

Then eUNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

LENOVO (UNITED STATES) INC.,

Petitioner

v.

INTELLECTUAL VENTURES II, LLC,

Patent Owner

IPR2024-00124

U.S. Patent No. 7,325,140 B2

**DECLARATION OF DR. GUEVARA NOUBIR
IN SUPPORT OF PATENT OWNER'S RESPONSE**

TABLE OF CONTENTS

I. INTRODUCTION AND SCOPE OF WORK	1
II. QUALIFICATIONS	2
III. RELEVANT LEGAL STANDARDS	5
IV. SPECIFICATION OF THE ‘140 PATENT	8
A. Overview	8
B. Terminology	12
C. Considered Claims of the ‘140 Patent	15
V. LEVEL OF A PERSON HAVING ORDINARY SKILL IN THE ART	16
VI. PRIOR ART RELIED ON IN THE PETITION	17
A. Neufeld	17
B. IPMI	21
C. Lawrence	26
D. Syvanne	26
VII. GROUND 1	27
VIII. GROUND 5	31
IX. GROUND 7	35
X. CONCLUSION	36
XI. DECLARATION	37

I, Dr. Guevara Noubir, hereby declare:

I. INTRODUCTION AND SCOPE OF WORK

1. I am making this declaration at the request of Intellectual Ventures II, LLC (“Intellectual Ventures”) in the matter of the *Inter Partes* Review (“IPR”) proceedings before the United States Patent and Trademark Office (“USPTO”) of U.S. Patent No. 7,325,140 (“the ‘140 Patent”) to provide expert technical opinions with respect to the ‘140 Patent.

2. I am being compensated for my work in this matter at my standard hourly consulting rate for such consulting services. My compensation is not dependent on, and in no way affects, the substance of my statements in this Declaration or the outcome of this case. I do not have any financial interest in the outcome of this proceeding.

3. In preparation of this Declaration, I considered the following materials, the subject matter of which are all within the scope of my education and professional experience:

- a. U.S. Patent No. 7,325,140 (Ex. 1002, “the ‘140 Patent”)
- b. The Petition for *Inter Partes* Review of U.S. Patent No. 7,325,140
- c. Intellectual Ventures II LLC (“Patent Owner”) Preliminary Response
- d. U.S. Patent Appln. Pub. No. 2003/0226015 A1 (Ex. 1005, “Neufeld”)
- e. U.S. Patent No. 7,269,639 (Ex. 1043, “Lawrence”)

- f. IPMI – Intelligent Platform Management Interface Specification V1.5 (Rev. 1.1 dated Feb. 20, 2002) (Ex. 1020, “IPMI”)
- g. U.S. 2003/0097590 A1 (Ex. 1016, “Syvanne”)
- h. Expert Declaration of Dr. Steven Goldberg (filed with the POPR in this matter)

4. I have been asked to review and provide my opinion as to whether Claim 1 of the ‘140 Patent is obvious over U.S. Patent Appln. Pub. No. 2003/0226015 to Neufeld et al. (hereinafter “Neufeld”). In addition, I have been asked to review and provide my opinion whether Claim 1 of the ‘140 Patent is obvious over Intelligent Platform Management Interface Specification V1.5 (Rev. 1.1 dated Feb. 20, 2002) (hereinafter “IPMI”) in combination with U.S. Patent No. 7,269,639 to Lawrence (hereinafter “Lawrence”). Further, I have been asked to review whether Claim 12 of the ‘140 Patent is obvious over Neufeld in combination with U.S. Patent Appln. Pub. No. 2003/0097590 to Syvanne (hereinafter “Syvanne”). Finally, I have been asked to review whether Claim 12 of the ‘140 Patent is obvious over IPMI in combination with Lawrence and Syvanne.

II. QUALIFICATIONS

5. My name is Dr. Guevara Noubir. I am currently employed at Northeastern University as a Professor in the Khoury College of Computer Sciences.

6. Based on my qualifications, education, knowledge, expertise, experience, and work background, I believe I am qualified to offer opinions relating to the technology described in the '140 Patent. Exhibit 2007 is a copy of my current *curriculum vitae* ("CV"), detailing my education and experience. Additionally, the following overview of my background pertains to my qualifications for providing expert testimony in this matter.

7. I hold an M.S. in Computer Science (diplôme d'ingénieur) from the Ecole Nationale Supérieure d'Informatique et de Mathématiques Appliquées de Grenoble (ENSIMAG) and Institut National Polytechnique de Grenoble (INPG), France, with a specialization in networks and real-time systems. I hold a Ph.D. in computer science from the Swiss Federal Institute of Technology in Lausanne (EPFL), Switzerland. One of my primary fields of study in graduate school was algebraic methods for fault tolerance in the context of communications network protocols.

8. For three years, after my PhD, I worked as a senior research scientist in the Real-time Software and Networking Group at the Centre Suisse d'Electronique et de Microtechnique (CSEM) in Switzerland. My responsibilities included coordinating several European projects in wireless communications and network security.

9. I have been a Professor of computer science at Northeastern University since 2001, and a full Professor since 2011.

10. In my first full-time academic position (Assistant Professor) at Northeastern University, I taught in the areas of computer networks and wireless networks. I have taught courses on Internetworking and Wireless Networks since 2001. I introduced and began teaching courses in network security in 2003. In 2005, I began teaching a course in cryptography.

11. I have published and lectured extensively, primarily focusing in the areas of networks and wireless systems security, including new frontiers in networking, robustness and privacy in wireless and mobile systems, secure sharing of location information in wireless networks, preventing hacking of wireless networks, directing wireless security research, game theory in the wireless adversarial context, challenges to wireless security, heterogeneous wireless networks, wireless infrastructure, scalability and security of wireless networks, securing wireless sensor networks, and security challenges of multihop ad hoc networks.

12. I was named an IEEE Senior member in 2001 and have received numerous awards and honors for my research, including the US National Science Foundation CAREER Award in 2005, Google Faculty Research Award on Privacy in 2016, Northeastern University Excellence in Research and Creative Activity Award in 2018, best paper awards at ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec) in 2011 and 2018, and the IEEE Conference on Communications and Network Security best paper in 2016.

13. I led winning teams from Northeastern University in the DARPA Spectrum Collaboration Challenge (SC2) in 2017, 2018, and a finalist team in 2019.

14. I delivered keynote lectures at the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec 2019), and IEEE Conference on Communications and Network Security (CNS 2022).

15. I chaired the technical program committee of several academic conferences including the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec 2015) and IEEE Conference on Communications and Network Security (CNS 2015). I served on the editorial boards of ACM Transaction on Privacy and Security, IEEE Transactions on Mobile Computing, Elsevier Journal on Computer Networks, and IEEE Transaction on Information Forensics and Security.

III. RELEVANT LEGAL STANDARDS

16. The opinions I am expressing in this report involve the application of my knowledge and experience to evaluate certain references with respect to the '140 patent. My formal knowledge of patent law is no different than that of any lay person. Therefore, I have requested the attorneys from Volpe Koenig, who represent Intellectual Ventures, to provide me with guidance as to the applicable patent law in this matter. The paragraphs below express my understanding of how I must apply current principles related to patent validity to my analysis.

17. I understand that an IPR is a proceeding before the USPTO for evaluating the patentability of an issued patent's claims based on prior-art patents and printed publications.

18. I understand that the claims of a patent are reviewed from the point of view of a hypothetical person of ordinary skill in the art ("POSITA") at the time of the filing of the patent. The "art" is the field of technology to which a patent relates.

19. I understand that in determining whether a patent claim is anticipated or obvious in view of the prior art, the Patent Office must apply the *Phillips* standard to construe the claims.

20. I understand that, in this proceeding, Petitioners have the burden of proving that the challenged claims of the '140 Patent are unpatentable by a preponderance of the evidence. I understand that "preponderance of the evidence" means that a fact or conclusion is more likely true than not true.

21. I understand that there are two ways in which prior art may render a patent claim unpatentable. First, the prior art can "anticipate" the claim. Second, the prior art can render the claim "obvious" to a POSITA. I understand that for an invention claimed in a patent to be patentable, it must not be anticipated and must not be obvious based on what was known before the invention was made. For purposes of this Declaration only, I have applied the definition of a POSITA proposed in the Petition.

22. I understand that a patent claim may be invalid as “anticipated” under 35 U.S.C. § 102 if each and every limitation of the claim as construed by the court is found, either expressly or inherently, in a single prior art reference, such as a publication or patent, that predates the claimed subject matter. I also understand that a patent claim may be invalid as obvious under 35 U.S.C. § 103 if the differences between the subject matter claimed and the prior art are such that the claimed subject matter as a whole would have been obvious to a POSITA at the time the invention was made. I also understand that several factual inquiries underlie a determination of obviousness. These inquiries include (1) the scope and content of the prior art, (2) the level of ordinary skill in the field of the invention, (3) the differences between the claimed invention and the prior art, and (4) any objective evidence of non-obviousness (i.e., “secondary considerations”).

23. I also understand that, when a party alleges obviousness based on a combination of references, that party should identify a reason why a POSITA would have been motivated to combine the asserted references in the manner recited in the claims and to explain why a POSITA would have had a reasonable expectation of success in making such combinations.

24. I understand that an obviousness analysis permits the application of “common sense” in examining whether a claimed invention would have been obvious to a POSITA. For example, I understand that combining familiar elements according

to known methods and in a predictable way may suggest obviousness when such a combination would yield nothing more than predictable results. I also understand a party asserting obviousness should still provide a specific motivation to combine the references as recited in the claims and explain why one would have reasonably expected to succeed in doing so.

25. I understand that one indicator of non-obviousness is when a prior art reference leads in a different direction or discourages that combination, recommends steps that would not likely lead to the patent's result, or otherwise indicates that a seemingly inoperative device would be produced.

IV. SPECIFICATION OF THE '140 PATENT

A. Overview

26. The '140 Patent, titled "Secure Management Access Control for Computers, Embedded and Card Embodiment," describes a system for secure remote management of networked devices through a Secure Management Access Controller (SMACC). Ex. 1002, at 12:21. A key innovation is the integration of a Virtual Management Interface (VMI) in the SMACC, which securely virtualizes the user data interfaces, specifically the network interface, to enable simultaneous yet separated handling of user data and management data. Ex. 1002, at 12:17. The SMACC is embedded in a device to be monitored and ensures such separation of the user data

and management data, both within the device to be monitored and during transit. Ex. 1002, at 12:21.




Fig. 1

27. This system allows secure sharing of user data interfaces by introducing a virtual management layer that runs a VPN protocol on the VMI/SMACC, ensuring that management data is encrypted and logically isolated, even when transmitted over the same physical interfaces used for user traffic. Ex. 1002, at 6:5 and 6:34. The management interfaces are protected as the VMI/SMACC run the VPN and firewall, and therefore no user traffic can reach the management interface and all management traffic is encrypted. Ex. 1002, at 12:9. This is a subtle and technical contribution to

the virtualization of user data interfaces to provide secure access and separation of management and user data: “A VMI is also established for logically separating management traffic from user data when the in-band path is to be used for management data. The VMI is the interface between the SMACC chipset and the user data interfaces. The VMI utilizes existing and developing technology such as VPN to build secure tunnels **between the SMACC chipset and the management center** while utilizing the user data interfaces of the managed device and the user network to provide the transport of the management data cost effectively. The VPN technology provides the logical separation, confidentiality, and integrity of the management traffic while it is in transit.” Ex. 1002, at 6:22.




Fig. 3




Fig. 6

28. The ‘140 Patent also supports both in-band and out-of-band management, robust authentication/authorization mechanisms and dynamic switching between available interfaces. Ex. 1002, at 5:23-5:35.

29. Another key function of the ‘140 Patent is to integrate firewall mechanisms directly within the SMACC/VMI to restrict access and protect management operations: “Another primary function of the SMACC is to protect the management interfaces from attack. This is accomplished through a combination of firewall, VPN, and authentication and authorization applications. **The SMACC chipset implements the logic to support VPN tunnels to the management center,** thus protecting the management traffic between the management interface on the SMACC (VMI or SMACC interface) and the management center. The firewall functionality protects the SMACC chipset from access by unauthorized parties, both internal and external, and from unauthorized protocols.” Ex. 1002, at 6:41.

30. Additionally, the SMACC can act as a proxy management interface for co-located legacy devices, and it can be embedded or deployed via upgradeable interface cards. Ex. 1002, at 8:39.

31. Overall, the ‘140 Patent teaches how network devices can be securely and efficiently managed, through a novel approach in the SMACC/VMI of VPN tunneling and firewalling for secure user data interfaces, virtualization, and traffic separation.

B. Terminology

32. A “Private Network” is a type of computer network consisting of nodes (which can be computer hosts, network switches or routers) and links (e.g., ethernet, wi-fi). It is private in the sense that it uses non-routable private IP addresses. Private IP ranges are defined by RFC 1918 (e.g., 10.0.0.0/8, 192.168.0.0/16). Private IP ranges are routable (reachable) from the public internet. Figure 1, below, is an example of a Private Network with private IP addresses. Ex. 2008, at Page 1 (Section 2), and Page 5 (Section 3).




Figure 1. An example of private network with private IP addresses.

33. A “Virtual Private Network” (“VPN”) is a Private Network where some of the links are virtual, meaning they connect two nodes by tunneling traffic over other networks such as the Internet. This is done transparently to the nodes in the VPN. Figure 2, below, is an example of a VPN with two branches, each with private IP

addresses, connected to a VPN tunnel. In the example shown in Figure 2, the two branches are “connected” via a VPN tunnel that terminates at a router on each branch. Ex. 2009, Page 6. A POSITA would understand that this is a typical layout for a general purpose VPN with two branches. This example is explained in more detail below at ¶ 37 *infra*.




Figure 2. An example of virtual private network with two branches (each with private IP addresses) connected by a VPN tunnel.

34. Often, the term “VPN” refers to the networking technology that enables the seamless connection of multiple computers that are not directly connected by a physical network. With this technology, computers on the VPN appear to be on the same network/subnet. Ex. 2009, Page 5.

35. Figure 3 is an example of a Private Network in which a VPN client, hosted on a client computer, initiates a VPN connection and establishes a VPN tunnel that terminates at a VPN Gateway on a network. Ex. 2009, Page 8. In the illustrated example, the VPN gateway is on a distributed network that includes a number of workstations and a printer. The client computer, in this example, could be operated

by a network administrator performing maintenance on one of the workstations. This example is explained in more detail below in paragraph 37. In both cases, the VPN is initiated by a VPN client, hosted on a client computer, and terminates at a point on the network or branch (a router in Figure 2 and a VPN gateway in Figure 3).




Figure 3. An example of private network with a client remotely connecting using a VPN tunnel.

36. There are multiple flavors and protocols to build a VPN. Ex 2009, at Pages 5 and 7. For example, the VPN client software can create a virtual network interface (on the client computer) whose IP address belongs to the VPN subnet. In this example, the VPN client encrypts and encapsulates IP packets sent over this interface (adding an IP header), then routes them over untrusted networks to the VPN server/gateway that decapsulates and decrypts the original IP packet sent over the virtual interface and injects them into the network (Figure 3). Another example is when a company has two branches and makes them appear on the same virtual private

network by using VPN gateways establishing a VPN tunnel between the two branches (Figure 2). Such VPN can be established, for example, using IPSec in tunnel mode.

C. Considered Claims of the ‘140 Patent

37. Claim 1: A remote device management communication system for securely controlling access to management applications and communications to and from said management applications on network devices in a distributed computer network that includes one or more network services, one or more secure management access controllers, and one or more managed network devices, the remote device management system comprising:

at least one secure management access controller connected to one or more data bus of said managed network device for the communication of device management data;

an out-of-band access connection means for connecting said one or more network services or remote users with said secure management access controller for management of said network device; and

at least one **virtual management interface connection means for connecting said one or more network services or remote users with said secure management access controller;**

wherein said virtual management interface connection means **provides logical separation of management data from user data and**

utilizes user interfaces of said managed network element for connecting said one or more network services or remote users with said secure management access controller.

38. Claim 10: The system of claim 1 wherein user and management communications are transmitted to the at least one secure management access controller and the at least one controller further comprises a virtual management interface for separating the user communications from the management communications.

39. Claim 12: The system of claim 10 wherein **the virtual management interface includes a secure transmission tunnel to protect management communications and a firewall** to protect the at least one management access controller from unauthorized access.

V. LEVEL OF A PERSON HAVING ORDINARY SKILL IN THE ART

40. I have been advised that the application that ultimately issued as the ‘140 Patent was filed on October 20, 2004, and issued January 29, 2008.

41. I have been advised that a “person of ordinary skill in the art” (POSITA) is a hypothetical person to whom one could assign a routine task with reasonable confidence that the task would be successfully carried out. I have been advised that the relevant timeframe is prior to October 20, 2004. When I refer to “October 2004,” I am referring specifically to prior to October 20, 2004.

42. In my opinion, a person of ordinary skill in the art in October 2004 would have a bachelor's degree in computer science or electrical engineering and would have one or two years of experience in computer networking or equivalent educational and professional experience. I base this opinion on my experience in the networking industry in general, and with computer networking specifically, in October 2004. At that time, those thinking about or developing improved systems or ways for secure communications networks would have at least this level of experience and knowledge.

VI. PRIOR ART RELIED ON IN THE PETITION

A. Neufeld

43. Neufeld describes a system that customizes user authentication levels and/or access privileges based on various criteria such as connection type (e.g., direct, in-house, or remote) and time of access (e.g., business or non-business hours). Its main focus is on improving system security by implementing more rigorous and adaptive authentication and restricting privileges for remote or less secure access scenarios. It allows user roles and permissions to be customized flexibly, offering varying access to system functions ranging from non-critical to highly critical tasks. It states that this adaptive security model is especially useful for enterprise systems with remote management needs.

44. Neufeld's customized user authentication levels/access privileges are designed for use by a remote administrator attempting access to a managed server that

is part of networks of varying scales: “At one end of the spectrum, a small business or home network may include a few client computers connected to a common server, which may provide a shared printer and/or a shared internet connection. On the other end of the spectrum, a global company's network environment may require interconnection of hundreds or even thousands of computers across large buildings, a campus environment, or even between groups of computers in different cities and countries.” Ex. 1005, ¶ [0008]. It also considers clients remotely connecting through various means. Relative to this, Neufeld describes a possible way for a client to access the network, where the managed server is located, through a VPN or the Internet: “Further, the arrangements of servers and clients in a larger network environment could be connected in any of a large number of topologies that may include local area networks (“LANs”), wide area networks (“WANs”), and municipal area networks (“MANs”). **Clients may also have the ability to access these networks via a virtual private network (“VPN”) or via the Internet.**” Ex. 1005, ¶ [0009] (emphasis added). A POSITA would have understood this type of network access to be work as in the illustration of the VPN in Figure 3.

45. Neufeld also distinguishes between in-band and out-of-band communications. For example, Neufeld states in ¶ [0011]: “Remote management tools can communicate with a managed server using either (1) in-band communication or (2) out-of-band communication. In-band communication refers to communicating

with the server over a standard network connection, such as the managed server's normal Ethernet connection. In-band communication with the server is, accordingly, only possible when the server is able to communicate over its normal network connection. Practically speaking, **this limitation restricts in-band communication to times when the OS of the managed server is operational** (online).” Ex. 1005, ¶ [0011] (emphasis added). Neufeld, therefore, clearly states that in-band communication is limited to use when the OS of the managed server is running. It does not, however, discuss how in-band communication could work in Neufeld's system as Neufeld's invention is limited to discussions of system operation when the OS is not running (offline).

46. Neufeld also clearly states that it considers out-of-band communication as a powerful tool for server management: “Out-of-band communication, which is not performed across the managed server's normal connection to the network, is a much more powerful tool for server management.” Ex. 1005, ¶ [0012]. Consistent with this, Neufeld states: “The Ethernet interface 180 provides **the main communication interface** between the remote server management controller 116 and the outside world.” Ex. 1005, ¶ [0061] (emphasis added). In FIG. 3 of Neufeld, the Ethernet interface 180 is coupled to the Network M – the management Network M is, in fact, the only Network coupled to Neufeld's remote server management controller 116. Neufeld in fact directly states this: “[T]his connection is typically used to couple the

remote management controller 116 to the management network M.” Ex. 1005, ¶ [0061].

47. At first glance, paragraph [0043] of Neufeld may appear to suggest a contradiction by stating that “the remote management controller 116 connects to . . . the network N and/or the management network M . . . to provide functionality for accessing, interacting, and monitoring the managed server 2 from the remote console.” However, upon closer analysis, a POSITA would understand, particularly in light of FIGs. 2 and 3, that the remote management controller 116 is physically connected to internal system components, such as the PCI bus and keyboard controller, enabling it to snoop on system activity—such as capturing display output and keyboard inputs—without relying on or interfering with the server’s host operating system. This architecture underscores that the mention of network N or the management network M in paragraph [0043] does not imply that the remote management controller transmits management data through the NIC 110 or the server’s primary network path. To the contrary, Neufeld expressly states that the remote management controller communicates externally through a dedicated interface, such as Ethernet interface 180, and that the entire system is purposefully designed to operate independently of any software on the managed server. Thus, a POSITA would recognize that the reference to network N merely identifies that the remote management controller can **passively monitor** the data sent to or received

from the managed device over the network; it must be connected to that path to snoop on the traffic, but that connection does not equate to active participation in the data transmission or use of the primary network interface for management communication.

B. IPMI

48. IPMI (Intelligent Platform Management Interface) is a standardized interface designed for managing and monitoring computer systems independently of the operating system or system state. The IPMI v1.5 specification describes a set of standardized interfaces that enable out-of-band monitoring and control of system hardware, such as temperature sensors, power supplies, and fan speeds. It also allows for logging system events and issuing alerts. This functionality is accessible via multiple channels like LAN, serial/modem connections, or local interfaces and is managed through a component known as the Baseboard Management Controller (BMC). The architecture enables remote system diagnosis and recovery. IPMI Figure 1-2 provides a block diagram of the IPMI system and is reproduced below as Figure 4.

Figure 1-2, IPMI Block Diagram




Figure 4 IPMI Block Diagram [from IPMI version 1.5 interface specification].

49. The IPMI specification describes the out-of-band IPMI LAN interface mechanisms for transferring IPMI messages between the BMC and a remote management system (remote console) over an Ethernet LAN connection using UDP over IPv4: “There are two implementations that are likely to be used to deploy an IPMI LAN Interface using a shared LAN controller. The first implementation is using an embedded LAN controller, as shown in Figure 12-1 (reproduced below as Figure 5), and the second is using a LAN controller on an add-in card, as shown in Figure 12-2 (reproduced below as Figure 6). **Both examples show a LAN Controller that has the capability to detect UDP datagrams sent to a “management port”. Any datagrams received on that port are forwarded to a “side-band” interface that**

allows them to be delivered to, or retrieved by, the BMC. As Figure 12-1 shows, these incoming “platform management” datagrams may also be **delivered to system software in parallel with being delivered to the BMC**. The BMC **can use this same interface to inject datagrams onto the LAN**. These datagrams **are interleaved with the network packets that are generated by system software.**” It is important to note that the LAN Controller/”side-band” only allows the detection and forwarding of a subset of IP/UDP packets to the BMC and injection of BMC packets onto the LAN. LAN Controllers are not capable of running sophisticated software such as a VPN (e.g., IPSec).




Figure 12-2, PCI Management Bus Implementation




Figure 6 IPMI with LAN Controller implemented as a PCI add-in card connected to the BMC via a PCI Management Bus connection [from IPMI version 1.5 interface specification].

50. IPMI also describes ‘in-band’ Remote Access. In-band IPMI management occurs through the host operating system, using software agents that communicate with the Baseboard Management Controller (BMC) over local system interfaces. These are hardware-level communication channels between the CPU and the BMC that allow IPMI commands to be issued from within the operating system. “The BMC manages the interface between system management software and the platform management hardware, provides autonomous monitoring, event logging, and recovery control, and serves as the gateway between system management software and the IPMB and ICMB.” Ex. 1020, page 31, § 1.6.3. Further, “IPMI is best used in conjunction with **system management software running under the operating system**. This provides an enhanced level of manageability by **providing in-band access** to the IPMI management information and integrating IPMI with the additional

management functions provided by management applications and the OS. System management software and the OS can provide a more sophisticated control, error handling and alerting, than can be directly provided by the platform management subsystem.”. Ex. 1020, page 30, § 1020. It is important to note that for the in-band mode, **IPMI requires that the OS is operating and running the management application**. Furthermore, as shown in Figure 4, **communication over the system interface between the system management software and BMC uses IPMI Message format and not Internet Protocol (IP) packets**. Figure 1-1 of IPMI shows the system management software running under the operating system and is reproduced below as Figure 7.

Figure 1-1, IPMI and the Management Software Stack




Figure 7 IPMI in-band remote access system is in conjunction with system management software running under the operating system and via a PCI Management Bus connection [from IPMI version 1.5 interface specification]

C. Lawrence

51. Lawrence recites a method for managing a network device over a VPN. It recognizes that network device management is typically performed either out-of-band (using a separate network) or in-band (using the main network, but often without adequate security). It creates an in-band management system by “**dedicating one of the VPN data links** to the transfer of management commands; linking the dedicated VPN data link to a source of management commands; using the VPN to transfer management commands from the source of the management commands to the network device; and **transferring management commands** between the dedicated VPN data link and a router configuration management module.” Ex. 1043, Claim 1 (emphasis added). Lawrence uses a **dedicated** VPN link for management traffic not carrying user data. Furthermore, it processes the IP packets to extract and transfer the management commands to the router configuration management module.

D. Syvanne

52. Syvanne recites a mechanism to enhance the security of portable computing devices, such as laptops, by enabling a local/personal firewall to automatically adapt its behavior based on the device’s network location. The patent recites how to equip the firewall with distinct sets of security rules for different environments, such as home and foreign networks. The **firewall operates on Internet Protocol (IP) packets** and detects its current location using the IP address. It can

verify its current location by communicating with known nodes in the considered network, and activates the appropriate rule set without user intervention. It also includes mechanisms for centralized management of firewall rules and log files, ensuring automatic updates and centralized logging when the device reconnects to a trusted (home) network. The goal of this adaptive security approach is to provide a stronger protection when the portable computing device is in hostile environments without unnecessarily restricting functionality in secure ones.

VII. GROUND 1

53. In Ground 1, the Petitioner asserts that Claims 1-4, 6, 7, 10, 11, 13 and 16 are obvious under 35 U.S.C. § 103 over Neufeld. In my opinion, the asserted ground is deficient for several technical reasons.

54. Neufeld neither discloses, describes, nor suggests a “virtual management interface connection means,” as required by the challenged claims. Specifically, Neufeld fails to teach or suggest a mechanism that both (1) provides **logical separation of management data from user data**, and (2) **utilizes user interfaces of the managed network element to connect** one or more network services or remote users **with a secure management access controller (SMACC)**.

55. In support of the Claim feature requiring “**logical separation of management data from user data**,” the ‘140 Patent explicitly describes a SMACC chipset that has the logic to support VPN tunnels, or, more specifically, secure

transmission tunnels to protect management communications. Ex. 1002, at 6:38. In the '140 Patent, the SMACC can also include a firewall to protect the SMACC from unauthorized access. Ex. 1002, at 6:36. This could mean, for example, that the SMACC can host a VPN client that establishes VPN tunnels or that VPN tunnels established by a remote administrator computer can terminate at the SMACC. The firewall provides security against attacks on the SMACC, providing an extra level of security for the management data before the data is sent via the user interfaces.

56. The '140 Patent further explains how the virtual management interface (VMI) integrates with the SMACC to securely virtualize the user data interfaces, thereby allowing both management data and user data to be processed concurrently over the same physical hardware while remaining logically segregated. Ex. 1002, at 4:20-31. This virtualization architecture enables the system to treat the same network interface as multiple distinct logical interfaces, each governed by separate security and access rules, enforced by the SMACC. Ex. 1002, at 11:64-12:9; Ex 2010 at 11:64-12:9. This is consistent with virtualization concepts, including virtual network interfaces. Ex. 2010 at §§2.3 and 3.3.5. The virtualized user interfaces, in combination with the logic to support VPN tunnels, enable the SMACC to separate management data from user data, in locations where they are co-mingled within the device, and kept securely separated during transit to a remote administrator computer. Ex. 1002, at 3:13 and 4:20.

57. Neufeld's architecture, by contrast, lacks any discussion of virtualized user interfaces. Moreover, while Neufeld discloses a remote management controller in a CPU, the remote management controller lacks any logic to support VPN tunnels and has no capability to facilitate simultaneous and logically separated data handling. Instead, Neufeld only marginally mentions a general purpose VPN or Virtual Private Network as a way for a client to reach the private network in which the managed device is located (as illustrated in Figure 3). Ex. 1005, ¶ [0009] ("Clients may also have the ability to access these networks via a virtual private network ("VPN") or via the Internet.) This VPN is not hosted on Neufeld's remote management controller and does not terminate at the remote management controller or the NIC. In contrast, a general purpose VPN used for network access terminates at a point on the network, such as a router or VPN gateway, rather than at a point within the managed device. (See paragraphs 34, 36, 37, FIG. 2, FIG. 3 *supra*).

58. Accordingly, it is my opinion that Neufeld does not disclose, describe or suggest a virtual management interface that **utilizes user interfaces of the managed network element to connect** one or more network services or remote users **with a secure management access controller (SMACC)**, as required by Claim 1 of the '140 Patent (emphasis added).

59. Further, the remote management approach described in Neufeld is not compatible with in-band communication (as defined in Neufeld). More specifically,

Neufeld is focused on solutions for remote server management when the managed server is powered down (e.g., for a reboot), before the OS is up and running (e.g., for initial setup), or while the OS is offline. Ex. 1005, ¶ [0034]. However, Neufeld also acknowledges that in-band communication is not an option without the OS. Ex. 1005, ¶ [0011] (“Practically speaking, this limitation restricts in-band communication to times when the OS of the managed server is operational.”). Since the proposed solution in Neufeld operates before the OS is loaded, it cannot be using in-band communication, which is only usable when the OS of the managed server is operable.

60. Accordingly, in Neufeld, management data and user data are not co-mingled on the in-band path. This means that a VPN used for network access, as mentioned in Neufeld, would not inherently, or obviously, separate management data from user data, as required by Claim 1, because there is no co-mingled data to separate. And even if this were the case, which it is not, because this VPN terminates at a point on the network, and not within the device, even if user data and management data were co-mingled at another point in the system, such as the NIC, the VPN would not be able to separate it because it operates **outside of the managed server**.

61. For the reasons provided in ¶¶ 55-61 *supra*, it is my opinion that Neufeld does not teach, suggest, or even contemplate the use of a virtualized management interface that leverages the **user interfaces of the managed device** to provide **logically separated, simultaneous data channels** under control of a **secure**

management access controller. As a result, a POSITA would not be motivated to modify Neufeld in any way that would yield the claimed architecture. The virtual management interface connection means disclosed in the ‘140 Patent represents a fundamental departure from Neufeld’s architecture, offering a solution to problems that Neufeld does not identify and cannot address using its out-of-band model.

VIII. GROUND 5

62. In Ground 5, the Petitioner asserts that Claims 1-3, 6, 7, 10 and 11 are obvious under 35 U.S.C. § 103 over IPMI in combination with Lawrence. In my opinion, the asserted ground is deficient for several technical reasons.

63. Neufeld neither discloses, describes, nor suggests a “virtual management interface connection means,” as required by the challenged claims. Specifically, Neufeld fails to teach or suggest a mechanism that both (1) provides **logical separation of management data from user data**, and (2) **utilizes user interfaces of the managed network element to connect** one or more network services or remote users **with a secure management access controller (SMACC).**

64. I first note that the terms “in-band” and “out-of-band” are described differently in the ‘140 Patent versus IPMI. In IPMI, “in-band” remote access management occurs through the host operating system, using software agents that communicate with the Baseboard Management Controller (BMC) over local system interfaces. These are hardware-level communication channels between the CPU and

the BMC that allow IPMI commands to be issued from within the operating system. For the “in-band” mode, **IPMI requires the OS to be operating and running the management application.** Furthermore, as shown in Figure 4 above, **communication over the system interface between the system management software and the BMC uses IPMI Messages format and not Internet Protocol (IP) packets.** Accordingly, while this path is referred to as “in-band,” it cannot be said that communication of these IPMI messages “utilize[] user interfaces of said managed network element for connecting said one or more network services or remote users with said secure management access controller,” as recited in Claim 1 of the ‘140 Patent. Lawrence does not cure this deficiency of IPMI.

65. Furthermore, within the context of IPMI “in-band” remote access, there was no motivation nor practical way of integrating a VPN tunnel on the BMC. **A VPN tunnel carries Internet Protocol packets. On the other hand, the BMC communication with the remote client packets first goes through the operating system, to the system management software, which converts the application layer messages into the IPMI message format defined by the IPMI specification.** The messages received by the BMC over the system interface are not Internet Protocol packets and, therefore, any VPN tunnel would need to be terminated at the OS level and cannot be forwarded to the BMC.

66. IPMI also describes and illustrates “out-of-band” LAN interface mechanisms for transferring IPMI messages between the BMC and a remote management system (remote console) over an Ethernet LAN connection using a side-band interface. The Petition cites this “out-of-band” mode in IPMI, in combination with the VPN of Lawrence, against Claim 1. However, neither of these references teaches setup of a VPN tunnel at a LAN controller (such as the LAN controller taught in IPMI).

67. With respect to this feature of Claim 1, the ‘140 Patent explicitly teaches: “In a preferred exemplary embodiment of the current invention, **the VMI will utilize a Virtual Private Network (VPN) mechanism such as IPSec to encrypt, protect and logically separate the management traffic** before the traffic is passed to the user interfaces.” In contrast, a LAN controller, which the Petition cites to IPMI for, only allows the detection and forwarding of a subset of IP/UDP packets to the BMC and injection of BMC packets onto the LAN. LAN Controllers have very limited capabilities and are not capable of running sophisticated software such as a VPN (e.g., IPSec). A VPN server using IPSec is complex and involves resource heavy cryptographic primitives as well as complex manipulation of the Internet Protocol stack. This complexity far exceeds the capabilities of LAN controllers, and this was especially true in October 2004. A POSITA would not have thought or considered running a VPN server on a LAN controller as a realistic approach.

68. Specifically, VPN protocols (e.g., IPsec as described in the ‘140 Patent) require session management, cryptographic operations (e.g., AES, SHA, DH/ECDH), and packet inspection for access control. These tasks demand substantial processing power, memory, and secure key storage. In contrast, LAN controllers in October 2004 (and even today) were typically implemented as fixed-function hardware devices or microcontrollers with limited or no programmable execution environment, lacking the necessary support needed to run such functionality. Furthermore, the state-of-the-art placed VPN and firewall logic at higher levels in the network stack (e.g., on the host CPU). There is no teaching, suggestion, or motivation in the prior art that would lead a POSITA to offload VPN or firewall responsibilities to the LAN controller.

69. Similarly, adding a VPN module, as described by Lawrence, but also enabling capabilities, such as described in ‘140 Patent, [t]he SMACC can and should be configured to perform authentication and authorization of the remote administrator before allowing access to the management functions for the managed device. While the SMACC could be configured with a database of authorized administrators, it is also possible to utilize an ACS to authenticate and authorize a remote administrator. The SMACC utilizes its VMI to access network services such as an ACS in the network using a protocol such as RADIUS, TACACS+, or LDAP. This would allow centralized authentication and authorization of the administrators.”. Doing so would

have contradicted prevailing IPMI design conventions and introduced architectural complexities without recognized benefit.

IX. GROUND 7

70. In Ground 7, the Petitioner asserts that Claims 11 and 12 are obvious under 35 U.S.C. § 103 over IPMI in combination with Lawrence and Syvanne. Relative to Claim 12, which depends from both Claims 1 and 10, Ground 7 is deficient for the technical reasons provided below.

71. Within the context of IPMI (October 2004), there was no motivation, nor any practical way, of integrating a VPN and firewall on the LAN Controller. **Syvanne recites firewall methods that operate on Internet Protocol (IP) packets. On the other hand, the BMC communication with the remote client packets first goes through the operating system, to the system management software which converts the application layer messages into the IPMI message format defined by the IPMI specification.** The messages received by the BMC over the in-band system interface are not Internet Protocol packets, therefore any firewall or VPN tunnel would need to be terminated at the OS/application level.

72. On the out-of-band LAN interface, and as discussed above, a POSITA would not envision running a VPN tunnel on a LAN Controller. LAN Controllers are not capable of running sophisticated software such as a VPN (e.g., IPSec). At the time the application that issued into the '140 Patent was filed, a POSITA would not have

found it obvious, nor reasonably contemplated, to implement a VPN tunnel and firewall functionality on a LAN controller. LAN controllers at that time were designed and optimized solely for packet transmission and basic network operations, with no architectural provisions for hosting or executing complex protocol stacks such as those required for VPN termination and firewall enforcement.

73. Similarly, incorporating a VPN module, as described by Lawrence, would have contradicted prevailing IPMI design conventions and introduced architectural complexities without recognized benefit.

X. CONCLUSION

74. For the foregoing reasons, it is my technical opinion that Claims 1 and 12 of the '140 Patent would not have been obvious over the cited references and that the Petition fails to demonstrate a reasonable likelihood of prevailing on at least Grounds 1, 5 and 7 of the Petition.

XI. DECLARATION

75. I, Dr. Guevara Noubir, declare that all statements made herein on my own knowledge are true, and all statements made on information and belief are believed to be true. Further, I am aware that these statements are made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under 18 U.S.C. § 1001.h. I declare under penalty of perjury that the foregoing is true and correct.

Executed on May 14, 2025, in Boston, Massachusetts.



Dr. Guevara Noubir