

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

PALO ALTO NETWORKS, INC.,
Petitioner

v.

BT AMERICAS INC.,
Patent Owner

Case No. IPR2023-00889
U.S. Patent No. 7,895,641

**DECLARATION OF KEVIN JEFFAY, PH.D., IN SUPPORT OF PETITION
FOR *INTER PARTES* REVIEW OF U.S. PATENT NO. 7,895,641**

Mail Stop PATENT BOARD
Patent Trial and Appeal Board
U.S. Patent & Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450

PAN EX1003
Palo Alto Networks v. BT Americas
IPR2023-00889

TABLE OF CONTENTS

I.	Introduction	1
II.	Grounds of Unpatentability	3
III.	Qualifications.....	4
IV.	Materials Considered	8
V.	Legal Standards	10
	A. My Understanding of Claim Construction.....	10
	B. My Understanding of a Person of Ordinary Skill in the Art	11
	C. My Understanding of Obviousness	12
VI.	The '641 Patent.....	16
	A. Alleged Problem in the Prior Art	16
	B. Overview of the '641 Patent	18
	C. Prosecution History.....	23
	D. Previous <i>Inter Partes</i> Review	25
	E. Level of Ordinary Skill in the Art	26
	F. Claim Construction	27
VII.	Overview of the Prior Art	29
	A. Overview of Duvall	29
	B. Overview of Chu.....	37
	C. Overview of Trcka	39
	D. Overview of Ziese.....	43
	E. Overview of Cogger.....	46
	F. The prior art is analogous art.....	47
VIII.	Ground 1: Claims 1-7 and 15-17 are obvious over Duvall in view of Chu. ...	48
	A. A POSA would have been motivated to combine Duvall with Chu....	48
	B. Independent Claim 1	59
	1. [1-pre] “A system for operating a probe as part of a security monitoring system for a computer network”	60
	2. [1A] “a sensor coupled to collect status data from at least one monitored component of the network”	64

3.	[1B] “a filtering subsystem coupled to analyze status data to identify potentially security-related events represented in the status data, wherein the analysis includes filtering followed by an analysis of post-filtering residue, wherein the post-filtering residue is data neither discarded nor selected by filtering”	67
a)	Duvall discloses “a filtering subsystem coupled to analyze status data..., wherein the analysis includes filtering.”...	67
b)	Duvall discloses “an analysis of post-filtering residue, wherein the post-filtering residue is data neither discarded nor selected by filtering.”	70
c)	Duvall’s analysis is to “identify potentially security-related events represented in the status data.”	75
4.	[1C] “a communications system coupled to transmit information about the identified events to an analyst system associated with the security monitoring system”	77
5.	[1D] “a receiver for receiving feedback at the probe based on empirically-derived information reflecting operation of the security monitoring system”	81
6.	[1E] “a modification control system for dynamically modifying an analysis capability of the probe during operation thereof based on the received feedback”	83
C.	Claim 2	85
D.	Claim 3	86
E.	Claim 4	88
F.	Claim 5	88
G.	Claim 6	88
H.	Claim 7	93
I.	Claim 15	93
J.	Claim 16	95
K.	Claim 17	96
IX.	Ground 2: Claims 8-13, 16 are obvious over Duvall in view of Chu, and in further view of Trcka.	97
A.	A POSA would have been motivated to combine the Duvall-Chu system with Trcka.	97
B.	Claims 7, 8 and 12	99

C.	Claims 9 and 13	102
D.	Claim 10	103
E.	Claim 11	104
F.	Claim 16	105
X.	Ground 3: Claims 14-15 are obvious over Duvall in view of Chu and Trcka, and in further view of Ziese.	106
A.	A POSA would have been motivated to combine the Duvall-Chu-Trcka system with Ziese.	106
B.	Claim 14	110
C.	Claim 15	111
XI.	Ground 4: Claims 18-25 are obvious over Duvall in view of Chu, and in further view of Cogger.	112
A.	A POSA would have been motivated to combine the Duvall with Chu and Cogger.	112
B.	Independent Claim 18	117
1.	[18-pre] “A method of operating a secure operations center as part of a security monitoring system for a customer computer network”	117
2.	[18A] “creating an event record for information received about an identified potentially security-related event occurring on the network, wherein the potentially security-related event is identified by filtering followed by an analysis of post-filtering residue, wherein the post-filtering residue is neither discarded nor selected by the filtering”	119
3.	[18B] “correlating the event record with customer information and a symptom record”	120
4.	[18C] “using the correlated symptom record to link the event record to problem resolution information”	123
5.	[18D] “consolidating the event record, correlated customer information and symptom record, and linked problem resolution assistance information into a problem ticket”	125
6.	[18E] “providing the problem ticket to a security analyst console for analysis”	126
C.	Claim 19	127
D.	Claim 20	130

E. Claim 21 131

F. Claim 22 132

G. Claims 23 and 24 134

H. Claim 25 136

XII. Objective Indicia Do Not Support Patentability 137

XIII. Conclusion.....139

I. INTRODUCTION

I, Kevin Jeffay, Ph.D., declare as follows:

1. I am over the age of twenty-one (21) years, of sound mind, and capable of making the statements set forth in this Declaration. I am competent to testify about the matters set forth herein. All facts and statements contained herein within my personal knowledge are true and correct, and I believe those items stated herein on information and belief to be true and correct.

2. I have been retained on behalf of Palo Alto Networks, Inc. (“Petitioner”) for the above-captioned *inter partes* review proceeding to provide my expert opinions and expert knowledge. I understand that this proceeding involves U.S. Patent No. 7,895,641 (“the ’641 patent,” EX1001), which is titled “Method and system for dynamic network intrusion monitoring, detection and response” and is currently assigned to BT Americas Inc.

3. The ’641 patent discloses a system that “filters and analyzes [] collected data to identify potentially security-related events happening on the network.” EX1001, Abstract. Although the ’641 patent describes its alleged invention in the context of “monitoring, detecting and responding to computer network intrusions,” the patent and claims are broadly directed to “monitoring of any system,” “filtering the data from the monitoring for noteworthy events

happening in or affecting any system,” and “alerting human beings to any such noteworthy events.” *Id.*, 15:57-16:6.

4. I am familiar with the subject matter disclosed in the '641 patent as of its October 20, 2006 filing date and as of March 16, 2000, which I understand is the '641 patent's earliest possible priority date. I have not formed an opinion as to whether the '641 patent is entitled to the March 16, 2000 priority date, but I reserve the right to do so in the future, if necessary. That said, for the purposes of my opinions in this Declaration, I have been asked to treat March 16, 2000, as the time of invention of the '641 patent.

5. I have been asked to provide my independent technical review, analysis, insights, and opinions regarding the '641 patent and the references that form the basis for the grounds of unpatentability for claims 1-25 of the '641 patent.

6. I have reviewed and am familiar with the '641 patent and its file history. I have also reviewed and am familiar with the documents I cite in this Declaration. I understand that these documents have been provided as exhibits to the Petition that Petitioner is concurrently submitting. My analysis and opinions will focus on all of the challenged claims, claims 1-25 of the '641 patent.

7. I confirm that to the best of my knowledge the accompanying exhibits are true and accurate copies of what they purport to be and that an expert in the field

would reasonably rely on them to formulate opinions such as those set forth in this declaration.

8. I am being compensated at my customary rate of \$750 per hour for my work on this case. My compensation is not dependent upon my opinions, my testimony, or the outcome of this case.

II. GROUNDS OF UNPATENTABILITY

9. In forming my opinions about the '641 patent, I have considered the following grounds of unpatentability. Based on my review of the prior-art references that form the basis of these grounds, it is my opinion that claims 1-25 of the '641 patent would have been obvious to a person of ordinary skill in the art ("POSA") as of March 16, 2000.

Ground	References	Basis	Challenged Claims
1	Duvall in view of Chu	§ 103	1-7, 15-17
2	Duvall in view of Chu, and in further view of Trcka	§ 103	8-13, 16
3	Duvall in view of Chu and Trcka, and in further view of Ziese	§ 103	14-15
4	Duvall in view of Chu, and in further view of Cogger	§ 103	18-25

10. I have been asked to consider how a POSA would have understood the challenged claims in light of the disclosures of the '641 patent. I also have been asked to consider how a POSA would have understood the prior-art references Duvall, Chu, Trcka, Ziese, and Cogger. Further, I have been asked to consider and provide my technical review, analysis, insights, and opinions regarding whether a

POSA would have understood that the combinations of the prior-art references listed in the table above render obvious claims 1-25 of the '641 patent.

III. QUALIFICATIONS

11. In forming the opinions in this Declaration, I have considered and relied on my education, background, and experience. My experience and education are detailed in my *curriculum vitae* (“CV”). See EX1032. My CV lists my publications and identifies parties on behalf of whom I have provided expert testimony.

12. Currently, I am a tenured professor in the Department of Computer Science at the University of North Carolina at Chapel Hill (UNC), where I hold the position of Gillian T. Cell Distinguished Professor of Computer Science. I have been a faculty member at UNC since 1989 and from 2014-2022 I served as the Chairman of the Department.

13. I received a Ph.D. in computer science from the University of Washington in 1989. Before that, I received a M.S. degree in computer science from the University of Toronto in 1984, and a B.S. degree with Highest Distinction in mathematics from the University of Illinois at Urbana-Champaign in 1982.

14. I have been involved in the research and development of networked computing systems for over 40 years. During this time, I have performed research and taught in the areas of computer networks, distributed systems, real-time

systems, operating systems, multimedia networking, network traffic analysis, and network performance evaluation, among others. A major theme of my research has been the measurement and analysis of traffic patterns and traffic performance in communications networks with a particular focus on the nature and structure of web traffic. In my research and teaching, I have considered problems of the design and implementation of computer networks, including issues of multimedia networking, network performance measurement and management, network anomaly detection, and aspects of router design. In addition, I have also explored problems in the design and implementation of systems for distributed collaboration including hypertext systems such as the web.

15. Much of my research has been performed jointly with industry. For example, starting in 1991, in collaboration with IBM and Intel, my research group constructed and operated one of the first Internet videoconferencing systems. We also developed a data conferencing, “shared window system” that was similar to today’s Cisco’s WebEx and Zoom’s screen sharing products and services. The work extended to the development of collaboration tools based on the world-wide web.

16. Additional research projects included the development of tools and techniques to monitor, characterize, and control the traffic flowing in production networks. For example, in the late 1990s and early 2000s, my research evolved to consider router-based mechanisms for controlling the performance of network

traffic. In much of this research, my students and I built and instrumented network routers and performed large-scale experiments with this equipment. Part of the technology we developed for this research included means for classifying flows within routers and providing differential treatment to flows as a way of providing so-called “better than best effort” forwarding services.

17. The research further included the development of traffic management and scheduling policies and mechanisms to provide guarantees of quality-of-service to packet flows. In addition, part of the experimental infrastructure we developed to support this effort included constructing passive packet-capture devices to classify, filter, schedule, and monitor the performance of network flows in real time. Based on some of the results of this research, in 2003 my group at UNC won the most prestigious research award for original research in computer networking.

18. I have authored or co-authored over 100 articles in peer-reviewed journals, conference proceedings, texts, and monographs in the aforementioned areas of computer science and others. I have previously served as Editor-in-Chief for the journal *Multimedia Systems* and Associate Editor for the journal *Real-Time Systems*. I have also been an active participant in the Association for Computing Machinery (ACM) and the Institute of Electrical and Electronics Engineers (IEEE). Specifically, I have been a member of the ACM Special Interest Group on Data Communications (SIGCOMM) Steering Committee, the IEEE Technical Committee

on Real-Time Systems, and the ACM Special Interest Group on Multimedia. I was also previously involved in the Statistical and Applied Mathematical Sciences Institute (SAMSI) program on Network Modeling for the Internet.

19. I have edited and co-edited numerous published proceedings of technical conferences, and I have edited a book of readings in multimedia computing and networking (with Hong-Jiang Zhang), published by Morgan Kaufman. I am a co-author (with Long Le and F. Donelson Smith) of a monograph related to computer network protocols, and a co-author (with Jay Aikat and F. Donelson Smith) of a second monograph related to experimental computer networking.

20. I have served on numerous proposal review panels for the National Science Foundation and other international funding agencies in the aforementioned areas of computer science. I have served as a program chair or member of the technical program committee for over 100 professional, international, and technical conferences, workshops, and symposia.

21. I am a named inventor on four U.S. Patents. These patents are generally related to computer networking and the delivery of services over networks.

22. I have served as an expert witness and technical consultant in litigation and *inter partes* review matters concerning computer networks, distributed systems, the world-wide web, operating systems, multimedia networking, cellular and

wireline telephony, voice over IP (VoIP) telephony, datacenter networking, embedded systems, embedded software, and real-time systems, among others. I have testified in several trials, arbitrations, and claim construction hearings as an expert witness.

IV. MATERIALS CONSIDERED

23. In formulating my opinions, I have relied on my training, knowledge, and experience. Furthermore, I have specifically considered the following documents, in addition to any other documents cited in this Declaration.

<i>Exhibit #</i>	<i>Description</i>
1001	U.S. Patent No. 7,895,641 to Schneier <i>et al.</i> (“’641 patent”)
1002	Prosecution History of U.S. Patent No. 7,895,641 (“Prosecution History”)
1004	U.S. Patent No. 5,884,033 to Duvall <i>et al.</i> (“Duvall”)
1005	Yang-hua Chu, “Trust Management for the World Wide Web, M.I.T. (June 13, 1997) (“Chu”)
1007	U.S. Patent No. 6,088,804 to Hill <i>et al.</i> (“Hill”)
1008	Lane B. Warshaw <i>et al.</i> , “VenusIDS: An Active Database Component for Intrusion Detection,” Applied Research Laboratories / Department of Computer Sciences, The University of Texas at Austin, Annual Computer Security Application Conference (1999) (“Warshaw”)
1011	<i>Fortinet Inc. v. BT Americas, Inc.</i> , IPR2019-01324, Paper 9 (Institution Decision) (P.T.A.B. Jan. 7, 2020)
1012	<i>British Telecommunications PLC v. Fortinet, Inc.</i> , Civil Action No. 18-1018, Memorandum Order re: Claim Construction (D. Del. Aug. 5, 2021)
1013	<i>British Telecommunications PLC v. Fortinet, Inc.</i> , Civil Action No. 18-1018, Claim Construction Order (Aug. 18, 2021 D. Del.)
1014	U.S. Patent Pub. No. 2001/0039579 to Trcka <i>et al.</i> (“Trcka”)

<i>Exhibit #</i>	<i>Description</i>
1015	U.S. Patent No. 6,484,315 to Ziese (“Ziese”)
1017	Prosecution History of U.S. Patent No. 7,159,237 (“’237 Prosecution History”)
1026	Sinrod et al., “ <i>Cyber-Crimes: A Practical Approach to the Application of Federal Computer Crime Laws,</i> ” Santa Clara High Technology Law Journal (January 2000)
1027	Guttman et al., “ <i>Implementing Internet Firewall Security Policy,</i> ” National Institute of Standards and Technology (April 13, 1998)
1028	Rothfuss et al., “ <i>Go Ahead, Visit Those Websites, You Can’t Get Hurt ... Can You?,</i> ” Lawrence Livermore National Laboratory (June 3, 1997)
1029	Capture of Merriam Webster definition of “empirical,” www.merriam-webster.com/dictionary/empirical
1030	Coates, et al., “ <i>Melissa Virus Exposes High-Tech Vulnerability,</i> ” Chicago Tribune (March 30, 1999)
1031	Abie, “ <i>An Overview of Firewall Technologies,</i> ” Norwegian Computing Center (January 2000)
1032	Curriculum Vitae of Kevin Jeffay, Ph.D.
1033	U.S. Patent No. 6,859,783 to Cogger <i>et al.</i> (“Cogger”)
1034	<i>British Telecommunications PLC v. Fortinet, Inc.</i> , Civil Action No. 18-1018, Claim Construction Report and Recommendation (Apr. 15, 2021 D. Del.)
1036	U.S Patent No. 6,775,657 to Baker (“Baker”)
1037	Escamilla, T., “ <i>Intrusion Detection: Network Security Beyond the Firewall,</i> ” Wiley Press, 263-281 (1998) (excerpt)
1038	U.S. Patent No. 5,261,044 to Dev <i>et al.</i> (“Dev”)
1039	U.S. Patent No. 5,768,501 to Lewis (“Lewis”)

V. LEGAL STANDARDS

24. I have also relied upon various legal principles (as explained to me by Petitioner's counsel) in formulating my opinions. My understanding of these principles is summarized below.

25. I have been told the following legal principles apply to analysis of patentability. I also have been told that, in an *inter partes* review proceeding, a patent claim may be deemed unpatentable if it is shown by a preponderance of the evidence that the claim is rendered obvious by one or more prior-art references under 35 U.S.C. § 103.

26. I understand that documents and materials (such as printed publications or patents) that qualify as prior art can be used to render a patent claim obvious.

27. I understand that, once the claims of a patent have been properly construed, the second step in determining invalidity of a patent claim requires a comparison of the properly construed claim language to the prior art on a limitation-by-limitation basis.

A. My Understanding of Claim Construction

28. I understand that during an *inter partes* review proceeding, claims are to be construed in light of the specification as would be read by a POSA at the time of the purported priority date. I also understand that claims in an IPR are construed using the same claim construction standard that would be used to construe the claim in a civil action. I understand that claim terms are given their ordinary and

customary meaning as would be understood by a POSA in the context of the entire disclosure. A claim term, however, will not receive its ordinary meaning if the patentee acted as his own lexicographer and clearly set forth a definition of the claim term in the specification. In this case, the claim term will receive the definition set forth in the patent.

B. My Understanding of a Person of Ordinary Skill in the Art

29. I understand that a POSA is presumed to be aware of all pertinent art, thinks along conventional wisdom in the art, and is a person of ordinary creativity—not an automaton.

30. I have been asked to consider the level of ordinary skill in the field that someone would have had at the time the alleged invention was made. In deciding the level of ordinary skill, I considered the following:

- the levels of education and experience of persons working in the field;
- the types of problems encountered in the field; and
- the sophistication of the technology.

31. I understand that the relevant time for considering whether a claim would have been obvious to a POSA is at the time of the alleged invention, which I have assumed to be March 16, 2000.

32. Regardless, if I use “I” or “POSA” during my technical analysis below, all of my statements and opinions are to be understood to be based on how a POSA would have understood or read a document at the time of the alleged invention.

C. My Understanding of Obviousness

33. I understand that a patent claim is unpatentable if the claimed invention would have been obvious to a POSA at the time of the alleged invention. I understand that this means that even if all of the elements of the claim cannot be found in a single prior-art reference that would anticipate the claim, the claim can still be unpatentable.

34. It is my understanding that, to obtain a patent, a claimed invention must have been, as of its earliest possible priority date, non-obvious in view of the prior art in the field. I understand that a patent claim is obvious when the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious to a POSA at the time the alleged invention was made.

35. When considering the issues of obviousness, I have been told that I am to do the following:

- Determine the scope and content of the prior art;
- Ascertain the differences between the prior art and the claims at issue;
- Resolve the level of ordinary skill in the art; and

- Consider evidence of secondary indicia of non-obviousness (if available).

36. With respect to determining the proper scope of prior art to examine, I understand that for a prior-art reference to be properly used in an obviousness ground under 35 U.S.C. § 103, the prior-art reference must be analogous art to the claimed invention. I have been told that a reference is analogous art to the claimed invention only if: (1) the reference is from the same field of endeavor as the claimed invention; or (2) the reference is reasonably pertinent to the particular problem solved by the inventor.

37. I understand that factors relevant to determining the proper field of endeavor include the inventor's explanations of the subject matter (including the patent specification) and the claimed invention's structure and function. And I further understand that to be reasonably pertinent to the particular problem solved by the inventor, a prior-art reference must logically commend itself to the inventor's attention in considering his or her problem.

38. I have been told that an analogous reference may be modified or combined with other analogous references or with the POSA's own knowledge if the person would have found the modification or combination obvious. I have also been told that a POSA is presumed to know all relevant, analogous prior art, and the

obviousness analysis may take into account the inferences and creative steps that a POSA would employ.

39. In determining whether a prior-art reference would have been combined with another prior-art reference or other information known to a POSA, I have been told that the following principles may be considered:

- A combination of familiar elements according to known methods is likely to be obvious if it yields predictable results;
- The simple substitution of one known element for another is likely to be obvious if it yields predictable results;
- The use of a known technique to improve similar devices, products, or methods in the same way is likely to be obvious if it yields predictable results;
- The application of a known technique to a prior-art reference that is ready for improvement is likely obvious if it yields predictable results;
- Any need or problem known in the field and addressed by the reference can provide a reason for combining the elements in the manner claimed;
- A skilled artisan often will be able to fit the teachings of multiple references together like a puzzle; and
- The proper analysis of obviousness requires a determination of whether a POSA would have a “reasonable expectation of success”—not

“absolute predictability” of success—in achieving the claimed invention by combining prior-art references.

40. I understand that whether a prior-art reference renders a claim unpatentable as obvious is determined from the perspective of a POSA. I have also been told that, while there is no requirement for the prior art to contain an express suggestion to combine known elements to achieve the claimed invention, a suggestion to combine known elements to achieve the claimed invention may come from the prior art as a whole or individually, as filtered through the knowledge of one skilled in the art. In addition, I have been told the inferences and creative steps a POSA would employ are also relevant to the determination of obviousness.

41. I also understand that when a work is available in one field, design alternatives and other market forces can prompt variations of it, either in the same field or in another. I have also been told that if a POSA can implement a predictable variation and would see the benefit of doing so, that variation is likely to be obvious. I have been told that in many fields, there may be little discussion of obvious combinations, and in these fields market demand—not scientific literature—may drive design trends. I have been told that, when there is a design need or market pressure and there are a finite number of predictable solutions, a POSA has good reason to pursue those known options.

42. I have been told there is no rigid rule prescribing that a reference or combination of references must contain a “teaching, suggestion, or motivation” to combine references. But I also have been told that the “teaching, suggestion, or motivation” test can be a useful guide in establishing a rationale for combining elements of the prior art. I have been told this test poses the question as to whether there is an express or implied teaching, suggestion, or motivation to combine prior art elements in a way that realizes the claimed invention, and that it seeks to counter impermissible hindsight analysis.

VI. THE '641 PATENT

A. Alleged Problem in the Prior Art

43. The '641 patent is directed to “network security” and explains that “[m]ost computer and network security products focus on prevention.” EX1001, 1:7-9, 1:24-25. For instance, “[f]irewalls prevent unauthorized traffic from entering a company’s internal network; authentication mechanisms prevent unauthorized persons from logging on to a company's computers; and encryption prevents unauthorized persons from reading files.” *Id.*, 1:25-29. But “such products cannot be relied upon to work perfectly.” *Id.*, 1:29-30. The patent also states that “[s]ystem administrators normally do not have the time or ability to read through large amounts of constantly updated audit information, looking for attacks on their systems [and] also do not have the time to continuously monitor hacker activities [or] become experts on every kind of intrusion and to maintain that expertise.”

EX1001, 1:34-43. And while the '641 patent acknowledges that “automatic defenses may work against automated attacks,” such systems “are at a disadvantage against an intelligent attack.” *Id.*, 1:50-53.

44. The '641 patent explains that system administrators are unable to perfectly maintain network security because, for instance, they “normally do not have the time or ability to read through large amounts of constantly updated audit information, looking for attacks on their systems.” *Id.*, 1:36-39. Additionally, network administrators “do not have the time to continuously monitor hacker activities” or “to become experts on every kind of intrusion and to maintain that expertise.” *Id.*, 1:39-43.

45. Accordingly, the '641 patent posits that “[a]n effective monitoring, detection and response system should be designed not to replace a customer's system administrators but to augment their abilities.” *Id.*, 1:34-36. The '641 patent thus proposes “[a] monitoring, detection and response system that employs human intelligence, uses trained personnel in the loop, and takes advantage of security intelligence and other knowledge databases.” *Id.*, 1:44-47. Such systems “can provide customer system administrators the advice and coaching they need, when they need it, to help them repel attacks and maintain network integrity and uptime.” *Id.*, 1:44-50.


B. Overview of the '641 Patent

46. To address the above considerations, the '641 patent purports to extend traditional mechanisms that prevent unauthorized traffic by “provid[ing] a managed security monitoring service (the ‘MSM service’) that monitors a customer’s network activity using a probe or ‘sentry’ system, collects status data from monitored components, [and] filters or otherwise analyzes the collected data for activity possibly implicating security concerns.” EX1001, 1:59-65. The patent involves human analysts by “transmit[ting] information about such activity” to them to determine an “appropriate response.” *Id.*, 1:65-2:2. The '641 patent explains, “[t]he MSM service is not intended to replace but to supplement, and thereby render more effective, a customer’s existing preventative security products,” including traditional “firewalls, servers, [and] routers.” *Id.*, 2:2-8.

47. Figure 1 of the '641 patent depicts an “overview of the systems in an MSM service.” *Id.*, 4:44-45. The MSM service includes two portions—“components and systems that operate on the customer site (that is, within the customer's firewall) and components and systems that operate within the SOC [(secure operations center)] (that is, within the SOC firewall).” *Id.*, 4:45-49. The probe/sentry system 2000 of the customer site “monitors sensors attached to customer network 1000 for evidence of potential security-related events happening on network 1000.” *Id.*, 4:55-59. “Such sensors can include firewalls and intrusion detection systems 1010,

commercially available sensors and agents 1020, decoys and honeypots 1030 . . .

and custom sensors and agents 1040.” *Id.*, 4:59-64.




EX1001, FIG. 1.

48. “Probe/sentry system 2000 collects and filters (positively and/or negatively) or otherwise analyzes the constantly updated status data it receives from sensors.” EX1001, 5:25-27. The probe/sentry system 2000 sends “sentry messages” to the SOC, which contain “[n]oteworthy data possibly indicating an unauthorized

intrusion or other security-related event.” *Id.*, 5:44-49. Upon receipt of such sentry messages, the SOCRATES 6000 at the SOC “manages the problem resolution process,” such as by generating “‘problem tickets’ representing potential security-related happenings on the customer's network.” *Id.*, 7:24-30. “MSM service security analysts receive problem tickets for handling access information useful for problem investigation, resolution and response, and otherwise interface with SOCRATES 6000 via security analyst consoles 6010.” *Id.*, 7:38-41.

49. Figure 2 of the '641 patent provides “a system overview of an exemplary embodiment of a probe/sentry system.” *Id.*, 8:41-42.



EX1001, FIG. 2.

In this example, “[d]ata collected by sensors 1010, 1020, 1030 and 1040 . . . are collated by sensor data collator 2010.” EX1001, 8:48-51. “The data is first filtered by negative filtering subsystem 2020, which discards uninteresting information, and then by positive filtering subsystem 2030, which selects possibly interesting information” *Id.*, 8:51-54. The patent does not specify how discarding or selecting occurs, or how the system determines which information is interesting.

50. “Data neither discarded by negative filtering subsystem 2020 nor selected out as interesting by positive filtering subsystem 2030 form the ‘residue,’ which is sent to anomaly engine 2050 for further analysis.” *Id.*, 8:55-59. “Interesting” data is then forwarded to a gateway system 4000 via pipes 3000 for review by a human analyst. *Id.*, 8:67-9:2, 9:13-18.

51. In some embodiments, the MSM service “may be customized, either dynamically or offline, to accommodate network-specific needs and to reflect feedback received about the demonstrated efficacy of a real world response to an actual event.” *Id.*, 2:31-35. For instance, the SOC can send “updates” to the probe/sentry system “during actual operation” so that the probe sentry system can securely install the updates. *Id.*, 5:32-37.

52. It is important to note that, although an example use case of the ’641 patent is “network intrusion monitoring,” *see id.*, 1:18-20, 15:58-62, the patent and claims are not limited to addressing network intrusions. The patent explains,

“[t]hose skilled in the art will realize [] that the present invention is usable generally for (i) **monitoring of any system**; (ii) filtering the data from the monitoring **for noteworthy events happening in or affecting any system**; (iii) alerting human beings to any such noteworthy events; and (iv) providing for a response by human beings to any such alert.” *Id.*, 15:62-16:6.¹ Claim 1 of the ’641 patent reflects this broad applicability, reciting “a security monitoring system” and “identify[ing] potentially security-related events,” which need not relate specifically to network intrusion. Claim 18 similarly recites “a security monitoring system” and “creating an event record for information received about an identified potentially security-related event occurring on the network.”

53. Indeed, at the time of the earliest claimed priority for the ’641 patent, a POSA would have understood that intrusion detection would have been just one aspect of network security. For example, a POSA would have understood that “[n]etwork intrusion detection is a process that can identify and respond to misuse or policy violations on a network” by using “network equipment such as routers and firewalls.” EX1036,² 1:19-28. But, in addition to network intrusion detection, a POSA would have understood that network security would have also included,

¹ All emphasis added unless otherwise indicated.

² U.S. Pat. No. 6,775,657 to Baker (“Baker”), filed Dec. 22, 1999.

among other things, “virus checking on incoming file transfers” and “tracking malicious acts that are performed ‘on-site’ by company employees.” EX1014 (Trcka), ¶7.

C. Prosecution History

54. The ’641 patent was rejected twice during prosecution before being allowed. The Office rejected claims 45-61 (issued claims 1-17) on basis of obviousness-type double patenting over its parent, U.S. Patent No. 7,159,237 (EX1011, “’237 patent”), because of the similarity between the claims. EX1002, 136, 197. The Applicant overcame this rejection by filing a terminal disclaimer over the ’237 patent. *Id.*, 223.

55. During prosecution of the parent, the ’237 patent, the Office rejected the as-filed claims of the ’641 patent as being anticipated by U.S. Patent No. 5,796,942 to Esbensen. EX1017, 0207-208. In response, Applicant made one substantive amendment to the independent claims, as shown below.

b) <u>analyzing status data to identifying potentially security-related events represented in the ~~pertaining to said network by analyzing said status data, wherein the analysis includes filtering followed by an analysis of post-filtering residue, wherein the post-filtering residue is data neither discarded nor selected by filtering;~~</u>
--

EX1017, 0219 (excerpt).

56. Applicant focused solely on this claim amendment in its arguments for the allowance of the claims—Applicant did not argue that any other elements of the independent claims were novel. EX1017, 0226-27. The Examiner, however,

subsequently issued a new obviousness rejection, asserting that U.S. Patent No.

6,704,874 to Porras disclosed the amended “post-filtering residue” limitations. *Id.*,

237-38. In its next response, Applicant argued:

While Porras et al. disclose filtering formatted alerts and discarding alerts that do not match certain criteria of the user-specified filters, no mention is made of analyzing data that was not discarded and that was also not selected by the filtering process. Any analysis performed by Porras et al. is clearly done on alerts that were selected by the filtering process.

EX1017, 0257 (emphasis in original).

57. The Examiner disagreed and issued another obviousness rejection based on the same art. EX1017, 0282-89. Applicant argued that Porras disclosed “data that is selected by filtering” and “data [that] is discarded in the filtering process,” but Porras did not disclose any “data that could be considered post-filtering residue.” *Id.*, 0300-301. The Examiner agreed and issued a notice of allowance, stating that “Porras fails to disclose or suggest the claim limitation, wherein the post-filtering residue is data neither discarded nor selected by filtering,” because “Porras discloses that if the filter does not match the criteria the data is discarded.” *Id.*, 0326.

58. During prosecution of the ’641 patent, the Applicant also overcame rejections to 62-68 and 70 (issued claims 18-25) by making a similar amendment to the claims to add the “post-filtering residue” limitation:

creating an event record for information received about an identified potentially security-related event occurring on the network, wherein the potentially security-related event is identified by filtering followed by an analysis of post-filtering residue, wherein the post-filtering residue is neither discarded nor selected by the filtering;

EX1002, 218.

59. The Examiner subsequently allowed the claims of the '641 patent, stating that the prior art did not disclose the “post-filtering residue” limitation. EX1002, 235-238. But techniques for analyzing “post-filtering residue,” i.e., data neither discarded nor selected by filtering, were well known at the time of the '641 patent, as illustrated in Duvall and Chu, which were never before the Examiner during prosecution of the '641 patent.

D. Previous *Inter Partes* Review

60. I understand that the '641 patent was previously the subject of a petition for *inter partes* review in IPR2019-01324, which the Board denied. EX1011, 22. The petitioner in that proceeding—Fortinet Inc.—alleged that claims 1-25 of the '641 patent were obvious over Hill (EX1007) in view of Warshaw (EX1008). EX1011, 6-7.

61. Fortinet presented two theories for the “post-filtering residue” limitations of the independent claims—one based on Hill alone, and one based on Hill in view of Warshaw. EX1011, 13. The Board disposed of Fortinet’s first theory because Hill allegedly did not teach “data ... discarded ... by filtering.” EX1011, 17-18.

62. Regarding the second theory, the Board did not dispute Warsaw's alleged teaching of analyzing "post-filtering residue." EX1011, 18-19. The Board instead was not persuaded "that [Fortinet's] rationale for combining Hill and Warsaw is supported by sufficient rational underpinning." *Id.*, 18-21. The Board thus denied institution. *Id.*, 22.

E. Level of Ordinary Skill in the Art

63. A person of ordinary skill in the art ("POSA") at the time of the '641 patent would have had a B.S. degree in Computer Science, Computer Engineering, or an equivalent field, as well as at least 2-3 years of academic or industry experience in the design, analysis, and monitoring of computer networks, including issues of network security and network administration, or comparable industry experience.

64. I am well qualified to determine the level of ordinary skill in the art and to provide opinions from the perspective of a POSA. As I explain in Section III and further show in my CV, as of March 16, 2000, I had acquired extensive experience related to the design, analysis, and monitoring of computer networks, including issues of network security and network administration. Thus, as of this date, I had more than ordinary skill in the art. Throughout my career and the projects that I have been involved in, I have worked with many people having the qualifications of a POSA identified above. My interactions with those people included working with

graduate students who were POSAs, as well as collaborating with persons in industry and academia who were POSAs and interacting with alumni who were POSAs.

F. Claim Construction

65. I understand that certain claims of the '641 patent were previously construed in a district-court litigation between the Patent Owner (“British Telecommunications”) and Fortinet. *See British Telecommunications PLC v. Fortinet Inc.*, No. 1:18-cv-01018-CFC-MPT (D. Del.).³ Specifically, the district court construed the following terms:

- “*status data*”: “data extracted from or generated about the traffic or system processing it that is informative as to the status of the network and its components”;
- “*dynamically*”: “during actual operation, rather than offline”; and
- “*probe*”: “a component that collects data from one or more network components to which it is attached, filters or otherwise analyzes the data that

³ I understand that the district-court case between British Telecommunications and Fortinet was terminated on January 4, 2022 based on a stipulation of dismissal. D.I. 161. It is my understanding that the issues of infringement and invalidity were not decided.

has been collected, transmits noteworthy information, and receives feedback in order to update its capabilities of analysis.”

- “*information received about an identified potentially security-related event occurring on the network wherein the potentially security related event is identified by filtering followed by an analysis of post-filtering residue*”:
“information received from a probe about an identified potentially security-related event occurring on the network wherein the potentially security-related event is identified at the probe by filtering status data followed by an analysis of post-filtering residue.”
- “*customer information*”: “information about a customer.”
- “*problem ticket*”: “a consolidation of the event record, correlated customer information and symptom record, and linked problem resolution assistance information.”

EX1013, 1-2; *see also* EX1012. Additionally, British Telecommunications and Fortinet agreed that the phrase “*post-filtering residue, wherein the post-filtering residue is data neither discarded nor selected by filtering*” should be construed as “status data that undergoes negative and positive filtering, but is neither discarded by such negative filtering nor selected by such positive filtering.” EX1013, 3.

66. For this proceeding, it is my opinion that no claim terms require an explicit construction. That said, it is also my opinion that, as I show below, the

challenged claims are unpatentable under either the claims' ordinary and customary meaning as understood by one of ordinary skill in the art at the time of the invention in light of the specification or the district court's previous claim constructions.


VII. OVERVIEW OF THE PRIOR ART

67. Below I provide a summary of each prior-art reference relied upon for the grounds of unpatentability discussed herein. *See* Section II. I have been informed by counsel that each of these references are prior art to the '641 patent and were publicly available as of the earliest possible priority date of the '641 patent.

A. Overview of Duvall


68. **U.S. Patent No. 5,884,033** to Duvall *et al.* (EX1004) is directed to a computer network wherein “[a] client-based filtering system compares portions of incoming and/or outgoing messages to filtering information in a filter database, and determines whether to block or allow incoming and/or outgoing transmissions of messages in response to the comparison.” EX1004, Abstract. Duvall explains that, “[w]hile the Internet is very useful for obtaining information, many people are concerned that objectionable materials, especially sexually explicit materials, can be obtained too easily.” *Id.*, 1:12-15. To address these issues, Duvall set out “to provide a system that allows Internet users to filter material transmitted over the Internet.” *Id.*, 1:27-29.

69. Duvall's Figure 1 illustrates an example “network with a client computer for accessing the Internet.” *Id.*, 2:24-25.



EX1004, FIG. 1.

70. “A user uses a computer which serves as a client 10 to communicate with other computers over Internet 12.” EX1004, 2:34-35. In one embodiment, “a filtering system resides in client computer 10” and includes a “filter database 24,” as shown in Figure 2. *Id.*, 3:41-45.



EX1004, FIG. 2.

71. In another embodiment, “the filtering system can be provided from a server 30 that is on the client’s own network 40.” EX1004, 8:18-21. Duvall explains, “[s]uch a topology may be more desirable, for example, in a corporate network that has a limited size and a uniform set of filters for many users.” *Id.*, 8:21-23. “This version of the filtering system uses the same type of filter database as a client-based filtering system, but the filter database is located on server 30.” *Id.*, 8:23-26.

72. “The filter database has lists of filters, some of which are identified as either ALLOW filters or BLOCK filters for respectively allowing or blocking transmission.” *Id.*, 3:64-66. The filter database further includes “deferred action filters” for “defer[ring] the decision whether to allow or block.” *Id.*, 1:43-47, 4:27-30. An ALLOW filter *allows* data transmissions to pass through. *Id.*, 3:64-66. A BLOCK filter *blocks* data transmissions from passing through. *Id.* And a deferred action filter can specify additional analysis of data transmissions that are neither ALLOWED nor BLOCKED. *Id.*, 4:65-5:19.


73. Each filter entry in the filter database includes “fields for port number and IP address.” *Id.*, 3:66-4:1. The fields can specify a “specific port number, a specific IP address, and/or meta values ANY IP and ANY PORT, which stand for all possible IP addresses and all possible port numbers, respectively.” *Id.*, 4:1-5. “Each filter entry in the filter database also has a field for specifying an action to be

taken by the client if that filter were retrieved.” *Id.*, 4:12-14. ALLOW and BLOCK filters specify “direct actions,” while deferred action filters may specify “deferred actions.” *Id.*, 4:14-21, 4:27-30. “Direct actions indicate that the system should unconditionally allow or unconditionally block the transmission.” *Id.*, 4:15-17. Deferred actions are additional actions that may be taken with respect to data transmissions upon satisfaction of some additional criteria. *Id.*, 4:65-5:19. Further, deferred action filters “preferably have additional fields,” such as “(1) a keyword...; (2) a filter pattern to be compared to data in the message...; (3) a directional indicator (IN/OUT) for indicating incoming or outgoing transmissions; (4) a compare directive for the type of match; and (5) an action to be taken, typically to allow or block the transmission.” *Id.*, 5:8-15.

74. In a preferred embodiment, the filters are implemented in sequence so that “the system searches ALLOW filters first, BLOCK filters next, and deferred action filters last.” *Id.*, 4:22-5:1. “In this order, ALLOW filters [] take precedence over BLOCK filters.” *Id.*, 4:30-36. For data transmissions that do not match ALLOW or BLOCK filters in step 104 (i.e., residue), Duvall determines whether those data transmissions match a deferred action filter. *Id.*, 4:65-5:1, 5:8-29.

75. For one embodiment, Figure 3 depicts a method performed by the filtering system for “filter[ing] outgoing messages on the basis of the IP address of the host server that is being contacted by the client.” *Id.*, 4:38-40. At step 100, the

filtering system “detects when the client opens the data stream for a particular port and IP address [], and searches the filter database for matching filters.” *Id.*, 4:40-43. At step 102, the filtering system retrieves filters that match the particular port and IP address. *Id.*, 4:40-46. At step 104, the system searches the retrieved filters “to determine if any require immediate action” (i.e., direct action—allow or block transmission). *Id.*, 4:46-50. If a matched filter specifies to allow the transmission, at step 108 the filtering system “allows the transmission to proceed normally.” *Id.*, 4:50-52. If a matched filter specifies to block the transmission, at step 110 the system “takes action to block the transmission,” and, at step 112, “provides a ‘Blocked’ message to the user.” *Id.*, 4:52-55.



EX1004, FIG. 3.


76. If, at step 104 and 106, it is determined that no immediate action must be taken (i.e., if the IP address/port is not matched by an ALLOW or BLOCK filter, EX1004, 4:28-31, at step 116 the system “determine[s] whether a deferred action must be taken with respect to any of the retrieved filter.” *Id.*, 4:65-5:1. If there is no deferred action, in addition to determining if unconditional allowing or blocking is required in steps 104 and 106, the system “can default to allow the transmission” at

step 108, or the system “can default to block the transmission” at steps 110 and 112.

Id., 4:48-54, 5:1-3.

77. If, at step 116, there is a deferred action, then at step 130 (shown in Figure 4), the system compares the direction of the transmission to the directional indicators of the deferred filter entries. *Id.*, 5:19-22. “For each entry in a filter that matches the direction,” at step 132, the system compares “the beginning of the data stream . . . to the keyword in the filter entry.” *Id.*, 5:23-25. “If the key-words match,” at step 134, the system compares “the data . . . to the filter pattern of the filter(s) according to the compare directive” in the deferred filter entries. *Id.*, 5:25-28. “If a match is found at this level,” at step 138 the system performs “the specified action in the filter.” *Id.*, 5:28-29.

78. Duvall also explains that data transmissions that initially match a deferred action filter may not match the additional fields of that filter. *Id.*, 5:8-29. Specifically, this can occur at step 134 if the keywords don’t match, and/or at step 138 if a match is not found based on the compare directive (shown in the annotated version of Figure 4 below). Duvall does not explicitly disclose techniques for resolving these unknown data transmissions. Instead, Duvall simply suggests that “the system can default to allow the transmission (step 108 [of FIG. 3]), or it can default to block the transmission (steps 110 and 112).” *Id.*, 5:1-3.



EX1004, FIG. 4.

79. Duvall also discloses that the list of filters can be updated and/or modified in the filter database. *Id.*, 7:16-21. For instance, to account for new web sites that are added to the internet, “[t]he updating mechanism in the filtering system can ... cause[] the filtering system to download ... one of a number of new sets of filters.” *Id.*, 7:21-26. The updating mechanism can also “cause[] some or all of the existing filters in the filter database to be replaced with the newly downloaded

filters.” *Id.*, 7:26-28. As another example, “the filtering system has an editing Manager that allows the user to edit the database to add, delete, or modify filters in the database.” *Id.*, 8:2-5. “The editing filter allows the user to make custom changes if the user believes that the filtering system is filtering too much or too little.” *Id.*, 8:5-7.

B. Overview of Chu

80. **“Trust Management for the World Wide Web”** by Yang-hua Chu (EX1005) provides techniques for determining whether to block or allow unknown data through a network. EX1005, 9, 44. Chu explains that “precise and reliable ways to manage trust” are critical for various applications in computer networking, such as “electronic commerce and mobile programming.” *Id.*, 9. This is in part because “malicious parties may lure users into believing that a false piece of information is trustworthy.” *Id.* To address these issues, Chu developed a “trust management infrastructure, in which trust is specified, disseminated, and evaluated in parallel with the information infrastructure.” *Id.*

81. Specifically, Chu is directed to protecting users from untrustworthy information available on the World Wide Web and provides examples of trust-management issues, such as a user attempting to download software from a Web page. EX1005, 9, 23. Chu explains that the user may have “concerns that prevent her from downloading the game on her machine,” such as “[d]oes this game contain

a virus that would erase her hard drive? (security issue).” *Id.* Chu’s trust management infrastructure therefore implements various policy languages that allow users to make trust decisions about data transmitted over the Internet. *Id.*, 28, 35-43.

82. For example, Chu discloses sample policies addressing the question of “should I download the active content at this URL.” EX1005, 43-48. Chu’s exemplary policies make use of a “blacklist” (i.e., “sites or directories [the system] should not download codes from”) and a “whitelist” (i.e., “sites know[n] to be trustworthy”), designed to prevent transmissions from “Web sites that serve potentially dangerous” content. *Id.*, 44. An exemplary Chu policy (shown below) *blocks* downloads from URLs in the blacklist, and Chu’s policy *allows* downloads from URLs in the whitelist. *Id.* If a URL is unknown (i.e., it’s not in the blacklist or whitelist), the user is prompted for attention. *Id.*

Policy in English

Do not download the code if the URL is served from Harvard or CalTech Web servers. Download it automatically if served from MIT. Prompt me for my attention otherwise.

Profiles-0.92

```
(threshold-and
  2
  (not (url-match URL ("http://www.harvard.edu"
    "http://www.caltech.edu")))) ← Block filter
  (url-match URL ("http://web.mit.edu")) ← Allow filter
  unknown) ← Unmatched residue for human review
```

EX1005, 44 (annotated).

83. While Chu discloses that “[t]he blacklist and whitelist ensure good automation of the trust decision process if the lists are reasonably complete,” Chu recognizes that data transmissions not matching the blacklist or whitelist need further analysis and human intervention. EX1005, 44. Chu thus explains that a person can be prompted for attention when the transmission does not meet the criteria of the blacklist or whitelist. *Id.* (“User intervention is needed only when the given URL is in neither the blacklist nor the whitelist.”).


C. Overview of Trcka

84. **U.S. Patent Pub. No. 2001/0039579** to Trcka *et al.* (EX1014) is directed to a “network security and surveillance system” that “passively monitors and records the traffic present” on a computer network “without interrupting or otherwise interfering with the flow of the traffic.” EX1014, Abstract. Trcka explains that “[a] significant problem in the field of computer networks has been the inability to adequately protect private Internet-connected computer networks from security attacks.” *Id.*, ¶3. For instance, “connecting the internal network to the Internet can have devastating consequences in the absence of an adequate security mechanism.” *Id.*, ¶4. Examples of such devastating consequences include a “break-in by a hacker” that results in “deletion of important data or software files, the introduction of a virus to the network, and/or the public dissemination of confidential information.” *Id.*

85. Trcka explains that “firewalls” provide “limited protection against intrusions is a network firewall system” by “restrict[ing] the flow of traffic between two networks based on a pre-programmed access control policy.” *Id.*, ¶5. However, such firewalls “are generally only effective at protecting against known types of security attacks.” *Id.*, ¶6. “Once a user determines how to circumvent the firewall's access control policy, the firewall offers little or no protection.” *Id.* “Another problem with existing firewall systems is that they perform little or no virus checking on incoming file transfers, and, even when virus checking is performed, detect only known types of viruses.” *Id.*, ¶7.

86. Trcka addresses these issues by providing “a network security and analysis system which includes a variety of features for automatically and interactively monitoring and analyzing traffic on a LAN (local area network), WAN (wide area network), or other type of computer network.” *Id.*, ¶9. The system 60 typically “resides in a secure area of a corporation or other organization.” *Id.*, ¶80. “The system includes a controller 64 that connects to the network 30 using a network interface card (not shown).” *Id.*, ¶55. The controller 64 is coupled to data recorder 66 that “passively captur[es], at the data link level (or a lower level), all network traffic that is present at the network monitoring point 36,” for instance, using a “standard network interface card.” *Id.*, ¶¶38, 56. The system 60 also includes

a data playback unit 68 that is “used to play back previously-recorded archival recordings for analysis of the traffic sequences represented therein.” *Id.*, ¶60.



EX1014, FIG. 2.

87. The system 60 can be configured to “automatically detect security breaches, inbound viruses, and other anomalies during or shortly after (e.g., within several minutes or hours) of occurrence.” EX1014, ¶76. The system 60 further “enable[s] authorized users to interactively analyze and manipulate pre-recorded traffic data through a set of powerful analysis tools.” *Id.*, ¶77. “This function is accomplished through the analysis applications 100, which include functionality for performing such actions as displaying user-specified types of network events, conducting pattern searches of selected packet data, reconstructing transaction sequences, and identifying pre-defined network problems.” *Id.* To enable such

functionality, Trcka's system 60 includes a "graphical user interface (GUI) 104" through which "the user can launch and control the various analysis applications 100 through a common set of menus and controls." *Id.*, ¶79.

88. The applications 100 can include, for instance, "Audit, Problem Determination, Lost Data Recovery/Network Replay, Damage Assessment/Recovery, Security Violations, Network Operating Characteristics, Network Problem Solving, Special Non-Network Events, and Data Replay" applications. *Id.*, ¶109. An example user interface "from which the user can select individual analysis applications is depicted in FIG. 12." *Id.*; *see also id.*, ¶138.

FIG. 12

EX1014, FIG. 12.

D. Overview of Ziese


89. **U.S. Patent No. 6,484,315** to Ziese (EX1015) is directed to “dynamically distributing updates in a network.” EX1015, Title. Ziese explains that “[l]ike the Internet, private networks also have become common means for communicating important information,” and that such private networks may be distributed over “wide area networks (WANs)” that include “disparate network sites.” *Id.*, 1:25-31, 2:29-36. Ziese further explains, “[o]ne danger to reliable and secure network communications is posed by hackers or other unauthorized users disrupting or interfering with network resources.” *Id.*, 1:32-34. “For example, serious financial disruptions occur when hackers obtain financial account information or credit card information and use that information to misappropriate funds.” *Id.*, 1:37-40.

90. Thus, “network administrators” commonly use “[i]ntrusion detection systems” to “detect and identify unauthorized use of a computer network before the network resources and information are substantially disrupted or violated.” *Id.*, 1:41-46. Such systems “look for specific patterns in network traffic, known as intrusion signatures to detect malicious activity.” *Id.*, 1:46-48. “Detected intrusion signatures are reported to network administration.” *Id.*, 1:52-53.

91. However, “[a] problem with conventional intrusion detection systems is that when a new vulnerability, or type of attack on the network, is discovered, a new

intrusion signature must be generated and installed for each intrusion detection system.” *Id.*, 1:54-57. “Because new types of attacks appear more frequently than network administrators typically check with an intrusion detection provider for new signatures, networks often remain vulnerable to new types of attacks even though new signatures are available to identify and prevent such attacks.” *Id.*, 1:62-67.

92. To address this issue, Ziese discloses techniques “for dynamically distributing intrusion detection and other types of updates in a network that substantially eliminate or reduce disadvantages and problems associated with prior methods and systems.” *Id.*, 2:2-6. For example, “programs are automatically updated by downloading and distributing an update in response to an automated event,” and “[a]s a result, systems with a common program separately running at several sites may update each site with no or minimal operator interaction.” *Id.*, 2:39-44. Such updates may be distributed from a centralized update server, as shown in Figure 1 below.



EX1015, FIG. 1 (annotated).

93. Specifically, Figure 1 depicts a system 10 “for dynamically distributing updates in a network.” EX1015, 3:17-18. The system 10 includes a sensor update sever 30 that “includes new intrusion signatures generated by an intrusion detection service provider in response to new types of attacks.” *Id.*, 3:57-63. “The intrusion detection service provider generates the new signatures and provides them as the update 32 on a web page at the sensor update server 30 to allow customers to access the new signatures over the Internet 22.” *Id.*, 3:63-67. Customers are able to download the update 32 over the Internet 22 so that “the new signatures [are] added to ... the host 24.” *Id.*, 3:67-4:3. “In this way, the intrusion detection sensors 26 are

kept up-to-date and able to detect and report new types of network and/or host based attacks.” *Id.*, 4:3-6.

E. Overview of Cogger

94. **U.S. Patent No. 6,859,783 to Cogger *et al.*** (EX1033) is directed to “[a] system and method for opening and tracking trouble tickets over the public Internet.” EX1033, Cogger, Abstract. Cogger provides techniques for creating and transmitting information about detected “network event[s]” to “service organizations responsible for resolving the network event.” EX1033, 2:50-63. Specifically, Cogger provides “a trouble ticket management system for enabling an Internet enabled customer to generate a trouble ticket relating to a service provided by an enterprise to said customer. EX1033, 2:64-3:1. Cogger’s “trouble tickets” aggregate information concerning a “network event” for transmission to a “Customer Service Management System” for analysis and resolution by the service organization. EX1033, 2:64-4:3.

95. “Once a trouble ticket is opened, the customer workstation tracks the existing trouble tickets through a browser based graphical user interface.” EX1033, 2:57-60. “The graphical user interface provides current and historical status reports of the actions taken to resolve a network event and the service organizations responsible for resolving the network event.” EX1033, 2:60-63.

96. Cogger illustrates and example “trouble ticket” in Figure 8(g):




FIG. 8(g)

EX1033, FIG. 8(g).

As shown above, Cogger's "trouble ticket" can include various information about a network event, for example a description, customer information, time and date of occurrence, and priority. EX1033, 20:1-15. And Cogger discloses that this information is customizable and "will be different for different types of tickets." EX1033, 20:13-15.

F. The prior art is analogous art.

97. Each applied reference is in the same field of endeavor as the '641 patent and would have also been reasonably pertinent to problems faced by the inventor of the '641 patent. As described above, the '641 patent and the prior-art references (Duvall, Chu, Trcka, Ziese, Cogger) are all in the same field of endeavor,

dealing with computer network monitoring systems that monitor or analyze traffic to/from the Internet and/or the World Wide Web. Additionally, Duvall, Chu, Trcka, Ziese, and Cogger are all pertinent to the problems faced by the inventors—creating effective monitoring, detection, and response systems. Accordingly, it is my opinion that Duvall, Chu, Trcka, Ziese, and Cogger are analogous prior art.


VIII. GROUND 1: CLAIMS 1-7 AND 15-17 ARE OBVIOUS OVER DUVALL IN VIEW OF CHU.

98. It is my opinion that the combination of Duvall and Chu teaches, suggests, or renders obvious every limitation of claims 1-7 and 15-17, that a POSA would have been motivated to combine these prior-art references, and that a POSA would have had a reasonable expectation of success in combining these prior-art references. Accordingly, claims 1-7 and 15-17 would have been obvious to a POSA in view of the Duvall-Chu combination.

A. A POSA would have been motivated to combine Duvall with Chu.


99. Duvall is directed to “filtering messages transmitted between the Internet and a client computer” to ensure content that implicates security concerns does not reach recipients. EX1004, 1:7-24. Duvall’s techniques prevent illegal transmission by “allow[ing] Internet users to filter material transmitted over the Internet.” *Id.*, 1:27-29.

100. Duvall’s Figure 1 illustrates an example “network with a client computer for accessing the Internet.” EX1004, 2:23-24.



EX1004, FIG. 1.

101. With respect to Figure 1, Duvall discloses that “[a] user uses a computer which serves as a client 10 to communicate with other computers over Internet 12.” EX1004, 2:34-35. In an embodiment, “[a] filtering system resides in client computer 10” and includes a “filter database 24.” This is shown in Figure 2. *Id.*, 3:41-45.




EX1004, FIG. 2.

102. In another embodiment, “the filtering system can be provided from a server 30 that is on the client’s own network 40.” EX1004, 8:18-21. Duvall explains that “[s]uch a topology may be more desirable, for example, in a corporate network that has a limited size and a uniform set of filters for many users.” *Id.*, 8:21-23. Duvall also explains that “[t]his version of the filtering system uses the same type of filter database as a client-based filtering system, but the filter database is located on server 30.” *Id.*, 8:23-26.

103. Duvall discloses filtering techniques that include ALLOW filters, BLOCK filters, and deferred action filters. EX1004, 1:59-64, 3:64-66. An ALLOW filter *allows* data transmissions to pass through. *Id.*, 3:64-66. A BLOCK filter *blocks* data transmissions from passing through. *Id.* And a deferred action filter can specify additional criteria to filter data transmissions that are neither ALLOWED nor BLOCKED. *Id.*, 4:65-5:19.

104. Duvall teaches that its filters are implemented in sequence. Specifically, “the system searches ALLOW filters first, BLOCK filters next, and deferred action filters last.” *Id.*, 4:22-5:1. Duvall discloses that certain data transmissions may not match any ALLOW or BLOCK filters, leaving post-filtering residue. Referring to Figure 3 (reproduced below), Duvall’s system (1) at block 104 “check[s] to determine if any [filters] require immediate action, i.e., if unconditional allowing or blocking is required (steps 104 and 106)” and (2) if allowing and


blocking filters are not applied (i.e., there is residue data), “it is determined whether a deferred action must be taken” at block 116 (i.e., further analysis). *Id.* at 4:38.



EX1004, FIG. 3 (annotated).

105. The data transmissions that do not match the immediate ALLOW or BLOCK filters in step 104 (i.e., the residue data) are further analyzed by the system in “deferred filters.” EX1004, 5:8-29. This is shown by the example in FIG. 4 of Duvall, which discloses techniques for analyzing residue data. EX1004, 5:20-29. For instance, such analysis can include performing keyword matching and pattern

matching, to determine what to do with residue data. *Id.* The residue data can be blocked, discarded, replaced, or edited depending on the results of the matching. *Id.*, 5:30-51.



EX1004, FIG. 4.

106. In Duvall's system, residue data may not trigger a deferred action and the system may therefore default to allowing or blocking the residue data.

Specifically, Duvall discloses that “[i]f there is no deferred action, the system can

default to allow the transmission (step 108 [of FIG. 3]), or it can default to block the transmission (steps 110 and 112).” EX1004, 5:1-3.


107. A POSA would have recognized that applying a default action of allow or block to residue data could lead to inefficiencies. Indeed, Duvall itself provides that such default action could result in a **“filtering system [] filtering too much or too little.”** EX1004, 8:6-8. For example, a POSA would have recognized that a default action for residue data (i.e., data that matches neither an ALLOW or BLOCK filter at step 104 of Figure 3) can lead to allowable data being blocked, or more importantly, harmful data being allowed through the system. A POSA would have recognized that such results are disadvantageous. Duvall itself provides an “editing Manager” that “allows the user to make custom changes **if the user believes that the filtering system is filtering too much or too little.**” EX1004, 8:5-7. Thus, Duvall recognizes that its filtering techniques could potentially result in inaccuracies and provides a mechanism to combat this. A POSA with knowledge of Duvall’s system would have sought to ensure that the system correctly filters data on a consistent basis. Thus, a POSA would have been motivated to seek additional techniques for more accurately resolving the status of residue data transmissions to ensure transmissions are correctly blocked or allowed.

108. In my opinion, this would have led a POSA to search for and locate Chu, which is in the same field of endeavor and addresses similar problems as

Duvall. Specifically, Chu discloses techniques for determining whether to block or allow residue data. EX1005, 9, 44. Chu is directed to protecting users from untrustworthy information available on the Internet (same as Duvall) and provides examples of trust-management issues that may occur, for instance, when a user attempts to download software (e.g., a game) from a Web page. *Id.*, 9, 23. Chu explains that the user could have “concerns that prevent her from downloading the game on her machine,” such as “[d]oes this game contain a virus that would erase her hard drive? (security issue).” *Id.* To address these concerns, Chu discloses various trust policies allowing computer systems to make trust decisions with regard to user data (e.g., Web requests) transmitted over the Internet. *Id.*, 28, 35-43.

109. For example, Chu discloses example policies addressing the question of “should I download the active content at this URL,” which is similar to the problem of allowing Web requests in Duvall. EX1005, 43-48. Similar to Duvall, Chu’s example policies make use of a “blacklist” that specifies “sites or directories [the system] should not download codes from,” and a “whitelist” that specifies “sites know[n] to be trustworthy.” *Id.*, 44. The blacklist and whitelist are designed to prevent transmissions from “Websites that serve potentially dangerous” content. *Id.* Like Duvall’s BLOCK filter, an example Chu policy (shown below) *blocks* downloads from URLs included in the blacklist. Like Duvall’s ALLOW filter, Chu’s policy *allows* downloads from URLs included in the whitelist. EX1005, 44. If

a URL is unknown (i.e., it's not in the blacklist or whitelist), Chu's system notifies the user so that the user can review the URL. *Id.*, 44.



EX1005, 44 (annotated).

110. Thus, while Chu discloses that “[t]he blacklist and whitelist ensure good automation of the trust decision process if the lists are reasonably complete,” Chu also recognizes that data transmissions not matching the blacklist or whitelist need further analysis and human intervention. EX1005, 44. Chu explains that a user can be notified for intervention when the transmission does not meet the criteria of either the blacklist or whitelist. EX1005, 44 (“User intervention is needed only when the given URL is in neither the blacklist nor the whitelist.”).

111. In my opinion, a POSA would have been motivated to incorporate Chu's user intervention into Duvall's system for data transmissions not matching any of Duvall's ALLOW or BLOCK (i.e., positive or negative) filters at step 104, nor resolved by Duvall's post-filtering residue analysis, as depicted in FIG. 4. For

instance, motivation is provided by Duvall's express disclosure of "allow[ing] the user to make custom changes [to filters] if the user believes that the filtering system is filtering too much or too little." EX1004, 8:5-7. As I explained above, a POSA would have sought to ensure that Duvall's filtering system accurately and consistently filters data (i.e., that it does not filter too much or too little).

112. A POSA would have understood that, in Duvall's corporate environment, where a "uniform set of filters for many users" is provided, the person making the changes to these filters would be a person trained in network administration and security, not each individual employee at their own computer. EX1004, 8:18-26. For example, U.S. Patent No. 5,261,044 to Dev *et al.* ("Dev"), which issued on November 9, 1993, explains that "[n]etwork management systems," like Duvall's "editing Manager," "have been utilized in the past in attempts to address" a number of network-related issues. EX1038, 1:35-46. Dev explains that the person responsible for accessing and using the information from such network management systems would have been a "network administrator," not a typical user:

The network management system collected large volumes of information which required evaluation by *a network administrator*. . . . *He must be a networking expert* in order to understand the implications of a change in a network device parameter. *The administrator must also understand the topology* of each section of the network in order to understand what may have caused the change. In addition, *the administrator must sift through reams of information* and false alarms in order to determine the cause of a problem.

Id., 1:49-60; *see also* U.S. Pat. No. EX1039, 2:16-20 (“Older [network] management systems typically operated by collecting large volumes of information which then required evaluation by **a network administrator**, and thus placed a tremendous burden on and required **a highly-skilled network administrator**.”).

113. In addition, a POSA would have further understood that in Duvall’s corporate environment the responsibility for making changes to the filters via Duvall’s “editing Manager” would have been advantageously outsourced to a third-party network-security service because each individual corporation “may not be able to staff and maintain [its] own center of [network-security] competency.” EX1037, *Intrusion Detection: Network Security Beyond the Firewall*, 1998, Terry Escamilla (Wiley), ISBN 0-471-29000-9, 224-27. For example, a 1998 book explains that IBM previously provided an outsourced network-security service:

IBM’s Emergence Response Center offers a fee-based service with NetRanger. Instead of staffing your own team of security experts, you can use IBM’s strength in this area. **A network operations center** is staffed 7 X 24, and **a specific expert is assigned to your account**. When an event is detected, IBM’s security experts notify you and help you respond to the event. Up-front planning and response policy design also are available. As hackers become more sophisticated, outsourcing your network intrusion detection seems attractive because you may not be able to staff and maintain your own center of competency.

Id.

114. Combining Duvall and Chu as I have described would amount to nothing more than incorporating well-known techniques taught by Chu (providing

information to a human user for analysis) into Duvall's known system. Accordingly, it is my opinion that a POSA would have had a reasonable expectation of success implementing these modifications. As I have explained, Duvall already suggests reviewing system activity to determine whether "the system is filtering too much or too little," EX1004, 8:5-7, and further "provides a screen on the user's display" for viewing information about data transmissions. *Id.*, 4:61-64. Moreover, Duvall and Chu are both designed to filter transmissions based on similar information, such as a URL or IP address. *Compare* EX1004, 4:38-40 ("the filtering system can filter outgoing messages **on the basis of the IP address** of the host server that is being contacted by the client.") *with* EX1005, 44 ("This policy determines access based on [a] requested URL."). Thus, A POSA would have recognized that modifying Duvall's system to provide information to a user about data transmissions not matching any of Duvall's filters would require only minor changes to Duvall's overall process.

115. A POSA would have been further motivated to combine Duvall and Chu because both are in the same field of preventing unsecure data transmissions over the Internet generally and the Web specifically. EX1004, 1:7-24; EX1005, 9, 44; Section VII.F. Both references disclose techniques for determining whether to allow or block data transmissions, and Chu further discloses advantageous techniques for enabling a human user to properly address residue transmissions.

Thus, it is my opinion that a POSA would have been motivated to combine Duvall and Chu with a reasonable expectation of success, and that the results of the combination would have been predictable.

B. Independent Claim 1

116. As demonstrated below, it is my opinion that the Duvall-Chu combination renders obvious claim 1. I reproduce claim 1 below and assign the elements of claim 1 labels, such as [1-pre], [1A], [1B], *etc.*

Element	Claim Language
[1-pre]	A system for operating a probe as part of a security monitoring system for a computer network, the system comprising:
[1A]	a) a sensor coupled to collect status data from at least one monitored component of the network;
[1B]	b) a filtering subsystem coupled to analyze status data to identify potentially security-related events represented in the status data, wherein the analysis includes filtering followed by an analysis of post-filtering residue, wherein the post-filtering residue is data neither discarded nor selected by filtering;
[1C]	c) a communications system coupled to transmit information about the identified events to an analyst system associated with the security monitoring system;
[1D]	d) a receiver for receiving feedback at the probe based on empirically-derived information reflecting operation of the security monitoring system; and

Element	Claim Language
[1E]	e) a modification control system for dynamically modifying an analysis capability of the probe during operation thereof based on the received feedback.

1. [1-pre] “A system for operating a probe as part of a security monitoring system for a computer network”

117. To the extent the preamble is limiting, Duvall discloses the preamble of claim 1. Duvall discloses “*a security monitoring system for a computer network.*” Duvall is directed to “filtering messages transmitted between the Internet and a client computer” to ensure content that implicates security concerns does not reach recipients. EX1004, 1:7-24. For example, Duvall’s techniques prevent illegal transmission (e.g., under the “Communications Decency Act” (CDA)) of “objectionable material, especially sexually explicit materials,” by “allow[ing] Internet users to filter material transmitted over the Internet.” *Id.*, 1:12-20, 1:27-29. That is, Duvall’s system secures the client computer against receiving such objectionable material. What’s more, a POSA would have understood that such objectionable material often presented further security risks to recipients and recipients’ organizations. For instance, it was common for computer viruses to use objectionable materials to lure users into downloading the virus. *E.g.*, EX1026, Sinrod, 218 (describing the infamous “Melissa” virus); EX1030, Coates, 4 (same); EX1028, Rothfuss, 5 (describing a virus propagated via a website called “sexygirls.com”). By securing the client computer against objectionable materials,


Duvall's system also secures the client computer against receiving viruses, worms, malware, etc. associated with the objectionable materials.

118. The '641 patent recites that “exemplary implementations of the present invention [are] directed at monitoring, detecting, and responding to computer network intrusions.” EX1001, 15:59-62. More broadly, the '641 patent explains that “the present invention is usable generally for (i) monitoring of any system; (ii) filtering data from the monitoring for noteworthy events . . . (iii) alerting human beings to any such noteworthy events, and (iv) providing for a response by human beings to any such alert.” *Id.*, 15:62-67. In any event, if Patent Owner argues that monitoring of objectionable, illegal material is not expressly “security monitoring,” a POSA would have been motivated to also block other security-implicating content, such as network attacks, from reaching a user. For example, a POSA would have recognized that data transmissions from unknown or questionable sources can carry viruses or malware, as noted, for example, by Chu, which discusses concerns of downloading content from Internet sources, such as whether the downloaded content may “contain a virus that would erase her hard drive” or “secretly collect information from [the user's] computer.” EX1005, 23, *see also* EX1014, ¶4 (“Even an innocent act by a company employee, such as the downloading of a virus-ridden file from a Web site, can have devastating effects”); EX1026, 218; EX1030, 4; EX1028, 5. Duvall also discloses that its filtering system can be “incorporated into a

firewall,” such as in a “corporate network.” EX1004, 1:58-64, 8:18-23. A POSA would have understood that firewalls are commonly used for protecting networks against external attacks and intrusion, as well as for blocking unwanted traffic data. EX1027, Guttman, 3.

119. Duvall’s filtering system monitors transmissions for questionable content that should be blocked. *See* EX1004, 3:33-37, 5:8-15, 6:10-42. In addition, a POSA would have been motivated to block content that may carry viruses or malware. Implementing Duvall’s system to additionally block content could be achieved without structural modifications to Duvall’s system. Indeed, domains (e.g., URLs or IP addresses) believed to carry security-implicating content (e.g., viruses) would simply be included in Duvall’s blocking filters. *See, e.g.*, EX1004, 6:10-27. In other words, Duvall’s BLOCK filters would be implemented to include URLs and/or IP addresses matching URLs and/or IP addresses that provide such security-implicating content. This would be readily achievable by a POSA. Duvall already discloses “an updating mechanism” for enabling the filtering system to download “a number of new sets of filters.” *Id.*, 7:18-28. Duvall also discloses an “editing Manager that allows the user to edit the database to add, delete, or modify filters in the database.” *Id.*, 8:2-16. Duvall explains that such editing techniques were “known.” *Id.*, 8:10-11. Thus, Duvall explicitly discloses enabling users to select

which materials are blocked by the filtering system (i.e., security-implicating content).



EX1004, FIG. 1 (annotated).

120. Duvall’s “filtering system can be provided from a server 30 [(i.e., “a probe”)].” EX1004, 8:18-21; *see also id.*, 1:60-64, FIG. 1 (reproduced above).


Duvall’s server 30 is “a probe.” As noted in Section VI.F, the district court previously construed “a probe” as “a component that collects data from one or more network components to which it is attached, filters or otherwise analyzes the data that has been collected, transmits noteworthy information, and receives feedback in order to update its capabilities of analysis.” EX1013, 2. Duvall’s server 30 collects and analyzes data from other network components to which it is coupled, such as clients 10 in Figure 1, and as I discuss further for element [1A]. Duvall explains that

“[s]uch a topology may be [] desirable, for example, in a corporate network that has a limited size and a uniform set of filters for many users.” EX1004, 8:21-23.

Further, Duvall’s filtering system may be part of a firewall—i.e., it collects data from the firewall. EX1004, 1:59-64. And as I explain for elements [1C] and [1D] below, Duvall’s server 30, alone or in combination with Chu, transmits noteworthy information and receives feedback in order to update its capabilities of analysis. Accordingly, Duvall discloses a “*system for operating a probe as part of a security monitoring system for a computer network.*”

2. [1A] “a sensor coupled to collect status data from at least one monitored component of the network”

121. Duvall’s server 30 analyzes data received from clients 10, which reside on the same network (i.e., “*at least one monitored component of said network*”).



EX1004, FIG. 1 (annotated).

122. Specifically, Duvall discloses, “[a] user uses a computer which serves as a client 10 to communicate with other computers over Internet 12.” EX1004, 2:35-37. “Messages are transmitted over the Internet among computers with a transport protocol known as Transmission Control Protocol/Internet Protocol (TCP/IP).” EX1004, 2:42-44. Duvall also discloses that “[w]hen a message is transmitted, whether that message is incoming or outgoing with respect to the client computer, the filtering system compares the IP address and/or other information in the data stream to the filter entries stored in the database to determine whether some action needs to be taken.” *Id.*, 4:22-27. And Duvall’s “filtering system can be provided from a server 30 that is on the client’s own network.” *Id.*, 4:18-21. Thus, Duvall discloses “*collecting ... data from at least one monitored component of said network,*” for instance, by collecting information from data streams communicated to and from client computers 10. *Id.*, 4:22-27. The clients 10 are “*monitored components*” because the filtering system of the server 30 monitors and collects incoming and outgoing messages transmitted by or intended for the clients 10 to “determine[] whether to block or allow” the messages. *Id.*, Abstract, 4:22-27, 8:18-21. A POSA would have also further understood that Duvall’s server 30 acts as “*a sensor coupled to collect . . . data from at least one monitored component of said network*” because it collects data in that is in the process of being transmitted or received, e.g., “[w]hen a message is transmitted.” *Id.*, 4:22-27

123. The information received by Duvall’s filtering system is “*status data*.”

As I noted in Section VI.F, the district court construed “*status data*” as “data extracted from or generated about the traffic or system processing it that is informative as to the status of the network and its components.” EX1013, 1. Duvall explains that the received information of the data stream can include, among other things, “a particular port and IP address” with which a client is attempting to communicate, as well as protocol information, URLs, and associated commands (e.g., an HTTP “GET command”). EX1004, 4:39-42, 5:66-6:27; *see also id.*, 1:30-35 (“filtering system compares *portions of incoming and/or outgoing messages* with filtering information”). This information is “data extracted from . . . the traffic or system processing it.” EX1013, 1. The information is also “informative as to the status of the network and its components” because it provides information regarding the activity of clients 10, the type, protocol, destination, domains, and content of requests and responses flowing through the network, and the entities with which clients 10 are communicating (and vice versa).

124. Indeed, I understand that the district court found the same data disclosed in Duvall to be “*status data*,” explaining that the “items listed in Table 6 [of the ’641 patent], including **an IP address**, are examples of ‘status data.’” EX1034, 13 n.47. Specifically, Table 6 of the ’641 patent shows “information that might be included” in a “sentry message,” such as “interesting status data.” EX1001,

8:60-9:1. Table 6 includes entries for IP addresses of computing systems within the network (e.g., “Victim IP”—“IP address of the Victim device”). This is akin to Duvall’s “particular port and IP address” collected from the client 10.

3. [1B] “a filtering subsystem coupled to analyze status data to identify potentially security-related events represented in the status data, wherein the analysis includes filtering followed by an analysis of post-filtering residue, wherein the post-filtering residue is data neither discarded nor selected by filtering”


a) Duvall discloses “a filtering subsystem coupled to analyze status data..., wherein the analysis includes filtering.”

125. Duvall discloses “*a filtering subsystem*” residing on its server 30:

“Referring again to FIG. 1, . . . the filtering system can be provided from a server 30 that is on the client’s own network 40.” EX1004, 8:18-22. Duvall discloses that its filtering subsystem “*analyz[es] status data*,” and that this “*analysis includes filtering*.” Specifically, Duvall uses “direct action” filters, which “indicate that the system should unconditionally allow or unconditionally block [a] transmission.” EX1004, 4:12-21, 5:8-29.

126. For example, Duvall filters based on IP address (i.e., “*status data*”) and other information, such as a port number. Specifically, Duvall explains that “filters in the database preferably have fields for port number and IP address.” EX1004, 3:66-4:1. “When a message is transmitted, whether that message is incoming or outgoing with respect to the client computer, the filtering system compares the IP address and/or other information in the data stream to the filter entries stored in the


database to determine whether some action needs to be taken.” EX1004, 4:22-27;
see also id., 1:30-35 (“filtering system compares *portions of incoming and/or outgoing messages* with filtering information”). Duvall’s “filter database,” which is shown in Figure 2, “is located on server 30” in Duvall’s server-based embodiment.
Id., 8:18-26, FIG. 2.



EX1004, FIG. 2 (excerpt).

127. Duvall explains that “[t]he filter database has lists of filters, some of which are identified as either ALLOW filters or BLOCK filters for respectively allowing or blocking transmission.” EX1004, 3:64-66. “Each filter entry in the filter database [] has a field for specifying an action to be taken by the client if that filter were retrieved,” and “[t]hese actions are essentially divided into two groups, direct action or deferred action.” *Id.*, 4:12-15. Duvall also explains, “[t]he filters are preferably stored so that the system searches ALLOW filters first, BLOCK filters next, and deferred action filters last,” where the “ALLOW” and “BLOCK” filters are direct action filters. *Id.*, 4:27-29.

128. As described for Figure 3, which I have reproduced below, Duvall uses “*filtering*” to determine whether to allow or block data transmissions: “[t]he filtering system detects when the client opens the data stream for a particular port and IP address (step 100), and searches the filter database for matching filters.” EX1004, 4:37-42. For instance, the system retrieves matching filters by comparing the port and IP address of the data transmission. *Id.*, 4:39-49. Then “[t]he retrieved filters are checked to determine if any require immediate action, i.e., if unconditional allowing or blocking is required (steps 104, 106).” *Id.* If so, immediate action is required, and the action specified by the matching filter “is to allow the transmission, the system allows the transmission to proceed normally (step 108).” *Id.*, 4:49-51. If the action specified by the matching filter “is to block the transmission, the system takes action to block the transmission (step 110).” *Id.*, 4:51-54.




EX1004, FIG. 3.

b) Duvall discloses “an analysis of post-filtering residue, wherein the post-filtering residue is data neither discarded nor selected by filtering.”

129. The plain language of this element does not require any data to actually be “discarded” or “selected by filtering”—only that the “*post-filtering residue is data neither discarded nor selected by filtering.*” EX1001, Claim 1. As I explained in Section VI.F, the parties in the BT-Fortinet litigation agreed that “*post-filtering*

residue, wherein the post-filtering residue is data neither discarded nor selected by filtering” should be construed as “status data that undergoes negative and positive filtering, but is neither discarded by such negative filtering nor selected by such positive filtering.” EX1013, 4. It is my opinion that, under either this agreed-upon construction or the claim’s plain language, Duvall teaches “*an analysis of post-filtering residue, wherein the post-filtering residue is data neither discarded nor selected by filtering.*”

130. Duvall’s filtering system results in data that can be (i) discarded (i.e., blocked), (ii) selected (i.e., allowed), and (iii) neither discarded (blocked) nor selected (allowed). Specifically, Duvall discloses ALLOW filters and BLOCK filters, which are implemented in sequence “so that the system searches ALLOW filters first [and] BLOCK filters next.” EX1004, 4:27-30. These filters are shown in Duvall’s Figure 3 below.



EX1004, FIG. 3 (annotated).

131. Referring to Figure 3 above, at steps 104, 106, Duvall's system "checks for and retrieves any filters that match the particular IP address. The retrieved filters are checked to determine if any require immediate action, i.e., if unconditional allowing or blocking is required." EX1004, 4:46-50. Duvall explains that "[w]hen filter entries are retrieved, they are first scanned for entries that require direct action." *Id.*, 4:17-20. As I explained above, Duvall teaches that ALLOW filters are searched and scanned first, and then the BLOCK filters are searched and scanned

second. *Id.*, 4:27-30. If direct action is required, “these actions are carried out immediately.” *Id.*, 4:17-20.

132. That is, if an ALLOW or BLOCK filter matches the data transmission (i.e., if immediate action is required), then one of two things happens:


- If the data is allowed, “the system allows the transmission to proceed normally (step 108).” *Id.*, 4:50-52.
- If the data is blocked, then “the system takes action to block the transmission (step 110).” *Id.*, 4:52-54.

133. Consistent with the plain and ordinary meaning of the term in the art, data that is blocked is discarded data. Duvall explicitly states that “the system can block the incoming data by discarding it.” EX1004, 6:26-27; *see also id.*, 1:35-40 (“blocking options” include “discarding incoming data”), 5:35-40, claim 20. This data blocking comprises “negative filtering” in accordance with the BT-Fortinet litigants’ agreed-upon construction. Data that is allowed is selected data. *Id.*, 4:49-54. This data selection comprises “positive filtering” in accordance with the BT-Fortinet litigants’ agreed-upon construction.

134. If immediate action is required (data is neither blocked nor allowed at block 104 (i.e., the data is residue data)), the data is passed for further analysis. This residue data is status data that undergoes negative and positive filtering but is neither discarded by such negative filtering nor selected by such positive filtering—

i.e., “post-filtering residue.” This is because Duvall’s filtering system first scans the ALLOW filters to determine if the data transmission matches, and then scans the BLOCK filters to determine if the data transmission matches. If the data does not match any direct action filters (ALLOW or BLOCK), then no direct action is taken on the data. Such data is thus “*post-filtering residue*.”

135. The “*post-filtering residue*” is then analyzed using Duvall’s deferred action filters. EX1004, 4:27-30 (“The filters are preferably stored so that the system searches ALLOW filters first, BLOCK filters next, and deferred action filters last, and takes action based on the first matched filter.”), 4:65-5:7. This post-filtering analysis includes keyword matching and pattern matching, as depicted in Figure 4. *Id.*, 5:8-29.



EX1004, FIG. 4 (annotated).

136. Therefore, Duvall discloses the “*post-filtering residue*” elements of claim 1.

c) Duvall’s analysis is to “identify potentially security-related events represented in the status data.”

137. Duvall’s filtering is to “identify potentially security-related events represented in the status data” because the filtering ensures that content that implicates security concerns does not reach recipients. EX1004, 1:7-24. As I discussed for [1-pre], data transmissions blocked by Duvall’s filtering system, as

well as those not matching (or completely matching) any filters, are “*potentially security-related events represented in the status data*” because they may relate to requests for, or transmissions of, “indecent material,” which may be illegal (e.g., “outlaw[ed]”) and/or threaten the security of the requesting user or client device. *Id.*, 1:7-24; *see also id.*, 4:61-64 (“The system also provides a screen on the user’s display indicating that the transmission was blocked, and advising the user on how to get back to the state before the user tried to open the stream and send the message.”); EX1005, 23; EX1026, 218; EX1030, 4; EX1028, 5; *supra* Section VIII.B.1. Thus, Duvall performs analysis to identify these “*potentially security-related events*” through the additional analysis of residue data, as I discussed in the previous section. Section VIII.B.3(b).

138. If Patent Owner argues that identification of illegal or indecent material in a data transmission is not a “*potentially security-related event*” (which it is), it still would have been obvious to apply Duvall’s and Chu’s techniques to filter and identify data transmission representing “*potentially security-related events*.” This is taught, for example, by Chu’s blacklists, which contain lists of “Web sites that serve potentially dangerous active codes.” EX1005, 44; *see also id.*, 23 (discussing virus-ridden downloadable content); *see also* EX1014, ¶4 (“Even an innocent act by a company employee, such as the downloading of a virus-ridden file from a Web site, can have devastating effects”); EX1026, 218; EX1030, 4; EX1028, 5. Duvall’s and

Chu's filtering techniques would apply equally well in other security contexts, such as malware or intrusion detection, without needing any structural modifications. *See* Sections VIII.A, VIII.B.1. In such implementations, Duvall's filters would simply include, for example, IP addresses or other criteria (e.g., URLs) associated with malware or other potentially security-related content. *See* Section VIII.A. Applying Duvall's techniques in this manner would amount to nothing more than use of known technique to improve similar devices, methods, or products in the same way (e.g., to detect malware or a potential intrusion instead of objectionable material). *See id.*

4. [1C] “a communications system coupled to transmit information about the identified events to an analyst system associated with the security monitoring system”

139. Duvall discloses this element.

140. With respect to Figure 4, Duvall acknowledges that, in some cases, its *post-filtering residue* analysis may not be able to keyword match the data at step 134 or pattern match the data at step 138. EX1004, FIG. 4. In such case, Duvall discloses taking a default block or allow action. EX1004, 8:6-8; Section VIII.A. However, Duvall nonetheless recognizes that default blocking or allowing could result in “filtering too much or too little.” *Id.*

141. To address this concern, and to ensure that Duvall's system accurately and consistently filtered data, a POSA would have been led to search for and locate

Chu to optimize the handling of such unresolved residue data. *See* Section VIII.A.

When Chu's system encounters unresolved data, the system notifies a human user so that the user can intervene to review the unresolved data. This user intervention technique is illustrated in the following highlighted excerpt from Chu.

Profiles-0.92 is capable of taking the policy a step further by returning *unknown* if the requested URL is neither in the blacklist nor in the whitelist (sites know to be trustworthy):

Policy in English

Do not download the code if the URL is served from Harvard or CalTech Web servers. Download it automatically if served from MIT. Prompt me for my attention otherwise.

Profiles-0.92

```
(threshold-and
  2
  (not (url-match URL ("http://www.harvard.edu"
                      "http://www.caltech.edu")))
  (url-match URL ("http://web.mit.edu"))
  unknown)
```

The blacklist and whitelist ensure good automation of the trust decision process if the lists are reasonably complete. User intervention is needed only when the given URL is in neither the blacklist nor the whitelist.

EX1005, 44 (highlighted)

142. In the Duvall-Chu combination, the human intervention would be performed by an analyst system having people trained to manage Duvall's corporate network. Duvall's system is appropriate for a server-based system in, for example, "a corporate network that has limited sized and a *uniform set of filters for many users*." EX1004, 8:18-23. Because a corporate environment requires a "uniform set of filters," the decision of how to handle unresolved data would not be left to each user individually, but rather, to a group of people trained in making such decisions. *See above* ¶111; *see also* EX1038, 1:49-60 (explaining that "network management systems," like Duvall's "editing manager," "collect[] large volumes of information

which required evaluation by *a network administrator*"); EX1039, 2:16-20 (explaining that "network management systems," like Duvall's "editing manager," "typically operated by collecting large volumes of information which then required evaluation by . . . *a highly-skilled network administrator*"). This is evidenced by the fact that Duvall's "editing manager" is "password protected," indicating that the persons editing filters in a corporate network have access to the password and are trained to manage the system. EX1004, 8:8-10. Moreover, the responsibility for editing the filters in a corporate network would have been advantageously outsourced to a third-party network-security service because each individual corporation would not likely be able to staff and maintain its own center of network-security competency. *See* EX1037, 224 (explaining that IBM previously provided an outsourced network-security service). Thus, in the Duvall-Chu combination, information about the unresolved residue data would be sent to a trained network analyst for handling the unresolved residue data.

143. A POSA would have been motivated, with a reasonable expectation of success, to incorporate Chu's user intervention techniques into Duvall's system because, as Duvall itself recognizes, Duvall's residue analysis may result in a default allow/block, causing the system to filter too much or too little. EX1004, 8:3-10 (Duvall's "editing Manager" "allows the user to make custom changes [to the filters] if the user believes that the filtering system is filtering too much or too

little.”). A POSA would have understood that, in the corporate context of Duvall, the “user” of Duvall’s “editing manager” would be a trained professional associated with the corporate network and filtering system, not the user trying to access the illicit content—particularly because Duvall teaches that in the corporate context the filters are uniform across the users. *See above* ¶111; *see also* EX1038, 1:49-60 (explaining that “network management systems,” like Duvall’s “editing manager,” “collect[] large volumes of information which required evaluation by **a network administrator**”); EX1039, 2:16-20 (explaining that “network management systems,” like Duvall’s “editing manager,” “typically operated by collecting large volumes of information which then required evaluation by . . . **a highly-skilled network administrator**”); EX1037, 224 (explaining that IBM previously provided an outsourced network-security service employing its own team of security experts). Indeed, allowing the user who is trying to access illicit content to edit the filters would defeat the purpose of the filtering system. This group of trained professionals responsible for maintaining uniformity of the filters for a corporate network is “*an analyst system... associated with said security monitoring system.*” And as discussed above, Duvall discloses that the network administrator would be viewing the transmitted information on an “*analyst system associated with said security monitoring system,*” as required by the claim. EX1004, 4:61-64.

144. Moreover, even if not explicitly disclosed by Duvall, a POSA would have been motivated to assign the role of managing and editing filters to a group of people who have the expertise to properly evaluate whether “the filtering system is filtering too much or too little,” rather than an individual client user who may be motivated to undermine the filtering system. EX1004, 8:6-8; *see also above* ¶111; EX1038, 1:49-60; EX1039, 2:16-20.

145. Thus, the Duvall-Chu combination also teaches element [1C].

5. [1D] “a receiver for receiving feedback at the probe based on empirically-derived information reflecting operation of the security monitoring system”

146. The Duvall-Chu combination discloses this element. As I discussed for element [1C], in the Duvall-Chu combination, information about unresolved residue data is sent to an analyst for review. Duvall-Chu discloses feedback in the form of an “editing Manager” that allows a trained network analyst “to edit the database to add, delete, or modify filters in the database.” EX1004, 8:2-5. A POSA would have recognized that Duvall’s editing manager thus provides “*a receiver for receiving feedback . . . based on empirically-derived information reflecting operation of said security monitoring system*” because the information is only provided to the trained network analyst if the data is unresolved residue data.

147. That is, the data is provided to the network analyst only if the data did not match any allow/block filters and passed through subsequent residue analysis

(e.g., keywords and pattern matching) without any resolution. *Id.*, 5:8-18. Thus, decisions made by the trained network analyst (e.g., whether the filters should be modified) would be based on observable information about the data (e.g., transmission path, URL, etc.). The trained network analyst is enabled to make “custom changes” based on observation, experience, or belief “that the filtering system is filtering too much or too little.” *Id.*, 8:2-9. *See* EX1029 (defining “empirical” as “originating in or based on observation or experience” and “relying on experience or observation alone often without due regard for system and theory”).

148. Feedback is received “*at the probe*” because Duvall’s editing manager is part of the filtering system implemented on server 30. *See* EX1004, 8:2-5 (“the filtering system has an editing Manager”); *see also id.*, 8:18-21. Moreover, a POSA would have been motivated to send the feedback to the “*probe*” because it was known (i) to use sensing devices, like the claimed “*probe*,” to monitor network traffic and (ii) to send the most up-to-date information about security threats to those sensing devices. For example, U.S. Patent No. 6,775,657 to Baker (“Baker”) explains that it was known to use “sensing enabled devices,” like the claimed “*probe*,” in network intrusion detection to monitor network traffic:

Network intrusion detection is a process that can identify and respond to misuse or policy violations on a network. By placing ***sensing enabled devices*** at determined points on . . . the network, network traffic can be monitored and compared against patterns or ‘signatures’

that represent suspicious activity, misuse, or actual attacks. These devices can send alerts to the security management system and, under appropriate circumstances, send commands directly to network equipment such as routers and firewalls, reconfiguring them to deny access to the attacker.

EX1036, 1:19-28. And U.S. Patent No. 6,484,315 to Ziese (“Ziese”) explains that it was known to send updates/feedback to the sensors in an intrusion detection system to reduce vulnerability to new types of attacks:

A problem with conventional intrusion detection systems is that when a new vulnerability, or type of attack on the network, is discovered, a new intrusion signature must be generated and installed for each intrusion detection system. As a result, unless a network administrator frequently checks for new signatures developed by an intrusion detection provider and installs the new signatures *for each sensor* in his or her system, the system will remain vulnerable to the new types of attack. . . .

[To address this problem, Ziese] provides a method and system for dynamically distributing intrusion detection and other types of updates in a network that substantially eliminate or reduce disadvantages and problems associated with prior methods and systems.

EX1015, 1:54-2:6.

6. [1E] “a modification control system for dynamically modifying an analysis capability of the probe during operation thereof based on the received feedback”

149. As discussed for element [1D], the Duvall-Chu combination discloses “*receiving feedback*” from an analyst through Duvall’s “editing Manager,” which “allows the [network administrator] to edit the database to add, delete, or modify filters in the database.” Section VIII.B.5; EX1004, 8:2-6. Editing the filters in Duvall’s filter database “*modif[ies] an analysis capability of said probe . . . based*

on said feedback” because Duvall compares “information in the data stream to the filter entries stored in the database to determine whether some action needs to be taken.” EX1004, 4:22-27; *see also id.*, 4:27-30, 4:40-43.

150. That is, Duvall’s system searches filters in the database when determining whether to allow or block a data transmission. The system analyzes these filters by comparing the filters to the incoming or outgoing transmissions. *Id.*, 1:30-35 (“filtering system compares portions of incoming and/or outgoing messages with filtering information stored in a filter database”). As filters are added, deleted, or modified, Duvall’s analysis reflects these updates, such that the analysis capability of the server 30 would be modified. *Id.*, 8:3-16.

151. Duvall explains that “[b]ecause Internet sites are being added to the Internet at a fast rate . . . the filtering system preferably also has an updating mechanism to keep filters current” and that “the system can adapt as new sites and servers are added to the Internet.” EX1004, 7:18-21; 2:16-18. In view of the fast rate of updates, a POSA would have recognized that Duvall’s “*analysis capability of said probe*” is “*dynamically modified . . . during operation thereof*” because taking the system offline each time an update was required would be disadvantageous. *See above* ¶148; *see also* EX1015, 1:54-2:6.

152. To avoid this disadvantage, Duvall “provid[es] updates online” so that “the system can adapt as new sites and servers are added to the Internet.” EX1004,

2:16-18. Duvall’s filtering system accommodates dynamic updates because it searches “filter entries stored in the database” when performing its analysis, EX1004, 4:22-30, which would reflect changes to filters as they are edited. *Id.*, 4:22-43; 8:3-16.

153. Duvall explains that the “implementation and use of the editing manager are generally based on known database editing techniques.” *Id.*, 8:10-11. A POSA would have understood that such known database editing techniques implemented “during a session over the Internet” enable “*dynamically modifying an analysis capability*” of the server 30. Indeed, a main benefit of databases is that entries can be edited or added without taking the system offline, which prevents the filtering system from being down during updates. Since Duvall’s system “searches the filter database for matching filters” when a new data stream is opened, the “*analysis capability of said probe*” always reflects the latest and current set of filters. *Id.*, 4:40-43.

154. Accordingly, the Duvall-Chu combination discloses element [1E].

C. Claim 2

155. Claim 2 recites, “*wherein the identifying step includes performing a multi-stage analysis of the status data.*”

156. As I discussed for element [1B], Duvall performs “*a multi-stage analysis of the status data.*” Section VIII.B.3. For example, Duvall searches filters

in multiple stages: “The filters are preferably stored so that the system searches ALLOW filters first, BLOCK filters next, and deferred action filters last.” EX1004, 4:27-30. Duvall then evaluates status data in multiple stages. Duvall first “determine[s] if any require immediate action,” and then “determine[s] whether a deferred action must be taken with respect to any of the retrieved filter[s].” *Id.*, 4:38-50, 4:65-5:1, 5:8-29 (discussing additional analysis performed for deferred actions).

157. With respect to Figure 4, Duvall further discloses an iterative analysis of residue transmission data, providing “*a multi-stage analysis of the status data.*” Moreover, Duvall in view of Chu teaches an additional analysis stage involving a user (or trained network professional) to further evaluate residue data transmissions. Section VIII.B.3; EX1005, 44.

158. Accordingly, the Duvall-Chu combination renders obvious claim 2.

D. Claim 3

159. Claim 3 recites, “*wherein the multi-stage analysis includes performing a discrimination analysis on the status data.*”

160. The ’641 patent explains that “[n]egative filtering, positive filtering, and residue analysis are examples of data discrimination analyses, other types of which are well-known to those skilled in the art.” EX1001, 8:57-59. Thus, “discrimination analysis,” as used by the ’641 patent, is merely analysis that

differentiates certain data from other data. The Duvall-Chu combination teaches such analysis.

161. As discussed for element [1B], the filtering system of the Duvall-Chu combination differentiates between data that should be allowed through the network and data that should be blocked. *See* Section VIII.B.3. This is consistent with the '641 patent's description of "discrimination analysis," which explains that "[n]egative filtering, positive filtering, and residue analysis are examples of data discrimination analyses, other types of which are well-known to those skilled in the art." EX1001, 8:57-59. And the purpose of the '641 patent's "negative filtering" and "positive filtering" is to determine whether to select or discard certain data. EX1001, 8:45-50 ("[D]ata is first filtered by negative filtering subsystem 2020, which discards uninteresting information, and then by positive filtering subsystem 2030, which selects possibly interesting information ...").

162. The Duvall-Chu combined system performs negative filtering, positive filtering, and residue analysis, as described for [1B] of claim 1. Section VIII.B.3. For example, Duvall's "ALLOW" and "BLOCK" filters are used, respectively, to select out data to allow through the network and discard data that the system determines to block, preventing the data from passing through the network. EX1004, 3:64-66, 4:12-15, 4:27-30, 4:49-54. And Chu similarly discloses negative filtering,

positive filtering, and residue analysis in the form of a “blacklist,” “whitelist,” and user intervention for “unknown” data not meeting either list. EX1005, 44.

163. Accordingly, the Duvall-Chu combination renders obvious claim 3.

E. Claim 4

164. Claim 4 recites, “*wherein the discrimination analysis includes positive filtering.*”

165. The Duvall-Chu combination renders obvious claims claim 4 for the reasons discussed for claim 3.

F. Claim 5

166. Claim 5 recites, “*wherein the discrimination analysis includes negative filtering.*”

167. The Duvall-Chu combination renders obvious claim 5 for the reasons discussed for claim 3.

G. Claim 6

168. Claim 6 recites, “*wherein the multi-stage analysis includes analysis at the probe and analysis at a secure operations center configured to receive data from the probe.*”

169. As I discussed for claim 2, Duvall’s “multi-stage analysis includes analysis at the probe,” including evaluating status data in multiple stages. Section VIII.C; EX1004, 4:27-30, 4:39-49, 4:65-5:1, 5:8-29. The Duvall-Chu combination teaches analysis at a probe followed by analysis at an analyst system. Sections

VIII.B.3, VIII.C; EX1005, 44. As discussed in further detail below, a POSA would have understood that the analyst system could be on-site or at a secure portion of the corporate network or secure remote site, i.e., “*a secure operations center.*” *See also* EX1037, 277 (“Instead of staffing your own team of security experts, you can use IBM’s strength in this area. A network operations center is staffed 7 x 24, and a specific expert is assigned to your account. When an event is detected, IBM’s security experts notify you and help you respond to the event. Up-front planning and response policy design also are available.”).

170. Duvall’s “filtering system can be provided from a server 30 that is on the client’s own network,” which is “desirable, for example, in a corporate network.” EX1004, 8:18-26. Moreover, Duvall discloses that in the corporate setting a “uniform set of filters for many users” is provided and Duvall further discloses that these filters can be modified with an “editing manager” that is “preferably password” protected. *Id.*, 8:2-9, 8:21-23. Thus, a POSA would have recognized that a trained network professional, not each individual employee, would have been tasked with ensuring that the Duvall-Chu combination was not “filtering too much or too little.” *See, e.g.*, EX1038, 1:49-60 (explaining that “network management systems,” like Duvall’s “editing Manager,” “collect[] large volumes of information which required evaluation by *a network administrator*”); EX1039 (explaining that “network management systems,” like Duvall’s “editing manager,” “typically

operated by collecting large volumes of information which then required evaluation by . . . *a highly-skilled network administrator*”); EX1037, 224 (explaining that IBM previously provided an outsourced network-security service employing its own team of security experts).

171. Further, Duvall’s filtering system can “be incorporated into a firewall [or] gateway.” EX1004, 1:60-64; EX1027, 2 (“A firewall is a safeguard one can use to control access between a trusted network and a less trusted one” and “A firewall enforces a security policy”); EX1031, Abie, 2 (“traditional network firewalls prevent unauthorized access and attacks by protecting the points of entry into the network”). A POSA would have recognized that accessing a firewall or gateway can be, and most typically is, accessed remotely, e.g., from a separate “network operations center.” *See, e.g.*, EX1037, 277. Further Duvall discloses that the filtering system is accessed through an “editing manager” that is “preferably password protected.” EX1004, 8:2-9, 8:21-23. Thus, Duvall recognizes the need for security. A POSA would have recognized that the security analyst system would likewise be secure, including communications between the analyst system and the firewall (*e.g.*, through HTTPS, SSL, TLS, etc.).

172. Moreover, it was common in the industry at the time to provide off-site security (*e.g.* hosted by a third party). *See, e.g.*, EX1037, 277. Indeed, Duvall discloses off-site operations where the “updating mechanism in the filtering system

can locate a particular HTTP or FTP update server 32 over the Internet . . . [T]he updating mechanism causes the filtering system to download from the update server . . . filters.” EX1004, 7:21-29. Customers of Duvall’s system are “given an opportunity to obtain a subscription . . . to obtain further [filter] updates.” *Id.*, 7:35-40. Thus, a POSA would have recognized that the trained professionals need not be located at the corporate firewall, but can be located remotely (possibly as a third party subscription service), in their own secure location (“*secure operations center*”) such as a network operations center (“NOC”). *See, e.g.*, EX1038, 1:47-49 (explaining that “network management systems typically operated *by remote access* to and monitoring of information from network devices”), 5:7-26 (explaining, with reference to FIG. 2, that a network management system monitors a number of buildings, including “building 52 remotely located (in a different city, state or country) from buildings 42 and 48”). For example, an analyst system would receive residue data, evaluate that data, update the respective filters, and provide those updates to the corporation per the terms of the subscription. *See, e.g.*, EX1039, 1:14-2:3 (describing situations where “a service company . . . provides network management services to several companies”); EX1037, 224 (explaining that IBM previously provided an outsourced network-security service employing its own team of security experts). A POSA would have recognized the updating mechanism would have used encryption to protect its transmissions (e.g., via HTTPS, SSL,

TLS, etc.). *See, e.g.*, EX1014, ¶14 (“One benefit of encrypting the packet stream is that it significantly hinders the covert manipulation . . . by unauthorized users (intruders, company employees, etc.).”).

173. Finally, to the extent not expressly disclosed in Duvall, a POSA would have been motivated to ensure that the trained network professionals managing Duvall’s filtering system were doing so as part of a “*secure operations center*.” Indeed, it was known that “network management systems,” like Duvall’s “editing Manager,” were “typically operated by *remote access* to and monitoring of information from network devices.” EX1038, 1:47-49; *see also id.*, 5:7-26 (describing an example where a network management system monitors “[a] building 52 remotely located (in a different city, state or country) from buildings 42 and 48”); EX1037, 224 (explaining that IBM previously provided an outsourced network-security service employing its own team of security experts). And it was also known that a user’s ability to access a network management system could be “limited by the access rights assigned by the network administrator.” EX1039, 6:17-19. Providing only limited, remote access to network management systems, like Duvall’s “editing Manager,” would ensure that unauthorized users could not edit, add, or remove filters that could potentially allow malicious traffic to enter the corporate network.

174. Accordingly, the Duvall-Chu combination renders obvious claim 6.

H. Claim 7

175. Claim 7 recites, “*wherein the identifying step includes aggregating and synthesizing the status data at the probe.*” The ’641 patent does not provide any details regarding such “*aggregating and synthesizing.*” That said, Duvall teaches this element in at least as much detail as the ’641 patent. As set forth above for element [1A], Duvall’s IP addresses comprise “status data.” Duvall also explains that, in the server environment, there may not be “ready access to the Internet’s domain name services.” EX1004, 8:28-33. To address this issue, Duvall discloses a “domain name table” for “rapidly converting a domain name into . . . an IP address.” *Id.* The “domain name table” includes a list of 32-bit values (IP addresses) that each correspond to a domain name. *Id.*, 8:34-38. “This list is searched associatively for an entry which matches the domain name being looked up.” *Id.*, 8:36-37. Duvall’s “domain name table” at the server (“*probe*”) aggregates the IP addresses (“*status data*”) by collecting these IP addresses in a list. And Duvall’s “domain name table” at the server (“*probe*”) synthesizes the IP addresses (“*status data*”) by making this list associatively searchable. Thus, Duvall teaches claim 7.

I. Claim 15

176. Claim 15 recites, “*instantaneously self-tuning the probe based on previously collected status data.*”

177. The Duvall-Chu combination teaches this element. As explained for [1E] of claim 1, Duvall’s filtering system always searches the latest and current set

of filters when performing its filtering analysis. Section VIII.B.6; EX1004, 4:22-43, 8:3-16. Thus, when filters are updated in Duvall's system, the analysis is "instantaneously self-tun[ed]" because Duvall is programmed to use the updated filters in subsequent analysis.⁴ Additionally, as discussed for elements [1B]-[1D], the "self-tuning" in the Duvall-Chu combined system is "based on previously collected status data" because filter edits are based on the user's analysis of previously collected information associated with potentially security-related events. See Sections VIII.B.3-5. Specifically, Duvall explains that the network administrator can "edit the database to add, delete, or modify filters in the database ... if the user believes that the filtering system is filtering too much or too little." EX1004, 8:2-15. Thus, the network administrator edits the database based on previous actions regarding previously collected status data (requested IP addresses). The updated filters ("self-tuning") are implemented based on these edits.

178. Accordingly, the Duvall-Chu combination renders obvious claim 15.

⁴ I note that, although software can perform operations quickly, these operations cannot literally be "*instantaneous*" because the software must execute some instruction or a series of instructions to perform an operation. The claims therefore must encompass a timeframe reasonable for software to perform the recited "*self-tuning*" operation.

J. Claim 16

179. Claim 16 recites, “*wherein the dynamic modifying step includes consideration of non-real-time information from ongoing security research efforts.*”

180. The Duvall-Chu combination teaches this element. As discussed for element [1D], the Duvall-Chu combination teaches “receiving feedback” in the form of “an editing Manager that allows the user to edit the database to add, delete, or modify filters in the database.” Section VIII.B.5; EX1004, 8:3-6. Duvall further discloses the “editing Manager” “allows the user to make custom changes if the user believes that the filtering system is filtering too much or too little.” EX1004, 8:3-8. Thus, a POSA would have understood that the Duvall-Chu combination prompts a user (e.g., a trained network professional) to consider past network activity (i.e., “non-real-time information”) to determine how to edit the filters. A POSA would have further understood that the consideration of past network activity is part of “ongoing security research efforts” because the activity analyzed by the user includes “potentially security-related events.” See Section VIII.B.3. Said differently, it would have been obvious for trained network professionals to consider this type of past network activity (“*non-real-time information*”) in the context of the Duvall-Chu combination because fee-based services for analyzing such information were well known well before the ’237 patent’s earliest priority date. *See, e.g.*, EX1037, 277

(describing a “fee-based service” in which “IBM’s security experts” could help respond to a previously detected security event).

K. Claim 17

181. Claim 17 recites, “*wherein the receiving feedback step occurs in substantially real time.*”

182. As discussed for [1C] and [1D] of claim 1, Duvall in view of Chu discloses prompting a user when a data transmission does not match any of Duvall’s filter criteria, and receiving feedback through Duvall’s editing manager. Sections VIII.B.4-5. Chu teaches that these prompts are received in real time to ensure untrusted content is not downloaded. EX1005, 44. And a POSA would have recognized that Chu’s user would respond to the prompt “in substantially real time” to ensure the system knew whether to block or allow the requested download. Similarly, a POSA would have been motivated to respond to a prompt regarding unknown data in Duvall’s system “in substantially real time” to ensure the filtering system does not continue to improperly block or allow data. Duvall itself suggests providing feedback “in substantially real time,” teaching the benefit of ensuring the filters are accurate to avoid “the filtering system [] filtering too much or too little.” EX1004, 8:5-7. And Duvall further provides mechanisms for alerting users of events of interest in real time, such as “provid[ing] a screen on the user’s display indicating that [a] transmission was blocked.” *Id.*, 4:61-64.

183. Accordingly, the Duvall-Chu combination renders obvious claim 17.

IX. GROUND 2: CLAIMS 8-13, 16 ARE OBVIOUS OVER DUVALL IN VIEW OF CHU, AND IN FURTHER VIEW OF TRCKA.

A. A POSA would have been motivated to combine the Duvall-Chu system with Trcka.

184. Duvall discloses an “editing Manager” that “allows the user to make custom changes if the user believes that the filtering system is filtering too much or too little.” EX1004, 8:3-8. Thus, Duvall teaches that a user (e.g., trained professional) reviews network activity to ensure the filters accurately allow and block data transmissions. *Id.* But neither Duvall nor Chu expressly discloses implementation details as to how the user performs this review. Duvall explains that such database updating techniques were “known.” *Id.*, 8:10-15. As evidenced by Duvall’s disclosure, a POSA with knowledge of with network security and filtering mechanisms would have understood numerous ways for performing this review. However, to the extent a POSA would have needed any additional guidance in practicing the teachings of the Duvall-Chu combination, a POSA would have sought references expressly disclosing such “known” techniques, such as Trcka.

185. Trcka is directed to “[a] network security and surveillance system [that] passively monitors and records the traffic present on a local area network, wide area network, or other type of computer network, without interrupting or otherwise interfering with the flow of the traffic.” EX1014, Abstract. Trcka’s system further includes “[a] set of analysis applications and other software routines allow[ing]

authorized users to interactively analyze the low-level traffic recordings to evaluate network attacks, internal and external security breaches, network problems, and other types of network events.” *Id.*; *see also id.*, ¶¶15, 16 (“The analysis applications can [] be used to view, analyze and process the traffic data.”), ¶53 (“[T]he system includes a powerful set of analysis tools . . . for allowing authorized users to perform interactive, off-line analyses of recorded traffic data.”). Trcka further provides a “graphical user interface (GUI)” through which “the user can launch and control the various analysis applications 100 through a common set of menus and controls.” *Id.*, ¶79.

186. In my opinion, a POSA would have been motivated to incorporate Trcka’s techniques for recording and viewing past network activity in Duvall-Chu’s system to enable the user to analyze whether “the filtering system is filtering too much or too little,” as taught by Duvall. EX1004, 8:5-7. Duvall already suggests that the user, such as a trained network professional, should have access to past network activity in order to accurately update filters. EX1004, 8:2-7. Trcka’s analysis applications would enable the user to interact with such information as part of the user’s analysis in ensuring that the filtering system is not filtering too much or too little. Indeed, Duvall performs filtering based on data contained in network packets, such as “port number and IP address,” EX1004, 3:64-4:55. Trcka similarly captures and presents packet-level data, including “headers of all packets transferred over the

network,” EX1014, ¶¶38-39. Duvall and Trcka thus operate upon similar traffic data, indicating that Trcka’s analysis techniques would apply equally well to the data in Duvall-Chu’s system.


187. Implementing Trcka’s analysis techniques in the Duvall-Chu system amounts to nothing more than combining prior-art elements (i.e., Duvall-Chu’s filtering system and editing manager and Trcka’s data-analysis applications) according to known methods to yield predictable results. Thus, it is my opinion that a POSA would have had a reasonable expectation of success in making such implementations. Duvall already discloses processing and storage mechanisms, such as those implemented on server 30, as well as a “display” from which the user can view information. EX1004, 3:42-46, 4:61-64, 8:18-26. Thus, a POSA would not encounter unreasonable technical hurdles in implementing applications for interacting with past network activity, as in Trcka, in the system of Duvall-Chu.

188. Therefore, a POSA would have been motivated, with a reasonable expectation of success, to incorporate Trcka’s techniques for user analysis of network activity into the system of Duvall-Chu.

B. Claims 7, 8 and 12

189. Claim 7 recites, “*wherein the identifying step includes aggregating and synthesizing the status data at the probe.*” Claims 8 and 12 recite, “*wherein the identifying step includes cross-correlating data across the monitored components.*”

190. Trcka discloses both “*aggregating and synthesizing the status data at the probe*” and “cross-correlating data” across different monitored devices. For example, Trcka’s “Audit” and “Problem Determination” applications correlate data from different devices to “particular types of network problems” occurring within a specific timeframe. EX1014, ¶¶112, 116. A user can then view this information on screen, as shown for instance in Figure 19. EX1014, ¶¶142-43, FIG. 19. This correlates network activity from different devices to rejected logon attempts. *Id.* As shown below, the collected “*status data*” is displayed in aggregate, and in a meaningful, synthesized format.



EX1014, FIG. 19.

191. As another example, Trcka’s “Network Operating Characteristics Application” cross-correlates activity from different devices for indications of “congestion,” “traffic throughput,” and “network outages.” EX1014, ¶¶124-25. Similarly, Trcka’s “Network Problem Solving Application” cross-correlates activity from different devices to detect “network-level problems” such as “transmissions to nonexistent entities.” EX1014, ¶¶126-27. Specifically, the Network Problem Solving Application “utilizes the search engine and other API utilities to identify and generate a report of a predefined set of network-level problems” based on data gathered from multiple probes in the network. *Id.* A POSA would have understood this to be a cross-correlation of data across monitored components. Such metrics are derived by “*aggregating and synthesizing*” the collected “*status data*,” generating a metrics representative of the collected “*status data*” as a whole. Additionally, a POSA would have understood that Trcka’s aggregation and synthesis occurs “*at the probe*” where the data would be stored. *See* Section IX.A. Indeed, a POSA would have understood that this would be the most logical place to perform the aggregation and synthesis, as it would minimize time to retrieve the data and reduce the amount of data transmitted to an analyst’s computer.

192. Accordingly, Duvall in view of Chu and Trcka renders obvious claims 8 and 12.

C. Claims 9 and 13

193. Claims 9 and 13 recite, “*wherein the identifying step includes analyzing the frequency of occurrence of the events.*”


194. Trcka discloses “analyzing the frequency of occurrence of” potentially security-related events, such as “repeated unsuccessful logon attempts from an outside client” and “frequent transfers of files to outside entities.” EX1014, ¶¶122-23. Trcka further allows the user to view the frequency of potential security-related events in its “Audit” and “Problem Determination” applications, including for example, events related to “NETWORK ERRORS,” “FILE ACTIVITY/ACCESS,” and user “TRANSACTIONS.” EX1014, ¶142. Using these techniques taught by Trcka, a POSA would have been motivated to analyze the frequency of events in Duvall’s system, such as data requests transmitted to blocked or unknown URLs, to make the network administrator or analyst aware of such events and ensure that filters could be properly updated, thus preventing the system from “filtering too much or too little.” EX1004, 8:5-7. And a POSA would have had a reasonable expectation of success in doing so, as discussed in Section IX.A.

195. Accordingly, Duvall in view of Chu and Trcka renders obvious claims 9 and 13.

D. Claim 10

196. Claim 10 recites, “*further comprising after the step (c), a secure operations center coupled to perform further computer-based analysis and to receive data from the probe.*”

197. As discussed for claim 6, a POSA would have understood that Duvall’s “corporate network” provides a “secure operations center” within which users can operate. Section VIII.G. And as discussed in Section IX.A, Trcka discloses “computer-based analysis” in the form of “analysis applications,” as shown in Figure 7 below. EX1014, ¶¶15, 16, 53, 79.



EX1014, FIG. 7 (annotated).

198. For the reasons discussed in Section IX.A, a POSA would have been motivated to incorporate Trcka’s “computer-based analysis” techniques—recording and viewing past network activity using analysis applications, which execute on a computer—into the system of Duvall-Chu to enable the user to analyze whether “the filtering system is filtering too much or too little.” EX1004, 8:5-7; EX1014, ¶¶15, 16, 53, 79. Therefore, Duvall in view of Chu and Trcka teaches “*a secure operations center coupled to perform further computer-based analysis and to receive data from the probe.*”

E. Claim 11

199. Claim 11 recites, “*wherein the computer-based analysis includes aggregating, synthesizing, and presenting alerts on an ensemble basis.*”

200. Trcka discloses that the analysis applications for viewing past network activity include a “Report Generator Module.” EX1014, ¶113. Trcka explains that “[t]he Report Generation Module [] manipulates the data to match a format pre-selected by the user” and “can then be delivered to a printer or display, or can be saved to a file.” *Id.*, ¶104. Reports in Trcka’s system can include information such as “transaction activity,” “unauthorized accesses to restricted files and databases,” “statistics on congestion, peak loads, traffic throughput, network outages, and utilization,” and “frequent transfers of files to outside entities.” *Id.*, ¶¶112, 123, 125.

This report data is “aggregate[ed]” and “synthesiz[ed]” because it is data generated or synthesized from raw network activity and presented in a meaningful format.

201. By viewing past network activity in Duvall-Chu-Trcka’s system, as discussed in Section IX.A, the combination discloses “aggregating, synthesizing, and presenting alerts on an ensemble basis.” For example, such techniques would allow trained network professionals to view aggregated and synthesized information derived from raw packet data regarding blocked or unknown data transmissions, thereby alerting the trained network professionals to all such activity at the same time (i.e., “on an ensemble basis”). EX1014, ¶104. And this operation is consistent with the description in the ’641 patent, which discusses a “report generation module” that generates reports of past network activity. EX1001, 8:7-31.

202. Therefore, Duvall in view of Chu and Trcka renders obvious claims 11 and 36.

F. Claim 16

203. Claim 16 recites, “*wherein the dynamic modifying step includes consideration of non-real-time information from ongoing security research efforts.*”

204. Trcka discloses “consideration of non-real-time information” in its analysis applications, for example a user “selectively viewing and generating audit trails from the **archived** traffic data,” allowing the user to “specify such parameters as **start time/date**, **end time/date**, the type of event/s of interest.” EX1014, ¶112.

As I discussed in Section IX.A, a POSA would have been motivated to incorporate Trcka's techniques for recording and viewing past network activity (i.e., "non-real-time information") in Duvall-Chu's system to enable the user to analyze whether "the filtering system is filtering too much or too little." EX1004, 8:5-7.

205. Further, a POSA would have understood that the analysis of this data in Duvall-Chu-Trcka's system is part of "ongoing security research efforts" because the network administrator or analyst is researching past data to identify potential security issues that have not previously been identified. Again, these ongoing research efforts allow trained network professionals to determine whether Duvall's "filtering system is filtering too much or too little," and to update filters accordingly. *Id.*, 8:5-7.

206. Therefore, Duvall in view of Chu and Trcka renders obvious claims 16 and 41.


X. GROUND 3: CLAIMS 14-15 ARE OBVIOUS OVER DUVALL IN VIEW OF CHU AND TRCKA, AND IN FURTHER VIEW OF ZIESE.

A. A POSA would have been motivated to combine the Duvall-Chu-Trcka system with Ziese.

207. Ziese is directed to "dynamically distributing updates in a network." EX1015, 1:7-9. Ziese explains that "[l]ike the Internet, private networks also have become common means for communicating important information," and that such private networks may be distributed over "wide area networks (WANs)" that include "disparate network sites." *Id.*, 1:25-30, 2:29-36. Ziese further explains,

“[o]ne danger to reliable and secure network communications is posed by hackers or other unauthorized users disrupting or interfering with network resources.” *Id.*, 1:32-34. Thus, many security systems make use of “intrusion signatures to detect malicious activity.” *Id.*, 1:46-48. However, “[b]ecause new types of attacks appear more frequently than network administrators typically check with an intrusion detection provider for new signatures, networks often remain vulnerable to new types of attacks even though new signatures are available to identify and prevent such attacks.” *Id.*, 1:62-67.

208. To address this issue, Ziese discloses techniques “for dynamically distributing intrusion detection and other types of updates in a network that substantially eliminate or reduce disadvantages and problems associated with prior methods and systems.” EX1015, 2:2-6. For example, “programs are automatically updated by downloading and distributing an update in response to an automated event,” and “[a]s a result, systems with a common program separately running at several sites may update each site with no or minimal operator interaction.” *Id.*, 2:39-44. Such updates may be distributed from a centralized update server, as shown in Figure 1 below.



EX1015, FIG. 1 (annotated).

209. Although Duvall discloses an “editing manager” to edit filters when “the filtering system is filtering too much or too little,” EX1004, 8:2-7, in the case of a corporate network, Duvall expressly discloses only a single “server 30” on which the filtering system is running, *id.*, 8:18-26. Duvall and Ziese both explain that clients may be distributed widely beyond a traditional local area network (LAN). *See* EX1004, 2:35-39 (“A number of **clients are typically linked** to a network 14, such as a local area network (LAN) **or a broad-based commercial service.**”); EX1015, 1:26-31 (“Private networks” include “wide area networks (WANs).”), 2:29-32 (“[T]he downloaded update may be **distributed to disparate network sites**

... .”), 2:41-44 (“[S]ystems with a **common program separately running at several sites** may update each site with no or minimal operator interaction.”).

210. In the case of a distributed corporate network, a POSA would have understood that Duvall’s “filtering system” would employ multiple “server[s] 30” to serve local clients within each network site. EX1004, 8:18-21 (“[T]he filtering system can be provided from **a server 30 that is on the client’s own network 40.**”). This arrangement would ensure that the filtering system can apply filters to data transmissions before the transmission is allowed to enter a public network, such as the Internet.

211. Implementing Duvall’s “editing manager” on each local server would require common filter updates (i.e., updates that would apply to all clients within a corporation) to be made on each respective server. Ziese’s techniques solve this problem by allowing updates to be made centrally and automatically distributed to “disparate network sites” operating the same program. EX1015, 2:25-46. A POSA would have thus been motivated to modify Duvall-Chu-Trcka’s system to allow trained network professionals for a distributed corporate network to centrally analyze and update filters, and to have the filters automatically distributed to individual network sites.

212. And a POSA would have had a reasonable expectation of success incorporating Ziese’s automatic update distribution techniques into Duvall-Chu-

Trcka's system. Duvall already suggests the use of centralized updates (similar to Ziese) by disclosing an "[u]pdating [m]echanism" for distributing filter updates to its filtering system "over the Internet." EX1004, 7:15-67. Duvall explains that "the updating mechanism causes the filtering system to download from the update server one of a number of new sets of filters, and preferably causes some or all of the existing filters in the filter database to be replaced with the newly downloaded filters." *Id.*, 7:25-29; *see also id.*, 7:19-22 ("Because Internet sites are being added to the Internet at a fast rate, [] the filtering system preferably also has an updating mechanism to keep the filters current."). Therefore, incorporation of Ziese's techniques in the Duvall-Chu-Trcka system would amount to nothing more than use of known technique (i.e., Ziese's automatic distribution of updates) to improve similar methods or products (i.e., Duvall-Chu-Trcka's filtering system) in the same way.

B. Claim 14

213. Claim 14 recites, "*wherein the computer-based analysis includes cross-probe correlation.*"

214. As discussed above, in the case of a distributed corporate network, a POSA would have understood that Duvall's "filtering system" would employ multiple "server[s] 30" to serve local clients within each network site. Section X.A; EX1004, 8:18-21. And as discussed for claims 8 and 12, Duvall in view of Chu and

Trcka discloses analyzing and correlating network activity across different devices within the network. Section IX.B. In Duvall-Chu-Trcka-Ziese's system in which analysis of network activity is performed from a central location, a POSA would have understood that the analyzed network activity would span the multiple "server[s] 30" (i.e., "different probes"), and thus data would be correlated not only across monitored devices, but also across the multiple "different probes."

215. Therefore, Duvall in view of Chu, Trcka, and Ziese renders obvious claim 14.

C. Claim 15

216. Claim 15 recites, "*instantaneously self-tuning the probe based on previously collected status data.*"

217. Duvall in view of Chu, Trcka, and Ziese further renders claim 15. As discussed above, Ziese discloses techniques in which "programs are automatically updated by downloading and distributing an update in response to an automated event." EX1015, 2:39-41; Section X.A. This automated process ensures that updates are distributed when they are available, rather than waiting for a user to initiate the update process. EX1015, 1:54-67.

218. By automatically distributing updates in Duvall-Chu-Trcka-Ziese's system—for example, in response to a trained network professional completing a filter edit or addition—Duvall's filtering system would "instantaneously" have

access to the most up-to-date filters. And as discussed for [1E] of claim 1, when filters are updated in Duvall's filtering system, the analysis is "instantaneously self-tun[ed]" because Duvall is programmed to use the updated filters in subsequent analysis. Section VIII.B.6; EX1004, 4:22-43; 8:3-16. Additionally, as discussed for elements [1B]-[1D] of claim 1, this "self-tuning" is "based on previously collected status data" because filter edits are based on the user's analysis of previously collected information associated with potentially security-related events. See Sections VIII.B.3-5.

219. Therefore, Duvall in view of Chu, Trcka, and Ziese renders obvious claim 15.

XI. GROUND 4: CLAIMS 18-25 ARE OBVIOUS OVER DUVALL IN VIEW OF CHU, AND IN FURTHER VIEW OF COGGER.

A. A POSA would have been motivated to combine the Duvall with Chu and Cogger.

220. As I discussed in section VIII.A, a POSA would have been motivated to incorporate Chu's user intervention teachings into Duvall's system to more accurately resolve the status of residue data transmissions to ensure transmissions are correctly blocked or allowed. A POSA would have understood, however, that professionals managing a particular network (e.g., a corporate network) may not always have the knowledge or resources to determine whether a data transmission should be blocked or allowed. For example, Duvall acknowledges that "Internet sites are being added to the Internet at a fast rate." EX1004, 7:16-29. Because of

this, a corporation may not have the resources to keep its filtering up to date. *Id.*

This would have motivated a POSA to use a third party having more expertise and resources in identifying what data transmissions to allow and block, as Duvall itself suggests. EX1004, 7:16-29 (“Because Internet sites are being added to the Internet at a fast rate, [] the filtering system preferably also has an updating mechanism to keep the filters current. The updating mechanism in the filtering system can locate a particular HTTP or FTP update server 32 over the Internet.”), 7:35-38. A POSA would have thus been motivated to adapt Duvall in view of Chu to enable a third party having that expertise to process residue data.

221. Cogger discloses a “system and method for opening and tracking trouble tickets over the public Internet.” EX1033, Cogger, Abstract. Cogger discloses techniques for creating and transmitting information about detected “network event[s]” to “service organizations responsible for resolving the network event.” EX1033, Cogger, 2:50-63. Cogger’s “trouble tickets” aggregate information concerning a “network event” for transmission to a “Customer Service Management System” for analysis and resolution by the service organization. EX1033, 2:64-4:3.

222. Cogger illustrates an example “trouble ticket” in Figure 8(g):




FIG. 8(g)

EX1033, FIG. 8(g)

223. As shown above, Cogger's "trouble ticket" can include various information about a network event, for example a description, customer information, time and date of occurrence, and priority. EX1033, 20:1-15. And Cogger discloses that this information is customizable and "will be different for different types of tickets." EX1033, 20:13-15.

224. A POSA would have been motivated to incorporate Cogger's techniques for generating and transmitting trouble tickets in Duvall-Chu's system for use in prompting Cogger's network administrators to resolve residue data transmissions. Transmitting "trouble tickets" with sufficient information about residue data transmissions to Cogger's network administrators would ensure that the service can appropriately evaluate and resolve such data transmissions. This would

enable Cogger's network administrators to update the corporation's filter database when appropriate by adding new filters for newly-discovered or updated Internet sites. This would enable the system to block or allow traffic to/from a newly-discovered site using the new filters. Cogger's trouble tickets would also provide a tracking mechanism in the combined system so that Duvall's network administrators of each corporate network could assure residue data is resolved in a timely manner. EX1033, 2:50-33.

225. Moreover, Cogger's "trouble tickets" can include "[r]emarks" that "enable the customer to describe the trouble in his/her own words." EX1033, 20:50-64. Allowing the client user or network administrator in Duvall-Chu's system to include remarks regarding a residue data transmission would allow the network administrators of the third-party service to better diagnose whether the residue data transmission should be allowed or blocked. For instance, the remarks would provide greater flexibility in expressing the characteristics of the residue data. This would enable the third-party service to more efficiently generate "new sets of filters" based on the received trouble tickets, furthering Duvall's express goal of "keeping the filters current." EX1004, 7:18-28. I also note that enabling users to send trouble tickets for residue data transmissions is consistent with Chu's teaching of enabling "user intervention" to resolve data transmissions that do not meet the criteria of its blacklist or whitelist. Cogger's network administrators would have more expertise

as Cogger describes a system intended to support many disparate customers, leading to a better database of troubleshooting solutions. That is, Cogger's network administrators would have experience with a wide array of data transmissions and would have greater expertise in reviewing and handling such data.

226. Incorporating Cogger's techniques into Duvall-Chu's system would amount to nothing more than incorporating well-known techniques taught by Cogger (generating and transmitting tickets about network events) into Duvall-Chu's known filtering system. Thus, it is my opinion that a POSA also would have had a reasonable expectation of success in implementing these modifications. Duvall already discloses that its filtering system is connected to the update service via the Internet, and Chu already discloses prompting a user with expertise to resolve unknown data. EX1004, 7:21-23; FIG. 1; EX1005, 44. Duvall's server also already implements a "screen on the user's display" for indicating the results of, for instance, a blocked transmission, and for enabling the user/network administrator to "fill[] out an on-screen [] form." EX1004, 4:61-64, 7:33-39. Thus, a POSA would not have encountered unreasonable technical issues generating and transmitting information about residue data for resolution by a network administrator in Cogger's system.

B. Independent Claim 18

227. It is my opinion that independent claim 18 is obvious for the reasons I set forth below.

1. [18-pre] “A method of operating a secure operations center as part of a security monitoring system for a customer computer network”

228. As discussed for element [1-pre] of claim 1, Duvall-Chu discloses “*a security monitoring system for a ... computer network.*” Section VIII.B.1. As discussed in Section XI.A, Duvall further discloses a filter update service as part of its “*security monitoring system*” that services multiple subscribers (i.e., “*customer[s]*”) “over the Internet.” EX1004, 7:16-29. And a POSA would have been motivated to incorporate Cogger’s “trouble ticket” techniques into Duvall-Chu’s system to prompt a network administrator, as disclosed in Cogger, to resolve residue data transmissions. Section XI.A.

229. Cogger further discloses “*operating a secure operations center*” in the form of a “Customer Service Management System” (CSM) that receives trouble tickets transmitted from customers. EX1033, 3:65-66. For example, Cogger discloses that communication between customers and the CSM is “*secure*” because it provides “**secure web servers and back end services** to provide applications that establish user sessions . . . and communicate with adaptor programs to simplify the **interchange of data across the network.**” EX1033, 5:46-50; *see also id.*, 7:55-60 (data is communicated “by establishing a **secure** TCP messaging session with one of

the DMZ networkMCI Interact Web servers 24 via an Internet **secure communications path** 22 established, preferably, with a **secure sockets SSL version of HTTPS.**”). Therefore, the CSM, which receives and resolves trouble tickets, provides “*a secure operations center*” for communicating with customers.

230. In the system of Duvall-Chu-Cogger, a POSA would have implemented Cogger’s CSM as part of Duvall-Chu’s network-administrator resolution pathway for receiving and analyzing tickets transmitted from subscriber networks. Section XI.A. Thus, like Cogger’s CSM, Duvall-Chu-Cogger’s residue data review system provides “*a secure operations center as part of a security monitoring system for a customer computer network.*”

231. Moreover, a POSA would have been motivated to ensure that the personnel who operate Duvall-Chu-Cogger’s service do so in a secure manner. Because the update service would receive network activity from private customer networks (e.g., “a corporate network,” EX1004, 8:21-23), a POSA would have recognized the need to ensure such information is handled securely and kept private. Indeed, such techniques were well-known in the art, and their utility was recognized, for instance, in Cogger. *E.g.*, EX1033, 5:46-50; see also *id.*, 7:55-60, 8:32-35. Therefore, a POSA would have understood that Duvall-Chu-Cogger’s service would be operated as part of a “*secure operations center.*”

2. [18A] “creating an event record for information received about an identified potentially security-related event occurring on the network, wherein the potentially security-related event is identified by filtering followed by an analysis of post-filtering residue, wherein the post-filtering residue is neither discarded nor selected by the filtering”

232. As discussed for element [1B], Duvall discloses identifying a “potentially security-related event occurring on the network, wherein the potentially security-related event is identified by filtering followed by an analysis of post-filtering residue, wherein the post-filtering residue is neither discarded nor selected by the filtering.” Section VIII.B.3. As discussed for element 1[C], Duvall-Chu also discloses transmitting information about the identified event to a trained network professional. A POSA would have understood that the information received by the trained network professional is “an event record” because it is information about the network event, such as receipt of a residue (i.e., unmatched) data transmission. Thus, the event record relates to information received about the residue data transmission (potentially security-related event).

233. Moreover, as discussed in Section XI.A, a POSA would have been motivated to incorporate such information into a “trouble ticket,” as taught by Cogger, to efficiently and coherently provide the information to a trained network professional (e.g., Cogger’s network administrators). The information as recorded in the trouble ticket thus also provides “an event record.” The trouble ticket is created

for information received about a residue data transmission, which, as I explained above, is a potentially security-related event.

3. [18B] “correlating the event record with customer information and a symptom record”

234. When a customer opens a new trouble ticket in Cogger, customer information may be “prepopulat[ed]” to minimize the amount of required user input. EX1033, 3:36-53. Figure 8(g) of Cogger shows the “customer information” as part of the trouble ticket, which is illustrated as “customer tab 287(b).” EX1033, 20:6-13.

Service Inquiry

File Edit View Tools Help

Create Query Find Refresh Help Cust Service

Details for Ticket Number 9803270000014

Ticket Customer Circuit Call

Ticket No.: 9803270000014

Product: DATA Service:

Trouble Desc.: NEXTGEN TKT UPFRONT 992 Open Date:

MCI Priority: 1

Occurred Date (GMT): 1998-03-27 Occurred Time (GMT): 20:04

Occurred Date (Local): Occurred Time (Local):

Originating Org: WEM Service ID: MDC895410468

OK Clone Modify Remarks Activities

Code

547

547

547

547

547


547

EX1033, FIG. 8(g) (annotated).

A POSA would have been motivated to include customer information as taught by Cogger in the trouble ticket transmitted to Duvall-Chu-Cogger’s service to inform the service of the corporate network on which the event occurred. This would better

enable the service to “validate the user,” as taught by Duvall. EX1004, 7:39-40. This would also enable the update service to determine the origin of the trouble ticket so that the service could contact the customer as needed to resolve the data transmission. This would also enable the update service to analyze the trouble ticket in relation to previous traffic on the network (e.g., based on previous trouble tickets provided by the network). This is consistent with Duvall’s disclosure that the update server maintains a “user data file” for different customers. EX1004, 7:41-51.

235. Cogger further discloses that its trouble tickets include “*a symptom record*.” For example, Cogger discloses, “[t]he first step in the creation of a trouble ticket is to identify the Event-Type/Call-Type,” and “[t]o do that the client front end asks the customer the problem/identifier type and then [to] narrow down the problem.” EX1033, 21:23-31. An example is illustrated in Cogger’s Figure 8(k), which populates problem information (e.g., “Description of Problem”) based on customized questions and answers. *Id.*, 21:55-22:6.



EX1033, FIG. 8(k).


236. Cogger explains that “[t]he questions that need to be asked to open a ticket vary by product and problem type.” EX1033, 22:58-59. A POSA would have understood that such information relates to “*symptom[s]*” of a particular problem. *See id.* A POSA would have also been motivated to create a similar “*symptom record*” in Duvall-Chu-Cogger’s system as part of the trouble ticket to relay information about the type of problem encountered. For example, such information related to the residue data might be communication protocol, URL, file type, etc. This would provide useful symptom information for Duvall-Chu-Cogger’s service in analyzing the data. And as I described above, this would also allow the update

server to update the user data file or other user file for recording the customer's traffic history (e.g., by logging trouble tickets provided by the customer). *See* EX1004, 7:41-52.

237. Accordingly, Duvall-Chu-Cogger discloses “*correlating the event record with customer information and a symptom record*” within a trouble ticket.

4. [18C] “using the correlated symptom record to link the event record to problem resolution information”

238. Cogger discloses “*problem resolution information*” in the form of “remarks” added to a trouble ticket. EX1033, 20:33-62. These remarks “enable the customer to describe the trouble in his/her own words or add any comments” and “can aid in the resolution of a problem” (i.e., “*problem resolution information*”). EX1033, 20:59-62.



EX1033, FIG. 8(i).

239. In Duvall-Chu-Cogger's system, the remarks included as part of the trouble ticket would allow a network administrator to add further commentary for use by Duvall-Chu-Cogger's service in determining how to resolve the received event (i.e., whether and how filters need to be updated). For instance, this could include remarks stating that a residue URL appears to be for a legitimate company, indicating to the network administrators that this is likely an allowable URL. Therefore, such remarks "can aid in the resolution of a problem." As another example, the trouble ticket could implement remarks indicating the network

administrator's knowledge of or past interaction with a company indicated in the URL. This would provide added context to aid the update service in updating current filters and/or generating new filters.

240. Cogger further discloses “*using the correlated symptom record to link the event record to problem resolution information*” in at least two ways. First, as I discussed for [18B], creation of the “*symptom record*” is “[t]he first step in the creation of a trouble ticket.” EX1033, 21:23-31, 21:39-47. The “*symptom record*” is used to open the ticket. *Id.* Thus, the “*symptom record*” allows the event information and remarks to be linked within the ticket. Second, Cogger discloses that the “questions and corresponding answers” associated with the “*symptom record*” “may be entered into [the] remarks section of the trouble ticket.” 3:58-64. Therefore, the “*symptom record*” is used to create remarks (i.e., “*problem resolution information*”) that are linked to the event information within the trouble ticket.

241. Therefore, Duvall-Chu-Cogger discloses element [18C].

5. [18D] “consolidating the event record, correlated customer information and symptom record, and linked problem resolution assistance information into a problem ticket”

242. As discussed for elements [18A]-[18C], Duvall-Chu-Cogger discloses “consolidating the event record, correlated customer information and symptom record, and linked problem resolution assistance information into a problem ticket,” in the form of a “trouble ticket.” Sections XI.A, XI.B.2-4. Indeed, the Duvall-Chu-

Cogger trouble ticket indicates an event record (e.g., residue data transmission). The trouble ticket also specifies and correlates customer information and a symptom record. EX1033, 22:58-59 (“[t]he questions that need to be asked to open a ticket”). The trouble ticket also specifies problem resolution assistance information (e.g., customer remarks), which is consolidated into the trouble ticket for transmission to the update server.

6. [18E] “providing the problem ticket to a security analyst console for analysis”

243. Cogger discloses that “the trouble ticket is submitted to [a] Customer Service Management System.” EX1033, 3:65-66. “When the CSM accepts the trouble ticket, a trouble ticket number is assigned.” *Id.*, 3:66-67. As discussed for [18-pre], in the Duvall-Chu-Cogger system, Cogger’s CSM would be implemented as part of Duvall-Chu’s system to provide an outside service for assessing residue data—*i.e.*, by resolving tickets. EX1033, 2:24-30; Section XI.A.


244. A POSA would have recognized that a security analyst is required to resolve a ticket, by viewing event information and updating filters as appropriate. Section XI.A. As discussed for element [1C], Duvall already teaches providing an “*analyst system*” to network administrators to view information about network events. Section VIII.B.4. A POSA would have been motivated to provide a similar “*analyst console*” to analysts operating Duval-Chu-Cogger’s service in order to analyze event information received in a trouble ticket and update filters as

appropriate. Indeed, Cogger teaches receiving trouble tickets electronically (*see* [18-pre]). A POSA would have thus understood analysts operating Duval-Chu-Cogger's third-party service would view ticket information from an "*analyst console*."


C. Claim 19

245. Claim 19 recites, "correlating the event record with a pre-existing event record stored on an event database within the secure operations center; and linking the event record to an open problem ticket associated with the pre-existing event record."


246. Cogger discloses that a "trouble ticket... may exist as a **data record in a database**." EX1033, 2:18-20. Thus the event record included in the trouble ticket is "a pre-existing event record stored on an event database within the secure operations center." Cogger further discloses that a customer can search for and edit open tickets, as shown in the following interface screens:



EX1033, FIG. 8(f).



EX1033, FIG. 8(e).



EX1033, FIG. 8(c).

247. When an open ticket is selected, a customer can edit details, including adding additional remarks to the ticket. EX1033, 20:50-59. A POSA would have understood that in Duvall-Chu-Cogger’s system, certain events are related. For example multiple residue data transmissions to or from a URL with the same top-level domain may be related. In such a case, a POSA would have been motivated to incorporate the related events in the same ticket for analysis by the third-party service. This would enable related events to be analyzed and resolved together. Cogger also provides a mechanism for doing this by adding the related event information as a new remark within an existing ticket, i.e., “linking the event record to an open problem ticket associated with the pre-existing event record.” EX1033, 20:50-59.

D. Claim 20

248. Claim 20 recites, “wherein the problem resolution assistance information links the symptom to a diagnosis.”

249. As discussed for [18B], Cogger discloses a series of questions used to determine “symptom[s]” of the problem. EX1033, 21:55-22:6, 22:58-59; Section XI.B.3. Cogger further discloses that these questions are dynamic and used to diagnose the problem: “QuestionTrees [] are essentially a **set of questions for a particular product and problem type**,” and “[w]ithin each group of Questions, there is one question that is designated as the decision point which is **used to determine the next group of Questions** that need to be presented to the user.” EX1033, 23:15-16, 24:20-23. A POSA would have thus understood that the particular series of questions and corresponding answers link symptoms of the problem to a diagnosis.


250. Additionally, Cogger discloses that “[b]oth the questions and the corresponding answers may be input into a remarks section of the trouble ticket.” EX1033, 3:62-64. “[R]emarks are comments added to a ticket for historical purposes and can aid in the resolution of the problem.” *Id.*, 20:33-34. Cogger thus discloses that “the problem resolution information,” (i.e., the remarks) “links the symptom to a diagnosis.” And as I explained above, in the Duvall-Chu-Cogger combination, the service resolves the received event (e.g., whether and how filters

need to be updated) based on “remarks” section of the trouble ticket. Thus, the combination of Duvall, Chu, and Cogger renders this claim obvious.

E. Claim 21

251. Claim 21 recites, “wherein the problem resolution assistance information links the symptom to a technology.”

252. As shown below, Cogger’s tickets include a “[p]roduct type” to which the problem relates. EX1033, 21:27-32; FIG. 8(k).



EX1033, FIG. 8(k) (annotated).

253. A POSA would have understood that, in the Duvall-Chu-Cogger system, the product types would relate to the customer devices sending or receiving

data. For instance, such product types could relate to the type of machine, browser, and operating system associated with the network event. This understanding is consistent with the '641 patent, which explains that “Technology” may refer to “[a]n operating system, software product or hardware product.” EX1001, 19:56-58.

254. Because the “*problem resolution assistance information*” (i.e., the remarks) include “*the symptom*,” as I discussed above for claim 20, the “*problem resolution assistance information links the symptom to a technology*” within the trouble ticket.

F. Claim 22

255. Claim 22 recites, “correlating the event record with a vulnerability; and using the correlated vulnerability to link the event record to problem resolution assistance information.”

256. Cogger discloses that a “priority” is assigned to each trouble ticket. EX1033, 20:9-12; FIGs. 8(c), 8(g).





FIG. 8(c)

EX1033, FIG. 8(c) (annotated).

A POSA would have understood that such a priority indicates a “*vulnerability*,” defining the importance of resolving a particular network event. Higher priorities represent higher vulnerabilities. Higher priority network events should be resolved before lower priority network events because the higher priority event represent higher vulnerabilities. Thus, Cogger discloses “*correlating the event record with a vulnerability*” within the trouble ticket.

257. Cogger further discloses “*using the correlated vulnerability to link the event record to problem resolution assistance information*” because trouble tickets can be searched by priority. EX033, 58:14-15 (“wherein said trouble ticket search criteria includes an indication of priority of said trouble ticket”).



EX1033, FIG. 8(f) (annotated).


258. The results of the search query link the “*event record*,” as indicated by the ticket number, to the remarks contained within the ticket, i.e., “*problem resolution assistance information*.” See, e.g., EX1033, 19:55-57 (“The customer may then view the results of a query, i.e., the details, remarks or activities of a Ticket chosen from a list of Tickets.”).

G. Claims 23 and 24

259. Claim 23 recites, “wherein the problem resolution assistance information links the vulnerability to a diagnosis.” Claim 24 recites, “wherein the problem resolution assistance information links the vulnerability to a technology.”

260. Cogger discloses that the remarks within a trouble ticket (“problem resolution assistance information”) include diagnosis information, as discussed for claim 20. Section XI.D. Cogger discloses that “the problem resolution assistance information links the vulnerability to a diagnosis” because the ticket priority (“vulnerability”) is linked to the remarks as part of the ticket.

261. Cogger further discloses that ticket remarks can include information about “a technology,” such as the “CIRCUIT STATUS” or “DATA SPEED.”
EX1033, FIG. 8(h).



EX1033, FIG. 8(h).

262. In Duvall-Chu-Cogger’s system, a POSA would have been motivated to include relevant technology information in the remarks to help “describe the

trouble in his/her own words.” EX1033, 20:59-62. For instance, Duvall discloses performing “pattern match[ing]” based on a “protocol” determined from the analyzed data transmission. EX1004, 6:5-9. In the event that the data is residue, such protocol information would be useful to include in the trouble ticket to assist with the analysis performed by Duvall-Chu-Cogger’s third-party service. By including such information in the ticket remarks, Duvall-Chu-Cogger discloses “*the problem resolution assistance information links the vulnerability to a diagnosis.*”

H. Claim 25

263. Claim 25 recites, “updating the problem resolution assistance information after the potentially security-related event is resolved.”

264. Cogger discloses that remarks may be “added on the ticket either by the customer or by the trouble ticket processing organization.” EX1033, 20:36-39.

Cogger further discloses that the customer is responsible for closing the ticket after a resolution:

Since only the organization that opened the ticket has authority to close it, **once a ticket has been resolved** the ticket is referred out to the customer’s organization. **If the customer is not satisfied with the problem resolution**, that customer may refer the ticket back to the ticket resolution enterprise organization.

EX1033, 22:44-49.

265. A POSA therefore would have been motivated to include the resolution information from the ticket processing organization in the remarks so that the customer could view the resolution. This process would apply to customers of

Duvall's filtering service and would be beneficial to inform network administrators of the resolution for residue or blocked data transmissions. Indeed, Cogger teaches this operation as the remarks include "comments added to a ticket for historical purposes," e.g., memorializing the resolution. Accordingly, Duvall-Chu-Cogger discloses "*updating the problem resolution assistance information after the potentially security-related event is resolved.*"

XII. OBJECTIVE INDICIA DO NOT SUPPORT PATENTABILITY

266. I understand that the obviousness analysis includes considering evidence of secondary indicia of nonobviousness (if available). I understand that any secondary indicia of nonobviousness must have a nexus, or a close connection, to the novel aspects of the claimed invention. I understand that secondary indicia of nonobviousness include the following: commercial success, long-felt but unmet need, failure of others, copying, praise in the industry, and unexpectedly superior results as compared with the closest prior art.

267. Considering the information I have reviewed, my general knowledge of the field, and publications within the field and related fields, I am unaware of any objective evidence that might support the patentability of the claims of the '641 patent. Further, from my review of the prosecution history of the '641 patent, the applicant never raised secondary indicia of nonobviousness. I am not otherwise aware of any publicly available evidence of secondary indicia of nonobviousness.


Should any alleged secondary indicia of nonobviousness be presented, I reserve the right to provide an opinion thereon.

XIII. CONCLUSION

268. In signing this declaration, I recognize that this declaration will be filed as evidence in an *inter partes* review before the Patent Trial and Appeal Board of the United States Patent and Trademark Office. I also recognize that I may be subject to cross-examination in the case and that cross-examination will take place within the United States. If cross-examination is required, I will appear for cross-examination within the United States during the time allotted.

269. I hereby declare that all statements made herein are made of my own knowledge and are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

Executed on this 27th day of April, 2023 in Chapel Hill, NC.



Kevin Jeffay, Ph.D.