

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

NETSKOPE, INC.,

Petitioner,

v.

FORTINET, INC.,

Patent Owner.

PTAB Case No. IPR2023-00458

Patent No. 9,280,678 B2

**DECLARATION OF WENKE LEE, PHD, IN SUPPORT OF PETITIONER
NETSKOPE, INC.'S REPLY TO PATENT OWNER'S RESPONSE TO
PETITION FOR *INTER PARTES* REVIEW OF U.S. PATENT NO. 9,280,678 B2**

I. Claim Construction

A. The “creating” and “configuring” steps are not part of the encrypting the content but are part of the operations to produce a searchable encrypted content

1. Claim 1, recites the following operations:

encrypting, by the gateway device, using cryptographic key information defined by the policy, content of the file to produce a searchable encrypted file by:

- (a) dividing the file into a plurality of chunks;
- (b) creating namespaces for one or more of the plurality of chunks; and
- (c) configuring the namespaces of the one or more chunks such that content of the file is encrypted in a manner that makes it searchable.

’678 Patent, 19:62-20:3.¹

2. Claim 1 specifies the process of transforming the content of the file into a searchable encrypted file by encrypting the content to produce a searchable encrypted file by performing operations (a), (b) and (c). ’678 Patent, 19:62-20:3. This wording, particularly the use of the phrase “to produce a searchable encrypted

¹ Identifiers (a), (b) and (c) have been added.

file by,” is intended to outline how the encrypted searchable file is produced, rather than detailing the individual steps contained within the “encrypting” operation itself. This arrangement positioning of these steps in the claim, positioned closer to the resulting action rather than the encryption process, suggests that these steps are for generating the encrypted searchable file, and not as subcomponents of the encryption. In other words, by placing these steps close to the outcome—a searchable encrypted file—highlights their role as instrumental in the creation process rather than being elements of the encryption.

3. Patent Owner misinterprets the “encrypting” file content as it necessitates the creation and configuration of namespaces. This misinterpretation contradicts by the example provided in Appendix A of the ’678 patent that describes a “cloud provider-agnostic searchable cloud data storage solution” (Dragonfruit) (Abstract), which includes four major steps: Setup, Encryption, Token, and Query. Ex. 1001, Appendix A, Section IV. In direct contradiction to what Patent Owner has asserted, the “Encryption” is solely dedicated to the “encryption of each individual record,” resulting in the generation of a cipher text vector. In other words, this Encryption step, as described, is a straightforward process that does not involve any of creation or configuration of namespaces. It is focused exclusively on transforming plain text into encrypted data to ensure

confidentiality. In fact, it does not mention namespaces at all as part of Encryption. Ex. 1001, Appendix A, Section IV.

4. The Patent Owner and their expert witness, Dr. Black, have presented conflicting statements regarding the process of encrypting file content and the creation of a namespace. For example, the Patent Owner Response (POR) and Dr. Black (Ex, 2003) both identically state: “These *policies therefore create a namespace* that defines the naming convention for the file chunks.” POR, 20; Ex. 2003, ¶ 71. This statement indicates that the policy creates the namespace, which means namespace creation is done in advance of the encryption (or other operations). This statement by Dr. Black and Patent Owner contradicts their later explanation that Cidon does not describe creating a namespace because: “... it is quite clear that the set of available names—SHA2 signatures—is simply defined by the Cidon reference in advance. *Id.* Thus, the system does not create the namespace during the encryption process because it was created in advance.” POR, 22.

5. Dr. Black further states that “[h]ence, *naming the folder with the encrypted filename* and storing the file chunks in that folder creates a namespace scope for the file chunks.” POR, 20; Ex. 2003, ¶ 71). However, using an encrypted *filename* to name a folder is not part of encrypting the contents of the file. In particular, the encryption process applies algorithms to data content,

altering its appearance to make it unreadable without the appropriate decryption key. This process is concerned solely with the data's content and not with how the data is labeled or organized within a storage system. On the other hand, the use of encrypted filenames for folders are organizational tools, which help in identifying, organizing, and retrieving data within a storage system. While employing encrypted filenames as identifiers may add a layer of obfuscation, it does not alter the encrypted state of the file's contents. Patent Owner's expert, Dr. Black, concurs with this view, acknowledging that encrypting the contents of a file does not directly influence the filename. Black Tr. 54:15-17 ("I agree with you if you're saying that the contents of the file, if you encrypt that, that's not influencing directly the filename.").

6. It is my opinion that Patent Owner's position that imports additional steps to "encrypting the file content" is not supported by the '678 patent. A POSITA, reading the 678 patent, would not have known how "encrypting the file *content*" would have resulted in creation and configuration of namespaces because they are not described in the '678 patent. To the contrary, and as I explained above, the description of "Encryption" step in Appendix A of the '678 patent clearly explains only the encryption of the content.

B. The “policy” recited in the claims is not “a single policy”

7. The ’678 claims recite a “policy” and not a “single policy.” Moreover, Patent Owner’s interpretation as a “single policy” is at odds with the ’678 patent. The ’678 patent explicitly states “sub-policies” (noted in the plural): “User-assigned *policy can further include key management sub-policies* across cloud providers and local devices.” ’678 Patent, 9:40-46 (emphasis added). Based on at least this disclosure, a POSITA would have interpreted the claimed policy as a policy that can include sub-policies, and is therefore a conglomeration of individual policies, each of addressing different functions, such as the “file storage policy” (*id.* at 2:18-22) that defines, e.g., access rights and the policy that defines “cryptographic key information” (*id.* at 10:1-6).

II. The Cidon Grounds

A. Cidon’s creating identifiers constituted creating namespaces

8. The claims recite “creating namespaces for one or more of the plurality of chunks.” ’678 Patent, 19:66-67, 20:66-67. The ’678 Patent mentions “namespaces” only four times, but does not describe how to create (or configure) a namespace, other than creating a name for it. *See* ’678 Patent, 6:44-47, 9:19-26, 10:11-16, FIG. 3. In fact, claim 1 recites “creating namespaces for one or more of the plurality of chunks,” which means it suffices to create *a namespace for one file chunk*, which clearly does not require creating a range of names or values. In other words, to create a namespace for a file chunk, there is no requirement for

establishing a range of names or values. Creating a namespace can include simple cases, such as naming a file (or chunk) with a specific identifier, or naming each of multiple chunks with corresponding identifiers.

9. Cidon described creating a unique identifier for each of file segments (chunks) based on the content of a file segment, using SHA2 signature computation. Cidon, ¶¶305, 383, FIG. 13.




FIG. 13 (Annotated)

Cidon, FIG. 13.

10. SHA2 is a cryptographic secure hash generation algorithm that provides a fixed-length output, for example, a 256-bit output. Its fixed-length

output for creating hash values that are unpredictable and unique, making it virtually impossible to generate the same hash value from two different inputs. This fixed-length output assigned to identifiers has a format similar to that of other example namespaces in the field, such as a number in the range of 0 to $2^{256}-1$. Ex. 1023, 344. Thus, Cidon's use of SHA2 to generate identifiers for file chunks and assigning these identifiers to the file chunks constituted creating namespaces for those file chunks because these identifiers are a "set of names available for naming" those file chunks. Ex. 1014 (Dictionary of Computer and Internet Terms).

11. Even if it were required to "create" namespaces as part of the "encrypting" process, Cidon's creating and assigning identifiers were part of the encryption process. This is at least because Cidon's selection of SHA-2 and then running the SHA-2 algorithm to produce identifiers constituted creating namespaces (as discussed above) and the identifiers became part of encrypted file 1455 along with encrypted file chunks. Cidon, FIG. 13. In this regard, Cidon's operations can be compared and analogized to the operations described in the '678 Patent, as shown in the annotated figures below. *Compare* '678 Patent, FIG. 3A *with* Cidon, FIG. 13.




Fig. 3A

FIG. 13 (Annotated)

'678 Patent, FIG. 3A

Cidon, FIG. 13.

B. Shikfa's mapping of keywords to identifiers in an index constituted "configuring the namespaces"

12. A POSITA would have understood that configuration encompasses any of arrangement, utilization, or integration in a particular way to achieve a specific functionality. Ex. 1026 (definition of "configure" as "to set up for operation esp. in a particular way"). In other words, "configuring" the namespaces does not require modifying them, but can be arranging them or setting them up for operation, especially in a particular way. *Id.* Shikfa's maps the identifiers (namespaces) to keywords in indexes in a particular way to reference encrypted

file chunks. Therefore, in the combination, Shikfa described configuring namespaces to make the encrypted content searchable.

C. Cidon’s policy was assigned to a “user”

13. Patent Owner’s argument that Cidon’s encryption policies were solely defined on documents but not assigned to users (POR, 32) overlooks the flexibility in Cidon’s policy framework. Cidon described an example where an encryption policy applied to a file’s ancestry tree, but this example does not preclude assigning the policy at the user level. Cidon, ¶¶207. In fact, Cidon’s policy specified access and sharing permissions, copy control, and encryption suggests that the policy was user-specific. *Id.*, ¶¶207-213. Specifically, Cidon in paragraphs 208 through 212 explains what such policies may comprise—placement in directories, access permissions, copy control, and encryption—and indicates that the policies can be tailored to the needs and roles of specific users. The nature of these policy components suggests a user-centric design. For instance, access and sharing permissions (paragraph 210) are elements of user-specific policies, as they directly govern who can interact with the document and how. Most importantly, paragraph 217 states that users (not administrators) can be alerted to policy violations and given the authority to approve the violation or modify the policy as necessary. This level of individual administration and interaction with the policy directly involves the user in the policy governance process. Thus, the granular control and

responsiveness of Cidon’s policy framework, as specifically outlined in paragraphs 207-213, teach policies being specific to individual users.

III. The Auradkar Grounds

A. Auradkar’s tags and Chiueh’s document IDs satisfy the “namespaces” limitation under Patent Owner’s interpretation of “namespace”

14. Patent Owner proposes a construction of term “namespace” as “the set of names available in a particular context.” Even under this construction, Auradkar’s tags and Chiueh’s document IDs satisfy the “namespaces” limitation. This is because they were unique within the system’s context (e.g., among encrypted XML segments or encrypted pages) for identifying individual segments (e.g., encrypted XML segment, or encrypted page). Their uniqueness ensures that each segment is accorded a specific tag or ID that differentiates it from others. This unique identification system constitutes namespaces’ functionality, which is to name things (e.g., file segments) and identify them within the system. Indeed, this functionality of Auradkar’s tags and Chiueh’s document IDs as namespaces is affirmed by their use in facilitating search operations among encrypted payloads or pages, a process relying on the functionality provided by namespaces—ensuring every encrypted payload or page can be identified uniquely by their identifiers. Pet., 94-97; Auradkar, ¶¶125, 150, FIGS. 12, 19; Chiueh, ¶¶20, 40.

15. Patent Owner's argument that a filename is merely a single item within a set and is not a namespace fails to consider the function of namespaces and how they are implemented within a system. POR, 44. While a filename *can* be considered a local name within a broader namespace, creating an identifier (based on a filename) and associating it with a particular file segment in the Auradkar-Chieuh combination created *a namespace for that file segment*. Namely, these identifiers are namespaces for differentiating between objects or entities that may have identical local filenames. Put differently, namespaces came into existence as these identifiers were being created and assigned to particular file segments, as these identifiers now became a unique set of names within the context of the encrypted file(s).

B. Chiueh's mapping of keywords to document identifiers in an index constituted "configuring the namespaces"

16. As explained above, a POSITA would have understood that configuration encompasses any of arrangement, utilization, or integration in a particular way to achieve a specific functionality. Ex. 1026 (definition of "configure" as "to set up for operation esp. in a particular way"). Thus, a POSITA would also have understood that Chiueh's use of document IDs with the index pointing to encrypted pages in a particular way effectively configured these IDs as part of a searchable encryption scheme to make the encrypted content searchable. Indeed, this aligns with the descriptions in Appendix A of the '678 patent, which

outlines the creation of a hierarchical series of indices utilizing the SHA-3 hashing algorithm. Exhibit 1001, Appendix A, Section V. This process described in Appendix A begins with the initial hashing of all cipher-texts, followed by the production of a secondary hash for each cipher-text, and this secondary hash serves as an identifier for a specific table that consolidates numerous cipher-texts. *Id.* These tables are located with the cloud service provider, and where the “cloud provider is queried for the table, which is then downloaded, loaded into memory, and then searched for the original cipher-text.” *Id.* In this method, the encrypted file becomes searchable through the use of hierarchical indices that are compiled in a table. *Id.* These indices are then employed to pinpoint the desired ciphertext, in a manner similar to Chiueh’s method of constructing an index that linked encrypted keywords to the respective document ID(s) of the encrypted documents, thereby enabling the search of the document content using those keywords. Pet., 102-104; Chiueh, ¶¶21-23, FIG. 2.

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; that these statements were made with knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under 18 U.S.C. § 1001; and further that such willful false statements may jeopardize the validity of

the application or any patent issued thereon. I declare under penalty of perjury under the laws of the United States of America that the foregoing is true and correct.

Executed on April 9, 2024 in Atlanta, GA, USA.

A handwritten signature in black ink, appearing to read 'Wenke Lee', written over a horizontal line.

Wenke Lee, Ph.D.