

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

KINGSTON TECHNOLOGY COMPANY, INC., KINGSTON TECHNOLOGY
CORPORATION, AND KINGSTON DIGITAL, INC.,
Petitioners,

v.

VERVAIN, LLC,
Patent Owner.

IPR2025-00616
U.S. Patent No. 10,950,300

**PATENT OWNER'S PRELIMINARY RESPONSE ON THE MERITS AND
OTHER NON-DISCRETIONARY CONSIDERATIONS**

TABLE OF CONTENTS

	Page
EXHIBIT LIST	v
I. INTRODUCTION	1
II. OVERVIEW OF THE CHALLENGED PATENT AND THE CHALLENGED CLAIMS	2
A. SLC and MLC Flash	3
B. Address Table.....	4
C. Data Integrity Tests	5
D. Independent Claims.....	10
1. Claim 1	11
2. Claim 12	13
III. PERSON OF ORDINARY SKILL IN THE ART	15
IV. CLAIM CONSTRUCTION	15
V. OVERVIEW OF THE CITED PRIOR ART	17
A. Gavens (Ex. 1045).....	18
1. First Embodiment (Figs. 14A-14B).....	21
2. Second Embodiment (Figs. 16A-16C).....	24
B. Secondary References Related To Random Access Volatile Memory	26
1. Gorobets 912 (Ex. 1049).....	27
2. Paley (Ex. 1051).....	28
3. Harari (Ex. 1057)	30
4. Gorobets PCT (Ex. 1058)	31

VI.	REASONS FOR DENYING INSTITUTION.....	32
A.	Ground 1: The Petition Does Not Establish That Claims 1 and 12 Would Have Been Obvious Over Gavens.....	32
1.	[1d] and [12d]: “at least one <u>SLC nonvolatile memory element</u> that can be mapped into the SLC memory space”	32
2.	[1e] and [12e]: “at least one <u>random access volatile memory</u> ”	38
3.	[1g] and [12f]: “the controller controlling access of the MLC and SLC nonvolatile memory elements and the random access volatile memory for storage of data therein, the controller, in at least a Write access operation to the MLC nonvolatile memory element, operable to store data in the MLC nonvolatile memory element and <u>retain such stored data in the random access volatile memory</u> ”	44
4.	[1h] and [12g]: “the controller performing a data integrity test on stored data in the MLC nonvolatile memory element after at least a Write access operation performed thereon by <u>comparing the stored data to the retained data in the random access volatile memory</u> ”	49
B.	The Petition Does Not Establish That Dependent Claims 2-11 Would Have Been Obvious Over Gavens.....	58
VII.	CONCLUSION.....	58

TABLE OF AUTHORITIES

	Page(s)
CASES	
<i>Acon Inc. v. Amo Dev., LLC</i> , IPR2021-00858, 2022 WL 17585777 (P.T.A.B. Dec. 12, 2022).....	34
<i>Amazon.com, Inc., et al. v. Avago Technologies General IP (Singapore)</i> <i>PTE. LTD.</i> , IPR2017-01112, Paper 11 (P.T.A.B. Aug. 23, 2017).....	43
<i>Application of Ratti</i> , 270 F.2d 810 (C.C.P.A. 1959)	37, 57
<i>Callaway Golf Co. v. Acushnet Co.</i> , 576 F.3d 1331 (Fed. Cir. 2009)	30, 31
<i>Captioncall, LLC v. Ultratec, Inc.</i> , No. IPR2020-01215, 2021 WL 278381 (P.T.A.B. Jan. 27, 2021)	41, 54
<i>Gaham v. John Deere Co. of Kansas City</i> , 383 U.S. 1 (1966).....	42, 45
<i>In re De Seversky</i> , 474 F.2d 671 (C.C.P.A. 1973)	30
<i>In re Kahn</i> , 441 F.3d 977 (Fed. Cir. 2006)	34, 46
<i>In re Stepan Co.</i> , 868 F.3d 1342 (Fed. Cir. 2017)	34
<i>Masimo Corp. v. Apple Inc.</i> , IPR2023-00745, Paper 44 (P.T.A.B. Oct. 15, 2024).....	43
<i>Nearmap US, Inc. v. Eagle View Technologies, Inc.</i> , IPR2024-00716, Paper 9 (P.T.A.B. Oct. 9, 2024).....	43
<i>Netflix, Inc. v. DivX, LLC</i> , 84 F.4th 1371 (Fed. Cir. 2023)	43
<i>Phillips v. AWH Corp.</i> , 415 F.3d 1303 (Fed. Circ. 2005) (<i>en banc</i>)	16

<i>Samsung Elecs. Co., Ltd. v. Mojo Mobility Inc.,</i> IPR2023-01099, 2024 WL 457731 (P.T.A.B. Feb. 6, 2024)	37
<i>Toyota Motor Corp. v. Cellport Systems, Inc.,</i> IPR2015-00633, 2015 WL 4934778 (P.T.A.B. Aug. 14, 2015).....	16
<i>Vervain, LLC v. Phison Electronics Corp.,</i> 1:24-cv-00259 (W.D. Tex.)	16
<i>Vivid Techs., Inc. v. Am. Sci. & Eng’g, Inc.,</i> 200 F.3d 795 (Fed. Cir. 1999)	16

STATUTES

35 U.S.C.	
§ 103.....	17
§ 312(a)(3)	42
§ 313.....	1
§ 314(a)	1

OTHER AUTHORITIES

37 C.F.R.	
§ 42.6(a)(3)	34
§ 42.22(a)(2)	34, 43
§ 42.107.....	1
§ 100(b).....	16

EXHIBIT LIST

Exhibit No.	Description	Previously Submitted
2001	Declaration of Dr. Sunil Khatri from <i>Phison Electronics Corporation v. Vervain, LLC</i> , IPR2025-00215, Ex. 2001 (P.T.A.B. May 14, 2025) (expert report filed with Patent Owner’s Preliminary Response)	X
2002	<i>Vervain, LLC v. Western Digital Corp. et al.</i> , No. 6:21-cv-488-ADA, Dkt. 41 (W.D. Tex. Jan. 24, 2022) (Claim Construction Order)	X
2003	<i>Vervain, LLC v. Kingston Technology Company, Inc., et al.</i> , No. 1:24-cv-254-ADA, Dkt. 19 (W.D. Tex. Jul. 15, 2024) (agreed scheduling order)	X
2004	<i>Vervain, LLC v. Kingston Technology Company, Inc., et al.</i> , No. 1:24-cv-254-ADA, Dkt. 56 (W.D. Tex. Jan. 28, 2025) (order denying motion to transfer venue)	X
2005	<i>Vervain, LLC v. Phison Electronics Corporation</i> , No. 1:24-cv-259-ADA, Dkt. 16 (W.D. Tex. Jul. 15, 2024) (agreed scheduling order)	X
2006	“Judge Albright Patent FAQ,” https://www.txwd.uscourts.gov/for-attorneys/judge-albright-courtroom-faq/ . Accessed December 18, 2024.	X
2007	Judge Albright’s Standing Order Governing Proceedings (OGP) 4.4—Patent Cases, Filed January 23, 2024.	X
2008	Preliminary Claim Constructions for <i>Vervain, LLC v. Phison Electronics Corporation</i> , No. 1:24-cv-259-ADA and <i>Vervain, LLC v. Kingston Technology Company, et al.</i> , No. 1:24-cv-254-ADA – email from Mr. Brown received on February 5, 2025.	X
2009	Markman Hearing Transcript dated February 6, 2025, for <i>Vervain, LLC v. Phison Electronics Corporation</i> , No. 1:24-	X

Exhibit No.	Description	Previously Submitted
	cv-259-ADA and <i>Vervain, LLC v. Kingston Technology Company, et al.</i> , No. 1:24-cv-254-ADA.	
2010	<i>Vervain, LLC v. Western Digital Corp. et al.</i> , No. 6:21-cv-00488-ADA, Dkt. 180 (W.D. Tex. Jul. 26, 2023) (Redacted Copy of Order on the Pending Motions and Motions <i>In Limine</i>).	X
2011	<i>Phison Electronics Corp. v. Vervain, LLC</i> , PGR2024-00047, Paper 9 (P.T.A.B. Mar. 17, 2025) (Decision Denying Institution of Post-Grant Review of U.S. Patent No. 11,830,546)	X
2012	<i>Phison Electronics Corp. v. Vervain, LLC</i> , PGR2024-00048, Paper 9 (P.T.A.B. April 28, 2025) (Decision Denying Institution of Post-Grant Review of U.S. Patent No. 11,854,612)	X
2013	Docket Sheet for <i>Vervain, LLC v. Kingston Technology Company, Inc., et al.</i> , No. 1:24-cv-254-ADA (W.D. Tex.). Accessed May 19, 2025.	X
2014	Excerpts from Kingston’s Final Invalidity Contentions served on February 28, 2025 in <i>Vervain, LLC v. Kingston Technology Company, Inc., et al.</i> , No. 1:24-cv-254-ADA (W.D. Tex.)	X
2015	“Interim Processes for PTAB Workload Management Memorandum,” issued March 26, 2025, https://www.uspto.gov/sites/default/files/documents/InterimProcesses-PTABWorkloadMgmt-20250326.pdf . Accessed May 13, 2025.	X
2016	Docket Sheet for <i>Vervain, LLC v. Phison Electronics Corporation</i> , No. 1:24-cv-259-ADA (W.D. Tex.). Accessed May 19, 2025.	X

Exhibit No.	Description	Previously Submitted
2017	<i>Vervain, LLC v. Western Digital Corp. et al.</i> , No. 6:21-cv-00488-ADA, Dkt. No. 183, Pretrial Hearing Redacted Transcript (W.D. Tex. Aug. 7, 2023).	
2018	Chen et al., <i>Ultra MLC Technology Introduction</i> , Advantech Technical White Paper (Oct. 5, 2012) (“Chen”)	
2019	Excerpts from Micheloni et al., <i>Inside NAND Flash Memories</i> (1 st ed. 2010) (“Micheloni”)	
2020	<i>Vervain v. Western Digital Corp. et al.</i> , No. 6:21-cv-00488-ADA, Dkt. 180 (W.D. Tex. Jul. 26, 2023) (Redacted Copy of Order on the Pending Motions and Motions <i>In Limine</i>).	

I. INTRODUCTION

Vervain, LLC (“PO” or “Vervain”) submits this Preliminary Response in accordance with 35 U.S.C. § 313 and 37 C.F.R. § 42.107, responding to the Petition for *Inter Partes* Review (IPR) (Paper 1) of U.S. Patent No. 10,950,300 (“300 Patent” or “the Challenged Patent”) filed by Kingston Technology Company, Inc., Kingston Technology Corporation, and Kingston Digital, Inc. (“Kingston” or “Petitioners”). The Petition includes a single ground of unpatentability challenging claims 1-12 (“Challenged Claims”). This ground relies on single-reference obviousness, based on U.S. Patent No. 8,634,240 (“Gavens”) (Ex. 1045).

The Board should deny institution because Petitioner’s ground is meritless. It fails to reconcile critical differences between the Challenged Claims and Gavens—namely, the Challenged Claims require retaining stored data in RAM, a feature not taught or suggested by Gavens. The Board has already denied institution of petitions including similar challenges based on this same deficiency. *See* Ex. 2011, 76-79 (denying institution of PGR2024-00047); Ex. 2012, 42-44 (denying institution of PGR2024-00048). As these prior decisions reflect, Petitioner has not shown a reasonable likelihood of prevailing on any ground. *See* 35 U.S.C. § 314(a).

Accordingly, PO requests that the Board deny institution of the Petition.

II. OVERVIEW OF THE CHALLENGED PATENT AND THE CHALLENGED CLAIMS

The Challenged Patent, entitled “Lifetime Mixed Level Non-Volatile Memory System” was filed on June 12, 2018, and has an effective filing date of July 19, 2011. Ex. 1001. Dr. G. R. Mohan Rao is the sole named inventor of the Challenged Patent.

At a high level, the Challenged Patent describes a reliable flash memory storage system combining both single-level cell (SLC) and multi-level cell (MLC) non-volatile memories.¹ Ex. 1001, Abstract; Ex. 2001, ¶37. Prior to the Challenged Patent, Dr. Rao recognized that “MLC NAND flash SSDs are slowly replacing and/or coexisting with SLC NAND flash in newer SSD systems” because “MLC flash memory is less expensive than SLC flash memory[] on a cost per bit basis.” Ex. 1001, 3:36-39, 5:51-54. However, while “MLC NAND flash enjoys greater density than SLC NAND flash” it comes “at the cost of a decrease in access speed and lifetime (endurance).” *Id.*, 3:41-44. As a result, various hybrid systems combining SLC and MLC (among others) are taught in the Challenged Patent, to combine the benefits of both types of non-volatile flash storage at a low cost. *Id.*, 3:67-4:2.

¹ Flash memory is a specific type of non-volatile memory, where data is stored in “blocks” of “pages.” Ex. 1001, 2:52-3:2; Ex. 2001, ¶27.


The Challenged Patent addresses improvements and solutions for managing the writing of data optimally for improved reliability and lifetime (endurance) of such hybrid memory systems. *Id.*, 3:62-4:2. Specifically, the Challenged Claims are directed to techniques for efficiently using SLC and MLC flash to improve the overall performance of the memory. *Id.*, claim 1. For example, if certain data is used more frequently, then it is transferred to higher-performance SLC. *Id.* By doing so, the number of errors is reduced, and overall endurance of the memory is increased. *Id.*, 3:67-4:2.

A. SLC and MLC Flash

SLC memory stores 1 bit per cell, and MLC memory stores more than 1 bit per cell. Ex. 1001, 2:17-20; Ex. 2001, ¶¶32-34. As noted above, there are advantages and disadvantages to SLC and MLC flash. In general, SLC is faster and less prone to errors, but requires more space and power to store a given amount of data. Ex. 1001, 1:58-62. The opposite is true of MLC. MLC flash is slower and more prone to errors, but stores data more densely with less power consumption. *Id.*, 3:41-44.

SLC and MLC flash memories both use floating gate transistors. *Id.*, 3:52. They store a charge in the floating gate of each transistor (cell), which changes the threshold voltage of the transistor. Ex. 2001, ¶32. The memory uses the threshold

voltage to determine what bit, or bits, were stored in the transistor. The MLC cell in the figure below illustrates threshold voltages for a 2-bit MLC cell.



Ex. 2018, 3.

The primary difference between SLC and MLC is what data each threshold voltage represents. Ex. 2001, ¶33. With SLC flash, the transistor stores only a 1 or 0, so a wide range of threshold voltages can be allotted to a single bit. Ex. 1001, 2:59-65. This allows for faster and more reliable memory access. On the other hand, MLC flash must be slowly and carefully programmed to a narrower, more precise range of threshold voltages, with each threshold voltage range representing a specific pair of bits (*see* figure above, which shows four pairs of bits—11, 10, 01, and 00—corresponding to smaller ranges of threshold voltages compared to the SLC). Ex. 1001, 3:38-41.

B. Address Table

To provide wear leveling, garbage collection, and bad block management, a translation layer is used to map logical addresses to actual physical memory


locations. Ex. 2019, 40-42; Ex. 1001, 3:3-3:35; *see also* Ex. 2001, ¶35. As part of this translation layer, “tables are widely used in order to map sectors and pages from logical to physical.” Ex. 2019, 40-42; Ex. 1001, 3:19-3:23. These tables map logical blocks to physical blocks. *Id.* Using a “block” or similar granularity is important, since flash memory is arranged so that when erasing and rewriting data, a whole block is “erased together.” Ex. 2019, 22; Ex. 1001, 2:59-3:2. Dr. Rao explained that “[t]he address ranges within the translation table will assume some minimum quantum, such as, for example, one block....” Ex. 1001, 5:54-58, Figs. 3a-3b. Dr. Rao further explained that memory is written and mapped on the granularity of a “quantum.” *Id.* During operation of the flash memory, logical addresses are frequently remapped to new physical locations. *Id.*, 3:19-3:52, 4:26-4:36, 5:57-67. Over time, a particular logical address may be mapped or associated with many different physical locations (blocks). And multiple logical addresses may point to the same block over time, so there is not a one-to-one correspondence between the logical addresses and the blocks over time.

C. Data Integrity Tests

As mentioned above, when data is stored in MLC memory, it is more prone to errors, and some data is more prone to errors than other data. Ex. 2001, ¶36. One reason for this is that the threshold voltage intervals for MLC memory are smaller than the intervals for SLC memory, and thus, errors can occur when writing or

reading the data. *Id.*, ¶36. Errors can also be caused by the data stored in neighboring cells. *Id.* A data integrity test checks the integrity of the data (*i.e.*, whether errors have occurred). This test can be run immediately after data is written, or at a later time. If the test reveals a problem such as corrupt data, the data can be remapped to SLC (which is less error-prone) or, alternatively, other MLC blocks, and the address table is modified accordingly. Ex. 1001, 3:13-3:35, 4:30-36.

The Challenged Patent claims a particular way of writing the data to the MLC and performing the data integrity test. Dr. Rao recognized that there were advantages in using a dynamic random access memory (DRAM) to assist with both the write access operation and data integrity test. As seen in Figure 1 below, the controller (purple) controls a DRAM (red).




Ex. 1001, Fig. 1 (annotations added²), 5:37-38.

The DRAM is used during a write access operation to the MLC flash (highlighted blue). The data to be written to the MLC flash is maintained (or retained) in the DRAM. Ex. 1001, 5:59-61. If time permits, the controller performs a data integrity test on the data that was stored in the MLC. *Id.*, 5:61-64. If the

² Unless otherwise noted, PO added coloring to Figures.

stored data fails the data integrity test, the address range for the stored data is remapped to a new location (*e.g.*, the SLC flash). *Id.*, 5:64-67.

Figures 3A and 3B show in more detail how the DRAM is used during the write access operation and data integrity test.



Id., Fig. 3A.

The write access operation begins at step 100. *Id.*, 6:17-22. In step 102, the data to be written to the MLC is read from the DRAM into the controller memory. *Id.* In step 104, the controller obtains the logical and physical address ranges for the data to be written to the MLC. *Id.*, 6:22-25. In step 106, the controller combines the

data from the DRAM with data from the MLC. *Id.*, 6:25-27. In step 108, the controller erases the physical address range where the data is being written. *Id.*, 6:27-28. Finally, in step 110, the combined data is written to the MLC flash. *Id.*, Fig. 3A (step 110), 6:28-30.

The data integrity test begins at the bottom of Figure 3A and continues to Figure 3B (below). In step 112, the data that was stored in the MLC (“the stored data”) is read into the controller memory. *Id.*, 6:30-32. In step 114, the stored data is compared to the data that was retained in the DRAM (“the retained data”). *Id.*, 6:33-38.




FIG. 3b

Id., Fig. 3B.

If there is a match, the data was successfully stored in the MLC. *Id.*, 6:38-40.

If not, the retained data is written to another location (*e.g.*, SLC). *Id.*, 6:40-44.

D. Independent Claims

The Challenged Patent includes two independent claims: Claims 1 and 12.

1. Claim 1

Claim 1 of the Challenged Patent requires a system for storing data comprising: (1) non-volatile MLC memory; (2) non-volatile SLC memory; and (3) random access volatile memory (highlighted red below). The random access volatile memory is a key feature of the invention and it permeates many of the elements of the claim ([1e]-[1h]). As discussed above, the random access volatile memory is particularly important to the “data integrity test” in [1h].

Claim 1	
[1Pre]	A system for storing data comprising:
[1a]	memory space containing volatile memory space and nonvolatile memory space, wherein the nonvolatile memory space includes both multilevel cell (MLC) memory space and single level cell (SLC) memory space;
[1b]	at least one controller to operate memory elements and associated memory space;
[1c]	at least one MLC nonvolatile memory element that can be mapped into the MLC memory space;
[1d]	at least one SLC nonvolatile memory element that can be mapped into the SLC memory space;
[1e]	at least one random access volatile memory ;

[1f]	an FTL flash translation layer, wherein the at least one controller, or FTL, or a combination of both maintain an address table in one or more of the memory elements and random access volatile memory ;
[1g]	the controller controlling access of the MLC and SLC nonvolatile memory elements and the random access volatile memory for storage of data therein, the controller, in at least a Write access operation to the MLC nonvolatile memory element, operable to store data in the MLC nonvolatile memory element and retain such stored data in the random access volatile memory ;
[1h]	the controller performing a data integrity test on stored data in the MLC nonvolatile memory element after at least a Write access operation performed thereon by comparing the stored data to the retained data in the random access volatile memory ;
[1i]	wherein the address table maps logical and physical addresses adaptable to the system, wherein the mapping is performed as necessitated by the system to maximize lifetime, and wherein the mapping maps blocks, pages, or bytes of data in either volatile or nonvolatile, or both, memories; and
[1j]	wherein a failure of the data integrity test performed by the controller results in a remapping of the address space to a different physical range of addresses and transfer of data corresponding to the stored data to those remapped physical addresses from those determined to have failed the data integrity test to achieve enhanced endurance.

In [1g], the controller stores data in the MLC non-volatile memory and retains that stored data in the random access volatile memory. The controller performs a data integrity test to determine whether the data has any errors ([1h]). To do so, the controller compares the “stored data” in the MLC to the “retained data” in the random access volatile memory ([1h]).

2. Claim 12

Claim 12 is similar to Claim 1, and also requires a system for storing data comprising: (1) non-volatile MLC memory; (2) non-volatile SLC memory; and (3) random access volatile memory.

Claim 12	
[12Pre] (same as [1Pre])	A system for storing data comprising:
[12a] (similar to [1a])	memory space containing volatile memory space and nonvolatile memory space, wherein the nonvolatile memory space includes both multilevel cell (MLC) space and single level cell (SLC) space;
[12b] (similar to [1b])	at least one controller to operate memory elements and associated memory space, and to maintain an address table in one or more of the memory elements;
[12c] (similar to [1c])	at least one MLC nonvolatile memory element that can be mapped into the nonvolatile memory space;

[12d] (similar to [1d])	at least one SLC nonvolatile memory element that can be mapped into the nonvolatile memory space;
[12e] (same as [1e])	at least one random access volatile memory;
[12f] (same as [1g])	the controller controlling access of the MLC and SLC nonvolatile memory elements and the random access volatile memory for storage of data therein, the controller, in at least a Write access operation to the MLC nonvolatile memory element, operable to store data in the MLC nonvolatile memory element and retain such stored data in the random access volatile memory;
[12g] (same as [1h])	the controller performing a data integrity test on stored data in the MLC nonvolatile memory element after at least a Write access operation performed thereon by comparing the stored data to the retained data in the random access volatile memory;
[12h] (same as [1i])	wherein the address table maps logical and physical addresses adaptable to the system, wherein the mapping is performed as necessitated by the system to maximize lifetime, and wherein the mapping maps blocks, pages, or bytes of data in either volatile or nonvolatile, or both, memories; and
[12i] (same as [1j])	wherein a failure of the data integrity test performed by the controller results in a remapping of the address space to a different physical range of addresses and transfer of data corresponding to the stored data to those

	remapped physical addresses from those determined to have failed the data integrity test to achieve enhanced endurance.
--	---

As shown above, Claims 1 and 12 share many limitations, and are substantially the same for purposes of the arguments made in this Patent Owner Preliminary Response. Particularly, Limitations [12d], [12e], [12f], and [12g] are the same as limitations [1d], [1e], [1g], and [1h], respectively.

III. PERSON OF ORDINARY SKILL IN THE ART

For purposes of this Preliminary Response only, PO adopts Petitioner’s definition of a person of ordinary skill in the art (POSA). Petition, 23; Ex. 2001, ¶¶22-26.³

IV. CLAIM CONSTRUCTION

Claim terms are afforded “their ordinary and customary meaning,” which is “the meaning that the term would have to a person of ordinary skill in the art in

³ PO submits the declaration of Dr. Sunil Khatri—an expert in the field of the Challenged Patent—previously filed with the patent owner preliminary response in IPR2025-00215. Ex. 2001, ¶¶1-26. “Kingston’s petition [in IPR2025-00616] and the petition in the Phison 300 IPR [IPR2025-00215] are substantively identical[.]” Paper 3 (Motion for Joinder to Phison’s IPR2025-00215 to assume an “understudy” role), 1-2.

question at the time of the invention.” *Phillips v. AWH Corp.*, 415 F.3d 1303 (Fed. Circ. 2005) (*en banc*); 37 C.F.R. § 100(b). The Board, however, only construes claim terms when necessary to resolve an underlying controversy. *Toyota Motor Corp. v. Cellport Systems, Inc.*, IPR2015-00633, 2015 WL 4934778, at *8 (P.T.A.B. Aug. 14, 2015).

Petitioner proposes construction of eight terms: “MLC non-volatile memory,” “SLC non-volatile memory,” “elements” or “memory element,” “memory space,” “controller/FTL,” “mapping...to maximize lifetime,” “remapping...and transfer...to achieved enhanced endurance,” and “data integrity test.” Petition, 23-27. PO disagrees with all of Petitioner’s proposed constructions, but construction of these terms is not necessary to resolve the parties’ dispute regarding whether institution should be granted. *See Vivid Techs., Inc. v. Am. Sci. & Eng’g, Inc.*, 200 F.3d 795, 803 (Fed. Cir. 1999). Particularly, even if Petitioner’s proposed constructions are applied, Petitioner’s grounds still fail, as explained below. Accordingly, PO does not contest Petitioner’s proposed constructions for these terms for the limited purpose of this Preliminary Response.

Further, most of Petitioner’s proposals mirror those already presented in the parallel district court litigation between Vervain and Phison (*Vervain, LLC v. Phison Electronics Corp.*, 1:24-cv-00259 (W.D. Tex.) (the “Parallel Litigation”)), where they were fully briefed and, during a *Markman* hearing on February 6, 2025, **already**

rejected by the Court. See Exs. 2008-2009. Specifically, the Court rejected Petitioner's proposed constructions, instead adopting PO's proposed constructions, for "MLC non-volatile memory," "SLC non-volatile memory," "memory element," "memory space," "controller" and "FTL," "the mapping is performed as necessitated by the...system to maximize lifetime," "remapping...and transfer of data...to achieved enhanced endurance," and "data integrity test." See Ex. 2008, 3-5 (providing preliminary constructions for these terms prior to the *Markman* hearing); Ex. 2009 (hearing transcript from the *Markman* hearing, where the Court adopted all its preliminary constructions); *see also* Ex. 2002 (Claim Construction Order from an earlier case involving the Challenged Patent in which Petitioner was not a party).


V. OVERVIEW OF THE CITED PRIOR ART

Petitioner asserts that Claims 1-12 are obvious over U.S. Patent No. 8,634,240 ("Gavens") (Ex. 1045) in view of the knowledge of a POSA. See Petition, 27-58. In its single-reference § 103 ground, Petitioner also cites references that Gavens allegedly incorporates by reference, including U.S. Patent Appl. Pub. No. 2011/0153912 ("Gorobets 912") (Ex. 1049), U.S. Patent Appl. Pub. No. 2010/0172180 ("Paley") (Ex. 1051), U.S. Patent No. 5,936,971 ("Harari") (Ex. 1057), and PCT Pub. No. WO 03/027828 ("Gorobets PCT") (Ex. 1058).

A. Gavens (Ex. 1045)

Gavens discusses a problem that occurs with flash memory (a type of non-volatile memory)—as it ages, its error rate increases. Ex. 1045, 3:56-4:5; Ex. 2001, ¶¶53-68. To ensure data integrity in such situations, the memory uses a resource-intensive error correction code (ECC) to correct errors. Ex. 1045, 3:56-4:5. Gavens explains that this poses a problem because, “for most of the life time of the device [(before the device ages and the error rate increases)], the ECC is only marginally utilized, resulting in its large over-heads being wasted[.]” *Id.*, 3:67-4:2.

To perform ECC, Gavens describes a flash memory device including **controller 102** (highlighted purple) and **memory chip 100** (highlighted teal).



Ex. 1045, Fig. 1 (annotated), 8:13-39; Ex. 2001, ¶54. The memory chip includes input/output (I/O) circuits, **data latches 430**, sense modules, and a non-volatile flash **memory array 200**. *Id.*, 4:13-16, 8:13-39. In contrast to the Challenged Patent, there is no mention of RAM in the memory chip 100, controller chip 102, or anywhere else in the flash memory device 90. In fact, the only time that Gavens mentions RAM is in the “Background of the Invention,” when Gavens is

distinguishing its purported invention from RAM. *Id.*, 1:26-29 (“Unlike RAM...that is also solid-state memory, flash memory is non-volatile”).

Before writing a page of data to the memory array, “ECC is computed for the data page” by the ECC Processor 62. *Id.*, 12:51-63. The page of write-data then consists of two portions: the “user data” and the computed ECC. *Id.*

When a page is read from the memory array, it is “shifted out” (or “toggled”) to **controller 102**. There, first an error detection code (EDC) “indicates the existence of any errors in the read data page[,]” and if there are errors, “the data page’s existing ECC is compared to a second version of the ECC computed on the read data [(the “user data”)].” *Id.*, 12:64-13:8.⁴

Gavens teaches that using ECC in this way—where the ECC must be programmed to correct the “worst-case” number of bits—is burdensome in terms of processing time as well as chip area on the controller:

Using ECC to correct a worst-case number of error bits will consume a great amount of processing time. The more bits it has to correct, the more computational time is required. The memory performance will

⁴ This EDC/ECC relies entirely on data that is read from the non-volatile memory array. Thus, unlike the claimed invention of the Challenged Patent, this EDC/ECC does not involve a comparison of data read from the non-volatile memory with data *retained* in volatile memory (such as RAM). Ex. 2001, ¶56 n.1.

be degraded. Additional dedicated hardware may be implemented to perform the ECC in a reasonable amount of time. ***Such dedicated hardware can take up a considerable amount of space on the controller ASIC chip.*** Moreover, for most of the life time of the device, ***the ECC is only marginally utilized***, resulting in its large over-heads being wasted and realizing no real benefits.

Ex. 1045, 3:48-4:2; *see also id.*, 5:15-30, 13:9-16.


To address this problem, Gavens discloses the embodiments of Figures 14A-14B and 16A-16C, which describe the **memory array 200** of Figure 1 in greater detail. *See id.*, 8:13-15 (“Fig. 1 illustrates...a memory device in which the features of the present invention are embodied.”).

These embodiments rely on a non-volatile memory with a first SLC portion (which is less error prone) and a second MLC portion (which is more error prone). Gavens explains that optimizing where data is stored (in the SLC or MLC portion) based on the expected number of error bits “allows a smaller and more efficient...ECC...to be designed for correcting a smaller number of error bits, thereby improving the performance and reducing the cost of the memory.” *Id.*, 4:28-31.

1. First Embodiment (Figs. 14A-14B)

Gavens describes a first embodiment with respect to Figures 14A-14B. *See* 1045, 17:24-26 (referring to “the embodiment described in Fig. 14A and Fig. 14B”).

For this embodiment, as shown in Figure 14B below, Gavens describes a memory array that relies on a “**lower density**” first portion 410 (*i.e.*, SLC) and a “**higher density**” second portion 420 (*i.e.*, MLC). *Id.*, 4:13-19, 16:14-28; Ex. 2001, ¶¶60-62. The MLC portion is more error prone than the SLC portion but can be used “for efficient storage.” Ex. 1045, 4:16-19, 16:41-45. Gavens explains that data can first be written to the MLC portion, then “read back in a post-write read operation to check for excessive error bits.” *Id.*, 4:19-21; *see also id.*, Fig. 14A, 16:41-56.



Id., Fig. 14B (annotated); Ex. 2001, ¶60.


Gavens explains that the post-write error detection process can check for errors in one of two ways: “either by comparison with the original copy which may be cached or by checking the EDC portion of the ECC.” Ex. 1045, 16:46-49. While Gavens at 16:46-49 does not specify how the “original copy” is cached, Gavens later explains that “[i]n the preferred embodiment” (where the antecedent basis of “the” refers to the immediately preceding “preferred embodiment” of Figures 14A-14B) the “data is cached in [a] first section of the first portion” of non-volatile memory (*i.e.*, a section of the non-volatile SLC memory). *Id.*, 18:1-8; *see also id.*, 16:14-16. Gavens does not describe caching the “original copy” in volatile memory.

Once the error rate is checked, the device determines whether to keep the data in MLC or instead move it to SLC. If the stored data has an acceptable number of error bits, then the stored copy remains in the MLC and becomes the “valid” copy. *Id.*, 16:41-56. “If the error bits exceeded a predetermined amount, the data is” instead stored in “the less error-prone” SLC portion. *Id.*, 4:21-22; *see also id.*, Fig., 14B, 17:7-20, 18:1-8. To rewrite the data, Gavens teaches that the memory may rely on the “cached” copy of the original data or “correcting the errors bits [from the copy stored in the MLC] with the ECC.” *Id.*, 17:12-14; *see also id.*, 18:1-8.

2. Second Embodiment (Figs. 16A-16C)

Figures 16A-16C depict a second embodiment of Gavens. *See* Ex. 1045, 7:26-28 (referring to “the embodiment described in Fig. 16A to Fig. 16C”); Ex. 2001, ¶¶63-68. For this embodiment, Gavens teaches splitting the first portion (*i.e.*, the SLC portion) into a first section “for caching” and a second section “for rewrit[ing]” data. *Id.*, Fig. 16A, 18:9-23. Gavens emphasizes that this approach limits reliance on the controller 102. *See* Ex. 1045, Fig. 1, 3:60-4:2, 18:36-39.

As shown in Figure 16A below, Gavens explains that the lower density First Portion (*i.e.*, SLC) includes two sections—a “first section” 411 for caching write data, and a “second section” 412 for rewriting data from the higher density Second Portion (*i.e.*, MLC). *Id.*, 18:9-23.



Id., Fig. 16A (annotated); Ex. 2001, ¶64.

Incoming data is written both to the MLC and as a “cached copy” to the first section of the SLC. Ex. 1045, 18:21-29. Detecting the number of error bits is then “accomplished by comparison [of the data written to MLC] with the cached copy.” *Id.*, 18:40-44.

Gavens specifies that this comparison may be performed at the **data latches 430** (highlighted orange in Figure 1 above) on the **flash memory chip 100**. *Id.*, 16:46-49, 18:30-39. Gavens strongly recommends using **data latches 430** because

it eliminates the need to “toggle” the data out to the **controller chip 102**, as would be required if using ECC Processor 62 for EDC-based error detection. *Id.*, 18:36-39; *see also id.*, 3:60-4:2. Gavens explains that “much time can be saved” by eliminating the need to toggle data to the controller chip, thus increasing the performance of the device. *Id.*

If there is an acceptable number of error bits, the MLC portion becomes the “valid” copy. *Id.*, 18:45-52, Fig. 16B. If there is an unacceptable number of error bits, the “cached copy” in the “first section” of the SLC is rewritten to the “second section” of the SLC. *Id.*, 18:53-67, Fig. 16C.

By caching the data in the non-volatile memory array and using the latches to compare the data, Gavens lowers overhead and reduces reliance on the controller chip. *Id.*, 18:40-67.


Accordingly, unlike PO’s claimed invention, Gavens does not disclose caching or retaining data outside of the non-volatile memory array, or using data retained outside of the non-volatile memory array for detecting errors.

B. Secondary References Related To Random Access Volatile Memory

In its ground, Petitioner cites several references that Gavens allegedly incorporates by reference—Gorobets 912, Paley, Harari, and Gorobets PCT. Petitioner cites these references’ disclosure related to the use of random access volatile memories.

1. Gorobets 912 (Ex. 1049)

As shown in Figure 11 below, Gorobets 912 describes a flash memory device with a non-volatile **SLC portion** and a non-volatile **MLC portion**. Ex. 1049, ¶¶104, 109; *see also* Ex. 2001, ¶¶69-71. The **SLC portion** acts as “a cache for data to be written to the **[MLC] memory**.” *Id.*; *see also id.*, ¶108 (explaining the utility of the non-volatile SLC cache).



Ex. 1049, Fig. 11 (annotated); Ex. 2001, ¶70.

The flash memory device also relies on a controller, including a generic volatile **RAM-based cache**. Gorobets 912 explains that this **cache** is “high-speed storage for *temporarily* storing data being passed between a high-speed and slower-speed component of the system.” Ex. 1049, ¶109 (emphasis added); *see also id.*, ¶130.

Gorobets 912 does not describe a volatile cache that stores data after it is written to the non-volatile memory—in contrast to the Challenged Patent, data is not *retained* in volatile memory. *See id.*, ¶¶24, 109, 130, 144.

2. Paley (Ex. 1051)

Paley is closely related to Gorobets 912⁵ and, as shown in Figure 15 below, also describes a non-volatile memory with **SLC** and **MLC** portions, where the **SLC portion** is used as a cache to aggregate data before it is written to **MLC**. Ex. 1049, ¶192 (“A second portion forms a cache for data to be written to the main memory.”), ¶193 (“In the preferred embodiment, individual cells in the cache portion are each

⁵ Paley and Gorobets 912 share multiple named inventors. *Compare* Ex. 1051, Cover (listing inventors, including “Sergey...Gorobets”) *with* Ex. 1049, Cover (listing inventors, including “Sergey...Gorobets”). Additionally, much of these references’ disclosure is substantially the same. *Compare* Ex. 1051, Figure 15 *with* Ex. 1049, Figure 11 (other than differences in numbering, these figures are the same).

configured to store one bit of data while the cells in the main memory portion each stores more than one bit of data.”); *see* Ex. 2001, ¶¶72-75.




FIG. 15

Ex. 1051, Fig. 15 (annotated); *see also* Ex.1049, Fig. 11 (Gorobets 912’s Figure 11 is substantially the same as Paley’s Figure 15); Ex. 2001, ¶73.

Paley further describes where the controller includes a generic volatile **RAM-based cache**. Paley—just like Gorobets 912—explains that this **cache** is “high-speed storage for *temporarily* storing data being passed between a high-speed and slower-speed component of the system.” Ex. 1051, ¶197 (emphasis added); Ex. 1049, ¶109 (same). Thus, unlike the Challenged Claims, both Gorobets 912 and Paley describe a memory system including a **RAM-based cache** that only *temporarily* stores data being passed between components.

3. Harari (Ex. 1057)

Petitioner also cites Harari, which Gorobets 912 and Paley allegedly incorporate by reference. Petition, 28 n.13 (“Harari...was incorporated by reference by Paley ‘180 at ¶ [0027]”), 35. This assertion is false. Gorobets 912 and Paley merely state: “Using RAM in a write cache operating with a flash memory has been disclosed in U.S. Patent No. 5,936,971 to Harari et al.” Ex. 1051, ¶27; Ex. 1049, ¶24 (same). This is the only paragraph in which these references mention Harari. Neither reference clearly states that Harari is “incorporated by reference.” “To incorporate matter by reference, a host document must contain language ‘clearly identifying the subject matter which is incorporated and where it is to be found’; a ‘mere *reference* to another application, or patent, or publication is not an *incorporation* of anything therein....’” *Callaway Golf Co. v. Acushnet Co.*, 576 F.3d 1331, 1346 (Fed. Cir. 2009) (quoting *In re De Seversky*, 474 F.2d 671, 674 (CCPA 1973)) (emphasis in original).

That aside, Harari describes a RAM “cache buffer” “for *temporarily* holding write data files.” Ex. 1057, Fig. 8 (including “cache buffer 705”), 13:29-30 (emphasis added). It acts as a “buffer” because it holds data before it is written to a non-volatile memory. *See id.*, 14:1-38. When “the cache memory 705...start[s] to fill up[,]” it preferentially writes (“archive[s]”) infrequently written files to the flash non-volatile memory. *Id.*, 14:26-30. “In this way, only infrequently used data files

are written into the Flash memory while frequently used data files are trapped in the cache memory.” *Id.*, 14:22-25. Harari’s “cache buffer” therefore “reduce[s] the number of writes to the Flash...memory[.]” *Id.*, Abstract; *see also* Ex. 2001, ¶76.

4. Gorobets PCT (Ex. 1058)

Petitioner also alleges that Gorobets 912 and Paley incorporate by reference Gorobets PCT. Petitioner, however, provides no basis for this claim. *See* Petition, 35 (referencing “note 14” as supporting its contention that Gorobets PCT is incorporated by reference, but Petitioner’s footnote 14 never mentions Gorobets PCT). Nor can they. Gorobets 912 and Paley merely state—like with Harari—that “WO 03/027828 [(Gorobets PCT)] disclose[s] a memory system....” Ex. 1049, ¶19; Ex. 1051, ¶23 (same). These references do not state, with the required specificity, that Gorobets PCT is incorporated by reference. *See Callaway Golf*, 576 F.3d at 1331.

Notwithstanding, Gorobets PCT—just like Gorobets 912, Paley, and Harari—describes a volatile RAM-based buffer used to *temporarily* store data. Ex. 1058, 11:10-13 (describing “SRAM block 30 [that] stores all *temporary* information within the controller”) (emphasis added). Data “being transferred” from system components to a flash memory “is *buffered* in the SRAM 30.” *Id.*, 11:21-23 (emphasis added); *see also* Ex. 2001, ¶77.

Thus, unlike the Challenged Claims, Gorobets 912, Paley, Harari, and Gorobets PCT disclose RAM-based caches that *temporarily* store (buffer) data before it is written to flash. Ex. 2001, ¶77.

VI. REASONS FOR DENYING INSTITUTION

The Petition fails to establish a reasonable likelihood that it will prevail in its single ground. *See infra* Section VI.A.

A. Ground 1: The Petition Does Not Establish That Claims 1 and 12 Would Have Been Obvious Over Gavens

As discussed below, the Petition’s analysis for Ground 1 (obviousness over Gavens) is fundamentally flawed with respect to multiple limitations of Claims 1 and 12 of the Challenged Patent, and those flaws propagate to each challenged dependent claim. Each of the below flaws warrants denial of institution.

1. [1d] and [12d]: “at least one SLC nonvolatile memory element that can be mapped into the SLC memory space”

Limitation [1d] requires, in part, “at least one *SLC nonvolatile memory element*.” Petitioner’s obviousness analysis is deficient and fails to show why the modification to Gavens proposed by Petitioner would have been obvious to a POSA. *See* Petition, 33-34; Ex. 2001, ¶¶79-87.

For this limitation, Petitioner applies its own construction for “SLC non-volatile memory”: “non-volatile memory cells arranged with circuitry *incapable* of storing multiple logical pages in a single physical page of cells.” Petition, 25

(emphasis added). Applying this construction, Petitioner suggests that limitation [1d] requires a “*native* SLC flash chip.” Petition, 33 (“Gavens does not disclose the use of a ‘native’ SLC flash memory module or chip...because the Gavens memory module or chip 100 is capable, with its circuitry, of programming more than one logical page into a physical page[.]”). Petitioner contends that Gavens’s use of “pseudo-SLC” flash—that is, flash that is capable of operating in an MLC mode *or* SLC mode—is insufficient under its construction. *Id.* While PO disagrees with this construction, even if it is applied—where Petitioner presents its ground solely under its own construction for “SLC non-volatile memory”—Petitioner’s ground fails.


First, Petitioner fails to adequately explain why it would have been obvious to substitute one of Gavens’s SLC memory *portions* (operating in pseudo-SLC) with a “‘native’ SLC flash chip.” Petition, 33-34. Instead, Petitioner offers only a conclusory assertion, citing paragraph 92 of its expert report. *See* Petition, 33 (asserting, without explanation, that “[a]s explained in [Sechen ¶ 92], it would be obvious to a POSITA to substitute a ‘native’ SLC flash memory module for the Gavens first portion, with a similar, if not wider margin of error than binary mode operation of MLC (see [Sechen ¶ 49]) and achieve similar, if not better results with the same NAND flash controller...with assurance of success.”) (emphasis added). But “arguments may not be incorporated from an expert’s declaration wholesale into the Petition,” and Petitioner provides no rationale for its obviousness theory in its

ground for limitation [1d] itself. *Acon Inc. v. Amo Dev., LLC*, IPR2021-00858, 2022 WL 17585777, at *16 (P.T.A.B. Dec. 12, 2022); *see also* 37 C.F.R. § 42.6(a)(3).

For example, Petitioner fails to explain why a POSA would have been motivated to make the suggested modification. Simply alleging that this “substitut[ion]” would yield “similar” results does not answer why a POSA would have been motivated to pursue it. Petition, 33; *see In re Stepan Co.*, 868 F.3d 1342, 1346 n.1 (Fed. Cir. 2017) (“there must be a motivation to make the combination”). “[O]bviousness grounds cannot be sustained by mere conclusory statements; instead, there must be some articulated reasoning with some rational underpinning to support the legal conclusion of obviousness.” *In re Kahn*, 441 F.3d 977, 988 (Fed. Cir. 2006); *see also* 37 C.F.R. § 42.22(a)(2) (requiring an “explanation of the significance of the evidence”). Thus, Petitioner fails to demonstrate why it would have been obvious to modify Gavens to use a “native SLC flash chip.”

Second, contrary to Petitioner’s conclusory assertion, and as illustrated in the annotated version of Gavens’s Figure 1 below, a POSA would not have been motivated to substitute one of Gavens’s SLC portions—which, along with the MLC portion, makes up the memory array of one of Gavens’s memory chips—with an entirely separate “**native SLC flash chip**.” *See* Ex. 2001, ¶83; *see also* Section V.A (describing Gavens, including its “memory chip” comprising the “memory array” of Figure 1, where each memory array includes a “first portion” (SLC) and “second

portion” (MLC) as shown in Figures 14 and 16). As shown in the modified version of Gavens’s Figure 1 below—altered to include the separate “SLC Array”—Petitioner’s proposed modification would result in **one memory chip consisting entirely of MLC**, and **another memory chip consisting entirely of SLC**.



Ex. 1045, Figure 1 (altered to include a second memory chip comprising the “SLC array”).

Gavens, however, explains that each of its memory chips are already capable of operating with one portion as MLC, and another portion as SLC (albeit, not “native” SLC). See Ex. 1045, Abstract, 4:13-31, 6:17-25, 15:62-16:12, 16:17-25, Fig. 14, Fig. 16. That is the essence of Gavens’s invention—in a single memory

chip, taking advantage of the efficiencies of MLC storage, by relying on a “less error-prone” SLC **portion** when an MLC portion has too any errors. *See id.*, 4:13-31, 4:18-22 (“Data is written to the second portion for efficient storage. Afterwards, the data is readback in a post-write read operation to check for excessive error bits. If the error bits exceeded a predetermined amount, the data is rewritten or kept at the less error-prone first portion.”), 15:66-16:13, 16:41-45.

Additionally, Gavens explains that the SLC portion is also used to cache data, as described in its first and second embodiments (*see supra* Section V.A), to avoid needing to “toggle[]” data out to **controller 102** when checking data written to the MLC portion for errors. *See* Ex. 1045, 18:21-39, 18:36-39. This is because “**data latches**” **within the memory chip** (which includes *both the MLC and SLC portions*) enable data comparison without relying on **external controller 102**. *See id.*, 18:30-39, Fig. 1 (showing “Data Latches” 430 as part of the Memory Chip, separate from the Controller 102). Gavens explains: “By not making the comparison at the controller, the data does not have to be toggled out to the controller, [and] much time can be saved. Fig. 1 shows the data latches 430, **which is on-chip**, for the data comparison to take place.” Ex. 1045, 18:36-39 (emphasis added).

Replacing the SLC portion of the *same* memory chip with a *separate* SLC chip connected to the controller 102 would negate this benefit, as data cached in a **separate SLC chip** would need to be toggled out to the **controller** for comparison

with data stored in the **MLC chip**. Ex. 2001, ¶86. Therefore, Gavens teaches away from Petitioner’s proposed modification, and a POSA would not have been motivated to substitute an SLC portion of a memory chip with a separate “‘native’ SLC flash chip,” as this modification would undermine Gavens’s intended advantage—efficiently utilizing higher-density MLC memory. *See* Ex. 2001, ¶86. Petitioner’s proposed modification “would require a substantial reconstruction and redesign of the elements shown in [Gavens] as well as a change in the basic principles under which [Gavens] was designed to operate” and cannot be correct. *Samsung Elecs. Co., Ltd. v. Mojo Mobility Inc.*, IPR2023-01099, 2024 WL 457731, at *10 (P.T.A.B. Feb. 6, 2024) (denying institution) (quoting *Application of Ratti*, 270 F.2d 810, 813 (C.C.P.A. 1959)).

Furthermore, Limitations [12d] and [1d] are similar, and Petitioner relies on the same ground for both limitations. *See* Petition, 56. Accordingly, the Petition’s analysis of limitation [12d] fails for the same reasons as those given above for limitation [1d].

2. [1e] and [12e]: “at least one random access volatile memory”

For limitation [1e], Petitioner fails to show how Gavens and its allegedly incorporated references—Gorobets 912, Paley, Harari, and Gorobets PCT⁶—disclose wherein the system for storing data includes “at least one random access volatile memory element.” Ex. 2001, ¶¶88-93; Petition, 34-36. Petitioner fails to clearly identify which aspect of the prior art is being mapped to this limitation. Instead, Petitioner has taken a scattershot approach, ambiguously referencing unrelated disclosures from Gorobets 912, Paley, Harari, *and* Gorobets PCT. Petitioner fails to show how the allegedly incorporated references’ disclosure of RAM-based buffers applies to or can be combined with Gavens.

First, Petitioner does not allege that Gavens itself discloses the “at least one random access volatile memory element.” Petition, 34-36. Nor can they. Gavens does not disclose or suggest that its system includes RAM. Gavens explicitly distinguishes its teachings from systems that use RAM, explaining how “nonvolatile storage” is “*unlike* RAM.” Ex. 1045, 1:26-29 (“Unlike RAM (random access memory) that is also solid-state memory, flash memory is non-volatile, and retaining its stored data even after power is turned off.”). This is consistent with Gavens’s

⁶ As explained above in Sections V.B.3-V.B.4, Petitioner is incorrect in its assertion that Gorobets 912 and Paley incorporate by reference Harari and Gorobets PCT.

description of its system in Figure 1, as shown below, which includes only the ***non-volatile*** “Memory array 200.” *See* Ex. 1045, 1:17-20 (“This application relates to the operation of re-programmable ***non-volatile memory systems*** such as semiconductor flash memory, and, more specifically, to handling and efficient managing of errors in memory operation.”) (emphasis added), 4:13-16 (“According to a general aspect of the invention,...a ***flash*** memory having an array of memory cells is configured with a first portion and a second portion.”) (emphasis added). Therefore, “Gaven’s invention is directed specifically to non-volatile memory” and “a [POSA] would [not] have found Gavens to teach or suggest [a memory system] that includes RAM.” Ex. 2012, 42 (decision denying institution of PGR2024-00048); *see also* Ex. 2012, 76 (decision denying institution of PGR2024-00047, where the Board determined that “Gavens does not disclose ‘random access volatile memory’ according to claim 1, e.g. in a ‘system’ with, among other things, a ‘controller’ and a ‘bank of nonvolatile memory’”).




FIG. 1

Ex. 1045, Fig. 1.

Second, apparently conceding that Gavens itself does not disclose the “random access volatile memory element” required by limitation [1e], Petitioner points to disclosure from four different references—Gorobets 912, Paley, Harari, and Gorobets PCT—that Gavens allegedly incorporates by reference. Petition, 34-36. Petitioner vaguely contends that these references distinct teachings support its claim that it was “well-known” “to use...RAM...in a NAND flash controller as a staging or cache location for data to be written to flash memory.” *Id.*, 34. Petitioner ignores, however, that the cited embodiments from these references—each allegedly

disclosing a RAM-based cache—are *distinct from the embodiments of Gavens*. Ex. 2001, ¶91. For one, as explained above, Gavens never discloses the use of random access volatile memory. Petitioner provides no explanation how the distinct disclosure of the allegedly incorporated references applies to or can be combined with the previously cited disclosure of Gavens. Petitioner cannot pick and choose between distinct elements from different disclosures. *See Captioncall, LLC v. Ultratec, Inc.*, No. IPR2020-01215, 2021 WL 278381, at *7 (P.T.A.B. Jan. 27, 2021) (denying institution when petitioner relied on “distinct” embodiments from a reference and another reference that was incorporated by reference, explaining that “[a]nticipation requires all elements arranged as in the claim without the need to pick and choose elements from different disclosures”).⁷

⁷ Phison (the party who filed the petition in IPR2025-00215—the petition which Petitioner Kingston now seeks to join) made the same error in PGR2024-00048, which challenged U.S. Patent No. 11,854,612—a continuation of the 300 Patent challenged in this Petition. There, Phison cited Gorobets 912’s teaching of RAM for a similar limitation, and the Board found that, “As Patent Owner correctly notes, ‘Petitioner...does not state whether Gorobets’s ‘RAM’ is being used to modify Gavens’s system, and why a POSA would have made any modification.’” Ex. 2012, 43.

Finally, to the extent Petitioner is attempting to rely on an obviousness analysis for limitation [1e], that also fails.⁸ Petitioner provides no coherent obviousness analysis. Petitioner does not explain what aspect of the claim limitation is missing in Gavens and, correspondingly, does not explain what from Gavens must be modified to arrive at the claim language. And Petitioner also does not state whether any of the allegedly incorporated references' RAM is being used to modify Gavens's system, and why a POSA would have been motivated to make any modification. *See* Petition, 34-36. Petitioner fails to perform a proper obviousness analysis, which requires, *inter alia*, ascertaining the “differences between the prior art and the claims at issue.” *Graham v. John Deere Co. of Kansas City*, 383 U.S. 1, 17-18 (1966). The Petition's obviousness analysis is insufficient and must be rejected.

Thus, Petitioner has not demonstrated with particularity how limitation [1e] maps to “at least one random access volatile memory element” disclosed in the prior art, as required by 35 U.S.C. § 312(a)(3). The Petition lacks an “explanation of the

⁸ Petitioner vaguely references what a “POSITA would...understand” but in no way suggests that it is relying on obviousness for limitation [1e]. Petitioner concludes its ground with: “Thus, Gavens and its incorporated references ***disclose*** claim element [1e].” Petition, 36 (emphasis added).

significance of the evidence.” 37 C.F.R. § 42.22(a)(2). Petitioner’s scattershot approach—citing disparate teachings from the allegedly incorporated references without clearly identifying which *one* they rely on—must be rejected. *See Masimo Corp. v. Apple Inc.*, IPR2023-00745, Paper 44 (final written decision), 36 (P.T.A.B. Oct. 15, 2024) (determining that the challenged claim is not unpatentable, explaining: “In essence, Petitioner invites the Board to piece together a single theory that meets all claim elements from these disparate theories that each require distinct structures....”) (citing *Amazon.com, Inc., et al. v. Avago Technologies General IP (Singapore) PTE. LTD.*, IPR2017-01112, Paper 11, at 13 (P.T.A.B. Aug. 23, 2017) (“Petitioner is inviting us to assume the role of archeologist of the record, which invitation we decline.”)); *Nearmap US, Inc. v. Eagle View Technologies, Inc.*, IPR2024-00716, Paper 9, 18 (P.T.A.B. Oct. 9, 2024) (denying institution of *inter partes* review, explaining that “[t]he question before us is not whether we could come up with a way to understand how the asserted prior art could render the challenge claims unpatentable, but rather did Petitioner provide the arguments and evidence to show that the challenged claims are unpatentable.”); *Netflix, Inc. v. DivX, LLC*, 84 F.4th 1371, 1379 (Fed. Cir. 2023) (“While the Board should review a petition holistically, it is not obligated to cobble together assertions from different sections of a petition or citations of various exhibits in order to infer every possible permutation of a petitioner’s argument.” (cleaned up)).

Furthermore, Limitations [12e] and [1e] are the same, and Petitioner relies on the same ground for both limitations. *See* Petition, 57. Accordingly, the Petition’s analysis of limitation [12e] fails for the same reasons as those given above for limitation [1e].

3. [1g] and [12f]: “the controller controlling access of the MLC and SLC nonvolatile memory elements and the random access volatile memory for storage of data therein, the controller, in at least a Write access operation to the MLC nonvolatile memory element, operable to store data in the MLC nonvolatile memory element and retain such stored data in the random access volatile memory”

Petitioner’s obviousness analysis is also deficient for limitation [1g], at least because Petitioner fails to demonstrate how Gavens and its allegedly incorporated references—Gorobets 912, Paley, Harari, and Gorobets PCT—disclose or suggest random access volatile memory” that “retain[s] such stored data.” Petition, 40-41; *see also* Ex. 2001, ¶¶94-103.

First, like its deficient analysis for limitation [1e], Petitioner again fails to adequately explain how Gavens and its allegedly incorporated references disclose or suggest a RAM element. As explained above with respect to limitation [1e], Gavens itself does not disclose a RAM element, and Petitioner does not explain why a POSA would have been motivated to combine the teachings of the allegedly incorporated references with Gavens’s distinct disclosure. *See supra* Section VI.A.2; *see also*

supra V.B.1-V.B.4 (describing the allegedly incorporated references’ disclosure of RAM-based caches for temporarily storing data).

For limitation [1g], Petitioner merely cites back to its analysis of limitation [1e], and conclusorily asserts that “Gavens discloses a controller” whose “operation obviously involves control of the RAM data cache disclosed by Gavens, *e.g.*, incorporated Paley ‘180 ¶ [0197][.]” Petition, 40. But Petitioner again fails to provide any meaningful obviousness analysis, including any explanation of how or why the distinct teachings of Paley—or any other allegedly incorporated reference—would have been combined with Gavens. *See Graham*, 383 U.S. at 17-18.

Second, limitation [1g] further requires that “in at least a Write access operation to the MLC nonvolatile memory element, [the controller is] operable to ***store data in the MLC nonvolatile memory element and retain such stored data in the random access volatile memory.***” In other words, limitation [1g] requires that, during a “write access operation,” the controller is operable to store data in MLC memory while ***retaining*** that data in RAM.

Petitioner ambiguously asserts that a POSA would understand this operation to be “satisfied by Gavens by ‘storing’ data in the NVM that was and is cached (temporarily stored or ‘retained’) in the disclosed RAVM.” Petition, 40-41. However, Petitioner fails to explain how a POSA would have reached this conclusion, including on what basis a POSA would have understood Gavens to

suggest that data cached in RAM is “**retained**” within the meaning of the claim. *See In re Kahn*, 441 F.3d at 988. Petitioner appears to assert that RAM that caches or “temporarily store[s]” data in any way necessarily “retains” that data as required by the claim. Petition, 40-41, 41. But that is incorrect. Limitation [1g] requires that the “stored data” be “retain[ed]” in the RAM in the sense that it is also operable to be stored in the MLC memory. Ex. 2001, ¶98.

Petitioner also refers to its analysis for the subsequent limitation [1h] (Petition, 41), but that ground fares no better. Petitioner’s ground for limitation [1h] similarly does not provide a sufficient obviousness rationale to cure this deficiency. *See* Petition, 41-42 (focusing on the claimed “data integrity test” and Gavens’s teaching of data “comparison”); Ex. 2001, ¶99.

Notwithstanding the fact that Petitioner failed to provide the requisite obviousness analysis, a POSA would not have found it obvious that Gavens in combination with any of the allegedly incorporated references, such as Paley, renders obvious a controller operable to **retain** data in RAM when it is written to MLC. *See* Ex. 2001, ¶¶100-103. Instead, the teachings of the allegedly incorporated references cited by Petitioner store data **temporarily**, where that data is buffered **until** it is written to MLC. For example, Petitioner cites ¶197 of Paley but, in this paragraph, Paley explicitly states: “A cache is high-speed storage for **temporarily storing data** being **passed between** a high-speed and a slower-speed

component of the system.” Ex. 1051, ¶197; Petition, 40 (citing ¶197 of Paley). Paley further suggests that a “high-speed volatile RAM” in a controller can be used as a cache to temporarily store data being passed to non-volatile memory, including SLC or MLC. *See* Ex. 1051, ¶197; Ex. 2001, ¶100.

Gorobets 912 includes the same disclosure. *See* Ex. 1049, ¶109 (“A Cache is high-speed storage for temporarily storing data being passed between a high-speed and a slower-speed component of the system.”). Gorobets 912 goes on to explain that this RAM, by design, is of “relatively limited capacity,” so once data is “transferred out” to the non-volatile memory, the contents of the RAM are cleared to “receive more data” being transferred—in other words, this RAM is not operable to retain data that is written to non-volatile memory. *See* Ex. 1049, ¶144 (“As data has been transferred out of RAM at 721, this has opened up the RAM, which is relatively limited capacity, to receive more data from the host....”); Ex. 2001, ¶100.

Additionally, the Board already considered Gorobets 912 in denying institution of a related petition filed by Phison—PGR2024-00047, challenging U.S. Patent No. 11,830,546 which is a continuation of the 300 Patent challenged here—and determined, congruent with PO’s current arguments, that “Gorobets [912] ...relate[s] to using RAM to *temporarily* store data *before transferring the data* to non-volatile memory[.]” Ex. 2011, 78 (emphasis added).

Thus, Paley and Gorobets 912 make clear that—contrary to Petitioner’s assertion—the RAM-based cache only “temporarily” stores data when it is being “passed” to the non-volatile memory. Harari and Gorobets PCT—while not explicitly cited in Petitioner’s grounds for limitations [1g] or [1h]—suffer from the same flaw, and similarly disclose RAM-based caches/buffers for *temporarily* storing data. *See* Ex. 1057, Fig. 8, 13:29-30; Ex. 1058, 11:10-13; *see also supra* Sections V.B.3-V.B.4; Ex. 2001, ¶101.

Accordingly, even assuming *arguendo* that a POSA would have combined a RAM-based cache for temporarily storing data—purportedly taught by each of the allegedly incorporated references—into Gavens’s cited system, Petitioner fails to demonstrate why a POSA would have understood the controller to be operable to *retain* data in that RAM when it is stored in MLC, as is required by limitation [1g]. Ex. 2001, ¶¶100-103.

Thus, Petitioner fails to show that Gavens discloses or suggests the features of limitation [1g].


Furthermore, Limitations [12f] and [1g] are the same, and Petitioner relies on the same ground for both limitations. *See* Petition, 57. Accordingly, the Petition’s analysis of limitation [12f] fails for the same reasons as those given above for limitation [1g].

4. [1h] and [12g]: “the controller performing a data integrity test on stored data in the MLC nonvolatile memory element after at least a Write access operation performed thereon by comparing the stored data to the retained data in the random access volatile memory”

The issues with Petitioner’s obviousness analysis are further compounded into its ground for limitation [1h]. Here, Petitioner asserts that Gavens’s disclosure of a “‘post write read’...check for errors ‘by comparison with the original copy *which may be cached*’” would have been understood to refer to a RAM-based cache, as required by limitation [1h]. Petition, 41-42 (quoting Ex. 1045, 16:46-49) (emphasis altered from that applied by Petitioner). A POSA would not have shared Petitioner’s understanding—Gavens’s “cache[]” that retains the “original copy” (Ex. 1045, 16:46-49) refers to a *non-volatile cache*, not a RAM-based cache, and Petitioner fails to show that a POSA would have interpreted Gavens differently based on the allegedly incorporated references. *See* Ex. 2001, ¶¶104-114.

First, Petitioner cannot show that Gavens itself discloses a RAM-based cache used in the post-write error check. *See* Ex. 1045, 16:46-49 (describing a post-write error detection process that can check for errors “by comparison with the original copy which may be cached”). Gavens instead explains that the “original copy” is cached in SLC non-volatile memory. *See supra* Section V.A; *see also* Ex. 2001, ¶106.

Specifically, for its ground, Petitioner cites teachings from the “first embodiment” of Gavens, including Figure 14; an overview of this embodiment is provided in Section V.A.1 above. *See* Petition, 41 (citing Gavens’s Figures 14 and 15, and its disclosure at 16:13-17:42). With Figure 14B, as illustrated below, Gavens describes a memory including an **SLC portion** which is less error prone, and **MLC portion** which is more error prone but may be used “for [more] efficient storage.” Ex. 1045, 4:16-19, 16:41-45; *see also id.*, 4:13-19, 16:14-28. Gavens explains that data written to the MLC portion can be “read back in a post-write read operation to check for excessive error bits.” *Id.*, 4:19-21. Gavens at 16:46-49, as cited by Petitioner, refers to this post-write error check of data written to the MLC portion.



Ex. 1045, Fig. 14B (annotated).


Gavens at 16:46-49 specifies that the post-write error check may be performed “by [a] comparison with the original copy which may be cached[.]” Ex. 1045, 16:46-49. At 16:46-49, Gavens does not specify how the “original copy” is cached, but Gavens later explains that “[i]n the preferred embodiment” (where the antecedent basis of “the” refers to the immediately preceding “preferred embodiment” of Figures 14A-14B) the “data is cached in [a] first section of the first portion” of non-

volatile memory (*i.e.*, a section of the non-volatile SLC memory). *Id.*, 18:1-8; *see also id.*, 16:14-16. Thus, Gavens describes where the “original copy” in non-volatile memory—not volatile memory, such as RAM, as Petitioner contends. Ex. 2001, ¶108.

Additionally, Gavens further explains with the embodiment of Figure 16 why caching the original data in non-volatile memory improves the system. Like in the embodiment of Figure 14 discussed above, in the embodiment of Figure 16, Gavens describes caching data in ***non-volatile*** memory (“first section 411”) where it is ***retained*** for use in error detection. *See* Ex. 1045, 18:21-44; *see also supra* Section V.A.2 (describing Gavens’s second embodiment of Figure 16). Gavens’s non-volatile cache retains data even after the data is written to a “high[er] density second portion” of the non-volatile memory. *See* Ex. 1045, 18:24-29. The retained data in the non-volatile cache can then be compared to the data written to the “higher density second memory” to check if the data was written correctly. *See id.*, 18:30-44; Ex. 2001, ¶109.

Gavens explains that data used for error detection should be cached in the non-volatile memory (*i.e.*, the “First section 411” of Figure 16) to avoid needing to toggle data to the controller for error detection. Ex. 1045, 18:36-39 (“By not making the comparison at the controller, the data does not have to be toggled out to the controller, much time can be saved.”). Specifically, as illustrated in Gavens’s Figure

1 below, when data is cached in the SLC portion of the **non-volatile memory array 200** at the time it is written to the MLC portion of that same **array**, the **on-chip data latches 430** (part of **memory chip 100**) can be used for the post-write error check—thus avoiding the need to toggle the written data out to the **controller 102**. Gavens explains that caching data in the non-volatile memory improves performance by reducing reliance on the controller chip. *See id.*, 18:36-67; Ex. 2001, ¶109.



Ex. 1045, Fig. 1 (annotated).

Thus, a POSA would not have understood Gavens itself to teach caching data in RAM for the “comparison” cited by Petitioner. Gavens instead teaches that, for the cited “comparison,” data is cached in the non-volatile SLC memory. Ex. 2001, ¶110.

Second, the teachings of the allegedly incorporated references, including Gorobets 912 and Paley (as cited by Petitioner), do not justify Petitioner’s incorrect interpretation of Gavens. Petitioner’s ground for limitation [1h] relies only on what “Gavens and its incorporated references *disclose*”—Petitioner does not purport to rely on obviousness for this limitation. Petition, 42 (emphasis added). Petitioner, however, again cites the distinct teachings of Gorobets 912 and Paley to argue how they believe a POSA would have “understood” Gavens. But, as was also addressed in countering Petitioner’s ground for limitation [1e] above, Petitioner again provides no explanation how the unrelated disclosure from these references applies to or can be combined with the previously cited disclosure of Gavens. Petitioner cannot cite the RAM-based caches taught by Gorobets 912 and Paley as if they are part of Petitioner’s cited embodiment of Gavens. *See Captioncall*, 2021 WL 278381, at *7.

Notwithstanding, to the extent Petitioner was attempting to invoke obviousness—which it did not—Petitioner again fails to provide the requisite obviousness analysis to explain, for example, why Gavens would have been modified or re-interpreted based on the teachings of Gorobets 912 and Paley to rely

on a RAM-based cache for the disclosed data comparison. A POSA would have instead understood that Gavens’s non-volatile cache should not be replaced with the RAM-based caches taught by Gorobets 912 and Paley. Ex. 2001, ¶112.

Specifically, as explained above, the non-volatile, SLC cache of Gavens cited by Petitioner is used to **retain data** so it can be used in the subsequent post-write error-check. *See, e.g.*, Ex. 1045, 16:46-49. In contrast, as explained above in Section VI.A.3 (addressing Petitioner’s ground for limitation [1g]), the RAM-based caches of Gorobets 912 and Paley cited by Petitioner do not retain data, and instead only “temporarily” store when it is being transferred to non-volatile memory. *See* Ex. 1049, ¶¶109, 144; Ex. 1051, ¶197.

Thus, unlike the non-volatile cache of Gavens, Gorobets 912’s and Paley’s RAM-based based caches **do not retain data** and, instead, act only as write buffers. *See, e.g.*, Ex. 1049, ¶144 (explaining that the RAM-based cache is of “relatively limited capacity,” so as soon as data is “transferred out of [the] RAM” to the non-volatile memory, the RAM “open[s] up” the space). A POSA would not have been motivated to replace Gavens’s non-volatile cache with the RAM-based caches taught by Gorobets 912 or Paley, as these caches perform entirely different functions—**retaining** data (as in Gavens, and as required by the claims) versus **temporarily** storing data (as in Gorobets 912 and Paley). *See* Ex. 2001, ¶112; *see also* Ex. 2011, 78-79 (in its decision denying institution of PGR2024-00047, the Board determined

that Petitioner’s cited portions of Gorobets 912 “say ‘nothing about retaining data’ in RAM ‘for later use in a data integrity test performed by the controller’” and that “Paley’s disclosure...does not cure the deficiencies in Gavens and Gorobets [912]”).

Moreover, Gorobets 912 and Paley describe RAM-based caches potentially located in the controller. *See* Ex. 1049, ¶109 (“Typically high-speed volatile RAM are employed as cache as in a host cache 82 and/or in controller cache 102 of the memory controller”); Ex. 1051, ¶197 (same); *see also* Ex. 1049, ¶130. For example, Figure 11 of Gorobets 912, and the substantively identical Figure 15 of Paley, show the “Cache (RAM) 102” in the memory controller. In contrast, Gavens teaches *against* using the controller for its error detection process. As described above, Gavens explains that data used for error detection should be cached in the non-volatile memory *to avoid having to toggle data to the controller*. Ex. 1045, 18:36-39 (“By not making the comparison at the controller, the data does not have to be toggled out to the controller, much time can be saved.”); *see also supra* Section V.A. Gavens explains that caching data in the non-volatile memory increases performance by reducing reliance on the controller chip. *See id.*, 18:36-67. Relying on a cache in the controller (as may be required if the RAM-based cache of Gorobets 912 or Paley was used) rather than a cache in the non-volatile memory would contradict the purpose of Gavens’s invention. Ex. 2001, ¶113.

Petitioner provides no rationale to overcome this clear teaching away. Thus, to the extent Petitioner presented an obviousness ground—which it did not—replacing Gavens’s non-volatile cache with a RAM-based cache taught by Gorobets 912 or Paley “would require a substantial reconstruction and redesign of the elements shown in [Gavens] as well as a change in the basic principles under which [Gavens] was designed to operate” and cannot be correct. *Application of Ratti*, 270 F.2d at 813; *see also* Ex. 2011, 71, 76 (decision denying institution of PGR2024-00047, finding that “Petitioner does not explain how the teachings of Gorobets [912] and Paley ‘combine with the cited teachings of Gavens’ to yield” a similar claim limitation requiring “comparing the stored data in the controller memory in the given one of the MLC nonvolatile memory elements to the retained data that was associated with the stored data in the random access volatile memory”).

Thus, Petitioner fails to show that Gavens discloses or suggests the features of limitation [1h].

Furthermore, Limitations [12g] and [1h] are the same, and Petitioner relies on the same ground for both limitations. *See* Petition, 57. Accordingly, the Petition’s analysis of limitation [12g] fails for the same reasons as those given above for limitation [1h].

**B. The Petition Does Not Establish That Dependent Claims 2-11
Would Have Been Obvious Over Gavens**

Petitioner's challenges to dependent claims 2-11 in its single ground also lack merit, at least because these claims depend from claim 1, for which Petitioner's ground is deficient for the reasons stated above. *See* Petition, 45-54.

VII. CONCLUSION

For the foregoing reasons, Petitioner has not established that there is a reasonable likelihood that it will prevail on any ground with respect to even one of the Challenged Claims. Accordingly, institution should be denied.

Date: June 20, 2025

Respectfully submitted,

/Alan Whitehurst/

Alan Whitehurst

Reg. No. 43,263

awhitehurst@mckoolsmith.com

Arvind Jairam

Reg. No. 62,759

ajairam@mckoolsmith.com

Christopher McNett

Reg. No. 64,489

cmcnett@mckoolsmith.com

Christian Dorman

Reg. No. 79,118

cdorman@mckoolsmith.com

McKool Smith, P.C.

1717 K St. NW, Suite 1000

Washington, DC 20006

Telephone: (202) 370-8300

Facsimile: (202) 370-8344

Counsel for Patent Owner

CERTIFICATE OF WORD COUNT

Pursuant to 37 C.F.R. §42.24(d), Patent Owner hereby certifies that the word count for the foregoing **PATENT OWNER'S PRELIMINARY RESPONSE ON THE MERITS AND OTHER NON-DISCRETIONARY CONSIDERATIONS** is 10,554 words, which is within the limit set in 37 C.F.R. §42.24(b)(1). This word count does not include the table of contents, table of authorities, exhibit list, certificate of word count, and certificate of service. This word count was prepared using Microsoft Word.

Date: June 20, 2025

By: /Alan Whitehurst/
Alan Whitehurst, Reg. No. 43,263

CERTIFICATE OF SERVICE

I hereby certify, in accordance with 37 C.F.R. § 42.205, that the foregoing
**PATENT OWNER'S PRELIMINARY RESPONSE ON THE MERITS AND
OTHER NON-DISCRETIONARY CONSIDERATIONS** was served via
electronic mail on June 20, 2025 on the following counsel of record for Petitioner:

Cono A. Carrano
ccarrano@akingump.com
Ryan S. Stronczer
rstronczer@akingump.com
Brock F. Wilson
bfwilson@akingump.com
Akin Gump Strauss Hauer & Feld LLP
2001 K St. NW
Washington, DC 20006
Telephone: (202) 887-4000
Facsimile: (202) 887-4288

George Andrew Lever Rosbrook
arosbrook@akingump.com
Akin Gump Strauss Hauer & Feld LLP
112 E. Pecan Street, Suite 1010
San Antonio, TX 78205
Telephone: (210) 781-7026

AG-KINGSTON-VERVAIN@akingump.com

Dated: June 20, 2025

By: /Alan Whitehurst/
Alan Whitehurst, Reg. No. 43,263