

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

MICROSOFT CORPORATION,
Petitioner,

v.

TOUTVIRTUAL, INC.,
Patent Owner.

Case IPR2026-00297
U.S. Patent No. 12,079,667

PATENT OWNER'S PRELIMINARY RESPONSE

TABLE OF CONTENTS

	Page
I. INTRODUCTION	1
II. OVERVIEW OF '667 PATENT AND THE ASSERTED GROUNDS.....	3
A. Overview of the '667 Patent.....	3
B. The '667 Patent Addressed the Industry's Need for Cross-Platform Virtualization Management.....	6
C. The Asserted Grounds	7
D. The Relevant Asserted Art	7
1. Overview of Hashimoto (Ex. 1005).....	7
2. Overview of Hirschfeld (Ex. 1007)	9
3. Overview of Lubbers (Ex. 1006)	11
III. LEVEL OF ORDINARY SKILL IN THE ART	12
IV. CLAIM CONSTRUCTION	13
V. INSTITUTION SHOULD BE DENIED BECAUSE MICROSOFT'S CLAIM CONSTRUCTION POSTURE IS INCONSISTENT ACROSS FORUMS.....	15
VI. INSTITUTION SHOULD BE DENIED BECAUSE MICROSOFT FAILS TO DEMONSTRATE THAT ANY CHALLENGED CLAIM IS UNPATENTABLE.....	17
A. Ground 1: Independent Claim 1 is Patentable Over Hashimoto in View of Lubbers.....	17
B. Ground 2: Independent Claims 9 and 18 Are Patentable Over Hashimoto in View of Lubbers and Hirschfeld	20

1.	Hashimoto Combined with Hirschfeld Does Not Reach the Claimed Inventions of Independent Claims 9 and 18.....	20
2.	Hashimoto-Lubbers-Hirschfeld Combination Does Not Satisfy Selecting Virtual Assets for Storage Based on an Allocation of Data on the Virtual Data Storage Assets (Claim 18).....	24
VII.	CONCLUSION.....	24

TABLE OF AUTHORITIES

	Page(s)
Cases	
<i>ACTV, Inc. v. Walt Disney Co.</i> , 346 F.3d 1082 (Fed. Cir. 2003)	13
<i>Am. Airlines, Inc. v. Intell. Ventures I LLC</i> , No. IPR2025-01055, Paper 11 (P.T.A.B. Nov. 21, 2025).....	16
<i>Ford Motor Co. v. Autoconnect Holdings LLC</i> , No. IPR2026-00171, Paper 11 (P.T.A.B. May 13, 2026)	15
<i>Revvo Techs., Inc. v. Cerebrum Sensor Techs., Inc.</i> , No. IPR2025-00632, Paper 20 (P.T.A.B. Nov. 3, 2025).....	1, 17
<i>Tesla, Inc. v. Intell. Ventures II, LLC</i> , No. IPR2025-00340, Paper 18 (P.T.A.B. Nov. 5, 2025).....	15
Statutes	
35 U.S.C. § 112.....	15, 16

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	Intentionally Omitted
2022	Intentionally Omitted
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	Intentionally Omitted
2025	Intentionally Omitted
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

I. INTRODUCTION

The Board should deny institution on three independent bases, each of which establishes that Petitioner (Microsoft) has not raised a reasonable likelihood that any claim of the '667 Patent is unpatentable. At every level—claim construction, element-by-element mapping, and the proposed combination of prior art references—Microsoft's Petition is internally incoherent.

First, Microsoft's Petition construes disputed claim terms and maps them to prior art while contending that the same terms are indefinite (and therefore incapable of construction) in parallel district court proceedings. The Petition offers no explanation for this inconsistency. Under *Revvo* and related authority, that unexplained inconsistency alone is grounds for denial.

Second, Microsoft has not demonstrated that the Hashimoto, Hirschfeld, and Lubbers, alone or in combination, discloses “selecting one or more of the virtual assets for storage of the set of data based on an allocation of data on the virtual storage assets,” as recited in Claims 1 and 18. Microsoft’s proposed construction of “allocation of data”—which purportedly encompasses “the *mapping of physical storage capacity to logical storage units* according to specific application critical rules”—ignores that the selection of a virtual asset must be based on the current allocation of data “*on the virtual storage assets.*” The Board should deny institution because Microsoft does not argue that the prior art teaches selecting a virtual asset

for data storage based on how data is already allocated across the selectable virtual assets.

Third, the Board should confirm the patentability of Claims 9-20 because Microsoft's motivation to combine Hashimoto with Hirschfeld is architecturally inconsistent. Hirschfeld's server cloud manager ("SCM") (the alleged "control center application") decouples virtual asset management from per-host physical-server control, which is the feature of Hashimoto's architecture that the Petition relies on to meet the requirement that "virtual assets are managed by virtualization layer instances running on each of the physical resources." Microsoft has not explained why a POSITA would be motivated to retain Hashimoto's virtualization layer when implementing Hirschfeld's SCM or demonstrated that a POSITA would have had a reasonable expectation of success building the proposed hybrid system.

ToutVirtual respectfully requests that the Board deny institution at least because Microsoft's grounds depend on claim terms Microsoft simultaneously contends are indefinite and amenable to construction, prior art mappings that do not meet the plain claim language, and reference combinations that are architecturally inconsistent.

II. OVERVIEW OF '667 PATENT AND THE ASSERTED GROUNDS

A. Overview of the '667 Patent

The '667 Patent is directed to an apparatus and method for automatic configuration management of virtualization environments across a computer network. The '667 Patent distinguishes between physical resources (physical devices such as servers, storage devices, and routers) and virtual assets (software instantiations of computer systems deployed on those physical resources, including virtual machines, virtual storage assets (e.g., logical units or databases), virtual routers, and virtual desktops). *Id.* at 6:43-48, 9:51-62, 20:55-21:16.

The '667 Patent's 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). The control center application can manage the virtualization environment regardless of the underlying processor, operating system, or virtualization platform:

The system, through the Control Center, provides active management of the virtualization environment by initiating automatic responses to operational situations that influence dynamic demands on the physical resources and virtual assets. In this way, multiple virtual environments can be managed through a single user interface of a Control Center application 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.

Id. at 3:11-4:6; *see also id.* at 3:53-57, 6:49-7:3.

The control center management console can include one or more modular components—including an asset manager, a provisioning manager, a dynamic application router, an optimizer, a performance manager, and a capacity planning manager—that collectively inventory, provision, optimize, and monitor virtual assets across the network. *Id.* at 4:7-19, 7:44-9:50, Fig. 2. The control center communicates with each physical resource directly or through an optional control agent installed on the host that acts as a universal adapter between the control center and the host's virtualization layer that creates and manages the virtual assets. *Id.* at 9:51-10:2, 10:30-54, Figs. 3-4.

Of particular relevance to the challenged claims, the '667 Patent discloses and claims “selecting one or more virtual assets for storage of data based on the current allocation of data on those assets.” *Id.*; *see also id.* at Claim 18. For example, the control center application can control how virtual assets are allocated, prioritized, and reconfigured based on real-time performance metrics—including utilization and data allocation—for each virtual asset. *Id.* at 17:60-18:7, 23:3-44. Thus, the selection of a virtual storage asset for data storage can be driven by the current allocation of data on the virtual storage assets, not merely by static hardware-level partitioning rules as disclosed in the asserted art, including Lubbers.

This integrated, platform-independent architecture addressed a pressing need in the virtualization industry for unified, cross-platform management.

B. The '667 Patent Addressed the Industry's Need for Cross-Platform Virtualization Management

ToutVirtual's approach solved an identified need in the industry for a unified, cross-platform virtualization management tool to confront the proliferation of platform-dependent hypervisors. Ex. 2011 at 2. While only "7% of respondents [in a survey conducted in mid-2008] had only one virtualization supplier in their test and development and production environments," virtualization platforms, like VMware, "focuse[d] on boosting its capabilities for managing its own hypervisors." *Id.* at 2, 3. The industry, therefore, identified a need to move away from vendor-specific management tools, which required separate interfaces for each platform or platform type. *Id.* at 2 ("[F]rom an IT perspective you don't want multiple management tools, repositories, teams of users. You want a single pane of glass.").

The industry received ToutVirtual's VirtualIQ with significant 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). For example, the industry players lauded the fact that ToutVirtual's VirtualIQ technology includes a unified management console that 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. 2014 at 1-2 ("VirtualIQ provides a 'single pane of glass' across physical and virtual infrastructure . . .").

C. The Asserted Grounds

The Petition asserts three grounds, two of which challenge Independent Claims 1, 9 and 18 (bolded):

Ground	Claims	Basis/Asserted Reference(s)
1	1-5, 7, 8	§103: Hashimoto, Lubbers
2	9-17, 18-20	§103: Hashimoto, Lubbers, Hirschfeld
3	6	§103: Hashimoto, Lubbers, Khandekar

Pet. at 1. Because Grounds 1 and 2 challenge the independent claims, the dependent claim ground, which depends on Independent Claim 1, also fails. Therefore, the arguments addressed herein regarding Ground 1 apply equally to dependent claim Ground 3.

D. The Relevant Asserted Art

1. Overview of Hashimoto (Ex. 1005)

Hashimoto's server/storage virtualization architecture is built on a control terminal (blue, below) communicating with hypervisors (red, below), which deploy virtual machines (VMs) (purple, below), and are located on each physical host. *See* Hashimoto at Abstract, ¶¶ 59, 84, 178-179, Fig. 2




FIG. 2

Id. at Fig. 2 (annotated). The control terminal “sends and receives control signals and control information” to the physical hosts, including instructions to create VMs.

Id. at ¶¶ 85, 100. The hypervisors on the physical hosts receive these instructions and execute them by logically partitioning the physical resources of the physical hosts to create and run VMs. *Id.* at ¶¶ 59, 84, 178-179.


Thus, Hashimoto’s system manages virtual machines with hypervisors installed on each physical host, in contrast with Hirschfeld (*infra*), which discloses

moving the virtual asset management layer from the physical host to a logical cloud system.

2. Overview of Hirschfeld (Ex. 1007)

Hirschfeld is built on a virtualized logical server cloud in which logical servers are abstracted from the physical servers and managed at the cloud. Hirschfeld at ¶¶ 2, 34, Fig. 2. In Hirschfeld, each logical server is managed by the SCM, which is a management layer that is removed from the physical host, and which allows user access to the logical servers on the cloud. Pet. at 49 (Hirschfeld describes a “‘virtualized logical server cloud’ in which users access logical servers through a portal-like user interface without needing to know the access information of the underlying physical servers.” (citing Hirschfeld at ¶¶ 8, 31, 66)).

Hirschfeld discloses a bifurcated management architecture: the SCM (blue) manages the logical servers on the cloud, and the per-host virtualization software (purple) manages each physical server (green).



Hirschfeld at Fig. 2 (annotated). The SCM thus “*manages*” the *logical servers* which are abstracted from each physical server (unlike in Hashimoto where management is through the hypervisors on each physical host). *Id.* at ¶ 46 (emphasis added); *see also id.* at ¶ 40 (“The SCM 102 is shown as a separate entity interfaced to the shared network 211 and the physical servers 108.”).

In order to implement its cloud management approach, Hirschfeld’s SCM implements a keyword/video/mouse (KVM) module on each logical server, each physical server, and the SCM itself. Hirschfeld at ¶¶ 11, 30-31. Accordingly, as FIG. 5C reproduced shows, each logical server includes a KVM resource 305 that generates the video and audio output a user receives, and the virtualization software on each underlying physical server implements a virtual KVM 321 that virtualizes

that KVM resource and interfaces it to the SCM through its internal KVM client 519 (blue):




FIG. 5C

Id. at Fig. 5C. The SCM's KVM service 520 passes the data to the user's remote KVM client 107. Hirschfeld at ¶¶ 43, 55.

3. Overview of Lubbers (Ex. 1006)

Lubbers discloses a virtualized storage system that partitions physical storage based on “variable levels of data protection.” Pet. at 36 (citing Lubbers at ¶ 128). The framework of Lubbers is built entirely around physical storage segments, redundancy levels, and drive-placement constraints. Lubbers at ¶¶ 55, 57, 60, 62, 65, 66, 69, 75-81. A Redundant Store (“RStore”), a fundamental unit of storage, is defined by a redundancy set descriptor (“RSD”) that specifies the physical member

drives, RAID type, and physical segment assignments. *Id.* at ¶¶ 62, 65. “RAID information, or more generically data protection information, describes the type of data protection, if any, that is implemented in [a] particular RStore,” for example, RAID-type protection or drive separation. *Id.* at ¶¶ 55, 62.

To implement this technique, Lubbers discloses implementing data protection rules that govern how physical storage is bound to the RStore, dictate which physical drives participate, and determine how data is distributed across the physical drives for redundancy purposes. This redundancy and data protection partitioning provides a fail-safe scheme that manages the capacity and performance of large storage volumes at the hardware level. *Id.* at ¶¶ 57, 69. This “data protection” technique improves redundancy to ensure that data survives a physical hardware failure, but it is facially unrelated to selecting a virtual asset for storing data based on how data is allocated across available virtual assets, as claimed. Pet. at 36 (The NSC stores rules “used to select the level of data protection.”); Lubbers at ¶¶ 55, 65 (NSC provides a redundancy layer and selects physical storage segments “based on the desired protection.”), ¶ 79 (describing RAID-5 placement constraints relate to segments on separate drives/enclosures to satisfy protection scheme).

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. Additional education could substitute for professional experience, and professional experience could substitute for education.” Pet. at 9 (citation omitted). ToutVirtual reserves the right to propose its own definition of a POSITA if trial is instituted.

IV. CLAIM CONSTRUCTION

The Board should not apply Microsoft’s “allocation of data” construction because it is inconsistent with the plain language of the claim. *See* Pet. at 4-5. Microsoft proposes that “allocation of data” encompasses “the mapping of physical storage capacity to logical storage units according to specific application critical rules.” *Id.* This construction, however, is unrelated to the surrounding claim language which concerns “an allocation of data *on the virtual storage assets.*” ’667 Patent at Claim 1 (emphasis added); *see also id.* at Claim 18; *ACTV, Inc. v. Walt Disney Co.*, 346 F.3d 1082, 1088 (Fed. Cir. 2003) (“[T]he context of the surrounding words of the claim also must be considered in determining the ordinary and customary meaning of those terms.”).

Microsoft’s incorrect construction is plucked from an embodiment that uses the term “allocate” in the context of “how storage in physical servers can be allocated by the virtual server.” Pet. at 4 (quoting ’667 Patent at 20:58-60) (emphasis omitted).

Just because the term “allocate” is used in this description does not, however, mean that “allocation of data” refers to “the mapping of physical storage capacity to logical storage units” when the term is read in the context of the surrounding claim language. That language confirms that the claimed allocation of data refers to the storage of data in a virtual asset “based on an allocation of data on the virtual storage assets,” which is facially unrelated to how that data is mapped to the underlying physical resources.

The '667 Patent's figure 18 supports the plain meaning interpretation of the term “allocation of data” because it illustrates and describes how virtual storage assets are selected based on the current allocation state. As explained in Section II.A above, this embodiment illustrates the process of checking real-time performance metrics (e.g., utilization data) of each virtual asset against ACRs and, based on the current allocation state, selecting which virtual asset should receive new data. *See also* '667 Patent at 23:3-44, Fig. 18.

Microsoft's interpretation of “allocation of data” ignores this critical language when arguing that the term encompasses mapping physical storage to logical storage based on specific application critical rules, defies the plain meaning of the claim. *See* Pet. at 4-5. The embodiments described with respect to figure 12 concern the configuration and management of physical storage allocated into logical units. '667 Patent at 5:17-19, 20:58-65 (figure 12 exemplifies “how storage in physical servers

can be allocated by the virtual server,” that is “a mapping between actual Physical Resources . . . to a collection of Logical Units”). It does not involve selecting virtual storage assets based on how data is currently allocated on the virtual storage assets, as claimed.

In view of the foregoing, the Board should reject Microsoft’s “allocation of data” construction and apply the plain meaning of the claim, which as discussed below, the Hashimoto-Lubbers does not meet.

V. INSTITUTION SHOULD BE DENIED BECAUSE MICROSOFT’S CLAIM CONSTRUCTION POSTURE IS INCONSISTENT ACROSS FORUMS

The Board should deny institution because Microsoft has taken *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.¹ *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

¹ The bolded language represents claim terms that Microsoft expressly construes in this proceeding. *See* Pet. at 3-5.

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: obtaining a set of data to be stored the virtual storage assets on any of the physical resources .	Indefinite under 35 U.S.C. § 112 (¶ 2)
Claim 1: selecting one or more of the virtual assets for storage of the set of data based on an allocation of data on the virtual storage assets.	Indefinite under 35 U.S.C. § 112 (¶ 2)
Claim 18: selecting one or more of the virtual assets for storage of the set of data based on an allocation of data on the virtual data storage assets.	Indefinite under 35 U.S.C. § 112 (¶ 2)
Claim 1: causing storage of the set of data at the selected virtual storage assets on any of the physical resources .	Indefinite under 35 U.S.C. § 112 (¶ 2)
Claim 18: causing storage of data, by the control center application, at the selected virtual data storage assets on each of the one or more physical resources .	Indefinite under 35 U.S.C. § 112 (¶ 2)
Claim 9: at least two of the virtual assets on the one or more of the physical resources include virtual storage assets	Indefinite under 35 U.S.C. § 112 (¶ 2)

Ex. 2027 at 96 (emphasis added).

Notably, the claim limitations Microsoft alleges are indefinite in the district court forum contain terms that Microsoft expressly construes in the Petition and are bolded in the table above. *See* Pet. at 3-5. 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. Denial is, therefore, necessary because Microsoft "takes alternative positions before the Board and a district court" without "explain[ing] 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).

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

A. Ground 1: Independent Claim 1 is Patentable Over Hashimoto in View of Lubbers

The combination of Hashimoto and Lubbers does not teach or suggest "selecting one or more of the virtual assets for storage of the set of data based on an allocation of data on the virtual storage assets," as recited in Independent Claims 1 and 18. Microsoft concedes that "Hashimoto leaves the details of how to partition the physical storage into virtual storage to the POSITA implementing the system," and relies on Lubbers' partitioning techniques. Pet. at 17. However, these partitioning techniques control how data is stored in *physical resources* based on data protection rules, not which virtual asset(s) should be chosen for storing data

based on how data is currently allocated across the selectable virtual assets, as claimed.

The plain claim language requires selecting virtual storage assets based on the current allocation of data on the *virtual storage assets*. *See* § IV, *supra*. As described in the specification, the system evaluates possible virtual assets for data storage based on real-time performance metrics (utilization/allocation data) associated with each virtual storage asset. '667 Patent at 23:3-44, Fig. 18. Microsoft's proposed construction of the "allocation of data," ignores this aspect of the claims. *See* Pet. at 5 ("[A] POSITA would have understood that the claimed 'allocation of data' encompasses 'the mapping of physical storage capacity to logical storage units according to specific application critical rules.'").

The Board should deny institution because Microsoft does not assert that the combination of Hashimoto and Lubbers discloses "selecting one or more of the virtual assets for storage of the set of data based on an allocation of data on the virtual storage assets" as properly interpreted. Microsoft relies on Lubbers' hardware level protection rules (RStores) for the "based on an allocation of data on the virtual storage assets." Pet. at 37-38; *see also* Lubbers at ¶ 19 ("RStores preferably implement a selected [level] of data protection."), ¶ 79. These rules, however, operate at the hardware level—concerning physical drive placement, parity distribution, and redundancy levels—rather than selecting virtual assets for storing

data based on how data is currently allocated on the virtual storage assets. *See* Lubbers at ¶¶ 57, 69; § II.C.3, *supra*.

Microsoft's own statements reveal the gap. Microsoft states that the rules dictate "how the data must be allocated on the virtual storage assets." Pet. at 37-38. The statement stands for how new data is distributed based on data protection levels—not which virtual storage asset is selected based on the current allocation of data on the virtual storage assets. *Id.* at 36 (discussing "various levels of protection and allow for construction of virtual disks with variable levels of data protection") (quoting Lubbers at ¶ 128) (emphasis omitted). Because how data is allocated based on hardware protection rules is not which virtual storage asset should receive the data based on the current allocation state on the virtual storage assets, Lubbers does not meet selecting based on an allocation of data on the virtual storage assets, as claimed.

In view of the foregoing, the Hashimoto-Lubbers combination does not disclose "selecting virtual assets for storage based on an allocation of data on the virtual storage assets."

B. Ground 2: Independent Claims 9 and 18 Are Patentable Over Hashimoto in View of Lubbers and Hirschfeld

1. Hashimoto Combined with Hirschfeld Does Not Reach the Claimed Inventions of Independent Claims 9 and 18

Independent Claims 9 and 18 are patentable at least because Microsoft has not established that a POSITA would have combined Hashimoto and Hirschfeld to reach the claimed invention, which requires (1) “virtualization layer instances on each of the physical resources” that manages virtual assets and (2) a control center application that obtains a user request and retrieves from or stores data to a virtual storage asset.

Even assuming that Hashimoto discloses the claimed virtualization layer and Hirschfeld teaches the claimed control center application (which ToutVirtual does not concede), the combination does not reach the claims because Hirschfeld's management system moves the virtual asset management layer off of the physical resource (as required by the claims). Microsoft has not explained why or how a POSITA would have retained Hashimoto's virtualization layer when implementing Hirschfeld's SCM as the alleged control center computer or demonstrated that a POSITA would have had a reasonable expectation of success building the proposed hybrid system.

Hirschfeld's SCM architecture works by removing the virtualization layer from the physical host system and implementing it as part of a managed cloud

environment. Hirschfeld's "[c]omplete logical separation is facilitated by the addition of the SCM, which is an automated multi-server management layer" that sits above the physical hosts. Hirschfeld at ¶ 65. Because the SCM maintains all status and instance information for the logical servers—including the persistent attributes that establish each logical server's identity—the management function must reside at the cloud layer rather than on any individual physical host. *Id.* at ¶¶ 46, 47. Only by operating at that layer can the SCM deliver its core benefit: ensuring that "the user does not have to know or provide any physical server information to access one or more logical server(s)" and that "the user experiences no change in access approach even when a logical server is reallocated to a different physical server." *Id.* at ¶ 66; *see also id.* at ¶ 40 ("The SCM 102 is shown as a separate entity interfaced to the shared network 211 and the physical servers 108.').

Microsoft cannot, therefore, have it both ways. In order to implement Hirschfeld's SCM to achieve its alleged benefits, the combined system could not retain Hashimoto's hypervisor, which runs on each physical server. *See, e.g.,* Pet. at 27-29 (relying on Hashimoto's hypervisor running on each physical server for the claim limitation "virtual assets are managed by virtualization layer instances running on each of the physical resources"). These two architectures are fundamentally incompatible because Hashimoto's control terminal manages VMs through a *hypervisor located at each physical host*, whereas Hirschfeld's server cloud

manager (“SCM”) manages logical servers (VMs) *operating at the cloud layer*.

Compare § II.C.1 with § II.C.2, supra.

This is not a plug and play of Hirschfeld's SCM/portal architecture into Hashimoto's control terminal. *See* Pet. at 52-53 (illustrating SCM box functionality overlaid onto Hashimoto's system). In fact, Microsoft's combination is flawed specifically because it concludes with no further explanation that Hashimoto's control terminal “provides a similar system control function as the SCM.” *Id.* at 55. It does not. To the contrary, to achieve Hirschfeld's SCM's user access functionality would introduce structural changes to Hashimoto that the Petition neither recognizes nor accounts for. Significantly, unlike Hashimoto, Hirschfeld “completely separate[s] logical servers from physical servers” such that users do not “experience[e] any change in access approach”—which is a purported benefit of the combination. *Id.* at 50. Therefore, to configure “Hirschfeld's SCM/portal architecture” in Hashimoto would introduce at least two structural changes in Hashimoto: (1) abstracting Hashimoto's VMs to a logical server cloud managed by the SCM, and (2) adding KVM functionality to the logical servers on the cloud communicating with the SCM.

Without that restructuring, which brings the combination outside the scope of the claims, there is no interface through which a user's device can access logical server data as Microsoft's Hashimoto combination demands. Pet. at 57 (“[I]n

Hashimoto-Lubbers-Hirschfeld, the virtual machine is controlled via the SCM/control terminal . . . “[i]n this manner, a user of any remote terminal 105 may login to and access any of the logical servers 106”. . . . Additionally, “the SCM 102 includes a user interface (I/F) 104 that provides a web page portal or other access mechanism”” such that a user “may login to and access any of the logical servers 106 of the virtualized logical server cloud 101”); *id.* at 60 (“[A] terminal 105 executes a keyboard/video/mouse (KVM) client 107 that transmits keyboard and mouse commands to a selected logical server 106.”). Thus, implementation of Hirschfeld’s SCM architecture would place Hashimoto’s VMs in a logical server cloud managed by the SCM—as this is the way the SCM allows a user to interface with the logical servers and Microsoft’s combination does not show otherwise (in fact, it is entirely missing analysis on *how* Hirschfeld’s SCM architecture would be implemented).

Accordingly, the proposed combination presents an irreconcilable dilemma: if Hirschfeld's SCM architecture is implemented in Hashimoto—as Microsoft's combination requires—the resulting system abstracts virtual asset management away from the physical hosts, thereby defeating the claim requirement that “virtual assets [be] managed by virtualization layer instances on each of the physical resources.” Conversely, if Hashimoto's hypervisor-based architecture is retained, the system lacks the SCM portal and KVM infrastructure through which a user can

request retrieval of data from, or storage of data to, the virtual storage assets, as independently claimed. Either way, the Hashimoto-Hirschfeld combination does not—and cannot—render obvious Challenged Claims 9 and 18.

2. Hashimoto-Lubbers-Hirschfeld Combination Does Not Satisfy Selecting Virtual Assets for Storage Based on an Allocation of Data on the Virtual Data Storage Assets (Claim 18)

The Hashimoto-Lubbers combination fails to disclose Claim 18's selecting based on "an allocation of data on the virtual data storage assets" for the same reason it fails to disclose Claim 1. *See* § VI.A, *supra*. While Ground 2 includes the addition of Hirschfeld, Hirschfeld is not relied on for the selection limitation. *See* Pet. at 75 ("See [1.4], which is unaffected by the addition of Hirschfeld"). Therefore, Ground 2's Hashimoto-Lubbers-Hirschfeld combination fails for the same reason its Hashimoto-Lubbers combination of Ground 1 fails.

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 requests that the Board deny institution of *inter partes* review of the '667 Patent.

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 4,710 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 (IPR42342-0157IP1@fr.com)
<i>Documents served</i>	PATENT OWNER'S PRELIMINARY RESPONSE
<i>Persons Served</i>	Fish & Richardson P.C.: W. Karl Renner Usman Khan David Holt

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