UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

LENOVO (UNITED STATES) INC.,

Petitioner,

v.

INTELLECTUAL VENTURES II, LLC,

Patent Owner

Case No. IPR2024-00124
U.S. Patent No. 7,325,140

PATENT OWNER'S RESPONSE

TABLE OF CONTENTS

PATENT OWNER'S EXHIBIT LIST iv
I. INTRODUCTION 1
II. '140 PATENT OVERVIEW 2
A. Tradeoffs Between In-Band and Out-of-Band Network Access 4
B. Secure Management Access Control for Computer Chipset (SMACC) 5
1. Interfaces Available at the SMACC 7
2. SMACC Functionality 9
C. Summary of the '140 Patent 10
III. PRIOR ART RELIED ON IN THE PETITION 12
A. Neufeld 12
1. Remote Management System 14
2. Remote Management Controller 15
3. Security From External Attacks 16
B. Intelligent Platform Management Interface Specification V1.5 (IPMI) 18
C. Lawrence 21
IV. THE PETITION FAILS TO DEMONSTRATE A REASONABLE LIKELIHOOD OF SUCCESS ON THE MERITS 22
A. Ground 1 – Claims 1-4, 6, 7, 10, 11, 13 and 16 Are Not Obvious Over Neufeld Pursuant To 35 U.S.C. § 103 22
1. The Petition Fails to Establish that Neufeld Discloses or Renders Obvious the "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" Element of Claim 1 23
2. The Petition Fails to Show that Neufeld Discloses or Renders Obvious the Limitation of Claim 1 Requiring Virtual Management Interface Connection Means to Provide Logical Separation of Management Data from User Data 27
3. Petitioner's Legal Reasoning is Deficient Under Federal Circuit Precedent 35
4. Conclusion 37

B. Ground 2 – Dependent Claims 11 and 12 are Not Obvious over Neufeld and Syvanne 38
C. Ground 3 - Dependent Claim 14 is Not Obvious Over Neufeld and Fowler Pursuant To 35 U.S.C. § 103 39
D. Ground 4 - Dependent Claim 13 is Not Obvious Over Neufeld and Fowler Pursuant To 35 U.S.C. § 103 39
E. Ground 5 - Claims 1-3, 6, 7, 10 And 11 Are Not Obvious Over IPMI and Lawrence Pursuant To 35 U.S.C. § 103.....39
1. The Claim Feature "At Least One Virtual Interface Connection Means [That] Provides Logical Separation Of Management Data From User Data" Would Not Be Obvious From IPMI And Lawrence.....40
a) IPMI's LAN Controller Does Not Separate Data, But Rather Detects Packets Addressed to the BMC.....41
b) Petitioner Has Not Established that It Would Have Been Obvious to Modify IPMI, With the Teachings of Lawrence, to Arrive at the Claimed Invention Nor Would the Combination Have Been Obvious to a POSITA 42
F. Ground 6 - Dependent Claim 4 is Not Obvious Over IPMI, Lawrence and Leedy Pursuant To 35 U.S.C. § 103 47
G. Ground 7 - Dependent Claims 11 And 12 Are Not Obvious Over IPMI, Lawrence And Syvanne Pursuant To 35 U.S.C. § 103.....48
H. Ground 8 - Dependent Claims 13 and 14 Are Not Obvious Over IPMI, Lawrence and Fowler Pursuant To 35 U.S.C. § 103 49
I. Dependent Claim 16 is Not Obvious Over IPMI, Lawrence and RFC1492 Pursuant To 35 U.S.C. § 103 49
V. PETITIONER HAS NOT MET ITS BURDEN OF PROVING UNPATENTABILITY OF THE CHALLENGED CLAIMS 50
VI. CONCLUSION.....51

PATENT OWNER'S EXHIBIT LIST

Ex. 2001	Declaration of Dr. Steven Goldberg
Ex. 2002	Resume of Dr. Steven Goldberg
Ex. 2003	Virtual private networks, Dennis Fowler, Morgan Kaufmann Publishers, Inc., 1999
Ex. 2004	Virtual Private Networks, Charlie Scott, Paul Wolfe, Mike Erwin, O'Reilly, 1999, Page 8
Ex. 2005	Networking Explained, Second Edition, Michael Gallo and William Hancock, Digital Press, 2002, Page 188.
Ex. 2006	Declaration of Dr. Guevara Noubir
Ex. 2007	Resume of Dr. Guevara Noubir
Ex. 2008	Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, E. Lear, "Address Allocation for Private Internets", Internet Engineering Task Force (IETF), Network Working Group, February 1996. doi:10.17487/RFC1918.
Ex. 2009	Vijay Bollapragada, Mohamed Khalid, Scott Wainner, "IPSec VPN Design", Cisco Press, 2005.
Ex. 2010	Paul Barham, Boris Dragovic, Keir Fraser, Steven Hand, Tim Harris, Alex Ho, Rolf Neugebauer, Ian Pratt, and Andrew Warfield, "Xen and the art of virtualization", In Proceedings of the nineteenth ACM symposium on Operating systems principles (SOSP '03, 2003. https://doi.org/10.1145/945445.945462
Ex. 2011	Mark Thorson, "System Management Mode Explained", Microprocessor Report 6 (1992): 14-16.

I. INTRODUCTION

Petitioner challenges claims 1-4, 6, 7, 10-14, and 16 ("Challenged Claims") of the US 7,325,140 (" '140 Patent") on nine grounds. The '140 Patent has 16 total claims, with only claim 1 being independent. In Ground 1, Petitioner challenges independent claim 1, and certain dependent claims, as allegedly obvious over US Patent Appln. Pub. No. 2003/0226015 to Neufeld et al. (hereinafter "Neufeld"). In Ground 5, Petitioner challenges independent claim 1, and certain dependent claims, as allegedly obvious over Intelligent Platform Management Interface Specification V 1.5 (hereinafter, "IPMI") and US Patent No. 7,269,639 to Lawrence (hereinafter "Lawrence"). The remainder of the grounds address dependent claims only. Dependent claims 5, 8, 9, and 15 are not challenged.

In Ground 1, Petitioner contends that claim 1 is obvious over Neufeld. This argument is, however, fundamentally flawed due to a mischaracterization of Neufeld's disclosure. Petitioner argues that a VPN, used only to connect to a network on which a managed server resides, inherently separates management data from user data because management data and user data are co-mingled on the in-band path; however, in Neufeld, only the out-of-band path is actively used for server management. Accordingly, this argument fails as its inherency argument is based on an incorrect assumption about how Neufeld's system works. Moreover, Petitioner's position relies almost exclusively on the conclusory and unsupported

opinions of its expert, all of which are grounded in this fundamentally incorrect understanding of Neufeld's disclosure.

In Ground 5, Petitioner alleges that Claim 1 is obvious over IPMI, a technical specification, in view of Lawrence, an issued U.S. Patent. In its argument, the Petition alleges that a LAN controller, which is used in the technical specification to control access to a side-band connection/system bus, is actually separating management data from user data and that a POSITA would have looked to Lawrence for implementation details. This is a complete mischaracterization of the reference. The only place separation of management data from user data is explicitly mentioned in any art mentioned in the Petition is in the specification of the '140 Patent itself.

For at least the reasons set forth above, the Petitioner has not met its burden of producing evidence in the Petition to support a conclusion of unpatentability and the Petition should, therefore, be denied.

II. '140 PATENT OVERVIEW

The '140 Patent is generally directed toward remote management of network elements in distributed computer networks. *See* Ex. 1002, 1:19-23. Distributed computer networks are computer networks in which at least some, and usually most, of the network elements are not in the same geographic location. *See id* at 1:27-31. This means that network administrators responsible for network maintenance are typically located some geographical distance away from the network elements they

are responsible for maintaining. *See id.* Network administrators, however, still need to be able to quickly access network elements when they are not working properly in order to minimize disruption to network users. *See id.* at 1:31-38. This limits network administrators to *remotely* accessing network elements (in most instances) to make repairs (rather than requiring network administrators to physically access them). *See id.* at 1:39-45. As the '140 Patent explains, at the time the patent application was filed, businesses were commonly using at least some form of remote management to maintain and troubleshoot network elements and servers. *See id.* at 1:43-45.

The '140 Patent describes two general approaches to remote network management that were commonly used at the time the application that issued as the '140 Patent was filed: in-band network access and out-of-band network access. *See id.* at 1:45-52; Ex. 2006, ¶ 28. In-band network access is a type of network access in which a remote administrator "communicates with the system using *the same network as the user data for the managed system*." Ex. 1002. at 1:47-49 (emphasis added). Out-of-band network access, on the other hand, is a type of network access in which a remote administrator "communicates with the system *using a means other than the network utilized by the user data of the managed system*." *Id.* at 1:50-52 (emphasis added).

A. Tradeoffs Between In-Band and Out-of-Band Network Access

In-band network access is generally less expensive than out-of-band network access at least because in-band access uses network elements that are otherwise already being used to communicate user data. *See id.* at 1:47-50. In order to minimize costs, the '140 Patent proposes to implement in-band network access in a manner in which user traffic and management traffic would otherwise be "co-mingled at multiple locations throughout the management path" – including both *within the device itself* – and during transport. *See id.* at 18. This cost saving approach, however, makes securing the management data difficult. *See id.* at 3:21-23. "In-band management methods often depend . . . on the security of the network the element is a part of to protect the management traffic," which may "provide a sufficient level of protection" from attacks initiated outside of the network but "will not provide a sufficient level of protection" from attacks initiated from inside the network. *Id.* at 2:12-19.

Regardless of the expense, out-of-band network access is typically still needed for remote network maintenance at least because in-band access is not available when the in-band network is not working and upon initial configuration of the managed device when the operating system (OS) is not yet up and running. *See id.* at 2:30-42. Out-of-band access also provides a generally more secure means of network connection and, therefore, may be exclusively used for performing certain

maintenance functions involving highly sensitive communications, such as password recovery. *See id.* at 2:42-45.

Like in-band network access, implementation of out-of-band network access requires some tradeoff between expense and security. The most secure means of implementing out-of-band network access is to build a dedicated network for out-of-band access, but this approach is cost prohibitive. *See id.* at 2:17-25, 45-48. Some form of public shared network would be the least costly means of implementing out-of-band network access, but security of public shared networks is a major concern and some potential solutions undesirably by-pass perimeter security, such as firewalls and access lists. *See id.* at 2:48-63.

B. Secure Management Access Control for Computer Chipset (SMACC)

The '140 Patent addresses the aforementioned cost and security problems by embedding a Secure Management Access Control for Computer Chipset (SMACC) in the managed device. *See id.* at 4:7-11, 5:50-52, FIG. 1. "A primary function of the SMACC is to provide for the separation of management data from user data *both within the device being managed and while the management information is in transit*." *Id.* at 6:5-8 (emphasis added). The SMACC is illustrated in FIG. 3 of the '140 Patent, which is reproduced below for convenience. In an embodiment, the SMACC "connects to the managed system via a Peripheral Component Interconnect

(PCI) bus," which it may use to communicate device management data with the managed system/device. *Id.* at 9:15-19; *see id.* at 6:8-14, 10:65-11:8, FIGs. 28-29.

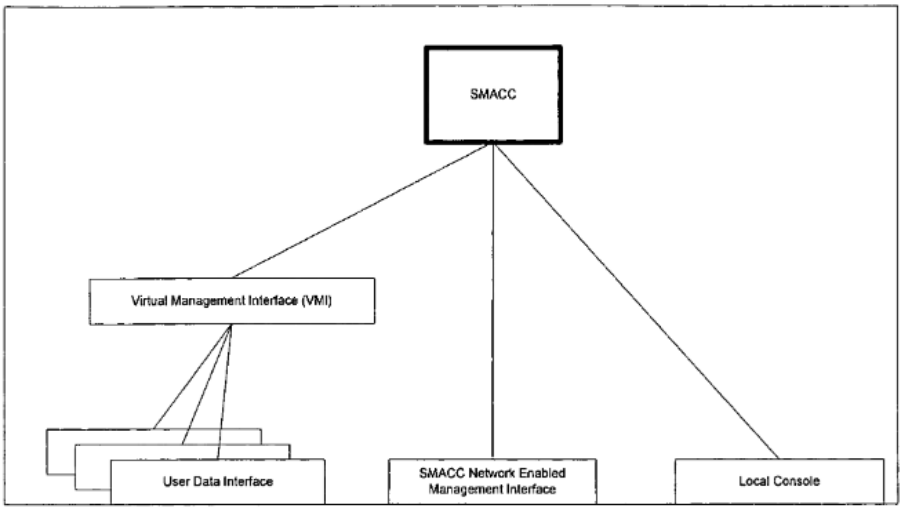

Fig. 3

As can be seen in FIG. 1 of the '140 Patent (reproduced below for convenience), the SMACC is located within a managed device in a distributed computer network. The SMACC is communicatively coupled both to the CPU Bus and the System Bus of the managed network device. *See* Ex. 1002, FIG. 1. The managed network device is configured for communication using user data interfaces via the System Bus, as shown. *See id.*

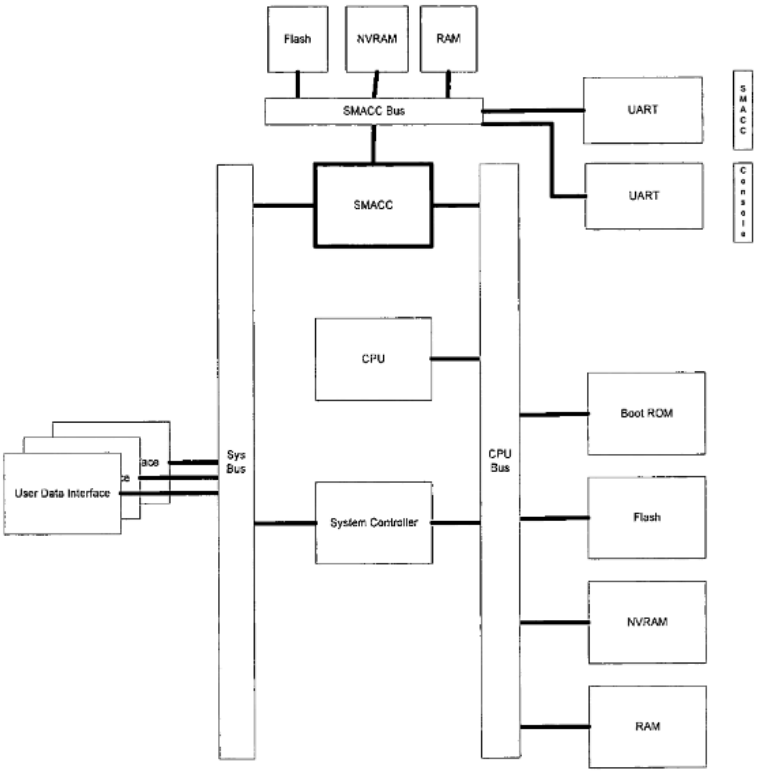

Fig. 1

1. Interfaces Available at the SMACC

"[A]s shown in FIG. 3, the SMACC has two types of interfaces for transporting management traffic between the SMACC and the management center": a SMACC Network Enabled Management Interface (also referred to in the '140 Patent as "SMACC interface") and a Virtual Management Interface (VMI). *Id.* at 11:60-63; *see id.* at 11:64, 12:1. The SMACC interface is for out-of-band network access and "provides a physically separate interface for the management network connection rather than the user interfaces on the managed device." *Id.* at 12:1-4.

The VMI is for in-band network access and provides security for the management data.

One way in which the VMI provides security is by "logically separat[ing] management traffic from user data when they are going to utilize the same interfaces on the managed device and the same physical network." *Id.* at 12:7-10. As is shown in FIG. 3 above, the VMI logically resides between the SMACC processor and the user data interfaces. In one described embodiment, the VMI uses a Virtual Private Network (VPN) mechanism "to encrypt, protect and logically separate the management traffic *before it is passed to the user interfaces." See id.* at 12:9-13. The VPN may be set up by the SMACC by running "a suitable Virtual Private Network (VPN) protocol such as IPSec over the VMI." *See id.* at 19:9-11. As stated in the '140 Patent:

> 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.

Id. at 6:26-33 (emphasis added); *see* Ex. 2006, ¶ 57.

The SMACC provides additional security (i.e., in addition to the VMI/VPN) of the management data within the managed device, "even when transmitted over the same physical interfaces used for user traffic." Ex. 2006, ¶ 27. "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. 2006, ¶ 26 (*citing* the '140 Patent, Ex. 1002, at 12:17). "Another key function of the '140 Patent is to integrate firewall mechanisms *directly within the SMACC/VMI* to restrict access and protect management operations." Ex. 2006, ¶ 29 (*citing* the '140 Patent, Ex. 1002, at 6:41) (emphasis added).

2. SMACC Functionality

The SMACC essentially acts as a secure hub for both: (1) network administrator communications (e.g., for maintenance/repairs) and (2) network services.

An administrator that needs to remotely perform maintenance/repairs on a managed device first needs to access a management center. *See* Ex. 1002, 8:20-22. The remote administrator may then connect "to the SMACC via the secure tunnel from the management center to the SMACC" via either the user data interfaces (for in-band access) or the SMACC interface (for out-of-band access). *See id.* at 8:20-27. "The remote administrator will be able to utilize any of these interfaces to connect from the management center to the SMACC." *Id.* at 8:27-29. The tunnels between the management center and the SMACC terminate *within* the SMACC itself (e.g., at the VMI). *See, e.g., id.* at 6:26-28, 12:10-13, 19:7-11.

Devices need access to "a set of services provided through a network" for effective management of a large number of devices. *See id.* at 14:56-58. In the '140 Patent, these services can be provided, for example, through the management center. *See id.* at 14:58-60. "A SMACC enabled device will be able to be configured to securely access these network services over the user interfaces utilizing the VMI or over the SMACC interfaces." *Id.* at 14:61-64. As further explained in the '140 Patent specification:

> The SMACC interface is a network enabled interface and therefore is able to support the full range of protocols that are typically utilized for management of a device such as SNMP, Trivial File Transfer Protocol (TFTP), File Transfer Protocol (FTP), DNS, SysLog, Telnet, Secure Shell (SSH), or HTTP. The SMACC interface will also control access to the management functions of the device and protect the management traffic in transit through the utilization of VPN technology such as IPSec. The SMACC interface can be utilized to connect to a central location that provides for these network services as well as for remote administrator access.

Id. at 12:24-35.

C. Summary of the '140 Patent

The '140 Patent seeks to solve the longstanding problem of securely connecting devices with both remote administrators (e.g., for device maintenance/repairs) and needed network services, such as network management application, NTP servers, Log Servers, ACS, etc., by embedding a SMACC in the device itself. *See id.* at 5:50-52, 12:24-35, 14:66-15:2, 19:11-16, FIG. 1. The SMACC enables the communication of such management data using both an in-band

path and an out-of-band path and provides both security and cost savings, particularly for the in-band path; management data, which can inexpensively be co-mingled with user data and communicated using the device's user interfaces, is secured and separated from the user data both within the device and while in transit, via tunnels that terminate at *the SMACC itself*. *See id.* at 6:26-33, 11:60-12:4, 12:10-13, 19:7-11. The management center both provides the network services and enables administrators to connect to it to communicate management data with the device via the tunnels. *See id.* at 8:20-22, 14:56-60. Therefore, the SMACC can act as a hub for both receiving all of these types of management traffic and ensuring such traffic is securely communicated between the device and the management center (and vice versa).

This is not necessarily the least expensive option for securing management traffic as running a VPN on a device processor (i.e., at the SMACC) would not have been cheap at the time the application that issued into the '140 Patent was filed. However, this solution was far less expensive than using solely the out-of-band path for communication of management data, and particularly with respect to building a dedicated network for remotely accessing devices for such purposes. *See id.* at 2:12-18, 45-48. This solution enables user data and management traffic to be inexpensively co-mingled but securely separated *within the device* for communication *outside of the device* via secure tunnels, reducing its vulnerability to

both internal and external attacks. *See id.* at 6:26-33, 11:60-12:4, 12:10-13, 19:7-11. In the words of the '140 Patent specification itself:

> The SMACC provides for the separation of management data from user data both in the device being managed and while the management data is in transit. Since there is a very strong need for securing remote management and for separating management data from user data, *yet this has not been done primarily because it has been considered too expensive to implement, this invention is clearly not obvious to one of ordinary skill in this area*. The invention combines existing building blocks along with additional new and innovative features and concepts to solve the shortcomings in remote management in an integrated and affordable solution.

Id. at 9:26-36 (emphasis added).

III. PRIOR ART RELIED ON IN THE PETITION

The Petition advances two counts against independent Claim 1 of the '140 Patent. In Count 1, the Petitioner asserts that Claim 1 is obvious under 35 U.S.C. § 103 over a single reference, U.S. Patent Appln. Pub. No. 2003/0226015 to Neufeld et al. (hereinafter "Neufeld"). In Count 5, the Petitioner asserts that Claim 1 of the '140 Patent is obvious under 35 U.S.C. § 103 over IPMI – Intelligent Platform Management Interface Specification V1.5 (Rev. 1.1 dated 2/20/2020) (hereinafter "IPMI") in combination with U.S. patent No. 7,269,639 to Lawrence (hereinafter "Lawrence).

A. Neufeld

The Petition relies on Neufeld as its primary (and only) reference for Count 1. Neufeld is focused on solutions for enabling secure, remote access to a managed server *independent of the managed server's OS*. *See* Ex. 1005, ¶¶ [0010], [0011], [0034], [0037], [0038]; Ex. 2006, ¶ 45. This provides the advantage of enabling remote access to a managed server for initial setup (i.e., before the OS is up and running) and when the OS is down – times when it is really important for a remote administrator to be able to connect to the managed server for setup, maintenance and repairs.

In Neufeld, security is implemented via a security manager, provided at the managed server, that limits access privileges, imposes higher levels of access control, and/or enhances security settings for remote user access to the managed device based on factors such as the type of attempted access (e.g., direct, in-house, remote access, remote access using the internet, remote access using a VPN, etc.) and/or time of day of the attempted access (e.g., during or outside of normal work hours). *See id.* at ¶¶ [0076], [0084]-[0088]; Ex. 2006, ¶¶ 43-44. The security manager is located in a remote management controller, which can be "powered and capable of operation whether or not the managed server 2 is powered up (turned on) or online." *See* Ex. 1005, ¶ [0052]. Accordingly, in Neufeld, devices can be securely and remotely reached via the out-of-band path before the device's operating system is up and

running, while the device's operating system is down, and even when the device is powered off. *See* Ex. 2006, ¶ 45.

1. Remote Management System

FIG. 1 of Neufeld (re-produced below for convenience) illustrates a remote management system. The remote management system includes a remote management PC 5 communicatively coupled to a managed server 2 via a network N and a management network M. *See* Ex. 1005, ¶ [0034]. As illustrated, the managed server 2 includes a CPU 3 connected to a monitor 4. *See id.*

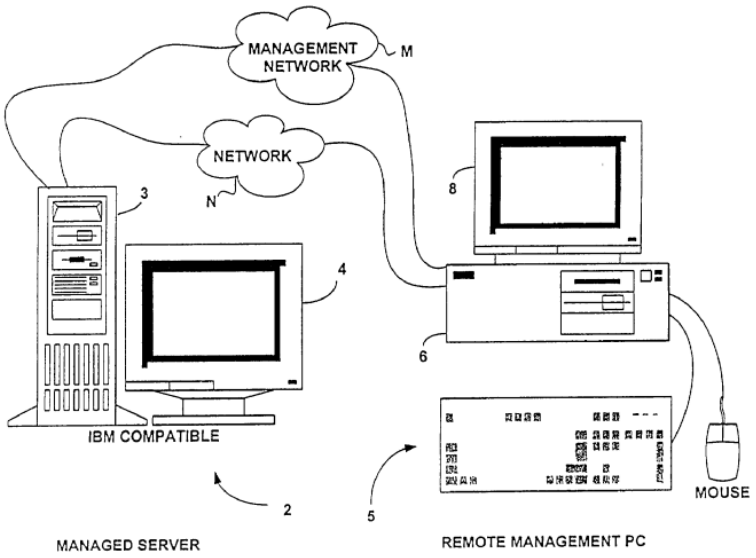

FIG. 1

FIG. 2 of Neufeld (re-produced below for convenience) is a block diagram of the CPU 3 of the managed server 2. As illustrated in FIG. 2, the managed server 2 contains a remote management controller 116, which provides "functionality for accessing, interacting [with], and monitoring the managed server 2 from the remote

console 5 . . . ." *Id.* at ¶ [0043]. The remote management controller 116 is communicatively coupled to a PCI bus 16. *See id.* at ¶ [0043], FIG. 2. "Further attached to the PCI bus 16 is a video graphics controller 114 and one or more communications devices, such as a network interface controller ("NIC") 110." *Id.* at ¶ [0045]. The network N, coupled to the NIC 110 in FIG. 2 (and not the remote management controller 116), is the managed server's normal communication network for communicating user data. As shown in FIG. 2, Network M is directly coupled to the remote management controller. *See id.* at FIG. 2; Ex. 2006, ¶ 46.

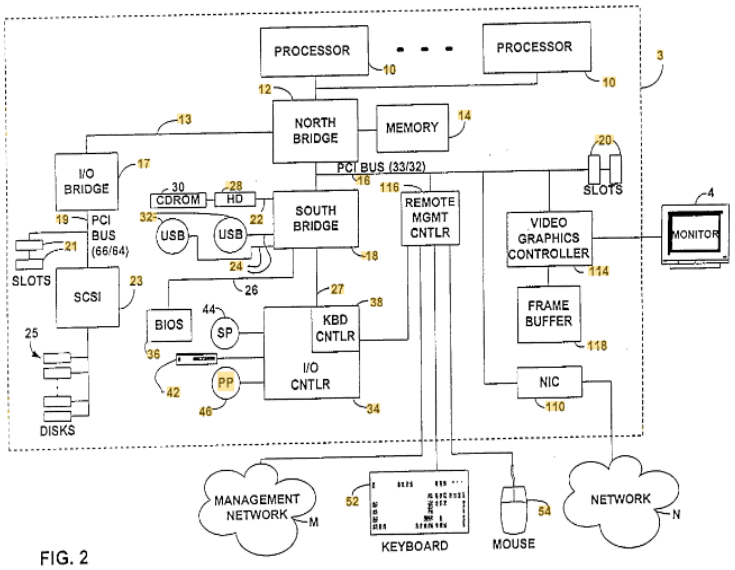

FIG. 2

2. Remote Management Controller

FIG. 3 of Neufeld (reproduced below for convenience) is a block diagram of the remote management controller 116. In FIG. 3, the PCI bus 172 is considered to be an extension of the PCI bus 16 of FIG. 2. *See* Ex. 1005, ¶ [0058]. "For example, FIG. 3 illustrates the remote server management controller 116 being coupled to the

video graphics controller 114, and thus its associated frame buffer 118 and display 4, via the PCI bus 172." *Id.*

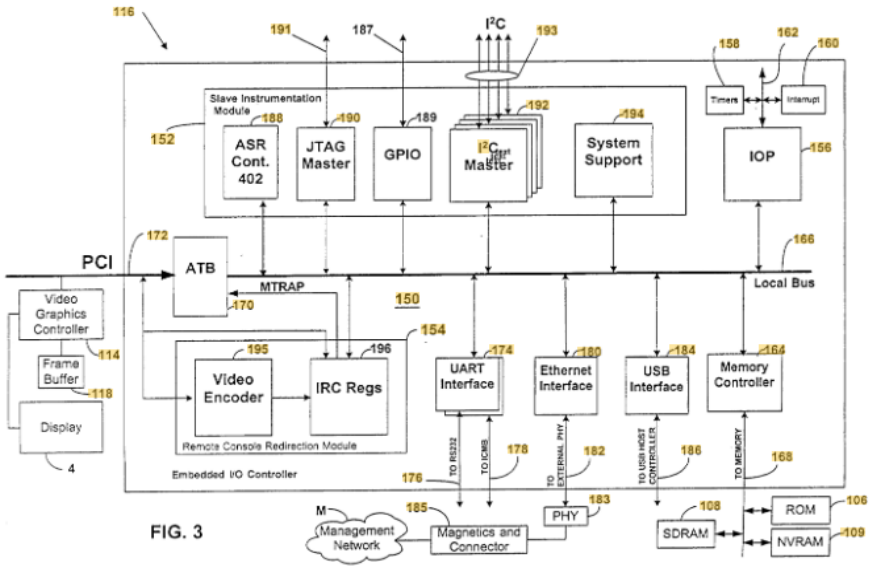

FIG. 3

3. Security From External Attacks

Neufeld further seeks to enable better authentication of remote users attempting to access managed servers, authenticating the user rather than the device. *See* Ex. 1005, ¶ [0076]; Ex. 2006, ¶¶ 43-44. For this purpose, a security manager is "stored and operated on the remote management controller 116" and used to grant network privileges to a user, and/or subject the user to certain authentication requirements and security settings, depending on the time of day the user attempts to access the network, the location and IP address from which the user attempts to access the network, and/or the connection type the user attempts to use to access the network. *See id.*

> As shown in these examples, the more remote the connection, the more rigorous the authentication that may be required. Likewise, connections made during normal working hours may require less rigorous authentication tha[n] connections made during non-business hours. Also, similarly, the range of a particular user's privileges may be restricted depending upon the type and time of connection. In other words, a particular user or group of users may have a first feature set if accessing the managed server 2 from a relatively secure location during normal business hours, while the particular user or group of users may have other feature sets if accessing the managed server 2 from other locations and/or at other times.

Ex. 1005, ¶ [0084].

Tables 1, 2, and 3 collectively provide examples of IP Address/Connector Types, times of day, authentication requirements, security settings, and privileges that can be granted to and/or imposed on a person attempting to access a network. Examples of IP Address/Connector types include direct connection, in-house connection, remote direct, remote (Internet), or remote (VPN). Examples of times of day include business hours (8 am to 5 pm) or after business hours (5 pm to 8 am). Examples of authentication requirements include password, smartcard, biometric, or smartcard/biometric. Examples of security settings include encryption or no encryption. Examples of privileges include full, non-critical only, non-critical, moderate and non-critical only, or none. Ex. 1005, Tables 1-3; Ex. 2006, ¶ 43.

TABLE 2

Network Administrator

User ID	IP Address/ Connector Type	Time of Day	Authentication	Security Settings	Privileges
1 Net-admin 1 (group:admin 1)	Direct	8 am–5 pm	Password	No encryption	Full
2 Net-admin 1 (group:admin 1)	In-house	8 am–5 pm	Smartcard	Encryption	Full
3 Net-admin 1 (group:admin 1)	Remote	8 am–5 pm	Biometric	Encryption	Moderate and non-critical only
4 Net-admin 1 (group:admin 1)	Direct	5 pm–8 am	Biometric	No encryption	Full
5 Net-admin 1 (group:admin 1)	In-house	5 pm–8 am	Biometric	Encryption	Moderate and non-critical only
6 Net-admin 1 (group:admin 1)	Remote	5 pm–8 am	Biometric	Encryption	Non-critical only

TABLE 4

User

User ID	IP Address/ Connector Type	Time of Day	Authentication	Security Settings	Privileges
1 Joe-user (group:user 1)	Direct	8 am–5 pm	NA	NA	None
2 Joe-user (group:user 1)	In-house	8 am–5 pm	Password	Encryption	Non-critical
3 Joe-user (group:user 1)	Remote	8 am–5 pm	Smartcard/ Biometric	Encryption	Non-critical
4 Joe-user (group:user 1)	Direct	5 pm–8 am	NA	NA	None
5 Joe-user (group:user 1)	In-house	5 pm–8 am	Smartcard/ Biometric	Encryption	Non-critical
6 Joe-user (group:user 1)	Remote (Internet)	5 pm–8 am	NA	NA	None
7 Joe-user (group:user 1)	Remote (VPN)	5 pm–8 am	Smartcard/ Biometric	Encryption	Non-critical

B. Intelligent Platform Management Interface Specification V1.5 (IPMI)

IPMI is a standardized protocol developed by Intel that is used to remotely manage and monitor servers, even when the main OS is down or the system is powered off. *See* Ex. 2006, ¶ 48. To accomplish this, IPMI primarily leverages out-of-band communication[1], providing several different dedicated management ports, including several auxiliary IPMB connectors for communications via remote management cards, an Intelligent Chassis Management Bus (ICMB) bridge, and a

[1] While IPMI refers to this as "out-of-band," this is more akin to "in-band" communication as the term is commonly used (and vice versa). *See* Ex. 2006, ¶ 65.

serial connector. *See id.* at ¶ 49. These management ports can be seen in FIG. 1-2 of IPMI (reproduced below). While IPMI is primarily designed for out-of-band management, there is a form of in-band communication available. *See id.* at ¶ 50.

Figure 1-2, IPMI Block Diagram

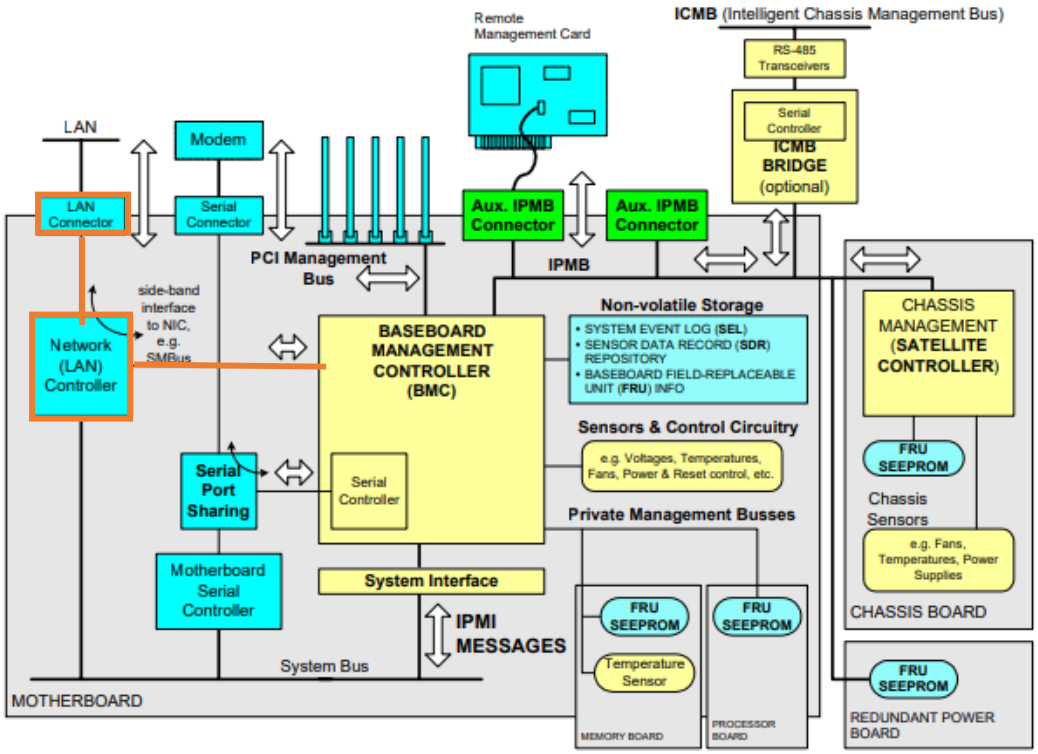

"At the heart of the IPMI architecture is a microcontroller called the *Baseboard Management Controller* or, BMC."

> The BMC provides the intelligence behind Intelligent Platform Management. 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 the ICMB.

Ex. 1020, 9:1-6; *see* Ex. 2006, ¶¶ 49-50.

In terms of in-band-communication (referred to as out-of-band in the IPMI specification), as can be seen from Figure 1-2 of IPMI, the BMC shares a network interface with the managed system's network traffic. *See* Ex. 2006, ¶ 49. The communication path for such in-band-communication is highlighted in orange in FIG. 1-2 above. FIG. 12-1 of IPMI (reproduced below) is a block diagram of this interface (referred to in IPMI as the embedded LAN Controller Implementation).

Figure 12-1, Embedded LAN Controller Implementation

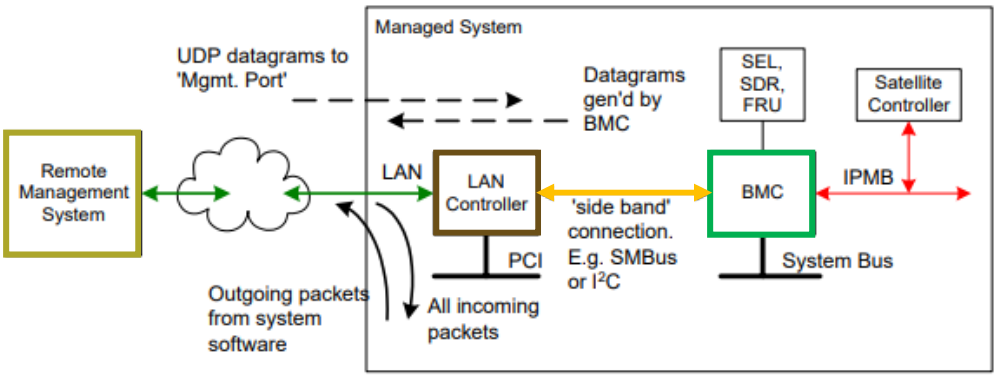

Using the embedded LAN Controller Implementation, IPMI messages may be transferred "between the BMC [green highlighting] and a remote management system (remote console) [gold highlighting] over an Ethernet LAN connection using UDP under Ipv4. The UDP datagrams are formatted to contain IPMI request and response messages . . . ." Ex. 1020, 106:1-4; *see* Ex. 2006, ¶ 49. The LAN controller [brown highlighting] "has the ability to detect UDP datagrams sent to a 'management port.' Any datagrams received on that port are forwarded on a 'side-

band' interface [orange highlighting] that allows them to be delivered to, or received by, the BMC." Ex. 1020, 106:10-12; *see* Ex. 2006, ¶ 49. "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." Ex. 1020, 106:14-15; *see* Ex. 2006, ¶ 49.

C. Lawrence

Lawrence is generally directed toward routers for secure in-band management for packet data networks. More specifically, FIG. 7 (reproduced below) shows an in-band management configuration for a router 602. The router 602 includes two main units: a generic routing and forwarding module 612, which "receives packets and performs basic routing and forwarding of data packets," and a router configuration management module, 610, which receives "management traffic to perform functions, such as adding a link or customer, tearing down a link, or configuring a new card for a router." Ex. 1043, 5:4-6, 15-16; *see* Ex. 2006, ¶ 51. An internal management VPN link 711 "connects management VPN module 655 to router configuration module 610," such that "management traffic can be sent to router configuration management module 610 internally via management VPN module 655 and internal management VPN link 711." Ex. 1043, 5:42-47 *see* Ex. 2006, ¶ 51.

> Each of the VPN modules 655 through 658 can perform VPN-related encapsulation, routing, and forwarding functions on packets through router 602. The encapsulation process entails adding extra headers to each data packet to identify VPN membership and access; furthermore,

during the encapsulation process, the data packet may be encrypted using IP security (IPSec) process or another data encryption process. In this way, the packets are kept private and secure.

Ex. 1043, 5:22-30 *see* Ex. 2006, ¶ 51.

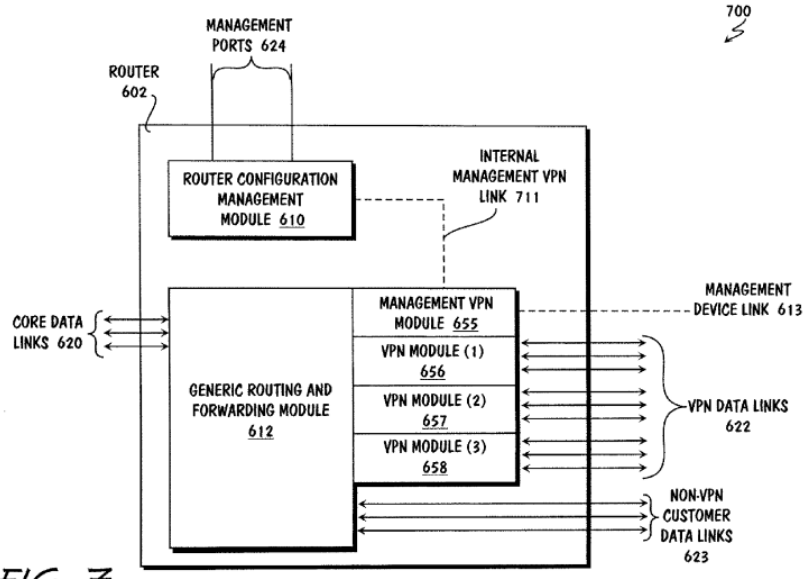

FIG. 7

IV. THE PETITION FAILS TO DEMONSTRATE A REASONABLE LIKELIHOOD OF SUCCESS ON THE MERITS

A. Ground 1 – Claims 1-4, 6, 7, 10, 11, 13 and 16 Are Not Obvious Over Neufeld Pursuant To 35 U.S.C. § 103

In Ground 1, the Petition relies on a single reference, Neufeld, in challenging Claims 1-4, 6, 7, 10, 11, 13 and 16 as allegedly obvious under 35 U.S.C. § 103. This allegation is deficient for a number of reasons. First, the VPN in Neufeld is used for access to a *management network only* and therefore cannot be said to inherently (or obviously) logically "separat[e] management data from user data" because user data and management data, in Neufeld, are never co-mingled, no matter how much

Neufeld is contorted or read in isolation without consideration of other contradictory text in the publication. Second, the VPN in Neufeld does not "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 required by Claim 1. Finally, the Petition does not meet the threshold evidentiary burden required to establish a prima facie of obviousness for Claim element 1(c).

1. The Petition Fails to Establish that Neufeld Discloses or Renders Obvious the "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" Element of Claim 1

Claim 1 of the '140 Patent recites in relevant part:

> . . .at least one virtual management interface connection means . . .;
>
> wherein said virtual management interface connection means . . . *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*.

Ex. 1002, 22:48-49, 52-57 (emphasis added).

The plain language of the claim requires that the *virtual management interface (VMI) connection means terminates within the managed network element itself*, specifically using the user interfaces of that element to facilitate communication with the secure management access controller (SMACC). *See also* Ex. 2006, ¶ 55, 57. This is not merely a connection *to* the network; it is a connection *through* and *within*

the managed element via its user interfaces (e.g., NIC). *See id.* The specification reinforces this interpretation. As explained:

> In a preferred exemplary embodiment, 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. The managed device is designed such that management traffic is not allowed to be sent directly to the user interfaces. All management traffic must go through one of the management interfaces (the VMI or the SMACC interface). The VMI allows any of the user interfaces to be utilized for management traffic while still protecting the management interface and traffic."

Ex. 1002, 12:9–20 (emphasis added). Consistent with this, FIG. 3 of the '140 Patent depicts the VMI as residing between the SMACC processor and the user data interfaces, confirming that the connection is internal to the managed device and logically isolated within it. *See* Section II.B. *supra*; Ex. 2006, ¶ 27, 29, 55, 56, FIG. 2, FIG. 3.

In contrast, Petitioner relies on Neufeld and alleges that VPN tunneling functionality, which it alleges is used for communication between the remote management PC and Neufeld's remote server management controller 116, satisfies this claim element. Specifically, Petitioner contends that:

> The VPN tunneling functionality dedicated to communications of management data from remote management PC 5 also performs the function of connecting a remote user on remote management PC 5 with Neufeld's remote server management controller 116 for communication of management data *using the same user interface (Neufeld's NIC 110 connected to network N) over which non-management user data is communicated to the managed server 2*.

Petition, 34:8–14 (emphasis).

This assertion is unsupported. Neither Petitioner nor its expert (whose Declaration, at ¶ 115, merely repeats the conclusory language from the Petition) cites any portion of Neufeld disclosing that a VPN tunnel connects through or terminates at the managed server's user interfaces in the manner required by the '140 Patent. Rather, Neufeld refers only to *general-purpose VPNs* used for *network access*, not internal communication paths within the managed device:

- "Clients may also have the ability to access these networks via a virtual private network ('VPN') or over the Internet." (Ex. 1005, ¶ [0009]).
- "In this situation, the system administrator must typically use the Internet or a virtual private network to access the system." (*Id.* at ¶ [0022]).
- "[A] user . . . may attempt different types of remote connections, such as via the Internet or via a virtual private network. . . . However, all users, or certain groups of users, may be permitted access during normal and/or off business hours via a virtual private network with proper authentication." (*Id.* at ¶ [0088]).

As explained by Dr. Noubir, a POSITA would have understood that such *general-purpose VPNs* terminate outside the managed device—typically at a *VPN gateway or network router*—and do *not* utilize or connect through the user interfaces of the managed device itself. Ex. 2006, ¶¶ 32-36, 57.

Accordingly, Petitioner has failed to show that Neufeld discloses or renders obvious the claim limitation requiring that the VMI "utilizes user interfaces of said managed network element" to connect services or remote users with the remote management controller 116, which it equates with the "secure management access controller" of Claim 1. *See* Ex. 2006, ¶ 58.

Moreover, to retrofit Neufeld's architecture to comply with this claim feature, a POSITA would need to substantially *redesign the system*, such as by moving the VPN endpoint inside Neufeld's management controller 116 and rearchitecting its interfaces to match the claimed configuration. This would not only be *contrary to the teachings of Neufeld*, but also *economically and technically prohibitive*, as recognized in the '140 Patent itself. *See* Ex. 1002, 9:26-36. Such a modification amounts to *hindsight reconstruction* and would not have been obvious under 35 U.S.C. § 103. Indeed, the Federal Circuit has consistently held that modifications requiring a *complete overhaul* of the prior art system weigh *against* a finding of obviousness. *See In re Ratti*, 270 F.2d 810, 813 (C.C.P.A. 1959) (*concluding that* A prima facie case of obviousness cannot be established when the "suggested combination of references would require a substantial reconstruction and redesign of the elements shown in [the reference] as well as a change in the basic principle under which the [reference] construction was designed to operate."); *Trivascular,*

Inc. v. Samuels, 812 F.3d 1056, 1068–69 (Fed. Cir. 2016) (finding that substantial changes to the prior art "teach away" from obviousness).

Accordingly, the Petition fails to meet its burden to demonstrate that Claim 1, including element 1(c), would have been obvious in view of Neufeld.

2. The Petition Fails to Show that Neufeld Discloses or Renders Obvious the Limitation of Claim 1 Requiring Virtual Management Interface Connection Means to Provide Logical Separation of Management Data from User Data

Claim 1 of the '140 Patent recites in relevant part:

> . . .at least one virtual management interface connection means provides logical separation of management data from user data....

Ex. 1002, 22:48-49, 52-57 (emphasis added).

The Petition asserts that this limitation is satisfied by a general-purpose VPN referenced in Neufeld, which it claims would have inherently or obviously provided "logical separation" of management data from user data via tunneling because "the VPN tunneling necessarily differentiates traffic for the dedicated VPN from other user traffic." *See* Petition at 32:6–34:14; Ex. 2006, ¶ 57; Institution Decision at 20-22. This might be the case if the VPN were used to make an in-band connection for transport of management data, an assumption made in the Petition; however, this is not how Neufeld's system works. *See* Ex. 2006, ¶¶ 47, 59. This argument not only lacks evidentiary support but is fundamentally inconsistent with the actual teachings

of Neufeld, which describe a system architecture where *management data is not communicated over network N at all*. *See id.*

The Petition fails to cite any passage in Neufeld for the "separation of management data from user data" claim element showing that management data is transmitted over the in-band network N. *See* Petition, 32:6-34:14. Moreover, the Petitioner fails to realize that *all management data in Neufeld is transmitted over the management network M*. *See id.* As made in clear in Neufeld: "[i]n-band communication with the server is . . . 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." *See* Ex. 1005 at ¶ [0011]. Indeed, the management network M is the only network coupled to any of the communication interfaces of the remote management controller 116 in Neufeld – in both FIGs. 2 and 3, the Network N is coupled to the PCI bus, outside of the remote management controller 116.

Petitioner's assertion that VPN tunneling over the in-band NIC could logically separate management and user data, therefore, fails at least because *logical separation requires the presence of both data types in the same communication channel*, such that they must be distinguished and isolated. *See* Ex. 2006, ¶ 61. In Neufeld, the in-band channel carries *only user data*, and the management data is sent through a separate out-of-band path. As such, even if a VPN were used, there is *no*

evidence of logical separation being performed because there is *no mixed traffic to separate*.

Accordingly, because Neufeld does not disclose or suggest using the in-band network to transmit management data and fails to describe any logical separation of user and management traffic over that path, the Petition fails to establish that the limitation requiring a virtual management interface connection means that "provides logical separation of management data from user data" is disclosed or rendered obvious. Accordingly, this limitation is not met, and the Petition should be denied as to Claim 1.

Neufeld does state, in an isolated instance in paragraph [0043], that "remote management controller 116 connects to … the network N **and/or a management network M** to provide functionality for accessing, interacting, and monitoring the managed server 2 from the remote console 5." Ex. 1005, ¶[0043]; Institution Decision at 18; Petition; 27:4-7 (*leaving out key "and/or management network M" information from the original sentence in Neufeld*). Relying on this disclosure of Neufeld, the Institution Decision states: "one of ordinary skill in the art would have found it obvious to implement Neufeld's VPN for the in-band connection used for non-management user traffic (i.e., the NIC 110 and network N) to provide *a dedicated VPN for communications of management data between the managed server 2 and the remote management PC 5*." Institution Decision at 21 (emphasis

added). Paragraph [0043], however, is about passive snooping on the PCI bus, and not any active exchange of management data between the remote management controller and the managed server. *See* Ex. 2006, ¶ 47. In fact, the RMC 116 operates independently of the managed server's main data path, using its direct connections to the PCI bus and input devices solely to *observe* system activity without interfering with or participating in normal bus communications. *See id.*; Ex. 1005, ¶ [0047]. Neufeld makes clear that the remote management controller 116 communicates externally through a dedicated out-of-band channel—"The Ethernet interface 180 provides the main external communication interface between the remote server management controller 116 and the outside world." Ex. 1005, ¶0061. Thus, any suggestion that the remote management controller receives encrypted in-band management commands over the NIC 110 is inconsistent with the disclosed architecture and the express teachings of Neufeld.

More specifically, paragraph [0043] of Neufeld is describing the hardware-level connections that enable the remote management controller 116 to perform passive out-of-band monitoring of the server—not active in-band communication. Paragraph [0043] serves to inform a person of ordinary skill in the art that the remote management controller 116 is physically connected to system components, such as the PCI bus and keyboard controller, which allow it to snoop on system activity (e.g., intercepting display data or keyboard inputs) without relying on or interfering with

the host operating system. The mention of network N or a management network M does not imply that the RMC uses the NIC 110 to transmit management data through the server's primary network path; to the contrary, the publication makes clear that the remote management controller's external communications are handled through a dedicated interface (e.g., Ethernet interface 180), *and that the entire system is designed to operate independently of any software on the managed server*. Instead, the reference to network N is simply identifying that the remote management controller is able to snoop on the data being sent/received by/to the managed device via the network. It could not snoop on the traffic unless it was connected. Thus, Paragraph [0043] describes the hardware context for the remote management controller's out-of-band functionality—not a software-dependent in-band communication path.

By being attached to the PCI bus, the remote management controller is positioned to observe bus-level transactions, including configuration and memory operations involving components like the video controller, allowing it to capture display data for redirection to a remote user. The RMC's connection to the keyboard controller further enables it to snoop keyboard input—monitoring scan codes transmitted from a physical keyboard to the system's 8042 controller—so that it can emulate or replicate those keystrokes remotely. These hardware-level connections enable the RMC to collect input/output state information non-intrusively, *without*

requiring cooperation from the host operating system, and are foundational to its out-of-band management capabilities, such as remote console control and system diagnostics.

Accordingly, to the extent Petitioner is twisting Paragraph [0043] to describe that RMC 116 communicates with the remote console 5 via NIC 110, that twisted interpretation is fundamentally at odds with the operating system and software independence required in Neufeld's system. For the RMC 116 to send data over the PCI 32/33 over NIC 110 to Network N would require the involvement of the operating system. For RMC 116 to transmit data over the PCI bus to NIC 110 and onto Network N would ordinarily require the host operating system to be active and managing the network stack—contrary to the core purpose of enabling out-of-band communication. Paragraph [0043] simply identifies the hardware components to which the RMC 116 is connected, such as the PCI bus, keyboard controller, and peripheral devices, to inform a person of ordinary skill in the art that the RMC is positioned to passively monitor—or snoop on—the internal activities of the managed server 2, independent of host software.

To the extent Network N can *ever* be used for communication of management data, which seems to be contradicted by Neufeld's own statements about the limitations on the operability of the Network N for in-band use, including statements that teach away from such use, the "dedicated VPN for communications of

management data between the managed server 2 and the remote management PC 5" that one of ordinary skill may contemplate is schematically depicted below by yellow highlighting in Neufeld's annotated FIG. 1.

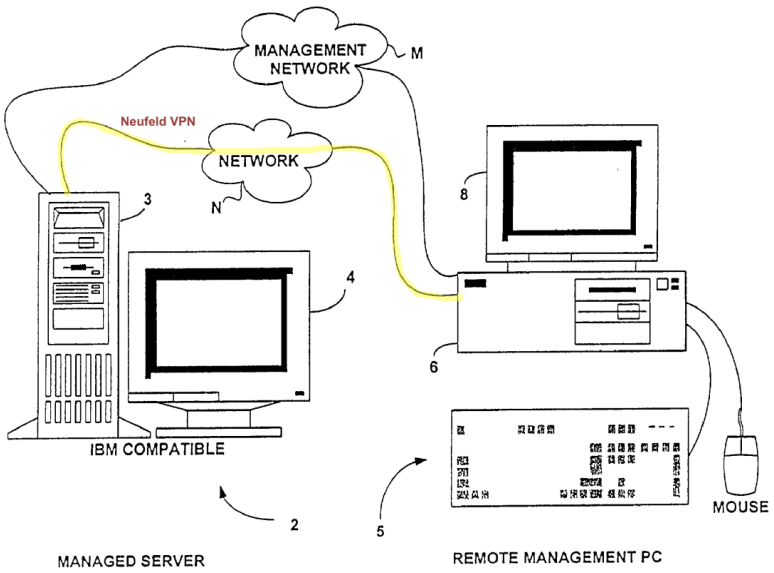

FIG. 1

Ex. 1005, FIG. 1 (annotated).

Consistent with the tunneling aspect thereof, and assuming *arguendo* that management data can or should in fact be transmitted over the Network N based on Neufeld's own teachings, such a VPN would theoretically separate management data (from the remote management PC 5) from all other network data passing through the network N (highlighted in yellow above). Consistent with a conventional VPN, the VPN tunnels would not extend, in any event, beyond the NIC 110 (and in fact would not even extend beyond a VPN Gateway or Router on the network). As such, any user data and management data co-mingled along a connection path from the

Router or VPN Gateway to the remote management controller 116, and especially within the CPU 3, will remain co-mingled, as shown in Neufeld's annotated FIG. 2 below.

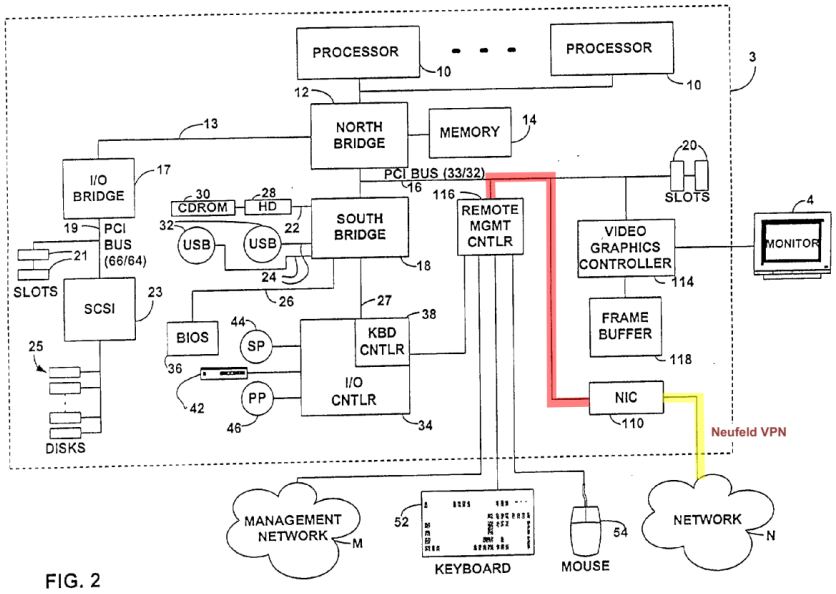

FIG. 2

Ex. 1005, FIG. 2 (annotated).

This is in stark contrast to what the '140 Patent teaches, which is:

(i) a virtual management interface that:

(A) is located between the user interface and the SMACC; and

(B) separates user data from management data both within the device and along the entire path at least between the VMI and the remote management device.

In the '140 Patent, in fact, the VMI virtualizes the user data interfaces "to ensure simultaneous yet separate handling of user data and management data." Ex. 2006, ¶¶ 26-27, 56. This would ensure the management data is kept separate, even

when communicated over same physical hardware as user data, which is *not possible* using Neufeld's architecture. *See id.*

In light of Neufeld's disclosure and the knowledge of a POSITA, a POSITA would not have known to provide a VPN between Neufeld's NIC 110 and the remote management controller 116 to provide separate management data from user data, so that only management data may be provided to the remote management controller 116. For this additional reason, Neufeld does not teach or suggest "virtual management interface connection means [that] provides logical separation of management data from user data and utilizes user interfaces of said managed network element," as claim 1 recites.

3. Petitioner's Legal Reasoning is Deficient Under Federal Circuit Precedent

In addition to the technical deficiencies discussed above, Petitioner's legal theory under 35 U.S.C. § 103 is fundamentally flawed. The Federal Circuit has held that 35 U.S.C. § 103 "does not permit [stitching] together an obviousness finding from discrete portions of prior art references without considering the references as a whole." In re Enhanced Sec. Research, LLC, 739 F.3d 1347, 1355 (Fed. Cir. 2014). Accordingly, "[i]t is impermissible within the framework of section 103 to pick and choose from any one reference only so much of it as will support a given position, to the exclusion of other parts necessary to the full appreciation of what such reference fairly suggests to one of ordinary skill in the art." In re Wesslau, 353 F.2d

238 (CCPA 1965). . Like *In re Enhanced Sec. Research* and *In re Wesslau*, this is an attempt, by the Petitioner, to piece together discrete concepts, some of which are mentioned only in a single sentence and in an entirely different cocept. More specifically, here, Petitioner cites only to brief background and security related descriptions in Neueld, which discuss VPNs in the context of access control policies—not in relation to network level architecture, internal VPN termination, or data segreation. Accordingly, Petitioner has not met its burden of establishing a prima facie of obviousness because the discrete concepts cited to in the Petition are simply not combinable to produce entire limitations of claim 1.

Petitioner's entire argument for claim element 1(c) relies on paragraphs [0009], [0022], and [0088] of Neufeld, which describe VPN use as a *remote access method*—not as a mechanism for logical separation or management channel routing. Yet Petitioner extrapolates from these disclosures the entire functionality of the claimed virtual management interface, including the critical features of *logical separation* and *interface utilization*, based on little more than general assertions from its expert (Dr. Meldal). *See* Petition, 32:22–34:14; Ex. 1003 at ¶¶ 110–115. These statements are entirely conclusory. Neither the Petition nor Dr. Meldal provides *any contemporaneous references, technical explanation*, or *reasoned analysis* showing how or why a POSITA would understand a general-purpose VPN as described in Neufeld to perform the claimed functions. This lack of evidence is fatal. As the

Federal Circuit held in *Arendi S.A.R.L. v. Apple Inc.*, 832 F.3d 1355, 1362 (Fed. Cir. 2016), the Board *cannot rely on "common sense" or generalized assertions* to supply a missing limitation in the absence of substantial evidence—particularly "where the missing limitation is not unusually simple or the technology not particularly straightforward." This case fits squarely within that rule. The '140 Patent deals with *sophisticated device-level architecture and data channel isolation*, not the kind of trivial functionality that can be assumed without proof.

Likewise, in *TQ Delta, LLC v. Cisco Sys., Inc.*, 942 F.3d 1352, 1360 (Fed. Cir. 2019), the Federal Circuit reiterated that *expert opinions must provide "reasoned analysis,"* not just ipse dixit conclusions. Here, Dr. Meldal simply parrots the Petition's unsupported assumptions without tying his statements to anything in Neufeld or any independent technical source. This falls short of the evidentiary threshold required to establish obviousness under 35 U.S.C § 103.

4. Conclusion

For the foregoing reasons, Ground 1 of the Petition fails to meet the evidentiary and legal threshold required under 35 U.S.C. § 103. Petitioner's reliance on Neufeld in its single-reference obviousness theory is legally insufficient and technically unsupported. The asserted reference does not disclose or render obvious critical limitations of Claim 1—most notably, the use of user interfaces within the managed network element to facilitate communication with a secure management

access controller, and the provision of logical separation of management data from user data. Attempting to retrofit Neufeld's architecture to approximate these functionalities would require a substantial redesign that contradicts its core teachings, rendering such a modification impermissible under established Federal Circuit precedent. Furthermore, the Petition lacks any reasoned explanation or substantial evidence justifying the alleged modifications, relying instead on conclusory expert assertions that amount to impermissible hindsight. As emphasized in *Arendi*, *TQ Delta*, and *TriVascular*, such conjecture cannot supply missing claim elements or establish a prima facie case of obviousness. Accordingly, the Board should reject Ground 1 and confirm the patentability of Claims 1–4, 6, 7, 10, 11, 13, and 16 of the '140 Patent.

B. Ground 2 – Dependent Claims 11 and 12 are Not Obvious over Neufeld and Syvanne

The Petition does not allege, nor does Syvanne teach, "at least one virtual management interface connection means [that] provides logical separation of management data from user data and utilizes user interfaces of said managed network element for connection said one or more services or remote users with said secure management access controller," as recited in independent Claim 1. Accordingly, because Syvanne does not cure the shortcomings of Neufeld, the Petitioner has failed to meet its burden with regard to Ground 2 for at least the reasons discussed with regard to Ground 1 above.

C. Ground 3 - Dependent Claim 14 is Not Obvious Over Neufeld and Fowler Pursuant To 35 U.S.C. § 103

The Petition does not allege, nor does Fowler teach, "at least one virtual management interface connection means [that] provides logical separation of management data from user data and utilizes user interfaces of said managed network element for connection said one or more services or remote users with said secure management access controller," as recited in independent Claim 1. Accordingly, because Fowler does not cure the shortcomings of Neufeld, the Petitioner has failed to meet its burden with regard to Ground 3 for at least the reasons discussed with regard to Ground 1 above.

D. Ground 4 - Dependent Claim 13 is Not Obvious Over Neufeld and Fowler Pursuant To 35 U.S.C. § 103

The Petition does not allege, nor does Fowler teach, "at least one virtual management interface connection means [that] provides logical separation of management data from user data and utilizes user interfaces of said managed network element for connection said one or more services or remote users with said secure management access controller," as recited in independent Claim 1. Accordingly, because Fowler does not cure the shortcomings of Neufeld, the Petitioner has failed to meet its burden with regard to Ground 4 for at least the reasons discussed with regard to Ground 1 above.

E. Ground 5 - Claims 1-3, 6, 7, 10 And 11 Are Not Obvious Over IPMI and Lawrence Pursuant To 35 U.S.C. § 103

In Ground 5, the Petition argues that Claim 1 is obvious over an alleged combination of IPMI and Lawrence. This allegation is deficient at least because the claim feature "at least one virtual interface connection means [that] provides logical separation of management data from user data," would not be obvious from IPMI and Lawrence.

1. The Claim Feature "At Least One Virtual Interface Connection Means [That] Provides Logical Separation Of Management Data From User Data" Would Not Be Obvious From IPMI And Lawrence

IPMI is a standardized interface used for remotely managing and monitoring computer hardware. The Petitioner's argument largely centers around the BMC, which provides management functions. IPMI uses User Datagram Protocol (UDP) for sending IPMI packets. If IPMI is implemented in a system that makes use of out-of-band communication for network maintenance, a network LAN controller will detect UDP datagrams addressed to the BMC and send them to the BMC over a side band connection. Similarly, the BMC can inject UDP datagrams onto the side-band connection. Ex. 2006, ¶¶ 49, 67.

Lawrence is directed toward a VPN. The VPN in Lawrence has the following functionality: performing "VPN-related encapsulation, routing, and forwarding functions on packets passing through router 602." Ex. 1043, 5:23-24. Lawrence further mentions that "a VPN is a private data network that makes use of the

currently implemented network by using a tunneling protocol for security purposes." Ex. 1043, 2:35-37.

In its obviousness contention, the Petitioner argues as follows:

> Since IPMI discloses in-band communication from the physical LAN connector to the LAN Controller, a POSITA would have found it obvious to incorporate Lawrence's VPN module 655 into IPMI's LAN Controller to separate the IPMI UDP datagrams from other user traffic and send the IPMI UDP datagrams over the side-band interface to the BMC. In the IPMI-Lawrence combination, those packets are processed by Lawrence's dedicated management VPN module 655....

Petition, page 64 (citations removed).

a) IPMI's LAN Controller Does Not Separate Data, But Rather Detects Packets Addressed to the BMC

The Petition's main argument for Ground 5 is that IPMI discloses a LAN controller that "separate[s] the IPMI UDP datagrams from other user traffic and send[s] the UDP datagrams over the side-band interface to the BMC." Petition, page 64. "In the IPMI-Lawrence combination, those packets are processed by Lawrence's dedicated management VPN module 655." *Id.* This argument is based on a fundamental misunderstanding of how LAN controllers work.

LAN controllers function as intermediaries in network communication by managing data packet flow between devices and a computer's main processor. Ex. 2006, ¶ 67. Acting as "traffic cops," LAN controllers inspect the address of each received packet to determine its intended destination. Using embedded logic, the

LAN controller evaluates the packet's address against the network configuration and decides where to route it. This ensures that only relevant packets reach the appropriate bus or device, preventing unnecessary congestion and maintaining efficient data flow. *See id.*

Accordingly, the basic functionality of LAN controllers is to route packets based on their addresses, and the LAN controller in IPMI does that without the addition of a VPN. In fact, if IPMI's LAN controller was combined with a VPN, as Petitioner suggests, and then encountered VPN packets, it would route them exactly the same as every other packet because, from the perspective of the LAN controller, the packets have the same destination address, and hence, are indistinguishable. There is simply nothing missing from IPMI that would have caused a POSITA to go looking for Lawrence. Therefore, as the Institution Decision acknowledges, "one of ordinary skill in the art would [not] have had a rational reason to combine IPMI and Lawrence's teachings as Petitioner does." Institution Decision at 29.

b) Petitioner Has Not Established that It Would Have Been Obvious to Modify IPMI, With the Teachings of Lawrence, to Arrive at the Claimed Invention Nor Would the Combination Have Been Obvious to a POSITA

To establish obviousness under 35 U.S.C. § 103, a petitioner must demonstrate that the claimed invention would have been obvious to a POSITA at the time of the invention. This requires showing a reason or motivation to combine or

modify prior art references with a reasonable expectation of success, supported by evidence or reasoning. The analysis must be based on the scope and content of the prior art, differences between the prior art and the claimed invention, and the level of ordinary skill in the art. *Graham v. John Deere Co.*, 383 U.S. 1, 1966. As explained by the Federal Circuit, "[t]o satisfy its burden of proving obviousness, a petitioner cannot employ mere conclusory statements." *In re Magnum Oil Tools Int'l, Ltd.*, 829 F.3d 1364, 1380 (Fed. Cir. 2016). Further, "[a]n obviousness determination requires finding that a POSITA would have been motivated to combine or modify the teachings in the prior art and would have had a reasonable expectation of success in doing so." *Endo Pharm., Inc. v. Actavis LLC*, 922 F.3d 1365, 1373 (Fed. Cir. 2019) (emphasis added). Accordingly, "[w]hether there was a motivation to combine references, and a reasonable expectation of success in doing so to meet the limitations of the claimed invention are questions of fact" that must be established by the Petitioner. *Samsung Elecs. Co. v. Elm 3DS Innovations, LLC*, 925 F.3d 1373, 1380 (Fed. Cir. 2019)

With respect to a reason or motivation to modify the reference, the Petition gives the following reason:

> This implementation of IPMI"s side-band interface using Lawrence's VPN technology is the identical structure for this limitation disclosed in the 140 patent, namely a VPN. Both structures perform the identical claimed functions – connecting remote users (e.g., on IPMI's remote management system) with the BMC and providing logical separation of IPMI management

> data from user data utilizing the same LAN interface of the managed system and the network connected thereto.

Petition, 65:3-9.

As explained the Federal Circuit, "the inventor's own path itself never leads to a conclusion of obviousness; that is hindsight. What matters is the path that the person of ordinary skill in the art would have followed, as evidenced by the pertinent prior art." *Otsuka Pharm. Co., v. Sandoz, Inc.*, 678 F.3d 1280, 1296 (Fed. Cir. 2012). Therefore, obviousness is not established under 35 U.S.C §103 when "[i]t is only with the benefit of hindsight that a person of skill in the art would have had a reasonable expectation of success in view of the asserted references." *OSI Pharms., LLC v. Apotex Inc*., 939 F.3d 1375, 1385 (Fed. Cir. 2019). Accordingly, a finding of obviousness under 35 U.S.C. § 103 cannot be based on hindsight; rather, it must be supported by a reasonable expectation of success from the perspective of a person of ordinary skill in the art, without relying on the inventor's own path. Therefore, the obviousness allegation in this case fails to meet the required standard and should not be sustained.

Additionally, in *In re Ratti*, the Court held that a combination of references is improper if it renders the combined references inoperable for their intended purpose. 270 F.2d 810, 813, 47 C.C.P.A. 1059 (1959)). The Federal Circuit emphasized that a modification leading to an incompatible or inoperative result fails to support a finding of obviousness. In *In re Ratti*, the combination of prior art references would

have required substantial alterations that changed the fundamental operation of the devices involved, thus making the combination non-functional. Similarly, the United States Patent and Trademark Office, Manual of Patent Examining Procedure (MPEP)[2] explains that *In re Gordon*, 733 F.2d 900, 221 USPQ 1125 (Fed. Cir. 1984) held that "[i]f a proposed modification would render the prior art invention being modified unsatisfactory for its intended purpose, then there is no suggestion or motivation to make the proposed modification." (MPEP §2143.01). Accordingly, as explained in *In re Ratti* and *In re Gordon*, a proposed combination of references is improper under 35 U.S.C. §103 when the combination would render the references inoperable for their intended purpose.

The Petitioner's proposed combination of IPMI and Lawrence would lead to a device that is not operable for its intended purpose. The LAN controller in IPMI works by reading the addresses of received packets and routing them to the appropriate entity (in this case, the BMC). Adding a VPN into the mix, allegedly to separate management data from user data, would completely change the operation of the LAN controller from routing data within a network to separating management data from user data. Therefore, Petitioner's proposed combination would not result merely in an implementation of the LAN controller but would only serve to change

[2] "The MPEP has no binding force … but is entitled to notice so far as it is an official interpretation of statutes or regulations with which it is not in conflict." Litton Sys. v. Whirlpool Corp., 728 F.2d 1423, 1439 (Fed. Cir. 1984)

the functionality of IPMI's LAN controller.

The Institution Decision states that the Board does "not find that [the LAN controller and VPN] Functions are Incompatible, But Rather the Combination Petitioner Proposes Would Allow for Both." Institution Decision at 28. However, the Petition proposes a combination in which the LAN controller somehow runs the VPN. *See* Petition, 64:7-11 (*stating in relevant part "a POSITA would have found it obvious to include Lawrence's VPN and tunneling technology in IPMI's Network (LAN) controller in order to implement IPMI's side-band interface"*). This combination is not possible given the severe limitation of LAN capabilities, and would "introduce[] architectural complexities without recognized benefit. Ex. 2006, ¶¶ 67-69. Dr. Noubir provides ample evidence of this in declaration:

> . . . [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.
>
> 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.
>
> 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.

Ex. 2006, ¶¶ 67-69.

The critical flaws in Petitioner's argument establish that the alleged combination of IPMI and Lawrence fails to teach or suggest the separation of management data from user data required by the challenged claims. Accordingly, for at least the reasons provided above, the Petition has failed to meet its burden regarding Ground 5.

F. Ground 6 - Dependent Claim 4 is Not Obvious Over IPMI, Lawrence and Leedy Pursuant To 35 U.S.C. § 103

The Petition does not allege, nor does Leedy teach, "at least one virtual management interface connection means [that] provides logical separation of

management data from user data," as recited in independent Claim 1. Accordingly, because Leedy does not cure the shortcomings of IPMI and Lawrence, the Petitioner has failed to meet its burden with regard to Ground 6 for at least the reasons discussed with regard to Ground 5 above.

G. Ground 7 - Dependent Claims 11 And 12 Are Not Obvious Over IPMI, Lawrence And Syvanne Pursuant To 35 U.S.C. § 103

The Petition does not allege, nor does Syvanne teach, "at least one virtual management interface connection means [that] provides logical separation of management data from user data," as recited in independent Claim 1. Additionally, as explained by Dr. Noubir:

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

Ex. 2006, ¶¶ 71-73.

Accordingly, because Syvanne does not cure the shortcomings of IPMI and Lawrence, and also because one of ordinary skill in the art would not have been motivated to modify IPMI, as suggested in the Petition, the Petitioner has failed to meet its burden with regard to Ground 7.

H. Ground 8 - Dependent Claims 13 and 14 Are Not Obvious Over IPMI, Lawrence and Fowler Pursuant To 35 U.S.C. § 103

The Petition does not allege, nor does Fowler teach, "at least one virtual management interface connection means [that] provides logical separation of management data from user data," as recited in independent Claim 1. Accordingly, because Fowler does not cure the shortcomings of IPMI and Lawrence, the Petitioner has failed to meet its burden with regard to Ground 8 for at least the reasons discussed with regard to Ground 5 above.

I. Dependent Claim 16 is Not Obvious Over IPMI, Lawrence and RFC1492 Pursuant To 35 U.S.C. § 103

The Petition does not allege, nor does RFC1491 teach, "at least one virtual management interface connection means [that] provides logical separation of management data from user data," as recited in independent Claim 1. Accordingly,

because RFC1491 does not cure the shortcomings of IPMI and Lawrence, the Petitioner has failed to meet its burden with regard to Ground 9 for at least the reasons discussed with regard to Ground 5 above.

V. PETITIONER HAS NOT MET ITS BURDEN OF PROVING UNPATENTABILITY OF THE CHALLENGED CLAIMS

The Petition is required to include "[a] full statement of the reasons for relief requested, including a detailed explanation of the significance of the evidence including material facts," 37 C.F.R. § 42.22(a)(2), which must include a statement of "[h]ow the challenged claim is to be construed" and "[h]ow the construed claim is unpatentable." *Id.* § 42.104(b)(3)-(4). For the IPR to be instituted, the Petitioner must "show[] that there is a reasonable likelihood that the Petitioner would prevail with respect to at least 1 of the claims challenged in the petition." 35 U.S.C. § 314. As stressed by the Federal Circuit:

> It is of the utmost importance that petitioners in the IPR proceedings adhere to the requirement that the initial Petition identify with particularity the evidence that supports the grounds for the challenge to each claim. . . . [T]he expedited nature of IPRs bring with it an obligation for petitioners to make their case in their Petition to institute. *Wasica Finance GmbH v. Continental Automotive Systems, Inc., 853 F.3d 1272,* 1286−87 (Fed. Cir. 2017).

Accordingly, in an IPR, the patent challenger has the burden of producing evidence in the Petition to support a conclusion of unpatentability.

In the present case, Petitioner mis-construes the cited references, placing emphasis on features of these references that have little to do with the inventions described therein. And all Petitioner provides as support for its wild assertions are mere conclusions, with no evidentiary support other than the identical conclusory statements of their expert. Accordingly, for all the reasons provided herein, the Board should reject Grounds 1-10 of the Petition and confirm the patentability of Claims 1–4, 6, 7, 10-14, and 16 of the '140 Patent.

VI. CONCLUSION

For the above reasons, Patent Owner respectfully requests that the Board deny the Petition and confirm patentability of Claims 1–4, 6, 7, 10-14, and 16 of the '140 Patent.

Respectfully submitted,

Dated: May 14, 2025

By: /Brandon R. Theiss/
Brandon R. Theiss
Registration No. 70,507
Counsel for Patent Owner
Intellectual Ventures II LLC

CERTIFICATE OF COMPLIANCE WITH WORD COUNT

Pursuant to 37 C.F.R. § 42.24(d), I certify that this Patent Owner Response complies with the type-volume limits of 37 C.F.R. § 42.24(b)(1) because it contains 11,018 words, excluding the parts of this Patent Owner's Response that are exempted by 37 C.F.R. § 42.24(a), according to the word processing system used to prepare this Patent Owner's Response.

Respectfully submitted,

Dated: May 14, 2025

By: /Brandon R. Theiss/
Brandon R. Theiss
Registration No. 70,507
Counsel for Patent Owner
Intellectual Ventures II LLC

CERTIFICATE OF SERVICE

The undersigned certifies that pursuant to 37 C.F.R. § 42.6(e), a copy of the foregoing **PATENT OWNER'S RESPONSE** was served via email (as consented to by counsel) on May 14, 2025 to lead and backup counsel of record for Petitioners as follows:

James Heintz
Jim.heintz@us.dlapiper.com

Gianni Minutoli
gianni.minutoli@us.dlapiper.com

DLA Piper LLP
One Fountain Square
11911 Freedom Drive
Suite 300
Reston, VA 20190-5602
Service email: DLA-LenovoIPR-IV@us.dlapiper.com

Dated: May 14, 2025

By: /Brandon R. Theiss/
Brandon R. Theiss
Registration No. 70,507
Counsel for Patent Owner
Intellectual Ventures II LLC