Trials@uspto.gov
571-272-7822

Paper 29
Entered: July 10, 2015

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

FINJAN, INC.,
Petitioner,

v.

FIREEYE, INC.,
Patent Owner.

Case IPR2014-00492
Patent 8,171,553 B2

Before BRYAN F. MOORE, LYNNE E. PETTIGREW, and FRANCES L. IPPOLITO, *Administrative Patent Judges.*

IPPOLITO, *Administrative Patent Judge.*

FINAL WRITTEN DECISION
Inter Partes Review
35 U.S.C. § 318(a) and 37 C.F.R. § 42.73

I. INTRODUCTION

Finjan, Inc. filed a Corrected Petition ("Pet.") on March 21, 2014, requesting an *inter partes* review of claims 1–30 of U.S. Patent No. 8,171,553 B2 ("the '553 patent"). Paper 4. Patent Owner FireEye, Inc. filed a Preliminary Response ("Prelim. Resp.") to the Petition. Paper 7. On July 25, 2014, we instituted an *inter partes* review of claims 1, 3–8, 12–14, 16–20, and 22–30 on the following grounds of unpatentability alleged in the Petition (Paper 8, "Dec."):

A. Claims 1, 5, 7, 17, 22, and 25–27 are unpatentable under 35 U.S.C. § 103 over Kaeo[1] and Venezia[2];

B. Claims 6, 8, 12–14, 16, 18, and 19 are unpatentable under 35 U.S.C. § 103 over Kaeo, Venezia, and Chen[3];

C. Claims 1, 3–5, 7, 17, and 22–28 are unpatentable under 35 U.S.C. § 103 over Kaeo and Liljenstam[4]; and

D. Claims 18, 20, 29, and 30 are unpatentable under 35 U.S.C. § 103 over Kaeo, Liljenstam, and Dunlap[5].

[1] Merike Kaeo, *Designing Network Security*, Cisco Press (2nd ed. Nov. 2003) (Ex. 1006, "Kaeo").

[2] Paul Venezia, *NetDetector Captures Intrusions*, InfoWorld Issue 27 (July 14, 2003) (Ex.1005, "Venezia").

[3] Peter M. Chen and Brian D. Noble, *When Virtual Is Better Than Real*, Department of Electrical Engineering and Computer Science, University of Michigan (May 21, 2001) (Ex. 1009, "Chen").

[4] Michael Liljenstam et al., *Simulating Realistic Network Worm Traffic for Worm Warning System Design and Testing*, Institute for Security Technology studies, Dartmouth College (Oct. 27, 2003) (Ex. 1007, "Liljenstam").

[5] George W. Dunlap et al., *ReVirt: Enabling Intrusion Analysis through Virtual-Machine Logging and Replay*, Proceeding of the 5th Symposium on

After institution of trial, Patent Owner filed a Patent Owner Response ("PO Resp.," Paper 20) and Petitioner filed a Reply thereto ("Reply," Paper 23). An oral argument was held on March 31, 2015. The transcript of the oral hearing has been entered into the record. Paper 28, "Tr."

We have jurisdiction under 35 U.S.C. § 6(c). This Final Written Decision is issued pursuant to 35 U.S.C. § 318(a) and 37 C.F.R. § 42.73.

Petitioner has shown, by a preponderance of the evidence, that claims 1, 3–7, 17, 18, 20, and 22–30 of the '553 patent are unpatentable. Petitioner has not shown, by a preponderance of the evidence, that claims 8, 12–14, 16, and 19 are unpatentable.

A. Related Proceedings

Petitioner indicates that the parties are involved in a related proceeding, *Finjan, Inc. v. FireEye, Inc.*, No. 4:13-cv-03133-SBA, filed in the United States District Court for the Northern District of California. Paper 6, 1.

The parties also are involved in Case IPR2014-00344, directed to U.S. Patent No. 8,291,499 B2 ("the '499 patent"), which shares a common disclosure with the '553 patent.

B. The '553 Patent

The '553 patent describes an authorized activity capture or detection system that analyzes copied network data with a heuristic to determine if the copied network data has the characteristics of a computer worm. *See* Ex. 1001, Claim 1. If the compared network data has a characteristic of a computer worm, the system flags the compared network data for replay in an analysis environment. *Id.*

Operating Systems Design and Implementation, USENIX Association (Dec. 9–11, 2002) (Ex. 1008, "Dunlap").

Figure 7 of the ’553 patent is reproduced below.

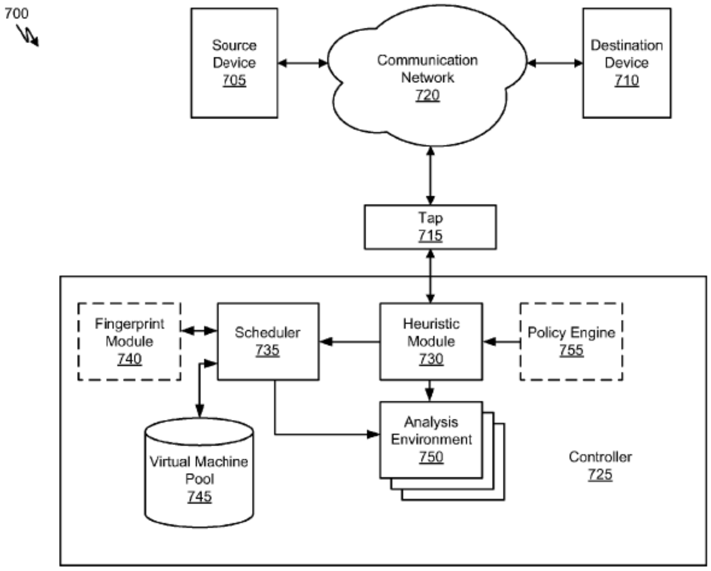

FIG. 7

Figure 7 depicts an embodiment of an unauthorized activity detection system described in the ’553 patent. Unauthorized activity detection system 700 includes source device 705, destination device 710, and tap 715, each of which is coupled to communication network 720. *Id.* at 26:21–26. Tap 715 is further coupled to controller 725. *Id.* at 26:25–26. In operation, tap 715 monitors network data and provides a copy of the network data to controller 725. *Id.* at 26:35–37.

Figure 7 also shows controller 725, which “can be any digital device or software that receives network data from the tap 715.” Ex. 1001, 27:1–2. “In some embodiments, controller 725 is contained within computer worm sensor 105.” *Id.* at 27:2–4. Controller 725 also may be contained within separate traffic analysis device 135 or may be a stand-alone digital device. *Id.* at 27:4–6. Controller 725 can comprise virtual machine pool 745, analysis environment 750, heuristic module 730, and policy engine 755. *Id.*

at 27:6–9. “[V]irtual machine pool 745 is configured to store virtual machines [and] . . . can be any storage capable of storing software.” *Id.* at 28:51–52. Additionally, “analysis environment 750 simulates transmission of the network data between the source device 705 and the destination device 710 to analyze the effects of the network data upon the destination device 710.” *Id.* at 28:59–62. Heuristic module 730 can receive copied network data from tap 715 and apply heuristic and/or probability analysis to determine if the network data contains suspicious activity. *Id.* at 27:12–15.

C. Illustrative Claim

Of the challenged claims, claims 1, 8, 17, and 28 are independent. Claim 1, reproduced below, is illustrative of the subject matter of the ’553 patent:

> 1. An unauthorized activity capture system comprising:
>
> a tap configured to copy network data from a communication network; and
>
> a controller coupled to the tap and configured to receive the copy of the network data from the tap, analyze the copy of the network data with a heuristic to determine if the copy of the network data has one or more characteristics of a computer worm, flag at least a portion of the copy of the network data as suspicious by flagging the at least a portion of the copy of the network data for replay in an analysis environment based upon the heuristic determination that the at least a portion of the analyzed copy of the network data has one or more characteristics of a computer worm, and replay transmission of the suspicious, flagged network data copied from the communication network to a destination device.

Ex. 1001, 31:60–32:8.

II. ANALYSIS

A. Claim Construction

During a review before the Patent Trial and Appeal Board (“Board”), we construe claims in an unexpired patent in accordance with the broadest reasonable interpretation in light of the specification of the patent in which they appear. 37 C.F.R. § 42.100(b); *see In re Cuozzo Speed Techs., LLC,* 778 F.3d 1271, 1278–82 (Fed. Cir. 2015) ("Congress implicitly adopted the broadest reasonable interpretation standard in enacting the AIA," and "the standard was properly adopted by PTO regulation."); *see* Office Patent Trial Practice Guide, 77 Fed. Reg. 48,756, 48,766 (Aug. 14, 2012). Under the broadest reasonable interpretation standard, claim terms are given their "ordinary and customary meaning" as would be understood by one of ordinary skill in the art in the context of the entire disclosure. *In re Translogic Tech., Inc.,* 504 F.3d 1249, 1257 (Fed. Cir. 2007). An inventor may rebut that presumption by providing a definition of the term in the Specification with "reasonable clarity, deliberateness, and precision." *In re Paulsen,* 30 F.3d 1475, 1480 (Fed. Cir. 1994). In the absence of such a definition, limitations are not to be read from the Specification into the claims. *In re Van Geuns,* 988 F.2d 1181, 1184 (Fed. Cir. 1993).

1. flag or flagging (Claims 1, 8, 17, and 28)

For the purposes of our Decision to Institute, we determined that the broadest reasonable interpretation of the terms "flag" and "flagging" is "identify" and "identifying." Dec. 6–7 (adopting our analysis in the Decision to Institute (Paper 17) for the same term at issue in IPR2014-00344, in which the '499 patent shares the same disclosure as the '553 patent). Neither party disputes this interpretation. Pet. 5; PO Resp. 12. Based on the complete record now before us, we discern no reason to change this interpretation; thus, we adopt our previous analysis and interpret "flag" and "flagging" to mean "identify" and "identifying," respectively.

2. virtual machine pool (claims 6, 14, and 19)

In the Decision instituting trial, we construed “virtual machine pool” to include “any storage capable of storing one or more virtual machines.” Dec. 6–7. Patent Owner contests this construction and argues that the “notion that ‘any storage’ is a virtual machine pool would mean that any hard drive is a virtual machine pool regardless of whether it stores potential virtual machines.” PO Resp. 12 n.1.

We do not agree with Patent Owner’s arguments. Our construction in the Decision to Institute does not include “any storage,” as Patent Owner suggests, but “storage *capable* of storing one or more virtual machines.” Dec. 6–7 (referring to our discussion of “virtual machine pool” in the Decision to Institute (Paper 17) for IPR2014-00344). This construction is consistent with the Specification, which states that “virtual machine pool 745 can be any storage capable of storing software” and “virtual machine pool 745 is configured to store virtual machines.” Ex. 1001, 28:50–52.

Thus, under the broadest reasonable interpretation, we construe “virtual machine pool” to mean “any storage capable of storing one or more virtual machines.”

3. analysis environment (claims 1, 8, 17, and 28)

In the Decision to Institute, we determined, based on the preliminary record, that the term “analysis environment” means “an environment in which analysis of the effect of the network data upon a destination device is performed.” Dec. 6–7.

In Patent Owner’s Response, Patent Owner disagrees with our construction because it “permits an analysis environment to be a passive location or one in which a human being performs analysis.” PO Resp. 13. Patent Owner asserts that a person of ordinary skill in the art would not consider an “analysis environment” to be an environment where analysis is performed by either the analysis environment or some other actor. *Id.* at 14. Patent Owner adds that the "analysis environment" is described throughout the '553 patent as an actor rather than merely a passive component enabling actions by others. *Id.* at 14–15 (citing Ex. 1001, 29:24–25, 30:3, 30:8, 31:17–19).

Although the '553 patent provides examples where analysis environment 750 "determines," "simulates," "can react," or "replays," as noted by Patent Owner in the cited sections above (Ex. 1001, 29:24–25, 30:3–4, 30:8, 31:17–19), the Specification also indicates these descriptions of analysis environment 750 are non-limiting examples that disclose "some embodiments" or "one embodiment." Ex. 1001, 29:22–23, 29:36–37, 30:65–67. This disclosure in the Specification is consistent with the language of the challenged claims, which do not require explicitly that the analysis environment actively perform any action. For example, claim 1 requires that the recited controller "flag at least a portion of the copy of the network data as suspicious by flagging the at least a portion of the copy of the network data for replay *in* an analysis environment." (Emphasis added). In claim 1, the analysis environment provides a location for replay, which does not require that the network data is replayed *by* the analysis environment. Further, looking back to the Specification, the '553 patent provides that "in accordance with one embodiment of the present invention . . . the analysis environment 750 *replays* transmission of the network data." Ex. 1001, 30:65–66, 31:17–18 (emphasis added). This disclosure along with the express language of claim 1 indicates that, in the context of the '553 patent, a distinction exists between an analysis environment that provides a location for replaying data versus an analysis environment that itself performs replay.

Thus, we do not find that the recited "analysis environment" requires that the environment perform the analysis. Although claims are interpreted in light of the specification, limitations from the specification are not read into the claims. *In re Van Geuns*, 988 F.2d 1181, 1184 (Fed. Cir. 1993). The claim language does not require or mention the analysis environment performing an analysis. Moreover, even assuming Patent Owner is correct that the '553 patent only describes the analysis environment as actively performing analysis, claims generally are not limited to any particular embodiment disclosed in the specification, even where only a single embodiment is disclosed. *Innova/Pure Water, Inc. v. Safari Water Filtration Sys., Inc.*, 381 F.3d 1111, 1117 (Fed. Cir. 2004); *see, e.g., Silicon Graphics, Inc. v. ATI Techs., Inc.*, 607 F.3d 784, 792 (Fed. Cir. 2010) ("A construing court's reliance on the specification must not go so far as to import limitations into claims from examples or embodiments appearing only in a patent's written description . . . unless the specification makes clear that the patentee . . . intends for the claims and the embodiments in the specification to be strictly coextensive." (internal quotation marks omitted)).

Accordingly, we construe "analysis environment" to mean "an environment in which analysis of the effect of the network data upon a destination device is performed." *See* Dec. 6.

4. virtual switch (claim 20)

In the Decision to Institute, we determined that under the broadest reasonable interpretation, the term "virtual switch" means "software that is configured to mimic the performance of a switch." Dec. 6. The parties do not dispute this construction. PO Resp. 12; Reply 2–5. Based on the complete record now before us, we discern no reason to change this construction; we adopt our previous analysis for this non-disputed claim term.

5. *replay transmission of the suspicious, flagged network data copied from the communication network to a destination device (claim 1);*

   replaying transmission of the flagged at least a portion of the analyzed copied network data which was copied from the communication network to a destination device to identify unauthorized activity based on playback of the flagged suspicious at least a portion of the analyzed copy of the network data (claim 17);

   replay transmission of the flagged suspicious at least a portion of the analyzed copied network data copied from the network to a destination device (claim 28)

In distinguishing the challenged claims over the asserted prior art, Patent Owner argued at the oral hearing that the replay/replaying phrases (shown above) recited in independent claims 1, 17, and 28 require replay of data to a destination device.

> JUDGE IPPOLITO: Before you do, I would just like to go back to my original question about the claim construction that you are proposing for the replaying step. I just want to get on the record what exactly are you using for support for that claim construction that the replaying is done to a destination device as opposed to replaying transmission that originally was to a destination device.
>
> MR. McCOMBS: Your Honors, the only discussion in the entire specification of the patent is, when there is a replaying transmission, that it is done to a virtual machine. And that's described in the specification at column 29, line 36 through 42, and then at column 29, line 56 through 60.
>
> What is happening is the replaying, it is a simulation of a transmission where it is a virtual machine that is simulating the destination device. That's the only time that there is any replaying done in the patent in an analysis environment to a

> destination device, which is a virtual machine that is simulating the original destination device.
>
> There is not a discussion of actually replaying transmission back out onto the communications network to some original destination. That is never described in the patent.

Tr. 48:3–24.

We understand Patent Owner's reading of these claim phrases to be that replaying the transmission of data requires replaying the transmission to a destination device such as a virtual machine. However, we do not agree that this is the broadest reasonable interpretation of these phrases. For example, claim 1 recites "replay *transmission* of . . . data copied from the communication network to a destination device." The term "replay" appears logically and grammatically to apply to the term "transmission," which immediately follows "replay." Further, the term "transmission" is modified by the following phrase "of the suspicious, flagged network data," which describes the "transmission" as a transmission of suspicious, flagged network data. Claim 1 further describes the "data" as "copied from the communication network to a destination device." Thus, we read the phrase "copied from the communication network to a destination device" as applying to "data," and not requiring that "replay" occurs to a destination device. Similarly, we read the corresponding language in claim 28 as applying the phrase "copied from the network to a destination device" to the "copied network data" and not to "replay transmission." Additionally, for claim 17, we read the claim language "which was copied from the communication network to a destination device" to apply to "copied network data" and not to "replaying transmission."

Our reading of the claim language is consistent with the disclosure of the ’553 patent. The ’553 patent uses the term “destination device” to describe original destination device 710 that receives the transmission of data from Source Device 705 via Communication Network 720. Ex. 1001, 26:18–43, 29:56–60, Figs. 7, 10. Further, the sections of the ’553 patent cited by the Patent Owner do not support Patent Owner’s proposed interpretation of these phrases. Column 29, lines 36 through 42 do not refer to a destination device. Column 29, lines 56 through 60 disclose that virtual machine 815 *simulates* destination device 710. In other words, the ’553 patent does not teach that virtual machine 815 is “a destination device,” instead it teaches that a virtual machine may simulate or mimic the original destination device. *Id.* at 29:56–60, Fig. 10. Additionally, we note that to the extent Patent Owner contends the recited replay phrases require replay to a virtual machine, the claim language does not recite a virtual machine.

6. Other Claim Terms

Patent Owner further proposes constructions for claim terms “determine” and “determination.” PO Resp. 15–16. Nonetheless, based on the evidence of record, these terms do not require express construction for the purposes of this Decision.

B. Claims 1, 5, 7, 17, 22, and 25–27 – Obviousness over Kaeo (Ex. 1006) and Venezia (Ex. 1005)

Petitioner argues that claims 1, 5, 7, 17, 22, and 25–27 are unpatentable under 35 U.S.C. § 103(a) over Kaeo and Venezia. Pet. 20–60. As explained in further detail below, having considered the arguments and evidence presented, we are persuaded that Petitioner has shown, by a preponderance of the evidence, that claims 1, 5, 7, 17, and 25–27 are unpatentable over Kaeo and Venezia. We are not persuaded of the same for claim 22.

1. Relevant Legal Principles

A claim is unpatentable under 35 U.S.C. § 103(a) if the differences between the claimed subject matter and the prior art are such that the subject matter, as a whole, would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. *KSR Int'l Co. v. Teleflex Inc.*, 550 U.S. 398, 406 (2007). The question of obviousness is resolved on the basis of underlying factual determinations including (1) the scope and content of the prior art; (2) any differences between the claimed subject matter and the prior art; (3) the level of skill in the art; and, (4) where in evidence, so-called secondary considerations, including commercial success, long-felt but unsolved needs, failure of others, and unexpected results. *Graham v. John Deere*, 383 U.S. 1, 17–18 (1966) ("the *Graham* factors"). The level of ordinary skill in the art usually is evidenced by the references themselves. *See Okajima v. Bourdeau*, 261 F.3d 1350, 1355 (Fed. Cir. 2001); *In re GPAC Inc.*, 57 F.3d 1573, 1579 (Fed. Cir. 1995); *In re Oelrich*, 579 F.2d 86, 91 (CCPA 1978).

For an obviousness analysis, prior art references "must be 'considered together with the knowledge of one of ordinary skill in the pertinent art.'" *In re Paulsen*, 30 F.3d 1475, 1480 (Fed. Cir. 1994) (quoting *In re Samour*, 571 F.2d 559, 562 (CCPA 1978)). Moreover, "it is proper to take into account not only specific teachings of the reference but also the inferences which one skilled in the art would reasonably be expected to draw therefrom." *In re Preda*, 401 F.2d 825, 826 (CCPA 1968). That is because an obviousness analysis "need not seek out precise teachings directed to the specific subject matter of the challenged claim, for a court can take account of the inferences and creative steps that a person of ordinary skill in the art would employ." *KSR*, 550 U.S. at 418; *see In re Translogic Tech., Inc.*, 504 F.3d at 1259.

2. Level of Ordinary Skill in the Art

The parties agree that a person of ordinary skill in the art would have the following education and/or experience: a recent degree in a field such as computer science or computer networking and two or more years of experience in the computer networking or computer security industry. PO Resp. 16–17; Ex. 1003 ¶ 33. "Alternatively, in lieu of recent formal education, a person of ordinary skill in the art would have had six or more years of relevant experience in the computer networking or computer security industry." PO Resp. 16–17. This level of ordinary skill in the art is consistent with the ordinary skill reflected in the prior art of record, which is directed to computer networking and computer security systems. For example, Venezia and Kaeo both disclose intrusion-detection-systems. Ex. 1005; Ex. 1006.

With this level of ordinary skill in mind, we now turn to the analysis of the differences between the asserted prior art references and the subject matter recited in the claims-at-issue.

3. Summary of Venezia (Ex. 1005)

Venezia discloses the performance of NetDetector, an intrusion-detection-system ("IDS"). Ex. 1005, 1. Venezia states that "[r]ather than simply capturing the packet headers of monitored data streams, and examining them for possible attacks, the NetDetector stores every packet, from header to payload, in an indexed database." *Id*. Venezia adds that NetDetector notifies an administrator of an attack and allows the administrator to playback or "reconstruct the attack, keystroke by keystroke, packet by packet." *Id.* Venezia further indicates that NetDetector relies on Snort, an open source IDS, for intrusion detection. *Id.* at 2. Snort is described as being able to "monitor all traffic or a selected segment (based on filtering rules) on any given interface." *Id.* Venezia also states that "it's possible to select a specific time frame or capture and reprocess that traffic stream through the IDS engine." *Id.* Venezia explains that once an attack or signature has been identified, every packet comprising that event is available. *Id.*

4. Summary of Kaeo (Ex. 1006)

Kaeo describes various design options for network security, including intrusion detection systems based on statistical analysis and rule-based methods. Ex. 1006, 361. Kaeo indicates that the rule-based analysis method "uses rules that characterize known security attack scenarios and raise an alarm if observed activity matches any of its encoded rules." *Id.* "This analysis can also detect intruders who exhibit specific patterns of behavior known to be suspicious or in violation of site security policy." *Id.* Kaeo adds that most rule-based systems are user configurable so that the user can define her own rules based on her own corporate environment. *Id*. Kaeo also describes network intrusion detection systems with cable taps that serve as "[p]assive Ethernet taps . . . where 'copies' of the frames are sent to a second switch dedicated to IDS sensors." *Id.* at 362, Fig. 8-2. Additionally, Kaeo teaches that "honey pots" are locations to send suspected traffic to/from an attack. *Id.* at 363. The data then can be collectively analyzed to mitigate some possible attacks. *Id.*

5. Analysis

a. Claims 1 and 17

Petitioner contends that Kaeo and Venezia teach or suggest all the limitations of claims 1, 5, 7, 17, 22, and 25–27. Pet. 20–59. We have reviewed the Petition, the Patent Owner's Response, and Petitioner's Reply, as well as the evidence discussed in each of those papers, and are persuaded that Petitioner has shown, by a preponderance of the evidence, that claims 1 and 17 would have been obvious based on Kaeo and Venezia. Our discussion below focuses on the limitations of independent claim 1, which are illustrative and largely overlap with limitations recited in independent claim 17. However, to the extent the limitations of independent claim 17 require separate treatment, those limitations are discussed separately below. Additionally, dependent claims 5, 7, 22, and 25–27 are discussed in a following section.

Claim 1 recites "a tap configured to copy network data from a communication network" and "a controller coupled to the tap and configured to receive the copy of the network data from the tap." Petitioner asserts that Kaeo's disclosure of cable taps or a SPAN/mirror port coupled to a network intrusion detection system meets these limitations. Pet. 30–32. We find Petitioner has shown sufficiently that Kaeo teaches these limitations.

Claim 1 further recites a controller that is configured to "analyze the copy of the network data with a heuristic to determine if the copy of the network data has one or more characteristics of a computer worm." We are persuaded by Petitioner's assertion that Kaeo's disclosure of a network intrusion detection system that performs IDS rule-based analysis and statistical analysis satisfies this limitation. Pet. 33–35.

Additionally, claim 1 requires that the controller is configured to

> flag at least a portion of the copy of the network data as suspicious by flagging the at least a portion of the copy of the network data for replay in an analysis environment based upon the heuristic determination that the at least a portion of the analyzed copy of the network data has one or more characteristics of a computer worm, and replay transmission of the suspicious, flagged network data copied from the communication network to a destination device.

For these limitations, Petitioner asserts that Venezia's NetDetector "stores every packet, from header to payload, in an indexed database," which "not only permits an administrator to be notified when an attack has occurred but also to reconstruct the attack, keystroke by keystroke, packet by packet, and determine the exact commands issued by the attacker, in addition to any files or other data that was transmitted to or from the compromised system." Pet. 14 (citing Ex. 1005, 1). Petitioner adds that Venezia further teaches that once NetDetector has identified a particular attack or signature, every packet comprising that event is available in raw packet form with the option to replay the session just as it was recorded. *Id.*

Patent Owner argues that Venezia does not disclose "flagging . . . for replay" required in claim 1, because NetDetector's replay occurs at the option of an administrator and does not occur automatically after NetDetector identifies network data matching an attack signature. PO Resp. 19 (citing Ex. 2009 ¶¶ 69–71). Patent Owner adds that the replay decision is made by a human administrator and NetDetector does not have the ability to decide whether or not to replay data. *Id.* Patent Owner further argues that Venezia's examples of replay involve data that was not identified as suspicious. *Id.* at 19–20. Specifically, there is no indication of an attack for the replay of an AOL Instant Messenger ("AIM") session (Ex. 1005) or the replay discussed in the Niksun white paper (Ex. 1012)[6]. *Id.*

We do not agree with Patent Owner's arguments. First, as written, claim 1 requires "flagging . . . for replay," but does not indicate expressly that the replay occurs *automatically* after flagging. Further, Patent Owner has not explained sufficiently how the claim language requires automatic replay otherwise.

Second, we also do not agree that the "flagging" limitation excludes a replay decision made by a human administrator. Claim 1 requires that the recited controller is configured to "flag . . . data as suspicious by flagging . . . the network data for replay." However, claim 1 does not recite that the controller (or any other component) must decide whether or when the replay occurs.

Third, we are not persuaded that Venezia does not teach or suggest the replay of data that has been identified as suspicious. Specifically, as Petitioner argues, Venezia describes the replay of the AIM session as an example of how data is replayed once it has been recorded. This example of replay is given in the context of having first identified an attack prior to replay. Ex. 1005, 2 ("once a particular attack or signature has been identified, every packet comprising that event is available both in raw packet form."). Moreover, Petitioner points to Venezia's teaching that an administrator can reconstruct an attack, keystroke by keystroke, packet by packet, after being notified of an attack. Pet. 15 (citing Ex. 1005, 1). Thus, we find that Venezia teaches that once an attack event has been identified,

[6] The Petition refers Ex. 1012 ("Niksun"), titled "Network Security – NetDetector Intrusion Forensic System," as further describing the operation of the NetDetector system disclosed in Venezia. *See* Pet. 35.

the data for that event is recorded such that it can be replayed as described in the example of the recorded AIM session. *Id.*

Next, Patent Owner argues that Venezia does not teach an "analysis environment" because (1) an administrator, rather than the environment, performs the analysis rather than the environment; and (2) Venezia's NetDetector does not provide any ability to replay the packets to a destination device and then reconstruct or otherwise display the effect of those packets on the destination device. PO Resp. 20–24 (citing Ex. 2009 ¶¶ 76–79).

As discussed, we find that the language of claim 1 does not exclude manual analysis and, further, does not require the analysis environment to perform the analysis under the broadest reasonable construction of "analysis environment." *See supra* Claim Construction. Further, we do not agree that the claim term "analysis environment" requires replay to a destination device. As discussed above, we construe "analysis environment" to mean "an environment in which analysis of the effect of the network data upon a destination device is performed." The "analysis environment" provides an environment for analyzing the effects on a destination device, but does not require that the data must be replayed to a destination device. Further, as discussed in the context of the of the claim phrase "replay transmission of the suspicious, flagged network data copied from the communication network to a destination device" (e.g., claim 1), the '553 patent consistently describes "destination device" as the device receiving data in the original transmission, and not a separate device receiving the replayed data (e.g., virtual machine). Ex. 1001, 26:18–43, 29:56–60, Figs. 7, 10. Thus, we do not agree with Patent Owner that the term "analysis environment" requires replay to a destination device.

Patent Owner further asserts that Venezia's ability to reconstruct an attack does not teach replay in an analysis environment because NetDetector's reconstruction does not contain information about the effect of packets on a destination device after an administrator replays them. PO Resp. 23–24. Petitioner responds that Venezia's ability to reconstruct an attack teaches playback or replay because the screenshot on page 1 of Exhibit 1005 shows analysis of "retrieve" operations upon a "compromised server." Reply 7–8 (citing Ex. 1005, 1; Ex. 1031 ¶¶ 65–66). Petitioner's declarant, Dr. Trent Jaeger, testifies that "NetDetector's attack reconstruction screenshot . . . demonstrates that NetDetector also evaluates and shows analysis of the effects of the data (e.g., operations resulting from an attacker's transmitted commands) upon the destination device (e.g., a compromised server)." Ex. 1031 ¶ 65. Dr. Jaeger further testifies that the screenshot on page 1 of Ex. 1005 shows the reconstructed commands from the attacker (i.e., "User Interaction") resulted in computer operations ("Retrieve smurf.c" and "Retrieve newones") being executed upon a compromised server. *Id.*

We are persuaded by Petitioner's assertion that Venezia's reconstruction satisfies the limitation "replay in an analysis environment." Referring to the screenshot on page 1 of Exhibit 1005, Venezia teaches

> NetDetector dissects attacks and allows administrators to reconstruct them. Here we see that an attacker used FTP to pull the files 'smurf.c' and 'newones' to a compromised server. By clicking on the file names, we can even view the contents of the transmitted files.

Ex. 1005, 1 (screenshot on left). Further, in describing reconstruction, Venezia teaches that

> NetDetector stores every packet, from header to payload, in an

> indexed database. This not only permits an administrator to be notified when an attack has occurred but also to *reconstruct the attack, keystroke by keystroke, packet by packet, and determine the exact commands issued by the attacker, in addition to any files or other data that was transmitted to or from the compromised system*.

Id. (emphasis added).

Additionally, with respect to claim 17, Patent Owner argues that Kaeo and Venezia fail to disclose

> replaying transmission . . . to identify unauthorized activity based on playback of the flagged suspicious at least a portion of the analyzed copy of the network data.

Ex. 1001, 33:19–24. *See* PO Resp. 24–26. Patent Owner asserts that the replay examples provided in Venezia and the NetDetector whitepaper (Ex. 1012) involve replay of data that is not suspicious. *Id.* at 24.

As Petitioner points out, Venezia discloses that attack reconstruction is based on "NetDetector's playback capability," and that the attack reconstruction feature is used to "determine the exact commands issued by the attacker" or identify "that an attacker used FTP to pull the files 'smurf.c' and 'newones' to a compromised server." Pet. 17; Reply 8–9. We agree with Petitioner that this described reconstruction of the attack, such as that shown in the screenshot on page 1 of Exhibit 1005, allows an administrator to "identify unauthorized activity based on the playback of the flagged . . . data."

Patent Owner further argues that Petitioner has not provided a reason for why a person of ordinary skill in the art would modify Venezia to (1) flag for replay and (2) include an analysis environment. PO Resp. 24–26. Nonetheless, Petitioner has explained sufficiently how Venezia teaches both flagging for replay and an analysis environment. Thus, Petitioner does not rely on a modification of Venezia to disclose these features. Additionally, we are persuaded that Petitioner has provided sufficient rationale to combine the teachings of Kaeo and Venezia. Pet. 22–23. These rationales include combining Venezia's replay features with Kaeo's IDS to minimize false positives in Kaeo's IDS through further verification of suspicious packets with Venezia's replay. *Id*.

Accordingly, based on the evidence and arguments presented, we find that Petitioner has shown, by a preponderance of the evidence, supported by articulated reasoning with rational underpinning, that claims 1 and 17 would have been obvious over Kaeo and Venezia.

a. Claims 5, 7, and 25–27

Petitioner provides detailed explanations of how each limitation of dependent claims 5, 7, and 25–27 is taught or suggested by the combination of Kaeo and Venezia. Pet. 42–43, 44–45, and 58–59. These claims are discussed below.

Claim 5 depends from claim 1 and further recites "wherein the controller further comprises a policy engine configured to flag the at least a portion of the analyzed copy of the network data as suspicious based on comparing the at least a portion of the analyzed copy of the network data to policies." Ex. 1001, 32:19–23. Claim 26 depends from claim 17 and recites "wherein analyzing the copied network data flags the at least a portion of the copied network data as suspicious by comparing the copied network data to policies within a policy engine." *Id.* at 34:7–11.

For both claims 5 and 26, Petitioner argues Kaeo discloses an intrusion detection system using a combination of statistical analysis to flag observed activity and rules-based analysis to flag observed activity matching encoded rules. Pet. 42–43, 59 (citing Ex. 1006, 361–363). Petitioner's arguments are persuasive. As Petitioner discerns, Kaeo teaches rule-based analysis "can also detect intruders who exhibit specific patterns of behavior . . . in violation of site security policy." Pet. 43 (citing Ex. 1006, 361).

Claim 7 depends from claim 1 and claim 27 depends from claim 17. Both recite "wherein the one or more characteristics of a computer worm include being configured to duplicate itself for propagation." Ex. 1001, 32:27–29, 34:12–14.

Petitioner argues that Kaeo's disclosure of the worm propagation satisfies this limitation. Pet. 44 (citing Ex. 1006, 355–57). Specifically, Kaeo teaches the Sapphire worm doubled in size every 8.5 seconds. *Id.* Based on the evidence and arguments presented, we find Petitioner's assertions persuasive.

Claim 25 depends from claim 17 and recites "wherein identifying the unauthorized activity includes identifying of a hacker associated with the network data." Ex. 1001, 34:4–6. Petitioner asserts Kaeo satisfies this limitation because it teaches the "[s]tatistical analysis may also detect intruders who exploit previously unknown vulnerabilities that cannot be detected by any other means . . . [and the] rule-based analysis method . . . can also detect intruders who exhibit specific patterns of behavior known to be suspicious." Pet. 58 (citing Ex. 1006, 361) (emphasis and parenthetical information omitted). We find Petitioner's arguments persuasive.

Accordingly, upon consideration of the Petition's analysis and supporting evidence, the Patent Owner's Response, Petitioner's Reply, and the evidence and arguments presented, we are persuaded that Petitioner has shown, by a preponderance of the evidence, that claims 5, 7, and 25–27 would have been obvious over Kaeo and Venezia.

b. Claim 22

Claim 22 depends from claim 17 and further recites "wherein the heuristic is configured to detect the network data sent to an unassigned internet protocol address." Ex. 1001, 33:45–47.

Upon consideration of the Petition, the Patent Owner's Response, and Petitioner's Reply, as well as the evidence discussed in each of those papers, we are not persuaded, by a preponderance of the evidence, that claim 22 would have been obvious based on Kaeo and Venezia. Specifically, we are not persuaded by Petitioner's assertion that Kaeo's disclosure of a Cisco router satisfies these limitations. *See* Pet. 40–41 (citing Ex. 1006, 656). Petitioner has not explained sufficiently in the Petition how or why one of ordinary skill in the art would have modified Kaeo's network intrusion detection system (Ex. 1006, 360–61) with features from Kaeo's separately described router (*id.* at 656).

C. Claims 6, 8, 12–14, 16, 18, and 19 – Obviousness over Kaeo, Venezia, and Chen (Ex. 1009)

Petitioner argues that claims 6, 8, 12–14, 16, 18, and 19 are unpatentable under 35 U.S.C. § 103(a) over Kaeo, Venezia, and Chen. Pet. 43–51, 54–57. As explained in further detail below, having considered the arguments and evidence presented, we are persuaded that Petitioner has shown, by a preponderance of the evidence, that claim 6 is unpatentable over Kaeo, Venezia, and Chen. We are not persuaded of the same for claims 8, 12–14, 16, 18, and 19.

1. Summary of Chen (Ex. 1009)

Chen is a position paper titled "When Virtual is Better than Real," which proposes that "the operating system and applications currently running on a real machine should relocate into a virtual machine." Ex. 1009,

1. Chen provides a virtual-machine structure that runs on the host machine. *Id*. at Fig. 1. Rather than running suspicious events on the real system, which risks compromising the system, Chen suggests safely conducting "this type of a test on a *clone* of the real system." *Id*. at 4. Chen adds that

> [v]irtual machines make it easy to clone a running system, and an intrusion preventer can use this clone to test how a suspicious input event would affect the real system. The clone can be run as a hot standby by keeping it synchronized with the real system (using primary-backup techniques), or it can be created on the fly in response to suspicious events. In either case, clones admit more powerful intrusion preventers by looking at the response of the system to the input event rather than looking only at the input event. Because clones are isolated from the real system, they also allow an intrusion preventer to run potentially destructive tests to verify the system's health. For example, an intrusion preventer could forward a suspicious packet to a clone and see if it crashes any running processes. Or it could process suspicious input on the clone, then see if the clone still responds to shutdown commands.

Id. Chen teaches that "[l]ike a network-based intrusion detector, virtual-machine-based intrusion detectors are separate from the guest operating systems and applications. Unlike network intrusion detectors, however, virtual-machine intrusion detectors can see all events occurring in the virtual machine they monitor." *Id*.

2. *Analysis*

a. *Claim 6*

Claim 6 depends from claim 1 and further recites "wherein the controller further comprises a virtual machine pool configured to store a virtual machine." Ex. 1001, 32:24–26.

Petitioner asserts that Chen's disclosure of running a clone as a "hot standby" or "on the fly" indicates that Chen stores a virtual machine in a virtual machine pool. Pet. 43–44. Petitioner's declarant, Dr. Jaeger, adds that "[t]he ability to run a virtual machine on hot standby or on the fly in response to suspicious events meets the claimed virtual machine pool because one of skill in the art understood that at least one virtual machine had to be stored in order to perform such functionality." Ex. 1003 ¶ 287. Petitioner further asserts that Chen teaches it was widely known to have multiple virtual machines running multiple operating systems as demonstrated by commercial products such as VMware and VirtualPC. Pet. 44 (citing Ex. 1009, 1). Additionally, Petitioner asserts that a combination with Chen would enhance Venezia's and Kaeo's systems by minimizing the risk of compromising the real system. *Id*. at 25–26; Ex. 1003 ¶¶ 273–284.

Patent Owner argues that Petitioner has not provided sufficient reason to combine the teachings of Kaeo, Venezia, and Chen. PO Resp. 36–38. Specifically, Patent Owner asserts that adding Chen's virtual machine-based intrusion detection would not enhance Venezia's replay by minimizing risk to the real system, because NetDetector and Chen operate in entirely different locations in the network which traditionally employ different security solutions. NetDetector monitors traffic for the entire network; Chen's intrusion preventer focuses on a single host. PO Resp. 37 (citing Ex. 2008 ¶¶ 65, 116). More specifically, Patent Owner argues that the combination of Venezia and Chen raises several unanswered questions as to how a person of ordinary skill in the art would incorporate Chen's virtual machine with Venezia's NetDetector. *Id.* (citing Ex. 2008 ¶¶ 154–164). For example, Patent Owner questions (1) the implications and costs of moving Chen's clone to "the network periphery"; (2) whether every host system needs to be virtualized and clonable at the periphery; (3) whether end host systems are protected when NetDetector "cannot block the original packets from being delivered"; (4) why the NetDetector/Chen combination would be moved to a host; and (5) whether initiating replay by a human administrator to protect a system as described by Petitioner is "really worth that price." *Id.* at 38–39.

Patent Owner's arguments are based on various combinations of the references that have not been asserted by the Petitioner for claim 6. Petitioner has not argued that Chen's clone must be added to a network periphery or that NetDetector must be physically incorporated into Chen's system. Rather, Petitioner asserts that it would have been obvious to one of ordinary skill in the art to modify Venezia to have a controller that includes a virtual machine pool configured to store a virtual machine because Chen teaches the use of virtual machine-based intrusion detection and storage for storing multiple virtual machines. *See* Pet. 25–26, 43 (citing Ex. 1003 ¶¶ 285–288).

Additionally, we are not persuaded by Patent Owner's questions regarding the costs and practicality of modifying the NetDetector or Chen system as Petitioner proposes. For example, Patent Owner relies on the testimony of its declarant, Dr. Tal Garfinkel, who testifies that for the asserted combination

> [t]he person of skill would either need to replay traffic from the IDS to the clone, or they would need to copy the clone over to the NIDS, and then send the traffic. Both options involve significant amounts of novel engineering. If the person of skill sent flagged data to the Chen clone, they would need to modify Venezia to, at the appropriate time, permit the administrator to communicate to the Chen virtual machine monitor to make a clone of the "real" virtual machine available (we don't know how to do this without a substantial performance impact on the real destination host; it's not even clear if this could be made

> practical, and certainly not with the hardware constraints present in 2003).

Ex. 2008 ¶ 159. However, the mere existence of disadvantages resulting from a modification does not refute the obviousness of the modification, especially when the prior art indicates that the modification also offers an advantage. Tradeoffs regarding features, costs, manufacturability, or the like, do not necessarily prevent the combination. *See Medichem, S.A. v. Rolabo, S.L.*, 437 F.3d 1157, 1165 (Fed. Cir. 2006) ("[A] given course of action often has simultaneous advantages and disadvantages, and this does not necessarily obviate motivation to combine." (citations omitted)); *Winner Int'l Royalty Corp. v. Wang*, 202 F.3d 1340, 1349 n.8 (Fed. Cir. 2000) ("The fact that the motivating benefit comes at the expense of another benefit, however, should not nullify its use as a basis to modify the disclosure of one reference with the teachings of another. Instead, the benefits, both lost and gained, should be weighed against one another.").

Further, to the extent that Patent Owner argues the asserted references are non-enabling, we note that although "a prior art reference cannot anticipate a claimed invention if the allegedly anticipatory disclosures cited as prior art are not enabled," *In re Antor Media Corp.*, 689 F.3d 1282, 1289 (Fed. Cir. 2012) (citations and internal quotation marks omitted), a non-enabling reference may qualify as prior art for the purpose of determining obviousness, *Symbol Techs., Inc. v. Opticon, Inc.*, 935 F.2d 1569, 1578 (Fed. Cir. 1991) (citations omitted); *see Geo. M. Martin Co. v. Alliance Mach. Sys. Int'l LLC*, 618 F.3d 1294, 1302 (Fed. Cir. 2010) ("Under an obviousness analysis, a reference need not work to qualify as prior art; it qualifies as prior art, regardless, for whatever is disclosed therein." (citations and internal quotation marks omitted)); *Beckman Instruments, Inc. v. LKB Produkter AB*, 892 F.2d 1547, 1551 (Fed.Cir.1989) ("Even if a reference discloses an inoperative device, it is prior art for all that it teaches." (citations omitted)).

Further, upon consideration of the evidence and arguments presented, we determine that Petitioner has articulated sufficient reasons for why one of ordinary skill in the art would combine the teachings of Kaeo, Venezia, and Chen. These include that the combination enhances the capability of each system to minimize the risk of compromising the real system by using a clone to test inputs, and the combination would have yielded nothing more than predictable results because Chen's virtual machine already received network data in the form of suspicious packets and each prior art reference in the combination is directed towards intrusion detection and thus the references share a similar goal. *See* Pet. 25–26.

Accordingly, we are persuaded that Petitioner has shown, by a preponderance of the evidence, that claim 6 would have been obvious over Kaeo, Venezia, and Chen.

b. Claim 8

Claim 8 is directed to an unauthorized activity capture system that includes a controller configured to "configure a replayer to replicate the at least a portion of the analyzed copy of the network data which contains suspicious activity to the virtual machine." Ex. 1001, 32:30–50.

Petitioner argues that Chen satisfies this limitation because it discloses (1) "*forwarding suspicious packets to a virtual machine to see if it crashes any running processes*" and (2) cooperative logging by virtual machines. Pet. 47 (citing Ex. 1009, 3–4). Patent Owner disagrees that Chen teaches a configurable replayer. PO Resp. 32–33. Patent Owner argues that although Chen discloses "an intrusion preventer can *forward* a suspicious packet to a clone and see if it crashes any running processes," this disclosure does not describe configuration of a replayer. *Id.* at 33 (citing Ex. 2008 ¶ 136). Patent Owner further argues that Chen's description of cooperative logging also does not teach how to configure a replayer (e.g., NetDetector) to replay data to Chen. *Id.* Moreover, Patent Owner contends Petitioner's arguments conflate the Chen logging server with the Chen intrusion preventer without explaining why the two can be combined. *Id.* (citing Ex. 2008 ¶ 141).

Upon consideration of the evidence and arguments presented, we determine Petitioner's arguments are not persuasive. Although Chen discloses *forwarding* suspicious packets to a clone, Petitioner has not explained sufficiently how forwarding packets teaches or suggests "*configuring* a replayer to transmit . . . data to the virtual machine." Moreover, Petitioner does not explain how Chen's logging disclosure (Ex. 1009, 2–3) applies to configuring a replayer. At the oral hearing, Petitioner was asked to explain its reliance on Chen for a similar configuration limitation recited in claim 20 of the '499 patent.

> JUDGE IPPOLITO: For Chen what are you relying on for the configuring step? I assume you are talking about claim 20 for the '499 patent, for example.
>
> MR. HANNAH: Thank you, yes. For the configuring step it would be loading up, creating on-the-fly, in response to the suspicious events, the loading of the virtual machine. So that is retrieving a virtual machine, loading it up and then, also, causing the replay to happen is by sending those packets to the clone and watching what the clone does.
>
> So that is actually replaying the packet within the virtual environment, looking at its behavior to determine whether it crashes or not. So Chen is very descriptive in terms of how it works in terms of the virtual machine and is spot-on to what is described in the '499 and the '553.
>
> JUDGE IPPOLITO: Well, I recall in the petition that

> there was some reliance on what is described as cooperative logging, is that correct, for Chen? I want to get a better sense of how that fits into your argument.
>
> MR. HANNAH: Sure. The cooperative logging was one example of how you could have these messages be sent to each other in terms of the packets, and just showing, it was showing that it can retrieve and be able to manipulate and look at network traffic in terms of packets specifically. And that's referred to on page 8 of Chen.

Tr. 33:5–34:4.

We agree with Petitioner that Chen runs a virtual machine on hot standby and can create a clone of the real system "on the fly" in response to suspicious events, which allows the virtual machine to be retrieved to receive flagged data. However, we are not persuaded that the *retrieval* of the virtual machine teaches the step of *configuring* a *replayer,* such as Venezia's NetDetector, to transmit data to a virtual machine. Moreover, although Petitioner points to cooperative logging as showing how messages can be sent by the Chen system, Petitioner does not explain sufficiently how the disclosed logging relates to configuring a replayer. We decline to speculate on how the cited disclosure in Chen supports Petitioner's position.

Additionally, we note that in Petitioner's Reply, Petitioner states Patent Owner's arguments are conclusory and "cannot overcome the evidence cited by the Petition nor the Board's analysis of that evidence." Reply 14 (citations omitted). For clarity, we reiterate that in an *inter partes* review, the burden is on the Petitioner to show, by a preponderance of the evidence, that a claim is unpatentable. Our discussion in the Decision to Institute was based on a preliminary record at that stage of the proceeding and any decision to institute review of a claim under any ground does not create a presumption of unpatentability or absolve the Petitioner of the burden ultimately to satisfy the required showing for unpatentability. Based on the complete record before us, we are not persuaded Petitioner has met its burden for claim 8.

Accordingly, we are not persuaded that Petitioner has shown, by a preponderance of the evidence, that claim 8 would have been obvious over Kaeo, Venezia, and Chen. Claim 12–14 and 16 depend from claim 8. Ex. 1001, 32:59–67, 33:5–7. For the same reasons, we are not persuaded that claims 12–14 and 16 are unpatentable over Kaeo, Venezia, and Chen.

c. Claims 18 and 19

Claim 18 depends from claim 17 and further recites, *inter alia,* "configuring a replayer to transmit the flagged at least a portion of the analyzed copied network data to the virtual machine." Ex. 1001, 33:25–36.

Again, Petitioner argues that Chen teaches configuration of a replayer because Chen discloses that "an intrusion preventer could forward a suspicious packet to a clone and see if it crashes any running processes" and cooperative logging can be used to replay data. Pet. 56. Patent Owner also contends this disclosure in Chen does not teach or suggest "configuring a replayer." PO Resp. 32–34.

For the same reasons discussed above with respect to claim 8, Petitioner's arguments are not persuasive. Specifically, we are not persuaded that *forwarding* suspicious packets to a clone or cooperative logging teaches or suggests configuration of a replayer.

Accordingly, we find that Petitioner has not shown, by a preponderance of the evidence, that claim 18 would have been obvious over Kaeo, Venezia, and Chen. Claim 19 depends from claim 18. For the same reasons, we are not persuaded that claim 19 is unpatentable over Kaeo, Venezia, and Chen.

D. Claims 1, 3–5, 7, 17, and 22–28 – Obviousness based on Kaeo and Liljenstam (Ex.1007)

Petitioner argues that claims 1, 3–5, 7, 17, and 22–28 are unpatentable under 35 U.S.C. § 103(a) over Kaeo and Liljenstam. Pet. 30–45, 52–54, 58–60. As explained in further detail below, having considered the arguments and evidence presented, we are persuaded that Petitioner has shown, by a preponderance of the evidence, that claims 1, 3–5, 7, 17, and 22–28 are unpatentable over Kaeo and Liljenstam.

1. Summary of Liljenstam (Ex. 1007)

Liljenstam discloses a worm simulation model that generates realistic input traffic for a working prototype worm detection and tracking system, the Dartmouth ICMP BCC: System/Tracking and Fusion Engine (DIB:S/TRAFEN) system. Ex. 1007, 1. Liljenstam describes the DIB:S/TRAFEN system as capable of detecting and classifying "active Internet worms in their earliest stages of propagation." *Id*. at 2. To do so, the DIB:S/TRAFEN system collects copies of ICMP type 3 (ICMP-T3) unreachable messages. *Id*. at 4. Liljenstam explains that worms spread by randomly probing IP addresses and that "[t]his random scanning, however, will probe many unassigned IP addresses . . . that are not associated with a reachable computer." *Id*. at 2. "[R]outers that receive a packet destined for an unreachable IP address will drop the packet and return an *ICMP Destination Unreachable* (ICMP Type 3) message to the original originator." *Id*. These ICMP-T3 messages include the source and destination IP addresses. *Id*.

To collect the ICMP-T3 messages, the DIB:S system includes a select group of participating or instrumented routers that forward all the ICMP-T3 messages that they generate to an analysis station. Ex. 1007, 2, 5, Fig. 3. Instrumented routers in the Internet send copies of ICMP-T3 messages to the DIB:S system, which correlates and analyzes the data. *Id*. at 5, Fig. 3. As the ICMP-T3 messages arrive at the DIB:S analysis station, they are sorted and analyzed according to the embedded source and destination addresses and ports. *Id*. DIB:S generates a scan alert for worm detection when a single source machine uses the same protocol to contact the same port on target machines within a certain time interval. *Id*. at 2. TRAFEN detects whether there is an exponential increase in the number of alerts for the same port and protocol, which most likely indicates a propagating worm. *Id*.

Liljenstam discloses that the DIB:S/TRAFEN system performance was evaluated by feeding simulated ICMP-T3 unreachable traffic into the working DIB:S/TRAFEN prototype. Ex. 1007, 7. All packets arriving to the DIB:S system in the model are dumped to a file in binary tcpdump format. *Id*. at 4. Using the tcpreplay tool, the packet streams are replayed into the real DIB:S system to simulate the ICMP packets observed during the attack. *Id*. at 5. The system analyzes the ICMPs to identify scanning activity, and correlates that scanning activity to track worm infection. *Id*.

2. *Analysis*

a. *Claims 1, 17, and 28*

Below we discuss independent claim 1, which is illustrative of the subject-matter of independent claims 17 and 28.

Claim 1 recites "a tap configured to copy network data from a communication network" and "a controller coupled to the tap and configured to receive the copy of the network data from the tap." Petitioner asserts that Kaeo's disclosure of cable taps or a SPAN/mirror port coupled to a network intrusion detection system meets these limitations. Pet. 30–32. We find Petitioner has shown sufficiently that Kaeo teaches these limitations.

Claim 1 further recites a controller that is configured to "analyze the copy of the network data with a heuristic to determine if the copy of the network data has one or more characteristics of a computer worm." We are persuaded by Petitioner's assertion that Kaeo's disclosure of a network intrusion detection system that performs IDS rule-based analysis and statistical analysis satisfies this limitation. *See* Pet. 33–35.

Additionally, claim 1 requires that the recited controller is configured to:

> flag at least a portion of the copy of the network data as suspicious by flagging the at least a portion of the copy of the network data for replay in an analysis environment based upon the heuristic determination that the at least a portion of the analyzed copy of the network data has one or more characteristics of a computer worm, and replay transmission of the suspicious, flagged network data copied from the communication network to a destination device.

For these limitations, Petitioner asserts that Liljenstam discloses collecting copies of ICMP-T3 messages that are suspicious of worm scanning activity and organizing such packets into a binary tcpdump format to be replayed into the real DIB:S system for simulation of the ICMP packets observed during the attack. Pet. 36–38. Petitioner points to Liljenstam's collection of ICMP messages from instrumented routers and Liljenstam's worm simulation model where packets arriving to the DIB:S system are dumped into tcpdump format and replayed by a tcpreplay tool into the DIB:S system to simulate the ICMP packets observed during the attack. *Id*.

Patent Owner responds that the combination of Kaeo and Liljenstam does not teach or suggest an "analysis environment" because (1) Liljenstam does not teach that the DIB:S/TRAFEN system can monitor or analyze the behavior of a destination device; and (2) Liljenstam does not teach how a computer or human being uses the DIB:S/TRAFEN system to analyze the effect of network traffic on a destination device. PO Resp. 43–47. Patent Owner further argues a person of skill in the art considering Liljenstam would understand that ICMP-T3 packets do not have any meaningful effect on computer behavior. *Id.* at 44–45 (citing Ex. 2008 ¶ 180). Dr. Garfinkel testifies that

> ICMP-T3 messages simply provide notice that a host was unreachable; they do not cause a computer to behave in any unexpected or unpredictable way. Second, even if the ICMP-T3 messages could impact the behavior of a destination device, DIB:S would not be able to observe that behavior, since it only sorts the ICMP-T3 packets and looks for patterns that suggest scanning behavior.

Ex. 2008 ¶ 180. Patent Owner also contends Petitioner's expert, Dr. Jaeger, has not explained consistently what in Liljenstam teaches an "analysis environment." PO Resp. 45–47.

In its Reply, Petitioner argues that the "flagging the at least a portion of the copy of the network data for replay in an analysis environment" limitation does not require a "meaningful effect on computer behavior." Reply 10. Petitioner argues that a person of ordinary skill in the art would understand that Liljenstam discloses a DIB:S analysis that includes replaying (e.g., tcpreplay tool) to simulate the ICMP packets observed during the attack and that "[t]his particular worm detection system collects copies of ICMP messages generated by random scanning and tries to recognize signatures of early worm propagation by correlating the collected messages." *Id.* at 10–11. Petitioner also asserts that Dr. Jaeger provided a range of environments in Liljenstam that satisfy an "analysis environment." *Id.* at 12 (citing *KSR*, 550 U.S. at 421 (parenthetical information omitted)).

Petitioner's arguments are persuasive. Our construction of "analysis environment" does not require the analysis environment to actively perform an analysis. Moreover, our construction also does not require that the "analysis environment" analyze a *meaningful* effect of the network data upon a destination device. Petitioner has explained sufficiently how the DIB:S system provides an "analysis environment" for analyzing the effect of the network data upon a destination device. Pet. 36–38; *see* Tr. 26:4–19. In particular, Liljenstam teaches that "[u]sing the tcpreplay tool . . . , we can then replay the packet stream into the real DIB:S system to simulate the ICMP packets observed during the attack." Ex. 1007, 4–5. Liljenstam further discloses that the "system analyzes the ICMPs to identify scanning activity, and correlates that scanning activity to track worm infection." *Id.* at 5, Fig. 3 (emphasis omitted).

Patent Owner further asserts that Kaeo and Liljenstam do not teach:

> flagging the at least a portion of the copy of the network data for replay in an analysis environment based upon *the heuristic determination* that the at least a portion of the analyzed copy of the network data has one or more characteristics of a computer worm,

as recited in claim 1 (Ex. 1001, 32:1–6). PO Resp. 48–49. Patent Owner argues that Petitioner relies on Kaeo for the determination limitation during the compare/comparing step, but relies on a completely different function in Liljenstam for the flagging limitation. *Id.*

We do not agree with Patent Owner. In its Reply, Petitioner explains that "the Petition shows that Liljenstam alone as well as the combination of

Liljenstam and Kaeo disclose that its flagging can be based on the claimed comparison." Reply 11. Petitioner argues that "Liljenstam states that it can perform a policy comparison before sending the ICMP-T3 message to the analysis station," or that it would have been obvious for the flagging operation to utilize Kaeo's compared copied network data teachings. *Id.* (citing Pet. 37); *see* Tr. 28:13–31:10. Upon consideration of the evidence and arguments of record, we find Petitioner's assertions are persuasive.

Patent Owner also argues that Petitioner's articulated reasons for the combination of Kaeo and Liljenstam are inadequate. PO Resp. 49–50. Specifically, Patent Owner asserts that "a person of skill in the art would not couple Liljenstam's DIB:S system to a tap or span port in order to see both sides of a network conversation and to reduce false positives" (citing Pet. 26–27) because "this would result in a substantial and probably overwhelming load on the link between the tap and the DIB:S station." *Id.* Patent Owner further asserts that "because DIB:S/TRAFEN only collects ICMP-T3 host unreachable messages, it would not know what to do with any other network data that was collected from a network tap." *Id.* at 50 (citing Ex. 2008 ¶ 180). Petitioner responds that even if the combination of Kaeo and Liljenstam results in an inefficient system, the system still renders claim 1 obvious, because the evidence presented establishes that adding Liljenstam's replay to Kaeo's system reduces false positives and adding Kaeo's tap to Liljenstam provides better visibility of a full-duplex conversation. Reply 12–13; Pet. 24–25.

Petitioner's arguments are persuasive. Again, tradeoffs regarding features, costs, manufacturability, or the like, do not necessarily prevent the combination. *See Medichem*, 437 F.3d at 1165; *Winner Int'l Royalty Corp.*, 202 F.3d at 1349 n.8.

Accordingly, we are persuaded that Petitioner has shown, by a preponderance of the evidence, that claims 1, 17, and 28 would have been obvious over Kaeo and Liljenstam.

b. Claims 5, 7, and 25–27

Petitioner also provides detailed explanations of how each limitation of claims 5, 7, and 25–27 is taught or suggested by the combination of Kaeo and Liljenstam. Pet. 42–45, 58–59. More specifically, Petitioner relies on the disclosure in Kaeo, discussed above in connection with the obviousness ground based on Kaeo and Venezia, to demonstrate Kaeo teaches or suggests the required limitations. As discussed above, we are persuaded by Petitioner's arguments regarding the teachings of Kaeo. Accordingly, we determine that Petitioner has shown, by a preponderance of the evidence, that claims 5, 7, and 25–27 would have been obvious over Kaeo and Liljenstam.

c. Claims 3, 4, and 22–24

Claim 3 depends from claim 1 and claim 22 depends from claim 17. Ex. 1001, 32:12–14, 33:45–47. Both claims recite "wherein the heuristic is configured to detect the network data sent to an unassigned internet protocol address." *Id.* Petitioner asserts Liljenstam's disclosure of detecting packets destined for unreachable IP addresses satisfies this limitation. Pet. 41–42 (citing Ex. 1007, 1). Petitioner further notes that Liljenstam teaches

> routers that receive a packet destined for an unreachable IP address will drop the packet and return an ICMP Destination Unreachable (ICMP Type 3) message to the packet originator. This ICMP-T3 message will include the original IP header and at least 8 bytes of the protocol header, which together will include the source and destination IP addresses and port numbers for both UDP and TCP packets.

Id. at 42 (emphasis and citation omitted). We find Petitioner's arguments persuasive.

Claim 4 depends from claim 1 and claim 23 depends from claim 17. Both claims recite "wherein the heuristic is configured to detect the network data sent to an unassigned port address." Ex. 1001, 32:13–15, 33:48–50. For this limitation, Petitioner relies on Kaeo's teaching that port numbers may be used to filter and recognize services. Pet. 42 (citing Ex. 1006, 336–37), 58. Petitioner further relies on the testimony of Dr. Jaeger, who states "one of skill in the art understood Kaeo's suggestion to detect the destination IP address to which network data is directed is for the claimed policy as it is described within the chapter for 'Design and implementation of the Corporate Security Policy.'" Ex. 1003 ¶ 152; *see* Pet. 42 (citing Ex. 1003 ¶¶ 151–153). Petitioner's arguments are persuasive.

Claim 24 depends from claim 17 and further recites "wherein identifying the unauthorized activity includes identifying malware associated with the network data." Ex. 1001, 34:1–3. Petitioner argues that Kaeo discloses malware as a type of attack or vulnerability that can be dealt with at the policy level. Pet. 49–50, 58 (referring to analysis of claim 12). Petitioner's arguments are persuasive.

Accordingly, for claims 3, 4, and 22–24, we determine that Petitioner has shown, by a preponderance of the evidence, that these claims would have been obvious over Kaeo and Liljenstam.

E. Claims 18, 20, 29, and 30 – Obviousness over Kaeo, Liljenstam, and Dunlap (Ex. 1008)

Petitioner argues that claims 18, 20, 29, and 30 are unpatentable under 35 U.S.C. § 103(a) over Kaeo, Liljenstam, and Dunlap. Pet. 54–57, 60. As explained in further detail below, having considered the arguments and evidence presented, we are persuaded that Petitioner has shown, by a preponderance of the evidence, that claims 18, 20, 29, and 30 would have been obvious over Kaeo, Liljenstam, and Dunlap.

1. Summary of Dunlap (Ex. 1008)

Dunlap is an article titled "ReVirt: Enabling Intrusion Analysis through Virtual-Machine Logging and Replay." Ex. 1008, 2. Dunlap describes ReVirt as a post-intrusion analysis system capable of encapsulating the target system inside a virtual machine. *Id.* at 3. "ReVirt is able to replay the complete, instruction-by-instruction execution of the virtual machine, even if that execution depends on non-deterministic events such as interrupts and user input. An administrator can use this type of replay to answer arbitrarily detailed questions about what transpired before, during, and after an attack." *Id*. "Replay can be conducted on any host with the same processor type as the original host. Replaying on a different host allows an administrator to minimize downtime for the original host." *Id*. at 8.

2. Analysis

a. Claims 18, 20, and 29

Claim 18 depends from claim 17 and further recites,

> wherein replaying the transmission of the flagged at least a portion of the analyzed copied network data comprises:
>
> > retrieving a virtual machine configured to receive the flagged at least a portion of the analyzed copied network data;
> >
> > configuring a replayer to transmit the flagged at least a portion of the analyzed copied network data to the virtual machine; and
> >
> > performing a simulation by transmitting the

> previously flagged at least a portion of the analyzed copied network data to the virtual machine.

Ex. 1001, 33:25–36. Claim 29 recites similar limitations for the claimed non-transitory computer readable medium of 28. *Id.* at 34:33–43. For these limitations, Petitioner asserts that Dunlap's ReVirt system, which uses virtual machine UMLinux, can receive Liljenstam's collection of packets for replay. Pet. 55. Petitioner adds that "Dunlap's replay performs a simulation because it re-creates the sent data and replays the complete instruction-by-instruction execution of the virtual machine." *Id*. at 56 (emphasis and citations omitted). Additionally, Petitioner's declarant, Dr. Jaeger, testifies that one of ordinary skill in the art would understand "[t]his simulation is demonstrated by . . . ReVirt's . . . TUN/TAP virtual Ethernet device that emulates the network card. Thus, the transmission of flagged copied network data to the virtual machine is over an emulated or virtual network." Ex. 1003 ¶ 482.

Further, claim 20 depends from claim 18 and recites "the flagged at least a portion of the analyzed copied network data is transmitted between the replayer and the virtual machine over a virtual switch." Ex. 1001, 33:39–42. Petitioner asserts that Dunlap's TUN/TAP virtual Ethernet device meets the limitations of a virtual switch. Pet. 57.

Patent Owner argues that claims 18, 20, 29, and 30 are not unpatentable over Kaeo, Liljenstam, and Dunlap for the same reasons it argues independent claims 17 and 28 are not unpatentable over Kaeo and Liljenstam. PO Resp. 53. However, for the reasons discussed above, we find Petitioner has shown, by a preponderance of the evidence, that claims 17 and 28 are unpatentable over Kaeo and Liljenstam.

Next, Patent Owner argues that Dunlap does not teach how or why one of ordinary skill in the art would configure ReVirt to replay to receive data from Liljenstam. PO Resp. 53. Patent Owner further argues that "it is not even clear what data Liljenstam has that would even be useful to ReVirt." *Id.* Patent Owner argues Liljenstam is only concerned with ICMP-T3 packets, whereas ReVirt uses network data and nondeterministic input. *Id.* at 53–54. Patent Owner further argues Petitioner does not provide "no reason why a person of skill in the art would" retrieve a virtual machine and configure it to replay ICMP-T3 data. *Id.* at 54.

We do not agree with Patent Owner's arguments. Petitioner's declarant, Dr. Jaeger, testifies that

> Liljenstam already teaches "flagged at least a portion of the analyzed copy of the network data" because Liljenstam's collects suspicious copied packets in binary tcpdump format for replay using TCPReplay. *See* Liljenstam at 4 ("Instrumented routers in the Internet send copies of ICMP type 3 messages to the DIB:S system which correlates and analyzes the data…..As shown in Figure 3, all packets arriving to the DIB:S system in the model are dumped to a file in binary tcpdump format.") (FJN 1007).
>
> At the same time, Dunlap's ReVirt replay logged records saved to disk in syslog format. Dunlap at 7 ("Log records are added and saved to disk in a manner similar to that used by the Linux syslogd daemon.") (FJN 1008). Notably, one of skill in art at the time understood that data in tcpdump files can be easily converted into syslog format. *See* Boubalos at 1 (providing an automated utility to "extract syslog packets from tcpdump files") (FJN 1013). Thus, no technical impediment existed for Dunlap's ReVirt to replay Liljenstam's collection of packets in binary tcpdump format.

Ex. 1003 ¶¶ 475–76 (emphasis omitted). We credit Dr. Jaeger's testimony and are persuaded Petitioner presents sufficient evidence to support a finding that a person of ordinary skill in the art would have been able to modify

Dunlap's ReVirt to replay Liljenstam's packets. Further, we are persuaded that Petitioner has articulated sufficient reasons for the combination of Dunlap and Liljenstam. For example, Petitioner asserts Dunlap's replay to a virtual machine enhances Liljenstam's replay by maintaining the integrity of the compromised system. Pet. 29–30.

Additionally, Patent Owner argues that Kaeo, Liljenstam, and Dunlap do not teach how to configure a replayer. PO Resp. 54. However, Petitioner argues that Dunlap teaches a virtual machine called UMLinux that "is implemented as a loadable module . . . [and] [t]he VMM module is called before and after each signal and system call to/from the virtual-machine process." Pet. 54–55 (citing Ex. 1008, 4). Petitioner further asserts Dunlap teaches an X-proxy that acts as a new X-client. *Id.* at 56. The X-proxy sends "the same display messages to the X-server as the virtual machine did during logging." *Id.* (citation omitted). Petitioner also relies on the testimony of its declarant, Dr. Jaeger, who testifies that

> Dunlap demonstrates that its replayer also transmits by "sending the data stream to the X server." *See* Dunlap at 10 ("The X proxy accomplishes this by receiving the packets being (re)sent from the replaying virtual machine, stripping off the Ethernet, IP, and TCP headers from these packets, reconstituting the X window data stream, and sending the data stream to the X server.").

Ex. 1003 ¶ 480 (emphasis omitted).

Patent Owner asserts that the X-proxy is not a replayer and the X-server is not a virtual machine (GUI). PO Resp. 55. Patent Owner explains that the purpose of the X-proxy is to act as a new X-client that sends display messages to the X-server, and that the X-server is outside the control of the ReVirt virtual machine. *Id.*

Although we agree with Patent Owner that Dunlap discloses the X-server can be outside of a virtual machine, we note that Petitioner also explains that "*Dunlap states that the X server can be located in a virtual machine:* '…we could move the <u>X server</u> into another <u>virtual machine.</u>'" Pet. 57 (quoting Ex. 1008, 7 (emphasis added)). Moreover, we understand Petitioner's argument to be that the described communication feature between the X-proxy and X-server teaches or suggests how one of ordinary skill in the art would configure a replayer to transmit data to a virtual machine. Tr. 81:1–22. Thus, we find Petitioner has explained sufficiently how the asserted references satisfy this limitation.

Patent Owner further argues that Kaeo, Liljenstam, and Dunlap do not teach or suggest "performing a simulation by transmitting the flagged . . . network data to the virtual machine" because Liljenstam does not teach transmitting data to the virtual machine and Dunlap's virtual machine cannot receive data from Liljenstam. PO Resp. 56–57. However, we find that Petitioner has explained sufficiently that Liljenstam discloses flagging ICMP-T3 messages/packets that could be transmitted to Dunlap's virtual machine for performing a simulation. Pet. 55–56.

Accordingly, we determine that Petitioner has shown, by a preponderance of the evidence, that claims 18, 20, and 29 are unpatentable over Kaeo, Liljenstam, and Dunlap.

b. Claim 30

Claim 30 depends from claim 28, and further recites "wherein the one or more characteristics of a computer worm include being configured to duplicate itself for propagation." Ex. 1001, 34:44–47. Petitioner argues that Kaeo's disclosure of the worm propagation satisfies this limitation. Pet. 44, 60. Specifically, Kaeo teaches the Sapphire worm doubled in size every 8.5

seconds. *Id.* Based on the evidence and arguments presented, we find Petitioner's assertions persuasive.

III. CONCLUSION

Petitioner has shown, by a preponderance of the evidence, that claims 1, 3–7, 17, 18, 20, and 22–30 of the '553 patent are unpatentable.

IV. ORDER

In consideration of the foregoing, it is hereby:

ORDERED that claims 1, 3–7, 17, 18, 20, and 22–30 of the '553 patent have been shown, by a preponderance of the evidence, to be unpatentable;

FURTHER ORDERED that claims 8, 12–14, 16, and 19 of the '553 patent have not been shown, by a preponderance of the evidence, to be unpatentable; and

FURTHER ORDERED that because this is a final written decision of the Board under 35 U.S.C. § 318(a), parties to the proceeding seeking judicial review of the decision must comply with the notice and service requirements of 37 C.F.R. § 90.2.

llw

PETITIONER:

James R. Hannah
Michael Lee
KRAMER LEVIN NAFTALIS & FRANKEL LLP
jhannah@kramerlevin.com
mhlee@kramerlevin.com

PATENT OWNER:

David L. McCombs
Thomas B. King
Gregory P. Huh
HAYNES AND BOONE, LLP
David.mccombs.ipr@haynesboone.com
ipr.thomas.king@haynesboone.com
gregory.huh.ipr@haynesboone.com