IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

OPTIMUM COMMUNICATIONS SERVICES, INC. Plaintiff, v. COGECO US FINANCE, LLC D/B/A BREEZELINE; and COGECO COMMUNICATIONS INC., Defendants.	)	C.A. No. __________ **TRIAL BY JURY DEMANDED**

COMPLAINT FOR PATENT INFRINGEMENT

1. Plaintiff Optimum Communications Services, Inc. ("Plaintiff" or "OCS") files this Complaint for patent infringement and demand for jury trial against Defendants Cogeco US Finance, LLC d/b/a Breezeline, and Cogeco Communications Inc. (collectively, "Cogeco" or "Defendants"), and alleges as follows:

NATURE OF THE ACTION

2. This is a civil action for patent infringement arising under the patent laws of the United States, 35 U.S.C. § 1, *et seq*., including 35 U.S.C. § 271, which gives rise to the remedies specified under 35 U.S.C. §§ 281 and 283–285.

THE PARTIES

3. Plaintiff OCS is a corporation organized under the laws of the State of Delaware, with its principal place of business at 344 Grove Street #242, Jersey City, New Jersey 07302. OCS and its predecessors and affiliates will individually and collectively be referred to as "OCS" herein.

4. On information and belief, Cogeco US Finance, LLC d/b/a Breezeline is a subsidiary of Cogeco Communications Inc., and is a company organized under the laws of the

State of Delaware, with its principal place of business at 3 Batterymarch Park, Suite 200, Quincy, Massachusetts 02169.

5. On information and belief, Cogeco Communications Inc., is a corporation organized under the laws of Canada, with its principal place of business at 1 Place Ville-Marie, Suite 3301, Montréal, Québec H3B 3N2.

JURISDICTION AND VENUE

6. This is a complaint including causes of action for patent infringement arising under 35 U.S.C. § 271, *et seq.* This Court has subject matter jurisdiction under 28 U.S.C. § 1331 (federal question jurisdiction), and 28 U.S.C. § 1338(a) (jurisdiction over patent actions).

7. This Court has personal jurisdiction over Defendants at least because each Defendant has purposefully availed itself of the privilege of conducting business within Delaware. Cogeco Communications Inc. has conducted business in Delaware through its subsidiary, Cogeco US Finance, LLC. Personal jurisdiction is proper under Delaware's Long-Arm Statute, 10 Del. C. § 3104, including because each Defendant has sufficient minimum contacts with this forum such that the maintenance of this suit does not offend traditional notions of fair play and substantial justice.

8. Defendant Cogeco Communications Inc. directs and controls Cogeco US Finance, LLC's provision of fiber optic broadband and other services within the United States, including in Delaware.

9. Venue is proper in this judicial district pursuant to 28 U.S.C. § 1400(b) as to Cogeco US Finance, LLC at least because it is incorporated in Delaware. Venue is proper as to Cogeco Communications Inc. because, as a non-resident of the United States, it may be sued in any judicial

district, 28 U.S.C. § 1391(c)(3), and because it has committed acts of infringement in this district and has purposefully directed its business here.

THE ASSERTED PATENTS

10. Mark Sandstrom is the sole and true inventor of U.S. Patent No. 7,333,511 (the "'511 Patent"), U.S. Patent No. 7,558,260 (the "'260 Patent"), U.S. Patent No. 10,567,474 (the "'474 Patent"), U.S. Patent No. 10,848,546 (the "'546 Patent") (collectively, "the Asserted Patents"), which generally relate to two fundamental innovations in communications network systems. The first innovation dynamically optimizes network capacity allocation to maximize data throughput. Prior to these inventions, traditional networks relied on fixed physical layer connections with constant bit rates, which can be inefficient for bursty, unpredictable packet-based traffic like the Internet. The patented technology addresses these limitations by, in certain non-limiting examples, enabling the use of dynamically channelizable multi-source-node data transport buses. These systems can utilize distributed algorithms at network nodes to monitor data volumes in real-time and dynamically adjust capacity allocation among various source nodes based on their current traffic demands.

11. The second innovation relates to a telecommunications network management system based on file transfer routines and actions based on the file contents. In one exemplary embodiment, instead of relying on intermediate transaction languages and data format conversions, this system transfers network management data (NMD) files directly between a Network Management System (NMS) server and remote Network Elements (NEs). In this embodiment, the remote NEs continuously exchange files with the server, automatically executing the operations defined by the configuration files, and subsequently returning updated status files to the NMS server. This decoupled, file-based approach can eliminate the processing overhead and bottlenecks

associated with conventional management systems, resulting in more scalable, transparent, and reliable network operations.

12. Mr. Sandstrom's innovations allowed networks to re-allocate capacity thousands of times per second, solving the efficiency gap between fixed physical-layer connections and "bursty" internet traffic, and to achieve enormous efficiencies in communications network capacity usage.

13. The '511 Patent, entitled "Dynamically Channelizable Packet Transport Network," was duly and legally issued by the United States Patent and Trademark Office ("USPTO") on February 19, 2008. The USPTO issued Certificates of Correction on June 10, 2008, correcting claim 24, and on February 5, 2013, correcting claim 29. OCS is the owner of the '511 Patent, with all substantial rights, including the exclusive right to enforce, sue, and recover damages for past and future infringements. A copy of the '511 Patent, together with its Certificates of Correction, is attached as Exhibit 1.

14. The '260 Patent, entitled "Byte-Timeslot-Synchronous, Dynamically Switched Multi-Source-Node Data Transport Bus System," was duly and legally issued on July 7, 2009. The USPTO issued a Certificate of Correction on September 22, 2009, correcting claim 11. OCS is the owner of the '260 Patent, with all substantial rights, including the exclusive right to enforce, sue, and recover damages for past and future infringements. A copy of the '260 Patent, together with its Certificate of Correction, is attached as Exhibit 2.

15. The '474 Patent, entitled "Direct Binary File Transfer Based Network Management System Free of Messaging, Commands and Data Format Conversions," was duly and legally issued on February 18, 2020. The USPTO issued a Certificate of Correction on December 27, 2022, correcting claims 1 and 5. OCS is the owner of the '474 Patent, with all substantial rights, including the exclusive right to enforce, sue, and recover damages for past and

future infringements. A copy of the '474 Patent, together with its Certificate of Correction, is attached as Exhibit 3.

16. The '546 Patent, entitled "Direct Binary File Transfer Based Network Management System Free of Messaging, Commands and Data Format Conversions," was duly and legally issued on November 24, 2020. OCS is the owner of the '546 Patent, with all substantial rights, including the exclusive right to enforce, sue, and recover damages for past and future infringements. A copy of the '546 Patent is attached as Exhibit 4.

17. OCS owns all rights, titles, and interest in the Asserted Patents.

OCS'S PIONEERING TECHNOLOGY

18. The '511 Patent – the earliest filed of the four Asserted Patents – explains that to provide "deterministic QoS for packet switched data traffic," providers traditionally relied on "dedicated Layer 1 or Layer 0 point-to-point" connections between sites. Exhibit 1 at 2:9–15. However, this "network architecture based on non-shareable connections is inefficient for bursty packet traffic, since the bandwidth reserved for the dedicated point-to-point connections is wasted when not used." *Id.* at 2:16–19. In such environments, "none of the point-to-point link bandwidth can be used to carry other potential billable traffic even when not needed by traffic between the two sites connected by the point-to-point link." *Id.* at 2:20–24. This creates a fundamental "trade-off between performance and cost advantage," as conventional networks either waste physical capacity to ensure QoS or utilize shared packet-switched resources that "are not able to provide deterministic QoS without compromising the bandwidth efficiency of packet-switching." *Id.* at 3:55–67.

19. The technical problem solved by the '511 Patent, in an example, can be illustrated by a "group of five POP-routers" that each require "10 Gbps of aggregate full-duplex traffic

forwarding capacity." *Id.* at 2:28–34. In a traditional "dedicated link-based network architecture," a "total of 20x10 Gbps=200 Gbps of physical network capacity is needed to support" only "50 Gbps of service interface capacity." *Id.* at 3:2–6. Alternatively, statically subdividing links into smaller, non-adaptive connections would "be blocking up to 10 Gbps-2.5Gbps=7.5 Gbps worth of billable traffic per POP, thus wasting up to 75% of the traffic forwarding capacity of the POP routers." *Id.* at 3:10–14.

20. The exemplary, non-limiting system described in the '511 Patent solves this problem by providing "real-time dynamic, traffic load adaptive allocation of transport network capacity to continuously maximize the data throughput of the transport network for dynamic packet traffic." *Id.*, Abstract. The exemplary system enables "multiple independent, i.e. private-like, site-to-site mesh subnetworks. . . over a common shared network infrastructure" that "effectively combines the QoS of dedicated L1 point-to-point circuits with the network resource usage efficiency of L2/L3 packet-switching." *Id.* at 4:18–29.

21. With reference to Figure 2, reproduced below, the non-limiting exemplary network system comprises a "dynamically L1-channelizable, logical packet transport bus 9" that "transports data packets, such as MPLS, IP, Frame Relay or Ethernet MAC packets, from its source nodes 4(*a*), 4(*b*), 4(*c*) and 4(*d*) to its destination node 4(*e*)." *Id.* at 5:25–26, 6:5–10. This bus (9) provides a "dedicated, packet-layer transparent, variable-bandwidth channel 6 . . . from each individual one of its source nodes 11 to its destination node 10." *Id.* at 6:18–21. These "channels 6 on the bus 9 are L1 connections" and involve "no packet-layer functionality, i.e. L2 or upper layer protocol functionality" between the source and destination nodes. *Id.* at 6:21–25.

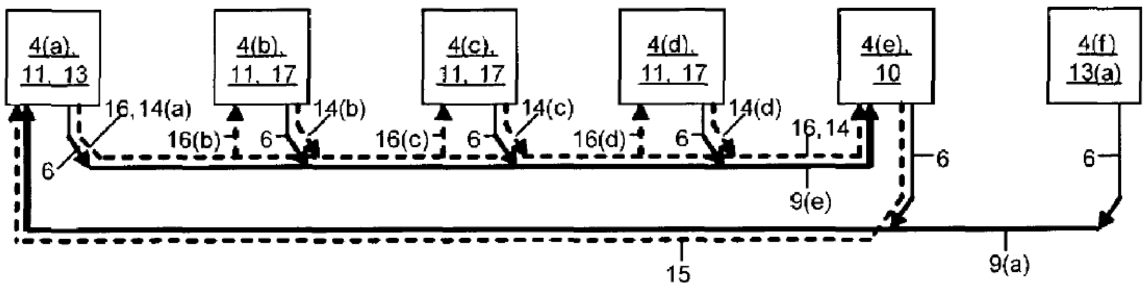

22. In this non-limiting example, the capacity of these source-node-specific channels (6) is "adjusted by a control process controlling the bus capacity allocation based on the traffic loads from the source nodes 11 toward the destination node 10." *Id.* at 6:27–29. In that studied example, each channel (6) uses "concatenation of the capacity allocation units to form a single logical variable-bandwidth packet transport channel." *Id.* at 6:57–67. The system recomputes concatenation coefficients for each "new process cycle," which in the studied example is "one millisecond i.e. eight SDH/SONET frame periods in duration," to "maximize the data throughput of the bus." *Id.* at 7:6–16.

23. Also, in the non-limiting example, a "bus control signaling scheme" is executed every cycle. *Id.* at 8:31–32, 53–55. As illustrated in Figure 2, above, "capacity allocation info 15 is transferred from the destination node 4(*e*). . . to its end-of-bus (EOB) 13 source node 4(*a*)." *Id.* at 8:34–36. The EOB node then "loops-back i.e. transmits. . . the capacity allocation info on the bus 9(*e*)" as a second-pass flow (16). *Id.* at 8:51–55. Downstream source nodes copy their assigned allocation, and "insert [their] capacity demand info 14 into an appropriate field." *Id.* at 9:5–10.

24. In an exemplary system, individual source nodes (11) determine their "capacity demand FIG. 14 by translating [their] amount of data bytes queued. . . into such an amount of … slots. . . that provides sufficient bandwidth in the bus process cycle (e.g. 1 ms) to transfer all the queued data…." *Id.* at 11:53–62. This "process of quantifying a demand FIG. 14 in capacity allocation units" enables source nodes to request as much bus capacity as they have received data

bytes to map on that bus. *Id.* at 11:64–65. The destination node (10) then "copies from the control payload the capacity demand info 14 inserted therein by the source nodes" to compute the next allocation. *Id.* at 9:49–55.

25. In the described non-limiting examples, the allocation algorithm distributes the "bus capacity for the next bus process cycle. . . so as to minimize the aggregate amount of unmatched demand 15 by the sources 11 of the bus 9, and to ensure fairness." *Id.* at 9:55–60. It assigns units "one bus capacity allocation unit at a time per each source node whose bus capacity demand FIG. 14. . . is greater than the amount of capacity allocation units so far allocated to it, until all the capacity allocation units are allocated, or until the capacity demand figures by all the source nodes 11 of the bus 9 have been matched, whichever comes first", and if, "there still remain unallocated bus capacity allocation units after the capacity demand 14 by each source node of the bus has been matched, the bus capacity allocation algorithm continues to allocate the remaining capacity allocation units among the source nodes of the bus according to its regular rotating order until all the capacity allocation units, i.e. the aggregate bus capacity for the next bus process cycle, have been allocated." *Id.* at 16:29–46. This "achieves the below two goals that improve the network data throughput and QoS: i) it minimizes the aggregate amount of unmatched demand 14 by the source nodes 11 of the bus 9 left after the allocation of the bus capacity, and ii) it ensure[s] fairness in capacity allocation among the source nodes 11 of the bus 9." *Id.* at 16:47–54.

26. The '260 Patent describes a non-limiting exemplary "self-optimizing digital communications network for transporting data among a set of nodes" based on "multi-source-node data transport channels" where "frame-slots are dynamically assigned among the individual source nodes. . . using channel access control signaling in signal frame overhead" as depicted in part in Figure 2, which is reproduced below. Exhibit 2, Abstract. The system "optimizes its capacity

allocation pattern according to the real-time data traffic load patterns among the nodes across the network" for "maximization of the effective network data throughput." *Id.* at 3:10–14, 4:1–3.

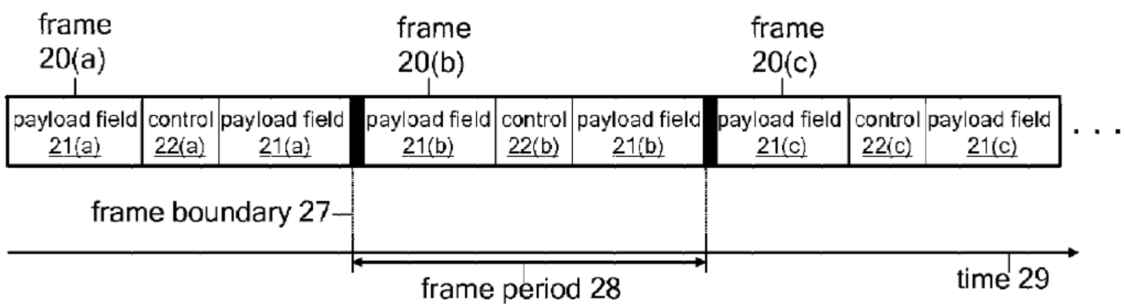

FIG. 2

27. In a non-limiting example, the solution includes a "specific channel-access-control signaling bit-fields within the overhead timeslots of signal frames" to "identify which individual source node of the channel is to transmit data on the channel on the following frame period." *Id.* at 3:28–32. Each "individual source node is identified with a unique channel access number called Local Node Identifier [LNI]." *Id.*, Abstract. A control information field, the "Active Node Identifier [ANI] carried in the overhead of each frame," selects the active node for the "next frame period" on the channel signal. *Id.* at 6:35–36. This "dynamic frame-slot switching method" is particularly efficient when channels are grouped into a "multi-channel data transport bus, on which the allocation of member channels. . . is optimized for every new bus frame period." *Id.* at 3:32–40.

28. In the exemplary system, to prevent data loss, "nodes along such channels access the channel synchronously using the same frame-timeslot phase." *Id.*, Abstract. This ensures that "the channel can have a different active source node even for every new frame period, without any channel down-time." *Id.* In a preferred embodiment, member channels are "time-slot-interleaved within their bus envelope." *Id.* at 5:3–5. This allows multiple "geographically distant and independently clocked source nodes" to "map data on the same communication channel" on every new frame-slot. *Id.* at 11:7–10.

29. In a non-limiting example described in the '260 Patent, the "bus capacity allocation . . . is computed for the bus by a bus access controller [BAC]" based on "bus capacity request [BCR] information gathered from the individual channel source nodes." *Id.* at 3:39–43. The BCR figure is "an integer number requested channels" determined by a source node based on its "current amount of data queued on a buffer for a future delivery." *Id.* at 11:51–53, 10:24–26. The BAC then "allocates and assigns the bus channels 8 among the set of bus source nodes 11 once every bus frame period 28." *Id.* at 11:52–55. The allocation algorithm assigns channels "one channel at a time, per each source whose BCR figure exceeds the number of channels so far assigned to it" until the bus is full or all requests are met. *Id.* at 11:58–64.

30. In the described non-limiting example, by utilizing existing overhead timeslots, the invention enables "dynamic optimization of network capacity allocation without adding extra overhead." *Id.* at 3:54–63. For standard 125-μs frame periods, the system optimizes the allocation "8000 times per second." *Id.* at 3:46–50. Through modifications to the "POH byte on each of the nine VC-3 rows of a VC-3 frame," it is possible to "identify a new active source node. . . for every new VC-3 row," allowing for "optimization of bus capacity allocation at frequency of 9/0.125 ms = 72,000 times per second." *Id.* at 9:40–50.

31. The ’546 Patent is a continuation of the application for the ’474 Patent, and they share the same disclosure. Exhibit 4, (63). The ’474 and ’546 Patents are directed to a network management system (NMS) for managing remote network elements (NEs) based on network management data (NMD) file transfer routine, whereby “[a] system user produces configuration files at the NMS server for NEs using a network management GUI, and the hardware of NEs automatically complete the network management operations indicated by the NMD files transferred to them from the NMS server and produce their status files to the NMS server. The network management GUI displays network status based on the latest NE status files at the NMS server.” Exhibit 3, Abstract. Conventional NMS techniques relied on “command and messaging based communications for transactions and information distribution,” which often leads to “exhaustive implementational complexity.” *Id.* at 1:50–54, 2:53–54. These systems are “prone to become overloaded during times of high loads” because their performance “degrades when the NMS capabilities are most urgently needed, e.g., during bursts of messaging . . . major network failures.” *Id.* at 2:17–20, 25–29. Furthermore, standard management transactions require “several stages of protocol, message, language and data format conversions,” which reduces “NMS transparency from NMS user IF [interface] to the network element HW [hardware].” *Id.* at 2:14–17, 35–37. This lack of transparency forced “network operator’s personnel to know about or deal with the peculiarities of any intervening messaging protocol.”. *Id.* at 6:12–15.

32. By utilizing “automatic routines for transferring binary NMD files,” the system “avoids the complexity and restrictions of intermediate messaging protocols or transaction languages”. *Id.*, Abstract. In this architecture, “network management actions occur as automatic consequences of the contents of the NMD files” being stored in the local memories of network elements. *Id.* at 4:12–14. This “direct binary file transfer” approach ensures that “[h]eavy load of

NMS or network event activity on one element of the system... does not negatively interfere with other elements," resulting in a system that is "highly reliable and scalable." *Id.* at 5:45–62.

DEFENDANTS' INFRINGEMENT OF THE ASSERTED PATENTS

33. Defendants have been and are presently infringing, and will continue to infringe, the Asserted Patents by, among other things, using the GPON, XGS-PON, NETCONF and YANG standards, and selling and offering for sale internet services that employ the GPON, XGS-PON, NETCONF and YANG standards.

34. Defendants directly infringe the Asserted Patents pursuant to 35 U.S.C. § 271(a), either literally or under the doctrine of equivalents, or both.

DEFENDANTS' INFRINGING METHODS

35. Defendants are internet service providers (ISPs) that provide fiber optic internet services to consumers using Passive Optical Networks (PONs). Exhibit 6. The PON architecture includes an optical line terminal (OLT) located at an ISP central office for internet communications over a fiber optic network to provide internet connectivity to a set of optical network units (ONUs) that are located near end-users as shown in the figure below. Exhibit 10 at 21, Fig. 7-5; Exhibit 23.

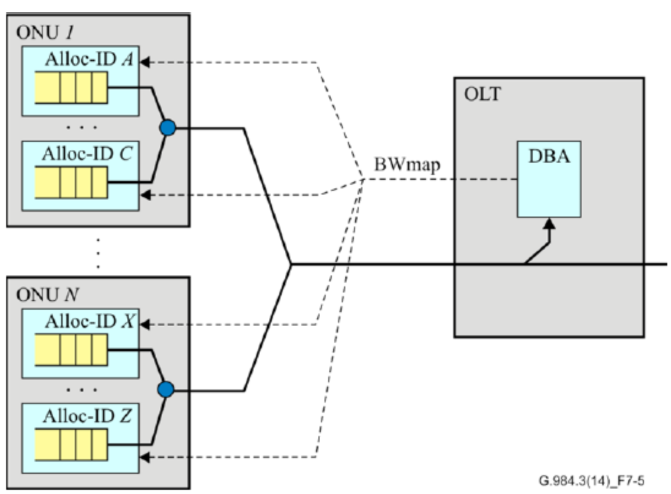

In some PON configurations, including those used by Defendants, the OLTs are controlled using the NETCONF protocol for OLT configuration and network management. Exhibit 7 at 1. Each

ONU is connected to one or more optical network terminals (ONT), with each ONT associated with an end-user and typically installed on-site in homes and businesses. Exhibit 8.

36. On information and belief, Defendants provide fiber optic-based internet services using switches that are compliant with the GPON and/or XGS-PON standards, and they use software-defined networking (SDN) controllers, OLTs, and switches that are compliant with NETCONF and YANG standards. For example, Defendants use SDN controllers and OLTs provided by telecommunications equipment manufacturer Adtran, Inc. ("Adtran") that are compliant with the GPON/XGS-PON standards, as well as switches provided by Adtran that are compliant with the NETCONF, YANG, and GPON/XGS-PON standards. Exhibits 5, 6. The Adtran SDX 6320 OLT is a representative example of an Adtran OLT. Exhibit 7. According to Adtran, the "SDX 6000 Series of software-defined access (SD-Access) OLTs consists of open and disaggregated access devices that support a broad range of PON standards, including 10G Combo PON, XGS-PON, and GPON." *Id.* at 1. The Adtran SDX 631x ONT is a representative example of an Adtran ONT, and is typically paired with the Adtran SDX 6320 OLT. Exhibit 8.

37. GPON is a point-to-multipoint telecommunications standard defined by ITU-T recommendation G.984.3. *See* Exhibit 10 at i (Summary). It utilizes single-mode optical fiber to distribute data between a provider's OLT and multiple end-user ONTs. *Id.*, § 6.3.3.

38. The Adtran OLTs and ONTs used by Defendants are configured to perform Dynamic Bandwidth Assignment according to either the GPON standard, G.984.3 or the XGS-PON standard. Exhibits 10, 11. Under the Traffic Management section of the OLT product brochure, "Dynamic Bandwidth Allocation" is specifically identified. Exhibit 7. This refers to the same DBA techniques described at sections 7.4 and A.9.3 of the GPON and XGS-PON standards, respectively. Exhibits 10, 11. Dynamic Bandwidth Assignment ("DBA," which is the same as

Dynamic Bandwidth Allocation) in gigabit-capable passive optical networks (G-PONs) is the process by which the OLTs reallocate the upstream transmission opportunities to the traffic-bearing entities within ONUs based on the dynamic indication of their activity status and their configured traffic contracts. Exhibit 10, § 7.4. "In comparison with static bandwidth assignment, the DBA mechanism improves the G-PON upstream bandwidth utilization by reacting adaptively to the ONUs['] bursty traffic patterns." *Id.* The practical benefits of DBA are twofold. "First, the network operators can add more subscribers to the PON due to more efficient bandwidth use. Second, the subscribers can enjoy enhanced services, such as those requiring variable rate with the peaks extending beyond the levels that can reasonably be allocated in a static fashion." *Id.*

39. DBA is a mandatory feature of the GPON and XGS-PON standards, meaning that any equipment claiming to be GPON or XGS-PON compliant must include the specified DBA functionality. The GPON standard specifies that "The G-PON OLT is required to support DBA" and that "Dynamic bandwidth assignment in G-PON encompasses the following functions. These functions shall apply on the level of individual Alloc-IDs that are associated with T-CONTs and provisioned bandwidth component parameters." *Id.*, §§ 7.4, 7.4.2. The XGS-PON standard specifies that "[t]he XGS-PON OLT shall support DBA for the efficient sharing of upstream bandwidth among the connected ONUs and the traffic-bearing entities within the individual ONUs based on the dynamic indication of their activity." Exhibit 11, § A.9.3.

40. About 70% of all fiber networks in the United States are GPON compliant. Exhibit 9 at 1. This figure includes XGS-PON compliant equipment. XGS-PON is the more modern (and higher speed) successor to the GPON standard. Exhibit 11 at 1.

41. NETCONF is a standard defined primarily in RFC 6241. Exhibit 12. Use of the NETCONF standard obviates manual provisioning or configuring of switches or routers

(collectively referred to hereafter as "switches"). *Id.*, § 1. Instead of a human logging into each switch in a network and typing commands, a central manager sends a Remote Procedure Call (RPC) including an XML file containing the desired configuration to one or more switches in the network. *Id.* As discussed below, the receiving switches read the file, apply the settings, and confirm it is done and/or the central manager sends another RPC to do the same.

42. YANG is a standard defined in RFC 7950. Exhibit 13. YANG is a data layer companion for NETCONF. *Id.* YANG is the modeling language used to describe the data. *Id.* It provides a standardized data structure that can be used universally with all switches in a network—whether provided by Adtran or another vendor—by ensuring they all use the common standard definitions for configuration and state data. As noted in the NETCONF standard, "[t]he YANG data modeling language [RFC6020] has been developed for specifying NETCONF data models and protocol operations," covering the Operations and the Content layers of the NETCONF stack. Exhibit 12, § 1.2.

43. The NETCONF/YANG approach has been widely adopted by ISPs for several reasons. It replaces manual, error-prone command line interface (CLI) provisioning with an automated process that is less prone to human error. *Id.*, § 1. The use of YANG data models enables Secure Zero Touch Provisioning (SZTP), allowing the network to automatically detect and provision ONUs as soon as they are connected to the fiber. Exhibit 17. This approach facilitates a more open, multi-vendor ecosystem. Exhibit 12, § 1; Exhibit 19. By using a centralized SDN controller that communicates via NETCONF, ISPs can manage hardware from different vendors—including Adtran, Huawei, ZTE and white label sources—using standardized interfaces. *Id.* This reduces vendor lock-in, as a single high-level orchestration script can be used to manage diverse OLT environments simultaneously. *Id.*

44. The Adtran Mosaic Cloud Platform (Mosaic CP) is NETCONF/YANG compliant. Mosaic CP utilizes NETCONF as a core management protocol. Exhibit 14 at 3. Mosaic CP includes a NETCONF interface, supporting thousands of concurrent sessions to ensure large-scale network scalability and simplified programmability. *Id.* The NETCONF standard indicates that "[t]he YANG data modeling language [RFC6020] has been developed for specifying NETCONF data models and protocol operations," covering the Operations and the Content layers of the NETCONF stack. Exhibit 12, § 1.2.

45. Most regional ISPs, including Defendants, either have transitioned or are currently transitioning from GPON networks to higher speed XGS-PON networks, which in particular are almost exclusively managed by NETCONF/YANG for the reasons mentioned above. Exhibits 20, 21.

COUNT ONE
DIRECT INFRINGEMENT OF U.S. PATENT NO. 7,333,511

46. OCS incorporates by reference the allegations in the preceding paragraphs as if fully set forth herein.

47. On information and belief, Defendants have infringed claims of the '511 Patent, including at least claim 29, in violation of 35 U.S.C. § 271(a) by using the claimed process.

48. Claim 29 of the '511 Patent (Exhibit 1) recites:

> 29. A process for optimizing capacity allocation within a data transport network system referred to as a network system, the network system comprising a set of ingress interfaces, an egress interface, and a capacity pool for transporting data packets from the set of ingress interfaces to the egress interface of the network system, the process having a repeating process cycle, each process cycle comprising the steps:
>
> optimizing, by the egress interface, allocation of said capacity pool among the individual ingress interfaces for transport of data packets from the ingress interfaces to the egress interface based at least in part on demand for

> network capacity by the individual ingress interfaces for transporting data from them to the egress interface of the network system,
>
> assigning, by the egress interface, units of capacity within said capacity pool to said individual ingress interfaces according to the optimizing of capacity allocation, and
>
> transporting data packets from the ingress interfaces to the egress interface based on the assigning of units of capacity within said capacity pool.

49. On information and belief, Defendants have directly infringed at least claim 29 of the ’511 Patent by using GPON and XGS-PON compliant switches such as those available from Adtran.

50. Defendants use GPON/XGS-PON compliant Adtran fiber equipment, as shown in Exhibits 5 and 6. Thus, at least a portion of Defendants’ fiber networks is GPON and/or XGS-PON compliant.

51. Adtran equipment is used in Defendants’ fiber networks, including Adtran optical network terminals (ONTs, the equipment or box often installed in customers’ premises) and optical line terminals (OLTs, a distribution switch that connects to many, sometimes thousands, of ONTs in individual homes). Exhibits 5, 6.

52. The Adtran SDX 6320 OLT is a representative example of an Adtran OLT. Exhibit 7. It is typically paired with the Adtran SDX 631x ONT. Exhibit 8. On information and belief, Defendants infringe the ’511 Patent by practicing the GPON and/or XGS-PON standards using the Adtran SDX 6320 OLT and Adtran SDX 631x ONT.

53. Defendants employ **“[a] process for optimizing capacity allocation within a data transport network system referred to as a network system, the network system comprising a set of ingress interfaces, an egress interface, and a capacity pool for transporting data packets from the set of ingress interfaces to the egress interface of the network system.”**

Defendants employ the process recited in claim 29 of the ’511 Patent by using GPON/XGS-PON compliant Adtran equipment or equivalent equipment from Adtran or other vendor(s). The G-PON and XGS-PON standards define requirements for transmission of data on optical fiber networks, which are **data transport network system[s]**. Exhibit 10, § 1; Exhibit 11, § 1. As shown in the GPON standard’s Figure 7-5 (reproduced below) and the XGS-PON standard’s Figure C.7.1, the described network consists of two or more ONUs, each of which may be an ONT device that is specific to a given subscriber or set of subscribers. Exhibit 10, §§ 3.2.15–17; Exhibit 11, §§ 3.1.26–28. An OLT device bridges the fiber connection provided by the service provider and the connections to the individual ONUs. *Id.* As shown in the aforementioned figures of the standards, the OLT includes a DBA process, and the OLT is connected to the ONUs through fiber optics. Exhibit 10, § 7.4, Exhibit 11, § C.7.2.1. The DBA is a method of **optimizing capacity allocation**. *Id.* Traffic from the subscriber to the fiber network moves from the left, through the Alloc-ID (allocation identifier) buffers, through the OLT and then on to the fiber network. Exhibit 10, §§ 3.2.12, 15–17; Exhibit 11, §§ 3.1.26–28. The OLT corresponds to the recited **egress interface**, while the ONUs or any of their subscriber-specific Alloc-ID buffers correspond to the two or more individual **ingress interfaces**.

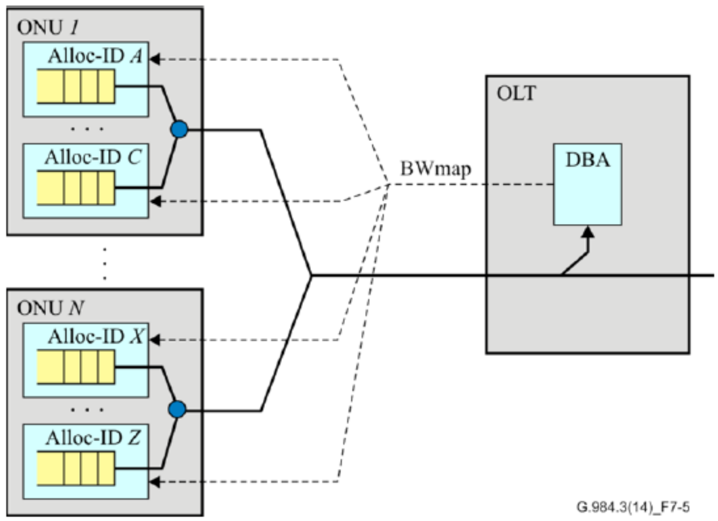

Exhibit 10, § 7.4, Exhibit 11, § C.7.2.1. As noted above, the DBA process allocates **a capacity pool** among the ONUs by circulating a "BWmap" (bandwidth map). *See* Exhibit 10, § 7.4, Fig. 7-5; Exhibit 11, § C.7.2.1, Fig. C.7.1. The bandwidth is allocated to a transmission of **data packets** from the ingress interfaces to the egress interfaces. *Id.*

54. Defendants employ a **"process having a repeating process cycle, each process cycle comprising the steps: optimizing, by the egress interface, allocation of said capacity pool among the individual ingress interfaces for transport of data packets from the ingress interfaces to the egress interface, based at least in part on demand for network capacity by the individual ingress interfaces for transporting data from them to the egress interface of the network system."** In GPON, the G-PON Transmission Convergence (GTC) frame period (125 µs) corresponds to the repeating process cycle. Exhibit 10, § 5.6. In XGS-PON, the same 125-µs frame is called the XGS-PON Transmission Convergence (XGS TC) frame. Exhibit 11, §§ 3.2.26 and C.6.1.2. The OLT controls the upstream capacity sharing among the ONUs and their Alloc-IDs by creating BWmaps for each new upstream GTC/XGS TC period. *Id.*; *see also* Exhibit 10, §§ 7.2, 7.4 and C.7.2.1. The DBA process in the OLT (**egress interface**) allocates bandwidth (portions of the **capacity pool**) for the ONUs **(ingress interfaces**) using the BWmap.

Exhibit 10, § 7.4; Exhibit 11, § C.7.2.1. The GPON and XGS-PON standards describe two methods for the DBA process to infer the Alloc-ID buffer occupancy levels to assess their respective demand levels for upstream capacity. Exhibit 10, § 7.4.3; Exhibit 11, § C.7.2.1. The first method is the Status Reporting (SR) DBA, where explicit Alloc-ID buffer occupancy reports are solicited by the OLT and submitted to the OLT by the ONUs (**based at least in part on demand for network capacity by the individual ingress interfaces for transporting data from them to the egress interface of the network system**, **the process having a repeating process cycle**). *Id.* The second method is the Traffic Monitoring (TM) DBA which is based on the OLT's observation of the idle G-PON Encapsulation Method (GEM) frames transmitted by a given ONU/Alloc-ID in place of upstream transmission opportunities granted to it (via BWmap) (also **based at least in part on demand for network capacity by the individual ingress interfaces for transporting data from them to the egress interface of the network system**). *Id.* Both SR and TM are alternatives for the DBA for its inference of the Alloc-ID buffer occupancy levels, and the rest of the DBA process (including BWmap creation and application) operates the same way based on either alternative for inference of Alloc-ID buffer load levels.

55. Defendants employ a process including **"assigning, by the egress interface, units of capacity within said capacity pool to said individual ingress interfaces according to the optimizing of capacity allocation."** The OLT (**egress interface**) via the DBA sends (**assign[s]**) the BWmap (**optimizing of capacity allocation**) to the ONUs (**ingress interfaces**). Exhibit 10, § 7.2.1; Exhibit 11, § C.7.2.1. The BWmap reflects **assign[ed] units of capacity** that are delimited by the start/end byte pointers (or a start byte pointer and a grant size field in the case of XGS-PON) in the BWmap for each Alloc-ID of each ONU. *Id.*; *see also* Exhibit 11, § C.6.1.6.

56. Defendants employ a process including **"transporting data packets from the ingress interfaces to the egress interface based on the assigning of units of capacity within said capacity pool."** The **data packets** are **transport[ed]** from Alloc-ID buffers from the ONUs (**ingress interfaces**) to the OLT (**egress interface**) via upstream GTC/XGS TC frames according to the BWmaps.

57. Defendants have caused Plaintiff damage by direct infringement of the '511 Patent.

58. In accordance with 35 U.S.C. § 287, Defendants have had actual notice and knowledge of the '511 Patent by no later than the filing of this Complaint.

59. On information and belief, Defendants continue, without license, to practice methods that infringe the '511 Patent.

60. Defendants' infringement of the '511 Patent has injured OCS in its business and property rights. OCS is entitled to recover monetary damages for the injuries arising from Defendants' infringement in an amount to be determined at trial.

61. The full extent of Defendants' infringement is not presently known to OCS. OCS makes this preliminary identification of infringing methods and infringed claims in this count without the benefit of discovery in this action, and expressly reserves the right to augment, supplement, and revise its identification based on additional information obtained through discovery or otherwise.

COUNT TWO
DIRECT INFRINGEMENT OF U.S. PATENT NO. 7,558,260

62. OCS incorporates by reference the allegations in the preceding paragraphs as if fully set forth herein.

63. On information and belief, Defendants have infringed claims of the '260 Patent, including at least claim 12, in violation of 35 U.S.C. § 271(a) by using the claimed process.

64. Claim 12 of the '260 Patent (Exhibit 2) recites:

> 12. A control process for dynamically sharing a digital communications channel bandwidth among multiple source nodes, said channel being able to transport data from multiple source nodes to a channel destination node and having a carrier signal consisting of a continuous sequence of signal frames, said process comprising:
>
> dynamically allocating the channel bandwidth among its multiple source nodes, by assigning channel signal frame transmission slots among the source nodes, through activating, via signal frame overhead information fields, individual source nodes of the channel to transmit data on the channel on future signal frames associated with said signal frame overhead information fields, and
>
> transmitting data on the channel by its individual source nodes on exactly those signal frames that were assigned to each individual source node via signal frame overhead information fields associated with said signal frames,
>
> wherein the process element of transmitting involves a sub-process performed by at least one of the set of channel source nodes that includes the below steps:
>
> I) monitoring specified overhead bitfields, referred to as Access Control Tags (ACTs), in each signal frame on the channel, and
>
> II) based at least in part on the step I), either transmitting or not transmitting data on the channel on the signal frame associated with a given ACT.

65. On information and belief, Defendants have directly infringed at least claim 12 of the '260 Patent by using GPON compliant switches such as those available from Adtran.

66. Defendants use GPON/XGS-PON compliant Adtran fiber equipment, as shown in Exhibits 5 and 6. Thus, at least a portion of Defendants' fiber networks is GPON and/or XGS-PON compliant.

67. Adtran equipment is used in Defendants' fiber networks, including Adtran optical network terminals (ONTs, the equipment or box often installed in customers' premises) and optical line terminals (OLTs, a distribution switch that connects to many, sometimes thousands, of ONTs). Exhibits 5, 6.

68. The Adtran SDX 6320 OLT is a representative example of an Adtran OLT. Exhibit 7. It is typically paired with the Adtran SDX 631x ONT. Exhibit 8. On information and belief, Defendants infringe the '260 Patent by practicing the GPON and/or XGS-PON standards using the Adtran SDX 6320 OLT and Adtran SDX 631x ONT or equivalent equipment.

69. Defendants employ **"[a] control process for dynamically sharing a digital communications channel bandwidth among multiple source nodes, said channel being able to transport data from multiple source nodes to a channel destination node and having a carrier signal consisting of a continuous sequence of signal frames."** The G-PON, and XGS-PON standards define requirements for transmission of data on optical fiber networks, which, in the upstream direction, is a network system with multiple source nodes and which uses an optical carrier signal. Exhibit 10, § 1; Exhibit 11, § 1. As shown in the GPON standard's Figure 7-5, and the XGS-PON standard's Figure C.7.1, the described network consists of two or more ONUs each of which may be an ONT device that is specific to a given subscriber or set of subscribers. Exhibit 10, §§ 3.2.15–17, Fig. 7-5; Exhibit 11 at §§ C.7.2.1, 5.10. As shown in the aforementioned figures of the standards, the OLT includes DBA, which is connected to the ONUs through fiber optics. Exhibit 10, § 7.4; Exhibit 11, § C.7.2.1. The bandwidth allocation is a method of **dynamically sharing a digital communications channel bandwidth among multiple source nodes**. *Id*. GPON relies on a 125-µs frame structure (GTC/XGS TC or frames) (**carrier signal consisting of a continuous sequence of signal frames**). Exhibit 10, § 5.6; Exhibit 11, § C.10.1.1. The OLT corresponds to the recited **channel destination node**, while the ONUs, or any of their subscriber-specific Alloc-ID buffers correspond to the two or more individual source nodes (**multiple source nodes**). Exhibit 10, § 7.4; Exhibit 11, § C.7.2.1.

70. Defendants' process includes **"dynamically allocating the channel bandwidth among its multiple source nodes, by assigning channel signal frame transmission slots among the source nodes, through activating, via signal frame overhead information fields, individual source nodes of the channel to transmit data on the channel on future signal frames associated with said signal frame overhead information fields."** The DBA process in the OLT (**destination node**) allocates bandwidth for the ONUs (**source nodes**) using the BWmap (**dynamically allocating the channel bandwidth among its multiple source nodes**). Exhibit 10, § 7.4; Exhibit 11, § C.7.2.1. The standards describe two methods for the DBA process for inference of the Alloc-ID buffer occupancy levels, for the DBA process to assess their respective demand levels for upstream capacity. Exhibit 10, § 7.4.3; Exhibit 11, § C.7.2.1. The first method is the Status Reporting (SR) DBA, where explicit Alloc-ID buffer occupancy reports are solicited by the OLT and submitted to the OLT by the ONUs. *Id.* The second method is the Traffic Monitoring (TM) DBA which is based on the OLT's observation of the idle GEM frames transmitted by a given ONU/Alloc-ID in place of upstream transmission opportunities granted to it (via BWmap). *Id.* Both SR and TM are alternatives for the DBA for its inference of the Alloc-ID buffer occupancy levels, and the rest of the DBA process (including BWmap creation and application) operates the same way based on such assessments of Alloc-ID buffer load levels. The OLT assigns specific start and stop times (**slots**) (or start time and grant size for XGS-PON) for when each ONU is allowed to send packets during allotted frames (**assigning channel signal frame transmission slots**). Exhibit 10, § 8.1.3.6; Exhibit 11, § C.8.1.1.2; *see also* Exhibit 10, § 7.4; *see also* Exhibit 11, § C.7.2.1. The header of packets sent by the OLT contains the BWmap (**through activating, via signal frame overhead information fields**). Exhibit 10, § 8.1.3 (physical control block (PCBd)); Exhibit 11, § C.8.1 (PHY frame). Every customer modem has a unique ID (**individual source**

nodes . . . associated with said signal frame overhead information fields). Exhibit 10, § 7.4.3; Exhibit 11, § C.7.2.1. The source node receives the BWmap, processes it, and transmits its data in a specific future window (**transmit data . . . on future signal frames**). Exhibit 10, § 10.4; Exhibit 11, § C.13.1.

71. Defendants' process includes **"transmitting data on the channel by its individual source nodes on exactly those signal frames that were assigned to each individual source node via signal frame overhead information fields associated with said signal frames."** The source nodes send packets in the specific window indicated in the BWmap (**transmitting data on the channel by its individual source nodes on exactly those signal frames that were assigned to each individual source node**). Exhibit 10, § 10.4; Exhibit 11, § C.13.1. As also discussed above, the headers of packets (FS frame headers) sent by the OLT contain the BWmap (**via signal frame overhead information fields associated with said signal frames**). Exhibit 10, § 8.1.3 (PCBd); Exhibit 11, § C.8.1 (FS frame header).

72. In Defendants' process, the **"element of transmitting involves a sub-process performed by at least one of the set of channel source nodes that includes the below steps: I) monitoring specified overhead bitfields, referred to as Access Control Tags (ACTs), in each signal frame on the channel, and II) based at least in part on the step I), either transmitting or not transmitting data on the channel on the signal frame associated with a given ACT."** The ONUs (**source nodes**) constantly read the header (PCBd) of every incoming frame (**monitoring specified overhead bitfields**). Exhibit 10, § 8.1.3 (PCBd); Exhibit 11, § C.8.1 (FS frame header). Alloc-IDs are 12-bit numbers (or 14-bit for XGS-PON) assigned to a specific ONU/ONT (or service on that ONU/ONT (**Access Control Tags (ACTs)**)). Exhibit 10, § 8.1.3.6.1; Exhibit 11, § C.8.1.1.2.1. The ONU must check the header of every single downstream

frame to see if it has permission to send data in the next cycle (**in each signal frame on the channel**). Exhibit 10, § 8.1.3 (PCBd); Exhibit 11, § C.8.1 (FS frame header). The source nodes send packets in the specific window indicated in the BWmap (**based at least in part on the step I), either transmitting or not transmitting data on the channel on the signal frame associated with a given ACT**). Exhibit 10, § 10.4; Exhibit 11, § C.13.1.

73. Defendants have caused Plaintiff damage by direct infringement of the ’260 Patent.

74. In accordance with 35 U.S.C. § 287, Defendants have had actual notice and knowledge of the ’260 Patent by no later than the filing of this Complaint.

75. On information and belief, each Defendant continues, without license, to practice methods that infringe the ’260 Patent.

76. Defendants’ infringement of the ’260 Patent has injured OCS in its business and property rights. OCS is entitled to recover monetary damages for the injuries arising from Defendants’ infringement in an amount to be determined at trial.

77. The full extent of Defendants’ infringement is not presently known to OCS. OCS makes this preliminary identification of infringing products and/or services and infringed claims in this count without the benefit of discovery in this action, and expressly reserves the right to augment, supplement, and revise its identifications based on additional information obtained through discovery or otherwise.

COUNT THREE
DIRECT INFRINGEMENT OF U.S. PATENT NO. 10,567,474

78. OCS incorporates by reference the allegations in the preceding paragraphs as if fully set forth herein.

79. On information and belief, Defendants have infringed claims of the ’474 Patent, including at least claim 9, in violation of 35 U.S.C. § 271(a) by using the claimed method.

80. Claim 9 of the ’474 Patent (Exhibit 3) recites:

> 9. A method for interaction among a set of system components connected through a wide area network, the method comprising steps as follows:
>
> producing, at a first one of the components, a series of management data units, with a given one of said data units providing an array of information fields including a certain information field whose contents are used to indicate an action on part of a second one of the components;
>
> transferring the series of data units from said first component over the network to said second component to maintain synchronization between respective management data contents held at said first and second components; and
>
> at said second component, processing one or more of the information fields in the given data unit, including performing the action as indicated by the contents of the certain information field,
>
> wherein memories at each of said first and second components comprise respective segments for configuration and status files, and the given management data unit includes information segments corresponding to contents of both the configuration and the status files at said first and second components,
>
> wherein, in a first instance of the method, in which the first component comprises a computer subsystem and the second component comprises a remote element, the certain information field, based on whose contents the performing of the action at the remote element is done, is included within the configuration file segment of the given management data unit,
>
> wherein, in a second instance of the method, in which the first component comprises the remote element and the second component comprises the computer subsystem, the certain information field, based on whose contents the performing of the action at the computer subsystem is done, is included within the status file segment of the given management data unit,
>
> and wherein the management data contents held at the computer subsystem are accessible to a user of the method in order to enable the user to remotely manage the remote element.

81. On information and belief, Defendants have directly infringed at least claim 9 of the ’474 Patent by using NETCONF and YANG compliant SDN controllers, OLTs, and switches such as those available from Adtran.

82. On information and belief, Defendants are using NETCONF/YANG compliant equipment (including Adtran's Mosaic CP or equivalent) in its fiber network at least because Defendants are using Adtran OLTs specifically designed for use with the Mosaic CP. Exhibit 5; Exhibit 7 at 1; Exhibit 14.

83. Defendants practice **"[a] method for interaction among a set of system components connected through a wide area network."** NETCONF is designed as a client-server protocol to operate over an IP network. Exhibit 12, § 1. It uses a secure transport layer (typically SSH over TCP) to connect the "Manager" (Client) to the "Network Device" (Server) (**method for interaction among a set of system components connected through a wide area network**). *Id.*, §§ 1, 1.1, 1.2. A central manager synchronizes with and reconfigures remote devices (OLTs) using structured XML data units. *Id.*, § 1. In this architecture, the OLT's memory is logically divided into two distinct segments: a configuration segment (containing read/write <candidate>, <running>, and read-only <intended> configuration data) and a status/operational segment (containing the current OLT hardware state). Exhibit 16, §§ 5.1.2–4, 5.3. Each OLT actively monitors incoming XML data units received from the central manager and performs a "mismatch" check before taking action. Exhibit 12, § 7.2. Each OLT compares the incoming configuration values against its current memory segment; it only triggers a physical hardware reconfiguration if the new values differ from the current state. *Id.* Once the change is applied, users or automation scripts verify the result by performing a <get> operation, which retrieves both the "running" (intended) configuration and the OLT's live operational status. *Id.*, § 7.7.

84. Defendants' method includes **"producing, at a first one of the components, a series of management data units, with a given one of said data units providing an array of information fields including a certain information field whose contents are used to indicate**

an action on part of a second one of the components." YANG defines data fields through a structure where each data node is modeled as a leaf, leaf-list, container, or list. Exhibit 13, § 7. Every defined field (**array of information fields**) has a corresponding namespace-qualified element, which allows for unambiguous automated validation of the payload against the model. *Id.*, §§ 3, 5.3. In a NETCONF <edit-config> operation (**one of . . . a series of management data units**), YANG data fields are used to construct a target configuration (**certain information field whose contents are used to indicate an action on part of a second one of the components**) the OLT (**second one of the components**) must apply. *Id.* In a NETCONF <edit-config> message, the data fields are formatted according to the YANG data structures. In an example where an <edit-config> command is used to change a user "Fred" to the "superuser" user type, the XML command may have the following structure:

```
<rpc message-id="102"
     xmlns="urn:ietf:params:xml:ns:netconf:base:1.0"
     xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0">
<edit-config>
     <target>
          <running/>
     </target>
     <config>
          <system xmlns="urn:example:config">
               <user>
                    <name>fred</name>
                    <type>superuser</type>
               </user>
          </system>
     </config>
</edit-config>
</rpc>
```

Exhibit 13 at 91. The nesting of system > user > name > type is a schema defined by YANG. *Id.* The <name> leaf is also specified by YANG. *Id.*

85. Defendants' method includes **"transferring the series of data units from said first component over the network to said second component to maintain synchronization between respective management data contents held at said first and second components."** The central manager (client) sends (**transfer[s]**) RPC messages (**series of management data units**) over the network to the OLTs (servers) (**second component**) to synchronize the configuration of the OLTs to a master or intended configuration held at the central manager (**to maintain synchronization between respective management data contents held at said first and second components**). *See* Exhibit 12, Abstract, § 1.

86. Defendants' method involves **"at said second component, processing one or more of the information fields in the given data unit, including performing the action as indicated by the contents of the certain information field."** YANG specifies the information fields in an RPC (**information fields in the given data unit**). Exhibit 12, § 1, 7.2; Exhibit 13, § 7.6. The OLT (**second component**) parses the incoming RPC's XML payload (**certain information field**) against the YANG schema and applies the configuration change to the hardware registers (**processing one or more of the information fields in the given data unit, including performing the action as indicated by the contents of the certain information field**). *Id.*

87. In Defendants' method, **"memories at each of said first and second components comprise respective segments for configuration and status files, and the given management data unit includes information segments corresponding to contents of both the configuration and the status files at said first and second components."** When the OLT receives a <get> request, the OLT retrieves its operational (or "live") state from the <operational> datastore (**status files**). Exhibit 16, § 5.3. The intended configuration (**configuration … files**) is stored in the

<intended> datastore. *Id.*, § 5.1.4. Collectively, these memories are **memories at … second components [that] comprise respective segments for configuration and status files**. *Id.* The central manager stores in memory a local copy of the OLT configuration and operational state, as is necessary to perform the standard <compare> operation. Exhibit 18, § 1. Collectively, the memory locations are **memories at … first … components [that] comprise respective segments for configuration and status files**. *Id.*

88. In **"a first instance of** [Defendants'] **method, in which the first component comprises a computer subsystem and the second component comprises a remote element, the certain information field, based on whose contents the performing of the action at the remote element is done, is included within the configuration file segment of the given management data unit."** YANG specifies the information fields in the RPC command (**the certain information field**). Exhibit 12 at § 1, 7.2; Exhibit 13 at § 7.6. The central manager (**first component**) must include a computer subsystem to be operational. *Id.* The OLT (**second component**) is remote from the central manager. *Id.* The central manager reconfigures OLTs by sending an RPC with a YANG-compliant XML payload specifying the desired configuration (**certain information field**). *Id.* Upon receiving an <edit-config> command, the OLT applies the specified configuration to its running datastore; changes to the running configuration cause the OLT to reconfigure its hardware accordingly (**based on whose contents the performing of the action at the remote element is done**). Exhibit 12, § 7.2. An <edit-config> command sent by the central manager includes YANG specified hierarchy (such as system > user > name > type in the example below) and data fields (such as <name> in the example below) (**included within the configuration file segment of the given management data unit**).

```
<rpc message-id="102"
    xmlns="urn:ietf:params:xml:ns:netconf:base:1.0"
    xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0">
<edit-config>
    <target>
        <running/>
    </target>
    <config>
        <system xmlns="urn:example:config">
            <user>
                <name>fred</name>
                <type>superuser</type>
            </user>
        </system>
    </config>
</edit-config>
</rpc>
```

Exhibit 13 at 91. The nesting of system > user > name > type is a schema defined by YANG. *Id.* The <name> leaf is also specified by YANG. *Id.*

89. In Defendants' method, **"in a second instance of the method, in which the first component comprises the remote element and the second component comprises the computer subsystem, the certain information field, based on whose contents the performing of the action at the computer subsystem is done, is included within the status file segment of the given management data unit."** YANG specifies the information fields in the RPC command (**the certain information field**). Exhibits 12, 13. The central manager (**first component**) is remote from the OLT. *Id.* The OLT (**second component**) must include a **computer subsystem** to be operational. *Id.* The OLT reconfigures itself upon receiving and parsing an XML instruction from the central manager (**the certain information field, based on whose contents the performing of the action at the computer subsystem is done**). *Id.* When the central manager reads the status of the OLT, the OLT returns configuration and state data in an XML reply message (**is included within the status file segment of the given management data unit**). Exhibit 12, § 7.7.

90. In Defendants' method, "**the management data contents held at the computer subsystem are accessible to a user of the method in order to enable the user to remotely manage the remote element.**" At the central manager, an admin (**user**) can access the YANG configuration and state data (**management data contents**) stored at the central manager managing all OLTs, the latter being **accessible** by a <get> operation. Exhibit 12 at 11. The NETCONF protocol performed with the YANG data structures **enables the user to remotely manage** a distributed array of OLTs (**remote element**). Exhibits 12, 13.

91. Defendants have caused Plaintiff damage by direct infringement of the '474 Patent.

92. On information and belief, each Defendant continues, without license, to practice methods that infringe the '474 Patent.

93. In accordance with 35 U.S.C. § 287, Defendants have had actual notice and knowledge of the '474 Patent by no later than the filing of this Complaint.

94. Defendants' infringement of the '474 Patent has injured OCS in its business and property rights. OCS is entitled to recover monetary damages for the injuries arising from Defendants' infringement in an amount to be determined at trial.

95. The full extent of Defendants' infringement is not presently known to OCS. OCS makes this preliminary identification of infringing products and/or services and infringed claims in this count without the benefit of discovery in this action, and expressly reserves the right to augment, supplement, and revise its identifications based on additional information obtained through discovery or otherwise.

COUNT FOUR
DIRECT INFRINGEMENT OF U.S. PATENT NO. 10,848,546

96. OCS incorporates by reference the allegations in the preceding paragraphs as if fully set forth herein.

97. On information and belief, Defendants have infringed claims of the ’546 Patent, including at least claim 1, in violation of 35 U.S.C. § 271(a) by using the claimed method.

98. Claim 1 of the ’546 Patent (Exhibit 4) recites:

> 1. An infrastructure management method comprising a set of actions:
>
> at a set of remote elements,
>
> (a) holding element-side copies of configuration data (CD) at a storage media, wherein the element-side copies are repeatedly synchronized with corresponding copies of the CD maintained as user accessible copies of the CD, and
>
> (b) performing operations, by digital logic, at least in part based on contents of the element-side copies of the CD;
>
> at a computer system,
>
> (a) holding the user accessible copies of the CD at a storage facility, and
>
> (b) enabling a user of the method, via a user interface, to access the user accessible copies of the CD to manage the set of remote elements; and
>
> via a network interface, repeatedly transferring the user accessible copies of the CD to the storage media for holding the element-side copies of the CD at the remote elements, to maintain synchronization between the accessed user accessible copies of the CD and the element-side copies of the CD,
>
> wherein (i) contents of the user accessible copies of the CD are generated automatically through derivation from one or more parameters defining a contract that the set of remote elements are deployed for, and (ii) the corresponding element-side copies of the CD drive contents of hardware logic control registers at the set of remote elements,
>
> and wherein the user interface enables the user to access, including view and modify, both the user accessible copies of the CD as well as the parameters defining the contract that the set of remote elements are deployed for, from which parameters the user accessible copies of the CD as well as the corresponding element-side copies of the CD, and

consequently, the contents of the registers at the remote elements, are derived through automatic generation.

99. On information and belief, Defendants have directly infringed at least claim 1 of the '546 Patent by using NETCONF and YANG compliant SDN controllers, OLTs and switches such as those available from Adtran, or equivalents. Exhibits 5, 6.

100. Defendants practice **"[a]n infrastructure management method."** NETCONF uses a central manager (central manager or "client") that synchronizes with and reconfigures remote devices (OLTs) using RPCs that include XML data units (**[a]n infrastructure management method**). Exhibit 12, § 1; *see also* Exhibit 15. NETCONF uses a secure transport layer (typically SSH over TCP) to connect the "Manager" (Client) to the "Network Device" (Server). Exhibit 12, §§ 1, 1.1, 1.2. A central manager monitors and reconfigures remote devices (OLTs) using structured XML data units. *Id.*, § 1. In this architecture, the OLT's memory is logically divided into two distinct segments: a configuration segment (containing read/write "intent" data) and a status/operational segment (containing the current OLT hardware state). Exhibit 16, §§ 5.1.4, 5.3. Each OLT actively monitors incoming XML data units received from the central manager and performs a "mismatch" check before taking action. Exhibit 12, § 7.2. Each OLT compares the incoming configuration values against its current memory segment; it only triggers a physical hardware reconfiguration if the new values differ from the current state. *Id.* Once the change is applied, users or automation scripts verify the result by performing a <get> operation, which retrieves both the "running" (intended) configuration and the OLT's live operational status. *Id.*, § 7.7.

101. Defendants' method includes, **"at a set of remote elements, (a) holding element-side copies of configuration data (CD) at a storage media, wherein the element-side copies are repeatedly synchronized with corresponding copies of the CD maintained as user**

accessible copies of the CD, and (b) performing operations, by digital logic, at least in part based on contents of the element-side copies of the CD." The OLT stores its operational (or "live") state in the <operational> datastore and the intended configuration is stored in the <intended> datastore (**holding element-side copies of configuration data (CD) at a storage media**). Exhibit 16, §§ 5.1.4, 5.3. The central manager (where the **user accessible copies of the CD** are stored) requests that specific YANG-modeled data fields be sent at a fixed interval (e.g., every 60 seconds) (**wherein the element-side copies are repeatedly synchronized with corresponding copies of the CD maintained as user accessible copies of the CD**). Exhibit 22, § 3.1. The central manager receives a continuous, predictable heartbeat of operational data, which is essential for maintaining a copy of the network configuration and status without the high overhead of repeated <get> requests. Exhibit 12 at 11. The OLT reconfigures itself (by changing values in hardware registers) based on the "intended" configuration included in the <intended> datastore (**performing operations, by digital logic, at least in part based on contents of the element-side copies of the CD**). Exhibit 12, § 1, 7.2; Exhibit 13, § 7.6.

102. Defendants' method includes, **"at a computer system, (a) holding the user accessible copies of the CD at a storage facility, and (b) enabling a user of the method, via a user interface, to access the user accessible copies of the CD to manage the set of remote elements."** The OLT maintains its operational (or "live") state in the <operational> datastore and the intended configuration (**configuration data** file) in the <intended> datastore. Exhibit 16, §§ 5.1.4, 5.3. The central manager (**at a computer system, at a storage facility**) stores in memory a local copy of the OLT configuration and operational state (**holds copies of the CD**), as is necessary to perform the standard <compare> operation. Exhibit 18, § 1. The copies are **user-accessible** by an admin through the NETCONF and YANG compliant Mosaic CP or equivalent.

Exhibit 14, Exhibit 18. The NMDA standard specifies that management framework for users to access and compare these parallel datastores, allowing for the detection of discrepancies between the intended configurations and the actual OLT configuration (**to access the user accessible copies of the CD to manage the set of remote elements**). Exhibit 18, §§ 1, 3.

103. Defendants' method includes, **"via a network interface, repeatedly transferring the user accessible copies of the CD to the storage media for holding the element-side copies of the CD at the remote elements, to maintain synchronization between the accessed user accessible copies of the CD and the element-side copies of the CD."** The central manager repeatedly sends the RPC's <edit-config> message including an XML payload containing the copies of the CD from the central manager to the OLTs (**remote elements**) which the OLTs store (in remote-element-side memory) as the intended configuration (**repeatedly transferring the user accessible copies of the CD to the storage media for holding the element-side copies of the CD at the remote elements**). Exhibits 12, 13. This is to ensure synchronicity between the configurations specified in the "master" held at the central manager and the remote OLTs (**to maintain synchronization between the accessed user accessible copies of the CD and the element-side copies of the CD**). *Id.*

104. In Defendants' method, **"(i) contents of the user accessible copies of the CD are generated automatically through derivation from one or more parameters defining a contract that the set of remote elements are deployed for, and (ii) the corresponding element-side copies of the CD drive contents of hardware logic control registers at the set of remote elements."** The contents of the user-accessible configuration datastores are generated automatically through derivation by the central manager, which transforms high-level contract parameters (e.g., bandwidth or QoS specifications) into the specific XML/YANG-modeled

configuration data required to reach a desired operational state (**contents of the user accessible copies of the CD are generated automatically through derivation from one or more parameters defining a contract that the set of remote elements are deployed for**). Exhibit 14 at 3 ("Access Domain Orchestrator - Provides the business logic for creation and provisioning of services or applications"; *see also* https://sdn.systemsapproach.org/access.html ("control state flows down (e.g., assigning subscribers to particular QoS classes) and monitoring state flows up"). Once this derived intent is transmitted via an RPC operation like <edit-config>, the config is stored in the OLT's <running/> configuration datastore (**element-side copies**). Exhibit 12, §§ 5.1, 7.2. The OLT compares the received configuration with the current operational state and accordingly makes changes to the contents of hardware registers (**(ii) the corresponding element-side copies of the CD drive contents of hardware logic control registers at the set of remote elements**). *Id.*

105. In Defendants' method, **"the user interface enables the user to access, including view and modify, both the user accessible copies of the CD as well as the parameters defining the contract that the set of remote elements are deployed for, from which parameters the user accessible copies of the CD as well as the corresponding element-side copies of the CD, and consequently, the contents of the registers at the remote elements, are derived through automatic generation."** A central manager user with full admin privileges has read/write access to the customer service contract parameters (e.g., max connection speed) and the master configuration data (**wherein the user interface enables the user to access, including view and modify, both the user accessible copies of the CD as well as the parameters defining the contract that the set of remote elements are deployed for**). Exhibit 14 at 3 ("Access Domain Orchestrator - Provides the business logic for creation and provisioning of services or

applications”). Using this and other information, the central manager automatically generates a “master” configuration file for configuration data to be sent to each OLT for local storage (**from which parameters the user accessible copies of the CD as well as the corresponding element-side copies of the CD . . . are derived through automatic generation**). *Id.* The OLT compares the received configuration with the current operational state and accordingly makes changes to the contents of hardware registers (**and consequently, the contents of the registers**). *Id.*

106. Defendants have caused Plaintiff damage by direct infringement of the ’546 Patent.

107. In accordance with 35 U.S.C. § 287, Defendants have had actual notice and knowledge of the ’546 Patent by no later than the filing of this Complaint.

108. On information and belief, each Defendant continues, without license, to practice methods that infringe the ’546 Patent.

109. Defendants’ infringement of the ’546 Patent has injured OCS in its business and property rights. OCS is entitled to recover monetary damages for the injuries arising from Defendants’ infringement in an amount to be determined at trial.

110. The full extent of Defendants’ infringement is not presently known to OCS. OCS makes this preliminary identification of infringing products and/or services and infringed claims in Count Four without the benefit of discovery in this action, and expressly reserves the right to augment, supplement, and revise its identifications based on additional information obtained through discovery or otherwise.

PRAYER FOR RELIEF

WHEREFORE, OCS prays for the following judgments and relief:

(a) A judgment that Defendants have infringed and are infringing the Asserted Patents;

(b) An award of damages adequate to compensate OCS for Defendants' patent infringement, and an accounting to adequately compensate OCS for the infringement, including, but not limited to, lost profits and/or a reasonable royalty;

(c) An award of pre-judgment and post-judgment interest at the maximum rate allowed by law;

(d) An order finding that this is an exceptional case and awarding OCS its costs, expenses, disbursements, and reasonable attorneys' fees related to each Defendant's patent infringement under 35 U.S.C. § 285 and all other applicable statutes, rules and common law; and

(e) Such other further relief, in law or equity, as this Court deems just and proper.

DEMAND FOR JURY TRIAL

In accordance with Rule 38 of the Federal Rules of Civil Procedure, OCS hereby demands a jury trial on all issues triable before a jury.

Dated: July 20, 2026

Of Counsel:
Cook Alciati
Greg H. Gardella
Michael Stimson
Lee Goldberg
GARDELLA ALCIATI
80 M Street SE
1st Floor
Washington, D.C. 20003
(703) 556-9600
calciati@ga-iplaw.com
ggardella@ga-iplaw.com
mstimson@ga-iplaw.com
lgoldberg@ga-iplaw.com

YOUNG CONAWAY STARGATT
& TAYLOR, LLP

/s/ Pilar G. Kraman
Pilar G. Kraman (No. 5199)
Colin A. Keith (No. 7074)
Rodney Square
1000 North King Street
Wilmington, DE 19801
(302) 571-6600
pkraman@ycst.com
ckeith@ycst.com

Attorneys for Plaintiff
Optimum Communications Services, Inc.