored using an Intrusion Detection System (IDS); however, it is preferable to use an Intrusion Prevention System (IPS) to enable real-time responses. While monitoring only one of the three locations is recommended, customers monitoring all three points have the best visibility enabling detection of malicious activity or misconfiguration of components.

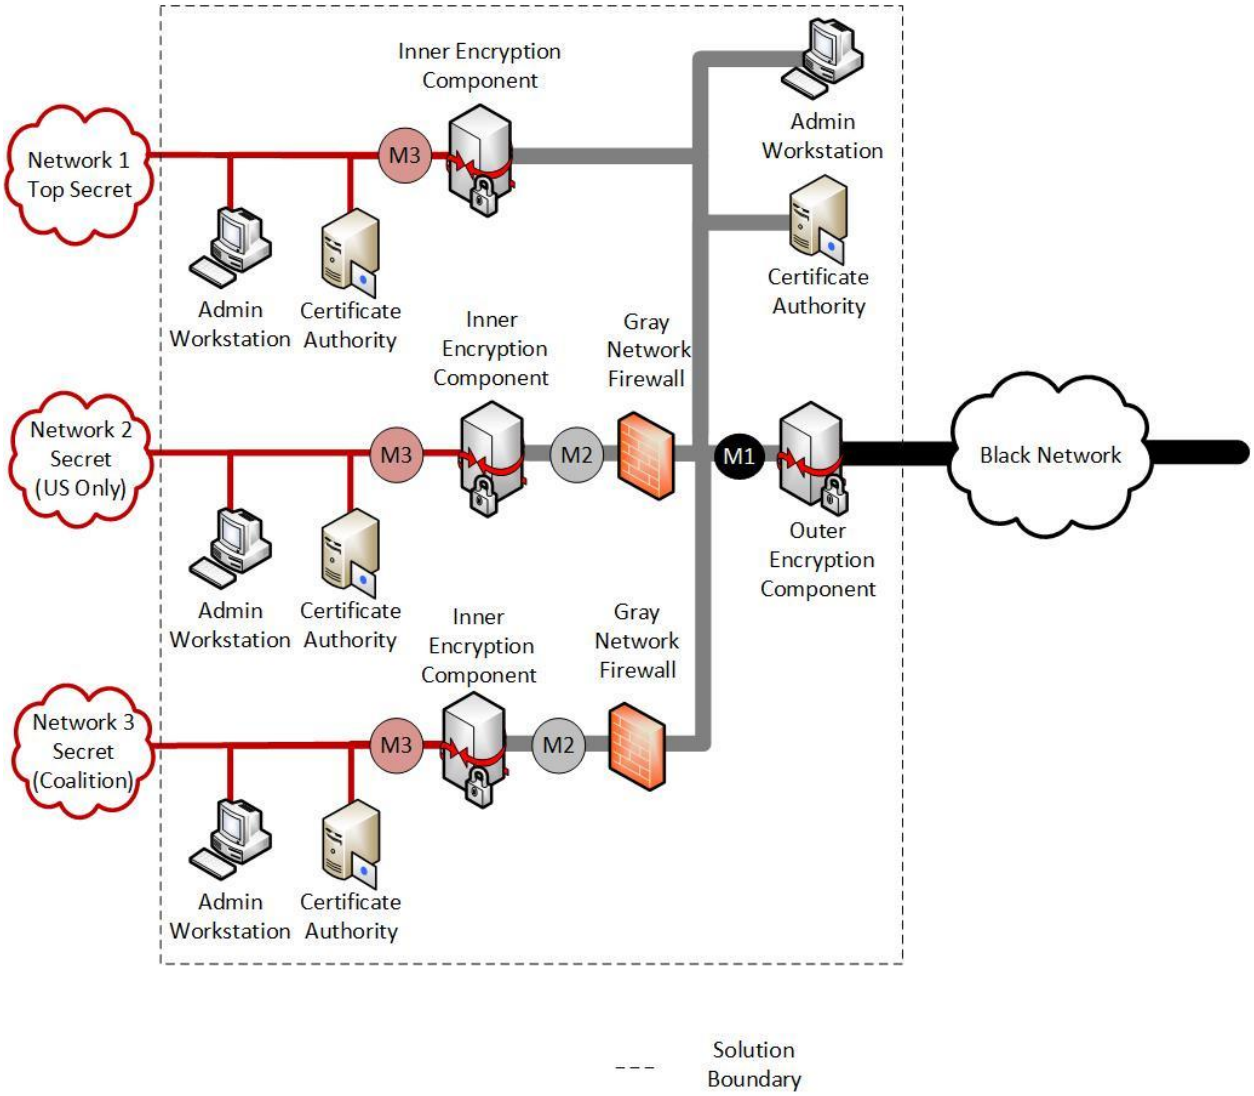


Figure 10. MSC Solution Continuous Monitoring

Figure 10 depicts three locations that customers can select to implement network monitoring capabilities. There are several alternatives for deploying the IDS at one or all of the Monitoring Points (M1, M2, and/or M3). IDSs can ingest traffic from network taps, Switched Port Analyzers (SPANs), or in line with the solution. Similarly, an IPS can be placed either inline or in promiscuous mode. When operating in promiscuous mode, the IPS analyzes traffic at M1, M2, or M3 and issues commands to firewalls (standalone or integrated) to block traffic flows. The following paragraphs define each of the three Monitoring Points. These descriptions detail the analysis and alerts which would be generated by the IDS. If a customer decides to implement an IPS, then it should be configured to block that traffic flow and also send an alert.

Monitoring Point 1 (M1) is located on the internal side of the Outer Encryption Component. At a minimum, a M1 IDS is configured to send alert upon detection of any traffic which should have been blocked by the Outer Encryption Component. These alerts indicate a failure of the filtering or encryption functions of the Outer Encryption Component and are either indicative of an improper configuration or a potential compromise. Normal traffic at M1 includes IPsec, MACsec, control plane traffic, and management traffic. It is important that customers understand what ports and protocols are used by each of the Inner Encryption Components to ensure that filters are properly configured and IDSs are well tuned. Nearly all traffic traversing M1 is encrypted either with IPsec, MACsec, or SSHv2, which prevents the ability to perform deep packet inspection. Management of a M1 IDS occurs from a Gray network.

Monitoring Point 2 (M2) is located between the Gray Network Firewall and the Inner Encryption Component when networks of different classification levels share the Outer Encryption Component. At a minimum, a M2 IDS is configured to send an alert upon detection of any traffic which should have been blocked by the Gray Network Firewall, especially traffic from the network with the higher classification. These alerts indicate a failure of the filtering functions by the Gray Network Firewall and are either indicative of an improper configuration or a potential compromise. Normal traffic at M2 includes IPsec, control plane traffic, and management traffic. It is important that customers understand what ports and protocols are used by each of the Inner Encryption Components to ensure that firewall filters are properly configured and IDSs are well tuned. Nearly all traffic traversing M2 is encrypted either with IPsec, MACsec or SSHv2, which prevents the ability to perform deep packet inspection. Management of a M2 IDS occurs from the Gray Management Services.

Monitoring Point 3 (M3) is located between the Inner Encryption Component and the Enterprise/Red network. At a minimum, a M3 IDS is configured to send an alert upon detection of any traffic which should have been blocked by the Inner Encryption Component. These alerts indicate a failure of the filtering function of the Inner Encryption Component. Of the three monitoring points, M3 is the most difficult to define a normal baseline, but in many implementations, monitoring at M3 allows for deep packet inspection since traffic is not encrypted. Management of the M3 IDS occurs from the Red Management Services.

Monitoring Multiple Points: Although this CP only requires monitoring of one out of the three points, customers are encouraged to monitor all three locations. To ensure that no bypass of the Outer or Inner Encryption Components occurs, customers can implement three IDS(s)/IPS(s). Implementation of three separate components to monitor each point ensures that malicious traffic cannot inadvertently be transferred to the Enterprise/Red network, but it also increases cost and complexity for managing the solution. This approach also prevents correlation of data as it traverses throughout the solution.

Alternatively, to allow correlation of data at multiple points this CP allows a one-way tap located at M1 or M2 to feed the IDS located in an enclave of the Red Network that is isolated from the Enterprise/Red network. The one-way tap used must physically only allow data flow from M1 or M2 to the Enterprise/Red network. Additionally, the selected one-way tap must be approved by the AO.

Movement of network traffic from M3 to the Gray network is explicitly prohibited. Prior to implementing one-way taps as part of the MSC Solution, AOs need to be aware of the Residual Risks associated with transferring data from Gray networks to a Red network of a higher classification.

6.2 MONITORING LOG DATA

SIEMs are not a mandatory components of the MSC Solution. However, customers are still required to analyze logs from all Solution Components on at least a weekly basis.

SIEMs collect, aggregate, correlate, and analyze log data from Solution Components and provide alerts to Auditors when anomalous behavior is detected.

To protect the integrity of the data, all logs sent to the SIEM should be encrypted with SSHv2, TLS, IPsec or MACsec. At a minimum the SIEM is configured to send alerts upon receiving a log entry for blocked packets from an Encryption Component or Gray Network Firewall.

To allow correlation of data from both Gray and Red components, this CP allows a one-way tap located in the Gray network to feed Gray Log Data to a Red SIEM located in an enclave of the Red network that is isolated from the Enterprise/Red network. The one-way tap used must physically only allow log data to flow from Gray Components to the Enterprise/Red network. Additionally, the selected one-way tap must be approved by the AO.

7 KEY MANAGEMENT

One of the most difficult parts of any solution is determining how the key management will be implemented in a secure manner. In the MSC Solution, certificates are used for VPN Gateways and CAKs are used for MACsec Devices.

7.1 CERTIFICATES

This section provides details on issuing, renewing, rekeying and revoking certificates for VPN Gateways.

7.1.1 CERTIFICATE ISSUANCE, RENEWAL AND REKEY

In this solution, the only certificates necessary are for the device authentication certificates on VPN Gateways at the end of each IPsec tunnel. The certificates and private keys used in the solution are considered Controlled Unclassified Information (CUI), because they are only used for mutual device authentication, not for traffic encryption.

If IPsec is securing both tunnels, no single CA can provide keys to both Inner and Outer VPN Gateways. The CAs for Outer VPN Gateways are located on a Gray Management network, connected to an Outer VPN Gateway. A locally-run CA may need to be stood up to key Outer VPN Gateways, requiring that a CA product be selected from the NSA-approved CSfC Component List for Outer tunnel PKI. In addition, a Certificate Policy/Certification Practice Statement (CPS) document must be created or tailored for each CA used in the solution. It is then the AO's responsibility to approve the use of a CA. If an Outer tunnel

CA is an Enterprise CA already running on the necessary Gray Management network, no additional approval is necessary for use of the CA.

CAs for Inner VPN Gateways are located on Red networks, which may allow for use of existing Enterprise CAs already operational on the Red network. For networks in which an existing Enterprise CA is not available, the use of a locally-run CA on the Red network is an acceptable alternative. If an Inner tunnel CA is an Enterprise CA already running on a Red network, no additional approval is necessary for use of this CA. For example, a solution may use an Enterprise CA (such as a CNSS-approved CA, which follows CNSS Instruction (CNSSI) 1300 under the National Security Systems (NSS) PKI Root CA, to issue certificates to an Inner VPN Gateway. If, however, an Inner tunnel CA uses a locally-run CA on a Red network, the approval process given in the preceding paragraph for an Outer tunnel CA applies and must be followed.

Each VPN Gateway has at least one CA signing certificate (sometimes referred to as a Trust Anchor), which is used by the VPN Gateway to authenticate to other VPN Gateways in the solution. If centralized management is used throughout the solution, there will be only one CA signing certificate in each VPN Gateway. Otherwise, one CA signing certificate is installed in each Inner VPN Gateway for each Inner Tunnel CA used in the system. Similarly, one CA signing certificate will be installed in each Outer VPN Gateway for each Outer Tunnel CA used in the system.

Each VPN Gateway will contain a private key that corresponds to a certificate issued by its CA, and one or more CA signing certificates as described above. Each VPN Gateway will also contain revocation information. The private key may be locally generated and must be adequately protected. Both Inner and Outer tunnel PKIs shall use Elliptic Curve Digital Signature Algorithm (ECDSA) over the curve P-384 with SHA-384 or RSA 3072 bit or greater signatures within X.509 certificates. The algorithms and elliptic curves that are approved for use in the VPN Gateways are found in Table 5 (see Section 11.2).

The MSC Solution described here requires certificates to establish the secure tunnels between VPN Gateways. Without certificates, the network cannot function. Thus, an out-of-band method must be used to issue the initial certificates to the VPN Gateways. Subsequent rekeying, however, should take place over the network through this solution prior to the current key's expiration. The key validity period for certificates issued by locally-run CAs cannot exceed 14 months, while the key validity period for certificates issued by an Enterprise CA are inherited from the Enterprise CA certificate policy. Updates to CRLs are distributed to VPN Gateways within 24 hours of CRL issuance.

7.1.2 External Distribution of Certificate Revocation Lists

Part of the security of the MSC solution depends on the certificate-based mutual authentication that occurs between two VPN Gateways establishing an IPSec tunnel. One step of this mutual authentication entails checking whether the certificate used by the other VPN Gateway has been revoked, which requires each VPN Gateway to have access to a current CRL. As the number of sites interconnected through the MSC Solution increases, out-of-band CRL distribution becomes increasingly burdensome

and error-prone. Although VPN Gateways may retrieve the latest CRL directly from the appropriate CA, for Remote Sites this requires first establishing an IPsec tunnel to the Central Management Site where the CAs are located. These additional IPsec tunnels increase the time needed to establish an IPsec tunnel between two Remote Sites. Furthermore, the Remote Site's VPN Gateways still require out-of-band CRL distribution to be able to check for revocation of the certificates used by VPN Gateways at the Central Management Site, since the IPsec tunnel to the Central Management Site must be established before the CRLs can be obtained from the CAs.

To avoid these issues, this CP permits the distribution of CRLs on the external side of the VPN Gateways, which allows the VPN Gateways to retrieve the current CRL without first establishing an IPsec tunnel. A CDP or OCSP Responder resides on a different network than the CA that produced the CRL it hosts. An Outer CDP or OCSP Responder resides on a Black network and hosts a CRL created by the CA on a Gray network. Similarly, an Inner CDP or OCSP Responder resides on a Gray network and hosts a CRL created by the CA on a Red network. Because the CDP or OCSP Responder and its CA reside on different networks, a one-way transfer mechanism is needed to periodically distribute the current CRL from the CA to the CDP or OCSP Responder; the details of the one-way transfer mechanism are left to a solution's AO.

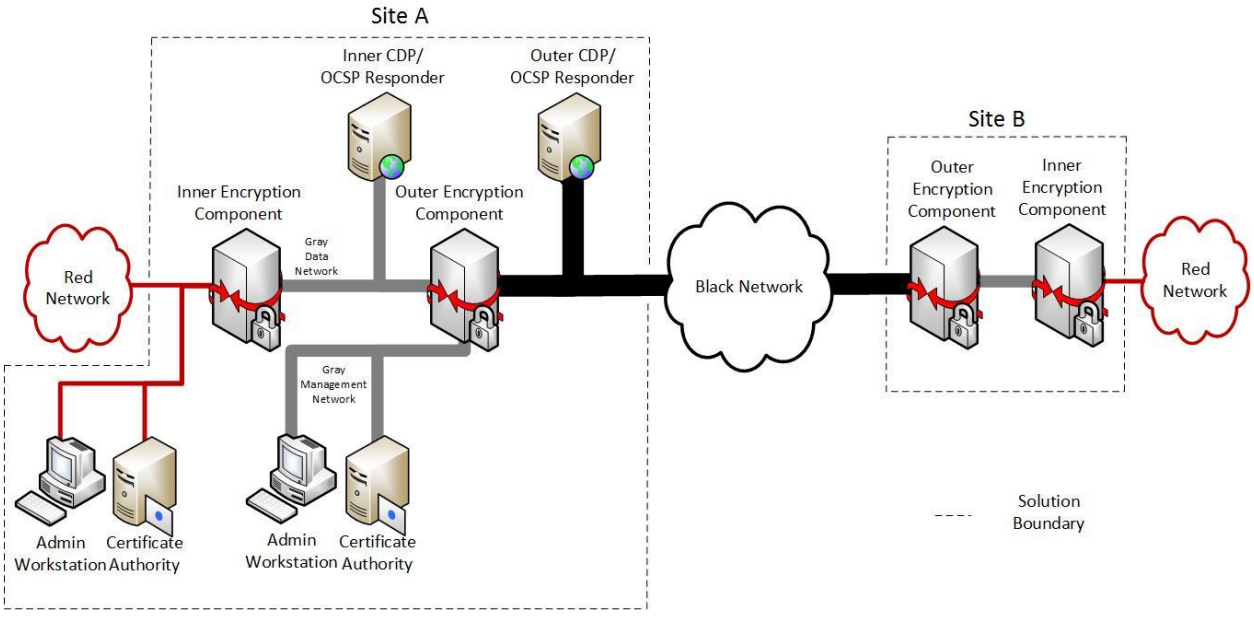


Figure 11. MSC Solution using CRL Distribution Points and OCSP Responders

Figure 11 illustrates the placement of CDPs and OCSP Responders to make CRLs accessible to Remote Sites on the Black network before IPsec tunnel establishment. During negotiation of the outer encryption tunnel, the Outer VPN Gateways contact the Outer CDP or OCSP Responder on the Black network to download the latest CRL produced by the Outer CA. Similarly, during negotiation of the inner

encryption tunnel, the Inner VPN Gateways contact the Inner CDP or OCSP Responder on the Gray network to download the latest CRL produced by the Red CA.

The use of CDPs and OCSP Responders on the external side of the VPN Gateways requires that the contents of the CRLs hosted on them are unclassified, since the CDPs and OCSP Responders are located on networks that the design of the MSC Solution treats as unclassified. If a solution owner's classification authority decides that its CRLs are classified, then its MSC Solutions would be unable to make use of external CDPs and OCSP Responders.

For solutions that support networks of different security levels (see Section 4.3.2.2), a single Inner CDP or OCSP Responder may be used to host the CRLs for the Inner VPN Gateways of multiple Red networks.

A solution owner may choose to implement zero, one, or multiple CDPs or OCSP Responders on Black and Gray networks, based on their expected utility in facilitating CRL distribution to Remote Sites. Having multiple redundant CDPs or OCSP Responders on the same network improves the availability of CRL distribution, since a VPN Gateway only needs to be able to contact one CDP or OCSP Responder to obtain the CRL. Conversely, in a small-scale solution, manual out-of-band distribution of CRLs may be more cost-effective than deploying and maintaining CDPs or OCSP Responders.

7.2 Connectivity Association Keys

This section provides details on issuing, renewing, rekeying and revoking CAKs, the PSKs used by MACsec Devices for authentication. CAK management is a critical function. If a CAK is compromised, all MACsec Devices using that CAK need to be updated with a new CAK. This is different from a certificate-based solution in that revocation of any given certificate only impacts the device associated with that certificate. CAK management can be provided by Enterprise services (e.g., Key Management Infrastructure (KMI)) or via locally-operated services.

Enterprise CAK management services are those services that can be provided on a Department-wide or Agency-wide scale, and include services such as CAK generation; CAK distribution; CAK accounting and CAK compromise reporting. A department-wide program that can provide these types of services is the NSA KMI. The KMI also allows agencies to locally manage CAKs using a specialized Management Client (MGC).

Currently, the MGC does not have a commercial standard interface. However, customers can acquire or develop a custom application that can receive key from KMI and reformat it for secure distribution to and installation on MACsec Devices.

If an Enterprise CAK is not feasible, customers can deploy their own locally-operated, NSA-approved KGC to generate and manage CAKs for MACsec Devices.

The customer is responsible for developing a Key Management Plan (KMP) for the CAKs used in the MSC Solution. The KMP must be approved by the National Cryptographic Solutions Management Office (NCSMO). The KMP addresses the following with regard to the CAKs:

- CAK technical specifications, including the format for the CAK to be installed onto the MACsec Devices;
- Operational environment for the MSC Solution;
- MSC Solution Components that use and manage CAKs, including planned volumes;
- Processes for CAK ordering, generation, distribution and storage, including physical and technical controls;
- Accounting processes and controls for CAKs;
- Compromise management and recover (how service is restored in the event of CAK compromise);
- CAK archival and recovery, if necessary (how CAKs are recovered in case of loss);
- Interoperability requirements with other systems and solutions; and
- Roles and responsibilities for managing CAKs.

In addition to a KMP, customers may need to obtain an IAD MD-110 waiver since commercial equipment typically allow CAKs to exist in red form during part of the CAK life-cycle management process.

Customers are strongly encouraged to contact their Client Advocate early in the process to obtain the NCSMO's assistance in developing the KMP, selecting a NSA-approved KGC, or obtaining an IAD MD-110 waiver.

7.2.1 Connectivity Association Key Issuance, Renewal and Rekey

Each MACsec Device has at least one CAK, which is used by the MACsec Device to authenticate with another MACsec Device. A different CAK is required for every pair of MACsec Devices establishing an encryption tunnel. Also, if both layers of the solution use MACsec, different CAKs are required for the inner and outer encryption tunnels.

CAKs must be securely distributed to the MACsec Devices. Secure distribution can be achieved via technical means (e.g., encryption) or procedural controls. A pre-placed PSK Encryption Key (PEK) can be used to encrypt CAKs. A separate PEK should be used for each site.

CAK Administrators (CAKA) install the CAKs into the MACsec Devices. CAKAs account for CAKs to ensure their location and use is known at all times. In the case of compromise, CAKAs need to be able to

determine where all instances of a given CAK exist and update that CAK in accordance with compromise reporting and recovery procedures.

CAKs require periodic updating to limit the amount of operational exposure for the CAKs. CAKs must be updated at least every 30 days. Likewise, PEKs must be updated at least every 90 days.

7.2.2 Connectivity Association Key Compromise Recovery

CAK compromise recovery is a critical function within the MSC Solution. Therefore, good accounting records are necessary to know which CAKs are installed on which MACsec devices. When a CAK is compromised, the CAK must be updated immediately at both sites. CAK compromise recovery procedures are documented in the KMP.

8 THREATS

This section details how the required components work together to provide overall security in the solution. Figure 2 through Figure 9 show the boundary of the MSC Solution for each high-level design covered by this CP.

An assessment of security was conducted on each of the high-level designs described in this CP while making no assumptions regarding use of specific products for any of the defined components. There are several different threats to consider when evaluating the risk of transporting data over secure or unsecure networks. By examining these threats, the organization can have a better understanding of the risks they are accepting by implementing the solution and how these risks affect the Confidentiality, Integrity, and Availability of the network, systems, and data. To obtain the classified risk assessment associated with this CP, please contact the NSA via your Client Advocate.

8.1 Passive Threats

This threat refers to internal or external actors attempting to gain information from the network without changing the state of the system. Threat actions include collecting or monitoring traffic (e.g., traffic analysis or sniffing the network) passing through a network to gain useful information through data analysis.

The security against a passive attack targeting the data in transit across the Black network is provided by the layered encryption tunnels using IPsec or MACsec. To mitigate passive attacks, two layers of CNSA Suite encryption, Advanced Encryption Standard (AES), are employed to provide confidentiality for the solution. Use of AES is approved to protect classified information, meeting IAD and CNSSP-15 guidance for adequate confidentiality. The Inner and Outer Encryption Components that are used to set up the two tunnels must be independent in a number of ways (see Section 10). Due to this independence, the adversary should not be able to exploit a single cryptographic implementation to compromise both tunnels.

The use of one or more Outer CDPs or OCSP Responders to distribute unencrypted CRLs on a Black network potentially allows a passive threat actor with access to the Black network path between an Outer CDP/OCSP Responder and Outer VPN Gateway to obtain a copy of the CRL issued by an Outer CA. However, the content of the CRL is primarily limited to a list of serial numbers of revoked certificates, the date and time when each certificate was revoked, and a high-level reason why each certificate was revoked (such as key compromise or cessation of operation). The CRL does not specify what certificates are still valid, nor does it identify the physical or network locations of any components in the solution. The CRL also does not reveal any information about certificates issued by anything other than the Outer CA. If a solution owner's AO considers the limited information in a CRL too sensitive to distribute on the Black network, the solution owner can choose not to implement Outer CDPs/OCSP Responders and rely on other means to distribute CRLs to Outer VPN Gateways in a timely fashion.

8.2 External (Active) Threats

This threat refers to outsiders gaining unauthorized access to a system or network, exfiltration of sensitive Red network data, or degradation of availability of the system or network. Threat actions include introducing viruses, malware, or worms with the intention to compromise the network or exfiltrate data, or to analyze the design of the network or system for future attacks. Adversaries could gain access to an Encryption Component, and then exploit or compromise other devices on the network. DoS or Distributed DoS (DDoS) attacks compromise availability of the system, degrading/disrupting secure communication across a Black network. Further external threat actions would include social engineering attacks to assist attackers with gaining additional access to a network for the purpose of compromising a system or network, traffic injection or modification attacks, or replay attacks.

8.2.1 Rogue Traffic

One method for detecting rogue traffic from an external attack as it attempts to pass through one or both Encryption Components is by having the port filtering native to each Encryption Component enabled and configured to audit and log any traffic that is not of the format described in the configuration. This will allow the Auditors and/or the Security Administrators to detect whether the Outer Encryption Component has been breached, thus providing an early warning of a potential intrusion. It will also provide detection of misconfigured Outer Encryption Components.

Another method for detecting a potential intrusion into the solution is requiring automated configuration change detection on Red and Gray Management networks to ensure Encryption Component configurations are not changed without the knowledge of Auditors and Security Administrators. Auditors also ensure through the audit logs that all configuration changes are valid. This will counter attacks that take advantage of Encryption Component misconfigurations.

CDPs and OCSP Responders are protected from rogue traffic by implementing port filtering on the server. Rogue traffic to CDPs and OCSP Responders can be further mitigated by implementing a firewall or other packet filtering device between the CDP/OCSP Responder and the rest of the network.

8.2.2 MALWARE AND UNTRUSTED UPDATES

Administration Workstations and CAs for Inner Encryption Components shall be distinct from the Administration Workstations and CAs for a Gray network. This separation minimizes the potential for malware on a single device to impact components supporting both Inner and Outer Encryption tunnels.

Each individual component of this solution has the capability to perform trusted updates through verification of a signature or hash to ensure that the update is from a reliable source, such as signed by the vendor. This mitigates threats of malicious users trying to push updates or code patches that affect the security of the component (and therefore system). The source of all updates and patches should be verified before installation occurs.

8.2.3 DENIAL OF SERVICE

DoS attack risks cannot be completely mitigated. MSC Solutions in compliance with this CP are required to drop all packets that are not Internet Key Exchange (IKE), Encapsulating Security Payload (ESP), MACsec Key Agreement (MKA), Extensible Authentication Protocol-Transport Layer Security (EAP-TLS) or other approved protocols on the appropriate interfaces, which significantly reduces the potential of flooding attacks. For customers requiring more protection against these attacks, one option is the use of a perimeter router between the Outer Encrypting Component and Black networks to filter traffic before it reaches the Outer Encryption Component. Another option for customers requiring more protection is to add additional filtering based on specifics like known network IP addresses to filter traffic from devices not included in this solution or leasing private lines for the Black network. Other mitigations are acceptable and up to the AO to approve their use.

A single Encryption Component failure is likely to result in a DoS condition. One assumption underlying this solution is that high assurance of availability is not required. If availability is critical for the customer, protection against DoS attacks can be achieved through network redundancy and instituting DoS response procedures when loss of availability is detected.

When using Outer CDPs or OCSP Responders to distribute CRLs to Outer VPN Gateways on a Black network, a sustained DoS attack on the CDPs/OCSP Responders could prevent Outer VPN Gateways from receiving updated CRLs. If the CRLs cached at the other VPN Gateways then expire, they would be unable to establish IPsec tunnels due to the inability to check the revocation status of certificates during the mutual authentication process. Deploying multiple Outer CDPs or OCSP Responders reduces the likelihood of a successful DoS attack on this part of the solution, since as long as even one Outer CDP/OCSP Responder is available, Outer VPN Gateways will be able to retrieve CRL updates.

Additionally, a solution using CDPs or OCSP Responders should still have procedures in place for out-of-band CRL distribution to use in the event that all Outer CDPs/OCSP Responders become unavailable.

8.2.4 SOCIAL ENGINEERING

It is the responsibility of the customer to define the appropriate policies and training necessary to protect against Social Engineering attacks. In addition, these types of attacks generally take advantage of other attacks detailed in this section and are already discussed.

8.3 INSIDER THREATS

This threat refers to an authorized or cleared person or group of people with physical or logical access to the network or system who may act maliciously or negligently, resulting in risk exposure for the organization. This threat could include poorly trained employees, curious employees, disgruntled employees, escorted personnel who gain access to the equipment, dishonest employees, or those that have the means and desire to gain escalated privileges on the network.

Threat actions include insertion or omission of data entries that result in a loss of data integrity, unintentional access to an unauthorized system or network, unwillingly or unknowingly executing a virus or malware, intentionally exposing the network and systems to viruses or malware, cross-contaminating a system or network with data from a higher classification to a lower classification (e.g., Secret data to an Unclassified network or system), or malicious or unintentional exfiltration of classified data. Typically, the threat from insiders has the potential to cause the greatest harm to an organization, and insider attacks are also the hardest to monitor and track.

To mitigate insider threats, separation of roles within the solution is required (see Section 130). In addition, logging and auditing of security critical functionality (see Section 11.10) is required. Also, strong authentication of the Security Administrator and Auditor are required for access to ensure accountability of these individuals. Finally, outbound filters on Encryption Components and firewalls are configured to block traffic leaving the internal network that does not go through the encryption tunnels. An IDS is also recommend on, at a minimum, the Black, Gray, or Red network to help identify unusual or suspicious traffic that could result from a failure, misconfiguration, or attack on Inner or Outer Encryption Components.

8.4 SUPPLY CHAIN THREATS

A critical aspect of the U.S. Government's effectiveness is the dependability, trustworthiness, and availability of the Information and Communication Technology (ICT) components embedded in the systems and networks upon which the ability to perform U.S. Government missions rely. The supply chains for those ICT components are the underpinnings of those systems and networks, and supply chain attacks are attempts to proactively compromise those underpinnings.

Unfortunately, the supplier cannot always provide guarantees of a safe delivery of a component; they are only able to provide assurances based on their reliance on established procedures and processes they have developed. In a single change of hands, the component may be introduced to potential threats and compromises on many levels.

The supply chain threat refers to an adversary gaining access to a vendor, retailer, reseller, or shipper and then attempting to insert or install a modification or a counterfeit piece of hardware into a component that is destined for a U.S. Government customer in an effort to gain information or cause operational issues. This threat also includes the installation of malicious software on components of the solution. This threat is difficult to identify, and is increasingly more difficult to prevent or protect against since vendors build products containing components manufactured by subcontractors. It is often difficult to determine where different pieces of components are built and installed within the supply chain.

Threat actions include manufacturing faulty or counterfeit parts of components that can be used to disrupt system or network performance, leaving open back doors in hardware that allow attackers easy ways to attack and evade monitoring, as well as easy ways to steal data or tamper with the integrity of existing/new data. Supply Chain attacks may occur during development and production, updates, distribution, shipping, at a warehouse, in storage, during operations, or disposal. For this reason, it is imperative that all components selected for use in CSfC solutions are subject to the applicable Supply Chain Risk Management (SCRM) process to reduce the risk of acquiring compromised components.

Each component that is selected from the CSfC Components List shall go through a Product Supply Chain Threat Assessment to determine the appropriate mitigations for the intended application of the component per the organization's AO-approved Product Supply Chain Threat Assessment process (see CNSSD 505 Supply Chain Risk Management (SCRM) for additional guidance).

Doctrinal requirements are placed on Product Selection, Implementers, and System Integrators of these solutions to minimize the threat of supply chain attacks. To further mitigate Supply Chain Threats implementing organizations should use the following guidance:

- Establish an ICT SCRM program which conforms to applicable policy based on external and organizational requirements and constraints. The ICT SCRM program should be integrated into the organizational business and mission processes.
- Assess all aspects of the performance of potential vendors, not only the product quality, cost, and performance, but also supply chain risk factors of vendor selection. These risk factors include political ties to foreign governments, citizenship of employees, partner affiliations, employee clearance levels, and location of suppliers and sub-suppliers.
- Ensure that each component selected from the CSfC Components List go through a Product Supply Chain Threat Assessment to determine the appropriate mitigations for the intended

application of the component (see CNSSD 505 Supply Chain Risk Management and Intelligence Community Directive (ICD) 731 Supply Chain Risk Management).

- Conduct a Criticality Analysis by which mission-critical functions and components are identified and prioritized with respect to improving acquirer practices (see Defense Acquisition Guidebook, Chapter 13).

Supply chain risk management is a critical consideration in acquiring commercial products. Even after selecting components from the CSfC Components List and utilizing a rigorous acquisition process, an AO must perform due diligence when integrating commercial components for mission operations.

8.5 INTEGRATOR THREATS

This threat refers to an integrator who has unrestricted access to all components within the solution prior to the customer purchasing and implementing the solution within their system. This is different than a Supply Chain threat in that these integrators have access to all components to be used in the solution, rather than only those being procured from a particular vendor.

Threat actions could include installing or configuring components in a manner that places the organization at risk for attack or open to an unknown vulnerability that may not be detected through normal tests, scans, and security counter-measures.

To mitigate this threat, integrators are required to be cleared to the highest level of data protected by the MSC Solution. To further reduce the integrator threat, a customer may wish to use multiple integrators, such that no one integrator has access to all components of the solution. More information on the NSA's list of trusted integrators can be found on the NSA CSfC Website in the "Criteria for CSfC Integrators" section at this link: https://www.nsa.gov/resources/everyone/csfc.

9 REQUIREMENTS OVERVIEW

The following five sections (Sections 10 through 15) specify requirements for implementations of MSC Solutions compliant with this CP. However, not all requirements in the following sections will apply to each compliant solution.

9.1 THRESHOLD AND OBJECTIVE REQUIREMENTS

In some cases, multiple versions of a requirement may exist in this CP. Such alternative versions of a requirement are designated as being either a Threshold requirement or an Objective requirement.

- A Threshold (T) requirement specifies a feature or function that provides the minimal acceptable capability for the security of the solution.

- An Objective (O) requirement specifies a feature or function that provides the preferred capability for the security of the solution.

In general, when separate Threshold and Objective versions of a requirement exist, the Objective requirement provides a higher degree of security for the solution than the corresponding Threshold requirement. However, in these cases meeting the Objective requirement may not be feasible in some environments or may require components to implement features that are not yet widely available. Solution owners are encouraged to implement the Objective version of a requirement, but in cases where this is not feasible solution owners may implement the Threshold version of the requirement instead. These Threshold and Objective versions are mapped to each other in the "Alternatives" column. Objective requirements that have no related Threshold requirement are marked as "Optional" in the "Alternatives" column.

In most cases there is no distinction between the Threshold and Objective versions of a requirement. In these cases, the "Threshold / Objective" column indicates that the Threshold equals the Objective (T=O).

Requirements that are listed as Objective in this CP may become Threshold requirements in a future version of this CP. Solution owners are encouraged to implement Objective requirements where possible to facilitate compliance with future versions of this CP.

9.2 Requirements Designators

Each requirement defined in this CP has a unique identifier consisting of the prefix "MSC," a digraph that groups related requirements together (e.g., "KM"), and a sequence number (e.g., 11). Table 2 lists the digraphs used to group together related requirements and identifies the sections in which those requirement groups can be found.

Table 2. Requirement Digraphs

Digraph	Description	Section	Table
PS	Product Selection Requirements	Section 10	Table 3
SR	Overall Solution Requirements	Section 11.1	Table 4
VG	VPN Gateway Requirements	Section 11.2	Table 6
MD	MACsec Device Requirements	Section 11.3	Table 8
IR	Additional Requirements for Inner Encryption Components	Section 11.4	Table 9
OR	Additional Requirements for Outer Encryption Components	Section 11.5	Table 10
PF	Port Filtering Requirements for Solution Components	Section 11.6	Table 11
CM	Configuration Change Detection Requirements	Section 11.7	Table 12
DM	Device Management Requirements	Section 11.8	Table 13
MR	Continuous Monitoring Requirements	Section 11.9	Table 14
AU	Auditing Requirements	Section 11.10	Table 15

Digraph	Description	Section	Table
KM	Key Management Requirements	Section 11.11	Table 16 through Table 23
FW	Gray Firewall Requirements	Section 11.12	Table 24
GD	Requirements for the Use and Handling of Solutions	Section 12.1	Table 25
	Role-Based Personnel Requirements	Section 13	Table 27
RP	Incident Reporting Requirements	Section 12.2	Table 26
TR	Testing Requirements	Section 14.1	Table 28

10 REQUIREMENTS FOR SELECTING COMPONENTS

In this section, a series of requirements are given for maximizing the independence between the components within the solution. The requirements in Table 3 will increase the level of effort required to compromise this solution.

Table 3. Product Selection (PS) Requirements

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-PS-1	The products used for any VPN Gateways shall be chosen from the list of IPsec VPN Gateways on the CSfC Components List.	T=O	
MSC-PS-2	The products used for any MACsec Device shall be chosen from the list of MACsec Ethernet Encryptors on the CSfC Components List.	T=O	
MSC-PS-3	The products used for any Firewalls shall be chosen from the list of Stateful Traffic Filtering Firewalls (TFFW) on the CSfC Components List.	T=O	
MSC-PS-4	The products used for any CAs shall either be chosen from the list of CAs on the CSfC Components List or the CAs shall be pre-existing Enterprise CAs (e.g., DoD PKI, IC PKI).	T=O	
MSC-PS-5	Intrusion Prevention Systems (IPS) shall be chosen from the list of IPS on the CSfC Components List.	O	Optional
MSC-PS-6	The Inner Encryption Component and the Outer Encryption Component shall either: come from different manufacturers, where neither manufacturer is a subsidiary of the other; or be different products from the same manufacturer, where NSA has determined that the products meet the CSfC criteria for implementation independence. Differences between Service Packs (SP) and version numbers for a particular vendor's operating system (OS) do not provide adequate diversity.	T=O	

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-PS-7	The cryptographic libraries used by the Inner Encryption Component and Outer Encryption Component shall either: come from different manufacturers, where neither manufacturer is a subsidiary of the other; or be different libraries from the same manufacturer, where NSA has determined that the libraries meet the CSfC criteria for implementation independence.	O	Optional
MSC-PS-8	If the solution contains an Inner CA and an Outer CA, the cryptographic libraries shall either: come from different manufacturers, where neither manufacturer is a subsidiary of the other; or be different libraries from the same manufacturer, where NSA has determined that the libraries meet the CSfC criteria for implementation independence.	O	Optional
MSC-PS-9	Gray Network Firewalls and Inner Encryption Components shall either: come from different manufacturers, where neither manufacturer is a subsidiary of the other; or be two different products from the same manufacturer, where NSA has determined that the two products meet the CSfC criteria for implementation independence.	T=O	
MSC-PS-10	The Inner Encryption Component and Outer Encryption Component shall use physically separate components, such that no component is used for more than one function.	T=O	
MSC-PS-11	The Inner Encryption Component and Gray Network Firewall shall use physically separate components, such that no component is used for more than one function.	T=O	
MSC-PS-12	If the solution contains an Inner CA and an Outer CA, the CAs shall either: come from different manufacturers, where neither manufacturer is a subsidiary of the other; or be different products from the same manufacturer, where NSA has determined that the products meet the CSfC criteria for implementation independence; or use an Enterprise PKI approved by the AO.	T=O	

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-PS-13	Each component that is selected out of the CSfC Components List shall go through a Product Supply Chain Threat Assessment to determine the appropriate mitigations for the intended application of the component per the organization's AO-approved Product Supply Chain Threat Assessment process. (See CNSSD 505 SCRM for additional guidance.)	T=O	
MSC-PS-14	Components shall be configured to use the NIAP-certified evaluated configuration.	T=O	

11 CONFIGURATION REQUIREMENTS

Once the products for the solution are selected, the next step is setting up the components and configuring them in a secure manner. This section consists of generic guidance on how to configure the components of the MSC Solution.

CPs provide architecture and configuration information that allows customers to select COTS products from CSfC Components Lists for their solution and then to properly configure those products to achieve a level of assurance sufficient for protecting classified data. CSfC Components Lists consist of eligible COTS products identified by model/version numbers that have met appropriate Protection Profile requirements.

This section contains requirements applicable to the MSC Solution components. In this section, a series of overarching architectural requirements are given for maximizing the independence between the components within the solution. This independence will increase the level of effort required to compromise this solution.

The products that are approved for use in this solution will be listed on the CSfC Components List on the IAD/CSfC website (https:// www.nsa.gov/resources/everyone/csfc). No single commercial product shall be used to protect classified information. The only approved methods for using COTS products to protect classified information in transit is through an approved CP.

Once the products for the solution are selected, each product shall go through a Product Supply Chain Threat Assessment to determine the appropriate mitigations for the intended application of the component per the organization's AO-approved Product Supply Chain Threat Assessment process. (See CNSSD 505 Supply Chain Risk Management (SCRM) for additional guidance.)

11.1 Overall Solution Requirements

Table 4 provides the overall solutions requirements for this CP.

Table 4. Overall Solution Requirements (SR)

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-SR-1	Network services provided by control plane protocols (such as DNS and NTP) shall be located on the inside network (i.e., Gray network for Outer Encryption Component and Red network for Inner Encryption Component).	T=O	
MSC-SR-2	Sites that need to communicate shall ensure that each tunnel's Encryption Components selected by each site are interoperable.	T=O	
MSC-SR-3	The time of day on the Inner Encryption Components and Red Management Services shall be synchronized to a time source located in the Enterprise/Red network.	T=O	
MSC-SR-4	The time of day on the Outer Encryption Components, Gray Network firewall, and Gray Management Services shall be synchronized to a time source located in the Gray Management network.	T=O	
MSC-SR-5	Default accounts, passwords, community strings, and other default access control mechanisms for all Solution Components shall be changed or removed.	T=O	
MSC-SR-6	All components shall be properly configured in accordance with local policy and applicable U.S. Government guidance. In the event of conflict between the requirements in this CP and local policy, this CP takes precedence.	T=O	
MSC-SR-7	All physical paths within a Gray network between two Inner Encryption Components for Red networks of different classification levels shall include an Outer Encryption Component or a Gray Network Firewall.	T=O	
MSC-SR-8	All physical paths within a Gray network between a CA, an Administration Workstation, or a CDP/OCSP Responder and an Inner Encryption Component for a Red network whose classification level is lower than the highest classification of data protected by the solution shall include an Outer Encryption Component or a Gray Network Firewall.	T=O	
MSC-SR-9	The only approved physical paths leaving the Red network shall be through a MSC Solution in accordance with this CP or via an AO-approved solution for protecting data in transit.	T=O	

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-SR-10	Solution Components shall receive virus signature updates as required by the local agency policy and the AO.	T=O	
MSC-SR-11	When multiple Inner Encryption Components share an Outer Encryption Component, they shall be placed in parallel.	T=O	
MSC-SR-12	Inner Encryption Components shall not perform switching or routing for other Encryption Components	T=O	

11.2 VPN Gateway Requirements

This section addresses requirements for VPN Gateways. Table 5 identifies the algorithms approved for IPsec encryption. Table 6 provides requirements for VPN Gateways.

Table 5. IPsec Encryption (Approved Algorithms for Classified)

Security Service	Algorithm Suite	Specifications
Confidentiality (Encryption)	AES-256	FIPS PUB 197 IETF RFC 6239 IETF RFC 6379 IETF RFC 6380 IETF RFC 6460
Authentication (Digital Signature) (Threshold – Unclassified Only)	RSA 2048	FIPS PUB 186-4
Authentication (Digital Signature) (Objective) (Threshold – All Classified NSS)	RSA 3072 or ECDSA over the curve P-384 with SHA-384	FIPS PUB 186-4 FIPS PUB 186-4 IETF RFC 6239 IETF RFC 6380 IETF RFC 6460
Key Exchange/ Establishment	ECDH over the curve P-384 (DH Group 20) or RSA 3072 or DH 3072	NIST SP 800-56A IETF RFC 6239 IETF RFC 6379 IETF RFC 6380 IETF RFC 6460 NIST SP 800-56A
Integrity (Hashing)	SHA-384	FIPS PUB 180-4 IETF RFC 6239 IETF RFC 6379 IETF RFC 6380 IETF RFC 6460

Table 6. VPN Gateway (VG) Requirements

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-VG-1	The proposals offered by VPN Gateways in the course of establishing the IKE Security Association (SA) and the ESP SA for Inner and Outer Tunnels shall be configured to offer algorithm suite(s) containing only CNSA Suite algorithms (see Table 5).	T=O	
MSC-VG-2	Default, self-signed or proprietary device certificates, which are frequently pre-installed by the vendor, for any VPN Gateway shall not be used for establishing SAs.	T	MSC-VG-3
MSC-VG-3	Default, self-signed or proprietary device certificates, which are frequently preinstalled by the vendor, for any VPN Gateways shall be removed.	O	MSC-VG-2
MSC-VG-4	A unique device certificate shall be loaded onto each VPN Gateway along with the corresponding CA (signing) certificate.	T=O	
MSC-VG-5	The private key stored on VPN Gateways shall not be accessible through an interface.	T=O	
MSC-VG-6	A device certificate shall be used for VPN Gateway authentication during IKE.	T=O	
MSC-VG-7	VPN Gateway authentication shall include a check that the certificate is authorized, which can include a Certificate Revocation List (CRL), OCSP Responder, or whitelist.	T=O	
MSC-VG-8	The VPN Gateway authentication shall include a check that the certificate is not expired.	T=O	
MSC-VG-9	Both IPsec tunnels shall use IKEv2 (IETF RFC 5996) key exchange.	T=O	
MSC-VG-10	All VPN Gateways shall use Cipher Block Chaining for IKE encryption.	T=O	
MSC-VG-11	All VPN Gateways shall use Cipher Block Chaining for ESP encryption with an HMAC for integrity.	T	MSC-VG-12
MSC-VG-12	All VPN Gateways shall use Galois Counter Mode for ESP encryption.	O	MSC-VG-11
MSC-VG-13	All VPN Gateways shall set the IKE SA lifetime to at most 24 hours.	T=O	
MSC-VG-14	All VPN Gateways shall set the ESP SA lifetime to at most 8 hours.	T=O	

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-VG-15	Inner VPN Gateways shall only authenticate and establish an IPsec tunnel with one another if their Red networks operate at the same security level (as defined in this CP).	T=O	
MSC-VG-16	All VPN Gateways shall re-authenticate the identity of the VPN Gateway at the other end of the established tunnel before rekeying the IKE SA.	T=O	

11.3 MACSEC DEVICE REQUIREMENTS

This section addresses requirements for MACsec Devices. Table **7** identifies the algorithms approved for MACsec encryption. Table 8 provides requirements for MACsec Devices.

Table 7. MACsec Encryption (Approved Algorithms for Classified)

Security Service	Algorithm Suite	Specifications
Confidentiality (Encryption)	AES-GCM-256 AES-GCM-XPN-256	FIPS PUB 197 IEEE 802.1AEbw IETF RFC 6239 IETF RFC 6379 IETF RFC 6380 IETF RFC 6460
Key Wrap	AES-CMAC-256	FIPS PUB 197 NIST SP 800-38F

Table 8. MACsec Device (MD) Requirements

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-MD-1	MACsec Devices shall use AES Key Wrap in CMAC mode for key distribution and AES GCM for MACsec; both shall use cryptographic key sizes of 256 bits.	T=O	
MSC-MD-2	MACsec Devices shall authenticate using Pre-Shared Keys (PSKs), known as Connectivity Association Keys (CAKs).	T=O	
MSC-MD-3	CAKs shall be AES 256 bits and generated on a KGC.	T=O	
MSC-MD-4	MACsec Devices shall have the length of the Connectivity Association Key Name (CKN) set to a minimum of 8 bytes (64 bits).	T=O	

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-MD-5	For each pair of MACsec Devices establishing an encryption tunnel, one of the two shall be configured to be the Key Server by setting this value to 0 (zero). The other MACsec Device shall have its Key Server value set to 1.	T=O	
MSC-MD-6	MACsec Devices shall enable data delay protection for MKA.	T=O	
MSC-MD-7	MACsec Devices shall have an MKA Lifetime Timeout limit set to 6.0 seconds and Hello Timeout limit set to 2.0 seconds.	T=O	
MSC-MD-8	MACsec Devices shall have the replay window set to 2 or as low as possible given the nature of the Black network being traversed.	T=O	
MSC-MD-9	MACsec Devices shall require all traffic on an external facing port to be encrypted (e.g., must-secure).	T=O	
MSC-MD-10	MACsec Device configuration files, whether printed or electronically copied, shall be protected to the highest classification of the MACsec Device's CAK.	T=O	
MSC-MD-11	MACsec Devices shall have the Confidentiality Offset set to 0 (zero).	T=O	
MSC-MD-12	If a standalone device is required to provide encapsulation of MACsec traffic between an Inner MACsec Device and an Outer Encryption Component, the standalone device shall be considered a Solution Component when satisfying requirements in Section 11.1.	T=O	

11.4 Additional Requirements for Inner Encryption Components

Additional requirements for Inner Encryption Components are identified in Table 9.

Table 9. Additional Requirements for Inner Encryption Components (IR)

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-IR-1	Inner VPN Gateways shall use Tunnel mode IPsec or Transport mode IPsec with an associated IP tunneling protocol (e.g., Transport Mode IPsec with GRE).	T=O	
MSC-IR-2	Sizes for packets or frames leaving the external interface of Inner Encryption Components shall be configured to reduce fragmentation and impact performance. This requires proper configuration of the Maximum Transmission Unit (MTU) (for IPv4 or MACsec) or Path MTU (PMTU) (for IPv6) and should consider Black network and Outer Encryption Component MTU/PMTU values to achieve this.	O	Optional
MSC-IR-3	Inner Encryption Components shall not allow any packets received on an interface connected to a Red network to bypass encryption and be forwarded out through an interface connected to a Gray network.	T=O	
MSC-IR-4	Inner Encryption Components shall not allow any packets received on an interface connected to a Gray network to bypass decryption and be forwarded out through an interface connected to a Red network.	T=O	

11.5 Additional Requirements for Outer Encryption Components

Additional requirements for Outer Encryption Components are identified Table 10.

Table 10. Additional Requirements for Outer Encryption Components (OR)

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-OR-1	Outer VPN Gateways shall use Tunnel mode IPsec.	T=O	
MSC-OR-2	Outer Encryption Components shall not allow any packets received on an interface connected to a Gray network to bypass encryption and be forwarded out through an interface connected to a Black network.	T=O	
MSC-OR-3	All traffic received by the Outer Encryption Components on an interface connected to a Gray network, with the exception of Control Plane traffic, shall have already been encrypted once.	T=O	

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-OR-4	Outer Encryption Components shall not allow any packets received on an interface connected to a Black network to bypass decryption.	T=O	
MSC-OR-5	The Outer Encryption Components shall not permit split-tunneling.	T=O	
MSC-OR-6	Outer Encryption Components shall not perform routing.	T=O	

11.6 PORT FILTERING REQUIREMENTS FOR SOLUTION COMPONENTS

Requirements for port filtering for Solution Components are identified in Table 11.

Table 11. Port Filtering (PF) Requirements for Solution Components

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-PF-1	All Solution Components shall have all network interfaces restricted to the smallest address ranges, ports, and protocols possible.	T=O	
MSC-PF-2	All Solution Components shall have all unused network interfaces disabled.	T=O	
MSC-PF-3	For all Outer VPN Gateway interface connected to a Black network, traffic filtering rules shall be applied to both inbound and outbound traffic, such that only IKE, ESP, and control plane protocols (as defined in this CP) approved by organization-defined policy are allowed.	T=O	
MSC-PF-4	For all Outer MACsec Device interface connected to a Black network, traffic filtering rules shall be applied to both inbound and outbound traffic, such that only EAP-TLS and control plane protocols (as defined in this CP) approved by organization-defined policy are allowed.	T=O	
MSC-PF-5	For all interfaces connected to a Gray network, traffic filtering rules shall be applied to both inbound and outbound traffic, such that only IKE, IPsec, MKA, MACsec, and control plane protocols (as defined in this CP) approved by organization-defined policy are allowed. All packets not explicitly allowed shall be blocked.	T=O	
MSC-PF-6	Any service or feature that allows an Outer Encryption Component to contact a third party server (such as one maintained by the manufacturer) shall be blocked.	T	MSC-PF-7

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-PF-7	Any service or feature that allows an Outer Encryption Component to contact a third party server (such as one maintained by the manufacturer) shall be disabled.	O	MSC-PF-6
MSC-PF-8	Each Encryption Component shall only accept management traffic on the physical ports connected to its management network.	T=O	
MSC-PF-9	Multicast messages received on external interfaces of Outer Encryption Components shall be dropped.	T=O	
MSC-PF-10	For solutions using IPv4, an Outer VPN Gateways using IPsec shall drop all packets that use IP options.	T=O	
MSC-PF-11	For solutions using IPv4, each VPN Gateway shall only accept packets with Transmission Control Protocol (TCP), User Data Protocol (UDP), Encapsulating Security Payload (ESP), or ICMP in the IPv4 Protocol field and drop all other packets.	T=O	
MSC-PF-12	For solutions using IPv6, each VPN Gateway shall only accept packets with ESP, TCP, UDP, or ICMPv6 in the IPv6 Next Header field and drop all other packets.	T=O	
MSC-PF-13	The Gray network interfaces of Outer Encryption Components shall allow IKE and IPsec, or MACsec traffic, as appropriate, that is between two Inner Encryption Components protecting networks of the same classification level or that is being used for management of the Gray network.	T=O	
MSC-PF-14	The Gray network interfaces of Outer VPN Gateways shall allow HTTP traffic between Inner VPN Gateways and Inner CDPs/OCSP Responders.	T	MSC-PF-15 and MSC-PF-16
MSC-PF-15	The Gray network interfaces of Outer VPN Gateways shall allow HTTP GET requests from Inner VPN Gateways to Inner CDPs for the URL of the CRL needed by the Inner VPN Gateway, and block all other HTTP requests.	O	MSC-PF-14
MSC-PF-16	The Gray network interfaces of Outer VPN Gateways shall allow HTTP responses from Inner CDPs/OCSP Responders to Inner VPN Gateways that contain a well-formed CRL per IETF RFC 5280, and block all other HTTP responses.	O	MSC-PF-14
MSC-PF-17	The Gray network interfaces of Outer Encryption Components shall only permit packets whose source and destination IP addresses match the external interfaces of Inner Encryption Components that support Red networks of the same classification level.	T=O	

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-PF-18	The Gray network interfaces of Outer Encryption Components shall block all packets whose source address does not match a list of addresses or address ranges known to be reachable from the interface on which the packet was received.	T=O	
MSC-PF-19	The Gray network interfaces of Outer Encryption Components shall allow management and control plane protocols (as defined in this CP) that have been approved by policy.	T=O	
MSC-PF-20	The Gray network interfaces of Outer Encryption Components shall deny all traffic that is not explicitly allowed by requirements MSC-PF-8, MSC-PF- 13, MSC-PF-14, MSC-PF-15, or MSC-PF-19.	T=O	
MSC-PF-21	CDPs/OCSP Responders shall only allow inbound HTTP traffic.	T=O	

11.7 Configuration Change Detection Requirements

Table 12 defines the requirements for Configuration Change Detection.

Table 12. Configuration Change Detection (CM) Requirements

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-CM-1	A baseline configuration for all components shall be maintained by the Security Administrator and be available to the Auditor.	T=O	
MSC-CM-2	An automated process shall ensure that configuration changes are logged.	T=O	
MSC-CM-3	Log messages generated for configuration changes shall include the specific changes made to the configuration.	T=O	
MSC-CM-4	All Solution Components shall be configured with a monitoring service that detects all changes to configuration.	T=O	

11.8 Device Management Requirements

Table 13 defines the requirements for Device Management.

Table 13. Device Management (DM) Requirements

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-DM-1	Administration Workstations shall be dedicated for the purposes given in this CP and shall be physically separated from workstations used to manage non-CSfC solutions.	T=O	
MSC-DM-2	Administration Workstations shall physically reside within a protected facility where CSfC solution(s) are managed.	T=O	
MSC-DM-3	Administration Workstations shall connect from an internal port. Specifically, the Inner Encryption Component shall be managed from the Red network and the Outer Encryption Component shall be managed from the Gray network.	T=O	
MSC-DM-4	A separate LAN or VLAN on the Enterprise/Red network shall be used exclusively for all management of Inner Encryption Components and Solution Components within the Red network.	T=O	
MSC-DM-5	A separate LAN or VLAN on the Gray network shall be used exclusively for all management of Outer Encryption Components and Solution Components within the Gray network.	T=O	
MSC-DM-6	The Gray Management network shall not be directly connected to Non-secure Internet Protocol Router Network (NIPRNet) or any other Unclassified network not dedicated to the administration of CSfC solutions.	T=O	
MSC-DM-7	All components shall be configured to restrict the IP address range for the network administration device to the smallest range possible. Note that locally managing Solution Components is also acceptable.	T=O	
MSC-DM-8	All administration of Solution Components shall be performed from an Administration Workstation remotely using one of SSHv2, IPsec, MACsec, or TLS 1.2 or later version, or by managing the Solution Components locally.	T=O	
MSC-DM-9	Security Administrators shall authenticate to Solution Components before performing administrative functions.	T	MSC-DM-10

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-DM-10	Security Administrators shall authenticate to Solution Components with CNSA Suite compliant certificates before performing administrative functions remotely.	O	MSC-DM-9
MSC-DM-11	Security Administrators shall initiate certificate signing requests for Solution Components as part of their initial keying within the solution.	T=O	
MSC-DM-12	Administration Workstations that interact with the Certificate Authority for the Outer VPN Gateways must be located on the Gray network.	T=O	
MSC-DM-13	VPN Gateways shall obtain certificates through the use of PKCS #10 and #7 requests.	T	MSC-DM-14
MSC-DM-14	Devices shall use Enrollment over Secure Transport (EST) as detailed in IETF RFC 7030 for certificate management.	O	MSC-DM-13
MSC-DM-15	The same Administration Workstation shall not be used to manage Inner Encryption Components and Outer Encryption Components.	T=O	
MSC-DM-16	Outer Encryption Components and Solution Components within the Gray network shall forward log entries to a SIEM on the Gray Management network within 10 minutes.	T=O	
MSC-DM-17	Inner Encryption Components and Solution Components within the Red network shall forward log entries to a SIEM on the Red Management network within 10 minutes.	T=O	
MSC-DM-18	All logs forwarded to a SIEM on the Gray Management network shall be encrypted using SSHv2, IPsec, MACsec, or TLS 1.2 or later.	O	Optional
MSC-DM-19	All logs forwarded to a SIEM on a Red Management network shall be encrypted using SSHv2, IPsec, MACsec, or TLS 1.2 or later.	O	Optional
MSC-DM-20	Outer Encryption Components shall only be managed by Security Administrators cleared to at least the highest level of classification of each Red network supported by the Outer Encryption Component at the physical site the Outer Encryption Component is located.	T=O	
MSC-DM-21	When managing Solution components over the Black network, the management traffic shall be encrypted with algorithms IAW Table 5 and Table 7.	T=O	

11.9 CONTINUOUS MONITORING REQUIREMENTS

Continuous monitoring requirements are identified in Table 14.

Table 14. Requirements for Continuous Monitoring (MR)

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-MR-1	Traffic from the Black, Gray, or Red networks shall be monitored from an Intrusion Detection System (IDS).	T	MSC-MR-2
MSC-MR-2	Traffic from the Black, Gray, or Red networks shall be monitored from an Intrusion Protection System (IPS).	O	MSC-MR-1
MSC-MR-3	An IDS shall be deployed between the Outer Encryption Component and Gray Network Firewall (M1), or between the Gray Network Firewall and the Inner Encryption Component (M2), or on the internal side of the Inner Encryption Component (M3).	O	MSC-MR-4 MSC-MR-5 MSC-MR-6
MSC-MR-4	An IDS shall be deployed between the Outer Encryption Component and Gray Network Firewall (M1), and between the Gray Network Firewall and the Inner Encryption Component (M2), and on the internal side of the Inner Encryption Component (M3).	O	MSC-MR-3 MSC-MR-5 MSC-MR-6
MSC-MR-5	An IPS shall be deployed between the Outer Encryption Component and Gray Network Firewall (M1), or between the Gray Network Firewall and the Inner Encryption Component (M2), or on the internal side of the Inner Encryption Component (M3).	O	MSC-MR-3 MSC-MR-4 MSC-MR-6
MSC-MR-6	An IPS shall be deployed between the Outer Encryption Component and Gray Network Firewall (M1), and between the Gray Network Firewall and the Inner Encryption Component (M2), and on the internal side of the Inner Encryption Component (M3).	O	MSC-MR-3 MSC-MR-4 MSC-MR-5
MSC-MR-7	Each IDS in the solution shall be configured to send alerts to the Security Administrator.	O	MSC-MR-8
MSC-MR-8	Each IPS in the solution shall be configured to send alerts to the Security Administrator.	O	MSC-MR-7
MSC-MR-9	Each IDS in the solution shall be configured with rules that block and generate alerts upon detection of any unauthorized destination IP addresses.	O	MSC-MR-10
MSC-MR-10	Each IPS in the solution shall be configured with rules that block and generate alerts upon detection of any unauthorized destination IP addresses.	O	MSC-MR-9
MSC-MR-11	Each IDS in the solution shall be configured with rules that block and generate alerts upon detection of any unauthorized source IP addresses.	O	MSC-MR-12

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-MR-12	Each IPS in the solution shall be configured with rules that block and generate alerts upon detection of any unauthorized source IP addresses.	O	MSC-MR-11
MSC-MR-13	A SIEM component shall be placed within the Gray network unless devices are configured to push events to an Enterprise/Red network SIEM through an AO-approved one-way tap.	T=O	
MSC-MR-14	A SIEM shall be configured to send alerts to the Auditor when anomalous behavior is detected (i.e., blocked packets from the Outer Encryption Component or Gray Network Firewall).	T=O	
MSC-MR-15	The Gray SIEM shall collect logs from the Outer Encryption Component, Gray Network Firewall, and any components located within the Gray Management Services.	T=O	
MSC-MR-16	Logs sent to the Gray SIEM shall be encrypted with SSHv2, IPsec, MACsec, or TLS 1.2 or later.	T=O	
MSC-MR-17	One-way taps deployed as part of the solution shall be approved for use by the AO.	T=O	
MSC-MR-18	One-way taps deployed as part of the solution shall only allow monitoring data to flow from Monitoring Point 2 (M2) and/or Monitoring Point 1 (M1) to an enclave at the Red level that is isolated from the Red/Enterprise network.	T=O	

11.10 AUDITING REQUIREMENTS

Auditing requirements for the MSC Solution are identified in Table 15.

Table 15. Auditing (AU) Requirements

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-AU-1	Encryption Components shall log establishment of an encryption tunnel.	T=O	
MSC-AU-2	Encryption Components shall log termination of an encryption tunnel.	T=O	
MSC-AU-3	Solution Components shall log all actions performed on the audit log (off-loading, deletion, etc.).	T=O	
MSC-AU-4	Solution Components shall log all actions involving identification and authentication.	T=O	

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-AU-5	Solution Components shall log attempts to perform an unauthorized action (read, write, execute, delete, etc.) on an object.	T=O	
MSC-AU-6	Solution Components shall log all actions performed by a user with super-user or administrator privileges.	T=O	
MSC-AU-7	Solution Components shall log escalation of user privileges.	T=O	
MSC-AU-8	Solution Components shall log generation, loading, and revocation of certificates.	T=O	
MSC-AU-9	Solution Components shall log changes to time.	T=O	
MSC-AU-10	Solution Components shall log when packets received on Gray network interfaces are dropped or blocked.	T=O	
MSC-AU-11	Solution Components shall log the results of built-in self-tests.	T=O	
MSC-AU-12	MACsec Devices shall log creation and updates of Secure Association Keys.	T=O	
MSC-AU-13	MACsec Devices shall log administrator lockout due to excessive authentication failures.	T=O	
MSC-AU-14	MACsec Devices shall log detected replay attempts.	T=O	
MSC-AU-15	Each log entry shall record the date and time of the event.	T=O	
MSC-AU-16	Each log entry shall include the identifier of the event.	T=O	
MSC-AU-17	Each log entry shall record the type of event.	T=O	
MSC-AU-18	Each log entry shall record the success or failure of the event to include failure code, when available.	T=O	
MSC-AU-19	Each log entry shall record the subject identity.	T=O	
MSC-AU-20	Each log entry shall record the source address for network-based events.	T=O	
MSC-AU-21	Each log entry shall record the user and, for role-based events, role identity, where applicable.	T=O	
MSC-AU-22	VPN Gateways shall log the failure to download the CRL from a CDP or OCSP Responder.	T=O	
MSC-AU-23	VPN Gateways shall log if the version of the CRL downloaded from a CDP or OCSP Responder is older than the current cached CRL.	T=O	
MSC-AU-24	VPN Gateways shall log if signature validation of the CRL downloaded from a CDP or OCSP Responder fails.	T=O	
MSC-AU-25	Auditors shall compare and analyze collected network flow data against the established baseline on at least a weekly basis.	T=O	

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-AU-26	Locally-run CAs shall comply with the audit and archival requirements defined in IETF RFC 3647 Sections 4.5.4 and 4.5.5, respectively.	T=O	
MSC-AU-27	Locally-run CAs shall comply with periodic audit and assessment requirements defined in IETF RFC 3647 Section 4.8.	T=O	
MSC-AU-28	Audits and assessments for CAs shall be performed by personnel who are knowledgeable in the CAs' operations, as well as the CAs' CP and CPS requirements and processes, respectively.	T=O	
MSC-AU-29	KGCs that delivery CAK management services for MSC Solutions are to comply with audit and assessment requirements defined by the customer's operational security doctrine and enterprise KGC (if applicable).	T=O	
MSC-AU-30	Audits and assessments are to be performed by personnel who are knowledgeable in the KGCs' operations, as well as the KGCs' audit requirements and processes, respectively.	T=O	

11.11 Key Management Requirements

This section details key management requirements for the MSC Solution. General requirements are identified, followed by requirements specific to certificates and CAKs.

11.11.1 General Requirements

General key management requirements are identified in Table 16.

Table 16. General Key Management (KM) Requirements

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-KM-1	Certificate and CAK management services for the Inner tunnel shall be provided through the Red network.	T=O	
MSC-KM-2	Certificate and CAK management services for the Outer tunnel shall be provided through either the Gray or Red network.	T=O	
MSC-KM-3	CAK management services (enterprise or locally-owned) shall be connected to the local Red network.	O	Optional

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-KM-4	If the CAK Management Services operate at the same classification level as a Red network, a non-CDS Controlled Interface shall be used to control information flow between the CAK management services and the Red network.	T=O	
MSC-KM-5	If the CAK Management Services operate at a different classification level than a Red network or Gray network, a Controlled Interface that is also a CDS shall be used to control information flow between the CAK management services and the Red network or Gray network.	T=O	
MSC-KM-6	If multiple Red enclaves exist in the MSC Solution and the Outer CA resides in the Red network, the Outer CA must reside in the Red network with the highest classification level.	T=O	
MSC-KM-7	All device certificates issued by the Gray and Inner CAs, and their corresponding private keys, shall be treated as CUI (or higher as determined by the AO).	T=O	
MSC-KM-8	All certificates issued by the Outer and Inner CAs for the MSC Solution shall be Non-Person Entity (NPE) certificates.	T=O	
MSC-KM-9	Authentication certificates issued by the Gray and Inner CAs for the Solution shall be X.509 v3 certificates as defined in ITU-T Recommendation X.509.	T=O	
MSC-KM-10	CAKs issued to Outer Encryption Components are Unclassified, but are to be treated as CUI.	T=O	
MSC-KM-11	CAKs issued to Inner Encryption Components are classified to the level of the Red network.	T=O	
MSC-KM-12	Enterprise certificate and CAK management services shall be used to the greatest extent possible.	O	Optional
MSC-KM-13	The key sizes and algorithms for CA certificates and authentication certificates issued to VPN Gateways and Administrative Device Components shall be as illustrated in Table 5.	T=O	
MSC-KM-14	All public/private key pairs and certificates for VPN Gateways shall be used for authentication only.	T=O	
MSC-KM-15	All CAKs generated by or issued to an Encryption Component are to be used in strict accordance with approved protocols identified in this CP.	T=O	
MSC-KM-16	CAs shall not escrow private keys.	T=O	
MSC-KM-17	Outer and Inner CAs shall not have access to private keys used in the Solution Components.	T=O	

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-KM-18	A locally-run CA supporting an Inner Tunnel VPN Gateway shall be physically separated from a locally-run CA supporting an Outer Tunnel VPN Gateway.	T=O	
MSC-KM-19	The Outer and Inner CAs shall each operate in compliance with a Certificate Policy and Certification Practice Statement (CPS) that are formatted in accordance with Internet Engineering Task Force (IETF) Request for Comments (RFC) 3647.	T=O	
MSC-KM-20	KGCs and the MSC Solution being supported shall operate in compliance with a NSA-approved Key Management Plan (KMP).	T=O	
MSC-KM-21	MSC Solutions using CAKs are to obtain an IAD MD 110 waiver if the CAKs exist in red form during any part of the CAK life-cycle management process.	T=O	
MSC-KM-22	CAs shall run anti-virus software.	T=O	
MSC-KM-23	Authentication certificate profiles for the Gray and Inner CAs for the MSC Solution shall comply with IETF RFC 5280.	T=O	
MSC-KM-24	Private keys associated with on-line, locally run Outer and Inner CAs shall be protected using Hardware Security Modules (HSMs) validated to at least FIPS 140-2 Level 2. "On-line" means the CA is always powered on.	T=O	

11.11.2 CERTIFICATE ISSUANCE REQUIREMENTS

Requirements for issuing certificates are provided in Table 17.

Table 17. Certificate Issuance Requirements

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-KM-25	Gray and Red Management Services Components shall be initially keyed and loaded with certificates within a physical environment certified to protect the highest classification of the MSC Solution.	T=O	
MSC-KM-26	Outer and Inner CAs shall use Public Key Cryptographic Standard (PKCS)#10 and PKCS#7 to issue authentication certificates to VPN Gateways, and Gray and Red Management Services Components.	T	MSC-KM-27

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-KM-27	Outer and Inner CAs shall use IETF RFC 7030 Enrollment over Secure Transport (EST) to issue authentication certificates to VPN Gateways, and Gray and Red Management Services Components.	O	MSC-KM-26
MSC-KM-28	Certificate signing requests for Gray and Red Management Services Components shall be submitted to the CA in accordance with the CA's Certificate Policy (CP) and Certification Practices Statement (CPS).	T=O	
MSC-KM-29	Outer and Inner CAs shall issue certificates in accordance with their Certificate Policies and CPSs.	T=O	
MSC-KM-30	Certificate Policies and CPSs for non-Enterprise, locally run CAs shall ensure the CAs issue certificates within a defined and limited name space and assert: Unique Distinguished Names (DNs) Appropriate key usages A registered policy Object Identifier (OID)	T=O	
MSC-KM-31	Outer and Inner CAs shall assert at least one CRL Distribution Point (CDP) Uniform Resource Locator (URL) in certificates issued to Solution Infrastructure VPN Gateways, and Gray and Red Management Services Components. The CDP URL specifies the location of the CAs' CRLs.	T=O	
MSC-KM-32	The key validity period for certificates issued by non-Enterprise, locally run CAs to Solution Components shall not exceed 36 months.	T=O	
MSC-KM-33	Inner CAs shall only issue certificates to Inner VPN Gateways and Red Network Components of MSC Solutions.	T=O	
MSC-KM-34	Outer CAs shall only issue certificates to Outer VPN Gateways and Gray Network Components of MSC Solutions.	T=O	
MSC-KM-35	The Inner VPN Gateway shall only trust the Inner CA used for its network.	T=O	
MSC-KM-36	The Outer VPN Gateway shall only trust the Outer CA used within the solution.	T=O	
MSC-KM-37	The key validity period for certificates issued by Locally-run CAs shall not exceed 14 months.	T=O	
MSC-KM-38	New certificates shall be issued as needed in accordance with local policy.	T=O	

11.11.3 CERTIFICATE RENEWAL AND REKEY REQUIREMENTS

Requirements for renewing and rekeying certificates are provided in Table 18.

Table 18. Certificate Renewal and Rekey Requirements

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-KM-39	Certificate renewal or rekey shall occur prior to a certificate expiring.	T=O	
MSC-KM-40	If rekeying of the VPN Gateways is not completed prior to expiration of keys, they shall be rekeyed through the same process as initial keying.	T=O	
MSC-KM-41	Certificate renewal or rekey shall be performed in accordance with the CA's Certificate Policy and CPS.	T=O	
MSC-KM-42	Outer and Inner CAs shall issue renewed/rekeyed authentication certificates to Solution Components using PKCS#10 and PKCS#7.	T	MSC-KM-43
MSC-KM-43	Outer and Inner CAs shall issue renewed/rekeyed authentication certificates to Solution Components using EST (RFC 7030).	O	MSC-KM-42

11.11.4 CERTIFICATE REVOCATION REQUIREMENTS

Requirements for revoking certificates are provided in Table 19.

Table 19. Certificate Revocation Requirements

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-KM-44	CRL profiles shall comply with IETF RFC 5280.	T=O	
MSC-KM-45	Outer and Inner CAs shall revoke a certificate issued to Solution Components when the binding between the subject information and public key within the certificate issued is no longer considered valid.	T=O	
MSC-KM-46	Outer and Inner CAs shall make certificate revocation information available in the form of CRLs signed by the CAs.	T=O	
MSC-KM-47	Procedures for requesting certificate revocation shall comply with the CA's Certificate Policy and Certification Practices Statement.	T=O	

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-KM-48	Certificate Policies and CPSs for non-Enterprise, locally run CAs shall ensure revocation procedures address the following: Removal of a revoked infrastructure device (e.g., VPN Gateway) from the network Re-establishment of a Solution Component whose certificate was revoked Revocation of an authentication certificate if simultaneous use of the certificate is detected from different IP addresses	T=O	
MSC-KM-49	Enterprise CAs shall create and publish CRLs in accordance with the Enterprise CAs' Certificate Policies and CPSs.	T=O	
MSC-KM-50	Non-enterprise, locally run CAs shall publish new CRLs at least once every 28 days.	T=O	
MSC-KM-51	Non-enterprise, locally run CAs shall create a new CRL within one hour of a certificate being revoked.	T=O	
MSC-KM-52	Solution Components shall have access to new certificate revocation information within 24 hours of the CA creating a new CRL.	T=O	
MSC-KM-53	CRLs shall expire no later than 31 days after their issue date.	T=O	
MSC-KM-54	Non-enterprise, locally run CAs shall ensure that newly created CRLs are published at least 7 days prior to the expiration of the current CRLs.	T=O	
MSC-KM-55	The MSC Solution shall provide certificate revocation status information via an Online Certificate Status Protocol (OCSP) Responder on the Red and Gray network that is compliant with IETF RFC 6960.	O	Optional
MSC-KM-56	Certificate revocation status messages delivered by an OCSP Responder shall be digitally signed and compliant with IETF RFC 6960.	O	Optional
MSC-KM-57	Outer and Inner CAs shall make CRLs available to authorized CRL Distribution Points (CDPs), so that the CRLs can be accessed by Solution Components.	T=O	
MSC-KM-58	CRLs hosted by CDPs shall not contain extensions other than what is specified in IETF RFC 5280.	T=O	
MSC-KM-59	CRLs hosted on Inner CDPs shall be signed by the associated Red Network CA.	T=O	
MSC-KM-60	CRLs hosted on Outer CDPs shall be signed by the associated Gray Network CA.	T=O	
MSC-KM-61	CDPs shall only issue CRLs over port 80 (HTTP).	T=O	

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-KM-62	CRLs shall be transferred via an AO-approved one-way transfer mechanism from Red Network CAs to associated Inner CDP servers.	T=O	
MSC-KM-63	CRLs shall be transferred via an AO-approved one-way transfer mechanism from Gray Network CAs to associated Outer CDP servers.	T=O	
MSC-KM-64	Newly issued CRLs shall be transferred to CDP servers at least 4 days prior to the expiration of the current CRLs.	T=O	
MSC-KM-65	VPN Gateways shall attempt to download the latest CRL from a CDP at least once every 24 hours.	T=O	
MSC-KM-66	CDPs shall only accept traffic on port 80 and ports used for remote management traffic.	T=O	
MSC-KM-67	CDPs shall only accept connections from known VPN Gateway or Administration Workstation addresses or address ranges.	T=O	
MSC-KM-68	If an integrity check of a CRL pulled from a CDP fails, then VPN Gateways shall use the current cached CRL.	T=O	
MSC-KM-69	If a CDP is offline or contains an invalid CRL, then Inner and Outer VPN Gateway CRLs shall be manually updated prior to the expiration of the current CRLs.	T=O	
MSC-KM-70	Red Network CAs shall set the CRL Distribution Points extension of the certificates it generates for the MSC Solution to the list of URLs hosted by Inner CDPs from which Inner VPN Gateways can download the CRL.	T=O	
MSC-KM-71	Gray Network CAs shall set the CRL Distribution Points extension of the certificates it generates for the MSC Solution to the list of URLs hosted by Outer CDPs from which Outer VPN Gateways can download the CRL.	T=O	

11.11.5 CAK Generation and Distribution Requirements

Requirements for generating and distributing CAKs are provided in Table 20.

Table 20. CAK Generation and Distribution Requirements

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-KM-72	Generation of CAKs shall be performed by an NSA-approved Key Generation Component (KGC). NSA-approved means: a) a component from the CSfC Approved Products List; or b) a component approved for the CSfC solution by IADIR; or c) an already approved enterprise service.	T=O	

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-KM-73	Centralized generation, distribution and management of CAKs shall be performed by a dedicated KGC.	T=O	
MSC-KM-74	An Inner KGC shall generate CAKs for Inner MACsec Devices and may generate CAKs for Outer MACsec Devices.	T=O	
MSC-KM-75	If an Inner KGC generates CAKs for Outer MACsec Devices, the CAKs shall be transferred via an AO-approved one-way transfer mechanism	T=O	
MSC-KM-76	An Outer KGC may only generate CAKs for Outer MACsec Devices.	T=O	
MSC-KM-77	CAKs shall be AES 256 bits.	T=O	
MSC-KM-78	CAKs shall not be exposed in plaintext form until they are ready to be installed onto MACsec Devices. Installation of CAKs may be performed via file transfer or text input.	T=O	
MSC-KM-79	CAKs shall be protected from unauthorized disclosure when they are distributed outside of a controlled boundary or over unprotected communications channels through appropriate manual distribution procedures and methods, as defined in the KMP.	T	MSC-KM-80
MSC-KM-80	CAKs shall be protected from unauthorized disclosure when they are distributed outside of a controlled boundary or over unprotected communications channels through the use of pre-placed symmetric PSK Encryption Keys (PEKs) or an approved key distribution protocol.	O	MSC-KM-79
MSC-KM-81	PEKs shall be AES 256 bits.	T=O	
MSC-KM-82	The classification of pre-placed PEKs shall be the same as the classification of the CAKs that are encrypted with the pre-placed PEKs.	T=O	

11.11.6 CAK USAGE REQUIREMENTS

Requirements for using CAKs are provided in Table 21.

Table 21. CAK Usage Requirements

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-KM-83	CAKs are to only be used with the MACsec protocol.	T=O	
MSC-KM-84	CAKs and PEKs are to be stored within an approved cryptographic boundary within a Solution Component.	T=O	

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-KM-85	CAKs and PEKs exported from a Solution Component are to be protected from unauthorized disclosure through manual procedure protection methods.	T	MSC-KM-86
MSC-KM-86	CAKs and PEKs exported from a Solution Component are to be protected from unauthorized disclosure through encryption.	O	MSC-KM-85
MSC-KM-87	A compromised CAK/PEK is to never be used in the MSC Solution.	T=O	

11.11.7 CAK UPDATE REQUIREMENTS

Requirements for updating CAKs are provided in Table 22.

Table 22. CAK Update Requirements

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-KM-88	The same CAK shall be used in only one pair of MACsec Devices that are establishing an encryption tunnel.	T=O	
MSC-KM-89	CAKs shall be updated every 30 days, or as defined by the KMP.	T=O	
MSC-KM-90	PEKs are to be updated every 90 days, or as defined by the KMP.	T=O	

11.11.8 CAK COMPROMISE RECOVERY REQUIREMENTS

Requirements for recovering from compromised CAKs are provided in Table 23.

Table 23. CAK Compromise Recovery Requirements

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-KM-91	The KMP shall document the CAK/PEK compromise recovery process, to include: • Removal of a compromised infrastructure device (e.g., MACsec Devices) from the network, and • Re-establishing a MACsec Device after its CAK is compromised.	T=O	
MSC-KM-92	Accounting procedures need to support CAK and PEK compromise recovery to ensure all copies of compromised CAKs and PEKs are identified and updated.	T=O	
MSC-KM-93	CAKs/PEKs are to be updated immediately if they are considered compromised.	T=O	

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-KM-94	If a CAK/PEK is considered compromised, a compromise notification shall be submitted to the KGC along with a request to update the CAK/PEK.	T=O	
MSC-KM-95	If a CAK/PEK is compromised, the procedures for CAK/PEK compromise reporting as defined by the applicable KMP shall be followed.	T=O	
MSC-KM-96	If a compromised device is to be reused, that device must go through the initial CAK issuance process.	T=O	

11.12 Gray Network Firewall Requirements

Table 24 provides requirements for Gray Network Firewalls that are used when networks of different classification share an Outer Encryption Component.

Table 24. Gray Network Firewall (FW) Requirements

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-FW-1	Gray Network Firewalls shall permit IKE, IPsec, EAP-TLS and MACsec traffic between two Inner Encryption Components protecting networks of the same classification level.	T=O	
MSC-FW-2	Gray Network Firewalls shall allow HTTP traffic between Inner VPN Gateways and Gray CDP/OCSP Responder.	T	MSC-FW-3 and MSC-FW-4
MSC-FW-3	Gray Network Firewalls shall allow HTTP GET requests from Inner VPN Gateways to Inner CDPs/OCSP Responders for the URL of the CRL needed by the Inner VPN Gateway, and block all other HTTP requests.	O	MSC-FW-2
MSC-FW-4	Gray Network Firewalls shall allow HTTP responses from Inner CDPs/OCSP Responders to Inner VPN Gateways that contain a well-formed CRL per IETF RFC 5280, and block all other HTTP responses.	O	MSC-FW-2
MSC-FW-5	Gray Network Firewalls shall only accept management traffic on the physical ports connected to the Gray Management network.	T=O	
MSC-FW-6	Gray Network Firewalls shall only permit packets whose source and destination IP addresses match the external interfaces of Inner Encryption Components that support Red networks of the same classification level.	T=O	
MSC-FW-7	Gray Network Firewalls shall block all packets whose source address does not match a list of addresses or address ranges known to be reachable from the interface on which the packet was received.	T=O	

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-FW-8	Gray Network Firewalls shall allow control plane traffic (NTP, DHCP, DNS).	T=O	
MSC-FW-9	Gray Network Firewalls shall deny all traffic that is not explicitly allowed by requirements MSC-FW-1, MSC-FW- 2, MSC-FW-3, MSC-FW-4, MSC-FW-5 or MSC-FW-8.	T=O	

12 REQUIREMENTS FOR SOLUTION OPERATION, MAINTENANCE, AND HANDLING

12.1 REQUIREMENTS FOR THE USE AND HANDLING OF SOLUTIONS

The requirements in Table 25 shall be followed regarding the use and handling of the solution.

Table 25. Requirements for the Use and Handling of Solutions

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-GD-1	All Solution Components shall be physically protected as classified devices, classified at the level of the network with the highest classification in the solution or in any other MSC Solutions with which it is interconnected.	T=O	
MSC-GD-2	Only authorized and appropriately cleared (or escorted) administrators and security personnel shall have physical access to the Solution Components.	T=O	
MSC-GD-3	All components of the solution shall be disposed of as classified devices, unless declassified using AO-approved procedures.	T=O	
MSC-GD-4	Acquisition and procurement documentation shall not include information about how the equipment will be used, to include that it will be used to protect classified information.	T=O	
MSC-GD-5	The solution owner shall allow, and fully cooperate with, NSA or its authorized agent to perform an IA compliance audit (including, but not limited to, inspection, testing, observation, interviewing) of the solution implementation to ensure it meets the latest version of this CP.	T=O	
MSC-GD-6	The AO will ensure that a compliance audit shall be conducted every year against the latest version of this CP as part of the annual solution re-registration process.	T=O	

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-GD-7	Results of the compliance audit shall be provided to and reviewed by the AO.	T=O	
MSC-GD-8	Customers interested in registering their solution against this CP shall register with NSA and receive approval prior to the AO authorization to operate.	T=O	
MSC-GD-9	The implementing organization shall complete and submit an MSC CP requirements compliance matrix to their respective AO.	T=O	
MSC-GD-10	Registration and re-registration against this CP shall include submission of CP registration forms and compliance matrix to NSA.	T=O	
MSC-GD-11	When a new approved version of the MSC CP is published by NSA, the AO shall ensure compliance against this new CP within 6 months.	T=O	
MSC-GD-12	Solution implementation information, which was provided to NSA during solution registration, shall be updated annually (in accordance with Section 14.3) as part of the annual re-registration process.	T=O	
MSC-GD-13	Audit log data shall be maintained for a minimum of 1 year.	T=O	
MSC-GD-14	The amount of storage remaining for audit events shall be assessed quarterly to ensure that adequate memory space is available to continue recording new audit events.	T=O	
MSC-GD-15	Audit data shall be frequently off-loaded to a backup storage medium.	T=O	
MSC-GD-16	A set of procedures shall be developed by the implementing organization to provide guidance for identifying and reporting security incidents associated with the audit events to the proper authorities and to the data owners.	T=O	
MSC-GD-17	The implementing organization shall develop a continuity of operations plan for auditing capability which includes a mechanism or method for determining when the audit log is reaching its maximum storage capacity.	T=O	
MSC-GD-18	The implementing organization shall develop a continuity of operations plan for auditing capability which includes a mechanism or method for off-loading audit log data for long-term storage.	T=O	

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-GD-19	The implementing organization shall develop a continuity of operations plan for auditing capability which includes a mechanism or method for responding to an overflow of audit log data within a product.	T=O	
MSC-GD-20	The implementing organization shall develop a continuity of operations plan for auditing capability which includes a mechanism or method for ensuring that the audit log can be maintained during power events.	T=O	
MSC-GD-21	Strong passwords shall be used that comply with the requirements of the AO.	T=O	
MSC-GD-22	Security critical patches shall be tested and subsequently applied to all components in the solution in accordance with local policy and this CP.	T=O	
MSC-GD-23	Local policy shall dictate how the Security Administrator will install patches to Solution Components.	T=O	
MSC-GD-24	Solution components shall comply with local TEMPEST policy.	T=O	
MSC-GD-25	All hardware components shall be tracked through an AO-approved inventory management process that identifies each component as part of a CSfC Solution.	T=O	

Additional MSC-GD requirements can be found in Section 13.

12.2 Requirements for Incident Reporting

Table 26 lists requirements for reporting security incidents to NSA to be followed in the event that a solution owner identifies a security incident which affects the solution. These reporting requirements are intended to augment, not replace, any incident reporting procedures already in use within the solution owner's organization. It is critical that Security Administrators (SAs), Certificate Authority Administrators (CAAs), CAK Administrators (CAKAs), and Auditors are familiar with maintaining the solution in accordance with this CP. Based on familiarity with the known-good configuration of the solution, personnel responsible for the operations and maintenance of the solution will be better equipped to identify reportable incidents.

For the purposes of incident reporting, "malicious" activity includes not only events that have been attributed to activity by an adversary, but also any events that are unexplained. In other words, an activity is assumed to be malicious unless it has been determined to be the result of known non-malicious activity.

Table 26 only provides requirements directly related to the incident reporting process. See Section 11.9 for requirements supporting the detection of events that may reveal that a reportable incident has occurred.

Table 26. Incident Reporting Requirements (RP)

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-RP-1	Solution owners shall report confirmed incidents meeting the criteria in MSC-RP-3 through MSC-RP-14 within 24 hours of detection via Joint Incident Management System (JIMS) or contacting NSA as specified in the CSfC Registration Letter issued for the solution.	T=O	
MSC-RP-2	At a minimum, the organization shall provide the following information when reporting security incidents: • CSfC Registration Number • Point of Contact (POC) name, phone, email • Alternate POC name, phone, email • Classification level of affected solution • Name of affected Network(s) • Affected component(s) manufacturer/ vendor • Affected component(s) model number • Affected component(s) version number • Date and time of incident • Description of incident • Description of remediation activities • Is Technical Support from NSA requested? (Yes/No)	T=O	
MSC-RP-3	Solution owners shall report a security failure in any of the CSfC Solution Components.	T=O	
MSC-RP-4	Solution owners shall report any evidence of a compromise or spillage of classified data caused by a failure of the CSfC solution.	T=O	
MSC-RP-5	For Gray Network interfaces, solution owners shall report any malicious inbound and outbound traffic.	T=O	
MSC-RP-6	Solution owners shall report any evidence of an unauthorized device/user gaining access to the classified network via the solution.	T=O	
MSC-RP-7	Solution owners shall report if a Solution Component sends traffic with an unauthorized destination address.	T=O	
MSC-RP-8	Solution owners shall report any malicious configuration changes to the components.	T=O	

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-RP-9	Solution owners shall report any unauthorized escalation of privileges to any of the CSfC Solution Components.	T=O	
MSC-RP-10	Solution owners shall report any evidence of malicious physical tampering with Solution Components.	T=O	
MSC-RP-11	Solution owners shall report any evidence that one or both of the layers of the solution failed to protect the data.	T=O	
MSC-RP-12	Solution owners shall report any significant degradation of services provided by the solution.	T=O	
MSC-RP-13	Solution owners shall report malicious discrepancies in the number of connections established by the Outer Encryption Component.	T=O	
MSC-RP-14	Solution owners shall report malicious discrepancies in the number of connections established by the Inner Encryption Component.	T=O	

13 ROLE-BASED PERSONNEL REQUIREMENTS

The roles required to administer and maintain the solution are defined below, along with doctrinal requirements for these roles.

Security Administrator – The Security Administrator shall be responsible for maintaining, monitoring, and controlling all security functions for the entire suite of products composing the MSC Solution. Security Administrator duties include but are not limited to:

1) Ensuring that the latest security-critical software patches and updates (such as Information Assurance Vulnerability Alerts (IAVAs)) are applied to each product.

2) Documenting and reporting security-related incidents to the appropriate authorities.

3) Coordinating and supporting product logistic support activities including integration and maintenance. Some logistic support activities may require that the Security Administrator escort uncleared personnel.

4) Employing adequate defenses of auxiliary network devices to enable proper and secure functionality of the MSC Solution.

5) Ensuring that the implemented MSC Solution remains compliant with the latest version of this CP.

Certificate Authority Administrator (CAA) – The CAA shall be responsible for maintaining, monitoring, and controlling all security functions for the CA products. CAA duties include but are not limited to:

1) Administering the CA, including authentication of all components requesting certificates.

2) Maintaining and updating the CRL.

CAK Administrator (CAKA) – The CAKA shall be responsible for maintaining, monitoring, and controlling all security functions for the KGC products. CAKA duties include but are not limited to:

1) Administering the KGC, including authentication of all components requesting CAKs and PEKs.

2) Maintaining and updating the CAK revocation list.

Auditor – The Auditor shall be responsible for reviewing the actions performed by the Security Administrator, CAA, or CAKA and events recorded in the audit logs to ensure that no action or event represents a compromise to the security of the MSC Solution. Auditor duties include but are not limited to:

1) Reviewing, managing, controlling, and maintaining security audit log data.

2) Documenting and reporting security-related incidents to the appropriate authorities.

3) The Auditor will only be authorized access to Outer and Inner administration components.

Solution Integrator – In certain cases, an external integrator may be hired to implement a MSC Solution based on this CP. Solution Integrator duties may include but are not limited to:

1) Acquiring the products that compose the solution.

2) Configuring the MSC Solution in accordance with this CP.

3) Documenting, testing, and maintaining the solution.

4) Responding to incidents affecting the solution.

Additional policies related to the personnel that perform these roles in a MSC Solution are identified in Table 27.

Table 27. Role-Based Personnel Requirements

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-GD-26	The Security Administrator, CAAs, CAKAs, Auditor, and Solution Integrators shall be cleared to the highest level of data protected by the MSC Solution. When an Enterprise CA/KGC is used in the solution, the CAA/CAKA already in place may also support this solution, provided they meet this requirement.	T=O	

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-GD-27	The Security Administrator, CAA, CAKA, and Auditor roles shall be performed by different people.	T=O	
MSC-GD-28	All Security Administrators, CAAs, CAKAs, and Auditors shall meet local Information Assurance (IA) training requirements.	T=O	
MSC-GD-29	The CAA(s)/CAKA(s) for the Inner tunnel shall be different individuals from the CAA(s)/CAKA(s) for the Outer tunnel.	T=O	
MSC-GD-30	The Security Administrator(s) for the Inner Encryption Components and supporting components on Red networks shall be different individuals from the Security Administrator(s) for the Outer Encryption Components and supporting components on Gray networks.	T=O	
MSC-GD-31	Administrators shall periodically inspect the physical attributes of infrastructure hardware for signs of tampering or other unauthorized changes.	O	Optional
MSC-GD-32	The Auditor shall review all logs specified in this CP at least once a week.	T=O	
MSC-GD-33	Security Administrators shall initiate the certificate revocation process prior to disposal of any Solution Component.	T=O	
MSC-GD-34	Auditing of the Outer and Inner CA operations shall be performed by individuals who were not involved in the development of the CP and CPS, or integration of the MSC Solution.	T=O	
MSC-GD-35	Auditing of the Key Generation Component operations shall be performed by individuals who were not involved in the development of the KMP, or integration of the MSC Solution.	T=O	

14 INFORMATION TO SUPPORT AO

This section details items that likely will be necessary for the customer to obtain approval from the system AO. The customer and AO have obligations to perform the following:

- The customer, possibly with support from a System Integrator, instantiates a solution implementation that follows the NSA-approved CP.
- The customer has a testing team develop a Test Plan and perform testing of the MSC Solution (see Section 14.1).

- The customer has system certification and accreditation performed using the risk assessment information referenced in Section 14.2.
- The customer provides the results from testing and system certification and accreditation to the AO for use in making an approval decision. The AO is ultimately responsible for ensuring that all requirements from this CP have been properly implemented in accordance with this CP.
- The customer registers the solution with NSA and re-registers yearly to validate its continued use as detailed in Section 14.3.
- Customers who want to use a variant of the solution detailed in this CP will contact their NSA/IAD Client Advocate to determine ways to obtain NSA approval.
- The AO will ensure that a compliance audit shall be conducted every year against the latest version of the MSC CP, and the results shall be provided to the AO.
- The AO will ensure that certificate revocation information is updated on all the Solution Components in the MSC Solution in the case of a compromise.
- The AO will ensure that any Layer 2 or Layer 3 control plane protocols that are used in the solution are necessary for the operation of the network and that local policy supports their use.
- The AO will report incidents affecting the solution in accordance with Section 12.2.

The system AO maintains configuration control of the approved solution implementation over the lifecycle of the solution. Additionally, the AO shall ensure that the solution remains properly configured with all required security updates implemented.

14.1 Solution Testing

This section provides a framework for a Test and Evaluation (T&E) plan and procedures to validate the implementation of a MSC Solution. This T&E will be a critical part of the approval process for the AO, providing a robust body of evidence that shows compliance with this CP.

The security features and operational capabilities associated with the use of the solution shall be tested. The following is a general high-level methodology for developing the Test Plan and Procedures and for the execution of those procedures to validate the implementation and functionality of the MSC Solution. The entire solution, to include each component described in Section 5, is addressed by this Test Plan.

1) Set up the baseline network and configure all components.
2) Document the baseline network configuration. Include product model and serial numbers, and software version numbers at a minimum.

3) Develop a Test Plan for the specific implementation using the test requirements from Section 15. Any additional requirements imposed by the local AO should also be tested, and the Test Plan shall include tests to ensure that these requirements do not interfere with the security of this solution as described in this CP.

4) Perform testing using the Test Plan derived in Step 3. Network testing will consist of both Black Box testing and Gray Box testing. A two-person testing approach should be used to administer the tests. During test execution, security and non-security related discrepancies with the solution shall be documented.

5) Compile findings, to include comments and vulnerability details as well as possible countermeasure information, into a Final Test Report to be delivered to the AO for approval of the solution.

The following testing requirement has been developed to ensure that the MSC Solution functions properly and meets the configuration requirements from Section 11. Testing of these requirements should be used as a minimum framework for the development of the detailed Test Plan and Procedures.

Table 28. Test (TR) Requirements

Req #	Requirement Description	Threshold / Objective	Alternative
MSC-TR-1	The organization implementing the Capability Package shall perform all tests listed in Section 15.	T=O	

14.2 Risk Assessment

The Risk Assessment of the MSC Solution presented in this CP focuses on the types of attacks that are feasible against this solution and the mitigations that can be employed. Customers should contact their NSA/IAD Client Advocate to request this document, or visit the Secret Internet Protocol Router Network (SIPRNet) CSfC site for information. The process for obtaining the Risk Assessment is available on the SIPRNet CSfC website. The AO shall be provided a copy of the NSA Risk Assessment for their consideration in approving the use of the solution.

14.3 Registration of Solutions

All customers using CSfC solutions to protect information on National Security Systems shall register their solution with NSA prior to operational use. This registration will allow NSA to track where MSC Solutions are instantiated and to provide the AOs at those sites with appropriate information, including any significant vulnerabilities that may be discovered in components or high-level designs approved for these solutions. The CSfC solution registration process is available at https://www.nsa.gov/resources/everyone/csfc.

Solution registrations are valid for one year from the date the solution registration is approved, at which time customers are required to re-register their solution to continue using it. Approved CPs will be reviewed twice a year, or as events warrant. Registered users of this CP will be notified when an updated version is published. When a new version of this IAD-approved CP is published, customers will have six months to bring their solutions into compliance with the new version of this CP and re-register their solution (see requirement MSC-GD-11). Customers are also required to update their registrations whenever the information provided on the registration form changes.

15 TESTING REQUIREMENTS

This section contains the specific tests that allow the Security Administrator or System Integrator to ensure they have properly configured the solution. As defined in Section 9, to comply with this CP, a solution must at minimum implement all Threshold requirements associated with each of the capabilities it supports, and should implement the Objective requirements associated with those capabilities where feasible. These tests may also be used to provide evidence to the AO regarding compliance of the solution with this CP. Note that the details of the procedures are the responsibility of the final developer of the solution test plan in accordance with AO-approved network procedures. The AO is ultimately responsible for ensuring that all requirements from this CP have been properly implemented.

15.1 PRODUCT SELECTION

This section contains a procedure to verify that the components in this CP were selected to ensure independence in several important features.

Requirements being tested: MSC-PS-1 through MSC-PS-14, MSC-SR-2,

Procedure Description:

1) For each VPN Gateway, perform the following:

    a) Verify that the Inner and Outer VPN Gateways are on the list of IPsec VPN Gateways on the CSfC Components List. (MSC-PS-1)

2) For each MACsec Device, perform the following:

    a) Verify that the Inner and Outer MACsec Devices are on the list of MACsec Ethernet Encryptors on the CSfC Components List. (MSC-PS-2)

3) For the Inner Encryption Component and the Outer Encryption Component, perform the following:

a) Verify that the Inner Encryption Component and Outer Encryption Component either come from different independent manufacturers or that NSA has determined that sufficient implementation independence exists. (MSC-PS-6).

b) Verify that the cryptographic libraries used by the Inner Encryption Component and Outer Encryption Component either come from different independent manufacturers or that NSA has determined that sufficient implementation independence exists. (MSC-PS-7)

c) Verify the Inner Encryption Component and Outer Encryption Component are running on physically separate hardware platforms. (MSC-PS-10)

4) For each CA, perform the following:

a) Verify the Inner and Outer tunnel CAs came from the list of CAs on the CSfC Components List or are Enterprise CAs. (MSC-PS-4)

b) Verify that the cryptographic libraries used by the Inner and Outer CAs either come from different independent manufacturers or that NSA has determined that sufficient implementation independence exists. (MSC-PS-8)

c) Verify that the Inner and Outer CAs either come from different independent manufacturers or that NSA has determined that sufficient implementation independence exists. (MSC-PS-12)

5) For each Gray Network Firewall and Inner Encryption Component, perform the following:

a) Verify that the Gray Network Firewall and Inner Encryption Component either come from different independent manufacturers or that NSA has determined that sufficient implementation independence exists. (MSC-PS-9)

b) Verify the Gray Network Firewall and Inner Encryption Component are running on physically separate hardware platforms. (MSC-PS-11)

6) For each Gray Network Firewall, perform the following:

a) Verify that the Gray Network Firewalls are on the list of Firewalls on the CSfC Components List. (MSC-PS-3)

7) For each IPS, perform the following:

a) Verify that the IPS is on the list of IPSs on the CSfC Components List. (MSC-PS-5)

8) For all components in the solution:

a) Verify that each component has gone through a Product Supply Chain Threat Assessment. (MSC-PS-13)

b) Verify that the components are configured to use the NIAP-certified evaluated configuration. (MSC-PS-14)

9) For sites requiring interoperability, ensure that VPN Gateways selected for each tunnel can be configured to communicate using the requirements specified in this Capability Package. (MSC-SR-2)

Expected Result:

The results of the inspection should reveal that the MSC Solution components conform to this CP.

15.2 Overall Solution

This section contains a procedure to create an accurate record of the physical components composing the MSC Solution (including workstations, VPN Gateways, CAs, and wiring). The test will also ensure that the physical implementation of the MSC Solution matches one of the high-level designs given in the VPN Capability Package.

Requirements being tested: MSC-SR-1, MSC-SR-3 through MSC-SR-9, MSC-SR-11

Procedure Description:

1) Verify that the physical location of any network services for the Outer Encryption Component is located on the appropriate Gray Management network. Similarly, verify these components for the Inner Encryption Component are located on the appropriate Red network. (MSC-SR-1)

2) Verify that the time of day on the Inner Encryption Components and Red Management Services are synchronized with a time source located in the Red Network. (MSC-SR-3)

3) Verify that the time of day on the Outer Encryption Components, Gray Network Firewall, and Gray Management Services are synchronized with a time source located in the Gray Network. (MSC-SR-4)

4) Verify that all default accounts, passwords, community strings and other default access controls have been either changed or removed. (MSC-SR-5)

5) If Red networks of different classification levels share an Outer Encryption Component, verify different ports are used on the Outer Encryption Component or a Gray Network Firewall is located before the Red network of a lower classification level. (MSC-SR-7)

6) If Red networks of different classification levels share an Outer Encryption Component, verify the Gray network between the CA, Administration Workstations, or CDP/OCSP Responder and Inner Encryption Component for a Red network of a lower classification level are on a different port on the Outer Encryption Component or behind a Gray Network Firewall. (MSC-SR-8)

7) Ensure there are no wireless or physical connections to the solution that are not included in this CP, which may allow traffic to leave a Red or Gray network in a manner that does not go through the MSC Solution (or an NSA-certified encryptor). (MSC-SR-9)

8) Verify that when using multiple Inner Encryption Components, those components are placed in parallel. (MSC-SR-11)

Expected Result:

For Step 1, Gray Management network traffic should be separate from Red Management network traffic. For Steps 2 and 3, the time of day should be synchronized to the appropriate sources. For Step 4, all defaults should be changed or removed. For steps 5 and 6, the Gray Network Firewall should be appropriately located or separate ports on the Outer Encryption Component should be used. For step 7, there should be no extraneous wireless or physical connections allowing data to leave Red or Gray networks besides through the MSC Solution (or an NSA-certified encryptor). For Step 8, Inner Encryption Components are accurately placed.

15.3 VPN GATEWAY CONFIGURATIONS

This section contain procedures to ensure that the configurations for all the VPN Gateways in the MSC Solution follow the requirements given in this CP.

Requirements being tested: MSC-VG-1 through MSC-VG-4, MSC-VG-6 through MSC-VG-15, MSC-VG-17

Procedure Description:

1) For each VPN Gateway in the solution, perform the following:

   a) Obtain the current configuration for the VPN Gateway.

   b) Verify a unique device certificate is loaded with the corresponding CA signing certificate. (MSC-VG-4)

   c) Verify a device certificate from a CA included in the MSC Solution is listed in the configuration for authentication. (MSC-VG-6)

   d) Ensure the corresponding CA signing certificate and certificate revocation information are on the VPN Gateway. (MSC-VG-7)

   e) Verify the requirements MSC-VG-1 through MSC-VG-3 and MSC-VG-8 through MSC-VG-15 are configured properly.

   f) Verify that the VPN Gateways are configured to re-authenticate the identity of the VPN Gateways at the other end before rekeying the IKE SA. (MSC-VG-17)

Expected Result:

For Step 1, all VPN Gateways should be configured properly according to the requirements found in this CP.

15.4 MACsec Device Configurations

This section contains a procedure to ensure that the configurations for all the MACsec Devices in the MSC Solution follow the requirements given in this CP.

Requirements being tested: MSC-MD-1 through MSC-MD-11

Procedure Description:

1) For each MACsec Device, perform the following:
    a) Obtain the current configuration for the MACsec Device.
    b) Verify AES Key Wrap in CMAC mode is used for key distribution and AES GCM is used for encryption; also verify key sizes for both are set to 256 bits. (MSC-MD-1)
    c) Verify authentication is being performed with a CAK. (MSC-MD-2)
    d) Verify the CAK is 256 bits and confirm it was generated in accordance with the KMP. (MSC-MD-3)
    e) Verify the CKN is set to a minimum of 64 bits. (MSC-MD-4)
    f) Verify the Key Server value is set to 0 (zero) if the MACsec Device is the designated Key Server for the MACsec pair; otherwise ensure the value is set to 1. (MSC-MD-5)
    g) Ensure data delay protection for MKA is enabled. (MSC-MD-6)
    h) Ensure the MKA Lifetime Timeout limit is set to 6.0 seconds and the Hello Timeout limit is set to 2.0 seconds. (MSC-MD-7)
    i) Verify the replay window is set to 2; otherwise confirm it is the lowest possible number given the nature of the Black network being traversed. (MSC-MD-8)
    j) Verify all traffic on the external facing ports require encryption. (MSC-MD-9)
    k) Verify that printed or electronic versions of the configuration files are protected to the highest classification of the MACsec Device's CAK. (MSC-MD-10)
    l) Ensure the Confidentiality Offset is set to 0 (zero). (MSC-MD-11)

Expected Result:

For Step 1, all MACsec Devices should be configured properly according to the requirements found in this CP.

15.5 Inner and Outer Encryption Component Configurations

This section contains procedures to ensure that the configurations for all the Inner and Outer Encryption Components in the MSC Solution follow the requirements given in this CP.

Requirements being tested: MSC-SR-3 through MSC-SR-6, MSC-SR-11, MSC-SR-12, MSC-DM-1, MSC-DM-4 through MSC-DM-6, MSC-DM-8 through MSC-DM-11, MSC-DM-14, MSC-DM-15, MSC-DM-21, MSC-IR-1 through MSC-IR-4, MSC-OR-1, MSC-OR-3 through MSC-OR-6

Procedure Description:

1) For each Encryption Component, perform the following:

   a) Ensure that default accounts, passwords, community strings, and other default access control mechanisms are changed or removed. (MSC-SR-5)

   b) Ensure that all components are configured in accordance with local policy and applicable U.S. Government guidance, or, in the event of conflict between this CP and local policy, this CP takes precedence. (MSC-SR-6)

   c) Ensure that Inner Encryption Components are not performing switching or routing for other Encryption Components. (MSC-SR-12)

   d) Ensure the time of day on the Inner Encryption Component and Red Management Services matches the current time. This should be within a small margin of error, to be determined by the AO. (MSC-SR-3)

   e) Ensure the time of day on the Outer Encryption Component, Gray Network Firewall and Gray Management Services matches the current time. This should be within a small margin of error, to be determined by the AO. (MSC-SR-4)

2) For each Inner Encryption Component, use the configuration from 1a and perform the following:

   a) Log into any Inner VPN Gateways and verify that they are configured to use Tunnel or Transport mode IPsec with an associated IP Protocol (e.g., GRE). (MSC-IR-1)

   b) Log into any Inner VPN Gateways and verify that the MTU (for IPv4) or the PMTU (for IPv6) has been configured to an appropriate size. (MSC-IR-2)

c) Using a packet analyzer tool on the Inner Encryption Component, verify that traffic leaving the external interface going to the Outer Encryption Component is encrypted. (MSC-IR-3)

d) Using a packet analyzer tool on the Inner Encryption Component, verify that traffic coming through the external interface of the Inner Encryption Component is decrypted. (MSC-IR-4)

e) Verify a separate LAN or VLAN is established on the Enterprise/Red network and using a packet sniffer, inspect traffic within the Enterprise/Red Network to ensure it is being used exclusively for all management of Inner Encryption Components and Solution Components within the Red network. (MSC-DM-4)

3) For each Outer Encryption Component, use the configuration from 1a and perform the following:

a) Log into the Outer VPN Gateways and verify that they are configured to use Tunnel mode IPsec. (MSC-OR-1)

b) Using a packet analyzer tool on the Outer Encryption Component, verify that traffic leaving the external interface going to the Black Network is encrypted. (MSC-OR-3)

c) Using a packet analyzer tool on the Outer Encryption Component, verify that traffic coming through the external interface of the Outer Encryption Component is decrypted. (MSC-OR-4)

d) Verify that any Outer Encryption Components are not configured to perform split tunneling or routing. (MSC-OR-5, MSC-OR-6)

e) Verify that a separate LAN or VLAN is established on the Gray network and using a packet sniffer, inspect traffic within the Gray network to ensure it is being used exclusively for all management of Outer Encryption Components, Gray Network firewall, and Solution Components within the Gray network. (MSC-DM-5)

4) For all device administration, verify that requirements MSC-DM-1, MSC-DM-8, MSC-DM-21, and MSC-SR-11 are configured properly.

5) For each administration workstation, ensure the Security Administrator is required to authenticate to the component before being granted access. (MSC-DM-9)

6) For each administration workstation, ensure the Security Administrator is required to authenticate to Solution Components using CNSA Suite compliant certificates. (MSC-DM-10)

7) Ensure that certificate signing requests are initiated by the Security Administrator as part of their initial keying within the solution. (MSC-DM-11)

8) Ensure that devices that use EST as detailed in RFC 7030 for certificate management. (MSC-DM-14)

9) Ensure that the same Administration Workstation is not used to manage both Inner and Outer Encryption Components. (MSC-DM-15)

10) Ensure that requirement MSC-DM-6 has been configured properly.

Expected Results:

For Steps 1 through 3, the Inner and Outer Encryption Components should be configured properly according to the requirements found in this CP. For Steps 4 through 10, all administration devices should be configured properly according to the requirements found in this CP.

15.6 PORT FILTERING

This section contains procedures to ensure that the port filtering configurations for the MSC Solution follow the requirements in this CP.

Requirements being tested: MSC-PF-1 through MSC-PF-7 and MSC-PF-9 through MSC-PF-12

Procedure Description:

1) Perform the following steps on each of the Solution Components:

   a) Log into the component.

   b) Verify through the configuration file that network interfaces are restricted to the smallest address range, ports, and protocols. (MSC-PF-1)

   c) Verify through the configuration file that all unused network interfaces are disabled. (MSC-PF-2)

2) For the Outer Encryption Components, perform the following:

   a) Obtain the current configuration for the Outer Encryption Component.

   b) Verify that the requirements of MSC-PF-3, MSC-PF-4, MSC-PF-6, MSC-PF-7, and MSC-PF-9 through MSC-PF-12 are met.

3) For the Inner Encryption Components, perform the following:

   a) Obtain the current configuration for the Inner Encryption Component.

   b) Verify the requirement MSC-PF-5 is met.

Expected Results:

Port filtering is performed correctly and in compliance with requirements in this CP.

15.7 Configuration Change Detection

This section contains a procedure to ensure that changes made to any of the MSC Solution configurations are detected by the Configuration Change Detection tool.

Requirements being tested: MSC-CM-1 through MSC-CM-4

Procedure Description:

1) The following steps shall be performed for each of the Solution Components.
    a) Log into the Solution Component.
    b) Compare the current version of the Solution Component's configuration with the stored baseline and ensure the current version matches the stored configuration. (MSC-CM-1)
    c) Make a change to the configuration, preferably something that is not fundamental to the security of the MSC Solution.
    d) Look in the audit log to determine if a log entry has been generated about the configuration change and that the changes from 1c are recorded. (MSC-CM-2, MSC-CM-3)
    e) Inspect the monitoring service to verify that the service has detected a change in configuration. (MSC-CM-4)

Expected Result:

The Auditor will validate the baseline configuration was stored in Step 1b. In Step 1d, there should be a log entry created for the configuration change in the audit log including the actual configuration change. Lastly, if there was a configuration change, a monitoring service will detect a change in the configuration.

15.8 Continuous Monitoring

This section contains procedures for ensuring traffic is monitored for and alerts generated for potential unauthorized/malicious traffic. It also contains procedures for ensuring a SIEM is in place to collect logs and that it is configured correctly.

Requirements being tested: MSC-MR-1 through MSC-MR-18

Procedure Description:

1) Ensure that an IDS/IPS is deployed to monitor traffic in at least one of three locations (MSC-MR-1 through MSC-MR-6):

a) Between the Outer Encryption Component and Gray Network Firewall (M1)

b) Between the Gray Network Firewall and the Inner Encryption Component (M2)

c) On the internal side of the Inner Encryption Component (M3)

2) Ensure that each IDS/IPS in the solution is configured to send alerts to the Security Administrator and, where possible, block malicious traffic. (MSC-MR-7, MSC-MR-8)

3) Ensure that each IDS/IPS in the solution is configured with the rules that will generate alerts and, where possible, block traffic for any unauthorized source and destination IP addresses. (MSC-MR-9 through MSC-MR-12)

4) Ensure that a SIEM is implemented. (MSC-MR-13 through MSC-MR-16)

a) Ensure the SIEM is implemented in the Gray network.

b) Otherwise, if the SIEM is implemented within the Enterprise/Red network, ensure devices are configured to push events to an Enterprise/Red SIEM and through an AO-approved one-way tap.

c) Send packets expected to be blocked by the Outer Encryption Component or Gray Network Firewall. Ensure the SIEM sends alerts to the Auditor when anomalous behavior such as this is detected.

d) Ensure that logs from the Outer Encryption Component, Gray Network Firewall, and any other components located within the Gray Management Services are collected on the Gray SIEM.

e) Ensure these logs are encrypted with SSHv2, IPSec, MACsec, or TLS v1.2 or later.

5) Ensure that any one-way taps are deployed as per MSC-MR-17 and MSC-MR-18.

a) Ensure that any one-way taps deployed as part of the solution are approved for use by the AO.

b) Ensure that the SIEM implemented at the Red level that collects black and/or gray monitoring data sent through any one-way tap is deployed in an enclave isolated from the Red/Enterprise network.

c) Ensure that monitoring data flowing from M2 and/or M1 can transit to the SIEM if implemented at the Red level.

d) Attempt to send other data through the one-way taps to determine if this data is blocked.

Expected Results:

For Steps 1 – 3, an IDS or IPS is in place to monitor, block where possible, and send alerts as appropriate. For Steps 4 and 5, a SIEM shall be implemented either in the Gray network or the Red/Enterprise network via one-way taps, as approved by the AO, and only monitoring data that will be able to transit through these taps.

15.9 AUDITING

This section contains procedures for ensuring audit events are detected, the proper information is logged for each event, and there is a procedure detailed in the CPS documentation for auditing each CA and in the KMP for auditing each KGC.

Requirements being tested: MSC-AU-1 through MSC-AU-9, MSC-AU-11, MSC-AU-15 through MSC-AU-25, MSC-DM-16 through MSC-DM-19

Procedure Description:

1) Examples for testing the ability of each Encryption Component to audit and log audit events specified in this CP are given below. Verify that for each event logged, the applicable data regarding the event is recorded for the log entry in accordance with Section 11.10.

   a) All actions performed by a user with superuser privileges (auditor, administrator, etc.) and any escalation of user privileges. (MSC-AU-6, MSC-AU-7)

      i) Log in as an administrator to the Encryption Component.

      ii) Perform a variety of administrator actions on the Encryption Component.

      iii) Verify a log entry was created for each action taken in Step ii that required superuser privileges and also states the escalation of privileges.

      iv) Revert back to the baseline configuration, eliminating the changes made in Step ii.

      v) Repeat the above with the Auditor role.

   b) Changes to time. (MSC-AU-9)

      i) Log in as a Security Administrator to the Encryption Component.

      ii) Modify the system time on the Encryption Component by at least 1 hour.

      iii) Verify a log entry was created due to the change in system time and by whom.

      iv) Revert the system time back to the accurate time of day.

c) Log into and out of the MSC Solution as a normal user and send traffic to the Red Network. Then log into the central log server as an Auditor, and inspect the audit entry for the following: (MSC-DM-16, MSC-DM-17)

   i) Verify that the log on as a normal user is logged and has an identifiable code for the type of event. (MSC-AU-4, MSC-AU-17)

   ii) Verify that the log entry identifies the subject accessing the solution. (MSC-AU-19)

   iii) Verify that the log entry identifies the event. (MSC-AU-16)

   iv) Verify that the log entry includes the time, date, and the time zone offset. (MSC-AU-15)

d) Establish and terminate an IPsec tunnel. Verify in the logs, that these two events were logged. Repeat for a MACsec tunnel. (MSC-AU-1, MSC-AU-2)

e) Ensure all built-in self-test results have been recorded in the audit log which may indicate failures in cryptographic functionality. (MSC-AU-11)

   i) Completely power down the Encryption Component.

   ii) Power the Encryption Component back up so that the automatic self-tests are run.

   iii) Verify a log entry was created due to running the self-test.

f) Log into a Solution Component as a Security Administrator and delete previously recorded audit log. Verify the log recorded this deletion. (MSC-AU-3)

g) As the Certificate Administrator, log into the audit log and attempt to delete a log entry. Verify this action is recorded with a failure code. (MSC-AU-5, MSC-AU-18)

h) Verify a log entry was created for the attempted unauthorized action.

2) Verify the source address for all audit log entries is recorded. (MSC-AU-20, MSC-AU-21)

3) Verify that all logs forwarded to a log server on a Gray Management network are configured to be encrypted while in transit using SSHv2, IPSEC, MACsec, or TLS 1.2 or later with the appropriate CNSA Suite algorithm supported by the solution. (MSC-DM-18)

4) Verify that all logs forwarded to a log server on a Red Management network are configured to be encrypted while in transit using SSHv2, IPsec, MACsec, or TLS 1.2 or later with the appropriate CNSA Suite algorithm supported by the solution. (MSC-DM-19)

5) Verify the procedure required by MSC-AU-25 is currently in place by the implementing organization and is followed.

6) Verify VPN Gateways log the failure to pull the CRL from the respective CDP/OCSP Responder. (MSC-AU-22)

   a) CDP Servers and OSCP Responders shall remove all CRLs.

   b) Connecting VPN Gateways shall attempt to pull the CRL from the respective CDP/OCSP Responder.

   c) Review the VPN Gateways audit logs to verify that a log report is generated from failure to pull the CRL.

7) Verify VPN Gateways log if the version of the CRL on the CDP/OCSP Responder is older than the current cached CRL. (MSC-AU-23)

   a) Load the CDP/OCSP Responder with CRLs that are older than the current cached CRLs on the VPN Gateway.

   b) Have the VPN Gateway attempt to pull the CRLs.

   c) Review the VPN Gateway audit logs to verify that a log report is generated.

8) Verify the VPN Gateway logs if signature validation of CRL on the CDP/OCSP Responder fails. (MSC-AU-24)

   a) Load the CDP/OCSP Responder with CRLs that contain invalid signature.

   b) Have the VPN Gateway pull the CRLs.

   c) Review the VPN Gateway audit logs to verify that a log report is generated due to an invalid CRL signature.

9) For all Solution Components, install appropriate certificates, generated by the approved CA, and configure the solution so that components use the certificates for authentication.

   a) Verify an entry to the audit log has been created due to certificate loading and generation (MSC-AU-8)

   b) Initiate a revocation of certificates for the Solution Components.

   c) Verify that an entry in the audit log has been created due to certificate revocation. (MSC-AU-8)

Expected Result:

For Step 1, all occurrences of auditable events given should generate an entry in the audit log. For Step 2, the source address should be the Encryption Component's loopback address. For Steps 3-4, all logs

forwarded on Red Management and Gray Management networks should be encrypted with the appropriate protocols. For Steps 5-7, there should be an audit log entry created for each activity. For Step 8, a log should be generated for generation and revocation of certificates.

15.10 KEY MANAGEMENT

This section contains procedures to ensure all CAs, certificates, KGCs and CAKs used in the MSC Solution are following the requirements in this CP.

15.10.1 CERTIFICATE AUTHORITIES AND CERTIFICATES

This section contains procedures to ensure all CAs and certificates used in the MSC Solution are following the requirements in this CP.

Requirements being tested: MSC-KM-1, MSC-KM-2, MSC-KM-5, MSC-KM-7, MSC-KM-9, MSC-KM-13, MSC-KM-14, MSC-KM-16 through MSC-KM-18, MSC-KM-24 through MSC-KM-34, MSC-KM-39, MSC-KM-41 through MSC-KM-45, MSC-KM-47 through MSC-KM-52, MSC-KM-54 through MSC-KM-56, and MSC-KM-61, MSC-AU-26 through MSC-AU-28, MSC-PF-21

Procedure Description:

1) Perform the following to validate the correct deployment of CAs:
    a) For CAs, verify the configuration of the MSC Solution to ensure that Outer CAs deliver services through either the Gray or Red networks and Inner CAs only deliver services through the Red network (MSC-KM-1, MSC-KM-2)
    b) For CAs, verify that the Outer and Inner CAs are physically separate from one another. (MSC-KM-18)
    c) For Locally-run CAs that operate on-line, verify the CAs use FIPS 140-2 Level 2 or higher HSMs to protect the CAs' private signing keys. (MSC-KM-24)
    d) For all CAs, verify that the CA does not have access to any Solution Components' private keys. (MSC-KM-17)
2) Perform the following to validate the structure of certificates issued by CAs:
    a) Obtain a sample set of test certificates issued to Solution Components.
    b) Verify the Distinguished Name in each certificate identifies an NPE. (MSC-KM-30)
    c) Verify the Key Usage extension in the certificate only asserts "digitalSignature". (MSC-KM-14)

d) Verify the certificates are compliant with the data standard for Version 3 certificates defined in ITU-T Recommendation X.509. (MSC-KM-9)

e) Verify the certificates are compliant with IETF RFC 5280 profile requirements. (MSC-KM-31)

f) Verify the certificates are compliant with key sizes and algorithms specified in this CP. (MSC-KM-13)

3) Perform the following to validate correct policy implementation as it relates to CAs and certificates issued by the CAs:

a) Verify the CA has a Certificate Policy and a CPS in place that is compliant with IETF RFC 3647, and that the CA operates in accordance with the policy and CPS. (MSC-KM-29, MSC-AU-26, MSC-AU-27, MSC-AU-28)

b) Verify the Certificate Policy states device authentication certificates issued by the CA, along with corresponding private keys, are considered CUI, and user private keys are classified to the level determined by the AO. (MSC-KM-7)

4) Perform the following steps to validate the certificate issuance capability of the MSC Solution:

a) Verify that a physical environment is identified to initially load keys and certificates onto Solution Components, where the environment is certified to protect information at the highest classification level of the Red network. (MSC-KM-25)

b) Verify that the key and certificate provisioning for Solution Components ensures private keys are never escrowed. (MSC-KM-16)

c) Generate a public/private key pair for the Outer VPN Gateway that complies with the key size and algorithm requirements in this CP.

d) Generate a certificate request for the Outer VPN Gateway, and ensure the request complies with PKCS#10.

e) Submit the certificate request to the Outer CA and verify that the CA returns a signed certificate request using PKCS#7. (MSC-KM-26)

f) Repeat steps 4c through 4e for each Inner VPN Gateway.

g) If the MSC Solution supports IETF RFC 7030 (EST), verify that the certificate request, response and installation process complies with EST. (MSC-KM-27)

h) Verify that the certificate request and issuance processes comply with the Outer and Inner CAs' Certificate Policies and CPSs. (MSC-KM-28, MSC-KM-29)

i) For locally run CAs, verify the Certificate Policies and CPSs to ensure certificates issued by the CAs: 1) enforce unique DNs; 2) assert key usages as defined by MSC-KM-5; and 3) assert a registered policy OID.

j) For locally run CAs, examine the contents of a sample set of issued certificates to ensure that the certificates assert: 1) unique DNs; 2) key usages as defined by MSC-KM-5; and 3) a registered policy OID. (MSC-KM-30)

k) For all CAs, examine the contents of a sample set of issued certificates and ensure that at least one valid CDP is asserted in the CDP extension of the certificates. (MSC-KM-31)

l) For locally run CAs, ensure the validity periods asserted in certificates issued by the CAs do not exceed 36 months for Solution Components. (MSC-KM-32)

m) Verify the Inner CAs can only issue certificates to Inner VPN Gateways and Red network components. (MSC-KM-33)

n) Verify the Outer CAs can only issue certificates to Outer VPN Gateways and Gray network components. (MSC-KM-34)

5) Perform the following steps to validate the certificate renewal and rekey capability:

a) Verify that the certificate renew and rekey processes comply with the Outer and Inner CAs' Certificate Policies and CPSs. (MSC-KM-41)

b) Verify that the Outer and Inner CAs' Certificate Policies and CPSs require certificate renew and rekey be performed prior to a certificate expiring, and require Solution Components go through the initial certificate issuance process if the certificate is expired. (MSC-KM-39)

c) Generate a new public/private key pair for the Outer VPN Gateway that complies with the key size and algorithm requirements in this CP.

d) Generate a certificate renew and rekey request for the Outer VPN Gateway, and ensure the request complies with PKCS#10.

e) Submit the certificate request to the Outer CA, and verify that the CA returns a signed certificate using PKCS#7. (MSC-KM-42)

f) Repeat steps 5c through 5e for each Inner VPN Gateway.

g) If the MSC Solution supports IETF RFC 7030 (EST), verify the certificate renew and rekey request, response and installation process complies with EST. (MSC-KM-43)

6) Perform the following steps to validate the certificate revocation and CDP capabilities:

a) Verify the Outer and Inner CAs' Certificate Policies and CPSs define requirements and procedures for revoking certificates, where certificate revocation is required when the binding between the subject information and public key within the certificate is no longer considered valid. (MSC-KM-45)

b) Verify the Outer and Inner CAs' Certificate Policies and CPSs define requirements and procedures for requesting the revocation of certificates. (MSC-KM-47)

c) For locally run CAs, verify the Outer and Inner CAs' Certificate Policies and CPSs define certificate revocation requirements and procedures that address: 1) removal of a revoked Solution Component from the MSC Solution; and 2) re-establishment of a Solution Component after certificate revocation is performed. (MSC-KM-48)

d) Verify that the Outer and Inner CAs have the capability to generate CRLs after certificate revocation functions are performed. (MSC-KM-51)

e) Obtain CRLs from the Outer and Inner CAs and ensure their structures are compliant with the data standard for Version 2 CRLs defined in ITU-T Recommendation X.509, and with the CRL profile standard defined by IETF RFC 5280. (MSC-KM-44)

f) Obtain CRLs from the Outer and Inner CAs and upload them onto the CDPs.

g) Verify that Solution Components can access the CDPs and download the CRLs issued by the Outer and Inner CAs via HTTP. (MSC-KM-61, MSC-PF-21)

h) For Enterprise CAs, verify that the Certificate Policies and CPSs define requirements and procedures for publishing CRLs. (MSC-KM-49)

i) For locally run CAs, verify that the Certificate Policies and CPSs define requirements and procedures for: 1) publishing new CRLs at least once every 28 days; 2) creating a new CRL within one hour of a certificate being revoked; and 3) publishing a newly created CRL at least 7 days before the expiration of the current CRL. (MSC-KM-50, MSC-KM-51, MSC-KM-54)

j) Verify the MSC Solution has procedures defined to transfer new CRLs to CDPs within 24 hours of the CRLs being created. (MSC-KM-52)

k) For MSC Solutions that support OCSP to provide certificate revocation status information, verify the OCSP Responders are deployed on the Gray and Red networks to deliver OCSP responses in accordance with IETF RFC 6960. (MSC-KM-55)

l) Generate an OCSP request from the connecting Outer VPN Gateway and send the request to the OCSP Responder operating in the Black network.

m) Generate an OCSP response from the OCSP Server in the Black network and deliver it to the connecting Outer VPN Gateway.

n) Examine the OCSP response and verify that it is digitally signed and compliant with IETF RFC 6960. (MSC-KM-56)

Expected Results:

CAs, CDPs, OCSP Responders and certificates are correctly deployed and in compliance with requirements in this CP.

15.10.2 Key Generation Components and Connectivity Association Keys

This section contains procedures for ensuring the KGCs and CAKs used within the solution comply with the requirements of this CP.

Requirements being tested: MSC-KM-1, MSC-KM-2, MSC-KM-4, MSC-KM-5, MSC-KM-10, MSC-KM-11, MSC-KM-15, MSC-KM-20, MSC-KM-21, MSC-KM-72, MSC-KM-74 through MSC-KM-83, MSC-KM-85 through MSC-KM-92, MSC-KM-94, MSC-KM-95

Procedure Description:

1) Perform the following to validate the correct deployment of CAK Management Services and KGCs:

   a) For CAK Management Services (including KGCs), verify the configuration of the MSC Solution to ensure that Outer CAK Management Services are provided through either the Gray or Red networks and Inner CAK Management Services are provided by the Red network. (MSC-KM-1, MSC-KM-2)

   b) If CAK Management Services operate at the same classification level as a Red network, verify a Controlled Interface is used to control information flow between the CAK Management Services and the Red network. (MSC-KM-4)

   c) If CAK Management Services operate at a different classification level than a Red network or Gray network, verify a CDS Controlled Interface is used to control information flow between the CAK Management Services and the Red or Gray networks. (MSC-KM-5)

   d) Verify KGCs and the MSC Solution are operating in compliance with an NSA-approved Key Management Plan. (MSC-KM-20)

   e) Verify an IAD MD 110 waiver was obtained, if necessary. (MSC-KM-21)

   f) Verify CAKs are generated by a NSA-approved KGC. (MSC-KM-72)

g) Verify an Inner KGC generates CAKs for Inner MACsec Devices or Outer MACsec Devices. (MSC-KM-74)

h) Verify an Outer KGC only generates CAKs for Outer MACsec Devices. (MSC-KM-76)

i) If an Inner KGC generates a CAK for Outer MACsec Devices, verify an AO-approved one-way transfer mechanism is used. (MSC-KM-75)

j) Confirm that CAKs are 256 bits. (MSC-KM-77)

k) Confirm that CAKs are not exposed in plaintext form until they are ready to be installed onto MACsec Devices. (MSC-KM-78)

l) Verify the KMP identifies the methods for protecting CAKs from unauthorized disclosure while being distributed and at any time while exported from MACsec Devices, and confirm the policies and procedures are followed. (MSC-KM-79, MSC-KM-80, MSC-KM-85, MSC-KM-86)

m) If PEKs are used to protect CAKs, confirm that PEKs are 256 bits and are protected to the same classification as the CAKs being encrypted. (MSC-KM-81, MSC-KM-82)

n) Confirm that CAKs are only used with the MACsec protocol. (MSC-KM-83)

o) Confirm the KMP states that compromised CAKs and PEKs are never to be used in the MSC Solution. (MSC-KM-87)

2) Perform the following to validate the updating of CAKs and PEKs:

a) Verify the KMP states that the same CAK may only be used by a pair of MACsec Devices establishing an encrypted tunnel. (MSC-KM-88)

b) Log into a MACsec Device and verify the CAK is set to expire in 30 days. (MSC-KM-89)

c) Repeat Step 2a for every MACsec Device.

d) Verify that the KMP states how often PEKs are to be updated and confirm the process and procedures are followed. (MSC-KM-90)

3) Perform the following to validate CAK compromise and recovery:

a) Confirm the KMP documents the CAK/PEK compromise recovery process, including removal of compromised MACsec Devices and re-establishing MACsec Devices after its CAK is compromised. (MSC-KM-91)

b) Confirm the KMP documents accounting procedures to support CAK and PEK compromise recovery. (MSC-KM-92)

c) Confirm the KMP documents procedures for compromise notification and reporting. (MSC-KM-94, MSC-KM-95)

Expected Results:

KGCs, CAKs and PEKs are correctly deployed and in compliance with requirements in this CP.

15.11 GRAY NETWORK FIREWALL

This section contains procedures for ensuring that the placement of Gray Network Firewalls within the solution comply with the requirements of this CP.

15.11.1 GRAY NETWORK FIREWALL FILTERING RULES

This section contains a procedure for ensuring that the filtering rules on Gray Network Firewalls are configured so that the only encrypted traffic allowed through the firewall is between Inner Encryption Components that are allowed to establish encryption tunnels with one another.

Requirements being tested: MSC-AU-10, MSC-FW-1, MSC-FW-6, MSC-FW-9

Procedure Description:

1) For each Inner Encryption Component within the solution (hereafter referred to as Inner Encryption Component A):

   a) For each other Inner Encryption Component within the solution that protects a Red network of the same security level:

      i) Attempt to establish an IPsec/MACsec connection to it from Inner Encryption Component A.

      ii) Verify that the IPsec/MACsec connection was established. (MSC-FW-1, MSC-FW-6)

   b) For each Inner Encryption Component within the solution that protects a Red network of a different security level (hereafter referred to as Inner Encryption Component B):

      i) Identify the first Gray Network Firewall on the physical path from Inner Encryption Component A and Inner Encryption Component B. If no such Firewall exists, skip the remainder of Step 1(b).

      ii) Place a packet sniffer on the interface of the Gray Network Firewall facing Inner Encryption Component B.

      iii) Attempt to establish an IPsec VPN connection from Inner Encryption Component A to Inner Encryption Component B.

      iv) Verify that the IPsec/MACsec connection was not established.

v) Verify that the packet sniffer did not record any IKE, IPsec, EAP-TLS or MACsec packets with a source address of Inner Encryption Component A and a destination address of Inner Encryption Component B. (MSC-FW-9)

vi) Verify that the Gray Network Firewall logs contain an event for an IKE, IPsec, EAP-TLS or MACsec packet with a source address of Inner Encryption Component A and a destination address of Inner Encryption Component B. (MSC-AU-10)

Expected Result:

In Step 1(a), the Gray Network Firewall allows IKE, IPsec, EAP-TLS, and MACsec traffic between pairs of Inner Encryption Component that are allowed to establish encryption tunnels with one another. In Step 1(b), the Gray Network Firewall denies IKE, IPsec, EAP-TLS, and MACsec traffic between pairs of Inner Encryption Component that protect Red networks of different classification levels.

15.11.2 GRAY NETWORK FIREWALL HTTP FILTERING RULES

This section contains a procedure for ensuring that the filtering rules on Gray Network Firewalls are configured so that the only HTTP traffic allowed through the firewall is from an Inner VPN Gateway to an Inner CDP.

Requirements being tested: MSC-AU-10, MSC-FW-2, MSC-FW-9

Procedure Description:

1) For each Inner VPN Gateway within the solution:

   a) For each Inner CDP within the solution:

      i) Attempt to have the Inner VPN Gateway download the current CRL from an Inner CDP.

      ii) Verify that the download was successful. (MSC-FW-2)

      iii) Identify the first Gray Network Firewall on the physical path from the Inner VPN Gateway to the Inner CDP. If no such Gray Network Firewall exists, skip the remainder of Step 1(a).

      iv) Place a packet sniffer on the interface of the Gray Network Firewall facing the Inner VPN Gateway.

      v) From the Inner CDP, attempt to make an HTTP request to the Inner VPN Gateway.

      vi) Verify that the request failed.

      vii) Verify that the packet sniffer did not record any packets with a source address of the Inner CDP and a destination address of the Inner VPN Gateway. (MSC-FW-9)

viii) Verify that the Gray Network Firewall logs contain an event for a packet with a source address of the Inner CDP and a destination address of the Inner VPN Gateway. (MSC-AU-10)

b) For every other device on the Gray network that is not an Inner CDP:

i) Identify the first Gray Network Firewall on the physical path from the Inner VPN Gateway to the other device. If no such Gray Network Firewall exists, skip the remainder of Step 1(b).

ii) Place a packet sniffer to the interface of the Gray Network Firewall facing the other device.

iii) Attempt to have the Inner VPN Gateway download a CRL from the other device.

iv) Verify that the download failed.

v) Verify that the packet sniffer did not record any packets with a source address of the Inner VPN Gateway and a destination address of the other device. (MSC-FW-8)

vi) Verify that the Gray Network Firewall logs contain an event for a packet with a source address of the Inner VPN Gateway and a destination address of the other device. (MSC-AU-10)

Expected Results:

In Step 1(a), the Gray Network Firewall allows HTTP requests from the Inner VPN Gateway to the Inner CDP, but not from the Inner CDP to the Inner VPN Gateway. In Step 1(b), the Gray Network Firewall denies HTTP requests from the Inner VPN Gateway to devices that are not Inner CDPs.

15.11.3 GRAY NETWORK FIREWALL MANAGEMENT

This section contains a procedure for ensuring that Gray Network Firewalls can only be managed from the Administration Workstation on the Gray Management network.

Requirements being tested: MSC-AU-10, MSC-FW-5

Procedure Description:

1) For each Gray Network Firewall within the solution:

a) From the Administration Workstation on the Gray Management network, attempt to connect to the Gray Network Firewall's remote management interface.

b) Verify that the connection attempt was successful. (MSC-FW-5)

c) For each physical network interface on the Gray Network Firewall except the one through which the Administration Workstation connects:

i) From a device reachable from the physical network interface, attempt to connect to the Gray Network Firewall's remote management interface.

ii) Verify that the connection attempt failed. (MSC-FW-5)

iii) Verify that the Gray Network Firewall logs contain an event for a packet with a source address of the selected device and a destination address of the Gray Network Firewall. (MSC-AU-10)

Expected Results:

In Step 1(b), the Gray Network Firewall allows management traffic from the Administration Workstation. In Step 1(c), the Gray Network Firewall blocks attempts to access the management interface through other physical network interfaces.

15.11.4 GRAY NETWORK FIREWALL ADDRESS SPOOFING

This section contains a procedure for ensuring that Gray Network Firewalls detect spoofing of source addresses in traffic sent through it.

Requirements being tested: MSC-AU-10, MSC-FW-7

Procedure Description:

1) For each Gray Network Firewall within the solution:

a) For each physical network interface on the Gray Network Firewall:

i) Select a device on the network connected to that interface of the Gray Network Firewall. Hereafter the device will be called Device A.

ii) Select a device on the network connected to an interface of the Gray Network Firewall that Device A is not connected to. Hereafter the device will be called Device B.

iii) Select a device on the network connected to an interface of the Gray Network Firewall that Device A is not connected to, and that Device B is allowed to communicate with. Hereafter the device will be called Device C.

iv) Place a network sniffer between the Gray Network Firewall and Device C.

v) Configure Device A to use the IP address of Device B.

vi) Attempt to send traffic from Device A (spoofing Device B's IP address) to Device C, of a type that Device B is allowed to send to Device C.

vii) Verify that the packet sniffer did not observe any packets with a source address of Device B and a destination address of Device C. (MSC-FW-7)

viii) Verify that the Gray Network Firewall logs contain an event for a packet received on the physical interface through which Device A connects, with a source address of Device B and a destination address of Device C. (MSC-AU-10)

Expected Results:

Each Gray Network Firewall detects the use of spoofed addresses and does not allow packets with spoofed source addresses from passing through, even if non-spoofed traffic from that source address would be allowed.

15.11.5 GRAY NETWORK FIREWALL HTTP DEEP PACKET INSPECTION

This section contains a procedure for ensuring that the deep packet inspection performed by the Gray Network Firewalls is configured so only the specific types of HTTP traffic desired between Inner VPN Gateways and Inner CDPs is allowed.

Requirements being tested: MSC-AU-10, MSC-FW-3, MSC-FW-4, MSC-FW-9

Procedure Description:

1) For each Inner VPN Gateway within the solution:

   a) For each Inner CDP within the solution:

      i) Attempt to have the Inner VPN Gateway download the current CRL from the Inner CDP.

      ii) Verify that the download was successful. (MSC-FW-3, MSC-FW-4)

      iii) Identify the first Gray Network Firewall on the physical path from the Inner VPN Gateway to the Inner CDP. If no such Gray Network Firewall exists, skip the remainder of Step 1(a).

      iv) Replace the CRL on the Inner CDP with a text file.

      v) Place a packet sniffer on the interface of the Gray Network Firewall facing the Inner VPN Gateway.

      vi) Attempt to have the Inner VPN Gateway download the current CRL from the Inner CDP.

      vii) Verify that the request failed.

      viii) Verify that the packet sniffer did not record any packets with a source address of the Inner CDP and a destination address of the Inner VPN Gateway that contains an HTTP response payload. (MSC-FW-4)

ix) Verify that the Gray Network Firewall logs contain an event for an improper HTTP response payload with a source address of the Inner CDP and a destination address of the Inner VPN Gateway. (MSC-AU-10)

x) Restore the CRL on the Inner CDP.

xi) Move the packet sniffer to the interface of the Gray Network Firewall facing the Inner CDP.

xii) Attempt to have the Inner VPN Gateway download a CRL from the Inner CDP using an incorrect URL.

xiii) Verify that the request failed.

xiv) Verify that the packet sniffer did not record any packets with a source address of the Inner VPN Gateway and a destination address of the Inner CDP. (MSC-FW-3)

xv) Verify that the Gray Network Firewall logs contain an event for an improper HTTP request with a source address of the Inner VPN Gateway and a destination address of the Inner CDP. (MSC-AU-10)

xvi) Attempt to have the Inner VPN Gateway issue an HTTP POST request for the URL of the CRL on the Inner CDP.

xvii) Verify that the request failed.

xviii) Verify that the packet sniffer did not record any packets with a source address of the Inner VPN Gateway and a destination address of the Inner CDP. (MSC-FW-3)

xix) Verify that the Gray Network Firewall logs contain an event for an improper HTTP request with a source address of the Inner VPN Gateway and a destination address of the Inner CDP. (MSC-AU-10)

b) For every other device on the Gray network that is not an Inner CDP, follow the procedure for Step 1(b). (MSC-AU-10, MSC-FW-9)

Expected Results:

In Step 1(a), the Gray Network Firewall allows only HTTP traffic between the Inner VPN Gateway and Inner CDP that consists of a GET request for the appropriate CRL and a response containing the CRL. In Step 1(b), the Gray Network Firewall denies HTTP requests from the Inner VPN Gateway to devices that are not Inner CDPs.

15.12 Incident Reporting Guidance

This section ensures that procedures are followed regarding incident reporting to NSA in the event a solution owner identifies a security incident which affects the solution.

Requirements being tested: MSC-RP-1 through MSC-RP-14

Procedure Description:

1) Verify the procedures given in MSC-RP-1 through MSC-RP-14 were/are followed and are currently in place.

Expected Results:

For Step 1, all of these procedures have been followed or are in place.

15.13 Implementation of Guidance

This section ensures there are procedures in place and/or that procedures were followed regarding the procurement of products and use of the MSC Solution. It also ensures the personnel are in place to manage and administer this solution following the guidelines given in this CP.

Requirements being tested: MSC-GD-1 through MSC-GD-35, MSC-SR-10

Procedure Description:

1) Verify the procedures for obtaining virus signature updates as required by local agency policy and the AO were/are followed and/or are in place. (MSC-SR-10)

2) Verify the procedures given in MSC-GD-1 through MSC-GD-4, and MSC-GD-13 through MSC-GD-25 were/are followed and/or are currently in place.

3) Verify the solution owner understands that he/she shall allow and fully cooperate with an NSA-ordered IA compliance audit of this solution implementation. (MSC-GD-5)

4) Verify the solution owner and AO are aware that a compliance audit will be conducted every year. (MSC-GD-6)

5) Verify the AO is aware that they shall receive the results of the compliance audit and are responsible for reviewing the completed audit. (MSC-GD-7)

6) Verify the customer is aware that when they are interested in registering their solution against this CP, that NSA must grant them an approval prior to the AO authorizing the solution for operation. (MSC-GD-8)

7) Verify the customer completes and submits the compliance matrix to the AO. (MSC-GD-9)

8) Verify the customer is aware that registration and re-registration against this CP includes submission of CP registration forms and compliance matrix to NSA. (MSC-GD-10)

9) Verify the solution owner and AO are aware that when new versions of the MSC CP are published by NSA they will have 6 months to bring their solution into compliance with the new version. (MSC-GD-11)

10) Verify the solution owner and AO are aware that they shall provide updated solution information to NSA on a yearly basis. (MSC-GD-12)

11) Verify the personnel requirements given in MSC-GD-26 through MSC-GD-35 are met by the personnel supporting this implementation of the MSC Solution.

Expected Result:

For Steps 1-10, all of these procedures have been followed or are in place. For Step 11, assigned personnel meet the stated personnel requirements.

15.14 SOLUTION FUNCTIONALITY

This section contains a procedure for ensuring the implementing organization complies with the testing requirements.

Requirements being tested: MSC-TR-1

Procedure Description:

1) The implementing organization's AO will inspect the test report to ensure all testing requirements have been met. (MSC-TR-1)

Expected Result:

The report will ensure the implementing organization complies with this CP.

APPENDIX A. GLOSSARY OF TERMS

Accreditation – The official management decision given by a senior agency official to authorize operation of an information system and to explicitly accept the risk to agency operations (including mission, functions, image, or reputation), agency assets, or individuals, based on the implementation of an agreed-upon set of security controls. (NIST 800-37)

Assurance – Measure of confidence that the security features, practices, procedures, and architecture of an information system accurately mediates and enforces the security policy. (CNSSI 4009)

Audit – The activity of monitoring the operation of a product from within the product. It includes monitoring of a product for a set of pre-determined events. Each audit event may indicate rogue behavior, or a condition that is detrimental to security, or provide necessary forensics to identify the source of rogue behavior.

Audit Log – A chronological record of the audit events that have been deemed critical to security. The audit log can be used to identify potentially malicious activity that may further identify the source of an attack, as well as potential vulnerabilities where additional countermeasures or corrective actions are required.

Availability – Assurance that the system and its associated assets are accessible and protected against Denial of Service attacks, as well as available when the user needs them and in the form needed by the user.

Black Box Testing – Testing the functionality of a component of the solution, such that testing is limited to the subset of functionality that is available from the external interfaces of the box during its normal operational configuration without any additional privileges (such as given to the Security Administrator or Auditor).

Black Network – A network that contains classified data that has been encrypted twice.

Capability Package – The set of guidance provided by NSA that describes recommended approaches to composing COTS components to protect classified information for a particular class of security problem. This package will point to potential products that can be used as part of this solution.

Central Management Site – A site within a MSC Solution that is responsible for remotely managing the Solution Components located at other sites.

Certification – The technical evaluation of a system's security features, performed as a part of and in support of the approval/accreditation process that establishes the extent to which a particular computer system's design and implementation meet a set of specified security requirements.

Certification and Accreditation (C&A) – A comprehensive assessment of the management, operational, and technical security controls in an information system, made in support of security accreditation, to determine the extent to which the controls are implemented correctly, operating as intended, and producing the desired outcome with respect to meeting the security requirements for the system. In conjunction with the official management decision given by a senior agency official to authorize operation of an information system and to explicitly accept the risk to agency operations (including mission, functions, image, or reputation), agency assets, or individuals, based on the implementation of an agreed-upon set of security controls. (NIST 800-37).

Certificate Authority (CA) – An authority trusted by one or more users to create and assign certificates. [ISO9594-8]

Certificate Policy – A named set of rules that indicate the applicability of a certificate to a particular community and/or class of application with common security requirements. For example, a particular Certificate Policy might indicate applicability of a type of certificate to the authentication of parties engaging in business-to-business transactions for the trading of goods or services within a given price range. [IETF RFC 3647]

Committee on National Security Systems Policy No. 15 (CNSSP-15) – Policy specifies which public standards may be used for cryptographic protocol and algorithm interoperability to protect National Security Systems (NSS).

Confidentiality – Assurance that the data stored in, processed by, or transmitted by the system are protected against unauthorized disclosure, and confidence that only the appropriate set of individuals or organizations would be provided the information.

Control Plane Protocol – A routing, signaling, or similar protocol whose endpoints are network infrastructure devices such as VPN Gateways or routers. Control plane protocols carry neither user data nor management traffic.

CRL Distribution Point (CDP) – A web server that hosts a copy of a CRL issued by a CA for VPN Gateways to download.

Cross Domain Solution (CDS) – A form of controlled interface that provides the ability to manually and/or automatically access and/or transfer information between different security domains. [CNSSI 4009]

Data Plane Protocol – A protocol that carries the data being transferred through the solution.

Encryption Component – Either a VPN Gateway or a MACsec Device.

External Interface – The interface on an Encryption Component that connects to the outer network (i.e., the Gray network on the Inner Encryption Component or the Black network on the Outer Encryption Component).

Federal Information Processing Standards (FIPS) – A set of standards that describe the handling and processing of information within governmental agencies.

Gray Box Testing – The ability to test functionality within a component of the solution, such that full management privileges are granted (i.e., knowing passwords for security administrator and Auditor and access to the capabilities associated with those privileges). In addition, the use of any and all testing equipment and/or testing software used inside and outside the developed solution is available.

Gray Network – A network that contains classified data that has been encrypted once.

Gray Network Firewall – A stateful traffic filtering firewall placed on the Gray network to provide additional separation between flows of singly-encrypted data of different classification levels.

Independently Managed Site – A site within a MSC Solution whose Solution Components are locally managed and that does not remotely manage other sites' Solution Components.

Internal Interface – The interface on an Encryption Component that connects to the inner network (i.e., the Gray network on the Outer Encryption Component or the Red network on the Inner Encryption Component).

Locally Managed Device – A device that is being managed by the direct connection of the Administration Workstation to the device in a hardwired fashion (such as a console cable).

Malicious – Any unauthorized events that are either unexplained or in any way indicate adversary activity.

Management Plane Protocol – A protocol that carries either traffic between a system administrator and a component being managed, or log messages from a Solution Component to a log server or similar repository.

Protection Profile – A document used as part of the certification process according to the Common Criteria. As the generic form of a security target, it is typically created by a user or user community and provides an implementation independent specification of information assurance security requirements.

Public Key Infrastructure (PKI) – Framework established to issue, maintain, and revoke public key certificates.

Red Network – A network that contains unencrypted classified data.

Remotely Managed Device – A device that is being managed by any other method besides that given in the definition of a Locally Managed Device.

Remote Site – A site within a MSC Solution whose Solution Components are remotely managed by a Central Management Site.

Security Level – The combination of classification level, list of compartments, dissemination controls, and other controls applied to the information within a network.

APPENDIX B. ACRONYMS

Acronym	Definition
ACL	Access Control List
AES	Advanced Encryption Standard
AM	Advisory Memorandum
AO	Authorizing Official
ARP	Address Resolution Protocol
BFD	Bidirectional Forwarding Detection
BGP	Border Gateway Protocol
C&A	Certification and Accreditation
CA	Certificate Authority
CAA	Certificate Authority Administrator
CAK	Connectivity Association Key
CAKA	Connectivity Association Key Administrator
CDP	CRL Distribution Point
CDS	Cross Domain Solution
CKN	Connectivity Association Key Name
CMAC	Cipher-based Message Authentication Code
CNSA	Commercial National Security Algorithm [Suite]
CNSS	Committee on National Security Systems
CNSSI	Committee on National Security Systems Instruction
CNSSP	Committee on National Security Systems Policy
COTS	Commercial Off-the-Shelf
CP	Capability Package
CP	Certificate Policy
CPS	Certification Practice Statement
CRL	Certificate Revocation List
CSfC	Commercial Solutions for Classified
CUI	Controlled Unclassified Information
DDoS	Distributed Denial of Service
DH	Diffie-Hellman
DHCP	Dynamic Host Configuration Protocol
DM	Device Management
DNS	Domain Name System
DoD	Department of Defense
DoS	Denial of Service
DSA	Digital Signature Algorithm
EAP	Extensible Authentication Protocol
EAP-TLS	Extensible Authentication Protocol-Transport Layer Security
ECDH	Elliptic Curve Diffie-Hellman
ECDSA	Elliptic Curve Digital Signature Algorithm
EoMPLS	Ethernet over Multiprotocol Label Switching

Acronym	Definition
ESP	Encapsulating Security Payload
EST	Enrollment Over Secure Transport
FIPS	Federal Information Processing Standards
GCM	Galois Counter Mode
GOTS	Government Off-the-Shelf
GRE	Generic Routing Encapsulation
HMAC	Host-based Message Authentication Code
HSM	Hardware Security Module
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure
IA	Information Assurance
IAD	Information Assurance Directorate
IAVA	Information Assurance Vulnerability Alerts
IC	Intelligence Community
ICD	Intelligence Community Directive
ICMP	Internet Control Message Protocol
ICT	Information and Communication Technology
ID	Identification
IDS	Intrusion Detection System
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IKE	Internet Key Exchange
IP	Internet Protocol
IPS	Intrusion Prevention System
IPsec	Internet Protocol Security
IPv4	Internet Protocol Version 4
IPv6	Internet Protocol Version 6
IS-IS	Intermediate System to Intermediate System
JIMS	Joint Incident Management System
KGC	Key Generation Component
KM	Key Management
KMI	Key Management Infrastructure
L2TPv3	Layer 2 Tunneling Protocol Version 3
MACsec	Media Access Control Security
MGC	Management Client
MSC	Multi-Site Connectivity
MLD	Multicast Listener Discovery
MTU	Maximum Transmission Unit
NIAP	National Information Assurance Partnership
NIST	National Institute of Standards and Technology
NPE	Non-Person Entity
NSA	National Security Agency

Acronym	Definition
NSS	National Security Systems
NTP	Network Time Protocol
O	Objective
OCSP	Online Certificate Status Protocol
OID	Object Identifier
OS	Operating System
OSI	Open Systems Interconnection
OSPF	Open Shortest Path First
PKCS	Public Key Cryptographic Standard
PKI	Public Key Infrastructure
PMTU	Path Maximum Transmission Unit
PSK	Pre-Shared Key
PTP	Precision Time Protocol
QoS	Quality of Service
RFC	Request for Comment
RIP	Routing Information Protocol
RSA	Rivest Shamir Adelman algorithm
S3	Secure Sharing Suite
SA	Security Association
SAK	Secure Association Key
SCRM	Supply Chain Risk Management
SHA	Secure Hash Algorithm
SIEM	Security Information and Event Management
SIPRNet	Secret Internet Protocol Router Network
SP	Service Packs
SSH	Secure Shell
SSHv2	Secure Shell Version 2
T	Threshold
T&E	Test and Evaluation
TCP	Transmission Control Protocol
TFFW	Traffic Filtering Firewall
TLS	Transport Layer Security
UDP	User Datagram Protocol
VLAN	Virtual Local Area Network
VPN	Virtual Private Network
XPN	eXtended Packet Number

APPENDIX C. REFERENCES

CNSSI 1300	*CNSSI 1300, National Security Systems Public Key Infrastructure X.509 Certificate Policy*	December 2014
CNSSI 4009	*CNSSI 4009, National Information Assurance (IA) Glossary Committee for National Security Systems.* http://www.cnss.gov/Assets/pdf/cnssi_4009.pdf	April 2015
CNSSP 11	*CNSS Policy (CNSSP) Number 11, National Policy Governing the Acquisition of Information Assurance (IA) and IA-Enabled Information Technology Products*.	June 2013
CNSSP 15	*CNSS Policy (CNSSP) Number 15, National Information Assurance Policy on the Use of Public Standards for the Secure Sharing of Information Among National Security Systems Committee for National Security Systems*	October 2012
CNSS AM IA 02-15	*Use of Public Standards for the Secure Sharing of Information Among National Security Systems*	July 2015
CNSSD 505	*CNSS Directive (CNSSD) Number 505, Supply Chain Risk Management (SCRM)*	March 2012
FIPS 140-2	*Federal Information Processing Standard 140, Security Requirements For Cryptographic Modules National Institute for Standards and Technology FIPS Publication* http://csrc.nist.gov/publications/fips/fips140-2/fips1402.pdf	May 2001
FIPS 180-4	*Federal Information Processing Standard 180-4, Secure Hash Standard (SHS)*	August 2015
FIPS 186	*Federal Information Processing Standard 186-4, Digital Signature Standard (DSS)*	July 2013
FIPS 197	*Federal Information Processing Standard 197, Advanced Encryption Standard (AES)*	November 2001
FIPS 201-1	*Federal Information Processing Standard 201, Personal Identity Verification (PIV) of Federal Employees and Contractors National Institute for Standards and Technology FIPS Publication* http://csrc.nist.gov/publications/fips/fips201-1/FIPS-201-1-chng1.pdf	June 2006
RFC 2409	*IETF RFC 2409 The Internet Key Exchange (IKE)*. D. Harkins and D. Carrel.	November 1998
RFC 3647	*IETF RFC 3647 Internet X.509 Public Key Infrastructure Certificate Policy and Certification Practices Framework* Internet Engineering Task Force	November 2003
RFC 4252	*IETF RFC 4252 The Secure Shell (SSH) Authentication Protocol*. T. Ylonen and C. Lonvick.	January 2006
RFC 4253	*IETF RFC 4253 The Secure Shell (SSH) Transport Layer Protocol*. T. Ylonen and C. Lonvick.	January 2006

RFC 4254	*IETF RFC 4254 The Secure Shell (SSH) Connection Protocol*. T. Ylonen and C. Lonvick.	January 2006
RFC 4256	*IETF RFC 4256 Generic Message Exchange Authentication for the Secure Shell Protocol (SSH)*. F. Cusack and M. Forssen.	January 2006
RFC 4302	*IETF RFC 4302 IP Authentication Header*. S. Kent	December 2005
RFC 4303	*IETF RFC 4303 IP Encapsulating Security Payload*. S. Kent	December 2005
RFC 4307	*IETF RFC 4307 Cryptographic Algorithms for Use in the Internet Key Exchange Version 2 (IKEv2)*. J. Schiller	December 2005
RFC 4308	*IETF RFC 4308 Cryptographic Suites for IPsec*. P. Hoffman	December 2005
RFC 4754	*IETF RFC 4754 IKE and IKEv2 Authentication Using the Elliptic Curve Digital Signature Algorithm (ECDSA)*. D. Fu and J. Solinas.	January 2007
RFC 5246	*IETF RFC 5246 The Transport Layer Security (TLS) Protocol Version 1.2*. T. Dierks and E. Rescorla.	August 2008
RFC 5280	*IETF RFC 5280 Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*. D. Cooper, et. al.	May 2008
RFC 5759	*IETF RFC 5759 Suite B Certificate and Certificate Revocation List (CRL) Profile*. J. Solinas and L. Zieglar.	January 2010
RFC 5996	*IETF RFC 5996 Internet Key Exchange Protocol Version 2 (IKEv2)*. C. Kaufman, et. al.	September 2010
RFC 6239	*IETF RFC 6239 Suite B Cryptographic Suites for Secure Shell (SSH)*. K. Igoe.	May 2011
RFC 6379	*IETF RFC 6379 Suite B Cryptographic Suites for IPsec*. L. Law and J. Solinas.	October 2011
RFC 6380	*IETF RFC 6380 Suite B Profile for Internet Protocol Security (IPsec)*. K. Burgin and M. Peck.	October 2011
RFC 6460	*IETF RFC 6460 Suite B Profile for Transport Layer Security (TLS)*. M. Salter and R. Housley.	January 2012
RFC 6818	*IETF RFC 6818 Updates to the Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*. P. Yee	January 2013
RFC 7030	*IETF RFC 7030 Enrollment over Secure Transport*. M. Pritikin, P. Yee, and D. Harkins.	October 2013
SP 800-53	*NIST Special Publication 800-53 Rev. 4, Security and Privacy Controls for Federal Information Systems and Organizations*. Joint Task Force Transformation Initiative.	May 2013
SP 800-56A	*NIST Special Publication 800-56A Rev. 2, Recommendation for Pair-Wise Key Establishment Schemes Using Discrete Logarithm Cryptography*. E. Barker, et. al.	May 2013

SP 800-56B	*NIST Special Publication 800-56B Rev. 1, Recommendation for Pair-Wise Key Establishment Schemes Using Integer Factorization Cryptography*. E. Barker, et. al.	October 2014
SP 800-56C	*NIST Special Publication 800-56C, Recommendation for Key Derivation through Extraction-then-Expansion*. L. Chen.	November 2011
SP 800-131A	*NIST Special Publication 800-131A Rev. 1, Recommendation for Transitioning of Cryptographic Algorithms and Key Lengths*. E. Barker.	November 2015
SP 800-147	*NIST Special Publication 800-147, BIOS Protection Guidelines*. D. Cooper, et. al.	June 2011

Michael Franz

444 Information & Computer Science Building
University of California, Irvine
Irvine, CA 92697-3435
franz@uci.edu

Litigation Experience

Ongoing

- *Netskope, Inc. v. Fortinet, Inc.,.* Civil Action No. 3:22-cv-01852-TLT (N.D. California). Retained by *Perkins Coie, LLP,* San Diego, California., on behalf of *Netskope;* since November 2022.
- *CUPP Cybersecurity, LC, et al. v. Trend Micro, Inc.,* Civil Action Nos. 3:18-cv-01251 and 3:20-cv-03206 (N.D. Texas). Retained by *DLA Piper,* East Palo Alto, California, on behalf of defendant *Trend Micro, Inc.;* since December 2021. I was **deposed** in this case.

Completed

- *Netskope, Inc. v. Bitglass, Inc.,* Civil Action No. 3:21-cv-00916-EMC (N.D. California). Retained by *Perkins Coie, LLP,* San Diego, California., on behalf of *Netskope;* February 2021 – December 2022. Based on my expert testimony, the Patent Trial and Appeal Board confirmed that every single one of the challenged claims of the contested patents was unpatentable.
- Consulting services to *Orrick, Herrington & Sutcliffe LLP,* Menlo Park, California, in the matter *Express Mobile, Inc. v. eBay, Inc.,* Civil Action No. 6:20-cv-00802-ADA (W.D. Texas), on behalf of defendant *eBay, Inc.;* May 2021.
- *Uniloc v. Netflix, Inc., Hulu, LLC, and Roku, Inc.,* Civil Action No. 8:18-cv-02055-GW-DFM and 8:18-cv-02150-GW-DFM (C.D. California). Retained by *Perkins Coie, LLP,* San Diego, California., on behalf of defendants *Netflix, Hulu, and Roku.;* June 2019 – February 2021. The Patent Trial and Appeal Board granted institution of an Inter Partes Review based in part on my expert report (IPR2020-00041); the claims in question were ultimately invalidated.
- *Ameranth v. Starbucks Corp,* Civil Action No. 3:11-cv-01810-DMS-WVG (S.D. California). Retained by *Perkins Coie, LLP,* San Diego, California., on behalf of defendant *Starbucks.;* June 2017 - August 2018. I was **deposed** in this case. The case ended when the court granted a summary judgement invalidating the patent in suit.
- *Vehicle IP, LLC v. AT&T Mobility, LLC, Cellco Partnership, Networks in Motion, Inc., TeleCommunication Systems, Inc., and Telenav Inc.,* Civil Action No. 00-1007-LPS-CJB (D. Delaware). Retained by *Perkins Coie, LLP,* San Diego, California., on behalf of defendants *TeleCommunication Systems, Inc. and Networks In Motion, Inc.;* January 2016 – June 2017. I was **deposed** in this case. Defendants received a favorable disposition by the court.
- Consulting services to *Smith, Gambrell & Russell, LLP,* Washington, D.C., in the matter *Callwave Communications, LLC v. Verizon Services Corp., Cellco Partnership d/b/a Verizon Wireless and Google Inc.,* Civil Action No. 12-cv-1704-RGA (D. Delaware) on behalf of defendant *Verizon;* January – March 2016.
- *Finjan, Inc. v. Proofpoint, Inc. et al.,* Case No. 5:13-cv-05808-BLF (N.D. California). Retained by *Quinn Emanuel Urquhart & Sullivan, LLP,* San Francisco, California, on behalf of defendant *Proofpoint, Inc.;* September 2015 – May 2016. I was **deposed** in this case, after expert reports had been exchanged; the case settled before trial.

- *Enterprise v. Apple Inc.,* C.A. No. 14-00765 (D. Delaware) and "In re Communications or Computing Devices and Components Thereof," Investigation No. 337-TA-925 before the U.S. Patent and Trademark Office. Retained by *Ropes & Gray LLP,* New York, New York, on behalf of defendant *Apple, Inc.;* January – May 2015. The case was settled during initial work on my expert report.

- *REC Software USA, Inc. v. Huawei Devices USA, Inc., et al.,* Case No. 2:14-cv-01050 (W.D. Washington). Retained by *Lee Tran & Liang LLP,* Los Angeles, California, on behalf of defendant *Huawei, Inc.;* December 2014 – February 2015. The case was settled while I was working on my expert report.

- Consulting services to *Cooley LLP,* Broomfield, Colorado, in the matter *Finjan, Inc. v. Websense, Inc.,* Case No. 5:15-cv-04398 (N.D. California), on behalf of defendant *Websense, Inc.;* July 2014.

- *LBS Innovations LLC v. BP America, Inc. et al.,* Case No. 2:11-cv-00407-JRG (E.D. Texas). Retained by *Perkins Coie LLC,* San Diego, California, on behalf of defendant *Microsoft Corporation* and its customers *American Express, Choice Hotels International, Inc., Conoco Phillips, Costco Wholesale Corporation, Hyatt, JP Morgan Chase, KeyCorp, Sally Beauty Supply LLC, Starbucks Corporation, Target Corporation, The Kroger Co., U.S. Bancorp, and Wells Fargo Bank, N.A.;* August – November 2013. I was **deposed** in this case, after expert reports and rebuttal export reports had been exchanged; shortly afterwards, the case settled.

- *EMG Technology, LLC v. Google, Inc. et al.,* Case No. 6:12-cv-00439-MHS (E.D. Texas). Retained by *K&L Gates LLP,* Palo Alto, California, on behalf of defendant *Google, Inc.;* July 2013. The case unexpectedly settled while I was working on my expert report.

- *Porto Technology, Co., Ltd. et al. v. Cellco Partnership d/b/a Verizon Wireless,* Case No. 3:12-cv-678 (E.D. Virginia). Retained by *Wiley Rain LLP,* Washington, D.C., at the suggestion of and on behalf of defendant; December 2012 – March 2013. The case was dismissed while I was working on my expert report.

- *Openwave Systems Inc. v. Apple Inc. et al.,* "In the Matter of Certain Devices for Mobile Data Communication," U.S. International Trade Commission, Docket No. 337-TA-2843. Retained by *Irell & Manella LLP,* Los Angeles, California, on behalf of respondent (defendant) *Research In Motion, Inc.;* November 2011 – October 2012. The complaint was withdrawn by the complainant in October 2012, after expert reports and rebuttal export reports had been exchanged.

- Consulting services to *Morrison & Foerster LLP,* Los Angeles, California, in the matter *Tse v. Apple, Inc. et al.,* on behalf of defendant *Apple, Inc.;* June – August 2012.

- *Alfred B. Levine v. Samsung Telecommunications America, LLC et al.* Case No. 2:09-CV-372-MHS (E.D. Texas). Retained by *Morrison & Foerster LLP,* Los Angeles, California, on behalf of defendant *Cellco Partnership d/b/a Verizon Wireless;* April – September 2012. My technical expert report played a substantial role in reaching a settlement agreement that covers Verizon and that led to Verizon's dismissal from the lawsuit.

- Consulting services to *Nixon & Vanderhye,* Attorneys At Law, Arlington, Virginia, in a matter related to virtual machines; August – November 2009.

- *Vehicle IP, LLC v. General Motors Corporation et al.*; Civil Action No. 07-C-0345-S (W.D. Wisconsin). Retained by *Morrison & Foerster LLP,* Los Angeles, California, on behalf of defendants *Networks in Motion, Inc. and Cellco Partnership d/b/a Verizon Wireless*; September 2007 – April 2009. This case was dismissed in summary judgement partially based on my expert report. Because of the tight schedule of this case with overlapping actions, I was **deposed** before the summary judgement was passed down.

- *Newport Corp. v. WareNet, Inc.*, 2002 WL 1732452 (California App. 4th District). Retained by *Call & Jensen,* Newport Beach, California, on behalf of the defendant; January 2002 – January 2003. I was **deposed** in this case; shortly afterwards, the case settled.