

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

MICROSOFT CORPORATION,
Petitioner,

v.

TOUTVIRTUAL, INC.,
Patent Owner.

Case IPR2026-00295
U.S. Patent No. 11,656,915

PATENT OWNER'S PRELIMINARY RESPONSE

TABLE OF CONTENTS

I.	INTRODUCTION	1
II.	OVERVIEW OF THE ’915 PATENT AND THE ASSERTED GROUNDS	3
A.	The ’915 Patent Solved an Industry Need for a Cross-Platform Virtualization Management Tool	3
B.	The Asserted Grounds	8
C.	The Relevant Asserted Art	10
1.	Khandekar (Grounds 1A and 1B): It Provides Host-Level, Platform-Dependent Management	10
2.	Hirschfeld (Grounds 2A and 2C): its Virtualization Software Runs on the Operating System and Manages only the Physical Resources of Each Host.....	11
3.	Vasilevsky (Ground 1D): is Directed to Grid Computing	14
III.	LEVEL OF ORDINARY SKILL IN THE ART	15
IV.	CLAIM CONSTRUCTION	16
V.	THE BOARD SHOULD DENY INSTITUTION BECAUSE MICROSOFT’S CLAIM CONSTRUCTION POSTURE IS INCONSISTENT ACROSS FORUMS	19
VI.	INSTITUTION SHOULD BE DENIED BECAUSE MICROSOFT FAILS TO DEMONSTRATE THAT ANY CHALLENGED CLAIM IS UNPATENTABLE.....	21
A.	Khandekar (Grounds 1A and 1B): Independent Claim 1 is Patentable Over Khandekar.....	21
1.	Khandekar Does Not Satisfy Microsoft’s Own Construction of “Managed”	21

2.	Khandekar Does Not Satisfy Microsoft’s Own Construction of “Control Computer Connection” (Claim 1)	23
B.	Khandekar-Vasilevsky (Ground 1D): Independent Claim 11 and the Dependent Claims Are Patentable Over the Khandekar-Vasilevsky System	25
1.	Khandekar and Vasilevsky Are Directed to Different Architecture Designs Which Are Incompatible with One Another.....	26
a)	Khandekar’s Architecture Includes Physical Resources That Each Deploy Multiple Virtual Machines.....	26
b)	Vasilevsky’s Grid Computing Architecture Provides a Single Virtualization System	26
2.	Modifying Khandekar with Vasilevsky Would Have Rendered Khandekar Inoperable for its Intended Purpose	29
3.	The Khandekar-Vasilevsky System Does Not Satisfy “Virtual Assets [that] Are Managed by Virtualization Layer Instances Running on Each Physical Resource”	32
C.	Hirschfeld Combinations (Grounds 2A and 2C): Independent Claims 1 and 11 Are Patentable Over the Hirschfeld Combinations.....	33
1.	Hirschfeld’s Virtualization Software Runs on the Operating System Layer, Not the Physical Resource	34
2.	A POSITA Would Not Have Used a Smith-Like VMM In Hirschfeld Because its SCM Manages its Logical Servers	37
3.	Microsoft’s Hirschfeld-Smith Mapping Does Not Satisfy Microsoft’s Own Construction of “Managed”	39
VII.	CONCLUSION.....	40

TABLE OF AUTHORITIES

	Page(s)
Cases	
<i>Am. Airlines, Inc. v. Intell. Ventures I LLC</i> , No. IPR2025-01055, Paper 11 (P.T.A.B. Nov. 21, 2025).....	20
<i>Facebook, Inc. v. Sound View Innovations, LLC</i> , No. IPR2017-00985, Paper 30 (P.T.A.B. Oct. 19, 2017).....	17
<i>First Alert, Inc. v. Walter Kidde Portable Equip., Inc.</i> , No. IPR2023-00898, Paper 14 (P.T.A.B. Dec. 13, 2023)	32
<i>Ford Motor Co. v. Autoconnect Holdings LLC</i> , No. IPR2026-00171, Paper 11 at 2-3 (P.T.A.B. May 13, 2026)).....	20
<i>In re Gordon</i> , 733 F.2d 900 (Fed. Cir. 1984)	29
<i>Lenovo (U.S.) Inc. v. Neodron Ltd.</i> , No. IPR2020-00688, Paper 8 (P.T.A.B. Sept. 15, 2020).....	32
<i>NTP, Inc. v. Research in Motion, Ltd.</i> , 418 F.3d 1282 (Fed. Cir. 2005)	17
<i>Phillips v. AWH Corp.</i> , 415 F.3d 1303 (Fed. Cir. 2005)	35
<i>Revvo Techs., Inc. v. Cerebrum Sensor Techs., Inc.</i> , No. IPR2025-00632, Paper 20 (P.T.A.B. Nov. 3, 2025).....	21
<i>Southwall Techs., Inc. v. Cardinal IG Co.</i> , 54 F.3d 1570 (Fed. Cir. 1995)	36
<i>Tesla, Inc. v. Intell. Ventures II, LLC</i> , No. IPR2025-00340, Paper 18 (P.T.A.B. Nov. 5, 2025).....	20
<i>Winner Int'l Royalty Corp. v. Wang</i> , 202 F.3d 1340 (Fed. Cir. 2000)	37

Statutes

35 U.S.C. § 112.....	20
----------------------	----

Other Authorities

37 C.F.R. § 42.104(b)(3).....	17, 39
37 C.F.R. § 42.104(b)(4).....	17, 22, 39, 40

PATENT OWNER'S EXHIBIT LIST

EXHIBIT	DESCRIPTION
2001	First Amended Docket Control Order, <i>ToutVirtual, Inc. v. Microsoft Corporation</i> , No. 25-cv-00943-RWS-RSP, Dkt. No. 35 (E.D. Tex. Feb. 23, 2026)
2002	Declaration of Vipul Pabari in Support of Patent Owner's Request for Discretionary Denial PROTECTIVE ORDER MATERIAL
2003	Declaration of James Hannah in Support of Patent Owner's Request for Discretionary Denial
2004	Docket Navigator webpage, Time to Trial in the Eastern District of Texas, dated January 1, 2024 through May 9, 2026
2005	Excerpts from U.S. District Court – Federal Case Management Statistics, Median Time from Filings to Trial (E.D. Tex.), 12-Month Periods Ending December 31, 2025, available at https://www.uscourts.gov/data-news/data-tables/2025/12/31/statistical-tables-federal-judiciary/t-3
2006	Summons in a Civil Action, <i>ToutVirtual, Inc. v. Microsoft Corp.</i> , No. 25-cv-00943-RWS-RSP, Dkt. No. 5 (E.D. Tex. Sep. 8, 2025)
2007	Excerpts from Defendant's Preliminary Invalidity Contentions, <i>ToutVirtual, Inc. v. Microsoft Corp.</i> , No. 25-cv-00943-RWS-RSP (E.D. Tex.), served on March 12, 2026
2008	U.S. Patent No. 8,799,431
2009	U.S. Patent No. 9,444,762
2010	ToutVirtual – VirtualIQ Pro 3 – Virtual Intelligence Best Practices

EXHIBIT	DESCRIPTION
2011	Reuters webpage, "Virtualization management takes center stage," dated August 26, 2009, available at https://www.reuters.com/article/business/environment/virtualization-management-takes-center-stage-idUS1669466711/
2012	ToutVirtual press release, "ToutVirtual Announces VirtualIQ Pro 3," dated January 13, 2009, available at https://www.toutvirtual.com/company/pr/VirtualIQ-Pro-January-13-2009.pdf
2013	Network Computing webpage, "ToutVirtual Unveils New Software," dated May 30, 2007, available at https://www.networkcomputing.com/network-management/toutvirtual-unveils-new-software
2014	PR.com press release, "ToutVirtual to Offer Product Suite via BT Engage IT," dated June 1, 2012, available at https://www.pr.com/press-release/417010
2015	Computer World webpage, "Computerworld Horizon Awards 2006 Honorees," dated August 21, 2006, available at https://www.computerworld.com/article/1649654/computer-world-horizon-awards-2006-honorees.html
2016	Network World webpage, "10 virtualization companies to watch," dated August 20, 2007, available at https://www.networkworld.com/article/824596/data-center-10-virtualization-companies-to-watch.html
2017	Network Computing webpage, "ToutVirtual Offered with Microsoft," dated April 5, 2007, available at https://www.networkcomputing.com/network-infrastructure/toutvirtual-offered-with-microsoft
2018	ToutVirtual press release - "ToutVirtual Announces Support for Microsoft Windows Server 2008 Hyper-V Virtualization," dated October 27, 2008

EXHIBIT	DESCRIPTION
2019	ToutVirtual, Inc. presentation, “Virtual Problems – Real Solutions,” dated November 2006 PROTECTIVE ORDER MATERIAL
2020	E-mail from V. Pabari to T. Bailey and M. Neil, dated November 15, 2006 PROTECTIVE ORDER MATERIAL
2021	Excerpts of File History for U.S. Patent No. 8,799,431
2022	Excerpts of File History for U.S. Patent No. 11,656,915
2023	Order, <i>Maxell, Ltd. v. Samsung Elecs. Co.</i> , No. 5:25-cv-00052-RWS, Dkt. No. 64 (E.D. Tex. Dec. 15, 2025)
2024	U.S. Patent No. 10,810,050
2025	Petition for <i>Inter Partes</i> Review, <i>Microsoft Corp. v. ToutVirtual, Inc.</i> , No. IPR2026-00296, Paper 2 (P.T.A.B. Mar. 10, 2026)
2026	Default Protective Order
2027	Additional excerpts from Defendant's Preliminary Invalidity Contentions, <i>ToutVirtual, Inc. v. Microsoft Corp.</i> , No. 25-cv-00943-RWS-RSP (E.D. Tex.), served on March 12, 2026
2028	Excerpts from VMWare's GSX Server 3 Administration Guide (Version 3.0), https://dn790007.ca.archive.org/0/items/manualzilla-id-7300686/7300686.pdf

I. INTRODUCTION

The Board should deny institution because Petitioner's ("Microsoft's") Petition grounds fail at three levels.¹ First, Microsoft advances inconsistent claim construction positions between the PTAB and district court forums with no explanation for why it should be allowed to do so. In the co-pending district court litigation, Microsoft alleges that several claim limitations are indefinite but here expressly construes terms in these limitations and measures the asserted art against their claim scope. Second, Microsoft's asserted art mappings are inconsistent with the constructions Microsoft itself proposed. Third, a POSITA would not have made the combinations Microsoft proposes because its references are structurally incompatible with one another.

Grounds 1A and 1B (Khandekar, Independent Claim 1): this ground fails because Microsoft's Khandekar mappings are, on their face, inconsistent with its own proposed constructions for several claim terms, including managed and control computer connection, advanced in this proceeding and the co-pending IPR2026-

¹ Microsoft advances nine grounds, five of which—Grounds 1A, 1B, 1D, 2A and 2C—challenge Independent Claims 1 and 11. Pet. at 1-2.

00296 (“-00296 IPR”).² Because of these internal inconsistencies, Microsoft fails to show how Khandekar meets the independent claims under its own proposed constructions.

Ground 1D (Khandekar-Vasilevsky, Independent Claim 11): this ground fails because a POSITA would not have modified Khandekar with Vasilevsky because the two references are directed to mutually exclusive architecture designs. Khandekar deploys multiple VMs on each physical server while Vasilevsky abstracts multiple physical servers into a single virtualized computing system. Vasilevsky's management capabilities are achieved through this abstraction. Therefore, to incorporate Vasilevsky's management capabilities into Khandekar, as proposed, would require abstraction of Khandekar's physical servers which would render Khandekar inoperable for its intended purpose. Indeed, the Office found a similar Vasilevsky combination inoperable for its intended purpose. That finding applies equally here.

Grounds 2A and 2C (Hirschfeld Combinations, Independent Claims 1 and 11): this ground fails because Hirschfeld fails the requirements of the independent claims: “virtual assets are managed by virtualization layer instances

² The -00296 IPR challenges related U.S. Patent No. 10,810,050 (“’050 Patent,” Ex. 2024).

running on each of the physical resources.” *First*, its virtualization software runs on the operating system layer, not on the physical resources, which places it squarely within the same architectural deficiency the Applicant expressly distinguished during prosecution when overcoming prosecution art on identical grounds. *Second*, Hirschfeld’s management architecture is bifurcated between the virtualization software (the alleged “virtualization layer”) and the separate Server Cloud Manager (“SCM”). The SCM manages the logical servers (the alleged “virtual assets”), not the virtualization software as the claims require. The additions of Smith and Vasilevsky do not cure Hirschfeld’s deficiencies.

II. OVERVIEW OF THE ’915 PATENT AND THE ASSERTED GROUNDS

A. The ’915 Patent Solved an Industry Need for a Cross-Platform Virtualization Management Tool

The ’915 Patent describes methods and apparatus for automatic management of computer network virtualization environments. ’915 Patent at 3:45-65, 4:11-25, 6:38-59. By the late 2000s, the proliferation of different hypervisors had driven the industry need for a unified, cross-platform virtualization management tool. Ex. 2011 at 2. For example, VMware, a virtualization vendor, only “focuse[d] on boosting its

capabilities for managing its own hypervisors.” *Id.* at 3³. Yet only “7% of respondents [in a survey conducted in mid-2008] had only one virtualization supplier in their test and development and production environments.” *Id.* at 2. The industry thus was moving away from vendor-specific management since it “prefer[s] different types of hypervisors for different types of applications” which “require[] that you have a single view, and from an IT perspective you don’t want multiple management tools, repositories, teams of users. You want a single pane of glass.” *Id.* at 2.

ToutVirtual’s VirtualIQ answered that need by delivering a unified management console that can automatically manage virtualization environments across different vendors of platforms and platform-types, and was met with industry praise. Paper 9 (“PO Discretionary Brief”) at 9-11; *see also* Ex. 2002 at ¶ 5 (citing Ex. 2010), ¶ 4 (citing Ex. 2011 at 2, Ex. 2013 at 2), ¶ 7 (citing Ex. 2013 at 2, Ex. 2012 at 1, Ex. 2014 at 2). Specifically, ToutVirtual’s VirtualIQ technology includes a unified management console that, *inter alia*, carries out automatic provisioning, server virtualization assessment, asset management, performance management, and capacity management in cross-platform environments. Ex. 2012 at 1; *see also* Ex.

³ VMware is the assignee of Khandekar asserted in Grounds 1A and 1B, which as explained above suffers from these very issues. *See* Khandekar at Cover.

2014 at 1-2 (“VirtualIQ provides a ‘single pane of glass’ across physical and virtual infrastructure . . .”).

ToutVirtual's VirtualIQ technology is built directly on the '915 Patent's management architecture platform. *See* Ex. 2019 at 6. The '915 Patent—and by extension VirtualIQ—discloses and claims a control center that manages “multiple virtual environments through a single user interface . . . even where the underlying CPU of the system physical resources are different from that of the control center, even where the operating systems of the control center, physical resources, and virtual assets are different, and even where the virtualization environments being managed are different from each other.” '915 Patent at 3:39-4:3. This environment-level, cross-platform management is achieved through control center 202 (red), which serves as the unified management console; the control agent 324 (blue), which resides on each host 300 and acts as the universal adapter between the control center and the host's virtualization layer; and the virtualization layer 322 (green) that runs on the host 300, as reproduced below:




FIG. 3

Id. at Figs. 2 and 3 (annotated). This management architecture is the foundational platform for ToutVirtual’s VirtualIQ technology.

The control center serves as the centralized management console for one or more virtualization environments. *Id.* at 7:5-17. The control center comprises a collection of management components—including asset manager 204, provisioning manager 206, dynamic application router 208, optimizer 210, performance manager 212, and capacity planning manager 214 (as shown above)—each of which provides a discrete management function across the different virtualization environments. *Id.* at 4:4-8, 7:36–9:44, Fig. 2.

The control center collectively executes these management components across virtualization environments and does so “without regard to processors, operating systems, virtualization platforms, and application software of the

virtualization environment”—environment-level, platform-independent management. *Id.* at 4:16-39; *id.* at 6:42-63 (“The present invention provides automated management of network-based virtual environments through Control Center software that is **platform-independent**”) (emphasis added). For example, the control center can manage across different types of virtualization platforms—server (physical hosts to virtual machines), storage (physical disks to logical units), and network (physical routers to virtual network paths)—or across different vendors of each platform type (VMWare, Xen, Microsoft, or IBM), all without regard to the processors, operating systems, virtualization platforms, and application software of the virtualization environment. *Id.* at 6:42-63, 18:53-19:3 (“the Control Center virtualization management application being described can also be used to manage other virtual assets, such as user desktops (user interfaces), data storage . . . and network traffic devices.”), 20:19-22 (server virtualization management), 20:40-47 (storage virtualization management), 21:4-7 (network router virtualization management), Figs. 11-14.

To accomplish this management, the control center communicates with and instructs the virtualization layer running on each host through control agent 324, as one embodiment details. *Id.* at 10:27-33. The control agent is installed on each host to “facilitate[] communication between the external control center components 204-214 and the virtualization layer 322 of the host machine.” *Id.* at 10:33-36. That is,

the control agent “provides a common interface to the variety of virtualization software that might be available on the host machine”—for example, virtualization layers provided by VMWare, Xen, Microsoft, or IBM.” *Id.* at 10:33-41. In this way, the control center computer thus instructs the virtualization layer to carry out the management operations on each host, and does so regardless of the platform present (“platform-independent”). *Id.* at 6:43-64; *see also id.* at 3:46-65, 4:16-39.

Independent Claims 1 and 11 capture the core of the management architecture. Independent Claim 1 recites, *inter alia*: a “control computer” that is connected to the physical resources and virtual assets and through a control computer connection enables an end user to access the virtual assets. ’915 Patent at Claim 1. Likewise, Independent Claim 11 recites, *inter alia*, a “control computer” that enables “provisioning of a virtual asset of the one or more virtual assets” through a virtualization layer that manages the virtual assets. *Id.* at Claim 11.

B. The Asserted Grounds

The Petition asserts nine grounds, four of which challenge Independent Claims 1 and 11 (bolded):

Ground	Claims	Basis/Asserted Reference(s)
1A/1B	1,2	§§102/103: Khandekar
1C	3-5	§103: Khandekar, Bandhole
1D	6-8, 11-14	§103: Khandekar, Vasilevsky
1E	9-10	§103: Khandekar, Fujino
2A	1, 2	§103: Hirschfeld, Smith

Ground	Claims	Basis/Asserted Reference(s)
2B	3-5	§103: Hirschfeld, Smith, Bandhole
2C	6-8, 11-14	§103: Hirschfeld, Smith, Vasilevsky
2D	9-10	§103: Hirschfeld, Smith, Fujino

Pet. at 1-2.

As discussed in Section VI below, Microsoft's **Khandekar** (Grounds 1A and 1B), **Khandekar-Vasilevsky** combination (Ground 1D), and the **Hirschfeld** combinations (Grounds 2A and 2C) fail to show a reasonable likelihood that Microsoft would prevail with respect to any claim because Khandekar, Khandekar-Vasilevsky, and the Hirschfeld combinations entail structurally different architectures from the claimed invention's management architecture.

Notably, the Office has been prosecuting the '915 Patent family portfolio for roughly two decades, during which time the same or substantially the same art and arguments asserted here were previously before the Office. PO Discretionary Brief at 13-20. The Petition grounds are essentially repackages of art and arguments that were previously before the Office and overcome.

Because Grounds 1A, 1B, 1D, 2A, and 2C challenge the independent claims, the dependent claim grounds, which depend from the independent claims, also fail. Therefore, the arguments addressed herein regarding the independent claim grounds apply equally to the dependent claim grounds.

C. The Relevant Asserted Art

1. Khandekar (Grounds 1A and 1B): It Provides Host-Level, Platform-Dependent Management

VMware—and by extension Khandekar (a VMware patent)—is focused on the management of its own virtualization environment. Ex. 2011 at 3; Khandekar at Cover. The industry's "real goal," however, was for a unified, cross-platform virtualization management. Ex. 2011 at 2; § II.A, *supra*. Khandekar is thus the vendor-specific architecture that the industry moved away from.

A closer look at Khandekar's architecture reveals its host-level management, tethered to VMware-specific platform.⁴ Each host is installed with a VMware-specific Virtual Machine Monitor ("VMM"), which is a VMware-specific platform for its server virtualization. Khandekar at 7:7-27, 13:17-25. VMware provides the VMM that is installed on every host to ensure VMware compatibility and

⁴ Khandekar's VMware-dependency is illustrative of its broader platform dependency—it is tethered to a vendor's server virtualization platform and does not extend to storage, network, or other virtualization platform types. *See, e.g.*, Khandekar at 5:25-49, 9:26-35, 11:20-35 (provisioning architecture directed to server virtualization). Thus, while ToutVirtual addresses herein Khandekar's vendor-specific dependency, that dependency also extends to platform type.

configuration when creating and loading virtual machines. *Id.* at 7:7-9, 13:17-25 (“[A] VMM is already pre-installed on each host . . . so as to ensure compatibility and proper configuration,” with products “from VMware” providing that VMM). Every management operation depends on the VMware-specific VMM at each host. Every resource request a virtual machine (“VM”) issues—for memory, processor time, storage, or network access—passes through the VMware VMM on that host, which schedules and handles all requests by its VM for machine resources as well as various faults and interrupts. *Id.* at 7:18-47, 13:17-35.


The VMware-specific platform dependency extends upward through the provisioning architecture operational sequence. It is platform dependent. For example, the provisioning server's heuristics engine reaches the physical hosts (1000-1003) through the vendor-specific VMM interface to determine the host with enough memory and disk space to host the VM about to be deployed on the host. *Id.* at 20:16-58. Because the provisioning architecture communicates only with its VMware-specific platform, its provisioning operations are limited to the vendor-specific platform.

2. Hirschfeld (Grounds 2A and 2C): its Virtualization Software Runs on the Operating System and Manages only the Physical Resources of Each Host

Hirschfeld's architecture is, on its face, distinct from Independent Claims 1 and 11 that impose two requirements of a virtualization layer: (1) that it runs on the


physical resource; and (2) manages the virtual assets. '915 Patent at Claims 1 and 11. As a preliminary matter, during prosecution, the Applicant expressed that virtualization software running on an OS layer “cannot equate to” a virtualization layer running on the physical resource. Ex. 1002 at 91-92.

Hirschfeld's architecture does not meet either claim requirement because its virtualization software runs on the operating system (OS) layer and manages only the physical resources of each physical server. Hirschfeld's virtualized logical server cloud architecture is built on the concept that logical servers exist as abstractions from the physical servers. Hirschfeld at 1:15-20, 7:63-8:11, Fig. 2. Each physical server 108 includes three distinct architectural layers: a hardware layer 210, an OS layer 212, and a software layer 214. *Id.* at 8:39-9:46, Fig. 2. Hirschfeld's virtualization software 229 (purple, below) resides in the software layer 214 and runs on the OS layer (blue, below), as shown below.



Id., Fig. 2 (annotated), *see also id.* at 9:35-46. Because the virtualization software runs on the OS layer—not the physical resource—it does not meet claim requirement 1.

Hirschfeld also does not meet claim requirement 2 because its virtualization software’s management is expressly limited to the physical resources: it “manages ***physical resources*** of the logical servers 106 linked thereto.” Hirschfeld at 9:44-46 (emphasis added). The logical servers are managed by the server cloud manager (“SCM”) 102, a component located on a separate plane. *Id.* at Abstract, 11:45-46, 12:45-49. As shown in figure 4 below, the logical server (“LS”) includes a logical server agent (“LSA”) 411 that runs on the logical server’s OS and extends control of the logical server to SCM 102. *Id.* at 11:27-30; *see also id.* at 11:45-46 (“The SCM 102 manages all of the LSAs 411”).




Id. at Fig. 4.

3. Vasilevsky (Ground 1D): is Directed to Grid Computing

Unlike Khandekar and Hirschfeld, Vasilevsky is directed to grid computing architecture that abstracts multiple physical servers into a single virtualized system. Ex. 1002 at 93 (citing Vasilevsky ¶ 57); *id.* at 95. Vasilevsky has previously been before the Office and the Applicant addressed and overcame its grid computing architecture, as characterized below:

Vasilevsky sets out to abstract a grid of multiple computers to generate a single virtualized system that can execute applications and operating systems across resources of the multiple computers that creates an illusion of a single system.

Id. at 93; *see also* Vasilevsky at ¶¶ 13, 57. Vasilevsky's grid computing architecture—where multiple computers 107A to 107C are abstracted to form virtualized system 101—is reproduced below:



Id. at Fig. 1.

III. LEVEL OF ORDINARY SKILL IN THE ART

For purposes of this Preliminary Response only, ToutVirtual uses Microsoft's proposed definition of a person of ordinary skill in the art ("POSITA"), which is a person who had "(1) at least a bachelor's degree in computer science, computer engineering, electrical engineering, or a related field, and (2) at least two years of industry experience in virtualization systems. MS-1003, ¶¶20-22. Additional education could substitute for professional experience, and professional experience could substitute for education." Pet. at 8. ToutVirtual reserves the right to propose its own definition of a POSITA if trial is instituted.

IV. CLAIM CONSTRUCTION

For purposes of this Preliminary Response only, ToutVirtual applies Microsoft's proposed constructions. *See* Pet. at 3-8.⁵ Microsoft represents that its constructions “demonstrate subject matter that falls within the metes and bounds of these terms.” *Id.* at 3. Microsoft further represents that “[t]he mappings included in Section III are not inconsistent with these constructions.” *Id.*

In the co-pending -00296 IPR, Microsoft also expressly construes the term “managed”—the same term recited in the corresponding claim limitation of the '915 Patent. Ex. 2025 (“’050 Pet.”) at 5-6; *see also* Ex. 2024 ('050 Patent) at Claim 1. There, Microsoft also relies on the same Khandekar mappings and takes the same affirmative representation that its mappings are “not inconsistent with” its constructions. '050 Pet. at 3, 13, 29 (relying on Khandekar's VMM); Pet. at 21 (same). Because the '915 and '050 Patents share the same specification and recite the same claim term in the same limitation, and because Microsoft advances the same Khandekar mapping, Microsoft's “managed” construction applies here for

⁵ ToutVirtual does not adopt Microsoft's constructions and expressly reserves the right to propose constructions if trial is instituted. ToutVirtual only applies these constructions here for the limited purpose of demonstrating that Microsoft's mappings are internally inconsistent with its own constructions.

consistency purposes. *NTP, Inc. v. Research in Motion, Ltd.*, 418 F.3d 1282, 1293 (Fed. Cir. 2005) (the same claim term in related patents sharing a common specification must be interpreted consistently across those patents); *Facebook, Inc. v. Sound View Innovations, LLC*, No. IPR2017-00985, Paper 30, at 9-10 (P.T.A.B. Oct. 19, 2017) (holding that the petitioner's express claim construction position in a petition "is a bell that cannot be unrung," and a petitioner that "simultaneously argue[s] different claim constructions under the same claim construction standard . . . without explanation, may raise even greater concerns").

By providing express construction under 37 C.F.R. § 42.104(b)(3), and by affirmatively representing that its mappings are "not inconsistent with these constructions" for § 42.104(b)(4) purposes, Microsoft holds itself to these constructions. These constructions identify the subject matter that Microsoft contends falls within the claim scope and to which its mappings are purportedly consistent. Any mapping therefore that fails to reach these constructions is "inconsistent with these constructions." Pet. at 3.

As explained below, Microsoft's asserted art mappings for the "managed" and "control computer connection" are inconsistent with Microsoft's own constructions set forth in the table below:⁶

Term	Construction
"control computer connection"	"a communication pathway where the control computer acts as an intermediary or gateway in the data path, such that the end user's access to the virtual assets must be routed directly through the control computer itself." <i>See Pet. at 6-7.</i>
"managed"	The "management of computer network virtualization environments is provided by performing one or more functions," including: (1) identification and management of network resources and virtual assets, (2) provisioning of virtual assets in response to network workflow demands, (3) dynamic deployment of virtual assets across the computer network, (4) performance measurement and reporting of resources and virtual assets, and (5) planning and forecasting of resource demands and asset utilization of the virtualization environment, such that the functions are carried out without regard to processors, operating systems, virtualization platforms, and application software of the virtualization environment. MS-1001, 4:16-39. A POSITA would have understood that the claimed "managed" and "manages" at least

⁶ Because Microsoft's "managed" construction captures the "without regard to" term which Microsoft construed in the -00296 IPR, ToutVirtual also sets forth Microsoft's construction for this term in the table. However, "without regard to" will not be individually addressed herein.

Term	Construction
	encompasses the management features noted in the passage above. MS-1003, ¶63. <i>See</i> '050 Pet. at 5-6.
“Without regard to processors, operating systems, virtualization platforms, and application software of the virtualization environment”	A POSITA would have understood that the claimed “without regard to . . .” at least encompasses “irrespective of any.” See MS-1001, 4:11-36, 6:42-63; MS-1003, ¶63. <i>See</i> '050 Pet. at 3.

Notably, Microsoft’s “managed” construction captures the “management features” in the “management of computer network virtualization environments” (*at the environment-level*) that are “carried out without regard to.” Its “without regard to” construction—“irrespective of any”—in turn captures the control center’s “*platform-independent*” disclosure under which the control center operates irrespective of whichever “processors, operating systems, virtualization platforms, and application software” are present. '050 Pet. at 3 (citing '050 Patent at 6:42-63), *id.* at 5-6. These constructions thus set the standard for Microsoft’s mappings that follow.

V. THE BOARD SHOULD DENY INSTITUTION BECAUSE MICROSOFT’S CLAIM CONSTRUCTION POSTURE IS INCONSISTENT ACROSS FORUMS

The Board should deny institution on the grounds that Microsoft takes *unexplained, inconsistent* claim construction positions in the co-pending district

court proceeding, particularly alleging indefiniteness under 35 U.S.C. § 112, ¶ 2 of the claim limitations in the table below. *Tesla, Inc. v. Intell. Ventures II, LLC*, No. IPR2025-00340, Paper 18 at 3 (P.T.A.B. Nov. 5, 2025) (informative) (denying institution and observing “Petitioner’s statement that it cannot raise indefiniteness challenges in an *inter partes* review is not a sufficient explanation”); *Ford Motor Co. v. Autoconnect Holdings LLC*, No. IPR2026-00171, Paper 11 at 2-3 (P.T.A.B. May 13, 2026) (denying institution where the petitioner failed to address the “terms that it argues are indefinite in the district court”); *Am. Airlines, Inc. v. Intell. Ventures I LLC*, No. IPR2025-01055, Paper 11 at 12 (P.T.A.B. Nov. 21, 2025) (denying institution where “[c]ontrary to its position in this case, in district court, Petitioner argues that several limitations in the Challenged Claims . . . are indefinite, and *not* subject to reasonable interpretation”).

Claim Limitation	Microsoft’s District Court Position
Claim 1: the virtual assets are managed by virtualization layer instances running on each of the physical resources.	Indefinite under 35 U.S.C. § 112 (¶ 2)
Claim 1: accessing, through the control computer connection, the virtual assets by the end user.	Indefinite under 35 U.S.C. § 112 (¶ 2)
Claim 9: providing the virtual assets further comprise providing a logical network resource.	Indefinite under 35 U.S.C. § 112 (¶ 2)
Claim 11: virtualization layer instances running on each of the one or more physical resources.	Indefinite under 35 U.S.C. § 112 (¶ 2)

Ex. 2027 (Microsoft’s invalidity contentions) at 90, 95-96.

Notably, the claim limitations Microsoft alleges are indefinite in the district court forum contain terms that Microsoft expressly construes in this instant Petition and the -00296 IPR petition, including “control computer connection” and “without regard to” identified above. It is impossible to reconcile Microsoft’s apparent position that it cannot understand these claim limitations in the district court forum, but here propose constructions and attempt to show how its mappings fall within the claim scope. “[W]hen a petitioner takes alternative positions before the Board and a district court, that petitioner should, at a minimum, explain why alternative positions are warranted.” *Revvo Techs., Inc. v. Cerebrum Sensor Techs., Inc.*, No. IPR2025-00632, Paper 20 at 3-4 (P.T.A.B. Nov. 3, 2025) (precedential) (emphasis added). Microsoft makes no attempt to explain why alternative positions are warranted here. Institution therefore should be denied.

VI. INSTITUTION SHOULD BE DENIED BECAUSE MICROSOFT FAILS TO DEMONSTRATE THAT ANY CHALLENGED CLAIM IS UNPATENTABLE

A. Khandekar (Grounds 1A and 1B): Independent Claim 1 is Patentable Over Khandekar

1. Khandekar Does Not Satisfy Microsoft’s Own Construction of “Managed”

Khandekar does not satisfy Independent Claim 1, which requires that virtual assets are *managed* by the virtualization layer. See ’915 Patent at Claim 1. Microsoft’s managed construction covers environment-level management features

carried out “irrespective of any” (platform-independent). *See* § IV, *supra*. Microsoft’s own construction defines what it believes falls within the claim scope, and Microsoft’s own mapping shows Khandekar’s structural inconsistencies with that scope.

While Microsoft maps Khandekar’s VMM to the term, it fails to show how its mapping is consistent with its own construction. *See* 37 C.F.R. § 42.104(b)(4). For the managed limitation, Microsoft identifies Khandekar’s VMM and asserts that the VMM “manage[s] the virtual machines . . . on that host.” Pet. at 21 (Grounds 1A/1B). Microsoft, however, provides no explanation of how Khandekar’s VMM performs the management features at the environment-level nor does so independent of the platform as its construction demands. That is not a simple oversight; it is a gap that cannot be bridged because Khandekar is structurally incapable of satisfying Microsoft’s own construction.

Khandekar is a VMware patent, and its VMM platform is designed to operate exclusively within a VMware server virtualization environment, binding every operation to the host-level. *See* § II.C.1, *supra* (describing Khandekar’s VMM as a VMware-specific server virtualization platform). Because Khandekar’s vendor-specific VMM performs every operation at the host-level—e.g., scheduling and handling requests for VMs on its host—it is incapable of meeting what Microsoft’s own construction demands: environment-level, platform-independent management

features. *See* Khandekar at 7:18-47, 13:17-35. That is the structural gap that produces the inconsistency. Microsoft's Khandekar's VMM mapping shows host-level, vendor-dependency which is structurally inconsistent with its own claim construction election.

In view of the foregoing, because Microsoft's Khandekar mapping is inconsistent with its own construction, Khandekar cannot satisfy the managed limitation of Independent Claim 1.

2. Khandekar Does Not Satisfy Microsoft's Own Construction of "Control Computer Connection" (Claim 1)

Khandekar does not satisfy Microsoft's own "control computer connection" construction. Microsoft construes "control computer connection" to at least refer to "a communication pathway where the control computer acts as an intermediary or gateway in the data path, such that the end user's access to the virtual assets must be routed directly through the control computer itself." Pet. at 7. While Microsoft affirmatively represents that its Khandekar mapping is "not inconsistent with" this construction, that is untrue. *See* § IV, *supra*.

Microsoft's Khandekar user access mapping is inconsistent with this construction because it conflates two fundamentally different Khandekar pathways: the provisioning request pathway and the VM access pathway. *See* Pet. at 21-23. Microsoft identifies the user's browser-based connection to Khandekar's provisioning server through which the user submits configuration selections,

receives a task identifier, and is notified when the VM is created. *Id.*; *see also id.* at 24 (“query the status of the provisioning task”). But submitting provisioning request tasks is not accessing the VM. The VM user access path is distinct from the provisioning path.

Khandekar itself makes this distinction plain: once the provisioning server completes its work, the user “is so informed” and only then does the user “access the VM using *remote access software*.” Khandekar at 14:27-38 (emphasis added); *see also id.* at 13:14-27 (“The user can then access the newly created VM via *any known remote access software*, such as the *remote console software available from VMware, Inc.*”) (emphases added). This access is direct. At the time of Khandekar’s filing, VMware’s remote access software functioned as a client-side access tool that, when installed on the client workstation, connected the user directly to the VMware VMM running on the host (e.g., a server hosting virtual machines). *See* Ex. 2028 at 50-53 (the VMware Virtual Machine Console installed on a client workstation connects to the server host over port 902). That direct path—from the user’s workstation to the VM on the host—does not involve the provisioning server. *Id.* The provisioning server exits the picture entirely at the moment of notification, at which point the remote access software on the client workstation directly connects with the VM on the host. The provisioning server is thus not “an intermediary or gateway” in that data path. Pet. at 7.

Because Khandekar's provisioning server does not "act as an intermediary or gateway in the data path," Microsoft's Khandekar mapping is inconsistent with its own construction, and therefore Khandekar cannot satisfy the "control computer connection" of Independent Claim 1.


B. Khandekar-Vasilevsky (Ground 1D): Independent Claim 11 and the Dependent Claims Are Patentable Over the Khandekar-Vasilevsky System

For this ground, Microsoft proposes a combination in which Khandekar's provisioning environment incorporates "Vasilevsky's management capabilities (e.g., manager 212)" and "techniques." Pet. at 32-33. However, for the reasons discussed below, a POSITA would not have modified Khandekar with Vasilevsky because they are directed to different architectures and importing the teachings from Vasilevsky into Khandekar renders Khandekar inoperable for its intended purpose. But even assuming *arguendo* that a POSITA would have made the modification, the Khandekar-Vasilevsky system would not satisfy "virtual assets are managed by virtualization layer instances running on each of the one or more physical resources."

1. Khandekar and Vasilevsky Are Directed to Different Architecture Designs Which Are Incompatible with One Another

a) Khandekar's Architecture Includes Physical Resources That Each Deploy Multiple Virtual Machines

Khandekar's architecture employs physical hosts (servers) to deploy *multiple virtual machines* on each host. Khandekar's Figure 2, reproduced below, depicts physical hosts (green, below) that each deploy one or more VMs (blue, below):




Khandekar at Fig. 2 (annotated). As described in Section II.C.1 above, each host includes a VMM that operates at the host level to carry out provisioning operations on each host.

b) Vasilevsky's Grid Computing Architecture Provides a Single Virtualization System

In contrast to Khandekar, Vasilevsky describes a grid computing architecture where physical resources are abstracted to make a *single virtualized system*. See


Vasilevsky at ¶¶ 13, 17, 57, Figs. 1-2; *see also* Ex. 1002 at 93. Figure 1, reproduced below, depicts Vasilevsky's architecture that abstracts multiple physical systems (e.g., servers) 107A, 107B, and 107C (green, below) into one single virtualized system 101 (blue, below):



Vasilevsky at Fig. 1 (annotated); *see also id.* at ¶ 58.

Unlike Khandekar's architecture that deploys **multiple** virtual machines on each host, Vasilevsky's grid computing deploys a **single** virtualized system distributed across multiple hosts. *See, e.g.,* Ex. 1002 at 93 (citing Vasilevsky at ¶ 57). Moreover, unlike Khandekar that **manages at the host level** (per each host's virtual machine monitor (VMM)), Vasilevsky **manages at the system level** (per a distributed virtual machine monitor (DVMM)). *See* Vasilevsky at Abstract, ¶¶ 13, 65, 80. Specifically, because the physical hosts (green, below) have been abstracted into a single virtualized system (blue, below), Vasilevsky's manager 212 (red,

below) interacts with DVMM 203 (purple, below) of the virtualized system through a management agent on each node.



Vasilevsky at Fig. 2 (annotated). Accordingly, in this architecture design, manager 212 and its agents operate upon and within the single virtualized system abstracted from the physical resources.

Thus, Vasilevsky's and Khandekar's architectures are incompatible with one another: Vasilevsky's grid computing deploys a single virtualized system by abstracting physical resources and manages at the system level, whereas Khandekar deploys multiple virtual machines by expanding each physical resource and manages at the host level.

**2. Modifying Khandekar with Vasilevsky Would Have
Rendered Khandekar Inoperable for its Intended Purpose**

A POSITA would not have modified Khandekar with “Vasilevsky’s management capabilities (e.g., manager 212)” and “techniques” because Vasilevsky would have rendered Khandekar inoperable for its intended purpose of deploying multiple VMs on each physical server. *In re Gordon*, 733 F.2d 900, 902 (Fed. Cir. 1984) (holding that a proposed modification that renders the prior art “unsatisfactory for its intended purpose” fails to support a conclusion of obviousness); *see also* § VI.B.1, *supra*. The Office previously considered a Vasilevsky and Khandekar-like combination (Singh), and that combination was deemed inoperable. That same conclusion also applies for Microsoft’s Vasilevsky combination.

Whether Vasilevsky-Singh (prosecution) or Khandekar-Vasilevsky (petition), the combinations are inoperable because Vasilevsky is an entirely different architecture than Khandekar and Singh. Vasilevsky is directed to “grid computing system in which multiple computers are virtualized to generate a single virtualized computing system.” Ex. 1002 at 93 (citing Vasilevsky, ¶ 57); *id.* at 91. That is, “multiple computers [] generate a single virtualized system.” *Id.* at 93 (citing Vasilevsky at Fig. 1). During prosecution, the Examiner relied on Singh because Vasilevsky does not include a “virtualization layer instances running on each of the one or more physical resources.” *Id.* at 95 (modifying “Vasilevsky by the teaching of Singh”) (quoting Office Action, p.8). Singh’s architecture has physical servers

that “each compris[e] one or more virtual machines.” *Id.* at 94. But, as the Applicant argued, the modification—requiring each physical server to include multiple virtual machines in Vasilevsky—“would prevent the abstraction of each computer . . . in Vasilevsky to generate a single virtualized system.” *Id.* at 95. This would render Vasilevsky inoperable. *Id.*

Substituting Khandekar for Singh does not change the inoperability conclusion. Just like Singh, Khandekar's architecture has each physical server deploying multiple virtual machines. *Compare* Ex. 1002 at 94 (citing Singh's figure 1) *with* § VI.B.1.a., *supra* (annotation of Khandekar's architecture showing physical resources (green) with multiple virtual machines (blue)). Vasilevsky operates in the inverse. Its management architecture—manager 202, agents, DVMM—operates as an integral component of a single virtualized system that is abstracted from multiple physical servers. Ex. 1002 at 93 (citing Vasilevsky at Fig. 1).

Microsoft itself recognizes Vasilevsky's management architecture when stating that Vasilevsky capabilities and techniques are “implemented via a management server/agent arrangement in which the management server writes configuration information . . . that management agents read to configure distributed server/virtual server execution.” Pet. at 32 (citing Vasilevsky (MS-1005) at ¶¶ 85, 88, 155). That is, for Vasilevsky's manager 212 to carry out its management capabilities, it requires its corresponding DVMM and agents that operate within the

single virtualized system. Thus, for “Khandekar’s provisioning server 800 [to be] configured to include (or operate in conjunction with) a Vasilevsky-compatible manager”—and extend “*Vasilevsky’s management capabilities*”/“*virtualization-management techniques*” to the “Khandekar-Vasilevsky System”—Khandekar’s physical servers would need to be abstracted into a single virtualized computing system. Pet. at 32-34, 36. This would render Khandekar inoperable for its intended purpose of provisioning multiple virtual machines on each physical server.

The incompatibility is symmetric—Vasilevsky’s architecture abstracts multiple physical servers into one single virtualized computing system, while Khandekar expands each individual physical server into multiple virtual machines. These two architectural principles are mutually exclusive: introducing the teachings of one into the other destroys the other’s core function. Microsoft does not consider the distinct architectural differences between the references and how each independently operates. Instead, Microsoft simply, and perfunctorily states that because the “combination amounts to using known virtualization-management techniques,” “Khandekar’s provisioning server environment would have been incorporated [with] Vasilevsky’s management capabilities (e.g., manager 212).” Pet. at 33, 36. This conclusory statement alone is insufficient to show how Khandekar would be modified with Vasilevsky’s management techniques, much less address the different architectures and incompatibility between the two. *See,*

e.g., First Alert, Inc. v. Walter Kidde Portable Equip., Inc., No. IPR2023-00898, Paper 14 at 15 (P.T.A.B. Dec. 13, 2023) (denying institution because the petitioner did “not explain how a person of ordinary skill in the art would have ‘modif[ied]’” the primary reference based on the teachings of the secondary reference); *Lenovo (U.S.) Inc. v. Neodron Ltd.*, No. IPR2020-00688, Paper 8 at 17-18 (P.T.A.B. Sept. 15, 2020) (denying institution because the petition did not “provide any detail as to how a person of ordinary skill in the art would have modified Barkan’s leads 160 to incorporate external connections such as those in Tagg”).

**3. The Khandekar-Vasilevsky System Does Not Satisfy
“Virtual Assets [that] Are Managed by Virtualization Layer
Instances Running on Each Physical Resource”**

The Khandekar-Vasilevsky system would not satisfy “virtual assets [that] are managed by virtualization layer instances running on each physical resource.” Although for this limitation, Microsoft relies on Khandekar’s VMM, which cannot satisfy the “manage” limitation for the reasons set forth above, the modification would change the management architecture altogether. Pet. at 21 (addressing Claim [1.3]); *id.* at 38-39 (referring to its Claim [1.3] analysis). With this modification, the Khandekar-Vasilevsky system would also fail to meet this limitation but for different reasons.

The Khandekar-Vasilevsky system would manage at a single virtualization system abstracted from the physical resources. But the claims require that virtual

assets are managed by virtualization layer instances running on each physical resource. '915 Patent at Claims 1, 11. As discussed above, Vasilevsky's management techniques are carried out by a DVMM and management agents that operate within the single virtualization system. *See* § VI.B.1.b, *supra*. This operation extends to the Khandekar-Vasilevsky system. Because Vasilevsky's manager is not an extractable modular component that operates in isolation, the "Vasilevsky-compatible manager" would necessarily manage through the DVMM and management agents. *See* Pet. at 34-35 (Microsoft implementing Vasilevsky's manager at Khandekar's provisioning server). That management is thus through the DVMM which is distributed across the single virtualization system, not the VMM on each physical resource (as in Khandekar).

Accordingly, even if a POSITA would have made the combination, the "Khandekar-Vasilevsky System" would not have satisfied the claimed "virtual assets [that] are managed by virtualization layer instances running on each physical resource."

C. Hirschfeld Combinations (Grounds 2A and 2C): Independent Claims 1 and 11 Are Patentable Over the Hirschfeld Combinations

The Hirschfeld combinations (Smith and Smith-Vasilevsky) do not cure Hirschfeld's architecture deficiencies identified in Section II.C.2, above. Like Hirschfeld, the Hirschfeld combinations do not teach a virtualization layer that runs

on the physical resources (claim requirement 1); nor that the virtualization layer manages the virtual assets (claim requirement 2). '915 Patent at Claims 1 and 11.

First, Hirschfeld's virtualization software 229 running on the OS (requirement 1 deficiency) extends to the Hirschfeld combinations because Microsoft relies on Hirschfeld's virtualization software as the "virtualization layer." **Second**, Hirschfeld's requirement 2 deficiency also remains because, although Microsoft proposes incorporating a Smith-like VMM into Hirschfeld's virtualization software 229 to "manag[e] multiple logical servers on a single physical server" (*see* Pet. at 55), but Hirschfeld's SCM already manages the logical server at the cloud level. Microsoft overlooks this dispositive fact. **Finally**, setting aside these deficiencies, Microsoft fails to show how its Hirschfeld combination mapping is consistent with its own construction of "managed" as set forth in the -00296 IPR, which Microsoft is likewise held to here. *See* § IV, *supra*.

1. Hirschfeld's Virtualization Software Runs on the Operating System Layer, Not the Physical Resource

Hirschfeld's claim requirement one failure extends to the Hirschfeld combinations because Microsoft identifies "Hirschfeld's virtualization software 229" as the "virtualization layer instance running on that physical resource." Pet. at 64-65 (identifying "Hirschfeld's virtualization software 229" as the "virtualization layer instance") (Element [1.3]) (Ground 2A)); *see also id.* at 79 (Ground 2C) (referring to Element [1.3]). But Hirschfeld's virtualization software 229 runs on

the operating system (OS) rather than “running on the same physical server hardware,” as claimed. *Id.* at 55-56. The prosecution record draws the architectural boundary of the claim: virtualization software running on the OS layer “cannot equate to” virtualization layer running on the physical resource. Ex. 1002 at 91-92.


During prosecution, the Examiner applied the Singh reference against this claim limitation, which was overcome, in part, because its virtualization software runs on the operating system, not the physical resources. *See* Ex. 1002 at 91-92; *Phillips v. AWH Corp.*, 415 F.3d 1303, 1317 (Fed. Cir. 2005) (explaining that “the prosecution history provides evidence of how the PTO and the inventor understood the patent.”). The Applicant argued that Singh provides “virtual system software [that] runs on the operating system” which “cannot equate” to “virtualization layer instances running on each of the physical resources.” Ex. 1002 at 91-92.

Singh explicitly fails to teach or suggest the virtual assets are managed by virtualization layer instances running on each of the physical resources, as recited in claim 1 Singh requires that the server runs an operating system, ***and that the virtual system software runs on the operating system. Such a disclosure cannot equate to the virtual assets are managed by virtualization layer instances running on each of the physical resources,*** as recited in claim 1 as amended

Id. (emphasis modified). The prosecution record clearly laid out the architectural boundary the virtualization layer runs on—the physical resources, not the OS—and

thereby provided the scope of the claim. *Southwall Techs., Inc. v. Cardinal IG Co.*, 54 F.3d 1570, 1576 (Fed. Cir. 1995) (“The prosecution history limits the interpretation of claims so as to exclude any interpretation that was disclaimed during prosecution.”). Hirschfeld falls outside the claim scope because its virtualization software 229 runs on the OS, just like the Singh reference.

Each of Hirschfeld’s physical servers 108 includes three distinct architectural layers, a hardware layer 210 (red, below), an OS layer 212 (blue, below), and a software layer 214 with the virtualization software 229 (purple, below), as shown below. Hirschfeld at 8:39-41, 9:35-46, Fig. 2. The virtualization software (purple) runs on the OS layer (blue).



Id. at Fig. 2 (annotated).

Because Hirschfeld's virtualization software runs on the OS layer, "[s]uch a disclosure cannot equate to" the virtualization layer running on each of the physical resources. Ex. 1002 at 92.


2. A POSITA Would Not Have Used a Smith-Like VMM In Hirschfeld Because its SCM Manages its Logical Servers

The Hirschfeld combinations also do not resolve Hirschfeld's claim requirement 2 failure because Hirschfeld's SCM manages its logical server at the cloud layer—with no need for a Smith-like VMM. Microsoft's theory is grounded in the belief that Hirschfeld leaves open implementation details for logical server management configurations and therefore proposes incorporating a Smith-like VMM into Hirschfeld's virtualization software 229. Pet. at 55-56, 64-65 (proposing a Smith-like VMM to provide "mechanism for managing multiple logical servers on a single physical server.") (Ground 2A), 79 (Ground 2C referring to its Ground 2A analysis). But Hirschfeld leaves no logical server management implementation details open.

That is dispositive because the alleged problem motivating the combination is entirely illusory. Hirschfeld's novelty is its server cloud manager (SCM) and its management configuration. *See* § II.C.2, *supra*. There is thus no motivation to implement a Smith-like VMM into virtualization software 229 when the SCM already performs logical server management at the cloud layer. *Winner Int'l Royalty Corp. v. Wang*, 202 F.3d 1340, 1349 (Fed. Cir. 2000) (affirming a finding of no

motivation to combine because there was no problem with the existing prior art mechanism).

Hirschfeld discloses a bifurcated management architecture: (1) the SCM (blue, below) manages the logical servers on the cloud, and (2) the per-host virtualization software (purple, below) manages each physical server (green, below) it is located on. *See* § II.C.2, *supra*.



Hirschfeld at Fig. 2 (annotated). As figure 2 shows above, the virtualization software is on the OS layer in each physical server, and its management is expressly limited to the physical resources: “*manages physical resources* of the logical servers 106 linked thereto.” *Id.* at 9:44-46. The server cloud manager (SCM) is a separate entity that sits on a separate architectural plane and “*manages the logical servers* on the cloud. *Id.* at 12:45-49 (emphasis added); *see also id.* at 9:58-10:8 (“The SCM 102

is shown as a separate entity interfaced to the shared network 211 and the physical servers 108.”). Notably, in the -00296 IPR, Microsoft’s analysis is grounded in logical server management by the SCM. *See* ’050 Pet. at 55-56 (“***The SCM*** establishes and maintains ‘status and instance information’ for the logical servers”), 60-61 (expressing that the ***SCM*** “***manages/s/*** [the] physical and logical attributes of ***logical servers.***”), 63 (“***SCM-managed*** data structures”) (emphases added).

Accordingly, because Microsoft’s allegation that Hirschfeld lacks implementation details present an illusory, non-existent problem, a POSITA would not have been motivated to incorporate Smith’s VMM into Hirschfeld’s virtualization software 229. Hirschfeld’s claim requirement two deficiency thus remains intact. Accordingly, the Hirschfeld combinations do not satisfy virtual assets managed by the virtualization layer—regardless of Microsoft’s “managed” construction (discussed *infra*).

3. Microsoft’s Hirschfeld-Smith Mapping Does Not Satisfy Microsoft’s Own Construction of “Managed”

While Microsoft expressly construes “managed” under § 42.104(b)(3), it fails to show how its Hirschfeld-Smith mapping is consistent with its own construction as required under § 42.104(b)(4). *See* § IV, *supra*. As detailed in the Khandekar section, Microsoft’s construction of “managed” entails environment-level, platform-independent management. *See* § VI.A.1, *supra*. Despite Microsoft having committed its mappings to that construction, it fails to show how its Hirschfeld-

Smith virtualization software mapping is consistent with its own construction. Like for its Khandekar analysis, Microsoft does not address a single element of its construction in its Hirschfeld analysis.

Instead, Microsoft advances a single, bare-bone conclusion that its Hirschfeld-Smith-like virtualization software manages the logical servers. Pet. at 56 (concluding that the “Smith-like VMM manages multiple logical operating systems run on a physical server to support logical servers) (emphasis omitted). Microsoft does not even attempt to show how its mapping is consistent with the “management features” from its own construction that are performed at the environment-level and carried out irrespective of any (platform-independent), which it holds itself to. See § IV, *supra*. Microsoft thus cannot meet its burden under § 42.104(b)(4).

VII. CONCLUSION

For at least the foregoing reasons, Microsoft fails to demonstrate that there is a reasonable likelihood that it would prevail with respect to any challenged claim. Therefore, ToutVirtual respectfully requests that the Board deny institution of *inter partes* review of the '915 Patent.

Patent Owner's Preliminary Response
IPR2026-00295 (U.S. Patent No. 11,656,915)

Respectfully submitted,

Date: June 16, 2026

/James Hannah/

James Hannah (Reg. No. 56,369)
Kristopher Kastens (Reg. No. 57,517)
Elisa Santos (Reg. No. 83,399)
HERBERT SMITH FREEHILLS
KRAMER (US) LLP
333 Twin Dolphin Drive, Suite 700
Redwood Shores, CA 94065
Telephone: (650) 752-1700
james.hannah@hsfkramer.com
kris.kastens@hsfkramer.com
elisa.santos@hsfkramer.com

Jeffrey H. Price (Reg. No. 69,141)
Jeffrey Eng (Reg. No. 63,189)
Jenna Fuller (Reg. No. 74,212)
HERBERT SMITH FREEHILLS
KRAMER (US) LLP
1177 Avenue of the Americas
New York, NY 10036
Telephone: (212) 715-9100
jeffrey.price@hsfkramer.com
jeffrey.eng@hsfkramer.com
jenna.fuller@hsfkramer.com

Attorneys for Patent Owner
ToutVirtual, Inc.

CERTIFICATE OF COMPLIANCE WITH 37 C.F.R. § 42.24

The undersigned hereby certifies that the portions of the foregoing **PATENT OWNER'S PRELIMINARY RESPONSE** has 7,345 words in compliance with the word limit set forth in 37 C.F.R. § 42.24. This word count was prepared using the Microsoft Word word-processing system used to prepare this paper.

Dated: June 16, 2026

/James Hannah/
James Hannah (Reg. No. 56,369)

Attorneys for Patent Owner

CERTIFICATE OF SERVICE

The undersigned certifies, in accordance with 37 C.F.R. § 42.6(e), service was made on Petitioners as detailed below.

<i>Date of service</i>	June 16, 2026
<i>Manner of service</i>	Electronic Mail (IPR42342-0156IP1@fr.com)
<i>Documents served</i>	PATENT OWNER'S PRELIMINARY RESPONSE
<i>Persons Served</i>	Fish & Richardson P.C.: W. Karl Renner Usman Khan Ryan Chowdhury

/James Hannah/
James Hannah
Registration No. 56,369
Lead Counsel for Patent Owner