

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

Inter Partes Review

Cisco Systems, Inc.,
Petitioner,
v.

Croga Innovations Ltd.,
Patent Owner.

IPR2024-01283
U.S. Patent No. 11,223,601)

**DECLARATION OF DR. ERIC COLE IN SUPPORT OF PATENT
OWNER RESPONSE**

TABLE OF CONTENTS

I.	INTRODUCTION	1
II.	QUALIFICATIONS.....	2
III.	SCOPE OF WORK	9
IV.	LEVEL OF ORDINARY SKILL IN THE ART	9
V.	LEGAL PRINCIPLES.....	12
	A. Anticipation.....	12
	B. Obviousness	14
VI.	OVERVIEW OF THE CHALLENGED PATENT.....	17
	A. Assumed Effective Filing and Invention Dates	17
	B. The Challenged Patent	18
	C. General Discussion of the Challenged Patent and the Prior Art	Error! Bookmark not defined.

I, Eric Cole, of Ashburn, Virginia, declare that:

I. INTRODUCTION

1. Croga Innovations Ltd. has retained me as an independent expert in connection with IPR2024-01283 to provide testimony regarding the technology involved in U.S. Patent No. 11,223,601 (the “’601 Patent,” or the “Challenged Patent”) and how the inventions described and claimed therein compare with the asserted prior art. I have been asked to provide my technical review, analysis, insights, and opinions regarding the references set forth in the above-referenced IPR by Petitioner Cisco Systems.

2. I am being compensated for my work in this matter at my standard hourly rate for consulting services at \$700 per hour. My compensation is not dependent on the outcome of the case and does not affect the substance of my statements in this Declaration. I have no financial interest in the matter or the patents at issue.

3. I have reviewed the specification and the claims of the ’601 patent, as well as the prosecution history for the patent. I am familiar with the technology involved therein, having practical, hands-on experience in the field for over 30 years.

II. QUALIFICATIONS

4. My academic credentials include a master's degree in computer science and a doctorate in information security, and I have worked in the cyber and technical information security industry for over twenty-five years. The details of my education and experience are summarized in my curriculum vitae ("CV"), which is attached as Exhibit 2004. My CV also includes a list of my publications over the past ten years.

5. I have worked in computer security for over twenty-five years. To briefly summarize my career, I have: 1) been responsible for the security of sensitive government projects; 2) managed security at major private sector companies; 3) consulted for major private sector companies; 4) been actively involved in designing, building, and deploying intrusion detection systems; 5) worked on protecting and securing critical information including trade secrets; 6) continually taught and written about cybersecurity.

6. I am the founder and president of Secure Anchor Consulting, LLC, located in Ashburn, Virginia. Secure Anchor Consulting provides cybersecurity consulting services and leads research and development initiatives to advance information systems security. Many of the clients I work with are responsible for

protecting sensitive client and/or consumer information and implementing effective security.

7. At Secure Anchor Consulting, I also actively work on incidents, and I have a detailed understanding of measures that organizations can take to both prevent and detect breaches. In addition, I have worked with organizations to design, build, and implement security measures that would be classified as both fixes and compensating controls, and I have detailed knowledge of the effort that is required to implement various remediation measures.

8. I worked for the U.S. government for eight years as an employee and have held various contracting jobs with government agencies, which involved working with classified information. I have worked for a wide range of government organizations, including the Federal Bureau of Investigation (FBI), the National Security Agency (NSA), the Central Intelligence Agency (CIA), the Department of Energy, the Department of Defense (DOD), the Treasury Department, the Secret Service, and the Nuclear Regulatory Commission (NRC). I have held various top-secret security clearances with the DOD, the CIA, and the NRC. I also was involved in funded research from the Defense Advanced Research Projects Agency (DARPA), an agency of the DOD responsible for the development of new intrusion

detection technologies for use by the military. During this time, I have reviewed, edited, and created numerous security policies. On several occasions, I have also reviewed the reasonableness of measures intended to protect and secure Confidential Information through various types of documentation to include policies, procedures, standards, and baselines.

9. While serving as a Senior Officer for the CIA as Program Manager / Technical Director for the Internet Program Team with the Office of Technical Services, I implemented the Internet Program Team that designs, develops, tests, and deploys internet security products in three- to six-month intervals. In this role, I received a letter of appreciation from the Director of Central Intelligence and six Exceptional Performance Awards.

10. As a member of the Information Security Assessment Team with the Office of Security at the CIA, I evaluated and performed security assessments of network operating systems to identify potential vulnerabilities and solutions. I also designed a large-scale auditing system with an automated review capability and worked on several virus investigations for the Office of Security.

11. In my role as Chief Information Officer for the American Institutes for Research, I repaired and developed IT infrastructures for various organizations. I

was also responsible for building out the entire data classification program to ensure that critical and proprietary data was properly protected and secured. I was also one of the principal leads on research funding from DARPA. In this role, I was responsible for understanding the current state of the art in intrusion detection and designing and building a new cybersecurity system that could catch attacks that were targeting organizations.

12. As Chief Scientist and Senior Fellow for Lockheed Martin, I performed research and development in information systems security. I also specialized in evaluating and designing secure network design, perimeter defense, vulnerability discovery, penetration testing, and intrusion detection systems. At Lockheed Martin, I served as a technical advisor in high-profile security projects for government clients including the DOD, the FBI Sentinel case management systems, Department of Homeland Security Enterprise Acquisition Gateway for Leading Edge solutions, Jet Propulsion Labs, Hanford Labs, and FBI Information Assurance Technology Infusion programs.

13. As Chief Technical Officer for McAfee, I executed the strategy for technology platforms, partnerships, and external relationships to establish product vision and achieve McAfee's goals and business strategies. In this capacity, I

worked closely with groups tasked with the development of intellectual property. While at McAfee, I worked with many business customers responsible for protecting clients' information, helping them implement effective security solutions.

14. I was a fellow and instructor for twenty years with the SANS Institute, a research and education organization consisting of information security professionals. SANS is the largest source for information security training and security certifications in the world. I am an author of several security courses, such as SEC401-Security Essentials and SEC501-Enterprise Defender. SEC401-Security Essentials is a required course for agents of the FBI and the NSA. These courses have been taken by over 30,000 people, and I have taught more than 50% of these students via live and online training. The courses that I have written have multiple modules and lessons on security policies, employee agreements, and protection of intellectual property. As one of the top-rated instructors at SANS, I have delivered keynote addresses around the world.

15. I am a member of the European InfoSec Hall of Fame, a professional membership awarded by nomination and election by a panel of industry experts.

16. I hold a Certified Information System Security Professional (CISSP) certification and was the creator of the Global Information Assurance Certification (GIAC) Security Essentials Certification (GSEC).

17. I am a contributing author of “Securing Cyberspace for the 44th President” and served as a Commissioner on cybersecurity for President Obama. My Nine books on cybersecurity include “Online Danger,” “Network Security Bible - 2nd Edition,” “Advanced Persistent Threat,” and “Insider Threat,” which have become recognized as industry-standard sources. They are used in many classrooms and colleges around the world as authoritative texts on cybersecurity. I have also published several articles in the field.

18. The focus of my work over the last ten years has been on data protection for commercial organizations. I am very familiar with best practices in cybersecurity and have developed many of the standard practices.

19. Today, my consulting efforts help organizations properly protect client information and relate to various alternative methods of security that can be implemented. This includes hardening of systems, securing servers, proper configuration, patching, and compensating controls. In cases where a fix cannot be made because it would negatively impact the business, I have extensive experience

in understanding various compensating controls as an alternative to minimize the risk to an acceptable level and how to implement those controls.

20. I routinely help clients develop and implement company-wide policies, training, and best practices to prevent potential cybersecurity incidents. My work addresses threats from both internal and external sources, including accidental and intentional incidents.

21. I have led, and been involved in working on, responses to many cybersecurity incidents. This includes performing root cause analysis, verifying the reasonableness of security measures, and validating the implementation of additional remediation measures.

22. In my consulting work and in the various positions that I have held, I have been responsible for protecting organizations to minimize and prevent data breaches, in addition to responding to breaches post detection. I have a very detailed understanding of what is and is not a reasonable level of security based on the hundreds of organizations with which I have worked and the various security task forces with which I have been involved.

III. SCOPE OF WORK

23. I have been asked to consider and provide my opinions on the claims of the Challenged Patent, including the technical subject matter and the application of various references that predate the priority date of the Challenged Patent. In particular, I have been asked to consider what one of ordinary skill in the art would have understood from the claims and specification of the Challenged Patent, and whether certain references disclose or suggest the features recited in the claims of the Challenged Patent.

24. I am compensated at a rate of \$700/hour for my work. My compensation is in no way contingent on the nature of my findings, the presentation of my findings in testimony, or the outcome of any proceeding.

25. My opinions set forth below are based on my education, training, experience, and the content of the references considered.

IV. LEVEL OF ORDINARY SKILL IN THE ART

26. My opinions are provided based on what a person of ordinary skill in the art in the technical field of the inventions of the Challenged Patent would have understood at the time of filing of the Challenged Patent.

27. I have been informed and understand that the content of a patent and prior art should be interpreted the way a person of ordinary skill in the art would have interpreted those references at the time of filing of the patent. I understand that the person of ordinary skill is a hypothetical concept referring to one who understands the technology and associated conventional wisdom at the time.

I was asked to provide an opinion as to the level of one of ordinary skill in the art pertinent to the subject matter set forth in the Challenged Patent at the time of the priority date. Around September 28, 2017 (the filing date of provisional patent application 62/564,371), it is my opinion that a person of ordinary skill in the art in the field of the Challenged Patent would have been a person with a Bachelor's degree in computer science, computer engineering, or electrical engineering and at least two years of experience in computer networking, virtual machines and network system security. Throughout this declaration, I use the term "POSITA" to refer to a person with these or similar qualifications. To have similar qualifications, an individual with additional education or additional industrial experience could still be of ordinary skill in the art if that additional education or experience compensates for a deficit in one of the other aspects of the requirements stated above. The Petition and Dr. Shenoy assert that a POSITA

would have “been someone knowledgeable and familiar with network security and application isolation techniques as those are applied in the context of network security available in the 2010s. Such a POSITA would have a bachelor’s degree in computer science, computer engineering, electrical engineering, or equivalent training, and approximately two years of experience working in the field of network security. Additional work experience can substitute for specific educational background, and vice versa.” Ex. 1003 at ¶ 21. While I disagree with the experience they assert is required, I note that, applying their asserted level of skill, my opinions contained herein do not change.

28. The basis for my familiarity with the level of ordinary skill is my interaction with large numbers of workers in the computer security industry who were at this level of skill. In reaching this opinion as to the hypothetical person of ordinary skill in the art, I have considered the types of problems encountered in the art, the prior art solutions to those problems, the rapidity with which innovations are made, the sophistication of the technology, and the educational level and professional capabilities of workers in the field. I further note that, as of September 28, 2017, I was at least at this level of skill.

V. LEGAL PRINCIPLES

A. Anticipation

29. I understand and have been informed that in order for an inventor to be entitled to a patent, the invention must be “new,” and the inventor must not have lost his or her right to a patent by delaying filing of a patent application directed to the invention. There are a number of ways that an invention can be “anticipated”—that is, not new.

30. One such way is when the invention was already patented or described in a printed publication anywhere in the world before the date of invention by the inventor or more than one year before the earliest effective filing date of the patent application. Additionally, a patented invention may not be new if it was on sale or publicly used more than one year before the filing of an application for patent. A patented invention is also not new if it was “known” or used by others in the United States before it was invented by the patent applicants. To invalidate a patent claim based on knowledge, I understand that the knowledge must be accessible to the public. If the subject matter covered by a patent claim is found within the “four corners” of a prior art reference, the claim is said to be anticipated by that reference.

31. A patent claim can be said to be anticipated if each and every limitation of the patent claim is found either expressly or inherently in a single item of prior art. While a prior art reference need not use the same words as a patent claim to anticipate the claim, the prior art reference must describe the requirements of the claim with sufficient clarity such that a POSITA would have been able to make and use the claimed invention based on the reference and his or her knowledge in the applicable technical field.

32. The disclosure of a feature of a claimed invention can be either “express,” meaning that the text, figures, or other content of a reference actually teaches the aspects of the patent claim, or the disclosure can be “inherent.” An express disclosure does not need to use the same words as the patent claim being considered and does not need to be depicted in exactly the same way, but the disclosure must actually convey the claimed invention to a POSITA using some combination of words, figures, data, or other portions of the patent. I have been informed and understand that in order to establish that an element of a claim is “inherent” in the disclosure in a prior art reference, it must be clear to a person skilled in the art that the missing element is an inevitable part of what is explicitly described in the reference, and that it would have been recognized as necessarily present by a

person skilled in the art. Inherency cannot be established by probabilities or possibilities, and the mere fact that something may result from a given set of circumstances is not enough to establish inherency.

33. Finally, I further understand that to anticipate a claim, the elements of the claim must be found in the prior art reference “arranged as in the claim.” I have been informed that this means that the elements of the prior art reference may not be recombined in a way that is not taught in the prior art to arrive at the invention; instead, the prior art must teach the claimed subject matter in the arrangement required by the claims.

34. I understand that any disclosure that anticipates a claim also renders that claim obvious.

B. Obviousness

35. I have been informed and understand that subject matter claimed in a patent is obvious under 35 U.S.C. § 103 if a person of ordinary skill in the art at the time the alleged invention was made would have had reason to combine or modify the disclosures of one or more prior art references to arrive at the claimed subject matter. A claim may be unpatentable even if each and every claim limitation is not present or disclosed in a single prior art item.

36. I have been informed and understand that, under the doctrine of obviousness, a claim is unpatentable if the differences between the invention 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 the subject matter pertains. A person of ordinary skill in the art is presumed to have knowledge of the relevant prior art at the time of the claimed invention.

37. I have been informed and understand that obviousness is based on the scope and content of the prior art, the differences between the prior art and the claim, the level of ordinary skill in the art, and secondary indicia of obviousness and non-obviousness to the extent such indicia exist.

38. I have been informed and understand the determination of whether the asserted claims would have been obvious to a person of ordinary skill in the art and, therefore, invalid is not governed by any rigid test or formula. A determination that a claim is obvious is, instead, based on a common-sense determination that the claimed invention is merely a combination of known limitations to achieve predictable results. Any of the following rationales are acceptable justifications to conclude that a claim would have been obvious: (1) the claimed invention is simply a combination of known prior art methods to yield predictable results; (2) the

claimed invention is a simple substitution of one known element for another to obtain predictable results; (3) the claimed invention uses known techniques to improve similar devices (methods or products) in the same way; (4) the claimed invention applies a known technique to a known device (method or product) ready for improvement to yield predictable results; (5) the claimed invention was “obvious to try”—choosing from a finite number of identified, predictable solutions with a reasonable expectation of success; (6) there is known work in one field of endeavor that may prompt variations of it for use in either the same field or a different one based on design incentives or other market forces if the variations would have been predictable to one of ordinary skill in the art; or (7) there is some teaching, suggestion, or motivation in the prior art that would have led one of ordinary skill in the art to modify the prior art reference to combine prior art teachings to arrive at the claimed inventions.

39. I have been informed and understand that a claim may be obvious in light of a single reference, without the need to combine references, if the elements of the claim that are not found in the reference can be supplied by the common sense or knowledge of one of ordinary skill in the art or taught in different areas of the single reference.

40. I have been informed and understand that an analysis of whether a claimed invention is obvious must not rely on a hindsight combination of prior art. The analysis must proceed in the context of the time of the invention or claimed priority date and consider whether the invention as a whole would have been obvious to a person of ordinary skill in the art, taking into consideration any interrelated teachings of the prior art, the effects of demands known to the design community or present in the marketplace, and the background knowledge possessed by a person having ordinary skill in the art, all in order to determine whether there was an apparent reason to combine any known elements in the fashion claimed by the patent at issue.

VI. OVERVIEW OF THE CHALLENGED PATENT

A. Assumed Effective Filing and Invention Dates

41. The '601 Patent is entitled "Network Isolation for Collaboration Software" and was filed on September 25, 2018. *See* Ex. 1001. The '601 Patent claims priority to patent application 62/564,371, which was filed on September 28, 2017. Consistent with Dr. Shenoy, *see* Ex. 1003, ¶ 38, I assume the priority date for the Challenged Patent is the filing date of the provisional patent application, September 28, 2017.

B. The '601 Patent

42. As noted in the Background section of the Challenged Patent, networked computer systems are subject to security issues presented by malware. *See* Ex. 1001, 1:13-21. Malware may be delivered to a computer system through malicious code that may be unintentionally downloaded when accessing a website, which may result in security losses, such as loss of data, installation of additional malware, and loss of command and control of the infected host computer system. *See* Ex. 1001, 1:22-34.

43. The Challenged Patent further describes the security issues presented by malware and its associated executable toolset. More specifically, the toolset can (1) perform reconnaissance of the infected system to determine what computing resources are accessible and what network resources are accessible, (2) determine system privileges that can be used to access and appropriate files/data stored on the computer or files/data accessible over a network; and/or (3) attempt infection or access of other network resources accessible from the infected computer system. *See* Ex. 1001, 1:45-55.

44. To provide a computing security infrastructure to prevent/limit the possibility of malware toolsets accomplishing any of the above malevolent acts, the

Challenged Patent discloses and claims a combination of (1) an isolated computing environment that uses the host operating system, Ex. 1001, 5:33-43; (2) an internal isolation firewall used to isolate the isolated computing environment, Ex. 1001, 6:63-7:35; and (3) a host-based firewall that isolates the system from communicating with one or more other devices on the network, Ex. 1001, 6:51-62.

45. The “*internal isolation firewall*” is recited in each of the independent claims of the Challenged Patent. Claim 1 recites “a processor configured to . . . implement an isolated computing environment that uses the host operating system . . . and . . . isolate the isolated computing environment from the workspace using an internal isolation firewall.” Ex. 1001, 22:42-63. Similarly, claim 9 recites “implementing an isolated computing environment that uses the host system [and] isolating the isolated computing environment from the workspace using an internal isolation firewall. Ex. 1001, 23:38-24:2.

VII. JEFFRIES DOES NOT DISCLOSE OR RENDER OBVIOUS “A PROCESSOR CONFIGURED TO . . . ISOLATE THE ISOLATED COMPUTING ENVIRONMENT FROM THE WORKSPACE USING AN INTERNAL ISOLATION FIREWALL”

46. For the “*internal isolation firewall*” limitations of claims 1 and 9, the Petition and Dr. Shenoy rely on U.S. Patent No, 10,885,189 to Jeffries et al. (“Jeffries”). See Petition at 63-65, 68; Ex. 1003, ¶¶ 159-164.

47. The Petition and Dr. Shenoy contend that Jeffries' container manager 118 and secure kernel correspond to an "*internal isolation firewall*" because they perform the same functions of the internal isolation firewall described by '601 patent, including preventing communication between the isolated computing environment and the workspace." Petition at 65 (citing Ex. 1001 at 1:67-2:2, 3:38-40, 6:63-66, 13:15-23, 14:7-14); Ex. 1003, ¶ 163. I disagree.

48. Dr. Shenoy contends that container manager 118 in Jeffries "enforces isolation between the container 122 and the remainder of the system." Ex. 1003 at ¶ 160. He relies upon a disclosure in Jeffries that states "container manager 118 is configured to quarantine the container 122 and prevent the container 122 from further accessing resources. By blocking the container 122 from further accessing resources, the container manager 118 prevents malicious activity occurring in the container from gaining file system access." Ex. 1003 at ¶ 160 (citing Jeffries at 16:29-34).

49. Dr. Shenoy also contends that "a secure kernel helps to enforce the seaparation between the host operating system 102 and the container 122. Ex. 1003 at ¶ 161. He relies upon a disclosure in Jeffries that states "The communicative connection between the host operating system 102 and the container 122 implements

a secure kernel to maintain integrity of information pass from the container 122 to the host operating system 102 . . . Accordingly, when malicious software attempts to access off-limit data or memory locations, the secure kernel causes an exception or interrupt to occur.” Ex. 1003 at ¶ 161 (citing Ex. Jeffries at 12:7-19).

50. I do not agree that the container manager 118 and/or secure kernel in Jeffries disclose or render obvious “*an internal isolation firewall*” as claimed in the ’601 patent. In addition, I do not agree that Dr. Shenoy provided sufficient evidence to conclude that container manager 118 and secure kernel, as disclosed in Jeffries, “*isolate the isolated computing environment*” as claimed in the ’601 Patent or that it would have been obvious to do so.

51. The secure kernel in Jeffries is aimed at analyzing activity to determine a threat level of malicious activity. Jeffries says “[i]n response to detecting resource access by the application, a security event is generated describing malicious activity that occurs from the accessing the resource. This security event is analyzed to determine a threat level of the malicious activity. If the threat level does not satisfy a threat level threshold, the host operating system allows the application to continue accessing resources and continues to monitor resource access. When the threat level satisfies the threat level threshold, the operating system takes corrective action to

prevent the malicious activity from spreading beyond the isolated container.” Ex. 1005 at 1 (Abstract). In Jeffries, “a determination of whether potentially malicious activity indicated by the security event satisfies a threat level threshold is used by the container manager **118** to either allow an application **124(1)**, . . . , **124(n)** to continue accessing resources or take corrective action to mitigate the malicious activity.” Ex. 1005 at 11:18-23.

52. As to the types of behaviors that could be monitored, Jeffries states “if a user starts visiting numerous websites that the user has not previously visited, this behavior is logged as one or more security events and compared against a threat level threshold to determine whether corrective action is needed. Similarly, if an application begins downloading and installing software that the application has not previously downloaded and installed, this activity is logged in one or more security events and compared against a threat level threshold to determine whether corrective action is needed.” Ex. 1005 at 5:25-33.

53. Jeffries also states “in response to determining that a threat level of malicious activity described by a received security event satisfies a threat level threshold, the host operating system takes corrective action to mitigate a negative effect of the malicious activity. For instance, the host operating system may

quarantine the container to prevent the container from accessing any resources or data outside the container. Alternatively or additionally, the host operating system may identify a resource causing the malicious activity and add the identified resource to a list of untrusted resources so that other computing devices are aware of possible negative effects.” Ex. 1005 at 5:47-58.

54. None of the above described behaviors teach or suggest that Jeffries’ container manager 118 or secure kernel include or use an “*internal isolation firewall*.” First, the secure kernel in Jeffries merely generates a security event that describes that activity and, in some instances, creates an or interrupt. As a POSITA would readily understand, an exception or interrupt are notifications to the processor that an event has occurred. They do not block access to resources. They merely inform the processor of an event. Second, the “corrective action” taken by container manager 118 in Jeffries when “an activity satisfies a threat level threshold” is not isolation using a firewall. In particular, “quarantin[ing] the container to prevent the container from accessing any resources or data outside the container” is not equivalent to isolating an isolated computing environment using an internal isolation firewall, because a firewall functions to allow certain actions and disallow others. The wholesale restriction in Jeffries, which prevents the container from accessing

any resources or data outside the container, is an entirely different process than what the '601 patent describes and claims. As a POSITA would understand, a firewall (and, specifically, the claimed "*internal isolation firewall*" in the '601 Patent) isolates the ICE by filtering specific actions, operations, or data so that some of those actions, operations, or data are allowed to pass through the firewall while others are not.

55. In addition, "add[ing] the identified resource to a list of resources so that other computing devices are aware of possible negative effects" also does not "isolate the isolated computing environment from the workspace using an internal isolation firewall." Creating a list of untrusted resources does not allow or block any actions from occurring. It merely makes other computing devices aware of those resources.

56. Jeffries also includes "terminating the container 122 and rebuilding a new container" as a remedy for malicious behavior. Jeffries at 27:34-35. This is also not isolating an isolated computing environment using an internal isolation firewall. I note that Dr. Shenoy appears to recognize the difference between a firewall's filtering function and the "correction actions" taken in Jeffries. In his discussion of claim 5, Dr. Shenoy states "using an isolated container or virtual machine to protect

against received malware typically requires terminating the container or machine, along with its contents, entirely to prevent the spread of malware. *See* Jeffries, 16:41-42 (“terminating the container 122 and rebuilding a new container”), 27:34-35. Doing so results in the loss of unsaved or unfinished work and is a lengthy process leading to frustration of the user. A POSITA would have recognized that advantage of **additionally filtering** packets received by the containers/virtual machines to minimize the instances in which an isolated container or virtual machine must be terminated and rebuilt.” Ex. 1003 at ¶ 248 (emphasis added).

57. The use of the terms “firewall” and “internal isolation firewall” in the ’601 Patent is consistent with the plain meaning of those terms as understood by a POSITA in view of the ’601 Patent specification and file history. First, the specification describes that “internal isolation firewall may . . . **allow**[] access to the host system to perform one or more of a cut or copy of data, a paste of data, a printing of data to a local printer, a file transfer, and/or the like” and it “may **prevent** one or more other transfers of data between the sandboxed computing environment and the workspace.” Ex. 1001 at 7:20-35 (emphasis added). The ’601 Patent also says “[t]he isolated computing environment may be isolated from the workspace by an internal isolation firewall [that] . . . may **prevent** communication between the isolated

computing environment and the workspace, . . . [or] **allow** communication between the isolated computing environment and the workspace.” Ex. 1001 at 1:67-2:6 (emphasis added).

58. Second, “the internal isolation firewall 218 may allow a predefined set of processes to be executed within the sandboxed computing environment 200 . . . [and] may prevent execution . . . of any process outside of the predefined set of processes.” Ex. 1001 at 13:17-24. Thus, the internal isolation firewall is able to allow certain predefined processes to execute while preventing others from executing within the sandboxed computing environment.

59. Third, “[t]he sandbox container process may restrict access to certain resources by processes running in the sandboxed computing environment. The sandboxed container process may allow the workspace and the sandboxed computing environment to share access to the host operating system and host operating resources (e.g., including the kernel and libraries) while enforcing the segregation of the memory spaces access by the workspace and the sandboxed computing environment.” Ex. 1001 at 5:62-6:3.

60. Fourth, “[t]he internal isolation firewall may allow certain actions to be performed with approval of the user of the host system. For example, the internal

isolation firewall may prompt the user to allow communication between the workspace and the isolated computing environment. For example, an application and/or process operating within the sandboxed computing environment may be allowed access to the host system to perform one or more of a cut or copy of data, a paste of data, a printing of data to a local printer, a file transfer, and/or the like.” Ex. 1001 at 7:20-29.

61. Fifth, “[t]he internal isolation firewall 218 may be configured to prevent establishment of a communication channel for malware to interact with an untrusted device (e.g., such as the host computer system 234).” Ex. 1001 at 14:14-18.

62. Sixth, “[t]he sandbox container process (e.g., the internal isolation firewall 318) may prevent the data on the host computer system 312 and/or the network 330 from being reported back to the host computer system 310, another device that may have sent the malware to the host computer system 312, and/or any other remote devices.” Ex. 1001 at 17:21-26.

63. Seventh, “the respective internal isolation firewalls may prevent the host computer system 410 from accessing any of the sandboxed computing environments operating on the host computer system 412.” Ex. 1001 at 21:17-21.

Thus, the '601 Patent specification describes an internal isolation firewall serving to prevent access to a sandboxed computing environment by a host computer system.

64. Thus, the '601 Patent specification describes an internal isolation firewall that allows certain specific actions (e.g., access to specific host resources, transfers of certain data between the sandboxed computing environment and the workspace, execution of predetermined processes within the sandboxed computing environment, etc.) while disallowing others based on predetermined rules or policies. Those actions can involve data transmission from the isolated computing environment to a workspace, or data transmission to the isolated computing environment from a workspace.

65. This understanding of isolation using a firewall as process of filtering—i.e., selectively allowing or blocking—specific actions is also supported by the intrinsic record. For example, U.S. Patent Application Publication 2014/0344912 to Chapman et al. (listed under References Cited in Ex. 1001 at 2) discloses that a firewall's "primary objective is to monitor and control the incoming and outgoing network traffic by analyzing the data packets and determining whether each packet should be allowed through or not, based on a predetermined rule set" Ex. 2006 at [0004]. Similarly, U.S. Patent Application Publication 2009/0249472 to Litvin et al.

(listed under References Cited in Ex. 1001 at 2) states “[a] firewall . . . checks incoming and/or outgoing packets of data against an existing list of characteristics that determine whether the packets should be allowed to continue to their destination or should be blocked.” Ex. 2007 at [0009]. And, U.S. Patent No. 7,150,040 to Brownell (listed under References Cited in Ex. 1001 at 1) states “[f]irewalls are used to screen traffic between . . . networks.” Ex. 2008 at 1:29-37.

66. I also note that Dr. Shenoy opines as to the “host-based firewall” limitation of claim 5, that a “*host-based firewall*” that “*isolates the system from communicating with one or more other devices on the network*” is disclosed by a component in Cucinotta (Ex. 1008) that “isolates the host computer system by intercepting all incoming and outgoing packets” and “selectively allow[ing] packets based on security policies.” Ex. 1003 at ¶¶ 255-259; Pet. 98. Dr. Shenoy also appears to equate a “firewall” with a “filter,” when he states “[i]n the combination, a host-based filter, such as the host-based firewall component described in Cucinotta, would be implementd with Jeffries’ host operating system 102.” Ex. 1003 at ¶ 258.

67. Applying this understanding of “isolation,” “firewall,” and “internal isolation firewall,” a POSITA would readily understand that Jeffries’ container manager 118 and secure kernel do not “*isolate the isolated computing environment*

from the workspace using an internal isolation firewall” as claimed in the ’601 Patent.

68. Because every challenged claim requires “*a processor configured to . . . isolate the isolated computing environment from the workspace using an internal isolation firewall*” (recited in claim 1 and its dependents) or “*isolating the isolated computing environment from the workspace using an internal isolation firewall*” (recited in claim 9 and its dependents), I disagree that Petitioner or Dr. Shenoy have shown that any of the challenged claims are rendered obvious.

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

Date: May 22, 2025



Dr. Eric Cole