

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

CISCO SYSTEMS, INC.
Petitioner,

v.

UMBRA TECHNOLOGIES, LTD.
Patent Owner.

U.S. Patent No. 10,574,482
Issue Date: February 25, 2020
Title: MULTI-PERIMETER FIREWALL IN THE CLOUD

Inter Partes Review No.: IPR2024-00498

**DECLARATION OF DR. A.L. NARASIMHA REDDY IN SUPPORT OF
PETITION FOR *INTER PARTES* REVIEW OF
U.S. PATENT NO. 10,574,482**

Mail Stop “PATENT BOARD”

Patent Trial and Appeal Board
U.S. Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450

TABLE OF CONTENTS

	Page
I. INTRODUCTION	3
II. MY BACKGROUND AND QUALIFICATIONS	5
III. RELEVANT LEGAL STANDARDS.....	9
IV. BACKGROUND	12
A. The '482 Patent	12
B. Claim Construction.....	16
C. Priority Date of Claims	16
D. Person of Ordinary Skill in the Art (POSITA).....	17
V. STATE OF THE ART	18
1. firewalls	18
2. distributed Firewalls in the Cloud	22
VI. OVERVIEW	22
VII. CHALLENGES.....	23
A. Identification of Challenges.	23
B. Challenge 1: Claims 1-26.....	23
1. The Sheih Reference	23
2. The Potti Reference	31
3. Detailed Application of Shieh and Potti to the Challenged Claims	32
VIII. SUMMARY OF THE '482 PATENT	63
APPENDIX A – MATERIALS RELIED UPON.....	3

APPENDIX B – <i>CURRICULUM VITAE</i>	4
--	---

I, Dr. A.L. Narasimha Reddy, declare as follows:

I. INTRODUCTION

1. I am over the age of eighteen (18) and otherwise competent to make this declaration.

2. I have been retained by Cisco Systems, Inc. (Petitioner) as an independent expert consultant in this proceeding before the United States Patent and Trademark Office. Although I am being compensated at my rate of \$650 per hour for the time I spend on this matter, no part of my compensation depends on the outcome of this proceeding, and I have no other interest in this proceeding. To the best of my knowledge, I have no financial interest in Client.

3. This Petition for *inter partes* review involves U.S. Patent No. 10,574,482 (“the ’482 Patent”) (EX-1001), entitled “MULTI-PERIMETER FIREWALL IN THE CLOUD” and lists Carlos Eduardo Oré and Joseph E. Rubenstein, as the inventors.

4. For the purposes of this *inter partes* review as discussed later, I have been instructed to assume that the effective filing date of the Claims of the ’482 Patent challenged by the Petitioner in this *inter partes* review is no earlier than April 7, 2015, the filing date of U.S. Provisional Patent Application No. 62/144,293.

5. I understand that according to USPTO records, the '482 Patent is currently assigned to UMBRA Technologies Ltd. ("UMBRA" or "Patent Owner").

6. The '482 Patent is directed to a communication network having a distributed firewall system in the cloud. The distributed firewall system provides a layered defense by pairing a first access point server with a first perimeter firewall that performs stateful packet inspection, and then a second access point server with a second perimeter firewall that does deep packet inspection. Packets traversing the network are subject to stateful packet inspection at the first perimeter firewall, which prevents some traffic from passing to the second access point server. The second perimeter firewall then performs deep packet inspection on the traffic that passed from the first access point server to the second access point server. The deep packet inspection is performed on a cloned copy of the traffic in order to avoid interrupting the flow of traffic through the network. I am familiar with the technology described in the '482 Patent as of the earliest possible priority date of April 7, 2015.

7. In preparing this Declaration, I have reviewed the '482 Patent (EX-1001), the file history of the '482 Patent (EX-1002), and each of the documents cited herein, and I have considered these documents in light of the general knowledge in the art as of April 7, 2015. In formulating my opinions, I have relied upon my experience in the relevant art. I have also considered the viewpoint of a

person of ordinary skill in the art (“POSITA”) in the relevant field, as of April 7, 2015, who I describe below (§IV.D). For convenience, the materials I considered in arriving to my opinion are listed in Appendix A.

8. The opinions and comments formulated during this assessment are based on observations and information available at the time of the investigation. The findings presented herein are made to a reasonable degree of scientific certainty. I have made every effort to accurately and completely investigate all areas of concern identified during our investigation.

9. I have been asked to provide my technical expertise, analysis, insights and opinions regarding the ’482 Patent and relevant references that form the basis of the arguments set forth in the accompanying Petition for *inter partes* review of the ’482 Patent. As described in detail below, I offer the following opinion in this Declaration:

- A POSITA would have found Claims 1-26 of the ’482 Patent are unpatentable as obvious over Shieh and Potti.

II. MY BACKGROUND AND QUALIFICATIONS

10. I am currently the Truchard Foundation Chair Professor of Electrical and Computer Engineering at Texas A&M University in College Station, Texas. I have over 25 years of experience in a wide variety of technologies and industries

relating to data communications, storage systems, and distributed systems, including packet-switched network communications.

11. My academic credentials include a Bachelor of Technology degree in Electronics and Electrical Communications Engineering from the Indian Institute of Technology, Kharagpur, India, in August 1985. I then received a Master of Science and a Ph.D. in Computer Engineering from the University of Illinois at Urbana-Champaign in May 1987 and August 1990, respectively.

12. I have worked for over 25 years in the field of Electrical Engineering. My primary focus and research interest has been on computer networks, storage systems, multimedia systems, and computer architecture. I have authored and co-authored over a hundred technical papers and several book chapters related to several of these interests, including on such topics as multipath routing, route control, high-speed networks, network congestion, packet management, quality of service regulation, network security, network modeling, differentiated services, storage system enhancements, caching strategies, and multimedia system enhancements to name a few examples. I am listed as an inventor on six U.S. Patents in the field of multi-node communication networks and computer systems.

13. My employment history following my graduation from the University of Illinois at Urbana-Champaign began at the IBM Almaden Research Center in

San Jose, California in 1990. At IBM, I worked on projects related to disk arrays, multiprocessor communication, hierarchical storage systems and video servers.

14. In 1995, I joined the faculty of the department of Electrical Engineering at Texas A&M University initially as an Associate Professor and was later promoted to full, tenured Professor. At Texas A&M, during 2016-21, I was the Associate Agency Director for Strategic Initiatives and Centers for the Texas A&M Engineering Experiment Station (TEES), which engages in engineering and technology-oriented research and educational collaborations. Further, I also served as Associate Dean for Research in the College of Engineering during 2016-21.

15. At Texas A&M, I have taught dozens of courses related to computer networking and communications, as well as computer architecture, multimedia systems and networks, topics in networking security, multimedia storage and delivery, as well as networking for multimedia applications. I have also served on various committees for the benefit of the scientific community and the Texas A&M University community.

16. I am a member of a number of professional societies, including the Institute of Electrical and Electronic Engineers (IEEE), where I have been elevated to an IEEE Fellow, and the Association for Computing Machinery (ACM). I have been responsible for chairing or co-chairing numerous conferences and programs, as well as presenting research at major IEEE and ACM conferences. For example,

I served as program co-chair for the 2008 5th International Conference on Broadband Communications, Networks and Systems, panels co-chair for the 2008 3rd International Conference on Communication Systems Software & Middleware, and panel chair of the IEEE Conference of High Performance Computer Architecture.

17. My presentations include a Keynote speech at International Conference on Information Technology-New Generations in 2013, a Keynote speech at IEEE International Symposium on Computers and Communications 2010, and several invited talks including at Georgia Tech (2013), COMSNETS Conference (2013), Int. Conf. on Networking and Communications (2012), Samsung (2011), Korea University (2011), Aijou University (2011), Catedra Series at University of Carlos III, Madrid (2009), Thomson Research, Paris (2009), Telefonica Research, Barcelona (2009), and a Distinguished Seminar at IBM Austin Research Lab (2008).

18. I have received multiple awards in the field of networks and computer architecture. I received the NSF Career Award from 1996-2000. I received an outstanding professor award by the IEEE student branch at Texas A&M during 1997-1998, an outstanding faculty award by the department of Electrical and Computer Engineering during 2003-2004, a Distinguished Achievement award for teaching from the former students' association of Texas A&M University, and a

citation “for one of the most influential papers from the 1st ACM Multimedia conference.

19. My curriculum vitae, which includes a more detailed summary of my background and experience, is attached as Appendix B.

III. RELEVANT LEGAL STANDARDS

20. I am not a lawyer and do not provide any legal opinions in this Declaration. Although I am not a lawyer, I have been informed and understand that certain legal standards are to be applied by technical experts in forming opinions regarding the meaning and validity of patent claims. I have been asked to provide my opinions regarding whether the claims of the '482 Patent are anticipated or would have been obvious to a person having ordinary skill in the art at the time of the alleged invention, in light of the prior art.

21. I have been informed and understand that a patent claim is not patentable under 35 U.S.C. § 103 if the differences between the patent claim and the prior art are such that the claimed subject matter as a whole would have been obvious at the time the claimed invention was made to a person having ordinary skill in the relevant art. Obviousness, as I have been informed and understand, 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, to the extent that they exist, certain objective indicia of non-obviousness.

22. I understand that objective indicia can be important evidence regarding whether a patent is obvious or nonobvious, if it has an appropriate nexus to the claimed invention, i.e., is a result of the merits of a claimed invention (rather than the result of design needs or market-pressure advertising or similar activities). Such indicia include: commercial success of products covered by the patent claims; a long-felt need for the invention; failed attempts by others to make the invention; copying of the invention by others in the field; unexpected results achieved by the invention as compared to the closest prior art; praise of the invention by the infringer or others in the field; the taking of licenses under the patent by others; expressions of surprise by experts and those skilled in the art at the making of the invention; and the patentee proceeded contrary to the accepted wisdom of the prior art.

23. I have been informed that whether there are any relevant differences between the prior art and the claimed invention is to be analyzed from the view of a person of ordinary skill in the relevant art at the time of the invention. As such, my opinions below as to a person of ordinary skill in the art are as of the time of the invention, even if not expressly stated as such; for example, even if stated in the present tense.

24. In analyzing the relevance of the differences between the claimed invention and the prior art, I have been informed that I must consider the impact, if

any, of such differences on the obviousness or non-obviousness of the invention as a whole, not merely some portion of it. The person of ordinary skill faced with a problem is able to apply his or her experience and ability to solve the problem and also look to any available prior art to help solve the problem.

25. I have been informed that a precise teaching in the prior art directed to the subject matter of the claimed invention is not needed. I have been informed that one may take into account the inferences and creative steps that a person of ordinary skill in the art would have employed in reviewing the prior art at the time of the invention. For example, if the claimed invention combined elements known in the prior art and the combination yielded results that were predictable to a person of ordinary skill in the art at the time of the invention, then this evidence would make it more likely that the claim was obvious. On the other hand, if the combination of known elements yielded unexpected or unpredictable results, or if the prior art teaches away from combining the known elements, then this evidence would make it more likely that the claim that successfully combined those elements was not obvious.

26. I have been informed and understand that there are recognized, exemplary, rationales for combining or modifying references to show obviousness of claimed subject matter. Some of the rationales include the following:
combining prior art elements according to known methods to yield predictable

results; simple substitution of one known element for another to yield predictable results; use of a known technique to improve a similar device (method or product) in the same way; applying a known technique to a known device (method or product) ready for improvement to yield predictable results; choosing from a finite number of identified, predictable solutions, with a reasonable expectation of success; known work in one field of endeavor 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 are predictable to one of ordinary skill in the art; and some teaching, suggestion, or motivation in the prior art that would have led one of ordinary skill to modify the prior art reference or to combine prior art teachings to arrive at the claimed invention.

IV. BACKGROUND

A. The '482 Patent


27. Claims 1-26 of the '482 Patent (the “Challenged Claims”) are unpatentable as obvious over the prior art. The claims are directed to a communication network having distributed firewall system in the cloud that was well known in the art before the filing of the '482 Patent.

28. The distributed firewall system provides a layered defense by pairing a first access point server with a first perimeter firewall that performs stateful packet inspection, and then a second access point server with a second perimeter

firewall that does deep packet inspection. Packets traversing the network are subject to stateful packet inspection at the first perimeter firewall, which prevents some traffic from passing to the second access point server. The second perimeter firewall then performs deep packet inspection on the traffic that passed from the first access point server to the second access point server. The deep packet inspection is performed on a cloned copy of the traffic in order to avoid interrupting the flow of traffic through the network. This claimed architecture functionality already existed in the prior art.

29. For example, the prior art Shieh reference is directed to a communication network having a distributed firewall system in the cloud. Shieh includes a parallel disclosure to that of the '482 Patent in that it is directed to the same problem and provided the same solution as the '482 Patent, the only difference being that Shieh uses different words to describe the identical concepts described and claimed in the '482 Patent.


30. As illustrated below, the architecture shown in FIG. 12 of the '482 Patent is substantially identical to the architecture shown in FIG. 3 of Shieh with common elements being colored with the same color.



EX-1001 FIG. 12

FIG. 12

31.



EX-1004 FIG. 3

FIG. 3

32. The first access point server (blue) and first perimeter firewall (light blue) for performing stateful packet inspection in the '482 Patent (FIG. 12) are called the security processing module 309 (blue) and service processing module 312 (light blue), respectively, in Shieh (FIG. 3). The second access point server (red) and second perimeter firewall (pink) for performing deep packet inspection in the '482 Patent (FIG. 12) are called the security processing module 310 (red) and service processing module 313 (pink), respectively, in Shieh (FIG. 3). The egress ingress point device (green) and endpoint device (purple) of the '482 Patent are called virtual input/output module 304 (green) and server 324 (purple), respectively.

33. As described in more detail below, Sheih discloses in its cloud based distributed firewall system, a myriad of firewall functionality, including network address translation, stateful inspection and deep patent inspection. EX-1004 3:5-14.

34. The only feature recited in the Challenged Claims but not expressly disclosed in Sheih was the well-known concept of performing deep packet inspection *on a copy* of the packets traversing the network. Indeed, a POSITA understood that performing deep packet inspection on a copy of the packets traversing the network was a well-known method that avoided interrupting the

flow of traffic through the network while the time and resource intensive deep packet inspection was performed, as demonstrated by the prior art Potti reference.

35. Because there are no material differences between the architectures disclosed in the '482 Patent and the combination of Shieh and Potti, the Challenged Claims are unpatentable as obvious over Shieh and Potti.

B. Claim Construction

36. I have been informed and understand that in order to properly evaluate the '482 Patent, the terms of the claims must first be interpreted.

37. I understand that Petitioner proposes that each claim term in the Challenged Claims be given its plain and ordinary meaning in this proceeding, and that no specific construction of any claim term is required because the prior art relied on in this Petition meets each of the claim terms under any reasonable construction. The plain meaning is the ordinary and customary meaning to those skilled in the art.

C. Priority Date of Claims

38. For the purpose of this *inter partes* review, I have been instructed to assume that the effective filing date of each of the Challenged Claims is no earlier than the filing date of U.S. Provisional Application Nos. 62/144,293. Solely for the purpose of this IPR, I have been instructed to assume the priority date is April 7, 2015.

D. Person of Ordinary Skill in the Art (POSITA)

39. I am familiar with the knowledge and capabilities of one of ordinary skill in the art. Unless otherwise stated, my testimony below refers to the knowledge of one of ordinary skill in the art as of April 7, 2015, the earliest possible effective filing date of the '482 Patent.

40. I have been informed and understand that a Person of Ordinary Skill in the Art ("POSITA") is a hypothetical person who is presumed to be aware of all pertinent prior art, thinks along conventional wisdom in the art, and is a person of ordinary creativity.

41. With respect to the '482 Patent, a POSITA in April 2015 would have been familiar with communication networks having a distributed firewall system in the cloud. A POSITA would have a working knowledge of the function performed by firewalls including stateful inspection and deep packet inspection. A POSITA would have gained knowledge of these concepts through a mixture of training and work experience, such as by having a Bachelor's degree in computer science and two years of experience; or by obtaining a Master's degree in computer science, but having only one year of experience; or by having no formal education but experience in storage systems of at least four years.

42. A POSITA may work as part of a multi-disciplinary team and draw upon not only his or her own skills, but also take advantage of certain specialized skills of others on the team to solve a given problem.

V. STATE OF THE ART

43. The following section describes the state of the art for distributed firewalls as of April 2015. The prior art references, and the discussions of what was known to a POSITA, provide the factual support for the general description of the state of the art at the time of the invention, and provide additional motivation to modify or combine the references. Accordingly, these references should be considered by the Board.

1. Firewalls

44. By April 2015, firewall technology was very well developed and decades old. Firewalls provided a variety of services to networks in terms of security. They provided for network address translation (NAT), they helped establish virtual private networks (VPN), and filtered traffic that did not conform to the network's stated security policy. *EX-1006* p. 1.

45. Firewalls control the flow of traffic to protect a communication network from security threats at defined boundary points, whether at the perimeter between a public and private network, or between different areas of a network. *EX-1007* pp. 42, 91.

46. A POSITA understood that a computer network security design typically uses a layered defense approach of multiple firewalls so that if one layer is compromised the entire network is not exposed to the security threat. *EX-1007* pp. 83-84.

47. Known firewalls included stateful inspection firewalls and deep packet inspection firewalls. By the mid-to-late 1990s, stateful inspection had become an industry standard. Stateful inspection provided for the analysis of packets at the network layer as well other layers (typically the transport layer in the OSI model but the firewall looked at layers above that as well) in order to assess the overall packet. By combining information from various layers (transport, session, and network) the firewall was better able to understand the protocol it was inspecting. This also provided for the ability to create virtual sessions in order to track connectionless protocols such as UDP-based applications as well as RPC-based applications. *EX-1006* p. 2.

48. Stateful firewalls keep track of the state of a connection: whether the connection is in initiation, data transfer, or termination state. *EX-1007* p. 53-54. Stateful firewalls use a state table as a mechanism to keep track of the state of a connection. *EX-1007* p. 59.

49. Deep packet inspection (DPI) has also been around for over twenty years. Deep packet inspection looked within the application payload of a packet or

traffic stream and made decisions on the significance of that data based on the content of that data. The engine that drives deep packet inspection typically included a combination of signature-matching technology along with heuristic analysis of the data in order to determine the impact of that communication stream.

EX-1006 p. 2.

50. Deep packet inspection could be used to detect and prevent cyberattacks, combat malware, optimize servers to reduce overhead, and analyze user behavior. Deep packet inspection could be inline, where a packet was inspected and if it satisfied the configured security policies the packet would be forwarded to its destination. Or the packet inspection could be offline, where a copy of the packet is sent for further analysis, while the original packet is forwarded to its destination. *EX-1008* pp. 617-618.

51. Offline analysis permitted further analysis for traffic patterns and implementing advanced security functions without slowing down the original traffic, including IP logging, shunning (blocking packets) and anomaly detection. *EX-1008* pp. 650-669.

52. For example, deep packet inspection could be used to identify malicious content, and then an IP address as the source of that malicious content, and then the IP address would be sent to other firewalls to block or drop packets from this IP address. *EX-1008* pp. 659-666.

53. The sharing of threat information between firewalls was also well-known and was sometimes referred to as “attack information advertising” and was critical to ensure that firewalls on the perimeter take countermeasures to prevent malicious packets from entering the network. *EX-1009* p. 1915.

54. It was well-known to use multiple layers of firewalls because it defended the network more efficiently and thus provided a safer environment. For example, it was known to layer firewalls for network defense from the lightest to the heaviest solution. A packet filter firewall was placed on the perimeter of the network and the Internet, typically in a router which required little capacity and could handle large amounts of data. Stateful firewalls were then placed behind the packet filter to maintain the session tables and protected against attacks let through by the packet filter. Because some traffic was blocked by the stateless (packet filter) firewall, it lightened the load on the stateful firewall. Deep packet inspection firewalls, used to examine the content for undesired patterns or messages specified by the network administrator, were then placed in front of the servers requiring highest security. *EX-1010* p. 7. A POSITA understood that because some traffic was blocked by the stateful firewall, it lightened the load on the deep packet inspection firewall.

55. It was also known that deep packet inspection could be integrated with a stateful firewall with intrusion detection and prevention systems. The

stateful firewall examined the received packet and the intrusion detection system would inspect the message carried. The DPI provided significant benefit when compared to a traditional stateful firewall system because the message is also inspected for fraud patterns. *EX-1010* p. 8.

2. Distributed Firewalls in the Cloud

56. Cloud computing generally refers to the on-demand availability of computer system resources, especially data storage and computing power. Large clouds often have functions distributed over multiple locations, and rely on sharing of resources to achieve coherence and typically use a pay-as-you-go model. For over a decade, as more computing resources were located in the cloud, firewall security services were also moved to the cloud, and the virtual firewall usually became the best choice because the cloud was meant to offer on-demand service for the user with the pay-as-you-go-model. *EX-1010* pp. 24-27.

VI. OVERVIEW

57. As detailed below, I explain how each limitation recited in the Challenged Claims are found in the prior art. The Challenged Claims are directed to a communication network having distributed firewall system in the cloud. The distributed firewall system provides a layered defense by pairing a first access point server with a first perimeter firewall that performs stateful packet inspection, and then a second access point server with a second perimeter firewall that does

deep packet inspection. Packets traversing the network are subject to stateful packet inspection at the first perimeter firewall, which prevents some traffic from passing to the second access point server. The second perimeter firewall then performs deep packet inspection on the traffic that passed from the first access point server to the second access point server. The deep packet inspection is performed on a cloned copy of the traffic in order to avoid interrupting the flow of traffic through the network.

VII. CHALLENGES

A. Identification of Challenges.

58. Based on my understanding of the law of obviousness discussed above, Claims 1-26 of the '482 Patent are invalid as being obvious over obvious over Shieh in view Potti.

B. Challenge 1: Claims 1-26

1. The Sheih Reference

59. Shieh is directed to providing network security using distributed firewalls in the cloud. FIG. 1 is a block diagram illustrating an example of network configuration according to one embodiment of the invention. *EX-1004* 2:60-62.




FIG. 1

EX-1004 Fig. 1

60. Network access device 204 (green), which may be a router, gateway, switch access point, etc., provides an interface at the perimeter between a public network 203 (red) and private network 205 (blue). As shown in FIG. 1, network 203 may be an external network such as a wide area network or internet while network 205 represents a local area network (LAN). EX-1004 2:62-67; FIG. 1.

61. Nodes 206-207 go through gateway device 204 (labeled “Network Access Device(s) 204” in FIG. 1 above) in order to reach nodes 201-202, or vice versa. Any of nodes 201-202 and 206-207 may be a client device (e.g., a desktop, laptop, smartphone, gaming device) or a server. EX-1004 2:67-3:4.

62. According to one embodiment, network access device 204 is associated with a distributed firewall 212 (purple) that includes various firewall processing modules (pink), for example, each being executed within a virtual machine (VM). In one embodiment, each firewall module is responsible for performing one or more firewall functions, but it does not include all of the firewall functions of a firewall. Examples of the firewall functions include, but are not limited to, network address translation (NAT), virtual private network (VPN), deep packet inspection (DPI), and/or anti-virus, etc. *EX-1004* 3:5-14.

63. In one embodiment, some of the firewall processing modules are located within network access device 204 (e.g., firewall modules 209) and some are located external to network access device 204 (e.g., firewall modules 210 maintained by firewall processing node(s) 211, which may be a dedicated firewall processing machine). All of the firewall modules 209-210 are managed by firewall controller 208 (orange), which may be located within network access device 204, or external to network access device 204. *EX-1004* 3:14-22.

64. Shieh describes that each of the firewall modules may be implemented in a virtual machine, which may be located in the cloud. *EX-1004* 3:3:23-30; 9:12-15. Shieh describes that using such modules in a firewall cloud provides a security design that is best-fit for the emerging cloud computing, and provides great scalability and system availability. *EX-1004* 9:33-36.

65. Shieh describes a more detailed embodiment of the distributed firewall system in FIG. 3:




FIG. 3

EX-1004 Fig. 3

66. Distributed firewall 212 includes, four different types of modules: virtual I/O modules 301-304 (green), security processing modules 309-311 (blue), service processing modules 312-313 (purple), and central controller 208 (orange). All modules can run on the same virtual machine, or on different virtual machines, or on same or different physical hosts. Some or all of modules 301-304 and 309-313 may be executed within a virtual machine. Dependent upon the specific configuration, each of modules 301-304 and 309-313 may be executed by a

respective virtual machine. In other configurations, multiple of modules 301-304 and 309-313 may be executed by the same virtual machine. *EX-1004* 7:11-30.

67. Shieh describes that a virtual machine is software hosted on a server for performing the functions for the different types of modules. *EX-1004* 3:41-43; 11:33-36.

68. Shieh describes that the modules could be distributed in different networks, even on different locations, as long as the modules can reach the module that is next in terms of processing and the central controller. In one embodiment, virtual I/O modules and corresponding security processing modules are in a public cloud and the central controller is in a private cloud. This configuration may provide the flexibility to secure and control packets coming from the public cloud, and allow central controller to have an overall view of traffic from the Internet as well as from an internal network. *EX-1004* 9:9-18.

69. Each of virtual I/O modules 301-304 (green) receives packets from any of servers 321-324 of LAN 320 and sends packets to external network 315 outside of the firewall. In one embodiment, each of I/O modules 301-304 keeps a local cache (e.g., caches 305-308) storing location(s) of a security processing module(s) (e.g., security processing modules 309-311) for each connection session. A cache maintained by each I/O module contains a forwarding table mapping certain connection sessions to any of security modules 309-311. Upon receiving a

packet, an I/O module performs a packet classification to find out the associated connection and forwards the packet to the corresponding security processing module identified by the forwarding table. *EX-1004* 7:32-45.

70. In one embodiment, each of security processing modules 309-311 (blue) performs major security processing functions, such as, for example, network address translation (NAT), virtual private network (VPN), deep packet inspection (DPI), and/or anti-virus, etc. A security processing module receives packets and runs the packets through one or more various security functions in the module for security processing. There could be several security modules and each handles the same or different security functions. If the packets need to go through another security or service processing, the module sends the packets to the other modules. Optionally, it can run the packets through a load balancing mechanism to distribute the load to multiple modules. If a module is the last processing module in the chain to process the packets, it can forward the packets back to the virtual I/O module to send out, or send the packet out directly to its destination if it's configured to do so. *EX-1004* 7:53-67.

71. In one embodiment, each of service processing modules 312-313 (red) performs one or more of the functions of security processing module, such as, for example, NAT, VPN, DPI, and/or anti-virus, etc. However, the service processing module is different from the security processing module in that it only

receives and sends packets to the same security processing module. If the tasks cannot be done in a security processing module, for example, due to a resource limitation, system load, or the requirement of a different operation system, the packets can be forwarded to one or more of service processing modules 312-313 for further processing. The packets then are sent back to the same security processing module for the next security function processing. To further share the system load, any of security processing modules 309-311 can load balance the computational-intensive services using multiple service processing modules. *EX-1004* 8:1-16.

72. If a module is the last processing module in the chain to process the packets, it can forward the packets back to the virtual I/O module to send out, or send the packet out directly to its destination (purple) if it is configured to do so. *EX-1004* 7:64-67.

73. In one embodiment, central controller 208 (orange) is the central place to control forwarding of the packets amongst I/O modules 301-304, security processing modules 309-311, and service processing modules 312-313. When a virtual I/O module receives a packet, according to one embodiment, it forwards the packet to central controller 208 if it cannot find an existing connection in its local cache. When central controller 208 receives the packet, it decides which of security processing modules 309-311 is able to process the packets, and then

forwards the packets to the designated security processing module. It also instructs the virtual I/O module to create the local cache to store connection state information so the subsequent packets of the same connection session do not need to be forwarded to central controller 208: rather, they can be directly forwarded to the proper security processing module identified in the cache. *EX-1004* 8:17-33.

74. Shieh describes several benefits and improvements of using a distributed firewall system in the cloud. For example, by dividing a firewall into different modules, it allows putting virtual I/O and security processing functions at the best locations to protect the network entrance, while keeping the central control and monitoring functionality at the central controller. *EX-1004* 8:33-37.

75. Second, it enhances the scalability of the system since all modules can be expanded to multiple instances to share the system load. *EX-1004* 8:37-39.

76. Third, the distributed firewall can employ a significantly large amount of CPU and memory resources for service processing and protect the networks at multiple geometric locations. *EX-1004* 9:19-23.

77. Fourth, the central controller decides which security processing module is capable of processing particular connection, and is able to start a new security processing at the place deemed best for packet processing. As a result, the location of the packet I/O is not limited on a single appliance. The I/O modules can be placed anywhere as virtual machines. *EX-1004* 9:24-30.

78. Fifth, the security processing power is significantly higher as packets and connections can be load balanced to any number of the security processing modules, and the modules could be added or deleted dynamically. *EX-1004* 9:30-33.

79. Lastly, using such modules in a firewall cloud provides a security design that is best-fit for the emerging cloud computing, and provides great scalability and system availability. *EX-1004* 9:33-36.

2. The Potti Reference

80. Potti is generally directed to enforcing service level agreements in a network infrastructure element, by giving some packets priority over other packets based on the contents of the packets. *EX-1005* Abstract, ¶0344.

81. Potti discloses that to determine whether one or more conditions of a service level agreement are violated, deep packet inspection of packets is required. Because deep packet inspection can be time consuming and interrupt traffic flow, Potti discloses receiving packets in the network element, forwarding the packets to a next hop, and concurrently copying the packets without modification and without interrupting the flow of traffic through the network element. The copied packets are subjected to deep packet inspection in which values of the IP header, TCP header and application-layer payload are examined. *EX-1005* ¶0345.

82. Performing DPI on copied packets rather than packets in the traffic flow is consistent with the State of the Art section described above, that offline packet inspection allows further analysis without interrupting traffic flow.

3. Detailed Application of Shieh and Potti to the Challenged Claims

a. Claim 1

[1.0] A multi-perimeter firewall system located in a cloud and forming part of a global virtual network, comprising:

83. To the extent the preamble is limiting, Shieh discloses a distributed *firewall system located in a cloud* using virtual machines located at several perimeters.

84. Shieh discloses network access device 204 is associated with a distributed firewall 212 that includes various firewall processing modules, for example, each being executed within a virtual machine (VM). *EX-1004 3:5-7.*

85. Further, Shieh discloses that all of the firewall modules 209-210 are managed by firewall controller 208, which may be located within network access device 204, or external to network access device 204, such as, for example, in a public cloud associated with network 203, or in a private cloud associated with network 205. Controller 208 and firewall processing modules 209-210 collectively are referred to as distributed firewall 212. *EX-1004 3:20-26.*

86. Shieh discloses that the security processing modules can be the endpoints of virtual private networks (*virtual network*) capable of performing firewall functionality such as VPN processing. *EX-1004* 3:5-14.

87. Further, Shieh discloses that each of network access devices, including the firewall modules, maintains a persistent connection such as secure connections or tunnels with a controller or management entity. *EX-1004* 4:14-18.

88. Shieh discloses in FIG. 3 that all of the firewall modules are located functionally at perimeters between an internal network and an external network (*multi-perimeter firewall system*). See also FIG 2 (internal network 205 and external network 203).

89. Shieh discloses that by dividing a firewall into different modules, and placing them at various geographic locations in the network, it increases security at the network entrance, while keeping the central control and monitoring functionality at the central controller, and enhances the scalability of the system. *EX-1004* 8:33-42; 9:9-18.

90. A POSITA understood that each of the four types of modules described by Shieh would be implemented in virtual machines distributed in the cloud at various geographic locations and would be considered part of a *global virtual network*.


91. This is consistent with the '482 Patent's disclosure that the firewalls are virtualized in the cloud and the perimeters are located between the client and the internet at each firewall. *EX-1001* 11:17-25; 13:64-14:4.

92. Shieh's disclosure is also consistent with PO's infringement contentions, e.g., "By delivering security and network services together from the cloud, organizations will be able to securely connect any user or device to any application without having to install and maintain the network management and security infrastructure, i.e., a multi-perimeter firewall system located in a cloud and forming part of a global virtual network." *EX-1011* p. 3.

[1.1] an egress ingress point device

93. Shieh discloses this limitation. Any one of the virtual input/output (I/O) module (with cache) (green) of FIG. 3¹ can be the *egress ingress point device*:

¹ This annotated FIG. 3 will be used to illustrate each of the limitations of claim 1 ("Annotated FIG. 3").



EX-1004 FIG. 3

FIG. 3

94. As illustrated in Annotated FIG. 3, virtual I/O modules are an *ingress point device* and an *egress point device* for packets from and to the external network 315², respectively, and an *ingress point device* and an *egress point device* for packets to and from the internal network/LAN 320, respectively.

95. For example, each of virtual I/O modules 301-304 receives packets from any of servers 321-324 of LAN 320 and sends packets to external network 315 outside of the firewall. In one embodiment, each of I/O modules 301-304

² External Network 315 is mislabeled in FIG. 3 as External Network 320. EX-1004

keeps a local cache (e.g., caches 305-308) storing location(s) of a security processing module(s) (e.g., security processing modules 309-311) for each connection session. A cache maintained by each I/O module contains a forwarding table mapping certain connection sessions to any of security modules 309-311. *EX-1004* 7:31-41; FIG. 3.

96. Shieh's disclosure is consistent with the '482 Patent's disclosure that the egress ingress point device can link to both the open internet or to an organizations cloud based assets. *EX-1001* 13:17-19.

***[1.2] a first access point server in communication with
the egress ingress point device;***

97. Shieh discloses that any of the security processing modules 309-311 can be the *first access point server*.

98. For example, with respect to Annotated FIG. 3, security processing module 309 (blue) can be the *first access point server* in communication with virtual I/O 304 (*egress ingress point device* (green)).³ *EX-1004* FIG. 6, 9:9-12 ("In

³ As illustrated in FIG 3, Shieh discloses several different security processing modules in communication with several service processing modules, and also in communication with several virtual I/O modules. In the claims analysis of the Petition, specific examples are provided, but a POSITA understood that there are

one embodiment, firewall modules 300A-300B could be distributed in different networks, even on different locations, as long as the modules can reach the module that is next in terms of processing and the central controller.”). Shieh describes that the security processing module 309 may perform firewall services. *EX-1004* 7:53-67.

99. Shieh describes that one of the advantages is that the distributed firewall can employ a significantly large amount of CPU and memory resources for service processing and protect the networks at multiple geometric locations. The central controller decides which security processing module is capable of processing particular connection. *EX-1004* 9:19-26.

100. Shieh discloses that the security processing modules can perform VPN processing which a POSITA understood to be an endpoint of a secure connection or tunnel of a virtual network (*access point servers*). *EX-1004* 7:53-55; 7:64-67; 4:14-19.

101. Shieh’s disclosure is consistent with the disclosure of the ’482 Patent which describes that the access point servers are endpoints of a virtual tunnel and may perform firewall services as well. *EX-1001* 2:46-48; 4:43-45; 5:9-13.

many other such examples that can be made from the various combinations of security processing modules, service processing modules and I/O modules.

***[1.3] a second access point server in communication
with the first access point sever;***

102. Shieh discloses this limitation for the same reason as with respect to limitation [1.2].

103. For example, the *second access point* can be security processing module 310 (red). Although not shown in Annotated Figure 3, Shieh discloses that the security processing modules can send packets to another security processing module:

A security processing module receives packets and runs the packets through one or more various security functions in the modules for security processing. There could be several security modules and each handles the same or different security functions. If the packets need to go through another security or service processing, the module sends the packets to the other modules.

EX-1004 7:53-63.

***[1.4] an endpoint device in communication with the
second access point server;***

104. Shieh discloses this limitation. Shieh further discloses that destination nodes can be a destination server. *EX-1004* Abstract; 3:2-4.

105. Shieh discloses that security processing module 310 (red) (*second access point*) can be in communication with a destination server 4 (purple)

(*endpoint device*) directly, or through virtual I/O module 304 (green) (*ingress egress point device*):

If a module is the last processing module in the chain to process the packets, it can forward the packets back to the virtual I/O module to send out, or send the packet out directly to its destination if its configured to do so.

EX-1004 7:64-67.

106. For example, in Annotated FIG. 3, security processing module 310 (*second access point device*) can send the packet directly to the destination server 4 of LAN 320 (*endpoint device*). Shieh's disclosure is consistent with the '482 Patent's disclosure that an endpoint can be the destination host server. *EX-1001* 4:53-55.

[1.5] a first perimeter firewall in communication with the first access point server, wherein the first perimeter firewall performs stateful packet inspection to prevent at least some traffic from passing from the first access point server to the second access point server; and

107. Shieh discloses this limitation. Shieh discloses that service processing module 312 (light blue) (*first perimeter firewall*) performs one or more of the security functions for the security processing module 310 (red) (*first access point server*). *EX-1004* 8:1-13.

108. As shown in Annotated FIG. 3, service processing module 312 is located functionally at a perimeter between an internal network 320 and an external network 315. *EX-1004* 7:31-52.

109. Shieh discloses that that the distributed firewalls can be located at multiple boundaries in the cloud between a client location and their link to the internet. With reference to Annotated FIG. 3, a first perimeter (dashed black line) can be located in a cloud where the security processing module 309 and the service processing module 312 are located in the same location, and a second perimeter (dashed black line) can be located in a cloud where the security processing module 310 and the service processing module 313 are located in a different location so long as the two security processing modules are in communication with each other. *EX-1004* 9:9-11.

110. This is consistent with the '482 Patent's disclosure that the firewalls are virtualized in the cloud and the perimeters are located between the client and the internet at each firewall. *EX-1001* 11:17-25; 13:64-14:4.

111. As discussed in the State of the Art section, a well-known security function that had become an industry standard was stateful inspection. *EX-1006* p. 2. Stateful inspection provided for the analysis of packets at the network layer to identify packets belonging to the same flow or session. *EX-1006* p. 2.

112. Although, Shieh does not use the term stateful inspection, a POSITA understood that it would be included in Shieh's generic description that the firewall modules perform "firewall functions" because stateful inspection was an industry standard firewall function. *EX-1004* 3:11-14.

113. Further, Shieh describes that security processing was done to packets belonging to a session that is stateful, and that information regarding each session including the state of the session was saved in a session table. *EX-1004* 4:63-5:3.

114. Shieh further discloses that a security device would perform the security inspection on packets at the beginning of the flow or session, and at a certain point, the security device decides that it no longer needs to inspect further packets of the same session. *EX-1004* 5:38-42. A POSITA understood that using a session table to make the decision to no longer inspect further packets would be as a result of the security inspection being a stateful inspection.

115. Shieh discloses that the security processing done at service processing module 312 is different from the security processing module 309 in that it only receives and sends packets to the same security processing module 309. After processing by service processing module 312 (*first perimeter firewall*), the packets then are sent back to the same security processing module 309 (*first access point server*) for the next security function processing. *EX-1004* 8:1-13. Thus, if service processing module 312 identified traffic that posed a security risk, it would prevent

the traffic from passing from the security processing module 309 (*first access point server*) to the security processing module 310 (*second access point server*).

116. In addition, a POSITA would have found it obvious to modify the teachings of Shieh to include stateful packet inspection at the service processing module 312 (*first perimeter firewall*). For example, as discussed in the State of the Art section, it was well-known to use stateful packet inspection before deep packet inspection as part of a layered defense. *EX-1007* pp. 83-84.

117. As discussed below, the service processing module 313 (*second perimeter firewall*) performs deep packet inspection, and a POSITA would be motivated to include a stateful inspection at service processing module 312 prior to the deep packet inspection for several reasons.

118. First, as explained in the State of the Art, stateful inspection provided for the analysis of packets at the network layer as well other layers in order to assess the overall packet. By combining information from various layers (transport, session, and network) the firewall was better able to understand the protocol it was inspecting. This also provided for the ability to create virtual sessions in order to track connectionless protocols such as UDP-based applications as well as RPC-based applications. *EX-1006* p. 2.

119. Second, because some traffic is blocked by the stateful inspection, it lightens the load on the traffic subject to deep packet inspection. *EX-1010* p. 7.

120. Third, stateful inspection had become an industry standard and would be simple to implement since Shieh already maintained a session table which identified the state of the connection, which is required for stateful inspection.

121. Thus, a POSITA would be motivated to modify the teachings of Shieh to include stateful inspection in service processing module 312 prior to the deep packet inspection performed in security processing module 313 and would be confident of its success.

[1.6] a second perimeter firewall in communication with the second access point server, wherein the second perimeter firewall performs deep packet inspection to prevent at least some traffic from passing from the second access point server to the end point device,

122. Shieh discloses this limitation. Shieh discloses that service processing module 313 (pink) (*second perimeter firewall*) performs one or more of the security functions for the security processing module 310 (red) (*second access point server*). *EX-1004* 8:1-13.

123. As discussed above with respect to limitation [1.5], and as shown in Annotated FIG. 3, service processing module 313 is located functionally at a second perimeter (dashed black line) between an internal network 320 and an external network 315. *EX-1004* 7:31-52.

124. Shieh discloses that the security processing done at service processing module 313 is different from the security processing module 310 in that it only receives and sends packets to the same security processing module 310 (*in communication with the second access point server*). After processing by service processing module 313 (*second perimeter firewall*), the packets then are sent back to the same security processing module 310 (*second access point server*) for the next security function processing. *EX-1004* 8:1-13.

125. Shieh discloses that service processing module 313 (*second perimeter firewall*) performs one or more firewall functions including deep packet inspection (DPI). After processing by service processing module 313 (*second perimeter firewall*), the packets then are sent back to the same security processing module 310 (*second access point server*) for the next security function processing. *EX-1004* 8:1-13.

126. Thus, if service processing module 313 identified traffic that posed a security risk, it would prevent the traffic from passing from the security processing module 310 (*second access point server*) to destination server 324 (*end point device*).

[1.7] wherein at least one of the egress ingress point device, the first access point server, the second access point server, the endpoint device, the first perimeter firewall, or the second perimeter firewall comprises hardware,

127. Shieh discloses this limitation. FIG. 7 illustrates a method for performing firewall operations using the disclosed distributed firewall. *EX-1004* 1:55-57.

128. Shieh discloses that method 700 may be performed by processing logic that may include software, hardware, or a combination of both. *EX-1004* 9:39-41.

129. Further, Shieh discloses that a virtual machine is software hosted on a server (*hardware*) for performing the functions for the different types of modules (e.g., *egress ingress point device, the first access point server, the second access point server, the endpoint device, the first perimeter firewall, or the second perimeter firewall*). *EX-1004* 3:41-43; 11:33-36.

[1.8] wherein the deep packet inspection is performed on a cloned copy of traffic that flows through the second perimeter firewall.

130. Shieh discloses deep packet inspection, but does not expressly state performing deep packet inspection on a cloned copy of traffic that flows through the service processing module 313 (*second perimeter firewall*).

131. However, as discussed in the State of the Art section, a POSITA understood that deep packet inspection (DPI) could be performed inline or offline on a cloned copy of traffic that flows through the firewall. Thus, a POSITA understood Shieh's broad disclosure to include both inline and offline DPI.

132. To the extent that a POSITA would not understand Shieh to disclose performing DPI on a cloned copy of traffic, it would be obvious for a POSITA to modify the teachings of Shieh to include DPI on a cloned copy of traffic.

133. As discussed in the State of the Art section, performing deep packet inspection on a cloned copy of traffic was well-known and provided several benefits.

134. For example, Potti discloses that, because deep packet inspection can be time consuming and interrupt traffic flow, allowing packets to be forwarded to their destination while cloned copies are inspected allows traffic to flow without interrupting the flow of traffic through the network element. *EX-1005* ¶0345. This is particularly important to packets such as voice or streaming packets that require low latency.

135. Shieh specifically identifies using rules to govern how packets are routed and that paths through the network can be set up to optimize for speed, fewest number of hops or reduced latency. *EX-1004* 4:38-44. A POSITA understood that this speed and/or low latency requirement could be fulfilled by performing DPI on cloned copies.

136. Second, using deep packet inspection of cloned copies of traffic allows further analysis of packets to discover trailing trends that may not have been identified using deep packet inspection on the traffic in real-time.

137. Further, a POSITA understood that implementing a security processing function in Shieh that performed deep packet inspection of cloned traffic for traffic that required low latency as taught by Potti could be implemented successfully using well-known techniques.

b. Claim 2

[2] The multi-perimeter firewall system according to claim 1, wherein at least one of the access point servers is configured to perform firewall services.

138. Shieh discloses that security processing module 309 (*first access point server*) and security processing module 310 (*second access point server*) are configured to perform firewall functions, including “for example, NAT [network address translation], VPN [virtual private network], DPI [deep packet inspection], and/or anti-virus, etc.” *EX-1004* 7:53-63.

c. Claim 3

[3] The multi-perimeter firewall system according to claim 1, wherein the first perimeter firewall is in communication with the second perimeter firewall through a communication path.

139. Shieh discloses this limitation for the reasons discussed with respect to limitation [1.3], [1.5], [1.6]. Specifically, the service processing module 312 (*first perimeter firewall*) is in communication with security processing module 313

(*second perimeter firewall*) through a communication path through security processing module 309 (*first access point server*). *EX-1004* 7:53-63; 8:1-13.

140. Shieh further discloses that packets can be forwarded through the central controller 208 to the appropriate security processing modules. *EX-1004* 7:45-52.

141. Shieh further discloses that the controller could set up paths through the network optimized for speed, fewest number of hops or reduced latency, among other characteristics. Using OpenFlow takes control of how traffic flows through the network out of the hands of the infrastructure, the switches and routers, and puts it in the hands of the network owner (such as a corporation), individual users or individual applications. *EX-1004* 4:44-48.

142. Thus, Shieh also discloses a communication path between the service processing modules through the security processing modules and central controller.

d. Claim 4

[4] The multi-perimeter firewall system according to claim 3, wherein the communication path between the first perimeter firewall and the second perimeter firewall is a network path.

143. Shieh discloses this limitation in FIG. 2A which shows that firewall modules 209A-209C may be part of a distributed firewall connected through a network path (solid lines).

144. For example, when firewall module 209B receives a packet from node 206, it may forward the packet to firewall processing node 211A for content inspection using a network path. *EX-1004* 5:5-15.

145. Further, Shieh discloses the service processing modules 312 and 313 (*first perimeter firewall and second perimeter firewall*) could be located in different networks and thus would communicate through a network path. *EX-1004* 9:9-18.

146. Additionally, as discussed with limitation [3], Shieh discloses the employment of OpenFlow that allows paths to be set up between different entities on the network. *EX-1004* 4:44-48.

e. **Claim 5**

[5] The multi-perimeter firewall system according to claim 3, wherein the communication path between the first perimeter firewall and the second perimeter firewall is a network back channel.

147. The '482 Patent describes a back channel mechanism that offers a channel into devices which are not reachable from the open internet. *EX-1001* 6:8-11.

148. Shieh discloses this limitation in FIG. 2A which shows that firewall modules 209A-209C may be part of a distributed firewall connected through a secured tunnel (*network back channel*) (dashed line).

149. For example, when firewall module 209B receives a packet from node 206, it may forward the packet to firewall processing node 211A for content inspection and forwards the packet to controller 208 for routing information using secure tunnel 260. Each of the firewall processing nodes 211A-211B may further maintain a persistent connection or tunnel with controller 208, for example, using the OpenFlow communication protocol. *EX-1004* 5:5-25.

150. Thus, Shieh discloses a network back channel connecting the firewalls through the firewall controller.

f. **Claim 6**

[6] The multi-perimeter firewall system according to claim 3, wherein the first perimeter firewall and the second perimeter firewall share threat information including at least one of heuristic patterns, signatures of known threats, known malicious source IP addresses, or attack vectors.

151. Shieh discloses that the distributed firewalls exchange information and configuration, which a POSITA understood would include threat information such as a source IP address related to a malicious attack.

152. For example, Shieh discloses that all of the firewall modules 209-210 are managed by firewall controller 208. *EX-1004* 3:20-26.

153. For example, each of network access devices 204A-204C maintains a persistent connection such as secure connections or tunnels 260 with a controller or management entity 208 for exchanging management messages and configurations,

or distributing routing information to network access devices 204A-204C, etc. *EX-1004* 4:14-32.

154. Shieh discloses that an administrator 265 configures, for example, via a controller or a management entity 208, a network access device (e.g., network access devices 204A-204C) to set up a set of filtering rules concerning whether and/or what types of packets should be forwarded to a security device and which of the security devices (e.g., security devices 211A-211B) for security inspection. Controller 208 is configured to manage multiple network access devices 204A-204C and/or multiple security devices 211A-211B. *EX-1004* 5:25-34.

155. As discussed in the State of the Art, one of the purposes for performing deep packet inspection is to identify malicious content, and the source IP address associated with a received IP packet or packets having malicious content. A POSITA understood that Shieh's teaching that a firewall controller exchanges management messages and configurations with firewall modules would include source IP address identified by a firewall performing deep packet inspection that would be sent to other firewall devices as part of the packet filtering rules. *EX-1009* pp. 659-666.

156. Shieh's disclosure of exchanging messages and configurations between security modules through a controller is consistent with the '482 Patent's

disclosure that threat information is shared via a central control server. *EX-1001* 2:53-58.

157. In addition, to the extent that Shieh does not expressly state that the firewall modules share threat information, a POSITA would be motivated to modify the teachings of Shieh to include the well-known feature of “attack information advertising”, as discuss in the State of the Art section, to ensure that firewalls on the perimeter take countermeasures to prevent malicious packets from entering the network. *EX-1009* p. 1915.

158. A POSITA understood that one of the purposes of a firewall performing deep packet inspection is to identify threats, such as the source IP address of a packet that contains malicious content. A POSITA would be motivated to modify the teachings of Shieh to share the identified source IP address with other firewalls as part of the filtering rules to be used to block or drop packet from the identified source IP address of malicious content in order to protect the network. This would result in the benefit of traffic being blocked at a more outer perimeter (e.g., as traffic enters from the external network) as part of a well-designed layered defense approach and defense in depth. A POSITA would be confident of the success of this modification as it would be implementing the layered defense that was discussed in the State of the Art section.

g. Claim 7

[7] The multi-perimeter firewall system according to claim 6, wherein the first perimeter firewall and the second perimeter firewall share threat information with a central control server.

159. Shieh discloses this for the same reason as discussed with respect to Claim 6.

h. Claim 8

[8] The multi-perimeter firewall system according to claim 1, wherein the deep packet inspection is performed on flow through traffic.

160. As discussed with respect to limitation [1.6], Shieh discloses that service proceeding module performs deep packet inspection, and a POSITA understood there is no need to perform DPI on cloned traffic for traffic that is not speed and latency sensitive. Thus, a POSITA understood that the same firewall performs DPI on cloned traffic could perform DPI on flow through traffic.

i. Claim 9

[9] The multi-perimeter firewall system according to claim 1, wherein at least one of the firewalls includes a cloud firewall load balancer and wherein the cloud firewall load balancer can allocate cloud firewall resources on demand.

161. Shieh discloses that any of security processing modules 309-311 can load balance the computational-intensive services using multiple service processing modules. *EX-1004* 8:13-16; 7:62-64.

162. Shieh discloses load balancing across any number of the security processing modules in the cloud, which provides a security design that is best-fit for the emerging cloud computing, and provides great scalability and system availability. *EX-1004* 9:28-36.

163. Further, Shieh discloses that each of service processing modules 312-313 performs one or more of the functions of security processing module 309-311, which a POSITA understood would include load balancing. *EX-1004* 8:1-3.

164. Shieh's disclosure is consistent with the '482 Patent's disclosure that cloud based firewalls are dynamically scalable via cloud firewall load balancing mechanisms that can bring more resources online as needed. *EX-1001* 11:20-22.

j. **Claim 10**

[10] The multi-perimeter firewall system according to claim 1, wherein the egress ingress point device receives traffic directly from the Internet.

165. As shown in Annotated FIG. 3, Shieh discloses that virtual I/O module 304 (*egress ingress point device*) receives traffic directly from the external network 315, which a POSITA understood includes the Internet. *EX-1004* FIG. 3; 2:65-67.

k. **Claim 11**

[11] The multi-perimeter firewall system according to claim 1, wherein the first perimeter firewall is coupled to a multi-honed backbone.

166. Assuming Claim 11 is referring to a “multi-homed”, a POSITA understood it to mean that the connection to the internet is provided through at least two links for redundancy.

167. Shieh does not provide the details of the connection to the internet but using a multi-homed backbone was common and well-known, and FIG. 3 shows that multiple virtual I/O modules, which communicate with each other, are each in communication with the internet, which a POSITA would consider to be a multi-homed backbone.

168. Further, a POSITA understood that cloud routinely multi-home their services and provide a load balancer in front of the virtual I/O modules to ensure availability.

I. Claim 12

[12] The multi-perimeter firewall system according to claim 1, wherein the at least some traffic prevented from passing from the first access point server to the second access point server includes traffic with recognized malicious headers.

169. As described in the State of the Art section, filtering data according to a source IP address (*recognized malicious headers*) is a very well-known firewall function, which can be implemented as a filtering rule in service processing module 312 (*first perimeter firewall*) as described in limitation [1.5] to prevent traffic from passing from the first security processing module 309 (*first access*

point server) to the second security processing module 310 (*second access point server*).

170. Further, Shieh provides IP address based specific services with rules governing the traffic from a specific IP address. *EX-1004* 4:38-41. Thus a POSITA understood that packets known to be malicious from a source IP address would be subject to a rule to be dropped.

m. Claim 13

[13.0] A multi-perimeter firewall system located in a cloud and forming part of a global virtual network, comprising:

171. Shieh discloses this limitation for the same reasons as discussed with respect to limitation [1.0].

[13.1] an egress ingress point device;

172. Shieh discloses this limitation for the same reasons as discussed with respect to limitation [1.1].

[13.2] a first access point server in communication with the egress ingress point device;

173. Shieh discloses this limitation for the same reasons as discussed with respect to limitation [1.2].

[13.3] a second access point server in communication with the first access point sever;

174. Shieh discloses this limitation for the same reasons as discussed with respect to limitation [1.3].

[13.4] an endpoint device in communication with the second access point server;

175. Shieh discloses this limitation for the same reasons as discussed with respect to limitation [1.4].

[13.5] a first perimeter firewall in communication with the first access point server, wherein the first perimeter firewall performs stateful packet inspection to prevent at least some traffic from passing from the first access point server to the second access point server using a first firewall mechanism; and

176. Shieh discloses this limitation for the same reasons as discussed with respect to limitation [1.5].

177. The *first firewall mechanism* is the stateful packet inspection. During prosecution, the examiner rejected [issued claim 13] which recited that the first firewall used a first firewall mechanism and the second firewall used a second different, firewall mechanism. *EX-1002* pp. 42-44.

178. In response, the applicant amended [issued Claim 13] to include the limitation that the first firewall “performs stateful packet inspection” and the second firewall “performs deep packet inspection.” *EX-1002* pp. 63-64.

179. Thus, the *first firewall mechanism* is specifically recited to be stateful packet inspection, which is different from the *second firewall mechanism*, which is specifically recited to be deep packet inspection.

[13.6] a second perimeter firewall in communication with the second access point server, wherein the second perimeter firewall performs deep packet inspection to prevent at least some traffic from passing from the second access point server to the end point device using a second, different firewall mechanism,

180. Shieh discloses this limitation for the same reasons as discussed with respect to limitation [1.6] and [13.5]. The *second, different firewall mechanism* is deep packet inspection.

[13.7] wherein at least one of the egress ingress point device, the first access point server, the second access point server, the endpoint device, the first perimeter firewall, or the second perimeter firewall comprises hardware,

181. Shieh discloses this limitation for the same reasons as discussed with respect to limitation [1.7].

[13.8] wherein the deep packet inspection is performed on a cloned copy of traffic that flows through the second perimeter firewall.

182. Shieh discloses this limitation for the same reasons as discussed with respect to limitation [1.8].

n. Claim 14

[14] The multi-perimeter firewall system according to claim 13, wherein at least one of the access point servers is configured to perform firewall services.

183. Shieh discloses this limitation for the same reasons as discussed with respect to Claim 2.

o. Claim 15

[15] The multi-perimeter firewall system according to claim 13, wherein the first perimeter firewall is in communication with the second perimeter firewall through a communication path.

184. Shieh discloses this limitation for the same reasons as discussed with respect to Claim 3.

p. Claim 16

[16] The multi-perimeter firewall system according to claim 15, wherein the communication path between the first perimeter firewall and the second perimeter firewall is a network path.

185. Shieh discloses this limitation for the same reasons as discussed with respect to Claim 4.

q. Claim 17

[17] The multi-perimeter firewall system according to claim 15, wherein the communication path between the first perimeter firewall and the second perimeter firewall is a network back channel.

186. Shieh discloses this limitation for the same reasons as discussed with respect to Claim 5.

r. Claim 18

[18] The multi-perimeter firewall system according to claim 15, wherein the first perimeter firewall and the second perimeter firewall share threat information including at least one of heuristic patterns, signatures of known threats, known malicious source IP addresses, or attack vectors.

187. Shieh discloses this limitation for the same reasons as discussed with respect to Claim 6.

s. Claim 19

[19] The multi-perimeter firewall system according to claim 18, wherein the first perimeter firewall and the second perimeter firewall share threat information with a central control server.

188. Shieh discloses this limitation for the same reasons as discussed with respect to Claim 7.

t. Claim 20

[20] The multi-perimeter firewall system according to claim 13, wherein the deep packet inspection is performed on flow through traffic.

189. Shieh discloses this limitation for the same reasons as discussed with respect to Claim 8.

u. Claim 21

[21] The multi-perimeter firewall system according to claim 13, wherein at least one of the firewalls includes a cloud firewall load balancer and wherein the cloud firewall load balancer can allocate cloud firewall resources on demand.

190. Shieh discloses this limitation for the same reasons as discussed with respect to Claim 9.

v. **Claim 22**

[22] The multi-perimeter firewall system according to claim 13, wherein the second perimeter firewall is between the endpoint device and the second access point server.

191. Shieh discloses this limitation for the same reasons as discussed with respect Claim [1.6]. Service processing module 313 (*second perimeter firewall*) performs DPI after receiving a packet from security processing module 310 (*second access point server*) and before it is forwarded to destination server 324 (*endpoint device*). See Annotated FIG. 3.

w. **Claim 23**

[23] The multi-perimeter firewall system according to claim 13, wherein the first perimeter firewall is between the first access point server and the second access point server.

192. Shieh discloses this limitation for the same reasons as discussed with respect Claim [1.5]. Service processing module 312 (*first perimeter firewall*) performs SPI after receiving a packet from security processing module 309 (*first*

access point server) and before it is forwarded to security processing module 310 (second access point server). See Annotated FIG. 3.

x. **Claim 24**

[24] The multi-perimeter firewall system according to claim 13, wherein the egress ingress point device receives traffic directly from the Internet.

193. Shieh discloses this limitation for the same reasons as discussed with respect to Claim 10.

y. **Claim 25**

[25] The multi-perimeter firewall system according to claim 13, wherein the first perimeter firewall is coupled to a multi-honed backbone.

194. Shieh discloses this limitation for the same reasons as discussed with respect to Claim 11.

z. **Claim 26**

[26] The multi-perimeter firewall system according to claim 13, wherein the at least some traffic prevented from passing from the first access point server to the second access point server includes traffic with recognized malicious headers.

195. Shieh discloses this limitation for the same reasons as discussed with respect to Claim 12.


VIII. SUMMARY OF THE '482 PATENT

196. The '482 Patent is directed to network security which protects the flow of traffic by the strategic positioning of distributed firewall devices placed at multiple perimeters in the cloud. *EX-1001* 1:15-19. The '482 Patent describes that it is an improvement upon the standard use of firewalls through increasing their utility value by extending perimeters into the cloud. *EX-1001* 5:56-58.

197. The '482 Patent describes several advantages of a cloud based firewall systems. For example, locating a multi-perimeter firewall in the cloud allows identifying and isolating problem traffic at a point upstream from the location of a traditional firewall, and thus identifies problem traffic before it reaches the “last mile” connection network that tends to be slower than the backbone network. *EX-1001* 11:4-16. Further, cloud based firewalls are dynamically scalable via cloud firewall load balancing mechanisms that can bring more resources online as needed. *EX-1001* 11:20-22. In addition, the '482 Patent describes that a cloud based firewall offers the advantage of running different firewall mechanisms, such as stateful packet inspection and deep packet inspection, simultaneously. *EX-1001* 11:43-45.

198. FIG. 12 of the '482 Patent as annotated (Annotated FIG. 12)

illustrates one embodiment of a multi-perimeter firewall in the cloud for protecting a flow of traffic through a network.



EX-1001 FIG. 12

FIG. 12

199. The network may comprise an egress ingress point device (green), a first access point server (dark blue) and second access point server (red), and an endpoint device (Purple). A first firewall (light blue) is in communication with the first access point server and may prevent network traffic from flowing through the

first access point server. A second firewall (pink) is in communication with the second access point server and may prevent network traffic from flowing through the second access point server. *EX-1001* 2:36-45.

200. The '482 Patent describes that in one embodiment, the first firewall performs stateful packet inspection, and the second firewall performs deep packet inspection. *EX-1001* 2:59-61. The '482 Patent describes that when performing stateful packet inspection, the firewall examines the headers of packets to identify offending information while it ignores the payload of the packets. The '482 Patent explains that the advantages of stateful packet inspection over deep packet inspection are that it is fast and that it consumes lower processor, RAM and resources. The disadvantage is that the firewall does not look at the content within the payload of the packet(s). *EX-1001* 9:67-10:4.

201. On the other hand, deep packet inspection looks beyond and examines the content, e.g. payload of the packets. The advantages are that this examination is more comprehensive in providing visibility into content both within the payload of not just one packet but from a compilation payload from a series of multiple packets. The disadvantage of DPI is that it is significantly slower than stateful packet inspection and also that it consumes considerably more processor, RAM, and other resources than stateful packet inspection. *EX-1001* 10:5-15.

202. The '482 Patent describes that stateful packet inspection firewalls are devices through which the traffic flows, while deep packet inspection firewall devices can either be flow through or can analyze cloned copies of traffic offering the option for deep packet inspection functionality as a trailing indicator. If harmful traffic is detected, it can be subsequently blocked once identified. A benefit of the communication between devices is that information about bad traffic sources identified by deep packet inspection firewalls can be messaged to stateful packet inspection firewalls to be blocked there. *EX-1001* 12:65-13:6. The '482 Patent describes the known benefit of analyzing cloned traffic as offering trailing results of time- and resources-expensive detection operations such as deep packet inspection without impeding the speed of the flow of traffic. *EX-1001* 13:24-27.

203. The '482 Patent describes that the distributed firewalls can be located at multiple boundaries in the cloud between a client location and their link to the internet. With reference to Annotated FIG. 12, the second perimeter 12-M2 (solid black line) is a boundary in the data center in close proximity to the second access point server (SRV_AP 12-300), and the first perimeter is another boundary at either the same data center as the second access point server (SRV_AP 12-300) or at another location in close proximity to the first access point server (SRV_AP 12-302), possibly in another region. *EX-1001* 13:64-14:4.

204. The '482 Patent describes that in one embodiment, an access point server may also perform firewall services. *EX-1001* 2:46-48.

205. The '482 Patent describes that in some embodiments, the first firewall and the second firewall share threat information including at least one of heuristic patterns, signatures of known threats, known malicious source IP addresses, or attack vectors. The threat information may be shared via a central control server. *EX-1001* 2:49-58.

206. The '482 Patent describes that in some embodiments, at least one of the firewalls includes a cloud firewall load balancer that can allocate cloud firewall resources on demand. *EX-1001* 2:65-67.

CONCLUSION AND DECLARATION

The findings and opinions set forth in this declaration are based on my research and analysis to date. I may continue my examinations in view of additional documents or other factual evidence over the course of this proceeding that may necessitate supplementing and/or refining my opinions. I reserve the right to add to, alter, or delete my opinions and my Declaration upon discovery of additional information. I reserve the right to make such changes as necessary.

In signing this declaration, I understand that the declaration will be filed as evidence in a contested case before the Patent Trial and Appeal Board of the United States Patent and Trademark Office. I acknowledge that I may be subject to cross-examination in the case and that cross-examination will take place within the United States. If cross-examination is required of me, I will appear for cross-examination within the United States during the time allotted for cross-examination.

I declare that all statements made herein on 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 the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

[SIGNATURE ON FOLLOWING PAGE]

Dated: 26th January 2024

By A. L. Narasimha Reddy

Dr. A.L. Narasimha Reddy

APPENDIX A – MATERIALS RELIED UPON

<i>Exhibit #</i>	<i>Description</i>
1001	U.S. Patent No. 10,574,482 (“the ’482”)
1002	Prosecution history of the ’482 (“’482 File History”)
1004	U.S. Pat. No. 8,955,093 (“Shieh”)
1005	U.S. Pat. App. Pub. No. 2007/0156919 (“Potti”)
1006	Firewall Evolution- Deep Packet Inspection, Security Focus, Ido Dubrawsky (2003) accessed at https://web.archive.org/web/20031206194534/http://www.securityfocus.com/printable/infocus/1716 (Dec. 6, 2003).
1007	Cisco Router Firewall Security, Richard A. Deal, Cisco Press, 2005
1008	Cisco ASA, All-in-One Firewall, IPS, and VPN Adaptive Security Appliance, Jazib Frahim, Cisco Press, Second Edition, 2010
1009	A 10Gb/s Firewall System for Network Security in Photonic Era, IEICE Trans. Commun., Vol. E88-B, No. 5 May 2005
1010	Security as a Service for Hybrid Clouds, Markku Vajaranta, Master of Science Thesis, May 2014
1011	Exhibit 5, Infringement Contentions for ’482 Patent, UMBRA Technologies Inc. , v. Cisco Systems Inc., USDC WD Tex, Case 1:23-cv-00903 Document 1-5, Filed 8/02/23
1012	U.S. Pat. Publ. 2012/0210417

APPENDIX B – *CURRICULUM VITAE*

NARASIMHA REDDY
Dept. of Elec. & Comp. Engg.
Texas A & M University
333J Wisenbaker, College Station, TX 77843-3259
<https://cesg.tamu.edu/faculty/narasimha-reddy/>
<https://scholar.google.com/citations?user=en&user=NbHZDW4AAAAJ>
reddy@tamu.edu, 979-845-7598 (O), 979-845-2630 (FAX)

Educational Record:

University of Illinois (Urbana) : Ph.D. in Electrical and Computer Engineering. Aug. 1990.

University of Illinois (Urbana) : M.S. in Electrical and Computer Engineering. May 1987.

Indian Institute of Technology (Kharagpur, India) : Bachelor of Technology (Honors) in Electronics and Electrical Communication Engineering. May 1985.

Thesis Topic:

Ph.D Thesis : Issues in IO architectures for high-performance/ multiprocessor systems.

M.S. Thesis : Algorithm based Fault Detection Techniques for Specific Applications.

B.Tech Thesis (Project) : Studies on Residue Number System for Signal Processing

Honors and Awards:

Cited for one of the most influential papers of first ACM Multimedia Conference

Catedra de Excelencia, University of Carlos III, Madrid, Spain, Aug. 2009-Feb.2010.

Distinguished Achievement Award for teaching (College level), AFS, TAMU, 2005.

Ford Fellow, College of Engg., TAMU 2004, TEES Fellow, College of Engg., TAMU, 2001, Neely Faculty Fellow, College of Engg., TAMU, 1999.

Outstanding Professor Award, Dept. of ECE, TAMU, 2003-04, 2014-15.

Outstanding Professor Award, IEEE Student branch, TAMU, 1997-1998.

NSF Career Award (1996-2000).

5 Patents and an IBM Plateau Invention Achievement Award.

IBM Graduate Fellowship (1988-90).

Presidential Silver Medal, Indian Institute of Technology, Kharagpur, India (May 86).

Research Interests:

Storage Systems, Cybersecurity, Computer Networks and Computer Architecture.

Work Experience:

Associate Dean for Research, College of Engineering, Texas A&M Univ., Apr. 2016 -Present

Associate Agency Director for Strategic Initiatives and Centers, TEES, Aug. 2016 - Present.

Assistant Agency Director for National and Global Initiatives, TEES, June 2014 - Apr. 2016.

J.W. Runyon Professor, Dept. of ECE., Texas A&M University, Oct. 2008-Present.

Professor, Dept. of Elec. and Compu. Engg, Texas A & M University, Sep. 2004 -Sep. 2008.

Associate Professor, Dept. of Elec. Engg, Texas A & M University, Aug. 1995 - Aug. 2004.

Research Staff Member, IBM Almaden Research Center: Aug. 1990 - Aug. 1995.

Professional Activities:

- **Fellow** of the IEEE Computer Society, member of ACM SIGARCH.
- Associate Editor, IEEE/ACM Transactions on Networking, 2014-2016.
- Member, Advisory Board of Center for Hyper Scale Micro Data Center, South Korea.
- External Evaluator, Faculty Promotion Committee, University of Jordan, 2012.
- Program Committee Member, INFLOW Workshop, 2016.
- Program Committee Member, NVM Workshop, 2011, 2015, 2016.
- Program Committee Member, IEEE ANTS Conf., 2014.
- Program Committee Member, Parallel Data Storage Workshop, 2012-2014.
- Member, Information Technology Research Academy Strategy Formulation, India, Oct. 2011
- Review Chair, Los Alamos National Lab Research Information Institute program, Apr. 2011
- Program Committee Member, USENIX HotStorage Conf., 2010-2012.
- External Evaluator, Faculty Search Committee, ICS-FORTH, Crete, Greece, 2009
- Program Co-chair, IEEE Broadnets 2008
- Area Lead, IEEE Infocom 2008
- Program Committee Member, USENIX FAST Conf. 2009
- Panels Co-Chair, COMSWARE 2008
- Program Committee Member, ICDCS 2008, 2013
- Program Committee Member, IEEE ICNP 2008
- Program Committee Member, COMSNETS 2009-2013, 2015
- Program Committee Member, COMSWARE 2007-2008
- Program Committee Member, Int. Workshop Visualization for Computer Security, 2006
- Program Committee Member, CoNEXT Conference, 2006, 2007, 2009.
- Program Committee Member, Int. Workshop on QOS, 2006-2011.
- Program Committee Member, IEEE Infocom, 2005-2015.
- External Reviewer, Faculty promotion Committee, KAIST, Korea, 2003.
- Technical Consultant, EMC Corp., 2001-2003.
- Educational Content Expert on Computer Networking, Jones International University.

- Member, Technical Advisory Board, Netgeo Inc., Mountain View.
- Technical Adviser, TBD Networks.
- Proposal Reviewer, Netherlands NSF.
- Program Committee member, ACM Multimedia Conference, 2002.
- Program Committee member, MCW-2001, Multimedia on Clusters of Workstations, 2001
- Program Committee member, IEEE Int. Conf. on Multimedia Comp. and Systems, 1999
- Panel Chair, IEEE Conf. High Performance Computer Architecture, 1999.
- Program Committee member, High Performance Computing, 1999
- Program Committee member, Int. Conf. on Parallel Processing, 1998.
- Tutorial Chair, IEEE Conf. on Multimedia Computing Systems, 1998.
- Panel Chair, Int. Conf. on Parallel Processing, 1995.
- Program Committee member, Int. Conf. on Parallel Processing, 1995.
- Program track coordinator, Hawaii Int. Conf. on Systems and Sciences, 1994.

Patents

- N. Annapareddy, C-C Chou, C. Tirumalasetty, P. Gratz and A. Abouelwafa, "Faster Page Handling in Computer Systems", Provisional Patent Application, May 2020.
- N. Annapareddy, J. Brady, D. Finney, M. Hartung, M. Ko, J. Menon, and D. Nowlen, Conflict resolution in multi-node communication network, U.S. Patent number: 6,101,194, received Aug. 8, 2000.
- N. Annapareddy, J. Brady, D. Finney, R. Freitas, M. Ko, and M. Rayfield, Method and System for message status reporting in a multi-node communication network, U.S. Patent number: 5,717,862, received Feb. 10, 1998.
- N. Annapareddy, D. Finney, M. Jenkins, L. Kessler, A. Lang, S. Liang, D. Mora, D. Plomgren, P. Urbisci and A. Walls, Interconnection network for a multi-nodal data processing system which exhibits incremental scalability, US Patent number: 5,602,044, received Feb. 11, 1997.
- N. Annapareddy, J. Brady and D. Finney, Adaptive and Dynamic message routing system for multinode wormhole networks, US Patent number: 5,602,839, Received Feb. 11, 1997.
- N. Annapareddy, J. Brady, D. Finney, R. Freitas, M. Hartung, M. Ko, N. Mendelson, J. Menon, D. Nowlen and S-Y Tzou, System and method using chaines structured queues for ordering of message delivery between connected nodes wherein unsuccessful message portion is skipped and retried, US Patent number: 5,577,211, Received Nov. 19, 1996.

Archival Publications

- Mahesh, P., Tiwari, A., Jin, C., Kumar, P.R., Reddy, A. L. N., Bukkapatnam, S., Gupta, N., and Karri, R., "A Survey of Cybersecurity of Digital Manufacturing", Accepted for publication in Proceedings of IEEE, October 2020.

- Tiwari, A., Reddy, A. L. Narasimha, and Bukkaparnam, S., “Cybersecurity assurance in the emerging Manufacturing-as-a-Service (MaaS) paradigm: A lesson from the video streaming industry, *ASTM Journal on Smart and Sustainable Manufacturing Systems*, SSMS-2020-0066.R1, October 2020.
- Wen, F., Qin. M.¹, Gratz, P., and Reddy, A. L. Narasimha, “Hardware memory management for future mobile hybrid memory systems”, accepted for publication in *IEEE Trans. on CAD*, Nov. 2020
- Chou, C-C1, Jung, J., Reddy, A. L. Narasimha, Gratz, P., and Voight, D., “Virtualize and Share nonvolatile memories in user space”, *CCF Transactions on High Performance Computing*, Jan 2020
- Moon, S.¹, and Reddy, A. L. Narasimha, “Does RAID improve lifetime of SSD Arrays?”, accepted for publication at *ACM Transactions on Storage*., Apr. 2016.
- Fedorov, V.¹, Qiu, S.¹, Reddy, A. L. Narasimha and Gratz, P., “ARI: Adaptive LLC-memory traffic management”, *ACM Transactions on Architecture and Code Optimization (TACO)*, Dec. 2013.
- Wu, X¹, Qiu, S.¹, and Reddy A. L. N., “SCMFS: A file system for Storage Class Memory and its extensions”, accepted for publication in *ACM Transactions on Storage*, Jan. 2013.
- Wu, X¹ and Reddy A. L. N., “A novel approach to manage hybrid storage system”, *Journal of communications*, vol.7., July 2012.
- Shin, S., Gu, G., Reddy, A. L. N., “A large-scale study of Conficker”, *IEEE Trans. Information Forensics and Security*, June 2012.
- Yadav, S¹, Reddy, A. K.¹, Reddy, A. L. Narasimha and Ranjan, S., “Detecting algorithmically generated malicious domain names”, *ACM/IEEE Trans. on Networking*, 2012.
- Lee, Y-O.¹ and Reddy, A. L. Narasimha, “Algorithms for disjoint path construction for failure recovery and multipath routing”, to appear in *Journal on Computer Networks*, 2012.
- Garrison, J.¹, and Reddy, A. L. Narasimha, “Umbrella File System: Storage management across heterogeneous devices”, *ACM Transactions on Storage*, vol.5, no.1, Mar. 2009.
- Kim, S.¹, and Reddy, A. L. Narasimha, “Statistical Techniques for detecting Traffic Anomalies through Packet header data” *IEEE/ACM Transactions on Networking (ToN)*, July 2008.
- Liu, Y.¹, and Reddy, A. L. N., ”Multihoming route control among a group of multihomed stub networks”, *Journal on Computer Communications*, Nov. 2007.
- Bhandarkar, S.¹, Reddy, A. L. N., Allman, M., Blanton, E., ”Improving the robustness of TCP to non-congestion events”, *IETF(Internet Engineering Task Force) RFC 4653*, August 2006.
- Kang, S¹., Reddy, A. L. N., “An Approach to Virtual Allocation in Storage Systems”, in *ACM Transactions on Storage*, Nov. 2006, Vol.2, pages:371-399.
- Kim, S¹., Reddy, A. L. N., “Image based Anomaly detection: Design, Algorithms and Evaluation”, *IEEE Journal on Selected Areas of Comm.*, vol. 24.10, October 2006, 1942-1954.

- Bhandarkar, S.¹, Reddy A. L. N., “LTCP: A layered framework for improving TCP’s performance in high-speed networks”, in ACM CCR, Jan. 2006.
- Kwon, D.W., Ko, M., Vannucci, M., Reddy, A. L. N., and Kim, S.¹, “Wavelet methods for the detection of anomalies and their application to network traffic analysis”, Quality and Reliability Engineering International (QREI) 2006, 22, pages: 1-17.
- Reddy, A. L. N., Wyllie, J., and Wijayarathne, R.¹, “Scheduling for multimedia servers”, invited submission for the premier issue of ACM Transactions on Multimedia, March 2005. Cited for being one of the most influential papers from ACM Multimedia conference 1993.
- Kim, S.¹, and Reddy, A. L. N., “Statistical techniques for detecting traffic anomalies through packet header data”, accepted for publication in ACM/IEEE transactions on Networking, January 2007.
- Garg, A.¹, and Reddy, A. L. N., “Mitigating Denial of service attacks through QOS regulation”, in the special issue of Journal of Microprocessors and Microsystems on secure platforms, May 2004.
- Bhandarkar, S.¹, Sadry, N.¹, Reddy A. L. N., Vaidya, N., “TCP-DCR: A novel protocol for tolerating wireless channel errors”, IEEE Transactions on Mobile Computing, March 2005.
- Zhao, Z.¹, Ametha. J., Swaroop. D. and Reddy A. L. N., “A method for estimating non-responsive traffic at a router”, ACM/IEEE Trans. on Networking, Aug. 2004.
- Yeom, I.¹ and Reddy A. L. N., “Adaptive Marking for Assured Forwarding Service” IEEE Communication Letters, Oct. 2002.
- Yeom, I.¹ and Reddy A. L. N., “Providing consistent delay differentiation” accepted by IEICE Transactions on Communications, Sept. 2002.
- Ma, X.¹ and Reddy A. L. N., “Multiview Storage Systems”, IEEE Trans. on Parallel and Distributed Systems, Apr. 2003.
- Yeom, I.¹ and Reddy A. L. N., “Marking for QOS Improvement”, Elsevier Journal of Computer Communications, pages 35-50, vol. 24, no.1, Jan 2001.
- Yeom, I.¹ and Reddy, A. L. N., “Modeling TCP behavior in a differentiated services network”, ACM/IEEE Transactions on Networking, vol. 9, no.1, pages 31-46, Feb. 2001.
- Ma, G.¹, Khaleel, A.¹ and Reddy, A. L. N., “Evaluation of storage systems based on network-attached disks”, IEEE Transactions on Parallel and Distributed Systems, September 2000, 956-968.
- Wijayarathne, R.¹ and Reddy, A. L. N., “Providing QOS guarantees for disk I/O”, Proc. of ACM/Springer Journal on Multimedia Systems, Jan 2000, vol.8, no.1, pages:57-68.
- Yeom, I.¹ and Reddy, A. L. N. , “ENDE: An End-to-end Network Delay Emulator” , Journal on Multimedia Tools and Applications, Kluwer Publishers, vol. 14, no.3, June 2001.
- Reddy, A. L. N. and Upfal, E., “Real-time communication scheduling in a multicomputer video server”, *Journal on Parallel and Distributed Computing*, Sep. 1999., vol.58, pages: 425-445.

- Reddy, A. L. N. and Wyllie, J., “I/O issues in a multimedia system”, *IEEE Computer*, Mar. 1994, vol.27, no.3, pages:69-74.
- Reddy, A. L. N., Chandy, J. and Banerjee, P., “Design and evaluation of gracefully degradable disk arrays”, *Journal of Parallel and Distributed Computing*, Jan. 1993, vol.17, pages: 28-40.
- Reddy, A. L. N. and Banerjee, P., “Algorithm Based Fault Detection Techniques for Signal Processing Applications”, *IEEE Transactions on Computers*, Oct. 1990, vol.39, no.10, pages: 1304-1308.
- Reddy, A. L. N. and Banerjee, P., “Design, Analysis and Simulation of IO Architectures for Hypercube Multiprocessors” *IEEE Transactions on Parallel and Distributed Systems*, April 1990, vol.1, no.2, pages:140-151.
- Reddy, A. L. N. and Banerjee, P., “An evaluation of multiple-disk IO systems” *IEEE Transactions on Computers*, Dec. 1989, vol. 38, no.12, pages:1680-1690.
- Reddy, A. L. N. and Banerjee, P., “A Study of Parallel Disk Organizations” *ACM Computer Architecture News*, Sept. 1989.
- Abraham, J., Banerjee, P., Chen, C.Y., Fuchs, W.K., Kuo, S. Y. and Reddy, A.L.N., “Fault Tolerance Techniques for Systolic Arrays”, *IEEE Computer*, July 1987, vol.20, no.7, pages:65-85.
- Chakraborti, N.B., Soundararajan, J.S. and Reddy, A. L. N., “An Implementation of Mixed-Radix Conversion for Residue Number Applications”, *IEEE Transactions on computers*, August 1986, vol. C-35, no.8, pages:762-764.

Publication of books or book chapters

- Bhandarkar, S.¹, and Reddy, A. L. N., “TCP in Wireless Networks”, invited book chapter in *Encyclopedia on wireless and mobile communications*, CRC Press, Dec. 2007.
- I. Yeom¹ and A. L. Narasimha Reddy, “Differentiated Services”, invited book chapter in *Wiley’s Encyclopedia on Communications*, 2002.
- A. L. Narasimha Reddy and J. Wyllie, “Disk Scheduling in a Multimedia I/O System”, reproduced as a book chapter in *Readings in Multimedia Computing and Networking*, Morgan Kaufmann Publishers, Editors: Kevin Jeffay and HongJiang Zhang, Summer 2001.
- Reddy, A. L. N. and Wijayarathne R., “Scheduling for integrated service in multimedia systems”, invited book chapter in *Multimedia Handbook*, CRC Press, Dec. 1997.
- Reddy, A. L. N. and Haskin R., “Video Servers”, invited book chapter in *Communications Handbook*, CRC Press, 1996.
- Reddy, A. L. N., “Improving the interactive responsiveness in a video server”, invited book chapter in *Multimedia Information Storage and Management*, Kluwer Academic Publishers, Dec. 1996.
- Reddy, A. L. N., “Scheduling in Multimedia Systems”, invited book chapter in *Design and Applications of Multimedia Systems*, Kluwer Academic Publishers, Aug. 1995.

Rigorously Refereed Conference Publications

- Qin, M.1, Reddy, A. L. Narasimha, Gratz, P., Pitchumani, R., Ki, Y-S., “KVRAID: High Performance, Write Efficient, Update Friendly Erasure Coding Scheme for KV-SSDs”, accepted at ACM Systor Conference., June 2021.
- Wen, F., Qin, M.1, Gratz. P., and Reddy, A. L. Narasimha, “OpenMem: Hardware/Software Cooperative Management for Mobile Memory System”, accepted at ACM/IEEE Design Automation Conference, Dec. 2021.
- Wen, F., Qin, M.1, Gratz. P., and Reddy, A. L. Narasimha, “FPGA-based Hybrid Memory Emulation System”, accepted at Field Programmable Logic and applications Conference, Aug. 2021.
- Qin, M.1, Lee. J-H., Pitchumani, R., Ki, Y-S., Reddy, A. L. Narasimha, Gratz, P., “A Generic FPGA Accelerator for Minimum Storage Regenerating Codes”, Proceedings of the 25th Asia and South Pacific Design Automation Conference, January, 2020.
- Chou, C-C.,1, Chen, Y., Milojicic, D., Reddy, A. L. Narasimha, Gratz, P., “Optimizing Post-Copy Live Migration with System-Level Checkpoint Using Fabric-Attached Memory, 2019 IEEE/ACM Workshop on Memory Centric High Performance Computing (MCHPC), Denver, CO, USA, 2019, pp. 16-24.
- Chou, Chih Chieh¹, Jung, J¹, Reddy, A.L. Narasimha, Gratz, Paul V., and Voight, Doug, “vNVML: An Efficient User Space Library for Virtualizing and Sharing Non-Volatile Memories”, IEEE 35th International Conference on Mass Storage Systems and Technology, 2019.
- Zheng, Z.¹, Webb, A.¹, Reddy, A. L. Narasimha, Bettati, Riccardo, “IoTAegis: A Scalable Framework to Secure the Internet of Things”, Invited Paper at IEEE ICCCN, July 2018.
- Zheng, Z.1, Reddy, A. L. Narasimha, “Safeguarding Building Automation Networks: THE-Driven Anomaly Detector Based on Traffic Analysis”, IEEE ICCCN (Invited Paper), July 2017.
- Fedorov, V¹., Kim, J., Qin, M¹., Reddy, A. L. Narasimha and Gratz, P., “Speculative Paging for future NVM and SSD”, Proc. of MEMSYS Conference, Oct. 2017.
- Zhang, Z¹, Jin, S., Bettati, R., Reddy, A. L. Narasimha, “Securing Cyber-Physical Systems with Adaptive Commensurate Response”, Proc. of IEEE CNS, Oct. 2017.
- Zhang, Z¹, and Reddy, A. L. Narasimha, “Safeguarding Building Automation Networks: THE-Driven Anomaly Detector Based on Traffic Analysis”, Invited Paper at ICCCN 2017, July 2017.
- Zhang, Z¹, and Reddy, A. L. Narasimha, “Towards Improving Data Validity of Cyber-Physical Systems through Path-Redundancy”, 3rd ACM Cyber-Physical Systems Security Workshop, April 2017.
- Kim, J., Pugsley, S., Gratz, P., Reddy, A. L. N., Wilkerson, C., and Chisti, Z., “Path Confidence based Lookahead Prefetching”, MICRO Conf., Oct. 2016. **Best Paper Nomination.**
- Webb, A.¹ and Reddy, A. L. Narasimha, “Finding proxy users at the service using anomaly detection”, IEEE Conf. on Communications and Network Security, Oct. 2016.

- Moon, S¹., and Reddy, A. L. Narasimha, “Adaptive policies for balancing performance and lifetime of mixed SSD arrays through workload sampling”, IEEE Mass Storage Symposium, April 2016
- Fedorov, V.¹, Reddy, A. L. Narasimha and Gratz, P., “Shared Last-Level Caches and The Case for Longer Timeslices”, MEMSYS Conference, Oct. 2015
- Kim, J., Fedorov, V.¹, Gratz, P., and Reddy, A. L. Narasimha, “Dynamic Memory Pressure Aware Ballooning”, MEMSYS Conference, Oct. 2015
- Fedorov, V.¹, Qiu, S.¹, Reddy, A. L. Narasimha and Gratz, P., “ARI: Adaptive LLC-memory traffic management”, HiPEAC Conf., Jan. 2014.
- Moon, S¹ and Reddy, A. L. Narasimha, “Don’t let RAID raid the lifetime of your SSD array”, Proc. of USENIX HotStorage Conf., June 2013.
- Qiu, S¹, and Reddy, A. L. Narasimha, “NVMFS: A hybrid file system for improving random write in NAND-flash SSD”, Proc. of Mass Storage Systems Symposium, May 2013.
- Lee, Y-O¹, and Reddy, A. L. Narasimha, “Multipath Routing for Reducing Network Energy”, IEEE Greencom Conf., Sep. 2012.
- Qui, S¹., Reddy, A. L. Narasimha, “Exploiting superpages in a nonvolatile memory file systems”, IEEE Mass Storage Symposium, Apr. 2012.
- Moon, S¹., Reddy, A. L. Narasimha, “Write amplification due to ECC on flash memory or leave those bit errors alone”, IEEE Mass Storage Symposium, Apr. 2012.
- Wu, X¹., Reddy, A. L. Narasimha, “SCMFS: A file system for storage class memory”, ACM/IEEE Supercomputing Conf. SC11, Nov. 2011.
- Yadav S.¹, Reddy, A. L. Narasimha, “Winning with DNS Failures: Strategies for faster botnet detection”, SecureComm Conf., Sep. 2011.
- Lee, Y-O¹, Rengarajan, B., and Reddy, A. L. Narasimha, “Reaching approximate Wardrop equilibrium at reduced cost of link state updates” , COMSNETS Conf., Jan 2011.
- Qian,B¹., and Reddy, A. L. Narasimha, “Measurement and performance study of PERT for on-demand video streaming” , Pfldnet workshop, Nov. 2010.
- Yadav, S¹, Reddy, A. K.¹, Reddy, A. L. Narasimha and Ranjan, S., “Detecting algorithmically generated malicious domain names” , Proc. of Internet Measurement Conference, Nov. 2010.
- Wu, X¹., Reddy, A. L. Narasimha, “Exploiting concurrency to improve latency and throughput in a hybrid storage system”, **Best Paper Nomination**, IEEE MASCOTS Conf., Aug. 2010.
- Devakota, P¹., Reddy, A. L. Narasimha, “Performance of Quantized Congestion Notification in TCP Incast Scenarios of Data centers”, IEEE MASCOTS Conf., Aug. 2010.
- Lee, Y-Oh¹., Reddy, A. L. Narasimha, “Disjoint Multi-Path Routing and Failure Recovery”, Proc. of ICC, May 2010.

- Yadav, S¹., Reddy, A. L. Narasimha, “DNS Anomaly detection with NS and AS information”, IASTED Conf. on Network Security, Nov. 2010.
- Wu, X¹., Reddy, A. L. Narasimha, “Managing storage space in a flash and disk hybrid storage system”, IEEE MASCOTS Conf., Oct. 2009.
- Podduturi, V., Kim, Y., Shakkottai, S., Reddy, A. L. Narasimha, “MultiTrack: A delay and cost aware P2P Overlay architecture”, ACM Sigmetrics, June 2009.
- Kotla, K¹., Reddy, A. L. N., “Making a delay-based protocol adaptive to heterogeneous environments”, Proc. of IwQOS, June 2008.
- Kang, S¹., Reddy, A. L. N., “User-centric data migration in networked storage systems”, Proc. of IPDPS, April 2008.
- Bhandarkar, S¹., Reddy, A. L. N., Zhang, Y., Loguinov, D., “Emulation of AQM at end hosts”, Proc. ACM Sigcomm, Aug. 2007.
- Bhandarkar, S¹., Reddy, A. L. N., “Robustness to packet reordering in high-speed networks”, Proc. of Pfdnet, Feb. 2007.
- Al-Najjar, C¹., Reddy, A. L. N., “A service provider’s approach to improving performance of aggregate voice-over-IP traffic”, Proc. of IWQOS, June 2006.
- Kim, S¹., Reddy, A. L. N., “Netviewer: A Network traffic visualization and analysis tool”, Proc. of USENIX LISA Conference, Dec. 2005, 185-196.
- Liu, Y¹, and Reddy A. L. N., “Routing among a group of multi-homed stub networks”, Proc. of Globecom, Dec. 2005.
- Bhandarkar, S¹., Reddy, A. L. N., “LTCP: A layered approach to TCP for high-speed networks”, PFLDNet, Feb. 2005.
- Kim, S¹., Reddy, A. L. N., “Modeling network traffic as images”, ICC, May 2005.
- Kim, S¹., Reddy, A. L. N., “A study of analyzing network traffic as images in real-time”, IEEE Infocom, March 2005.
- Kim, S¹., Reddy, A. L. N., “Real-time detection and containment of network attacks using QOS regulation”, Proc. of ICC, May 2005.
- Kang, S¹., Reddy, A. L. N., “Virtual Allocation: A scheme for flexible storage allocation” OASIS workshop at ASPLOS, Oct. 2004.
- Liu, Y¹, and Reddy A. L. N., “A fast rerouting scheme for OSPF/IS-IS networks”, Proc. of ICCCN, Oct. 2004.
- Achanta, P.¹, and Reddy A. L. N., “LRU-FQ: Design and Evaluation of a Partial State Router”, Proc. of ICC 2004, June 2004.
- Zhao, Z¹., Reddy, A. L. N., “Impact of bandwidth-delay product and non-responsive flows on queue management schemes”, Proc. of ICC 2004, June 2004.
- Bhandarkar, S¹., Reddy, A. L. N., “TCP-DCR: Making TCP robust to non-congestion events”, Proc. of Networking 2004, May 2004.

- Kim, S¹., Reddy, A. L. N., and Vannucci, M., “Detecting traffic anomalies at source through aggregate analysis of packet header data”, Proc. of Networking 2004, May 2004.
- Kim, S¹., Reddy, A. L. N., and Vannucci, M., “Detecting traffic anomalies using discrete wavelet transform”, Proc. of Int. Conf. on Information Networking, Feb. 2004.
- Ma, X¹., and Reddy, A. L. N., “SPIRAL: A novel third party transfer scheme for network attached devices”, IEEE Mass Storage Symposium, April 2003.
- Zhao, Z¹., Ametha. J., Swaroop. D. and Reddy A. L. N., “A method for estimating non-responsive traffic at a router”, ACM Conf. on Sigmetrics, June 2002.
- Garg, A¹., and Reddy, A. L. N., “Mitigating Denial of service attacks through QOS regulation”, IwQoS workshop, May 2002.
- Wijayarathne, R¹. and Reddy, A. L. N., “System support for providing integrated services from networked multimedia storage servers”, ACM Conf. on Multimedia, October 2001, 270-279.
- Smitha¹ and Reddy, A. L. N., “LRU-RED: An active queue management scheme to contain high bandwidth flows at a congested router”, in IEEE Globecom Conference, Nov. 2001.
- Smitha¹, Kim, I¹., and Reddy, A. L. N., “Identifying long term high rate flows at a router”, in Proc. of High Performance Computing, Dec. 2001.
- Yeom, I¹., and Reddy, A. L. N., “Adaptive marking for aggregated flows”, in Proc. of IEEE Globecom, Nov. 2001.
- Gopalakrishnan, S¹. and Reddy, A. L. N., “SACRIO: An active buffer management scheme for Differentiated Service Networks”, Proc. of NOSSDAV, June 2001.
- Ma, X¹. and Reddy, A. L. N., “MVSS: MultiView Storage Systems”, Proc. of Int. Conf. on Distributed Computer Systems (ICDCS), Apr. 2001.
- Khaleel, A¹. and Reddy, A. L. N., “Evaluation of Request and data distribution policies in a clustered server”, Proc. of High Performance Computing, Dec. 1999.
- Yeom, I¹. and Reddy, A. L. N., “Realizing throughput guarantees in a differentiated services network”, Proc. of IEEE Int. Conf. on Multimedia Computing and Systems, June 1999.
- Wijayarathne, R¹., Reddy, A. L. N., “Integrated QOS management for disk I/O”, Proc. of IEEE Int. Conf. on Multimedia Computing and Systems, June 1999.
- Yeom, I¹. and Reddy, A. L. N., “Impact of marking strategy on aggregated flows in a differentiated-services network”, Proc. of IEEE International Workshop on QOS (IwQOS 99), June 1999.
- Tong, D¹. and Reddy, A. L. N., “QOS Enhancement with Partial Sate”, Proc. of IEEE International Workshop on QOS (IwQOS 99), June 1999.
- Gupta, S¹. and Reddy, A. L. N., “IPRP: A Client-oriented IP Redirection mechanism”, Proc. of IEEE Infocom Conf., March 1999.
- Wijayarathne, R¹. and Reddy, A. L. N., “Techniques for improving the throughput of VBR streams”, Proc. of ACM/SPIE Multimedia Computing and Networking, Jan. 1999.

- Ma, G¹. and Reddy, A. L. N., “An evaluation of storage systems based on network-attached disks”, Proc. of ICPP Conf., Aug. 1998.
- Reddy, A. L. N., “Effectiveness of caching policies for a web server”, Proc. of High Performance Computing, Dec. 1997.
- Reddy, A.L.N., “Caching strategies for a multimedia server”, Proc. of IEEE Int. conf. on Multimedia Computing and Systems, June 1997.
- Reddy, A. L. N., “Improving Latency in an Interactive Video server”, Proc. of SPIE conf. on Multimedia Computing and Networking, Feb. 1997.
- Reddy, A.L.N., “Movie and communication scheduling on a multiprocessor video server system”, Proc. of *IEEE Int. Conf. on Multimedia Comput. and Systems*, May. 1995.
- Reddy, A.L.N., Wyllie, J., “Disk Scheduling in a Multimedia I/O System”, Proc. of *1st ACM Conf. on Multimedia*, Aug. 1993. Included as a chapter in “A guided tour of multimedia systems and applications”, IEEE Press, 1995. Translated into Japanese by *Nikkei Electronics*, a Japanese technical magazine. To be included as “one of the most influential papers” in ACM Transactions on Multimedia, Applications and Communications.
- Chandy, J., and Reddy, A.L.N., “Failure evaluation of disk array organizations”, Proc. of *Int. Conf. on Distributed Computing Systems*, May 1993.
- Reddy, A.L.N. and Freitas, R., “Fault-tolerance of adaptive routing algorithms in multicomputers”, Proc. of *4th IEEE Symp. on Parallel and Distributed Processing*, Dec. 1992.
- Reddy, A.L.N., “A Study of IO System Organizations”, Proc. of *ACM Int. Symp. on Computer Architecture*, May 1992.
- Reddy, A.L.N., “Reads and Writes: When IOs aren’t the same”, Proc. of Hawaii Int. Symp. on Systems and Sciences, Jan. 1992.
- Reddy, A.L.N., Banerjee, P. and Chen, D.K., “Compiler Support for Parallel I/O Operations”, Proc. of *Int. Conf. on Parallel Processing*, Aug. 1991.
- Reddy, A. L. N. and Banerjee, P., “Gracefully Degradable Disk Arrays”, Proc. of *Int. Symp. on Fault Tolerant Computer Systems*, June 1991.
- Reddy, A. L. N. and Banerjee, P., “A Study of IO Behavior of Scientific Applications on A Multiprocessor”, Proc. of *ACM Int. Symp. on Computer Architecture*, June 1990
- Reddy, A. L. N. and Banerjee, P., “Performance evaluation of multiple-disk IO systems”, Proc. of *Int. Conf. on Parallel Processing*, Aug. 1989.
- Reddy, A. L. N. and Banerjee, P., “IO issues in hypercubes”, Proc. of *ACM Int. Conf. on Supercomputing*, June 1989.
- Reddy, A. L. N., Banerjee, P. and Abrham, S.G., “IO embedding in Hypercubes”, Proc. of *Int. Conf. on Parallel Processing*, August 1988.
- Reddy, A. L. N. and Banerjee, P., “A Fault-Secure Dictionary Machine”, Proc. of *3rd International Conference on Data Engineering*, Feb. 1987.

Other Publications

- Kim, J., Gratz, P., and Reddy, A. L. Narasimha, “Lookahead Prefetching with Signature Path”, ACM ISCA Workshop on Data Prefetch Competition, June 2015
- Huang, B¹., Wang, S¹., and Reddy, A. L. Narasima, “Who is Hiring Whom: A New Method in Measuring Graduate Programs”, ASEE conference, June 2015.
- Singh, A.¹., and Reddy, A. L. Narasimha, “Multi path PERT”, **Invited Paper** at Int. Conf. on Computer Communications and Networks, July 2013.
- Amaleswaram, A.¹., Reddy, A. L. N., Yadav, S.¹, Gu. G., Yang, C., “CATS: Characterizing Automation of Twitter Spammers”, **Invited Paper** at Int. COMSNETS conf., Jan. 2013.
- Wu, X.¹., and Reddy, A. L. N., “Data organization in a hybrid storage system”, **Invited Paper** at Int. Conf. on Computing, Networking and Communications, ICNC’12, Feb. 2012.
- Yin, Z.¹., Hasari, C., Al-Nuweiri, H., and Reddy, A. L.N., “Improving the Performance of Delay Based Protocol in Delivering Real Time Media via Early Retransmission”, IEEE ICT Conf., Apr. 2011.
- Yin, Z.¹., Al-Nuweiri, H., and Reddy, A. L.N., “Performance of Early retransmission scheme in streaming media”, IEEE ICT Conf., Apr. 2010.
- Kim, S¹., and Reddy, A. L. N., “An evaluation of the effectiveness of measurement-based anomaly detection techniques”, ADSN workshop, held in conjunction with ICDCS, June 2006.
- Vellanki, S¹., and Reddy, A. L. N., “Improving service availability during link failure transients through alternate routing”, Tech. Report, Texas A & M University, Feb. 2003.
- Ma, X¹., and Reddy, A. L. N., “Implementation and Evaluation of an active storage system prototype”, Workshop on novel uses of System Area Networks, Feb. 2002.
- Reddy, A.L.N., “An evaluation of multiprocessor video server system”, *IBM Confidential report*, June 1994.
- Menon, J., Reddy, A.L.N., Treiber, K. and Tzou, S.Y., “Daisy: A Novel Hierarchical Storage Management System, Functional Specification”, *IBM Confidential Report*, 77 pages, Jan 1994.

Recent Presentations:

- **Tutorial** on Network Security at Int. Conference on Information Security, Hyderabad, India, Dec. 2014.
- **Invited talk** at IEEE Communications Society, Hyderabad, India, August 2014.
- **Invited talk** at University of Texas at San Antonio, Nov. 2013.
- Presentation at Flash Summit, Aug. 2013
- **Invited talk** at Georgia Tech., May 2013.

- **Keynote talk** at Int. Conf. on Information Technology: New Generations, Las Vegas, April 2013.
- **Invited talk** at COMSNETS Conf., Bangalore, India, Jan. 2013.
- **Invited talk** at Int. Conf. on Networking and Communications, Maui, Feb. 2012.
- **Invited talk** at NVRAMOS 2011, Jeju island, South Korea, Nov. 2011.
- **Invited talks** at Korea University, Ajou University, Samsung and ADD, Korea, Nov. 2011.
- **Keynote talk** at IEEE Int. Symp. on Computers and Communications, Italy, June 2010.
- **Catedra Series talk** at University of Carlos III, Dec. 2009, Madrid, Spain.
- Invited talk at Thomson Research, Paris, France, October 2009.
- Invited talk at Telefonica Research, Barcelona, Spain, September 2009.
- IBM Austin Research Lab **Distinguished Seminar series**, May 1 2008.
- Invited talk at University of Illinois, November 2006.
- Invited talk at Cisco, January 2005.
- Invited talk at University of Birmigham, UK, June 2004.
- Invited talk at Cisco Systems, San Jose, June 2002.
- Invited talk at Juniper Networks, Mountain View, June 2002.
- Invited talk at Univ. of California at Riverside, June 2002.
- Invited talk at Nortel Networks, Santa Clara, June 2000.
- Invited talk at Juniper Networks, Mountain View, June 2000.
- Talk at Internet Engineering Task Force, Washington D.C., Oct. 1999.
- Invited talk at EMC Corp., Boston, Aug. 1999.
- Tutorial on “Scheduling for Multimedia Applications”, presented at IEEE Int. Conf. on Multimedia Computing and Systems, June 1999.
- Tutorial on High-Performance I/O at IEEE/ACM Int. Conf on Computer Architecture, Atlanta, May 1999.
- Invited talk at AT & T Research Center, Florham Park, March 1999.
- Invited talk at Bell Labs Research Center, Murray Hill, March 1999.
- Invited talk at EMC Corp., Boston, Nov. 1997.

Funding:

- “Cooprative Research Agreement”, Army Futures Command, Co-PI, 2019-24.

- “Multi-layer Cybersecurity and Situational Awareness to Enhance Resiliency in Qatars Power Grid, Qatar National Research Foundation, 2020-2025, co-PI.
- “Protecting Building Automation Networks from insider threats”, 2016, QNRF, lead PI.
- EMC, DELL, HPE and Samsung membership participation in NSF I/UCRC Center.
- “NSF I/UCRC Center on Intelligent Storage”, 2014, NSF, lead PI.
- “Equipment Donation”, Viking Technologies, 2013.
- “Emerging memory architectures for big memory applications”, National Science Foundation, 2013-2016, co-PI.
- “Exploiting social communication channels against cyber criminals”, National Science Foundation, 2012-2015.
- “Client interfaces to object storage systems”, Argonne National Lab, 2012-2015.
- “Detecting malicious activities by analyzing social communication contexts”, Qatar National Research Foundation, 2012-2015.
- “Network traffic analysis for botnet detection”, Narus Inc., Research gift, 2009-2010.
- “Design and Analysis of new generation of protocols for triple-play networks”, QNRF, Jan. 1. 2008 - Dec. 31 2010, Co-PI.
- “Dynamic allocation and data distribution in networked storage systems”, National Science Foundation, Oct. 2007 - Sept. 2010, sole PI.
- “Active Data Systems”, National Science Foundation, Sept. 2006 -Aug. 2009, sole PI.
- NSF REU grant, co-PI, 2004-2007.
- Network Elements based on partial state, Intel Corp., gift, Feb. 2003.
- Texas A & M University, “High Availability Infrastructure Laboratory”, TITF Award, Sept. 2002 - Dec. 2003, PI with R. Bettati, M. Vannucci and D. Swaroop.
- “Secure Networked Storage Systems”, National Science Foundation, Research Instrumentation grant, Sept. 2002- Sept. 2005.
- Texas Advanced Technology Program, “Network Architectures based on partial state”, Jan. 2001 - Dec. 2002.
- “Fellowships for Computer Engineering”, Department of Education, Aug. 2001- Aug. 2004, PI with R. Mercer and K. Watson.
- “QOS Enhancement with Partial state”, National Science Foundation, June 2001-June 2004, sole PI.
- “Multiview Storage Systems”, National Science Foundation, May 2001 - May 2004, sole PI.
- Equipment Donation, 3Ware Inc., Nov. 2000.

- “Scalable request-specific file service”, National Science Foundation, Aug. 1999 - Aug. 2002, sole PI.
- Research gift, EMC Corp., 1998.
- “Scalable Multimedia Storage Server”, Texas Higher Education Board, January 1996 - Dec. 1997, sole PI.
- “Scalable Storage Server for Multimedia and Scientific Applications”, National Science Foundation, 1996-2000, sole PI.
- “Development of Novel Systems Software for Multimedia and High-performance computing”, National Science Foundation, 1997-2000, PI along with R. Mercer, P. Cantrell, G. Choi and M. Lu.
- Equipment grant, Tandem Corporation, March 1997.
- Software donation, Microsoft Corporation, April 1997.

Teaching:

Redesigned the course ECEN 449 “Microprocessor Systems Design”, Fall 2007.

Developed a new course ELEN 627 “Multimedia Systems and Networks”.

Developed a freshman undergraduate course ENGR111 “Introduction to Electrical and Computer Engineering”.

Developing a new graduate course on “Network systems security”. Taught in Spring 2004, 2006.

Courses Taught: Undergraduate Introduction to Electrical and Computer Engineering, Undergraduate Computer Architecture, Graduate Network Systems security, Graduate Computer Architecture, Graduate Computer Networks, Graduate Multimedia Storage and Delivery, Graduate Parallel Computer Systems, Graduate I/O and Memory Systems.

Student Advising:

- Chris Baldwin, Ph.D. candidate, Aug 2021-, Storage systems.
- Susav Shrestha, Ph.D. candidate, Aug 2021-, Memory systems.
- Chandrahas Tirumalasetty, Ph.D. candidate, Aug 2019-present, Storage Systems.
- Rishabh Singla, Ph.D. candidate, Aug 2017-Present, Network Security.
- Mian Qin, Ph.D. candidate, Aug 2016-Aug 2021, Storage Systems.
- Chih-Chieh Chou, Ph.D., Aug 2014-Jan. 2020, “Optimizing Emerging Memory Systems for Performance”, next at Amazon.
- Zhiyuan Zheng, Ph.D., Aug. 2013 -Dec. 2017, “Protecting Cyber-Physical Systems With Special Emphasis On Building Automation Networks”, next at Pinterest.
- Jaemin Jung, Postdoc, Oct. 2016-18, Storage systems, next at Samsung Research.

- Allen Webb, Ph.D. candidate, Jan. 2013-17, “Addressing Insider Threats From Smart Devices”, next at Google.
- Viacheslav Fedorov, Ph.D., 2010 - 2016, “Memory Management for emerging memory technologies”, later at NXP.
- Jongmin Gim, Postdoc, 2013-2015, Storage Systems Software, Later at Samsung Research Labs, California.
- Sangwhan Moon, Ph.D., 2010 - 2015, “Reliability of SSD Storage Systems”, later at Riverbed Technology.
- Sheng Qui, Ph.D. 2013 “Integration of nonvolatile memory into storage hierarchy”, later at Samsung Research Labs.
- Yong Oh Lee, Ph.D., May 2012, “Improving efficiency and effectiveness of multipath routing in computer networks”, later at Samsung.
- Jagadeesh Dyaberi, Ph.D., Aug. 2011, “Networking and Storage support for on-demand video delivery”, Purdue University, External thesis committee member.
- Sandeep Yadav, Ph.D., 2012, “Scalable techniques for anomaly detection”, Dec. 2012, later at Dumballa.
- Xiaojin Wu, Ph.D., Flash and Disk Hybrid storage systems, Aug. 2011, later at Marvell Semiconductor.
- John Garrison, Ph.D. , “Umbrella File System: Managing heterogeneous storage devices File systems”, Aug. 2009, later at DOD Labs, Tuscaloosa, Alabama.
- Sukwoo Kang, Ph.D., “Flexible space allocation and management in networked storage systems”, May 2007, later at IBM Almaden Research Center.
- Sumitha Bhandarkar, Ph.D., “Congestion control algorithms for TCP in emerging networks”, Aug. 2006, later at Motorola.
- Yong Liu, Ph.D., “Dynamic Routing protocols in multihomed networks”, Dec. 2006, later at EMC Corp.
- Seong Soo Kim, Ph.D., “Real-time analysis of aggregate network traffic”, May 2005, later at Samsung, Seoul, Korea.
- Zhili Zhao, Ph.D., May 2004, “Estimation and control of nonresponsive traffic”, later at Intel.
- Xiaonan Ma, Ph.D., Aug. 2003, “Augmenting functionality at storage systems”, later at IBM Almaden Research Center, San Jose.
- Ravi Wijayarathne, Ph.D., Dec. 2000, “PRISM: A file server architecture for providing integrated services”, later at Quantum, San Jose.
- Ikjun Yeom, Ph.D., Dec. 2000, “Bandwidth assurance in differentiated services networks”, later a faculty member at KAIST, South Korea.
- Daksh Vasistha, M.S., 2015-17, “ Detecting Anomalies in Controller Area Network for Automobiles”, next at Facebook

- Ayush Ruia, M.S., 2015, “Virtualization of nonvolatile memory”, later at Microsoft.
- Bolun Huang, M.S. 2014, “Sentiment based classification of tweeters and university programs”, later at Microsoft.
- Shyam Sambasivan, M.S. “Enhancements to SQLite library to improve performance on mobile platforms”, Aug. 2013, later at Cisco.
- Ankit Singh, M.S., “Multipath Probabilistic Early Response TCP”, Aug. 2012, later at Amazon.
- Cindy Cao, M.S., “SCMFS Performance Enhancement and Implementation on Mobile Platform”, Aug. 2012, later at Microsoft.
- Amit Anand, M.S., “Proactive spammer detection on Twitter”, Aug. 2012, later at Amazon.
- Bin Qian, M.S.. Aug. 2011, “Evaluation of Probabilistic Early Response TCP (PERT) for video streaming and extension with ACK coalescing”, later at Amazon Data Services.
- Ashwath Reddy, M.S., Aug. 2010, Detecting networks employing algorithmically generated names, later at Microsoft.
- Prajjwal Devakota, M.S. Aug. 2010, Performance of Quantized Congestion Notification in TCP Incast in Data Centers, later at Bloomberg.
- Chaitanya Tadi, M.S., May 2009, IP Prefix Hijack Prevention Service, later at Intel.
- Marcus Dunn, M.S., May 2009, A new I/O scheduler for Solid State Devices, later at Cisco.
- Kiran Kotla, M.S., “Adapting the delay-based protocol PERT to heterogeneous environments”, Aug. 2008, later at Cisco.
- Camelia Al-Najjar, M.S., August 2005, thesis on “An approach for improving performance of aggregate voice over IP traffic”, pursuing Ph.D. at Case Western University.
- Saurabh Jain, M.S., May 2005, thesis on “LTCP-RC: RTT compensation technique to scale highspeed protocols to high RTT links”, later at Cisco.
- Sridhar Vellanki, M.S., Feb. 2003, thesis on “Improving service availability through alternate routing”, later at Microsoft.
- Sadry, Nauzad, M.S., Dec. 2002, thesis on “Evaluation of Delayed Congestion Response protocols for tolerating channel errors in wireless networks”, later at Netezza.
- Achanta, Phani, M.S., Dec. 2002, thesis on “Design and Evaluation of Partial State Routers”, later at IBM.
- Smitha, M.S. student, May 2001, thesis on “An active queue management scheme to contain high bandwidth flows at a congested router”, later at Intel India.
- Sumitha Bhandarkar, M.S., August 2001, thesis on “Delayed Congestion Response Protocols”, Ph.D. at TAMU.
- Aman Garg, M.S., August 2001, thesis on “Mitigating Denial of Service using QOS regulation”, later at Tipping Point Technologies, Austin.

- Sai Gopalakrishnan, M.S., May 2000, later at Cisco, thesis on “TRIO: Active resource management in differentiated services networks”.
- Qian Ren, M.S., May 2000, later at Lucent, thesis on “MPEG-aware disk storage system”.
- Marcus Grande, M.S., Dec. 1999, later at IBM, thesis on “Multiported Storage Devices”.
- Adnan Khaleel, M.S., Aug. 1999, later at Motorola, thesis on “Request and data distribution on a clustered server”.
- Deying Tong, M.S., May 1999, later at Cisco, thesis on “Partial state based approaches to QOS enhancement”.
- Sumit Gupta, M.S., Aug. 1998, later at Sun Labs, thesis on “IPRP: A client-oriented IP redirection mechanism and its applications”.
- Ashwin Madhwaraj, M.S., May 1998, later at Cisco, thesis on “Adaptive video transport over networks”.
- Gang Ma, M.S., May 1998, later at EMC Corp., thesis on “Evaluation of storage systems based on network attached storage”.
- Mike Nichols, M.S., May 1998, later at UT Law School, thesis on “Personal informal retrieval agent”.
- Ikjun Yeom, M.S., May 1998, later a faculty member at KAIST, Korea, thesis on “ENDE: An End-to-end Network Delay Emulator”.

Service:

- Member, University IT Research Committee, 2018.
- Chair, Turbo Lab Director Search Committee, 2017.
- Chair, College Cybersecurity Search Committee, 2016-Present.
- Member, University Chief Communications Officer Hiring Committee, 2015-16.
- Member, University level Assoc. Former Students Awards Committee, 2016.
- Member, TAMUQ FASC Committee, 2013.
- Chair, Department Growth Committee, 2013.
- Chair, College Taskforce on National Security, 2013.
- Member, College Research Strategy Subcommittee, 2012.
- Member, Computer Science and Engineering Head Search Committee, 2010-11, 2013-14.
- Member College Awards Committee 2006-2008. 2010-2012.
- Actively participating in improving undergraduate research experience. Co-PI of NSF REU grant. Advised 7 undergraduate students in the last few years.

- Member, Computer Engineering Curriculum Committee, 2003-2012. Led the restructuring of Undergraduate curriculum in Computer Engineering.
- Member, Department Tenure & Promotion Committee, 2004-2006, 2010-2012.
- Member, Electrical and Computer Engineering Department Head Search Committee, 2004.
- Member, Department Advisory committee 2003-2004.
- Chair and Member, department Awards committee 2004-2006.
- Member, Department Awards committee, 1997-2003, 2008-2009.
- Member, Department Distinguished Speakers committee, 1996-1999, 2006-2008.
- Chair and Member, Department Distinguished Speakers committee, 2010-present.
- Member, Department undergraduate studies committee, 1997-2000.
- Member, Department web site design committee, 2000-2004.
- Group leader, Computer Engineering, 2005-2014.

Administrative Experience:

Group Leader, Computer Engineering, 2005-2014:

- Led evaluation, development and rearchitecting the Computer Engineering program, a joint program across two departments of Electrical and Computer Engineering and Computer Science Engineering.
- Led the hiring of about 10 Computer Engineering faculty.
- Led the development of ENGR 111 Freshman Introduction of Electrical and Computer Engineering, primarily to improve retention in Electrical and Computer Engineering programs. Increased retention rates from 45% to about 80%.
- Led the development of several other courses ECEN 449 Microprocessor Systems Design, ECEN 627 Multimedia Systems and Networks, ECEN689 Computer Networks Security.
- Rebuilt several existing courses including ECEN 350 Computer Architecture and ECEN 676 Advanced Computer Architecture, ECEN 602 Computer Networks.
- Contributed to the national ranking of both undergraduate and graduate Computer Engineering programs to top 20 programs in the country.

TEES Assistant Agency Director for National and Global Initiatives, 2014- April 2016:

- Developed and Established a TEES office in Yukatan, Mexico.
- Nurtured relationships in Brazil, France and India. The France relationship resulted in a collaborative center with ENSAM.
- Established realtionships with Kuwait Oil Company which resulted in a \$15 million research contract.

- Developed a strategy to engage with National Labs and Federal Agencies.

**Associate Dean for Research, College of Engineering and TEES Associate Agency
Director for Centers and Strategic Initiatives, April 2016-Present:**

- Established strategic focus on several emerging areas of research including Autonomy, Synthetic Biology, Soft Matter, Cyber Security and Quantum Computing.
- Developed a strategy for supporting large multi-PI proposals and increased large proposal funding from 5-6 per year to about 25 per year.
- Increased TEES Research funding from about \$110 million/year to \$210 million/year from 2016-20.
- Increased Federal funding percentage from 38% to 68% of total TEES funding.
- Nurtured relationships with Army Research Lab (ARL) and facilitated the establishment of ARL-south relationship with A & M. Direct impact: ARL Cooperative Agreements in Synthetic Biology, Electric Power, Materials and Manufacturing, Autonomy and Cybersecurity with about 15 ARL people stationed at Texas A& M.
- Supported establishment of \$65 million ARL/Army Futures Command (AFC) Cooperative agreement on Hypersonics, Autonomous vehicles and Networks.
- Contributed to \$26 million Distributed Autonomy related project funded by Army Research Lab.
- Nurtured and furthered collaborations with Air Force Research Lab, resulting in the establishment of AFRL fellowships at A & M. Direct Impact: AFRL \$4 million Minority Leaders Scholarships in Material Discovery and several projects with sensors, human performance and information directorates.
- Established relationships with Army Engineering and Research Development Center (ERDC). Resulted in \$3 million contract. Exploring further options for joint projects.
- Establishing close collaborative ties with DOE/DOD National Labs. Resulted in Master Research Agreements and joint projects with several labs including Idaho National Lab(INL), Lawrence Livermore National Lab(LLNL), Pacific Northwest National Lab (PNNL) and others.
- Contributed to successful \$40 million proposal for "Offshore Energy Safety Center 2.0" proposal to Department of Interior and Department of Energy.
- Started and organized "TEES Tech. Talk" series to facilitate hosting Labs and Agency researchers and program managers.
- Organized National Labs Day to facilitate interactions of graduate students with National Labs.
- Championed "Remote Ph.D. program for national labs".
- Established "Research Impact Awards" to facilitate a discussion and recognition of impactful research within the College.

- Established External Awards Support office to promote faculty for external professional awards.
- Supported young faculty applications to NSF, DOD/DOE/DARPA CAREER and Young Investigator awards. Increased NSF CAREER awards from 2/year to 8/year.
- Established NSA Masters Program.
- Expanded a weekly "Research Opportunities" bulletin to cover broader research opportunities.
- Interacted with the Research Deans of other colleges and supported inter-college collaborations.
- Established procedures and metrics for evaluating centers within college and reorganized centers after a thorough review of the existing centers. Managed about 30 Centers and Institutes.
- Established support to help graduate student applications for NSF and NDSEG Fellowships.

¹Reddy's Students.