

**IN THE UNITED STATES DISTRICT COURT
FOR THE WESTERN DISTRICT OF TEXAS
WACO DIVISION**

VIRTUAL CREATIVE ARTISTS, LLC,

Plaintiff,

v.

META PLATFORMS, INC.,

Defendant.

C.A. No. 6:22-cv-265

JURY TRIAL DEMANDED

PATENT CASE

ORIGINAL COMPLAINT FOR PATENT INFRINGEMENT

Plaintiff Virtual Creative Artist, LLC files this Original Complaint for Patent Infringement against Meta Platforms, Inc. and would respectfully show the Court as follows:

I. THE PARTIES

1. Plaintiff Virtual Creative Artists, LLC (“VCA” or “Plaintiff”) is a Delaware limited liability company, having a principal place of business located at 268 Bell Canyon Road, Bell Canyon, CA 91307.

2. On information and belief, Defendant Meta Platforms, Inc. (“Defendant”) is a corporation organized and existing under the laws of Delaware. Defendant has a place of business at 300 W 6th Street, Austin, TX 78701. Defendant has a registered agent at Corporation Service Company, 211 E. 7th Street Suite 620, Austin, TX 78701.

II. JURISDICTION AND VENUE

3. This action arises under the patent laws of the United States, Title 35 of the United States Code. This Court has subject matter jurisdiction of such action under 28 U.S.C. §§ 1331 and 1338(a).

4. On information and belief, Defendant is subject to this Court’s specific and general personal jurisdiction, pursuant to due process and the Texas Long-Arm Statute, due at least to its

business in this forum, including at least a portion of the infringements alleged herein at 300 W 6th Street, Austin, TX 78701. Defendant also signed a lease for 589,000 square feet across 33 floors in a soon to be completed skyscraper in downtown Austin, Texas. Defendant employs over 2,000 people in Austin, Texas, with plans to hire 400 more people.

5. Without limitation, on information and belief, within this state, Defendant has used the patented inventions thereby committing, and continuing to commit, acts of patent infringement alleged herein. In addition, on information and belief, Defendant has derived revenues from its infringing acts occurring within Texas. Further, on information and belief, Defendant is subject to the Court's general jurisdiction, including from regularly doing or soliciting business, engaging in other persistent courses of conduct, and deriving substantial revenue from goods and services provided to persons or entities in Texas. Further, on information and belief, Defendant is subject to the Court's personal jurisdiction at least due to its sale of products and/or services within Texas. Defendant has committed such purposeful acts and/or transactions in Texas such that it reasonably should know and expect that it could be haled into this Court as a consequence of such activity.

6. Venue is proper in this district under 28 U.S.C. § 1400(b). On information and belief, Defendant has a place of business in this district at 300 W 6th Street, Austin, TX 78701. On information and belief, from and within this District Defendant has committed at least a portion of the infringements at issue in this case.

7. For these reasons, personal jurisdiction exists and venue is proper in this Court under 28 U.S.C. § 1400(b).

III. COUNT I
(PATENT INFRINGEMENT OF UNITED STATES PATENT NO. 9,501,480)

8. Plaintiff incorporates the above paragraphs herein by reference.

9. On November 22, 2016, United States Patent No. 9,501,480 (“the ‘480 Patent”) was duly and legally issued by the United States Patent and Trademark Office. The ‘480 Patent is titled “Revenue-Generating Electronic Multi-Media Exchange and Process of Operating Same.” A true and correct copy of the ‘480 Patent is attached hereto as Exhibit A and incorporated herein by reference.

10. VCA is the assignee of all right, title, and interest in the ‘480 Patent, including all rights to enforce and prosecute actions for infringement and to collect damages for all relevant times against infringers of the ‘480 Patent. Accordingly, VCA possesses the exclusive right and standing to prosecute the present action for infringement of the ‘480 Patent by Defendant.

11. The invention relates to the field of creating and distributing media content, in particular, creating media content based upon submissions received on an electronic media exchange. At the time of the original invention in 1998, there was an Internet-centric problem that required a technical solution—how to develop a computer system that would allow remote contributors of electronic content to share and collaborate their content to develop new media content. The claimed invention, which predates modern crowdsourcing solutions, offers a unique, unconventional, and specially configured combination of “subsystems” in which to address the Internet-centric problem.

12. As set forth in the claims, the claimed invention has a collection of unconventional and particularly configured subsystems, including:

- “an electronic media submissions server subsystem,”
- “an electronic multimedia creator server subsystem,”
- “an electronic release subsystem,”
- “an electronic voting subsystem,” and
- their corresponding specialized databases.

13. Each of these subsystems are configured in a very specific, non-generic, unconventional, and non-routine manner to offer the novel and non-obvious claimed invention. For example, claim 1 requires an “electronic media submissions database,” which is a subsystem that receives media submissions from Internet users. This is not a generic database but rather a scalable database that must be able to receive, store, and manage multiple petabytes of multimedia data received from users all over the world. This is one of the many specialized databases required in the claim. In fact, the specification discloses the use of a sophisticated database management system known in the art at the time that was capable of handling data at this level. (Ex. A at 7:53-56). This type of database management system cannot operate on a generic computing system but rather requires specialized hardware and software.

14. As another example, the claim requires a specifically configured “electronic media submission server subsystem.” This subsystem is defined as specifically having:

- “one or more data processing apparatus,”
- “an electronic media submission database stored on a non-transitory medium,”
and
- “a submissions electronic interface.”

The “submissions electronic interface” is further specifically “configured” [1] “to receive electronic media submissions from a plurality of submitters over a public network, and [2] store the electronic media submissions in the electronic media submission database.” Further, “the electronic media submissions database” in this subsystem is further required to “store[] [1] data identifying the submitter and [2] data indicating content for each electronic media submission.” Collectively, the level of detail included in this very particular, well-defined, and unconventional subsystem makes clear that the claim includes substantially more than an abstract idea or merely performing an abstract idea on a computer.

15. Similarly, the claim also requires a separate specifically configured “an electronic multimedia creator server subsystem.” The claim specifically defines how this second subsystem interacts with other components including being “operatively coupled to the electronic media submissions server subsystem.” The claim also specifically defines this subsystem as “having”:

- “one or more data processing apparatus” and
- “an electronic creator multimedia database stored on a non-transitory medium.”

16. This subsystem is also specifically “configured [1] to select and [2] retrieve a plurality of electronic media submissions from the electronic media submissions database using an electronic content filter located on the electronic multimedia creator server.” The “filter” also includes a very specific algorithm of “being based at least in part on at least one of the one or more user attributes to develop multimedia content to be electronically available for viewing on user devices.” Even more detail is provided by requiring that “the identification of the submitter [be] maintained with each selected and retrieved submission within the multimedia content.” Here again, collectively, the level of detail included in this very particular, well-defined, and unconventional subsystem makes clear that the claim includes substantially more than an alleged abstract idea or merely performing an alleged abstract idea on a computer.

17. The claim also includes “an electronic release subsystem,” which is well defined and not conventional or routine. The claim defines how this subsystem is “operatively coupled to the electronic multimedia creator server subsystem.” The claim also defines the components of this subsystem as having “one or more data processing apparatus” and being particularly “configured to make the multimedia content electronically available for viewing on one of more user devices.” These details, collectively, also make this very particular, well-defined, and

unconventional subsystem substantially more than an abstract idea or performing an abstract idea on a computer.

18. The claim also requires “an electronic voting subsystem,” which is well-defined, specific, and unconventional. This claimed subsystem has “one or more data processing apparatus” and is specifically “configured to enable a user to electronic vote for or electronically rate an electronically available multimedia content or an electronic media submission within a respective electronically available multimedia content.”

19. Claim 1 is a specific and discrete implementation. For example, the claim requires an “electronic content filter” located at the server, remote from end users, and customizable based on user attributes. As another example, the “electronic voting subsystem” at the time of the invention was novel and inventive and added sufficient inventive contributions to avoid a risk of preempting creating and distributing media content. It is possible to create and distribute media content without ever having to include a “voting” subsystem on what components should be included in such media content. The detailed configuration “to enable a user to vote for or electronically rate an electronically available multimedia content or an electronic media submission within a respective electronically available multimedia content” has the level of particularity that avoids any risk of preemption.

20. Furthermore, the very particular and specifically configured “electronic media creator subsystem” not only provides a detailed and unique physical structure and interrelationship with other claimed components, but also includes a very specific configuration that is not conventional or routine. The claim make clear the interrelationship of the “electronic multimedia creator server subsystem” with respect to “the electronic media submission server subsystem” which must be “operatively coupled” thereto. The claim also provides detail on how the

“electronic media creator subsystem” is “configured” “to select and retrieve a plurality of electronic media submissions from the electronic media submission database using an electronic filter.” The claim also provides detail on how the “electronic filter” is “based at least in part on at least one of the one or more user attributes” and specifies that “the identification of the submitter is maintained with each selected and retrieved submission within the multimedia content.”

21. **Direct Infringement.** Upon information and belief, Defendant has been directly infringing claim 1 of the ‘480 Patent in Texas, and elsewhere in the United States, by employing a computer-based system using <https://www.facebook.com/> (“Accused Instrumentality”) (e.g., <https://www.facebook.com/>).

22. The Accused Instrumentality is a computer-based system comprising an electronic media submissions server subsystem having one or more data processing apparatus and an electronic media submissions database stored on a non-transitory medium and a submissions electronic interface configured to receive electronic media submissions from a plurality of submitters over a public network and store said electronic media submissions in said electronic media submissions database, wherein the electronic media submissions database further stores data identifying the submitter and data indicating content for each electronic media submission.

an electronic media submissions server subsystem having one or more data processing apparatus and an electronic media submissions database stored on a non-transitory medium and a submissions electronic interface configured to receive electronic media submissions from a plurality of submitters over a public network and store said electronic media submissions in said electronic media submissions database, wherein the electronic media submissions database further stores data identifying the submitter and data indicating content for each electronic media submission;

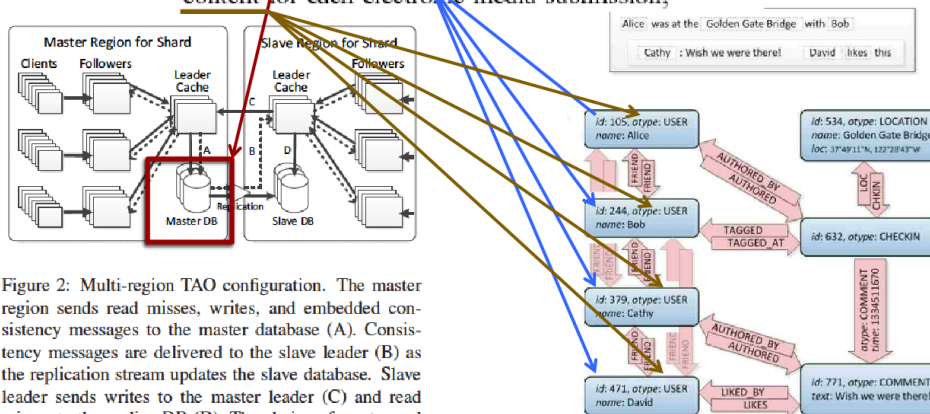


Figure 2: Multi-region TAO configuration. The master region sends read misses, writes, and embedded consistency messages to the master database (A). Consistency messages are delivered to the slave leader (B) as the replication stream updates the slave database. Slave leader sends writes to the master leader (C) and read misses to the replica DB (D). The choice of master and slave is made separately for each shard.

(E.g., *Mirror Worlds v. Facebook*, Appeal No. 18-2276, ECF 57-2 (“Challenges to Adopting Stronger Consistency at Scale”)).¹ In a previous case, Facebook’s expert, Nathan Bronson, described how the Accused Instrumentality comprises an electronic media submissions database, capable of storing data identifying the submitter and data indicating content for each electronic media submission. Bronson stated that,

The “objects” in the diagram above are shown by the seven blue boxes, and the “associations” are the pink arrows between those boxes. Each of the objects has the unique identifier (“id”) that I mentioned earlier, and an object type (“otype”) indicating what kind of object it is. The four objects shown on the left are USER objects for four Facebook users (Alice, Bob, Cathy, and David). The three objects on the right represent, respectively, the LOCATION object for the Golden Gate Bridge, the CHECKIN object that Alice created when she “checkedin” at that location, and a COMMENT object showing that Cathy wrote a textual comment reacting to Alice’s check-in. All of the actions taken by Alice, Cathy, and David are reflected in the diagram as a series of objects and the associations between them. The pink arrows reflect the type of association between the objects; for example, the bottom pair of arrows shows that David “LIKES”

¹ In this Complaint, claim language, underlining, colored arrows, and colored boxes are added unless otherwise noted.

(*Mirror Worlds v. Facebook*, Appeal No. 18-2276, ECF 57-2).




Figure 2: Multi-region TAO configuration. The master region sends read misses, writes, and embedded consistency messages to the master database (A). Consistency messages are delivered to the slave leader (B) as the replication stream updates the slave database. Slave leader sends writes to the master leader (C) and read misses to the replica DB (D). The choice of master and slave is made separately for each shard.

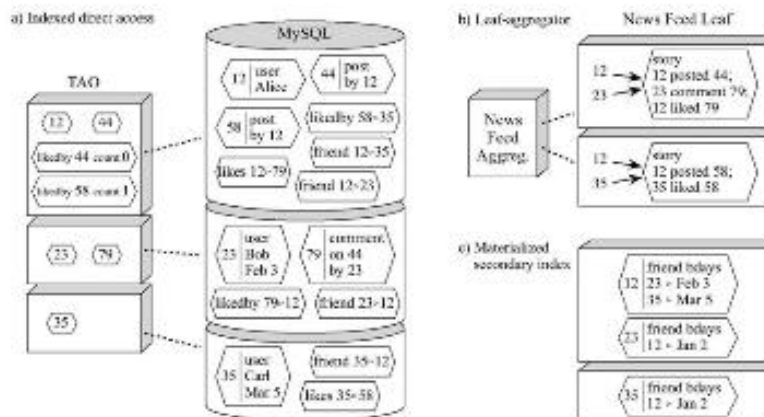


Figure 1: A hypothetical data layout for two posts by Alice, with interactions by friends Bob and Carl. MySQL and TAO (a) shard nodes by key and edges by their source. In News Feed (b) information is aggregated and indexed by all actors. Materialized secondary indexes (c) answer queries that would touch many TAO shards.

(E.g., *Mirror Worlds v. Facebook*, Appeal No. 18-2276, ECF 57-2).

2.1 Integrating Across Stateful Services

Facebook’s back-end infrastructure is comprised of many different services, each using a data placement and communication structure optimized for its workload. Figure 1 shows how data may be duplicated with different sharding strategies in different services.

The TAO caching layer [10], for example, uses multi-tenancy and hierarchical caching to tolerate very high query rates for simple graph queries. Another system materializes secondary indexes to support specific queries that would have a high fanout or require multiple rounds in TAO. Unicorn [14] uses document sharding and a leaf-aggregator model to handle the dynamic queries from Facebook’s Search. The News Feed back-end denormalizes all of the information associated with a candidate story so that ranking can be performed locally. Facebook’s deployment of memcache [21, 35] uses the flexible mcrouter proxy [30] to effect a wide variety of replication and locality strategies. Facebook’s Messenger service uses a data store optimized for efficient access to recent messages [20], but often embeds links to TAO nodes. In total, there are hundreds of services running in production at Facebook, many of which maintain a copy of portions of the social graph.

Facebook’s architecture of cooperating heterogeneous services is different from the monolithic service that most research designs assume. This difference alone is not fundamental—Facebook can be externally viewed as a single service even though internally it is comprised

(E.g., *Mirror Worlds v. Facebook*, Appeal No. 18-2276, ECF 57-2).

23. The Accused Instrumentality contains a user database comprising one or more user attributes. As shown below, the Accused Instrumentality contains more than one user and their attributes

Facebook has more than a billion active users who record their relationships, share their interests, upload text, images, and video, and curate semantic information about their data [2]. The personalized experience of social applications comes from timely, efficient, and scalable access to this flood of data, the *social graph*. In this paper we introduce TAO, a read-optimized graph data store we have built to handle a demanding Facebook workload.

24. The Accused Instrumentality contains an electronic multimedia creator server subsystem operatively coupled to the electronic media submissions server subsystem, having one or more data processing apparatus and an electronic creator multimedia database stored on a non-transitory medium, configured to select and retrieve a plurality of electronic media submissions from the electronic media submissions database using an electronic content filter located on the electronic multimedia creator server, said filter being based at least in part on at least one of the one or more user attributes to develop multimedia content to be electronically available for viewing on user devices, wherein the identification of the submitter is maintained with each selected and retrieved submission within the multimedia content. As shown below, the Accused Instrumentality uses an electronic content filter to maintain the identification of the submitter with each selected and retrieved submission within the multimedia content.

an electronic multimedia creator server subsystem operatively coupled to the electronic media submissions server subsystem, having one or more data processing apparatus and an electronic creator multimedia database stored on a non-transitory medium, configured to select and retrieve a plurality of electronic media submissions from the electronic media submissions database using an electronic content filter located on the electronic multimedia creator server, said filter being based at least in part on at least one of the one or more user attributes to develop multimedia content to be electronically available for viewing on user devices, wherein the identification of the submitter is maintained with each selected and retrieved submission within the multimedia content;

2 Background

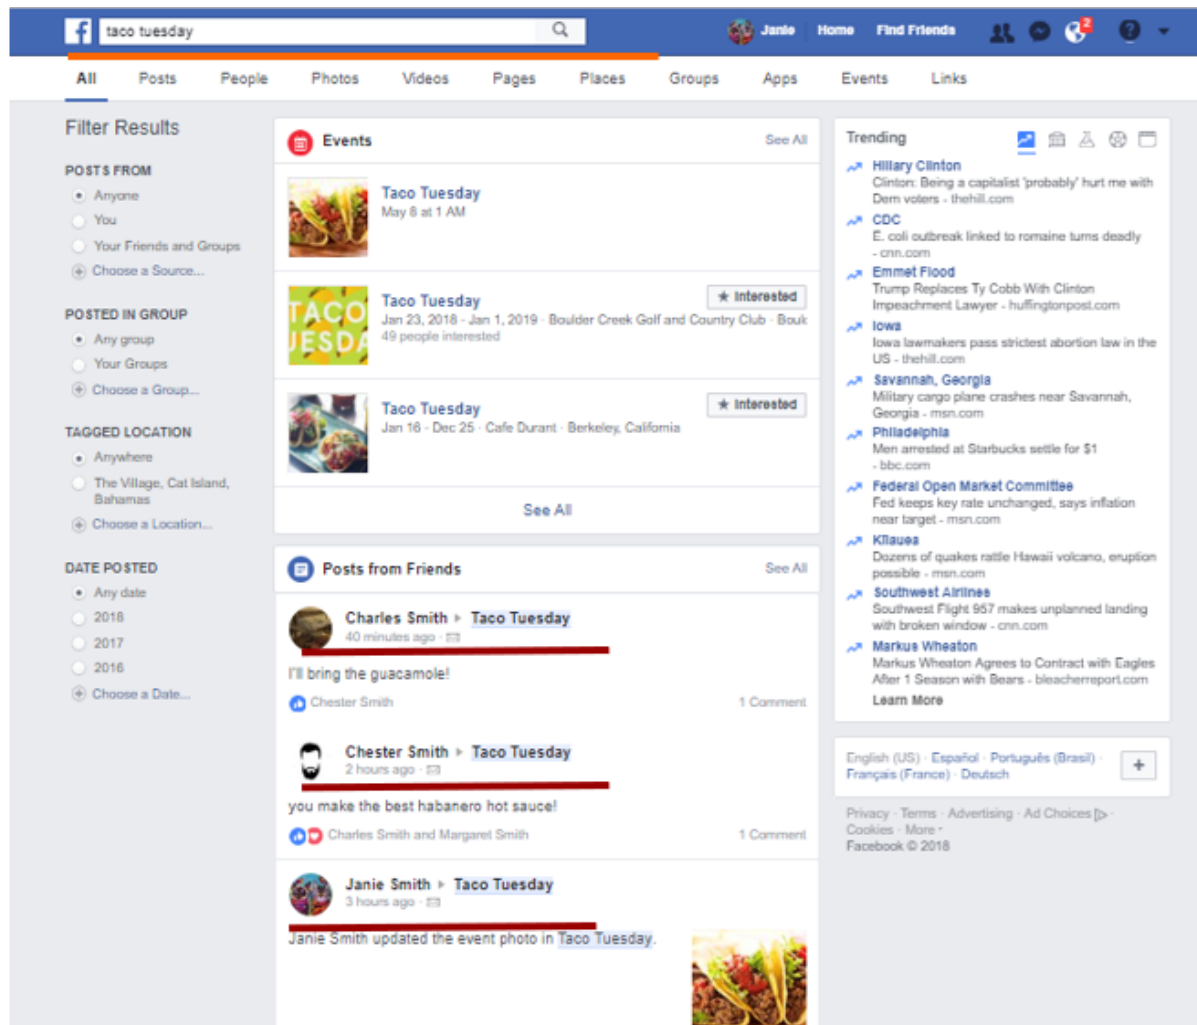
A single Facebook page may aggregate and filter hundreds of items from the social graph. We present each user with content tailored to them, and we filter every item with privacy checks that take into account the current viewer. This extreme customization makes it infeasible to perform most aggregation and filtering when content is created; instead we resolve data dependencies and check privacy each time the content is viewed. As much as possible we pull the social graph, rather than pushing it. This implementation strategy places extreme read demands on the graph data store; it must be efficient, highly available, and scale to high query rates.

3.1 Objects and Associations

TAO *objects* are typed nodes, and TAO *associations* are typed directed edges between objects. Objects are identified by a 64-bit integer (*id*) that is unique across all objects, regardless of object type (*otype*). Associations are identified by the source object (*id1*), association type (*atype*) and destination object (*id2*). At most one association of a given type can exist between any two objects. Both objects and associations may contain data as key→value pairs. A per-type schema lists the possible keys, the value type, and a default value. Each association has a 32-bit time field, which plays a central role in queries¹.

Object: (*id*) → (*otype*, (key → value)*)

Assoc.: (*id1*, *atype*, *id2*) → (time, (key → value)*)



Aggregator: This Aggregator also collects and ranks information.

Appx2446. As shown in the figure reproduced above, the Aggregator receives data from the front-end (" [REDACTED] ") that specifies parameters to retrieve information that a particular user may be interested in, including a [REDACTED] and [REDACTED] values. Appx2422; Appx1496; Appx1501-1502 (Koskinen

Decl. ¶¶ 73, 88-90) (discussing the data fields contained in this communication to the Aggregator).

(E.g., *Mirror Worlds v. Facebook*, Appeal No. 18-2776, Facebook Brief, ECF 34, at pp. 17-18).

25. As shown below, the Accused Instrumentality contains an electronic release subsystem operatively coupled to the electronic multimedia creator server subsystem, having one

or more data processing apparatus and configured to make the multimedia content electronically available for viewing on one or more user devices.

an electronic release subsystem operatively coupled to the electronic multimedia creator server subsystem, having one or more data processing apparatus and configured to make the multimedia content electronically available for viewing on one or more user devices; and




Figure 2: Multi-region TAO configuration. The master region sends read misses, writes, and embedded consistency messages to the master database (A). Consistency messages are delivered to the slave leader (B) as the replication stream updates the slave database. Slave leader sends writes to the master leader (C) and read misses to the replica DB (D). The choice of master and slave is made separately for each shard.

26. The Accused Instrumentality contains an electronic voting subsystem having one or more data processing apparatus and configured to enable a user to electronic vote for or electronically rate an electronically available multimedia content or an electronic media submission within a respective electronically available multimedia content.

7. The example web page in Figure 1 above describes a number of actions by users that Facebook deemed to be relevant to Janie Smith, and thus, presented them on her News Feed. For example, Janie Smith posted a photo of quilting with an accompanying textual comment “I could make so many quilts!” Margaret Smith and Charles Smith commented on Janie’s post, with Margaret Smith writing “It would be so much fun!!”, Charles Smith writing “You make the warmest quilts!”, and Margaret Smith replying “Rüdiger looks adorable in those quilts!”

8. Aside from presenting photos, News Feed supports a variety of other user activities. For example, other users of Facebook who are “Friends” of Janie Smith can write status updates about what is on their mind, upload videos, share links to other websites, or identify their current physical location (“check-in”). Each of these activities may appear as stories on Janie Smith’s News Feed.

9. Facebook users can also react to content that has already been posted. For example, as shown in Figure 1, the photo is accompanied by a thumbs-up “Like” button that allows Margaret Smith and Chester Smith to indicate that they “liked” the photo.

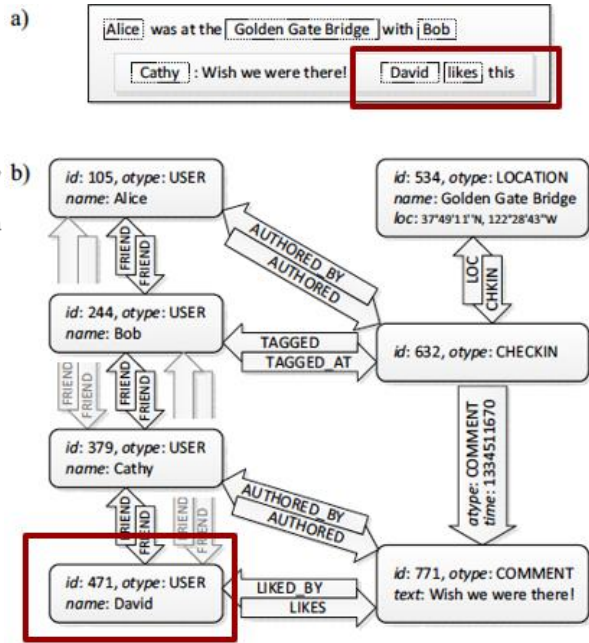
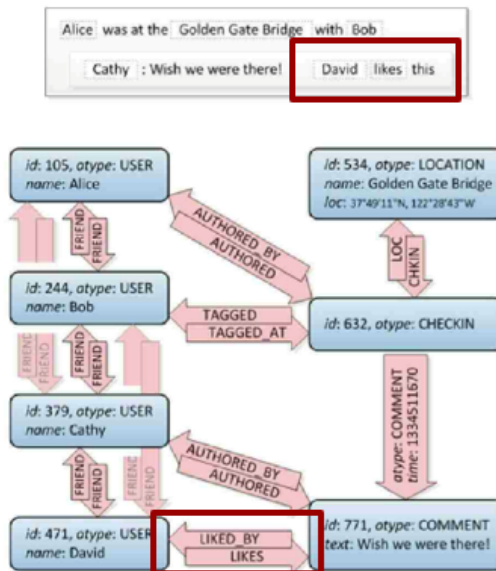


Figure 1: A running example of how a user’s checkin might be mapped to objects and associations.

(E.g., *Mirror Worlds v. Facebook*, Appeal No. 18-2276, ECF 57-2, Appx1104 (Declaration of Nathan Bronson)).

25. The “objects” in the diagram above are shown by the seven blue boxes, and the “associations” are the pink arrows between those boxes. Each of the objects has the unique identifier (“id”) that I mentioned earlier, and an object type (“otype”) indicating what kind of object it is. The four objects shown on the left are USER objects for four Facebook users (Alice, Bob, Cathy, and David). The three objects on the right represent, respectively, the LOCATION object for the Golden Gate Bridge, the CHECKIN object that Alice created when she “checkedin” at that location, and a COMMENT object showing that Cathy wrote a textual comment reacting to Alice’s check-in. All of the actions taken by Alice, Cathy, and David are reflected in the diagram as a series of objects and the associations between them. The pink arrows reflect the type of association between the objects; for example, the bottom pair of arrows shows that David “LIKES”



(E.g., *Mirror Worlds v. Facebook*, Appeal No. 18-2276, ECF 57-2, Appx1113 (Declaration of Nathan Bronson)).

2.1 Integrating Across Stateful Services

Facebook’s back-end infrastructure is comprised of many different services, each using a data placement and communication structure optimized for its workload. Figure 1 shows how data may be duplicated with different sharding strategies in different services.

The TAO caching layer [10], for example, uses multi-tenancy and hierarchical caching to tolerate very high query rates for simple graph queries. Another system materializes secondary indexes to support specific queries that would have a high fanout or require multiple rounds in TAO. Unicorn [14] uses document sharding and a leaf-aggregator model to handle the dynamic queries from Facebook’s Search. The News Feed backend denormalizes all of the information associated with a candidate story so that ranking can be performed locally. Facebook’s deployment of memcache [21, 35] uses the flexible mcrouter proxy [30] to effect a wide variety of replication and locality strategies. Facebook’s Messenger service uses a data store optimized for efficient access to recent messages [20], but often embeds links to TAO nodes. In total, there are hundreds of services running in production at Facebook, many of which maintain a copy of portions of the social graph.

Facebook’s architecture of cooperating heterogeneous services is different from the monolithic service that most research designs assume. This difference alone is not fundamental—Facebook can be externally viewed as a single service even though internally it is comprised

(E.g., *Mirror Worlds v. Facebook*, Appeal No. 18-2276, ECF 57-2, Appx2086 (“Challenges to Adopting Stronger Consistency at Scale”)).

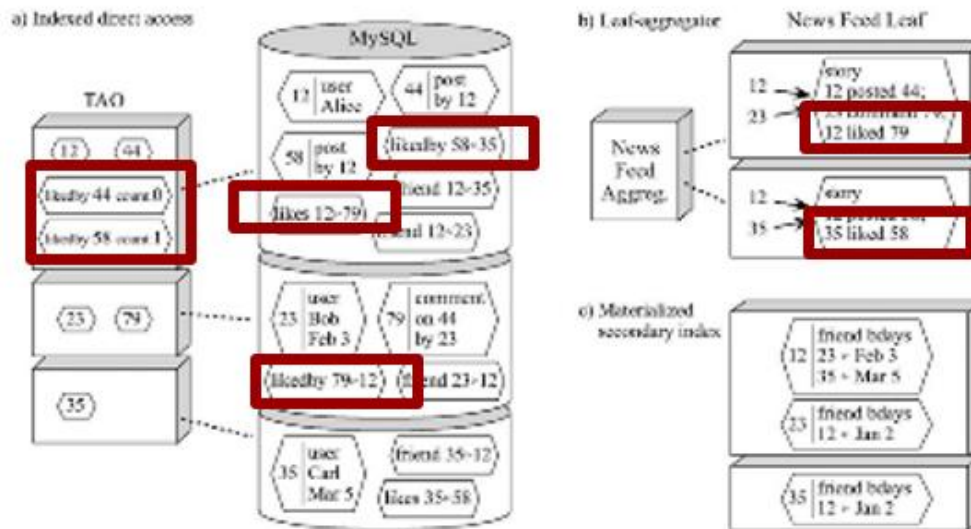


Figure 1: A hypothetical data layout for two posts by Alice, with interactions by friends Bob and Carl. MySQL and TAO (a) shard nodes by key and edges by their source. In News Feed (b) information is aggregated and indexed by all actors. Materialized secondary indexes (c) answer queries that would touch many TAO shards.

(E.g., *Mirror Worlds v. Facebook*, Appeal No. 18-2276, ECF 57-2, Appx2086 (“Challenges to Adopting Stronger Consistency at Scale”)).

27. Plaintiff has been damaged as a result of Defendant’s infringing conduct. Defendant is thus liable to Plaintiff for damages in an amount that adequately compensates Plaintiff for such Defendant’s infringement of the ‘480 Patent, *i.e.*, in an amount that by law cannot be less than would constitute a reasonable royalty for the use of the patented technology, together with interest and costs as fixed by this Court under 35 U.S.C. § 284.

28. On information and belief, Defendant has been aware of the patent family of the ‘480 patent since at least May 2012.

29. On information and belief, to the extent marking is required, VCA has complied with all marking requirements.

IV. COUNT II
(PATENT INFRINGEMENT OF UNITED STATES PATENT NO. 10,339,576)

30. Plaintiff incorporates the above paragraphs herein by reference.

31. On July 2, 2019, United States Patent No. 10,339,576 (“the ‘576 Patent”) was duly and legally issued by the United States Patent and Trademark Office. The ‘576 Patent is titled “Revenue-Generating Electronic Multi-Media Exchange and Process of Operating Same.” A true and correct copy of the ‘576 Patent is attached hereto as Exhibit B and incorporated herein by reference.

32. VCA is the assignee of all right, title, and interest in the ‘576 Patent, including all rights to enforce and prosecute actions for infringement and to collect damages for all relevant times against infringers of the ‘576 Patent. Accordingly, VCA possesses the exclusive right and standing to prosecute the present action for infringement of the ‘576 Patent by Defendant.

33. The application leading to the ‘576 patent was filed September 9, 2016, which was a continuation of application no. 13/679,659, which issued as United States Patent No. 9,477,665, which was a continuation of application no. 14/308,064 which issued as the ‘480 Patent. (Ex. B at cover).

34. The ‘576 Patent shares the identical specification as the ‘480 patent and therefore VCA incorporates the background and discussion of the invention in Paragraphs 11-20.

35. Claim 1 involves a system for generating multimedia content. The claim requires, among other things, the automatic generation of multimedia content for view on a plurality of user devices. The claim requires that the content be generated in a very specific way by applying an electronic filter to a plurality of electronic media submissions stored on one or more database, the filter having criteria associated with one or more users. This allows automatic generation of multimedia content in a much quicker and easier fashion based on specific user criteria. There is nothing abstract about this very particular, unconventional, and non-routine system for the

generation of multimedia content as specifically claimed and there is no risk of preempting creating and distribution contention generally, or even within the context of the Internet.

36. **Direct Infringement.** Upon information and belief, Defendant has been directly infringing claim 1 of the '576 Patent in Texas, and elsewhere in the United States, by employing a computer-based system using <https://www.facebook.com/> (“Accused Instrumentality”) (e.g., <https://www.facebook.com/>).

37. The Accused Instrumentality is a computer-based system for generating multimedia content comprising (a) an electronic media submissions server subsystem. The Accused Instrumentality comprises one or more data processing apparatus, one or more database stored on a non-transitory medium; and a submissions electronic interface configured to receive a first electronic media submission from a first user of a plurality of users over a public network and store said first electronic media submission in said one or more database with at least a second electronic media submission received from a second user of the plurality of users, where the second user is not the first user.

2.1 Integrating Across Stateful Services

Facebook’s back-end infrastructure is comprised of many different services, each using a data placement and communication structure optimized for its workload. Figure 1 shows how data may be duplicated with different sharding strategies in different services.

The TAO caching layer [10], for example, uses multi-tenancy and hierarchical caching to tolerate very high query rates for simple graph queries. Another system materializes secondary indexes to support specific queries that would have a high fanout or require multiple rounds in TAO. Unicorn [14] uses document sharding and a leaf-aggregator model to handle the dynamic queries from Facebook’s Search. The News Feed back-end denormalizes all of the information associated with a candidate story so that ranking can be performed locally. Facebook’s deployment of memcache [21, 35] uses the flexible mcrouter proxy [30] to effect a wide variety of replication and locality strategies. Facebook’s Messenger service uses a data store optimized for efficient access to recent messages [20], but often embeds links to TAO nodes. In total, there are hundreds of services running in production at Facebook, many of which maintain a copy of portions of the social graph.

Facebook’s architecture of cooperating heterogeneous services is different from the monolithic service that most research designs assume. This difference alone is not fundamental—Facebook can be externally viewed as a single service even though internally it is comprised

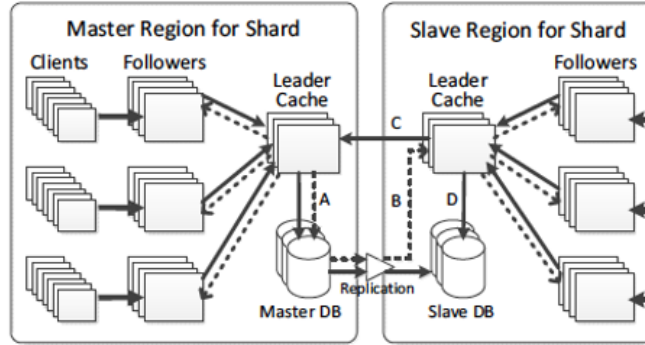


Figure 2: Multi-region TAO configuration. The master region sends read misses, writes, and embedded consistency messages to the master database (A). Consistency messages are delivered to the slave leader (B) as the replication stream updates the slave database. Slave leader sends writes to the master leader (C) and read misses to the replica DB (D). The choice of master and slave is made separately for each shard.

(*Mirror Worlds v. Facebook*, Appeal No. 18-2276, ECF 57-2, Appx2086 (“Challenges to Adopting Stronger Consistency at Scale”)).

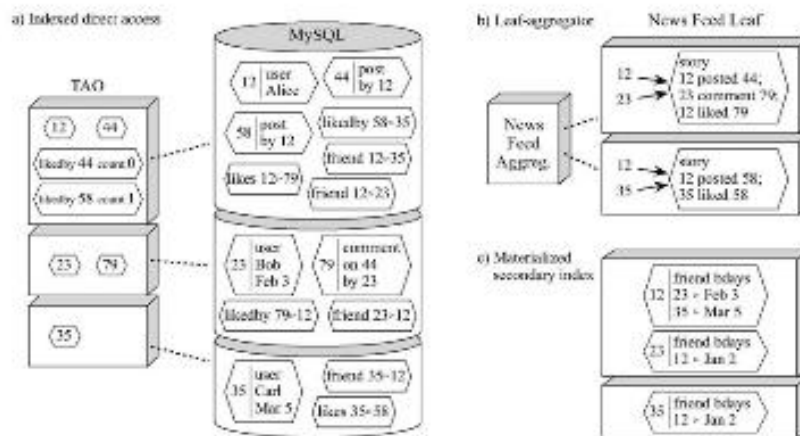


Figure 1: A hypothetical data layout for two posts by Alice, with interactions by friends Bob and Carl. MySQL and TAO (a) shard nodes by key and edges by their source. In News Feed (b) information is aggregated and indexed by all actors. Materialized secondary indexes (c) answer queries that would touch many TAO shards.

(*E.g., Mirror Worlds v. Facebook*, Appeal No. 18-2276, ECF 57-2, Appx2086 (“Challenges to Adopting Stronger Consistency at Scale”)).

- (a) an electronic media submissions server subsystem including:
- (1) one or more data processing apparatus,
 - (2) one or more database stored on a non-transitory medium; and
 - (3) a submissions electronic interface configured to receive a first electronic media submission from a first user of a plurality of users over a public network and store said first electronic media submission in said one or more database with at least a second electronic media submission received from a second user of the plurality of users, where the second user is not the first user,
- wherein the first electronic media submission includes:
- (i) data identifying the first user,
 - (ii) data identifying date and time associated with receipt of the first electronic media submission, and
 - (iii) data indicating content of the first electronic media submission.




Figure 2: Multi-region TAO configuration. The master region sends read misses, writes, and embedded consistency messages to the master database (A). Consistency messages are delivered to the slave leader (B) as the replication stream updates the slave database. Slave leader sends writes to the master leader (C) and read misses to the replica DB (D). The choice of master and slave is made separately for each shard.

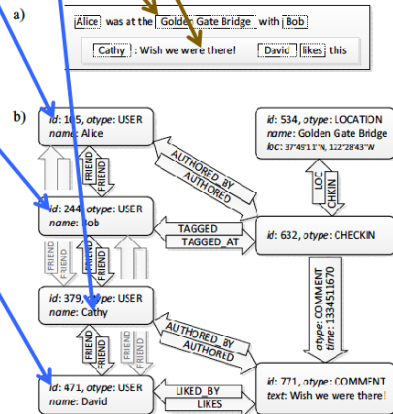
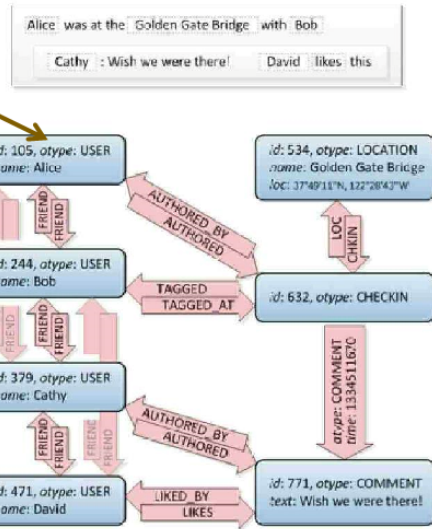


Figure 1: A running example of how a user's checkin might be mapped to objects and associations.

- (a) an electronic media submissions server subsystem including:
- (1) one or more data processing apparatus,
 - (2) one or more database stored on a non-transitory medium; and
 - (3) a submissions electronic interface configured to receive a first electronic media submission from a first user of a plurality of users over a public network and store said first electronic media submission in said one or more database with at least a second electronic media submission received from a second user of the plurality of users, where the second user is not the first user,
- wherein the first electronic media submission includes:
- (i) data identifying the first user,
 - (ii) data identifying date and time associated with receipt of the first electronic media submission, and
 - (iii) data indicating content of the first electronic media submission,



(E.g., *Mirror Worlds v. Facebook*, Appeal No. 18-2276, ECF 57-2, Appx1113 (Declaration of Nathan Bronson)).

38. The Accused Instrumentality comprises wherein the first electronic media submission includes data identifying the first user, data identifying date and time associated with receipt of the first electronic media submission, and data indicating content of the first electronic media submission. As Nathan Bronson, Facebook’s expert in a different case explained,

The “objects” in the diagram [below] are shown by the seven blue boxes, and the “associations” are the pink arrows between those boxes. Each of the objects has the unique identifier (“id”) that I mentioned earlier, and an object type (“otype”) indicating what kind of object it is. The four objects shown on the left are USER objects for four Facebook users (Alice, Bob, Cathy, and David). The three objects on the right represent, respectively, the LOCATION object for the Golden Gate Bridge, the CHECKIN object that Alice created when she “checkedin” at that location, and a COMMENT object showing that Cathy wrote a textual comment reacting to Alice’s check-in. All of the actions taken by Alice, Cathy, and David are reflected in the diagram as a series of objects and the associations between them. The pink arrows reflect the type of association between the objects; for example, the bottom pair of arrows shows that David “LIKES”

(*E.g., Mirror Worlds v. Facebook*, Appeal No. 18-2276, ECF 57-2, Appx1113 (Declaration of Nathan Bronson)).

39. The Accused Instrumentality contains (b) the one or more databases comprising criteria associated with one or more users of the plurality of users stored therein. As shown below, the Accused Instrumentality consists of at least one database.

Facebook has more than a billion active users who record their relationships, share their interests, upload text, images, and video, and curate semantic information about their data [2]. The personalized experience of social applications comes from timely, efficient, and scalable access to this flood of data, the *social graph*. In this paper we introduce TAO, a read-optimized graph data store we have built to handle a demanding Facebook workload.

40. The Accused Instrumentality comprises (c) an electronic multimedia creator server subsystem operatively coupled to the electronic media submissions server subsystem, including one or more second data processing apparatus, and an electronic content filter configured to apply criteria associated with at least one user of the plurality of users to obtain a plurality of electronic media submissions from the one or more database and to develop multimedia content to be electronically available for viewing on at least one user device associated with the first user wherein data identifying a respective user is maintained for each electronic media submission within the multimedia content; and.

(c) an electronic multimedia creator server subsystem operatively coupled to the electronic media submissions server subsystem, including:
 (1) one or more second data processing apparatus, and
 (2) an electronic content filter configured to apply criteria associated with at least one user of the plurality of users to obtain a plurality of electronic media submissions from the one or more database and to develop multimedia content to be electronically available for viewing on at least one user device associated with the first user,
wherein data identifying a respective user is maintained for each electronic media submission within the multimedia content; and

2 Background

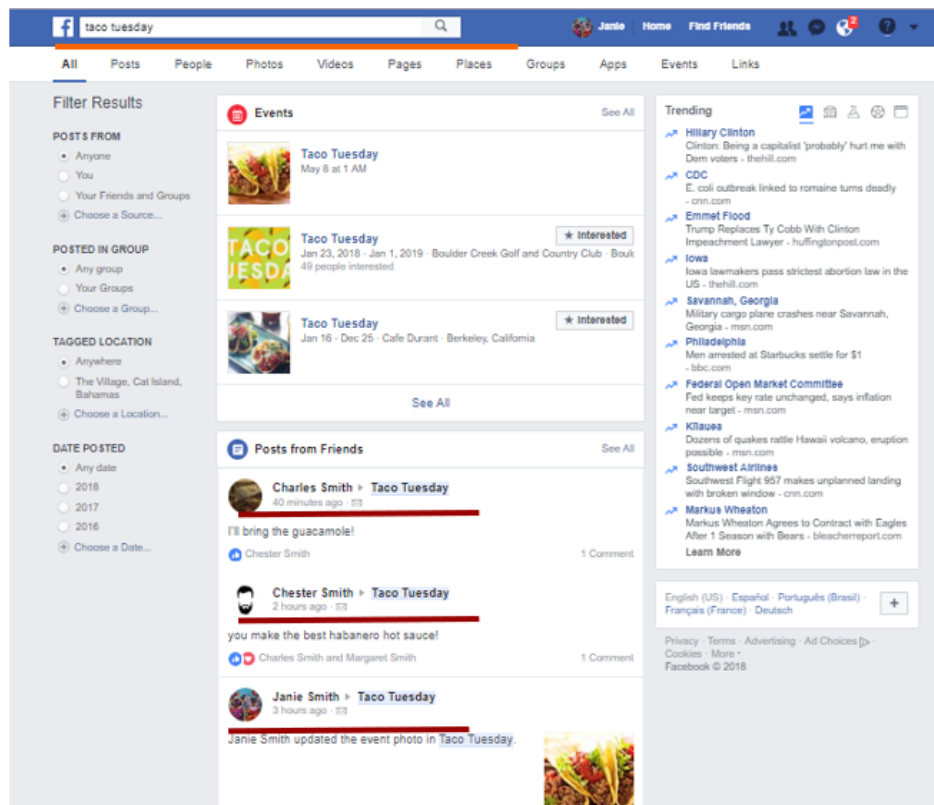
A single Facebook page may aggregate and filter hundreds of items from the social graph. We present each user with content tailored to them, and we filter every item with privacy checks that take into account the current viewer. This extreme customization makes it infeasible to perform most aggregation and filtering when content is created; instead we resolve data dependencies and check privacy each time the content is viewed. As much as possible we pull the social graph, rather than pushing it. This implementation strategy places extreme read demands on the graph data store; it must be efficient, highly available, and scale to high query rates.

3.1 Objects and Associations

TAO *objects* are typed nodes, and TAO *associations* are typed directed edges between objects. Objects are identified by a 64-bit integer (*id*) that is unique across all objects, regardless of object type (*otype*). Associations are identified by the source object (*id1*), association type (*atype*) and destination object (*id2*). At most one association of a given type can exist between any two objects. Both objects and associations may contain data as key→value pairs. A per-type schema lists the possible keys, the value type, and a default value. Each association has a 32-bit time field, which plays a central role in queries¹.

Object: (*id*) → (*otype*, (key → value)*)

Assoc.: (*id1*, *atype*, *id2*) → (time, (key → value)*)



41. As explained by Nathan Bronson, “Facebook also offers search features that allow users to search for items on Facebook and to filter the results based on various criteria. For example, a simple search of ‘Taco Tuesday’ using Janie Smith’s account yields the event she created using the Events feature, as well as other matches.” (*E.g., Mirror Worlds v. Facebook*, Appeal No. 18-2276, ECF 57-2, Appx1109-10 (Declaration of Nathan Bronson)).

Aggregator: This Aggregator also collects and ranks information. Appx2446. As shown in the figure reproduced above, the Aggregator receives data from the front-end (“[REDACTED]”) that specifies parameters to retrieve information that a particular user may be interested in, including a [REDACTED] [REDACTED] and [REDACTED] values. Appx2422; Appx1496; Appx1501-1502 (Koskinen Decl. ¶¶ 73, 88-90) (discussing the data fields contained in this communication to the Aggregator).

(E.g., *Mirror Worlds v. Facebook*, Appeal No. 18-2276, ECF 34, at pp. 17-18).

42. The Accused Instrumentality includes an electronic release subsystem operatively coupled to the electronic multimedia creator server subsystem, including one or more third data processing apparatus and configured to make the multimedia content electronically available for viewing on a plurality of user devices.

(d) an electronic release subsystem operatively coupled to the electronic multimedia creator server subsystem, including one or more third data processing apparatus and configured to make the multimedia content electronically available for viewing on a plurality of user devices.




Figure 2: Multi-region TAO configuration. The master region sends read misses, writes, and embedded consistency messages to the master database (A). Consistency messages are delivered to the slave leader (B) as the replication stream updates the slave database. Slave leader sends writes to the master leader (C) and read misses to the replica DB (D). The choice of master and slave is made separately for each shard.

43. Plaintiff has been damaged as a result of Defendant's infringing conduct. Defendant is thus liable to Plaintiff for damages in an amount that adequately compensates Plaintiff for such Defendant's infringement of the '576 Patent, *i.e.*, in an amount that by law cannot be less than would constitute a reasonable royalty for the use of the patented technology, together with interest and costs as fixed by this Court under 35 U.S.C. § 284.

44. On information and belief, Defendant has been aware of the patent family of the '576 patent since at least May 2012.

45. On information and belief, to the extent marking is required, VCA has complied with all marking requirements.

V. JURY DEMAND

Plaintiff, under Rule 38 of the Federal Rules of Civil Procedure, requests a trial by jury of any issues so triable by right.

VI. PRAYER FOR RELIEF

WHEREFORE, Plaintiff respectfully requests that the Court find in its favor and against Defendant, and that the Court grant Plaintiff the following relief:

- a. Judgment that one or more claims of United States Patent No. 9,501,480 have been infringed, either literally and/or under the doctrine of equivalents, by Defendant;
- b. Judgment that one or more claims of United States Patent No. 10,339,576 have been infringed, either literally and/or under the doctrine of equivalents, by Defendant;
- c. Judgment that Defendant account for and pay to Plaintiff all damages to and costs incurred by Plaintiff because of Defendant's infringing activities and other conduct complained of herein, and an accounting of all infringements and damages not presented at trial;
- d. That Plaintiff be granted pre-judgment and post-judgment interest on the damages caused by Defendant's infringing activities and other conduct complained of herein; and
- e. That Plaintiff be granted such other and further relief as the Court may deem just and proper under the circumstances.

March 11, 2022

Respectfully Submitted,

/s/ David R. Bennett

David R. Bennett

(Admitted to the U.S. Dist. Ct. for the W.D. Texas)

Direction IP Law

P.O. Box 14184

Chicago, IL 60614-0184

(312) 291-1667

dbennett@directionip.com

Attorney for Plaintiff

Virtual Creative Artists, LLC