

**IN THE UNITED STATES DISTRICT COURT  
FOR THE DISTRICT OF DELAWARE**

SEQUOIA TECHNOLOGY, LLC,

Plaintiff,

v.

HITACHI LTD. AND HITACHI VANTARA  
CORPORATION,

Defendants.

C.A. No. \_\_\_\_\_

**JURY TRIAL DEMANDED**

**COMPLAINT FOR PATENT INFRINGEMENT**

This is an action for patent infringement in which Plaintiff Sequoia Technology, LLC (“Sequoia”) demands trial by jury and makes the following allegations against Defendants Hitachi Ltd. and Hitachi Vantara Corporation (collectively, “Hitachi” or “Defendants”):

**NATURE OF THE ACTION**

1. This is a patent infringement action arising under the Patent Laws of the United States, 35 U.S.C. §§ 1, *et seq.* including without limitation, 35 U.S.C. §§ 271, 281.

**PARTIES**

2. Plaintiff Sequoia Technology, LLC is a Delaware Limited Liability Company.

3. Defendant Hitachi, Ltd. is a Japanese corporation organized and existing under the laws of Japan, with its principal place of business located at 6-6, Marunouchi 1-chome, Chiyoda-ku, Tokyo, 100-8280, Japan. On information and belief, Hitachi may be served pursuant to the provisions of the Hague Convention or through its subsidiary Hitachi Vantara Corporation.

4. Defendant Hitachi Vantara Corporation is a Delaware corporation. Hitachi Vantara Corporation's registered agent for service of process in Delaware is Corporation Service Company at 251 Little Falls Drive, Wilmington, DE 19808.

#### **JURISDICTION AND VENUE**

5. This Court has subject matter jurisdiction pursuant to 28 U.S.C. §§ 1331 and 1338(a).

6. This Court has personal jurisdiction over Defendants because, among other reasons, Defendant Hitachi Vantara Corporation is incorporated under the laws of the State of Delaware, Defendants have done business in this District, have committed and continue to commit acts of patent infringement in this District, and have harmed and continue to harm Sequoia in this District, by, among other things, using, selling, offering for sale, and importing infringing products and services in this District.

7. Venue is proper in this District under 28 U.S.C. §§ 1391(b)-(c) and 1400(b) because, among other reasons, Defendant Hitachi Vantara Corporation is incorporated in Delaware, Defendants are subject to personal jurisdiction in this District, and have committed and continue to commit acts of patent infringement in this District. On information and belief, Defendants conduct substantial business directly and/or through third parties or agents in this judicial district by selling and/or offering to sell the infringing products and/or by conducting other business in this judicial district. Furthermore, Plaintiff has been harmed by Defendants' conduct, business transactions and sales in this district.

#### **BACKGROUND**

8. Sequoia Technology, LLC is the exclusive licensee of the United States Patent No. 6,718,436 ("the '436 Patent" or "the Patent-in-Suit") that issued on April 6, 2004 and is

titled “Method for managing logical volume in order to support dynamic online resizing and software RAID and to minimize metadata and computer readable medium storing the same.” A true and correct copy of the ’436 Patent is attached as Exhibit A. Sequoia has all the substantial rights to sue for infringement and collect past and future damages for the infringement of the ’436 Patent.

9. The ’436 Patent was invented by Chang-Soo Kim, Gyoung Bae Kim and Bum Joo Shin of the Electronics and Telecommunications Research Institute (“ETRI”). ETRI is the national leader in Korea in the research and development of information technologies. Since its inception in 1976, ETRI has developed new technologies in DRAM computer memory, CDMA and 4G LTE cellular phone communications, LCD displays, as well as large-scale computer storage, the technology at issue in this case. ETRI employs over 2,034 research/technical staff, of whom 94% hold a post-graduate degree and 50% have earned a doctoral degree in their technological field. Over the last five years, ETRI has applied for a total of 16,917 patents, has contributed 8,337 proposals that have been adopted by international and domestic standard organizations, and has published over 1,282 articles in peer-reviewed technology publications.

10. Defendants manufacture, provide, use, sell, offer for sale, import, and/or distribute infringing products and services (“Accused Products, Systems and/or Services”). The Accused Products, Systems and/or Services include “Logical Volume Manager (“LVM”) Dynamic Resizing Products, Systems and/or Services.” For example, LVM Dynamic Resizing Products, Systems and/or Services include without limitation products and services with operating systems Red Hat Enterprise Linux Operating Systems versions 4 and later, and CentOS version 6 and later. For example, LVM Dynamic Resizing Products, Systems and/or Services include without

limitation the list of products listed in Exhibit B. Defendants support and encourage others to use its products and services in an infringing manner, including its customers, as set forth herein.

11. Sequoia, as the exclusive licensee of the '436 Patent, seeks the damages owed for Hitachi's use of ETRI's pioneering technology.

12. The '436 Patent has been cited by at least 152 issued patents as relevant prior art.

**COUNT I**  
**(INFRINGEMENT OF U.S. PATENT NO. 6,718,436)**

13. Sequoia incorporates and re-alleges every allegation set forth above, as though fully set forth herein.

14. Hitachi has directly infringed, contributed to the infringement of, and/or has induced the infringement of one or more claims of the '436 Patent in violation of 35 U.S.C. § 271 by, at least, making, using, supplying, distributing, importing, exporting, selling and/or offering for sale in the United States or by intending that others make, use, supply, distribute, import, export, sell, and/or offer for sale in the United States LVM Dynamic Resizing Products, Systems and/or Services (as exemplified, but not limited to, the Accused Products, Systems and/or Services in Paragraph 10 incorporated herein) that practice and/or are covered by one or more claims of the '436 Patent.

15. Hitachi's LVM Dynamic Resizing Products, Systems and/or Services infringe, either directly or under the doctrine of equivalents, at least claims 1-3 of the '436 Patent. Hitachi makes, uses, sells, offers for sale, supplies and/or distributes within the United States and/or imports and/or exports the Accused Products, Systems and/or Services and thus directly infringes the '436 Patent under 35 U.S.C. § 271(a).

16. As just one non-limiting example, set forth below (with claim language in italics) is a description of infringement of exemplary claim 1 of the '436 patent. Sequoia reserves the

right to modify this description, including on the basis of information it obtains during discovery:

*a method for managing a logical volume in order to support dynamic online resizing and minimizing a size of metadata, said method comprising steps of:* To the extent the preamble is limiting, the Accused Products, Systems and/or Services include a method for managing a logical volume in order to support dynamic online resizing and minimizing a size of metadata.


*See, e.g., Red Hat Enterprise Linux 5 Logical Volume Manager Administration LVM*

*Administrator's Guide Edition 1, at p. 12, reproduced below:*

## **1.2. Logical Volumes**

Volume management creates a layer of abstraction over physical storage, allowing you to create logical storage volumes. This provides much greater flexibility in a number of ways than using physical storage directly. With a logical volume, you are not restricted to physical disk sizes. In addition, the hardware storage configuration is hidden from the software so it can be resized and moved without stopping applications or unmounting file systems. This can reduce operational costs.

17. *“a) creating the logical volume by gathering disk partitions in response to a request for creating the logical volume in a physical storage space;”* The Accused Products, Systems and/or Services support the gathering plurality of disk partitions in physical storage space. *See, e.g., Red Hat Enterprise Linux 5, Logical Volume Manager Administration LVM Administrator's Guide Edition 1, at p. 14 with Figure 1.1., reproduced below:*



**Figure 1.1. LVM Logical Volume Components**

See, e.g., [https://access.redhat.com/documentation/en-us/red\\_hat\\_enterprise\\_linux/5/html/logical\\_volume\\_manager\\_administration/lv](https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/5/html/logical_volume_manager_administration/lv) reproduced below:

“A logical volume will have an attribute of "M" if it is a mirrored volume that was created without an initial synchronization, and it will have an attribute of "m" if it was created with initial synchronization. The following command displays the attributes of a mirrored logical volume named lv that was created without initial synchronization.”

```
# lvs vg
LV VG Attr LSize Pool Origin Snap% Move Log Copy% Convert
lv vg Mwi-a-m- 5.00g lv_mlog 100.00
```

18. “b) generating the metadata including information of the logical volume and the disk partitions forming the logical volume and storing the metadata to the disk partitions forming the logical volume;” The Accused Products include this element. See, e.g., Red Hat Enterprise Linux 5, LVM Administrator’s Guide at pp. 17 and 18 with Figure 2.1:

## 2.1. Physical Volumes

The underlying physical storage unit of an LVM logical volume is a block device such as a partition or whole disk. To use the device for an LVM logical volume the device must be initialized as a physical volume (PV). Initializing a block device as a physical volume places a label near the start of the device.

By default, the LVM label is placed in the second 512-byte sector. You can overwrite this default by placing the label on any of the first 4 sectors. This allows LVM volumes to co-exist with other users of these sectors, if necessary.


An LVM label provides correct identification and device ordering for a physical device, since devices can come up in any order when the system is booted. An LVM label remains persistent across reboots and throughout a cluster.

The LVM label identifies the device as an LVM physical volume. It contains a random unique identifier (the UUID) for the physical volume. It also stores the size of the block device in bytes, and it records where the LVM metadata will be stored on the device.

The LVM metadata contains the configuration details of the LVM volume groups on your system. By default, an identical copy of the metadata is maintained in every metadata area in every physical volume within the volume group. LVM metadata is small and stored as ASCII.

Currently LVM allows you to store 0, 1 or 2 identical copies of its metadata on each physical volume. The default is 1 copy. Once you configure the number of metadata copies on the physical volume, you cannot change that number at a later time. The first copy is stored at the start of the device, shortly after the label. If there is a second copy, it is placed at the end of the device. If you accidentally overwrite the area at the beginning of your disk by writing to a different disk than you intend, a second copy of the metadata at the end of the device will allow you to recover the metadata.

For detailed information about the LVM metadata and changing the metadata parameters, see [Appendix D, LVM Volume Group Metadata](#).



**Figure 2.1. Physical Volume layout**

19. “(c) *dynamically resizing the logical volume in response to a request for resizing, and modifying the metadata on the disk partitions forming the logical volume;*” The Accused Products, Systems and/or Services include this element. *See, e.g.,* [https://access.redhat.com/documentation/en-us/red\\_hat\\_enterprise\\_linux/5/html/logical\\_volume\\_manager\\_administration/lv:](https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/5/html/logical_volume_manager_administration/lv:)

```
# lvextend -L +5G vg/lv --nosync
Extending 2 mirror images.
Extending logical volume lv to 5.02 GiB
Logical volume lv successfully resized
```

See, e.g., Red Hat., Red Hat Enterprise Linux 5, Logical Volume Manager Administration LVM Administrator's Guide Edition 1, at pp. 12 and 114:

## 1.2. Logical Volumes

Volume management creates a layer of abstraction over physical storage, allowing you to create logical storage volumes. This provides much greater flexibility in a number of ways than using physical storage directly. With a logical volume, you are not restricted to physical disk sizes. In addition, the hardware storage configuration is hidden from the software so it can be resized and moved without stopping applications or unmounting file systems. This can reduce operational costs.

Logical volumes provide the following advantages over using physical storage directly:

- » Flexible capacity  
When using logical volumes, file systems can extend across multiple disks, since you can aggregate disks and partitions into a single logical volume.
- » Resizeable storage pools  
You can extend logical volumes or reduce logical volumes in size with simple software commands, without reformatting and repartitioning the underlying disk devices.

## LVM Volume Group Metadata

The configuration details of a volume group are referred to as the metadata. By default, an identical copy of the metadata is maintained in every metadata area in every physical volume within the volume group. LVM volume group metadata is stored as ASCII.

A metadata area is a circular buffer. New metadata is appended to the old metadata and then the pointer to the start of it is updated.

20. “(d) calculating and returning a physical address corresponding to a logical address of the logical volume by using mapping information of the metadata containing information of the physical address corresponding to the logical address;” The Accused Products, Systems and/or Services include this element. See, e.g., The Linux Volume Manager, <http://www.redhat.com/magazine/009jul05/features/lvm2/>, including Figure 2 reproduced below:



[Figure 2. Mapping logical extents to physical extents](#) shows how the logical volumes are mapped onto physical volumes. Each PV consists of a number of fixed-size physical extents (PEs); similarly, each LV consists of a number of fixed-size logical extents (LEs). (LEs and PEs are always the same size, the default in LVM 2 is 4 MB.) An LV is created by mapping logical extents to physical extents, so that references to logical block numbers are resolved to physical block numbers. These mappings can be constructed to achieve particular performance, scalability, or availability goals.




Figure 2. Mapping logical extents to physical extents

21. “wherein the metadata includes a disk partition table containing information of a disk partition in which the metadata is stored;” The Accused Products, Systems and/or Services include a disk partition table containing information of a disk partition in which the metadata is stored. *See, e.g.,* Red Hat Enterprise Linux 5, LVM Administrator’s Guide at pp. 17, 18 and 114:

**2.1.1. LVM Physical Volume Layout**

[Figure 2.1, “Physical Volume layout”](#) shows the layout of an LVM physical volume. The LVM label is on the second sector, followed by the metadata area, followed by the usable space on the device.




Figure 2.1. Physical Volume layout

### D.1. The Physical Volume Label

By default, the **pvcreate** command places the physical volume label in the 2nd 512-byte sector. This label can optionally be placed in any of the first four sectors, since the LVM tools that scan for a physical volume label check the first 4 sectors. The physical volume label begins with the string **LABELONE**.

The physical volume label Contains:

- › Physical volume UUID
- › Size of block device in bytes
- › NULL-terminated list of data area locations
- › NULL-terminated lists of metadata area locations

Metadata locations are stored as offset and size (in bytes). There is room in the label for about 15 locations, but the LVM tools currently use 3: a single data area plus up to two metadata areas.

22. “a logical volume table for maintaining the information of the logical volume by storing duplicated information of the logical volume onto all disk partitions of the logical volume;” The Accused Products, Systems and/or Services include this element. See, e.g., Red Hat Enterprise Linux 5, LVM Administrator’s Guide at p. 49.

```
# lvsdisplay /dev/new_vg/lvol0
--- Logical volume ---
LV Name           /dev/new_vg/lvol0
VG Name           new_vg
LV UUID           LBy1Tz-sr23-0jsI-LT03-nHLC-y8XW-EhC178
LV Write Access   read/write
LV snapshot status source of
                  /dev/new_vg/newvgsnap1 [active]
LV Status         available
# open            0
LV Size           52.00 MB
Current LE        13
Segments          1
Allocation        inherit
Read ahead sectors 0
Block device      253:2
```

23. “an extent allocation table for indicating whether each extent in the disk partition is used or not used;” The Accused Products, Systems and/or Services include this element. See, e.g., <https://www.redhat.com/archives/linux-lvm/2009-February/msg00090.html>.

```

pvdisplay -m reports:

--- Physical volume ---
PV Name /dev/sda2
VG Name VolGroup00
PV Size 931.32 GB / not usable 7.11 MB
Allocatable yes (but full)
PE Size (KByte) 32768
Total PE 29802
Free PE 0
Allocated PE 29802
PV UUID 0uKTxg-SCd9-NA1Z-BYt5-e2v1-fQ72-nrXywf

--- Physical Segments ---
Physical extent 0 to 29801:
Logical volume /dev/VolGroup00/LogVol01
Logical extents 0 to 29801

--- Physical volume ---
PV Name /dev/sdb1
VG Name VolGroup00
PV Size 931.51 GB / not usable 11.19 MB
Allocatable yes
PE Size (KByte) 32768
Total PE 29808
Free PE 2
Allocated PE 29806
PV UUID Mlvnwl-YJ1B-s7Fk-Kz7F-mf5G-1Fvk-BKXIN6

--- Physical Segments ---
Physical extent 0 to 319:
Logical volume /dev/VolGroup00/LogVol00
Logical extents 0 to 319
Physical extent 320 to 29743:
Logical volume /dev/VolGroup00/LogVol01
Logical extents 29802 to 59225
Physical extent 29744 to 29805:
Logical volume /dev/VolGroup00/LogVol02
Logical extents 0 to 61
Physical extent 29806 to 29807:
FREE

```

24. *“a mapping table for maintaining a mapping information for a physical address space corresponding to a logical address space which is a continuous address space equal in size of storage space to an entirety of said logical volume.”* The Accused Products, Systems and/or Services include this element. Metadata holds an extent mapping table. *See, e.g., Logical Volume Manager Administration, Red Hat Enterprise Linux 5, at pp. 19, 115, and 117, reproduced below:*

### 2.3.1. Linear Volumes

A linear volume aggregates space from one or more physical volumes into one logical volume. For example, if you have two 60GB disks, you can create a 120GB logical volume. The physical storage is concatenated.

Creating a linear volume assigns a range of physical extents to an area of a logical volume in order. For example, as shown in [Figure 2.2, "Extent Mapping"](#) logical extents 1 to 99 could map to one physical volume and logical extents 100 to 198 could map to a second physical volume. From the point of view of the application, there is one device that is 198 extents in size.

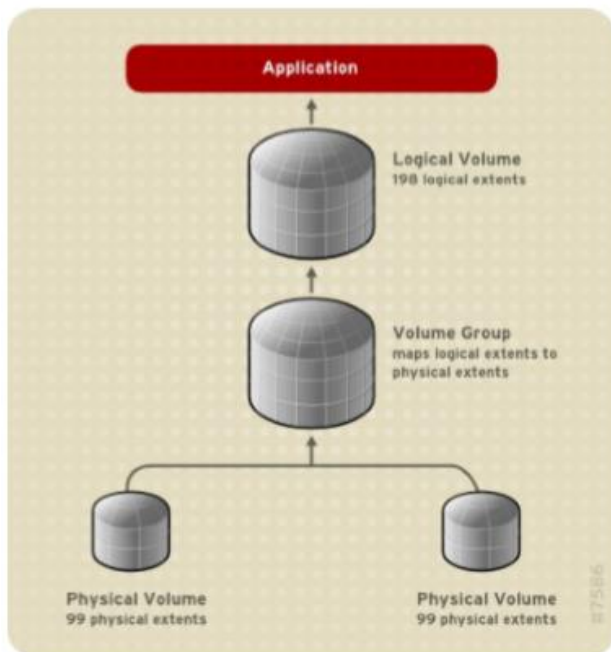


Figure 2.2. Extent Mapping

## D.2. Metadata Contents

The volume group metadata contains:

- Information about how and when it was created
- Information about the volume group:

The volume group information contains:

- Name and unique id
- A version number which is incremented whenever the metadata gets updated
- Any properties: Read/Write? Resizeable?
- Any administrative limit on the number of physical volumes and logical volumes it may contain
- The extent size (in units of sectors which are defined as 512 bytes)
- An unordered list of physical volumes making up the volume group, each with:
  - Its UUID, used to determine the block device containing it
  - Any properties, such as whether the physical volume is allocatable
  - The offset to the start of the first extent within the physical volume (in sectors)
  - The number of extents
- An unordered list of logical volumes, each consisting of
  - An ordered list of logical volume segments. For each segment the metadata includes a mapping applied to an ordered list of physical volume segments or logical volume segments

```

logical_volumes {
    mylv {
        id = "GhUYSF-qVM3-rzQo-a6D2-o0aV-LQet-Ur90F9"
        status = ["READ", "WRITE", "VISIBLE"]
        segment_count = 2

        segment1 {
            start_extent = 0
            extent_count = 1280      # 5 Gigabytes

            type = "striped"
            stripe_count = 1         # linear

            stripes = [
                "pv0", 0
            ]
        }
        segment2 {
            start_extent = 1280
            extent_count = 1280      # 5 Gigabytes

            type = "striped"
            stripe_count = 1         # linear

            stripes = [
                "pv1", 0
            ]
        }
    }
}

```

25. On information and belief, Hitachi had notice and/or knowledge of the '436 Patent since at least May 1, 2007, when Hitachi Ltd. cited the '436 Patent on an Information Disclosure Statement in its own patent application (Application No. 11/104,487) that led to issued U.S. Patent No. 7,213,105. *See Exhibit C.* The '436 Patent has also been cited on a number of other applications that led to issued U.S. Patents assigned to Hitachi. *See, e.g.,* U.S. Patent Nos. 9,626,129; 8,266,375; 8,028,127; 7,930,474; 7,752,390; 7,624,232; 7,334,084; 7,139,875; and 7,447,832.

26. Hitachi has committed acts of infringement without license or authorization. On information and belief, Hitachi acted with objective recklessness by proceeding despite an objective high likelihood that its actions constituted infringement of a valid patent. Hitachi knew or should have known that its actions would cause direct and indirect infringement of the '436 Patent.

27. Hitachi is also liable under 35 U.S.C. § 271(b) for actively inducing infringement and continuing to actively induce infringement. Hitachi actively induces and continues to induce its customers, distributors, end-users, vendors including customer-support and/or manufacturers to infringe the '436 Patent. On information and belief, Hitachi possessed a specific intent to induce infringement, and in fact did induce infringement, by engaging in affirmative acts such as by selling and causing the accused products, systems and/or services to be manufactured, by providing user guides, installation or instruction manuals, and other training materials, by advertising and solicitation and otherwise providing sales-related materials, and by instructing and/or demonstrating to customers, distributors, end-users, vendors including customer-support and/or manufacturers the normal operation of the accused products, systems and/or services that

infringe the '436 Patent. Hitachi is aware and/or willfully blind that these affirmative acts infringe and/or would induce infringement of the '436 Patent, of which it had knowledge.

28. Hitachi is also under 35 U.S.C. § 271(c) for contributing to and continuing to contribute to the infringement of the '436 Patent by, among other things, providing LVM Dynamic Resizing capability in its Accused Products, Systems and/or Services and by encouraging, at a minimum, customers, distributors, end-users, vendors including customer-support and/or manufacturers in this District and elsewhere, to infringe the '436 Patent. By importing, exporting, manufacturing, distributing, selling, and/or providing the accused products, systems and/or services for their intended use to customers, distributors, end-users, vendors including customer-support and/or manufacturers, Hitachi has, in the past and continues to contribute to the infringement of one or more claims of the '436 Patent. The Accused Products, Systems and/or Services, are material to the inventions claimed in the '436 Patent, have no substantial non-infringing uses, and are known by Hitachi (on information and belief) to be especially made or especially adapted for use in infringing the '436 Patent, and which are otherwise not staple articles of commerce suitable for substantial non-infringing use. Hitachi is aware and/or willfully blind that these affirmative acts infringe and/or constitute contributory infringement of the '436 Patent, of which it had knowledge.

29. Hitachi is liable for indirect infringement, i.e. both inducement and contributory infringement, based on the direct infringement that is the result of activities performed by customers, distributors, end-users, vendors including customer-support and/or manufacturers who use all elements or perform all steps of one or more claims of the '436 Patent. For example, end users of Hitachi's Accused Products, Systems and/or Services infringe, either directly or under the doctrine of equivalents, one or more claims of the '436 Patent (*e.g.*, claims 1-3) by

configuring and dynamically resizing logical volumes. At a minimum, Hitachi is liable for the indirect infringement of claims 1-3 of the '436 Patent.

30. As a result of Hitachi's infringement of the '436 Patent, Sequoia has suffered damages and will continue to suffer damages.

31. Hitachi will continue to infringe unless this Court enjoins Hitachi and its agents, servants, employees, representatives and all others acting in active concert with it from infringing the '436 Patent.

32. Sequoia has been damaged as a result of Hitachi's infringing conduct. Hitachi is, thus, liable to Sequoia in an amount that adequately compensates Sequoia for Hitachi's infringement, which, by law, cannot be less than a reasonable royalty, together with interest and costs as fixed by this Court under 35 U.S.C. § 284.

#### **DEMAND FOR JURY TRIAL**

Pursuant to Fed. R. Civ. P. 38(b), Sequoia Technology, LLC hereby demands trial by jury on all issues raised by the Complaint.

#### **PRAYER FOR RELIEF**

WHEREFORE, Sequoia respectfully requests the following relief:

- a) A judgement that Hitachi has infringed the Patent-in-Suit;
- b) An injunction barring Hitachi and its officers, directors, agents, servants, employees, affiliates, attorneys, and all others acting in privity or in concert with them, and their parents, subsidiaries, divisions, successors and assigns, from further acts of infringement of the Patent-in-Suit; alternatively, a judicial decree that Hitachi pay an ongoing royalty in an amount to be determined for continued infringement after the date of judgment;



- c) An award of damages adequate to compensate for Hitachi's infringement of the Patent-in-Suit, and in no event less than a reasonable royalty for Hitachi's acts of infringement, including all pre-judgment and post-judgment interest at the maximum rate permitted by law;
- d) An award of trebled damages under 35 U.S.C. § 284;
- e) A declaration that this case is exceptional under 35 U.S.C. § 285;
- f) An award of Sequoia's costs and attorney's fees under 35 U.S.C. § 285 and other applicable law; and
- g) Any other remedy to which Sequoia may be entitled.

Dated: July 31, 2018

BAYARD, P.A.

OF COUNSEL:

Nate Dilger  
Deepali Brahmhatt  
One LLP  
4000 MacArthur Blvd.,  
East Tower, Suite 500  
Newport Beach, CA 92660  
ndilger@onellp.com  
dbrahmbhatt@onellp.com

/s/ Stephen B. Brauerman  
Stephen B. Brauerman (#4952)  
600 N. King Street, Suite 400  
Wilmington, DE 19801  
(302) 655-5000  
sbrauerman@bayardlaw.com

*Attorney for Plaintiff  
Sequoia Technology LLC*

John Lord  
One LLP  
9301 Wilshire Blvd.  
Penthouse Suite  
Beverly Hills, CA 90210  
jlord@onellp.com