

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. IPR2024-00109

U.S. Patent No. 8,474,016

**EXPERT REPORT OF DR. PRABHAT MISHRA
IN SUPPORT OF PATENT OWNER'S RESPONSE**

TABLE OF CONTENTS

I.	INTRODUCTION AND QUALIFICATIONS	1
II.	SUMMARY OF MATERIALS REVIEWED AND CONSIDERED.....	14
III.	LEGAL STANDARDS	15
IV.	LEVEL OF SKILL IN THE ART AND PERSPECTIVE APPLIED IN THIS DECLARATION	20
V.	CLAIM CONSTRUCTION.....	21
VI.	DISCUSSION OF CERTAIN PRIOR ART RELIED ON IN PETITION	22
VII.	SPECIFICATION OF ‘016 PATENT	27
	A. Background.....	27
	B. Objects Of The Invention	29
	C. Summary Of The Invention.....	29
	D. Detailed Description Of The Invention	31
VIII.	PROSECUTION HISTORY OF ‘016 PATENT.....	38
	A. First Office Action.....	38
	B. Second Office Action	40
	C. Third Office Action - Restriction Requirement	41
	D. Fourth Office Action	42
	E. Advisory Action	43
	F. Fifth Office Action.....	43
	G. Sixth Office Action.....	44
	H. Response After Final Rejection and Notice of Allowance	45
IX.	CLAIMS OF ‘016 PATENT.....	46
X.	PRIOR ART RELIED ON IN THE PETITION.....	49
	A. Neufeld	49
	1. Background Set Forth In Neufeld.....	49
	2. Figure 1 of Neufeld.....	51

3. Figure 2 of Neufeld.....	52
4. Figure 3 of Neufeld.....	53
5. Petitioner Acknowledges That Neufeld Contains An Error In Paragraph [0047] But Fails To Explain The Rationale Supporting Petitioner’s Proposed “Correction” Of The Error	54
6. Neufeld’s Remote Server Management Controller Is Configured To Snoop On Activities Occurring At The Managed Server	58
B. Intelligent Platform Management Interface Specification V1.5 (IPMI)	61
C. Huckins	64
D. Intelligent Platform Management Interface Specification v2.0 (IPMI v2.0)	64
XI. GROUND 1 FAILS TO DEMONSTRATE THE UNPATENTABILITY OF THE CHALLENGED CLAIMS BY A PREPONDERANCE OF THE EVIDENCE.....	64
A. Neufeld’s Managed Server Does Not Receive “In Band” Remote Management Commands	65
B. Petitioner Cannot Identify a “Second Bus” in Neufeld That Is Distinct from the “First Bus”	70
C. Petitioner’s Attempt to Identify the PCI Bus as the “Second Bus” Misreads Neufeld’s Out-of-Band Architecture and Mischaracterizes the Nature of Passive Snooping.....	74
D. Petitioner Misconstrues Neufeld’s Out-of-Band Architecture by Mischaracterizing the RMC’s Passive Monitoring as In-Band Command Reception	75
XII. GROUNDS 2–8 FAIL BECAUSE THE ADDITIONAL REFERENCES DO NOT CURE NEUFELD’S DEFICIENCIES	80
XIII. GROUND 9 FAILS TO DEMONSTRATE THE UNPATENTABILITY OF THE CHALLENGED CLAIMS BY A PREPONDERANCE OF THE EVIDENCE.....	80

A.	Petitioner’s Attempt To Characterize The First Bus And Bus Controller As Part Of The “Network Device” Contradicts The Plain Language And Structural Separation Expressly Recited In Claim 1 ...	80
B.	Petitioner’s Reliance On PCI-Based Intra-Controller Communication Is Fundamentally Flawed, As It Contradicts Established PCI Architecture Principles And Introduces Unjustified Complexity With No Articulated Technical Benefit.....	83
C.	Petitioner’s Reliance On In-Band Encryption Is Misplaced, As The Cited Art Does Not Disclose Encrypted Requests Over The Out-Of-Band I²C First Bus	90
XIV.	Grounds 10-14 fail because the additional references do not cure the deficiencies in Ground 9.....	95
XV.	Declaration	96

I. INTRODUCTION AND QUALIFICATIONS

1. I, Prabhat Mishra, 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. 8,474,016 (“the ‘016 Patent”) to provide expert technical opinions with respect to the ‘016 Patent.

2. I am more than eighteen years of age, and I am a citizen of the United States, currently residing in the State of Florida.

3. I am being compensated for my work in this matter at my standard hourly rate for such consulting services. My compensation in no way depends on the outcome of this case, and I have no financial interest in the outcome.

4. In formulating my opinions, I have relied upon my knowledge, training, and experience in relevant art. My qualifications, along with a more detailed summary of my background and publications, are stated more fully in my *curriculum vitae*, which has been provided as Exhibit 2008. Here, I provide a concise summary of my qualifications.

5. I am a Professor in the Department of Computer Information Science and Engineering and a UF Research Foundation Professor at the University of Florida (UF).

6. I have also held positions as a visiting chair professor at the Indian Institute of Science (IISc, Bangalore), as a senior R&D engineer at Synopsys, as senior software engineer at Silicon Automation Systems, and as a software design engineer at Texas Instruments. I have also worked as a summer researcher at various companies, including Intel, Motorola, and Freescale (NXP).

7. I have coauthored five books, co-edited five monographs, 35 book chapters, and 250+ journal and refereed conference articles.

8. I have been recognized as a Fellow of the Institute of Electrical and Electronics Engineers (IEEE) and as a Fellow the American Association for the Advancement of Science (AAAS) for my contributions to system-on-chip validation and design automation for embedded systems. I am also a Distinguished Member of the Association for Computing Machinery (ACM).

9. I am currently a senior associate editor for ACM Transactions on Design Automation of Electronic Systems. I am also currently serving as an associate editor of ACM Transactions on Embedded Computing Systems (TECS). In the past, I have been an associate editor for IEEE Transactions on Very Large Scale Integration (VLSI) Systems, Springer Journal of Electronic Testing, IEEE Design & Test, and IET Computer & Digital Techniques. I have also served as the deputy editor-in-chief of the IET Computer & Digital Techniques for four years.

10. I am currently the technical program chair of International Conference on hardware/software Codesign and System Synthesis (CODES+ISSS), which is one of the top three conferences in Embedded Systems Week (ESWEEK). I am also serving as the vice program chair for the IEEE International Symposium on Hardware Oriented Security and Trust (HOST), one of the premier conferences in hardware security. I was a program chair for the International Conference on VLSI Design (VLSID) and the International Conference on Embedded Systems. My work has received 8060+ citations with an h-index of 48 (based on Google Scholar).

11. I have co-authored many publications related to the design of secure and trustworthy systems. Relevant examples of such publications include:

- Aruna Jayasena, Richard Bachmann and Prabhat Mishra, CISELeaks: Information Leakage Assessment of Cryptographic Instruction Set Extension Prototypes, IEEE Transactions on Information Forensics & Security (TIFS), volume 20, pages 1551-1565, January 2025.
- Hasini Witharana, Hansika Weerasena and Prabhat Mishra, Formal Verification of Virtualization-based Trusted Execution Environments, IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD), volume 43, issue 11, pages 4262-4273, November 2024.

- Zhixin Pan and Prabhat Mishra, TD-Zero: Automatic Golden-Free Hardware Trojan Detection using Zero-Shot Learning, IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD), volume 43, issue 7, pages 1998-2011, July 2024.
- Aruna Jayasena, Emma Andrews and Prabhat Mishra, TVLA*: Test Vector Leakage Assessment on Hardware Implementations of Asymmetric Cryptography Algorithms, IEEE Transactions on Very Large Scale Integration (VLSI) Systems (TVLSI), volume 31, issue 9, pages 1269-1279, July 2023.
- Aruna Jayasena and Prabhat Mishra, Scalable Detection of Hardware Trojans using ATPG-based Activation of Rare Events, IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD), volume 42, issue 12, pages 4450-4462, June 2023.
- Zhixin Pan, Jennifer Sheldon and Prabhat Mishra, Hardware-Assisted Malware Detection and Localization using Explainable Machine Learning, IEEE Transactions on Computers (TC), volume 71, issue 12, pages 3308-3321, February 2022.
- Adib Nahiyani, Farimah Farahmandi, Prabhat Mishra, Domenic Forte and Mark Tehranipoor, Security-aware FSM Design Flow for Identifying and Mitigating Vulnerabilities to Fault Attacks, IEEE

Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD), May 2018.

12. I have co-authored many publications related to design and verification of computer architectures. Relevant examples of such publications include:

- Aruna Jayasena and Prabhat Mishra, HIVE: Scalable Hardware-Firmware Co-Verification using Scenario-based Decomposition and Automated Hint Extraction, IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD), volume 43, issue 10, pages 3278-3291, April 2024.
- Abhijit Das, John Jose and Prabhat Mishra, Data Criticality in Multi-Threaded Applications: An Insight for Many-Core Systems, IEEE Transactions on Very Large Scale Integration (VLSI) Systems (TVLSI), 29(9), pages 1675-1679, 2021.
- Subodha Charles, Alif Ahmed, Umit Ogras and Prabhat Mishra, Efficient Cache Reconfiguration using Machine Learning in NoC-based Many-Core CMPs, ACM Transactions on Design Automation of Electronic Systems (TODAES), 24(6), Article 60, September 2019.
- Alif Ahmed, Yuanwen Huang and Prabhat Mishra, Cache Reconfiguration using Machine Learning for Vulnerability-aware Energy

Optimization, ACM Transactions on Embedded Computing Systems (TECS), 8(2), Article 15, March 2019.

- Ujjwal Gupta, Chetan Patil, Ganapati Bhat, Prabhat Mishra and Umit Ogras, DyPO: Dynamic Pareto Optimal Configuration Selection for Heterogeneous MpSoCs, ACM Transactions on Embedded Computing Systems (TECS), 16(5), pages 123:1-123:20, October 2017.
- Prabhat Mishra, Aviral Shrivastava, and Nikil Dutt, Architecture Description Language (ADL)-driven Software Toolkit generation for Architectural Exploration of Programmable SOC's, ACM Transactions on Design Automation of Electronic Systems (TODAES), volume 11, no 3, pages 626-658, July 2006.
- Prabhat Mishra, Mahesh Mamidipaka and Nikil Dutt, Processor-Memory Co-Exploration using an Architecture Description Language, ACM Transactions on Embedded Computing Systems (TECS), volume 3, number 1, pages 140-162, February 2004.

13. I have co-authored many publications related to design and optimization of embedded systems. Relevant examples of such publications include:

- Yuanwen Huang and Prabhat Mishra, Vulnerability-aware Energy Optimization for Reconfigurable Caches in Multitasking Systems,

IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD), 38(5), pages 809-821, May 2019.

- Xiaoke Qin, Weixun Wang and Prabhat Mishra, TCEC: Temperature- and Energy-Constrained Scheduling in Real-Time Multitasking Systems, IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD), 31(8), pages 1159-1168, August 2012.
- Weixun Wang, Prabhat Mishra and Ann-Gordon Ross, Dynamic Cache Reconfiguration for Soft Real-Time Systems, ACM Transactions on Embedded Computing Systems (TECS), volume 11, issue 2, article 28, 31 pages, July 2012.
- Weixun Wang and Prabhat Mishra, System-Wide Leakage-Aware Energy Minimization using Dynamic Voltage Scaling and Cache Reconfiguration in Multitasking Systems, IEEE Transactions on Very Large Scale Integration (VLSI) Systems (TVLSI), volume 20, issue 5, pages 902 - 910, May 2012.
- Xiaoke Qin, Chetan Murthy and Prabhat Mishra, Decoding-aware Compression of FPGA Bitstreams, IEEE Transactions on Very Large Scale Integration (VLSI) Systems (TVLSI), 19(3), pages 411-419, March 2011.

- Xiaoke Qin and Prabhat Mishra, A Universal Placement Technique of Compressed Instructions for Efficient Parallel Decompression, IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD), volume 28, no 8, pages 1224-1236, August 2009.
- Seok-Won Seong and Prabhat Mishra, Bitmask-Based Code Compression for Embedded Systems, IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD), volume 27, no 4, pages 673-685, April 2008.

14. I have co-authored many publications related to secure and energy-efficient on-chip communication. Relevant examples of such publications include:

- Hansika Weerasena and Prabhat Mishra, Breaking On-Chip Communication Anonymity using Flow Correlation Attacks, ACM Journal on Emerging Technologies in Computing Systems (JETC), volume 20, issue 4, pages 1-27, October 2024.
- Hansika Weerasena and Prabhat Mishra, Security of Electrical, Optical and Wireless On-Chip Interconnects: A Survey, ACM Transactions on Design Automation of Electronic Systems (TODAES), volume 29, issue 2, pages 1-41, February 2024.
- Aruna Jayasena, Binod Kumar, Subodha Charles, Hasini Witharana and Prabhat Mishra, Network-on-Chip Trust Validation using Security

Assertions, Springer Journal of Hardware and Systems Security (HASS), volume 6, pages 79-94, December 2022.

- Chamika Sudusinghe, Subodha Charles, Sapumal Ahangama and Prabhat Mishra, Eavesdropping Attack Detection using Machine Learning in Network-on-Chip Architectures, IEEE Design & Test (D&T), volume 39, issue 6, pages 28-38, August 2022.
- Subodha Charles, Vincent Bindschaedler and Prabhat Mishra, Digital Watermarking for Detecting Malicious Intellectual Property Cores in NoC Architectures, IEEE Transactions on Very Large Scale Integration (VLSI) Systems (TVLSI), volume 30, issue 7, pages 952-965, July 2022.
- Subodha Charles and Prabhat Mishra, A Survey of Network-on-Chip Security Attacks and Countermeasures, ACM Computing Surveys (CSUR), 54(5), article 101, pages 1-36, 2021.
- Subodha Charles and Prabhat Mishra, Reconfigurable Network-on-Chip Security Architecture, ACM Transactions on Design Automation of Electronic Systems (TODAES), 25(6), article 53, pages 1-25, 2020.
- Subodha Charles, Yangdi Lyu and Prabhat Mishra, Real-time Detection and Localization of Distributed DoS Attacks in NoC based SoCs, IEEE

Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD), 39(12), pages 4510-4523, 2020.

15. One of my recent inventions received an Invention of the Year Award from the University of Florida. I am a named inventor of 13 US patents (15 more patents pending) corresponding to the areas of my professional work. Accordingly, I have a general understanding of the U.S. patent prosecution process and of the novelty and non-obviousness requirements for patentability. Examples of patents where I am a named inventor include:

- United States Patent Number 11,829,475. Prabhat Mishra and Zhixin Pan. Hardware Trojan Detection using Reinforcement Learning, September 3, 2024.
- United States Patent Number 11,775,693. Prabhat Mishra, Zhixin Pan and Jennifer Sheldon. Hardware Trojan Detection Using Path Delay Based Side-Channel Analysis and Reinforcement Learning, October 3, 2024.
- United States Patent Number 11,829,475. Prabhat Mishra, Zhixin Pan and Jennifer Sheldon. Hardware-Assisted Malware Detection using Explainable Machine Learning, November 28, 2023.
- United States Patent Number 12,282,543. Prabhat Mishra, Subodha Charles, and Vincent Bindschaedler. Securing On-Chip Communication using Digital Watermarking, April 22, 2025.

- United States Patent Number 11,580,265. Prabhat Mishra and Yandi Lyu. Delay-based Side-channel Analysis for Trojan Detection, February 14, 2023.
- United States Patent Number 11,593,298. Prabhat Mishra and Subodha Charles. Reconfigurable Network-on-Chip Security Architecture, February 28, 2023.
- United States Patent Number 11,770,399. Subodha Charles and Prabhat Mishra. Lightweight and Trust-aware Routing in NoC based SoC Architectures, September 26, 2023.
- United States Patent Number 11,552,782. Prabhat Mishra, Subodha Charles and Yangdi Lyu. Securing System-on-Chip using Incremental Cryptography, January 10, 2023.
- United States Patent Number 11,797,667. Prabhat Mishra, Subodha Charles and Yangdi Lyu. Real-Time Detection and Localization of DoS Attacks in NoC based SoC Architectures, October 24, 2023.
- United States Patent Number 11,579,185. Prabhat Mishra and Yangdi Lyu. Maximization of Side-Channel Sensitivity for Trojan Detection, February 14, 2023.

- United States Patent Number 11,568,046. Prabhat Mishra and Yangdi Lyu. Trigger Activation by Repeated Maximal Clique Sampling, January 31, 2023.
- United States Patent Number 11,144,648. Swarup Bhunia, Prabhat Mishra and Jonathan Cruz. Trojan Insertion Tool, October 12, 2021.
- United States Patent Number 7533294. Prabhat Mishra and Nikil Dutt. Functional Coverage driven Test Generation for Validation of Pipelined Processors, May 12, 2009.

16. I believe that my extensive industry experience and educational background qualify me as an expert in the relevant field of the '016 Patent which includes secure management of embedded computing devices.

17. I am knowledgeable of the relevant skill set (or level of skill) that would have been possessed by a hypothetical person of ordinary skill in the art ("POSITA") at the time of the claimed invention of the '016 patent, which I have treated as October 2003 for purposes of this proceeding. I have the necessary expertise to provide my informed opinions in this IPR, and I am an expert in the art of the '016 Patent, based on my experience as a whole including the following:

- I received my undergraduate, master's, and doctoral degrees in computer science and engineering, computer and information technology, and information and computer science, respectively;
- I have extensive experience in designing secure and trustworthy computing systems. I have been conducting state-of-the-art research in the use of secure and trustworthy computing for a variety of disciplines, including embedded systems, computer architecture, and on-chip communication for the last two decades;
- I have been recognized by both IEEE and AAAS as a fellow for my contributions in system-on-chip verification and design automation of embedded systems;
- I have extensive experience in hardware/software co-design and verification in both academia and industry; and
- I have supervised 15+ doctoral students in computer science and related fields during the last 20+ years and six of them are tenure-track/tenured faculty members in top-tier universities (e.g., University of Florida, University of Texas at Dallas, etc.). The other ones are currently employed by companies, such as Apple, Amazon, Google, Nvidia, and Intel.

II. SUMMARY OF MATERIALS REVIEWED AND CONSIDERED

18. The opinions contained in this Declaration are based on the documents I reviewed and my knowledge and professional judgment. In forming the opinions expressed in this Declaration, while drawing on my experience, I reviewed every document cited in my declaration including the following documents:

	Description
Paper 2	Petition For Inter Partes Review
Paper 7	Institution Decision
EX1001	United States Patent No. 8,474,016 (“the ‘016 Patent”)
EX1003	Declaration of Dr. Sigurd Meldal (“Meldal Decl.”)
EX1004	Prosecution history of the ‘016 patent (‘016 File History)
EX1005	U.S. Pat. Pub. No. 2003/0226015 (“Neufeld”)
EX1006	U.S. Pat. No. 7,085,385 (“Frantz”)
EX1007	U.S. Patent Publication No. 2003/0097503 (“Huckins”)
EX1008	“PCI-to-PCI Bridge Architecture Specification”, Rev. 1.1 (“PCI Bridge Spec”)
EX1009	“Datasheet PCF8584 I2C-bus controller” (“PCF8584”)
EX1020	“IPMI – Intelligent Platform Management Interface Specification V1.5” (Rev. 1.1 dated 2/20/2002) (“IPMI”)
EX2009	“I2C Bus Specification” (Version 2.1, January 2000)
EX2010	“A Basic Guide to I2C” (November 2002, Joseph Wu, Texas Instruments)
EX2011	“PCI Local Bus Specification” (Revision 2.3, March 29, 2002)

	Description
EX2012	Tanenbaum, A. S., & Austin, T. “Structured Computer Organization 6/e”. Pearson Higher Education.

19. My opinions are additionally guided by my appreciation of how a person of ordinary skill in the art (POSITA) would have understood the claims of the ‘016 patent at the time of the invention, which I have been asked to assume is October, 2003. For the purpose of this declaration, I will use the definition of POSITA as provided by Dr. Meldal (Ex 1003) and described in Section IV.

III. LEGAL STANDARDS

20. I am not an attorney. However, I have been informed by counsel of certain legal standards and have applied those standards, described below, in formulating my opinions.

21. I understand that the claims of a patent define the limits of a patentee’s exclusive rights. I understand that to determine the scope of the claimed invention, courts may construe claim terms, the meaning of which the parties may dispute. I understand that claim terms should generally be given their ordinary and customary meaning as understood by a person of ordinary skill in the art at the time of filing of the patent application in light of the specification.

22. I understand that claims must be construed in light of and consistent with the intrinsic evidence. In this context, I understand that intrinsic evidence includes the

claims themselves, the written disclosure in the specification, and the patent's prosecution history. I understand that the specification is always highly relevant to the claim construction analysis and often is the single best guide to the meaning of a disputed term. I understand that extrinsic evidence may also be considered when construing claims and may include, for example, technical dictionaries, technical publications and books, treatises, and expert testimony, though the intrinsic record will be more persuasive as to the meaning of a particular term.

23. I understand that one or more prior art references can render a patent claim obvious to a person of ordinary skill in the art if the differences between the subject matter set forth in the patent claim and the prior art are such that the subject matter of the claim would have been obvious at the time the claimed invention was made. In analyzing obviousness, I understand that it is important to consider the scope of the claims, the level of skill of a person of ordinary skill in the relevant art, the scope and content of the prior art, the differences between the prior art and the claims, and "secondary considerations" of non-obviousness, which I also understand can include, for example, evidence of commercial success of the invention, evidence of a long-felt need that was solved by an invention, evidence that others copied an invention, or evidence that an invention achieved an unexpected result.

24. 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.

25. 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 understand that the mere fact that a certain thing may result from a given set of circumstances is not sufficient to establish inherency. Rather, the limitation at issue necessarily must be present in order to be inherently disclosed by the reference. Therefore, inherency cannot be established by the mere fact that something may happen. Instead, to be inherent, the alleged feature must necessarily be present in the cited reference.

26. 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”).

27. 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.

28. 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. And I 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.

29. 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.

30. I understand that the claimed subject matter can be obvious when a person of ordinary skill in the art would have been motivated to combine pre-existing elements, and that combination yields no more than what a person of ordinary skill in

the art would have expected. I also understand that, in assessing whether a claim is obvious, one must consider whether the claimed improvement is more than the predictable use of prior art elements according to their established functions. I understand that there need not be a precise teaching in the prior art directed to the specific subject matter of a claim because one can take account of the inferences and creative steps that a person of ordinary skill in the art would employ. However, I further understand that obviousness cannot be based on the hindsight combination of components selectively culled from the prior art.

31. I understand that for a reference to qualify as a relevant prior art reference in an obvious analysis, the reference must be “analogous”—meaning it is from the same field of endeavor as the patented technology, and if not, then the reference must be reasonable pertinent to the particular problem facing the inventors of the patent at issue. I also understand that a reference is reasonably pertinent if a person of ordinary skill in the art would have reasonably consulted those references and applied their teachings in seeking a solution to the problem facing the inventors.

32. I understand that, in this proceeding, Petitioners have the burden of proving that the challenged claims of the ‘016 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.

IV. LEVEL OF SKILL IN THE ART AND PERSPECTIVE APPLIED IN THIS DECLARATION

33. I understand that certain issues relating to validity must be judged from the perspective of a person of ordinary skill in the art (“POSITA”). I further understand that factors that may be considered in determining the level of skill of a person of ordinary skill in the art include, but are not limited to, (1) the type of problems encountered in the art; (2) prior art solutions to those problems; (3) rapidity with which innovations are made; (4) sophistication of the technology; and (5) the educational level of active workers in the field.

34. I have been informed and understand that a person of ordinary skill in the art is a hypothetical person who is presumed to be aware of all pertinent prior art, thinks along conventional wisdom in the art, and is a person of ordinary creativity. Importantly, this hypothetical person is of ordinary skill in the art to which the claimed subject matter pertains.

35. With respect to the ‘016 patent, as per Dr. Meldal’s declaration, “A person of ordinary skill in the art (“POSITA”) at the time of the alleged invention of the 016 patent would have had a Bachelors’ degree in electrical or computer engineering or a comparable field of study, plus approximately two to three years of professional experience in the field of computer and/or network design and security or other relevant industry experience. Additional graduate education could substitute

for professional experience, and significant experience in the field could substitute for formal education.” EX1003, ¶56.

36. I accept Dr. Meldal’s definition of a POSITA for purposes of these proceedings, and have applied the above-defined perspective of a person of ordinary skill in the art, as of October, 2003 in forming and expressing my opinions. Unless otherwise stated, my testimony below refers to the knowledge of one of ordinary skill in the art at that relevant time.

37. While my own level of skill exceeded that of the ordinary level of skill as of October, 2003, I am well-acquainted with the actual performance and capabilities of a person of ordinary skill in the art as defined above. For example, during the relevant timeframe, I was a senior graduate (Ph.D.) student and supervising and evaluating the performance of such individuals, such as junior graduate students who were performing research in this area. Furthermore, my participation in professional organizations and in conferences and other events exposes me to the state of the industry and the level of experience of individuals who, like I did, had at least the ordinary level of skill.

V. CLAIM CONSTRUCTION

38. In my analysis, I have applied the plain and ordinary meanings of the claim terms in the ‘016 Patent, as a person of ordinary skill in the art would have understood them in the context of that patent in October, 2003. And it is my opinion

that the prior art set forth in the Petition fails to disclose or render obvious the subject matter of the challenged claims under any reasonable construction and interpretation of the challenged claims' terms, including their plain meanings, as well as the constructions Petitioner's expert, Dr. Meldal, proposes.

VI. DISCUSSION OF CERTAIN PRIOR ART RELIED ON IN PETITION

39. I2C (Inter-Integrated Circuit) is a two-wire, serial communication protocol that allows devices (master and slave) to communicate with each other over a shared bus. It uses two wires, Serial Data (SDA) and Serial Clock (SCL), to transmit and receive data (Ex. 2009). As shown in the figure below, two devices can communicate using a simple 2-wire bus. EX2009, 8. Both SDA and SCL are bi-directional lines, connected to a positive supply voltage via a current-source or pull-up resistor. Data on the I2C-bus can be transferred at rates of up to 100 kbit/s in the standard mode. The number of interfaces connected to the bus is solely dependent on the bus capacitance limit of 400 pF. EX2009, 8. It is designed for short-distance communication within an embedded system. It is not used for long-distance communication since longer distance increases the bus capacitance. Moreover, signal noise and interference can degrade the I2C signals, leading to communication errors. The I2C bus neither requires encryption nor provides any support for encryption, authentication, or any security functionality.




Fig.3 Connection of Standard- and Fast-mode devices to the I²C-bus.

40. A POSITA would have known that I2C bus was commonly used in embedded systems for connecting sensors, displays, and other peripherals. In fact, major semiconductor companies still use I2C bus to connect simple components for short-range communication. For example, “A Basic Guide to I2C” from Texas Instruments observes “I2C is a two-wire serial communication protocol using a serial data line (SDA) and a serial clock line (SCL). The protocol supports multiple target devices on a communication bus and can also support multiple controllers that send and receive commands and data. Communication is sent in byte packets with a unique address for each target device.” Ex. 2010, 3. This guide also provides examples of how to connect with I2C bus from Digital-to-Analog Converter (DAC) (EX2010, 9) and Analog-to-Digital Converter (ADC) (EX2010, 11). Therefore, a POSITA would have known that such a simple bus is ideal for connecting components in embedded systems, such as sensors and actuators. In other words, a POSITA will not think of

utilizing I2C bus as the “first bus” (modified Figure 12-1, page 73 of EX1003) that has significant communication requirements.

41. The PCI Bus Specification (Ex. 2011, Revision 2.3, March 29, 2002) was the most widely used PCI bus protocol (state-of-the-art) at the time of the ‘016 patent (October 2003). As per the specification, “The bus is intended for use as an interconnect mechanism between highly integrated peripheral controller components, peripheral add-in cards, and processor/memory systems.” Ex. 2011, 15. The following figure shows an overview of PCI-based system in the PCI Bus Specification. For example, here are the sample steps for playing a video in the display attached to the graphics card. Assume that the video data is in the main memory in a specific format (e.g., MP4). The processor utilizes the bridge controller to initialize the PCI bus. The processor then sends commands to configure the graphics card’s registers to set display modes and other settings. The processor next fetches the video data from the main memory using the memory controller. Next, the processor transfers the video data to the graphics card memory (e.g., frame buffer) via the PCI bus. The hardware in the graphics card (e.g., graphics processing unit) then reads the video data from the frame buffer in the graphics card and generates the video signals for the display monitor.




Figure 1-2: PCI System Block Diagram

42. A POSITA will understand that the PCI bus is designed for a processor to communicate with peripherals, such as audio/video devices, hard disk, graphics cards, etc. The same understanding is also portrayed by widely used undergraduate textbooks, such as *Structured Computer Organization* by Andrew S Tanenbaum (sixth edition). Section 2.4.1 of the textbook describes PCI buses and typical usage. (EX2012). Unlike PCIe, which uses point-to-point connections, legacy PCI is a shared bus, meaning all devices on the bus share the same physical wiring and signaling time. The PCI bus provides a shared communication medium for the processor to interact with PCI devices and not primarily designed for point-to-point communication between PCI devices. Although PCI devices can communicate through Direct Memory Access (DMA), it will involve raw data and can lead to slower performance. Moreover, PCI has a limited bandwidth compared to PCIe, which can be a bottleneck for data intensive operations. Therefore, a POSITA would not have thought of

utilizing PCI bus as the “second bus” (modified Figure 12-1, page 73 of EX1003) to enable point-to-point communication.




Figure 2-31. A typical PC built around the PCI bus. The SCSI controller is a PCI device.

VII. SPECIFICATION OF '016 PATENT

A. Background

43. The '016 Patent is directed to the needs of network administrators to access network elements for regular maintenance or “to perform trouble shooting and resolving problems.” EX1001, 1:39-43. In discussing the state of the art, the '016 Patent identified two classes of remote management systems, namely, “in-band” systems and “out-of-band” systems. As explained in the '016 Patent:

... A number of means have been developed to provide for remote management of these systems. Remote management of the elements can be provided in-band (the remote administrator communicates with the system using the same network as the user data for the managed system) or out-of-band (the remote administrator communicates with the system using a means other than the network utilized by the user data of the managed system). *Id.*, 1:53-60 (emphasis added).

44. Next, the '016 Patent identified some of the pros and cons of “in-band” systems and “out-of-band” systems:

... One of the primary purposes of the remote management connection is to assist the remote administrator or technician in troubleshooting network problems. With in-band management, if a network problem has hindered connectivity to a network element, management connectivity to that element could be lost when it is needed the most. An out-of-band management solution is more likely to allow the administrator or technician to still remotely access the network element to troubleshoot and resolve the network problem in a timely manner. *Id.*, 2:35-44.

45. While out-of-band systems had advantages, the '016 Patent explained that the cost of dedicated and secure out-of-band systems was generally prohibitive:

... While a dedicated and secure out-of-band network would be the most preferable solution for out-of-band management from a security standpoint, the cost of such a solution is generally prohibitive. While some form of public shared network, such as the Public Switched Telephone Network (PSTN) or an Integrated Services Digital Network (ISDN) provides the most cost effective solution for a diverse out-of-band connection, the security of such solutions is a major concern. *Id.*, 2:52-60.

46. The ‘016 Patent also explained that “most protocols used for assisting in the remote management of network elements do not provide for the confidentiality or integrity of the information being transmitted between the remote administrator and the network element For instance, the protocol most frequently utilized for remote login to network elements (Telnet) transmits traffic in the clear (anyone who can tap into or sniff the network can capture and understand the traffic).” *Id.*, 3:4-14.

47. Next, the ‘016 Patent discussed that in order to control the costs of in-band remote management solutions, “user traffic and management traffic are being commingled at multiple locations throughout the management path. ... User traffic and device management traffic comes in over the same user interface, uses the same memory and buffers, and is processed by the same processor. The commingling of user traffic and management traffic can compromise the security of the device management.” *Id.*, 3:20-29.

48. Given all of the concerns with existing remote management solutions, the ‘016 explained that “[i]t would take a number of different devices configured to

work together to address most of the concerns. This would require a number of additional devices in an environment where rack space is very expensive.” *Id.*, 4:4-8.

B. Objects Of The Invention

49. After discussing the shortcomings of existing solutions, the ‘016 Patent identified several “objects of the invention.” In connection with “in-band” remote management, the ‘016 Patent stated:

[An] object of the invention is to define a virtual management interface for controlling management traffic that will flow over the in-band interfaces. The virtual management interface provides for logical separation of the management data from the user data even when the management data and the user data will transit the same physical network. *Id.*, 4:23-28 (emphasis added).

* * *

Yet another object of the invention is to be able to securely manage the device through in-band connections to the virtual management interface, the network enabled management connection, or the console port. *Id.*, 5:13-16 (emphasis added).

50. Other objects of the invention included providing for the “confidentiality and integrity of the information transmitted between the remote administrator and the management interfaces.” *Id.*, 4:49-52.

C. Summary Of The Invention

51. The disclosure of the ‘016 Patent includes numerous embodiments, many of which are not claimed in the ‘016 Patent. Aspects of the Summary of the

Invention pertinent to the claims of the '016 Patent described, *inter alia*, a secure chipset with a separate processor for implementing remote management functions:

The present invention addresses the foregoing problems, as well as other problems, by providing an exemplary embodiment of a Secure Management Access Control for Computer Chipset (SMACC) for inclusion in devices that are to be enabled for remote management. In this preferred exemplary embodiment, the SMACC functions are implemented on a separate processor with separate flash and memory; ...

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. Within the device, the SMACC sets up a separate processor for receiving management information and interacting with the control functions of the device. Remote management functions will pass through the SMACC processor. *Id.*, 5:48-6:11 (emphasis added).

52. Aspects of the Summary of the Invention pertinent to the claims of the '016 Patent also described an in-band path that used both the user data interface and the user network of the managed device for remote management of such device:

... A VMI [virtual management interface] 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. *Id.*, 6:21-29 (emphasis added).

53. In addition, the aspects of Summary of the Invention described the use of encryption to preserve the confidentiality and integrity of the management traffic:

... The VPN tunnels also utilize encryption to provide for the confidentiality and integrity of the management traffic utilizing the network. *Id.*, 6:66-7 (emphasis added).

D. Detailed Description Of The Invention

54. Figure 1 (reproduced below with annotations) illustrates one embodiment of “the positioning of the SMACC chipset within the managed device.”

Id., 11:1-2.




Fig. 1

In the annotated version of Figure 1 above, the components of the “SMACC chipset” are positioned in the red dotted box, and the components of the “managed device” are positioned in the blue dotted box. The SMACC chipset includes, *inter alia*, a separate SMACC processor (thick black box labelled “SMACC”) and a separate SMACC bus (green box). The managed device includes, *inter alia*, a system CPU (pink box), a CPU bus (purple box), a system bus (e.g., a PCI bus) (orange box), and a user data interface (brown box). The separation between the components of the

SMACC chipset and the components of the managed device is also illustrated in Figure 2 (reproduced below with colored annotations), which illustrates the SMACC chipset (red box) and managed device (blue box) each comprising a separate circuit with a separate power connection from a power supply.





Fig. 2

In discussing Figure 2, the ‘016 Patent refers to the managed device as “the SMACC enabled device,” and differentiates between the SMACC processor (thick black box in figure above) and the “main processor of the SMACC enabled device” (pink box above labelled “CPU”). *Id.*, 11:30-33.

55. The Detailed Description explained that Figures 1, 14 and 28-29 illustrated different architectures for implementing the invention:

FIG. 1 illustrates the positioning of the SMACC chipset within the managed device. In a preferred exemplary embodiment, the SMACC processor has its own flash, Non-Volatile Random Access Memory (NVRAM), and Random Access Memory (RAM) for operating, however, an alternative exemplary embodiment as illustrated in FIG. 14 is also possible in which the SMACC processor utilizes the same flash, NVRAM, and RAM chips as the main processor ... In these exemplary embodiments, it may be necessary to utilize a controller chip for the SMACC processor to access the necessary bus as illustrated in FIGS. 28 and 29. *Id.*, 11:1-13 (emphasis added).

56. The architectures shown in Figures 28 and 29 (reproduced below with annotations) correspond to examples of subject matter claimed in the '016 Patent.



57. Figures 28 and 29, the SMACC processor, SMACC Bus and Bus Controller (red boxes) are components of the SMACC chipset used to remotely manage the “managed device.” In the annotated figures, a component of the managed device (either its System PCI Bus in Figure 28 or its System CPU Bus in Figure 29)

is shown in the blue boxes. While in Figure 28, the System (PCI) Bus is the only illustrated component of the managed device, a person of ordinary skill would have understood that the System PCI Bus was part of the managed device as a whole which included other components such as the CPU and user data interface of the managed device illustrated in Figures 1 and 2. For example, in Figure 2 (reproduced again below), the user data interface (brown box) of the managed device is shown as directly connected to the System Bus (e.g., System PCI Bus) (orange box) of the managed device, and the CPU (pink box) of the managed device is shown as having an indirect connection to the System Bus via the System Controller:





Fig. 2

58. Set forth below is an illustration created by Patent Owner in which original Figure 28 (which illustrated only the System PCI Bus (orange box) in the

managed device) has been supplemented to include two other components of the managed device, namely, the user data interface (brown box) which, as in Figures 1 and 2, is shown as directly connected to the System Bus (e.g., System PCI Bus) (orange box) of the managed device, and the CPU (pink box) of the managed device which, as in Figures 1 and 2, is illustrated with an indirect connection to the System PCI Bus:



Support for the illustration above in which original Figure 28 is shown in conjunction with the user data interface and CPU of the managed device can be found throughout the specification. See, e.g., *id.*, 5:52-6:11 (describing implementation of SMACC functions on a separate processor); 9:10-17 (discussing connection of SMACC chipset to managed device including connection via a Peripheral Component Interconnect (PCI) bus); 10:39-42 (describing SMACC processor that uses controller chip to access the System PCI bus of the managed device in Figure 28); 10:43-47 (describing SMACC processor that uses controller chip to access the System CPU bus of the managed device in Figure 29); 11:9-13 (describing use by the SMACC processor of controller chip to access the necessary bus); and 13:5-7 (explaining that “it is likely that the SMACC card would not be on the same bus as the main CPU.”)

59. The Detailed Description also disclosed embodiments where a remote administrator remotely manages the managed device using in-band communications, such that management traffic goes over the same physical network as user data and passes through the same user data interface (brown box in figures shown above) of the managed device:

The Virtual Management Interface (VMI) logically separates management traffic from user data when they are going to utilize the same interfaces on the managed device and the same physical network. ... The VMI allows any of the user interfaces to be utilized for management traffic while still protecting the management interface and traffic. *Id.*, 12:10-24 (emphasis added).

* * *

The use of the VMI for connecting to the SMACC processor for management of the device allows the management data to [be] logically isolated from user data, but still physically use the user data network. ... The normal operational method of reaching the SMACC processor would be through the in-band network to the VMI, and only then a path to the VMI is not available would the secondary network need to be utilized. Since the secondary network often has lower performance characteristic, it is likely that the in-band path through the VMI will also provide for a faster connection when the in-band data network connections are available. *Id.*, 17:64-18:10 (emphasis added).

VIII. PROSECUTION HISTORY OF ‘016 PATENT

60. 60. The ‘016 Patent matured from U.S. Patent Appl. No. 11/946,976, filed November 29, 2007. As originally filed, independent claim 1 was generic in that it covered both in-band and out-of-band management systems. EX1004, 33. In addition, original claim 1 did not require a management processor (e.g., the SMACC processor) to be distinct from the processor (e.g., CPU) of the managed device. *Id.*

A. First Office Action

61. In the First Office Action, the Examiner rejected all of the claims as anticipated by U.S. Patent No. 6,859,841 (Narad). *Id.*, 116-126. In response, the Applicant amended the independent claims to require, *inter alia*, that there be a management processor (e.g., the SMACC processor) “separate from” the processors

of the managed device (referred to in the amended claim as “said computer system’s” processors.) *Id.*, 164. In the remarks, the Applicant distinguished Narad as follows:

... Applicant’s claims have been amended to make it clear that applicant is referring to a management apparatus dedicated to device management and implemented with its own dedicated processor. ...

Narad column 1, lines 51-57 mentions network infrastructure applications in general and column 4, lines 12-39 discloses methods of accelerating network infrastructure applications, but neither disclosure deals with any specific application such as controlling device management or device management functions communicating with administrators or services in the network. *Id.*, 174-175 (emphasis added).

B. Second Office Action

62. In the Second Office Action, the Examiner rejected claim 1 as obvious over Narad in view of U.S. Patent 6,532,507 (Falik). *Id.*, 196-198. The Examiner acknowledged that Narad was “silent on the capability of showing one or more management processors for controlling access to said device management function,” but reasoned that this limitation 1 was obvious over Falik. *Id.*, 198. In response, the Applicant amended claim 1 and added new claim 34, which recited:

34. (New) A management device, comprising:
a processor configured to control one or more functions of a network device based on management requests from a remote administrator, wherein the network device is configured [to] handle data requests received via a network interface of the network device, and wherein the network device is configured to receive, via the network interface, an encrypted form of the management requests; and
a bus controller configured to receive the encrypted form of the management requests from a system bus of the network device, and to convey the encrypted form of management requests to the processor;
wherein the processor is configured to decrypt the encrypted form of the management requests. *Id.*, 243.

63. In contrast to original claim 1, new claim 34 was directed to in-band remote management. More specifically, claim 34 recited a “network device” (the managed device) having its own network interface (“a network interface of the network device”) and its own system bus (“a system bus of the network device”), wherein the management requests from the remote administrator were “received via [the] network interface of the network device.” Claim 34 was limited to in-band

remote management because the claim required the management requests to be received via the network interface of the network device (the managed device). In the remarks, the Applicant explained how new claim 34 distinguished over the cited references. *Id.*, 247-248.

C. Third Office Action - Restriction Requirement

64. The Third Office Action was a three-way Restriction Requirement. *Id.*, 264-265. In response, the Applicant elected the third group of claims identified in the Restriction Requirement, i.e., claims 34-41. *Id.*, 281. At the same time, the Applicant amended claim 34 to require that the “bus controller” be “coupled to a system bus of the network device [the managed device]”:

34. (Currently Amended) An apparatus management—device, comprising:

a processor configured to control one or more functions of a network device having a network interface based on management requests from a remote administrator, wherein the network device is configured to receive handle data requests and an encrypted form of management requests received via [[a]] the network interface of the network device, and wherein the management requests are from a remote administrator ~~network device is configured to receive, via the network interface, an encrypted form of the management requests~~; and

a bus controller coupled to a system bus of the network device, wherein the bus controller is configured to receive the encrypted form of the management requests from [[a]] the system bus of the network device, and to convey the encrypted form of the management requests to the processor;

wherein the processor is configured to decrypt the encrypted form of the management requests. *Id.*, 289-290.

D. Fourth Office Action

65. In the Fourth Office Action, the Examiner rejected claim 34 as obvious over U.S. Patent No. 6,954,864 (Schelling) in view of U.S. Patent No. 7,120,692 (Hesselink). *Id.*, 289-290. In response, the Applicant further amended claim 34 as follows:

34. (Currently Amended) An apparatus, comprising:

a processor configured to control one or more functions of a network device having a network interface, wherein the network device is configured to receive data requests and an encrypted form of management requests via the network interface, wherein the management requests are from a remote administrator; ~~and~~


a management bus; and

a bus controller coupled to the processor via the management bus and coupled to a system bus of the network device distinct from the management bus, wherein the bus controller is configured to receive the encrypted form of the management requests from the system bus, and to convey the encrypted form of the management requests to the processor via the management bus;

wherein the processor is configured to decrypt the encrypted form of the management requests. *Id.*, 334.

66. Elements of amended claim 34 can be mapped to Figure 28 of the '016 Patent as shown below. Figure 28 is on the left, and corresponding elements from amended claim 34 are on the right. Here, claim 34 was amended to recite a “management bus” (purple box) that was “distinct from” the “system bus” (orange box) of the “network device.” Claim 34 was also amended to require that the bus controller be coupled to “the processor via the management bus” (yellow

connections) and also coupled to the “system bus” of the network device (green connection).



E. Advisory Action

67. The Examiner issued an Advisory Action stating that the last set of claim amendments raised new issues that would require further consideration and/or search, and asserting that there was no support in the specification for “a management bus.” *Id.*, 349-350.

F. Fifth Office Action

68. The Applicant filed a Request for Continued Examination and resubmitted the previous claim amendments. *Id.*, 353-366. In that submission, the

Applicant also explained where the “a management bus” was supported in the specification:

... the Examiner asserts that the newly recited “management bus” is not supported by the Specification. Applicant respectfully disagrees. Paragraph [0092], for example, indicates that, in some embodiments, “it may be necessary to utilize a controller chip for the SMACC processor to access the necessary bus as illustrated in FIGS. 28 and 29.” Figs 28 and 29 depict this controller chip and the processor as being coupled via an SMACC interface bus. With respect to the term “management,” Applicant notes that this term is supported by the SMACC acronym, which is defined in paragraph [0040] as “Secure Management Access Control for Computer Chipset (SMACC)” (emphasis added). Still further, paragraph [0042] describes the SMACC “processor [as] receiving management information and interacting with the control functions of [a] device.” Applicant therefore submits that the term “management bus” is supported by at least these portions of the Specification. *Id.*, 362 (emphasis added).

69. The Examiner responded with a further non-final rejection, rejecting all claims on the ground that the specification failed to provide support for “a management bus.” *Id.*, 375-376.

G. Sixth Office Action

70. In the next response, the Applicant amended claim 34 as shown below by replacing “a management bus” with “a first bus,” and replacing “a system bus” with a “second bus.”

34. (Currently Amended) An apparatus, comprising:
a processor configured to control one or more functions of a network device having a network interface, wherein the network device is configured to receive data requests and an encrypted form of

management requests via the network interface, wherein the management requests are from a remote administrator;

a ~~management~~ first bus; and

a bus controller coupled to the processor via the ~~management~~ first bus, wherein the bus controller is also coupled to a ~~system~~ second bus of the network device that is distinct from the ~~management~~ first bus, wherein the bus controller is configured to receive the encrypted form of the management requests from the ~~system~~ second bus, and to convey the encrypted form of the management requests to the processor via the ~~management~~ first bus;

wherein the processor is configured to decrypt the encrypted form of the management requests. *Id.*, 388.

71. In response to the Examiner's written description rejection, the Applicant stated:

Applicant respectfully submits that a "first bus" and a "second bus" are supported by the Specification. For example, Fig. 28 depicts an "SMACC Bus" in addition to a "System PCI Bus"; Fig. 29 depicts an "SMACC Bus" in addition to a "System CPU Bus." Various forms of the "SMACC Bus" are also depicted with other buses in Figs. 1, 2, 9-14, 17-20, and 23-27. Paragraph [0092] recites "[i]n these exemplary embodiments, it may be necessary to utilize a controller chip for the SMACC processor to access the necessary bus as illustrated in FIGS. 28 and 29." Paragraph [0107] recites "it is likely that the SMACC card would not be on the same bus as the main CPU." Accordingly, these portions support the claims recitation of the "first bus" and "second bus." *Id.*, 393.

72. The Examiner responded with a Final rejection, in which claim 34 was rejected as obvious over Schelling (US 6,954,864), in view of Watanabe et al (US 7,058,740), and further in view of Hesselink (US 7,120,692). *Id.*, 404-406.

H. Response After Final Rejection and Notice of Allowance

73. In response to the Final Rejection, the Applicant amended claim 34 as shown below to further recite that “the network device includes a processor configured to facilitate operation of the network device” (e.g., the network device includes a CPU), and that “the processor of the apparatus [e.g., the SMACC processor] is distinct from the processor included in the network device.” *Id.*, 428.

74. The Examiner responded with a Notice of Allowance. *Id.*, 440-441.

75. Upon issuance, amended claim 34 of the Application became claim 1 of the ‘016 Patent.

IX. CLAIMS OF ‘016 PATENT

76. Claim 1 of the ‘016 Patent recites:

1. An apparatus, comprising:


a processor configured to control one or more functions of a network device having a network interface, wherein the network device is configured to receive data requests and an encrypted form of management requests via the network interface, wherein the management requests are from a remote administrator;

a first bus; and


a bus controller coupled to the processor via the first bus, wherein the bus controller is also coupled to a second bus of the network device that is distinct from the first bus, wherein the bus controller is configured to receive the encrypted form of the management requests from the second bus, and to convey the encrypted form of the management requests to the processor via the first bus;

wherein the processor is configured to decrypt the encrypted form of the management requests, wherein the network device includes a processor configured to facilitate operation of the network device, and wherein the processor of the apparatus is distinct from the processor included in the network device.

77. Set forth below is a diagram illustrating the hardware components recited in claim 1 and how such components are interrelated as specified in the claim. In connection with each component, corresponding language from claim 1 is set forth.



78. In addition to specifying the physical relationships between the hardware components illustrated above, claim 1 specifies that management requests issued by a network administrator traverse an in-band path to reach the network device, and thereafter traverse a specific path through the hardware components. Set forth below is a diagram illustrating the path traversed by the data requests (solid blue line) and management requests (solid red line) recited in the claim, along with the corresponding language from the claim.



79. Consistent with the portions of the Summary of the Invention and Detailed Description of the Invention discussed above, claim 1 of the '016 Patent is directed to an *in-band* remote management system because the management requests pass through the *same user data* interface as user data requests, and the management requests therefore go over the *same physical network* as the user data requests. See, e.g., EX1001, 6:21-29 (discussing use of a VMI "for logically separating management traffic from user data when the *in-band* path is to be used for management data. ... The VMI ... utiliz[es] *the user data interfaces of the managed device and the user network to provide the transport of the management data cost effectively*"; 12:22-24 ("The VMI allows any of the *user interfaces to be utilized for management traffic* while still protecting the management interface and traffic"); and 17:64-18:5 ("The

use of the VMI for connecting to the SMACC processor for management of the device allows the management data to [be] logically isolated from user data, but still physically use the user data network. ... The normal operational method of reaching the SMACC processor would be through the in-band network to the VMI...")

X. PRIOR ART RELIED ON IN THE PETITION

A. Neufeld

80. Neufeld is the primary reference for Grounds 1-8. As discussed more fully below, in contrast to the Challenged Claims which are directed to in-band remote management system, Neufeld is directed to an out-of-band remote management system.

1. Background Set Forth In Neufeld

81. In discussing the “Background of the Related Art,” Neufeld disparages in-band systems in which communication with the server is “only possible when the server is able to communicate over its normal network connection” in favor of out-of-band communication which Neufeld characterized as “a much more powerful tool for server management” which does not require the managed server to be able to communicate over its normal network communication, or to even be online. EX1005, [0011]-[0014] (emphasis added). Neufeld’s stated goal is to provide for “lights out” management, which means the capability to provide remote management even when

the managed server is unable to communicate over its normal network communication, or even offline -- a capability that is only possible with an out-of-band remote management system. *Id.*, [0011]-[0014] and [0076].

82. Having concluded that out-of-band systems were “a much more powerful tool than in-band systems,” Neufeld went on to identify a shortcoming in out-of-band systems related to security that needed to be addressed:

In a typical remote management system, a user (typically, a member of the network management team) can initiate an out-of-band session with the dedicated server management computer hosted in the managed server via a remote console application program being executed on a client computer. *Id.*, [0017]-[0018] (emphasis added).

* * *

... it is believed to be desirable to alter authentication services and possibly to restrict access rights and privileges in some instances to improve security. For example, if a network administrator logs onto the system from a physically secure location where the network administrator's identification can be determined for the network administrator even to access the facility, the amount of user authentication and the number of restrictions placed upon a user in such a location need not be as onerous as user authentication procedures and restrictions placed upon users attempting to access the server from locations outside of the physically secure facility. ...

The present invention may be directed to certain issues discussed above. *Id.*, [0022]-[0023].

83. The specification of Neufeld disclosed an out-of-band remote management system where, *inter alia*, access privileges varied depending on how far

“the access point is removed from a direct connection in the secure physical location of the managed server.” *Id.*, [0083].

2. Figure 1 of Neufeld

84. Figure 1 of Neufeld (reproduced below with annotations) depicts “a connection diagram of a managed server and a remote management console according to one embodiment.” *Id.*, [0025].





FIG. 1

In Figure 1, “a managed server 2 [is] connected to a remote console 5 by a network N [labelled in blue].” *Id.*, [0034]. A network identification card (NIC 110) in the managed server 2 serves to connect to the managed server to network N (blue). *Id.*,

Figure 2 and [0045]. A remote server management controller 116 co-located with managed device 2 separately connects to a management network M (labelled in red) using an Ethernet interface 180. *Id.*, Figures 2-3 and [0061]. An administrator operating management PC 5 sends out-of-band management communications over management network M (labelled in red) to the remote server management controller 116 co-located with managed device 2 via the Ethernet interface. *Id.*, [0062] (“The user may, thus, engage in out-of-band communication with the remote server management controller 116 for the purpose of diagnosing, correcting and/or preventing problems with the managed server 2.”)

3. Figure 2 of Neufeld

85. Figure 2 of Neufeld (reproduced below with annotations) depicts “a block diagram of the managed server according to the embodiment of FIG. 1.” *Id.*, [0026]. Consistent with Figure 1, Figure 2 depicts two distinct network connections. Specifically, out-of-band management communications (from, e.g., remote management PC 5) are sent over the management network M (labelled in red) to the remote server management controller 116 co-located with managed device 2. *Id.*, [0062]. Consistent with Figure 1, the management network M (labelled in red) is shown in Figure 2 as being different from network N (blue) which connects to the managed server 2 via NIC 110.



4. Figure 3 of Neufeld

86. Figure 3 of Neufeld (reproduced below with annotations) depicts “a block diagram of the remote management controller of FIG. 2.” *Id.*, [0027]. Consistent with the earlier figures, Figure 3 illustrates Management Network M (red) as a separate network that connects to the Remote Management Controller (RMC) 116 via an Ethernet interface 180. *Id.*, [0061] (“The Ethernet interface 180 provides the main external communication interface between the remote server management controller 116 and the outside world. ... [T]his connection is typically used to couple the remote management controller 116 to the management network M.”)

<u>Actual Wording of Neufeld</u>	<u>Petitioner's Rendition</u>
The remote management controller 116 also includes circuitry to route keystrokes to the keyboard controller 38 from either the local keyboard 52 or from the remote console 5 via the modem or NIC 110 which may be coupled to the network <u>M</u> . EX1005, [0047] (emphasis added).	The remote management controller 116 also includes circuitry to route keystrokes to the keyboard controller 38 from either the local keyboard 52 or from the remote console 5 via the modem or NIC 110 which may be coupled to the network <u>M [sic, N]</u> . Petition, 20 and 25 (emphasis added).

88. Thus, Petitioner acknowledges that the subject sentence contains an error. I agree that there is an error in the subject sentence of Neufeld because, *inter alia*, the language of the sentence is inconsistent with the figures; however, I disagree that a POSITA would have corrected the error in the manner proposed by Petitioner, i.e., by replacing “M” with “N” at the end of the sentence.

89. Neufeld ties together Figures 1, 2 and 3 by stating that “FIG. 2 is a block diagram of the managed server of FIG. 1,” and that “FIG. 3 is a block diagram of the remote management controller of FIG. 2.” EX1005, [0026]-[0027] (emphasis added). These figures collectively depict a system where the Management Network M is connected to the remote Management Controller (RMC) 116 via Ethernet interface 180 (or UART or ICMB), and separately depict Network N being connected to the managed device 2 via Network Interface Card (NIC) 110. However, the subject sentence in Neufeld mixes up this correspondence between the network M and N and their respective connections to Ethernet Interface 180 and NIC 110 when it states that

“[t]he remote management controller 116 also includes circuitry to route keystrokes to the keyboard controller 38 from ... NIC 110 which may be coupled to the network M.” EX1005, [0047] (emphasis and coloring added). I agree that the subject sentence contains an error.

90. There are at least two approaches for correcting the subject sentence to align the correspondence between networks M and N and their respective connections to Ethernet Interface 180 and NIC 110 with the figures shown in Neufeld. One approach is to change “M” to “N” as argued in the Petition. The other approach -- which is much more consistent with Neufeld as a whole -- is to change “NIC 110” to “Ethernet Interface 180.” These two alternatives are set forth below:

Erroneous Sentence From Neufeld	Petitioner’s Proposal For Correction	Patent Owner’s Proposal For Correction
The remote management controller 116 also includes circuitry to route keystrokes to the keyboard controller 38 from either the local keyboard 52 or from the remote console 5 via the modem or <u>NIC 110</u> which may be coupled to the network <u>M</u> . EX1005, [0047] (emphasis and coloring added).	The remote management controller 116 also includes circuitry to route keystrokes to the keyboard controller 38 from either the local keyboard 52 or from the remote console 5 via the modem or <u>NIC 110</u> which may be coupled to the network M <u>N</u> . Petition, 20 and 25 (emphasis and coloring added)	The remote management controller 116 also includes circuitry to route keystrokes to the keyboard controller 38 from either the local keyboard 52 or from the remote console 5 via the modem or ~~NIC 110~~ <u>Ethernet Interface 180</u> which may be coupled to the network <u>M</u> .

91. Petitioner provides no rationale why the subject sentence should be corrected in the manner suggested in the Petition. Given Neufeld's singular focus as an out-of-band remote management solution, a POSITA would have understood that remote management commands destined for the RMC 116 in Neufeld would have necessarily remained on Neufeld's Management Network M. Moving Neufeld's remote management commands to Network N as suggested by Petitioner would transform Neufeld from an out-of-band solution to an in-band solution, which makes no sense in the context of the reference as a whole which began by disparaging in-band systems in which communication with the server is "only possible when the server is able to communicate over its normal network connection" in favor of out-of-band communication which Neufeld characterized as "a much more powerful tool for server management." EX1005, [0011]-[0014] (emphasis added); *see also*, [0059] (discussing the establishment of "out-of-band communication sessions"); [0060] (discussing use of "an external communication device to engage in an out-of-band communication session with remote server management controller 116"); [0061] ("The Ethernet interface 180 provides the main external communication interface between the remote server management controller 116 and the outside world. ... For example, this connection is typically used to couple the remote management controller 116 to the management network M"); and [0062] ("The user may, thus, engage in out-of-band communication with the remote server management controller 116 for the

purpose of diagnosing, correcting and/or preventing problems with the managed server 2.”).

6. Neufeld’s Remote Server Management Controller Is Configured To Snoop On Activities Occurring At The Managed Server

92. Neufeld repeatedly emphasizes the role of RMC 116 in passively “snooping” system activity—such as display configurations and keyboard input at the managed server 2. EX1005, [0047]-[0048]. This snooping functionality is described in the text of Neufeld as well as in U.S. Pat. No. 5,898,861, which is incorporated by reference into Neufeld in connection with Neufeld’s description of its snooping functionality. Paragraph [0047] of Neufeld describes in part how the RMC 116 “snoops” on the managed server 2:

The remote management controller 116, as described in more detail below, includes circuitry for snooping the PCI bus for configuration transactions between the processor 10 and the video graphics controller 114 to determine configuration and mode information, such as whether the video graphics controller is in text or graphics mode. More specifically, the remote management controller 116 snoops indexed input/output (I/O) ports of the video graphics controller 114 to provide a set of shadow registers corresponding to mode information. These I/O ports are particularly helpful for legacy video graphics array (VGA) compatibility mode. In addition, the shadow registers of the remote management controller 116 provide a set of registers for the I/O processor 156 to access independently of the operating system running on processor 10, thereby preventing any conflicts that could arise if both processors were trying to access the indexed I/O ports simultaneously. The remote management controller 116 also snoops and stores configuration information sent by the processor 10 to the video graphics controller 114.. ... EX1005, [0047] (emphasis added).

93. A POSITA would have understood that “snooping” refers to the remote management controller’s ability to passively monitor activity on the system’s PCI bus and extract information without interfering with normal operations of the managed server. For example, when the managed server’s main processor communicates with the video graphics controller to update display settings or write screen data to memory, the remote management controller 116 “snoops” (listens in) on the data being sent across the bus to detect whether the video controller is operating in text or graphics mode and capture content written to the video frame buffer. The captured information allows the RMC 116 to reconstruct the server’s video output and transmit it to a remote console, enabling administrators to view the server’s screen even if the operating system of the managed device has crashed or has not yet loaded. Importantly, the snooping process is non-intrusive: the RMC 116 does not initiate bus transactions or modify any data—it simply observes communication. This passive design ensures that RMC 116 can operate independently of the managed server’s operating system and not interfere with its normal functions, making it well-suited for out-of-band remote management.

94. The transmission to a remote administrator of information obtained from snooping is not equivalent to the reception by the managed device 2 of management requests transmitted from the remote administrator (as recited in claim 1 of the ‘016 Patent), because these communications travel in opposite directions.

Claim 1 of the ‘016 Patent recites management requests that are received by the network device “from a remote administrator,” which travel in the opposite direction of snooping information that is transmitted by the network device and received by the remote administrator.

95. The “snooping” functionality described in paragraph [0047] of Neufeld includes the “snooping” of keystrokes, as underscored by the final sentence of that paragraph, which provides:

... This keyboard functionality is more fully explained in U.S. Pat. No. 5,898,861, entitled “Transparent Keyboard Hot Plug.” EX1005, [0047].

U.S. Pat. No. 5,898,861 (“the ‘861 patent”) uses the term “snooping” over 50 times. *See, e.g.*, EX2006, 10:13-18 (“... serial data is simply passed directly between the primary keyboard 157 and the 8042 keyboard controller 156. It is, however, “snooped” by the virtual keyboard 204, so the state of the primary keyboard 157 is thus tracked by the SMM 212 software through system management interrupt from the IRC 170 to the CPU 50.”)

96. A POSITA would have understood that the term “snooping” as used in the ‘861 Patent, refers to the capability of a system component—specifically, an integrated remote console (IRC)—to passively monitor communication between a physical keyboard and a standard 8042 keyboard controller without interfering with that communication. *See, e.g.*, EX2006, 10:8-18 and 22:9-21. The IRC monitors, or “snoops,” the serial scan codes transmitted over the clock and data lines by the

physical keyboard to track the keyboard's operational state (e.g., key presses, LED indicators, repeat rates). *Id.* This enables the system to preserve and emulate the keyboard's behavior via a virtual keyboard. As with Neufeld, the snooping described in the '861 Patent is non-intrusive; the IRC does not alter or delay the signals it observes.


97. The snooping functionality of the '861 Patent aligns with the broader snooping framework described in Neufeld, where RMC 116 observes system bus traffic. Both disclosures utilize snooping to facilitate out-of-band management, enabling remote access and control even when the operating system of the managed device is non-functional. The '861 Patent thus provides foundational support for understanding how keyboard snooping operates within Neufeld's remote management architecture.

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

98. IPMI is a standardized protocol used to remotely manage and monitor servers, even when the main OS is down or the system is powered off. To accomplish this, IPMI leverages out-of-band communication, providing several different dedicated management ports, including auxiliary IPMB connectors for communications via remote management cards, an Intelligent Chassis Management Bus (ICMB) bridge, and a serial connector. 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.

Figure 1-2, IPMI Block Diagram




99. “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. EX1020, 31.

100. In terms of in-band-communication, as can be seen from FIGs. 1-2 of IPMI, the BMC shares a network interface with the managed system's network traffic. The communication path for such in-band-communication is highlighted in orange in FIGs. 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



101. 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....” EX1020, 128. 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.” *Id.* “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.” *Id.*

C. Huckins

102. Huckins discloses that a non-PCI bus controller 126 is coupled on a PCI bus 131 to an I/O Controller hub 130. The non-PCI bus controller 126 is capable of receiving and transmitting signals to and from the system 100 on the PCI bus 131.

D. Intelligent Platform Management Interface Specification v2.0 (IPMI v2.0)

103. IPMI v2.0 is a standardized protocol that proposed additions to the IPMI specification. RMCP+ is the working name for an enhanced protocol transferring IPMI messages and other types of payloads to an IPMI-based BMC over IP. EX1023, 7.

XI. GROUND 1 FAILS TO DEMONSTRATE THE UNPATENTABILITY OF THE CHALLENGED CLAIMS BY A PREPONDERANCE OF THE EVIDENCE

104. As explained above, claim 1 specifies that management requests issued by a network administrator traverse an in-band path to reach the network device.¹ By

¹ As discussed above, claim 1 of the ‘016 Patent is directed to an in-band remote management system because the claim requires that the management requests pass through the same user data interface as user data requests, and such management requests therefore go over the same physical network as the user data requests. *See also*, e.g., EX1001, 6:21-29, 12:22-24 and 17:64-18:5.

contrast, Neufeld is directed to an out-of-band remote management system. Petitioner's assertion that Neufeld discloses in-band management requests rests on Petitioner's edited rendition of the erroneous sentence at the end of paragraph [0047] of Neufeld, discussed above. However, Petitioner has failed to explain why a POSITA reading Neufeld as a whole (without knowledge of the '016 Patent) would have interpreted the erroneous sentence as referring to in-band communications, when the singular focus of Neufeld's solution was an out-of-band system.

A. Neufeld's Managed Server Does Not Receive "In Band" Remote Management Commands

105. Petitioner's assertion that Neufeld discloses or suggests the transmission of in-band remote management commands from the remote management PC to the managed server is directly at odds with the express teachings of the reference. As described repeatedly and unequivocally throughout the specification, Neufeld is directed to an out-of-band remote management system in which the RMC 116 operates independently of the host operating system and transmits commands to external consoles via dedicated interfaces such as Ethernet interface 180, UART, or ICMB. EX1005, [0059]-[0062]. The only mention of in-band communication in Neufeld appears in the background section, where it is contrasted with out-of-band

communication to highlight the limitations of in-band communication and underscore the advantages of the out-of-band solution proposed in the reference. *Id.*, [0011]-[0012]; *see also*, Section X(A)(1), *supra*. Petitioner’s attempt to characterize Neufeld as supporting in-band management is inconsistent with the entire thrust of Neufeld.

106. Petitioner states, “Neufeld’s managed server 2 is configured to receive both data requests and in-band management requests over a Network N from a remote management PC 5 (a “remote administrator”),” Petition, 18. However, the aspect of this statement relating to “in-band” management is predicated on Petitioner’s edited rendition of the last sentence of [0047] of Neufeld. As discussed above, Petitioner’s chosen rendition of the subject sentence is fundamentally opposed to the solution disclosed in Neufeld. For example, Neufeld states:

The embedded I/ O controller 150 provides a plurality of communication interfaces that can be employed to establish out-of-band communication sessions with the remote server management controller 116. EX1005, [0059].

* * *

A user may, accordingly, use an external communication device to engage in an out-of-band communication session with the remote server management controller 116 via the UART interface 176 or the ICMB interface 178. *Id.*, [0060].


* * *

... a user may connect remotely to the remote server management controller 116 via the Ethernet interface 180. Such a connection may be made, for example, using a remote console application running on a client computer anywhere on the

network that includes managed server 2. The user may, thus, engage in out-of-band communication with the remote server management controller 116 for the purpose of diagnosing, correcting and/or preventing problems with the managed server 2. *Id.*, [0062].

107. There is no discussion in Neufeld's solution of control being performed based on transmission of "in-band" management requests to the RMC 116. Instead, Neufeld is directed to remote management of a server using out-of-band management requests.

108. As shown in the annotated version of Figure 2 reproduced below, the RMC 116 (purple box) is directly connected to the Management Network M (red).



Neufeld also explains that the RMC 116 is connected to its own network interface, Ethernet Interface 180, which is depicted as distinct from the server's primary NIC

110 (blue box) of Figure 2. The Ethernet Interface 180 allows the RMC 116 to communicate with a remote console over a Management Network M. This architecture ensures that administrators can access, monitor, and control the server via the RMC even when the host OS is unbootable or the primary network interface is non-functional. EX1005, [0034]. Neufeld describes this aspect of its solution as follows:

The logic of the remote server management controller 116 is broken down into three main functional blocks. The first of these three functional blocks is an embedded I/O controller 150, which is essentially an independent computer system that is integrated within the managed server 2. The second and third functional blocks of the remote server management controller 116 are a slave instrumentation module 152 and a remote console redirection module 154. As described below, the embedded I/O controller 150 monitors and controls a wide range of conditions in the managed server 20 via the slave instrumentation module 152 and the remote console redirection module 154. EX1005, [0053] (emphasis added).

109. A POSITA reading Neufeld would have understood that RMC 116 includes its own I/O processor, memory, and firmware, enabling it to emulate keyboard and mouse inputs, redirect video output, and process security policies autonomously. *Id.*, [0034] (discussing operation “independent of an operating system”), [0037] (discussing operation “without reliance or interference with the operating system”) and [0053] (characterizing components of RMC 116 as “essentially an independent computer system.”) RMC 116 can receive commands and authentication credentials directly from a remote console without relying on the

managed server's software stack. EX1005, [0034]. This is the essence of out-of-band management: providing a side channel that remains operational regardless of the state of the main system.

110. The only discussion in Neufeld of “in-band” communication is in the “Background of the Related Art” where “in-band” solutions are disparaged in favor of “out-of-band” solutions such as the one offered by Neufeld. EX1005, [0011]-[0012] (quoted in Section X(A)(1), *supra*). While Neufeld mentions in-band solutions as existing in the related art, Neufeld goes on to emphasize the advantages of out-of-band solutions and then focuses solely on proposing its own out-of-band solution. EX1005, [0011]-[0012]; [0059]-[0062]; *see also*, Sections X(A)(1)-(5), *supra*. Specifically, Neufeld acknowledges that in-band communication refers to accessing the server over its standard network interface, and that doing so is only possible when the operating system is online and the network stack is functional. This form of management is described as inherently limited, because it cannot be used when the server is down or otherwise unable to communicate over its primary network path. EX1005, [0011].

111. In contrast to server management based on in-band communication, Neufeld discloses server management based on out-of-band communication, which enables server management even when the OS is not operational. *Id.*, [0034]. Neufeld describes various alternative channels for server management based on out-of-band

access—such as via a modem, serial port, dedicated Ethernet interface, or other communication paths that bypass the server’s normal network stack. *Id.*, [0059]-[0062]. As discussed above, Neufeld’s RMC 116 is designed for independent out-of-band communication, and Neufeld’s RMC 116 remains operational even when the server being managed is powered off or unresponsive. Thus, while Neufeld acknowledges in-band communication as a known technique, a POSITA would have understood that Neufeld’s solution is not directed to in-band server management and instead focuses exclusively on remote server management using out-of-band management requests.

112. Since Neufeld fails to disclose server management using in-band management requests, Neufeld fails to meet the limitations of claim 1 of the ‘016 Patent specifying that management requests issued by an administrator traverse an in-band path to reach the network device.

B. Petitioner Cannot Identify a “Second Bus” in Neufeld That Is Distinct from the “First Bus”


113. As explained above, the better approach for correcting the erroneous sentence in paragraph [0047] of Neufeld (and the one that would have been taken by a POSITA) is to change “NIC 110” to “Ethernet Interface 180.” This is the better approach because it aligns with Neufeld as a whole, which is focused on transmission of out-of-band remote management requests to RMC 116. When “NIC 110” is

changed to “Ethernet Interface 180”, Petitioner’s argument that Neufeld has the “first bus” and “second bus” as recited in claim 1 evaporates.

114. As explained above, during prosecution, claim 34 (which became claim 1 of the ‘016 Patent) was amended to recite a “first bus” (e.g., the SMACC bus) that was “distinct from” the “second bus” (e.g., the PCI system bus) of the “network device.” See, Section VIII(D)-(G), *supra*. When the subject sentence in paragraph [0047] of Neufeld is read in concert with the rest of the reference (as proposed by Patent Owner) such that commands from Neufeld’s remote management PC are received by the RMC 116 via an out-of-band communication that goes through Ethernet interface 180 and is therefore communicated internally within the RMC 116 to ATB 170 through the local bus 166, Petitioner’s attempt to satisfy the “second bus” limitation by relying on either the PCI bus 172 or the internal local bus 166 of Neufeld fails. EX1005, [0056]-[0057].

115. More specifically, since internal local bus 166 is part of the RMC 116, internal local bus 166 cannot be a “second bus of the network device” distinct from the first bus, as required by the claims. Contrary to the requirement that the “second bus” be a bus “of the network device,” local bus 166 is completely inaccessible to the managed sever 2 (which the Petition maps to “the network device.”) Neufeld explains that “the internal local bus 166 exists to allow communication between and among the logical components of the embedded I/O controller 150.” EX1005, [0056].


As shown below, in Neufeld commands that are received by the Ethernet interface 180 are communicated to ATB 170 via local bus 166.



116. A POSITA reading Neufeld would have understood that the ATB 170 receives the user commands via an internal bus 166 that allows for communication within the remote server management controller 116. However, the internal bus 166 cannot be the “second bus” recited in the claims because the claims expressly require “a second bus of the network device that is distinct from the first bus.” In the case of Neufeld, the commands are not received from a second bus of the network device (i.e., managed server 2) and are instead received via a local bus 166 that is internal to RMC 116 and completely inaccessible to the managed server 2.

117. In addition, the internal bus 166 cannot be the “second bus” because the claim expressly requires “*a second bus of the network device that is distinct from the*

first bus.” Specifically, the Petition alleges that element 1.d (“a first bus”) is met because “Neufeld discloses an “internal local bus 166” that is part of the remote management controller 116.” Petition, 22. In addition, the Petition provided the following annotated version of Figure 3 (reproduced below).



118. However, claim 1 of the ‘016 Patent directly prohibits the “first bus” and the “second bus” from being the same bus. Therefore, the local bus 166 cannot satisfy both the “first bus” and the “second bus” recited in claim 1 of the ‘016 Patent.

119. When the erroneous sentence in paragraph [0047] of Neufeld is read in concert with the rest of the reference (as would have been done by a POSITA), Petitioner’s argument that Neufeld has the “first bus” and “second bus” as recited in claim 1 falls apart.

C. Petitioner’s Attempt to Identify the PCI Bus as the “Second Bus” Misreads Neufeld’s Out-of-Band Architecture and Mischaracterizes the Nature of Passive Snooping

120. The Petition alleges that the PCI bus 172 is the second bus because “the ATB 170 receives the encrypted management requests from the remote management PC 5 through the NIC 110 via the PCI bus 172.” Petition, 25. Petitioner relies on its proposed edit to the erroneous sentence in paragraph [0047] of Neufeld for this proposition, which, as discussed above, is inconsistent with Neufeld as a whole. Moreover, Petitioner’s assertion that management commands are received by RMC 116 over PCI bus 172 reflects a fundamental misunderstanding of the architecture disclosed in Neufeld. A POSITA would have understood that that the only reason Neufeld connects RMC 116 to PCI bus 172 is to allow RMC 116 to perform the passive snooping function described above. See, e.g., Section X(A)(6), *supra*. Rather than describing the receipt of management commands via the PCI bus 172, Neufeld repeatedly emphasizes the role of RMC 116 in passively “snooping” system activity—such as display configurations and keyboard input. See, Section X(A)(6), *supra*. Petitioner’s attempt to recast passive observation of activities as reception of commands from a remote administrator should be rejected. As discussed above, the snooping described in Neufeld has absolutely nothing to do with receiving management requests from a remote administrator, as recited in claim 1 of the ‘016 Patent. See, e.g., Section X(A)(6), *supra*.

D. Petitioner Misconstrues Neufeld’s Out-of-Band Architecture by Mischaracterizing the RMC’s Passive Monitoring as In-Band Command Reception

121. Petitioner fundamentally misunderstands the architecture disclosed in Neufeld by conflating the passive, out-of-band monitoring functions of RMC 116 with the active reception of encrypted management requests from a remote administrator. 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. EX1005, [0037] (“without reliance or interference with the operating system.”) Neufeld makes clear that RMC 116 receives management requests from an external, remote administrator 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.” EX1005, [0061]. Thus, any suggestion that the RMC receives encrypted in-band management commands over the NIC 110 is inconsistent with Neufeld’s disclosed solution.

122. The Petitioner is erroneously conflating receiving an “encrypted form of the management requests” with parts of the managed server that are monitored by the RMC 116. Specifically, the Petitioner is paradoxically alleging that the RMC 116 snoops on the managed server 2 to receive the “the management requests” to deliver

back to the managed server 2. However, this “Rube Goldberg” architecture is not taught by Neufeld. Instead, Neufeld teaches that “The Ethernet interface 180 provides the main external communication interface between the remote server management controller 116 and the outside world.” EX1005, [0061].

123. Relying upon paragraph [0043] of Neufeld, the Petition alleges, “Neufeld discloses and/or renders obvious to a POSITA that the remote management controller 116 receives in-band communications including management requests from the remote management PC 5 via the NIC 110.” Petition, 24-25. However, this paragraph of Neufeld states:

[0043] Further attached to the PCI bus **16** is a remote management controller **116**. The remote management controller **116** connects to the keyboard controller **38**, the network N and/or a management network M, a keyboard **52**, and a mouse **54** to provide functionality for accessing, interacting, and monitoring the managed server 2 from the remote console 5 as will be more fully described below. EX1005, [0043] (emphasis added).

124. Paragraph [0043] of Neufeld is describing the hardware-level connections that enable the RMC 116 to perform passive out-of-band monitoring of the server—not active in-band communication. Paragraph [0043] would have informed a POSITA that the RMC 116 is physically connected to system components such as the PCI bus and keyboard controller, which allows 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 NIC 110 to receive management requests through the server's primary network path; to the contrary, the Neufeld makes clear that the RMC'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 RMC 116 is able to snoop on the data being sent/received by the managed device via the network N. RMC 116 could not snoop on such traffic unless it was connected to Network N. Thus, Paragraph [0043] of Neufeld describes the hardware context for the RMC's out-of-band functionality—not a software-dependent in-band communication path.

125. Paragraph [0043] of Neufeld is simply describing that RMC 116 connects to the PCI bus in order to access, interact with and monitor the managed sever 2. This underscores the ability of RMC 116 to passively monitor, or “snoop,” system activity by virtue of its direct connections to both the PCI bus and the keyboard controller. By being attached to the PCI bus, the RMC 116 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.

126. 126. Neufeld makes clear that its management system operates independent of an operating system (“OS”) running on the managed server. For example, Neufeld states:

The managed server 2 includes special circuitry and software for capturing, analyzing, compressing and transmitting video activity to the remote console 5 independent of an operating system (“OS”). The special circuitry and software operate without regard to the existence or type of OS present on the managed server 2. Therefore, the present technique may be useful for accessing, interacting, and/or monitoring the managed server 2 from the remote console 5 even before its OS has been loaded. More specifically, the video displayed on monitor 4 is capable of being viewed on a monitor 8 independent of the OS. EX1005, [0034].

127. Neufeld further explains:

In the operation of the present technique, video data is captured, analyzed, compressed, and transmitted to the remote console 5 by circuitry and software in the managed server 2 without reliance or interference with the operating system. *Id.*, [0037].

128. Thus, a POSITA would have understood that a core tenant of Neufeld is that the remote administrator is able to access the managed server 2 independent of any software system including the operating system running on the managed server.

129. The RMC 116 requires the host server to run software in order to send data over NIC 110 because unlike Ethernet interface 180, the RMC 116 lacks direct, independent control of the NIC hardware. In typical server architectures, NIC 110 is initialized and managed by the host operating system, which configures the network interface, assigns IP addresses, and handles all protocol-level communication. If the RMC attempts to use this same NIC, it must either pass its management data through the host's networking stack or rely on a specialized sideband interface, such as NC-SI, that supports out-of-band communication. Without such hardware support, the RMC cannot independently transmit data over NIC 110, as it would lack access to essential network services like ARP resolution, routing, and protocol handling. In such cases, software on the host system is needed to coordinate or tunnel RMC communications through the NIC.

130. As a result, the RMC 116 would necessarily depend on host software to transmit data through NIC 110 to network N, rendering such communication in-band rather than out-of-band. This dependency is fundamentally at odds with the goal of Neufeld's out-of-band management, which is to maintain remote access even when the host operating system is down or unresponsive.

131. To the extent Petitioner is twisting paragraph [0043] of Neufeld to allege that RMC 116 receives management requests from the remote console 5 via NIC 110, that twisting 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.


XII. GROUNDS 2–8 FAIL BECAUSE THE ADDITIONAL REFERENCES DO NOT CURE NEUFELD'S DEFICIENCIES

132. Grounds 2 through 8 fail for the same reasons set forth in response to Ground 1. Each of these grounds relies on Neufeld as the primary reference, but as explained above, Neufeld does not disclose or suggest all of the limitations required by claim 1. The additional references cited in Grounds 2 through 8 are not alleged, nor do they purport, to disclose these missing limitations. Because Grounds 2 through 8 are premised on the same flawed interpretation of Neufeld as Ground 1, Grounds 2-8 fail to demonstrate unpatentability of the Challenged Claims.

XIII. GROUND 9 FAILS TO DEMONSTRATE THE UNPATENTABILITY OF THE CHALLENGED CLAIMS BY A PREPONDERANCE OF THE EVIDENCE

A. Petitioner's Attempt To Characterize The First Bus And Bus Controller As Part Of The "Network Device" Contradicts The Plain Language And Structural Separation Expressly Recited In Claim 1

133. In Grounds 9-15, the Petition relies on IMPI (EX1020) as the primary reference. Petitioner’s identification of “the components of the platform corresponding to the first bus and bus controller” as the claimed “network device” is fundamentally inconsistent with the language of claim 1. Repeated below is a diagram illustrating the hardware components recited in claim 1, and how such components are interrelated as specified in the claim. In connection with each component, the corresponding language from claim is set forth.



134. Claim 1 explicitly distinguishes between two separate entities: (i) the *apparatus*, which comprises a processor, a first bus, and a bus controller, and (ii) the *network device*, which includes a second processor, a network interface, and a second bus. The claim further recites that the bus controller is coupled to both the first bus (of

the apparatus) and a second bus (of the network device), and that the processor of the apparatus is “*distinct from* the processor included in the network device.” This structural separation was required by the Examiner during prosecution, and is not merely implied—rather, it is expressly recited and functionally significant. For example, the bus controller serves as a communication bridge between the apparatus and the network device, receiving encrypted management requests from the network device’s second bus and conveying them to the apparatus processor over the first bus. This architectural delineation cannot be ignored without rewriting the claim.

135. The Petition states “the components of the platform corresponding to the first bus and bus controller identified below are the claimed “network device,” and the LAN Connector and/or the Network (LAN) Controller is the claimed “network interface” because they are the interface to the LAN as show below in Figure 1-2. Petition, 55. Thus, the Petition is advancing a mapping where the network device includes a first bus and a bus controller -- which fails to satisfy the claim language.

136. Petitioner’s attempt to treat the “first bus and bus controller” as part of the network device improperly collapses the two distinct systems into a single entity, thereby erasing the detailed architectural boundaries that were defined over the course of numerous Office Actions during the prosecution of the ‘016 Patent. *See*, Section VIII, *supra*.

137. Moreover, Petitioner’s mapping effectively nullifies the final clause of claim 1, which expressly recites that “the processor of the apparatus is distinct from the processor included in the network device.” Treating the apparatus and the network device as overlapping or identical structures is incompatible with this limitation and renders it superfluous.

138. Because the Petition relies on an interpretation of the claim that is contrary to its plain and unambiguous language, and because the Petition fails to identify structures in IPMI that satisfy the recited claim limitations—particularly the distinction between the apparatus and the network device—the Petition fails to show that IPMI anticipates or renders obvious the Challenged Claims.


B. Petitioner’s Reliance On PCI-Based Intra-Controller Communication Is Fundamentally Flawed, As It Contradicts Established PCI Architecture Principles And Introduces Unjustified Complexity With No Articulated Technical Benefit.

139. Petitioner’s proposed combination, which relies on the use of a PCI bus within the LAN controller to transmit management data to the BMC, is fundamentally at odds with both the architecture of the PCI bus and common design practices recognized by those of ordinary skill in the art. Petition, 64-65. As IPMI and Huckins both make clear, the PCI bus is not designed to facilitate direct communication between devices without coordination by a centralized controller such as the I/O Hub. Indeed, the PCI bus architecture requires that all inter-device communication be


mediated by the CPU or system chipset, rendering any inter-device communication over the PCI bus not only inefficient but architecturally incoherent. The Petition offers no rationale for why a skilled artisan would modify a LAN controller to use a PCI bus for internal communication—an approach that would unnecessarily complicate the design, reduce throughput, consume valuable PCI bandwidth, and increase power consumption. Such a configuration would defy the efficiencies sought in network controller design and be inconsistent with established principles of hardware architecture. Absent any articulated motivation or technical benefit for incorporating a PCI bus in this manner, the Petition fails to provide a persuasive reason why a person of ordinary skill would adopt the proposed combination. Accordingly, Petitioner’s reliance on this configuration cannot support a finding of obviousness.

140. The IPMI reference explains that “There are two implementations that are likely to be used to deploy an IPMI LAN Interface using a shared LAN controller. The first implementation uses an embedded LAN controller, as shown in *Figure 12-1*, and the second uses a LAN controller on an add-in card, as shown in *Figure 12-2*.”

Figure 12-1, Embedded LAN Controller Implementation



141. The Petition (at 64) relies upon the Embedded LAN Controller Implementation and provides a version of Figure 12-1, including elements from Huckins annotated in green.



142. To support the proposed combination, the Petition alleges:


In this modified version of Figure 12-1, the Network (LAN) Controller is implemented using known Ethernet Interface logic with additional Sideband logic that recognizes UDP datagram packets

directed toward the management port which carry IPMI management messages. EX1003, ¶199. These packets can be differentiated from packets carrying normal user data based on IP addresses and/or UDP port numbers. *Id.* Those packets are transmitted over the PCI bus (the second bus of limitation [1.f] as discussed below) to a PCI-I2C bus controller bridge of the type described in Huckins. Petition, 64-65.

143. However, the Petition offers no explanation why a POSITA would be motivated to design a LAN controller that utilizes a PCI bus connection to communicate data within the LAN controller. The PCI bus, or Peripheral Component Interconnect bus, is a parallel data bus standard used in computer systems to connect peripheral devices—such as network interface cards, sound cards, and storage controllers—to the central processing unit (CPU) and main memory. It provides a shared communication pathway where multiple devices can exchange data with the processor and memory using a common set of address, data, and control lines. Accordingly, in the PCI bus architecture, devices typically do not communicate directly with each other as suggested by the Petitioner. Instead, for components to exchange data with each other, the data needs to be coordinated by a centralized controller.


144. Huckins confirms this understanding of the PCI bus architecture. For example, Huckins explains that “The PCI protocol is commonly associated with a type of the local I/O bus. Some examples are **the audio controller**, firmware hub (flash BIOS) 128, data storage 124, legacy I/O controller containing user input and keyboard interfaces, a serial expansion port such as Universal Serial Bus (USB), wireless

transceivers, and a network controller 134.” EX1007, [0026]. As illustrated in Figure 1 of IPMI (annotated below), although the network controller and the audio controller, are both PCI devices, the devices cannot communicate directly with each other over the PCI bus.



145. Instead, for the network controller and the audio controller to communicate, the data must traverse the I/O Controller Hub 130. Therefore, Huckins confirms that in the PCI architecture devices cannot communicate directly and instead require the coordination of a centralized controller. Accordingly, it would be nonsensical for a LAN controller to utilize a PCI bus architecture to communicate data

within the LAN controller as suggested by the Petition. As illustrated below, such a configuration would require external coordination by the I/O controller hub 130.



146. As illustrated above, in the Petition's proposed combination that requires the PCI architecture, the data would have to be received by the network controller, routed to the Non-PCI Device (I2C controller) over the PCI bus by the I/O Control Hub, and then onward to the BMC. This external coordination increases the LAN controller's complexity, decreases the LAN controller's throughput, decreases the PCI

bus's bandwidth and increases power consumption. All of these are completely superfluous when communicating data within the same LAN controller.

147. The Petition has failed to explain why the data received by the LAN controller and destined for the BMC would ever leave the LAN controller and be placed on the PCI bus. Further, the Petition has failed to show why a POSITA would graft on the additional complexity of using a PCI bus to communicate data within a single component.

148. As a result, the Petition's proposed use of the PCI bus to facilitate internal communication within a single LAN controller reflects a fundamental misunderstanding of both the architecture and intended use of the PCI bus. As confirmed by both the IPMI specification and Huckins, the PCI bus is designed to support communication between peripheral devices and the central system, not for routing data within a single integrated component. Requiring a LAN controller to offload management data onto the PCI bus only to route it back internally via the I/O Controller Hub introduces unnecessary design complexity, reduces system performance, and consumes additional bandwidth and power, all without technical justification. The Petition offers no credible explanation for why a POSITA would adopt such an inefficient and counterintuitive design choice, especially when simpler and more direct alternatives exist.

C. Petitioner’s Reliance On In-Band Encryption Is Misplaced, As The Cited Art Does Not Disclose Encrypted Requests Over The Out-Of-Band I²C First Bus

149. Petitioner contends that limitation 1.h—”and to convey the encrypted form of the management requests to the processor via the first bus”—is rendered obvious by the combination of IPMI, Frantz, and Huckins, relying solely on prior arguments related to limitations 1.b and 1.e. However, as detailed below, the Petition’s reliance on encryption in the context of in-band communication over a network interface is entirely misplaced when applied to the internal communication path identified as the “first bus.” The Petition offers no analysis or evidence suggesting that encryption would be applied to the internal, out-of-band “side-band” connection within the embedded LAN controller, nor does it cite any reference that teaches or suggests encrypting data transmitted over such a connection. Instead, the Petition identifies the first bus as an I²C interface, a protocol that, as confirmed by industry standards, is a low-level, unencrypted communication protocol used for short-range signaling between chips on the same circuit board.

150. The entirety of the Petitioner’s mapping of limitation [1.h] is reproduced below.

IPMI combined with Frantz and Huckins renders this limitation obvious to a POSITA for the reasons discussed above in connection with limitations [1.b] and [1.e] of this Ground 9. EX1003, ¶204. Petition, 67.

151. For completeness, paragraph 204 of the declaration of Petitioner's expert (EX1003) provides:

204. IPMI combined with Frantz and Huckins renders this limitation obvious to a POSITA for the reasons discussed above in connection with limitations [1.b] and [1.e] of this Ground 9.

152. With respect to limitation 1.b, the Petition states:


[1.b] wherein the network device is configured to receive data requests and an encrypted form of management requests via the network interface. Petition, ix.

153. With regard to limitation 1.b in Ground 9, the Petition states:

A POSITA would have recognized that encrypting the payloads of packets would improve security, particularly in view of the teaching in Frantz of the desirability of making such in-band communication of management requests secure due to the sensitive nature of such management information and that a hacker's chances of gaining access to secure data on the managed server 2 would be improved if the hacker had access to such management information, and the use of encryption in order to do so. Petition, 59.

154. While this discussion is clearly related to why a POSITA would encrypt the management requests that are received over the network interface as required by 1.b., this discussion has no connection to the internal "side band" connection that the Petition identifies as the first bus, as shown below.

Figure 12-1, Embedded LAN Controller Implementation



155. The “side-band” connection is the epitome of an out-of-band signal as it is only accessible from within the embedded LAN controller. Accordingly, Petitioner’s argument regarding the desirability of making an in-band communication secure from hackers is completely irrelevant to an out-of-band signal that would be virtually impossible for the hacker to access.


156. The Petitioner further argues with regard to limitation 1.b that:

Additionally, other references from the relevant time period such as the “Alert Standard Format” (“ASF”) specification DSP0114, which defines the interface between a management console and a managed device using RMCP communications, discloses the possibility of designing-in security mechanisms within the RMCP packet structure, which a POSITA would recognize as including encryption in Section 2.4. EX1037, 7. DSP0114 also allows for customization of the data block based on the IANA Enterprise Number in the RMCP packet. EX1037, 19. Petition, 59-60.

157. Here, as identified by the Petition, the ADF specification “defines the interface between a management console and a managed device using RMCP

communications.” Petition, 59-60. The Petition does not allege, nor does the ADF have anything to do with, the internal communications within the management console or the managed device. Instead, the standard defines how the management console and the managed device communicate over a public network.

158. Similarly, with regard to limitation 1.e in Ground 9, the Petition makes no effort to explain why communication along the “side-band” connection of IPMI is encrypted. In fact, the Petition’s discussion of limitation 1.e is completely silent with regard to the word “encrypt” or “encryption.” See, Petition, 62-65. At most, as illustrated below, the Petition alleges that the first bus is the I2C bus.



159. However, the I2C bus is not inherently encrypted. Instead, the I2C, which stands for Inter-Integrated Circuit, is a low-speed, two-wire serial communication protocol developed by Philips (now NXP Semiconductors) in the 1980s. It is widely used for short-distance communication between integrated circuits on the same circuit

board, particularly in embedded systems and system management applications. The protocol relies on two lines: the Serial Data Line (SDA), which carries the data, and the Serial Clock Line (SCL), which carries the clock signal. Communication on the I²C bus is always initiated by a master device—typically a microcontroller or CPU—which controls the clock and selects the slave device to communicate with using a unique 7-bit or 10-bit address. Data is exchanged in bytes, and each byte is followed by an acknowledgment bit to confirm successful transmission.

160. I²C and SMBus are low-level, lightweight serial communication interfaces designed primarily for short-range, intra-board communication between chips—such as sensors, memory modules, and microcontrollers—on the same motherboard or hardware platform. Because of their simplicity and minimal overhead, these protocols do not include any built-in features for encryption, authentication, or integrity checking. They were developed under the assumption that the bus environment is physically secure and isolated from external threats.

161. Therefore, Petitioner has not alleged, nor does I²C necessarily require that the communication to be encrypted. Instead, I²C relies upon the fact that out-of-band communication within an embedded device is physically secure and isolated from external threats which would render any additional layer of security superfluous. Therefore, Petitioner has failed to show that IPMI inherently or expressly teaches a

system that is configured “to convey the encrypted form of the management requests to the processor via the first bus,” as recited in claim 1.

162. As a result, Petitioner fails to establish that the combination of IPMI, Frantz, and Huckins renders limitation: “to convey the encrypted form of the management requests to the processor via the first bus” (1.h) obvious. Petitioner’s arguments focus on encrypting management requests transmitted over in-band network interfaces and do not address the internal communication path identified as the “first bus.” The Petition provides no explanation or evidence to support why encryption would be applied to the internal “side-band” connection within the embedded LAN controller, which is inherently out-of-band. Because the cited references do not teach or suggest encrypting management requests transmitted over the first bus, the Petition does not demonstrate that this limitation is rendered obvious.

XIV. GROUNDS 10-14 FAIL BECAUSE THE ADDITIONAL REFERENCES DO NOT CURE THE DEFICIENCIES IN GROUND 9.

163. Grounds 10 through 15 fail for the same reasons explained in response to Ground 9. Each of these grounds relies on the IPMI Specification as the primary reference, but IPMI fails to disclose the architectural separation between the claimed apparatus and the network device, or the conveyance of encrypted management requests via a first bus that is distinct from a second bus of the network device. The additional references cited in Grounds 10 through 15—are not alleged to disclose, and do not disclose, these missing features. Because these additional references do not

cure the deficiencies in IPMI, Grounds 10-15 fail to demonstrate unpatentability of the Challenged Claims.

XV. DECLARATION

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

Date: May 14, 2025

By: Prabhat Mishra
Dr. Prabhat Mishra