

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

JUNIPER NETWORKS, INC.,
Petitioner,

v.

MONARCH NETWORKING SOLUTIONS LLC,
Patent Owner.

U.S. Patent No. 8,693,369

“Method of routing a data packet in a network and an associated device”

IPR2024-00175

**SUPPLEMENTAL DECLARATION OF DR. SAMRAT
BHATTACHARJEE IN SUPPORT OF PETITIONER’S REPLY
TO PATENT OWNER’S RESPONSE**

I.	INTRODUCTION.....	1
II.	GROUND 1: BISWAS TEACHES CLAIMS 8 AND 12, WHICH ARE DIRECTED TO INBOUND—NOT OUTBOUND—TRANSMISSIONS	1
III.	GROUND 2 AND 4: FANGMAN’S NAPPT USES PORT RANGES TO ROUTE INBOUND DATA PACKETS.....	9
IV.	GROUND 3 AND 5: A POSITA WOULD HAVE FOUND IT OBVIOUS TO APPLY KING’S OUTBOUND TEACHINGS TO INBOUND TRANSMISSIONS.....	22
V.	CONCLUSION	26

PETITIONER'S EXHIBIT LIST

EX1001	U.S. Patent No. 8,693,369 to Mohamed Boucadair (“’369 patent”)
EX1002	Prosecution history of U.S. Patent No. 8,693,369 (“’369 Prosecution History”)
EX1003	Declaration of Dr. Samrat Bhattacharjee under 37 C.F.R. §1.68
EX1004	<i>Curriculum Vitae</i> of Dr. Bhattacharjee
EX1005	Declaration of Dr. Sylvia Hall-Ellis under 37 C.F.R. §1.68
EX1006	P. Richter et al., “A Primer on IPv4 Scarcity,” (Feb. 27, 2015), available at https://arxiv.org/pdf/1411.2649.pdf
EX1007	H. Yeom et al., “IP Multiplexing by Transparent Port-Address Translator,” 10th USENIX System Administration Conference (LISA X) (Sept. 29, 1996), available at https://www.usenix.org/conference/lisa-96/ip-multiplexing-transparent-port-address-translator
EX1008	RFC 3022 (Traditional NAT), dated January 2001, available at https://datatracker.ietf.org/doc/html/rfc3022
EX1009	U.S. Patent Application Publication No. 2002/0016826 to Johansson et al. (“Johansson”)
EX1010	A. Reitz, “The Dynamic Port Reservation Protocol,” Univ. Ill. at Urbana-Champaign (2001), available at http://redefine.dyndns.org/~andyr/documents_and_writings/thesis/docs/areitz_msthesis.pdf
EX1011	M. Shin et al., “Using a Single IPv4 Global Address in DSTM,” Internet Engineering Task Force (“IETF”) (Feb. 22, 2001), available at https://www.ietf.org/archive/id/draft-shin-dstm-single-ipv4-00.txt (“Shin”)
EX1012	U.S. Patent No. 7,356,045 to Satapati et al. (“Satapati”)
EX1013	U.S. Patent Application Publication No. 2003/0154306 to Perry (“Perry”)

EX1014	RFC 950 (Internet Standard Subnetting Procedure), dated August 1985, available at https://datatracker.ietf.org/doc/html/rfc950
EX1015	U.S. Patent No. 7,624,195 to Biswas et al. ("Biswas-195")
EX1016	WO Publication No. 01/02965 A2 to Herbst et al. ("Herbst")
EX1017	RFC 7597 (Mapping of Address and Port with Encapsulation), dated July 2015, available at https://datatracker.ietf.org/doc/html/rfc7597
EX1018	RFC 4632 (Classless Inter-domain Routing), dated August 2006, available at https://datatracker.ietf.org/doc/html/rfc4632
EX1019	U.S. Patent Application No. 10/873,547 to Tornetta et al. ("Tornetta")
EX1020	U.S. Patent No. 7,873,038 to LaVigne et al. ("LaVigne")
EX1021	U.S. Patent Application No. 11/699,765 to McAlister ("McAlister")
EX1022	U.S. Patent No. 8,019,889 to Biswas et al. ("Biswas")
EX1023	U.S. Patent No. 7,394,809 to Kumar et al. ("Kumar")
EX1024	U.S. Patent No. 7,068,647 to Fangman et al. ("Fangman")
EX1025	U.S. Patent No. 8,032,639 to King et al. ("King")
EX1026	RFC 791, Internet Protocol Specification (September 1981), available at https://datatracker.ietf.org/doc/html/rfc791
EX1027	U.S. Patent No. 5,014,221 to Mogul ("Mogul")
EX1028	U.S. Patent No. 5,163,131 to Row et al. ("Row")
EX1029	July 20, 2006 Capture of RFC 791, Wayback Machine, available at https://web.archive.org/web/20060720182319/http://www.ietf.org/rfc/rfc791.txt
EX1030	Michael Witt, A Comment on Current Source Routing Techniques (1987), available at https://dl.acm.org/doi/pdf/10.1145/36727.36728

EX1031	Andrew Tanenbaum, <i>Computer Networks</i> (5th Ed. 2011)
EX1032	Monarch's Opposition to Juniper's Motion to Transfer, <i>Monarch Networking Solutions LLC v. Juniper Networks, Inc.</i> , No. 1:23-cv-670-TSE-LRV (E.D. Va. Oct. 3, 2023), ECF No. 52
EX1033	PTO Reexamination Operational Statistics, available at https://www.uspto.gov/learning-and-resources/statistics/reexamination-information
EX1034	Complaint, <i>Monarch Networking Solutions LLC v. Juniper Networks, Inc.</i> , No. 1:23-cv-670-TSE-LRV (E.D. Va. May 23, 2023), ECF No. 1
EX1035	Unified EPR (Control No. 90/015,315), filed on October 23, 2023
EX1039	Supplemental Declaration of Dr. Samrat Bhattacharjee under 37 C.F.R. §1.68
EX1040	Y. Miyazaki et al., <i>A Proposal for a NAT Traversal System That Does Not Require Additional Functions at Terminals</i> , IEEE (Oct. 2007), available at https://ieeexplore.ieee.org/abstract/document/4428885
EX1041	W. Jiang et al., <i>Towards Junking the PBX: Deploying IP Telephony</i> , Columbia Univ. (2001), available at https://www.cs.columbia.edu/~hgs/papers/Jian0106_Junking.pdf
EX1042	T. Wallingford, <i>Switching to VoIP</i> , O'Reilly (2005).
EX1043	P. Engelstad & G. Egeland, <i>Peer-to-Peer IP Telephony</i> , 102 <i>Teletronikk</i> 1 (2006), available at https://www.telenor.com/wp-content/uploads/2012/05/T06_1.pdf
EX1044	Y. Yu et al., <i>An Efficient NAT Traversal for SIP and Its Associated Media Sessions</i> , Int'l Comp. Symposium (Dec. 15-17, 2004), available at http://dspace.fcu.edu.tw/bitstream/2377/1829/1/ce07ics002004000150.pdf

EX1045	F. Thernelius, <i>SIP, NAT, and Firewalls</i> , Royal Institute of Tech. Stockholm (May 2000), available at https://www.cs.columbia.edu/sip/drafts/Ther0005_SIP.pdf
EX1046	Andrew Tanenbaum, <i>Computer Networks</i> (4th Ed. 2003)

I. INTRODUCTION

1. I have been retained by Quinn Emanuel Urquhart & Sullivan, LLP on behalf of Juniper Networks, Inc. (“Juniper” or “Petitioner”) as an independent expert in this *inter partes* review before the Patent Trial and Appeal Board of the United States Patent and Trademark Office to review claims 1-17 (“Challenged Claims”) of U.S. Patent No. 8,693,369 (“the ’369 patent”).

2. In forming my opinions, I have considered and relied on my education, knowledge of the relevant field, knowledge of scientific and engineering principles, and my general experience in the field of art. I have reviewed the Patent Owner Response (“POR”), all exhibits, and the declaration of Mr. Minoli, as well as other materials and references cited in this declaration.

II. GROUND 1: BISWAS TEACHES CLAIMS 8 AND 12, WHICH ARE DIRECTED TO INBOUND—NOT OUTBOUND—TRANSMISSIONS

3. Claims 8 and 12 of the ’369 patent are directed to inbound transmissions, not outbound transmissions. In fact, the ’369 patent in general is directed to a routing device that receives an inbound data packet and then uses pre-assigned port ranges to determine where to route that inbound data packet. Claims 8 and 12 in particular explicitly limit themselves to inbound transmissions. These claims recite the following:

8. The method according to claim 7 of configuring a piece of terminal equipment, comprising a step of assigning to said piece of terminal


equipment an identifier of a routing device adapted to route a data packet comprising said primary address as the destination address to said piece of terminal equipment, and the configuration message further comprising said routing device identifier.

12. The configuration method according to claim 11, wherein, once an identifier of a routing device adapted to route a data packet comprising said primary address as the destination address to said piece of terminal equipment has been received, said method further comprises a step of sending the configuration data to said routing device.

EX1001, claims 8, 12.

4. As shown above, both claims require “a data packet comprising said primary address *as the destination address* to said piece of terminal equipment.”

EX1001, claims 8, 12 (emphasis added). The emphasized portion of the claims confirms that claims 8 and 12 are directed to transmissions that are inbound to the network containing terminal devices (*e.g.*, network RT in Figure 3) because the primary address serves as a destination address only for inbound transmissions that pass into the telecommunications network containing the terminal devices. For example, in the annotated figure below, a data packet from device **150** to terminal equipment **110** is an inbound data packet that contains the primary address as its destination address:



EX1001, Fig. 3 (annotations added).

5. For outbound transmissions, the primary address would instead serve as a source address, which the claims do not purport to cover, as shown in the annotated figure below:




FIG. 3

EX1001, Fig. 3 (annotations added).

6. Additionally, both claims clarify that the “routing device [is] adapted to route a data packet . . . *to said piece of terminal equipment.*” EX1001, claims 8, 12 (emphasis added). This emphasized language further confirms that these claims

are directed to data packets sent *to* the terminal equipment (*i.e.*, inbound data packets), not *from* the terminal equipment (*i.e.*, outbound data packets).

7. Thus, when claims 8 and 12 require “a data packet comprising said primary address as the destination address to said piece of terminal equipment,” they are limiting their scope to data packets *entering* the telecommunications RT network, meaning inbound transmissions.

8. Biswas teaches these claims, as discussed in my prior declaration. EX1003 ¶¶105-107, 120-123. In particular, gateway master **110** receives an inbound data packet that contains the global address as the destination IP address, replaces the global address with a local address, and then routes that inbound data packet to the correct edge router **104** and host **102**:




Figure 1

EX1022, Fig. 1 (annotations added).

9. Here, the global address must be able to identify gateway master **110** for inbound transmissions. A data packet containing an address as the destination address is routed to the device identified by that destination address. As shown in the figure above, data packets from the public network (*e.g.*, host **118**) containing the global address as the destination address are routed to the gateway master **110**. Upon receipt, “[t]he Gateway Master **110** [] translates global addresses of data

received from public network **112** into local addresses of data being output from the Gateway Master **110** to private network **108**” and then routes the data packet accordingly. EX1022, 5:52-6:24; *see also* EX1022, 5:55-57 (“The gateway master 110 is operable to handle communication between the ISP network 108 and public network 112.”). This confirms that the global address is “an identifier of [the] routing device,” gateway master **110**. Otherwise, gateway master **110** would not receive inbound data packets that use the global address as the destination address. In fact, Biswas explains that the gateway master acts as a NAT device, which would further confirm to a POSITA that the global address identifies the gateway master because conventional NATs used to connect to the Internet are also identifiable by a global address. EX1022, 5:57-58; *see also* EX1040, 2 (“The packet is encapsulated ***with the NAT router’s global address***, GA2, and sent to the NAT router.” (emphasis added)); EX1040, Fig. 1 (depicting a global node sending a data packet to a NAT using the NAT’s global address as the destination address).

10. Finally, in my opinion, Patent Owner’s explanation of how outbound transmissions work is incorrect. Patent Owner asserts that an internal host cannot send a data packet to the gateway master that contains the global address as the destination address. I disagree. Even within a private network, a host can use a global address to communicate with the gateway master. In fact, I have tested and confirmed this functionality myself by using the popular Ping Tool. Specifically, I

used a host on an internal network (*e.g.*, a host **102** in Biswas) to solicit Internet Control Message Protocol (“ICMP”) Echo Replies from both the internal and external addresses of a NAT device on the internal network (*e.g.*, the gateway master **110** in Biswas). The ICMP Echo Replies that I received from the NAT device were from both the internal and external interface addresses of the NAT, signifying that the external address *is* reachable by internal network hosts. This first figure shows that when I pinged the NAT device’s internal address (10.0.1.1), the host I used transmitted five packets using the NAT device’s internal address as the destination address and then received five response packets from the NAT device, confirming that the transmissions were successful:

```
% ping -c 5 10.0.1.1
PING 10.0.1.1 (10.0.1.1): 56 data bytes
64 bytes from 10.0.1.1: icmp_seq=0 ttl=64 time=1.214 ms
64 bytes from 10.0.1.1: icmp_seq=1 ttl=64 time=1.174 ms
64 bytes from 10.0.1.1: icmp_seq=2 ttl=64 time=1.345 ms
64 bytes from 10.0.1.1: icmp_seq=3 ttl=64 time=1.326 ms
64 bytes from 10.0.1.1: icmp_seq=4 ttl=64 time=1.100 ms

--- 10.0.1.1 ping statistics ---
5 packets transmitted, 5 packets received, 0.0% packet loss
round-trip min/avg/max/stddev = 1.100/1.232/1.345/0.092 ms
```

This second figure shows that when I pinged the NAT device’s external address (*i.e.*, the global address in Biswas) (173.79.60.190), the host I used transmitted five packets using the NAT device’s external address as the destination address and then

received five response packets from the NAT device, confirming that the transmissions were successful:

```
% ping -c 5 173.79.60.190
PING 173.79.60.190 (173.79.60.190): 56 data bytes
64 bytes from 173.79.60.190: icmp_seq=0 ttl=64 time=1.231 ms
64 bytes from 173.79.60.190: icmp_seq=1 ttl=64 time=1.271 ms
64 bytes from 173.79.60.190: icmp_seq=2 ttl=64 time=1.230 ms
64 bytes from 173.79.60.190: icmp_seq=3 ttl=64 time=1.748 ms
64 bytes from 173.79.60.190: icmp_seq=4 ttl=64 time=1.321 ms

--- 173.79.60.190 ping statistics ---
5 packets transmitted, 5 packets received, 0.0% packet loss
round-trip min/avg/max/stddev = 1.230/1.360/1.748/0.197 ms
```

This functionality works the same way now as it did before the '369 patent's priority date. Thus, I disagree with Patent Owner's assertion that internal hosts cannot send data packets to the gateway master using the global address as the destination address. As my real-world experiment confirmed, internal hosts *can* send data packets to the gateway master using the global address as the destination address.

III. GROUNDS 2 AND 4: FANGMAN'S NAPPT USES PORT RANGES TO ROUTE INBOUND DATA PACKETS

11. As explained next, IP telephones (as disclosed in Fangman) were incompatible with traditional NAT.¹ In my opinion, a POSITA would have

¹ I use the term "traditional NAT" to refer to the process of maintaining a connection-state table and using that table to route inbound data packets (or, in the words of Patent Owner, the process by which "the service gateway checks whether

understood, based on explicit disclosures in Fangman, that Fangman's system is operable only if it uses port ranges to route inbound data packets, in the same way described in the '369 patent.

12. Fangman invented Network Address Persistent Port Translation ("NAPPT")—and specifically its use of port ranges to route inbound transmissions—to overcome traditional NAT's incompatibility with IP telephones. Traditional NAT works by maintaining a connection-state table. When an internal host sends an outbound data packet, the NAT intercepts the data packet; modifies information in the data packet's header, like the source address and/or port number;

[a] port number in the destination packet matches the port number used by the IP telephone in a prior outgoing transmission"). Fangman uses the more general terms "NAT" or "NAPT" to refer to the generic process of translating a private IP address into a public IP address and vice versa, which might or might not involve translating addresses via a connection-state table. *See, e.g.*, EX1024, 8:52-54 ("The SG **170** may continue to replace the internal or private source IP address with the NATed external source IP address."); EX1024, 10:61-62 (explaining that the service gateway can "[p]rovide basic NAT translation (private to public IP) without changing port number"); EX1024, 11:37 (explaining that the service gateway can "[a]ssign [a] NAT address").

logs information from the data packet in the NAT's connection-state table; and then routes the data packet to the recipient, external host. When an external host sends an inbound data packet to an internal host in the private network, the NAT receives the data packet, consults its connection-state table to determine if information in the inbound data packet matches a prior outbound transmission, and then processes the inbound data packet. If information in the inbound data packet matches a prior outbound transmission (*i.e.*, the inbound transmission was solicited by an internal host), then the NAT routes the inbound data packet to the associated internal host. If information in the inbound data packet does not match a prior outbound transmission (*i.e.*, the inbound transmission was not solicited by an internal host), then the NAT discards the data packet.

13. Traditional NAT's inability to handle unsolicited inbound transmissions makes it incompatible with IP telephony systems, which must be able to handle unsolicited inbound transmissions so that telephones on the network can receive calls. As Fangman explains, "the resulting need to use Network Address Translation (NAT) and related technologies" and "the limitations of existing firewall technologies" were two of the "primary challenges" of IP telephony. EX1024, 1:36-46. "Complicating matters further, VoIP protocols often use pairs of ports for communication, initiated from both the inside *and outside* of the network.

Firewalls^[2] tend to support single port communication *only initiated from the inside*,” i.e., solicited inbound transmissions. EX1024, 2:12-27 (emphasis added).

14. This unsolicited-transmissions issue was well known prior to the '369 patent's invention. POSITAs recognized that “[i]nterworking with the corporate firewalls and Network Address Translators (NATs) is another challenge” to implementing IP telephony. EX1041, 8. “A side effect of NAT is that protocols that must use both outbound and inward sockets (like SIP and H.323) simply don't work.” EX1042, 342. “A peer that is located outside a NAT cannot contact a peer located inside the NAT” and “[i]f two peers are located behind different NATs, peer-to-peer communication is not possible.” EX1043, 2. Although IP telephony protocols “function well in the public network,” they “fail when traversing through [a] NAT” because “NAT only permits the devices in the private realm to initiate the sessions, so devices behind a NAT may be able to make SIP calls but not be able to receive them.” EX1044, 873. Put simply, “a problem occurs” with NAT “if a session is not initiated from the inside, because there will be no easy way to set up the translation table.” EX1045, 25.

² The service gateway in Fangman implements both the NAT and the firewall. EX1024, 8:44-57, 9:52-58.

15. Fangman describes this as a “[c]onsiderable problem[]” of NAT, (EX1024, 2:54-65), which it is: if Fangman were to use a process “similar to traditional NAT or PAT” that requires a “prior outgoing transmission,” as Patent Owner assumes, (POR 13, 22), then the IP telephones in Fangman’s network would not be able to receive unsolicited inbound calls. In fact, the telephones likely would not be able to make outbound calls either because, if the caller and recipient are both behind NATs, then the caller will not be able to place an outbound call to the recipient because the recipient’s NAT would discard the call as an unsolicited inbound transmission. *See* EX1043, 2 (“If two peers are located behind different NATs, peer-to-peer communication is not possible.”). In sum, Fangman explains, and the prior art confirms, that traditional NAT is incompatible with IP telephony because of this unsolicited-inbound-transmissions problem.³ Fangman therefore disavows traditional NAT and instead teaches an improvement to it, termed NAPPT. EX1024, 1:61-64, 2:66-67, 4:4-17.

16. In light of the foregoing, NAPPT must use port ranges to route inbound data packets because it cannot rely on prior outbound transmissions. Fangman

³ Fangman acknowledges that a solution to this problem exists, but Fangman disparages it as a “complex” and “convoluted process” that “introduc[es] substantial inefficiencies to the system.” EX1024, 2:54-65.

explicitly makes this point, explaining that “[i]nstead of dynamically assigning ports to IP telephones” as in traditional NAT, NAPPT “allocate[s] **a range** of port numbers for each IP telephone” and stores the port ranges in “a table of semi-permanent mappings.” EX1024, 8:45-67; *see also, e.g.*, EX1024, 3:25-37 (“The Service Gateway may store the DHCP lease information” that contains “the range of port numbers . . . for the IP telephone.”); EX1024, 11:14 (explaining that the service gateway can “[a]ssign port ranges for host, reserved in NAPT table.”). NAPPT then handles inbound transmissions by “receiving a data packet from the IP device” that includes “a public destination IP address, a destination port number **in the assigned range of port numbers**, and source information, where the public destination IP address and the destination port number may be used to uniquely identify the IP telephone.” EX1024, 4:18-25 (emphasis added). “The Service Gateway may then perform a Network Address Persistent Port Translation (NAPPT) on the data packet”—not traditional NAT—by consulting its table of port-range mappings, determining what “range of port numbers” the destination port falls into, and then routing the data packet based on the private address associated with that range of port numbers. EX1024, 4:25-27. Fangman accordingly explains that the service gateway stores “the assigned **ports** for the IP telephone” to perform the translation process. EX1024, 32:48-57 (emphasis added). By using port ranges, “the traditional problem of dynamic port numbers,” such as “[t]he problem of supporting application

layer protocols that use UDP . . . and session[s] initiated from outside the network[,] may be addressed by the SG **170**” via “[t]he perpetual NAT table maintained by the SG **170**.” EX1024, 14:9-14.

17. Fangman never says that it uses traditional NAT, and it also never says that it relies on prior outbound transmissions to route inbound data packets. To the contrary, Fangman contains several disclosures that confirm that Fangman relies on port ranges to route inbound data packets rather than “prior outgoing transmissions. For instance, in the example provided in Figures 10A and 10B, a service gateway **170B** receives an inbound data packet and performs NAT on that inbound data packet without first sending a prior outbound packet. *See* EX1024, 30:12-24, Figs. 10A-B. In this example, the service gateway’s database would not contain information from a “prior outgoing transmission” because there *is no* prior outgoing transmission. The only way the service gateway could route the data packet in this situation is by using NAT to route the inbound data packet based on stored port range.

18. Further, Fangman emphasizes that the only reason to maintain a connection-state table, as used in traditional NAT, is to guarantee redundancy. As Fangman explains, “the primary SG **170** *may* maintain a connection state table for all NAT sessions” in order to ensure “*full redundancy* for each function that the SG **170** performs.” EX1024, 10:2-11 (emphasis added). If a connection state table is

needed only for redundancy, then the service gateway must be using some process other than traditional NAT when routing inbound data packets.

19. Additionally, Fangman discloses another hardware component, a media gateway controller, that the system uses to set up calls between IP telephones. EX1024, 2:43-53, 4:51-5:28. To establish these calls, the media gateway controller must be able to identify private addresses for IP telephones based on a public IP address and port number—just like the service gateway. Fangman therefore explains that the media gateway controller can “use the public IP address *and the port range* to determine the private IP address, for example, by using look-up tables.” EX1024, 3:53-64 (emphasis added). Thus, there are multiple disclosures throughout Fangman that confirm Fangman uses port ranges to identify internal hosts, not traditional NAT as Patent Owner argues. *E.g.*, EX1024, 1:36-46 (“[T]here remain many challenges facing developers of VoIP equipment” including “the resulting need to use Network Address Translation (NAT) and related technologies.”), 1:61-64 (In traditional NAT “when a Media Gateway Controller (MGC) or Trunking Gateway (TG) receives communication from an end-node . . . considerable problems may result.”), 2:12-27 (“Complicating matters further, VoIP protocols often use pairs of ports for communication, initiated from both the inside and outside of the network.”), 2:45-67 (describing alternative solutions as “complex” and “convoluted”), 3:11-37 (“The Service Gateway may use an identifier, e.g., a vendor ID, of the IP telephone to

determine a range of port numbers to assign to the IP telephone.”), 3:53-64 (“The Media Gateway Controller may be operable to use the public IP address and the port range to determine the private IP address, for example, by using look-up tables.”), 4:4-27 (“In one embodiment, performing IP communications using the IP telephone may include the Service Gateway receiving a data packet from the IP telephone, including a private source IP address, a source port number in the assigned range of port numbers, and destination information associated with an IP device.”), 8:45-67 (“Instead of dynamically assigning ports to IP telephones, the SG 170 may allocate a range of port numbers for each IP telephone 120”), 10:1-11 (“In the background, the primary SG 170 may maintain a connection state table for all NAT sessions” to provide “full redundancy.”), 11:14 (“Assign port ranges for host, reserved in NAPT table.”), 14:9-14 (“The problem of supporting application layer protocols that use UDP, multiple and sequential port numbers, and sessions initiated from outside the network may be addressed by the SG 170. The perpetual NAPT table maintained by the SG 170 may solve the traditional problem of dynamic port numbers.”), 30:12-24 (“The second SG 170B may decrypt the data, as indicated in step 1062, perform NAPPT in step 1064” without first receiving a prior outbound transmission.), 32:48-57 (“the Service Gateway performing a network address persistent port translation (NAPPT) on the data packet, including storing translation information in a data structure, wherein the translation information comprises: the

private source IP address of the IP telephone; a public IP address of the Service Gateway; and the assigned ports for the IP telephone.”), Figs. 10A-B (depicting a service gateway performing NAPPT without the need for a prior outbound transmission).

20. I disagree with Mr. Minoli's statement that “Fangman confirms that it uses only a port number” for routing inbound data packets because it discloses “using . . . ***the destination port number*** to uniquely identify the IP telephone.” EX2014 ¶86 (emphasis in original). I understand that, here, Mr. Minoli is trying to construe the phrase “using . . . the destination port number” to mean “exclusively relies on the destination port number.” But I further understand that this is not the plain meaning of that phrase, much less how it is used in this disclosure. I understand that the phrase “use” means “to employ for some purpose.” *Use*, Dictionary.com, available at <https://www.dictionary.com/browse/use> (last accessed Oct. 14, 2024); *see also Use*, Black's Law Dictionary (14th ed. 2024) (“To employ for the accomplishment of a purpose.”). NAPPT “uses . . . the destination port number to uniquely identify the IP telephone” when it compares the destination port number to the stored port range to determine the recipient host. Thus, I disagree with Mr. Minoli's suggestion that “using . . . the destination port number to uniquely identify the IP telephone” means that Fangman does not use port ranges.

21. In addition to Fangman's explicit disclosures that confirm it teaches using port ranges to route inbound data packets, a POSITA would have generally understood that using port ranges to route inbound data packets was the preferred procedure. For example, a POSITA would have understood that maintaining a connection-state table to route inbound data packets would be an inefficient use of memory in Fangman's system. The service gateway is already "maintain[ing] a table of semi-permanent mappings of the external IP address and ports to the internal hosts and ports," and it can route inbound packets according to the port ranges stored in this table. EX1024, 8:49-51; *see also* EX1024, 3:25-37 ("The Service Gateway may store the DHCP lease information" that contains "the range of port numbers . . . for the IP telephone."); EX1024, 11:14 ("Assign port ranges for host, reserved in NAPT table."). A POSITA would have therefore recognized that maintaining a separate connection-state table would be redundant and needlessly consume the gateway's memory. The table of port ranges stored by the service gateway is sufficient by itself.

22. I understand that the Board, in its institution decision, interpreted Fangman as mapping individual "external IP addresses and ports" to individual "internal hosts and ports" such that "[w]hen the service gateway receives a data packet destined for the IP telephone, it uses the table to look up the private IP address

of the telephone based on the port number, not the port range.”⁴ Paper 12, 23-24. Respectfully, I disagree for the reasons discussed above: Fangman stored port range associations, not individual port numbers as in traditional NAT. This is especially evident in Fangman's claim language, which claims “the Service Gateway performing a network address persistent port translation (NAPPT) on the data packet, *including storing translation information in a data structure, wherein the translation information comprises*: the private source IP address of the IP telephone; a public IP address of the Service Gateway; and *the assigned ports* for the IP telephone.” EX1024, 32:48-57 (emphasis added).

23. Further, a POSITA would have recognized that maintaining or using a table that maps individual internal ports to individual external ports is highly inefficient. For example, Fangman could theoretically use the following table for routing inbound packets to a first terminal device, as the Board implies in its institution decision:

⁴ I understand that Patent Owner did not make the following argument, but Petitioner has asked me to analyze it in case the Board allows Patent Owner to make this new argument in sur-reply.

Private Address	Internal Port	Public Address	External Port
10.255.255.123	8,192	31.31.98.152	8,192
10.255.255.123	8,193	31.31.98.152	8,193
10.255.255.123	8,194	31.31.98.152	8,194
...	...	...	...
10.255.255.123	8,447	31.31.98.152	8,447

But a POSITA would have known that this implementation is inefficient because it stores a large amount of redundant data, and would have further known that the entire table above could be simplified to the following single row:

Private Address	Public Address	Assigned Port Range
10.255.255.123	31.31.98.152	8,192-8,447 ⁵

In fact, the latter table is exactly what Fangman claims: “the Service Gateway perform[s] network address persistent port translation (NAPPT)” by “storing translation information in a data structure, wherein the translation information comprises” (1) “the private source IP address of the IP telephone,” (2) “a public address of the Service Gateway,” and (3) “the assigned ports for the IP telephone.” EX1024, 32:48-57. If the Board concludes that Fangman teaches the individual-port table above rather than the port-range table, then at the very least a POSITA would have known that the individual-port implementation could be condensed into the port-range implementation, and a POSITA would have found it obvious to use the

⁵ If implementing a port mask instead of a port range, then this cell would contain the appropriate port mask.

port-range implementation over the individual-port implementation because it saves resources, like memory. This is especially true here, where Fangman already teaches using port ranges to identify internal hosts. *See, e.g.*, EX1024, (“The Media Gateway Controller may be operable to use the public IP address *and the port range* to determine the private IP address, for example, by using look-up tables.” (emphasis added)).


24. In sum, Fangman's NAPPT uses port ranges because, as POSITAs universally recognized, traditional NAT is incompatible with IP telephones and Fangman explicitly disavows traditional NAT in favor of port ranges.

IV. GROUNDS 3 AND 5: A POSITA WOULD HAVE FOUND IT OBVIOUS TO APPLY KING'S OUTBOUND TEACHINGS TO INBOUND TRANSMISSIONS

25. A POSITA would have understood that King's outbound disclosures can easily be applied to routing inbound transmissions and that King's disclosures actually teach toward that idea. As Patent Owner recognizes, King teaches using port ranges to identify devices in a private network for outbound transmissions:

King describes methods for an “aggregation device” to obtain user policy information from a management device, such as an authorization, authentication, and accounting (“AAA”) server. *Supra* § II.D.4. A key improvement is using a table of source address and source port assignments accessible by both the customer premises equipment (“CPE”) and the aggregation device. EX2014 ¶ 95. By knowing the source port range assigned to each device by the CPE, the aggregation device can identify a user device and use that identification to obtain policy information from an AAA server. *Id.*

The system described by King is illustrated in Figure 1:



EX1025, Fig. 1; EX2014 ¶ 96. Figure 1 shows a packet flowing from a first user device 100 to a CPE device 108, then an aggregation device 110, and onward to remote devices 102 and/or 116. EX1025, 1:34-59; EX2014 ¶ 96. When aggregation device 110 receives the packet, aggregation device 110 must know what policy to apply to the packet, such as a QoS policy or parental control policy. EX1025, 1:28- 33; EX2014 ¶ 96. Because aggregation device 110 includes a table with source port ranges assigned by CPE device 108, it can use the source port to identify the source device and obtain a policy for that device from AAA server 114. EX1025, 5:63-65, Figs. 6, 7; EX2014 ¶ 96. The packet can then be transmitted using that policy. EX1025, 5:51-57; EX2014 ¶ 96.

POR 29-30.

26. A POSITA would have understood that applying this exact process to inbound data packets would be trivial. Instead of comparing the source port of a received data packet to the stored port range, it would have been obvious to just compare the destination port. In terms of code, this amounts to simply changing one variable in a function: “received_source_port” to “received_destination_port.” A POSITA would have understood, for example, that a SQL query implementing King’s invention would simply change this

```
SELECT    PrivateAddress
FROM      TranslationTable
        WHERE    PublicAddress = received_public_address
        AND (PortMask & received_source_port) = 0;
```

to this

```
SELECT    PrivateAddress
FROM      TranslationTable
        WHERE    PublicAddress = received_public_address
        AND (PortMask & received_destination_port) = 0;
```

in order to perform the disclosed functionality on inbound data packets.

27. King teaches toward this very idea. The background section of King discusses Port Address Translation (“PAT”), which is a process for routing both outbound *and inbound* data packets. EX1025, 2:62-3:6. King explains that the problem with PAT, however, is that “it is not possible for a provider edge device such as an aggregation device to classify incoming packets as belonging to a particular host or user device without additional information being provided.” EX1025, 3:47-50. King therefore proposes supplementing or replacing PAT with its invention. EX1025, 3:67-4:2, 5:33-35. A POSITA would have known that, to replace or supplement PAT, King’s invention would be performed on both outbound and inbound data packets. King recognizes as much, when it explains that “[t]he aggregation device obtains an applies a policy to communication *between* the user device and the Internet”—not just *from* the user device *to* the Internet—to ensure “that an appropriate quality of service (QoS) such as bandwidth is applied and that

an appropriate billing method is applied” and/or “to exert parental control on content”—which are policies that would need to be applied to both outbound and inbound transmissions. EX1025, 1:21-33 (emphasis added), 3:53-59.

28. Further, a POSITA would have recognized that aggregation device **110** must know where to send inbound data packets, especially when servicing “multiple CPEs” that “reuse [] a common IP address.” EX1025, 2:29-34; *see also* EX1025, 1:25-28 (“One or more CPE’s communicate with a gateway to the network such as an aggregation device.”). A POSITA would have understood that aggregation device **110** could use NAT or PAT to route these inbound data packets—as in the prior-art method described in King, (EX1025, 2:13-39)—but a POSITA would have known that implementing NAT or PAT would needlessly waste resources—for example, memory—because the aggregation device **110** can already identify internal devices by comparing the received port number to the stored port-range assignments. A POSITA would have therefore known to use King’s port-range disclosure to route inbound data packets.

29. Patent Owner does not appear to explain how aggregation device **110** could route an inbound data packet to a user device without using the user device’s private address, *i.e.*, the identifier. The global address is not sufficient, especially if the aggregation device is servicing multiple CPEs and user devices, because the global address is not unique to any one of these devices. That would be like trying

to send a letter to a specific person—*e.g.*, “John Smith”—by addressing the letter to a general entity that is not unique to that person—*e.g.*, “Connecticut” or “Corporate Headquarters.” The aggregation device **110** would need to use something more specific, and a readily available, convenient and conventional piece of more specific information is the user device’s identifier.

V. CONCLUSION

30. For the reasons stated above, it is my opinion that the Challenged Claims would have been obvious at the time of the invention.

31. In signing this declaration, I recognize that the declaration will be filed as evidence in a contested case before the Patent Trial and Appeal Board of the United States Patent and Trademark Office. I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true. I further declare that these statements were made with the knowledge that willful false statements and the like are punishable by fine or imprisonment, or both, under 18 U.S.C. § 1001.

Respectfully submitted,

Date: October 22, 2024



Samrat Bhattacharjee, Ph.D