

infringement of all three patents. AWS is a Delaware corporation registered to do business in Illinois. (*Id.* ¶ 7.) The patents in suit disclose a distributed-data-storage technology that can be used for large-scale cloud storage. Kove has specifically identified two accused products: the Amazon S3 and the DynamoDB, which allow users to store data in the cloud. (*Id.* ¶¶ 27–28.) AWS has asserted numerous affirmative defenses and counterclaims, including counterclaims for declaratory judgment that each patent is invalid and not infringed. (See Answer and Counterclaims [129] at 37–54.)

The merits of these claims and defenses may turn on the interpretations of the disputed claim terms. The court held a claim construction hearing by videoconference on July 23, 2021, and now addresses construction of the following disputed terms: “location information,” “location server,” “client,” “based on a hash function used to organize the data location information across the plurality of data location servers . . . based on the hash function applied to the identifier string,” and “data pertaining to the entity.” (See Proposed Constructions, Ex. A to Updated Joint Claim Construction Chart [382] (hereinafter “Cl. Constr. Chart”).)¹

BACKGROUND

The inventions disclosed in the '640 and '170 Patents “relate[] generally to the storage and retrieval of information, and in particular, to a protocol for dynamic and spontaneous global search and retrieval of information across a distributed network regardless of the data format.” ('170 Patent, 1:30–33, Joint Appendix [221] at JA0015; '640 Patent, 1:26–30, JA0040.) The '170 Patent is a continuation of the '640 Patent, meaning that the '640 Patent is the “parent” patent and the '170 Patent is the “child” patent. (See '170 Patent, 1:6–9, JA0015.) The two patents share the same specification and drawings. The '640 and '170 Patents claim systems for managing or retrieving data stored in a distributed network and/or methods of operating such

¹ The parties report that they resolved their disputes about the following terms: “location,” “identifier/identifier string/identification string,” “hash table,” “data location sewer [sic] network,” and the preambles to claims 14 and 17 of the '978 patent. (See Agreed Upon Constructions, Ex. B to Cl. Constr. Chart.)

systems. The specification explains that the inventions disclose “an architecture in which information about where data associated with particular application entities can be managed and obtained independently of the data itself.” (’170 Patent, 3:15–18, JA0016; ’640 Patent, 3:15–18, JA0041.)

The preferred embodiment of the inventions is referred to as a “network distributed tracking protocol” or “NDTP.” (’170 Patent, 4:9–10, JA0016.) Figure 11 of the ’640 and ’170 Patents illustrates at a high level how the protocol works. (See ’170 Patent, Fig. 11, JA0013; ’640 Patent, Fig. 11, JA0038.)² A “client” (112) queries one or more location servers in the “NDTP server constellation” (110), seeking information about the location of data. An NDTP server responds to the query by providing location information, and the client uses that location information to retrieve the desired data from a “repository” (114). (See, e.g., ’170 Patent, 16:58–65, 17:3–9, JA0022–23.)




Fig. 11 NDTP Server Constellation Context

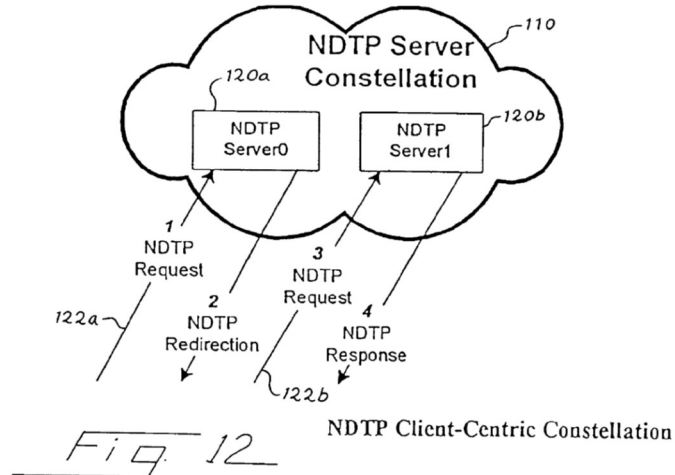
More specifically, an entity’s “identifier” is mapped to “locations” of repositories containing data pertaining to that entity. (See, e.g., ’170 Patent, 2:42–43, 4:8–26, JA0015, JA0016.) Particular “identifier/location mappings” are stored in “location servers” separately from the data itself. (*Id.*)

² The ’978 Patent, discussed below, contains a similar figure. (See ’978 Patent, Fig. 18, JA0067.)

The '640 and '170 Patents define “client” as “a network-attached component that initiates update or lookup of identifier/location mappings from an NDTP server with NDTP request messages.” ('170 Patent, 4:14–17, JA0016; '640 Patent, 4:11–14, JA0041.) In turn, an “NDTP server” is defined as “a network-attached component that maintains a set of identifier/location mappings that are modified or returned in response to NDTP request messages from NDTP clients.” ('170 Patent, 4:17–20, JA0016; '640 Patent, 4:14–18, JA0041.)

For example, suppose a doctor wants to request all medical records pertaining to a particular patient. That patient is an “entity” who is identified by an “identifier,” such as a Social Security number. The patient’s records are “data,” and the whereabouts of that data are “locations.” A patient’s data are associated with the patient’s identifier, and the identifier is mapped to data locations. These identifier/location mappings are stored on one or more “location servers.” A doctor can use a “client” to access patient data by querying one or more location servers with a patient’s identifier. The location server will return location information for data pertaining to the patient, and the doctor can then use that location information to retrieve the patient data from wherever it is stored. (See, e.g., '170 Patent, 2:14–26, 4:8–26.) From the doctor’s perspective, the process of requesting data and receiving that data is seamless, and the doctor may not perceive the intermediate step involving a location server.


Sometimes, a particular location server will not contain the necessary location information for the desired data. Querying every location server in the network would be inefficient, if not impossible, as the network grows to include many location servers. The inventions anticipate this problem and offer a solution, illustrated in Figure 12 of the '640 and '170 Patents. (See '170 Patent, Fig. 12, JA0013; '640 Patent, Fig. 12, JA0038.)



A client sends a request (arrow 1) to a location server (120a). If the first location server does not contain location information mapped to an entity's identifier, it will respond with a redirect message (arrow 2). The redirect message informs the client that the first location server did not contain the relevant location information, and it specifies a second location server that does contain such relevant location information. The client then queries (arrow 3) the second location server (120b), which returns the location information to the client.

The '978 Patent is a continuation-in-part of the '640 Patent. (See '978 Patent, 1:9–16, JA0072.) The '978 Patent “relates generally to the storage and retrieval of information, and in particular, to a system and method for managing global search and retrieval of information across a network.” (*Id.*) It builds off the '640 Patent by addressing the challenge of “scaling system capabilities in a manner sufficient to handle variable demand for resources” in a networked environment. (*Id.* at 2:11–13.) The specification explains that a redirect mechanism “permit[s] the distribution of identifiers to location mappings across members of an NDTP server cluster,” which consists of multiple location servers. (*Id.* at 15:53–55, JA0079.) “An advantage of distributing an NDTP server data set across independent mechanisms is that both capacity and transaction rate scale can be increased.” (*Id.* at 15:55–57.) The specification further explains that sets of identifier/location mappings can be transferred from one location server to another, as shown in Figure 22. (See *id.* at Fig. 22, JA0069.)

FIG. 22



In this illustration, Set B of identifier/location mappings (144) is originally stored on Server 1. If needed (for example, based on a performance criterion, such as Server 1's capacity), Set B could be transferred to a new location (148) on Server 2. The '978 Patent claims a system of multiple location servers for managing location information and providing location information in response to location queries, as well as methods of operating such systems.

Below, the court reproduces representative claims of the patents in suit.

A. The '640 Patent

The '640 Patent is titled "Network Distributed Tracking Wire Transfer Protocol." Claim 18 of the '640 Patent, an independent claim, recites:

A system for retrieving data location information for data stored in a distributed network, the system comprising:

a data repository configured to store data, wherein the data is associated with an identifier string;

a client responsive to a data query to query a data location server for location information associated with the identifier string;

a data location server network comprising a plurality of data location servers, at least one of the plurality of data location servers containing location information associated with the identifier string, wherein each of the plurality of data location servers comprises computer executable code configured to execute the following steps in response to receiving a data location request from the client:

if the data location server contains the location string associated with the identification string provided in the data location request, the data location

server transmits location information for use by the client to calculate a location of the data associated with the identification string;

if the data location server does not contain the location string associated with the identifier string, the location server transmits a redirect message to the client, wherein the redirect message contains redirect information for use by the client to calculate a location of a different data location server in the plurality of data location servers, wherein the different data location server contains the location string.

(’640 Patent, 22:41–23:2, JA0050–51.) Dependent Claim 24 of the ’640 Patent recites:

The system of claim 18, wherein the location information comprises a portion of a hash table distributed over the plurality of data location servers.

(*Id.* at 24:5–7, JA0051.)³

B. The ’170 Patent

The ’170 Patent, like the ’640 Patent, is titled “Network Distributed Tracking Wire Transfer Protocol.” Claim 1 of the ’170 Patent, an independent claim, recites:

A system for managing data stored in a distributed network, the system comprising:

a data repository configured to store a data entity, wherein an identifier string identifies the data entity; and

a data location server network comprising a plurality of data location servers, wherein data location information for a plurality of data entities is stored in the data location server network,

at least one of the plurality of data location servers includes location information associated with the identifier string,

each one of the plurality of data location servers comprises a processor and a portion of the data location information,

the portion of the data location information included in a corresponding one of the data location servers is based on a hash function used to organize the data location information across the plurality of data location servers,

³ The parties agree to the following construction of the term “hash table”: “a data structure that stores values in a table, where values are stored and retrieved by applying a hash function to an input and using the function result as an index into the table.” (Agreed Upon Constructions, Ex. B to Cl. Constr. Chart.) See *infra* Part E. The court discusses hash functions below. See *infra* Part D.

and each one of the data location servers is configured to determine the at least one of the plurality of data location servers based on the hash function applied to the identifier string.

(’170 Patent, 20:58–21:10, JA0024–25.) Dependent Claim 5 of the ’170 Patent recites:

The system of claim 1, wherein the location information includes a hash table and the hash table includes an association between a hash of the string identifier and at least one identifier of the at least one of the plurality of data location servers.

(*Id.* at 21:39–43, JA0025.) Claim 15 of the ’170 Patent, another independent claim, recites:

A method of handling location queries in a network, the network comprising a plurality of location servers including data location information, the method comprising:

correlating each one of a plurality of identifiers with at least one of a plurality of locations in the network, each one of the plurality of identifiers identifying a respective one of a plurality of data entities, wherein the data entities are stored in corresponding locations in the network;

receiving a location query from a client at one of the plurality of location servers, the location query requesting location information identifying a location of a data entity included in the data entities;

determining which of the plurality of location servers includes the location information;

sending a location response message to the client in response to determining the one of the plurality of location servers includes the location information, the location response message comprising the location information; and

sending a redirect message to the client in response to determining the one of the plurality of location servers fails to include the location information, the redirect message identifying which of the plurality of location servers includes the location information.

(*Id.* at 22:25–48, JA0025.)

C. The ’978 Patent

The ’978 Patent is titled “Method and Apparatus for Managing Location Information in a Network Separate from the Data to Which the Location Information Pertains.” Claim 1 of the ’978 Patent, an independent claim, recites:

A system having a plurality of location servers for managing location information and providing location information to location queries, the system comprising:

a first location server containing a first set of location information corresponding to at least one entity, the location information comprising an identifier and at least one location string associated with the identifier, wherein the identifier identifies an entity and the location string specifies a location of data pertaining to the entity[;]

a second location server comprising a second set of location information, wherein at least a portion of the second set of location information differs from the first set of location information; and

programming logic stored on each of the location servers responsive to a location query identifying a desired entity to return a location message, the location message comprising at least one location of data pertaining to the desired entity, if the location server receiving the location query contains location information for the desired entity.

(’978 Patent, 25:24–44, JA0084.) Claim 17 of the ’978 Patent, another independent claim, recites:

A method of scaling at least one of capacity and transaction rate capability in a location server in a system having a plurality of location servers for storing and retrieving location information, wherein each of the plurality of location servers stores unique set of location information of an aggregate set of location information, the method comprising:

providing a transfer protocol configured to transport identifier and location information, the location information specifying the location of information related to the identifier;

storing location information formatted according to the transfer protocol at a first location server;

receiving an identifier and a location relevant to the identifier at the first location server;

storing the received location in a location store at the first data location server, the location store comprising a plurality of identifiers, each identifier associated with at least one location, wherein the received location is associated with the received identifier in the location store; and

transferring a portion of the identifiers and associated locations to a second data location server when a performance criterion of the first location server reaches a predetermined performance limit.

(*Id.* at 27:19–43, JA0085.)

LEGAL STANDARD

A patent’s claims define the scope of the invention and the patentee’s right to exclude.

Phillips v. AWH Corp., 415 F.3d 1303, 1312 (Fed. Cir. 2005) (en banc). Where the meaning of a

claim is disputed, the court must determine its proper construction as a matter of law. *Markman v. Westview Instruments, Inc.*, 517 U.S. 370, 391 (1996). In construing a claim term, the court “look[s] to the words of the claim itself.” *Power Integrations, Inc. v. Fairchild Semiconductor Int’l, Inc.*, 711 F.3d 1348, 1361 (Fed. Cir. 2013). As a rule, courts should give claim terms their “ordinary and customary meaning,” which is “the meaning that the term would have to a person of ordinary skill in the art in question at the time of the invention.” *Phillips*, 415 F.3d at 1312–13 (internal quotation marks omitted). “[T]he person of ordinary skill in the art is deemed to read the claim term not only in the context of the particular claim in which the disputed term appears, but in the context of the entire patent, including the specification.” *Id.* at 1313.

Sometimes, a term’s ordinary meaning is equally apparent to a lay person as to a person of ordinary skill in the art (“POSITA”). See *id.* at 1314. In that circumstance, claim construction entails “little more than the application of the widely accepted meaning of commonly understood words.” *Id.* “[G]eneral purpose dictionaries may be helpful” for that exercise. *Id.* If a claim term “does not have an ordinary meaning, and its meaning is not clear from a plain reading of the claim,” the court should look to other sources of intrinsic evidence for guidance. *Power Integrations*, 711 F.3d at 1361; see also *Phillips*, 415 F.3d at 1314–19. The patent’s specification “is the single best guide to the meaning of a disputed term.” *Power Integrations*, 711 F.3d at 1361 (quoting *Vitronics Corp. v. Conceptoronic, Inc.*, 90 F.3d 1576, 1582 (Fed. Cir. 1996)); see also *Phillips*, 415 F.3d at 1315. The specification must include a “full, clear, concise, and exact” description of the claimed invention. *Phillips*, 415 F.3d at 1316 (quoting 35 U.S.C. § 112). If the specification reveals that the inventor has given a claim term a “special definition” that “differs from the meaning it would otherwise possess,” the “inventor’s lexicography governs.” *Id.* at 1316. The same is true if the specification shows that the inventor has limited the scope of the term. See *id.* Often, the specification “describes very specific embodiments of the invention.” *Id.* at 1323. But the Federal Circuit has “repeatedly warned against confining the claims to those embodiments.” *Id.*

The prosecution history is another helpful source of intrinsic evidence. See *id.* at 1317. The prosecution history “consists of the complete record of the proceedings before the [Patent and Trademark Office] and includes the prior art cited during the examination of the patent.” *Id.* Because the prosecution history “represents an ongoing negotiation between the PTO and the applicant,” however, “it often lacks the clarity of the specification and thus is less useful for claim construction purposes.” *Id.* Finally, in some circumstances the court is permitted to consider extrinsic evidence, so long as it is not used to contradict claim language that is “unambiguous in light of the intrinsic evidence.” *Id.* at 1318–19, 1324; see also *Vitronics Corp.*, 90 F.3d at 1583. Extrinsic evidence, which can include expert and inventor testimony, technical dictionaries, and treatises, is generally thought “less reliable than the patent and its prosecution history in determining how to read claim terms.” *Phillips*, 415 F.3d at 1317, 1318.

DISCUSSION

The parties’ disputes have narrowed over time. First, the parties have agreed upon the proper construction of the term “location”: “an encoding that is a member of a set of associations with an identifier in a location server, and that specifies where data pertaining to the entity identified by the identifier is stored.” (Agreed Upon Constructions, Ex. B to Cl. Constr. Chart.) They also agree that the terms “identifier,” “identifier string,” and “identification string” all mean “a unique encoding that identifies an individual entity, and with which zero or more location strings are associated in a location server.” (*Id.*) Kove has withdrawn its proposed construction of the term “redirect message” and agrees with AWS that it does not require construction. (See Kove Cl. Constr. Resp. [247] at 25.) There are no remaining disputes about definiteness. (See Cl. Constr. Hr’g [477] 17:7–11.) After the claim construction hearing, the parties informed the court that they have agreed upon the following construction of “hash table”: “a data structure that stores values in a table, where values are stored and retrieved by applying a hash function to an input and using the function result as an index into the table.” (Agreed Upon Constructions, Ex. B to Cl. Constr. Chart.) The court will now address the remaining disputed terms.

A. “Location information”

Claim Term	Plaintiff’s Proposed Construction	Defendant’s Proposed Construction
“location information” ’978 Patent: Claims 1, 3, 6, 10, 14, 17, 31 ’170 Patent: Claims 1–2, 6, 8, 15 ’640 Patent: Claims 17–18, 24	’978: “one or more identifiers and their associated locations” ’170/’640: “information pertaining to one or more locations of data and/or the identities of one or more location servers”	’978/’170/’640: “one or more identifiers and their associated locations”

AWS argues that the term “location information” should have the same meaning in all three patents in suit: “one or more identifiers and their associated locations.” For support, AWS points to two canons of claim construction: first, that courts should give claim terms their ordinary and customary meaning, *see Phillips*, 415 F.3d at 1312–13; and second, that the same terms in related patents should be construed in the same way, *see Omega Eng’g, Inc. v. Raytek Corp.*, 334 F.3d 1314, 1334 (Fed. Cir. 2003). Claim 1 of the ’978 Patent recites that “the location information compris[es] an identifier and at least one location string associated with the identifier.” (’978 Patent, 25:28–30, JA0084.) Based on this definitional language and similar usage of the term “location information” throughout the ’978 Patent, the parties agree that “location information” means (at least in the ’978 Patent) “one or more identifiers and their associated locations.” (*Id.*; *see also id.* at 2:21–28, 2:38–42, 4:55–57, 8:17–22, JA0072, JA0075.)⁴ The ’640 and ’170 Patents lack this definitional language contained in Claim 1 of the ’978 Patent—in fact, their specification does not use the term “location information” at all—but AWS argues that “location

⁴ As discussed above, if a patent’s specification reveals that the inventor has given a claim term a “special definition” that “differs from the meaning it would otherwise possess,” the “inventor’s lexicography governs.” *Phillips*, 415 F.3d at 1316. It is not clear to the court whether a claim, as opposed to an earlier part of a specification, can supply a definition for the purpose of the lexicography rule. The parties nevertheless apply this rule to the quasi-definitional language recited in Claim 1. (See Kove Cl. Constr. Resp. at 12; AWS Cl. Constr. Br. [219] at 9.).

information” should be construed the same way in all three patents because they are all related. (See AWS Cl. Constr. Br. [219] at 10.)

Kove agrees with AWS’s proposed construction of “location information” in the ’978 Patent, but Kove contends that the term should be construed differently in the ’640 and ’170 Patents. To reiterate, the ’640 Patent is the “parent” patent, the ’170 Patent is the “child” patent, and the ’978 Patent is a continuation-in-part of the ’640 Patent. The ’640 Patent and the ’170 Patent share the same specification, while the ’978 Patent has a different specification. Kove’s proposed construction of “location information” for the ’640 and ’170 Patents is broader: “information pertaining to one or more locations of data and/or the identities of one or more location servers.” According to Kove, the ’640 and ’170 Patents contemplate two types of “location information”: (1) the location of data stored in a distributed system, and (2) the identity of the location server that stores a particular data location, or “the location of the location.” (Kove Cl. Constr. Resp. at 12.) AWS’s proposed construction would cover only the first type of location information, not the second.

Kove points to intrinsic evidence in support of its broader construction. Claim 1 of the ’170 Patent, an independent claim, uses the phrase “*data* location information,” which Kove suggests refers specifically to the first type of location information. (*Id.* (emphasis in original).) Elsewhere in the same claim, the ’170 Patent refers to “location information,” which arguably has a broader meaning encompassing both types of location information. (See ’170 Patent, 20:66, JA0024.) By way of illustration, dependent Claim 5 of the ’170 Patent states:

The system of claim 1, wherein the *location information includes* a hash table and the hash table includes an association between a hash of the string identifier and *at least one identifier of the at least one of the plurality of data location servers.*

(’170 Patent, 21:26–30, JA0025 (emphases added).) Kove argues that the use of “location information” in Claim 5 refers to the second type, or the “location of the location.” And because Claim 5 depends on Claim 1, the reference to “location information” in Claim 1 must encompass both types of location information.

At the claim construction hearing, Kove clarified that the second type of location information includes “redirect information.” (See Kove Cl. Constr. Presentation at 7.) In other words, the second type of location information includes information about where to look for the desired data when the first location server does not contain that data. Kove points out that the specification of the ’640 and ’170 Patents describes the embodiment that is encompassed by Claim 5 of the ’170 Patent, whereas the ’978 Patent does not contain intrinsic evidence supporting this broader construction. (See ’170 Patent, 14:42–16:30, JA0021–22.) Accordingly, Kove contends, the court may construe the term “location information” differently in the ’640 and ’170 Patents than in the ’978 Patent.

AWS responds that Kove’s proposed construction for the ’640 and ’170 Patents “would allow Kove to expand its monopoly beyond what’s claimed, as ‘location information’ wouldn’t need to include a ‘location’ at all.” (AWS Cl. Constr. Reply [262] at 1.) From AWS’s perspective, redirect messages are a separate issue not relevant to the construction of “location information.” AWS points to the doctrine of claim differentiation. (*Id.* at 6 & n.20.) As relevant here, that doctrine provides that “the presence of a dependent claim that adds a particular limitation raises a presumption that the limitation in question is not found in the independent claim.” See *Acumed LLC v. Stryker Corp.*, 483 F.3d 800, 806 (Fed. Cir. 2007) (alteration omitted). “That presumption is especially strong when the limitation in dispute is the only meaningful difference between an independent and dependent claim, and one party is urging that the limitation in the dependent claim should be read into the independent claim.” *SunRace Roots Enter. Co., Ltd. v. SRAM Corp.*, 336 F.3d 1298, 1303 (Fed. Cir. 2003). According to AWS, claim differentiation dictates that a limitation in dependent Claim 5 of the ’170 Patent does not apply to independent Claim 1. AWS also notes that Kove’s construction is broader than Claim 5 itself, which contemplates a “hash table” with “an association between a hash of the string identifier” and “data location servers.” (AWS Cl. Constr. Reply at 6 (quoting ’170 Patent, 21:26–30, JA0025).) Claim 5 is also not one of the claims asserted in this litigation. Furthermore, AWS argues, Kove has not identified similar

claim language in the '640 Patent, so it is unclear why language from a dependent claim in the '170 Patent should inform the construction of a term in the '640 Patent.

Kove argues that to the extent there is a presumption that the same terms in related patents should be construed in the same way, that presumption is overcome here. See *IP Innovation v. Sony Elecs.*, No. 04 C 6388, 2005 WL 2035578, at *1, 5–7 (N.D. Ill. Aug. 18, 2005) (concluding that a claim term had a different meaning in a continuation-in-part patent than in the parent patent). Kove further argues that *Omega Engineering* is distinguishable because, in that case, the “patentee made a clear and unmistakable disclaimer of claim scope in its prosecution of the parent [] patent.” *Omega Eng’g*, 334 F.3d at 1334. AWS has cited no such prosecution disclaimer here. *NTP, Inc. v. Research in Motion, Ltd.* is similarly distinguishable, Kove argues, because the related patents in that case all “contain[ed] the same written descriptions.” See *NTP, Inc. v. Rsch. in Motion, Ltd.*, 418 F.3d 1282, 1289, 1293 (Fed. Cir. 2005). The '978 Patent, once again, has a different specification than the '640 and '170 Patents. And Kove’s construction conforms to the principles of claim differentiation, it contends, because usages of “location information” in independent claims encompass the narrower usages of “location information” in dependent claims. (Kove Cl. Constr. Presentation at 22.)

The court agrees with Kove’s proposed construction of “location information.” As a preliminary matter, the court is not convinced by AWS’s “plain meaning” argument for reading the term narrowly. (See AWS Cl. Constr. Br. at 9–10; AWS Cl. Constr. Reply at 5 & n.15.) AWS offers the analogy of a lawyer who calls chambers to ask where an upcoming hearing will occur. According to AWS, that lawyer is seeking the location of the hearing (presumably a courtroom number), not merely information “pertaining” to that location. But AWS does not explain why the precise piece of information that the lawyer *ultimately* seeks should narrow the scope of the category “location information.” Even if the best the clerk can do is to tell the lawyer to visit a specific whiteboard in the building’s lobby, which in turn will inform the lawyer of the courtroom number, the lawyer has still obtained location information from the clerk. The location information

has simply taken the form of an intermediate step—“the location of the location”—as Kove’s construction contemplates.

As the court understands AWS’s position, AWS believes that a narrower definition of a claim term should be imported from a continuation-in-part patent (here, the ’978 Patent) to the parent patent (the ’640 Patent). The court disagrees. *Omega Engineering* involved the reverse: in that case, the Federal Circuit applied a prosecution disclaimer from a parent patent to a continuation-in-part patent. See *Omega Eng’g*, 334 F.3d at 1333–34 (“[T]hat the ’678 patent is a continuation-in-part of the ’880 patent does not shield it from narrowing disclaimers made during the prosecution of a parent application.”).⁵ The opinion says nothing about applying a claim limitation from a continuation-in-part patent back to the parent patent. Moreover, AWS has not identified a prosecution disclaimer for the scope of the term “location information” in any of the patents in suit. By contrast, the patentee in *Omega Engineering* “made a clear and unmistakable disclaimer of claim scope in its prosecution of the parent ’880 patent,” leading the court to “presume, unless otherwise compelled, that the same claim term in the same patent or related patents carries the same construed meaning.” *Id.* at 1334.

The district court in *IP Innovation* derived the same principle from *Omega Engineering*. See *IP Innovation*, 2005 WL 2035578, at *2 (“[I]f the inventor unequivocally disavowed a particular meaning to obtain the patent or otherwise limited the scope of the invention during the course of the prosecution, ‘the doctrine of prosecution disclaimer attaches and narrows the ordinary

⁵ In its reply brief, AWS similarly relies on *NTP, Inc.*, 418 F.3d at 1293 (observing that when patents “derive from the same parent application and share many common terms, we must interpret the claims consistently across all asserted patents”). (See AWS Cl. Constr. Reply at 4 n.14.) But as in *Omega Engineering*, the context of the cited passage in *NTP* makes clear that the court was referring to prosecution disclaimers that limit the scope of related patents. See *NTP, Inc.*, 418 F.3d at 1293 (citing *Microsoft Corp. v. Multi-Tech Sys., Inc.*, 357 F.3d 1340, 1350 (Fed. Cir. 2004) (holding that statements made in prosecution of one patent are relevant to the scope of all sibling patents); *Laitram Corp. v. Morehouse Indus., Inc.*, 143 F.3d 1456, 1460 & n.2 (Fed. Cir. 1998) (noting that it was proper to consider the prosecution histories of two related re-examination patents originating from the same parent to determine the meaning of a term used in both patents)).

meaning of the claim congruent with the scope of the surrender.’” (quoting *Omega Eng’g*, 334 F.3d at 1324)). In *IP Innovation*, the district court gave different constructions to the same term in two related patents where the parent patent incorporated a prosecution disclaimer but the continuation-in-part patent contained significant differences in its specification. *Id.* at *6–7. The court acknowledged that there was a presumption “that a claim term carries the same meaning throughout a particular patent and related patents, including a continuation-in-part.” *Id.* at *7 (citing *Omega Eng’g*, 334 F.3d at 1334). “But this presumption may be overcome by evidence that the patentee clearly assigned different meanings to a term that appears in two related patents,” such as a different specification for the continuation-in-part patent. *Id.* Here, Kove has identified enough intrinsic evidence in the ’170 Patent to suggest that the claims contemplate two different types of location information. And because the ’170 Patent and ’640 Patent share the same specification, it is appropriate to apply this broader construction to both patents.

At the claim construction hearing, AWS argued that the court should instead consider how “location information” is used in independent Claim 15 of the ’170 Patent. (See AWS Cl. Constr. Presentation at 23.) Unlike Claim 5, Claim 15 is asserted in this litigation. Kove responded that whether a particular claim is being asserted in an infringement action is irrelevant for claim construction purposes. (Cl. Constr. Hr’g 34:8–12.) Neither party has directed the court to case law resolving this issue. Regardless, Claim 15 recites in relevant part:

A method of handling location queries in a network . . . the method comprising: . . . sending a location response message to the client in response to determining the one of the plurality of location servers includes the location information, *the location response messages comprising the location information*; and sending a redirect message to the client in response to determining the one of the plurality of location servers fails to include the location information, *the redirect message identifying which of the plurality of location servers includes the location information*.

(’170 Patent, 22:25–43, JA0025 (emphasis added).) Claim 15 uses the term “location information” when describing the contents of location response messages (which contain the exact location of data). But it does not use the term “location information” when describing the contents of redirect messages (which identify the server that contains the exact location of data). AWS therefore

contends that Claim 15 favors its construction of “location information,” which is limited to the exact location of data. But Kove’s proposed construction is not inconsistent with Claim 15. Through the inclusive “and/or” term, Kove’s construction allows for “location information” to mean the exact location of data, the identity of the server(s) containing the exact location, or both the exact location and the identity of the server(s) containing it. It is true that Claim 15 tracks most closely with the first of these options, but that does not nullify the other usages of “location information” in the ’640 and ’170 Patents.

In their claim construction briefing, neither party has identified claims in the ’640 Patent that might support their proposed constructions. The court itself observes that one of the asserted claims, independent Claim 18, claims “[a] system for retrieving data location information.” (’640 Patent, 22:41–23:2, JA0050–51.) In turn, dependent Claim 24 recites: “The system of claim 18, wherein the *location information comprises a portion of a hash table* distributed over the plurality of data location servers.” (’640 Patent, 24:5–7, JA0051 (emphasis added).) As the court reads that language, it lends support to Kove’s construction of “location information”: Claim 24 seems to contemplate that “location information” will include not only the exact location of data, but also “information pertaining to the one or more locations of data,” such as a hash table.

The court shares AWS’s concern that the use of two different constructions for the same term could be confusing for the jury. But AWS’s construction would also be problematic in that it ignores differences between the specification of the ’640 and ’170 Patents and the specification of the ’978 Patent. The court also rejects AWS’s argument that Kove’s construction violates the doctrine of claim differentiation with respect to Claims 1 and 5 of the ’170 Patent. Kove’s construction does not, as AWS seems to contend, import a limitation from a dependent claim to an independent claim. “[I]ndependent claims are presumed to have broader scope than their dependents.” *Acumed*, 483 F.3d at 806. It follows that the category of “location information” recited in Claim 1 encompasses the more specific type of “location information” referenced in Claim 5. In other words, “location information” includes not only the location of data stored in a

distributed system, but also “information pertaining to . . . the identities of one or more location servers.” (Cl. Constr. Chart at 1.)

The court adopts Kove’s proposed construction of “location information.”

B. “Location server”

Claim Term	Plaintiff’s Proposed Construction	Defendant’s Proposed Construction
“location server” ’978 Patent: Claims 1, 3, 6, 10, 14, 17, 23, 30, 31 ’170 Patent: Claims 1, 6, 8, 12, 15 ’640 Patent: Claims 17, 18, 24	“a network-attached component that maintains a set of identifier/location mappings that are modified or returned in response to location request messages from clients”	“a network-attached component that maintains a set of identifier/location mappings that are modified or returned in response to NDTP request messages from clients”

The parties generally agree as to the proper construction of “location server”: “a network-attached component that maintains a set of identifier/location mappings that are modified or returned in response to [] request messages from clients.” The only difference in their proposed constructions is that Kove prefers “*location* request messages,” while AWS would use “*NDTP* request messages.”⁶ As a reminder, NDTP is short for “network distributed tracking protocol,” the preferred embodiment of the patents in suit. (See, e.g., ’170 Patent, 4:9–10, JA0016.)

AWS argues that its construction matches the precise wording of the ’978 Patent’s specification: “An ‘NDTP server’ or a ‘server’ is a network-attached component that maintains a set of identifier/location mappings that are modified or returned in response to NDTP request messages from clients.” (’978 Patent, 4:34–38, JA0073; see also ’170 Patent, 4:17–20, JA0016; ’640 Patent, 4:14–18, JA0041.) According to AWS, Kove’s construction unnecessarily replaces “NDTP” with “location,” which “would invite further disputes as the case moves forward.” (AWS

⁶ AWS has withdrawn its request for a construction specifying that a “location server” “does not store data to which locations pertain.” (AWS Cl. Constr. Reply at 7 n.21.)

Cl. Constr. Reply at 7.) At the hearing, AWS pointed to *Martek Biosciences Corp. v. Nutrinova* for the proposition that “[w]hen a patentee explicitly defines a claim term in the patent specification, the patentee’s definition controls.” 579 F.3d 1363, 1381 (Fed. Cir. 2009) (citing *Phillips*, 415 F.3d at 1321). AWS also notes that Dr. Michael T. Goodrich, Kove’s expert witness, testified at his deposition that he could explain “NDTP” to a jury. (See Goodrich Dep. 82:22–83:4, Ex. 3 to AWS’s Cl. Constr. Reply [262-1].)⁷

Kove responds that the terms “NDTP server” and “location server” are used interchangeably in the specifications, but only “location server” appears in the claims themselves. (Kove Cl. Constr. Resp. at 6 n.6). Kove acknowledges that an NDTP server is functionally the same as a location server, and that “NDTP request messages” are therefore synonymous with “location request messages.” Still, Kove contends, the claims refer to a “location query” rather than an “NDTP query.” (Kove Cl. Constr. Sur-Reply [265] at 5 & n.3.) And Kove contends that using “NDTP” could confuse the jury. (See *id.* at 5.) Because NDTP is “not a term of art and has no plain meaning,” use of that term would counterproductively require further construction in order for a POSITA or a juror to make sense of it. (*Id.* at 5 (citing *Ignite USA, LLC v. Pac. Mkt. Int’l, LLC*, No. 09 C 03339, 2018 WL 2412375, at *6 (N.D. Ill. May 5, 2018) (“The purpose of claim construction is to help clarify the scope of the claims, not to add further confusion.”))).

The court agrees with Kove’s proposed construction for three reasons. First, now that the parties have agreed upon the construction of “location,” AWS’s concerns about further disputes over the term “location request messages” seem unwarranted. (See Agreed Upon Constructions, Ex. B to Cl. Constr. Chart.) Second, the court agrees with Kove that the jury could be unnecessarily confused by the use of “NDTP” in AWS’s proposed construction of “location server.”

⁷ Dr. Goodrich went on to say, however, that NDTP “is synonymous with location as an adjective, because NDTP server is synonymous with location server.” (Goodrich Dep. 83:22–84:2.) Similarly, he believes that “a [POSITA] would understand that NDTP server in this express definition is synonymous with location server. And so the appropriate thing, since location server is the claim term, is to continue using that adjective throughout the definition rather than introducing NDTP from a preferred embodiment into the claim.” (*Id.* at 84:20–85:5.)

Viewing the patents in light of their specifications, a POSITA would understand the term “NDTP request messages” as synonymous with “location request messages,” so a more straightforward construction is preferable. Finally, the specifications describe NDTP as a preferred embodiment, not the only embodiment, so despite the fact that the terms appear to be largely synonymous, it would be improper to limit the construction of “location server” with the use of “NDTP.” See, e.g., *Martek*, 579 F.3d at 1381 (“[P]articular embodiments appearing in the written description will not be used to limit claim language that has broader effect.” (quoting *Innova/Pure Water, Inc. v. Safari Water Filtration Sys., Inc.*, 381 F.3d 1111, 1117 (Fed. Cir. 2004))).

The court adopts Kove’s proposed construction of “location server.”

C. “Client”

Claim Term	Plaintiff’s Proposed Construction	Defendant’s Proposed Construction
“client” ’978 Patent: Claim 14 ’170 Patent: Claim 15 ’640 Patent: Claims 17–18, 24	“a network-attached component that initiates update or lookup of identifier/location mappings from a location server with location request messages”	Plain and ordinary meaning or, in the alternative: “a network-attached component that initiates update or lookup of identifier/location mappings from a NDTP server with NDTP request messages”

AWS contends that the claim term “client” needs no construction, such that the jury may refer to its plain and ordinary meaning. If the court determines that the term requires construction, AWS argues in the alternative that “client” be construed as “a network-attached component that initiates update or lookup of identifier/location mappings from a NDTP server with NDTP request messages.” This construction comes from the express definition of “NDTP client” or “client” in the specifications. (See ’170 Patent, 4:14–17, JA0016; ’640 Patent, 4:11–14, JA0041; ’978 Patent, 4:30–34, JA0073.)⁸ As with the term “location server,” the parties generally agree on the proper

⁸ The ’978 Patent’s definition of “NDTP client” or “client” differs slightly: “a network-attached component that initiates *add, delete, lookup and update* of identifier/location mappings” (’978 Patent, 4:30–34, JA0073 (emphasis added).) Dr. Goodrich, Kove’s expert,

construction of “client” with one exception: AWS prefers to use the modifier “NDTP” before “server” and “request messages,” while Kove would use “location.”

The court agrees with Kove that the term “client” requires construction. In ordinary speech, the word “client” could refer to “a person who engages the professional advice or services of another,” such as an attorney’s client; “one that is under the protection of another,” such as a nation state; or “a computer in a network that uses the services provided by a server.”⁹ One of the inventors similarly testified that “client” “only [has] meaning in context . . . of something. It doesn’t really mean anything in isolation.” (Bailey Dep. [220] 93:8–11; see *also* Goodrich Dep. 101:13–15 (similar).) Even more importantly, because the inventors acted as their own lexicographers by defining “NDTP client” or “client” in the specifications, their lexicography governs. See *Phillips*, 415 F.3d at 1316.

Next, the court turns to whether the construction of “client” should use “NDTP” or “location” as a modifier of “server” and “request messages.” As it did with respect to the term “location server,” the court adopts Kove’s construction of “client”: “a network-attached component that initiates update or lookup of identifier/location mappings from a location server with location request messages.” Kove’s construction matches the express definition from the specifications, with the minor clarification that “NDTP” is synonymous with “location” when used as an adjective modifying “server” and “request messages.” (See ’170 Patent, 4:14–17, JA0016; ’640 Patent, 4:11–14, JA0041; ’978 Patent, 4:30–34, JA0073.) By contrast, AWS’s proposed construction would improperly narrow the scope of the claims by importing the preferred embodiment (“NDTP”) from the specifications. Kove’s construction not only mirrors how a POSITA would understand the claims, but also should reduce the potential for juror confusion.

believes that a POSITA would understand “update” in the ’640 and ’170 Patents to include “add” and “delete,” so the addition of those two words in the ’978 Patent is not a substantive change. (See Goodrich Decl. [247-1] ¶ 31.)

⁹ *Client*, MERRIAM-WEBSTER.COM, <https://www.merriam-webster.com/dictionary/client> (last visited Dec. 14, 2021).

Relatedly, Kove contends that by using the term “network-attached component,” rather than “network-attached computer” or “machine,” the inventors intended for “client” to include both hardware and software. (See Goodrich Decl. ¶ 32; ’170 Patent, 2:30–45, JA0015; ’640 Patent, 2:28–36, JA0040.) For example, the specification of the ’978 Patent explains that “the client . . . may be any mechanism capable of communicating with the NDTP server.” (’978 Patent, 5:32–36, JA0074.) Kove points to several contemporaneous technical dictionaries in support of its interpretation. (See Kove Cl. Constr. Sur-Reply at 7.)¹⁰ Kove raised a similar argument in support of its proposed construction of “location server.” (See Kove Cl. Constr. Resp. at 14.) AWS disagrees and insists that the parties’ dispute over whether a “component” may be software is best resolved at summary judgment or at trial. (See AWS Cl. Constr. Br. at 14; AWS Cl. Constr. Reply at 8–9.)

The court concludes that it can resolve this dispute now, and concludes, further, that Kove is correct: a “client” or “component” may include software. Even *Merriam-Webster*, a non-technical dictionary, offers the following definition of “client”: “software that allows a computer to function as a client in a network.”¹¹ In order to reflect this understanding, the court construes the term “client” as “a network-attached component (*which may be software or hardware*) that initiates update or lookup of identifier/location mappings from a location server with location request messages.” With that addition, the court adopts Kove’s proposed construction of “client.”

¹⁰ The *IBM Dictionary of Computing* (1994) defines “component” as “[h]ardware or software that is part of functional unit.” (Ex. A to Kove Cl. Constr. Sur-Reply [265-1] at 127.) The *Microsoft Press Computer Dictionary* (1997) defines “component” as “[a]n individual modular software routine that has been compiled and dynamically linked, and is ready to use with other components or programs.” (Ex. B. to Kove Cl. Constr. Sur-Reply [265-2] at 106.)

¹¹ *Client*, MERRIAM-WEBSTER.COM, <https://www.merriam-webster.com/dictionary/client> (last visited Dec. 14, 2021).

D. “Based on a hash function . . .”

Claim Term	Plaintiff’s Proposed Construction	Defendant’s Proposed Construction
<p>“the portion of the data location information included in a corresponding one of the data location servers is based on a hash function used to organize the data location information across the plurality of data location servers, and each one of the data location servers is configured to determine the at least one of the plurality of data location servers based on the hash function applied to the identifier string”</p> <p>’170 Patent: Claim 1</p>	<p>“the portion of the data location information included in a corresponding one of the data location servers is based on a hash function that maps identifier strings to one or more of the data location servers, and each one of the data location servers is configured to determine the at least one of the plurality of data location servers based on the hash function applied to the identifier string”</p>	<p>“the portion of the data location information included in a corresponding one of the data location servers is based on a hash function applied to an identifier string in which the function result organizes the data location information by giving an index into a data structure that provides the location of an indicated location server, and each one of the data location servers is configured to determine at least one of the plurality of data location servers by applying the hash function to the identifier string in which the function result is an index into a data structure on the location server containing the location of an indicated location server for the identifier string”</p>

As noted earlier, the parties agree on the meaning of the term “hash table.” They disagree, however, about the meaning of the term “based on a hash function . . .” in Claim 1 of the ’170 Patent. (See ’170 Patent, 20:58–21:10, JA0024–25.) By way of background, a hash function is a function that converts inputs of varying lengths (such as a person’s name) to fixed-size outputs (such as an integer). (See Mitzenmacher Decl. [219-2] ¶ 10.) As Dr. Goodrich explains in one of his textbooks: “The use of such a function allows us to treat objects, such as strings [i.e., sequences of characters], as numbers. . . . Equipped with such a function, . . . we can apply the lookup table approach” to strings of varying size. (*Id.* ¶ 17 (quoting Michael T. Goodrich and Roberto Tamassia, *Algorithm Design and Applications* 192 (John Wiley & Sons 2015).) According

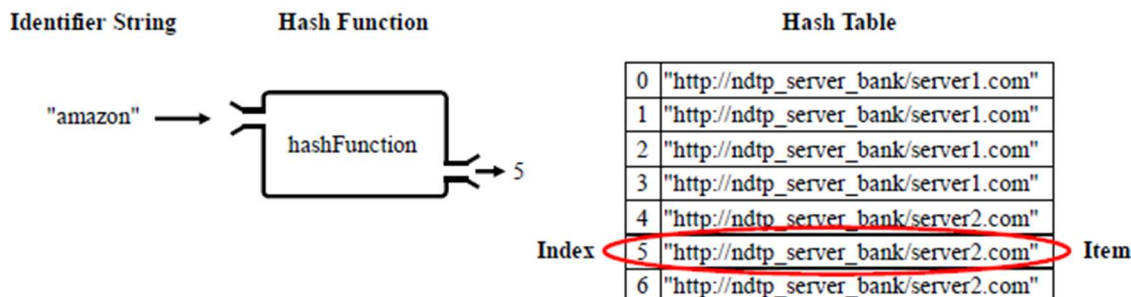
to a technical reference that was available when the patent application was filed, “hashing” is “[a] technique that is used for organizing tables to permit rapid searching or table lookup, and is particularly useful for tables in which items are added in an unpredictable manner.” (*Id.* ¶ 18 (quoting *Oxford Dictionary of Computing* 221 (4th ed. 1997)).) A POSITA would be aware of many possible ways to organize information that involve the application of hash functions. (*Id.* ¶ 24; see also Goodrich Decl. ¶ 36.)

The specification of the '170 Patent explains that “[t]he goal of NDTP is to support a network service that efficiently manages mappings from each individual key string, an identifier, to an arbitrary set of strings, locations.” ('170 Patent, 20:10–13, JA0024.) In short, hashing is a way to map identifiers to location servers more efficiently. The '170 Patent teaches that a client applies a function to an identifier, the result of which is the location information for data pertaining to an entity. (See *id.* at 15:4–15, 15:50–56, JA0022.) The specification discloses two variants for this process.¹²

In the first variant, a client applies “a well-known function that all NDTP server and client implementations know when they are programmed, and the [redirect] message carries a table of NDTP server URLs.” (*Id.* at 15:4–9 (emphasis added).) “The well-known function preferably applied is the hashpjw function,” which is a specific algorithm (or sequence of steps) for a hash function. (*Id.* at 15:15–16.) The specification provides source code for this algorithm. (See *id.* at 15:21–36.) After applying the hash function to an identifier, the function result is used as an index into a location server table. (*Id.* at 15:12–15.) AWS’s expert, Dr. Michael Mitzenmacher, created the following illustration to demonstrate how this variant works.

¹² A third redirection mechanism, “embedded redirection links,” is not relevant to the parties’ proposed constructions of the term “based on a hash function . . .” because it does not describe the use of a hash function. (See '170 Patent, 14:46–62, JA0021.)

Figure 2. Example Illustrating Patent Specification



(Mitzenmacher Decl. ¶ 15.) Kove does not dispute this characterization of the first variant, which Kove refers to as “indirect mapping.” (Kove Cl. Constr. Presentation at 51.)

Alternatively, the second variant “specifies that a *general function* description will be sent to the NDTP client in the [redirect] message. The NDTP client will apply this function to the identifier string and *the output of the function will be the NDTP server URL* to which to send NDTP requests for the particular string.” (’170 Patent, 15:4–9, 15:50–56, JA0022 (emphasis added).) In other words, the function result is not an index into a table but a URL to the location server itself. Kove refers to the second variant as “direct mapping.” (Kove Cl. Constr. Presentation at 49–50.)

The parties generally agree that in Claim 1 of the ’170 Patent, a location server applies a hash function to an identifier string, and the function result identifies the location server in a network that contains data pertaining to that identifier string. AWS’s construction would narrow the claim by requiring that the output of the hash function be an index into a data structure containing the identity of a location server. Put differently, AWS’s construction would cover only indirect mapping. Kove’s construction would not require that a hash function produce an index; rather, Kove argues, the claim is directed to “*any way of ‘organiz[ing]’ data location information using a hash function that would have been known to a skilled artisan.*” (Kove Cl. Constr. Resp. at 19 (emphasis in original) (quoting ’170 Patent, 21:5, JA0025).) Kove’s construction would therefore include both direct and indirect mapping.

Kove argues that its construction clarifies how data location information is organized while preserving the breadth of the claim. According to Kove, the hash function described in Claim 1 “maps identifier strings to . . . location servers,” but the claim does not require use of any specific algorithm (such as the hashpjw function) or any particular type of mapping (direct or indirect). In the absence of a “clear indication in the intrinsic record” that the patentees intended to limit Claim 1 to the specific algorithm disclosed in the specification, Kove says, the court should look to the ordinary and customary meaning of the term. See *GE Lighting Solutions, LLC v. AgiLight, Inc.*, 750 F.3d 1304, 1309 (Fed. Cir. 2014) (citation omitted). The district court in *Symantec Corp. v. Acronis, Inc.*, which involved patents related to a system for “fast incremental backup of a data storage device,” applied similar logic. See *Symantec Corp. v. Acronis, Inc.*, No. 12-cv-05331-JST, 2014 WL 230023, *3 (N.D. Cal. Dec. 13, 2013). By construing the term “hash function values” to mean “value generated by a hash function,” the court declined to import a limitation from the specification, which in the defendant’s view required that a value generated by a hash function be shorter than the input. *Id.* at *5–6. As the court noted, the specification merely stated that “a hash function is a ‘usually shorter value of a fixed length.’” *Id.* at *6. Adopting the defendant’s proposed construction “would therefore impose a limitation that the specification does not clearly encompass.” According to Dr. Goodrich, a POSITA would have understood that multiple hashing algorithms could satisfy the limitations of Claim 1. (See Goodrich Decl. ¶¶ 38–43.) Kove also points out that AWS’s construction would exclude a preferred embodiment—direct mapping—that is discussed in the specification. See *CSS Tech., Inc. v. Panduit Corp.*, 778 F. App’x 947, 950 (Fed. Cir. 2019) (“[A]n interpretation which excludes a disclosed embodiment from the scope of the claim is rarely, if ever, correct.” (citations and internal quotation marks omitted)). Finally, Kove notes that AWS’s construction is very long and would be hard for the jury to comprehend.

AWS responds that the claim language is ambiguous and that the construction AWS has proposed describes “the standard way of using hash functions to organize information at the time

of the patents.” (AWS Cl. Constr. Br. at 19.) AWS emphasizes that the words “*based on* a hash function” suggest the use of a hash table to organize the location information. In other words, a POSITA would understand “based on a hash function . . .” to mean indirect mapping with an index, not direct mapping without an index. AWS points to several technical dictionaries defining “hash function” or “hashing” in a way that requires the use of an index. (See Mitzenmacher Decl. ¶¶ 16, 18.)¹³ AWS also suggests that Kove’s construction would invalidate the claim for “lack of written description and enablement,” because the specification does not describe another way to use a *hash* function to organize location information. (AWS Cl. Constr. Br. at 20.) When discussing the first variant, the specification explains that the “well-known function preferably applied is the hashpjw function.” (’170 Patent, 15:15–16, JA0022.) But when discussing the second variant, the specification simply refers to “a general function.” (*Id.* at 15:51.)

These arguments do not persuade the court that the expression “based on a hash function . . .” must be limited to indirect mapping. (See Goodrich Decl. ¶ 45.) Claim 1 is not limited to any particular manner of organizing data based on a hash function. It merely requires a hash function “used to organize the data location information across the plurality of data location servers.” (’170 Patent, 21:4–6, JA0022.) As Dr. Goodrich has explained, a POSITA would understand this requirement to be satisfied by a hash function that “takes an identifier string as input and outputs the URL of one or more location servers rather than an index into a data structure that would contain such URL information”—in other words, it would be satisfied by direct mapping. (Goodrich Decl. ¶¶ 42–43.) Nowhere does Claim 1 require that a hash function’s output be an index into a table, as in indirect mapping. A hash function’s output may, alternatively, identify a server directly, as in direct mapping. Kove’s construction preserves that breadth by allowing for both direct and indirect mapping.

¹³ In fact, of the dictionaries cited, only one appears to use the word “index.” (See Mitzenmacher Decl. ¶ 18(c) (“A hash function is applied to the item’s key and the resulting hash value is used as an index to select one of a number of ‘hash buckets’ in a hash table.” (quoting *Blackie’s Dictionary of Computer Science* 106–107 (2013))).)

The specification supports this construction. As discussed above, the specification discloses two variants of a redirect message. In the first variant, the appropriate location server is selected by “applying a well-known function to the identifier string and using the function result as an index into the NDTP server table.” (’170 Patent, 15:13–15, JA0022.) The parties agree that this variant represents what Kove calls indirect mapping. In the second variant, a client applies a “general function” to the identifier, and “the output of the function will be the NDTP server URL.” (*Id.* at 15:51–56.) That description does not use the word “hash,” but a POSITA would understand it to encompass direct mapping with a hash function. (Goodrich Decl. ¶ 43.)

The court declines AWS’s invitation to read the entire second variant out of the claim. As noted, it is “rarely, if ever,” appropriate to interpret a claim in such a way as to “exclude[] a preferred embodiment” from that claim. *MBO Lab’ys, Inc. v. Becton, Dickinson & Co.*, 474 F.3d 1323, 1333 (Fed. Cir. 2007). True, as AWS notes, unlike the first variant, the second variant does not use the word “hash function” or provide the source code for a hash function (AWS Reply at 11–12), but this does not defeat the conclusion that the second variant also encompasses a method of organizing information that is “based on a hash function” As the Federal Circuit has explained, “it is improper to read limitations from a preferred embodiment described in the specification . . . into the claims absent a clear indication in the intrinsic record that the patentee intended the claims to be so limited.” *Liebel–Flarsheim Co. v. Medrad, Inc.*, 358 F.3d 898, 913 (Fed. Cir. 2004). AWS points to a preferred embodiment that constitutes indirect mapping, but it has not established any clear indication that direct mapping is excluded from Claim 1.

Finally, the court rejects AWS’s argument that the claim would be invalid for lack of written description and enablement if the term “based on a hash function” were given Kove’s construction. Because a patent is presumed to be enabled and adequately described, a challenger bears the burden of proving invalidity by clear and convincing evidence. See *Cephalon, Inc. v. Watson Pharms., Inc.*, 707 F.3d 1330, 1337 (Fed. Cir. 2013); *WBIP, LLC v. Kohler Co.*, 829 F.3d 1317, 1338 (Fed. Cir. 2016). AWS has not made this showing. “[A] patent

need not teach, and preferably omits, what is well known in the art.” *Hybritech Inc. v. Monoclonal Antibodies, Inc.*, 802 F.2d 1367, 1384 (Fed. Cir. 1986). The enablement requirement thus focuses primarily on the novel aspects of an invention. See *Genentech, Inc. v. Novo Nordisk A/S*, 108 F.3d 1361, 1366 (Fed. Cir. 1997). AWS emphasizes that there are many possible ways to organize information based on a hash function, but the use of a hash function is not the novel aspect of the invention, so that is not a basis for concern. After all, as Dr. Goodrich explained, a POSITA would be aware of various methods of organizing information “based on a hash function.” (Goodrich Decl. ¶¶ 36–37.) The patentees thus “rel[ie]d on information that is ‘well-known in the art’ for purposes of meeting the written description requirement.” *Bos. Sci. Corp. v. Johnson & Johnson*, 647 F.3d 1353, 1366 (Fed. Cir. 2011). They were “not required to describe in the specification every conceivable and possible future embodiment of [their] invention.” *Innogenetics, N.V. v. Abbott Lab’ys*, 512 F.3d 1363, 1370 (Fed. Cir. 2008); cf. *Pfizer Inc. v. Teva Pharms. USA, Inc.*, 555 F. App’x 961, 967 (Fed. Cir. 2014) (“In view of the finding that enantiomer separation methods are well-known and routine to a person of ordinary skill, we agree with the district court that the inventors were not required to provide a detailed recipe for preparing every conceivable permutation of the compound they invented to be entitled to a claim covering that compound.”).

E. “Hash table”

Claim Term	Plaintiff’s Proposed Construction	Defendant’s Proposed Construction
“hash table” ’978 Patent: Claim 6 ’170 Patent: Claims 2, 12 ’640 Patent: Claim 24	“a data structure that stores values in a table, where values are stored and retrieved based on the output of a hash function ”	“a data structure in which a hash function is applied to an input and the function result is used as an index to a table in which the desired output can be found ”

After the claim construction hearing, the parties informed the court that they have agreed upon the following construction of “hash table”: “a data structure that stores values in a table,

where values are stored and retrieved by applying a hash function to an input and using the function result as an index into the table.” (Agreed Upon Constructions, Ex. B to Cl. Constr. Chart.) The court therefore adopts the agreed upon construction.

F. “Data pertaining to the entity”

Claim Term	Plaintiff’s Proposed Construction	Defendant’s Proposed Construction
“data pertaining to the entity” '978 Patent: Claims 1, 31 '170 Patent: Claim 6	“data pertaining to a person or thing (<i>real, digital, or abstract</i>)”	“data pertaining to a person or thing <i>distinct from that data</i> ”

The three asserted claims containing this term all recite an “identifier” that identifies an “entity” and a location of data “pertaining to the entity.” (See '978 Patent, 25:27–33, 29:48–54, JA0084-85 (Claims 1 and 31); '170 Patent, 21:34–41, JA0025 (Claim 6).) The parties generally agree that the term “data pertaining to the entity” means “data pertaining to a person or thing,” but they differ as to qualifiers. Kove would clarify that the person or thing to which data pertains (i.e., the entity) may be “real, digital, or abstract,” while AWS would stipulate that the person or thing (i.e., the entity) is “distinct from” the data that pertains to it. Recall that the parties have agreed upon the following construction of “location”: “an encoding that is a member of a set of associations with an identifier in a location server, and that specifies where *data pertaining to the entity* identified by the identifier is stored.” (Agreed Upon Constructions, Ex. B to Cl. Constr. Chart (emphasis added).) The construction of “data pertaining to the entity” therefore implicates all of the asserted claims containing the word “location.”

Kove points to several pieces of evidence, both intrinsic and extrinsic, to prove that the person or thing to which data pertains (i.e., the entity) may be “real, digital, or abstract.” Not all this evidence is compelling; for example, Kove points out that the specifications of the '170 and '978 Patents explain that the inventions improve upon prior technology by retrieving data files in multiple forms, such as text, image, sound, or video files. (See '170 Patent, 1:47–57, JA0001;

'978 Patent, 1:48–58, JA0072.) This sheds little light on the appropriate construction, as it simply recognizes that *data* may comprise text, image, sound, or video files. That much is not contested. Kove is on firmer ground in citing to the provisional applications for the patents in suit, which list examples of entities including “Song,” “Customer,” “Software object,” or “Messaging group.” (See Provisional Application No. 60/277,408, Ex. F to Kove Cl. Constr. Resp. [247-7] at 6; Provisional Application No. 60/209,070, Ex. G to Kove Cl. Constr. Resp. [247-8] at 12.)¹⁴ One of the inventors even described “Content Delivery” as an application of the technology, in which the protocol assigns an identifier to the song “Hotel California” and copies of the song are stored in multiple locations. (Ex. F to Kove Cl. Constr. Resp. at 7.) Kove also points to a contemporaneous technical dictionary that defines “entity” as “a distinguishable object, either real or abstract, about which data are recorded.” (Dictionary of IEEE Standards Terms, Ex. E to Kove Cl. Constr. Resp. [247-6] at 3.)¹⁵

AWS does not challenge Kove’s understanding that an entity may be “real, digital, or abstract” other than contending that Kove’s clarification is “unnecessary” and “wouldn’t be helpful to a jury.” (Cl. Constr. Hr’g 103:20–24.) AWS focuses instead on its own proposed requirement that an entity be “distinct from” the data that pertains to it. According to AWS, the plain meaning of “data pertaining to the entity” implies a distinction between “data” and an “entity.” (See AWS Cl. Constr. Br. at 23–24.) AWS contends that its construction is consistent with all three patents’ specifications, which provide: “In networked environments where there are a large number of data repositories and any particular entity does not store data in all the repositories, a mechanism is needed that would permit queries to be directed only at data repositories with relevant

¹⁴ Kove has not presented evidence that these documents are, in fact, the provisional applications for the patents in suit, but AWS has not contested this assertion.

¹⁵ As AWS notes, the definition continues: “for example, a person such as a CUSTOMER, or a concept, such as SALES REVENUE, about which data is stored in a data structure.” (Ex. E to Kove Cl. Constr. Resp. at 3; see AWS Cl. Constr. Reply at 13–14.)

information.” (’170 Patent, 2:1–5, JA0015; ’640 Patent, 1:66–2:3, JA0040; see also ’978 Patent, 2:1–5, JA0072 (similar).) In the view of AWS, this passage implies that an “entity” *stores* “data,” not that an “entity” could also *be* “data.”

Kove responds that AWS’s construction “artificially narrows” the term “entity.” (Kove Cl. Constr. at 24.) In fact, Kove says, “an entity and data may, in practical terms, be the same.” (*Id.*) For example, if the entity in question is a film that is available on a video streaming platform, the “location of the data *is* the location of the entity, making entity and data effectively the same.” (*Id.*) Kove also cites stray uses of the terms “data entity” and “data object” in the ’170 Patent and a Patent Examiner’s statement, respectively. (Kove Cl. Constr. Resp. at 25.) It contends that these usages “recite ‘entity’ and ‘data’ as the same thing.” (*Id.*)

AWS replies that, by allowing “data” and an “entity” to be the same thing, Kove’s construction would render the phrase “pertaining to the entity” superfluous. (AWS Reply at 13.) See also *Wasica Fin. GmbH v. Continental Auto. Sys., Inc.*, 853 F.3d 1272, 1288 n.10 (Fed. Cir. 2017) (“It is highly disfavored to construe terms in a way that renders them void, meaningless, or superfluous.”). Relatedly, AWS urges that Kove’s construction would rewrite the claims such that an “identifier” would identify data rather than identifying an entity.¹⁶

The court adopts and combines both parties’ proposed clarifications of the term. As mentioned above, AWS has provided no substantive reason to reject Kove’s proposed clarification that an entity may be “real, digital, or abstract.” The court disagrees with AWS that this language would not be useful; specifying that an entity may be “real, digital, or abstract” helps clarify the breadth of the term. The court agrees, however, with AWS’s assertion that data is

¹⁶ In support of this point, AWS cites *Source Vagabond Sys. Ltd. v. Hydrapak, Inc.*, 753 F.3d 1291, 1301 (Fed. Cir. 2014). Although AWS does not provide an explanatory parenthetical, the court presumes that AWS is referring to the following passage: “[A] court may not rewrite a claim even if giving a disputed claim its plain meaning would lead to a ‘nonsensical result.’” *Id.* (citation omitted). But Kove does not appear to be advocating for rewriting the claim term to avoid absurd results or to preserve the validity of the claims. See *id.* (collecting cases).

distinct from the entity to which it pertains, and it therefore rejects Kove's conflation of these two concepts.

By way of illustration, imagine that a video streaming platform manages its data as taught by the patents in suit. Suppose further that the "entity" in question is an award-winning film that was first released on this platform ("*Award-Winning Film*"). The "data," in this example, would mainly be the millions of bytes of audiovisual information that a subscriber would temporarily download to view *Award-Winning Film*. The data might also include closed captions, subtitles, or alternative audio tracks (e.g., from other languages). The "identifier" could simply be the words of the film's title, perhaps with spaces and punctuation removed: "awardwinningfilm." A subscriber seeking to view *Award-Winning Film* would use a client to query a location server for the location of data associated with the identifier "awardwinningfilm." Whether the subscriber immediately received a location response message or instead first received a redirect message, the subscriber would eventually be directed to a repository (or repositories) containing data associated with the identifier "awardwinningfilm." That data, in turn, would allow the subscriber to view *Award-Winning Film*.

In this example, there is a sense in which, as Kove suggests, the "location of the data *is* the location of the entity." (Kove Cl. Constr. Resp. at 24.) Unlike a "real" entity such as a medical patient, a digital entity like *Award-Winning Film* cannot meaningfully be said to have any "location" other than the location of the data that constitutes it. But that does not mean that the entity (*Award-Winning Film*) and the data (the bytes of information) are "effectively the same." (*Id.*) The millions of bytes of information—the audio, video, text, and so on—are indeed just data. Although these pieces of data pertain to, and in the aggregate constitute, the entity *Award-Winning Film*, none of the pieces of data are themselves the "entity" that is identified by the identifier "awardwinningfilm." To help illustrate this point further, we can imagine two different ways in which the relevant data might be organized in the video streaming platform's distributed-data-storage system. First, the data that constitutes *Award-Winning Film* could be stored in unique

pieces across several data repositories. Even with the data scattered in this way, there would still be just one entity to which the data pertains. Conversely, the distributed-data-storage system could contain several duplicate sets of the data—each set independently constituting the entirety of *Award-Winning Film*. And here, too, there would be only one entity to which the data pertains. Thus, even in the case of a digital entity, the data and the entity are not “the same thing,” as Kove suggests.

To reflect this understanding, the court construes the term “data pertaining to the entity” as “data pertaining to a person or thing (real, digital, or abstract) distinct from that data.”

CONCLUSION

The claim terms in the '640 Patent, '170 Patent, and '978 Patent are construed as follows:

Claim Term	Construction
“location information”	'978: “one or more identifiers and their associated locations” '170/'640: “information pertaining to one or more locations of data and/or the identities of one or more location servers”
“location server”	“a network-attached component that maintains a set of identifier/location mappings that are modified or returned in response to location request messages from clients”
“client”	“a network-attached component (which may be software or hardware) that initiates update or lookup of identifier/location mappings from a location server with location request messages”
“based on a hash function used to organize the data location information across the plurality of data location servers . . . based on the hash function applied to the identifier string”	“the portion of the data location information included in a corresponding one of the data location servers is based on a hash function that maps identifier strings to one or more of the data location servers, and each one of the data location servers is configured to determine the at least one of the plurality of data location servers based on the hash function applied to the identifier string”
“data pertaining to the entity”	“data pertaining to a person or thing (real, digital, or abstract) distinct from that data”

ENTER:

Dated: December 17, 2021

A handwritten signature in black ink, appearing to read "Rebecca R. Pallmeyer", written over a horizontal line.

REBECCA R. PALLMEYER
United States District Judge