

1 SHEPPARD, MULLIN, RICHTER & HAMPTON LLP
2 HARPER S. BATTS, Cal. Bar No. 242603
3 hbatts@sheppardmullin.com
4 CHRISTOPHER S. PONDER, Cal. Bar No. 296546
5 cponder@sheppardmullin.com
6 JEFFREY LIANG, Cal. Bar No. 281429
7 jliang@sheppardmullin.com
8 LAI L. YIP, Cal. Bar No. 258029
9 lyip@sheppardmullin.com
10 1540 El Camino Real Suite 120
11 Menlo Park, California 94025-4111
12 Telephone: 650.815.2600

13 Attorneys for Advanced Micro Devices, Inc. and
14 Pensando Systems, Inc.

15 **UNITED STATES DISTRICT COURT**
16 **NORTHERN DISTRICT OF CALIFORNIA, OAKLAND DIVISION**

17 CONCURRENT VENTURES, LLC AND
18 XTREAMEDGE, INC.,

19 Plaintiffs,

20 v.

21 ADVANCED MICRO DEVICES, INC. AND
22 PENSANDO SYSTEMS, INC.,

23 Defendants.

Case No. 4:25-cv-09567-JST

**DEFENDANTS ADVANCED MICRO
DEVICES, INC. AND PENSANDO
SYSTEMS, INC.'S INVALIDITY
CONTENTIONS AND ACCOMPANYING
DOCUMENT PRODUCTION**

24 IPR2025-00478
25 Exhibit 3104

1 Pursuant to the Scheduling Order (Dkt. 134), Defendants Advanced Micro Devices, Inc.
2 and Pensando Systems, Inc. hereby disclose their Invalidity Contentions and Accompanying
3 Document Production (Pat. L.R. 3-3 and 3-4).

4 The Invalidity Contentions set forth below are based on information currently known to
5 Defendants. Discovery in this case is still in its early stages, and Defendants' investigation and
6 analysis of Plaintiffs' claims and Defendants' defenses are ongoing. Defendants also have not yet
7 obtained discovery from the named inventors identified on the Asserted Patents or the individuals
8 involved in prosecuting the respective applications leading to the Asserted Patents, nor have
9 Defendants completed discovery from third parties that may have evidence relating to prior art.
10 Defendants therefore reserve the right to supplement or modify these Invalidity Contentions based
11 on further discovery, evaluation of the scope and content of the prior art, or changes in any of
12 Plaintiffs' asserted claims or infringement contentions.

13 The Court has not yet issued any claim-construction rulings. Defendants' Invalidity
14 Contentions herein are based on the claim constructions apparently underlying Plaintiffs'
15 infringement contentions. Defendants' Invalidity Contentions are not intended to, and do not,
16 reflect any acquiescence or concession by Defendants as to the proper scope or construction of the
17 Asserted Claims. To the extent that these Invalidity Contentions reflect any claim construction
18 consistent with Plaintiffs' infringement contentions, no inference is intended, nor should any be
19 drawn, that Defendants agree with Plaintiffs' claim constructions, and Defendants expressly
20 reserve the right to contest such constructions. Moreover, Defendants' Invalidity Contentions
21 should not in any way be read as an admission that any of their products or technology infringes
22 any claim of any Asserted Patents; Defendants specifically deny such infringement.

23 The prior art references identified herein may disclose the limitations of the Asserted
24 Claims either explicitly or inherently, or may be relied upon to show the state of the art in the
25 relevant time frame. The discussion of the different references herein and in the attached claim
26 charts may reflect alternative applications of the prior art against the Asserted Claims—comments
27 on one item of prior art should not be construed as necessarily applying to any of the other prior
28 art discussed herein. For example, the suggested obviousness combinations are provided in the

1 alternative to Defendants' anticipation contentions and are not to be construed to suggest that any
2 reference included in the combinations is not by itself anticipatory. To the extent that Plaintiffs
3 contend that limitations of the Asserted Claims are not disclosed in the anticipatory prior art
4 identified herein, Defendants reserve the right to identify other references and combinations that
5 address those arguments, and reserve the right to so modify, amend, or supplement these Invalidity
6 Contentions.

7 Defendants' citations to specific portions of the prior art references relied upon herein
8 should not be construed to mean that other portions of the prior art references are not relevant to
9 the invalidity of the Asserted Claims. Defendants reserve the right to rely on the entirety of any or
10 all of the prior art references (whether or not claim-charted) as a basis for asserting invalidity of
11 the Asserted Patents, and/or as necessary to supplement their contentions with additional citations
12 and evidence.

13 **I. Patent L.R. 3-3(A)-(C): Identification of Prior Art That Anticipates or Renders**
14 **Obvious Each Asserted Claim**

15 In the below subsections, Defendants identify each item of prior art that anticipates and/or
16 renders obvious each Asserted Claim. The prior art identified herein also includes items that may
17 be relied upon to show knowledge or use by others under 35 U.S.C. § 102(a)(1), public use or
18 on-sale bar under § 102(a)(1), prior filing under § 102(a)(2), derivation or prior inventorship under
19 §§ 135 and 291, and the identity of any person or entity involved in and the circumstances
20 surrounding the making of the invention before the application for each Asserted Patent, including
21 relevant dates where presently known. The prior art cited and discussed below and in the attached
22 claim charts anticipates and/or renders obvious, alone or in combination with the below-disclosed
23 references, one or more of the Asserted Claims.

24 Defendants have provided limitation-by-limitation claim charts demonstrating where each
25 limitation of the Asserted Claims of the Asserted Patents is found in the prior art. These claim
26 charts are based on Defendants' current understanding of the terms and limitations of Plaintiffs'
27 Asserted Claims, and Defendants' current understanding of the claim constructions underlying
28 Plaintiffs' infringement contentions. To the extent any claim limitation is deemed not to be

1 disclosed or otherwise satisfied by a given item of prior art, Defendants contend that the
2 difference, if any, between the claim and such prior art would have been within the knowledge of
3 and obvious to one of ordinary skill in the relevant art at the time of the alleged invention, such
4 that the claimed invention would have been obvious both in light of the prior-art item alone and/or
5 in light of a combination of prior art.

6 Defendants have provided exemplary citations to each prior art item relied upon for each
7 limitation of each Asserted Claim. These citations represent neither every location where a
8 particular claim limitation is found in the prior art nor every instance in which such prior art
9 satisfies the claim limitation. Rather, Defendants reserve the right to rely on additional, or
10 different, portions or aspects of the prior art other than those specifically cited in the claim charts
11 provided herewith. Moreover, Defendants may rely on other portions or aspects of the prior art as
12 aids to understanding and interpreting the portions that are cited. Where Defendants cite a
13 particular figure in a prior art reference, such citation should be understood to encompass the
14 caption and description of the figure and any text relating to the figure, in addition to the figure
15 itself. Conversely, where a cited portion of text refers to a figure, the citation should be
16 understood to include the figure as well. Defendants may additionally rely on information
17 incorporated by reference in the identified prior art to provide additional support for Defendants'
18 contention for a given claim limitation.

19 Defendants reserve the right to supplement and/or amend these contentions based on the
20 herein-identified prior art considered individually or in combination with additional information
21 learned during discovery and/or in light of claim-construction developments in this case.
22 Defendants may also become aware of additional prior art as their investigation and discovery in
23 this case progress, and Defendants reserve the right to use any such prior art to demonstrate that
24 the Asserted Claims of the Asserted Patents are invalid.

25 Defendants contend that the prior-art references identified below are anticipatory
26 references under 35 U.S.C. §§ 102(a)(1)-(2). To the extent any of the anticipatory references were
27 embodied in actual systems, such systems may constitute prior art based upon which Defendants
28

1 will continue to develop further invalidity defenses as discovery and Defendants' investigation
2 progress.

3 As the Supreme Court has established, "[t]he combination of familiar elements according
4 to known methods is likely to be obvious when it does no more than yield predictable results."
5 *KSR Int'l Co. v. Teleflex, Inc.*, 550 U.S. 398, 416 (2007).

6 When a work is available in one field of endeavor, design incentives and other market
7 forces can prompt variations of it, either in the same field or a different one. If one
8 of ordinary skill can implement a predictable variation, § 103 likely bars its
9 patentability. For the same reason, if a technique has been used to improve one
device, and one of ordinary skill in the art would recognize that it would improve
similar devices in the same way, using the technique is obvious unless its actual
application is beyond his or her skill.

10 *Id.* at 417. "[I]n many cases a person of ordinary skill will be able to fit the teachings of multiple
11 patents together like pieces of a puzzle." *Id.* at 420. Indeed, it is sufficient that a combination of
12 elements was "obvious to try":

13 When there is a design need or market pressure to solve a problem and there are a
14 finite number of identified, predictable solutions, a person of ordinary skill has good
15 reason to pursue the known options within his or her technical grasp. If this leads to
16 the anticipated success, it is likely the product not of innovation but of ordinary skill
and common sense.

17 *Id.* at 421. "In that instance the fact that a combination was obvious to try might show that it was
18 obvious under § 103." *Id.* The Supreme Court has admonished against the application of "[r]igid
19 preventative rules that deny factfinders recourse to common sense...." *Id.* This is because
20 "[g]ranting patent protection to advances that would occur in the ordinary course without real
21 innovation retards progress and may, in the case of patents combining previously known elements,
22 deprive prior inventions of their value or utility." *Id.* at 419.

23 For each Asserted Patent, all prior art items relied upon herein are directed to the same
24 fields of endeavor. Accordingly, one of ordinary skill in the art would have recognized that the
25 results of the combinations were predictable, and would have been motivated to modify and/or
26 combine the prior art items for each Asserted Patent to arrive at the claimed inventions. At least
27 the following rationales support a finding of obviousness based on such combinations of prior art:

- 1 A. combining prior-art elements according to known methods to yield predictable
2 results;
- 3 B. simple substitution of one known element for another to obtain predictable results;
- 4 C. use of known techniques to improve similar devices (or methods) in the same way;
- 5 D. applying a known technique to a known device (or method) ready for improvement
6 to yield predictable results;
- 7 E. choosing from a finite number of identified, predictable solutions, with a
8 reasonable expectation of success (“obvious to try”);
- 9 F. prompting of variations of known work in one field of endeavor for use in either
10 the same field or a different one based on design incentives or other market forces
11 if the variations would have been predictable to one skilled in the art; and
- 12 G. some teaching, suggestion, or motivation in the prior art that would have led one
13 skilled in the art to modify the prior art or to combine prior-art teachings to arrive
14 at the claimed invention.

15 Thus, to the extent that the Asserted Claims are deemed not to be anticipated, Defendants
16 contend that the Asserted Claims are invalid as obvious under 35 U.S.C. § 103 in view of at least
17 the prior art listed herein and set forth in the attached Exhibits. Defendants’ contentions that the
18 Asserted Claims are obvious under 35 U.S.C. § 103 are in no way an admission or suggestion that
19 any prior art relied upon herein is not anticipatory under 35 U.S.C. § 102.

20 Defendants set forth certain specific exemplar combinations of prior art below in addition
21 to combinations identified in their claim charts, based on Defendants’ present understanding of the
22 Asserted Claims and the apparent claim constructions adopted by Plaintiffs. As discussed above,
23 Defendants contend that each prior-art item identified above is anticipatory under 35 U.S.C. § 102
24 and/or alone renders the claimed invention obvious under 35 U.S.C. § 103.

25 **A. U.S. Patent No. 8,924,596**

26 The following table identifies the prior art for the ’596 patent. Each item below constitutes
27 prior art under one or more provisions of 35 U.S.C. §§ 102(a)(1)-(2) by reason of its status as a
28 publication and/or patent disclosing, or a system on sale and/or in public knowledge or use in the
United States embodying the claimed invention before Plaintiffs’ asserted priority date of the ’596

1 patent. For each item of prior art, Defendants have attached a claim chart identifying where each
2 claim limitation is found in the cited prior art.

3 **Table A: Prior Art to the '596 Patent**

4	Prior Art to '596	Anticipation/ Obviousness	Claim Chart
5			
6	U.S. Patent Application Publication No. 2011/0161976 to Alexander et al. ("Alexander"); 7 filed on December 30, 2009; published on June 30, 2011.	A & O	Exhibit 596-01
8	U.S. Patent No. 6,229,813 to Buchko et al. ("Buchko"); filed on November 25, 1998; issued 9 on May 8, 2001.	A & O	Exhibit 596-02
10	U.S. Patent No. 7,743,197 to Chavan et al., ("Chavan"); filed on May 11, 2006; issued on 11 June 22, 2010.	A & O	Exhibit 596-03
12	U.S. Patent No. 8,559,439 to Dongare et al. ("Dongare"); filed on November 3, 2011; issued 13 on October 15, 2013.	O	Exhibit 596-04
14	U.S. Patent No. 6,173,307 to Drews ("Drews"); filed on August 20, 1998; issued on January 9, 15 2001.	O	Exhibit 596-05
16	U.S. Patent No. 8,051,227 to Gewirtz et al. ("Gewirtz"); filed on May 10, 2010; issued on 17 November 1, 2011.	O	Exhibit 596-06
18	U.S. Patent No. 6,055,579 to Goyal et al. ("Goyal"); filed on November 17, 1997; issued 19 on April 25, 2000.	O	Exhibit 596-07
20	U.S. Patent No. 5,204,954 to Hammer et al. ("Hammer"); filed on July 29, 1992; issued on 21 April 20, 1993.	A & O	Exhibit 596-08
22	U.S. Patent No. 7,206,901 to Kanai et al. ("Kanai"); filed on April 1, 2004; issued on April 23 17, 2007.	O	Exhibit 596-09
24	U.S. Patent No. 7,398,528 to Khawand et al. ("Khawand"); filed on November 13, 2004; 25 issued on July 8, 2008.	A & O	Exhibit 596-10
26			
27			
28			

Prior Art to '596	Anticipation/ Obviousness	Claim Chart
U.S. Patent Application Publication No. 2003/0149820 to Kolinummi et al. ("Kolinummi"); filed on January 15, 2003; published on August 7, 2003.	O	Exhibit 596-11
U.S. Patent No. 8,463,958 to Mannava et al. ("Mannava"); filed on August 8, 2011; issued on June 11, 2013.	O	Exhibit 596-12
U.S. Patent No. 6,810,462 to Matsunami et al. ("Matsunami"); filed on August 19, 2002; issued on October 26, 2004.	O	Exhibit 596-13
U.S. Patent No. 7,865,684 to Michaylov et al. ("Michaylov"); filed on June 27, 2005; issued on January 4, 2011.	O	Exhibit 596-14
U.S. Patent No. 8,028,284 to Ohkawa ("Ohkawa"); filed on October 31, 2006; issued on September 27, 2011.	O	Exhibit 596-15
U.S. Patent No. 5,805,930 to Rosenthal et al. ("Rosenthal"); filed on March 28, 1997; issued on September 8, 1998.	O	Exhibit 596-16
U.S. Patent No. 5,907,485 to Van Loo et al. ("Van Loo"); filed on March 31, 1995; issued on May 25, 1999.	A & O	Exhibit 596-17
ARM System(s), including but not limited to TCP/IP RTOS, ARM1156T2-S, Cortex-R5, Cortex-A, ARMv7, Corelink CCI-400, were sold, offered for sale, made known, and/or publicly used before December 6, 2013. ARM System(s) were sold, offered for sale, and/or made known to the public and third-party customers before this date, by December 5, 2013 at the latest. ARM System(s) were publicly used by ARM, the public, and third-party customers before this date, by December 5, 2013 at the latest.	A & O	Exhibit 596-18

Prior Art to '596	Anticipation/ Obviousness	Claim Chart
Intel/Fulcrum System(s), including but not limited to FM4000 24-Port 10G Ethernet L2/L3/L4 Switch/Router, FM4000 Series 10GbE Low Latency L2/L3 Switching Silicon, FocalPoint FM4000 Series 10G Enhanced Ethernet Switch-Router, FocalPointII A 240 Gb/s switch router, were sold, offered for sale, made known, and/or publicly used before December 6, 2013. Intel Fulcrum System(s) were sold, offered for sale, and/or made known to the public and third-party customers before this date, by December 5, 2013 at the latest. Intel/Fulcrum System(s) were publicly used by Intel/Fulcrum, the public, and third-party customers before this date, by December 5, 2013 at the latest.	A & O	Exhibit 596-19
NVIDIA/Mellanox System(s) including but not limited to RDMA Aware Networks, ConnectX-2 Single/Dual-Port Adapter Device with Virtual Protocol Interconnect, Mellanox OFED for Linux, Mellanox Technologies InfiniBand Fabrics, were sold, offered for sale, made known, and/or publicly used before December 6, 2013. NVIDIA/Mellanox System(s) were sold, offered for sale, and/or made known to the public and third-party customers before this date, by December 5, 2013 at the latest. NVIDIA/Mellanox System(s) were publicly used by the NVIDIA/Mellanox, the public, and third-party customers before this date, by December 5, 2013 at the latest.	A & O	Exhibit 596-20
U.S. Patent Application Publication No. 2005/0071529 to Borkowski et al. ("Borkowski"); filed on September 30, 2003; published on March 31, 2005	O	Exhibit 596-21
U.S. Patent No. 7,844,758 to Hughes ("Hughes"); filed on June 18, 2003; issued on November 30, 2010	A & O	Exhibit 596-22
U.S. Patent No. 6,938,253 to Kim ("Kim"); filed on May 2, 2001; issued on August 30, 2005	O	Exhibit 596-23

Prior Art to '596	Anticipation/ Obviousness	Claim Chart
U.S. Patent Application Publication No. 2006/0136640 to Tuan ("Tuan"); filed on April 28, 2005; published on June 22, 2006	O	Exhibit 596-24
IBM Cell System(s) including but not limited to IBM Cell Broadband Engine were sold, offered for sale, made known, and/or publicly used before December 6, 2013. IBM Cell System(s) were sold, offered for sale, and/or made known to the public and third-party customers before this date, by December 5, 2013 at the latest. IBM Cell System(s) were publicly used by IBM, the public, and third-party customers before this date, by December 5, 2013 at the latest.	A & O	Exhibit 596-25

In addition to the anticipation and obviousness explained in this section, the '596 patent is anticipated and/or rendered obvious by prior art as explained by Defendants' Petition for Inter Partes Review (assigned Case No. IPR2025-00478), which is incorporated herein by reference:

Grounds	Claims	Prior Art
1	1-2, 5-8, 11-14, 17-18	Dongare in view of Gewirtz
2	3-4, 9-10, 15-16	Ground 1 further in view of Matsunami
3	1-2, 6-8, 12-14, 18	Khawand in view of Rosenthal
4	1-2, 6-8, 12-14, 18	Ground 3 further in view of Ohkawa
5	3-4, 9-10, 15-16	Ground 4 further in view of Matsunami

1. Anticipation

For each item of prior art identified above as an anticipation reference, Defendants contend that the prior art item anticipates the one or more Asserted Claims of the '596 patent, including for the reasons illustrated in the accompanying claim chart and based on Plaintiffs' infringement allegations and contentions for the '596 patent.

2. Obviousness

A person of ordinary skill in the art would have been motivated to combine any reference in Table A with any other reference(s) in Table A. To the extent that any prior art identified in

1 Table A above is deemed to not anticipate an Asserted Claim, the prior art in Table A may be
2 combined with one or more prior art references in Table A to render the claim obvious. All of
3 these items of prior art are directed to and propose solutions similar to those discussed in the '596
4 patent, and is in the same field. All of these items of prior art are directed to multiprocessor
5 systems and/or computer architecture and propose solutions similar to those discussed in the '596
6 patent.

7 The following are exemplary combinations of prior art that would have been obvious to a
8 POSITA:

- 9 • **Alexander + Drews**
- 10 • **Alexander + Gewirtz**
- 11 • **Alexander + Goyal**
- 12 • **Alexander + Kolinummi**
- 13 • **Alexander + Ohkawa**
- 14 • **Alexander + Rosenthal**
- 15 • **Alexander + Borkowski**
- 16 • **Alexander + Hughes**
- 17 • **Alexander + Kim**
- 18 • **Alexander + Tuan**
- 19 • **Any of the above + Kanai**
- 20 • **Any of the above + Matsunami**
- 21
- 22

23 In addition to the above-mentioned motivations to combine, Alexander discloses a multi-
24 node processing system with queues, including a global command queue. *See* Abstract.
25 Alexander describes its technical field as “multiprocessor data processing systems” and discloses a
26 “method and system for completing parallel processing of work items of a single work set
27 distributed across multiple processing units of a multi-node data processing system.” ¶ [0002].
28

1 Alexander teaches a system with distributed processing across multiple processing units that
2 require queue management and flow control at both the global and local level. ¶¶ [0044], [0049].

3 Therefore, given the context provided by Alexander, a POSITA would have naturally
4 looked to other references that provide details related to multiprocessor systems, computer
5 architecture, queue management, reservation systems, and flow control. For example, a POSITA
6 would have been motivated to combine Alexander with Drews, which teaches a queue reservation
7 system based on a counter for multiprocessor systems. As an additional example, a POSITA
8 would have been motivated to combine Alexander with Gewirtz, which teaches a programmable
9 queue structure for multiprocessor systems with flow control based on a queue state register. As
10 an additional example, a POSITA would have been motivated to combine Alexander with Goyal,
11 which teaches queue management based on a centralized command processor for multiprocessor
12 systems. As an additional example, a POSITA would have been motivated to combine Alexander
13 with Kolinummi, which teaches a reservation system using a hardware-based semaphore for
14 multiprocessor systems. As an additional example, a POSITA would have been motivated to
15 combine Alexander with Ohkawa, which teaches queue management, queue structure, and queue
16 placement for multiprocessor systems. As an additional example, a POSITA would have been
17 motivated to combine Alexander with Rosenthal, which teaches queue management and flow
18 control based on a flow control register that indicates the amount of space available in the queue
19 for multiprocessor systems. As an additional example, a POSITA would have been motivated to
20 combine Alexander with Borkowski, which teaches the use of counting semaphores for use in
21 resource-locking including FIFOs. As an additional example, a POSITA would have been
22 motivated to combine Alexander with Hughes, which teaches a queue reservation system with a free
23 entry counter. As an additional example, a POSITA would have been motivated to combine
24 Alexander with Kim, which teaches a multiprocessing architecture that enables inter-processor
25 communication using a set of hardware semaphore registers. As an additional example, a
26 POSITA would have been motivated to combine Alexander with Tuan, which teaches a
27 multiprocessor system that uses hardware semaphores to arbitrate access to shared resources. A
28 POSITA would naturally have looked to these references and been motivated to combine them

1 with Alexander because, e.g., they provide implementation details and facilitate resource sharing,
2 optimize queue and reservation management, and reduce coordination overhead for accessing
3 shared resources in multiprocessor systems. This would have been simple combination of known
4 prior art element according to known methods to yield predictable results as well as a simple
5 substitution of known elements for another to obtain predictable results.

6 Additionally, Alexander discloses support for I/O operations. Therefore, a POSITA would
7 naturally have looked to references that support and provide features for various I/O operations,
8 including known storage systems. For example, a POSITA would have been motivated to
9 combine Alexander with Kanai or Matsunami, as both teach multiprocessor-based storage
10 systems, including support for NAS and SAN. SAN and NAS had been known and used in
11 networked settings and data centers for many years to provide data storage over networks. The
12 motivations to combine explained in the incorporated Petition (IPR2025-00487) also apply to
13 Alexander, and a POSITA would have been motivated to combine for similar reasons.

- 14 • **Buchko + Drews**
- 15 • **Buchko + Gewirtz**
- 16 • **Buchko + Goyal**
- 17 • **Buchko + Kolinummi**
- 18 • **Buchko + Ohkawa**
- 19 • **Buchko + Rosenthal**
- 20 • **Buchko + Burkowski**
- 21 • **Buchko + Hughes**
- 22 • **Buchko + Kim**
- 23 • **Buchko + Tuan**
- 24 • **Any of the above + Kanai**
- 25 • **Any of the above + Matsunami**
- 26
- 27
- 28

1 In addition to the above-mentioned motivations to combine, Buchko discloses a
2 multiprocessor system with distributed work among processors that share a queue. *See* Abstract.
3 Buchko describes its technical field as “digital systems wherein processed items are temporarily
4 stored in queues, typically having limited capacity, pending further processing and more
5 particularly to a method and mechanism for controlling queue size in such systems.” 1:6-12.
6 Buchko explains that “there is a requirement for an improved mechanism for controlling the queue
7 length or queue size in a multi task processor system or a multiprocessor system in which items
8 removed from queues continue to use up system resources during subsequent processing.” 2:36-
9 40. Buchko acknowledges the need for flow control: “It is then placed in one of the
10 aforementioned COS queues after the queue has been checked to determine whether there is
11 sufficient room. This is done by checking the queue counter or calculating the space between the
12 enqueue pointer and the virtual dequeue pointer. If there is not sufficient space the packet is
13 discarded or one of the queued packets is discarded to make room for the new packet, if
14 appropriate.” 6:43-48.

15 Therefore, given the context provided by Buchko, a POSITA would have naturally looked
16 to other references that provide details related to multiprocessor systems, computer architecture,
17 queue management, reservation systems, and flow control. For example, a POSITA would have
18 been motivated to combine Buchko with Drews, which teaches a queue reservation system based
19 on a counter for multiprocessor systems. As an additional example, a POSITA would have been
20 motivated to combine Buchko with Gewirtz, which teaches a programmable queue structure for
21 multiprocessor systems with flow control based on a queue state register. As an additional
22 example, a POSITA would have been motivated to combine Buchko with Goyal, which teaches
23 queue management based on a centralized command processor for multiprocessor systems. As an
24 additional example, a POSITA would have been motivated to combine Buchko with Kolinummi,
25 which teaches a reservation system using a hardware-based semaphore for multiprocessor systems.
26 As an additional example, a POSITA would have been motivated to combine Buchko with
27 Ohkawa, which teaches queue management, queue structure, and queue placement for
28 multiprocessor systems. As an additional example, a POSITA would have been motivated to

1 combine Buchko with Rosenthal, which teaches queue management and flow control based on a
2 flow control register that indicates the amount of space available in the queue for multiprocessor
3 systems. As an additional example, a POSITA would have been motivated to combine Buchko
4 with Borkowski, which teaches the use of counting semaphores for use in resource-locking
5 including FIFOs. As an additional example, a POSITA would have been motivated to combine
6 Buchko with Hughes, which teaches a queue reservation system with a free entry counter. As an
7 additional example, a POSITA would have been motivated to combine Buchko with Kim, which
8 teaches a multiprocessing architecture that enables inter-processor communication using a set of
9 hardware semaphore registers. As an additional example, a POSITA would have been motivated
10 to combine Buchko with Tuan, which teaches a multiprocessor system that uses hardware
11 semaphores to arbitrate access to shared resources. A POSITA would naturally have looked to
12 these references and been motivated to combine them with Buchko because, e.g., they provide
13 implementation details and facilitate resource sharing, optimize queue and reservation
14 management, and reduce coordination overhead for accessing shared resources in multiprocessor
15 systems. This would have been simple combination of known prior art element according to
16 known methods to yield predictable results as well as a simple substitution of known elements for
17 another to obtain predictable results.

18 Additionally, Buchko discloses the support of network operations. A POSITA would
19 naturally have looked to references that support and provide features for various network
20 operations, including network-based storage systems. For example, a POSITA would have been
21 motivated to combine Buchko with Kanai or Matsunami, as both teach multiprocessor-based
22 networked-storage systems, including support for NAS and SAN. SAN and NAS had been known
23 and used in networked settings and data centers for many years to provide data storage over
24 networks. The motivations to combine explained in the incorporated Petition (IPR2025-00487)
25 also apply to Buchko, and a POSITA would have been motivated to combine for similar reasons.

- 26 • **Chavan + Drews**
- 27 • **Chavan + Gewirtz**

- 1 • **Chavan + Goyal**
- 2 • **Chavan + Kolinummi**
- 3 • **Chavan + Ohkawa**
- 4 • **Chavan + Rosenthal**
- 5 • **Chavan + Borkowski**
- 6 • **Chavan + Hughes**
- 7 • **Chavan + Kim**
- 8 • **Chavan + Tuan**
- 9 • **Any of the above + Kanai**
- 10 • **Any of the above + Matsunami**
- 11
- 12

13 In addition to the above-mentioned motivations to combine, Chavan discloses an
14 intelligent network processor with multiple processors. *See* Abstract, 1:35-37. Chavan describes
15 its technical field as “Intelligent Network Processors (INPs) for use in computer network
16 systems.” 1:13-14. Chavan explains that its controller “creates queues in the external DDR
17 memory for storing data messages, and also writes to and reads from such queues.” 3:59-62.
18 Chavan also acknowledges the need for flow control and teaches a credit-based system: “Credits
19 are essentially the number of free buffers at a particular port for storing frames or packets.”
20 6:54-55.

21 Therefore, given the context provided by Chavan, a POSITA would have naturally looked
22 to other references that provide details related to multiprocessor systems, computer architecture,
23 queue management, reservation systems, and flow control. For example, a POSITA would have
24 been motivated to combine Chavan with Drews, which teaches a queue reservation system based
25 on a counter for multiprocessor systems. As an additional example, a POSITA would have been
26 motivated to combine Chavan with Gewirtz, which teaches a programmable queue structure for
27 multiprocessor systems with flow control based on a queue state register. As an additional
28 example, a POSITA would have been motivated to combine Chavan with Goyal, which teaches

1 queue management based on a centralized command processor for multiprocessor systems. As an
2 additional example, a POSITA would have been motivated to combine Chavan with Kolinummi,
3 which teaches a reservation system using a hardware-based semaphore for multiprocessor systems.
4 As an additional example, a POSITA would have been motivated to combine Chavan with
5 Ohkawa, which teaches queue management, queue structure, and queue placement for
6 multiprocessor systems. As an additional example, a POSITA would have been motivated to
7 combine Chavan with Rosenthal, which teaches queue management and flow control based on a
8 flow control register that indicates the amount of space available in the queue for multiprocessor
9 systems. As an additional example, a POSITA would have been motivated to combine Chavan
10 with Borkowski, which teaches the use of counting semaphores for use in resource-locking
11 including FIFOs. As an additional example, a POSITA would have been motivated to combine
12 Chavan with Hughes, which teaches a queue reservation system with a free entry counter. As an
13 additional example, a POSITA would have been motivated to combine Chavan with Kim, which
14 teaches a multiprocessing architecture that enables inter-processor communication using a set of
15 hardware semaphore registers. As an additional example, a POSITA would have been motivated
16 to combine Chavan with Tuan, which teaches a multiprocessor system that uses hardware
17 semaphores to arbitrate access to shared resources. A POSITA would naturally have looked to
18 these references and been motivated to combine them with Chavan because, e.g., they provide
19 implementation details and facilitate resource sharing, optimize queue and reservation
20 management, and reduce coordination overhead for accessing shared resources in multiprocessor
21 systems. This would have been simple combination of known prior art element according to
22 known methods to yield predictable results as well as a simple substitution of known elements for
23 another to obtain predictable results.

24 Additionally, Chavan discloses the support of I/O operations and network-based storage
25 such as SAN. A POSITA would naturally have looked to references that support and provide
26 features for various I/O operations and network-based storage, including storage systems. For
27 example, a POSITA would have been motivated to combine Chavan with Kanai or Matsunami, as
28 both teach multiprocessor-based storage systems, including support for NAS and SAN. SAN and

NAS had been known and used in networked settings and data centers for many years to provide data storage over networks. The motivations to combine explained in the incorporated Petition (IPR2025-00487) also apply to Chavan, and a POSITA would have been motivated to combine for similar reasons.

- **Dongare + Drews**
- **Dongare + Gewirtz**
- **Dongare + Goyal**
- **Dongare + Kolinummi**
- **Dongare + Ohkawa**
- **Dongare + Rosenthal**
- **Dongare + Borkowski**
- **Dongare + Hughes**
- **Dongare + Kim**
- **Dongare + Tuan**
- **Any of the above + Kanai**
- **Any of the above + Matsunami**

In addition to the above-mentioned motivations to combine, Dongare discloses a multiprocessor with multiple engines, input queues, and output queues. *See* Abstract. Dongare describes its technical field as “multi-engine processing environments. More particularly, the present disclosure relates to queue ordering in a multi-engine processing system.” 1:14-17. Dongare explains that “commands within a command input queue are related only in their sequence and are otherwise independent of other commands in the input queue. Hence, the commands within a particular command input queue can be processed independent of each other.” 5:27-32. Dongare also acknowledges the need for flow control: “The depth of the command output buffers can be chosen based on the bandwidth required by the system. The depth of a command output buffer refers to the number of processed commands or segments that can be

1 stored in the command output buffer. By having a dedicated command buffer for each engine,
2 each engine can directly access its command output buffer to write out processed commands
3 without any waiting time.” 12:20-27.

4 Therefore, given the context provided by Dongare, a POSITA would have naturally looked
5 to other references that provide details related to multiprocessor systems, computer architecture,
6 queue management, reservation systems, and flow control. For example, a POSITA would have
7 been motivated to combine Dongare with Drews, which teaches a queue reservation system based
8 on a counter for multiprocessor systems. As an additional example, a POSITA would have been
9 motivated to combine Dongare with Gewirtz, which teaches a programmable queue structure for
10 multiprocessor systems with flow control based on a queue state register. As an additional
11 example, a POSITA would have been motivated to combine Dongare with Goyal, which teaches
12 queue management based on a centralized command processor for multiprocessor systems. As an
13 additional example, a POSITA would have been motivated to combine Dongare with Kolinummi,
14 which teaches a reservation system using a hardware-based semaphore for multiprocessor systems.
15 As an additional example, a POSITA would have been motivated to combine Dongare with
16 Ohkawa, which teaches queue management, queue structure, and queue placement for
17 multiprocessor systems. As an additional example, a POSITA would have been motivated to
18 combine Dongare with Rosenthal, which teaches queue management and flow control based on a
19 flow control register that indicates the amount of space available in the queue for multiprocessor
20 systems. As an additional example, a POSITA would have been motivated to combine Dongare
21 with Borkowski, which teaches the use of counting semaphores for use in resource-locking
22 including FIFOs. As an additional example, a POSITA would have been motivated to combine
23 Dongare with Hughes, which teaches a queue reservation system with a free entry counter. As an
24 additional example, a POSITA would have been motivated to combine Dongare with Kim, which
25 teaches a multiprocessing architecture that enables inter-processor communication using a set of
26 hardware semaphore registers. As an additional example, a POSITA would have been motivated
27 to combine Dongare with Tuan, which teaches a multiprocessor system that uses hardware
28 semaphores to arbitrate access to shared resources. A POSITA would naturally have looked to

1 these references and been motivated to combine them with Dongare because, e.g., they provide
2 implementation details and facilitate resource sharing, optimize queue and reservation
3 management, and reduce coordination overhead for accessing shared resources in multiprocessor
4 systems. This would have been simple combination of known prior art element according to
5 known methods to yield predictable results as well as a simple substitution of known elements for
6 another to obtain predictable results.

7 Additionally, Dongare discloses the support of I/O operations, including disk-based
8 storage. A POSITA would naturally have looked to references that support and provide features
9 for various I/O operations, including storage systems. For example, a POSITA would have been
10 motivated to combine Dongare with Kanai or Matsunami, as both teach multiprocessor-based
11 storage systems, including support for NAS and SAN. SAN and NAS had been known and used
12 in networked settings and data centers for many years to provide data storage over networks. The
13 motivations to combine explained in the incorporated Petition (IPR2025-00487) also apply to
14 Dongare, and a POSITA would have been motivated to combine for similar reasons.

- 15 • **Hammer + Drews**
- 16 • **Hammer + Gewirtz**
- 17 • **Hammer + Goyal**
- 18 • **Hammer + Kolinummi**
- 19 • **Hammer + Ohkawa**
- 20 • **Hammer + Rosenthal**
- 21 • **Hammer + Borkowski**
- 22 • **Hammer + Hughes**
- 23 • **Hammer + Kim**
- 24 • **Hammer + Tuan**
- 25 • **Any of the above + Kanai**
- 26 • **Any of the above + Matsunami**
- 27
- 28

1 In addition to the above-mentioned motivations to combine, Hammer discloses a
2 multiprocessor I/O unit. *See* Abstract. Hammer describes its technical field as “control of work
3 flow between processes on a bus.” 1:27-31. Hammer’s system is a multiprocessor I/O system that
4 consists of “peer processors,” which causes “each processor [to] serve as many work requests as it
5 generates work requests.” 2:9-12. Hammer explains that the “bus hardware also controls
6 queueing of messages from the bus units in main storage queues.” Hammer also acknowledges
7 the need for flow control: “Messages are transferred between bus managers to control and re-
8 synchronize messages when a connection group does not have sufficient resources to service a
9 request intended for a process connected to that group. When a target bus unit cannot accept
10 additional work for a particular connection group, it returns a message with QUEUE FULL status
11 specified.” 10:59-65.

12 Therefore, given the context provided by Hammer, a POSITA would have naturally looked
13 to other references that provide details related to multiprocessor systems, computer architecture,
14 queue management, reservation systems, and flow control. For example, a POSITA would have
15 been motivated to combine Hammer with Drews, which teaches a queue reservation system based
16 on a counter for multiprocessor systems. As an additional example, a POSITA would have been
17 motivated to combine Hammer with Gewirtz, which teaches a programmable queue structure for
18 multiprocessor systems with flow control based on a queue state register. As an additional
19 example, a POSITA would have been motivated to combine Hammer with Goyal, which teaches
20 queue management based on a centralized command processor for multiprocessor systems. As an
21 additional example, a POSITA would have been motivated to combine Hammer with Kolinummi,
22 which teaches a reservation system using a hardware-based semaphore for multiprocessor systems.
23 As an additional example, a POSITA would have been motivated to combine Hammer with
24 Ohkawa, which teaches queue management, queue structure, and queue placement for
25 multiprocessor systems. As an additional example, a POSITA would have been motivated to
26 combine Hammer with Rosenthal, which teaches queue management and flow control based on a
27 flow control register that indicates the amount of space available in the queue for multiprocessor
28 systems. As an additional example, a POSITA would have been motivated to combine Hammer

1 with Borkowski, which teaches the use of counting semaphores for use in resource-locking
2 including FIFOs. As an additional example, a POSITA would have been motivated to combine
3 Hammer with Hughes, which teaches a queue reservation system with a free entry counter. As an
4 additional example, a POSITA would have been motivated to combine Hammer with Kim, which
5 teaches a multiprocessing architecture that enables inter-processor communication using a set of
6 hardware semaphore registers. As an additional example, a POSITA would have been motivated
7 to combine Hammer with Tuan, which teaches a multiprocessor system that uses hardware
8 semaphores to arbitrate access to shared resources. A POSITA would naturally have looked to
9 these references and been motivated to combine them with Hammer because, e.g., they provide
10 implementation details and facilitate resource sharing, optimize queue and reservation
11 management, and reduce coordination overhead for accessing shared resources in multiprocessor
12 systems. This would have been simple combination of known prior art element according to
13 known methods to yield predictable results as well as a simple substitution of known elements for
14 another to obtain predictable results.

15 Additionally, Hammer discloses the support of I/O operations. A POSITA would naturally
16 have looked to references that support and provide features for various I/O operations, including
17 storage systems. For example, a POSITA would have been motivated to combine Hammer with
18 Kanai or Matsunami, as both teach multiprocessor-based storage systems, including support for
19 NAS and SAN. SAN and NAS had been known and used in networked settings and data centers
20 for many years to provide data storage over networks. The motivations to combine explained in
21 the incorporated Petition (IPR2025-00487) also apply to Hammer, and a POSITA would have
22 been motivated to combine for similar reasons.

- 23 • **Khawand + Drews**
- 24 • **Khawand + Gewirtz**
- 25 • **Khawand + Goyal**
- 26 • **Khawand + Kolinummi**
- 27 • **Khawand + Ohkawa**
- 28

- 1 • **Khawand + Rosenthal**
- 2 • **Khawand + Borkowski**
- 3 • **Khawand + Hughes**
- 4 • **Khawand + Kim**
- 5 • **Khawand + Tuan**
- 6 • **Any of the above + Kanai**
- 7 • **Any of the above + Matsunami**

9 In addition to the above-mentioned motivations to combine, Khawand discloses a
10 multiprocessor system that delegates tasks among its processors through the use of command task
11 queues. *See* Abstract. Khawand describes its technical field as “mobile devices including an
12 efficient multiprocessor system and methods thereof.” 1:8-10. Khawand explains that “[e]ach
13 task [] comprises at least one queue [] for processing at least one service application.” 3:54-58.
14 Khawand also acknowledges the need for flow control: “In step 604, the queue [] of the second
15 processor [] is evaluated according to flow control criteria. ... the flow control criteria can include
16 a capacity threshold for each queue [].” 4:6-16.

17 Therefore, given the context provided by Khawand, a POSITA would have naturally
18 looked to other references that provide details related to multiprocessor systems, computer
19 architecture, queue management, reservation systems, and flow control. For example, a POSITA
20 would have been motivated to combine Khawand with Drews, which teaches a queue reservation
21 system based on a counter for multiprocessor systems. As an additional example, a POSITA
22 would have been motivated to combine Khawand with Gewirtz, which teaches a programmable
23 queue structure for multiprocessor systems with flow control based on a queue state register. As
24 an additional example, a POSITA would have been motivated to combine Khawand with Goyal,
25 which teaches queue management based on a centralized command processor for multiprocessor
26 systems. As an additional example, a POSITA would have been motivated to combine Khawand
27 with Kolinummi, which teaches a reservation system using a hardware-based semaphore for
28 multiprocessor systems. As an additional example, a POSITA would have been motivated to

1 combine Khawand with Ohkawa, which teaches queue management, queue structure, and queue
2 placement for multiprocessor systems. As an additional example, a POSITA would have been
3 motivated to combine Khawand with Rosenthal, which teaches queue management and flow
4 control based on a flow control register that indicates the amount of space available in the queue
5 for multiprocessor systems. As an additional example, a POSITA would have been motivated to
6 combine Khawand with Borkowski, which teaches the use of counting semaphores for use in
7 resource-locking including FIFOs. As an additional example, a POSITA would have been
8 motivated to combine Khawand with Hughes, which teaches a queue reservation system with a free
9 entry counter. As an additional example, a POSITA would have been motivated to combine
10 Khawand with Kim, which teaches a multiprocessing architecture that enables inter-processor
11 communication using a set of hardware semaphore registers. As an additional example, a
12 POSITA would have been motivated to combine Khawand with Tuan, which teaches a
13 multiprocessor system that uses hardware semaphores to arbitrate access to shared resources. A
14 POSITA would naturally have looked to these references and been motivated to combine them
15 with Khawand because, e.g., they provide implementation details and facilitate resource sharing,
16 optimize queue and reservation management, and reduce coordination overhead for accessing
17 shared resources in multiprocessor systems. This would have been simple combination of known
18 prior art element according to known methods to yield predictable results as well as a simple
19 substitution of known elements for another to obtain predictable results.

20 Additionally, Khawand discloses the support of I/O operations. A POSITA would
21 naturally have looked to references that support and provide features for various I/O operations,
22 including storage systems. For example, a POSITA would have been motivated to combine
23 Khawand with Kanai or Matsunami, as both teach multiprocessor-based storage systems,
24 including support for NAS and SAN. SAN and NAS had been known and used in networked
25 settings and data centers for many years to provide data storage over networks. The motivations
26 to combine explained in the incorporated Petition (IPR2025-00487) also apply to Khawand, and a
27 POSITA would have been motivated to combine for similar reasons.

28 • **Mannava + Drews**

- 1 • **Mannava + Gewirtz**
- 2 • **Mannava + Goyal**
- 3 • **Mannava + Kolinummi**
- 4 • **Mannava + Ohkawa**
- 5 • **Mannava + Rosenthal**
- 6 • **Mannava + Borkowski**
- 7 • **Mannava + Hughes**
- 8 • **Mannava + Kim**
- 9 • **Mannava + Tuan**
- 10 • **Any of the above + Kanai**
- 11 • **Any of the above + Matsunami**

12 In addition to the above-mentioned motivations to combine, Mannava discloses a
13 multiprocessor credit-based data processing system. *See* Abstract; Figure 3. Mannava describes
14 its technical field as “data processing systems, and in particular, to systems having a number of
15 initiator devices issuing transaction requests to a number of recipient devices and to the allocation
16 of resources within the recipient devices for executing these transaction requests.” 1:8-12.
17 Mannava explains that “[r]ecipient device [] comprises a buffer [] for storing pending transaction
18 requests that have been received prior to them being executed by the recipient device.” 8:24-27.
19 Mannava acknowledges the need for flow control: “Credit generator [] monitors to determine
20 when there are available storage locations.” 8:66-67.

21 Therefore, given the context provided by Mannava, a POSITA would have naturally
22 looked to other references that provide details related to multiprocessor systems, computer
23 architecture, queue management, reservation systems, and flow control. For example, a POSITA
24 would have been motivated to combine Mannava with Drews, which teaches a queue reservation
25 system based on a counter for multiprocessor systems. As an additional example, a POSITA
26 would have been motivated to combine Mannava with Gewirtz, which teaches a programmable

1 queue structure for multiprocessor systems with flow control based on a queue state register. As
2 an additional example, a POSITA would have been motivated to combine Mannava with Goyal,
3 which teaches queue management based on a centralized command processor for multiprocessor
4 systems. As an additional example, a POSITA would have been motivated to combine Mannava
5 with Kolinummi, which teaches a reservation system using a hardware-based semaphore for
6 multiprocessor systems. As an additional example, a POSITA would have been motivated to
7 combine Mannava with Ohkawa, which teaches queue management, queue structure, and queue
8 placement for multiprocessor systems. As an additional example, a POSITA would have been
9 motivated to combine Mannava with Rosenthal, which teaches queue management and flow
10 control based on a flow control register that indicates the amount of space available in the queue
11 for multiprocessor systems. As an additional example, a POSITA would have been motivated to
12 combine Mannava with Borkowski, which teaches the use of counting semaphores for use in
13 resource-locking including FIFOs. As an additional example, a POSITA would have been
14 motivated to combine Mannava with Hughes, which teaches a queue reservation system with a free
15 entry counter. As an additional example, a POSITA would have been motivated to combine
16 Mannava with Kim, which teaches a multiprocessing architecture that enables inter-processor
17 communication using a set of hardware semaphore registers. As an additional example, a
18 POSITA would have been motivated to combine Mannava with Tuan, which teaches a
19 multiprocessor system that uses hardware semaphores to arbitrate access to shared resources. A
20 POSITA would naturally have looked to these references and been motivated to combine them
21 with Mannava because, e.g., they provide implementation details and facilitate resource sharing,
22 optimize queue and reservation management, and reduce coordination overhead for accessing
23 shared resources in multiprocessor systems. This would have been simple combination of known
24 prior art element according to known methods to yield predictable results as well as a simple
25 substitution of known elements for another to obtain predictable results.

26 Additionally, Mannava discloses the support of I/O operations. A POSITA would
27 naturally have looked to references that support and provide features for various I/O operations,
28 including storage systems. For example, a POSITA would have been motivated to combine

1 Mannava with Kanai or Matsunami, as both teach multiprocessor-based storage systems, including
2 support for NAS and SAN. SAN and NAS had been known and used in networked settings and
3 data centers for many years to provide data storage over networks. The motivations to combine
4 explained in the incorporated Petition (IPR2025-00487) also apply to Mannava, and a POSITA
5 would have been motivated to combine for similar reasons.

- 6 • **Michaylov + Drews**
- 7 • **Michaylov + Gewirtz**
- 8 • **Michaylov + Goyal**
- 9 • **Michaylov + Kolinummi**
- 10 • **Michaylov + Ohkawa**
- 11 • **Michaylov + Rosenthal**
- 12 • **Michaylov + Borkowski**
- 13 • **Michaylov + Hughes**
- 14 • **Michaylov + Kim**
- 15 • **Michaylov + Tuan**
- 16 • **Any of the above + Kanai**
- 17 • **Any of the above + Matsunami**

18 In addition to the above-mentioned motivations to combine, Michaylov discloses a
19 messaging queue system for a multiprocessor system. *See* Abstract. Michaylov describes its
20 technical field as “managing message queues. Message queues can be used to provide
21 asynchronous communication protocol for access entities (e.g., servers, operating systems,
22 software modules, etc.) to exchange messages.” 1:5-8. Michaylov explains that the “queue
23 manager [] interacts with an input/output (I/O) manager [] that manages memory storage for a set
24 of message queues with corresponding queue data structures.” 3:24-31. Michaylov acknowledges
25 the need for flow control: “determining if a space is available in each queue for writing a
26 corresponding portion of the data, and if available, reserving the spaces in the queue.” 1:25-29.
27
28

1 Therefore, given the context provided by Michaylov, a POSITA would have naturally
2 looked to other references that provide details related to multiprocessor systems, computer
3 architecture, queue management, reservation systems, and flow control. For example, a POSITA
4 would have been motivated to combine Michaylov with Drews, which teaches a queue reservation
5 system based on a counter for multiprocessor systems. As an additional example, a POSITA
6 would have been motivated to combine Michaylov with Gewirtz, which teaches a programmable
7 queue structure for multiprocessor systems with flow control based on a queue state register. As
8 an additional example, a POSITA would have been motivated to combine Michaylov with Goyal,
9 which teaches queue management based on a centralized command processor for multiprocessor
10 systems. As an additional example, a POSITA would have been motivated to combine Michaylov
11 with Kolinummi, which teaches a reservation system using a hardware-based semaphore for
12 multiprocessor systems. As an additional example, a POSITA would have been motivated to
13 combine Michaylov with Ohkawa, which teaches queue management, queue structure, and queue
14 placement for multiprocessor systems. As an additional example, a POSITA would have been
15 motivated to combine Michaylov with Rosenthal, which teaches queue management and flow
16 control based on a flow control register that indicates the amount of space available in the queue
17 for multiprocessor systems. As an additional example, a POSITA would have been motivated to
18 combine Michaylov with Borkowski, which teaches the use of counting semaphores for use in
19 resource-locking including FIFOs. As an additional example, a POSITA would have been
20 motivated to combine Michaylov with Hughes, which teaches a queue reservation system with a
21 free entry counter. As an additional example, a POSITA would have been motivated to combine
22 Michaylov with Kim, which teaches a multiprocessing architecture that enables inter-processor
23 communication using a set of hardware semaphore registers. As an additional example, a
24 POSITA would have been motivated to combine Michaylov with Tuan, which teaches a
25 multiprocessor system that uses hardware semaphores to arbitrate access to shared resources. A
26 POSITA would naturally have looked to these references and been motivated to combine them
27 with Michaylov because, e.g., they provide implementation details and facilitate resource sharing,
28 optimize queue and reservation management, and reduce coordination overhead for accessing

1 shared resources in multiprocessor systems. This would have been simple combination of known
2 prior art element according to known methods to yield predictable results as well as a simple
3 substitution of known elements for another to obtain predictable results.

4 Additionally, Michaylov discloses the support of I/O operations. A POSITA would
5 naturally have looked to references that support and provide features for various I/O operations,
6 including storage systems. For example, a POSITA would have been motivated to combine
7 Michaylov with Kanai or Matsunami, as both teach multiprocessor-based storage systems,
8 including support for NAS and SAN. SAN and NAS had been known and used in networked
9 settings and data centers for many years to provide data storage over networks. The motivations
10 to combine explained in the incorporated Petition (IPR2025-00487) also apply to Michaylov, and
11 a POSITA would have been motivated to combine for similar reasons.

- 12 • **Van Loo + Drews**
- 13 • **Van Loo + Gewirtz**
- 14 • **Van Loo + Goyal**
- 15 • **Van Loo + Kolinummi**
- 16 • **Van Loo + Ohkawa**
- 17 • **Van Loo + Rosenthal**
- 18 • **Van Loo + Borkowski**
- 19 • **Van Loo + Hughes**
- 20 • **Van Loo + Kim**
- 21 • **Van Loo + Tuan**
- 22 • **Any of the above + Kanai**
- 23 • **Any of the above + Matsunami**
- 24
- 25

26 In addition to the above-mentioned motivations to combine, Van Loo discloses a flow
27 control method for uniprocessor and multiprocessor systems. *See* Abstract. Van Loo describes its
28 technical field as “a new method and apparatus for flow control in a packet-switched

1 microprocessor-based computer system (uniprocessor), a network of such computer systems, or a
2 multiprocessor system. Van Loo explains that its system “classifies system interconnect queues as
3 request queues, read-data queues, and write-data queues, and determines rules of transfer of both
4 requests and data as loosely coupled events.” Van Loo acknowledges the need for flow control:
5 “The elimination of (advance or overflow) feedback signals in the present flow control system
6 reduces the interface latency, since there is no extra handshake, no rescheduling or rearbitrating
7 for resources, and no retries by the master.” 3:38-41.

8 Therefore, given the context provided by Van Loo, a POSITA would have naturally looked
9 to other references that provide details related to multiprocessor systems, computer architecture,
10 queue management, reservation systems, and flow control. For example, a POSITA would have
11 been motivated to combine Van Loo with Drews, which teaches a queue reservation system based
12 on a counter for multiprocessor systems. As an additional example, a POSITA would have been
13 motivated to combine Van Loo with Gewirtz, which teaches a programmable queue structure for
14 multiprocessor systems with flow control based on a queue state register. As an additional
15 example, a POSITA would have been motivated to combine Van Loo with Goyal, which teaches
16 queue management based on a centralized command processor for multiprocessor systems. As an
17 additional example, a POSITA would have been motivated to combine Van Loo with Kolinummi,
18 which teaches a reservation system using a hardware-based semaphore for multiprocessor systems.
19 As an additional example, a POSITA would have been motivated to combine Van Loo with
20 Ohkawa, which teaches queue management, queue structure, and queue placement for
21 multiprocessor systems. As an additional example, a POSITA would have been motivated to
22 combine Van Loo with Rosenthal, which teaches queue management and flow control based on a
23 flow control register that indicates the amount of space available in the queue for multiprocessor
24 systems. As an additional example, a POSITA would have been motivated to combine Van Loo
25 with Borkowski, which teaches the use of counting semaphores for use in resource-locking
26 including FIFOs. As an additional example, a POSITA would have been motivated to combine Van
27 Loo with Hughes, which teaches a queue reservation system with a free entry counter. As an
28 additional example, a POSITA would have been motivated to combine Van Loo with Kim, which

1 teaches a multiprocessing architecture that enables inter-processor communication using a set of
2 hardware semaphore registers. As an additional example, a POSITA would have been motivated
3 to combine Van Loo with Tuan, which teaches a multiprocessor system that uses hardware
4 semaphores to arbitrate access to shared resources. A POSITA would naturally have looked to
5 these references and been motivated to combine them with Van Loo because, e.g., they provide
6 implementation details and facilitate resource sharing, optimize queue and reservation
7 management, and reduce coordination overhead for accessing shared resources in multiprocessor
8 systems. This would have been simple combination of known prior art element according to
9 known methods to yield predictable results as well as a simple substitution of known elements for
10 another to obtain predictable results.

11 Additionally, Van Loo discloses the support of I/O operations. A POSITA would naturally
12 have looked to references that support and provide features for various I/O operations, including
13 storage systems. For example, a POSITA would have been motivated to combine Van Loo with
14 Kanai or Matsunami, as both teach multiprocessor-based storage systems, including support for
15 NAS and SAN. SAN and NAS had been known and used in networked settings and data centers
16 for many years to provide data storage over networks. The motivations to combine explained in
17 the incorporated Petition (IPR2025-00487) also apply to Van Loo, and a POSITA would have
18 been motivated to combine for similar reasons.

- 19 • **ARM System(s) + Drews**
- 20 • **ARM System(s) + Gewirtz**
- 21 • **ARM System(s) + Goyal**
- 22 • **ARM System(s) + Kolinummi**
- 23 • **ARM System(s) + Ohkawa**
- 24 • **ARM System(s) + Rosenthal**
- 25 • **ARM System(s) + Borkowski**
- 26 • **ARM System(s) + Borkowski**
- 27 • **ARM System(s) + Hughes**
- 28

- 1 • **ARM System(s) + Kim**
- 2 • **ARM System(s) + Tuan**
- 3 • **Any of the above + Kanai**
- 4 • **Any of the above + Matsunami**

5 In addition to the above-mentioned motivations to combine, the ARM System(s) disclose
6 message queues and other queues for buffering data. *See* Exhibit 596-18. For example, on
7 information and belief, the ARM System(s) contain queues and support multiprocessors and
8 semaphores through, for example, the use of “load register exclusive” (LDREX) and “store
9 register exclusive” (STREX) instructions. *See id.* ARM microprocessors can be notified of queue
10 status. *See id.*

11 Therefore, given the context provided by ARM System(s), a POSITA would have naturally
12 looked to other references that provide details related to multiprocessor systems, computer
13 architecture, queue management, reservation systems, and flow control. For example, a POSITA
14 would have been motivated to combine ARM System(s) with Drews, which teaches a queue
15 reservation system based on a counter for multiprocessor systems. As an additional example, a
16 POSITA would have been motivated to combine ARM System(s) with Gewirtz, which teaches a
17 programmable queue structure for multiprocessor systems with flow control based on a queue
18 state register. As an additional example, a POSITA would have been motivated to combine ARM
19 System(s) with Goyal, which teaches queue management based on a centralized command
20 processor for multiprocessor systems. As an additional example, a POSITA would have been
21 motivated to combine ARM System(s) with Kolinummi, which teaches a reservation system using
22 a hardware-based semaphore for multiprocessor systems. As an additional example, a POSITA
23 would have been motivated to combine ARM System(s) with Ohkawa, which teaches queue
24 management, queue structure, and queue placement for multiprocessor systems. As an additional
25 example, a POSITA would have been motivated to combine ARM System(s) with Rosenthal,
26 which teaches queue management and flow control based on a flow control register that indicates
27 the amount of space available in the queue for multiprocessor systems. As an additional example,
28

1 a POSITA would have been motivated to combine ARM System(s) with Borkowski, which
2 teaches the use of counting semaphores for use in resource-locking including FIFOs. As an
3 additional example, a POSITA would have been motivated to combine ARM System(s) with
4 Hughes, which teaches a queue reservation system with a free entry counter. As an additional
5 example, a POSITA would have been motivated to combine ARM System(s) with Kim, which
6 teaches a multiprocessing architecture that enables inter-processor communication using a set of
7 hardware semaphore registers. As an additional example, a POSITA would have been motivated
8 to combine ARM System(s) with Tuan, which teaches a multiprocessor system that uses hardware
9 semaphores to arbitrate access to shared resources. A POSITA would naturally have looked to
10 these references and been motivated to combine them with ARM System(s) because, e.g., they
11 provide implementation details and facilitate resource sharing, optimize queue and reservation
12 management, and reduce coordination overhead for accessing shared resources in multiprocessor
13 systems. This would have been simple combination of known prior art element according to
14 known methods to yield predictable results as well as a simple substitution of known elements for
15 another to obtain predictable results.

16 Additionally, ARM System(s) disclose the support of I/O operations. A POSITA would
17 naturally have looked to references that support and provide features for various I/O operations,
18 including storage systems. For example, a POSITA would have been motivated to combine ARM
19 System(s) with Kanai or Matsunami, as both teach multiprocessor-based storage systems,
20 including support for NAS and SAN. SAN and NAS had been known and used in networked
21 settings and data centers for many years to provide data storage over networks. The motivations
22 to combine explained in the incorporated Petition (IPR2025-00487) also apply to the ARM
23 System(s), and a POSITA would have been motivated to combine for similar reasons.

- 24 • **Intel/Fulcrum System(s) + Drews**
- 25 • **Intel/Fulcrum System(s) + Gewirtz**
- 26 • **Intel/Fulcrum System(s) + Goyal**
- 27 • **Intel/Fulcrum System(s) + Kolinummi**
- 28

- 1 • **Intel/Fulcrum System(s) + Ohkawa**
- 2 • **Intel/Fulcrum System(s) + Rosenthal**
- 3 • **Intel/Fulcrum System(s) + Borkowski**
- 4 • **Intel/Fulcrum System(s) + Hughes**
- 5 • **Intel/Fulcrum System(s) + Kim**
- 6 • **Intel/Fulcrum System(s) + Tuan**
- 7 • **Any of the above + Kanai**
- 8 • **Any of the above + Matsunami**

9
10 In addition to the above-mentioned motivations to combine, the Intel/Fulcrum System(s)
11 disclose receive and transmit queues for scheduling packets and congestion management based on
12 status of the queues. *See* Exhibit 596-19. For example, the Intel/Fulcrum System(s) include a
13 switch with a shared memory that is accessible to multiple ports, and it discloses maintaining
14 receive and transmit queues for scheduling packets for transmission. *See id.* The Intel/Fulcrum
15 System(s) include congestion management that maintains information about the sizes of different
16 queues and determine if congestion occurs based on upper-bound limits on the queue sizes. *See*
17 *id.* When the status of a queue is changed, notifications regarding the status of the queue is sent to
18 one or multiple devices. *See id.*

19 Therefore, given the context provided by Intel/Fulcrum System(s), a POSITA would have
20 naturally looked to other references that provide details related to multiprocessor systems,
21 computer architecture, queue management, reservation systems, and flow control. For example, a
22 POSITA would have been motivated to combine Intel/Fulcrum System(s) with Drews, which
23 teaches a queue reservation system based on a counter for multiprocessor systems. As an
24 additional example, a POSITA would have been motivated to combine Intel/Fulcrum System(s)
25 with Gewirtz, which teaches a programmable queue structure for multiprocessor systems with
26 flow control based on a queue state register. As an additional example, a POSITA would have
27 been motivated to combine Intel/Fulcrum System(s) with Goyal, which teaches queue
28

1 management based on a centralized command processor for multiprocessor systems. As an
2 additional example, a POSITA would have been motivated to combine Intel/Fulcrum System(s)
3 with Kolinummi, which teaches a reservation system using a hardware-based semaphore for
4 multiprocessor systems. As an additional example, a POSITA would have been motivated to
5 combine Intel/Fulcrum System(s) with Ohkawa, which teaches queue management, queue
6 structure, and queue placement for multiprocessor systems. As an additional example, a POSITA
7 would have been motivated to combine Intel/Fulcrum System(s) with Rosenthal, which teaches
8 queue management and flow control based on a flow control register that indicates the amount of
9 space available in the queue for multiprocessor systems. As an additional example, a POSITA
10 would have been motivated to combine Intel/Fulcrum System(s) with Borkowski, which teaches
11 the use of counting semaphores for use in resource-locking including FIFOs. As an additional
12 example, a POSITA would have been motivated to combine Intel/Fulcrum System(s) with Hughes,
13 which teaches a queue reservation system with a free entry counter. As an additional example, a
14 POSITA would have been motivated to combine Intel/Fulcrum System(s) with Kim, which
15 teaches a multiprocessing architecture that enables inter-processor communication using a set of
16 hardware semaphore registers. As an additional example, a POSITA would have been motivated
17 to combine Intel/Fulcrum System(s) with Tuan, which teaches a multiprocessor system that uses
18 hardware semaphores to arbitrate access to shared resources. A POSITA would naturally have
19 looked to these references and been motivated to combine them with Intel/Fulcrum System(s)
20 because, e.g., they provide implementation details and facilitate resource sharing, optimize queue
21 and reservation management, and reduce coordination overhead for accessing shared resources in
22 multiprocessor systems. This would have been simple combination of known prior art element
23 according to known methods to yield predictable results as well as a simple substitution of known
24 elements for another to obtain predictable results.

25 Additionally, Intel/Fulcrum System(s) disclose the support of I/O operations. A POSITA
26 would naturally have looked to references that support and provide features for various I/O
27 operations, including storage systems. For example, a POSITA would have been motivated to
28 combine Intel/Fulcrum System(s) with Kanai or Matsunami, as both teach multiprocessor-based

1 storage systems, including support for NAS and SAN. SAN and NAS had been known and used
2 in networked settings and data centers for years to provide data storage over networks. The
3 motivations to combine explained in the incorporated Petition (IPR2025-00487) also apply to the
4 Intel/Fulcrum System(s), and a POSITA would have been motivated to combine for similar
5 reasons.

- 6 • **Nvidia/Mellanox System(s) + Drews**
- 7 • **Nvidia/Mellanox System(s) + Gewirtz**
- 8 • **Nvidia/Mellanox System(s) + Goyal**
- 9 • **Nvidia/Mellanox System(s) + Kolinummi**
- 10 • **Nvidia/Mellanox System(s) + Ohkawa**
- 11 • **Nvidia/Mellanox System(s) + Rosenthal**
- 12 • **Nvidia/Mellanox System(s) + Borkowski**
- 13 • **Nvidia/Mellanox System(s) + Hughes**
- 14 • **Nvidia/Mellanox System(s) + Kim**
- 15 • **Nvidia/Mellanox System(s) + Tuan**
- 16 • **Any of the above + Kanai**
- 17 • **Any of the above + Matsunami**

18 In addition to the above-mentioned motivations to combine, Nvidia/Mellanox System(s)
19 disclose maintaining queues and keeping track of the status of queues in the context of
20 networking. *See* Exhibit 596-20. For example, Nvidia/Mellanox System(s) disclose a network
21 architecture for communication among network adapters and I/O units attached to a switched
22 fabric, where remote direct memory access is supported. *See id.* Queue pairs, including a send
23 queue and a receive queue, are created and used for transferring data between nodes of a network.
24 *See id.* The receive queue may be shared among different nodes. *See id.* The Nvidia/Mellanox
25 System(s) further teach using a completion queue to track the status of send and receive queues
26 and using a completion channel to notify nodes as to such status. *See id.* The status of a queue
27
28

1 pair includes, for example, whether the receive queue is ready to receive, based on which the
2 network of nodes operate. *See id.*

3 Therefore, given the context provided by Nvidia/Mellanox System(s), a POSITA would
4 have naturally looked to other references that provide details related to multiprocessor systems,
5 computer architecture, queue management, reservation systems, and flow control. For example, a
6 POSITA would have been motivated to combine Nvidia/Mellanox System(s) with Drews, which
7 teaches a queue reservation system based on a counter for multiprocessor systems. As an
8 additional example, a POSITA would have been motivated to combine Nvidia/Mellanox
9 System(s) with Gewirtz, which teaches a programmable queue structure for multiprocessor
10 systems with flow control based on a queue state register. As an additional example, a POSITA
11 would have been motivated to combine Nvidia/Mellanox System(s) with Goyal, which teaches
12 queue management based on a centralized command processor for multiprocessor systems. As an
13 additional example, a POSITA would have been motivated to combine Nvidia/Mellanox
14 System(s) with Kolinummi, which teaches a reservation system using a hardware-based
15 semaphore for multiprocessor systems. As an additional example, a POSITA would have been
16 motivated to combine Nvidia/Mellanox System(s) with Ohkawa, which teaches queue
17 management, queue structure, and queue placement for multiprocessor systems. As an additional
18 example, a POSITA would have been motivated to combine Nvidia/Mellanox System(s) with
19 Rosenthal, which teaches queue management and flow control based on a flow control register
20 that indicates the amount of space available in the queue for multiprocessor systems. As an
21 additional example, a POSITA would have been motivated to combine Nvidia/Mellanox
22 System(s) with Borkowski, which teaches the use of counting semaphores for use in resource-
23 locking including FIFOs. As an additional example, a POSITA would have been motivated to
24 combine Nvidia/Mellanox System(s) with Hughes, which teaches a queue reservation system with
25 a free entry counter. As an additional example, a POSITA would have been motivated to combine
26 Nvidia/Mellanox System(s), which teaches a multiprocessing architecture that enables inter-
27 processor communication using a set of hardware semaphore registers. As an additional example,
28 a POSITA would have been motivated to combine Nvidia/Mellanox System(s) with Tuan, which

1 teaches a multiprocessor system that uses hardware semaphores to arbitrate access to shared
2 resources. A POSITA would naturally have looked to these references and been motivated to
3 combine them with Nvidia/Mellanox System(s) because, e.g., they provide implementation details
4 and facilitate resource sharing, optimize queue and reservation management, and reduce
5 coordination overhead for accessing shared resources in multiprocessor systems. This would have
6 been simple combination of known prior art element according to known methods to yield
7 predictable results as well as a simple substitution of known elements for another to obtain
8 predictable results.

9 Additionally, Nvidia/Mellanox System(s) disclose the support of I/O operations. A
10 POSITA would naturally have looked to references that support and provide features for various
11 I/O operations, including storage systems. For example, a POSITA would have been motivated to
12 combine Nvidia/Mellanox System(s) with Kanai or Matsunami, as both teach multiprocessor-
13 based storage systems, including support for NAS and SAN. SAN and NAS had been known and
14 used in networked settings and data centers for many years to provide data storage over networks.
15 The motivations to combine explained in the incorporated Petition (IPR2025-00487) also apply to
16 the Nvidia/Mellanox System(s), and a POSITA would have been motivated to combine for similar
17 reasons.

- 18 • **Hughes + Drews**
- 19 • **Hughes + Gewirtz**
- 20 • **Hughes + Goyal**
- 21 • **Hughes + Kolinummi**
- 22 • **Hughes + Ohkawa**
- 23 • **Hughes + Rosenthal**
- 24 • **Hughes + Borkowski**
- 25 • **Hughes + Hughes**
- 26 • **Hughes + Kim**
- 27 • **Hughes + Tuan**
- 28 • **Hughes + Kanai**

1 • **Hughes + Matsunami**

2 In addition to the above-mentioned motivations to combine, Hughes discloses a
3 multiprocessor computing system in which “multiple sources of transactions” “share a common
4 queuing structure which services those transactions,” including shared request queues used to
5 service requests to other processing nodes and system memory. Background; 1:43-50. Hughes
6 further teaches that the request queues are “responsible for accepting an . . . request . . . performing
7 actions necessary for completing the request, and returning appropriate responses back to the
8 requestors as necessary.” Background. Hughes explains that queue “complexity” can make “each
9 queue entry” consume “significant area,” so “limiting the total number of queue entries may be
10 required,” creating a need for “an efficient method of allocating and managing the available queue
11 resources.” Hughes, Background. To address this, Hughes teaches a request queue management
12 technique that keeps track of the number of free queue entries for each requestor processor and, at
13 initialization, “allots to each requestor” predetermined “hard” and “free” entries, with remaining
14 “un-allotted entries” forming “a free pool of entries.” Hughes, Abstract; Summary.

15 Therefore, given the context provided by Hughes, a POSITA would have naturally looked
16 to other references that provide details related to multiprocessor systems, computer architecture,
17 queue management, reservation systems, and flow control. For example, a POSITA would have
18 been motivated to combine Hughes with Drews, which teaches a queue reservation system based
19 on a counter for multiprocessor systems. As an additional example, a POSITA would have been
20 motivated to combine Hughes with Gewirtz, which teaches a programmable queue structure for
21 multiprocessor systems with flow control based on a queue state register. As an additional
22 example, a POSITA would have been motivated to combine Hughes with Goyal, which teaches
23 queue management based on a centralized command processor for multiprocessor systems. As an
24 additional example, a POSITA would have been motivated to combine Hughes with Kolinummi,
25 which teaches a reservation system using a hardware-based semaphore for multiprocessor systems.
26 As an additional example, a POSITA would have been motivated to combine Hughes with
27 Ohkawa, which teaches queue management, queue structure, and queue placement for
28 multiprocessor systems. As an additional example, a POSITA would have been motivated to

1 combine Hughes with Rosenthal, which teaches queue management and flow control based on a
2 flow control register that indicates the amount of space available in the queue for multiprocessor
3 systems. As an additional example, a POSITA would have been motivated to combine Hughes
4 with Borkowski, which teaches the use of counting semaphores for use in resource-locking
5 including FIFOs. As an additional example, a POSITA would have been motivated to combine
6 Hughes with Kim, which teaches a multiprocessing architecture that enables inter-processor
7 communication using a set of hardware semaphore registers. As an additional example, a
8 POSITA would have been motivated to combine Hughes with Tuan, which teaches a
9 multiprocessor system that uses hardware semaphores to arbitrate access to shared resources. A
10 POSITA would naturally have looked to these references and been motivated to combine them
11 with Hughes because, e.g., they provide implementation details and facilitate resource sharing,
12 optimize queue and reservation management, and reduce coordination overhead for accessing
13 shared resources in multiprocessor systems. This would have been simple combination of known
14 prior art element according to known methods to yield predictable results as well as a simple
15 substitution of known elements for another to obtain predictable results.

16 Additionally, Hughes discloses support for I/O operations interfacing with external devices
17 such as non-volatile memory devices. 1:10–22; 2:37–63. Therefore, a POSITA would naturally
18 have looked to references that support and provide features for various I/O operations, including
19 known storage systems. For example, a POSITA would have been motivated to combine Hughes
20 with Kanai or Matsunami, as both teach multiprocessor-based storage systems, including support
21 for NAS and SAN. SAN and NAS had been known and used in networked settings and data
22 centers for many years to provide data storage over networks. The motivations to combine
23 explained in the incorporated Petition (IPR2025-00487) also apply to Hughes, and a POSITA
24 would have been motivated to combine for similar reasons.

- 25 • **IBM Cell Broadband System(s) + Alexander**
- 26 • **IBM Cell Broadband System(s) + Buchko**
- 27 • **IBM Cell Broadband System(s) + Dongare**
- 28 • **IBM Cell Broadband System(s) + Drews**

- 1 • **IBM Cell Broadband System(s) + Gewirtz**
- 2 • **IBM Cell Broadband System(s) + Goyal**
- 3 • **IBM Cell Broadband System(s) + Hammer**
- 4 • **IBM Cell Broadband System(s) + Khawand**
- 5 • **IBM Cell Broadband System(s) + Kolinummi**
- 6 • **IBM Cell Broadband System(s) + Michaylov**
- 7 • **IBM Cell Broadband System(s) + Ohkawa**
- 8 • **IBM Cell Broadband System(s) + Van Loo**
- 9 • **IBM Cell Broadband System(s) + Borkowski**
- 10 • **IBM Cell Broadband System(s) + Hughes**
- 11 • **IBM Cell Broadband System(s) + Kim**
- 12 • **IBM Cell Broadband System(s) + Tuan**
- 13 • **Any of the above + Kanai**
- 14 • **Any of the above + Matsunami**

15 IBM Cell Broadband System(s) teaches a multiprocessor architecture with Power
16 Processing Engines (“PPEs”) and Synergistic Processing Units (“SPUs”) that communicate with
17 each other using hardware queues managed by a Memory Flow Control Unit (“MFCU”). As
18 explained in the accompanying claim chart, IBM Cell Broadband System(s) teaches queue
19 management, and there were teachings, suggestions, and motivations in the art for IBM Cell
20 Broadband System(s) for implementing queue management functionality, *e.g.*, for the MFCU,
21 using computer instructions running on a processor. *See, e.g.*, Day at PA043552-PA043628 at
22 11:25-42. A POSITA would therefore have been motivated to combine IBM Cell Broadband
23 System(s) with known queue management techniques that use software instructions on processors,
24 particularly prior art for multiprocessor systems and flow control. For example, a POSITA would
25 have been motivated to combine IBM Cell Broadband System(s) with Michaylov, which teaches
26 techniques for using software instructions, running on a processor, for queue management.
27 Michaylov, 13:9-44 (“The queue management functions described herein can be implemented
28 using software for execution on a computer.”). As other examples, a POSITA would have been

1 motivated to combine IBM Cell Broadband System(s) with Alexander, Buchko, Dongare, Drews,
2 Gewirtz, Goyal, Hammer, Khawand, Kolinummi, Ohkawa, Van Loo, Borkowski, Hughes, Kim,
3 and/or Tuan, which likewise taught using software instructions, running on a processor, for queue
4 management. A POSITA would naturally have looked to these references and been motivated to
5 combine them with IBM Cell Broadband System(s), for numerous reasons, to implement queue
6 management functionality, e.g., for the MFCU, using computer instructions running on a
7 processor. First, the Cell MFCU, Michaylov, Alexander, Buchko, Dongare, Drews, Gewirtz,
8 Goyal, Hammer, Khawand, Kolinummi, Ohkawa, Van Loo, Borkowski, Hughes, Kim, and Tuan
9 all manage queues in a multiprocessor system and/or for staging requests for memory operation.
10 Therefore, a POSITA looking at the operation of queue management in the Cell architecture,
11 including MFCU operation, would have been motivated to combine similar solutions from other
12 multiprocessor systems, such as Michaylov and the other listed references, for the same function
13 taught by each reference. Second, the Cell MFCU was highly configurable, allowing flexible
14 configuration of its operations. Therefore, a POSITA would have been motivated to increase
15 configurability and flexibility using software running on a processor to implement MFCU
16 functionality in IBM Cell Broadband System(s). Michaylov, Alexander, Buchko, Dongare,
17 Drews, Gewirtz, Goyal, Hammer, Khawand, Kolinummi, Ohkawa, Van Loo, Borkowski, Hughes,
18 Kim, and Tuan teach similar queue management functionality implemented using software
19 running on a processor. Therefore, a POSITA would have been motivated to combine these
20 references with IBM Cell Broadband System(s). Third, a POSITA would have been motivated by
21 several advantages of the above-described combinations, including a simplified implementation
22 for providing increased configurability, flexibility, and user-customizability. This would have
23 been simple combination of known prior art element according to known methods to yield
24 predictable results as well as a simple substitution of known elements for another to obtain
25 predictable results.

26 IBM Cell Broadband System(s) include a flow control register (e.g., queue status register)
27 that indicates the available space in a queue, and it teaches checking this register before writing to
28 its corresponding queue, to ensure there is sufficient space. After a write, the status register would

1 need to be updated to reflect the new amount of available space remaining. Therefore, a POSITA
2 would have been motivated to combine IBM Cell Broadband System(s) with known techniques
3 for flow control in multiprocessor systems, for updating status registers. Tuan teaches this by
4 decrementing the value stored at a flow control register before returning the result of a read to that
5 register. A POSITA would have been motivated to combine these teachings because they address
6 similar flow-control issues in multiprocessor systems, and the combination would have been a
7 straightforward application and use of a known technique to a known device ready for
8 improvement to yield predictable results. Applying these techniques also would have been
9 obvious to try, with a reasonable expectation of success, because there were only a finite number
10 of identified, predictable, and known solutions with respect to ordering the steps of decrementing
11 and returning. Indeed, a POSITA would have been motivated because the combination was a
12 simple way to update the flow control register, without separate synchronization or additional
13 write operations.

14 Additionally, IBM Cell Broadband System(s) disclose the support of I/O operations, with
15 various I/O ports for a variety of external devices. A POSITA would naturally have looked to
16 references that support and provide features for various I/O operations, including storage systems.
17 For example, a POSITA would have been motivated to combine IBM Cell Broadband System(s)
18 with Kanai or Matsunami, as both teach multiprocessor-based storage systems, including support
19 for NAS and SAN. SAN and NAS had been known and used in networked settings and data
20 centers for many years to provide data storage over networks. Moreover, IBM Cell Broadband
21 System(s) were designed to be used in a wide variety of applications, including for example in
22 servers, high-performance computing systems, and even video game consoles. Its architecture
23 was designed to process streams of data. Therefore, a POSITA would have been motivated to use
24 IBM Cell Broadband System(s) with SAN and NAS because both operate on streams of data in a
25 network environment. This would have been simple combination of known prior art element
26 according to known methods to yield predictable results as well as a simple substitution of known
27 elements for another to obtain predictable results. The motivations to combine explained in the
28

1 incorporated Petition (IPR2025-00487) also apply to the IBM Cell Broadband System(s) and a
2 POSITA would have been motivated to combine for similar reasons.

3 Defendants further incorporate by reference the invalidity grounds described in
4 Defendants' Petition for Inter Partes Review (assigned Case No. IPR2025-00487).

5 **B. U.S. Patent No. 9,529,767**

6 The following table identifies the prior art for the '767 patent. Each item below constitutes
7 prior art under one or more provisions of 35 U.S.C. §§ 102(a)(1)-(2) by reason of its status as a
8 publication and/or patent disclosing, or a system on sale and/or in public knowledge or use in the
9 United States embodying the claimed invention before Plaintiffs' asserted priority date of the '767
10 patent. For each item of prior art, Defendants have attached a claim chart identifying where each
11 claim limitation is found in the cited prior art.

12 **Table B: Prior Art to the '767 Patent**

13 Prior Art to '767	14 Anticipation/ Obviousness	Claim Chart
15 EMC VPLEX System(s), including but not 16 limited to EMC VPLEX VS1, EMC VPLEX 17 VS2, EMC VPLEX Local, EMC VPLEX Metro, 18 EMC VPLEX Geo, were sold, offered for sale, 19 made known, and/or publicly used before 20 December 6, 2013. EMC VPLEX System(s) 21 were sold, offered for sale, and/or made known to 22 the public and third-party customers before this date, by December 5, 2013 at the latest. EMC VPLEX System(s) were publicly used by Dell/EMC, the public, and third-party customers before this date, by December 5, 2013 at the latest.	A & O	Exhibit 767-01
23 IBM SAN Volume Controller System(s), 24 including but not limited to IBM SAN Volume Controller 2145-4F2, IBM SAN Volume 25 Controller 2145-CF8, IBM SAN Volume Controller 2145-8F2, IBM SAN Volume 26 Controller 2145-8F4, IBM SAN Volume Controller 2145-8G4, were sold, offered for sale, 27 made known, and/or publicly used before 28 December 6, 2013. IBM SAN Volume Controller System(s) were sold, offered for sale, and/or	A & O	Exhibit 767-02

Prior Art to '767	Anticipation/ Obviousness	Claim Chart
made known to the public and third-party customers before this date, by December 5, 2013 at the latest. IBM SAN Volume Controller System(s) were publicly used by IBM, the public, and third-party customers before this date, by December 5, 2013 at the latest		
Microsoft Server System(s), including but not limited to Microsoft Server 2012, Microsoft Server 2012 R2, Microsoft HPC Server 2008, Microsoft HPC Server 2008 R2, were sold, offered for sale, made known, and/or publicly used before December 6, 2013. Microsoft Server System(s) were sold, offered for sale, and/or made known to the public and third-party customers before this date, by December 5, 2013 at the latest. Microsoft Server System(s) were publicly used by Microsoft, the public, and third-party customers before this date, by December 5, 2013 at the latest	A & O	Exhibit 767-03
U.S. Patent Application Publication No. 2007/0050538 to Northcutt et al. ("Northcutt"); filed on August 24, 2006; published on March 1, 2007.	A & O	Exhibit 767-04
U.S. Patent No. 8,966,172 to Malwankar ("Malwankar"); filed on January 23, 2012; issued on February 24, 2015.	A & O	Exhibit 767-05
U.S. Patent Application Publication No. 2013/0073747 to Klughart ("Klughart"); filed on September 21, 2011; published on March 21, 2013.	O	Exhibit 767-06
U.S. Patent No. 9,129,068 to Yendigiri et al. ("Yendigiri"); filed on January 4, 2013; issued on September 8, 2015.	O	Exhibit 767-07
U.S. Patent No. 8,321,650 to Steinmetz et al. ("Steinmetz"); filed on November 26, 2007; issued on November 27, 2012.	O	Exhibit 767-08
U.S. Patent Application Publication No. 2015/0222705 to Stephens ("Stephens"); provisional application filed on September 6, 2012; published on August 6, 2015.	O	Exhibit 767-09

Prior Art to '767	Anticipation/ Obviousness	Claim Chart
U.S. Patent No. 9,424,224 to Voorhees et al. ("Voorhees '224"); filed on June 27, 2013; issued on August 23, 2016.	O	Exhibit 767-10
U.S. Patent Application Publication No. 2012/0124256 to Johnson et al. ("Johnson"); filed on November 15, 2010; published on May 17, 2012.	O	Exhibit 767-11
U.S. Patent No. 8,539,096 to Voorhees et al. ("Voorhees '096"); filed on September 26, 2003; issued on September 17, 2013.	O	Exhibit 767-12
U.S. Patent No. 6,996,741 to Pittelkow et al. ("Pittelkow '741"); filed on June 28, 2002; issued on December 27, 2006	A & O	Exhibit 767-13

1. Anticipation

For each item of prior art identified above as an anticipation reference, Defendants contend that the prior art item anticipates the Asserted Claims of the '767 patent, including for the reasons illustrated in the accompanying claim chart and based on Plaintiffs' infringement allegations and contentions for the '767 patent.

2. Obviousness

A person of ordinary skill in the art would have been motivated to combine any reference in Table B with any other reference(s) in Table B. To the extent that any prior art identified in Table B above is deemed to not anticipate an Asserted Claim, the prior art in Table B may be combined with one or more prior art references in Table B to render the claim obvious. All of these items of prior art are directed to data storage device management and propose solutions similar to those discussed in the '767 patent.

The following are exemplary combinations of prior art that would have been obvious to a POSITA:

- **Northcutt + Malwankar**
- **Northcutt + Pittelkow**

1 • **Northcutt + Steinmetz**

2 • **Northcutt + Stephens**

3 Northcutt discloses a scalable storage system that employs an “advanced storage system”
4 to represent physical storage devices as virtual drives. *See* Exhibit 767-04. The advanced storage
5 system acts as an intermediary between the host and the storage devices, connected in a cascaded
6 topology. *See* Northcutt, [0021], [0023], [0030]. A POSITA would have recognized that enabling
7 simultaneous bidirectional data transfer within such a cascaded storage system topology would be
8 desirable and beneficial, as it improves communication efficiency, load balancing, and fault
9 tolerance. Therefore, a POSITA would naturally have looked to other references that provide
10 details regarding full-duplex communication protocols for storage systems to achieve
11 simultaneous bidirectional data transfer. For example, a POSITA would have been motivated to
12 combine Northcutt with Malwankar (*See* Malwankar, 11:60-12:3, 14:38-49, 15:52-61, 15:62-
13 16:15) Pittelkow (*see* Pittelkow, 20:11-26, 12:56-13:4, 13:6-35), Stephens (*see* Stephens, [0100],
14 [0116], [0121], [0129]) or Steinmetz (*see* Steinmetz, 22:26-60, 23:23-32, 54:21-41), each of which
15 discloses full-duplex communication protocols for storage systems, such as Fibre Channel.

16 • **Northcutt + Malwankar + Yendigiri**

17 • **Northcutt + Malwankar + Voorhees ’224**

18 • **Northcutt + Stephens + Yendigiri**

19 • **Northcutt + Stephens + Voorhees ’224**

20 • **Northcutt + Pittelkow + Yendigiri**

21 • **Northcutt + Pittelkow + Voorhees ’224**

22 The combination of Northcutt and Malwankar discloses a data storage system that supports
23 various storage devices, including SATA devices. A POSITA would naturally have looked to
24 other references that provide details for other types of storage devices, particularly SAS devices,
25 as SAS interfaces support both SAS and SATA devices, allowing seamless integration of SATA
26 storage devices already supported by Northcutt. For example, a POSITA would have been
27
28

1 motivated to further combine Northcutt and Malwankar with Yendigiri (*see* Yendigiri, 2:19-33,
2 3:39-44, 4:41-43, 6:24-32) or Voorhees '224 (*see* Voorhees '224, 2:50-64, 8:63-9:11), each of
3 which is directed to SAS devices.

4 The combination of Northcutt and Stephens similarly discloses a data storage system that
5 supports various storage devices, including SATA devices. Therefore, for the same reason as
6 above, a POSITA would have been motivated to further combine Northcutt and Stephens with
7 Yendigiri or Voorhees '224.

8 The combination of Northcutt and Pittelkow similarly discloses a data storage system that
9 supports various storage devices, including SATA devices and legacy SCSI hard drives.
10 Therefore, for the same reason as above, a POSITA would have been motivated to further
11 combine Northcutt and Pittelkow with Yendigiri or Voorhees '224.

- 12 • **Northcutt + Steinmetz + Johnson + Voorhees '096**
- 13 • **Northcutt + Malwankar + Yendigiri + Johnson + Voorhees '096**
- 14 • **Northcutt + Malwankar + Voorhees '224 + Johnson + Voorhees '096**
- 15 • **Northcutt + Stephens + Yendigiri + Johnson + Voorhees '096**
- 16 • **Northcutt + Stephens + Voorhees '224 + Johnson + Voorhees '096**
- 17 • **Northcutt + Pittelkow + Yendigiri + Johnson + Voorhees '096**
- 18 • **Northcutt + Pittelkow + Voorhees '224 + Johnson + Voorhees '096**
- 19 • **Northcutt + Pittelkow + Voorhees '224 + Johnson + Voorhees '096**

20 The combination of Northcutt and Steinmetz provides a scalable data storage system that
21 supports SAS devices. A POSITA would naturally have looked to other references that provide
22 details for SAS device management to enhance the reliability and robustness of the system,
23 including details for device discovery in SAS topology. For example, both Voorhees '096 and
24 Johnson are directed to device discovery in SAS topology. Voorhees '096 discloses updating
25 routing tables in response to changes detected during a device discovery. *See* Voorhees '096, Fig.
26 2, 3:4-46, 6:61-7:6, 8:4-16, 8:17-30. Johnson discloses performing a device discovery upon
27 detecting a status change in the SAS topology and identifies events that could trigger such
28 changes. *See* Johnson, Fig. 4, [0017], [0022], [0023]. Therefore, a POSITA would have been

1 motivated to further combine Northcutt and Steinmetz with Voorhees '096 and Johnson to achieve
2 a robust and reliable storage system. For the same reason, a POSITA would have been motivated
3 to further combine Northcutt, Malwankar and Yendigiri with Voorhees '096 and Johnson, further
4 combine Northcutt, Malwankar and Voorhees '224 with Voorhees '096 and Johnson, further
5 combine Northcutt, Stephens and Yendigiri with Voorhees '096 and Johnson, and further combine
6 Northcutt, Stephens and Voorhees '224 with Voorhees '096 and Johnson.

7 A POSITA would have been motivated to further combine Northcutt, Pittelkow and
8 Yendigiri with Voorhees '096 and Johnson, because of same reason explained above.
9 Furthermore, because the combination of Northcutt, Pittelkow, and Yendigiri already establishes a
10 functional framework for device discovery within a SAN architecture and explicitly supports SAS
11 devices, it would have been a routine and obvious to incorporate the specific SAS discovery
12 mechanisms taught by Voorhees '096 and Johnson. Combining these references effectively
13 “connects the dots” by implementing SAS-specific discovery within a SAN environment that is
14 already configured for automated device discovery. For the same reasons, a POSITA would have
15 been motivated to further combine Northcutt, Pittelkow and Voorhees '224 with Voorhees '096
16 and Johnson.

- 17 • **Malwankar + Klughart**
- 18 • **Malwankar + Pittelkow**

19 Malwankar describes its technical field as “data storage” and discloses a processor-
20 agnostic PCIe-based data storage system that routes storage requests directly to storage devices.
21 *See* Exhibit 767-05. It teaches a plurality of storage controllers and assigning storage devices to
22 one or more controllers. *See* Malwankar, 1:13-47, 6:54-62. A POSITA would have found it
23 obvious that not all storage devices would be associated with each storage controller, and that in
24 the event that a storage controller receives a host command that directs to an unassociated storage
25 device, it would be desirable to communicate with the controller associated with the target storage
26 device and forward the command. Therefore, a POSITA would naturally have looked to other
27 references that provide details for inter-controller communication. For example, a POSITA would
28 have been motivated to combine Malwankar with Klughart, which has the same technical field as

1 “computer data storage” and discloses methods for communication and command transfer among
2 daisy-chained storage controllers. *See* Klughart, [0055], [0105], [0107], [0164]. Similarly, a
3 POSITA would have been motivated to combine Malwankar with Pittelkow, which also resides in
4 the field of storage area networks and teaches methods for inter-controller communication,
5 including over the Ethernet channel. *See* Pittelkow, 4:9-31, 12:56-13:4, 13:6-35, 20:11-26.

- 6 • **Malwankar + Yendigiri**
- 7 • **Malwankar + Steinmetz**
- 8 • **Malwankar + Voorhees ’224**
- 9 • **Malwankar + Klughart + Yendigiri**
- 10 • **Malwankar + Klughart + Steinmetz**
- 11 • **Malwankar + Klughart + Voorhees ’224**
- 12 • **Malwankar + Klughart + Voorhees ’224**
- 13 • **Malwankar + Pittelkow + Steinmetz**
- 14 • **Malwankar + Pittelkow + Voorhees ’224**

15 Additionally, Malwankar teaches converting storage-based requests into a format
16 compatible with the target device. *See* Malwankar, 1:58-61. A POSITA would have found it
17 obvious that the protocol conversion disclosed in Malwankar is not limited to PCIe based devices
18 and can be applied to other well-known storage devices, such as SAS devices. Therefore, a
19 POSITA would naturally have looked to other references that provide details for other types of
20 storage devices. For example, a POSITA would have been motivated to combine Malwankar with
21 Voorhees ’224, which discloses converting PCIe requests into SAS requests. *See* Voorhees ’224,
22 2:50-64, 8:63-9:11. Similarly, a POSITA would have been motivated to combine Malwankar with
23 Yendigiri (*see* Yendigiri, 2:19-33, 3:39-44, 4:41-43, 6:24-32) or Steinmetz (*see* Steinmetz, 53:36-
24 49, 53:64-54:3), each of which is directed to SAS devices. For the same reason, a POSITA would
25 have been motivated to further combine Malwankar and Klughart with Yendigiri, Steinmetz or
26 Voorhees ’224. Likewise, it would have been obvious to a POSITA to integrate the combined
27 disclosures of Malwankar and Pittelkow with Yendigiri, Steinmetz, or Voorhees ’224.
28

- 1 • **Malwankar + Yendigiri + Johnson + Voorhees '096**
- 2 • **Malwankar + Steinmetz + Johnson + Voorhees '096**
- 3 • **Malwankar + Voorhees '224 + Johnson + Voorhees '096**
- 4 • **Malwankar + Klughart + Yendigiri + Johnson + Voorhees '096**
- 5 • **Malwankar + Klughart + Steinmetz + Johnson + Voorhees '096**
- 6 • **Malwankar + Klughart + Voorhees '224 + Johnson + Voorhees '096**
- 7 • **Malwankar + Pittelkow + Yendigiri + Johnson + Voorhees '096**
- 8 • **Malwankar + Pittelkow + Steinmetz + Johnson + Voorhees '096**
- 9 • **Malwankar + Pittelkow + Voorhees '224 + Johnson + Voorhees '096**
- 10 • **Malwankar + Pittelkow + Voorhees '224 + Johnson + Voorhees '096**

11 The combination of Malwankar and Yendigiri provides a scalable data storage system that
12 supports SAS devices. A POSITA would naturally have looked to other references that provide
13 details for SAS device management to enhance the reliability and robustness of the system,
14 including details for device discovery in SAS topology. For example, both Voorhees '096 and
15 Johnson are directed to device discovery in SAS topology. Voorhees '096 discloses updating
16 routing tables in response to changes detected during a device discovery. *See* Voorhees '096, Fig.
17 2, 3:4-46, 6:61-7:6, 8:4-16, 8:17-30. Johnson discloses performing a device discovery upon
18 detecting a status change in the SAS topology and identifies events that could trigger such
19 changes. *See* Johnson, Fig. 4, [0017], [0022], [0023]. Therefore, a POSITA would have been
20 motivated to further combine Malwankar and Yendigiri with Voorhees '096 and Johnson to
21 achieve a robust and reliable storage system. For the same reason, a POSITA would have been
22 motivated to further combine Malwankar and Steinmetz with Voorhees '096 and Johnson, further
23 combine Malwankar and Voorhees '224 with Voorhees '096 and Johnson, further combine
24 Malwankar, Klughart, and Yendigiri with Voorhees '096 and Johnson, further combine
25 Malwankar, Klughart, and Steinmetz with Voorhees '096 and Johnson, and further combine
26 Malwankar, Klughart, and Voorhees '224 with Voorhees '096 and Johnson.
27
28

1 A POSITA would have been motivated to further combine Malwankar, Pittelkow and
2 Yendigiri with Voorhees '096 and Johnson, because of same reason explained above.
3 Furthermore, because the combination of Malwankar, Pittelkow, and Yendigiri already establishes
4 a functional framework for device discovery within a SAN architecture and explicitly supports
5 SAS devices, it would have been a routine and obvious to incorporate the specific SAS discovery
6 mechanisms taught by Voorhees '096 and Johnson. Combining these references effectively
7 “connects the dots” by implementing SAS-specific discovery within a SAN environment that is
8 already configured for automated device discovery. For the same reason, a POSITA would have
9 been motivated to further combine Malwankar, Pittelkow and Voorhees '224 with Voorhees '096
10 and Johnson, and combine Malwankar, Pittelkow, Steinmetz with Voorhees '096 and Johnson.

- 11 • **Pittelkow + Voorhees '224**
- 12 • **Pittelkow + Yendigiri**
- 13 • **Pittelkow + Steinmetz**

14 Pittelkow teaches a storage network implemented on the SAN fabric. *See* Exhibit 767-13.
15 A POSITA would have recognized that the SAN fabric architecture taught by Pittelkow is media-
16 agnostic and not limited to legacy parallel SCSI drives, but is equally capable of supporting other
17 storage devices, including SAS storage devices. Given the industry-wide transition from parallel
18 to serial storage protocols, a POSITA would have been naturally motivated to look to references
19 providing implementation details for these modern drive types. For example, a POSITA would
20 have been motivated to combine Pittelkow with Voorhees '224, which discloses converting host
21 requests into SAS requests. *See* Voorhees '224, 2:50-64, 8:63-9:11. Similarly, a POSITA would
22 have been motivated to combine Pittelkow with Yendigiri (*see* Yendigiri, 2:19-33, 3:39-44, 4:41-
23 43, 6:24-32) or Steinmetz (*see* Steinmetz, 53:36-49, 53:64-54:3), each of which is directed to SAS
24 devices.

- 25 • **Pittelkow + Klughart**
- 26 • **Pittelkow + Northcutt**
- 27 • **Pittelkow + Malwankar**

1 • **Pittelkow + Stephens**

2 Pittelkow teaches a system involving storage controllers. A POSITA would have found it
3 obvious to implement the prior art's teachings relevant to the claimed master controller(s) or slave
4 controller(s) using ASICs or FPGAs, well-known techniques as shown in Northcutt (*see*
5 Northcutt, [0205]), Malwankar (*see* Malwankar, [0171]), Stephens (*see* Stephens, [0129]), and
6 Klughart (*see* Klughart, [0059], [0259], [0261]). A POSITA would have found it obvious to
7 implement such functionality in the form of an application-specific integrated circuit for many
8 reasons including that an ASIC can provide better performance or be more efficient than an
9 implementation that does not use an ASIC (e.g., software implementation running on a general-
10 purpose processor). Similarly, a POSITA would have found it obvious to use an FPGA to achieve
11 similar hardware-level acceleration while maintaining the flexibility to reconfigure the controller
12 logic.

- 13 • **Pittelkow + Voorhees '224**
14 • **Pittelkow + Yendigiri**
15 • **Pittelkow + Steinmetz**
16 • **Pittelkow + Voorhees '224 + Northcutt**
17 • **Pittelkow + Voorhees '224 + Klughart**
18 • **Pittelkow + Yendigiri + Northcutt**
19 • **Pittelkow + Yendigiri + Klughart**
20 • **Pittelkow + Steinmetz + Northcutt**
21 • **Pittelkow + Steinmetz + Klughart**

22 The combination of Pittelkow and Voorhees '224 establishes a scalable data storage
23 architecture supporting SAS devices. Within such a networked environment, seamless
24 communication between storage controllers is required to coordinate complex storage-based
25 operations. While Pittelkow discloses several methods for inter-controller communication, it does
26 not explicitly detail a mechanism for decomposing high-overhead host commands into smaller,
27
28

1 distributed tasks that can be processed across multiple controllers. *See* Pittelkow, 4:9-31, 12:56-
2 13:4, 13:6-35, 20:11-26. A POSITA would have recognized that the performance of a distributed
3 system depends on efficient command execution and the ability to scale processing across the
4 controller fabric. Therefore, a POSITA would have naturally looked into other references that
5 provide details for resolving a single command into multiple commands to achieve distributed
6 operations across multiple controllers. For example, a POSITA would have been motivated to
7 further combine Pittelkow and Voorhees '224 with Northcutt (*see* Northcutt, [0004], [0018],
8 [0025], [0031], [0036]) or Klughart (*see* Klughart, [0105], [0107], [0139]), each of which teaches
9 one command transforming into more commands, to enhance the storage network system's
10 capability of managing distributed tasks efficiently. Similarly, a POSITA would have been
11 motivated to further combine Pittelkow and Steinmetz with Northcutt or Klughart, and further
12 combine Pittelkow and Yendigiri with Northcutt or Klughart.

- 13 • **Pittelkow + Yendigiri + Johnson + Voorhees '096**
- 14 • **Pittelkow + Steinmetz + Johnson + Voorhees '096**
- 15 • **Pittelkow + Voorhees '224 + Johnson + Voorhees '096**
- 16

17 The combination of Pittelkow and Yendigiri establishes a storage network architecture that
18 supports SAS devices and implements automated device discovery within a SAN environment.
19 Because Pittelkow already discloses the framework for general SAN discovery, a POSITA would
20 have found it obvious to incorporate the specific SAS discovery mechanisms required to manage a
21 SAS-based topology. Therefore, a POSITA seeking to optimize network initialization and
22 management would have naturally looked to specialized references disclosing SAS-specific
23 discovery protocols. For example, both Voorhees '096 and Johnson are directed to device
24 discovery in SAS topology. Voorhees '096 discloses updating routing tables in response to
25 changes detected during a device discovery. *See* Voorhees '096, Fig. 2, 3:4-46, 6:61-7:6, 8:4-16,
26 8:17-30. Johnson discloses performing a device discovery upon detecting a status change in the
27 SAS topology and identifies events that could trigger such changes. *See* Johnson, Fig. 4, [0017],
28 [0022], [0023]. Therefore, a POSITA would have been motivated to further combine Pittelkow

1 and Yendigiri with Voorhees '096 and Johnson to achieve a robust and reliable storage system.
2 For the same reason, a POSITA would have been motivated to further combine Pittelkow and
3 Steinmetz with Voorhees '096 and Johnson, further combine Pittelkow and Voorhees '224 with
4 Voorhees '096 and Johnson.

- 5 • **EMC VPLEX System + Yendigiri**
- 6 • **EMC VPLEX System + Steinmetz**
- 7 • **EMC VPLEX System + Voorhees '224**

8 EMC VPLEX System discloses a scalable storage virtualization system that integrates
9 diverse device storage arrays to enable high availability and seamless data mobility across storage
10 networks. *See* Ex. 767-01. A POSITA would have naturally looked into other references that
11 provide details for various storage devices to complement EMC VPLEX System's scalable and
12 heterogeneous storage virtualization framework, including SAS devices. For example, a POSITA
13 would have been motivated to combine EMC VPLEX System with Yendigiri (*see* Yendigiri, 2:19-
14 33, 3:39-44, 4:41-43, 6:24-32), Steinmetz (*see* Steinmetz, 53:36-49, 53:64-54:3), or Voorhees
15 '224 (*see* Voorhees '224, 2:50-64, 8:63-9:11), either of which is directed to SAS devices.

- 16 • **EMC VPLEX System + Yendigiri + Johnson + Voorhees '096**
- 17 • **EMC VPLEX System + Steinmetz + Johnson + Voorhees '096**
- 18 • **EMC VPLEX System + Voorhees '224 + Johnson + Voorhees '096**

19 The combination of EMC VPLEX System and Yendigiri provides a scalable data storage
20 system that supports SAS devices. A POSITA would naturally have looked to other references
21 that provide details for SAS device management to enhance the reliability and robustness of the
22 system, including details for device discovery in SAS topology. For example, both Voorhees '096
23 and Johnson are directed to device discovery in SAS topology. Voorhees '096 discloses updating
24 routing tables in response to changes detected during a device discovery. *See* Voorhees '096, Fig.
25 2, 3:4-46, 6:61-7:6, 8:4-16, 8:17-30. Johnson discloses performing a device discovery upon
26 detecting a status change in the SAS topology and identifies events that could trigger such
27 changes. *See* Johnson, Fig. 4, [0017], [0022], [0023]. Therefore, a POSITA would have been
28

1 motivated to further combine EMC VPLEX System and Yendigiri with Voorhees '096 and
2 Johnson to achieve a robust and reliable storage system. For the same reason, a POSITA would
3 have been motivated to further combine EMC VPLEX System and Steinmetz with Voorhees '096
4 and Johnson, and further combine EMC VPLEX System and Voorhees '224 with Voorhees '096
5 and Johnson.

- 6 • **EMC VPLEX System + Yendigiri + Northcutt**
- 7 • **EMC VPLEX System + Yendigiri + Klughart**
- 8 • **EMC VPLEX System + Steinmetz+ Northcutt**
- 9 • **EMC VPLEX System + Steinmetz + Klughart**
- 10 • **EMC VPLEX System + Steinmetz + Klughart**
- 11 • **EMC VPLEX System + Voorhees '224 + Northcutt**
- 12 • **EMC VPLEX System + Voorhees '224 + Klughart**

13 The combination of EMC VPLEX System and Yendigiri provides a scalable data storage
14 system that supports SAS devices. In such a system, one EMC VPLEX engine needs to
15 communicate with other EMC VPLEX engines within the cluster or across clusters to coordinate
16 storage operations. A POSITA would have recognized that such communication requires efficient
17 and scalable command execution mechanisms to manage distributed systems effectively.
18 Therefore, a POSITA would have naturally looked into other references that provide details for
19 resolving a single command into multiple commands to achieve distributed operations across
20 multiple EMC VPLEX engines. For example, a POSITA would have been motivated to further
21 combine EMC VPLEX System and Yendigiri with Northcutt (*see* Northcutt, [0004], [0018],
22 [0025], [0031], [0036]) or Klughart (*see* Klughart, [0105], [0107], [0139]), each of which teaches
23 one command transforming into more commands, to enhance EMC VPLEX System's capability
24 of managing distributed tasks efficiently. Similarly, a POSITA would have been motivated to
25 further combine EMC VPLEX System and Steinmetz with Northcutt or Klughart, and further
26 combine EMC VPLEX System and Voorhees '224 with Northcutt or Klughart.

- 27 • **IBM SAN Volume Controller System + Klughart**
- 28 • **IBM SAN Volume Controller System+ Northcutt**

1 IBM SAN Volume Controller System discloses a storage virtualization system that
2 manages and optimizes resources across multiple storage arrays, including SAS devices. *See* Ex.
3 767-02. In such a system, one SAN Volume Controller needs to communicate with other SAN
4 Volume Controllers to provide the host seamless access to virtualized storages. A POSITA would
5 have recognized the need for mechanisms to transform a single command into multiple
6 coordinated operations across nodes to achieve distributed operations and more efficient
7 functionality. For example, a POSITA would have been motivated to combine SAN Volume
8 Controller System with Northcutt (*see* Northcutt, [0004], [0018], [0025], [0031], [0036]) or
9 Klughart (*see* Klughart, [0105], [0107], [0139]), each of which teaches one command
10 transforming into more commands, to enhance SAN Volume Controller's capability of managing
11 distributed tasks efficiently.

12 For any of the system art (or the foregoing combinations that include system art), a
13 POSITA would have found it obvious to implement the prior art's teachings relevant to the
14 claimed master controller(s) or edge controller(s) using ASICs or FPGAs, well-known techniques
15 as shown in Northcutt (*see* Northcutt, [0205]), Malwankar (*see* Malwankar, [0171]), Stephens (*see*
16 Stephens, [0129]), and Klughart (*see* Klughart, [0059], [0259], [0261]). A POSITA would have
17 found it obvious to implement such functionality in the form of an application-specific integrated
18 circuit for many reasons including that an ASIC can provide better performance or be more
19 efficient than an implementation that does not use an ASIC (e.g., software implementation running
20 on a general-purpose processor). Similarly, a POSITA would have found it obvious to use an
21 FPGA to achieve similar hardware-level acceleration while maintaining the flexibility to
22 reconfigure the controller logic.

23 **C. U.S. Patent No. 10,873,753**

24 The following table identifies the prior art for the '753 patent. Each item below constitutes
25 prior art under one or more provisions of 35 U.S.C. §§ 102(a)(1)-(2) by reason of its status as a
26 publication and/or patent disclosing, or a system on sale and/or in public knowledge or use in the
27 United States embodying the claimed invention before Plaintiffs' asserted priority date of the '753
28

1 patent. For each item of prior art, Defendants have attached a claim chart identifying where each
 2 claim limitation is found in the cited prior art.

3 **Table C: Prior Art to the '753 Patent**

Prior Art to '753	Anticipation/ Obviousness	Claim Chart
4 Cray Supercomputer System(s), including but not 5 limited to Cray XC40, Cray XD1, and Cray XE, 6 were sold, offered for sale, made known, and/or 7 publicly used before July 15, 2018. Cray 8 Supercomputer System(s) were sold, offered for 9 sale, and/or made known to the public and third- 10 party customers before this date, by July 14, 2018 11 at the latest. Cray Supercomputer System(s) were 12 publicly used by HPE/Cray, the public, and third- 13 party customers before this date, by July 14, 2018 14 at the latest.	A & O	Exhibit 753-01
15 BreakingPoint Scalable System(s), including but 16 not limited to BreakingPoint Enterprise, 17 BreakingPoint FireStorm, BreakingPoint Elite, 18 were sold, offered for sale, made known, and/or 19 publicly used before July 15, 2018. 20 BreakingPoint Scalable System(s) were sold, 21 offered for sale, and/or made known to the public 22 and third-party customers before this date, by 23 July 14, 2018 at the latest. BreakingPoint 24 Scalable System(s) were publicly used by 25 BreakingPoint, the public, and third-party 26 customers before this date, by July 14, 2018 at the 27 latest.	A & O	Exhibit 753-02
28 SeaMicro Computing System(s), including but not limited to SeaMicro SM15000 Cloud Servers, were sold, offered for sale, made known, and/or publicly used before July 15, 2018. SeaMicro Computing System(s) were sold, offered for sale, and/or made known to the public and third-party customers before this date, by July 14, 2018 at the latest. SeaMicro Computing System(s) were publicly used by SeaMicro/AMD, the public, and third-party customers before this date, by July 14, 2018 at the latest.	A & O	Exhibit 753-03
U.S. Patent Application Publication No. 2013/0346667 to Stroud et al. ("Stroud"); filed on June 21, 2012; published on December 26, 2013	A & O	Exhibit 753-04

Prior Art to '753	Anticipation/ Obviousness	Claim Chart
U.S. Patent No. 9,306,816 to Jana et al. ("Jana"); filed on December 24, 2013; issued on April 5, 2016	O	Exhibit 753-05
U.S. Patent No. 10,673,783 to Forrest Lawrence Pierson ("Pierson"); filed on January 29, 2019; continuation of application No. 15/202,379 filed on July 5, 2016; parent application filed on September 9, 2013; claims priority to provisional application Nos. 62/188,731, filed on July 5, 2015 and 61/778,393, filed on March 12, 2013.	A & O	Exhibit 753-06
U.S. Patent Number. 9,148,389 to Da Silva et al. ("Da Silva"); filed November 12, 2012; issued on September 29, 2015	O	Exhibit 753-07
Harvey G. Cragon, <i>Memory Systems and Pipelined Processors</i> (Jones & Bartlett Publishers 1996).	O	Exhibit 753-08
Internet Protocol, Request for Comments (RFC) 791 (Jon Postel ed., Sept. 1981), https://www.rfc-editor.org/rfc/rfc791 .	O	Exhibit 753-09
U.S. Publication No. 2002/0087821 A1 to Saulsbury et al. ("Saulsbury"); published in July 4, 2002	O	Exhibit 753-10
U.S. Patent No. 5,155,833 to Cullison et al. ("Cullison"); issued on October 13, 1992	O	Exhibit 753-11
Dell PC System(s), including but not limited to Dell XPS 13 Notebook and Dell Precision 5820 Tower, were sold, offered for sale, made known, and/or publicly used before July 15, 2018. Dell PC System(s) were sold, offered for sale, and/or made known to the public and third-party customers before this date, by July 14, 2018 at the latest. Dell PC System(s) were publicly used by Dell, the public, and third-party customers before this date, by July 14, 2018 at the latest.	A & O	Exhibit 753-12

Prior Art to '753	Anticipation/ Obviousness	Claim Chart
IBM System(s), including but not limited to IBM z13, were sold, offered for sale, made known, and/or publicly used before July 15, 2018. IBM System(s) were sold, offered for sale, and/or made known to the public and third-party customers before this date, by July 14, 2018 at the latest. IBM System(s) were publicly used by IBM, the public, and third-party customers before this date, by July 14, 2018 at the latest.	A & O	Exhibit 753-13
U.S. Patent No. 9,317,429 to Ramanujan et al. ("Ramanujan"); filed September 30, 2011; issued April 19, 2016	O	Exhibit 753-14
Intel Optane System(s), including non-volatile memory devices and associated reference systems, were sold, offered for sale, made known, and/or publicly used before July 15, 2018. Intel Optane System(s) were sold, offered for sale, and/or made known to the public and third-party customers before this date, by July 14, 2018 at the latest. Intel Optane System(s) were publicly used by Intel, the public, and third-party customers before this date, by July 14, 2018 at the latest.	O	

In addition to the anticipation and obviousness explained in this section, the '753 patent is anticipated and/or rendered obvious by prior art as explained by Defendants' Petition for Inter Partes Review (assigned Case No. IPR2025-00486), which is incorporated herein by reference:

Grounds	Claims	Prior Art
1	1-3, 13, 14	Stroud and DaSilva
2	7, 9-10	Stroud and Cragon
3	7, 9-10	Stroud, Cragon, and Saulsbury
4	4, 8, 15	Stroud, DaSilva, Cragon, and Saulsbury

1. Anticipation

For each item of prior art identified above as an anticipation reference, Defendants contend that the prior art item anticipates the Asserted Claims of the '753 patent, including for the reasons

1 illustrated in the accompanying claim chart and based on Plaintiffs' infringement allegations and
2 contentions for the '753 patent.

3 2. Obviousness

4 A person of ordinary skill in the art would have been motivated to combine any reference
5 in Table C with any other reference(s) in Table C. To the extent that any prior art identified in
6 Table C above is deemed to not anticipate an Asserted Claim, the prior art in Table C may be
7 combined with one or more prior art references in Table C to render the claim obvious. All of
8 these items of prior art are directed to and propose solutions similar to those discussed in the '753
9 patent, and is in the same field. All of these items of prior art are directed to computer architecture
10 and/or dataflow processing, and propose solutions similar to those discussed in the '753 patent.

11 The following are exemplary combinations of prior art that would have been obvious to a
12 POSITA:

- 13 • **Cray Supercomputer System(s) + Pierson or Stroud**
- 14 • **Cray Supercomputer System(s) + RFC 791**
- 15 • **Cray Supercomputer System(s) + Da Silva**
- 16 • **Cray Supercomputer System(s) + Cragon**
- 17 • **Cray Supercomputer System(s) + Ramanujan**
- 18 • **Cray Supercomputer System(s) + Dell PC System(s)**
- 19 • **Cray Supercomputer System(s) + IBM System(s)**
- 20 • **Cray Supercomputer System(s) + Intel Optane System(s)**
- 21 • **Cray Supercomputer System(s) + Intel Optane System(s)**
- 22 • **Any of the above + Cullison or Saulsbury**
- 23 • **Any of the above + Cullison or Saulsbury**

24 The Cray Supercomputer System(s) discloses a supercomputer computer architecture with
25 swappable cards in a chassis-based networked topology for distributed computing. *See, e.g.*,
26 Exhibit 753-01. A POSITA would naturally have looked to other references that provide details
27 for networked computing environments, including details and features for chassis, cards, and
28 networks in data centers, such as the use of a chassis identifier, module identifier, board identifier,

1 instance identifier, and type identifier in packet headers. For example, a POSITA would have
2 been motivated to combine the Cray Supercomputer System(s) with Pierson, which teaches a
3 similar multi-card architecture for parallel, distributed processing in a networked environment,
4 with details and features for connecting cards using a backplane and other networks, and regarding
5 the configuration of chassis systems and the packet format and headers for systems having
6 multiple cards. *See* Pierson, Abstract, 3:35-4:2, 4:3-26, 7:1-32, 14:50-64, 15:2-13, 19:40-53,
7 20:46-55, 27:44-28:15, Figs. 1, 2, 5, 8. Likewise, a POSITA would have been motivated to
8 combine the Cray Supercomputer System(s) with Stroud, which teaches a multi-card chassis
9 architecture for network processing and simulation, with details and features regarding the
10 configuration of chassis systems and the packet format and headers for systems having multiple
11 cards. *See* Stroud, ¶¶ [0098]-[0102], [0322]-[0342] Figs. 3, 54, 55A, 55B, 56A. Similarly, a
12 POSITA would have been motivated to combine the Cray Supercomputer System(s) with Da
13 Silva, which teaches a multi-chassis system in a networked environment with details and features
14 (e.g., regarding packet format and headers) for routing packets in a multi-chassis architecture. *See*
15 Da Silva, Abstract, 4:23-33, 4:52-54, 5:13-18, 9:1-18, Figs. 4 and 6. A POSITA would naturally
16 have looked to these references and been motivated to combine their teachings—including details
17 and features for chassis, cards, and networks in data centers, such as the use of a chassis identifier,
18 module identifier, board identifier, instance identifier, and type identifier in packet headers—with
19 the Cray Supercomputer System(s) because, e.g., they provide additional implementation details
20 for chassis architectures in a networked setting; facilitate communication between cards and
21 processors, allowing more granular access to particular components; reduce network overhead in
22 distributed computing systems with multiple cards and chassis; and improve network efficiency in
23 such systems.

24 A POSITA would have similarly been motivated to combine Cray Supercomputer
25 System(s) with RFC 791. Cray Supercomputer System(s) discloses a supercomputer architecture
26 that sent packet communications throughout a network, which Cray Supercomputer System(s)
27 teaches or suggests is an IP network over Ethernet. *See, e.g.*, Exhibit 753-01. RFC 791 is the
28 widely-known standard governing the IP protocol. Therefore, a POSITA would have known to

1 use RFC 791 and been motivated to use the IP protocol with Cray Supercomputer System(s). For
2 decades, the IP protocol has been ubiquitous in computers and computer networks—including for
3 supercomputers such as Cray Supercomputer System(s). *See* RFC 791 at 1-3, 5-6, 8-9, and 11-23.
4 Therefore, a POSITA would have looked to RFC 791 and been motivated to combine it with Cray
5 Supercomputer System(s) packet-based network communications to facilitate packet-based
6 communications over IP and Ethernet.

7 Additionally, the Cray Supercomputer System(s) supports an assortment of memory
8 functions and types, including L1, L2, and L3 cache, DDR memory, storage arrays, flash memory,
9 and disk drives. *See, e.g.*, Exhibit 753-01. A POSITA would naturally have looked to references
10 that provide features for using different types of memory together. For example, a POSITA would
11 have been motivated to combine the Cray Supercomputer System(s) with Ramanujan, which
12 teaches a memory hierarchy and different configurable modes of memory operation. *See*
13 Ramanujan at 3:32-35, 7:45-54, 8:11-29, 8:30-37, 9:1-10, 9:27-31, 9:34-40, 10:31-38, 13:3-9,
14 17:59-18:6, 31:5-12, 31:13-19, 31:29-32, FIG. 1, FIG. 2, FIG. 5B, FIG. 6D. Likewise, Cragon
15 teaches multiple memory arrangements and modes of memory operation. *See* Cragon at 73, 91,
16 110, 191, 231-232, 257, 414-415, 491. A POSITA would have also been motivated to look at the
17 memory arrangements of existing computer systems at the time, including Dell PC System(s),
18 which teaches multiple memory arrangements, including for DDR3/DDR4; L1, L2, and L3 cache;
19 flash memory; hard disks; etc. (*see* Exhibit 753-12) and IBM System(s), which teaches multiple
20 memory arrangements, including DDR3; L1, L2, L3, and L4 cache; flash memory; hard disks; etc.
21 for performing analysis of large volumes of data and operational analytics of workloads
22 throughout a cloud network (*see* Exhibit 753-13). Additionally, given these various memory
23 arrangements, a POSITA would have been motivated to utilize the teachings of Saulsbury or
24 Cullison, which each teach the selective operation of memory in different configurations. *See*
25 Saulsbury, Abstract, ¶¶ [0003], [0005]-[0006], [0008], [0017], [0020], [0022], [0024]-[0026],
26 [0068]-[0070], Figs. 1 and 5; Cullison, 1:11-23, 1:41-63, 2:18-32, 2:36-48, 2:58-62, 2:63-3:21,
27 FIG. 1. A POSITA would naturally have looked to these references and been motivated to
28 combine them with the Cray Supercomputer System(s) because, e.g., the combinations would

1 provide for more flexible memory arrangements to account for different usage models and
2 applications; reduce memory latency for applications where latency a the bottleneck on
3 performance; improve memory throughput for applications where bandwidth is a bottleneck on
4 performance; improve efficiency and utilization of available memory; and increase performance of
5 distributed computing systems.

- 6 • **BreakingPoint Scalable System(s) + Pierson**
- 7 • **BreakingPoint Scalable System(s) + RFC 791**
- 8 • **BreakingPoint Scalable System(s) + Da Silva**
- 9 • **BreakingPoint Scalable System(s) + Cragon**
- 10 • **BreakingPoint Scalable System(s) + Ramanujan**
- 11 • **BreakingPoint Scalable System(s) + Dell PC System(s)**
- 12 • **BreakingPoint Scalable System(s) + IBM System(s)**
- 13 • **BreakingPoint Scalable System(s) + Intel Optane System(s)**
- 14 • **Any of the above + Cullison**

15 The BreakingPoint Scalable System(s) discloses a chassis-based architecture with
16 swappable blades for use in networked environments with various networking protocols for
17 distributed computing, including with streams of packet data. *See, e.g.,* Exhibit 753-02. A
18 POSITA would naturally have looked to other references that provide details for networked
19 computing environments, including details and features for chassis, cards/blades, and networks to
20 support data streams, such as the use of a chassis identifier, a module identifier, board identifier,
21 instance identifier, and type identifier in packet headers. For example, a POSITA would have
22 been motivated to combine the BreakingPoint Scalable System(s) with Pierson, which teaches a
23 multi-card architecture for parallel, distributed processing in a networked environment, with
24 details and features for connecting cards using a backplane and other networks, and regarding the
25 configuration of chassis systems and the packet format and headers for systems having multiple
26 cards. *See* Pierson, Abstract, 3:35-4:2, 4:3-26, 7:1-32, 14:50-64, 15:2-13, 19:40-53, 20:46-55,
27 27:44-28:15, Figs. 1, 2, 5, 8. Likewise, a POSITA would have been motivated to combine the
28 BreakingPoint Scalable System(s) with Stroud, which teaches a multi-card chassis architecture for

1 network processing and simulation, with details and features regarding the configuration of chassis
2 systems and the packet format and headers for systems having multiple cards. *See* Stroud, ¶¶
3 [0098]-[0102], [0322]-[0342] Figs. 3, 54, 55A, 55B, 56A. Similarly, a POSITA would have been
4 motivated to combine the BreakingPoint Scalable System(s) with Da Silva, which teaches a
5 scalable multi-chassis system in a networked environment with details and features (e.g.,
6 regarding packet format and headers) for routing packets in a multi-chassis architecture. *See* Da
7 Silva, Abstract, 4:23-33, 4:52-54, 5:13-18, 9:1-18, Figs. 4 and 6. A POSITA would naturally have
8 looked to these references and been motivated to combine their teachings—including details and
9 features for chassis, and cards/blades, such as the use of a chassis identifier, module identifier,
10 board identifier, instance identifier, and type identifier in packet headers—with the BreakingPoint
11 Scalable System(s) because, e.g., they provide additional implementation details for chassis
12 architectures in a networked setting; facilitate communication between cards and processors,
13 allowing more granular access to particular components; reduce network overhead in distributed
14 computing systems with multiple cards and chassis; and improve network efficiency in such
15 systems.

16 A POSITA would have similarly been motivated to combine Scalable System(s) with RFC
17 791. BreakingPoint Scalable System(s) discloses a chassis-based architecture that sent packet
18 communications throughout a network, which BreakingPoint Scalable System(s) teaches or
19 suggests is an IP network over Ethernet. *See*, e.g., Exhibit 753-02. RFC 791 is the widely-known
20 standard governing the IP protocol. Therefore, a POSITA would have known to use RFC 791 and
21 been motivated to use the IP protocol with BreakingPoint Scalable System(s). For decades, the IP
22 protocol has been ubiquitous in computers and computer networks- including for infrastructure
23 testing computer systems such as BreakingPoint Scalable System(s). *See* RFC 791 at 1-3, 5-6, 8-9,
24 and 11-23. Therefore, a POSITA would have looked to RFC 791 and been motivated to combine it
25 with BreakingPoint Scalable System(s) packet-based network communications to facilitate packet-
26 based communications over IP and Ethernet.

27 Additionally, the BreakingPoint Scalable System(s) supports an assortment of memory
28 functions and types, including RAM and hard disks. *See*, e.g., Exhibit 753-02. A POSITA would

1 naturally have looked to references that provide features for using different types of memory
2 together. For example, a POSITA would have been motivated to combine the BreakingPoint
3 Scalable System(s) with Ramanujan, which teaches a memory hierarchy and different configurable
4 modes of memory operation. *See* Ramanujan at 3:32-35, 7:45-54, 8:11-29, 8:30-37, 9:1-10, 9:27-
5 31, 9:34-40, 10:31-38, 13:3-9, 17:59-18:6, 31:5-12, 31:13-19, 31:29-32, FIG. 1, FIG. 2, FIG. 5B,
6 FIG. 6D. Likewise, Cragon, teaches multiple memory arrangements and modes of memory
7 operation (*See* Cragon at 73, 91, 110, 191, 231-232, 257, 414-415, 491). A POSITA would have
8 also been motivated to look at the memory arrangements of existing computer systems at the time,
9 including Dell PC System(s), which teaches multiple memory arrangements, including for
10 DDR3/DDR4; L1, L2, and L3 cache; flash memory; hard disks; etc. (*see* Exhibit 753-12) and IBM
11 System(s), which teaches multiple memory arrangements, including DDR3; L1, L2, L3, and L4
12 cache; flash memory; hard disks; etc. for performing analysis of large volumes of data and
13 operational analytics of workloads throughout a cloud network (*see* Exhibit 753-13). Additionally,
14 given these various arrangements, a POSITA would have been motivated to utilize the teachings
15 of Saulsbury or Cullison, which each teaches the selective operation of memory in different
16 configurations. *See* Saulsbury, Abstract, ¶¶ [0003], [0005]-[0006], [0008], [0017], [0020], [0022],
17 [0024]-[0026], [0068]-[0070], Figs. 1 and 5; Cullison, 1:11-23, 1:41-63, 2:18-32, 2:36-48, 2:58-
18 62, 2:63-3:21, FIG. 1. A POSITA would naturally have looked to these references and been
19 motivated to combine them with the BreakingPoint Scalable System(s) because, e.g., the
20 combinations would provide for more flexible memory arrangements to account for different
21 usage models and applications; reduce memory latency for applications where latency a the
22 bottleneck on performance; improve memory throughput for applications where bandwidth is a
23 bottleneck on performance; improve efficiency and utilization of available memory; and increase
24 performance of distributed computing systems.

- 25 • **SeaMicro Computing System(s) + Pierson**
- 26 • **SeaMicro Computing System(s) + RFC 791**
- 27 • **SeaMicro Computing System(s) + Da Silva**
- 28 • **SeaMicro Computing System(s) + Cragon**

- 1 • **SeaMicro Computing System(s) + Ramanujan**
- 2 • **SeaMicro Computing System(s) + Dell PC System(s)**
- 3 • **SeaMicro Computing System(s) + IBM System(s)**
- 4 • **SeaMicro Computing System(s) + Intel Optane System(s)**
- 5 • **Any of the above + Cullison**

6 The SeaMicro Computing System(s) discloses an energy-efficient, high-bandwidth
7 microserver architecture with swappable cards in a dense chassis configuration for networked,
8 distributed computing. *See, e.g.*, Exhibit 753-03. A POSITA would naturally have looked to
9 other references that provide details for networked computing environments, including details and
10 features for chassis, cards, and networks to provide “rapid, reliable, and scalable” infrastructure to
11 “application[s]; and enterprises,” such as the use of a chassis identifier, module identifier, board
12 identifier, instance identifier, and type identifier in packet headers. For example, a POSITA would
13 have been motivated to combine the SeaMicro Computing System(s) with Pierson, which teaches
14 a multi-card architecture for parallel, distributed processing in a networked environment, with
15 details and features for connecting cards using a backplane and other networks, and regarding the
16 configuration of chassis systems and the packet format and headers for systems having multiple
17 cards. *See* Pierson, Abstract, 3:35-4:2, 4:3-26, 7:1-32, 14:50-64, 15:2-13, 19:40-53, 20:46-55,
18 27:44-28:15, Figs. 1, 2, 5, 8. Likewise, a POSITA would have been motivated to combine the
19 SeaMicro Computing System(s) with Stroud, which teaches a multi-card chassis architecture for
20 network processing and simulation, with details and features regarding the configuration of chassis
21 systems and the packet format and headers for systems having multiple cards. *See* Stroud, ¶¶
22 [0098]-[0102], [0322]-[0342] Figs. 3, 54, 55A, 55B, 56A. Similarly, a POSITA would have been
23 motivated to combine the SeaMicro Computing System(s) with Da Silva, which teaches a scalable
24 multi-chassis system in a networked environment with details and features (e.g., regarding packet
25 format and headers) for routing packets in a multi-chassis architecture. *See* Da Silva, Abstract,
26 4:23-33, 4:52-54, 5:13-18, 9:1-18, Figs. 4 and 6. A POSITA would naturally have looked to these
27 references and been motivated to combine their teachings—including details and features for
28 chassis, and cards, such as the use of a chassis identifier, module identifier, board identifier,

1 instance identifier, and type identifier in packet headers—with the SeaMicro Computing
2 System(s) because, e.g., they provide additional implementation details for chassis architectures in
3 a networked setting; facilitate communication between cards and processors, allowing more
4 granular access to particular components; reduce network overhead in distributed computing
5 systems with multiple cards and chassis; and improve network efficiency in such systems.

6 A POSITA would have similarly been motivated to combine SeaMicro Computing
7 System(s) with RFC 791. SeaMicro Computing System(s) discloses in a dense chassis
8 configuration for networked, distributed computing, which SeaMicro Computing System(s)
9 teaches or suggests is an IP network over Ethernet. *See, e.g.,* Exhibit 753-03. RFC 791 is a widely-
10 known standard governing the IP protocol. Therefore, a POSITA would have known to use RFC
11 791 and been motivated to use the IP protocol with SeaMicro Computing System(s). For decades,
12 the IP protocol has been ubiquitous in computers and computer networks. *See* RFC 791 at 1-3, 5-
13 6, 8-9, and 11-23. Therefore, a POSITA would have looked to RFC 791 and been motivated to
14 combine it with SeaMicro Computing System(s) packet-based network communications to
15 facilitate packet-based communications over IP and Ethernet.

16 Additionally, the SeaMicro Computing System(s) supports an assortment of memory
17 functions and types, including L1, L2, and L3 cache, DRAM, SATA SSD drives, hard disk drives,
18 DDR3, and SODIMM. *See, e.g.,* Exhibit 753-03. A POSITA would naturally have looked to
19 references that provide features for using different types of memory together. For example, a
20 POSITA would have been motivated to combine SeaMicro Computing System(s) with
21 Ramanujan, which teaches a memory hierarchy and different configurable modes of memory
22 operation. *See* Ramanujan at 3:32-35, 7:45-54, 8:11-29, 8:30-37, 9:1-10, 9:27-31, 9:34-40, 10:31-
23 38, 13:3-9, 17:59-18:6, 31:5-12, 31:13-19, 31:29-32, FIG. 1, FIG. 2, FIG. 5B, FIG. 6D. Likewise,
24 Cragon teaches memory arrangements and modes of memory operation (*See* Cragon at 73, 91,
25 110, 191, 231-232, 257, 414-415, 491). A POSITA would have also been motivated to look at the
26 memory arrangements of existing computer systems at the time, including Dell PC System(s),
27 which teaches multiple memory arrangements, including for DDR3/DDR4; L1, L3, and L3 cache;
28 flash memory; hard disks; etc. (*See* Exhibit 753-12) and IBM System(s), which teaches multiple

1 memory arrangements, including DDR3; L1, L2, L3, and L4 cache; flash memory; hard disks; etc.
2 for performing analysis of large volumes of data and operational analytics of workloads
3 throughout a cloud network (*See* Exhibit 753-13). Additionally, given these various memory
4 arrangements, a POSITA would have been motivated to utilize the teachings of Saulsbury or
5 Cullison, which each teaches the selective operation of memory in different configurations. *See*
6 Saulsbury, Abstract, ¶¶ [0003], [0005]-[0006], [0008], [0017], [0020], [0022], [0024]-[0026],
7 [0068]-[0070], Figs. 1 and 5; Cullison, 1:11-23, 1:41-63, 2:18-32, 2:36-48, 2:58-62, 2:63-3:21,
8 FIG. 1. A POSITA would naturally have looked to these references and been motivated to
9 combine them with the SeaMicro Computing System(s) because, e.g., the combinations would
10 provide for more flexible memory arrangements to account for different usage models and
11 applications; reduce memory latency for applications where latency a the bottleneck on
12 performance; improve memory throughput for applications where bandwidth is a bottleneck on
13 performance; improve efficiency and utilization of available memory; and increase performance of
14 distributed computing systems.

- 15 • **Stroud + Da Silva**
- 16 • **Stroud + RFC 791**
- 17 • **Any of the above + Cragon**
- 18 • **Any of the above + Saulsbury or Cullison**
- 19

20 Stroud teaches a chassis-based architecture with swappable cards for networked
21 environments and distributed computing, including support for various networking protocols and
22 the generation and processing of streams of packet data. *See* Stroud, ¶¶ [0098]-[0102], [0322]-
23 [0342] Figs. 3, 54, 55A, 55B, 56A. Stroud further teaches details regarding packet headers sent
24 between chassis components. *See, e.g.,* Exhibit 753-04. Therefore, a POSITA looking at Stroud's
25 multi-chassis architecture would naturally have looked to other references that provide details for
26 networked computing environments, including details and features for packet routing in a multi-
27 chassis architecture. For example, a POSITA would have been motivated to combine Stroud with
28 Da Silva, which teaches a multi-chassis architecture with details and features, including packet

1 format and headers, for routing packets in a multi-chassis architecture. *See* Da Silva, Abstract,
2 4:23-33, 4:52-54, 5:13-18, 9:1-18, Figs. 4 and 6. A POSITA would naturally have looked
3 combine Da Silva with Stroud because, e.g., the combination would have provided additional
4 implementation details for chassis architectures in a networked setting; facilitated communication
5 between cards and processors, allowed more granular access to particular components; reduced
6 network overhead in distributed computing systems with multiple cards and chassis; and improved
7 network efficiency in such systems.

8 Stroud teaches the use of IP packets, which Plaintiffs have accused of infringement. Since
9 Stroud already teaches IP packets, a POSITA would have been motivated to combine Stroud with
10 the standards document for IP packets, RFC 791, which teaches details regarding the packet
11 format and headers. *See* RFC 791 at 13, 17-20, 29-30, 35, 38.

12 Additionally, Stroud supports an assortment of memory functions and types, including
13 flash memory, HDDR, DDR2, and DDR3. *See* Stroud, ¶¶ [0103], [0114], [0127], [0169], [0717],
14 Figs. 4-5, 14A-14B, 55A-55B. A POSITA would naturally have looked to references that provide
15 features for using different types of memory together. For example, a POSITA would have been
16 motivated to combine Stroud with Cragon, which teaches memory architectures and modes of
17 memory operation (*See* Cragon at 73, 91, 110, 191, 231-232, 257, 414-415, 491), and with
18 Saulsbury and/or Cullison, which each teaches the selective operation of memory in different
19 configurations. *See* Saulsbury, Abstract, ¶¶ [0003], [0005]-[0006], [0008], [0017], [0020], [0022],
20 [0024]-[0026], [0068]-[0070], Figs. 1 and 5; Cullison, 1:11-23, 1:41-63, 2:18-32, 2:36-48, 2:58-
21 62, 2:63-3:21, FIG. 1. A POSITA would naturally have looked to these references and been
22 motivated to combine them with Stroud because, e.g., the combinations would provide for more
23 flexible memory arrangements to account for different usage models and applications; reduce
24 memory latency for applications where latency a the bottleneck on performance; improve memory
25 throughput for applications where bandwidth is a bottleneck on performance; improve efficiency
26 and utilization of available memory; and increase performance of distributed computing systems.

- 27 • **Jana + RFC 791**
- 28 • **Jana + RFC 791 + Cragon**

- 1 • **Jana + RFC 791 + Ramanujan**
- 2 • **Jana + RFC 791 + Dell PC System(s)**
- 3 • **Jana + RFC 791 + IBM System(s)**
- 4 • **Jana + RFC 791 + Intel Optane System(s)**
- 5 • **Any of the above + Saulsbury or Cullison**

6 Jana discloses a chassis-based architecture with swappable cards for networked
7 environments and distributed computing, including support for various networking protocols and
8 the generation and processing of streams of packet data. *See* Jana, 4:51-5:45, 6:8-33, Figs. 2-3.
9 Jana teaches the use of IP packets, which Plaintiffs have accused of infringement. Since Jana
10 already teaches IP packets, a POSITA would have been motivated to combine Jana with the
11 standards document for IP packets, RFC 791, which teaches details regarding the packet format
12 and headers. *See* RFC 791 at 13, 17-20, 29-30, 35, 38.

13 Additionally, Jana supports an assortment of memory functions and types, including flash
14 memory and RAM. *See* Jana, 6:33-59, 7:5-26, FIG. 3. A POSITA would naturally have looked to
15 references that provide features for using different types of memory together. For example, a
16 POSITA would have been motivated to combine Jana with Cragon, Ramanujan, Dell PC
17 System(s), IBM System(s), or Intel Optane System(s), which each teaches memory architectures
18 and modes of memory operation (*see* charts 753-08, 753-12, 753-13, 753-14), and with Saulsbury
19 or Cullison, which each teaches the selective operation of memory in different configurations. *See*
20 Saulsbury, Abstract, ¶¶ [0003], [0005]-[0006], [0008], [0017], [0020], [0022], [0024]-[0026],
21 [0068]-[0070], Figs. 1 and 5; Cullison, 1:11-23, 1:41-63, 2:18-32, 2:36-48, 2:58-62, 2:63-3:21,
22 FIG. 1. A POSITA would naturally have looked to these references and been motivated to
23 combine them with Jana because, e.g., the combinations would provide for more flexible memory
24 arrangements to account for different usage models and applications; reduce memory latency for
25 applications where latency a the bottleneck on performance; improve memory throughput for
26 applications where bandwidth is a bottleneck on performance; improve efficiency and utilization
27 of available memory; and increase performance of distributed computing systems.

- 1 • **Pierson + Stroud**
- 2 • **Pierson + Stroud + Cragon**
- 3 • **Pierson + Stroud + Ramanujan**
- 4 • **Pierson + Stroud + Dell PC System(s)**
- 5 • **Pierson + Stroud + IBM System(s)**
- 6 • **Pierson + Stroud + Intel Optane System(s)**
- 7 • **Any of the above + Saulsbury or Cullison**

8 Pierson discloses a multi-card architecture for parallel, distributed processing in a
9 networked environment, with details and features for connecting cards using a backplane and other
10 networks. *See* Pierson, Abstract, 3:35-4:2, 4:3-26, 7:1-32, 14:50-64, 15:2-13, 19:40-53, 20:46-55,
11 27:44-28:15, Figs. 1, 2, 5, 8. A POSITA looking to implement a multi-card chassis architecture
12 for distributed computing or packet processing would naturally have looked to other references
13 that provide details for chassis-based architectures, cards, and networks in distributed computing.
14 For example, a POSITA would have been motivated to combine Pierson with Stroud, which
15 teaches a multi-card chassis architecture for network processing and testing, with details and
16 features regarding the configuration of chassis systems and the packet format and headers for
17 systems having multiple cards packet format and headers for routing packets in a network
18 environment including a multi-card chassis architecture. *See* Stroud, ¶¶ [0098]-[0102], [0322]-
19 [0342] Figs. 3, 54, 55A, 55B, 56A. A POSITA would naturally have looked to Stroud and been
20 motivated to combine Stroud with Pierson because, e.g., the combination would have provided
21 additional implementation details for chassis architectures in a networked setting; facilitated
22 communication between cards and processors, allowed more granular access to particular
23 components; reduced network overhead in distributed computing systems with multiple cards and
24 chassis; and improved network efficiency in such systems.

25 Additionally, Pierson supports an assortment of memory functions and types, including
26 cache, non-volatile memory, and DRAM. *See* Pierson 20:56-21:23, 21:23-27, 21:60-22:3, 22:20-
27 38, FIG. 4. A POSITA would naturally have looked to references that provide features for using
28

1 different types of memory together. For example, a POSITA would have been motivated to
2 combine Pierson with Cragon, Ramanujan, Dell PC System(s), IBM System(s), or Intel Optane
3 System(s), which each teaches memory architectures and modes of memory operation (*see* charts
4 753-08, 753-12, 753-13, 753-14), and with Saulsbury or Cullison, which each teaches the selective
5 operation of memory in different configurations. *See* Saulsbury, Abstract, ¶¶ [0003], [0005]-
6 [0006], [0008], [0017], [0020], [0022], [0024]-[0026], [0068]-[0070], Figs. 1 and 5; Cullison,
7 1:11-23, 1:41-63, 2:18-32, 2:36-48, 2:58-62, 2:63-3:21, FIG. 1. A POSITA would naturally have
8 looked to these references and been motivated to combine them with the Pierson because, e.g., the
9 combinations would provide for more flexible memory arrangements to account for different
10 usage models and applications; reduce memory latency for applications where latency is a
11 bottleneck on performance; improve memory throughput for applications where bandwidth is a
12 bottleneck on performance; improve efficiency and utilization of available memory; and increase
13 performance of distributed computing systems.

- 14 • **IBM System(s)**
- 15 • **IBM System(s) + RFC 791**
- 16 • **IBM System(s) + Intel Optane System(s)**
- 17 • **IBM System(s) + Dell PC System(s)**
- 18 • **IBM System(s) + Ramanujan**
- 19 • **Any of the above + Cullison or Saulsbury**

20 A POSITA would have been motivated to combine IBM System(s) with RFC 791. The
21 IBM System(s) discloses a computer server architecture, with teaching or suggestion that the
22 network uses IP packets over Ethernet. RFC 791 is the widely-known standard governing the IP
23 protocol. Therefore, a POSITA would have known to use RFC 791 and been motivated to use the
24 IP protocol with IBM System(s). For decades, the IP protocol has been ubiquitous in computers
25 and computer networks—including for servers and computers such as IBM System(s). *See* RFC
26 791 at 1-3, 5-6, 8-9, and 11-23. Therefore, a POSITA would have looked to RFC 791 and been
27 motivated to combine it with IBM System(s) packet-based network communications to facilitate
28 packet-based communications over IP and Ethernet.

1 Additionally, IBM System(s) support an assortment of memory functions and types,
2 including DDR3; L1, L2, L3, and L4 cache; flash memory; hard disks; etc. *See, e.g.*, Exhibit 753-
3 13. A POSITA would have looked to references that provide features for using different types of
4 memory together. For example, a POSITA would have been motivated to combine IBM System(s)
5 with Ramanujan, which teaches a memory hierarchy and different configurable modes of memory
6 operation. *See* Ramanujan at 3:32-35, 7:45-54, 8:11-29, 8:30-37, 9:1-10, 9:27-31, 9:34-40, 10:31-
7 38, 13:3-9, 17:59-18:6, 31:5-12, 31:13-19, 31:29-32, FIG. 1, FIG. 2, FIG. 5B, FIG. 6D. Likewise,
8 Cragon teaches multiple memory arrangements and modes of memory operation. *See* Cragon at
9 73, 91, 110, 191, 231-232, 257, 414-415, 491. A POSITA would have also been motivated to look
10 at the memory arrangements of existing computer systems at the time, including Dell PC
11 System(s), which teaches multiple memory arrangements, including for DDR3/DDR4; L1, L2, and
12 L3 cache; flash memory; hard disks; etc. *See* Exhibit 753-12. Additionally, given these various
13 memory arrangements, a POSITA would have been motivated to utilize the teachings of Saulsbury
14 or Cullison, which each teach the selective operation of memory in different configurations. *See*
15 Saulsbury, Abstract, ¶¶ [0003], [0005]-[0006], [0008], [0017], [0020], [0022], [0024]-[0026],
16 [0068]-[0070], Figs. 1 and 5; Cullison, 1:11-23, 1:41-63, 2:18-32, 2:36-48, 2:58-62, 2:63-3:21,
17 FIG. 1. A POSITA would naturally have looked to these references and been motivated to
18 combine them with the Cray Supercomputer System(s) because, e.g., the combinations would
19 provide for more flexible memory arrangements to account for different usage models and
20 applications; reduce memory latency for applications where latency a the bottleneck on
21 performance; improve memory throughput for applications where bandwidth is a bottleneck on
22 performance; improve efficiency and utilization of available memory; and increase performance of
23 distributed computing systems.

24 **D. U.S. Patent No. 10,985,943**

25 The following table identifies the prior art for the '943 patent. Each item below constitutes
26 prior art under one or more provisions of 35 U.S.C. §§ 102(a)(1)-(2) by reason of its status as a
27 publication and/or patent disclosing, or a system on sale and/or in public knowledge or use in the
28 United States embodying the claimed invention before Plaintiffs' asserted priority date of the '943

1 patent. For each item of prior art, Defendants have attached a claim chart identifying where each
2 claim limitation is found in the cited prior art.

3 **Table D: Prior Art to the '943 Patent**

4	Prior Art to '943	Anticipation/ Obviousness	Claim Chart
5	U.S. Patent No. 5,784,577 to Jacobson et al.		
6	("Jacobson"); filed on August 5, 1996; issued on	O	Exhibit 943-01
7	July 21, 1998.		
8	U.S. Patent No. 6,429,682 to Schultz et al.		
9	("Schultz"); filed on May 25, 2001; issued on	O	Exhibit 943-02
10	August 6, 2002.		
11	Marcello Pogliani et al., <i>Quality of Service</i>		
12	<i>Driven Runtime Resource Allocation in</i>	O	Exhibit 943-03
13	<i>Reconfigurable HPC Architectures</i> , in		
14	<i>Proceedings of the IEEE International</i>		
15	<i>Conference on Computational Science and</i>		
16	<i>Engineering (CSE)</i> (2016),		
17	https://ieeexplore.ieee.org/document/7982218 .		
18	U.S. Patent Application Publication No.		
19	2018/0027062A1 to Bernat et al. ("Bernat"); filed	O	Exhibit 943-04
20	on June 30, 2017; published on Jan 25, 2018.		
21	U.S. Patent Application Publication No.		
22	2015/0109024A1 to Abdelfattah et al.	O	Exhibit 943-05
23	("Abdelfattah"); filed on October 22, 2013;		
24	published on April 23, 2015.		
25	U.S. Patent No. 8,461,869 to Rahim et al.		
26	("Rahim"); filed on August 19, 2011; issued on	O	Exhibit 943-06
27	June 11, 2013.		
28	Cervero et al., <i>A Resource Manager for</i>		
	<i>Dynamically Reconfigurable FPGA-Based</i>	O	Exhibit 943-07
	<i>Embedded Systems</i> , in <i>Proceedings of the</i>		
	<i>Euromicro Conference on Digital System Design</i>		
	(2013),		
	https://ieeexplore.ieee.org/document/6628336 .		
	Cray Supercomputer System(s), including but not		
	limited to Cray XC40, Cray XD1, and Cray XE,	A & O	Exhibit 943-08
	were sold, offered for sale, made known, and/or		
	publicly used before July 15, 2018. Cray		
	Supercomputer System(s) were sold, offered for		

Prior Art to '943	Anticipation/ Obviousness	Claim Chart
sale, and/or made known to the public and third-party customers before this date, by July 14, 2018 at the latest. Cray Supercomputer System(s) were publicly used by HPE/Cray, the public, and third-party customers before this date, by July 14, 2018 at the latest.		
AWS EC2 System, including at least F1 instances, was sold, offered for sale, made known, and/or publicly used before July 15, 2018. AWS EC2 System was sold, offered for sale, and/or made known to the public and third-party customers before this date, by July 14, 2018 at the latest. AWS EC2 System was publicly used by Amazon Web Services, the public, and third-party customers before this date, by July 14, 2018 at the latest.	A & O	Exhibit 943-09
MS Cloud Platform Solutions System(s), including at least Open CloudServers, Azure cloud servers, Catapult v2, v3, and Project BrainWave Storm Peak. were sold, offered for sale, made known, and/or publicly used before July 15, 2018. MS Cloud Platform Solutions System(s) were sold, offered for sale, and/or made known to the public and third-party customers before this date, by July 14, 2018 at the latest. MS Cloud Platform Solutions System(s) were publicly used by Microsoft, the public, and third-party customers before this date, by July 14, 2018 at the latest.	A & O	Exhibit 943-10
UToronto SAVI System was sold, offered for sale, made known, and/or publicly used before July 15, 2018. UToronto SAVI System was sold, offered for sale, and/or made known to the public and third-party customers before this date, by July 14, 2018 at the latest. UToronto SAVI System was publicly used by the University of Toronto, the public, and third-party customers before this date, by July 14, 2018 at the latest.	A & O	Exhibit 943-11
IBM HPC System(s), including at least SuperVessel Cloud systems, Power8 CAPI systems, and cloudFPGA systems, were sold, offered for sale, made known, and/or publicly used before July 15, 2018. IBM HPC System(s) were sold, offered for sale, and/or made known to the public and third-party customers before this	A & O	Exhibit 943-12

Prior Art to '943	Anticipation/ Obviousness	Claim Chart
date, by July 14, 2018 at the latest. IBM HPC System(s) were publicly used by IBM, the public, and third-party customers before this date, by July 14, 2018 at the latest.		
Intel HARP/PAC System(s), including at least Intel HARP v1, Intel HARP v2, Intel PAC with Arria 10 GX FPGA, and MAP-80/100, were sold, offered for sale, made known, and/or publicly used before July 15, 2018. Intel HARP/PAC System(s) were sold, offered for sale, and/or made known to the public and third-party customers before this date, by July 14, 2018 at the latest. Intel HARP/PAC System(s) were publicly used by Intel/Altera, the public, and third-party customers before this date, by July 14, 2018 at the latest.	A & O	Exhibit 943-13
Xilinx Zynq Ultrascale+ System(s), including at least XCZU9EG, were sold, offered for sale, made known, and/or publicly used before July 15, 2018. Xilinx Zynq Ultrascale+ System(s) were sold, offered for sale, and/or made known to the public and third-party customers before this date, at least by Nov. 31, 2015. AMD-CV-0149659 (sales data for XCZU9EG). Xilinx Zynq Ultrascale+ System(s) were publicly used by Xilinx, the public, and third-party customers before this date, at least by Nov. 31, 2015.	A & O	Exhibit 943-14
U.S. Patent No. 10,776,142 to Seshadri et al. ("Seshadri"); filed on December 13, 2017; issued on September 15, 2020.	A & O	See IPR2025-00223
U.S. Patent Application Publication No. 2019/0095113A1 to Atsatt et al. ("Atsatt-113"); filed on September 28, 2017; published on Jan 25, 2018.	A & O	
U.S. Patent Application Publication No. 2018/0191642A1 to Biederman et al. ("Biederman"); filed on December 30, 2016; published on July 5, 2018.	O	
Naif Tarafdar et al., <i>Enabling Flexible Network FPGA Clusters in a Heterogeneous Cloud Data Center</i> , (Feb. 22–24, 2017),	O	

Prior Art to '943	Anticipation/ Obviousness	Claim Chart
https://tarafdar.github.io/files/papers/2017-02-22-enabling-flexible-network-fpga-clusters-in-a-heterogeneous-cloud-data-center.pdf .		
U.S. Patent No. 8,015,386 to Kulkarni et al. ("Kulkarni"); filed on March 31, 2008; issued on September 6, 2011.	O	
U.S. Patent Application Publication No. 2019/0041923A1 to Atsatt et al. ("Atsatt-923"); filed on December 27, 2017; published on February 7, 2019.	O	

In addition to the anticipation and obviousness explained in this section, the '943 patent is anticipated and/or rendered obvious by prior art as explained by Defendants' Petition for Inter Partes Review (assigned Case No. IPR2025-00223), which is incorporated herein by reference:

Grounds	Claims	Prior Art
1	1-5, 10-14	Seshadri
2	1-5, 9-14, 18	Seshadri further in view of Biederman
3	1-5, 10-14	Atsatt-113
4	1-5, 9-14, 18	Atsatt-113 further in view of Biederman
5	6, 15	Grounds 1-4 further in view of Tarafdar
6	7, 16	Grounds 1-4 further in view of Kulkarni
7	8, 17	Grounds 1-4 further in view of Atsatt-923

1. Anticipation

For each item of prior art identified above as an anticipation reference, Defendants contend that the prior art item anticipates the Asserted Claims of the '943 patent, including for the reasons illustrated in the accompanying claim chart and based on Plaintiffs' infringement allegations and contentions for the '943 patent.

2. Obviousness

A person of ordinary skill in the art would have been motivated to combine any reference in Table D with any other reference(s) in Table D. To the extent that any prior art identified in

1 Table D above is deemed to not anticipate an Asserted Claim, the prior art in Table D may be
2 combined with one or more prior art references in Table D to render the claim obvious. All of
3 these items of prior art are directed to PLDs and data processing systems and propose solutions
4 similar to those discussed in the '943 patent.

5 The following are exemplary combinations of prior art that would have been obvious to a
6 POSITA:

- 7 • **AWS EC2 System**
- 8 • **AWS EC2 System + Jacobson, Schultz, or Seshadri**
- 9 • **Any of the above + Cervero**
- 10 • **Any of the above + Pogliani or Bernat**
- 11 • **Any of the above + Abdelfattah**
- 12 • **Any of the above + Rahim**

13 For the AWS EC2 System, a POSITA would have naturally looked to other references that
14 provide details related to cloud computing and networked servers, virtual machines, the
15 management of distributed processing resources, and the composition and operation of FPGAs, as
16 explained below.

17 The AWS EC2 System discloses a PLD, such as an FPGA, for data processing in data-
18 center servers. *See* Ex. 943-09. For example, AWS EC2 Systems provide F1 instances that allow
19 users to configure and run custom logic designs on FPGAs in the AWS Elastic Compute Cloud.
20 *Id.* The F1 instance FPGAs have a Shell (SH) region and a Custom Logic (CL) region. *Id.* The
21 AWS EC2 System restricts access to the CL; to program the CL, users access the CL through API
22 calls, a command line interface (AWS CLI), a web console (AWS Console), or an SDK, among
23 others. Therefore, a POSITA would have been motivated to combine the AWS EC2 System with
24 known techniques for locking an FPGA against unauthorized access, including but not limited to
25 Jacobson, Schultz, or Seshadri. Jacobson teaches “using a predetermined code access” and “once
26 authorization is confirmed, the PLD is unlocked” for configuration. *See* Ex. 943-01. Schultz
27 teaches a “security control option [that] selectively restricts access to configuration and read
28

1 operations” and the use of “a frame mask register” “thereby allowing selective access to individual
2 groups of configuration data.” *See* Ex. 943-02. Seshadri—which a POSITA would have looked
3 to combine with AWS because it also comes from Amazon—teaches techniques for checking user
4 account permissions before allowing reconfiguration of client-configurable circuits. *See* IPR2025-
5 00223 (Ground 1). A POSITA would naturally have looked to these references and been
6 motivated to combine because it would have increased FPGA security in a networked environment
7 with multiple users loading custom logic designs on shared cloud resources. This would have
8 been the use of a known technique to improve similar devices in the same way.

9 As an additional example, the AWS EC2 System provided PLD/FPGA access, for example
10 via F1 instances, in multi-user environments, with users loading their own custom logic. *See* Ex.
11 943-09. Therefore, a POSITA would have been motivated to combine AWS EC2 with known
12 techniques for utilizing and managing multiple partitions on an FPGA, including support for
13 adjusting the size of partitions, including but not limited to Cervero. Cervero teaches the dynamic
14 allocation, resizing, and combination of multiple partitions within an FPGA. *See* Ex. 943-07. A
15 POSITA would have looked to such references and been motivated to combine because it would
16 have increased the utilization of FPGA resources and reduced costs for the cloud service provider
17 and users by allowing multiple users to load custom logic on the same FPGA, with partitions sized
18 appropriate to the user’s logic design.

19 As an additional example, the AWS EC2 System discloses PLDs, such as FPGAs, for data
20 processing in data-centers. Given the constraints on power, thermal budgets, and network
21 bandwidth in data centers, a POSITA would have looked to combine the AWS EC2 System with
22 known techniques for monitoring/performing analytics of resources and bandwidth utilization on
23 FPGAs, including but not limited to Pogliani or Bernat, which teach such techniques. *See* Ex.
24 943-03, Ex. 943-04. A POSITA would have looked to and been motivated to combine the AWS
25 EC2 System with such references because it would have improved the efficiency of FPGA
26 resource deployment in distributed computing and cloud settings, while improving and
27 maintaining quality of service.

1 Additionally, to further address data center constraints on power, thermal budgets, and
2 network bandwidth, a POSITA would have been motivated to combine the AWS EC2 System
3 with known techniques for scaling the voltage and clock frequency of the FPGA, including but not
4 limited to Rahim. Rahim teaches dynamically regulating frequency of discrete portions in an
5 FPGA for thermal and power management. *See* Exhibit Ex. 943-06. A POSITA would have
6 looked to and been motivated to combine the AWS EC2 System with such references because it
7 would have provided higher performance when needed while reducing power consumption and
8 heat generation when higher performance is not needed.

9 As an additional example, the AWS EC2 System includes FPGAs with various I/O ports
10 for communications. *See* Ex. 943-09. Therefore, a POSITA would have been motivated to
11 combine the AWS EC2 with known techniques for implementing and optimizing communication
12 ports on an FPGA, including but not limited to Abdelfattah. Abdelfattah teaches techniques for
13 using ports with adjustable data width and speed in FPGAs. *See* Ex. 943-05. A POSITA would
14 have looked to and been motivated to combine such references with the AWS EC2 System
15 because it would have provided a flexible means to carry out internal communications, scale
16 bandwidth when needed, and increase the performance of FPGAs.

17 Additionally, a POSITA would have been motivated to apply teachings from Biederman,
18 Tarafdar, Kulkarni, and/or Atsatt-923 to any of the above-described combinations involving the
19 AWS EC2 System. The Petition in IPR2025-00223 explained the motivation to combine
20 Biederman, Tarafdar, Kulkarni, and/or Atsatt-923 with known PLD/FPGA configurations,
21 including for use in data centers. The AWS EC2 System used PLD/FPGA configurations in data
22 centers for cloud computing. *See* Ex. 943-09. Therefore, the motivations to combine explained in
23 the incorporated Petition (IPR2025-00223) also apply to the AWS EC2 System, and POSITA
24 would have been motivated to combine for similar reasons.

25 The following are exemplary combinations of prior art that would have been obvious to a
26 POSITA:

- 27 • **MS Cloud Platform Solutions System(s)**
- 28 • **MS Cloud Platform Solutions System(s) + Jacobson, Schultz, or Atsatt-113**

- 1 • **Any of the above + Cervero**
- 2 • **Any of the above + Pogliani or Bernat**
- 3 • **Any of the above + Abdelfattah**
- 4 • **Any of the above + Rahim**

5 For the MSCPS System(s), a POSITA would have naturally looked to other references that
6 provide details related to cloud computing and networked servers, virtual machines, the
7 management of distributed processing resources, and the composition and operation of FPGAs, as
8 explained below.

9 The MSCPS System(s) discloses a PLD, such as an FPGA, for data processing in data-
10 center servers. *See* Ex. 943-10. For example, MSCPS System(s) provide FPGAs on NICs and
11 accelerator cards in cloud servers. *Id.* The FPGAs have a Shell region and a Role region. *Id.* The
12 MSCPS System(s) restricts access to the Role; to program the Role, users access reconfiguration
13 functions through the shell and must log in to servers hosting the cards. Therefore, a POSITA
14 would have been motivated to combine the MSCPS System(s) with known techniques for locking
15 an FPGA against unauthorized access, including but not limited to Jacobson, Schultz, or Atsatt-
16 113. Jacobson teaches “using a predetermined code access” and “once authorization is confirmed,
17 the PLD is unlocked” for configuration. *See* Ex. 943-01. Schultz teaches a “security control
18 option [that] selectively restricts access to configuration and read operations” and the use of “a
19 frame mask register” “thereby allowing selective access to individual groups of configuration
20 data.” *See* Ex. 943-02. Atsatt-113 teaches techniques for checking user account permissions
21 before isolating/deisolating a configuration port to allow reconfiguration of user partitions. *See*
22 IPR2025-00223 (Ground 3). Moreover, MSCPS uses Intel/Altera FPGAs in its servers, and
23 Atsatt-113 is a patent from Intel/Altera. A POSITA would naturally have looked to these
24 references and been motivated to combine because it would have increased FPGA security in a
25 networked environment with multiple users loading custom logic designs on shared cloud
26 resources. This would have been applying a known technique to a known device ready for
27 improvement to yield the predictable results of checking user permissions to increase security.
28

1 As an additional example, the MSCPS System(s) provided PLD/FPGA access in multi-
2 user environments with users loading their own custom logic, including for example on the
3 Azure/Catapult SmartNICs. *See* Ex. 943-10. Therefore, a POSITA would have been motivated to
4 combine MSCPS with known techniques for utilizing and managing multiple partitions on an
5 FPGA, including support for adjusting the size of partitions, including but not limited to Cervero.
6 Cervero teaches the dynamic allocation, resizing, and combination of multiple partitions within an
7 FPGA. *See* Ex. 943-07. A POSITA would have looked to such references and been motivated to
8 combine because it would have increased the utilization of FPGA resources and reduced costs for
9 the cloud service provider and users.

10 As an additional example, the MSCPS System(s) discloses PLDs, such as FPGAs, for data
11 processing in data-centers. Given the constraints on power, thermal budgets, and network
12 bandwidth in data centers, a POSITA would have looked to combine the MSCPS System(s) with
13 known techniques for monitoring/performing analytics of resources and bandwidth utilization on
14 FPGAs, including but not limited to Pogliani or Bernat, which teach such techniques. *See* Ex.
15 943-03, Ex. 943-04. A POSITA would have looked to and been motivated to combine the
16 MSCPS System(s) with such references because it would have improved the efficiency of FPGA
17 resource deployment in distributed computing and cloud settings, while improving and
18 maintaining quality of service.

19 Additionally, to further address data center constraints on power, thermal budgets, and
20 network bandwidth, a POSITA would have been motivated to combine the MSCPS System(s)
21 with known techniques for scaling the voltage and clock frequency of the FPGA, including but not
22 limited to Rahim. Rahim teaches dynamically regulating frequency of discrete portions in an
23 FPGA for thermal and power management. *See* Exhibit Ex. 943-06. Moreover, MSCPS uses
24 Intel/Altera FPGAs in its servers, and Rahim is a patent from Intel/Altera. A POSITA would have
25 looked to and been motivated to combine the MSCPS System(s) with such references because it
26 would have provided higher performance when needed while reducing power consumption and
27 heat generation when higher performance is not needed.

1 As an additional example, the MSCPS System(s) includes FPGAs with various
2 communication ports between the Shell and Role and between the Shell and external devices. *See*
3 Ex. 943-10. Therefore, a POSITA would have been motivated to combine the MSCPS System(s)
4 with known techniques for implementing and optimizing communication ports on an FPGA,
5 including but not limited to Abdelfattah. Abdelfattah teaches techniques for using ports with
6 adjustable data width and speed in FPGAs. *See* Ex. 943-05. A POSITA would have looked to and
7 been motivated to combine such references with the MSCPS System(s) because it would have
8 provided a flexible means to carry out internal communications, scale bandwidth when needed,
9 and increase the performance of FPGAs.

10 Additionally, a POSITA would have been motivated to apply teachings from Biederman,
11 Tarafdar, Kulkarni, and/or Atsatt-923 to any of the above-described combinations involving the
12 MSCPS System(s). The Petition in IPR2025-00223 explained the motivation to combine
13 Biederman, Tarafdar, Kulkarni, and/or Atsatt-923 with known PLD/FPGA configurations,
14 including for use in data centers. The MSCPS System(s) used PLD/FPGA configurations in data
15 centers for cloud computing. *See* Ex. 943-10. Therefore, the motivations to combine explained in
16 the incorporated Petition (IPR2025-00223) also apply to the MSCPS System(s), and POSITA
17 would have been motivated to combine for similar reasons.

18 The following are exemplary combinations of prior art that would have been obvious to a
19 POSITA:

- 20 • **SAVI System + Jacobson, Schultz, or Seshadri**
- 21 • **Any of the above + Cervero**
- 22 • **Any of the above + Pogliani or Bernat**
- 23 • **Any of the above + Abdelfattah**
- 24 • **Any of the above + Rahim**

25 For the SAVI System, a POSITA would have naturally looked to other references that
26 provide details related to cloud computing and networked servers, virtual machines, the
27
28

1 management of distributed processing resources, and the composition and operation of FPGAs, as
2 explained below.

3 The SAVI System discloses a PLD, such as an FPGA, for data processing in data-center
4 servers. *See* Ex. 943-11. For example, SAVI System provide FPGAs as accelerators as edge node
5 resources in a cloud. *Id.* The FPGAs have an Abstraction Shell region and a FPGA Kernel
6 region. *Id.* The SAVI System restricts access to the FPGA Kernel; to program user logic to a
7 Kernel, users access the Kernel through the virtual machine using, for example, Openstack APIs.
8 Therefore, a POSITA would have been motivated to combine the SAVI System with known
9 techniques for locking an FPGA against unauthorized access, including but not limited to
10 Jacobson, Schultz, or Seshadri. Jacobson teaches “using a predetermined code access” and “once
11 authorization is confirmed, the PLD is unlocked” for configuration. *See* Ex. 943-01. Schultz
12 teaches a “security control option [that] selectively restricts access to configuration and read
13 operations” and the use of “a frame mask register” “thereby allowing selective access to individual
14 groups of configuration data.” *See* Ex. 943-02. Seshadri teaches techniques for checking user
15 account permissions before allowing reconfiguration of client-configurable circuits. *See* IPR2025-
16 00223 (Ground 1). A POSITA would naturally have looked to these references and been
17 motivated to combine because it would have increased FPGA security in a networked environment
18 with multiple users loading custom logic designs on shared cloud resources.

19 As an additional example, the SAVI System discloses multiple users sharing space on a
20 PLD, such as an FPGA, including for example on the NetFPGA cards. *See* Ex. 943-11. The
21 SAVI System allows user logics to be loaded to partial reconfiguration areas on the same FPGA in
22 the cloud. *Id.* Therefore, a POSITA would have been motivated to combine SAVI with known
23 techniques for managing and adjusting multiple partitions on an FPGA, including but not limited
24 to Cervero. Cervero teaches the dynamic allocation, resizing, and combination of multiple
25 partitions within an FPGA. *See* Ex. 943-07. A POSITA would have looked to such references
26 and been motivated to combine because it would have increased the utilization of FPGA resources
27 and reduced costs for the cloud service provider and users.

1 As an additional example, the SAVI System discloses PLDs, such as FPGAs, for data
2 processing in data-centers such as the SAVI Testbed. Given the constraints on power, thermal
3 budgets, and network bandwidth in data centers, a POSITA would have looked to combine the
4 SAVI System with known techniques for monitoring/performing analytics of resources and
5 bandwidth utilization on FPGAs, including but not limited to Pogliani or Bernat, which teach such
6 techniques. *See* Ex. 943-03, Ex. 943-04. A POSITA would have looked to and been motivated to
7 combine the SAVI System with such references because it would have improved the efficiency of
8 FPGA resource deployment in distributed computing and cloud settings, while improving and
9 maintaining quality of service.

10 Additionally, to further address data center constraints on power, thermal budgets, and
11 network bandwidth, a POSITA would have been motivated to combine the SAVI System with
12 known techniques for scaling the voltage and clock frequency of the FPGA, including but not
13 limited to Rahim. Rahim teaches dynamically regulating frequency of discrete portions in an
14 FPGA for thermal and power management. *See* Exhibit Ex. 943-06. A POSITA would have
15 looked to and been motivated to combine the SAVI System with such references because it would
16 have provided higher performance when needed while reducing power consumption and heat
17 generation when higher performance is not needed.

18 As an additional example, the SAVI System includes FPGAs with various I/O ports
19 including PCIe and Ethernet. *See* Ex. 943-11. Therefore, a POSITA would have been motivated
20 to combine the SAVI System with known techniques for implementing and optimizing
21 communication ports on an FPGA, including but not limited to Abdelfattah. Abdelfattah teaches
22 techniques for using ports with adjustable data width and speed in FPGAs. *See* Ex. 943-05. A
23 POSITA would have looked to and been motivated to combine such references with the SAVI
24 System because it would have provided a flexible means to carry out internal communications,
25 scale bandwidth when needed, and increase the performance of FPGAs.

26 Additionally, a POSITA would have been motivated to apply teachings from Biederman,
27 Tarafdar, Kulkarni, and/or Atsatt-923 to any of the above-described combinations involving the
28 SAVI System. The Petition in IPR2025-00223 explained the motivation to combine Biederman,

1 Tarafdar, Kulkarni, and/or Atsatt-923 with known PLD/FPGA configurations, including for use in
2 data centers. The SAVI System used PLD/FPGA configurations in data centers for cloud
3 computing. *See* Ex. 943-09. Therefore, the motivations to combine explained in the incorporated
4 Petition (IPR2025-00223) also apply to the SAVI System, and a POSITA would have been
5 motivated to combine for similar reasons.

6 The following are exemplary combinations of prior art that would have been obvious to a
7 POSITA:

- 8 • **IBM HPC System(s)**
- 9 • **IBM HPC System(s) + Jacobson, Schultz, Seshadri, or Atsatt-113**
- 10 • **Any of the above + Cervero**
- 11 • **Any of the above + Pogliani or Bernat**
- 12 • **Any of the above + Abdelfattah**
- 13 • **Any of the above + Rahim**

14 For the IBM HPC System(s), a POSITA would have naturally looked to other references
15 that provide details related to high performance computing and networked servers, virtual
16 machines, the management of distributed processing resources, and the composition and operation
17 of FPGAs, as explained below.

18 The IBM HPC System(s) discloses a PLD, such as an FPGA, for distributing applications
19 to network-attached FPGAs in data centers. *See* Ex. 943-12. For example, IBM HPC System(s)
20 provide FPGAs in cloudFPGA and Power8+CAPI servers. *Id.* The FPGAs have a Shell region
21 and a Role region. *Id.* The IBM HPC System(s) restricts access to the Role; to program the Role,
22 users access reconfiguration functions through the FPGA Manager Core in the Shell. Therefore, a
23 POSITA would have been motivated to combine the IBM HPC System(s) with known techniques
24 for locking an FPGA against unauthorized access, including but not limited to Jacobson, Schultz,
25 Seshadri, or Atsatt-113. Jacobson teaches “using a predetermined code access” and “once
26 authorization is confirmed, the PLD is unlocked” for configuration. *See* Ex. 943-01. Schultz
27 teaches a “security control option [that] selectively restricts access to configuration and read
28

1 operations” and the use of “a frame mask register” “thereby allowing selective access to individual
2 groups of configuration data.” *See* Ex. 943-02. Seshadri teaches techniques for checking user
3 account permissions before allowing reconfiguration of client-configurable circuits. *See* IPR2025-
4 00223 (Ground 1). Atsatt-113 teaches techniques for checking user account permissions before
5 isolating/deisolating a configuration port to allow reconfiguration of user partitions. *See* IPR2025-
6 00223 (Ground 3). Moreover, IBM HPC provides Intel/Altera FPGAs as an option in its servers,
7 and Atsatt-113 is a patent from Intel/Altera. A POSITA would naturally have looked to these
8 references and been motivated to combine because it would have increased FPGA security in a
9 networked environment with multiple users loading custom logic designs on shared cloud
10 resources. This would have been applying a known technique to a known device ready for
11 improvement to yield the predictable results of checking user permissions to increase security.

12 As an additional example, the IBM HPC System(s) discloses virtualizing the FPGA to
13 allow multiple users to share space for Roles. *See* Ex. 943-12. The IBM HPC System(s) allows
14 custom user logics to be loaded to multiple roles on the same FPGA. *Id.* It provided PLD/FPGA
15 access in multi-user environments, with users loading their own custom logic. *Id.* Therefore, a
16 POSITA would have been motivated to combine IBM HPC with known techniques for utilizing
17 and managing multiple partitions on an FPGA, including support for adjusting the size of
18 partitions, including but not limited to Cervero. Cervero teaches the dynamic allocation, resizing,
19 and combination of multiple partitions within an FPGA. *See* Ex. 943-07. A POSITA would have
20 looked to such references and been motivated to combine because it would have increased the
21 utilization of FPGA resources and reduced costs for the cloud service provider and users.

22 As an additional example, the IBM HPC System(s) discloses PLDs, such as FPGAs, for
23 data processing in data-centers. Given the constraints on power, thermal budgets, and network
24 bandwidth in data centers, a POSITA would have looked to combine the IBM HPC System(s)
25 with known techniques for monitoring/performing analytics of resources and bandwidth utilization
26 on FPGAs, including but not limited to Pogliani or Bernat, which teach such techniques. *See* Ex.
27 943-03, Ex. 943-04. A POSITA would have looked to and been motivated to combine the IBM
28 HPC System(s) with such references because it would have improved the efficiency of FPGA

1 resource deployment in distributed computing and cloud settings, while improving and
2 maintaining quality of service.

3 Additionally, to further address data center constraints on power, thermal budgets, and
4 network bandwidth, a POSITA would have been motivated to combine the IBM HPC System(s)
5 with known techniques for scaling the voltage and clock frequency of the FPGA, including but not
6 limited to Rahim. Rahim teaches dynamically regulating frequency of discrete portions in an
7 FPGA for thermal and power management. *See* Exhibit Ex. 943-06. Moreover, IBM HPC
8 provides Intel/Altera FPGAs as an option its servers, and Rahim is a patent from Intel/Altera. A
9 POSITA would have looked to and been motivated to combine the IBM HPC System(s) with such
10 references because it would have provided higher performance when needed while reducing power
11 consumption and heat generation when higher performance is not needed.

12 As an additional example, the IBM HPC System(s) includes FPGAs with various ports for
13 communication within the FPGA and between FPGA Nodes. *See* Ex. 943-12. Therefore, a
14 POSITA would have been motivated to combine the IBM HPC System(s) with known techniques
15 for implementing and optimizing communication ports on an FPGA, including but not limited to
16 Abdelfattah. Abdelfattah teaches techniques for using ports with adjustable data width and speed
17 in FPGAs. *See* Ex. 943-05. A POSITA would have looked to and been motivated to combine
18 such references with the IBM HPC System(s) because it would have provided a flexible means to
19 carry out internal communications, scale bandwidth when needed, and increase the performance of
20 FPGAs.

21 Additionally, a POSITA would have been motivated to apply teachings from Biederman,
22 Tarafdar, Kulkarni, and/or Atsatt-923 to any of the above-described combinations involving the
23 IBM HPC System(s). The Petition in IPR2025-00223 explained the motivation to combine
24 Biederman, Tarafdar, Kulkarni, and/or Atsatt-923 with known PLD/FPGA configurations,
25 including for use in data centers. The IBM HPC System(s) used PLD/FPGA configurations in
26 data centers for cloud computing. *See* Ex. 943-12. Therefore, the motivations to combine
27 explained in the incorporated Petition (IPR2025-00223) also apply to the IBM HPC System(s),
28 and POSITA would have been motivated to combine for similar reasons.

1 The following are exemplary combinations of prior art that would have been obvious to a
2 POSITA:

- 3 • **Cray Supercomputer System(s)**
- 4 • **Cray Supercomputer System(s) + Jacobson, Schultz, or Seshadri**
- 5 • **Any of the above + Cervero**
- 6 • **Any of the above + Pogliani or Bernat**
- 7 • **Any of the above + Abdelfattah**
- 8 • **Any of the above + Rahim**

9 For the Cray Supercomputer System(s), a POSITA would have naturally looked to other
10 references that provide details related to cloud computing and networked servers, virtual
11 machines, the management of distributed processing resources, and the composition and operation
12 of FPGAs, as explained below.

13 The Cray Supercomputer System(s) discloses a PLD, such as an FPGA, for data processing
14 in high-performance computing systems. *See* Ex. 943-08. For example, Cray Supercomputer
15 System(s) provide FPGAs as “user-programmable accelerators” on blade. *Id.* The FPGAs have a
16 control region encompassing a communication core, a HyperTransport interface, and QDR
17 memory interfaces, and a user configurable region. *Id.* Therefore, given the user of FPGAs in a
18 multiuser environment, a POSITA would have been motivated to combine the Cray
19 Supercomputer System(s) with known techniques for locking an FPGA against unauthorized
20 access, including but not limited to Jacobson, Schultz, or Seshadri. Jacobson teaches “using a
21 predetermined code access” and “once authorization is confirmed, the PLD is unlocked” for
22 configuration. *See* Ex. 943-01. Schultz teaches a “security control option [that] selectively
23 restricts access to configuration and read operations” and the use of “a frame mask register”
24 “thereby allowing selective access to individual groups of configuration data.” *See* Ex. 943-02.
25 Seshadri teaches techniques for checking user account permissions before allowing
26 reconfiguration of client-configurable circuits. *See* IPR2025-00223 (Ground 1). A POSITA
27 would naturally have looked to these references and been motivated to combine because it would
28

1 have increased FPGA security in a networked environment with multiple users loading custom
2 logic designs on shared cloud resources. Indeed, there is teachings, suggestion and motivation in
3 the art for these combinations. For example, the Cray Supercomputer System(s) provide FPGAs
4 on NICs and accelerator cards in cloud servers, where users access reconfiguration functions
5 through the shell and must log in to servers hosting the cards. This would have been a
6 combination of prior art elements according to known methods to yield predictable results.

7 As an additional example, the Cray Supercomputer System(s) discloses custom logics
8 being loaded in parallel on a PLD, such as an FPGA cards. *See* Ex. 943-08. The Cray
9 Supercomputer System(s) allows for application acceleration by supporting multiple instances on
10 the same FPGA. *Id.* It also provided PLD/FPGA access in multi-user environments, with users
11 loading their own custom logic. *Id.* Therefore, a POSITA would have been motivated to combine
12 Cray Supercomputer System(s) with known techniques for managing and adjusting multiple
13 partitions on an FPGA, including but not limited to Cervero. Cervero teaches the dynamic
14 allocation, resizing, and combination of multiple partitions within an FPGA. *See* Ex. 943-07. A
15 POSITA would have looked to such references and been motivated to combine because it would
16 have increased the utilization of FPGA resources and reduced costs for the cloud service provider
17 and users.

18 As an additional example, the Cray Supercomputer System(s) discloses PLDs, such as
19 FPGAs, for data processing in data centers using, for example, its CrayXC or CrayXD systems.
20 Given the constraints on power, thermal budgets, and network bandwidth in data centers, a
21 POSITA would have looked to combine the Cray Supercomputer System(s) with known
22 techniques for monitoring/performing analytics of resources and bandwidth utilization on FPGAs,
23 including but not limited to Pogliani or Bernat, which teach such techniques. *See* Ex. 943-03, Ex.
24 943-04. A POSITA would have looked to and been motivated to combine the Cray
25 Supercomputer System(s) with such references because it would have improved the efficiency of
26 FPGA resource deployment in distributed computing and cloud settings, while improving and
27 maintaining quality of service.

1 Additionally, to further address data center constraints on power, thermal budgets, and
2 network bandwidth, a POSITA would have been motivated to combine the Cray Supercomputer
3 System(s) with known techniques for scaling the voltage and clock frequency of the FPGA,
4 including but not limited to Rahim. Rahim teaches dynamically regulating frequency of discrete
5 portions in an FPGA for thermal and power management. *See* Exhibit Ex. 943-06. A POSITA
6 would have looked to and been motivated to combine the Cray Supercomputer System(s) with
7 such references because it would have provided higher performance when needed while reducing
8 power consumption and heat generation when higher performance is not needed.

9 As an additional example, the Cray Supercomputer System(s) includes FPGAs with
10 various I/O ports including PCIe and Ethernet. *See* Ex. 943-08. Therefore, a POSITA would have
11 been motivated to combine the Cray Supercomputer System(s) with known techniques for
12 implementing and optimizing communication ports on an FPGA, including but not limited to
13 Abdelfattah. Abdelfattah teaches techniques for using ports with adjustable data width and speed
14 in FPGAs. *See* Ex. 943-05. A POSITA would have looked to and been motivated to combine
15 such references with the Cray Supercomputer System(s) because it would have provided a flexible
16 means to carry out internal communications, scale bandwidth when needed, and increase the
17 performance of FPGAs.

18 Additionally, a POSITA would have been motivated to apply teachings from Biederman,
19 Tarafdar, Kulkarni, and/or Atsatt-923 to any of the above-described combinations involving the
20 Cray Supercomputer System(s). The Petition in IPR2025-00223 explained the motivation to
21 combine Biederman, Tarafdar, Kulkarni, and/or Atsatt-923 with known PLD/FPGA
22 configurations, including for use in data centers. The Cray Supercomputer System(s) used
23 PLD/FPGA configurations in data centers for cloud computing. *See* Ex. 943-08. Therefore, the
24 motivations to combine explained in the incorporated Petition (IPR2025-00223) also apply to the
25 Cray Supercomputer System(s), and POSITA would have been motivated to combine for similar
26 reasons.

27 Defendants further incorporate by reference the invalidity grounds described in
28 Defendants' Petition for Inter Partes Review (assigned Case No. IPR2025-00223).

- 1 • **Xilinx Zynq Ultrascale+ System(s)**
- 2 • **Xilinx Zynq Ultrascale+ System(s) + Cray Supercomputer System(s), AWS**
- 3 **EC2 System, MS Cloud Platform Solutions System(s), IBM HPC System(s), or**
- 4 **Bernat**
- 5 • **Any of the above + Jacobson or Schultz**
- 6 • **Any of the above + Cervero**
- 7 • **Any of the above + Pogliani or Bernat**
- 8 • **Any of the above + Abdelfattah**
- 9 • **Any of the above + Rahim**

10 The Xilinx Zynq Ultrascale+ System(s) discloses a use case of its SoCs in advanced data-
11 center applications. *See* Ex. 943-14. Therefore, a POSITA would have been looked to combine
12 the Xilinx Zynq Ultrascale+ System(s) with known configurations of data-centers, where servers
13 typically included swappable pods or cards, for example as taught by Cray Supercomputer
14 System(s), AWS EC2 System, MS Cloud Platform Solutions System(s), IBM HPC System(s), and
15 Bernat. A POSITA would have been motivated to combine Xilinx Zynq Ultrascale+ System(s)
16 into servers with swappable pods or cards because it is suggested by Xilinx Zynq Ultrascale+
17 System(s), and it was known in the art to be advantageous for servers to have modular, expandable
18 designs, so that they can be manufactured at lower cost while allowing for greater flexibility. This
19 would have been a combination of prior art elements according to known methods to yield
20 predictable results.

21 The Xilinx Zynq Ultrascale+ System(s) discloses a system on chip (“SoC”) that includes
22 the programmable logic of an FPGA in one portion and a processing system (“PS”) in another
23 portion. *See* Ex. 943-14. The PS includes a configuration access port that allows software to
24 configure the programmable logic portion. *See id.* Therefore, a POSITA would have been
25 motivated to combine the Xilinx Zynq Ultrascale+ System(s) with known techniques for locking
26 an FPGA against unauthorized access, including but not limited to Jacobson, Schultz, Seshadri, or
27 Atsatt-113. For example, Jacobson teaches “using a predetermined code access” and “once
28 authorization is confirmed, the PLD is unlocked” for configuration. *See* Ex. 943-01. Schultz

1 teaches a “security control option [that] selectively restricts access to configuration and read
2 operations” and the use of “a frame mask register” “thereby allowing selective access to individual
3 groups of configuration data.” *See* Ex. 943-02. Moreover, Schultz is a patent from Xilinx
4 disclosing techniques and configurations for use in PLDs, such as the Xilinx Zynq Ultrascale+
5 System(s). Seshadri teaches techniques for checking user account permissions before allowing
6 reconfiguration of client-configurable circuits. *See* IPR2025-00223 (Ground 1). Atsatt-113
7 teaches techniques for checking user account permissions before isolating/deisolating a
8 configuration port to allow reconfiguration of user partitions. *See* IPR2025-00223 (Ground 3). A
9 POSITA would naturally have looked to these references and been motivated to combine because
10 it would have increased FPGA security when it is reprogrammed using software, which would
11 have been beneficial for workstation, networked environments, and a variety of different use
12 cases. Xilinx Zynq Ultrascale+ System(s) further suggests the combination to add such security
13 features by disclosing that security was a major concern in the design and operation of such
14 FPGA-based programmable logic devices. Ronaldo Architecture Specifications (AMD-CV-
15 0137948–AMD-CV-0138306), §10 (discussing architecture goals and security measures);
16 UG1228 Zynq UltraScale+ MPSoC Embedded Methodology Guide (AMD-CV-0119344–AMD-
17 CV-0119560), §8 (discussing security needs and security methodology). This would have been a
18 combination of prior art elements according to known methods to yield predictable results.

19 As an additional example, the Xilinx Zynq Ultrascale+ System(s) incorporated a PS to
20 allow greater flexibility and use cases, including for applications that needed heterogenous
21 processing using different types of hardware. Therefore, a POSITA would have been motivated to
22 combine Xilinx Zynq Ultrascale+ System(s) with known techniques for utilizing and managing
23 multiple partitions on an FPGA, including support for adjusting the size of partitions, including
24 but not limited to Cervero. Cervero teaches the dynamic allocation, resizing, and combination of
25 multiple partitions within an FPGA. *See* Ex. 943-07. A POSITA would have looked to such
26 references and been motivated to combine because it would have increased the utilization of
27 FPGA resources and allowed different sections of the programmable logic to be used for different
28 types of hardware acceleration, which furthers the goal of providing heterogenous computing

1 resources. This would have been the use of a known technique to improve similar devices in the
2 same way.

3 As an additional example, the Xilinx Zynq Ultrascale+ System(s) discloses a high-
4 performance SoC for data processing in data-centers. Given the constraints on power, thermal
5 budgets, and network bandwidth in the relevant timeframe—which had been widely recognized as
6 major constraints on computer performance by that time—a POSITA would have looked to
7 combine the Xilinx Zynq Ultrascale+ System(s) with known techniques for
8 monitoring/performing analytics of resources and bandwidth utilization on FPGAs, including but
9 not limited to Pogliani or Bernat, which teach such techniques. *See* Ex. 943-03, Ex. 943-04. A
10 POSITA would have looked to and been motivated to combine the Xilinx Zynq Ultrascale+
11 System(s) with such references because it would have improved the power and thermal constraints
12 on performance of the FPGA and increased efficiency of FPGA resources, while improving and
13 maintaining quality of service.

14 Additionally, to further address data center constraints on power, thermal budgets, and
15 network bandwidth, a POSITA would have been motivated to combine the Xilinx Zynq
16 Ultrascale+ System(s) with known techniques for scaling the voltage and clock frequency of the
17 FPGA, including but not limited to Rahim. Rahim teaches dynamically regulating frequency of
18 discrete portions in an FPGA for thermal and power management. *See* Exhibit Ex. 943-06. A
19 POSITA would have looked to and been motivated to combine the Xilinx Zynq Ultrascale+
20 System(s) with such references because it would have provided higher performance when needed
21 while reducing power consumption and heat generation when higher performance is not needed.

22 As an additional example, the Xilinx Zynq Ultrascale+ System(s) includes an SoC/FPGA
23 with various I/O ports for communications. Therefore, a POSITA would have been motivated to
24 combine the Xilinx Zynq Ultrascale+ System(s) with known techniques for implementing and
25 optimizing communication ports on an FPGA, including but not limited to Abdelfattah.
26 Abdelfattah teaches techniques for using ports with adjustable data width and speed in FPGAs.
27 *See* Ex. 943-05. A POSITA would have looked to and been motivated to combine such references
28 with the Xilinx Zynq Ultrascale+ System(s) because it would have provided a flexible means to

1 carry out internal communications, scale bandwidth when needed, and increase the performance of
2 FPGAs.

3 Additionally, a POSITA would have been motivated to apply teachings from Biederman,
4 Tarafdar, Kulkarni, and/or Atsatt-923 to any of the above-described combinations involving the
5 Xilinx Zynq Ultrascale+ System(s). The Petition in IPR2025-00223 explained the motivation to
6 combine Biederman, Tarafdar, Kulkarni, and/or Atsatt-923 with known PLD/FPGA
7 configurations. The Xilinx Zynq Ultrascale+ System(s) was an FPGA-based SoC. Therefore, the
8 motivations to combine explained in the incorporated Petition (IPR2025-00223) also apply to the
9 Xilinx Zynq Ultrascale+ System(s), and POSITA would have been motivated to combine for
10 similar reasons.

- 11 • **Intel HARP/PAC System(s)**
- 12 • **Intel HARP/PAC System(s) + Cray Supercomputer System(s), AWS EC2**
13 **System, MS Cloud Platform Solutions System(s), IBM HPC System(s), or**
14 **Bernat**
- 15 • **Any of the above + Jacobson or Schultz**
- 16 • **Any of the above + Cervero**
- 17 • **Any of the above + Pogliani or Bernat**
- 18 • **Any of the above + Abdelfattah**
- 19 • **Any of the above + Rahim**

20 The Intel HARP/PAC System(s) discloses accelerator cards/boards that were often used in
21 data-center applications. Therefore, a POSITA would have been looked to combine the Intel
22 HARP/PAC System(s) with known configurations of data-centers, where servers typically
23 included swappable pods or cards, for example as taught by Cray Supercomputer System(s), AWS
24 EC2 System, MS Cloud Platform Solutions System(s), IBM HPC System(s), and Bernat. A
25 POSITA would have been motivated to combine Intel HARP/PAC System(s) into servers with
26 swappable pods or cards because it is suggested by Intel HARP/PAC System(s), and it was known
27 in the art to be advantageous for servers to have modular, expandable designs, so that they can be
28

1 manufactured at lower cost while allowing for greater flexibility. This would have been a
2 combination of prior art elements according to known methods to yield predictable results.

3 The Intel HARP/PAC System(s) discloses accelerator card or boards that include an FPGA
4 with multiple Accelerator Functional Units (AFU) loaded to Partial Reconfiguration (PR)
5 partitions, encapsulated by an FPGA Interface Manager (FIM). *See* Ex. 943-13. The FIM
6 includes a port and bridge that controls access to configure the PR partitions, e.g., by loading
7 different AFUs. *See id.* Therefore, a POSITA would have been motivated to combine the Intel
8 HARP/PAC System(s) with known techniques for locking an FPGA against unauthorized access,
9 including but not limited to Jacobson, Schultz, Seshadri, or Atsatt-113. For example, Jacobson
10 teaches “using a predetermined code access” and “once authorization is confirmed, the PLD is
11 unlocked” for configuration. *See* Ex. 943-01. Schultz teaches a “security control option [that]
12 selectively restricts access to configuration and read operations” and the use of “a frame mask
13 register” “thereby allowing selective access to individual groups of configuration data.” *See* Ex.
14 943-02. Seshadri teaches techniques for checking user account permissions before allowing
15 reconfiguration of client-configurable circuits in FPGAs. *See* IPR2025-00223 (Ground 1). Atsatt-
16 113 teaches techniques for checking user account permissions before isolating/deisolating a
17 configuration port to allow reconfiguration of user partitions. *See* IPR2025-00223 (Ground 3).
18 Moreover, Intel HARP/PAC System(s) are Intel products that used Intel/Altera FPGAs and
19 Atsatt-113 is a patent from Intel/Altera. A POSITA would naturally have looked to these
20 references and been motivated to combine because it would have increased FPGA security when it
21 is reprogrammed using software, which would have been beneficial for workstation, networked
22 environments, and a variety of different use cases. This would have been a combination of prior
23 art elements according to known methods to yield predictable results.

24 As an additional example, the Intel HARP/PAC System(s) discloses the flexible
25 assignment of multiple PR partitions, in multi-user environments, with users loading their own
26 custom logic AFUs. *See* Ex. 943-13. Therefore, a POSITA would have been motivated to
27 combine Intel HARP/PAC System(s) with known techniques for utilizing and managing multiple
28 partitions on an FPGA, including support for adjusting the size of partitions, including but not

1 limited to Cervero. Cervero teaches the dynamic allocation, resizing, and combination of multiple
2 partitions within an FPGA. *See* Ex. 943-07. A POSITA would have looked to such references
3 and been motivated to combine because it would have increased the utilization of FPGA resources
4 and reduced costs for the cloud service provider and users by allowing multiple users to load
5 custom logic on the same FPGA, with partitions sized appropriate to the user's logic design.
6 Additionally, the Intel HARP/PAC System(s) discloses heterogenous processing platforms with
7 FPGAs and Xeon processors. The combination with Cervero would have allowed different
8 sections of the programmable logic to be used for different types of hardware acceleration, which
9 furthers the goal of providing heterogenous computing resources. This would have been the use
10 of a known technique to improve similar devices in the same way.

11 As an additional example, the Intel HARP/PAC System(s) discloses a card for data
12 processing in data-centers. Given the constraints on power, thermal budgets, and network
13 bandwidth in the relevant timeframe—which had been widely recognized as major constraints on
14 computer performance by that time—a POSITA would have looked to combine the Intel
15 HARP/PAC System(s) with known techniques for monitoring/performing analytics of resources
16 and bandwidth utilization on FPGAs, including but not limited to Pogliani or Bernat, which teach
17 such techniques. *See* Ex. 943-03, Ex. 943-04. Additionally, Bernat is a patent application from
18 Intel describing similar data processing platforms to the Intel HARP/PAC System(s). A POSITA
19 would have looked to and been motivated to combine the Intel HARP/PAC System(s) with such
20 references because it would have improved the power and thermal constraints on performance of
21 the FPGA and increased efficiency of FPGA resources, while improving and maintaining quality
22 of service.

23 Additionally, to further address data center constraints on power, thermal budgets, and
24 network bandwidth, a POSITA would have been motivated to combine the Intel HARP/PAC
25 System(s) with known techniques for scaling the voltage and clock frequency of the FPGA,
26 including but not limited to Rahim. Rahim teaches dynamically regulating frequency of discrete
27 portions in an FPGA for thermal and power management. *See* Exhibit Ex. 943-06. Moreover,
28 Rahim is an Intel/Altera patent describing techniques developed for application to Intel/Altera

1 FPGAs. A POSITA would have looked to and been motivated to combine the Intel HARP/PAC
2 System(s) with such references because it would have provided higher performance when needed
3 while reducing power consumption and heat generation when higher performance is not needed.

4 As an additional example, the Intel HARP/PAC System(s) includes an SoC/FPGA with
5 various I/O ports for communications. Therefore, a POSITA would have been motivated to
6 combine the Intel HARP/PAC System(s) with known techniques for implementing and optimizing
7 communication ports on an FPGA, including but not limited to Abdelfattah. Abdelfattah teaches
8 techniques for using ports with adjustable data width and speed in FPGAs. *See* Ex. 943-05. A
9 POSITA would have looked to and been motivated to combine such references with the Intel
10 HARP/PAC System(s) because it would have provided a flexible means to carry out internal
11 communications, scale bandwidth when needed, and increase the performance of FPGAs.

12 Additionally, a POSITA would have been motivated to apply teachings from Biederman,
13 Tarafdar, Kulkarni, and/or Atsatt-923 to any of the above-described combinations involving the
14 Intel HARP/PAC System(s). The Petition in IPR2025-00223 explained the motivation to combine
15 Biederman, Tarafdar, Kulkarni, and/or Atsatt-923 with known PLD/FPGA configurations. The
16 Intel HARP/PAC System(s) was an FPGA-based SoC. Therefore, the motivations to combine
17 explained in the incorporated Petition (IPR2025-00223) also apply to the Intel HARP/PAC
18 System(s), and POSITA would have been motivated to combine for similar reasons.

19 **E. U.S. Patent No. 10,944,634**

20 The following table identifies the prior art for the '634 patent. Each item below constitutes
21 prior art under one or more provisions of 35 U.S.C. §§ 102(a)(1)-(2) by reason of its status as a
22 publication and/or patent disclosing, or a system on sale and/or in public knowledge or use in the
23 United States embodying the claimed invention before Plaintiffs' asserted priority date of the '634
24 patent. For each item of prior art, Defendants have attached a claim chart identifying where each
25 claim limitation is found in the cited prior art.

Table E Prior Art to '634 Patent

Prior Art to '634	Anticipation/ Obviousness	Claim Chart
Verizon Communications Inc.'s networking devices ("Verizon Networking Devices"), including but not limited to Verizon switches, routers, gateways, and access points, were sold, offered for sale, made known, and/or publicly used before July 15, 2018. Verizon Networking Devices were sold, offered for sale, and/or made known to the public and third-party customers before this date, by July 14, 2018 at the latest. Verizon Networking Devices were publicly used by Verizon, the public, and third-party customers before this date, by July 14, 2018 at the latest.	A & O	Exhibit 634-01
Verizon Communications Inc.'s networking solutions ("Verizon Networking Solutions"), including but not limited to Verizon Managed SD WAN, Verizon Managed LAN Services, Verizon Managed SD Branch, and Verizon Secure Gateway Services, were sold, offered for sale, made known, and/or publicly used before July 15, 2018. Verizon Networking Solutions were sold, offered for sale, and/or made known to the public and third-party customers before this date, by July 14, 2018 at the latest. Verizon Networking Solutions were publicly used by Verizon, the public, and third-party customers before this date, by July 14, 2018 at the latest.	A & O	Exhibit 634-02
Verizon Communications Inc.'s Voice over Internet Protocol devices ("Verizon VoIP Products"), including but not limited to Fios Digital Voice, Verizon One Talk, VoIP Business Digital Voice, and Preferred Voice, were sold, offered for sale, made known, and/or publicly used before July 15, 2018. Verizon VoIP Products were sold, offered for sale, and/or made known to the public and third-party customers before this date, by July 14, 2018 at the latest. Verizon VoIP Products were publicly used by Verizon, the public, and third-party customers before this date, by July 14, 2018 at the latest.	A & O	Exhibit 634-03

Prior Art to '634	Anticipation/ Obviousness	Claim Chart
Google LLC's networking products ("Google Networking Products"), including but not limited to Google Fi, Google OnHub, Google Wifi, Google Voice, Google Meet, and Google Hangout, were sold, offered for sale, made known, and/or publicly used before July 15, 2018. Google Networking Products were sold, offered for sale, and/or made known to the public and third-party customers before this date, by July 14, 2018 at the latest. Google Networking Products were publicly used by Google, the public, and third-party customers before this date, by July 14, 2018 at the latest.	A & O	Exhibit 634-04
U.S. Patent No. 8,578,008 ("Sella"); filed on July 16, 2013; issued on November 5, 2013.	A & O	Exhibit 634-05
U.S. Patent No. 9,032,060 ("Blaisdell"); filed on September 14, 2011; issued on May 12, 2015.	A & O	Exhibit 634-06
U.S. Patent No. 7,415,038 ("Ullmann"); filed on March 29, 2001; issued on August 19, 2008.	O	Exhibit 634-07
U.S. Patent No. 7,742,499 ("Erskine"); filed on August 18, 2005; issued on June 22, 2010.	A & O	Exhibit 634-08
U.S. Patent No. 8,000,276 ("Scherzer"); filed on May 3, 2007; issued on August 16, 2011.	A & O	Exhibit 634-09
U.S. Patent No. 9,661,553 ("Gillett"); filed on July 14, 2015; issued on May 23, 2017.	O	Exhibit 634-10
Systems implementing the Congestion Manager developed by a group of researchers at Massachusetts Institute of Technology, Carnegie Mellon University, and other institutions ("Congestion Manager System") were sold, offered for sale, made known, and/or publicly used before July 15, 2018. The Congestion Manager System were sold, offered for sale, and/or made known to the public and third-party customers before this date, by July 14, 2018 at the latest. The Congestion Manager System were publicly used by its developers, the public, and third-party customers before this date, by July 14, 2018 at the latest.	A & O	Exhibit 634-11

Prior Art to '634	Anticipation/ Obviousness	Claim Chart
David G. Andersen et al., <i>System Support for Bandwidth Management and Content Adaptation in Internet Applications</i> , in Proceedings of the 4th USENIX Symposium on Operating Systems Design and Implementation (OSDI '00) (USENIX Ass'n 2000), https://www.usenix.org/legacy/publications/library/proceedings/osdi2000/full_papers/andersonbansal/andersonbansal.pdf .	A & O	Exhibit 634-12
RealNetworks' system and platform for delivering and streaming media content ("RealNetworks Media Delivery System"), including but not limited to Real Broadcast Network, RealSystem, RealServer, RealPlayer, Helix Server, and RealProducer, were sold, offered for sale, made known, and/or publicly used before July 15, 2018. RealNetworks Media Delivery System were sold, offered for sale, and/or made known to the public and third-party customers before this date, by July 14, 2018 at the latest. RealNetworks Media Delivery System were publicly used by RealNetworks, the public, and third-party customers before this date, by July 14, 2018 at the latest.	A & O	Exhibit 634-13

In addition to the anticipation and obviousness explained in this section, the '634 patent is anticipated and/or rendered obvious by prior art as explained by Defendants' Petition for Inter Partes Review (assigned Case No. IPR2025-00479), which is incorporated herein by reference:

Grounds	Claims	Prior Art
1	1-20	Erskine
2	1-20	Scherzer
3	1-20	Scherzer in view of Gillett

1. Anticipation

For each item of prior art identified above as an anticipation reference, Defendants contend that the prior art item anticipates the Asserted Claims of the '634 patent, including for the reasons

1 illustrated in the accompanying claim chart and based on Plaintiffs' infringement allegations and
2 contentions for the '634 patent.

3 2. Obviousness

4 A person of ordinary skill in the art would have been motivated to combine any reference
5 in Table E with any other reference(s) in Table E. To the extent that any prior art identified in
6 Table E above is deemed to not anticipate an Asserted Claim, the prior art in Table E may be
7 combined with one or more other prior art references in Table E to render the claim obvious. All
8 of these items of prior art are directed to and propose solutions similar to those discussed in the
9 '634 patent, and is in the same field. All of these items of prior art are directed to networking
10 systems and architectures, and propose solutions similar to those discussed in the '634 patent.

11 The following are exemplary combinations of prior art that would have been obvious to a
12 POSITA:

- 13 • **Sella + Ullmann**
- 14 • **Sella + Erskine**

15 In addition to the above-mentioned motivations to combine, Sella is "related to service
16 allocation in a data communication network" and discloses "a system [that] adjusts capacity on a
17 network." *See* Sella, 1:7-8, 1:37-61. Sella discloses "dynamically changing bandwidth
18 allocation." *See* Sella, 2:60-3:6. A POSITA would naturally have looked to other references that
19 provide details related to bandwidth allocation and management, Quality of Service (QoS), load
20 balancing, congestion control, adaptive routing, and/or network analytics and monitoring.

21 For example, a POSITA would have been motivated to combine Sella with Ullmann,
22 which teaches monitoring and managing network bandwidth. *See, e.g.*, Ullmann, 4:18-43.

23 As an additional example, a POSITA would have been motivated to combine Sella with
24 Erskine, which teaches "adaptive bandwidth management for a packet-based communications
25 network." *See, e.g.*, Erskine, 1:34-42.

- 26 • **Blaisdell + Scherzer**
- 27 • **Blaisdell + Gillett**

1 In addition to the above-mentioned motivations to combine, Blaisdell “relates to
2 computing devices and network selection. More specifically, it relates to network selection using
3 bandwidth availability.” *See* Blaisdell, 1:16-18. Blaisdell discloses receiving location data of
4 devices and enabling network connections based on bandwidth data. *See, e.g.*, Blaisdell, 2:37-47.
5 A POSITA would naturally have looked to other references that provide details related to
6 bandwidth allocation and management, Quality of Service (QoS), load balancing, congestion
7 control, adaptive routing, and/or network analytics and monitoring.

8 For example, a POSITA would have been motivated to combine Blaisdell with Scherzer,
9 which discloses obtaining location data of networking devices and establishing network
10 connections based on bandwidth data. *See, e.g.*, Scherzer, 2:42-3:43.

11 As an additional example, a POSITA would have been motivated to combine Blaisdell
12 with Gillett, which discloses establishing end-to-end network connections, determining locations
13 of networking devices, and measuring connection quality. *See, e.g.*, Gillett, Abstract, 9:57-10:5.

- 14 • **Erskine + Sella**
- 15 • **Erskine + Ullmann**

16 In addition to the above-mentioned motivations to combine, Erskine “relates generally to
17 communication systems, and more particularly to an apparatus and method for adaptive bandwidth
18 management for Voice-Over-IP (VOIP) networks.” *See* Erskine, 1:6-9. Erskine discloses “a
19 method of adaptive bandwidth management for a packet-based communications network.” *See,*
20 *e.g.*, Erskine, 1:34-42. A POSITA would naturally have looked to other references that provide
21 details related to bandwidth allocation and management, Quality of Service (QoS), load balancing,
22 congestion control, adaptive routing, and/or network analytics and monitoring.

23 For example, a POSITA would have been motivated to combine Erskine with Sella, which
24 discloses “adjust[ing] capacity on a network and “dynamically changing bandwidth allocation.”
25 *See, e.g.*, Sella, 1:7-8, 1:37-38, 2:60-3:6.

26 As an additional example, a POSITA would have been motivated to combine Erskine with
27 Ullmann, which teaches monitoring and managing network bandwidth. *See, e.g.*, Ullmann,
28 4:18-43.

1 • **Scherzer + Blaisdell**

2 • **Scherzer + Gillett**

3 Additionally, one of ordinary skill in the art would have been motivated to combine
4 Scherzer with any other reference(s) in Table E. In addition to the above-mentioned motivations
5 to combine, Scherzer “relates to wireless services and wireless access points.” *See* Scherzer,
6 1:14-15. Scherzer discloses obtaining location data of networking devices and establishing
7 network connections based on bandwidth data. *See, e.g.,* Scherzer, 2:42-3:43. A POSITA would
8 naturally have looked to other references that provide details related to bandwidth allocation and
9 management, Quality of Service (QoS), load balancing, congestion control, adaptive routing,
10 and/or network analytics and monitoring.

11 For example, a POSITA would have been motivated to combine Scherzer with Blaisdell,
12 which discloses receiving location data of devices and enabling network connections based on
13 bandwidth data. *See, e.g.,* Blaisdell, 2:37-47.

14 As an additional example, a POSITA would have been motivated to combine Scherzer
15 with Gillett, which discloses establishing end-to-end network connections, determining locations
16 of networking devices, and measuring connection quality. *See, e.g.,* Gillett, Abstract, 9:57-10:5.

- 17 • **Verizon Networking Devices + Sella**
18 • **Verizon Networking Devices + Sella + Ullmann**
19 • **Verizon Networking Devices + Sella + Erskine**
20 • **Verizon Networking Devices + Blaisdell**
21 • **Verizon Networking Devices + Blaisdell + Scherzer**
22 • **Verizon Networking Devices + Blaisdell + Gillett**
23 • **Verizon Networking Devices + Erskine**
24 • **Verizon Networking Devices + Erskine + Ullmann**
25 • **Verizon Networking Devices + Scherzer**
26 • **Verizon Networking Devices + Scherzer + Gillett**
27 • **Verizon Networking Devices + Congestion Manager System**
28

- 1 • **Verizon Networking Devices + Congestion Manager System + Sella**
- 2 • **Verizon Networking Devices + Congestion Manager System + Ullmann**
- 3 • **Verizon Networking Devices + Congestion Manager System + Erskine**
- 4 • **Verizon Networking Devices + Congestion Manager System + Blaisdell**
- 5 • **Verizon Networking Devices + Congestion Manager System + Scherzer**
- 6 • **Verizon Networking Devices + Congestion Manager System + Gillett**
- 7 • **Verizon Networking Devices + Andersen**
- 8 • **Verizon Networking Devices + Andersen + Sella**
- 9 • **Verizon Networking Devices + Andersen + Ullmann**
- 10 • **Verizon Networking Devices + Andersen + Erskine**
- 11 • **Verizon Networking Devices + Andersen + Blaisdell**
- 12 • **Verizon Networking Devices + Andersen + Scherzer**
- 13 • **Verizon Networking Devices + Andersen + Gillett**

14
15 Verizon Networking Devices are described in numerous documents including those listed
16 in Exhibit 634-01. One of ordinary skill in the art would have understood that these documents
17 describe Verizon Communications Inc.’s networking devices and would have been motivated to
18 combine the listed documents because they all describe Verizon’s networking devices, have the
19 same author, are from the same field, related to the same and/or interoperable methods and
20 systems, and/or attempt to solve the same problem.

21 In addition to the above-mentioned motivations to combine, Verizon Networking Devices
22 disclose methods, devices, and systems for establishing and managing networking connections
23 using various network parameters to optimize network performance. For example, Verizon
24 Networking Devices allow users “to connect all [their] devices for fast file transfers and smooth
25 media streaming” and provide “Quality of Service support for traffic prioritization.” *See, e.g.,*
26 Exhibit 634-01, Verizon 5 or 8 Port Gigabit Desktop Switch. A POSITA would naturally have
27 looked to other references that provide details related to bandwidth allocation and management,
28

1 Quality of Service (QoS), load balancing, congestion control, adaptive routing, and/or network
2 analytics and monitoring.

3 For example, a POSITA would have been motivated to combine Verizon Networking
4 Devices with Sella, which teaches dynamically changing bandwidth allocation.

5 As an additional example, a POSITA would have been motivated to combine Verizon
6 Networking Devices with Ullmann, which teaches monitoring and managing network bandwidth.

7 As an additional example, a POSITA would have been motivated to combine Verizon
8 Networking Devices with Erskine, which teaches adaptive bandwidth management for a packet-
9 based communications network.

10 As an additional example, a POSITA would have been motivated to combine Verizon
11 Networking Devices with Blaisdell, which teaches receiving location data of devices and enabling
12 network connections based on bandwidth data.

13 As an additional example, a POSITA would have been motivated to combine Verizon
14 Networking Devices with Scherzer, which teaches obtaining location data of networking devices
15 and establishing network connections based on bandwidth data.

16 As an additional example, a POSITA would have been motivated to combine Verizon
17 Networking Devices with Gillett, which teaches establishing end-to-end network connections,
18 determining locations of networking devices, and measuring connection quality.

19 As an additional example, a POSITA would have been motivated to combine Verizon
20 Networking Devices with Congestion Manager System, which teaches collecting network
21 characteristics, coordinating host- and domain-specific path information, sharing path properties
22 between different streams, and apportioning bandwidths to flows.

23 As an additional example, a POSITA would have been motivated to combine Verizon
24 Networking Devices with Andersen, which teaches collecting network characteristics,
25 coordinating host- and domain-specific path information, sharing path properties between different
26 streams, and apportioning bandwidths to flows.

- 27 • **Verizon Networking Solutions + Sella**
- 28 • **Verizon Networking Solutions + Sella + Ullmann**

- 1 • **Verizon Networking Solutions + Sella + Erskine**
- 2 • **Verizon Networking Solutions + Blaisdell**
- 3 • **Verizon Networking Solutions + Blaisdell + Scherzer**
- 4 • **Verizon Networking Solutions + Blaisdell + Gillett**
- 5 • **Verizon Networking Solutions + Erskine**
- 6 • **Verizon Networking Solutions + Erskine + Ullmann**
- 7 • **Verizon Networking Solutions + Scherzer**
- 8 • **Verizon Networking Solutions + Scherzer + Gillett**
- 9 • **Verizon Networking Solutions + Congestion Manager System**
- 10 • **Verizon Networking Solutions + Congestion Manager System + Sella**
- 11 • **Verizon Networking Solutions + Congestion Manager System + Ullmann**
- 12 • **Verizon Networking Solutions + Congestion Manager System + Erskine**
- 13 • **Verizon Networking Solutions + Congestion Manager System + Blaisdell**
- 14 • **Verizon Networking Solutions + Congestion Manager System + Scherzer**
- 15 • **Verizon Networking Solutions + Congestion Manager System + Gillett**
- 16 • **Verizon Networking Solutions + Andersen**
- 17 • **Verizon Networking Solutions + Andersen + Sella**
- 18 • **Verizon Networking Solutions + Andersen + Ullmann**
- 19 • **Verizon Networking Solutions + Andersen + Erskine**
- 20 • **Verizon Networking Solutions + Andersen + Blaisdell**
- 21 • **Verizon Networking Solutions + Andersen + Scherzer**
- 22 • **Verizon Networking Solutions + Andersen + Gillett**

23 Verizon Networking Solutions are described in numerous documents including those listed
24 in Exhibit 634-02. One of ordinary skill in the art would have understood that these documents
25 describe Verizon Communications Inc.'s networking solutions and would have been motivated to
26
27
28

1 combine the listed documents because they all describe Verizon's networking solutions, have the
2 same author, are from the same field, related to the same and/or interoperable methods and
3 systems, and/or attempt to solve the same problem.

4 In addition to the above-mentioned motivations to combine, Verizon Networking Solutions
5 disclose devices, methods, systems, and architectures for establishing and managing networking
6 connections using various network parameters to optimize network performance. For example,
7 Verizon Networking Solutions provide "[a]pplication assessments to determine flow, performance
8 thresholds and routing rules[,] [f]easibility assessments and recommendations to improve network
9 performance and control costs[, and] [u]nbiased analyses of all options, including a hybrid
10 network, adding additional bandwidth to the private network or other alternatives." *See, e.g.,*
11 Exhibit 634-02, Build a Wiser WAN - Managed SD WAN with Application Aware Routing. A
12 POSITA would naturally have looked to other references that provide details related to bandwidth
13 allocation and management, Quality of Service (QoS), load balancing, congestion control,
14 adaptive routing, and/or network analytics and monitoring.

15 For example, a POSITA would have been motivated to combine Verizon Networking
16 Solutions with Sella, which teaches dynamically changing bandwidth allocation.

17 As an additional example, a POSITA would have been motivated to combine Verizon
18 Networking Solutions with Ullmann, which teaches monitoring and managing network bandwidth.

19 As an additional example, a POSITA would have been motivated to combine Verizon
20 Networking Solutions with Erskine, which teaches adaptive bandwidth management for a packet-
21 based communications network.

22 As an additional example, a POSITA would have been motivated to combine Verizon
23 Networking Solutions with Blaisdell, which teaches receiving location data of devices and
24 enabling network connections based on bandwidth data.

25 As an additional example, a POSITA would have been motivated to combine Verizon
26 Networking Solutions with Scherzer, which teaches obtaining location data of networking devices
27 and establishing network connections based on bandwidth data.

28

1 As an additional example, a POSITA would have been motivated to combine Verizon
2 Networking Solutions with Gillett, which teaches establishing end-to-end network connections,
3 determining locations of networking devices, and measuring connection quality.

4 As an additional example, a POSITA would have been motivated to combine Verizon
5 Networking Solutions with Congestion Manager System, which teaches collecting network
6 characteristics, coordinating host- and domain-specific path information, sharing path properties
7 between different streams, and apportioning bandwidths to flows.

8 As an additional example, a POSITA would have been motivated to combine Verizon
9 Networking Solutions with Andersen, which teaches collecting network characteristics,
10 coordinating host- and domain-specific path information, sharing path properties between different
11 streams, and apportioning bandwidths to flows.

- 12 • **Verizon VoIP Products + Sella**
- 13 • **Verizon VoIP Products + Sella + Ullmann**
- 14 • **Verizon VoIP Products + Sella + Erskine**
- 15 • **Verizon VoIP Products + Blaisdell**
- 16 • **Verizon VoIP Products + Blaisdell + Scherzer**
- 17 • **Verizon VoIP Products + Blaisdell + Gillett**
- 18 • **Verizon VoIP Products + Erskine**
- 19 • **Verizon VoIP Products + Erskine + Ullmann**
- 20 • **Verizon VoIP Products + Scherzer**
- 21 • **Verizon VoIP Products + Scherzer + Gillett**
- 22 • **Verizon VoIP Products + Congestion Manager System**
- 23 • **Verizon VoIP Products + Congestion Manager System + Sella**
- 24 • **Verizon VoIP Products + Congestion Manager System + Ullmann**
- 25 • **Verizon VoIP Products + Congestion Manager System + Erskine**
- 26 • **Verizon VoIP Products + Congestion Manager System + Blaisdell**
- 27
- 28

- 1 • **Verizon VoIP Products + Congestion Manager System + Scherzer**
- 2 • **Verizon VoIP Products + Congestion Manager System + Gillett**
- 3 • **Verizon VoIP Products + Andersen**
- 4 • **Verizon VoIP Products + Andersen + Sella**
- 5 • **Verizon VoIP Products + Andersen + Ullmann**
- 6 • **Verizon VoIP Products + Andersen + Erskine**
- 7 • **Verizon VoIP Products + Andersen + Blaisdell**
- 8 • **Verizon VoIP Products + Andersen + Scherzer**
- 9 • **Verizon VoIP Products + Andersen + Gillett**

11 Verizon VoIP Products are described in numerous documents including those listed in
12 Exhibit 634-03. One of ordinary skill in the art would have understood that these documents
13 describe Verizon Communications Inc.’s VoIP products and would have been motivated to
14 combine the listed documents because they all describe Verizon’s VoIP products, have the same
15 author, are from the same field, related to the same and/or interoperable methods and systems,
16 and/or attempt to solve the same problem.

17 In addition to the above-mentioned motivations to combine, Verizon VoIP Products
18 disclose VoIP related devices, systems, and architectures for establishing and managing
19 connections using various network parameters to optimize network performance. For example,
20 Verizon VoIP Products facilitate end-to-end VoIP connections and provide “excellent voice
21 quality and optimized packet voice streaming over IP networks, enabling voice, fax, and data
22 traffic to be sent over the same IP network.” *See, e.g.,* Exhibit 634-03, Verizon T.38 FAX
23 Configuration Guide for AudioCodes MP-11x. A POSITA would naturally have looked to other
24 references that provide details related to bandwidth allocation and management, Quality of Service
25 (QoS), load balancing, congestion control, adaptive routing, and/or network analytics and
26 monitoring.

27 For example, a POSITA would have been motivated to combine Verizon VoIP Products
28 with Sella, which teaches dynamically changing bandwidth allocation.

1 As an additional example, a POSITA would have been motivated to combine Verizon
2 VoIP Products with Ullmann, which teaches monitoring and managing network bandwidth.

3 As an additional example, a POSITA would have been motivated to combine Verizon
4 VoIP Products with Erskine, which teaches adaptive bandwidth management for a packet-based
5 communications network.

6 As an additional example, a POSITA would have been motivated to combine Verizon
7 VoIP Products with Blaisdell, which teaches receiving location data of devices and enabling
8 network connections based on bandwidth data.

9 As an additional example, a POSITA would have been motivated to combine Verizon
10 VoIP Products with Scherzer, which teaches obtaining location data of networking devices and
11 establishing network connections based on bandwidth data.

12 As an additional example, a POSITA would have been motivated to combine Verizon
13 VoIP Products with Gillett, which teaches establishing end-to-end network connections,
14 determining locations of networking devices, and measuring connection quality.

15 As an additional example, a POSITA would have been motivated to combine Verizon
16 VoIP Products with Congestion Manager System, which teaches collecting network
17 characteristics, coordinating host- and domain-specific path information, sharing path properties
18 between different streams, and apportioning bandwidths to flows.

19 As an additional example, a POSITA would have been motivated to combine Verizon
20 VoIP Products with Andersen, which teaches collecting network characteristics, coordinating
21 host- and domain-specific path information, sharing path properties between different streams, and
22 apportioning bandwidths to flows.

- 23 • **Google Networking Products + Sella**
- 24 • **Google Networking Products + Sella + Ullmann**
- 25 • **Google Networking Products + Sella + Erskine**
- 26 • **Google Networking Products + Blaisdell**
- 27 • **Google Networking Products + Blaisdell + Scherzer**
- 28

- 1 • **Google Networking Products + Blaisdell + Gillett**
- 2 • **Google Networking Products + Erskine**
- 3 • **Google Networking Products + Erskine + Ullmann**
- 4 • **Google Networking Products + Scherzer**
- 5 • **Google Networking Products + Scherzer + Gillett**
- 6 • **Google Networking Products + Congestion Manager System**
- 7 • **Google Networking Products + Congestion Manager System + Sella**
- 8 • **Google Networking Products + Congestion Manager System + Ullmann**
- 9 • **Google Networking Products + Congestion Manager System + Erskine**
- 10 • **Google Networking Products + Congestion Manager System + Blaisdell**
- 11 • **Google Networking Products + Congestion Manager System + Scherzer**
- 12 • **Google Networking Products + Congestion Manager System + Gillett**
- 13 • **Google Networking Products + Andersen**
- 14 • **Google Networking Products + Andersen + Sella**
- 15 • **Google Networking Products + Andersen + Ullmann**
- 16 • **Google Networking Products + Andersen + Erskine**
- 17 • **Google Networking Products + Andersen + Blaisdell**
- 18 • **Google Networking Products + Andersen + Scherzer**
- 19 • **Google Networking Products + Andersen + Gillett**

22 Google Networking Products are described in numerous documents including those listed
23 in Exhibit 634-04. One of ordinary skill in the art would have understood that these documents
24 describe Google LLC's networking products and would have been motivated to combine the listed
25 documents because they all describe Google's networking products, have the same author, are
26 from the same field, related to the same and/or interoperable methods and systems, and/or attempt
27 to solve the same problem.

1 In addition to the above-mentioned motivations to combine, Google Networking Products
2 disclose methods, devices, systems, and architectures for establishing and managing connections
3 using various network parameters to optimize network performance. For example, Google
4 Networking Products provide “wired bridge mode and [r]outer functionality” and support features
5 such as “[m]ulti-user sharing a high-speed Internet connection” and “bandwidth control.” *See,*
6 *e.g.*, Exhibit 634-04, Google OnHub User Guide. A POSITA would naturally have looked to
7 other references that provide details related to bandwidth allocation and management, Quality of
8 Service (QoS), load balancing, congestion control, adaptive routing, and/or network analytics and
9 monitoring.

10 For example, a POSITA would have been motivated to combine Google Networking
11 Products with Sella, which teaches dynamically changing bandwidth allocation.

12 As an additional example, a POSITA would have been motivated to combine Google
13 Networking Products with Ullmann, which teaches monitoring and managing network bandwidth.

14 As an additional example, a POSITA would have been motivated to combine Google
15 Networking Products with Erskine, which teaches adaptive bandwidth management for a packet-
16 based communications network.

17 As an additional example, a POSITA would have been motivated to combine Google
18 Networking Products with Blaisdell, which teaches receiving location data of devices and enabling
19 network connections based on bandwidth data.

20 As an additional example, a POSITA would have been motivated to combine Google
21 Networking Products with Scherzer, which teaches obtaining location data of networking devices
22 and establishing network connections based on bandwidth data.

23 As an additional example, a POSITA would have been motivated to combine Google
24 Networking Products with Gillett, which teaches establishing end-to-end network connections,
25 determining locations of networking devices, and measuring connection quality.

26 As an additional example, a POSITA would have been motivated to combine Google
27 Networking Products with Congestion Manager System, which teaches collecting network
28

1 characteristics, coordinating host- and domain-specific path information, sharing path properties
2 between different streams, and apportioning bandwidths to flows.

3 As an additional example, a POSITA would have been motivated to combine Google
4 Networking Products with Andersen, which teaches collecting network characteristics,
5 coordinating host- and domain-specific path information, sharing path properties between different
6 streams, and apportioning bandwidths to flows.

- 7 • **Congestion Manager System + Sella**
- 8 • **Congestion Manager System + Sella + Ullmann**
- 9 • **Congestion Manager System + Sella + Erskine**
- 10 • **Congestion Manager System + Blaisdell**
- 11 • **Congestion Manager System + Blaisdell + Scherzer**
- 12 • **Congestion Manager System + Blaisdell + Gillett**
- 13 • **Congestion Manager System + Erskine**
- 14 • **Congestion Manager System + Erskine + Ullmann**
- 15 • **Congestion Manager System + Scherzer**
- 16 • **Congestion Manager System + Scherzer + Gillett**

17
18 The Congestion Manager System is described in multiple documents including those listed
19 in Exhibit 634-11. One of ordinary skill in the art would have understood that these documents
20 describe the Congestion Manager System and would have been motivated to combine the listed
21 documents because they all describe the Congestion Manager System, share authors, are from the
22 same field, related to the same and/or interoperable methods and systems, and/or attempt to solve
23 the same problem.

24 In addition to the above-mentioned motivations to combine, the Congestion Manager
25 System disclose devices, methods, systems, and architectures for establishing and managing
26 networking connections using various network parameters to optimize network performance. For
27 example, the Congestion Manager System includes a Congestion Manager (“CM”), which
28 “ensures proper congestion behavior and allows applications to easily adapt to network

1 congestion,” “maintains congestion parameters and exposes an API to enable applications to learn
2 about network characteristics, pass information to the CM, and schedule data transmissions,” and
3 “uses a window-based control algorithm, a scheduler to regulate transmissions, and a lightweight
4 protocol to elicit feedback from receivers.” *See, e.g.*, Exhibit 634-11, An Integrated Congestion
5 Management Architecture for Internet Hosts. A POSITA would naturally have looked to other
6 references that provide details related to bandwidth allocation and management, Quality of Service
7 (QoS), load balancing, congestion control, adaptive routing, and/or network analytics and
8 monitoring.

9 For example, a POSITA would have been motivated to combine Congestion Manager
10 System with Sella, which teaches dynamically changing bandwidth allocation.

11 As an additional example, a POSITA would have been motivated to combine Congestion
12 Manager System with Ullmann, which teaches monitoring and managing network bandwidth.

13 As an additional example, a POSITA would have been motivated to combine Congestion
14 Manager System with Erskine, which teaches adaptive bandwidth management for a packet-based
15 communications network.

16 As an additional example, a POSITA would have been motivated to combine Congestion
17 Manager System with Blaisdell, which teaches receiving location data of devices and enabling
18 network connections based on bandwidth data.

19 As an additional example, a POSITA would have been motivated to combine Congestion
20 Manager System with Scherzer, which teaches obtaining location data of networking devices and
21 establishing network connections based on bandwidth data.

22 As an additional example, a POSITA would have been motivated to combine Congestion
23 Manager System with Gillett, which teaches establishing end-to-end network connections,
24 determining locations of networking devices, and measuring connection quality.

- 25 • **Andersen + Sella**
- 26 • **Andersen + Sella + Ullmann**
- 27 • **Andersen + Sella + Erskine**
- 28 • **Andersen + Blaisdell**

- 1 • **Andersen + Blaisdell + Scherzer**
- 2 • **Andersen + Blaisdell + Gillett**
- 3 • **Andersen + Erskine**
- 4 • **Andersen + Erskine + Ullmann**
- 5 • **Andersen + Scherzer**
- 6 • **Andersen + Scherzer + Gillett**

7
8 Andersen discloses devices, methods, systems, and architectures for establishing and
9 managing networking connections using various network parameters to optimize network
10 performance. For example, Andersen teaches that a Congestion Manager (“CM”) “performs two
11 important functions”: “First, it enables efficient multiplexing and congestion control by integrating
12 congestion management across multiple flows. Second, it enables efficient application adaptation
13 to congestion by exposing its knowledge of network conditions to applications.” *See, e.g.*, Exhibit
14 634-12 (Andersen, 3). The CM “assigns [a] flow to the appropriate macroflow based on its
15 destination” and permits a client to transmit data based on available rate. *Id.* The CM also
16 provides an API that “specifies functions for managing state, for performing data transmissions,
17 for applications to inform the CM of losses, for querying the CM about network state, and for
18 constructing and splitting macroflows if the default per-destination aggregation is unsuitable for
19 an application.” *See, e.g.*, Exhibit 634-12 (Andersen, 4). A POSITA would naturally have looked
20 to other references that provide details related to bandwidth allocation and management, Quality
21 of Service (QoS), load balancing, congestion control, adaptive routing, and/or network analytics
22 and monitoring.

23 For example, a POSITA would have been motivated to combine Andersen with Sella,
24 which teaches dynamically changing bandwidth allocation.

25 As an additional example, a POSITA would have been motivated to combine Andersen
26 with Ullmann, which teaches monitoring and managing network bandwidth.

1 As an additional example, a POSITA would have been motivated to combine Andersen
2 with Erskine, which teaches adaptive bandwidth management for a packet-based communications
3 network.

4 As an additional example, a POSITA would have been motivated to combine Andersen
5 with Blaisdell, which teaches receiving location data of devices and enabling network connections
6 based on bandwidth data.

7 As an additional example, a POSITA would have been motivated to combine Andersen
8 with Scherzer, which teaches obtaining location data of networking devices and establishing
9 network connections based on bandwidth data.

10 As an additional example, a POSITA would have been motivated to combine Andersen
11 with Gillett, which teaches establishing end-to-end network connections, determining locations of
12 networking devices, and measuring connection quality.

- 13 • **RealNetworks Media Delivery System + Sella**
- 14 • **RealNetworks Media Delivery System + Sella + Ullmann**
- 15 • **RealNetworks Media Delivery System + Sella + Erskine**
- 16 • **RealNetworks Media Delivery System + Blaisdell**
- 17 • **RealNetworks Media Delivery System + Blaisdell + Scherzer**
- 18 • **RealNetworks Media Delivery System + Blaisdell + Gillett**
- 19 • **RealNetworks Media Delivery System + Erskine**
- 20 • **RealNetworks Media Delivery System + Erskine + Ullmann**
- 21 • **RealNetworks Media Delivery System + Scherzer**
- 22 • **RealNetworks Media Delivery System + Scherzer + Gillett**
- 23 • **RealNetworks Media Delivery System + Congestion Manager System**
- 24 • **RealNetworks Media Delivery System + Congestion Manager System + Sella**
- 25 • **RealNetworks Media Delivery System + Congestion Manager System + Ullmann**
- 26 • **RealNetworks Media Delivery System + Congestion Manager System +**
- 27 **Ullmann**

- 1 • **RealNetworks Media Delivery System + Congestion Manager System +**
2 **Erskine**
- 3 • **RealNetworks Media Delivery System + Congestion Manager System +**
4 **Blaisdell**
- 5 • **RealNetworks Media Delivery System + Congestion Manager System +**
6 **Scherzer**
- 7 • **RealNetworks Media Delivery System + Congestion Manager System + Gillett**
- 8 • **RealNetworks Media Delivery System + Andersen**
- 9 • **RealNetworks Media Delivery System + Andersen + Sella**
- 10 • **RealNetworks Media Delivery System + Andersen + Ullmann**
- 11 • **RealNetworks Media Delivery System + Andersen + Erskine**
- 12 • **RealNetworks Media Delivery System + Andersen + Blaisdell**
- 13 • **RealNetworks Media Delivery System + Andersen + Scherzer**
- 14 • **RealNetworks Media Delivery System + Andersen + Gillett**

15 RealNetworks Media Delivery System is described in numerous documents including
16 those listed in Exhibit 634-13. One of ordinary skill in the art would have understood that these
17 documents describe RealNetworks' system and platform for delivering and streaming media
18 content and would have been motivated to combine the listed documents because they all describe
19 RealNetworks' media delivery system, have the same author, are from the same field, related to
20 the same and/or interoperable methods and systems, and/or attempt to solve the same problem.

21 In addition to the above-mentioned motivations to combine, RealNetworks Media Delivery
22 System discloses devices, methods, systems, and architectures for establishing and managing
23 network connections based on various network parameters, optimizing network and media
24 delivery performance, and initiating and controlling transmission of data over the network. For
25 example, RealNetworks Media Delivery System includes intelligent routing software that routes
26 users "to the topographically closest or least busy splitter" where content can be accessed. *See,*
27 *e.g.,* Exhibit 634-13, Introducing the First Distributed Multi-Tier Broadcast Architecture for the
28 Internet. RealNetworks Media Delivery System also provides capabilities such as gathering

1 statistics about usage and quality of service, capacity sharing, fail over to alternative servers, and
2 adaptively changing the bandwidth of streams. *See, e.g.*, Exhibit 634-13. A POSITA would
3 naturally have looked to other references that provide details related to bandwidth allocation and
4 management, Quality of Service (QoS), load balancing, congestion control, adaptive routing,
5 and/or network analytics and monitoring.

6 For example, a POSITA would have been motivated to combine RealNetworks Media
7 Delivery System with Sella, which teaches dynamically changing bandwidth allocation.

8 As an additional example, a POSITA would have been motivated to combine
9 RealNetworks Media Delivery System with Ullmann, which teaches monitoring and managing
10 network bandwidth.

11 As an additional example, a POSITA would have been motivated to combine
12 RealNetworks Media Delivery System with Erskine, which teaches adaptive bandwidth
13 management for a packet-based communications network.

14 As an additional example, a POSITA would have been motivated to combine
15 RealNetworks Media Delivery System with Blaisdell, which teaches receiving location data of
16 devices and enabling network connections based on bandwidth data.

17 As an additional example, a POSITA would have been motivated to combine
18 RealNetworks Media Delivery System with Scherzer, which teaches obtaining location data of
19 networking devices and establishing network connections based on bandwidth data.

20 As an additional example, a POSITA would have been motivated to combine
21 RealNetworks Media Delivery System with Gillett, which teaches establishing end-to-end network
22 connections, determining locations of networking devices, and measuring connection quality.

23 As an additional example, a POSITA would have been motivated to combine
24 RealNetworks Media Delivery System with Congestion Manager System, which teaches
25 collecting network characteristics, coordinating host- and domain-specific path information,
26 sharing path properties between different streams, and apportioning bandwidths to flows.

27 As an additional example, a POSITA would have been motivated to combine
28 RealNetworks Media Delivery System with Andersen, which teaches collecting network

1 characteristics, coordinating host- and domain-specific path information, sharing path properties
2 between different streams, and apportioning bandwidths to flows.

3 **II. Identification of Limitations that Defendants Contend Lack Written Description,**
4 **Lack Enablement, or are Indefinite Under 35 U.S.C. § 112**

5 The Asserted Claims recite limitations that were already well known in the prior art (and
6 therefore are anticipated and obvious). If the claims are construed to recite some novel limitation
7 or combination of limitations, however, then they fail to describe or enable that limitation or
8 combination, including how it differs from what was already known.

9 Section 112 “contains a written description requirement separate from enablement.” *Ariad*
10 *Pharms., Inc. v. Eli Lilly & Co.*, 598 F.3d 1336, 1351 (Fed. Cir. 2010). The written description
11 requirement requires that “the disclosure of the application relied upon reasonably convey[] to
12 those skilled in the art that the inventor had possession of the claimed subject matter as of the
13 filing date.” *Id.* Indeed, “the specification must describe an invention understandable to [a]
14 skilled artisan and show that the inventor actually invented the invention claimed.” *Id.*

15 In addition to the written description requirement, Section 112 imposes an enablement
16 requirement, which is satisfied only when “one skilled in the art, after reading the specification,
17 could practice the claimed invention without undue experimentation.” *Sitrick v. Dreamworks,*
18 *LLC*, 516 F.3d 993, 999 (Fed. Cir. 2008). Further, “[t]he full scope of the claimed invention must
19 be enabled,” and “[a] patentee who chooses broad claim language must make sure the broad
20 claims are fully enabled.” *Id.*

21 Many of the Asserted Claims are so poorly drafted as to be indecipherable to anyone,
22 including those skilled in the art, and are therefore facially indefinite and invalid. To be valid, a
23 patent claim must “particularly point [] out and distinctly claim[] the subject matter which the
24 applicant . . . regards as his invention.” 35 U.S.C. § 112(b). The purpose of this definiteness
25 requirement is “afford clear notice of what is claimed, thereby apprising the public of what is still
26 open to them.” *Nautilus v. Biosig Instruments, Inc.*, 572 U.S. 898, 909 (2014) (internal quotations
27 omitted). Patent claims are invalid for indefiniteness when they fail to “inform those skilled in the
28 art about the scope of the invention with reasonable certainty.” *Id.* at 910.

1 Many of the Asserted Claims fail to satisfy the written-description, enablement, or
2 definiteness requirements of § 112 as explained herein.

3 **A. U.S. Patent No. 8,924,596**

4 For the claim terms “queue implemented in hardware” and “register” in the ’596 patent,
5 Defendants’ claim construction briefing explained how Plaintiffs’ claim interpretations and
6 proposed constructions rendered these limitations indefinite under 35 U.S.C. §112; that analysis
7 and explanation is hereby incorporated by reference. Dkts. 79, 81.

8 **1. “input queue implemented in hardware” / “output queue implemented**
9 **in hardware”**

10 Under Plaintiffs’ Infringement Contentions and apparent application of their interpretation
11 of the claim language, claims 1, 7, 13, and their dependents of the ’596 patent are invalid under 35
12 U.S.C. §112(b) because the claim language, when read in light of the intrinsic record, fails to
13 inform a person of ordinary skill in the art the scope of the claimed invention with reasonable
14 certainty with respect to “input queue implemented in hardware” and “output queue implemented
15 in hardware” claim limitations. Under Plaintiffs’ apparent application of these two claim terms to
16 their infringement theories, a person of ordinary skill in the art would not be reasonably certain as
17 the scope of the invention, because Plaintiffs’ infringement theories appear to blur hardware and
18 software implementations of queues by encompassing software within the scope of an “input
19 queue implemented in hardware” and “output queue implemented in hardware.” Plaintiffs’
20 infringement theories effectively read out of the claim limitation the words “implemented in
21 hardware” and fails to provide any bounds on the scope of the term “implemented in hardware,”
22 rendering both terms indefinite under 35 U.S.C. §112. The basis for this contention is disclosed in
23 Defendants’ claim construction briefing for the ’596 patent (Dkts. 79, 81), and is hereby
24 incorporated by reference.
25
26
27
28

1 **B. U.S. Patent No. 9,529,767**

2 **1. “abstraction protocol”**

3 Claims 1 and 16, and all claims depending from them, are invalid under 35 U.S.C. §112
4 because the patent fails to provide a sufficient written description of the claimed “abstraction
5 protocol,” lacks enablement, and fails to clearly define the term, rendering it indefinite.

6 For example, Claims 1 and 16 recite an “abstraction protocol” that comprises “a full-
7 duplex protocol supporting full command queuing for [] at least one storage media device.”
8 However, the specification fails to demonstrate that the inventors were in possession of the
9 specific structure of such a protocol at the time of filing, as it fails to provide sufficient detail
10 about how the “abstraction protocol” is implemented or how it can achieve these functionalities.
11 While the specification identifies features that the abstraction protocol may possess, such as being
12 “performance-centric” and “link-agnostic,” (’767 pat., 4:35-40), and “support[ing] common high-
13 level read and write access to a pool of storage media devices” (*id.*, 4:35-38), it fails to provide
14 sufficient details about the structure, operation or implementation of the “abstraction protocol.”
15 For example, the specification does not include concrete examples or diagrams illustrating how to
16 implement the “abstraction protocol.” Instead, it merely offers the circular explanation “[t]his
17 invention simplifies a host that accesses SATA and SAS storage media devices by abstracting the
18 SATA and SAS protocols with one full duplex protocol,” (*id.*, 3:65-67), which fails to inform with
19 reasonable certainty to one of ordinary skill in the art how the SATA and SAS protocols are
20 abstracted.

21 One of ordinary skill in the art at the time of the invention would have expected the
22 inventors to describe major aspects of a novel storage protocol’s design if they were in possession
23 of the protocol at the time of the claimed invention. The following examples are just some aspects
24 of a storage protocol that a POSITA would have expected to see described if the inventors
25 possessed such a protocol:

- 26 • Grammar of the protocol: The specification of the commands that are a part of the
27 protocol. This requires consideration of the native capabilities of the storage
28 devices supported by the protocol (e.g., the native command set). This also

1 includes the properties of individual commands included in the grammar (e.g.,
2 atomicity, consistency, idempotency, reliability, availability, and durability).

- 3 • Semantics of the protocol: The meaning of the commands.
- 4 • Data Structures: The specification of data types and structures, such as queues, and
5 handling of data structures.
- 6 • Message format and syntax: The specification of the content of messages,
7 formatting of messages, ordering of messages, etc.
- 8 • Error handling: The specification of error conditions, their properties (e.g., hard
9 and soft faults), how they are represented, error correction, and how handled.
- 10 • Allocation of responsibilities between the components of a system (e.g., the host,
11 the storage media device, and their associated interfaces and controllers).

12 The patent's specification does not adequately describe an abstraction protocol, nor does
13 the patent's specification identify any then-existing protocol that constitutes the claimed
14 abstraction protocol. However, Plaintiffs' infringement contentions identify NVMe over Fabrics
15 as the alleged abstraction protocol in the Accused Instrumentalities and states in part:

16 NVMe over Fabrics is an abstraction protocol, and relies on the underlying NVMe
17 Transport to ensure reliable delivery of NVMe commands and data. NVMe over
18 Fabrics also relies on RDMA. Defendants' implementation and use of NVMe over
19 Fabrics by itself and as implementing and using NVMe Transport and RDMA is a
20 claimed abstraction protocol. The underlying NVMe Transport (as well as RDMA)
21 serves as a part of the NVMe over Fabrics abstraction protocol, independent of any
22 specific physical interconnect characteristics. NVMe over Fabrics is a full-duplex
23 protocol and supports full command queuing for the at least one storage media
24 device. Further, NVMe over Fabrics is performance-centric, supports common
25 high-level read and write access to a pool of storage media devices, is link-
26 agnostic; and may be carried via a multiplicity of direct attach or networked
27 interfaces. NVMe-oF corresponds to the claimed abstraction protocol in that it
28 provides a performance-centric storage abstraction layer that exposes uniform high-
level read and write operations to a pool of storage media devices, independent of
the physical characteristics of the underlying media. NVMe-oF abstracts the
storage devices from the host and maintains efficient command queuing and
parallel execution semantics. NVMe-oF is link-agnostic, as it supports carriage
over multiple transport mechanisms, including RDMA-based fabrics and TCP/IP
networks, and is deployable over both direct-attach and networked interfaces.
Further, NVMe-oF permits heterogeneous storage implementations behind the
abstraction layer while presenting a consistent protocol interface to the host.

12/30/25 Inf. Contentions, Ex. B at 16-17 (citations omitted).

1 NVME is a storage protocol that was standardized by the NVMeExpress organization to
2 provide access to storage devices over PCIe links. NVME over Fabric is another specification
3 standardized by the NVMeExpress organization to provide encapsulation and transport of NVME
4 commands over fabrics instead of PCIe links. Each of the specifications are hundreds of pages
5 long and reflect over a hundred thousand hours of time spent by engineers to develop working
6 protocols. These standards themselves are further evidence of why a person of ordinary skill
7 would not have understood the applicant to be in possession of the claimed subject matter because
8 they illustrate aspects of a storage protocol that the applicants would have disclosed in the
9 specification had they been in possession of an abstraction protocol.

10 For example, the NVME specification illustrates the following aspects of a protocol that
11 the '767 patent does not disclose:

- 12 • Data structures
- 13 • Command sets
- 14 • Control architecture
- 15 • Error reporting and recovery

16 For example, the NVME over Fabrics specification illustrates the following aspects of a
17 protocol that the '767 patent does not disclose:

- 18 • Encapsulation and data transfers
- 19 • Commands
- 20 • Control architecture
- 21 • Discovery
- 22 • Authentication
- 23 • Transport (e.g., queueing, error recovery, etc.)

24 Moreover, the '767 patent specification does not enable one of ordinary skill in the art to
25 practice the full scope of the claimed invention without undue experimentation for a number of
26 reasons. As described above, there are a number of major aspects that must be specified for a
27 storage protocol and the patent specification does not describe any of them. In addition, the
28 NVME and NVME over Fabric standards, which Plaintiffs contend are examples of an abstraction

1 protocol, required over one hundred thousand hours of development effort to produce working
2 storage protocols. The patent specification does not provide any guidance to a person ordinary on
3 how to make any of these design decisions, let alone make design choices that achieve the various
4 properties the patent specification states the abstraction protocol must have.

5 “Abstraction protocol” is also indefinite. The term “abstraction protocol” is not precisely
6 defined and is used vaguely throughout the patent without clear boundaries, making the scope of
7 the claims ambiguous. The claims fail to inform a person of ordinary skill in the art about the
8 exact nature of the abstraction protocol or how it distinguishes from existing protocols, leaving
9 them to speculate the scope of the claimed invention. The basis for this contention is further
10 disclosed in Defendants’ claim construction briefing for the ’767 patent (Dkts. 79, 81), and is
11 hereby incorporated by reference.

12 **2. “full command queuing”**

13 Claims 1 and 16, and all claims depending from them, are invalid under 35 U.S.C. §112
14 because the patent fails to provide a sufficient written description of the claimed “full command
15 queueing” and fails to clearly define the term, rendering it indefinite. The claims recite “full
16 command queueing” as a feature of the abstraction protocol, however, the specification fails to
17 provide any explanation of what “full command queuing” entails, how it is implemented, or how it
18 differs from the existing queueing methods such as native command queuing and tagged command
19 queueing. The specification merely mentions “full command queueing” in a conclusory manner,
20 without offering any guidance as to its scope or operation. There is no description of how
21 commands are queued, managed, or executed under this full command queueing context. This
22 lack of detailed explanation further results in ambiguity, as the term “full command queuing”
23 could be subject to speculative interpretation, making it impossible for a person of ordinary skill in
24 the art to understand the full scope of the invention or determine the boundaries of the claimed
25 invention. As a result, the claims are invalid for lack of written description and indefiniteness
26 under 35 U.S.C. §112.

1 **3. “single high level command”**

2 Claims 7 and 19 are invalid under 35 U.S.C. §112 because the patent fails to provide a
3 sufficient written description of the claimed “single high level command” and fails to clearly
4 define the term, rendering it indefinite. The specification mentions the concept of a “single high
5 level command” in the context of the abstraction protocol, however, it fails to provide a detailed
6 explanation of what constitutes a “single high level command” or how it is structured. This lack
7 of detailed explanation further results in ambiguity, as the term “single high level command”
8 could be subject to speculative interpretation, making it impossible for a person of ordinary skill in
9 the art to understand the full scope of the invention or determine the boundaries of the claimed
10 invention. As a result, the claims are invalid for lack of written description and indefiniteness
11 under 35 U.S.C. §112.

12 **4. “the edge processing element associated with the edge controller**
13 **authors abstract protocol messages intended for the other edge**
14 **controller”**

15 Claim 1 and all claims depending from it are invalid under 35 U.S.C. §112 because the
16 patent fails to provide a sufficient written description of the claimed “the edge processing element
17 associated with the edge controller authors abstract protocol messages intended for the other edge
18 controller” and fails to clearly define the term, rendering it indefinite. The specification mentions
19 that the edge processing element “authors” abstract protocol messages, but it fails to explain or
20 define what authoring entails or how it is implemented. This lack of detailed explanation further
21 results in ambiguity, as “authors” could be subject to speculative interpretation, making it
22 impossible for a person of ordinary skill in the art to understand the full scope of the invention or
23 determine the boundaries of the claimed invention. In addition, the specification does not explain
24 or define what is meant by “*intended* for the other edge controller.” This lack of detailed
25 explanation further results in ambiguity, as “intended” could be subject to speculative
26 interpretation, making it impossible for a person of ordinary skill in the art to understand the full
27 scope of the invention or determine the boundaries of the claimed invention. As a result, the
28 claims are invalid for lack of written description and indefiniteness under 35 U.S.C. §112.

1 **5. “the master controller” and “the edge controller”**

2 Claims 8, 9, 11, 13, 14, 15 and 19 are invalid under 35 U.S.C. §112 because the patent fails
3 to clearly define the terms “the master controller” and “the edge controller” as recited in the
4 claims, rendering them indefinite. Claim 1 references “a master controller,” “one edge controller,”
5 “one other master controller,” and “one other edge controller.” Therefore, it is unclear which
6 specific master controller and edge controller are being referred to when claims 8, 9, 11, 13, 14, 15
7 and 19 recite “the master controller” or “the edge controller.” Such ambiguity makes it impossible
8 for a person of ordinary skill in the art to understand the full scope of the invention or determine
9 the boundaries of the claimed invention. As a result, the claims are invalid for indefiniteness
10 under 35 U.S.C. §112.

11 **C. U.S. Patent No. 10,873,753**

12 For the claim terms ““interleaved memory” and “swapped memory over a specified
13 address range” in the ’753 patent, Defendants’ claim construction briefing explained how
14 Plaintiffs’ claim interpretations and proposed constructions rendered these limitations indefinite
15 under 35 U.S.C. §112; that analysis and explanation is hereby incorporated by reference. Dkts.
16 79, 81.

17 **1. “a sandbox”**

18 Claims 3, 10, and 14 are invalid under 35 U.S.C. §112(b). Under Plaintiffs’ Infringement
19 Contentions and apparent application of their interpretation of the claim language, these claims are
20 invalid under 35 U.S.C. §112(b) because the claim language, when read in light of the intrinsic
21 record, fails to inform a POSITA of the scope of the claimed invention with reasonable certainty
22 with respect to the boundaries of the claim term “sandbox.”

23 The specification describes the sandboxes as being “for programmable electronic circuits”
24 and “user-defined functionality” on a PLD ’753 patent, 5:19-21, 12:60-61, 3:13-14 (“for user-
25 defined electronic circuits”), 4:29-30 (“for configurable modules”).

26 However, in contrast, Plaintiffs’ Infringement Contentions expand the scope of “sandbox”
27 beyond PLDs to encompass “P4 engines” in an ASIC that does not have user-definable
28 programmable circuits. ’753 Infringement Contentions at 53. The intrinsic record does not

1 provide any written description support for sandboxes that do not have user-definable
2 programmable circuits. Therefore, the intrinsic record fails to inform a POSITA with reasonable
3 certainty regarding the boundaries of “sandbox” under Plaintiffs’ apparent view of the claim,
4 particularly when Plaintiffs’ Infringement Contentions appear to take an expansive view that a
5 sandbox could encompass any arbitrary portion of an integrated circuit, which renders the claims
6 indefinite under 35 U.S.C. §112(b). The basis for this contention is further disclosed in
7 Defendants’ claim construction briefing for the ’943 patent (Dkt. 79); that analysis is hereby
8 incorporated by reference.

9 2. “memory superposition”

10 Under Plaintiffs’ Infringement Contentions and apparent application of their interpretation
11 of the claim language, claims 4, 7, 15, are invalid under 35 U.S.C. §112(b) because the claim
12 language, when read in light of the intrinsic record, fails to inform a person of ordinary skill in the
13 art the scope of the claimed invention with reasonable certainty, including with respect to the
14 boundaries of what is included or not included in “memory superposition.” For example,
15 Plaintiffs Infringement Contentions allege this claim limitation is satisfied by pointing to the
16 alleged presence of “various types of memory to optimize performance (e.g., various memory
17 types like SRAM, TCAM, DRAM memory cache, DDR4/50,” without addressing the claim’s
18 additional requirement of a selectable arrangement of the plurality of types of memory as
19 multilevel cache, interleaved memory, swapped memory over a specified address range, and
20 memory allocated to specified address ranges according to parameterized task needs. *See, e.g.,*
21 ’753 Infringement Contentions at 55. Therefore, Plaintiffs’ Infringement Contentions and
22 apparent application of their interpretation of the claim language renders the claims invalid for
23 indefiniteness under 35 U.S.C. §112 because a POSITA would not be able to ascertain with
24 reasonable certainty what is required for “memory superposition” or whether it can simply be
25 satisfied by the presence of various types of memory as Plaintiffs allege.

1 **3. “one or more swappable pods or cards” / “the plurality of swappable**
2 **pods or cards”**

3 Claims 1-4 of the '753 patent are invalid as indefinite under 35 U.S.C. § 112 because the
4 claim language, when read in light of the intrinsic record, fails to inform a POSITA of the scope of
5 the claimed invention with reasonable certainty, including with respect to the claim terms “one or
6 more swappable pods or cards,” and “the plurality of swappable pods or cards.” For example,
7 claim 1 “one or more swappable pods or cards” (which would encompass a single swappable pod
8 or card) and then subsequently refers to “the plurality of swappable pods or cards” (which requires
9 two or more swappable pods or cards). This internal inconsistency within the claims renders the
10 claims invalid under § 112 because the '753 patent fails to inform with reasonable certainty
11 whether these claims can be satisfied by a single swappable pod or card. The basis for this
12 contention is further disclosed in Defendants’ claim construction briefing for the '753 patent
13 (Dkts. 79, 81); that analysis is hereby incorporated by reference.

14 **D. U.S. Patent No. 10,985,943**

15 For the claim term “programmable logic device” in the '943 patent, Defendants’ claim
16 construction briefing explained how Plaintiffs’ claim interpretations and proposed constructions
17 rendered these limitations indefinite under 35 U.S.C. § 112; that analysis and explanation is hereby
18 incorporated by reference. Dkts. 79, 81.

19 **1. “sandboxes”**

20 Claims 1 and 10, and all claims depending from 1 and 10 recite “sandboxes.” Under
21 Plaintiffs’ Infringement Contentions and apparent application of their interpretation of the claim
22 language, these claims are invalid under 35 U.S.C. § 112(b) because the claim language, when read
23 in light of the intrinsic record, fails to inform a person of ordinary skill in the art the scope of the
24 claimed invention with reasonable certainty with respect to the boundaries of the claim term
25 “sandboxes.”

26 The claims state that the sandboxes comprise part of a PLD and have “user-definable
27 programmable electronic circuits of the PLD.” '943 patent, claim 1 (limitation [1c]). The
28 specification similarly describes the sandboxes as being “for programmable electronic circuits”

1 “for user-defined functionality.” ’943 patent, 3:8-16, 12:54-65, 4:27-30 (“for user-defined
2 electronic circuits”), 5:19-29 (“for configurable modules”).

3 However, in contrast, Plaintiffs’ Infringement Contentions expand the scope of
4 “sandboxes” beyond PLDs to encompass portions ASICs that do not have user-definable
5 programmable circuits of a PLD. The intrinsic record does not provide any written description
6 support for sandboxes that do not have user-definable programmable circuits. Therefore, the
7 intrinsic record fails to inform a POSITA with reasonable certainty regarding the boundaries of
8 “sandboxes” under Plaintiffs’ apparent view of the claim, particularly when Plaintiffs’
9 Infringement Contentions appear to take an expansive view that a sandbox could encompass any
10 arbitrary portion of an integrated circuit, which renders the claims indefinite under 35 U.S.C.
11 §112(b). The basis for this contention is further disclosed in Defendants’ claim construction
12 briefing for the ’943 patent (Dkt. 79), and is hereby incorporated by reference.

13 **E. U.S. Patent No. 10,944,634**

14 **1. “network connected device”**

15 Claims 15-20 recite “network connected device” and are invalid as indefinite under 35
16 U.S.C. § 112. The term “network connected device,” as used in claims 15-20, even when read in
17 light of the specification and prosecution history, fail to inform with reasonable certainty one of
18 ordinary skill in the art about the scope of the invention.

19 Claim 15 recites that a network connected device is “acting as a tuner server.” ’634 patent,
20 15:46-47. Claims 16 and 20 recite the “network connected device of claim 15.” ’634 patent, 15:9,
21 16:37. At least claims 15, 16, and 20 include limitations that require the network connected
22 device to function as a distinct device or to function circularly. For example, Claim 15 requires
23 the “network connected device” to “collect; at a first endpoint device coupled to the network;
24 parameter values for determination of bandwidth of network connections to further endpoint
25 devices,” where “the network connected device acting as a tuner server [is] distinct from the first
26 endpoint device.” ’634 patent, 15:42-48. This requires the network connected device that acts as
27 a distinct device to perform the collecting action. As another example, claim 16 requires the
28 “network connected device” to “send the parameter values for determination of bandwidth, from

1 the first endpoint device, to the tuner server,” i.e., to send parameter values to itself, which is
2 circular. *Id.*, 16:11. The claim also requires that the network connected device “send, from the first
3 endpoint device to the tuner server, connection information for the next network connection” and
4 “receive, at the first endpoint device, from the tuner server, the one or more parameter values for
5 the transmission bandwidth,” which is also circular. *Id.*, 16:15, 20. Similarly, claim 20 requires
6 the “network connected device” to “transmit data associated with the network traffic from the first
7 endpoint to the network connected device” and to “receive data for adjusting the transmission
8 bandwidth from the network connected device,” i.e., transmitting data to and receiving data from
9 itself, which is likewise circular. *Id.*, 16:41-42, 45-46. Because claim 15 and claims 16-20 that
10 depend on claim 15 require a network connected device to function as a distinct device or to
11 function circularly, a POSITA would not have understood their meaning. Thus, claims 15-20 are
12 invalid under § 112. The basis for this contention is disclosed in Defendants’ claim construction
13 briefing for the ’634 patent (Dkts. 79, 81), and is hereby incorporated by reference.

14 Moreover, claims 15-20 are invalid for lack of written description. For example, the
15 specification of the ’634 patent fails to describe how one or more processors at a tuner server can
16 collect parameter values, receive connection information, or receive data at an endpoint device
17 that is distinct from the tuner server. A POSITA reading the ’634 patent would not have
18 understood that its inventors had possession of processors at the tuner server performing the
19 aforementioned functionalities at the endpoint device. In their claim construction brief (Dkt. 80),
20 Plaintiffs argued that the network connected device “control[s]” the endpoint device to perform
21 operations. This argument contradicts the claim language and essentially requires redrafting the
22 claims. But even if one were to accept Plaintiffs’ argument, written description would still be
23 lacking. In particular, the specification lacks written description of control of the tuner server over
24 the endpoint device. For example, a POSITA would have expected that any control of the tuner
25 server over the endpoint device would require a mechanism for the tuner server to exert such
26 control. However, the specification does not describe the tuner server providing instructions or
27 commands to endpoint devices to perform the aforementioned operations, nor does the
28

1 specification describe communication protocols or mechanisms for the tuner server to control the
2 endpoint devices.

3 Claims 15-20 are also invalid for lack of enablement. The '634 patent does not include
4 any description that would have allowed a POSITA to implement processors at a network
5 connected device acting as a tuner server that collect parameter values, receive connection
6 information, or receive data at an endpoint device that is distinct from the tuner server. A
7 POSITA would not have known how such a configuration can be implemented and would not
8 have been able to implement the claimed invention in claims 15-20 without undue
9 experimentation, if at all.

10 **2. “matches a geographical area of a past network connection”**

11 Claims 1, 8, and 15, and all claims depending from them are invalid under 35 U.S.C. § 112
12 because the '634 patent fails to provide sufficient written description of the claimed determination
13 that a next network connection matches a geographical area of a past network connection. The
14 '634 patent describe examples of geographical area to include “an apartment complex,
15 neighborhood, or city to city or state to state network connection” '634 patent, 2:54-56. For
16 example, a POSITA would have understood that there was no commonly known way to determine
17 the apartment complex that a network connection is in or to determine whether two different
18 connections are in the same apartment complex. The '634 patent includes no description as to
19 how to determine two network connections have matching geographical areas. Given its lack of
20 relevant description, a POSITA reading the '634 patent would have understood that the inventors
21 did not have possession of this claimed feature.

22 All claims of the '634 patent are also invalid for lack of enablement. The '634 patent does
23 not include any description that would have informed a POSITA as to how to determine that two
24 network connections match in terms of geographical area, such as being in the same
25 neighborhood. A POSITA would not have understood any such determination to be practical or
26 reliable. At the very least, because the '634 patent does not disclose any method of determining
27 that a network connection matches the geographical area of another network connection, a
28

1 POSITA would not have been able to implement the claimed invention without undue
2 experimentation.

3 **3. “initiate/initiating the next network connection ... with a transmission**
4 **bandwidth based on the parameter values for the past network**
5 **connection”**

6 Claims 1, 8, and 15 and all claims depending from them are invalid under 35 U.S.C. § 112
7 because the '634 patent fails to provide sufficient written description of initiating the next network
8 connection with a transmission bandwidth based on the parameter values for the past network
9 connection. For example, a POSITA would have understood that commonly known and widely
10 used transport layer protocols, such as TCP, use dynamic and adaptively changing transmission
11 bandwidth. The transmission bandwidth of a network connection often depends upon factors that
12 are not fully controllable based on settings and parameters. The '634 patent does not describe how
13 to initiate a next network connection with a transmission bandwidth based on parameter values of
14 a past network connection. A POSITA reading the '634 patent would have understood that the
15 inventors did not have possession of this claimed feature.

16 All claims of the '634 patent are also invalid for lack of enablement. The '634 patent does
17 not include any description that would have informed a POSITA as to how to initiate a new
18 network connection with a transmission bandwidth based on parameter values of a past network
19 connection. A POSITA would have found various difficulties in attempting to implement this
20 feature under common schemes such as TCP, as this feature deviates from TCP's mechanism of
21 dynamically adjusting the transmission bandwidth using a mechanism other than one that is based
22 on parameter values of a past connections. A POSITA would not have been able to implement
23 this claimed feature without undue experimentation.

24 **4. “wherein the first endpoint device is contained within the network**
25 **connected device”**

26 Asserted Claim 17 of the '634 patent requires “wherein the first endpoint device is
27 contained within the network connected device.” *See* '634 patent, 16:25-26. Claim 17 depends
28 from claim 15, which recites that “the network connected device [is] distinct from the first
endpoint device.” The '634 patent, however, lacks written description for the claimed feature in

1 claim 17. For example, the '634 patent does not describe the structure of the network connected
2 device such that it is both distinct from the first endpoint device and contains the first endpoint
3 device. Furthermore, the '634 patent describes “offload[ing] the storage and bandwidth
4 determination (e.g., calculations) from the endpoints ... to the tuning server” ('634 patent,
5 4:42:44) but does not describe how that can be achieved when the network connected device,
6 acting as the tuner server, contains the first endpoint device. A POSITA reading the '634 patent
7 would have understood that the inventors did not have possession of claimed feature in claim 17.

8 Claim 17 is also invalid for lack of enablement. The '634 patent does not include any
9 description that would have informed a POSITA as to how to structure the network connected
10 device and the first endpoint device such that the former is distinct from, but contains, the latter.
11 In light of the absence of description, a POSA would not have been able to implement this claimed
12 feature without undue experimentation.

13 **III. Identification of Claims Defendants Contend Are Directed to Ineligible Subject** 14 **Matter Under 35 U.S.C. § 101**

15 **A. U.S. Patent No. 9,529,767**

16 Defendants contend that all claims of the '767 patent are invalid under 35 U.S.C. §101 for
17 claiming an abstract idea. The basis for this contention is disclosed in Defendants' motion to
18 dismiss for failure to state a claim under 35 U.S.C. § 101 for the '767 patent (Dkt. 156), which is
19 hereby incorporated by reference, and further disclosed herein.

20 Under *Alice* step one, the Asserted Claims of the '767 patent are directed to an abstract
21 idea. The '767 patent purports to solve a problem relating to the field of storage devices.
22 Describing the then state-of-the-art, the '767 patent describes how a host (a computer) uses an
23 interface protocol, such as Serial Advanced Technology Attachment (also known as Serial ATA or
24 SATA) or Serial Attached SCSI (also known as SAS), to interface with a storage media device
25 (e.g., hard drive or solid state drive). '767 pat. at 21-25. The '767 patent describes some of the
26 differences between the SATA and SAS interface protocols. *Id.* at 1:25-38. The '767 patent
27 observes that if a host needs to access a mixture of storage devices (e.g., some SATA drives and
28 some SAS drives), “the host must support and manage both protocols,” and that one consequence

1 is that “[t]he host software layers interfacing to SATA and SAS must support their high protocol
2 overhead.” *Id.* at 1:36-42.

3 The ’767 patent is directed toward its proposed solution, namely, to implement an
4 abstraction protocol that is used by the host to communicate with storage controllers, which in turn
5 translate the abstract, high-level commands from the host into native commands understood by the
6 storage device (e.g., SATA or SAS):

7 This invention simplifies a host that accesses SATA and SAS storage media devices
8 by abstracting the SATA and SAS protocols with one full duplex protocol that
9 supports full command queuing to each storage media device, whether SATA or
10 SAS. The invention is comprised of both hardware (system—architecture) and
software (method—abstraction protocol). The simplified abstraction protocol
provides high-level access commands to storage media devices.

11 *Id.* at 3:65-4:6.

12 The storage media devices SATA or SAS interfaces are pushed to the edge, away
13 from the host and toward the storage media device for which it supports. Pushing the
14 SAS/SATA protocol to the edge simplifies the host/controller since it does not
implement the legacy SAS/SATA protocol, it can do what fits its application best.

15 *Id.* at 4:10-16. However, the ’767 patent’s solution is directed to the abstract idea of translating
16 commands or instructions from one format to another. *See Automation Middleware Solutions, Inc.*
17 *v. Invensys Sys., Inc.*, Case No. 2:15-CV-898, 2017 WL 11573945 (E.D. Tex. Mar. 31, 2017)
18 (finding claims “directed to the abstract idea of ‘a middle translating layer’”).

19 Using a translation layer to translate commands or instructions in a higher-level format to
20 commands or instructions that are a lower-level format supported by a particular device is an
21 abstract idea that has long been applied in the field of computing, which includes storage
22 technology. For example, in an earlier era of computing, computer programs were written using
23 the instruction set defined by the computer’s manufacturer. Computers by different manufacturers
24 often used different and incompatible instruction sets. Even different computers made by the
25 same manufacturer used different and incompatible instructions sets. This led to the development
26 of programming languages that operate on a more abstract basis, such as FORTRAN, COBOL, C,
27 and C++. Instead of writing a computer program in the instruction set of a specific computer, the
28 programmer writes the program in a given programming language. A compiler then is used to

1 translate the program into the instruction set used by a particular computer. Java programming is
2 another application of this abstract idea. In Java, the programmer writes code in Java, and it is
3 compiled into a bytecode that can be understood by a Java runtime environment. When the
4 program is run on a computer, the Java runtime environment translates the bytecode into an
5 instruction set used by the computer.

6 Under *Alice* step 2, the Asserted Claims of the '767 patent do not contain an inventive
7 concept. The recited abstraction protocol is nothing more than the abstract idea itself. Instead, the
8 abstraction protocol is merely described in functional terms, or in terms of objectives sought to be
9 achieved from using an abstraction protocol:

10 [T]he present invention is related to a system and method for abstracting SATA
11 and/or SAS storage media devices via a full duplex queued command interface to
12 increase performance, lower host over head, and simplify scaling storage media
devices and system.

13 '767 pat. at 1:14-19.

14 [T]he master controller and the edge controller communicate via the abstraction
15 protocol, the abstraction protocol comprising a full-duplex protocol supporting full
command queuing for the at least one storage media device.

16 *Id.* at 2:36-40.

17 The abstraction protocol supports flow control between the host and the edge.

18 *Id.* at 4:30-32.

19 This abstraction protocol is performance-centric and supports common high-level
20 read and write access to a pool of storage media devices, each of which may have a
21 SATA or SAS interface. The abstraction protocol is link-agnostic and may be carried
22 via a multiplicity of direct attach or networked interfaces, including but not limited
23 to PCIe, Ethernet (e.g., 1GbE, 10GbE, 40GbE, or 100GbE), Infiniband, USB,
ThunderBolt, Firewire, and/or custom interfaces. The abstraction protocol link
bandwidth need not be related to the bandwidth of the individual storage media
devices SATA or SAS interface.

24 *Id.* at 4:35-45.

25 The abstraction protocol supports packet-switched routable fabric and error recovery.
26 Routing may be based on logical address, WWN (World Wide Number—a unique
27 global ID number), or physical address. Logical address or WWN may be mapped
28 to a physical address (e.g. by table lookup). This mapping may be static or
dynamically updated based upon a discovery process as storage media devices are
(i) inserted or removed from the system and/or (ii) powered on or off.

1 *Id.* at 4:55-63; *see also id.* at 3:65-4:6 and 4:10-16. The Federal Circuit has repeatedly stated that
2 “a claim is abstract if it describes a process in ‘result-based functional language’ and fails to
3 ‘sufficiently describe how to achieve these results in a non-abstract way.’” *Voip-Pal.Com, Inc. v.*
4 *Apple Inc.*, 411 F. Supp. 3d 926, 951 (N.D. Cal. 2019) (quoting *Two-Way Media Ltd. v. Comcast*
5 *Cable Comm’cs, LLC*, 874 F.3d 1329, 1337 (Fed. Cir. 2017); *see also Bridge & Post, Inc. v.*
6 *Verizon Commc’ns, Inc.*, 778 Fed. Appx. 882, 895 (Fed. Cir. 2019) (Bryson, CJ concurring) (“The
7 distinction between claims that recite functions or results (the ‘what it does’ aspect of the
8 invention) and those that recite concrete means for achieving particular functions or results (the
9 ‘how it does it’ aspect of the invention) is an important indicator of whether a claim is directed to
10 an abstract idea.”). “Indeed, the essentially result-focused, functional character of claim language
11 has been a frequent feature of claims held ineligible under § 101.” *Elec. Power Grp., LLC v.*
12 *Alstom S.A.*, 830 F.3d 1350, 1356 (Fed. Cir. 2016).

13 Notably, the specification and figures do not disclose a single example of the allegedly
14 inventive abstraction protocol or provide clear guidance for a POSITA to implement the allegedly
15 inventive abstraction protocol—highlighting that the abstraction protocol is little more than the
16 abstract idea itself.

17 The other recited elements of the claim (e.g., packet-switched networks, hosts, master
18 controllers, edge controllers, and storage media devices) do not provide an inventive concept. For
19 example, the ’767 patent does not purport to require any modifications to any packet-switched
20 networks or storage media devices, which are otherwise well-understood, routine and conventional
21 in the field of storage technology.

22 In the field of storage technology, a host (e.g., computer) was well-understood, routine and
23 conventional. The ’767 patent does not describe an inventive host, and the specification merely
24 describes that the host uses the abstraction protocol to issue commands to storage media devices
25 instead of using legacy protocols and commands (e.g., SATA or SAS). *Id.* at 3:65-4:6 and 4:10-
26 16.

27 In the field of storage technology, storage controllers were well-understood, routine, and
28 conventional. The ’767 patent claims master controllers and edge controllers, but does not

1 describe how they are inventive other than by functional description that they implement the
2 abstraction protocol:

3 A master controller 304, ... , contains (i) one or more master processing elements
4 306 that *carries out the processing of the master side of the abstraction protocol* and
5 communicates with at least one host 302 via a master controller interface 308
6 connected to the host 302 by a host link 320 and (ii) one or more control link
7 interfaces 310 through 312, which communicate with one or more edge controllers
8 316 through 318 via control links 320 and 322.... Together, the host 302, the master
9 controller 304, and edge controllers 316 through 318 form a routable packet-switched
10 network supported by an abstraction protocol.

11 *E.g., id.* at 4:64-5:13 (emphasis added).

12 The edge controller 402 contains (i) one or more edge processing elements 404 that
13 *carries out the processing of the edge side of the abstraction protocol* and
14 communicates with one or more master controllers 406 & 408 via control link
15 interfaces 410 & 412 communicating with a master controllers 406 & 408 via control
16 links 414 & 416; (ii) one or more SMD interfaces 418 & 420 which communicate
17 with one or more SMDs 422 & 424 via SATA and/or SAS links 426 and 428; (iii)
18 optionally, one or more forward link interfaces 430 communicating with other edge
19 controllers 432 via control links 448 used to forward those abstract protocol packets
20 from master controllers 406 or 408 (received via control link interfaces 410 & 412
21 that are not destined for the edge controller 402 that received them) to another edge
22 controller 432; and (iv) optionally, one or more master controllers 434 & 436
23 communicating with other edge controllers 438 & 440 via control links 442 & 444
24 when an edge processing element 404 wishes to *author abstract protocol messages*
25 to other edge controllers 438 & 440.

26 *E.g., id.* at 5:32-51 (emphasis added). Moreover, using a storage controller to translate commands
27 from one protocol/format into a protocol/format compatible with a storage media device was well-
28 understood, routine, and conventional. *See, e.g.,* Ex. 767-05 (Malwankar) at limitations [1d] and
[16e], Ex. 767-10 (Voorhees '224) at limitations [1d] and [16e].

29 **B. U.S. Patent No. 10,944,634**

30 Defendants contend that all claims of the '634 patent are invalid under 35 U.S.C. §101 for
31 claiming an abstract idea. The basis for this contention is disclosed in Defendants' motions to
32 dismiss for failure to state a claim under 35 U.S.C. § 101 for the '634 patent (Dkt. 23, 32, 156),
33 and is hereby incorporated by reference. In addition, the prior art identified above for the '634
34 patent is further evidence that the claimed elements are well-understood, routine, or conventional.

1 **IV. Document Production Accompanying Invalidity Contentions**

2 **Patent L.R. 3-4(a)** (“Source code, specifications, schematics, flow charts, artwork,
3 formulas, or other documentation sufficient to show the operation of any aspects or elements of an
4 Accused Instrumentality identified by the patent claimant in its Patent L.R. 3-1(c) chart.”):
5 Defendants have produced the source code for the Accused Instrumentalities for inspection at the
6 office of Sheppard Mullin in Menlo Park, California.

7 **Patent L.R. 3-4(b)** (“A copy or sample of the prior art identified pursuant to Patent L.R.
8 3-3(a) which does not appear in the file history of the patent(s) at issue. To the extent any such
9 item is not in English, an English translation of the portion(s) relied upon shall be produced.”):
10 Defendants have produced the prior art referenced in these invalidity contentions to Plaintiffs,
11 including in the documents bates labeled PA000001–PA040826.

12 Additional prior art referenced in these invalidity contentions is being produced to
13 Plaintiffs concurrently with service of these invalidity contentions labeled PA040827–PA068833,
14 AMD-CV-0065505–AMD-CV-0065651, AMD-CV-0110842–AMD-CV-0149658. Source code
15 for prior art has been made available for inspection.

16 **Patent L.R. 3-4(c)** (“All agreements that may be related to the accused instrumentality or
17 may be comparable to a license that would result from a hypothetical reasonable royalty
18 negotiation.”):
19 Defendants have produced responsive agreements to Plaintiffs, including documents bates labeled:
20 AMD-CV-0024567; AMD-CV-0030880; AMD-CV-0096930; AMD-CV-0096945; AMD-CV-
21 0097020; AMD-CV-0097025; AMD-CV-0097084; AMD-CV-0109074; AMD-CV-0109099;
22 AMD-CV-0109127; AMD-CV-0109152; AMD-CV-0109182; AMD-CV-0109334; AMD-CV-
23 0109351; AMD-CV-0109381; AMD-CV-0109392; AMD-CV-0109404; AMD-CV-0109417;
24 AMD-CV-0109837; AMD-CV-0109902; AMD-CV-0109906; AMD-CV-0110685; AMD-CV-
25 0110688; AMD-CV-0110693; AMD-CV-0110761; AMD-CV-0110786; AMD-CV-0110802;
26 AMD-CV-0110817.

27 Additional responsive agreements are being produced to Plaintiffs.
28

1 **Patent L.R. 3-4(d)** (“Documents sufficient to show the sales, revenue, cost, and profits for
2 accused instrumentalities identified pursuant to Patent L.R. 3-1(b) for any period of alleged
3 infringement.”):

4 Defendants have produced responsive documents showing sales, revenue, cost, and profits for
5 accused instrumentalities to Plaintiffs, including documents bates labeled: AMD-CV-0005255;
6 AMD-CV-0005256; AMD-CV-0008996; AMD-CV-0063058; AMD-CV-0063196; AMD-CV-
7 0063299; AMD-CV-0063375; AMD-CV-0063457; AMD-CV-0065662; AMD-CV-0065783;
8 AMD-CV-0065958; AMD-CV-0066066; AMD-CV-0066128; AMD-CV-0066252; AMD-CV-
9 0066341; AMD-CV-0066419; AMD-CV-0109441; AMD-CV-0109523; AMD-CV-0109626;
10 AMD-CV-0109698; AMD-CV-0109836.

11 **Patent L.R. 3-4(e)** (“All agreements that may be used to support the party denying
12 infringement’s damages case.”):

13 Defendants have produced responsive agreements to Plaintiffs, including documents bates labeled:
14 AMD-CV-0024567; AMD-CV-0030880; AMD-CV-0096930; AMD-CV-0096945; AMD-CV-
15 0097020; AMD-CV-0097025; AMD-CV-0097084; AMD-CV-0109074; AMD-CV-0109099;
16 AMD-CV-0109127; AMD-CV-0109152; AMD-CV-0109182; AMD-CV-0109334; AMD-CV-
17 0109351; AMD-CV-0109381; AMD-CV-0109392; AMD-CV-0109404; AMD-CV-0109417;
18 AMD-CV-0109837; AMD-CV-0109902; AMD-CV-0109906; AMD-CV-0110685; AMD-CV-
19 0110688; AMD-CV-0110693; AMD-CV-0110761; AMD-CV-0110786; AMD-CV-0110802;
20 AMD-CV-0110817; CV-AMD0000803; CV-AMD0000807; CV-AMD0000815; CV-
21 AMD0000819; CV-AMD0000827; CV-AMD0000835; CV-AMD0000843; CV-AMD0000847;
22 CV-AMD0000864; CV-AMD0000870; CV-AMD0000885; CV-AMD0000890; CV-
23 AMD0001280; CV-AMD0001482.

24 Additional responsive agreements are being produced to Plaintiffs.
25
26
27
28

1 DATED: February 13, 2026

2 SHEPPARD, MULLIN, RICHTER & HAMPTON LLP

3
4 By

/s/ Harper S. Batts

HARPER S. BATTS

CHRISTOPHER S. PONDER

JEFFREY LIANG

LAI YIP

Attorneys for Advanced Micro Devices, Inc. and
Pensando Systems, Inc.

1 **CERTIFICATE OF SERVICE**

2 I HEREBY CERTIFY that on this day, February 13, 2026, a true and correct copy of the foregoing
3 document, **DEFENDANTS ADVANCED MICRO DEVICES, INC. and PENSANDO**
4 **SYSTEMS, INC.'s INVALIDITY CONTENTIONS AND ACCOMPANYING**
5 **DOCUMENT PRODUCTION** was served by the indicated means to the persons at the addresses
6 below:

7 Michael Matulewicz-Crowley
8 Christine E. Lehman
9 Adam Adler
10 Courtland L. Reichman
11 Jennifer Estremera
12 Laura Carwile
13 Scott Cole
14 Ariane S. Mann

***Counsel for Plaintiffs Concurrent Ventures, LLC
and XtreamEdge, Inc.***

Via Email:

mmatulewicz-crowley@reichmanjorgensen.com
clehman@reichmanjorgensen.com
aadler@reichmanjorgensen.com
creichman@reichmanjorgensen.com
jestremera@reichmanjorgensen.com
lcarwile@reichmanjorgensen.com
scole@reichmanjorgensen.com
amann@reichmanjorgensen.com

15 **REICHMAN JORGENSEN LEHMAN**
16 **& FELDBERG LLP**

17 Dated: February 13, 2026

18 By /s/ Harper S. Batts

19 HARPER S. BATTS

20 *Attorneys for Defendants*
21 *Advanced Micro Devices, Inc. and*
22 *Pensando Systems, Inc.*