Federal Court Cour fédérale

Date: 20210331

Docket: T-448-17

Citation: 2021 FC 276

Ottawa, Ontario, March 31, 2021

PRESENT: Mr. Justice McHaffie

BETWEEN:

**GUEST TEK INTERACTIVE
ENTERTAINMENT LTD.**

**Plaintiff
Defendant by Counterclaim**

and

NOMADIX, INC.

**Defendant
Plaintiff by Counterclaim**

PUBLIC JUDGMENT AND REASONS

(Confidential Judgment and Reasons issued March 31, 2021)

TABLE OF CONTENTS

I. Overview

[1] Allowing multiple users access to a computer network, such as when guests and the public access a hotel’s network, raises a number of technical issues. This patent infringement action has to do with two of those issues: network security and bandwidth management. Guest Tek Interactive Entertainment Ltd claims its competitor, Nomadix, Inc, is infringing or inducing infringement of two of its patents by selling access gateway devices and licensing the software they run on.

[2] The first, Canadian Patent No 2,600,760 (the ’760 Patent), relates to a way of addressing security concerns when wireless devices, such as wireless laptops or phones, access a network via a wireless access node. Guest Tek alleges Nomadix has induced hotels to infringe certain claims of the ’760 Patent by offering its gateways and software to hotels, and providing instructions on how to use the gateways in a network in a way that will infringe the claims.

[3] The second, Canadian Patent No 2,750,345 (the ’345 Patent), relates to a way of allocating bandwidth between “zones” of users in accordance with the “user load” of those zones. Guest Tek alleges Nomadix itself infringes certain claims of the ’345 Patent by providing license keys to its software to Canadian purchasers, including hotels in Canada. It also alleges Nomadix induces the hotels to infringe the patent, again by offering its gateways and software, and providing instructions to use them in a way that infringes the identified claims.

[4] Nomadix, for its part, denies it infringes either patent. It also alleges, in defence and by counterclaim, that the asserted claims of the ’760 Patent and the ’345 Patent are invalid because they are anticipated or obvious in light of the prior art.

[5] The action was bifurcated. During this liability phase of the proceeding, the most significant issues pertained to the construction of the ’760 and ’345 Patents. Guest Tek and Nomadix proposed different constructions of various terms in the relevant claims of the patents, each supported by expert evidence and argument. On Guest Tek’s construction, Nomadix’s devices and software infringe the patents; on Nomadix’s construction, they do not.

[6] I conclude a person of ordinary skill in the art (POSITA) reviewing the patents in light of the common general knowledge (CGK) would understand some of the terms in the patent claims in accordance with the constructions Guest Tek proposed, and some as Nomadix proposed. On the basis of these constructions and my assessment of Nomadix’s software as informed by the expert evidence, I conclude Nomadix has not infringed or induced infringement of the asserted claims of either the ’760 Patent or the ’345 Patent.

[7] With respect to the ’760 Patent, Guest Tek has not established Nomadix induced infringement of the claims it asserts. Direct infringement by hotels in Canada was not made out, as Guest Tek did not show the existence of a network using a Nomadix gateway in which (a) a wireless access node was configured to receive encrypted wireless protocol traffic and to transmit *all packets* from wireless devices to the gateway; and (b) the gateway was configured to determine whether packets were directed to another *wireless* device on the network and transmit or drop packets accordingly. Nor did Guest Tek establish Nomadix had influenced Canadian

hotels to infringe to the point that, without the influence, any direct infringement would not have taken place.

[8] With respect to the ’345 Patent, Guest Tek has not established Nomadix infringed or induced infringement of the asserted claims. The software run on Nomadix’s gateways allocates bandwidth in a different manner than claimed in the ’345 Patent. It does not use a *quantum,* in the sense of a parameter that limits the amount of data dequeued from a packet queue, that is *dynamically adjusted* in accordance with the *tracked user load* associated with that queue. Even under the specific configurations and usage parameters in which Guest Tek argues the Nomadix gateways exhibit infringing behaviour, it does not have the essential elements of the system and method claimed in the ’345 Patent. There being no direct infringement of the patent through use of the Nomadix gateways, Nomadix cannot have induced any Canadian hotels to infringe.

[9] With respect to the counterclaim, I conclude Nomadix has not shown the asserted claims of either the ’760 Patent or the ’345 Patent to be invalid as having been anticipated or rendered obvious by the prior art. None of the prior art references cited by Nomadix disclose all of the essential elements of the respective claims of the patents. Nor are the differences between the identified prior art and the inventive concepts of the relevant claims steps that would have been obvious to the POSITA at the claim dates of the patents.

[10] The action and counterclaim are both therefore dismissed. If the parties are unable to agree on costs, they may make submissions in accordance with the schedule set out at the end of these reasons.

II. Background

A. *The Trial*

[11] This trial was conducted by videoconference pursuant to the order I issued on August 27, 2020: *Guest Tek Interactive Entertainment Ltd v Nomadix, Inc,* 2020 FC 860. Counsel, witnesses and party representatives attended by videoconference from their respective locations in Alberta, Quebec, Illinois, California, Texas, and Utah. At the outset of each witness' evidence, I instructed them on the protocol for giving evidence by videoconference. No objections were raised during the trial regarding compliance with that protocol or any difficulty arising from the conduct of the hearing by videoconference.

[12] The Court would like to thank the parties, their counsel, and the Court Registry staff, for their efforts and goodwill in conducting the trial by videoconference, which was a novel experience for all involved. Thanks to those efforts, the trial proceeded efficiently and with a minimum of technological or logistical difficulties. Parties were able to take advantage of the technology through displaying documents on screen, playing videos, and referring to electronic documents to which all participants had access.

[13] During the course of the proceeding, the parties entered into a protective agreement regarding the treatment of confidential and proprietary information. At trial, confidential evidence was heard *in camera* with the videoconference locked and closed to members of the public. On informal consent motion at the trial, the parties asked that the terms of their protective agreement be reflected in a protective order and that a confidentiality order be issued pursuant to

Rule 151 of the *Federal Courts Rules*, SOR/98-106. I issued a Protective and Confidentiality Order to this effect on October 30, 2020, which reflected the treatment of confidential information during the trial. Given the potential confidentiality issues, a draft confidential version of these reasons is being released to the parties to ensure any confidential information has been redacted before issuing a public version.

B. *The Parties*

[14] Guest Tek and Nomadix each provide network gateway equipment and supporting software to hotels. A gateway is a computer hardware device that connects computer networks, with information travelling from one network to the other passing through the gateway.

[15] Guest Tek is an Alberta company headquartered in Calgary. It offers its products under the brand name OneView, including its OneView Internet (OVI) solution. One of Guest Tek's employees is the inventor of the '345 Patent, while Guest Tek acquired the '760 Patent as part of a larger transaction from another company, iBAHN.

[16] Nomadix, a Delaware company headquartered in California, also sells gateway devices. Its gateways were originally sold under the name Universal Subscriber Gateway (USG), with models named Access Gateway (AG) and Edge Gateway (EG) being used more recently. As described in Nomadix's Access Gateway User Guide, the gateways enable public access service providers, such as hotels, to offer broadband internet connectivity to their customers:

Exhibit 105, p 3. The various Nomadix gateways run software known as the Nomadix Service Engine (NSE) software. A number of different model numbers are sold in each of the AG and

EG gateway series, and the NSE software has gone through and continues to go through various modifications. At issue in this proceeding are Nomadix models AG2400, AG2500, AG5600, AG5800, AG5900, and EG6000, running the NSE software versions 8.11 or higher (for the ’760 Patent) or 8.7 and higher (for the ’345 Patent). A change to Nomadix’s software shortly before trial in version 8.15.023 was the subject of an evidentiary dispute between the parties, discussed below at paragraphs [412] to [416].

[17] Nomadix does not sell its gateway products directly to end users in Canada. Rather, it has a contractual relationship with two distributors who are authorized to sell Nomadix gateways to the Canadian market: Exhibit 55. The distributors sell Nomadix gateways to resellers, who in turn sell to hotels: Transcript, pp 1264–1266. These distributors sometimes provide technical support on network and gateway configuration in connection with the sale. For the gateway to operate, an end user must register the gateway and enter into a license agreement directly with Nomadix: Transcript, pp 1367–1368, 1379–1380; Exhibit 63. While not in significant factual dispute, the parties took different positions on whether Nomadix’s role in selling gateways, configuring customer networks, and providing licenses and support could or did amount to infringement or inducing infringement. Given my conclusions on the construction of the patents, these issues are ultimately not determinative.

C. *The Witnesses*

[18] The witnesses called by the parties fell into three main categories: inventors, experts, and lay witnesses. I provide a brief summary of their evidence below and some general observations, and will refer to that evidence in greater detail as necessary during the course of these reasons.

(1) Inventors

[19] Guest Tek called the three living inventors of the ’760 Patent, **Brett Molen**, **Nichol Draper**, and **Jan DeHoop**, as well as the inventor of the ’345 Patent, **David Ong**. Each of the inventors provided background information regarding their education and experience, their work in the industry at the time of the inventions, and the context in which the inventions came about.

[20] Mr. Molen, Mr. Draper, and Mr. DeHoop, together with a fourth inventor, Richard Ehlers, were employed by a company Mr. Molen co-founded named STSN when they filed an application for a United States Patent in 2005. The inventors assigned their interest in the application and invention to STSN shortly after the US Patent application was filed: Exhibit 5. STSN subsequently changed its name to iBAHN. In 2013, iBAHN assigned to Guest Tek its interest in the invention and associated patents and applications, which by then included the Canadian application that became the ’760 Patent: Exhibit 10. Nomadix raises no issue with respect to inventorship, the assignments, or Guest Tek’s resulting ownership of the ’760 Patent.

[21] Mr. Ong is an employee of Guest Tek in Calgary, and developed what became the ’345 Patent while at Guest Tek in or around 2011. Again, Nomadix takes no issue regarding inventorship or Guest Tek’s ownership of the ’345 Patent.

[22] Some of the evidence elicited from the inventors by counsel, both in examination in chief and in cross-examination, strayed into areas that amounted to them giving their understanding of

the invention and the meaning of the terms used in the patents. Inventors' evidence regarding their understanding of the terms in a patent is not admissible for the purpose of construing the claims of the patent: *Nekoosa Packaging Corp v AMCA International Ltd*, [1994] FCJ No 1046, 56 CPR (3d) 470 (CA) at para 23; *Free World Trust v Électro Santé Inc*, 2000 SCC 66 at paras 61–66; *Bombardier Recreational Products Inc v Arctic Cat, Inc*, 2018 FCA 172 at paras 22–23, 51. I have not referred to this evidence in my construction of the patents herein.

(2) Experts

[23] The Court had the assistance of three experienced computer scientists called by the parties. Their various expert reports were filed on consent, and no objections were made to their respective qualifications.

[24] Guest Tek called two expert witnesses, one on each patent. **Dr. Peter Reiher** addressed the '760 Patent. Dr. Reiher is an Adjunct Professor in the Computer Science Department at the University of California, Los Angeles. He has taught courses at UCLA in computer science, including computer security and network security. He is an author on a large number of texts and papers, many of which deal with network security, one of his research interests. Dr. Reiher was tendered and accepted as qualified to give expert evidence in the areas of computer and computer network security, including wired and wireless networks; packet transmission and security issues in relation thereto; and security for mobile devices in wireless and wired networks. Dr. Reiher authored an initial report on construction and infringement [Reiher First Report]; a rebuttal report on validity and responding to construction issues [Reiher Second Report]; and a sur-reply report dealing with one non-infringement and construction issue [Reiher Third Report].

[25] **Dr. Peter Dordal** addressed the ’345 Patent. Dr. Dordal is an Associate Professor in the Computer Science Department at Loyola University Chicago, where he has taught courses in computer programming and computer networking. He is the author of *An Introduction to Computer Networks*, an online textbook published since 2012, and has contributed to other publications on computer networks. He was also System Administrator for departmental computing facilities at Loyola for almost 20 years, a position that involved management of university networks and software development. Dr. Dordal was tendered and accepted as qualified to give expert evidence in the areas of computer systems and programming of computers; and computer internet access equipment and management of bandwidth. Dr. Dordal filed an initial report on construction and infringement [Dordal First Report] and a rebuttal report on validity and responding to construction issues [Dordal Second Report].

[26] Nomadix called **Dr. Tal Lavian** to address both the ’760 and ’345 Patents. Dr. Lavian earned his Ph.D. in Computer Science from the University of California, Berkeley. He has taught as a Lecturer at UC Berkeley and acted as Visiting Scientist and as Industry Fellow at that university, teaching classes on wireless devices and smartphones. Dr. Lavian has worked in telecommunications, wireless, and networking technologies for over 30 years, including over a decade as Principal Scientist at his own network communications company. He is a named co-inventor on more than 100 issued patents. Dr. Lavian was tendered and accepted as qualified to give expert evidence in network communications; computer programming; mobile and wireless communications; computer networks; internet protocols; packet switching; and network design and architecture, including switches, access devices, edge devices, and gateways. Dr. Lavian filed a total of four reports, namely a first report on construction and validity in respect of each patent [Lavian First ’760 Report and Lavian First ’345 Report], and a second report on

infringement and responding construction issues in respect of each patent [Lavian Second ’760 Report and Lavian Second ’345 Report].

[27] The reports and oral evidence of the experts gave information regarding the relevant fields of endeavour, what was known to those working in those fields at the time, and their views on how the patents would be understood by a POSITA. Their evidence allowed the Court to put itself in the position of the POSITA for the purposes of construing the patents and assessing the parties’ respective arguments on infringement and invalidity: *Whirlpool Corp v Camco Inc*, 2000 SCC 67 at para 57.

[28] For the most part, I found all three experts were helpful and intended to fulfill their role of assisting the Court impartially on matters relevant to their expertise. At the same time, each expert on occasion put forward constructions and justifications that strained impartiality and moved towards advocacy. I refer to these elements in further detail as necessary in the course of these reasons.

[29] I also make the following general comments regarding the evidence of the three experts. It is now clearly recognized that counsel may assist in the preparation of an expert report, and that counsel’s involvement can even be beneficial in ensuring reports are framed in a way that is comprehensible and relevant: *Biogen Canada Inc v Taro Pharmaceuticals Inc*, 2020 FC 621 at para 71; *Moore v Getahun*, 2015 ONCA 55 at paras 55–64. In my assessment, Guest Tek’s reports suffered from too much involvement of counsel, while Nomadix’s might have benefited from somewhat more.

[30] It was clear—even before this was confirmed in cross-examination—that Dr. Reiher and Dr. Dordal's reports had significant input from counsel for Guest Tek: Transcript, pp 580–581, 584–590, 1022–1023, 1076–1077, 1080–1081, 1110–1111. For example, they each opined at some length, in similar or identical language, on factual issues unrelated to computer science, such as the extent to which Nomadix could be considered to have influenced Canadian hotels to infringe the patent through their user documentation and licenses: Reiher First Report, paras 493–524; Dordal First Report, paras 434–474. Dr. Reiher and Dr. Dordal's rebuttal reports also included very similar, and often argumentative, responses to Dr. Lavian's discussion of the CGK: Reiher Second Report, paras 16–19, 24–58; Dordal Second Report, paras 20–23, 30–40. They also each criticized Dr. Lavian for not following legal principles, again in argumentative and often similar or identical language: Reiher Second Report, paras 7–10, 60, 123, 262–265; Dordal Second Report, paras 6–13, 150–154. While the involvement of counsel in the preparation of reports is by no means objectionable in itself, as stated, Dr. Reiher and Dr. Dordal's reports suffered from the inclusion of what was effectively legal argument.

[31] Dr. Lavian's reports, on the other hand, might have benefited from greater guidance, as their structure and the issues they addressed were not always helpfully focused on the material issues in dispute. Dr. Lavian's discussions of the CGK, while thorough and generally useful, often veered into areas less relevant to the patents and issues in dispute. His analysis of validity issues tended to be broadly couched and less focused on the specific claims and the relevant analysis, and would have been of more assistance had it been more focused on the issues relevant to the obviousness assessment: Lavian First '760 Report, paras 10.1–10.6; Lavian First '345 Report, paras 9.1–9.7.

[32] Whether due to too much involvement of counsel or too little, each of the experts also undermined the strength of their own opinions through the nature of their criticisms of the other. For example, Dr. Reiher purported to give his opinion that Dr. Lavian "completely and utterly failed to follow" legal principles of claims construction (language from counsel that Dr. Reiher regretted), while Dr. Dordal argued Dr. Lavian "completely failed to follow" the steps in the obviousness analysis: Reiher Second Report, para 60; Transcript, pp 584–585; Dordal Second Report, para 150. Dr. Lavian criticized Dr. Reiher for, among other things, "struggl[ing] to build […] a lawyer's argument", and suggested Dr. Dordal's basic premise was "unreasonable and biased": Lavian Second '760 Report, para 122; Lavian Second '345 Report, para 21. While rigorous debate between experts should not be shied away from, and may even be beneficial, the nature of each expert's criticism of the other, particularly in their reports, was unhelpful.

[33] One surprising issue that arose out of Dr. Lavian's reports and oral evidence is also worth addressing as a preliminary matter. Dr. Lavian's First Report on each of the patents was entitled "Claim Construction and Validity." Each report said his mandate included providing his opinion on how a POSITA "would understand various words and phrases in the patents mentioned above and help the Court construe the claims": Lavian First '760 Report and Lavian First '345 Report, para 2.1.3. As is common, Dr. Lavian attached to each a list of legal principles that counsel provided for him to follow in his report, including principles of claims construction. In each, Dr. Lavian included a section entitled "Claim Construction" that provided a description of claim elements, his comments, and his opinion on whether the element was essential: Lavian First '760 Report, paras 3.2, 7.1–7.10 and Appendix TL-04; Lavian First '345 Report, paras 3.2, 8.1 and Appendix TL-03. Nonetheless, in his oral testimony, Dr. Lavian insisted on a number of

occasions that he had not undertaken a claims construction and “didn’t construe any claim”: Transcript, pp 1615–1616, 1847–1862, 1874–1878, 1881, 1945–1953. Rather, he said, he simply adopted the ordinary meaning of the terms and provided comments. His adamance that he had not construed the claims led to a certain inefficiency in cross-examining him on his “Claim Construction and Validity” reports.

[34] Leaving aside whether adopting the “ordinary meaning” is itself construing the claims, it is clear that Dr. Lavian’s reports provided a construction of the claims, or at least his understanding of how a POSITA would understand the claims and whether they would be considered essential. It appears, and I am prepared to accept, that the matter may simply be one of terminology, *i.e.*, that Dr. Lavian believed “construction” was a term of art that did not describe what he was doing, notwithstanding the use of the term in his report, or that he was perhaps concerned that construction is ultimately the Court’s domain: *Whirlpool* at para 57. In any event, I have taken Dr. Lavian’s evidence regarding the CGK and how a POSITA would understand the patents for the information it provides without significant concern as to whether he engaged in “claims construction” or not.

[35] The foregoing comments having been made, I reiterate that each expert ultimately appeared desirous of providing their expertise to help the Court understand the patents and the underlying technologies. This was particularly clear during their oral testimony. I have adopted the approach of accepting their technical expertise and opinions on this basis, considering their evidence where helpful and discounting that which was less helpful, without making any overall conclusion that one party’s expert was to be preferred over the other as a blanket matter.

(3) Lay Witnesses

[36] Guest Tek called three lay witnesses. **Arnon Levy** is the founder of Guest Tek. Formerly President and CEO of the company, he now sits as Chair of the Board of Directors and Founder. He gave an overview of Guest Tek's history and business, as well as the acquisition of the ’760 Patent and other intellectual property from iBAHN as part of a larger purchase agreement. **Zoey Sachdeva** was recently named Chief Information Officer (CIO) of Guest Tek, having worked at the company for about 13 of the last 16 years. She gave evidence primarily regarding brand standards adopted by hotels in respect of their networks, and in particular Marriott Hotels' Global Property Networking Standards (GPNS). **Kenneth Barnes** is the Chief Technology Officer (CTO) of Progress Residential, and was briefly a Guest Tek employee in the mid-2010s. He gave evidence primarily regarding his experiences while working as CIO of Omni Hotels & Resorts, a role that included responsibility for designing and assisting in developing guest high speed internet access (HSIA) standards for Omni Hotels.

[37] Nomadix called two lay witnesses, both Nomadix employees since 1999. **Jeremy Cook** is currently Product Manager of Nomadix Cloud. He gave evidence on the development and functionality of Nomadix's gateway devices and the NSE software, as well as its commercial sales and support structures. This included confidential evidence regarding Nomadix's license tracking system. **Vadim Olshansky** is currently the CTO of Edge Gateways for the company. He gave evidence on the development, architecture, and functionality of Nomadix's NSE software. His evidence included confidential descriptions and explanations of the source code for Nomadix's software, showing how the software achieves certain functionalities.

III. General Principles

[38] There was little dispute between the parties regarding the legal principles applicable to the issues of patent construction, infringement, and validity. As these principles are applicable to each of the patents at issue, I will set them out here before turning to their application.

A. *Claims Construction*

[39] The monopoly protected by a patent is defined by the claims: *Patent Act*, RSC 1985, c P-4, s 27(4). The appropriate approach to interpreting or construing those claims was recently reiterated and summarized by the Federal Court of Appeal in *Tearlab Corporation v I-MED Pharma Inc*, 2019 FCA 179, making reference to the Supreme Court of Canada's decisions in *Whirlpool*; *Free World Trust*; *Consolboard Inc v MacMillan Bloedel (Sask) Ltd*, [1981] 1 SCR 504; and *Teva Canada Ltd v Pfizer Canada Inc*, 2012 SCC 60 [*Teva*]. In *Tearlab*, Justice de Montigny for the Court of Appeal set out the main principles in the following terms:

> [31] The *Patent Act* promotes adherence to the language of the claims, which in turn promotes fairness and predictability. The words of the claims must, however, be read in an informed and purposive way, with a mind willing to understand. On a purposive construction, it will be apparent that some elements of the claimed invention are essential while others are non-essential. The interpretative task of the court, in claim construction, is to separate and distinguish between the essential and the non-essential elements, and to give the legal protection to which the holder of a valid patent is entitled only to the essential elements.
>
> [32] To identify these elements, the claim language must be read through the eyes of a POSITA, in light of the latter's common general knowledge. As noted in *Free World Trust*:
>
>> [51] …The words chosen by the inventor will be read in the sense the inventor is presumed to have intended, and in a way that is sympathetic to accomplishment of the inventor's purpose expressed or implicit in the text of the claims. However, if the inventor has misspoken or

> > otherwise created an unnecessary or troublesome limitation in the claims, it is a self-inflicted wound. The public is entitled to rely on the words used *provided* the words used are interpreted fairly and knowledgeably. [Emphasis in the original.]
>
> [33] Claim construction requires that the disclosure and the claims be looked at as a whole "to ascertain the nature of the invention and methods of its performance, … being neither benevolent nor harsh, but rather seeking a construction which is reasonable and fair to both patentee and public". Consideration can thus be given to the patent specifications to understand what was meant by the words in the claims. One must be wary, however, not to use these so as "to enlarge or contract the scope of the claim as written and … understood". The Supreme Court recently emphasized that the focus of the validity analysis will be on the claims; specifications will be relevant where there is ambiguity in the claims.
>
> [34] Finally, it is important to stress that claim construction must be the same for the purpose of validity and for the purpose of infringement.
>
> [Emphasis added; citations omitted.]

[40] As is often the case, there was considerably less dispute in this matter over the essentiality of the claim elements than over the interpretation of the essential elements. Unless a party maintains that a claim element is not essential, it will be considered essential: *Corlac Inc v Weatherford Canada Inc*, 2011 FCA 228 at para 26.

(1) Recourse to the disclosure in claims construction

[41] Paragraph 33 from *Tearlab*, reproduced above, prompts some further discussion, as the parties and experts each referred frequently to passages in the disclosure in support of their arguments on construction. Justice de Montigny reaffirms that claims construction *requires* that the disclosure and the claims must be "looked at as a whole." In support, he quotes the oft-cited

principle from *Consolboard* that "[w]e must look to the whole of the disclosure and the claims to ascertain the nature of the invention" and a reasonable and fair construction [emphasis added]: *Consolboard* at p 520, cited with approval in *Teva* at para 50; in *Monsanto Canada Inc v Schmeiser*, 2004 SCC 34 at para 18; in *Whirlpool* at para 49(g); and recently by the Court of Appeal in *Western Oilfield Equipment Rentals Ltd v M-I LLC*, 2021 FCA 24 at para 16.

[42] However, Justice de Montigny also cited the caution from *Whirlpool* that a Court may look to the disclosure and drawings to "understand" the words of a claim but not to "enlarge or contract the scope of the claim as written": *Tearlab* at para 33, citing *Whirlpool* at para 52. Of course, any construction given to the words in a claim will affect the scope of the claim: *Whirlpool* at para 49(h). I therefore take the rule against using the disclosure to "enlarge or contract" the claim as written to preclude adding words, elements, or limitations not found in the claim, or giving the words a meaning they cannot reasonably bear when interpreted in the context of the patent as a whole.

[43] Justice de Montigny also alludes to the principle that the specification will be relevant where there is "ambiguity" in the claim, citing *AstraZeneca Canada Inc v Apotex Inc*, 2017 SCC 36 at para 31. The Federal Court of Appeal in its 2016 decision in *Tadalafil* adopted the approach that "recourse" to the rest of the specification is unnecessary where the words in the claim are "plain and unambiguous": *Mylan Pharmaceuticals ULC v Eli Lilly Canada Inc*, 2016 FCA 119 at para 39 [*Tadalafil*], quoting The Hon RT Hughes, D Clarizio & N Armstrong, *Hughes and Woodley on Patents* (Toronto: LexisNexis Butterworths, 2005) (loose-leaf 2d ed) at p 312. The Court of Appeal concluded the trial judge in that case erred in even referring to the

disclosure when construing the claims, given the lack of ambiguity in the relevant claim, saying such reference was precluded:

> [39] In my view, the judge erred in referring to the specification when construing the claims of the ‘377 patent. The rules of patent construction preclude reference to the specification when the claims are clear, and also improper if it varies the scope of the claims: *Hughes and Woodley on Patents*, p. 312:
>
> > In construing a patent, the claims are the starting point. The claims alone define the statutory monopoly and the Patentee has a statutory duty to state, in the claims, what [t]he invention is for which protection is sought. In construing the claims, recourse to the rest of the specifications is (1) permissible to assist in understanding the terms used in the claims; (2) unnecessary where the words [are] plain and unambiguous and (3) improper to vary the scope or ambit of the claims.
> >
> > [Emphasis added.]

The final sentence of the foregoing passage from *Hughes and Woodley on Patents* is itself simply a restatement of the Court of Appeal’s earlier decision in *Dableh v Ontario Hydro*, [1996] 3 FC 751 (CA) at para 30. This passage from *Tadalafil* was referenced by the Court of Appeal recently without adverse comment about its applicability: *Hospira Healthcare Corp v Kennedy Trust for Rheumatology Research*, 2020 FCA 30 at paras 21–22.

[44] A similar approach was taken by the Court of Appeal shortly after *Tearlab* in *Tetra Tech EBA Inc v Georgetown Rail Equipment Company*, 2019 FCA 203. There, the Court of Appeal held that a Court may look to the whole of the specification to understand a claim term or “confirm” a construction arrived at upon consideration of the claims, but not to enlarge or contract the scope of the claim as written: *Tetra Tech* at paras 86, 103–104. The Court of Appeal criticized a construction by this Court that relied on expert testimony and a passage in the

disclosure, and not the language of the claim, stating that recourse to the disclosure should be undertaken only in cases of ambiguity: *Tetra Tech* at paras 91, 100–103.

[45] There is in my view some tension between the principle that claims construction *requires* the disclosure and the claims be looked at as a whole, and the notion that "recourse" to the disclosure is *only* permissible when the claims are ambiguous. The purpose of beginning the construction exercise with the disclosure, and requiring consideration of the disclosure and the claims as a whole, is presumably to recognize that the disclosure assists and influences the purposive understanding of the claim terms in their context: *Consolboard* at p 520; *Whirlpool* at paras 49(c)–(g), 52; *Tearlab* at para 33. If this is so, it is difficult to reconcile with the requirement that the POSITA—and the parties, experts, and Court—cannot refer (or have "recourse") to those very parts of the disclosure that assist in understanding the claim terms unless they first conclude the terms are ambiguous. As Justice Binnie recognized in *Whirlpool*, determining whether a term is ambiguous may itself require "careful review" of the disclosure: *Whirlpool* at para 52. I note that in *Monsanto*, the Supreme Court engaged in claims construction with reference to the disclosure without any prior determination that the claims were ambiguous: *Monsanto* at paras 17–19.

[46] Indeed, terms used in a claim will often be expressly discussed or defined in the disclosure, in keeping with the principle that the patentee may "act as his own lexicographer": *Kramer v Lawn Furniture Inc*, [1974] FCJ No 100, 13 CPR (2d) 231 at para 16. The rule against recourse to the disclosure presumably does not require the POSITA to find an ambiguity in a claim term before having recourse to the patentee's definition of the term in the disclosure, or the patentee's ability to define their lexicon would be significantly undermined.

[47] Ultimately, I conclude from the foregoing that the exercise of construction must consider both the disclosure and the claims, with the claims being purposively construed in the context of the patent as a whole and in light of the CGK of the POSITA. However, the focus remains on the language of the claims, which defines the scope of the monopoly. The disclosure should not be used to enlarge or contract the scope of the claims, particularly through the addition of words or limitations not found in the claims.

[48] As discussed below, in this case I conclude there are no terms where consideration of the disclosure suggests a construction different than that suggested by the face of the claim read in light of the patent as a whole. In two cases—the term *communications between* in the ’760 Patent, discussed at paragraphs [174] to [189] below, and the term *in proportion to* in the ’345 Patent, discussed at paragraphs [344] to [355]—reference to the disclosure is necessary to resolve an ambiguity in the claim itself. In all other cases, referring to the disclosure simply “confirms” the construction arrived at in considering the claims: *Tetra Tech* at para 86.

(2) Additional principles of claims construction

[49] Two further principles of claims construction are relevant in this proceeding, given the parties’ positions and some of the expert evidence. The first is that claims construction is “antecedent” to consideration of validity and infringement issues: *Whirlpool* at para 43. In other words, it is to be undertaken with an eye to understanding the patent as written, rather than with the goal of obtaining, or avoiding, a finding of infringement or invalidity. At the same time, it is appropriate for the Court to be aware of the issues on which the parties disagree—where the “shoe pinches”—so as to focus the construction exercise on relevant issues: *Cobalt Pharmaceuticals Company v Bayer Inc*, 2015 FCA 116 at para 83.

[50] The second principle involves a balance between two evocative terms well known in patent law: the Court should construe a patent with “judicial anxiety” to support a useful invention, but if the inventor has created troublesome limitations in the claims, this is a “self-inflicted wound” the Court will not cure: *Consolboard* at p 521, quoting *Hinks & Son v Safety Lighting Company* (1876), 4 Ch D 607 at p 612; *Free World Trust* at para 51. Again, these two principles balance the central importance of the words of the claim as drafted by the inventor with the need to read those words in the context of the patent as a whole and with a mind willing to understand.

B. *Prior Art and Common General Knowledge*

[51] Claims construction is undertaken through the eyes of the POSITA in light of the CGK: *Tearlab* at para 32; *Free World Trust* at paras 31(e)(i), 44, 51; *Whirlpool* at para 53. The Supreme Court of Canada has defined the CGK as “knowledge generally known by persons skilled in the relevant art at the relevant time”: *Apotex Inc v Sanofi-Synthelabo Canada Inc*, 2008 SCC 61 at para 37(2). *Sanofi-Synthelabo* addressed the CGK in the context of the test for anticipation, but the same definition applies to the CGK for construction purposes, although the relevant date of analysis may differ: *Tadalafil* at paras 23–25. The relevant date for purposes of construction of a patent is the date of publication, while the relevant date for purposes of invalidity is the claim date, which is the priority date if there is one, or the filing date if not: *Tadalafil* at para 31; *Whirlpool* at para 55; *Patent Act*, ss 28.1, 28.2, 28.3.

[52] The CGK has been described as a “subset” of the state of the art at the relevant time: *Hospira* at para 84. A piece of information is part of the CGK if “a skilled person would become

aware of it and accept it as 'a good basis for further action'": *Tadalafil* at para 24. Nonetheless, the CGK is not limited to the knowledge a POSITA would have memorized or know offhand. Rather it includes what the person "may reasonably be expected to know <u>and to be able to find out</u>" [emphasis added]: *Tetra Tech* at para 28. The CGK may include the information presented as background knowledge in the patent itself: *Valeant Canada LP/Valean Canada SEC v Generic Partners Canada Inc*, 2019 FC 253 at para 47. However, it is not limited to this information.

C. *Infringement*

[53] As noted above, part of the construction exercise is determining which elements of a claim are essential and which are non-essential. Infringement of a claim occurs only when all of the essential elements of the claim as purposively construed are present: *Free World Trust* at paras 31(f), 68, 75; *Western Oilfield (FCA)* at paras 48–49.

[54] The question of whether a party has "used" an invention in violation of the exclusive right granted by section 42 of the *Patent Act* is also influenced by the rule of purposive construction: *Monsanto* at para 33. Any act that interferes with the "full enjoyment of the monopoly granted to the patentee" by virtue of section 42 is prohibited: *Monsanto* at paras 34, 58. Nonetheless, while this may assist in assessing whether a defendant has "used" an invention as defined in the claims, it effectively brings the analysis back to the claims, as that is what defines the scope of the monopoly the patentee enjoys: *Whirlpool* at paras 18, 63. The issue thus remains whether each of the essential elements of the asserted claims are present, rather than just a broader analysis based on "full enjoyment" of the exclusive right to the invention.

[55] In terms of where the alleged infringement occurs, Nomadix pleaded that none of its acts, including its sale of gateways, occurred in Canada so they could not infringe the patents at issue. However, Nomadix did not pursue this jurisdictional argument at trial, focusing instead on whether there was evidence of infringement of the patents in Canada by any party. While I need not decide the issue, it seems to me that where software use is alleged to infringe a Canadian patent, it would be difficult for a foreign defendant to avoid liability on the basis of being outside Canada in circumstances where they directly license the use of the software in Canada via a cross-border transaction occurring over the internet.

D. *Inducing Infringement*

[56] Inducing infringement is simply a form of patent infringement rather than a distinct tort: *Hospira* at para 45; *Western Oilfield (FCA)* at para 60. The parties agree that allegations of inducing infringement are governed by the three-part test adopted in *Warner Lambert Co v Wilkinson Sword Canada Inc*, [1988] FCJ No 70, 19 CPR (3d) 402 (FCTD) and reiterated by the Federal Court of Appeal in *Corlac* at para 162:

> It is settled law that one who induces or procures another to infringe a patent is guilty of infringement of the patent. A determination of inducement requires the application of a three-prong test. <u>First</u>, the act of infringement must have been completed by the direct infringer. <u>Second</u>, the completion of the acts of infringement must be influenced by the acts of the alleged inducer to the point that, without the influence, direct infringement would not take place. <u>Third</u>, the influence must knowingly be exercised by the inducer, that is, the inducer knows that this influence will result in the completion of the act of infringement.
>
> [Emphasis added; citations omitted.]

[57] With respect to the first component of the test, “[d]irect infringement occurs when the direct infringer has performed all of the essential steps in the claimed invention”: *Western Oilfield (FCA)* at para 70. This does not necessarily require evidence coming directly from the direct infringer, but there must be evidence from which the Court can conclude on a balance of probabilities that direct infringement has occurred: *Western Oilfield Equipment Rentals Ltd v M-I LLC*, 2019 FC 1606 at paras 126, 129, aff’d *Western Oilfield (FCA)* at paras 67–68.

[58] Guest Tek argued the second requirement creates a “but for” test that asks whether the infringing conduct would have occurred but for the defendant’s conduct: *Western Oilfield (FC)* at paras 127, 130, aff’d *Western Oilfield (FCA)* at para 70. I agree the “without the influence” aspect of the second element of the *Warner Lambert/Corlac* test creates a “but for” test. But the test is whether the infringement would have occurred but for the *defendant’s influence*, and not simply but for the *defendant’s sale of a product* used by the direct infringer in the course of infringement. Again, proof of influence need not involve direct evidence from customers that they were induced to infringe by instructions given by the inducer, if this can be inferred from the inducer’s and the inducee’s conduct: *Western Oilfield (FC)* at paras 126, 130–131, aff’d *Western Oilfield (FCA)* at paras 67–69.

[59] Similarly, with respect to the knowledge component in the third element of the test, as Justice O’Reilly stated in *Western Oilfield (FC)*, “the alleged inducer simply has to know what the third party is likely to do in response to its influence”: *Western Oilfield (FC)* at para 133. The issue is not simply knowing what the third party is likely to do. It is knowing what the third party is likely to do *in response to the defendant’s influence*.

[60] Guest Tek drew the Court’s attention to a thoughtful article on liability for “indirect” patent infringement, *i.e.*, liability of those who do not themselves undertake all essential elements of a claim: N Siebrasse, “Contributory Infringement in Canadian Law” (2020) 35 Cdn Int Prop Rev 10. In the article, Professor Siebrasse considers both the current state of the law and policy considerations associated with liability for inducing infringement and contributory infringement. After discussing English and Canadian cases that have addressed indirect infringement—including *Slater Steel Industries Ltd v R Payer Co* (1968), 55 CPR 61 (Ex Ct), which he argues was wrongly decided, and *Windsurfing International Inc v Trilantic Corp* (1985), 8 CPR (3d) 241, [1985] FCJ No 1147 (CA)—he concludes a distinction should be drawn between the supply of products that have a substantial non-infringing use and those that have no non-infringing use or are especially adapted to infringe. In the latter case, he suggests that sale of such items should ground liability for indirect infringement: Siebrasse at p 26.

[61] While the discussion in Professor Siebrasse’s article is of interest, I agree with Nomadix that the Federal Court of Appeal has confirmed “contributory infringement” is not recognized as a cause of action in Canadian law: *Nycomed Canada Inc v Teva Canada Limited*, 2012 FCA 195 at para 3. In any event, I need not decide in this case whether the law of patent infringement ought to recognize the principles discussed by Professor Siebrasse. I conclude below Guest Tek has not shown that a third party in Canada directly infringed its patents. The questions of whether there are substantial non-infringing uses for the Nomadix gateways or whether they are especially adapted to infringe therefore become irrelevant or inapplicable. Further, Guest Tek’s Amended Statement of Claim only alleges Nomadix has induced infringement and does not plead Nomadix is liable for indirect or contributory infringement outside the context of inducing infringement: Amended Statement of Claim, paras 1(b), 12, 65–94.

E. *Validity*

[62] Subsection 43(2) of the *Patent Act* creates a presumption that an issued patent is valid. A party challenging a patent therefore has a burden to prove it is invalid on a balance of probabilities: *Whirlpool* at para 75. Nomadix's challenges in this case are brought on grounds of anticipation and invalidity. Each has a well-defined analytical framework established by the *Patent Act* and the Supreme Court of Canada.

(1) Anticipation

[63] A patent is invalid if it is not new, that is, if the invention it claims has been previously disclosed: *Patent Act*, ss 2("invention"), 28.2; *Eli Lilly Canada Inc v Novopharm Limited*, 2010 FCA 197 at para 43 [*Olanzapine*]. Establishing a prior art reference anticipates a claim involves a two-step inquiry, described at paragraphs 24 to 37 of *Sanofi-Synthelabo* as the requirement for "disclosure" and "enablement":

1) Does the prior art reference disclose subject matter which, if performed, would necessarily result in infringement of the claim?

2) Is the prior art reference sufficiently detailed to enable a POSITA to perform the claimed invention without the exercise of inventive ingenuity or undue experimentation?

[64] While the *Sanofi-Synthelabo* approach was based on language in the "old Act," *i.e.*, the version of the *Patent Act* applicable to applications filed prior to October 1, 1989, it remains applicable to the determination of anticipation under section 28.2 of the current *Patent Act*: *Sanofi-Synthelabo* at paras 15, 18–19; *Olanzapine* at paras 43–45; *Hospira* at para 66.

[65] The requirement that the prior art disclose subject matter which, if performed, would necessarily result in infringement means that a single prior art document must disclose each essential element of the claim to be considered anticipatory: *Eli Lilly Canada Inc v Mylan Pharmaceuticals ULC,* 2015 FC 125 at para 145; *Sanofi-Synthelabo* at para 28; *Free World Trust* at para 26. The disclosure need not be an exact description of the claimed invention, provided the POSITA can understand the disclosure without trial and error: *Sanofi-Synthelabo* at paras 23, 25, 32. At the second stage of assessing enablement, however, the question is not how the POSITA would understand the prior art, but whether they could work the invention. At this stage, the POSITA may use their CGK to supplement information in the piece of prior art and some trial and error is permitted provided it does not rise to an undue burden: *Sanofi-Synthelabo* at paras 27, 33, 37.

(2) Obviousness

[66] A patent is not valid if it is not inventive, that is, if the invention it claims would have been obvious to a POSITA: *Patent Act*, ss 2(“invention”), 28.3. Obviousness is assessed as of the claim date and with regard to information publicly disclosed before that date: *Patent Act*, s 28.3(*b*); *Hospira* at paras 85–86. Different timing rules apply where the disclosure is directly or indirectly from the applicant, an issue not relevant in this matter: *Patent Act*, s 28.3(*a*).

[67] The proper approach to obviousness was set out at paragraph 67 of *Sanofi-Synthelabo*, adopting the four-step inquiry as restated by Lord Justice Jacob in *Pozzoli SPA v BDMO SA*, [2007] EWCA Civ 588 at para 23:

> (1) (a) Identify the notional “person skilled in the art”;
>
> (b) Identify the relevant common general knowledge of that person;
>
> (2) Identify the inventive concept of the claim in question or if that cannot readily be done, construe it;
>
> (3) Identify what, if any, differences exist between the matter cited as forming part of the “state of the art” and the inventive concept of the claim or the claim as construed;
>
> (4) Viewed without any knowledge of the alleged invention as claimed, do those differences constitute steps which would have been obvious to the person skilled in the art or do they require any degree of invention?

[68] The fourth step involves the key inquiry of whether differences between the patent and the prior art would have been obvious. This can involve consideration of whether the solution identified by the patent would have been “obvious to try”: *Sanofi-Synthelabo* at paras 68–69.

[69] The *Sanofi-Synthelabo* test is not to be applied as a rigid rule: *Bristol-Myers Squibb Canada Co v Teva Canada Limited*, 2017 FCA 76 at paras 59–62. Rather, all factors relevant in the circumstances are considered in assessing whether the gap could have been bridged by the POSITA who has knowledge of the CGK but has no “scintilla of inventiveness or imagination”: *Bauer Hockey Ltd v Sport Maska Inc (CCM Hockey)*, 2020 FC 624 at para 144; *Hospira* at para 79. These factors may include the inventor’s course of conduct, the conduct of other industry participants, industry motivation, conventional wisdom, and commercial success: *Bauer Hockey* at paras 145–151; *Sanofi-Synthelabo* at paras 70–71.

IV. Canadian Patent 2,600,760

A. *Introduction*

[70] The ’760 Patent relates to security in wireless networks. Broadly stated, it covers a computer network configuration or architecture designed to improve security by controlling the flow of packets of data through the components of the network, forwarding or dropping packets based on certain criteria. While the different claims of the patent vary in certain elements, they all involve a *network* that includes a *gateway* and a *wireless access node*. The wireless access node receives packets from wireless devices and transmits them to the gateway, while the gateway makes the determinations on forwarding or dropping packets.

[71] The ’760 Patent has two types of claims: “network” claims and “method” claims. The network claims define configurations of the wireless access node and gateway. The method claims include steps for receiving and transmitting packets involving the wireless access node and the gateway. There are two independent claims in the patent. Claim 1 is a network claim from which all of the other network claims directly or indirectly depend. Claim 21 is a method claim from which all of the other method claims depend. Guest Tek asserts Nomadix is inducing infringement of network Claims 1, 4 (as it depends from Claim 1), 10 (as it depends from Claims 1 and 4), and 11 (as it depends from Claims 1 and 4); and method Claims 21, 30 (as it depends from Claim 21), and 31 (as it depends from Claims 21 and 30) of the ’760 Patent.

[72] The ’760 Patent was published on September 21, 2006, the material date for purposes of construction. The priority date, relevant to invalidity issues, is March 10, 2005. It issued on November 1, 2016.

B. *The Person of Ordinary Skill in the Art*

[73] The art to which the ’760 Patent is directed is computer network security, and in particular network security related to wireless devices accessing a network. The parties’ experts had somewhat differing views of the background of the person of ordinary skill in this art, with Dr. Lavian suggesting a slightly higher degree of academic and/or practical experience than Dr. Reiher: Lavian First ’760 Report, para 5.14; Reiher First Report, paras 67–68. Dr. Lavian also described fairly extensive knowledge and experience a POSITA would have, including significant experience designing and installing secure wireless networks: Lavian First ’760 Report, paras 5.15–5.16.

[74] In closing submissions, both parties confirmed their view that these differences had no material impact on the issues. I agree that little turns on the differences in the POSITA’s background as proposed by Drs. Reiher and Lavian. That said, the ’760 Patent deals with implementation of a particular network architecture and configuration of components in a practical setting. It makes reference to implementation of wireless solutions in accordance with standards published by the Institute of Electrical and Electronics Engineers (IEEE); the use of commercially available wireless access nodes and switches; the use of configurations such as Address Resolution Protocol (ARP) spoofing; and security issues including data theft, network penetration, unauthorized ARP requests, and denial of service (DoS) attacks. In my view, this suggests a POSITA would have more than just an academic knowledge of computer networking. Dr. Lavian’s description that a POSITA would have a bachelor’s degree plus some years’ work experience in computer networks appears to capture this. I agree with Dr. Reiher that work in the

computer networking field might provide equivalent knowledge to a degree. However, Dr. Reiher's suggestion of a mere three years' experience without any academic background in computer science, or a degree without experience, is insufficient to describe the POSITA.

[75] Regardless of their background, I agree with both experts that the POSITA would have significant knowledge of computer networking technology, including the integration of wireless devices into computer networks. That knowledge would include knowledge of security issues related to wireless communications and computer networks. While Dr. Lavian may have overstated his description of the "knowledge and experience" of a POSITA to some degree, I do not believe anything turns on this.

C. *The Common General Knowledge*

[76] The experts adopted very different approaches to describing the CGK of the POSITA. In his first report, Dr. Reiher provided only a brief high-level description of the CGK, describing general areas of knowledge: Reiher First Report, paras 69–73. Dr. Lavian, on the other hand, provided a fairly extensive list of documents to illustrate the CGK, including IEEE standards, a computer networks textbook, an academic survey paper, and commercial documents: Lavian First '760 Report, paras 6.1–6.11. Dr. Lavian also provided a detailed discussion of what he viewed to be the CGK in areas such as local area networks (LANs) including wireless LANs (WLANs), hierarchical network design, bridging, access control lists, and firewalls: Lavian First '760 Report, paras 6.12–6.29. In response, Dr. Reiher criticized Dr. Lavian's description of the CGK, noting there were very few people who would "possess detailed and intimate knowledge" in all the areas Dr. Lavian identified: Reiher Second Report, paras 24–58.

[77] As with the definition of the POSITA, the parties agree little turns on the differences of opinion on the CGK. There was only one construction issue on which the parties' submissions relied on this difference, namely the meaning of *whether communications between the first wireless computing device and the second wireless computing device are allowed*, addressed at paragraphs [174] to [189] below. For the reasons I give in that discussion, the extent to which a POSITA would know about and be focused on DoS attacks in particular is not ultimately relevant to the construction of this term. The parties' arguments on obviousness also did not turn on any differences as to the scope of the CGK.

[78] I therefore need not address the issue at length, but make the following observations. First, while Dr. Lavian's discussion of the CGK was helpful, its value was somewhat lessened by the fact that it seemed designed to underscore validity issues, including numerous comments on how aspects of Claim 1 of the '760 Patent would be known to the POSITA as part of the CGK or would be self-evident: Lavian First '760 Report, paras 6.2, 6.3, 6.6, 6.11, 6.20, 6.23.

[79] Second, Dr. Reiher took too narrow a view of the CGK. He appeared to limit the CGK to that which was "retained for instant recall," or known "by heart." The CGK is not limited to what has been memorized. It includes what the POSITA may reasonably be expected to know "and to be able to find out": *Tetra Tech* at para 28. Dr. Reiher recognized that a POSITA would have familiarity with the topics, and could use online searches to look up details regarding applicable protocols. For example, I consider that widely recognized standards such as the IEEE standards referred to by Dr. Lavian would have been within the CGK of the POSITA, even though it is unlikely the POSITA would know those documents from memory and may have to consult them to apply them in a particular context.

[80] Ultimately, the parties agreed the POSITA would have knowledge of a number of relevant areas of network design and management. This includes layered network architecture; wireless protocols for implementing WLANs; wireless encryption; Internet Protocol (IP) packet formats and their use in traffic forwarding and filtering through access control lists; networking equipment including the use of wireless access points, switches, routers, and gateways; and the use of proxy ARP to cause network traffic to be directed to a desired destination. I will address certain aspects of the CGK in further detail as they arise in context below.

[81] Finally, a brief note regarding the *Computer Networks* textbook that Dr. Lavian cited as part of the CGK and that Dr. Reiher recognized was "well-known and popular." Dr. Lavian cited the 4th edition of the text, which predates the publication and priority dates of the '760 Patent: AS Tanenbaum, *Computer Networks, 4th ed* (New Jersey: Prentice Hall PTR, 2003). However, the excerpt Dr. Lavian attached to his report was actually from the 5th edition, which postdates the patent: AS Tanenbaum and DJ Wetherall, *Computer Networks, 5th ed* (Boston: Prentice Hall, 2011); Lavian First '760 Report, Appendix TL-08. This apparent oversight was not raised until Guest Tek's closing written argument, where it argued the 5th edition is not citable. I conclude this is not an issue, as Dr. Lavian's report on the '345 Patent attached the 4th edition: Lavian First '345 Report, Appendix TL-09. As the 4th edition is cited by Dr. Lavian in his First '760 Report and is in the record, I will simply refer to the 4th edition. In any event, although the 5th edition postdates the publication date of the '760 Patent, its discussion is general and refers to previously known information rather than recently developed knowledge. Dr. Reiher did not raise any concern that the text's descriptions of basic aspects of computer networking were not part of the CGK at either the priority date or publication date of the '760 Patent. Nor did Guest Tek point to anything material from the 5th edition that would not have been known in 2005.

D. *Claims Construction*

[82] Claim 1 of the ’760 Patent claims the following (I have added the numbering):

> A network comprising:
>
> a gateway; and
>
> (1) a wireless access node coupled to the gateway and configured to receive first packets from a plurality of wireless computing devices attempting to access the network, each of the first packets corresponding to one of a plurality of traffic types, and at least one of the traffic types corresponding to an encrypted wireless protocol;
>
> (2) the wireless access node further configured to transmit all first packets received from the wireless computing devices to the gateway on the network regardless of destination addresses associated with the first packets;
>
> (3) the gateway configured to determine, for each packet of the first packets received from the wireless access node, whether the packet is from a first one of the wireless computing devices directed to any other of the wireless computing devices on the network with reference to at least a source address and a destination address associated with the packet;
>
> (4) the gateway further configured to transmit the packet to the destination address associated with the packet when the packet is not directed to any other wireless computing device on the network; and
>
> (5) when the packet is directed to a second wireless computing device on the network, the gateway further configured to determine whether communications between the first wireless computing device and the second wireless computing device are allowed, and to either forward the packet to the destination address associated with the packet when communications between the first and second wireless computing devices are allowed or prevent the packet from reaching the destination address when communications between the first and second wireless computing devices are not allowed.

[83] There is no dispute that Claim 1 claims a *network* that comprises a *gateway* and a *wireless access node*. Nor is there any material dispute that a *network* is at least two data-linked computers, and that a *wireless access node* is a hardware device that serves as an access point to receive wireless traffic from *wireless computing devices* such as phones or laptops (for brevity, I will refer to *wireless computing devices* simply as “wireless devices”).

[84] There were somewhat differing descriptions of what a *gateway* is, but the differences are immaterial to the infringement and validity issues. Dr. Reiher’s definition that a gateway is a “hardware device used to connect two or more networks” is consistent with how the term is used in the patent. I need not assess whether the term is broad enough to be equivalent to simply a “router,” a “layer 3 switch,” or a “wireless bridge” as Dr. Lavian suggests.

[85] Rather, the disputes between the parties lie in the claim’s descriptions of how the wireless access node and the gateway are configured. These configurations are set out in the five paragraphs of the claim structure I have numbered above. The first two refer to the configuration of the wireless access node, while the latter three refer to the configuration of the gateway. Each of these paragraphs in turn contains a number of relevant terms.

[86] On these issues, there was disagreement between the experts on a large number of terms in the claims. Both parties unfortunately insisted in closing arguments that their respective experts set out the correct construction of every term and declined to adopt the other’s construction of elements even where these did not lie at the heart of their arguments on invalidity or infringement. I will therefore address the terms on which the parties’ positions differ, but in doing so I will focus on the main issues the parties identified as particular points of contention: *Cobalt* at para 83.

[87] The parties’ closing arguments effectively focused on four construction issues each, although only two of these overlapped. Both parties addressed the terms *determine[…]whether the packet is from a first one of the wireless computing devices directed to any other of the wireless computing devices on the network* (addressed at paragraphs [140]-[168] below) and *whether communications between the first wireless computing device and the second wireless computing device are allowed* (paragraphs [175]-[189]). Guest Tek’s closing submissions additionally focused on the terms *coupled to the gateway* (paragraphs [90]-[93]) and *configured to transmit* (paragraphs [108]-[127]), while Nomadix additionally focused on *first packets* (paragraphs [94]-[101]) and *encrypted wireless protocol* (paragraphs [102]-[106]).

[88] The main disputes between the parties pertained to the interpretation of terms in the claims, rather than essentiality. Although the Reiher Second Report contained some comment on essentiality in response to Dr. Lavian’s report (which effectively concluded every element was essential), I agree with Nomadix that these comments are better viewed as disagreement on interpretation. Guest Tek made no submissions in either written or final argument that any elements of the asserted claims of the ’760 Patent were not essential. I agree the terms discussed below are essential to the claims of the patent, and would in any case apply the principle that a claim element will be considered essential if a party does not maintain that it is not essential: *Corlac* at paras 26–27.

[89] With this introduction, I turn to the specific terms in Claim 1 at issue, which for ease of reference I will address in accordance with the numbering added to the claim above, recognizing that the construction exercise is undertaken with reference to the claim as a whole.

(1) a wireless access node coupled to the gateway and configured to receive first packets from a plurality of wireless computing devices attempting to access the network, each of the first packets corresponding to one of a plurality of traffic types, and at least one of the traffic types corresponding to an encrypted wireless protocol

(a) *coupled to the gateway*

[90] The wireless access node must be *coupled to the gateway* to allow traffic received by the wireless access node to pass to the gateway. Dr. Lavian suggested that Claim 1 would not encompass a wireless access node connected to network switches which were in turn coupled to a gateway, *i.e.*, that the wireless access node had to be directly, and not indirectly, coupled to the gateway: Lavian Second ’760 Report, para 53 (p 18).

[91] I agree with Guest Tek that this construction is not supported by the language of the claim, which does not specify that the wireless access node must be directly coupled to the gateway. Nor did Dr. Lavian point to evidence that a POSITA would understand the term “coupled” as meaning exclusively a direct coupling. Dr. Lavian’s proposed construction also does not conform with a purposive reading of the claim. The purpose of the “coupling” is to allow traffic to be directed from the wireless access node to the gateway so that the gateway can make determinations about how to treat it. This is something that can be achieved regardless of the presence of intervening switches.

[92] Reference to the disclosure simply confirms this construction. Nothing in the patent suggests the importance of a direct connection between the wireless access node and the gateway without any intervening hardware. To the contrary, as Guest Tek points out, Figure 4 to the

’760 Patent, which is described in paragraphs [0042] and [0043] of the disclosure, shows a wireless access node (412) coupled to the gateway (402) via an intervening virtual local area network (VLAN) capable switch (404).

[93] I therefore conclude the term *coupled to the gateway* would be understood by a POSITA to require connection to the gateway, but not to require direct connection. In other words, the wireless access node must be coupled to the gateway, regardless of whether there is intervening hardware such as a switch.

(b) *first packets*

[94] A *packet* is a bundle of data configured in a form that can be recognized and routed across a computer network. The Internet Protocol provides for a packet format that includes a packet header. The packet header includes information such as, but not limited to, the source and destination addresses for the packet: Tanenbaum, pp 40, 206–207, 247; Reiher First Report, para 46.

[95] Claim 1 provides that the wireless access node is configured to receive *first packets* from wireless devices, and to transmit *all first packets* to the gateway regardless of destination addresses. I agree with Guest Tek that the admittedly somewhat unusual term *first packets* is a patent drafting device designed to distinguish the packets being referred to in Claim 1 from other packets, notably the *second packets* referred to in Claim 5. It thus simply refers to the packets that are sent from the wireless devices identified in the claim to the wireless access node.

[96] Dr. Lavian argued the term *first packets* should be given its “plain and ordinary meaning,” in which the word *first* conveys the sequence or order of the packets: Lavian Second ’760 Report, Appendix TL-21, paras 71–81. In other words, packets would be *first* packets because they came before other packets in terms of time or sequence. This led Dr. Lavian to suggest the first packets sent between a wireless device and a wireless access node would simply be those establishing a wireless connection: Lavian Second ’760 Report, paras 47, 53 (p 19).

[97] In my view, Dr. Lavian’s proposed construction—which Nomadix did not press in closing argument, although it still generally argued Dr. Lavian’s constructions should be adopted—cannot be accepted. There is no evidence there is a “plain and ordinary meaning” of the term *first packets* that would be understood by a POSITA at the time of publication or that it has any particular known meaning in the art. Nor do I agree, as Dr. Lavian suggests, that the English word *first* has a single meaning that necessarily evokes time or sequence.

[98] Even within Claim 1, there are two signs the term *first* is not being used to indicate time or sequence. Importantly, the wireless access node is configured to forward *all first packets* to the gateway. This is central to the claim, as the gateway then makes determinations about packet forwarding. The claim would make little sense if this function were limited to only a temporally or sequentially limited set of packets. A mind willing to understand reading the claim purposively would therefore not read it in this fashion: *Free World Trust* at para 44. As a further indicator, the term *first* is also used within Claim 1 to distinguish between two wireless devices: a *first wireless computing device* and a *second wireless computing device*. Since neither device

comes first or second in any sort of sequence, the terms can only be understood as a signifier to keep the two devices separate.

[99] Reviewing Claim 1 in the context of Claim 5 of the ’760 Patent confirms this reading. Claim 5 adds a *wired access node* to the network of Claim 1. The wired access node is configured to receive *second packets* from *wired computing devices*. Those second packets are then similarly transmitted to the gateway regardless of destination address and (in Claim 6) a determination is made regarding forwarding the packets. No reasonable meaning was proposed that gave a temporal or sequential meaning to the term *second* in Claim 5. Rather, *first* in Claim 1 and *second* in Claim 5 are simply used to distinguish packets from the wireless devices (*first packets*) and those from the wired devices (*second packets*).

[100] Guest Tek pointed to two Canadian cases and two American cases in which the terms “first” and “second” were used to distinguish between elements without invoking any temporal or other sequencing: *Western Oilfield (FC)* at paras 45, 65–66, aff’d *Western Oilfield (FCA)* at paras 23–26; *Hershkovitz v Tyco Safety Products Canada Ltd*, 2009 FC 256 at paras 55, 65; *3M Innovative Properties Co v Avery Dennison Corp*, 350 F.3d 1365, 1371 (Fed Cir 2003); *Free Motion Fitness, Inc v Cybex International, Inc*, 423 F.3d 1343, 1348 (Fed Cir 2005). To this, the Court would add reference to *Whirlpool*, in which the claims at issue referred to a “first agitator portion [or element]” and a “second agitator portion [or element]” to distinguish between the two portions, rather than to indicate any particular sequence between them: *Whirlpool* at paras 17, 20. In other contexts, “first” and “second” may well convey a sequence: see, *e.g.*, *Beloit Canada Ltd v Valmet-Dominion Inc*, [1997] 3 FC 497 at paras 22–23. What matters is the language of the

claim as used in context. Nonetheless, these cases do confirm that the terms “first” and “second” are used in the practice of patent drafting as a means of distinguishing elements rather than necessarily sequencing them.

[101] In addition to being skilled in the art, the hypothetical POSITA is a “skilled reader” of patents: see, *e.g.*, *Teva* at para 80. That skill includes knowledge of various patent drafting rules and conventions: *Bauer Hockey* at para 51. In my view, the use of the terms *first* and *second* in Claims 1 and 5 of the ’760 Patent derives from patent drafting conventions such as seen in the cases cited above. Reading Claim 1 through the eyes of the skilled reader, I conclude the *first packets* are simply those described in the claim as having been received by the *wireless access node* from a *wireless computing device*.

(c) *encrypted wireless protocol*

[102] At least one of the traffic types flowing from the wireless devices to the wireless access node must correspond to an *encrypted wireless protocol*. Dr. Reiher first described an encrypted wireless protocol as “a network protocol <u>designed for use over a wireless network</u> that uses some form of cryptography, most commonly to provide security services to information crossing the wireless network” [emphasis added]: Reiher First Report, para 97. He provided examples of Wired Equivalent Privacy (WEP), Wi-Fi Protected Access (WPA) and WPA2 protocols. He repeated this definition and examples in the Reiher Second Report.

[103] In his oral testimony, however, Dr. Reiher used a broader definition, referring to “a <u>protocol that is running over a wireless network</u> in which some or all of the data in that set of

communications across the wireless network have been encrypted" [emphasis added]: Transcript, pp 425–426. He noted that a common way of applying secrecy to a network with an open wireless access point is to use end-to-end cryptography. He referred in particular to the "secure socket layer" (SSL) protocol as "a standard way" to do this, suggesting SSL was an *encrypted wireless protocol*: Transcript, pp 428–429. In cross-examination, he referred to SSL as a "wired and wireless protocol" that was an encrypted wireless protocol "when used over wireless," but admitted it was not designed specifically for wireless: Transcript, pp 683–684. Dr. Lavian, for his part, responded in oral evidence that while SSL was an "encryption protocol," it was not an "encrypted wireless protocol": Transcript, pp 1656–1659.

[104] In my view, the POSITA reading Claim 1 would understand the term *encrypted wireless protocol* in accordance with Dr. Reiher's initial definition rather than his second broader one. It would thus be understood to mean encryption protocols designed specifically for wireless communications, including WEP, WPA, and WPA2. It would not include general encryption protocols that may happen to be used on wireless networks, such as SSL.

[105] The CGK known to the POSITA indicates that wireless encryption is a separate matter addressed in the IEEE 802.11 Standard: see ANSI/IEEE Std 802.11, 1999 Edition, s. 8.2; IEEE 802.11i Overview, February 9, 2005, Slides 12, 15; Tanenbaum at pp 25, 227, 242, 603–605. A POSITA would understand the claim specifies "wireless" in the term *encrypted wireless protocol* rather than simply referring to an encrypted protocol or encryption protocol. This indicates that protocols pertaining specifically to wireless encryption were intended. This was Dr. Reiher's initial reading: a protocol <u>designed</u> for use over a wireless network. His subsequent expansion of

the term, which appeared intended to react to a non-infringement argument raised by Dr. Lavian discussed below at paragraph [135], was unconvincing.

[106] I therefore conclude the first descriptive paragraph in Claim 1 of the ’760 Patent, as it would be understood by the POSITA, requires the network to include a wireless access node that is (i) coupled (directly or indirectly) to the gateway; and (ii) configured to receive from wireless devices packets of various traffic types, at least one of which corresponds to an encryption protocol designed for wireless traffic (such as WEP, WPA, or WPA2).

> (2) the wireless access node further configured to transmit all first packets received from the wireless computing devices to the gateway on the network regardless of destination addresses associated with the first packets

[107] Given the definition of *first packets* discussed above, this descriptive paragraph on its face requires the wireless access node to be configured to transmit to the gateway all the packets received from the wireless devices, regardless of the destination address of the packet or the sequence of their transmission. The parties’ arguments and the experts’ evidence raise two key issues with this paragraph: (a) whether the wireless access node can be considered *configured to transmit* the packets to the gateway if that result is achieved by or in combination with other devices, notably the gateway; and (b) the scope of the term *all first packets*.

(a) *configured to transmit*

[108] Dr. Reiher concluded that although Claim 1 refers to the *wireless access node* being *configured* to transmit all first packets to the gateway, this would be understood to mean that the

wireless access node is configured to perform this transmission "either fully of its own accord **or** because of actions performed by other devices on the network such as the gateway and the wireless computing devices" [Dr. Reiher's emphasis]: Reiher First Report, para 127. Dr. Lavian, on the other hand, concluded the term meant the wireless access node itself has a forwarding table (a table that tells the node where to send an incoming packet) configured to send all packets to the gateway: Lavian First '760 Report, para 8.1 (p 48). Dr. Lavian considered Dr. Reiher's interpretation not to reflect the reality of network devices, which transfer packets based on their own configuration and do not act in concert with other devices in the manner Dr. Reiher suggested: Lavian Second '760 Report, Appendix TL-21, paras 89–95.

[109] Dr. Reiher himself recognized "it would initially seem logical to take the opinion that, if a claim says a wireless access node is further configured to do something, the wireless access node should be configured to do that something by itself," through software or other programming in the wireless access node: Reiher First Report, para 100. However, Dr. Reiher concluded such a construction would not accord with the claims or disclosure of the '760 Patent. He provides two main reasons to depart from his initial reading: (i) the other use of the term *configured* in Claim 1, and (ii) the discussion of ARP spoofing in the '760 Patent. In my view, neither of these is persuasive. The former can be addressed fairly quickly, while the latter requires a more thorough consideration of ARP requests, ARP spoofing/proxy ARP, and the discussion of these processes in the '760 Patent.

(i) Other use of *configured* in Claim 1: *configured to receive*

[110] Claim 1 states that the wireless access node, in addition to being *configured to transmit* all first packets to the gateway, is *configured to receive* packets from a plurality of wireless devices. Dr. Reiher suggests a wireless access node could not be *configured to receive* packets without the wireless devices "playing along" by actually sending the packets. He therefore concludes the term *configured* in Claim 1 can indicate the wireless access node acts with the "help" of other devices, both for receiving packets from the wireless devices and transmitting them to the gateway: Reiher First Report, paras 102–103.

[111] I cannot accept this argument. In my view, the first use of the term *configured* in the claim does not require the "help" Dr. Reiher contends. Rather, the claim language specifies the wireless access node is *configured to receive* packets, not that it needs to actually receive such packets, or that the wireless devices necessarily have to transmit them. This is inherent in the nature of the claim as a "network" claim, which defines the elements of the network and how they are configured. Further, even if one considers the network in operation, I do not believe a POSITA would consider a wireless device that is sending packets to be "playing along" or somehow taking part in the action of receiving them. Rather, the wireless device sends wireless packets and is configured to do so; the wireless access node receives them and is configured to do so. While some communication between the devices may be necessary to establish a link and ensure the packets are sent, formatted, or encrypted correctly, this does not mean that the wireless device participates or assists in the act of receiving the packets. I therefore conclude that reference to the wireless access node being *configured to receive* does not support Dr. Reiher's construction of the term *configured*.

(ii) ARP spoofing

[112] Dr. Reiher’s second ground for construing the term *configured to transmit* as including actions by other network elements derives from the discussion of ARP spoofing in the disclosure and diagrams of the ’760 Patent. I am not convinced “recourse” to the disclosure is necessary to resolve any ambiguity in the term *configured* as it appears in Claim 1, particularly given Dr. Reiher’s indication as to what “initially seem[s] logical” in reading the claim. Nonetheless, in light of the discussion at paragraphs [41] to [47] above, I am conscious of the need to read the disclosure and the claims “as a whole,” so will consider the context of the disclosure and Dr. Reiher’s second ground. As noted, this requires an understanding of ARP requests and ARP spoofing/proxy ARP, which in turn requires an understanding of packet addressing protocols. The parties agree this information would have been known to the POSITA at the time of publication of the ’760 Patent. While the following discussion will no doubt be considered simplified from the perspective of the experts or the POSITA, I believe it is of sufficient detail to inform the construction of the claims of the ’760 Patent. The following understanding comes primarily from Dr. Reiher’s first report and the Tanenbaum text: Reiher First Report, paras 34–46, 109; Tanenbaum, pp 37–39; 344–345; Exhibit 114, Item 3.

[113] A computer device, including a wireless device or a gateway, has a unique Media Access Control (MAC) address. The MAC address identifies the computer at the “data link layer,” the second of seven layers in the Open Systems Interconnection (OSI) model of network architecture (being the physical; data link; network; transport; session; presentation; and application layers). The MAC address is a 48 bit number, usually written as six pairs of hexadecimal numbers, such

as 00:50:E8:00:03:AF (in this example, given by Dr. Reiher, the first three pairs represent Nomadix's vendor code, so the device with this MAC address would be a Nomadix device).

[114] An IP address, on the other hand, is a number assigned to a device that is connected to a network such as a LAN. An IP address identifies a computer at the network layer, and is typically a 32-bit number written in dotted decimal notation with four numbers between 0 and 255, such as 192.168.90.1.

[115] If a source device wishes to send a unicast packet (a packet intended for a specific destination, as opposed to a multicast or broadcast communication), it needs the data link layer MAC address of the destination. The packet source may know the destination's IP address, but not the MAC address. The source may send out an ARP request to resolve this. As explained to the Court, an ARP request may be viewed as the source computer sending a message out saying "Can anyone tell me the MAC address of the device associated with this IP address?" A device that knows the answer (typically the device with the MAC address in question) may send an ARP response, effectively providing the answer "This is the MAC address of the device associated with the IP address you requested."

[116] For various reasons, a different network device may respond to the ARP request. This is known as a "proxy ARP" response. A proxy ARP response may provide the "true" MAC address of the device whose IP address is requested, or may provide a different MAC address, such as its own.

[117] The use of proxy ARP responses may be for nefarious purposes, such as the diversion of traffic. This is known as “ARP spoofing.” Although there may be technical differences, the primary difference between proxy ARP and ARP spoofing appears to be in the intent and/or author of the conduct. A legitimate device pursuing valid networking goals engages in proxy ARP; an illegitimate device engaged in, for example, a security attack engages in ARP spoofing. The ’760 Patent uses the term “ARP spoofing,” but the parties agree that given how it is used, the preferable term would have been proxy ARP. The parties and experts agree nothing turns on this difference in terminology.

[118] Figure 2 of the ’760 Patent is a flowchart illustrating aspects of the network, the upper half of which appears as follows:

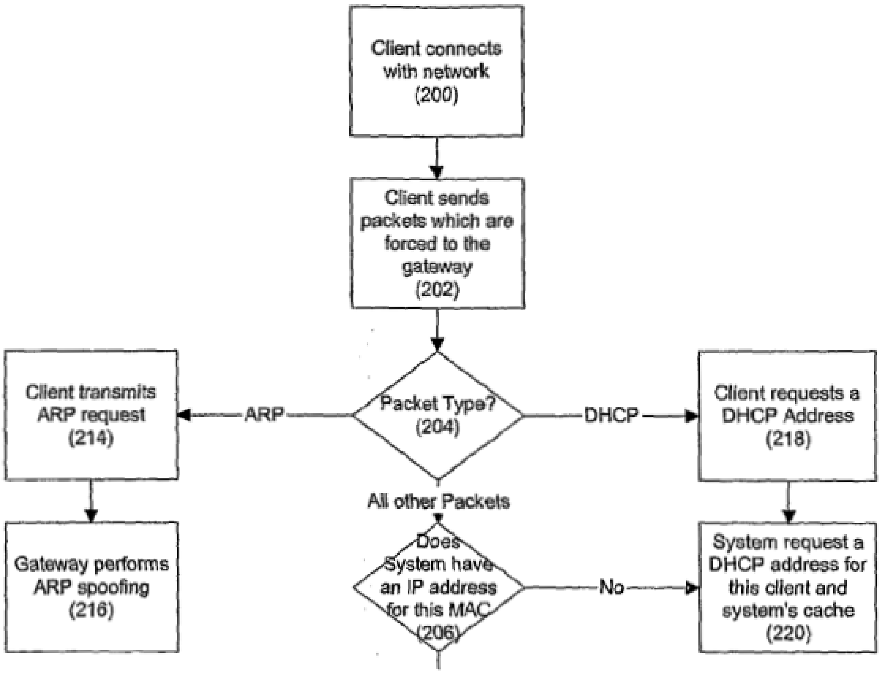
Client connects with network (200)
Client sends packets which are forced to the gateway (202)
Packet Type? (204)
ARP
DHCP
All other Packets
Client transmits ARP request (214)
Client requests a DHCP Address (218)
Gateway performs ARP spoofing (216)
Does System have an IP address for this MAC (206)
No
System request a DHCP address for this client and system's cache (220)

[Description of inserted diagram for accessibility: A partial flowchart is represented with 8 boxes connected by arrows. At the top, a rectangle labeled “Client connects with network (200)” has an arrow pointing down from it to a rectangle labeled “Client sends packets which are forced to the gateway (202).” From the latter rectangle an arrow points down to a diamond labeled “Packet Type? (204).” Three arrows lead from the left, right, and bottom points of the

"Packet Type?" diamond. The first arrow, pointing left, is labeled "ARP" and points to a rectangle labeled "Client transmits ARP request (214)," which is in turn connected by an arrow to a rectangle beneath it labeled "Gateway performs ARP spoofing (216)." The second arrow, pointing right from the "Packet Type?" diamond, is labeled "DHCP" and points to a rectangle labeled "Client requests a DHCP Address (218)," which is connected by an arrow to a rectangle beneath it labeled "System request a DHCP address for this client and system's cache (220)." The third arrow, pointing down from the "Packet Type?" diamond, is labeled "All other Packets" and points to a further diamond labeled "Does System have an IP address for this MAC (206)." From this further diamond, an arrow labeled "No" points right to the same rectangle labeled "System request a DHCP address for this client and system's cache (220)."]

[119] The disclosure of the '760 Patent describes this part of the flowchart as follows:

> [0029] FIG. 2 is a flowchart illustrating part of a session during which a client machine connects with a network […]. When a client machine connects with the network (200), e.g., by entering a wireless hotspot, <u>it starts sending packets which are forced to the network's gateway (202). If the client machine transmits an ARP request looking for the gateway on its home network (204, 214) and regardless of the client machine's settings the gateway (or an associated network device) performs ARP spoofing (216), returning its own MAC address instead of the requested gateway MAC address</u>. The client machine then starts sending packets to the network device as if it were the requested gateway. […]
>
> [Emphasis added]

[120] Thus, where a client connected to the network sends an ARP request, the gateway (or an associated device) sends a proxy ARP response identifying its own MAC address in place of the destination the client was looking for. Packets sent by the client device will thereafter be sent to the gateway (or associated device). Dr. Reiher reads this as being the method by which packets are "forced to the network's gateway," *i.e.*, the method by which, in the language of Claim 1, the packets are transmitted to the gateway on the network regardless of destination address. This, he contends, supports a construction that allows the wireless access node to "receive help" from other devices to transmit all first packets to the gateway: Reiher First Report, paras 106–113.

[121] The primary difficulty with this reliance on Figure 2 and paragraph [0029] of the ’760 Patent is that the flowchart and the text both describe a configuration in which packets are forced to the gateway *even before* a proxy ARP response is sent. Thus in the flowchart, the client connects to the network (200), and sends packets “which are forced to the gateway” (202), and only then does the distinction between packet types (ARP, DHCP (Dynamic Host Configuration Protocol), or “All other Packets”) occur. This process is repeated in the description in paragraph [0029]. The method of forcing packets to the gateway is therefore set out as being a different part of the process from the use of proxy ARP responses.

[122] Dr. Reiher appeared to recognize this in his examination in chief, noting that “for this invention to work properly as described in the patent, you must ensure that the wireless access point is configured, is set up, so that when it sees these ARP requests it says, well I’m always sending those to the gateway”: Transcript, p 435. In other words, proxy ARP responses cannot be the manner, or at least the only manner, in which packets are forced to the gateway. The wireless access node itself must be set up to send, at the very least, ARP requests to the gateway.

[123] This conclusion is reinforced by paragraph [0041] of the disclosure. That paragraph provides specific discussion of how an access node can be “configured to pass all of its incoming traffic to an associated gateway device.” This includes the use of distinct VLANs on every port; using separate physical hardware links; disabling port-to-port traffic on each device; or configuring a proprietary method to send the data directly to the gateway. In discussing how this is achieved, the ’760 Patent does not discuss the use of proxy ARP or ARP spoofing. Dr. Reiher recognizes that each of these are “implemented on the access node itself” and would not rely on

"assistance from other devices." However, he nonetheless concludes that these descriptions of how packets are forced to the gateway does not preclude such forcing from happening through the use of proxy ARP: Reiher First Report, paras 114–121.

[124] In my view, if the disclosure is to be reviewed in construing Claim 1 of the '760 Patent, that disclosure must be read in a manner designed to understand how the inventors intended to use the terms in the claim. Here, the inventors noted the importance of forcing all packets from the access node to the gateway, a step that occurs independent of, and before, the conduct of ARP spoofing. They also set out in some detail examples of how forcing packets can occur, each of which involves configuring the access node, and which do not include using proxy ARP as the forcing method. I therefore conclude that referring to the disclosure simply reinforces the language of Claim 1, namely that the wireless access node must itself be configured to transmit packets to the gateway.

[125] It is also significant that the proxy ARP process does not itself "configure" the wireless access node in any way. To the contrary, as seen in Figure 2, it is the client device that sends an ARP request. The gateway then sends a proxy ARP response to the ARP request. The result of that proxy ARP response, as the Court understands it, is that the client device (not the wireless access node) would then understand that the MAC address of the gateway is associated with the IP address it is looking to send its packet to. That association would be stored in the device's "ARP table": Reiher First Report, paras 46, 343. The device would therefore send future packets intended for that IP address to the gateway's MAC address. Nothing in the wireless access node is changed or configured by this process.

[126] Ultimately, Guest Tek's argument is that "[s]o long as the wireless access node transmits packets to the gateway, this limitation is met." I cannot agree. While a purposive construction must be given to the claims, the claims must nonetheless be construed as they are written: *Electro Santé* at para 31. Claim 1 states that the *wireless access node* is *configured* to transmit packets to the gateway. It does not say the wireless access node simply *transmits* packets to the gateway. Nor does it say that the *network* is configured so that packets are transmitted from the wireless access node to the gateway.

[127] I therefore conclude the POSITA would understand Claim 1 in the way that "initially seem[ed] logical" to Dr. Reiher, namely that the wireless access node must, itself, be configured so the packets it receives from wireless devices are transmitted to the gateway regardless of their destination address. At the same time, I agree with Dr. Reiher there is no limitation in Claim 1 on how the wireless access node may be configured to transmit all first packets to the gateway. It does not specify, for example, that this needs to be accomplished through the use of the forwarding tables Dr. Lavian describes: Reiher Second Report, para 85. As noted, paragraph [0041] of the disclosure of the ’760 Patent suggests a number of different methods.

[128] By way of foreshadowing, the foregoing discussion does not apply to the method of Claim 21. That claim does not require the wireless access node to be *configured* to transmit all first packets to the gateway; simply that it transmits them.

(b) *all first packets*

[129] As discussed above, the *first packets* are the packets received by the *wireless access node* from a *wireless computing device*. Claim 1 requires the wireless access node to be configured to transmit *all* such first packets to the gateway, regardless of the destination address associated with them. Given the security context of the ’760 Patent, the purpose of ensuring *all first packets* are transmitted to the gateway is so the gateway can perform its determination on whether to forward or drop the packets, as described in the remaining paragraphs of Claim 1 discussed below. If not all packets from the wireless devices are forwarded, then some of those packets might not be subjected to the security filtering process described in the remainder of the claim.

[130] Two issues arose with respect to the meaning of *all first packets*. The first concerns whether it includes data exchanged during the initial “handshake” or “key exchange” protocol by which a wireless device establishes a secure connection with a wireless access point. The second relates to whether the *all first packets* transmitted to the gateway must include packets corresponding to an encrypted wireless protocol.

[131] As alluded to above, Dr. Lavian suggested there was no transmission of *all first packets*, and thus no infringement of the ’760 Patent, if the packets exchanged in the handshake are not transmitted to the gateway: Lavian Second ’760 Report, paras 2, 32–33, 47–48, 53 (p 19). This implies a construction of *all first packets* that includes the information exchanged in this initial connection and authentication. I cannot accept this construction.

[132] As an initial matter, Dr. Reiher opined that the data exchanged in the key exchange protocol is in the form of “frames” rather than “packets” and therefore would not fall within the term *all first packets*: Transcript, pp 421–422, 438. In my view, this does not assist Guest Tek. The ’760 Patent does not distinguish between frames and packets. Dr. Reiher agreed in cross-examination that he would similarly term ARP requests “frames,” yet it is clear from Figure 2 they are among the “packets” that are to be forced to the gateway in the claimed network. Dr. Reiher himself referred to such frames as “packets” in his report with reference to Figure 2: Transcript, pp 703–704; Reiher First Report, para 91.

[133] Nonetheless, as Dr. Lavian and Dr. Reiher both stated, and the POSITA would know, the data exchanged between a wireless device and a wireless access node to establish communication at the physical layer are only ever exchanged between those two devices. They are not intended to be forwarded to a further node, nor are they ever forwarded: Transcript, pp 421–422, 1655–1656, 1998; Lavian Second ’760 Report, paras 47, 53 (p 19). While Dr. Lavian referred to this as a reason the Nomadix system could not infringe the patent, it would also exclude all other networks with a wireless access node and a gateway, such that no real-world network, including those described in the patent, would fall within the claim. In such a context, I cannot conclude a POSITA reading the patent purposively and with a mind willing to understand would read Claim 1 as requiring the forwarding of this data to the gateway. Further, the purpose of forwarding *all first packets* to the gateway is for the gateway to make a determination on whether to transmit or drop the packet based on the packet’s destination. Such a determination is inapplicable to the data initially exchanged between the wireless device and the wireless access node to establish communication.

[134] Indeed, Dr. Lavian's construction is to some degree at odds with comments he made in his first report. There, Dr. Lavian described the requirement that the wireless access node be "configured to receive first packets from a plurality of wireless computing devices" as being that "[t]he access points receive wireless traffic from the wireless nodes <u>that are associated and authenticated with it</u>": Lavian First '760 Report, para 8.1 (p 48). This language suggests that the association and authentication process occurs before the "wireless traffic" (*i.e.*, the first packets) is transmitted from the wireless device to the wireless access node. Further, as noted above, Dr. Lavian described the requirement that the wireless access node be configured to transmit all first packets to the gateway as contemplating a "forwarding table that is configured to send all packets to […] the gateway": Lavian First '760 Report, para 8.1 (p 48). Yet Dr. Lavian did not suggest such a forwarding table would ever forward the initial authentication packets.

[135] Dr. Lavian also suggested the requirement that the wireless access node *transmit all first packets received from the wireless computing devices to the gateway* required the wireless access node to transmit at least some traffic encrypted with an encrypted wireless protocol: Lavian Second '760 Report, paras 31–32 and Appendix TL-21, paras 150–160; Transcript, pp 1675–1679. Again, however, this is something that simply does not occur, since part of the function of a wireless access node is to decrypt encrypted wireless protocol traffic received from wireless devices. As Dr. Lavian testified, traffic is encrypted with an encrypted wireless protocol only for the purpose of transmission between the wireless device and the access point. After that, it is not encrypted with such encryption, and other devices such as a gateway would not be able to decrypt it if it were: Transcript, pp 1679–1680. As a result, while Dr. Lavian referred to this as another reason the Nomadix system could not infringe the patent, it would again also exclude

every other network with a wireless access node and a gateway, including those described in the patent.

[136] There is in any case nothing to support this construction in either the claim or the disclosure. The requirement that one of the traffic types correspond to an encrypted wireless protocol pertains to the packets *from* the wireless device. Nothing in Claim 1 requires the packet to be encrypted with the same encrypted wireless protocol when transmitted to the gateway, particularly where that would render the packet useless to the gateway.

[137] I therefore conclude the second descriptive paragraph in Claim 1 of the ’760 Patent, as it would be understood by the POSITA, requires the wireless access node, itself, to be configured to transmit to the gateway all of the packets it receives from the wireless devices, regardless of their destination addresses. However, this does not require the wireless access node to be configured to forward the data (whether referred to as packets or frames, which I need not decide) exchanged between a wireless device and the wireless access node for purposes of establishing a connection. Nor does it require any of the packets transmitted by the wireless access node to the gateway to be encrypted with an encrypted wireless protocol.

> (3) the gateway configured to determine, for each packet of the first packets received from the wireless access node, whether the packet is from a first one of the wireless computing devices directed to any other of the wireless computing devices on the network with reference to at least a source address and a destination address associated with the packet

[138] Having addressed the configuration of the wireless access node, Claim 1 then turns to the configuration of the gateway, addressed in the paragraphs I have numbered (3), (4), and (5). As

with the entirety of the claim, these paragraphs must be read together to properly understand their meaning, but I break them out for ease of reference. By way of overview, the parties agree the gateway must be configured to make a determination based on the source and destination of the packet [paragraph (3)]. If the packet is not directed to a wireless device on the network, the packet is transmitted [paragraph (4)]. If it is directed to a wireless device on the network, a further determination is required as to whether communications between the devices are allowed. If they are, the packet is transmitted; if not, the packet is dropped [paragraph (5)].

[139] Within this broad structure, however, the parties disagree on the construction of terms in each of the three paragraphs.

(a) *from a first one of the wireless computing devices directed to any other of the wireless computing devices on the network*

[140] The first described element of the configuration is that the gateway is to *determine […] whether the packet [received from the wireless access node] is from a first one of the wireless computing devices directed to any other of the wireless computing devices on the network*. This is done with reference to at least a *source address* and a *destination address* associated with the packet. The meaning of this determination is one of the significant points of dispute between the parties and was the subject of a separate report by Dr. Reiher, the Reiher Third Report.

[141] For the reasons discussed below, I conclude this term must be read as it appears on its face. That is to say, the gateway must be configured so it can and will determine whether a packet (a) comes from a wireless device on the network; and (b) is directed to any other of the wireless devices on the network. While this conclusion may appear innocuous or straightforward

given the language of the claim, it has significant impacts in the circumstances of this case, so it must be explained in further detail. Notably, this construction means that if a gateway is not configured to determine whether a packet is directed to a wireless device on the network, *as opposed to a wired device on the network or a device not on the network*, it does not meet this element of the claim.

[142] Dr. Reiher spent little time on the construction of this provision in his first report. He noted packets may contain a *source address*, namely a MAC address from which the packet is transmitted. He then made the following statement:

> Using consistently the definition of “first packets” as I have defined it to mean above, the meaning of this paragraph appears straightforward, and this paragraph of Claim 1 requires the gateway to determine, with reference to at least the source address and the destination address associated with the packet, whether the packet is from a wireless computing device on the “system” network and is intended to be sent to another wireless computing device on the “system” network.
>
> [Emphasis added; Reiher First Report, para 129.]

[143] Dr. Lavian was similarly fairly brief. His description of the paragraph was that “[t]he gateway (router) determines whether the source and destination address of the packet are on the same subnet [network]”: Lavian First ’760 Report, para 8.1 (p 49). He also noted in his discussion of the CGK that an access control list at a gateway could control which sources are allowed to communicate with which destinations based on source and destination IP addresses: Lavian First ’760 Report, para 6.11 (p 23).

[144] In his initial response to Dr. Lavian, Dr. Reiher took issue with Dr. Lavian's simplification of the language in the claim, underscoring that the determination described is that the gateway determines whether the packet is from one wireless device directed to any other wireless device on the network with reference to at least the source and destination address associated with the packet: Reiher Second Report, paras 91–93. Conversely, Dr. Lavian's response to Dr. Reiher only took issue with Dr. Reiher's use of the word "system" in his description: Lavian Second '760 Report, Appendix TL-21, paras 129–130.

[145] However, in addressing the issue of infringement, Dr. Lavian concluded Nomadix's gateway did not determine whether an incoming packet was coming from, or directed to, a wireless device as opposed to a wired device. He therefore concluded that this element of the claim was not met: Lavian Second '760 Report, paras 3, 22, 49–50, 53 (p 19).

[146] Dr. Reiher filed a sur-reply report to address this issue, which he had not anticipated: Reiher Third Report, paras 2–4. He opined that in his view, Claim 1 does not recite a requirement that there be a determination of whether a packet is from a wired or wireless device. Rather, in his view, "[t]he determination required is whether communication **is** or **is not** permitted between the devices where the devices are wireless devices" [Dr. Reiher's emphasis]: Reiher Third Report, paras 7–9. Claim 1, in his view, extended exclusively to wireless usage, and there was no requirement that the determination as to whether communication was allowed required an additional determination that the user was a wireless rather than wired user: Reiher Third Report, paras 10, 17–18.

[147] Guest Tek's position, relying on Dr. Reiher, is effectively that Claim 1 as a whole is only directed to "wireless-to-wireless" packets and does not address wired traffic: Transcript, pp 2213–2216. There is therefore no need for the gateway to determine, in Guest Tek's view, whether the packet is from or to a wireless as opposed to wired device.

[148] While I agree with Guest Tek that the *source* of the packets at issue in Claim 1 is of necessity a wireless device, I disagree that their *destination* is necessarily a wireless device. Nor do I agree Claim 1 is restricted purely to wireless-to-wireless packets such that no determination as to whether the destination is a wireless device needs to be made.

[149] With respect to the source of the packets, the determination at issue expressly pertains to, and only to, *each packet of the first packets received from the wireless access node*. As set out above, the *first packets* in question are those the wireless access node is configured to receive from a wireless device on the network. Thus, while the claim requires the gateway to be configured to determine whether the packet is *from a first one of the wireless computing devices [on the network]*, the *first packet* the determination is being made about is, by definition, from such a wireless device. If the packet were from a wired device, or a device not on the network, it would not be one of the *first packets received from the wireless node* and no determination would be called for.

[150] As a result, the need for the gateway to be configured to determine whether the packet is *from* a wireless device on the network is effectively redundant. I appreciate that generally speaking, a claim should be construed to avoid redundancy: *Ratiopharm Inc v Canada (Health)*,

2007 FCA 83 at para 33. This principle is typically applied as between claims (*i.e.*, to avoid claim redundancy), but in my view it has some value in assessing redundancy of elements within a claim. Nevertheless, I see no other way of reading the words *determine, for each packet of the first packets received from the wireless access node* other than that the determination in question applies to the *first packets* referred to as having been transmitted from the wireless access node, which are defined as being from a wireless device on the network. The remainder of the claim only refers to the destination of the packet, implicitly recognizing that the packet is necessarily from a wireless device on the network.

[151] The second aspect of the determination, namely where the packet is *directed to*, is the central part of the determination since it dictates the treatment of the packet by the gateway. A packet coming from a wireless device on the network may be directed to another wireless device on the network, to a wired device on the network, or to a destination outside the network. Contrary to Guest Tek's submissions, I see nothing in Claim 1 that limits the notion of *first packets received from the wireless access node* uniquely to wireless destinations. Rather, Claim 1 requires that the gateway be configured to determine whether the packet is directed to a wireless device on the network. Inherent in that determination is a distinction being made between packets *directed to a wireless computing device on the network* and those *not directed to a wireless computing device on the network*. This distinction is confirmed by the test in paragraph (4) of the claim.

[152] On Guest Tek's construction, the gateway need only be configured to determine whether the packet is directed to another *computing device on the network*, based on a reading that

Claim 1 deals only with wireless communications. I cannot accept this construction, for three reasons.

[153] First and foremost, the language of Claim 1 specifies that the determination is whether the packet is *directed to any other of the wireless computing devices on the network*. If the determination were simply whether the packet is directed to any other computing device on the network, the claim could have and should have said so. Instead, the claim specifies the determination is whether the packet is directed to a *wireless* device on the network. Similarly, there is no basis to assume all first packets, or any particular first packet, are necessarily directed to a wireless device, such that the only determination to be made is whether the wireless device is on or off the network. The principles of claims construction promote adherence to the language of the claims: *Free World Trust* at paras 31–43; *Tearlab* at para 31.

[154] Second, while the network of Claim 1 must include a *wireless access node* with which *wireless computing devices* may communicate, it may also include a wired access node with which wired computing devices may communicate. This is clear from Claim 5, in which the network of any of Claims 1 to 4 further comprises a *wired access node* configured to receive *second packets* from a plurality of *wired computing devices*. An independent claim cannot be given a construction that is inconsistent with the claims that depend from it: *Halford v Seed Hawk Inc*, 2004 FC 88 at para 91, aff'd 2006 FCA 275. This confirms that even in the network of Claim 1, a packet from a wireless device may be directed to (a) a wireless device on the network, (b) a wired device on the network, or (c) a device of either kind that is not on the network. The determination described in paragraph (3) of Claim 1, which must also be made in the network of

Claim 5, determines whether the packet is directed, in particular, to a *wireless* device on the network. The configuration in paragraphs (4) and (5) of Claim 1 stipulate how the gateway will act dependent on that determination.

[155] Third, the construction proposed by Dr. Reiher is inconsistent with his recognition in cross-examination as to how the gateway of Claim 1 would treat packets once a determination is made: Transcript, pp 675–679. This is discussed in further detail at paragraphs [171]–[172] below. Dr. Reiher recognized the gateway would forward a packet directed to a wired device on the network regardless of whether communications were allowed, while a packet to a wireless device on the network would be forwarded only if communications were allowed. This requires that the gateway distinguish between these types of traffic.

[156] I note that the determination in question is between packets *directed to* a wireless device on the network and packets *not directed to* a wireless device on the network. This does not necessarily involve a determination of whether the packet is directed to a wireless as opposed to a wired device. A gateway might, for example, keep a list of all "wireless devices on the network," namely those that have accessed the network via a wireless access node. The gateway would then only need to determine whether the destination of the packet was on that list to determine whether the packet was directed to a wireless device on the network. Any packet directed elsewhere, wired or wireless, would not fall in that category.

[157] I therefore conclude the third descriptive paragraph in Claim 1 of the '760 Patent, as it would be understood by the POSITA, requires that the gateway be configured to determine, with

reference to at least a source and destination address associated with the packet, whether each of the first packets received from the wireless access node is (a) from a wireless device on the network (although it would necessarily be so), and (b) to a wireless device on the network (as opposed to any other destination, whether a wired device on the network or a destination outside the network).

[158] In my view further reference to the disclosure is not needed. In any case, the discussion in the disclosure is consistent with the foregoing construction. In particular, the Summary of the Invention in the disclosure describes the invention at paragraph [0005] as involving "an end-to-end network architecture […] which enables a population of users having diverse machine configurations and connection capabilities to reliably and securely connect to the network and the Internet." While Claim 1 clearly deals with the "machine configuration" of wireless users connecting via a wireless access node, there is no basis to expect that the Claim is only dealing with those users' access to the wireless aspects of either the network or the internet as a whole.

[159] Similarly, Figures 1A and 1B, which each show exemplary network configurations, show computing devices accessing the internet at large, rather than only wireless destinations.
Figure 1B is simpler than 1A but is illustrative:

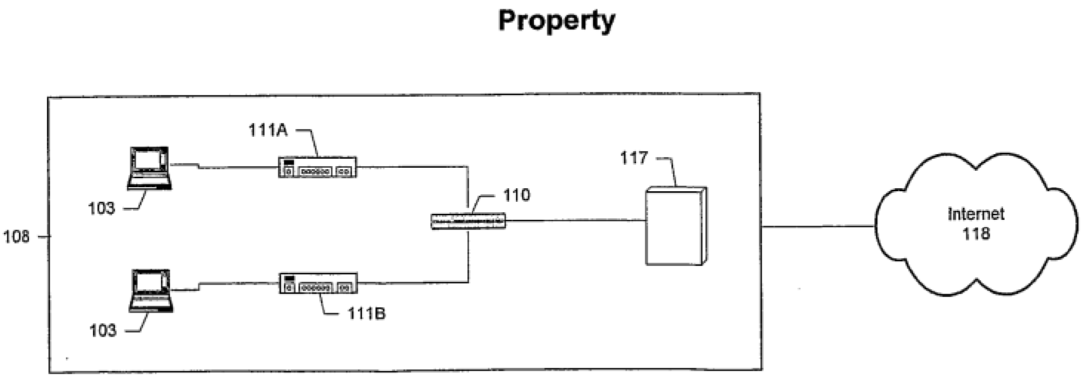

FIG. 1B

[Description of inserted diagram for accessibility: A schematic is labeled "Property." A large rectangle on the left labeled "108" is connected to a cloud on the right labeled "Internet 118." Within the large rectangle is a diagram of a simplified network in which an icon of a computer labeled 103 is connected to an icon of a hardware device labeled 111A, while a second icon of a computer labeled 103 is connected to an icon of a hardware device labeled 111B. The devices labeled 111A and 111B are both connected to an icon of a different device labeled 110, which is in turn connected to a box labeled 117.]

[160] In this figure, 111A is a wireless access node, 111B is a wired access node, and 110 is the gateway (117 is a firewall). In each case, the network configuration allows the users to access the internet as a whole. This might include the client "looking for the gateway on its home network," an example described in paragraph [0029] of the disclosure, discussed above at paragraph [119] of these reasons. No limitation is discussed in the disclosure in which access outside the network is limited to wireless destinations. There is therefore no reason in the disclosure to read Claim 1 as being limited to only packets sent by a wireless device to another wireless device, whether inside or outside the network.

[161] The disclosure of the patent also describes a network that *can* distinguish between wireless and wired devices on the network. Figure 4 in the '760 Patent is a diagram showing various devices connected to the network via different access nodes, including a wireless access

node and other nodes such as a Digital Subscriber Line Access Multiplexer (DSLAM) or a switch. With reference to Figure 4, paragraph [0043] of the ’760 Patent states:

> Referring to FIG. 4, all packets received from connected client devices (e.g., wired and wireless laptops, PDAs, etc.) by a network access node are tunneled to the gateway. <u>The gateway differentiates the packets by the tunnel in which they arrive. As discussed above, each tunnel associates the packet with its entry port into the system</u>. A wide variety of wireless and wired connection protocols may be supported. Examples of the different types of traffic which might have a dedicated tunnel (or the equivalent) include unencrypted, WPA, WPA2, AES, WEP, VoIP, or the traffic associated with a specific corporate entity.
>
> [Emphasis added.]

[162] By differentiating packets by the tunnel in which they arrive, the gateway can know whether devices that have communicated with the network are accessing it via a wireless access node. A gateway so configured would thus be able to tell whether a packet directed to a device on the network is directed to a wireless device, based on the tunnel to which the device is connected. This could also be done, according to paragraph [0040], through the different service set identifiers (SSIDs) used by different client devices based on the type of traffic they generate. Indeed, Claim 36 contemplates setting security options for a specific computing device based on the type of computing device and the access node to which it is connected to the network.

[163] As Nomadix points out, the disclosure of the ’760 Patent underscores security differences in wireless and wired communications, and focuses on wireless security concerns in particular. The first paragraph of the disclosure states the invention relates to “preventing unauthorized access to mobile devices in a wireless network.” Given this context, I believe the POSITA would

understand that making a determination between a packet directed to a wireless device on the network and other destinations was consistent with the purpose of the patent.

[164] Thus, to the extent reference is made to the disclosure of the ’760 Patent, the POSITA would see that distinguishing between wireless and wired devices was an aspect of the network contemplated by the inventors. This would simply confirm a reading of Claim 1 that treats the determination of whether a packet is *directed to any other of the wireless computing devices on the network* as including a distinction between wireless and wired devices on the network or, more precisely, between wireless devices on the network and any other device.

[165] I recognize other passages in the disclosure might suggest that the key determination to be made is whether communications between the source and destination devices are allowed, rather than whether the destination device is wired or wireless. In particular, paragraphs [0045] and [0046] refer to the determination process with reference to Figure 3 of the patent, in language that does not distinguish between wired and wireless devices:

> [0045] Once a client machine is authenticated (306), the gateway looks up the source and destination addresses in each packet (310) to determine whether any device on the network is attempting to improperly send packets <u>to any other device on the network</u>. If the source and destination of the packet are not both on the network (312) it is forwarded to the destination address (314) […].
>
> [0046] If, on the other hand, the source and destination of the packet header are determined to both be on the network (312), the packet may be an unauthorized attempt to communicate with another device on the network. <u>The gateway then determines whether communication between the two devices has been previously authorized</u> (316). If not, the gateway prevents the packet from reaching its destination, e.g., drops or redirects the packet (318). If, on the other hand, the communication has been authorized, the packet is forwarded to the destination (316) […].

[Emphasis added.]

[166] While these paragraphs suggest a similar treatment of packets from any device on the network to any other device on the network, I cannot conclude a POSITA reading Claim 1 in the context of this passage would construe the claim any differently. Notably, while the passage above refers to traffic coming from "any device on the network," Claim 1 is limited to packets coming from wireless devices, as all parties agree. Claim 1 therefore claims a particular set of limitations that are not set out in paragraphs [0045]–[0046]. Given the use of the word *wireless* in the phrase *directed to a wireless computing device on the network*, relying on these paragraphs to either effectively remove the word *wireless* or ignore potential traffic to wired devices would amount to an unjustified widening of the claim based on the disclosure: *Tetra Tech* at para 104.

[167] In light of this construction, I conclude a POSITA would interpret Claim 1 as requiring the gateway to be configured to forward all packets from a wireless device on the network that are directed to a wired device on the network. Given the differing security concerns for wireless and wired devices described in the patent, I cannot conclude this is necessarily a design flaw or a security risk the inventors did not intend or overlooked. To the extent it is, *i.e.*, that the inventors in drafting Claim 1 either failed to consider the treatment of a packet from a wireless device directed to a wired device on the network, or intended Claim 1 to only deal with wireless-to-wireless communication without saying so, I agree with Nomadix this would amount to a "self-inflicted wound" the Court should not seek to cure by reading out words of the claim or giving them a meaning they cannot bear: *Free World Trust* at para 51.

[168] As a final point with respect to this paragraph of Claim 1, I note that in reaching the foregoing construction I do not rely on Nomadix's argument based on section 53.1 of the *Patent Act*. During prosecution of the '760 Patent, Guest Tek distinguished a prior art reference cited by the examiner (identified as D1) by stating "it is clear that D1 does not make any determination about wireless to wireless device communication in order to allow outgoing traffic": Exhibit 8, p 4. Nomadix claims this shows Claim 1 was intended to distinguish between wired and wireless traffic. However, the D1 reference was not filed in evidence, so it is difficult to assess the point being made in the statement. On its face, the statement could be read as equally consistent with Guest Tek's argument that Claim 1 only addresses wireless-to-wireless communication and does not contemplate wireless-to-wired traffic at all. I therefore find the reference to be of no assistance, even if it could be taken to "rebut any representation" made by Guest Tek: *Patent Act*, s 53.1(1).

> (4) the gateway further configured to transmit the packet to the destination address associated with the packet <u>when the packet is not directed to any other wireless computing device on the network</u>

[169] Once the gateway has determined whether the packet is one *directed to any other wireless computing device on the network* or one *not directed to any other wireless computing device on the network*, the two final descriptive paragraphs in Claim 1 [which I have numbered (4) and (5)] address how the gateway is configured to deal with packets in each category. Paragraph (4) deals with packets that are *not directed to any other wireless computing device on the network*. It specifies, simply, that they are to be transmitted to their destination.

[170] Again, Guest Tek argues that Claim 1 only deals with source and destination devices that are wireless devices. For the reasons outlined above, I disagree. While the *first packets* necessarily come from a wireless device on the network, Claim 1 provides no limitation regarding the possible destinations of the packet.

[171] In cross-examination, Dr. Reiher agreed that under Claim 1, packets that are directed to wired devices would be directed to their destination address. In the following exchange, which I reproduce at some length given its importance, Nomadix's counsel put different scenarios to Dr. Reiher as to whether a packet would be transmitted to its destination in accordance with Claim 1:

> Q. So do you agree with me that these two paragraphs [the ones numbered (4) and (5) in these reasons] each present a different definition or a different condition?
>
> So the <u>first</u> one is concerned with <u>when the packet is not directed to any other wireless computing device on the network</u>, and the <u>second</u> is concerned with <u>when the packet is directed to a second wireless computing device on the network</u>?
>
> A. Those are different.
>
> Q. Okay. So you saw, in other words, two options, and we see, depending in which option we fall, how the gateway will react?
>
> A. Yes.
>
> Q. You said yesterday that the paragraphs -- the paragraph with the first option means that if a hotel guest is trying to communicate with a site outside the network, then the packet should be allowed to go through; is that correct?
>
> A. That is correct.
>
> Q. Okay. With the two, like, options that we've defined, I'd like to put to you some different scenarios, and I would want you to tell me in which of the two category it falls, okay?
>
> […]

So a packet directed to [a] wireless device on the same network would fall in the second category; right?

A. By saying “network,” you mean the hotel’s network; correct?

Q. Yes.

A. Yes. A -- in that -- to be very clear, if there is a packet from one wireless device in the hotel’s network to another wireless device on the hotel’s network, then that falls into the category of the second paragraph here.

Q. Okay, perfect. And if we talk about a packet directed to another network, in which category would it fall?

A. It would almost certainly fall into the first category. There may be a little bit of flexibility in terms of what is meant by the hotel network, but...

Q. Okay. I understand. But, in general, if you have a packet directed to another network, we can say that it is a packet not directed to any other wireless computing device on the network?

A. If we assume that, you know, we’re taking a reasonably broad view of what the hotel network is such as all of its access points.

Q. And if we talk about a packet directed to a wired computing device on the same network, in which category would it fall?

A. That would fall into the first category. It does not, in this particular claim, tell you that you should not deliver packets directed to wired networks.

Q. Okay, so since it’s not any other wireless computing device on the network, it would fall into the paragraph, what I would say is paragraph 4 of the claim?

A. The first of the two paragraphs you’ve discussed, yes.

Q. Okay. And the packet would be transmitted?

A. According to this claim, yes.

[Emphasis added; Transcript, pp 676–679.]

[172] Dr. Reiher thus agreed that the language of Claim 1 indicates a packet directed to either another network or to a wired device on the same network would be transmitted by the gateway since such packets would fall into the category of *not directed to a wireless computing device on the network*. In my view, this agreement was a fair and candid one that accords with the language of the claim as drafted. It does not, however, accord with the construction proposed by Guest Tek, in which Claim 1 does not contemplate (and indeed, excludes from consideration) packets directed to a wired device.

[173] In my view, the construction Dr. Reiher and Dr. Lavian ultimately agreed on is the one that accords with the language of the Claim as drafted and as it would be understood by a POSITA. I therefore conclude the fourth descriptive paragraph in Claim 1 of the ’760 Patent requires the gateway to be configured so that packets it has determined to not be directed to another wireless device on the network, which would include those directed to a wired device on the network, are transmitted to the destination address associated with the packet.

> (5) when the packet is directed to a second wireless computing device on the network, the gateway further configured to determine <u>whether communications between the first wireless computing device and the second wireless computing device are allowed</u>, and to either <u>forward the packet</u> to the destination address associated with the packet when communications <u>between the first and second wireless computing devices are allowed</u> or <u>prevent the packet from reaching the destination address</u> when communications between the first and second wireless computing devices are not allowed

[174] The final descriptive paragraph in Claim 1 addresses the other category of packets, those the gateway has determined to be directed to a second wireless device on the network. With respect to these, the gateway makes a further determination, namely whether *communications between the first wireless computing device* (the source of the packet, which I will call

“Computer 1” in the discussion below) *and the second wireless computing device* (the destination of the packet, which I will call “Computer 2”) *are allowed*. If communications between Computer 1 and Computer 2 are allowed, the packet is forwarded; if not, the packet is prevented from reaching the destination, *i.e.*, the packet is dropped.

[175] The primary dispute between the parties with respect to the construction of this paragraph pertains to the meaning of *communications between*, or more particularly, the determination of *whether communications between [the two wireless devices] are allowed*.

[176] Guest Tek contends, with the support of Dr. Reiher, that *communications between* refers to bidirectional communications rather than simply unidirectional communications. On this construction, *communications between* two computers are only *allowed* if Computer 1 is allowed to send packets to Computer 2 *and* Computer 2 is allowed to send packets to Computer 1. If Computer 1 were allowed to send packets to Computer 2, but Computer 2 were *not* allowed to send packets to Computer 1, then on Guest Tek’s construction, *communications between* those computers would not be allowed, such that a packet sent by Computer 1 to Computer 2 would be prevented by the gateway.

[177] Nomadix, on the other hand, views this construction as results-oriented, being designed to avoid Nomadix’s argument that its USG device anticipates the patent. Nomadix argues, with support from Dr. Lavian, that a POSITA would understand the term *communications between* as an ordinary term, either without making a distinction between bidirectional and unidirectional communication, or as encompassing both. Ultimately, on Nomadix’s construction, if Computer 1 is allowed to send packets to Computer 2, the gateway would determine the communication was

allowed, and the packet from Computer 1 would be forwarded to Computer 2, regardless of whether Computer 2 was also allowed to send packets to Computer 1.

[178] There is merit to both Guest Tek and Nomadix's constructions, and each finds some support in the text as purposively construed. On balance, for the reasons below, I adopt Nomadix's construction, and conclude a POSITA would understand the determination of whether communications between the first and second wireless devices are allowed to mean an assessment of whether the sending computer is allowed to communicate with the receiving computer.

[179] In my view, the strongest support for Guest Tek's position is in the use of the word "between" in the claim. Paragraph (5) of Claim 1 refers to assessing whether *communications between* the two computers are allowed, rather than whether *communications from the first computer to the second* are allowed. As Dr. Reiher points out, Claim 1 uses *from…to* language elsewhere to describe packets being directed from Computer 1 to Computer 2, suggesting a different meaning for *communications between*: Reiher First Report, paras 147–152.

[180] Conversely, the strongest support for Nomadix's position is the purpose of the claim. A critical feature of the claim, and the purpose for which packets are directed to the gateway from the wireless access node, is that the gateway makes a determination as to whether to forward a packet or not. Since that determination only relates to whether to send a packet from Computer 1 to Computer 2 or whether to drop it, whether Computer 2 is allowed to send a packet to Computer 1 is irrelevant. Indeed, since both Computer 1 and Computer 2 are by definition

wireless computing devices on the network, any packet Computer 2 might send back to Computer 1 would itself be a *first packet* that would be independently subject to the same determination by the gateway as to whether communications are allowed.

[181] Dr. Reiher gave his view that the POSITA's interpretation of *communications between* would be informed by their knowledge of the security risks of DoS attacks: Reiher First Report, paras 46–49, 142–144. In a DoS attack, computer security can be compromised by an attacker sending messages—a single packet with crafted contents or an overwhelming number of packets—that interfere with the operation of the target computer. Dr. Reiher opined that a POSITA would know that if unidirectional communications were allowed, with Computer 1 being allowed to send packets to Computer 2, but not *vice versa*, Computer 2 would remain vulnerable to DoS attacks from Computer 1. He suggests this supports a reading of *communications between* that checks whether bidirectional communications are allowed rather than just unidirectional communications.

[182] Nomadix argues Dr. Reiher was overly focused on DoS attacks, which are an area of his particular expertise and focus, but which are mentioned only once in the ’760 Patent. There is some merit to this criticism, particularly since the only mention of DoS attacks in the ’760 Patent involves protecting against them not through the determination process described in Claim 1, but through ARP control. At the same time, there is no requirement that an aspect of a POSITA's CGK be expressly referred to in the patent disclosure for it to have bearing on the construction of the patent.

[183] More importantly, though, the concern about DoS attacks does not in my view support one proposed construction or the other. As I understand Dr. Reiher’s evidence on DoS attacks, the concern is with packets being sent from a sending computer (Computer 1) to a receiving computer (Computer 2), where Computer 1 is the attacker. A determination as to whether Computer 1 is allowed to send to Computer 2 is all that is necessary to prevent such an attack. Whether Computer 2 can send to Computer 1 is irrelevant. In this regard, Dr. Reiher’s example of a situation in which Computer 2 is not permitted to send to Computer 1, but remains vulnerable to a DoS attack from Computer 1, is beside the point. The issue is whether the packet should be sent from Computer 1 to Computer 2. Nothing in Dr. Reiher’s explanation shows why it matters in this determination whether Computer 2 is able to send to Computer 1 or not.

[184] I also cannot accept Dr. Reiher’s argument based on reference to Claim 6. That claim requires the gateway to determine whether second packets *are directed to a specific one of the wired and wireless computing devices*. Dr. Reiher argues this shows that the inventors used “from…to” language, not “between” language, when unidirectional communications are at issue: Reiher First Report, paras 155–160. However, Claim 6 does not involve an equivalent determination as to whether communications are allowed. Rather, Claim 6 has a different structure, in which determination is made as to whether the packets are directed to a *specific* computing device, and preventing the packet if that is the case. I therefore do not find reference to Claim 6 assists in resolving the meaning of *communications between*.

[185] I do find some assistance in referring to Claims 2, 3, 22, and 23, which add limitations on how the gateway will determine whether communications between the devices are allowed. In

Claims 2 and 22, this determination is made “according to at least service set identifiers (SSIDs) respectively associated with the first and second wireless computing devices.” In Claims 3 and 23, it is “according to at least virtual local area networks (VLANs) identifiers respectively associated with the first and second wireless computing devices.” These claims specify that both Computer 1 and Computer 2 would have an identifier (either SSID or VLAN) associated with it, and these would be the basis for determining whether communications are allowed. In such a case, whether two computers are allowed to communicate might depend on whether, for example, they were on the same VLAN. This suggests the computers would both be allowed to communicate with each other, or neither would be allowed to communicate with the other.

[186] Additional assistance to resolve this construction issue is found in the disclosure of the ’760 Patent. In particular, paragraph [0033] discusses the assignment of a unique VLAN identifier to each network access node (including wireless and wired nodes). Paragraph [0035] then describes a scenario where the gateway would inhibit traffic between connected devices by recognizing that the source and destination addresses correspond to different VLANs. Again, as with Claims 3 and 23, this involves permissions being effectively granted at the VLAN level, with devices on one VLAN not being permitted to communicate with those on another VLAN, as described at paragraph [0041]. A similar approach using SSIDs based on traffic type is described at paragraph [0040].

[187] My review of the disclosure and claims of the ’760 Patent suggests the inventors were generally contemplating situations in which communications between two computers were either allowed or not allowed, regardless of direction. There is no discussion of any scenario in which

communications are allowed in one direction but not the other. In other words, the term *communications between* appears to have been used (rather than the *from…to* language used elsewhere in Claim 1) simply because the inventors only addressed scenarios in which permissions were the same in both directions.

[188] In my view, the expression *communications between the first and second wireless computing devices are allowed* can support either construction proposed by the parties: it could mean communications back and forth between the devices are allowed or communications from the first to the second computing device are allowed. I am therefore particularly persuaded that a POSITA adopting a purposive construction would look to *why* the determination in question is being made, that is, "the purpose it was intended to execute": *Whirlpool* at para 49(d). Here, the determination is being made to decide whether to forward a packet or not. The material assessment is therefore whether communications between the sending computer and the receiving computer, in the sense of communication from the sending computer to the receiving computer, are permitted. An assessment of whether communications from the receiving computer to the sending computer are allowed is irrelevant. This accords with the construction proposed by Nomadix.

[189] I therefore conclude the fifth descriptive paragraph in Claim 1 of the '760 Patent, as it would be understood by the POSITA, requires that when a packet from a wireless device on the network is directed to a second wireless device on the network, the gateway is configured to determine whether the communication from the first wireless device to the second is allowed, and to transmit or drop the packet accordingly.

[190] Putting the foregoing together, I conclude that the essential elements of Claim 1 of the ’760 Patent can be paraphrased as follows, with terms being understood as set out above:

(A) a *network* that includes a *wireless access node* that is *coupled* to a *gateway* and *configured* to (i) receive *packets* from *wireless computing devices* including at least some that are encrypted through an *encrypted wireless protocol*, and (ii) transmit *all* such packets to the gateway regardless of their destination address;

(B) the gateway being *configured* to (i) determine with reference to the destination address whether each packet is *directed to another wireless device on the network* or not, and (ii) if it is not, transmit the packet; and (iii) if it is, then (a) transmit the packet if *communications between* the wireless devices on the network are *allowed*, and (b) drop the packet if *communications between* the wireless devices are *not allowed*.

(6) Remaining Asserted Claims

[191] The other network claims asserted by Guest Tek raise no contentious issues of construction, and in particular no issues that are relevant to the parties’ arguments on infringement or validity.

[192] Claim 4 provides for a network of claim 1 or 2, with the added limitation that the *plurality of traffic types* referenced in claim 1 includes *WPA traffic, WPA2 traffic, AES traffic, WEP traffic, and VoIP traffic*. There was no dispute that these various types of network traffic represent either forms of encryption (WPA, WPA2, AES (Advanced Encryption Standard),

WEP) or traffic related to the application of internet telephony (Voice over Internet Protocol or VoIP).

[193] Claim 10 adds a limitation that the gateway is further configured to *perform network address translation* (NAT) to facilitate transmission of at least some of the first packets. I agree with Dr. Reiher that the POSITA would understand *network address translation* in the context of this claim of the ’760 Patent to mean the translation of non-routable IP addresses from a private set of IP addresses only valid on the network to routable IP addresses as the packets move from the local network to the internet: Reiher First Report, paras 198–199. Dr. Reiher and Dr. Lavian agreed that NAT was a common, though not universal, feature of gateway devices.

[194] Claim 11 adds a limitation in which the gateway is further configured to *supply locally valid network address[es]* to at least some of the wireless devices to facilitate transmission of at least some of the first packets. Again, I agree with Dr. Reiher that this term would be understood by the POSITA to mean the gateway is configured to provide an IP address usable on the local network (either a private IP address or from the normal IP address space) to a wireless device connecting to the network: Reiher First Report, paras 207–214. Dr. Reiher and Dr. Lavian agreed that this was a common feature of network configuration.

[195] The parties agreed that the assessment of infringement and validity of Claims 4, 10, or 11 would be the same as for Claim 1 as they relate to Nomadix’s actions and products. In other words, the parties’ infringement and validity arguments are “all or nothing” with respect to Claims 1, 4, 10, and 11.

[196] Claim 21 of the ’760 Patent is the independent method claim, which reads as follows:

> A method comprising:
>
> receiving, by a wireless access node, first packets from a plurality of wireless computing devices attempting to access a network, each of the first packets corresponding to one of a plurality of traffic types, and at least one of the traffic types corresponding to an encrypted wireless protocol;
>
> transmitting all first packets received from the wireless computing devices to a gateway on the network regardless of destination addresses associated with the first packets;
>
> determining by the gateway whether each packet is from a first one of the wireless computing devices directed to any other of the wireless computing devices on the network with reference to at least a source address and a destination address associated with the packet;
>
> transmitting the packet to the destination address associated with the packet when the packet is not directed to any other wireless computing device on the network; and
>
> when the packet is directed to a second wireless computing device on the network, determining whether communications between the first wireless computing device and the second wireless computing device are allowed, and either transmitting the packet to the destination address associated with the packet when communications between the first and second wireless computing devices are allowed or preventing the packet from reaching the destination address when communications between the first and second wireless computing devices are not allowed.

[197] As can be seen, much of the structure and language of Claim 21 parallels that of Claim 1. I agree with the parties that the terms as used in Claim 21 would be understood by the POSITA in the same way they would be understood in Claim 1 as set out above. However, in addition to the fact that Claim 21 is a method claim rather than a network or system claim, I note two differences between the substantive aspects of the claims.

[198] First, Claim 21 does not specify that the wireless access node is *coupled* to the gateway. While I have not accepted Dr. Lavian's suggestion that the wireless access node must be directly coupled to the gateway, I note that in any case, that issue would be irrelevant to the question of construction or infringement of Claim 21.

[199] Second, the method of Claim 21 does not specify the wireless access node is *configured to transmit* all first packets to the gateway regardless of destination address. Rather, the method simply requires that all first packets *be transmitted* [from the wireless access node] to the gateway regardless of destination address.

[200] Claims 30 and 31 add the same limitations to the method claim(s) that are added to the network claims in Claims 10 and 11, namely *performing network address translation* and *supplying locally valid network address[es]*. There is no dispute that these would be understood in the same way described above.

E. *Inducing Infringement*

[201] Having addressed the construction of the asserted claims of the '760 Patent, I turn to the allegation that Nomadix has induced Canadian hotels to infringe the patent.

[202] In closing written submissions, Guest Tek argued that Nomadix "uses the '760 Patent and is also responsible for inducing infringement" [emphasis added]: Plaintiff's Memorandum of Fact and Law, para 76. To the extent this is an argument Nomadix is itself infringing the '760 Patent, this goes beyond the action as Guest Tek pleaded it. The only infringement of the

’760 Patent pleaded is that Nomadix induced hotels in Canada to infringe the patent: Amended Statement of Claim, paras 1(b), 71–94. In any event, no evidence was led to establish Nomadix itself made, used, constructed, or sold to others to be used, a network that includes, in particular, a *wireless access node*, regardless of any other infringement issues. I will therefore focus on Guest Tek’s claim as pleaded.

[203] As set out above, to establish Nomadix has induced infringement of the asserted claims of the ’760 Patent, Guest Tek must show (1) there has been infringement of the claims by a direct infringer; (2) the acts of infringement were influenced by Nomadix’s acts to the point that, without the influence, direct infringement would not have taken place; and (3) Nomadix knew its influence would result in the acts that constitute infringement: *Corlac* at para 162. Given the bifurcation of this proceeding, what is at issue at this stage is not how much infringement has occurred or been induced, simply whether there has been infringement: *Hospira* at para 26.

(1) Infringement of the Claims by a Direct Infringer

[204] Guest Tek claims direct infringement of the ’760 Patent occurs when a Nomadix gateway is connected to a wireless access point in a network operated by a hotel. Guest Tek did not call any witnesses from Canadian hotels or from those authorized to sell and install Nomadix gateways in Canada to speak to network configuration in Canadian hotels. Rather, Guest Tek relied primarily on the functionality of the Nomadix gateway itself, as well as evidence from Guest Tek and Nomadix witnesses regarding their understanding of Canadian hotel installations. Particular focus was put on Marriott hotels, which must comply with the requirements in Marriott’s GPNS.

[205] There is no question that Nomadix gateways with the NSE software are installed in hotels in Canada. For the gateway to function, a hotel or other customer must enter a license agreement with Nomadix and obtain a functioning license key. Numerous such keys have been issued to Canadian hotel customers: Nomadix Confidential Exhibit 38. There is also no real dispute that there are hotels in Canada that have networks that comprise a Nomadix gateway and a wireless access node that can and does receive traffic from wireless devices: Transcript, pp 825–826, 829, 1422; Exhibit 37, p 121.

[206] Guest Tek argues that in at least some of these networks, all first packets from wireless devices are forced to the gateway using a process involving proxy ARP, and that such packets will be (a) forwarded to a destination wireless device on the network if bidirectional communications between them are allowed, which occurs when “intra-port communication” is disabled and both devices have the proxy ARP setting enabled; or (b) forwarded to their destination if the destination device is not on the network. Guest Tek states that the tests it performed on Nomadix gateway devices and the NSE software source code confirm Nomadix devices function in this way.

[207] Given the nature of Guest Tek’s allegations and Nomadix’s response to them, I will first review aspects of the NSE software’s functionality on the Nomadix gateway devices and the testing Guest Tek performed, before turning to whether the devices meet the elements of the asserted claims. While Nomadix gateway devices running the installed NSE software have a wide variety of different features and settings, the following discussion focuses on those relevant to Guest Tek’s allegations of infringement of the ’760 Patent. Nomadix agrees that the relevant

functionalities are the same across the AG and EG gateway models and NSE software versions (8.11 and higher) in respect of which Guest Tek asserts its claims: Exhibits 108, 109; Reiher First Report, paras 350–352.

(a) *NSE Software: Subscribers, Devices, proxy ARP, and intra-port communication*

[208] The NSE software classifies computers connected to the network as “subscribers” or “devices.” Since the terms “subscriber” and “device” are common and can have different meanings depending on context, I will adopt the parties’ convention of using the capitalized terms Subscriber and Device to refer to their status in the NSE software. When a new computer (including a wireless phone) joins the network, such as a guest in a hotel setting, they are by default designated as a Subscriber, rather than a Device. However, an administrator can change the classification of a computer by selecting that option in a menu: Transcript, pp 1487–1490; Exhibit 75, p 4.

[209] When a computer is designated as a Device, the NSE software allows the administrator to select the option “Proxy Arp For Device.” As discussed above, proxy ARP is a functionality by which a device will send a response to an ARP request providing its own MAC address as corresponding to the IP address queried in the ARP request. In the NSE software, when Proxy Arp For Device is enabled, that device is able to communicate with other Devices on the hotel network that have Proxy Arp For Device enabled: Transcript, p 1480. Subscribers cannot have Proxy Arp For Device enabled.

[210] Another setting available in the NSE software is “Subscriber intra-port communication.” Intra-port communication is communication between computers on the same port location, for example a conference room. When this setting is enabled for a port, Subscribers on that port can communicate with each other without intervention from the gateway. In such cases, the gateway does not respond to ARP requests from a Subscriber seeking MAC addresses of other Subscribers on the same port location: Exhibit 104, pp 152–153. The default setting in the NSE software for Nomadix gateways is for Subscriber intra-port communication to be disabled.

[211] The impact of the foregoing settings is that when Subscriber intra-port communication is disabled (as it is by default), communications between two computing devices on the network are prevented unless they are both designated as Devices and both have Proxy Arp For Device enabled. This is seen in testing undertaken by Guest Tek.

(b) *Testing using a Nomadix gateway*

[212] Guest Tek ran three tests of different versions of the NSE software running on different Nomadix AG and EG gateway devices: Reiher First Report, paras 251–347 and Appendices 3–4. In the tests, Guest Tek set up a network containing an Aruba wireless access point connected to a switch that was in turn connected to a Nomadix gateway. The gateway was also connected to a web server. Connected to a different port on the switch was an administrator laptop that (a) allowed the user to interface with and configure the NSE software; and (b) ran a program called Wireshark, which was able to monitor and record packet traffic flowing to and from the wireless access point as that traffic was mirrored onto the laptop’s port: Reiher First Report, para 254.

[213] The network setup used for the testing included turning off a feature on the Aruba wireless access point that would automatically allow two wireless devices connected to it to directly communicate with each other, so as “to isolate the user laptops from one another similar to as would be done in a hotel environment”: Reiher First Report, para 314; Lavian Second ’760 Report, para 35; Transcript, pp 465–466. Disabling port-to-port traffic on each access node by the means provided in the device is one of the ways the ’760 Patent suggests to ensure all traffic is passed to the gateway: ’760 Patent, para [0041].

[214] The tests had a number of steps, which were the same for each software/gateway configuration. Early steps in the test included wirelessly connecting two “user” laptops to the wireless access point using a WPA2 password and logging them into the Nomadix gateway to allow them to access the web server. The results of the tests were set out in Dr. Reiher’s first report and were not themselves contested, although there was dispute over what they signified. By way of summary, the tests showed that in each configuration:

- the user laptops were each able to connect with the web server and download a file;
- packets from each user laptop that were destined for the web server’s IP address were sent to the Nomadix device’s MAC address, and then transmitted to the web server by the gateway;
- the gateway performed dynamic address translation to correlate two Transmission Control Protocol (TCP) ports with the two wireless laptops so that traffic from the web server that was returned to the gateway was directed to the correct laptop;

- with Subscriber intra-port communication disabled and each user laptop designated as a Subscriber:
    - a “ping” request from each user laptop addressed to the other user laptop was transmitted to the gateway, but not then transmitted to the other laptop – neither laptop could ping the other;
    - the gateway responded to ARP requests sent by each user laptop in respect of the other by providing its own MAC address, performing a proxy ARP function;
- when each user laptop was designated as a Device and the Proxy Arp For Device option was enabled for each:
    - a ping request from each user laptop addressed to the other user laptop was transmitted to the gateway, and then transmitted to the other laptop – each laptop could ping the other;
    - the gateway responded to ARP requests sent by each user laptop in respect of the other by providing its own MAC address, performing a proxy ARP function;
- when one user laptop was designated as a Subscriber and one was designated as a Device with the Proxy Arp For Device option enabled, neither laptop could ping the other;
- when each user laptop pinged random IP addresses, the laptop sent out ARP requests which the gateway responded to with its own MAC address, causing the laptop to store the gateway’s MAC address in its ARP table in association with the IP addresses; and

- when the gateway was unplugged from the network, the laptops could not ping each other, as a result of the wireless access point having been configured as described in paragraph [213] above.

[215] Dr. Reiher pointed to aspects of the NSE source code that [redacted]. In particular, he identified truth tables that dictate when a Nomadix gateway will perform an ARP response [redacted]. I do not need to address these aspects of the source code in depth, as there was ultimately little dispute that the [redacted]. Rather, Nomadix argues that this behaviour does not constitute direct infringement by hotels, and is not in any event induced by Nomadix.

[216] I note that Guest Tek's testing did not test whether the gateway would send a proxy ARP response to an ARP request sent from a Device to another Device where one or the other did not have Proxy Arp For Device enabled. In such circumstances, the truth tables suggest that a proxy ARP response [redacted]: Reiher First Report, para 357; Transcript, pp 1709–1711.

[217] Having considered these aspects of the Nomadix gateway/NSE functionality, I turn to the essential elements of the claims of the '760 Patent to assess whether Guest Tek has established direct infringement of those claims.

(c) *Essential elements of Claim 1*

(i) *gateway, wireless access node, first packets*

[218] As noted above, there is little dispute that there are Canadian hotels that have in place a *network* comprising a Nomadix *gateway* and a *wireless access node* that is *coupled to the gateway*. For the reasons set out in construing the *coupled to the gateway* term, it matters not whether the wireless access node is connected directly to the gateway or via intermediate switches or other devices.

[219] I am also satisfied that hotels in Canada that have set up such a network would use a wireless access node that is *configured to receive first packets from a plurality of wireless computing devices attempting to access the network*. Indeed, this is the very purpose of having a wireless access node. Marriott’s GPNS, which is a global standard that would be used in Canadian Marriott hotels using a Nomadix gateway (of which there are several), confirms that their hotel networks are designed to accommodate guest wireless devices: Exhibit 51, p 5; Exhibit 54, pp 1, 3.

(ii) *encrypted wireless protocol*

[220] The evidence is less clear that there are networks in hotels in Canada in which the wireless access node is *configured to receive* traffic that includes traffic *corresponding to an encrypted wireless protocol* such as WEP, WPA, or WPA2. A wireless access point can be

configured to receive encrypted traffic such as WPA or to have an open network: Transcript, pp 137, 687–688.

[221] To establish this element, Guest Tek first points to Marriott's GPNS, which requires SSL encryption: Transcript, p 346. However, as I have construed this term, SSL would not be considered an *encrypted wireless protocol*. Guest Tek also asserts that Marriott requires WPA/WPA2 encryption for VoIP phones: Transcript, pp 344–345. However, the relevant section of the GPNS appears to only require WPA/WPA2 encryption "where a given property has deployed wireless VOIP": Exhibit 51, p 37; Transcript, p 1346. Guest Tek did not establish that there were Canadian Marriott hotels that had deployed wireless VoIP. If anything, the evidence pointed to the contrary: Transcript, pp 1346–1351.

[222] Dr. Reiher pointed to Guest Tek's testing, in which wireless devices were connected to the wireless access node via WPA2 encrypted wireless protocol: Reiher First Report, paras 257, 432–433. However, how Guest Tek chose to configure communications between its test laptops and its test wireless access node tells us nothing about how wireless access nodes in Canadian hotels that use Nomadix gateways are configured. Dr. Reiher also asserted that Nomadix's User Guides "instruct" the use of encrypted wireless protocols, but had to admit in cross-examination that the references he gave were simply to a Glossary of Terms found at the end of the User Guide: Reiher First Report, paras 434–435; Exhibit 104, p 269; Transcript, pp 688–689.

[223] Nonetheless, I am satisfied on a balance of probabilities that there are networks in hotels in Canada with a Nomadix gateway and a wireless access node configured to receive traffic that

includes WEP, WPA, or WPA2 traffic. Both experts recognized the particular concerns associated with wireless security and the potential for others to “listen in” on unencrypted wireless data: Transcript, pp 425–428; 1627–1628, 1654–1655. There was evidence that at least certain wireless access points do not need any special configuration in order to support wireless encryption protocols: Transcript, pp 288–290. While Dr. Lavian stated that in his experience, “many hotels in many places” have open (unencrypted) WiFi networks, he also recognized that a hotel’s wireless access node can be configured to operate using various traffic types including wireless encryption protocols: Lavian Second ’760 Report, para 53 (p 21). Given the number of Canadian hotels shown to use Nomadix gateways, I am prepared to infer that some number of these have a network that uses a wireless access node configured to receive encrypted wireless traffic. I note that in his report Dr. Lavian did not point to the absence of encrypted wireless traffic being sent from wireless devices to the wireless access node as a reason Claim 1 of the ’760 Patent was not infringed: Lavian Second ’760 Report, para 53 (p 18).

[224] That said, for there to be direct infringement of Claim 1, all of the essential elements of Claim 1 must co-exist in the same network. A number of Guest Tek’s arguments regarding other elements of the claim rely on evidence relating to Marriott hotels in particular, based on the Marriott GPNS. To the extent such evidence may show that these elements are present in Marriott hotels in particular, the other elements of Claim 1 must also be present in those Marriott hotels in order to establish infringement. While I am prepared to infer that there are hotels in Canada that have a Nomadix gateway with a wireless access node configured to receive traffic associated with an encrypted wireless protocol, I conclude the evidence does not establish that there are such networks specifically in Marriott hotels in Canada.

(iii) *all first packets*

[225] I am also not satisfied that Guest Tek has established there are hotels in Canada that have a network with a Nomadix gateway device in which the wireless access node is *configured to transmit all first packets received from the wireless computing devices to the gateway on the network regardless of destination addresses*. To establish this element, Guest Tek relies on its construction that the wireless access node may "receive help" from other devices, such that the forcing of packets to the gateway may effectively be done by operation of proxy ARP performed by the gateway. I have rejected this construction. The claim requires that the wireless access node *be configured* to transmit all first packets to the gateway, not that communication that effectively occurs between the gateway and the end user's wireless device *results* in the packets being addressed to the gateway.

[226] Further, setting up Guest Tek's test network involved configuring the wireless access node in a particular manner, namely by turning off a feature allowing wireless devices to communicate with each other. Dr. Reiher underscored that this did not, in itself, tell the wireless access node where to send packets, beyond not sending them to another wireless device: Transcript, pp 465–466; Reiher First Report, para 314. In my view, this evidence indicates that (a) some configuration of the wireless access node itself was necessary to obtain the results shown in Guest Tek's testing; and (b) this aspect of the configuration was not alone sufficient to result in all packets being forced to the gateway. Guest Tek does not point to any evidence showing hotels in Canada have undertaken such a configuration of wireless access nodes in networks that include a Nomadix gateway.

[227] Dr. Reiher’s report suggested that the configuration of the wireless access point in the testing was “as would be done in a hotel environment”: Reiher First Report, para 314. However, Dr. Reiher’s expertise was not in hotel networks in particular, and he admitted to having no firsthand knowledge of hotel configurations: Transcript, pp 592–593, 715–716. While the Marriott GPNS requires that traffic from a wireless guest room to a wireless or wired guest room be blocked, there was no evidence that would allow me to infer that networks in Marriott hotels would necessarily configure the wireless access node in the manner done in Guest Tek’s testing or in any other manner that transmits all packets to the gateway regardless of destination address: Exhibit 54, p 6; Transcript, pp 1357–1359. Similarly, while some access points may be pre-configured to not allow traffic between mobile devices, there was no evidence such devices or configurations were used in Canadian hotels: Transcript, pp 141, 145.

[228] I am also not satisfied Guest Tek’s testing showed *all first packets* were or would be forced to the gateway by operation of either the Nomadix gateway or a combination of the wireless access node and the Nomadix gateway. The testing did show that when the source and destination were both Subscribers, or both Devices with Proxy Arp For Device enabled, the gateway responded to ARP requests with a proxy ARP response. The wireless device subsequently addressed packets destined for the same IP address to the gateway’s MAC address. However, it did not show how or why the ARP requests from the user devices were themselves being forced to the gateway before the ARP response, something Dr. Reiher said was essential for the claimed system to work properly: Transcript, p 435. Nor did it show how or whether the gateway or the wireless access node would force to the gateway packets that were not addressed to the gateway’s MAC address, such as packets for which the sender’s computer already knows

the MAC address of the destination and therefore would not send an ARP request. It also did not show how or whether packets would be forced to the gateway in a circumstance where the NSE software's truth tables did not result in a proxy ARP response (such as where both source and destination are Devices but did not both have Proxy Arp For Device enabled).

(iv) *directed to any other of the wireless computing devices*

[229] I also conclude Guest Tek has not established that any Nomadix customer in Canada has its gateway *configured to determine, for each packet of the first packets received from the wireless access node, whether the packet is […] directed to any other of the wireless computing devices on the network*. Nomadix's gateway does not itself distinguish whether a packet is destined for a wireless or wired device on the network: Lavian Second '760 Report, paras 49–50. There was also no evidence Nomadix gateways on hotel networks in Canada are ever configured to determine—by differentiation based on the tunnel of traffic, the use of SSIDs, or otherwise—if devices are connected to the network via a wireless access node and to make distinctions on that basis. Guest Tek's argument with respect to infringement of this element of the claim relied only on the Nomadix gateway determining whether the packet is "from one subscriber device to another subscriber device," without any assessment of whether it is being sent to a wireless subscriber device: Reiher First Report, paras 445–447; Reiher Third Report, para 18.

[230] For the same reasons, I conclude Guest Tek has not established any Nomadix customer in Canada has its gateway configured as described in paragraphs (4) and (5) of Claim 1, namely to *transmit the packet to the destination address […] when the packet is not directed to any other wireless computing device on the network* and to transmit or drop depending on whether

communications are allowed *when the packet is directed to a second wireless computing device on the network.*

[231] As Dr. Reiher agreed in cross-examination, the gateway of Claim 1 is configured to treat a packet directed to a wired device on the network in the same way as a packet directed to a destination off the network, forwarding that packet regardless of whether communications are allowed. Packets directed to a wireless device on the network, on the other hand, are forwarded only when communications between the devices are allowed: Transcript, pp 675–679.

[232] In the network Guest Tek tested, packets directed to a destination off the network (the web server) were forwarded to their destination. However, there was no evidence of the Nomadix gateway being configured to forward all communications to wired devices on the network. To the contrary, the only evidence is that the Nomadix gateway would treat a packet directed to a wired device on the network in the same way as a packet directed to a wireless device on the network, unlike the gateway of Claim 1.

(v) *communications between*

[233] I do accept that the Nomadix gateway is configured to determine whether *communications between* the sending device and the destination device are *allowed* as required by paragraph (5) of Claim 1. The NSE software effectively prevents traffic from being transmitted to the destination device on the network unless communications between them are allowed. Communications between the devices are allowed when each is designated as a Device and has Proxy Arp for Device enabled. While this may be somewhat cumbersome as a method

for allowing communications between devices, there is no limitation in Claim 1 as to the mechanism by which the system indicates that communications are allowed. Further, contrary to Nomadix’s arguments, it does not in my view matter whether a Canadian hotel has ever actually “allowed” such communications by designating its guests’ wireless devices as Devices and enabling Proxy Arp For Device. What matters is that the Nomadix gateway is *configured* to make a determination regarding forwarding a packet or not depending on whether communications are allowed or not. The truth tables the NSE software implements perform that determination.

[234] However, as noted above, the Nomadix gateway performs this determination and forwards or drops packets in consequence not only *when the packet is directed to a second wireless computing device on the network*, but also when it is directed to a second wired computing device on the network. The result is that the gateway is not *configured to transmit the packet to the destination address […] when the packet is not directed to any other wireless computing device on the network*, as required by paragraph (4) of Claim 1. Rather, it transmits the packet to the destination address when the packet is not directed to any other wireless computing device on the network *only* when the destination is not on the network *or* when communications between the source and the destination on the network are allowed as a result of both being Devices with Proxy Arp for Device enabled.

[235] I therefore conclude Guest Tek has not established there has been direct infringement of Claim 1 of the ’760 Patent by Canadian hotels using a Nomadix gateway device running the NSE software.

(vi) Other asserted claims

[236] Claims 4, 10, and 11 are dependent on Claim 1. The foregoing conclusion therefore means there is also no direct infringement of the remaining asserted network claims of the ’760 Patent, even though the additional limitations in these claims are satisfied. The conclusion above regarding *encrypted wireless protocol* traffic means that the additional limitation of Claim 4 (*traffic types include WPA, WPA2, AES, WEP and VoIP traffic)* is present. Dr. Reiher’s evidence supports the conclusion that the additional limitations of Claim 10 (*network address translation*) and Claim 11 (*locally valid network address*) are also met and this conclusion was not challenged by Dr. Lavian or Nomadix: Reiher First Report, paras 472–479.

[237] With respect to the asserted method claims (Claims 21, 30, and 31), infringement requires performance of each step in the method: *Western Oilfield (FCA)* at para 48. I have described above the differences between the independent network claim of Claim 1 and the independent method claim of Claim 21. Claim 21 does not require the *wireless access node* to be *configured* to transmit all packets to the gateway as Claim 1 does. Therefore, while I concluded that direct infringement of Claim 1 was not established in part due to the configuration element, this conclusion does not apply to Claim 21. Nonetheless, my conclusions that the following essential elements of Claim 1 have not been shown to be met by a network comprising a wireless access node and a Nomadix gateway also apply to Claim 21 (in the list below, I use the language of Claim 21):

- *at least one of the traffic types corresponding to an encrypted wireless protocol*, to the extent that this element must be found in conjunction with all other elements in the same system (see paragraph [224] above);

- *transmitting all first packets [...] to a gateway* (paragraph [228]);
- *determining by the gateway whether each packet is […] to any other of the wireless computing devices on the network* (paragraph [229]);
- *transmitting the packet to the destination address associated with the packet when the packet is not directed to any other wireless computing device on the network* (paragraphs [230]–[234]).

[238] In the context of Claim 1, I concluded there is no requirement that a hotel actually allow *communications between* two devices, as the gateway need just be *configured* to make the claimed determination rather than actually make the determination. That conclusion is less clear in the method claim, which requires that the method include *either transmitting the packet to the destination address […] when communications between the first and second wireless computing devices are allowed or preventing the packet from reaching the destination address when communications between the first and second wireless computing devices are not allowed*. If no communications have ever been allowed between two wireless devices on a Nomadix network in Canada, which would require designating them as Devices and enabling Proxy Arp For Device, it is not clear this aspect of the claimed method has been performed. However, I need not decide this question given my determinations on the other elements of Claims 1 and 21.

[239] As for Claims 30 and 31, these claims add the same limitations to the method claims that Claims 10 and 11 add to the network claims. For the same reasons given, they add no further unmet limitations, but are not infringed because they are dependent on, and include the essential elements of, Claim 21.

[240] For the foregoing reasons, I conclude Guest Tek has not established direct infringement of the asserted claims of the ’760 Patent in any network in Canada using a Nomadix gateway device.

(2) Influence by Nomadix

[241] The second question in the test for inducement is whether Nomadix influenced hotels to undertake the infringing acts to the point that, without the influence, direct infringement would not have taken place. Evidently, if there is no direct infringement of the ’760 Patent by Canadian hotels using Nomadix devices, Nomadix cannot have influenced those hotels to infringe the patent. In addition, I conclude Guest Tek has not established Nomadix influenced Canadian hotels to undertake the steps Guest Tek alleges infringe the patent.

[242] I begin by observing that I attach little weight to the report of Dr. Reiher on the issue of influence. The existence and consequences of influence are largely factual issues on which no scientific expertise is needed. Dr. Reiher’s report on these issues does not apply his scientific expertise. Rather, it gives his understanding on issues such as whether diagrams in Nomadix’s user manuals and guides influence Canadian users, and the meaning and impact of contractual terms in Nomadix’s end user license agreement (EULA): Reiher First Report, paras 493–517. These are matters Dr. Reiher is in no better position to assess than the Court, and Dr. Reiher steps into the role of advocate in seeking to make arguments or determinations on essential factual and legal issues that fall outside his expertise.

[243] It is clear that Canadian hotels themselves desire wireless networks to provide wireless internet access to their guests. It is thus not Nomadix that “induces” them to create a wireless network. However, in my view this does not preclude Nomadix from being potentially liable for inducing infringement if they influence the hotels to build their wireless networks with the particular configuration set out in the ’760 Patent. In other words, the issue is not whether Nomadix induced Canadian hotels to build a network, or even a wireless network, but whether it induced them to build an infringing wireless network.

[244] I accept that Nomadix has influenced customer hotels in Canada to purchase Nomadix gateways to use them in their hotel networks. The fact that such a sale is indirect through a distributor does not affect this, particularly given the evidence of Nomadix’s direct involvement in marketing, software licensing, software updates, and support in Canada: Exhibit 55, pp 4–5; Exhibit 80; Transcript, pp 1360, 1367–1368, 1522–1525. I also accept that the purchase and use of a Nomadix gateway necessarily involves the use of the NSE software, which must be licensed directly from Nomadix for the gateway to function: Transcript, pp 1379–1381.

[245] However, as I have concluded above, Claim 1 of the ’760 Patent requires the *wireless access node* be *configured to transmit* to the gateway all packets received from wireless devices, regardless of the packet’s destination. Guest Tek has not established Nomadix influenced Canadian hotels to configure their wireless access nodes in this way or any particular way. As set out above, Guest Tek’s own test network included manual configuration of the Aruba wireless access point to prevent wireless devices from communicating directly with each other: Reiher First Report, para 314. Dr. Reiher stated that for the invention to work as described, the wireless access point had to be configured to always forward ARP requests to the gateway: Transcript,

p 435. But Guest Tek pointed to no Nomadix documentation or other evidence in which Nomadix encouraged users to configure the wireless access node in this manner. Rather, it relied exclusively on the gateway's use of proxy ARP as the means of forcing packets to the gateway.

[246] This concern applies equally to the method claim of Claim 21, even though the method does not require that the wireless access node be *configured* to transmit all first packets to the gateway. The first packets must still all be transmitted from the wireless access node to the gateway in Claim 21, a result that was only shown through some degree of configuration of the wireless access node. I am prepared to accept that the default configuration of the Nomadix gateways, which sends proxy ARP responses to ARP requests received from Subscribers, is sufficient to constitute influence to use this configuration on the gateway. However, this gateway configuration is not sufficient to achieve this essential element of the method of Claim 21, and there is no indication Nomadix influenced hotels to undertake the remaining steps necessary to configure a wireless access node to perform in a particular way.

[247] Nor is there evidence Nomadix influences Canadian hotels to build a network, or use a method, in which the wireless access node is configured to, or does, receive traffic corresponding to an *encrypted wireless protocol*. I have concluded above that Marriott's GPNS do not establish the likelihood of use of WPA or WPA2 encryption on a wireless network comprising a Nomadix gateway in a Marriott hotel. In any event, there is no indication Nomadix influenced Marriott to include any wireless encryption requirements in its GPNS. Nor is there any evidence Nomadix encouraged or otherwise influenced any customer to set their wireless access node up in a way to receive encrypted wireless protocol traffic, or even that Nomadix knew whether they did. It is

worth reiterating that the passages in Nomadix’s User Guides that Dr. Reiher suggested “instruct” the use of encrypted wireless protocols were simply references to the definitions of WEP and WPA in the glossary: Reiher First Report, paras 434–435; Exhibit 104, p 269; Transcript, pp 688–689. These give no indication Nomadix influenced hotels to use these encryption protocols in their networks. Indeed, the glossary is the only place in the User Guide where the terms are found.

[248] I also conclude Guest Tek has not established Nomadix influenced Canadian hotels to configure a Nomadix gateway so it could *determine whether* each first packet received from the wireless access node was *directed to any other wireless device on the network*. There was considerable discussion of the potential to allow communications between user devices by designating them as Devices and enabling Proxy Arp For Device, and whether Nomadix encouraged or suggested this. As I have noted, the ability to undertake such a designation is present in the Nomadix device and the NSE software, even if no hotel ever chooses to allow such communications. However, there was no evidence that Nomadix suggested or otherwise influenced Canadian hotels to create a network in which the gateway is able to make a determination between packets destined for a wireless device on the network and those not destined for a wireless device on the network. As described in paragraph [229], there is no evidence a Canadian hotel has configured their network to distinguish source and destination devices by mode of access, even assuming this can be done with a Nomadix gateway. There is similarly no evidence Nomadix influenced hotels to do this even if it were established that it had been done.

[249] For the same reasons, Guest Tek did not establish Nomadix influenced any Canadian customer to configure their gateway to transmit packets to their destination address when the packet is not directed to any other wireless device on the network and to only forward the packet to the destination when it is directed to another wireless device on the network when communications between the wireless devices are allowed.

[250] As the evidence does not establish Nomadix influenced Canadian hotels with respect to these issues, it equally does not establish any direct infringement would not have occurred but for that influence. I therefore conclude Guest Tek has not made out the second requirement of a finding of inducing infringement of the ’760 Patent. This conclusion applies to both the independent network Claim 1 and the independent method Claim 21, as well as each of the dependent claims.

(3) Knowledge

[251] As neither of the first two elements of the test for inducing infringement are met, I need not address the question of Nomadix’s knowledge that its influence would induce Canadian hotels to infringe. Since there has been no demonstrated direct infringement, and no demonstrated influence, there can be no conclusion that Nomadix knew that the influence would lead to the acts constituting infringement.

[252] I therefore conclude Guest Tek has not established Nomadix has induced infringement of the asserted claims of the ’760 Patent.

F. *Validity*

[253] Nomadix argues that if Guest Tek's construction of the ’760 Patent is accepted, then the patent was anticipated by its USG gateway device running version 3.08.105 of the NSE software. It also argues the ’760 Patent was anticipated or is rendered obvious by the prior art, notably an article by Henry Haverinen, and the CGK of the POSITA. For the reasons below, I conclude Nomadix has not met its onus to demonstrate the ’760 Patent is invalid.

[254] Both parties agreed that the validity of the ’760 Patent was an “all-or-nothing” proposition in that the validity or invalidity of all claims follows that of Claim 1. I will therefore focus the following discussion on Claim 1 of the ’760 Patent, but my conclusions apply equally to all of the claims.

(1) Anticipation

(a) *USG gateway with NSE version 3.08.105*

[255] Nomadix's first anticipation argument is conditional. It argues if the Court accepts Guest Tek's construction of the ’760 Patent and finds the AG and EG gateway devices running NSE version 8.11 or higher infringe the patent, then the same construction means its USG device anticipated the patent. In particular, Nomadix argues its USG device, which was commercialized before the priority date of the ’760 Patent, had the same proxy ARP function, Subscriber/Device distinction, and Proxy Arp For Device feature found in the current AG and EG devices.

[256] As set out above, I have not accepted Guest Tek's construction with respect to a number of claim elements. This includes the requirement of forcing all first packets to the gateway, and the requirement the gateway determine whether a first packet is directed to another wireless device on the network or to a destination that is not another wireless device on the network. I concluded these elements are not present in the Nomadix AG and EG gateway despite the presence of the proxy ARP function and the Proxy Arp For Device feature.

[257] As a result, the presence of these features on the USG device did not anticipate the '760 Patent, regardless of whether those features were fully disclosed and enabled. I therefore conclude the USG device did not anticipate the '760 Patent as it did not provide enabling disclosure of all of the essential elements of the '760 Patent.

[258] Even if I had accepted Guest Tek's construction, I would have nonetheless concluded the USG device did not anticipate the '760 Patent. Guest Tek argued the use of a Nomadix AG or EG gateway in a network with a wireless access node infringed the '760 Patent by virtue of (a) the gateway sending proxy ARP responses to any ARP request from a Subscriber or Device when Subscriber intra-port communication is disabled; (b) the gateway transmitting packets from a Subscriber or Device on the network to any destination not on the network; and (c) the gateway only transmitting packets directed to devices on the network when communications between the source and destination are allowed by designating both source and destination as a Device and enabling Proxy Arp For Device on both. The differences between the USG's ARP response protocol and that of the AG and EG devices, as set out below, mean it would not have anticipated the '760 Patent even on Guest Tek's construction.

[259] Guest Tek’s testing shows the USG running NSE 3.08.105 did not send proxy ARP responses to ARP requests sent by a device designated as a Subscriber to another Subscriber. The result is not all subsequent packets from such a wireless device would be addressed to and transmitted to the gateway. Therefore, even if one were to accept that proxy ARP in the AG and EG gateways results in *all first packets* being transmitted to the gateway, which I have not, the USG did not use this method in respect of all of the packets that the AG and EG gateways do, and thus did not do so for *all first packets*. For clarity, as set out at paragraph [228], Guest Tek has also not shown that the AG and EG gateways running NSE 8.7 or above result in *all first packets* being forced to the gateway. However, given the additional cases in which the USG also does not send a proxy ARP response, I would not conclude its use in a wireless network anticipated the ’760 Patent even if I had concluded that Nomadix had induced infringement.

[260] Conversely, I am not convinced by Guest Tek’s argument that the USG device can be distinguished from the ’760 Patent because the USG determined whether *communications between* devices were allowed based solely on whether unidirectional communication was allowed. The USG device forwarded a packet to a destination device where it was designated a Device with Proxy Arp For Device enabled, even if the source device was designated a Subscriber: Reiher Second Report, paras 185–228 and Appendix A. However, as I have construed the ’760 Patent, the determination as to whether *communications between* devices are *allowed* would be understood to mean a verification of whether the source is allowed to send to the destination, and not whether the destination is also allowed to send back to the source. Therefore the fact that the USG only assessed unidirectional rather than bidirectional communication would not have prevented it from anticipating if it had disclosed and enabled the other essential elements of the ’760 Patent.

[261] Given the foregoing, I need not address the parties' submissions on whether the commercial availability of the USG device and the ability to assess its performance using the same Wireshark set up used in Guest Tek's testing is sufficient to constitute enabling disclosure.

[262] Nomadix has therefore not shown the '760 Patent was anticipated by the use in a network comprising a wireless access node of the USG gateway running the NSE software version 3.08.105.

(b) *Haverinen*

[263] Nomadix also argues the '760 Patent was anticipated by H. Haverinen, "Improving User Privacy with Firewall Techniques on the Wireless LAN Access Point" (Paper delivered at the 13th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications, Pavilhao Atlantico, Lisboa, Portugal, September 15–18, 2002), vol. 2 PIMRC, pp 987–991 [Haverinen]: Lavian First '760 Report, Appendix TL-15. In its Amended Statement of Defence and Counterclaim, Nomadix claimed Haverinen rendered the '760 Patent obvious but did not plead anticipation by Haverinen. Nonetheless, Dr. Lavian gave his opinion that Haverinen anticipated the '760 Patent and Dr. Reiher responded to that opinion: Lavian First '760 Report, paras 9.1–9.6, 10.4; Reiher Second Report, paras 112–123. Guest Tek did not object to Nomadix's reliance on Haverinen for anticipation. Given the absence of objection and the lack of prejudice to Guest Tek, I will consider Nomadix's anticipation argument based on Haverinen, notwithstanding that it was not pleaded: *Georgetown Rail Equipment Company v Rail Radar Inc*, 2018 FC 70 at paras 71–75, rev'd on other grounds 2019 FCA 203.

[264] Haverinen addresses security issues arising when WLANs allow direct, unmediated communications between networked devices at the link layer. Haverinen identifies similar concerns as the ’760 Patent, noting that for users of public access networks such as those in hotels, internet access is useful but access to computers of other local users is undesirable. The paper describes techniques to “maintain communication between client computers and an access router while blocking direct communication between the client computers.”

[265] Haverinen proposes modifying the wireless access point in a network so it discards all broadcast ARP requests received from the WLAN interface, and sends a proxy ARP response to the request instead. This proxy ARP response would give the MAC address of “an appropriate router or server” such as the access router. Haverinen notes this would not prevent unicast traffic between network devices where they use other (non-ARP) means to learn each other’s MAC addresses, such as (a) manual configuration; or (b) a new type of resolution request message. It goes on to suggest other methods of blocking unicast traffic as an additional security measure, including a forwarding policy database at the access point with MAC addresses of routers and servers on the WLAN. This would allow the access point to transmit traffic to those routers and servers, but block unicast communications between WLAN clients. The access point could also maintain a list or range of allowed unicast IP addresses and use that list to filter unicast packets. Since this manual configuration of the access point is cumbersome, Haverinen also proposes an approach in which the access point checks its own database and asks other access points to determine whether the destination address is also a client, discarding the packet if so.

[266] I conclude Haverinen does not disclose several essential elements of Claim 1 of the ’760 Patent and is therefore not anticipatory.

[267] Dr. Lavian opined that the “main distinction” between Haverinen and the ’760 Patent is that Haverinen’s method “systematically blocks” communications between wireless devices instead of allowing them to communicate in some circumstances: Lavian First ’760 Report, para 10.4. However, he gave his understanding that the principles governing anticipation allowed a POSITA to supplement the information in Haverinen with the CGK.

[268] I agree with Guest Tek that Dr. Lavian’s approach to anticipation is incorrect. To anticipate, prior art must disclose the essential elements of a patent without supplementation by the CGK. Only in assessing whether the disclosure is enabling can the CGK be used to supplement the prior art: *Sanofi-Synthelabo* at paras 27, 33, 37. Nonetheless, I agree that Haverinen does not in fact have the distinction Dr. Lavian identified. Haverinen does refer to generally blocking unicast traffic between clients. But it also refers to communication between clients being possible through manual configuration or a new type of resolution request message: Reiher Second Report, para 127.

[269] At the same time, I also agree with Guest Tek that Haverinen does not disclose several essential elements of the claims of the ’760 Patent. In particular, Haverinen does not disclose at least the following essential elements:

- *wireless access node […] configured to transmit all first packets […] to the gateway*: Haverinen describes having the wireless access node respond to ARP requests and then drop the ARP request, such that ARP requests never get to the access router/gateway. In the ’760 Patent, ARP requests are among the packets forced to the gateway;

- *gateway configured to determine […] whether the packet is from a first one of the wireless computing devices directed to any other of the wireless computing devices*: Haverinen describes various determinations being made by the wireless access point, including whether a unicast packet is from a wireless client directed to another. It does not describe determinations being done at the gateway;

- *gateway configured to transmit the packet to the destination address […] when the packet is not directed to any other wireless computing device on the network*: Again, Haverinen describes the determination regarding forwarding and dropping packets by the access point, not a gateway; and

- *gateway further configured to determine whether communications between the first wireless computing device and the second wireless computing device are allowed*: Haverinen describes allowing communications between wireless clients on the network, but does so through manual configuration of devices or a novel resolution request. While it also refers to forwarding policies at the access node, this policy is described as listing MAC addresses of routers and servers on the WLAN, not listing permitted wireless devices. No determination of whether communications are allowed is described as occurring at the gateway.

[270] As Haverinen does not disclose all of the essential elements of the claims of the ’760 Patent, I need not address whether that disclosure is enabling. I conclude that Haverinen did not anticipate any claims of the ’760 Patent.

(2) Obviousness

[271] Nomadix's submissions on obviousness were limited, as were Dr. Lavian's statements in his report addressing validity. Neither addressed in detail the four-part approach to obviousness set out in *Sanofi-Synthelabo*. In essence, Nomadix argues the '760 Patent claims a combination of known techniques that would be obvious to the POSITA, and the differences between Haverinen and the '760 Patent would be obvious. In my view, neither argument is sustainable. I will again focus on Claim 1 as the parties agree the validity of that claim determines the validity of all claims.

[272] The first step of the *Sanofi-Synthelabo* approach is identifying the POSITA and their CGK. I have done so above at paragraphs [73] to [81] and in reference to relevant aspects of the CGK in the discussion at paragraphs [105] and [112] to [117].

[273] The second step involves identifying the inventive concept of the claim in question. Dr. Lavian described the inventive concept of Claim 1 as "a network architecture in which communications on a local network are forced to a gateway. The gateway can intercept and drop messages. The gateway is configured to determine which communication between two (2) devices are allowed": Lavian First '760 Report, para 10.3. Dr. Reiher disagreed slightly with this characterization, noting it did not reflect the requirement that packets forced to the gateway include those that would not normally go to the gateway, or the gateway's determination of whether each packet is from one wireless device to another wireless device on the network: Reiher Second Report, paras 156–159.

[274] Dr. Lavian and Dr. Reiher’s views are not in my opinion significantly different. I would characterize the “inventive concept” of Claim 1 as being to address certain security issues on a wireless network through a particular network architecture and configuration in which a wireless access node transmits all packets it receives from wireless devices to a gateway regardless of destination, and the gateway is configured to then transmit the packets if they are not addressed to another wireless device on the network or, if they are so addressed, to only transmit them if communications between the devices are allowed. Important in this inventive concept are the ideas that (a) the determinations regarding packet forwarding are made at the gateway; (b) packets directed to wireless devices on the network are treated differently than those that are not; and (c) the wireless access node transmits all packets to the gateway so it can make those determinations.

[275] The third step is to identify the differences between this concept and the prior art. There is no question that network systems containing wireless access nodes and gateways were well known to the POSITA as of the claim date. Similarly, general techniques of dropping or forwarding packets, access control lists, limiting client-to-client network traffic, and proxy ARP were known. The difference between Claim 1 and the prior art lies in taking known techniques and networking equipment, and implementing them in the particular network architecture and configuration described.

[276] The ultimate question at the fourth step of the *Sanofi-Synthelabo* approach is therefore whether it would have been obvious at the claim date to use the particular claimed network architecture and configuration to achieve the security result described. In particular, would it

have been obvious to distinguish between packets directed to other wireless devices on the network and those that are not, conduct an authorization check before transmitting packets directed to other wireless devices, perform these determinations at the gateway, and ensure all packets are forced to the gateway to allow the determinations to be made?

[277] I am not satisfied that Nomadix has met its burden of demonstrating it would have been obvious to the POSITA to design a network in this way to achieve the security improvements desired regarding communications between wireless devices accessing a wireless network. The mere assertion that it would have been obvious to combine known components from the CGK to arrive at what is recited in the patent is insufficient to show that the particular identified configuration was obvious. Nomadix provided little expert or factual evidence to show that making the particular packet forwarding determinations described, and making them at a centralized gateway after forcing all packets to that gateway, would be a solution the POSITA would have arrived at without inventiveness.

[278] The Haverinen reference indicates there was general recognition that communication between two wireless devices on a wireless network created a security issue. It also indicates this security issue could be solved by preventing such communications. However, the way in which Haverinen does this is materially different from the way the ’760 Patent does this, as described above. Contrary to Dr. Lavian’s evidence and Nomadix’s submission, I do not consider that the only difference between Haverinen and the ’760 Patent is allowing communication between wireless devices in some cases, and that this would “just involve a minor change to the forwarding table.”

[279] Haverinen’s solution is for communications between wireless devices to be prevented by the wireless access point. It also suggests this block could be circumvented when needed by programming individual wireless devices themselves. There is nothing in Haverinen that points to implementing traffic mediation at the gateway and forcing packets to the gateway for that purpose. To the contrary, Haverinen suggests cutting ARP packets off before they reach the gateway. While I am prepared to accept that simply moving a known functionality from one network component to another may in some cases be an obvious development, the ’760 Patent goes beyond that in defining particular traffic determinations, ensuring they occur at the gateway, and ensuring all packets are forced to the gateway for that purpose. Nomadix has not demonstrated this would be obvious to the non-inventive POSITA.

[280] Other factors also do not point to the invention being obvious. As noted, wireless networks with a centralized gateway performing a variety of functions pre-existed the priority date of the ’760 Patent. Wireless networks using a USG device are an example: Exhibit 107, p 13. However, even though Nomadix itself was manufacturing and selling a gateway designed for use in wireless networks, and had been improving its software continuously for years, there was no evidence Nomadix had developed or implemented the particular solution described in the ’760 Patent. There was also evidence hotels implemented a variety of different architectures for their wireless networks at the time: Transcript, pp 288–299. This points away from a finding that the solution of the ’760 Patent was obvious.

[281] The evidence regarding the inventors’ conduct in developing the invention of the ’760 Patent was equivocal. Mr. Molen described the invention as coming out of “brainstorming

ideas to solve this problem," but did not suggest any difficulty in arriving at the solution: Transcript, p 101. Mr. Draper said he worked with the other co-inventors "daily for many months on this invention," although it was unclear whether this referred to the invention itself or drafting the patent arising from the invention: Transcript, p 208. Mr. DeHoop stated the process included proving the concept and identifying equipment that could implement it: Transcript, p 272. I find the evidence regarding the inventors' conduct does not speak strongly either in favour of or against the solution arrived at being obvious.

[282] The evidence of commercial success or value was also limited. Mr. Draper indicated that customer feedback to their solution was supportive: Transcript, p 273. However, such evidence is not sufficient to materially affect the obviousness analysis. Similarly, while Mr. Levy gave the overall price of Guest Tek's purchase of assets from iBAHN and an estimate of the value of the patent portfolio in that transaction, there was no price attributed to the invention that became the '760 Patent that would allow any assessment of its commercial value to Guest Tek at the time: Transcript, pp 301–309.

[283] I therefore conclude Nomadix has not shown that the '760 Patent is invalid for being obvious over the prior art in light of the CGK.

G. *Conclusion*

[284] For the foregoing reasons, I conclude Guest Tek has not shown that Nomadix has infringed the asserted claims of the '760 Patent and Nomadix has not shown the '760 Patent to be invalid.

V. Canadian Patent 2,750,345

A. *Introduction*

[285] The ’345 Patent relates to the management of available bandwidth between users of a network. Multiple users of a network, such as hotel guests accessing the internet via the hotel’s network, may place competing demands on the network’s available bandwidth. At a high level, bandwidth management seeks to coordinate those demands by determining what data will be processed in what order. The ’345 Patent describes and claims a bandwidth management system in which multiple *queues* are used to regulate the flow of traffic. An *enqueuing module* puts the traffic into queues and a *dequeuing module* removes data from the queues in accordance with a particular dequeuing approach. Fundamental to the enqueuing and dequeuing approaches of the ’345 Patent are the concepts of *zones*, each of which is assigned a queue; *quantums*, which dictate the amount of data dequeued from a queue; and *user loads*, which relate to the amount of use in a zone and affect the quantum for the zone’s queue.

[286] The ’345 Patent has three types of claims: “system” claims, “method” claims, and one “computer-readable medium” or “software” claim. The system claims claim a bandwidth management system with a set of defined elements that have identified purposes. The method claims claim a method of allocating bandwidth in a system. The software claim claims software that, when executed by a computer, causes the computer to perform the method of any of the method claims. There are four independent claims in the patent. Claim 1 is an independent system claim from which Claims 2 to 18 directly or indirectly depend. Claims 19 and 20 are each independent system claims that have no dependent claims. Claim 21 is an independent method

claim from which the other method claims directly or indirectly depend. The software claim, Claim 39, also depends from the method claims. Guest Tek asserts Nomadix infringes system Claims 1, 3, 16 to 18 (as each depends from Claims 1 and 3), 19, 20, and the software Claim 39 (as it depends from Claims 21 and 23). Guest Tek also asserts Nomadix is inducing infringement of these claims plus method Claims 21, 23, and 36 to 38 (as each depends from Claims 21 and 23).

[287] The ’345 Patent was filed on August 24, 2011 with no claimed priority date, such that the filing date is the relevant date for invalidity purposes. The application was published on December 14, 2011, the material date for construction purposes. It issued on June 18, 2013.

B. *The Person of Ordinary Skill in the Art*

[288] The ’345 Patent is directed to the art of bandwidth management in computer networks. The parties’ experts, Dr. Dordal and Dr. Lavian, agreed a person of skill in this art would have knowledge of computer networking technology and computer programming. Dr. Dordal suggested this might be obtained either through a bachelor’s degree in computer science or engineering, or through three years’ work in the field of computer networking: Dordal First Report, para 57. Dr. Lavian considered the POSITA would have at least a bachelor’s degree plus four years of experience working in network bandwidth management: Lavian First ’345 Report, para 4.10.

[289] The parties agree nothing turns on this difference in described experience. In my view, the ’345 Patent describes the practical implementation of a bandwidth management system in a

setting such as a hotel or conference centre, including through the programming or configuration of the system. This suggests the POSITA would have some practical experience in addition to academic knowledge of networking. However, the system described pertains primarily to a particular enqueuing and dequeuing method based on allocation of bandwidth between users, which does not seem to call for the degree of experience Dr. Lavian suggests. In my view, a POSITA would be fairly described as either having a bachelor's degree in a computer science field including courses in networking, plus a year or two's experience in that field involving bandwidth management; or having equivalent work experience to gain the same level of knowledge and experience.

C. *The Common General Knowledge*

[290] As with the ’760 Patent, the two experts’ discussion of the CGK was markedly different. Dr. Dordal gave a brief description of the body of knowledge of a POSITA, which would include knowledge of “general principles” in a number of identified areas: Dordal First Report, paras 58–61. Dr. Lavian, on the other hand, set out a list of some 27 documents illustrative of the CGK, and provided a detailed background on issues such as queuing, quality of service, traffic shaping, congestion control, and the use of hierarchical token buckets: Lavian First ’345 Report, paras 5.1–5.128. As with his description of the CGK in respect of the ’760 Patent, Dr. Lavian’s description is helpful, but suffers by its apparent focus on validity issues, commenting on the extent to which the elements of the ’345 Patent were known and not inventive: Lavian First ’345 Report, paras 5.27, 5.33, 5.34, 5.81–5.85, 5.110, 5.127.

[291] Dr. Dordal criticized Dr. Lavian's description of the CGK as "excessive and unrealistic" and a "massive library of knowledge": Dordal Second Report, paras 30–31. As had Dr. Reiher, Dr. Dordal felt a POSITA would not "know and have retained in memory every detail" of documents such as the Tanenbaum text. He said even he would need to do online searches to come up to speed on these issues: Dordal Second Report, paras 34–38. In my view, by focusing on what would be memorized or would have to be looked up, Dr. Dordal took too narrow a view of the CGK, which includes what the POSITA may reasonably be expected to know "and to be able to find out": *Tetra Tech* at para 28. While this does not encompass all locatable prior art, it does encompass that which a POSITA would become aware of and accept as a good basis for further action: *Tadalafil* at para 24. In any event, despite his criticism, Dr. Dordal conceded in cross-examination that his main issues with Dr. Lavian's description of the CGK related to matters that had no bearing on the '345 Patent: Transcript, pp 1010–1012.

[292] Given its importance to the context of the '345 Patent, some further discussion of the CGK in respect of traffic and bandwidth management is merited before turning to claims construction. The following understanding comes primarily from Dr. Lavian's first report and from four documents he refers to, namely the Tanenbaum text; M Devera, "HTB Linux queuing discipline manual – user guide" (May 5, 2002); MA Brown, "Traffic Control HOWTO, v 1.0.2" (October 2006); and B Hubert et al, "Linux Advanced Routing & Traffic Control HOWTO, rev 1.1" (July 22, 2002): Lavian First '345 Report, paras 5.5, 5.23, 5.52–5.55, 5.65–5.77, 5.86–5.90 and Appendices TL-09, TL-28, TL-33, and TL-34.

(a) *Queues and queuing disciplines*

[293] Queues are locations (or buffers) containing data, such as packets, waiting to be processed. Data entering a queue is "enqueued" and data removed from the queue is "dequeued." Much of network traffic control is dependent on controlling how packets are scheduled to be enqueued and/or dequeued. A simple way of doing this is a "first in first out" (FIFO) model, in which packets are enqueued as they arrive and dequeued in the same order as quickly as they can be processed. For a variety of reasons, such as the nature of different types of internet traffic and the need to prioritize traffic from different sources, a number of more complex approaches to traffic management are used.

[294] A queuing discipline (qdisc) is a Linux component that acts as a scheduler, arranging or rearranging packets for output from a queue.

(b) *Fair queuing*

[295] "Fair queuing" is an algorithm for scheduling packets from multiple flows of traffic, such as packets relating to different traffic types, different users, or different groups of users. It involves multiple queues, one for each flow of traffic. Each queue will have a packet taken from it for transmission to its destination (or the next hop) in a round-robin sequence. In some cases, it may be preferable to give some queues, such as those carrying video traffic, greater bandwidth. In such cases, more bytes or packets can be taken from the prioritized queue through a process of "weighted fair queuing" (WFQ).

[296] “Stochastic Fair Queuing” (SFQ) is a fair queuing qdisc that uses a changing hashing algorithm to divide traffic into queues. In SFQ, the number of bytes that a stream can dequeue before the next queue gets a turn to dequeue is dictated by a parameter called the “quantum.” In other words, the quantum parameter controls how many bytes are released from each FIFO queue in round-robin fashion.

[297] It is worth noting that data is dequeued from a queue in packets. That is, a dequeuer will either dequeue a packet or not dequeue it, but will not dequeue half a packet: Transcript, p 1030. The “maximum transmission unit” (MTU) is the largest packet that can be sent on a network. On an Ethernet network, the maximum payload of a packet is 1500 bytes, and this figure is therefore the MTU on an Ethernet network: Lavian First ’345 Report, paras 5.86–5.89; Dordal First Report, para 211. Thus if the quantum is set to 1500 bytes (or 1514, to account for the Ethernet packet header) in an Ethernet network, one full-sized packet or, if small enough, multiple smaller packets, will be dequeued before turning to the next queue. In the SFQ qdisc, the quantum parameter defaults to 1 MTU and cannot be set below the MTU size: Lavian First ’345 Report, Appendix TL-27, s 5.1.2.

(c) *Classes in queuing*

[298] Some qdiscs involve classifying traffic into different classes or subclasses for different treatment. Each class contains a further qdisc. Class Based Queuing (CBQ) is one such Linux qdisc that uses classes.

[299] A hierarchical class structure can be created with classes at different levels, as in this example from the *HTB Linux queuing discipline manual – user guide*:

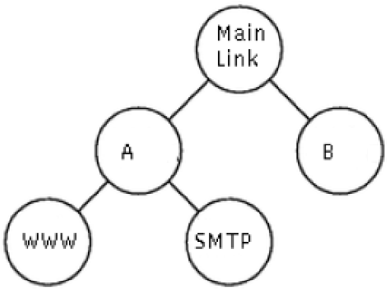

[Description of inserted diagram for accessibility: The diagram contains five circles arranged on three rows. On the top row is one circle labeled "Main Link." On the next row are two circles, each connected by a line to the top row circle, labeled "A" and "B" respectively. On the bottom row are two circles, each connected by a line to the circle labled "A," and labeled "WWW" and "SMTP" respectively.]

[300] In the diagram, A and B represent different users. WWW and SMTP represent A's web traffic and A's other traffic, respectively. The different types of traffic are represented by different classes, each with its own queue. Available bandwidth may be allocated between A and B, and A's allocation can be subdivided between their WWW and SMTP traffic. Classes at different levels are described as "parent" and "child" classes. A class with no parent is a "root" class. A class with no children is a "leaf" class. A class with a parent and a child is an "inner" class.

(d) *Linux based hierarchical token bucket*

[301] The "token bucket" algorithm is based on a metaphor of a bucket that can hold a certain amount of water (data). In a "leaky bucket," the water drips out of a hole in the bucket at a constant rate, regardless of how much water is in the bucket. In a token bucket, the rate at which data can leave the bucket can vary. Tokens, representing an ability to send a certain amount of

data to the network, are generated by a clock and added to the bucket up to the maximum capacity of the bucket.

[302] When a packet is transmitted, it cancels out one or more tokens equivalent to its size, such that when packets arrive they can be transmitted provided there are enough tokens in the bucket. If there are no tokens in the bucket, it must wait until another token is added before a packet can be transmitted. The token bucket algorithm thus allows saving of transmission capacity, permitting bursts of packets up to the maximum bucket size to be sent within a given unit of time. The size of the bucket (the number of tokens it can store) is an important parameter, because it dictates the maximum burst size.

[303] Hierarchical token bucket (HTB) is a qdisc that uses the concepts of tokens and buckets along with a class-based system. It contemplates multiple classes and subclasses that can have different token rates associated with them. In HTB, child classes can "borrow" tokens from their parents if they have exceeded their rate. HTB uses the concept of quantums discussed earlier, including in the borrowing of bandwidth. If several classes are competing for a parent's bandwidth, they will get it in proportion to their quantum.

[304] The '345 Patent refers to the Linux HTB and SFQ qdiscs as being well-known in the art, such that description of their operation could be omitted. Although he did not refer to it in his initial report, Dr. Dordal recognized that HTB was a well-known implementation of hierarchical fair queuing, which was "a foundational queuing strategy": Transcript pp 1010, 1113. He agreed the POSITA would have skill using Linux HTB or similar tools, which he recognized was a "core skill": Dordal Second Report, paras 61, 77; Transcript pp 864–866, 1008, 1010.

D. *Claims Construction*

[305] Claim 1 of the ’345 Patent claims the following (I have added the numbering):

> A bandwidth management system comprising:
>
> (1) a plurality of queues respectively corresponding to a plurality of zones;
>
> (2) an enqueuing module for receiving network traffic from one or more incoming network interfaces, determining a belonging zone to which the network traffic belongs, and enqueuing the network traffic on a queue corresponding to the belonging zone;
>
> (3) a quantum manager for dynamically adjusting values of a plurality of quantums, each of the queues having a respective quantum associated therewith; and
>
> (4) a dequeuing module for selectively dequeuing data from the queues and passing the data to one or more outgoing network interfaces;
>
> (5) wherein, when a selected queue has no guaranteed bandwidth rate or has already reached its guaranteed bandwidth rate, the dequeuing module dequeues at most an amount of data from the selected queue up to the quantum of the selected queue before dequeuing data from another of the queues; and
>
> (6) the quantum manager dynamically adjusts the quantums in proportion to tracked user load under each of the zones such that the quantum of the selected queue is higher than the other quantums while the zone to which the selected queue corresponds has higher user load than the other zones, and such that the quantum of the selected queue is lower than the other quantums while the zone to which the selected queue corresponds has lower user load than the other zones.

[306] The parties agree Claim 1 claims a system that includes *queues* that correspond to *zones*, an *enqueuing module*, a *quantum manager*, and a *dequeuing module*. Each of these components, and their function, is described in the paragraphs of the claim.

(1) a plurality of queues respectively corresponding to a plurality of zones

[307] There is no dispute between the parties that a *queue* is a lineup of data, such as packets, waiting to be processed. The system of the ’345 Patent has a number of queues, each of which is associated with one of a number of *zones*. While the experts initially disagreed on the meaning of the term *zones* in their reports, Nomadix did not pursue Dr. Lavian’s definition at trial, and largely agreed with the definition as proposed by Dr. Dordal.

[308] The disclosure of the ’345 Patent includes the following definition of *zone*:

> The definition of a “zone” according to the invention is flexible and depends upon application-specific design requirements. Generally speaking, zones represent manageable divisions of users, user devices, and/or other lower-level zones. Zones may be logically and/or physically separated from other zones. For example, different zones may be isolated on separate local area networks (LANs) or virtual LANs (VLANs) corresponding to separate rooms or areas of a hotel, or different zones may share a same LAN but correspond to different service levels of users within a hotel. The zones and tree-structure may be dynamic and change at any time as users associated with certain zones upgrade their Internet access or when meeting reservations begin and end, for example.
>
> [Emphasis added.]

[309] In my view, the POSITA looking at the “disclosure and claims […] as a whole” and giving consideration to the patent specification to understand what was meant by the words in the claim would recognize the term *zone* as used in the claims is being used in a particular way related to the bandwidth management systems the ’345 Patent describes: *Tearlab* at para 33. Neither expert suggested that the term *zone* has a meaning in networking that would render its use in the claim unambiguous.

[310] Dr. Dordal adopted a construction of *zones* consistent with the definition in the disclosure, namely divisions of users or user devices: Dordal First Report, paras 63–64. Dr. Lavian concluded the term *zones* would refer to a "location" in a geographic sense: Lavian First '345 Report, paras 5.114, 7.10 (p 73); Lavian Second '345 Report, paras 2, 152, 155, 156, 158. Nomadix in closing submissions did not insist on Dr. Lavian's limited geographic sense of *zones* and largely adopted the definition from the disclosure, but it did argue that in a hotel setting, *zones* would correspond to different physical rooms or areas in accordance with Dr. Lavian's approach: Transcript, pp 2362–2363.

[311] I see nothing in the '345 Patent that would limit the term in this way. To the contrary, the definition in the disclosure suggests different zones may "correspond to different service levels of users within a hotel" and a user may change zones simply by upgrading their internet access. Although the diagrams show zones primarily linked to rooms, I do not believe a POSITA considering the diagrams together with the written disclosure would conclude *zones* were limited to physical areas when the patented system was used in a hotel. Further, the fact that Claim 18 adds a limitation that each zone corresponds to one or more rooms of a hotel indicates that the zones of Claim 1 are not so limited: *Halford* at para 93.

[312] I conclude a POSITA would understand the term *zones* to mean manageable divisions of users, user devices, and/or other lower-level zones. Such zones may be geographically separated (such as rooms or areas) or may simply be logically separated (such as service levels of users).

> (2) an enqueuing module for receiving network traffic from one or more incoming network interfaces, determining a belonging zone to which the network traffic belongs, and enqueuing the network traffic on a queue corresponding to the belonging zone

[313] The parties do not disagree on the meaning of this term. In my view, a POSITA would understand the claim to require that the system include an *enqueuing module* component that is responsible for determining which zone incoming traffic belongs to, and placing it in a queue associated with that zone.

> (3) a quantum manager for dynamically adjusting values of a plurality of quantums, each of the queues having a respective quantum associated therewith

[314] This part of Claim 1 specifies that each queue in the system has a *quantum* associated with it. The *quantum manager* is a component of the system of Claim 1 that has the function of *dynamically adjusting values* of the quantums that are associated with each queue.

[315] As Dr. Dordal points out, the role or function of the *quantum* is specified later in Claim 1: Dordal First Report, para 67. In the paragraph I have numbered (5), the *dequeuing module* of the system "dequeues at most an amount of data from the selected queue up to the quantum of the selected queue before dequeuing data from another of the queues." Based on this and other references from the specification, Dr. Dordal proposes a construction of *quantum* as meaning the "maximum amount of data that can be dequeued from a queue before data will be dequeued from a different queue": Dordal First Report, paras 69–71.

[316] While Dr. Lavian phrased the definition of quantum somewhat differently in different places of his reports, I believe his definitions largely accord with Dr. Dordal's: Lavian First ’345 Report, para 7.10 (p 74); Lavian Second ’345 Report, paras 30–31, 176. Indeed, as Guest Tek points out, Dr. Lavian's Second Report includes a phrasing that essentially parallels that of Dr. Dordal, "[i]n the context of the ’345 patent, these quantums explicitly define an upper limit of data that can be dequeued from a given queue before the dequeuing module will move on to dequeue data from another queue": Lavian Second ’345 Report, para 176.

[317] Dr. Lavian opined that a POSITA would know that the term *quantum* is used in the HTB and SFQ qdiscs, and would understand that the term is being used in the same manner in the ’345 Patent: Lavian First ’345 Report, paras 5.73–5.81; Lavian Second ’345 Report, paras 30–31; Transcript, p 1741. I agree. The POSITA reading the claims of the ’345 Patent in light of their CGK, including the "core skill" of hierarchical fair queuing would, as Dr. Dordal acknowledged, understand that what was being described was a variation on hierarchical fair queuing: Transcript, p 1010. The POSITA reading the term *quantum* in Claim 1 would expect and understand it to be used in the way it is used in the art of network traffic management, with particular reference to its use as a parameter or variable in the HTB and SFQ qdiscs: Transcript, p 1069.

[318] As noted above, the inventor of the ’345 Patent refers expressly to the HTB and SFQ qdiscs. They note at page 14 of the disclosure that HTB can be used to implement the queues of the patent, and that there are advantages in doing so:

> The plurality of zone queues 402 described in this example may be implemented using a hierarchical token bucket (HTB) queuing

> discipline (qdisc) including an HTB for each queue 408. Using a HTB to implement each queue 408 allows for easy configuration of a rate, ceiling (cap), and quantum for each queue 408 as these parameters are already supported by HTB based queues.
>
> [Emphasis added.]

[319] I believe this is consistent with the definitions as given above. In these qdiscs, quantum is a parameter representing the number of bytes that will be dequeued from a queue "before the next queue gets a turn": Lavian First '345 Report, Appendix TL-34, s 9.2.3.1 (p 29). This also accords with the way the term *quantum* is described at pages 10 and 12 of the disclosure of the '345 Patent, namely "an amount of bytes that can be served at a single time from the zone and passed to a higher-level zone," "the maximum number of bytes that can be dequeued at one time," or "how many bytes can be dequeued at a single time by the dequeuing module."

[320] The parties' disagreement with respect to the term *quantum* lies less in the language of the definition, but in two aspects of the definition. First, while the parties agree the quantum represents a "maximum" or "upper limit" of bytes that can be dequeued at a time, they disagree on what this means. Nomadix argues the quantum represents the maximum for each dequeuing attempt. That is to say, the dequeuing module will dequeue up to the quantum of bytes from a queue. It will then turn to whichever queue is next dictated by the queuing strategy and dequeue up to the quantum of that queue. It will continue in this way, eventually returning to the original queue and again dequeuing up to the quantum of bytes. Guest Tek's interpretation, on the other hand, is that the maximum may represent not the maximum number of bytes on each dequeuing attempt but the maximum "average burst size, over time": Dordal First Report, paras 246–250, 380, 385–386; Transcript, pp 1044–1045.

[321] In my view, Nomadix's construction is the one that accords with the text of Claim 1, as it would be understood by the POSITA. Claim 1 states the dequeuing module "dequeues at most an amount of data from the selected queue up to the quantum of the selected queue before dequeuing data from another of the queues." This suggests an iterative process in which the dequeuing module dequeues an amount of data up to the quantum, then does the same for further queues in accordance with their respective quantums. Each time the dequeuing module dequeues from a queue, it takes data from the queue, using the quantum as the maximum. This accords with the meaning of *quantum* as it is used in the Linux qdiscs discussed above, which Guest Tek fairly conceded in closing argument was the source of the term: Transcript, pp 2118–2120.

[322] It is also worth noting the circumstances in which the full amount of the quantum may not be dequeued, *i.e.*, why it is a "maximum" and not simply the amount of data that will inavariably be dequeued from a queue. The '345 Patent describes such situations in its discussion of quantum at pages 10 and 12:

> When each zone is implemented as one or more queues enforcing rate and cap limits, the quantum indicates the maximum number of bytes that can be dequeued at one time. In some configurations, the cap limit may prevent the full quantum of bytes from being dequeued from a particular zone if this would cause the bandwidth provided to the zone to exceed its cap limit. In another configuration, as long as the data in the queue is sufficient, the quantum of bytes is always dequeued at once; however, the frequency of dequeuing the quantum of bytes may be reduced in order to limit the average bandwidth to the zone's cap.
>
> […]
>
> Each queue 408 has a respective quantum 409 that determines how many bytes can be dequeued at a single time by the dequeuing module 412. For example, with a round robin dequeuing strategy and assuming all zones have data to send, all guaranteed rates have been met, and no zone has exceeded its designated cap, the

> dequeuing module 412 cycles one-by-one to each queue at a particular tree-level and dequeues the queue’s quantum number of bytes.
>
> [Emphasis added.]

[323] Thus, the circumstances in which this maximum is not dequeued are limited. If there is less than the quantum of data in the queue, then evidently not the full quantum will be taken. This is clear both from the operation of dequeuing generally and from the reference in the disclosure to “as long as the data in the queue is sufficient” and “assuming all zones have data to send.” Less than the full quantum may also be taken where there are other limits in the configuration, such as bandwidth rates or caps, that result in the dequeuing module taking less than the quantum. Subject to such limitations, however, the dequeuing module would take from the queue the number of packets permitted by the quantum.

[324] These situations are incorporated into dependent Claim 7 of the ’345 Patent, which claims the bandwidth management system of prior claims “wherein, when all zones have data to send, all guaranteed bandwidth rates have been met, and no zone has exceeded its designated cap, the dequeuing module cycles through the queues in a round robin dequeuing strategy and dequeues each queue’s quantum of data.” The limitations of this claim are of course not to be read into Claim 1: *Halford* at para 93. However, they are useful in understanding the term *quantum* used in the claims, and in particular why the quantum is considered the “maximum amount” dequeued at one time, rather than simply the “amount” dequeued at one time.

[325] Second, and as a related matter, Dr. Dordal was of the view that the *quantum* need not be a specific value or variable stored in memory: Dordal First Report, paras 72–75. He notes that

while Claim 1 refers to the quantums for each queue being “dynamically adjusted,” it does not refer to them being stored in memory. He asserts instead that the quantum may simply be the observed maximum over time, stating that he construed the term quantum “in terms of observations, rather than specific numeric variable values”: Transcript, p 1023. In other words, Dr. Dordal took the *quantum* as potentially including simply the observed average burst size over time, rather than a number that is used by the system to impose a limit on the amount of data dequeued per dequeuing cycle. In my view, this is inconsistent with how a POSITA would understand the term *quantum* as used in the ’345 Patent in light of their knowledge of the CGK.

[326] In three telling passages from Dr. Dordal’s cross-examination, he appeared to recognize that a POSITA’s primary understanding of *quantum* would be as it is used in Linux HTB, *i.e.*, as a parameter set as an upper limit on the amount of data to be dequeued. However, he adjusted his definition in light of the fact that Nomadix’s system did not use such a quantum:

> Q. Now, I think we’re talking about quantum. For the person skilled in the art, what did a “value” mean? And what does it correspond to in source code?
>
> A. The way, in the sense I meant value there was purely the actual amount being dequeued. The Nomadix implementation doesn’t really set out a numeric value to be dequeued from one class before moving to another class. So we’ve adopted the more general description of just the amount dequeued.
>
> […]
>
> Q. Okay. And you thought it was necessary to add this “I further find” portion?
>
> A. I think -- again, the difficulty with interpreting all this, the claim language and applying it to the Nomadix device is that the obvious interpretation -- I don’t want to say “obvious” -- one interpretation of the claim language is to, you know, just say that the quantum is dequeued, that that’s what happens in Linux HTB,

> but because the Nomadix device fluctuates up and down, the amount dequeued fluctuates up and down, we went with the maximum, which I do believe is fully consistent with the claim language.
>
> Q. But isn't this, going back to the full quantum, doesn't that arise out of your analysis of the Nomadix system, that you thought it would be helpful to go back and further find this?
>
> A. If the Nomadix system had been based on HTB, it might not have -- this might not have been an issue, but it's not. So we looked at the -- this behaviour and looked at the claim -- looked at the claims in the '345 patent and came up with this claim interpretation, which is consistent with the -- which is completely supported by the '345 patent claim language and applies to the Nomadix innovation.
>
> […]
>
> Q. […] So why did you ask yourself three additional questions in face of this very plain language?
>
> A. If one looks at, again, Linux HTB, each class structure has a field called "quantum," and that is, you know, an explicit variable, and it is constantly changed -- well, it's not constantly changed, it's occasionally changed.
>
> Q. Mm-hmm.
>
> A. It does not meet the criteria of being dynamically adjusted. But in the Nomadix code, that explicit value of quantum, as a variable in the C programming language, or C++, is never declared. So we tried to create language that would cover this implicit definition of "quantum" as the maximum amount of data that can be dequeued.
>
> Q. When you say you "tried to create language," which language is it that you created?
>
> A. The idea that the quantum does not need to be stored in a variable.
>
> [Emphasis added; Transcript, pp 1022-1023, 1076–1077, 1081.]

[327] I agree with Nomadix that this approach of “trying to create language” to describe the elements of a claim in order that the claims capture a defendant’s conduct is wholly at odds with the correct approach to claims construction. Claims construction is antecedent to consideration of both validity and infringement issues, and is not to be undertaken with an eye on the allegedly infringing device: *Whirlpool* at paras 43, 49(a). Contrary to Guest Tek’s submission, what Dr. Dordal describes goes beyond merely focusing on where “the shoe pinches.” It is one thing to be aware of which terms or constructions are in dispute. It is quite another to try to “create language” to “come up with” an interpretation that will cover the impugned device.

[328] This is particularly so given Dr. Dordal’s recognition that “almost any POSITA” who was trying to implement the invention of the ’345 Patent would build something similar to Linux HTB “with a very explicit value for quantum”: Confidential Transcript (Oct 6, 2020), p 11. This is not to say that the first way a POSITA would try to implement a patent is the only way in which it may be infringed. However, it makes clear that Dr. Dordal’s assessment is that a POSITA would understand the term *quantum* in the ’345 Patent as an explicit parameter, and that the approach in which *quantum* may also encompass the observed average amount of dequeued data is a broadening of that meaning designed to capture Nomadix’s software. This significantly undermines Guest Tek’s position regarding their proposed construction of the term.

[329] In essence, the term *quantum* has a known meaning in the art as a variable that is set as a parameter to limit the amount of data dequeued by a dequeuer in various dequeuing strategies. The ’345 Patent uses that term, and expressly refers to the art in which the term is commonly used. The examples it gives all show use of the term quantum as a variable being stored and adjusted. Guest Tek and Dr. Dordal’s argument that the quantum may be something very

different—an unstored derived average value of the impact of how a particular system operates over time—because that definition would capture the Nomadix system and the ’345 Patent does not expressly exclude that option, is contrary to the principles of claim interpretation and is unpersuasive.

[330] As for the quantum manager’s function of *dynamically adjusting values* of the quantums, both parties recognized this involved the ongoing changing of the quantum values over time: Dordal First Report, para 75; Lavian First ’345 Report, para 7.10 (p 75); Transcript, p 876. The claim specifies it is the quantum manager that performs this dynamic adjustment of the quantum values rather than, for example, a manual change of the quantum parameter by an administrator. The basis on which the quantum manager adjusts the quantums is specified later in Claim 1 as being *in proportion to tracked user load*, a term discussed below.

[331] The primary disagreement with respect to the term *dynamically adjusting values* was whether the quantum value itself had to be stored in memory, a matter addressed above. The requirement that the quantum manager dynamically adjust the quantum values further underscores that a POSITA would understand the quantum to be a parameter used by the system to act as a limit on the amount of data dequeued from a queue each time it is that queue’s turn. It is not sufficient for the system to adjust other aspects of its functioning with the effect that average burst size over time displays the attributes of a quantum.

[332] Based on the foregoing I conclude the POSITA would understand this paragraph of Claim 1 to require that the system comprise a component (termed the *quantum manager*) that adjusts on an ongoing basis the value of a parameter (termed the *quantum*) that defines the upper

limit of data that will be dequeued from a given queue each time the dequeuing module dequeues from the queue before moving on to dequeue data from another queue.

> (4) a dequeuing module for selectively dequeuing data from the queues and passing the data to one or more outgoing network interfaces

[333] There is no dispute that the *dequeuing module* is a component of the system responsible for the dequeuing of data from the queues. Data from the queue of a lower level class may be dequeued into the queue of a higher level class, or will ultimately be dequeued to be passed along to an outgoing network interface.

> (5) wherein, when a selected queue has no guaranteed bandwidth rate or has already reached its guaranteed bandwidth rate, the dequeuing module dequeues at most an amount of data from the selected queue up to the quantum of the selected queue before dequeuing data from another of the queues

[334] As discussed above, this passage in Claim 1 describes the function of the quantum. While this passage refers to the dequeuing module dequeuing *at most an amount of data from the selected queue up to the quantum*, circumstances in which less than the quantum amount of data is dequeued are discussed above.

[335] This passage also specifies the quantum is used by the dequeuing module when a selected queue has no guaranteed bandwidth rate, or has already reached it. In other words, the quantum-based queuing approach only applies once a zone’s queue has reached any minimum guaranteed bandwidth rate it has. The POSITA would understand that other bandwidth management approaches would be used to ensure a given queue meets its guaranteed bandwidth rate. This is

consistent with the discussion in the disclosure of the ’345 Patent, which describes various examples that assume all guaranteed bandwidth rates have been met.

> (6) the quantum manager dynamically adjusts the quantums in proportion to tracked user load under each of the zones such that the quantum of the selected queue is higher than the other quantums while the zone to which the selected queue corresponds has higher user load than the other zones, and such that the quantum of the selected queue is lower than the other quantums while the zone to which the selected queue corresponds has lower user load than the other zones

[336] Claim 1 specifies the quantums are dynamically adjusted *in proportion to tracked user load*. This requirement raises three separate terms for construction: *user load, tracked user load,* and *in proportion to*.

(a) *user load*

[337] Dr. Lavian suggested the term *user load* in Claim 1 would be understood to mean simply the number of users in the zone: Lavian First ’345 Report, para 7.10 (p 75); Lavian Second ’345 Report, para 52. I agree with Guest Tek that this reading of *user load* is too narrow and inconsistent with the other claims of the patent.

[338] Claim 2 claims the system of Claim 1 wherein *user load* is tracked by *summing how many current users are in the zone*, which is consistent with Dr. Lavian’s definition. However, Claim 3 claims the system of Claim 1 wherein *user load* is tracked by *summing bandwidth caps of current users in the zone*, which goes beyond merely counting the number of users in the zone. As each dependent claim must fall within the scope of the independent claim, at least these two

approaches to *user load* must be encompassed within the term as used in Claim 1, and would be understood by the POSITA in this way: Transcript, p 881; *Halford* at para 91.

[339] The claims do not describe any other approaches to *user load*. Two other methods for determining *user load* in addition to those described in Claims 2 and 3 are described at page 22 of the disclosure, namely a summation of how much data current users have recently sent or received; and the amount of user traffic currently enqueued on one or more queues corresponding to the zone. The disclosure as a whole, including the examples given, therefore point to an assessment of *user load* that may simply be the number of *current users* in the zone (a term discussed further at paragraph [356] below), or may involve an assessment of the potential or actual demand for bandwidth traffic in the zone.

[340] In my view, the POSITA would understand the notion of *user load* to be a function of the number of *current users* and/or their demand for bandwidth, determined either as simply a count of the number of current users, as the sum of their bandwidth caps, or as a measure of the current or recent traffic in the zone.

(b) *tracked user load*

[341] Claim 1 refers not just to *user load* but to *tracked user load*. Similar to his approach to *quantum*, Dr. Dordal opined there was no requirement the user load be stored as a variable in memory, provided it is tracked in some manner: Dordal First Report, para 82. Dr. Lavian considered the question of storage to be an irrelevant factor, but questioned how a computer

could track something without storing it: Lavian Second ’345 Report, para 37; Transcript, pp 1761–1762.

[342] I agree with Dr. Lavian that the question of storage is somewhat of a distraction that is not raised in the claims or disclosure of the patent. Dr. Dordal is correct that nothing in the patent states that the *tracked user load* must be stored as a separate variable. In my view, what the claim requires is that the system keep track of the user load on an ongoing basis, that is, continue to monitor or track how many *current users* are in the zone and, if necessary for the user load being implemented, what their bandwidth cap is, or how much traffic they have recently used, or how traffic is enqueued on the queue for the zone: Transcript, p 1085.

[343] That said, I disagree that the notion of *tracked user load* can go so far as Dr. Dordal’s suggestion that “if the software performs actions that have the effect of allowing an inference of the user load, that’s sufficient” [emphasis added]: Transcript, p 880. Claim 1 calls for the quantum manager to dynamically adjust the quantums in proportion to the tracked user load of each zone. The tracked user load is thus an input that the quantum manager uses to perform its task of adjusting quantums. It is not sufficient to simply be able to infer a notional tracked user load based on the effects of other actions taken. The system must use the tracked user load to dynamically adjust the quantums.

(c) *in proportion to*

[344] Proportionality in mathematics means one variable is multiplicatively connected to another by a constant: Dordal First Report, para 76. To use Nomadix’s example, it can be

represented by the formula Y=cX, where “c” is a constant. Dr. Dordal suggested that in the context of Claim 1, the term *in proportion to* does not mean “directly proportional to” in a mathematical sense, but simply that the quantum is higher in zones with higher user loads and lower in zones with lower user loads: Dordal First Report, paras 76–79. Dr. Lavian, however, considered that proportionality necessarily involved a relationship that can be represented by a mathematical formula: Lavian Second ’345 Report, paras 35–36.

[345] A POSITA with the computer science background described by both experts would recognize the term *in proportion to* as invoking the notion of mathematical proportionality. However, the POSITA would be looking to understand whether the inventor meant the term in this purely mathematical sense in the context of the ’345 Patent. In doing so, they would look to the remainder of Claim 1, the other claims of the ’345 Patent, and the discussion of how quantum is calculated in the disclosure.

[346] There is clearly a relationship in the Claim 1 language between the term *in proportion to* and the two *such that* qualifications that follow: (a) *such that* the quantum of the selected queue is higher than the other quantums while the zone to which the selected queue corresponds has higher user load than the other zones, and (b) *such that* the quantum of the selected queue is lower than the other quantums while the zone to which the selected queue corresponds has lower user load than the other zones. On Dr. Dordal’s construction, these *such that* clauses effectively define the term *in proportion to* for the purposes of Claim 1 so that any relationship that meets the general description of the *such that* clauses is *in proportion to.*

[347] As Nomadix points out, if the *such that* clauses are all that is required of the term *in proportion to*, then the claim would have the same meaning if the words *in proportion to* were removed completely. They argue this redundancy should be avoided by giving *in proportion to* its mathematical meaning. Conversely, however, if the term *in proportion to* were given its mathematical sense, then the *such that* clauses would be largely redundant. Nomadix argues, without evidence on this point from either expert, that the *such that* clauses are there to specify the direction of the proportionality (*i.e.*, that “c” is a positive rather than negative number).

[348] While I agree that the *such that* clauses do specify direction, the positive proportionality is generally clear from the context of the ’345 Patent in any case. There is therefore an element of redundancy in the claim on either party’s proposed construction. I believe this is best explained by reading the *in proportion to* term and the *such that* clauses not in isolation but as mutually supportive and provided for clarity. Notably, the *such that* clauses involve a comparison between the quantum for a selected queue and other quantums for other queues, rather than simply the method of deriving a single quantum.

[349] A POSITA with an understanding of quantums in the context of bandwidth management would know that a larger quantum would generally result in more data being dequeued, and therefore a zone getting more effective bandwidth when it has a higher user load. This is in keeping with the context and language of the ’345 Patent and its claims. However, a POSITA would have to look further to the disclosure to understand how the inventor intended the term *in proportion to* to describe the relationship between user load and quantum.

[350] The ’345 Patent gives three examples of how quantum is calculated based on user load. In each, the quantum is a multiple of 1500, the MTU for an Ethernet network. Consistent with the CGK, the minimum quantum is the MTU of 1500 (were it less, no full-sized Ethernet packet could be dequeued). In each example in the patent, the quantum is “scaled” to the user load by deriving a scale factor and multiplying that by the user load. When not already equal to a multiple of the MTU, the scaled quantum is either rounded up to a multiple of the MTU (Figure 5), or rounded to the nearest multiple of the MTU (Figure 8).

[351] These calculation methods are claimed in dependent claims of the ’345 Patent. The approach of providing a minimum quantum equal to 1 MTU, and that of “scaling” the quantums based on minimum and maximum amounts related to the MTU, are seen in Claims 8, 10, and 12. The approach of rounding quantums to a multiple of the MTU is specified in Claim 14. These limitations should not be read into Claim 1, but the proportionality of Claim 1 must encompass these calculation methods.

[352] Guest Tek underscores that in these example calculations, there is not a strict proportionality between user load and quantum. In particular, it points to the fact that the 1 MTU minimum and the rounding to multiples of the MTU means that different user loads may result in the same quantum. I agree that these examples, as well as the dependent claims of the ’345 Patent, indicate that *in proportion to* does not convey strict mathematical proportionality, in which the quantum is invariably the simple result of multiplying the user load by a constant. Nomadix accepts this as well, but emphasizes that rounding and the 1 MTU minimum are the

only variations from mathematical proportionality, and that these notions are known in computer science: Transcript, pp 1888–1889; Lavian First ’345 Report, para 5.73.

[353] The variations to mathematical proportionality represented by the 1 MTU minimum and rounding the quantum to a multiple of the MTU are largely dictated by the traffic management context and would be understood by the POSITA to pertain to the dequeuing of data in packets. In my view, despite these variations, the POSITA would consider the use of *in proportion to* to convey that the primary underlying relationship between the quantum and the user load is that of proportionality. This would be understood from the use of the mathematical term *proportion* and would not be disturbed by the other context from the ’345 Patent. To the contrary, it would be strengthened by the discussion of scaling factors and the examples which invariably show a calculation of quantum based on a proportional relationship with user load.

[354] Nevertheless, recognizing that various dequeuing strategies and desired traffic management alternatives are discussed, the POSITA would understand that the invention would work in effectively the same manner even if the proportional relationship between quantums and user loads associated with various zones is not precise. In other words, a POSITA would consider that a system practiced the invention of Claim 1 if it used quantums that showed a generally proportional relationship to user load, and that a system would not avoid practising the patent simply through minor variation to precise mathematical proportionality.

[355] In summary, based on the foregoing discussion, I conclude Claim 1 of the ’345 Patent would be understood by a POSITA to claim a bandwidth management system comprising (1) a

plurality of *queues* corresponding to *zones* (manageable divisions or users, devices, or lower-level zones); (2) an *enqueuing module* that determines which zone incoming traffic belongs to and places it in the queue for that zone; (3) a *quantum manager* that adjusts on an ongoing basis the value of a parameter called the *quantum* that defines the upper limit of data that will be dequeued from a given queue each time the dequeuing module dequeues from the queue before moving on to dequeue data from another queue; (4) a *dequeuing module* for managing the dequeuing of data from queues; (5) the dequeuing module dequeues at most the *quantum* when a selected queue has no guaranteed bandwidth rate or has reached it; and (6) the quantum manager adjusts quantums on an ongoing basis in general positive *proportion* (subject to rounding and minimum values) to the *user load* of the zone that the system keeps ongoing track of based on the current users in the zone.

> (7) Claim 3: user load under the zone to which the selected queue corresponds is tracked by <u>summing bandwidth caps</u> of <u>current users</u> in the zone to which the selected queue corresponds

[356] <u>Claim 3</u> adds a limitation to the system of Claim 1 that the *user load* is tracked by *summing bandwidth caps* of *current users in the zone*. There is no dispute between the parties that *current users* in the zone can be counted by one of the various methods described in the patent, such as logged in users (Claim 4), active users (Claim 5), or users responding to a ping (Claim 6).

[357] In Claim 3, the user load represents the total of the bandwidth caps of those users, rather than simply the number of users as in Claim 2. A *bandwidth cap* is an upper limit of bandwidth that a user may receive. Typically, a higher level of service within a hotel will involve a higher

bandwidth cap. Greater weight is given in the dequeuing process to those users with higher bandwidth caps by determining user load based on not just the number of users but their bandwidth cap, and adjusting the quantum based on this approach to user load.

[358] Similar to his approach to *quantum* and *user load*, Dr. Dordal argued that *summing* the bandwidth caps need not involve a mathematical addition resulting in a specific value or numerical result: Dordal First Report, paras 90–95. In his testimony, he explained this meant "some action is taken that yields a result [that] will be representative of the sum of the bandwidth caps, but it doesn't need to be explicitly by sequential addition": Transcript, p 884.

[359] I agree with Dr. Lavian that a POSITA with a computer science background would view the word *summing* in Claim 3 and consider it to be the adding of numbers, without requiring further explanation: Lavian First '345 Report, para 7.10 (p 76); Lavian Second '345 Report, paras 39–40. To the extent the POSITA needed to look further to understand this term, which I question, they would first look at Claim 2, which refers to *summing how many* current users are in the zone, confirming that *summing* involves a numerical assessment. They might also look to the disclosure of the '345 Patent, which gives examples of summation in which totals are "incremented by one" or "decremented by one," again confirming a numerical addition.

[360] In my view, Dr. Dordal's approach that allows for some action that "yields a result that will be representative of the sum of the bandwidth caps" is not how a POSITA would understand the straightforward expression used in the claim. In this regard, Dr. Dordal's example of how a summation may occur without adding numbers—by analogy to visually assessing which of two

shoes had accumulated more sand at the beach by pouring the sand into two containers—does not helpfully explain how a computer might perform a summation of bandwidth caps for tracking user load without adding those bandwidth caps together.

[361] Dr. Dordal’s approach to the term *summation of bandwidth caps*, combined with his approach to the storage of variables for *quantum* and *user load*, appears designed to capture any system or method that achieves the result described in the claims of the ’345 Patent. However, claims “cannot be stretched to allow the patentee to monopolize anything that achieves the desirable result”: *Free World Trust* at para 32.

(8) Remaining Asserted System and Method Claims

[362] There is little dispute over the construction of the remaining system and method claims asserted by Guest Tek. This consists of independent Claims 19, 20, and 21; and dependent Claims 16 to 18 (as each depends from Claims 1 and 3), 23, and 36 to 38 (as each depends from Claims 21 and 23).

[363] Claims 16 and 17 add limitations to the bandwidth management system in which the selected zone has a maximum bandwidth cap (Claim 16) or a guaranteed bandwidth allotment (Claim 17), and the dequeuing module ensures these caps or minimums are met. Claim 18 adds a limitation to the bandwidth management system that each zone corresponds to one or more rooms of a hotel. No issue was raised with respect to the construction or application of these additional limitations.

[364] Claim 19 claims a bandwidth management system with the same functions as that of Claim 1. However, while Claim 1 comprises an *enqueuing module*, a *dequeuing module*, and a *quantum manager*, the system of Claim 19 comprises *one or more processors configured to* perform the various functions ascribed in Claim 1 to the enqueuing and dequeuing module and quantum manager: receiving network traffic, determining a zone to which the traffic belongs, enqueuing the traffic on the queue corresponding to the zone, dynamically adjusting values of a plurality of quantums, selectively dequeuing data from the queues, and passing the data to an outgoing network interface. The same limitations on dequeuing up to the quantum when any guaranteed bandwidth rates are met, and dynamically adjusting the quantum in proportion to tracked user load are also present. There is therefore no relevant functional difference between Claim 19 and Claim 1, as the parties agreed.

[365] Claim 20 similarly provides for a bandwidth management system with the same functionalities. However, rather than providing for one or more processors configured to perform the aspects of the system, it provides simply for the *means for* undertaking them. Again, the parties do not argue any relevant difference between Claim 20 and Claims 1 or 19.

[366] Claim 21 sets out a *method of allocating bandwidth in a system having a plurality of queues respectively corresponding to a plurality of zones*. The method has elements equivalent to the system of Claim 1, without specifying any particular module or manager that is responsible for performing the steps of the method.

[367] Claim 23 adds the same limitation to the method of Claim 21 that Claim 3 adds to the system of Claim 1, namely the user load being tracked by summing bandwidth caps. Similarly, Claims 36, 37, and 38 add the same limitations to the method that Claims 16, 17, and 18 add to the system claims.

(9) Computer-Readable Medium Claim

[368] Claim 39 claims "[a] *computer-readable medium* comprising *computer executable instructions* that *when executed by a computer* cause the computer to perform the method of any one of claims 21 to 38." The parties agree that Claim 39 effectively claims a storage medium with software that implements the method of one of the method claims. However, they disagree on whether it is an essential element of the claim that the instructions actually be executed, *i.e.*, that someone actually run the software.

[369] With respect to the medium element, Guest Tek suggests a definition equivalent to that accepted by Justice LeBlanc in *Bessette* for the term "computer readable storage medium," namely "a storage medium such as a USB flash drive, floppy disk, optical disk or magnetic tape, containing data stored in a computer readable format": *Bessette v Quebec (Attorney General)*, 2019 FC 393 at para 151. Nomadix does not take issue with such a definition.

[370] Guest Tek argues the computer-readable medium need only contain instructions that cause the computer to perform the method *when* they are executed. In other words, they say the instructions must have a particular function when (or if) executed, but do not need to be executed to fall within the claim. Nomadix responds that Claim 39 cannot be construed to cover any

computer software capable of performing the method of claims 21 to 38. They argue allowing this would cause the claim to extend beyond what was invented and even cover the Linux TC utility, which can be configured to implement the methods of the patent.

[371] Claim 39 as drafted only requires that the computer readable medium contain instructions that will perform the claimed method when executed. In my view, a POSITA with the identified background in computer science would understand this to mean the instructions do not actually need to be executed, as long as they would perform the method if they were executed. That said, if there is material configuration of the software that is necessary before the instructions will perform the method, then the instructions present on the medium will not "perform the method" unless that configuration is actually undertaken. In this regard, I believe a distinction may be drawn between an example of a software module that performs a claimed method and simply requires a user to enable the module, and software that theoretically could perform a claimed method if an extensive series of configurations and choices are made by the user. In my view, the POSITA would understand Claim 39 to claim the former, but not the latter.

[372] I therefore conclude a POSITA would consider Claim 39 to claim a *computer-readable storage medium*, such as a flash drive, computer disk, or tape, on which is recorded software that will perform the method of one of claims 21 to 38 without requiring significant configuration or modification from a user.

E. *Infringement*

[373] Guest Tek does not argue that Nomadix's gateways always infringe the asserted claims of the '345 Patent. Rather, it argues there is infringement when particular features of the NSE software are enabled and the network runs at particular traffic levels. In addition, as the potential for the gateway to run in this configuration is always present, Guest Tek argues this gives a "standby utility" that constitutes infringement regardless of whether a Canadian hotel has ever run its network in this configuration. Assessing whether the Nomadix gateways and NSE software may infringe the '345 Patent, or whether Nomadix is inducing infringement, requires consideration of different aspects of the NSE software's functionality on the Nomadix gateway devices than those relevant to the '760 Patent. I will consider these functionalities, the parties' testing of Nomadix gateways, and how the Nomadix software functions as set out in its source code before turning to an assessment of whether the essential elements of the asserted claims are present.

(1) NSE Software: CBQ, WFQ, and SUB

[374] The NSE software has a variety of options with respect to traffic management. Three of these are of particular relevance to Guest Tek's allegation that Nomadix infringes the '345 Patent: Class-Based Queuing (CBQ), Weighted Fair Queuing (WFQ), and Share Unused Bandwidth (SUB).

[375] CBQ is a "core" feature of the NSE software: Exhibit 105, p 6. As described in the NSE User Guide, it provides the ability to define multiple classes of users, to prioritize them, and to

set minimum and maximum bandwidth on a per-group basis: Exhibit 105, p 8. Users are assigned to a class and the rules are then applied to all users across the class: Exhibit 105, p 8.

[376] WFQ is an option that may be selected from within the Bandwidth Management interface of the NSE software. WFQ allocates bandwidth to individual users or groups in proportion to their bandwidth limits: Exhibit 105, p 21. The NSE User Guide describes it as "a fallback in an over-subscribed scenario."

[377] If WFQ is enabled, an administrator may also enable the SUB option. If this option is checked, any available unused bandwidth is distributed among users in proportion to the users' bandwidth caps. If unchecked, then users are held to their bandwidth cap limits: Exhibit 105, p 70.

[378] CBQ, WFQ, and SUB may be run concurrently, and these combinations are discussed in Nomadix's NSE User Guide: Exhibit 105, pp 10–11, 74–75.

(2) Testing Using Nomadix Gateways

(a) *Guest Tek testing*

[379] Guest Tek performed six tests in which devices, meant to simulate network users, were connected via a switch to a Nomadix gateway. The gateway was in turn connected to a server that included a web server. An administrator laptop was also connected to the switch to log in to the gateway and select settings in the software. The different tests involved simulating different

"rooms" with different numbers of users and bandwidth plans, each user downloading data from the web server. The observed bandwidth speeds in each room were recorded: Dordal First Report, paras 114–122 and Appendices 4–5. In all tests, the CBQ and WFQ options in the NSE software were enabled. The tests assessed the results of disabling and enabling SUB, adding different rooms at different times, changing the number of active users, changing bandwidth caps, and using subclasses: Dordal First Report, paras 123–201.

[380] Dr. Dordal summarized the results of these tests as being that when CBQ, WFQ, and SUB are enabled, the higher a class's user load, the more unused bandwidth is assigned to the class, and that as user loads change (either due to the number of users or their bandwidth caps), the relative amounts of extra bandwidth given to the class change in a proportional manner: Dordal First Report, para 202. Dr. Dordal's conclusion of proportionality was generally based on a directional approach, *i.e.*, that higher user load resulted in higher bandwidth sharing. In some cases, he also noted that the results showed direct (mathematical) proportionality, within a "reasonable margin of error," which he considered to be 10%.

[381] Dr. Lavian criticized Dr. Dordal's assessment of the test results on a number of grounds. Notably, he criticized Dr. Dordal's assessment that there was proportional bandwidth allocation behaviour on the basis of his own calculations: Lavian Second '345 Report, paras 54–92. Unfortunately, in many of these calculations, Dr. Lavian used the wrong numbers from the test results, an oversight that was not pointed out until trial and was not corrected: Transcript, pp 904–906, 913–914, 1905–1919. While Dr. Lavian tried to provide his updated views during cross-examination, I am not satisfied I can rely on the conclusions he made based on the

erroneous numbers. In addition, these aspects of Dr. Lavian’s criticism were also based on his reading of *user load* as limited to the counted number of current users, rather than potentially incorporating the bandwidth caps of those users, a construction I have rejected.

(b) *Nomadix testing*

[382] Nomadix also conducted a series of tests using a Nomadix gateway and simulated network traffic from different users. As with Guest Tek’s tests, Nomadix’s seven tests varied user numbers and bandwidth caps. They also changed whether the CBQ, WFQ, and SUB settings were enabled or disabled: Dordal First Report, paras 293–350 and Appendices 6–7.

[383] As Guest Tek points out, its claim is premised on the actions of Nomadix’s gateway when the CBQ, WFQ, and SUB features are enabled. As a result, many of Nomadix’s tests do not directly relate to Guest Tek’s allegations. Dr. Dordal concluded the Nomadix tests in which CBQ, WFQ, and SUB were enabled, and where user load was not saturated (a concept discussed below), actually confirmed the proportionality shown in the Guest Tek testing: Dordal First Report, para 357(e). Dr. Lavian did not comment on the Nomadix tests, other than those related to the ability to load web pages with low bandwidth caps. Nomadix ultimately did not rely on the results of its testing in its closing arguments.

(3) NSE Source Code and Approach to Traffic Management

[384] [redacted]

[redacted]: Confidential Transcript (Oct 5, 2020), p 6; Confidential Transcript (Oct 9, 2020), pp 37–38; Dordal First Report, para 206.

[385] [redacted]

[386] Since packets are dequeued from each slot as they pass the dequeuing position, the manner in which they are enqueued is important. [redacted]

[387] [redacted]: Dordal First Report, paras 213–221.

[388] When WFQ is enabled in the NSE software, the SUB feature can also be enabled. When SUB is enabled, [redacted]: Dordal First Report, paras 225–226. [redacted] [redacted]

[389] As set out above, when the CBQ functionality is enabled on the NSE software, an administrator can define multiple classes and subclasses with different priorities and different bandwidth minimums and maximums. As with other class-based queuing strategies, classes are associated with queues and classes can have parents and children. In Nomadix’s CBQ approach, each class and subclass can be configured with a relative priority, from 1 to 8 for top-level classes, or from 1 to 3 for subclasses, with 1 being the highest priority: Exhibit 103, pp 8–11, 74.

[390] [redacted] [redacted]: Confidential Exhibit 77, pp 9–10. [redacted] [redacted]

[redacted]: Confidential Exhibit 77, p 10; Confidential Transcript (Oct 9, 2020), pp 31–34. [redacted]

[391] [redacted]: Confidential Transcript (Oct 9, 2020), p 37. [redacted].

[392] When WFQ is not enabled, the CBQ function uses [redacted]: Confidential Transcript (Oct 9, 2020), pp 38–40, 53. When WFQ is enabled, each class [redacted]: Dordal First Report, para 222. [redacted]

[redacted]: Dordal First Report, para 224.

[393] When CBQ, WFQ, and SUB are all enabled, the CBQ prioritization approach is in place [redacted]

[394] Dr. Dordal's discussion of Nomadix's system made a distinction between two scenarios. The first is where [redacted]. Dr. Dordal refers to this as operating at an "unsaturated" level. Conversely, if the volume of traffic being enqueued [redacted], the queue is considered to be operating at a "saturated" level: Dordal First Report, paras 242–243.

[395] Dr. Dordal noted that when a queue is operating at an unsaturated level, the burst limit [redacted] is proportional to relative user loads, using that term to reflect both the number of users and

their bandwidth caps: Dordal First Report, paras 35–36, 247–248, 262, 283–285. This proportionality does not generally hold when the queues are operating at a saturated level, and to the extent it does, it is the result of some other unexplained mechanism: Dordal First Report, para 329; Transcript, pp 1067–1068, 1121–1124.

[396] Dr. Dordal provided calculations indicating that if a single user has a bandwidth of 4.92 Mbit/s or higher, the queue for that class will likely be saturated most of the time, as the spacing between that user's packets as they are enqueued will be [redacted] Dordal First Report, para 244. Saturation may also be reached where multiple users in a class have lower bandwidths that cumulatively amount to 4.92 Mbit/s, although the likelihood of saturation is not simply an addition of each user's bandwidth: Dordal First Report, para 245. All of Guest Tek's testing described above involved individual bandwidths caps well below 4.92 Mbit/s (ranging from 98 to 1213 kbit/s).

(4) Infringement of the Asserted Claims

[397] Guest Tek's theory, based on Dr. Dordal's evidence and its arguments regarding construction, is that a hotel network using a Nomadix gateway meets all of the elements of the '345 Patent when CBQ, WFQ, and SUB are enabled, and bandwidths are low enough that calendar queues are unsaturated. It argues there is sufficient evidence or inferences that may be drawn to establish this situation exists in at least one hotel in Canada. It also argues there is infringement even if this situation is not established, based on the language of the claims and the concept of stand-by utility: *Monsanto* at para 58(5). Guest Tek's position is that the licensing of

the NSE software to Canadian hotels amounts not just to inducing infringement, but to direct infringement by Nomadix of the asserted claims of the ’345 Patent.

[398] Nomadix argues that in addition to enabling CBQ, WFQ, and SUB, and operating at an unsaturated level, Guest Tek’s theory of infringement requires a hotel to assign all classes the same priority, as Dr. Dordal’s report assumed this and Guest Tek’s testing was done on this basis: Dordal First Report, paras 234, 237, 240 and Appendix 5; Transcript, pp 1012–1013. Nomadix argues there is no evidence that a Canadian hotel has ever had CBQ, WFQ, and SUB enabled, while simultaneously configuring all classes at the same priority and all individual user caps below 4.92 Mbit/s and otherwise meeting the requirements for unsaturated levels. In any event, it argues that the essential elements of the asserted claims are not met even when CBQ, WFQ, and SUB are enabled and the system is operating in an unsaturated state.

[399] For the reasons below, I conclude that even if a hotel is operating a Nomadix gateway in the specified conditions, there is no infringement of the asserted claims of the ’345 Patent. I therefore do not need to determine (a) whether Guest Tek has established whether a hotel in Canada has actually used a Nomadix gateway in this manner, (b) whether it would need to do so to make out a claim of infringement based on the language of the claims or a theory of stand-by utility, or (c) whether licensing the software is sufficient to result in direct infringement.

[400] In my view, a network system operating with the features and under the conditions described by Guest Tek fails to meet at least the following essential elements of the claims: (a) queues having a respective *quantum* associated therewith; (b) *dynamically adjusting values of*

the quantums in proportion to *tracked user load*; and (c) *dequeuing at most an amount of data* from the selected queue *up to the quantum* of the selected queue before dequeuing data from another of the queues.

(a) *quantum*

[401] Guest Tek does not contend that Nomadix's system uses a defined parameter equivalent to the quantum of the Linux HTB system, which acts as a limit on the amount a dequeuer will dequeue from a queue each time it dequeues from that queue: Transcript, p 1041. Rather, it [redacted] onto a calendar queue [redacted]: Dordal First Report, paras 371–374; Confidential Transcript (Oct 5, 2020), p 37.

[402] In my view, Guest Tek's allegation is unsustainable. For the reasons set out above at paragraphs [315] to [329], the POSITA would understand the *quantum* of the '345 Patent to be a parameter that sets an upper limit on the amount of data to be dequeued by the dequeuer provided there is sufficient data in the queue. Nomadix's system does not have or use such a parameter. Enqueuing data in a manner that yields similar results to the use of a quantum does not make the amount of data enqueued or that happens to be dequeued on average over time a *quantum* as Dr. Dordal suggests: Dordal First Report, paras 375, 380.

[403] This is so even if the effect is that bandwidth rates proportional to user load are obtained. As both parties agreed, the same result may be obtained in a number of different ways. The '345

Patent only claims one way of obtaining that result, namely through the use of an adjustable quantum that limits the amount of data dequeued. Nomadix’s system uses a different approach entirely. [redacted]. Rather, as described above, [redacted]: Lavian Second ’345 Report, paras 178–181. The rate of dequeuing is [redacted]. Defining whatever happens to achieve the result as effectively being the quantum, using a construction inconsistent with how the term is used in the art and designed for the purpose of finding infringement, is untenable. It amounts to an effort to claim the result rather than the means of achieving it: *Free World Trust* at para 32.

[404] Further, even on Dr. Dordal’s definition, what he identifies as the *quantum* in the Nomadix system does not represent an upper limit or maximum amount dequeued on each dequeuing attempt. Rather, as Dr. Dordal himself states[redacted]. The use of this average implies the dequeuer sometimes takes less than the amount Dr. Dordal says is the quantum, and sometimes takes more. Dr. Dordal agreed that “[t]he average [redacted]

[redacted]: Transcript, p 1023. This is not an upper limit on the amount dequeued on each attempt.

[405] I therefore conclude Guest Tek has not established that Nomadix's system or method comprises a *quantum* associated with each queue.

(b) *dynamically adjusts the quantums in proportion to tracked user load*

[406] For similar reasons, Nomadix's system does not dynamically adjust the value of quantums in proportion to tracked user load.

[407] Evidently, since the system has no quantums, it cannot dynamically adjust those quantums. In addition, Nomadix's system [redacted]. Dr. Dordal considers that since this impacts the amount of data dequeued, [redacted] is effectively *dynamically adjusted* by the enqueuer in proportion to the tracked user load: Dordal First Report, paras 384–387; Confidential Transcript (Oct 5, 2020), pp 9–10, 37, 40–41. Again, in my view, this is not the system or method claimed in the '345 Patent.

[408] The '345 Patent teaches and claims tracking user load, using that user load to adjust the quantum parameter on an ongoing basis, and applying that quantum as a limit when dequeuing packets. What Dr. Dordal describes is the reverse. What he identifies as the quantum being dynamically adjusted with user load in the Nomadix system is simply the result of a system that

[redacted]. That is not the system of *dynamically adjusting values* of the *quantums in proportion to tracked user load* that is claimed in the ’345 Patent.

> (c) *dequeuing at most an amount of data from the selected queue up to the quantum of the selected queue before dequeuing data from another of the queues*

[409] As indicated above, what Dr. Dordal defines as the quantum in the Nomadix system is not a limit imposed on the dequeuer for each dequeuing of a queue. [redacted]. However, this does not meet the essential element that the dequeuer dequeue *at most* the quantum amount of data from the queue before dequeuing data from another queue.

[410] As these essential elements of the claims are common to all of the system and method claims, I conclude Guest Tek has not established that any claims of the ’345 Patent are infringed through the use by a Canadian hotel of a Nomadix gateway, even if the CBQ, WFQ, and SUB features are enabled, and even if the system is operated such that calendar queues are unsaturated. To the extent Dr. Dordal showed that under certain conditions Nomadix’s system showed “infringing behaviour,” this amounts to no more than showing that the results of

operating Nomadix’s system sometimes coincided with the results that a system of the ’345 Patent would achieve by other means. This does not establish infringement.

[411] I similarly conclude that Claim 39, the computer-readable medium claim, is not infringed. Regardless of any issues relating to the need to configure the software and operate it in a particular range, I conclude the NSE software does not include instructions that when executed will perform the method of either Claims 21 or 23.

(5) Changes to Nomadix’s Source Code

[412] On September 22, 2020, less than a week before the commencement of trial, Nomadix wrote to the Court to advise it had released a change to the NSE software that it believed had implications regarding infringement of the ’345 Patent. In particular, Nomadix asserted that the revision, version number 8.15.023, resulted in its gateways operating only in the “saturated” range that Dr. Dordal viewed did not infringe the ’345 Patent. Guest Tek objected to evidence regarding this version being produced so shortly before trial, when there was limited opportunity for Guest Tek and Dr. Dordal to review and/or test the software.

[413] At a trial management conference on September 24, 2020, I ruled Nomadix could file evidence it had changed its software. However, if Dr. Dordal’s evidence at trial was that he could not opine on the impact of the revisions, there would be no argument on whether the new version infringed and Dr. Lavian would not be permitted to opine on the issue.

[414] At trial, Dr. Dordal stated that while the revision appeared intended to cause the Nomadix gateway to operate in the saturated rather than unsaturated range, he could not say with any assurance what would happen when the revision was implemented, and would need to do further testing: Transcript, pp 954–957, 1131–1132; Confidential Transcript (Oct 5, 2020), pp 59–62; Confidential Transcript (Oct 6, 2020), pp 17–19. I therefore confirmed my ruling at trial that Dr. Lavian was not permitted to opine on the impact of the changes: Confidential Transcript (Oct 15, 2020), pp 18–30.

[415] Guest Tek argued the result of this was that there was no evidence the source code changes to version 8.15.023 worked differently than the versions Dr. Dordal concluded fell within the scope of the ’345 Patent. I disagree. There was evidence, including from Dr. Dordal, that the software would work differently, to the extent that his current theory of infringement would not apply: Confidential Transcript (Oct 6, 2020), pp 17–19. However, in light of my rulings and Dr. Dordal’s evidence, there was no expert evidence as to whether version 8.15.023 nonetheless continued to infringe. Had I concluded the earlier versions of the NSE software infringed the asserted claims of the ’345 Patent, I would have made no conclusion, one way or the other, as to whether version 8.15.023 infringed the patent.

[416] However, the revisions to version 8.15.023 of the NSE software do not affect the basic functioning of the software, and thus do not affect the reasons I have concluded the software does not infringe the ’345 Patent. My conclusions on infringement therefore apply equally to this version of the software.

F. *Inducing Infringement*

[417] As there is no direct infringement of any claims of the ’345 Patent even under the conditions identified by Guest Tek, Nomadix cannot have induced infringement. I therefore need not address Nomadix’s arguments that it did not influence Canadian hotels to use CBQ, WFQ, and SUB in combination and to cause them to operate at an unsaturated level.

[418] I consider it worth noting that as with Dr. Reiher’s evidence on this issue, I would have had difficulty placing any weight on Dr. Dordal’s evidence regarding inducing infringement. Dr. Dordal purported to reach conclusions about whether Nomadix’s Quick Start Guide or the requirement to download a license key amounted to Nomadix influencing Canadian hotels to use Nomadix gateways in a particular way. He also gave his opinion on whether a press release and/or Nomadix’s User Guide influenced Canadian users to use features of the gateway in a particular way, and even on what Nomadix does and does not know about the manner in which their gateways are operated: Dordal First Report, paras 434–474. He did so in language that was frequently similar or identical to that used in Dr. Reiher’s report. The question of influence is a factual matter not within Dr. Dordal’s expertise, nor a matter on which the Court requires expert opinion. While one can contemplate situations in which expert opinion might be needed in order to show how particularly technical statements might be understood by a consumer so as to amount to influence, this is not such a situation. Dr. Dordal agreed that most of this section of his report was prepared by counsel: Transcript, pp 1110–1111. The role of experts is not to make legal arguments that are better left to counsel.

G. *Validity*

(1) Anticipation

[419] In closing submissions, Nomadix argued the asserted claims of the ’345 Patent were anticipated by MA Brown, “Traffic Control HOWTO, v 1.0.2” (October 2006) [Brown]. Nomadix’s Statement of Defence and Counterclaim does not plead that the ’345 Patent is invalid owing to anticipation by any piece of prior art, and does not refer to Brown. Nor does Dr. Lavian’s report on validity opine that the ’345 Patent was anticipated by any particular piece of prior art in a way that might have given Guest Tek fair notice of the argument. I therefore consider that it is not open to Nomadix to argue the ’345 Patent was anticipated by Brown.

[420] In any event, Nomadix concedes there is at least one aspect of Claim 1 of the ’345 Patent not explicitly described in Brown, namely that quantums are based on tracked user load. Anticipation requires a prior art reference disclose subject matter which, if performed, would necessarily result in infringement of the claim: *Sanofi-Synthelabo* at para 25; *Hospira* at para 66. Since Brown does not disclose each essential element of Claim 1, it cannot anticipate the patent, regardless of how “straightforward” Nomadix argues the missing element to be: *Free World Trust* at para 26.

(2) Obviousness

[421] Nomadix’s primary argument on invalidity is that the ’345 Patent is rendered obvious by the CGK regarding Linux traffic control (TC) and/or the CGK combined with the prior art including in particular Brown and the following references:

- M Devera, “HTB Linux queuing discipline manual – user guide” (May 5, 2002) [Devera];
- V Ramachandran, R Pandey & S-H G Chan, “Fair Resource Allocation in Active Networks” (October 2000) [Ramachandran];
- Canadian Patent No 2,366,781 [Chiussi];
- US Patent No 7,457,313 [Patrick];
- US Patent No 6,865,185 [Patel]; and
- US Patent Application No 2009/0144425 A1 [Marr].

[422] To assess this argument, it is again necessary to apply the four-part approach from *Sanofi-Synthelabo*. The first step, identifying the POSITA and their CGK, is done at paragraphs [288] to [304] above.

[423] The second step requires the identification of the inventive concept of the claim in question. As the parties agree the validity of Claim 1 determines the validity of all claims, I will focus the analysis on that claim. In my view, the inventive concept of the claim lies in managing bandwidth through the use of different queues for different groupings of users, and the dynamic adjustment of dequeuing quantums for those queues in proportion to the user load of the group. Both the ongoing or dynamic adjustment of the quantum and the relationship between quantum and user load are important to this inventive concept.

[424] The third step of the *Sanofi-Synthelabo* approach considers the differences between the prior art and the inventive concept. As the ’345 Patent itself recognizes and both experts agreed,

it was known in the art to use a parameter known as the *quantum* as a limit on the amount of data dequeued, and to use different quantums for queues assigned to different classes as a traffic management tool. Devera and Brown each set out the use of quantums in this manner: Devera, ss 2, 7; Brown, s 7.1. Each discusses the possibility of setting quantums manually, although each notes that the HTB qdisc itself computes values.

[425] In addition to setting quantums manually, they can be set and reset on an ongoing or dynamic basis. In Linux, setting or changing the quantum, either once or as part of an ongoing dynamic adjustment process, can be done through a command in the Linux TC utility: Transcript, pp 781, 1745, 1751. Chiussi describes a bandwidth management approach in which ongoing adjustments are made to certain quantums, although as part of a different strategy: Lavian First ’345 Report, paras 8.12–8.16; Dordal Second Report, paras 90–93. In other bandwidth management approaches, quantums for a class are determined based on a particular calculation, but then remain static, as in Patrick: Lavian First ’345 Report, para 8.31; Dordal Second Report, paras 119–122. Dr. Dordal recognized there were examples of the dynamic adjustment of quantums in the prior art, although he said there were “very few” and did not suggest the ’345 Patent to him: Dordal Second Report, para 87; Transcript, pp 1071–1072.

[426] Dr. Lavian contended that the prior art showed examples of bandwidth management systems in which bandwidth was allocated in accordance with user load. This included a reference to “service share” in Patrick and elements of the wireless system in Patel: Lavian First ’345 Report, paras 8.33, 8.36, 8.46–8.47. On my review of Patrick, I agree with Dr. Dordal that its discussion of service share does not represent consideration of the concept of user load as that

term is used in the ’345 Patent: Dordal Second Report, paras 123–124. Patel, however, does discuss bandwidth management based on dynamic estimates of bandwidth demand using information such as which groups of traffic flows are active and how many flows (users) each group has: Patel, col 9. This is encompassed within the meaning of *user load* as I have construed it in the ’345 Patent.

[427] While dynamic adjustment of the dequeuing quantum is seen in the prior art, the difference between the prior art and the ’345 Patent lies in making that dynamic adjustment to the quantums for different queues in proportion to the current (tracked) user load of the queue.

[428] The fourth step in the obviousness analysis is assessing whether the difference identified above is a step that would have been obvious to the POSITA, or whether it would have required a degree of invention. Based on the evidence before me, I conclude Nomadix has not satisfied its burden to show the step would have been obvious.

[429] I begin by noting that Dr. Lavian’s analysis of why the differences between the prior art and the ’345 Patent would have been obvious was of little assistance. For example, he asserts, with little explanation, that the existence and popularity of Linux traffic control tools “makes the Asserted Claims obvious, and provides ample motivation to combine the prior art references identified herein with a high expectation of success”: Lavian First ’345 Report, para 8.9. However, he does not say why there was a motivation to combine the prior art references, or what in the Linux traffic control tools makes it obvious to undertake the particular solution identified in the ’345 Patent, which is to dynamically adjust the quantum proportionally between

classes or zones based on a tracked user load. Similarly, while Dr. Lavian states the Linux code is “usable” to adjust the values of the quantums in proportion to the number of users in each class as it changes over time, he does not explain how or why someone without knowledge of the ’345 Patent would actually use it in that way: Lavian First ’345 Report, para 8.10.

[430] Dr. Lavian’s substantive analysis of obviousness was brief, and hampered by his stated difficulty in understanding the notion of the “inventive concept”: Lavian First ’345 Report, paras 9.3–9.7. Dr. Lavian states that the inventor’s use of the number of users to calculate quantum was “very basic” and could not “be elevated to a patentable invention,” a statement he made “based on the contents of the patent and the common general knowledge alone”: Lavian First ’345 Report, paras 9.4–9.5. He expanded on this to some degree in oral testimony, noting that the calculations involved in tracking user load were simple: Transcript, pp 1754–1755. However, in my view, the simplicity of the calculation that implements an idea does not determine its inventiveness. The issue is whether the idea itself, and in particular the differences between the inventive concept and the prior art, would have occurred to a POSITA with no “scintilla of inventiveness”: *Hospira* at para 79; *Sanofi-Synthelabo* at paras 67, 76–80.

[431] Notably, while Dr. Lavian appeared to suggest that dynamically adjusting quantums as a means of traffic management was itself well known, the only examples of this Nomadix was able to point to were the Chiussi patent and Guest Tek’s system: Lavian First ’345 Report, paras 8.12–8.16; Transcript, pp 810–812, 1071–1072. In this regard, I agree with Dr. Dordal that the Ramachandran article appears to involve the adjustment of a quantum that controls allocation of CPU resources, not the amount of data to be dequeued: Dordal Second Report, paras 98–99.

The examples of tracking the user load of a group or class, whether considered in terms of simply the number of active users in the class or a combination of the number of users and their bandwidth caps, and using this as the basis for distribution of available bandwidth were also few and far from clear. These factors decrease the chance that a POSITA would consider combining these features to come up with the invention of the ’345 Patent.

[432] Nomadix points to the evidence that it took Mr. Ong a couple of weeks to write the software to implement the bandwidth management system of the ’345 Patent: Transcript, pp 752, 769. If an inventor reached an invention “quickly, easily, directly and relatively inexpensively,” this may suggest that it is obvious: *Sanofi-Synthelabo* at para 71. However, in addition to the fact that the indicated time appeared to relate to the writing of code rather than developing the solution itself, the Court had little evidence of the speed of invention or solution in the field of computer networking and bandwidth management. It is therefore difficult to rely on the fact that the code took a couple of weeks to write as a significant factor in support of or against a finding of obviousness.

[433] There was also little evidence led with respect to other factors relevant to obviousness, such as commercial success or the conduct of industry participants. I think there is some relevance in the fact that while solving bandwidth management issues has been a field of endeavour and improvement for some time, there was little evidence demonstrating that either use of dynamically adjusted quantums or attributing excess bandwidth on the basis of user load in a particular class were widespread or common approaches to solving these issues. No other industry participant, including Nomadix, apparently adopted or considered the same or a similar

approach to combining these concepts that is described in the ’345 Patent. Nor was there any evidence of a particular motivation in the field that would point toward this particular solution.

[434] Dr. Lavian made brief reference to the “obvious to try” concept, and Nomadix argued in final argument that Mr. Ong’s testimony as to the inventive process and how long it took supported a conclusion that the solution was obvious to try. In my view, reference to the “obvious to try” test does not assist Nomadix in this case. The nature of the solution of the ’345 Patent lies in using the well-known concept of dequeuing quantums and adopting an approach in which quantums are dynamically adjusted in proportion to tracked user load. While coding of any software implementing this solution would no doubt have to be tested, the concept itself does not require testing to determine whether it would work. In other words, the invention lies in the idea of the particular solution, rather than in demonstrating that it works.

[435] This is consistent with Dr. Dordal’s observations that bandwidth management is not generally a field in which advances are won by experimentation. Instead, software designers plan algorithms that make “logical sense,” with testing focused on clearing bugs in the software for achieving the pre-conceived result: Dordal Second Report, para 169. Generally, the obvious to try test is appropriate in areas where advances are often won by experimentation: *Sanofi-Synthelabo* at para 68; *Hospira* at para 88. In my view, the assessment of obviousness in this case should be directed at whether an uninventive POSITA would have reached the system and method of the ’345 Patent, rather than whether, having reached that solution, it would be expected to work when implemented.

[436] Considering all of these factors, I conclude Nomadix has not shown that the difference between the prior art and the inventive concept of the particular solution described in the ’345 Patent is a step that would have been obvious to the POSITA. It has therefore not met its burden to demonstrate the ’345 Patent is invalid on grounds of obviousness.

H. *Conclusion*

[437] For the foregoing reasons, I conclude Guest Tek has not shown that Nomadix has infringed the asserted claims of the ’345 Patent and Nomadix has not shown the ’345 Patent to be invalid.

VI. Disposition and Costs

[438] Guest Tek has not demonstrated Nomadix has infringed or induced infringement of any asserted claims of either the ’760 Patent or the ’345 Patent. As there is no need to proceed to the second liability phase of the action, the action is dismissed. As Nomadix has not demonstrated any asserted claims of either the ’760 Patent or the ’345 Patent are invalid, the counterclaim is also dismissed.

[439] I encourage the parties to discuss and agree on costs. If they are unable to do so, they may make written submissions on costs in accordance with the following schedule:

- within 30 days of the date of judgment, Nomadix may file submissions not to exceed 15 pages, to which it may attach a bill of costs as an appendix;

- within 15 days of receipt of Nomadix’s submissions, Guest Tek may file submissions not to exceed 15 pages, to which it may attach as an appendix a bill of costs and/or a submission, not to exceed two pages, addressing specific line items in Nomadix’s bill of costs (if filed); and
- within 5 days of receipt of Guest Tek’s submissions, Nomadix may file reply submissions not to exceed 5 pages.

JUDGMENT IN T-448-17

THIS COURT'S JUDGMENT is that

1. The action is dismissed.

2. The counterclaim is dismissed.

3. The parties may make submissions on costs in accordance with the schedule given in the reasons.

"Nicholas McHaffie"

Judge

FEDERAL COURT

SOLICITORS OF RECORD

DOCKET:	T-448-17
STYLE OF CAUSE:	GUEST TEK INTERACTIVE ENTERTAINMENT LTD v NOMADIX, INC

TRIAL HELD BY VIDEOCONFERENCE FROM SEPTEMBER 28 TO OCTOBER 15, 2020 AND FROM OCTOBER 27 TO OCTOBER 28, 2020 FROM OTTAWA, ONTARIO (COURT); CALGARY, ALBERTA (PLAINTIFF); AND MONTREAL, QUEBEC (DEFENDANT)

PUBLIC JUDGMENT AND REASONS:	MCHAFFIE J.
DATED:	MARCH 31, 2021

APPEARANCES:

D. Doak Horne Patrick Smith Kevin Unrau	FOR THE PLAINTIFF/DEFENDANT BY COUNTERCLAIM
Bob H. Sotiriadis Camille Aubin Justin Freedin Antoine Jean	FOR THE DEFENDANT/PLAINTIFF BY COUNTERCLAIM

SOLICITORS OF RECORD:

Gowling WLG (Canada) LLP Calgary, Alberta	FOR THE PLAINTIFF/DEFENDANT BY COUNTERCLAIM
Robic, LLP Montreal, Quebec	FOR THE DEFENDANT/PLAINTIFF BY COUNTERCLAIM