



(19) **United States**

(12) **Patent Application Publication**
Beckman et al.

(10) **Pub. No.: US 2019/0228177 A1**
(43) **Pub. Date: Jul. 25, 2019**

(54) **HOSTED APPLICATION GATEWAY ARCHITECTURE WITH MULTI-LEVEL SECURITY POLICY AND RULE PROMULGATIONS**

(60) Provisional application No. 61/880,481, filed on Sep. 20, 2013.

Publication Classification

(71) Applicant: **Open Text SA ULC**, Halifax (CA)

(51) **Int. Cl.**
G06F 21/62 (2006.01)

(72) Inventors: **Gregory Beckman**, Baden (CA);
Robert Laird, Waterloo (CA);
Geoffrey Michael Obbard, Waterloo (CA)

(52) **U.S. Cl.**
CPC **G06F 21/6227** (2013.01)

(73) Assignee: **Open Text SA ULC**, Halifax (CA)

(57) **ABSTRACT**


(21) Appl. No.: **16/371,852**

A hosted application gateway server node may be communicatively coupled to backend systems, client devices, and database shards associated with database servers. Through the gateway server node, various services may be provided to managed containers running on client devices such that enterprise applications can be centrally managed. A sharding manager may manage relationships of database items across database shards. Each shard stores a copy of a table representing a split of a relationship. A shard ID mask is included in each item's ID. At query time, the shard ID can be extracted and used to query the correct database. This query routing mechanism allows navigation from one shard to another when multiple items are in a relationship (e.g., share the same resource such as a document). As such, embodiments can eliminate the need for APIs to join in data that span multiple shards.

(22) Filed: **Apr. 1, 2019**

Related U.S. Application Data

(63) Continuation of application No. 15/653,017, filed on Jul. 18, 2017, now Pat. No. 10,268,835, which is a continuation of application No. 14/534,623, filed on Nov. 6, 2014, now Pat. No. 9,747,466, which is a continuation-in-part of application No. 14/491,386, filed on Sep. 19, 2014, now Pat. No. 9,979,751.



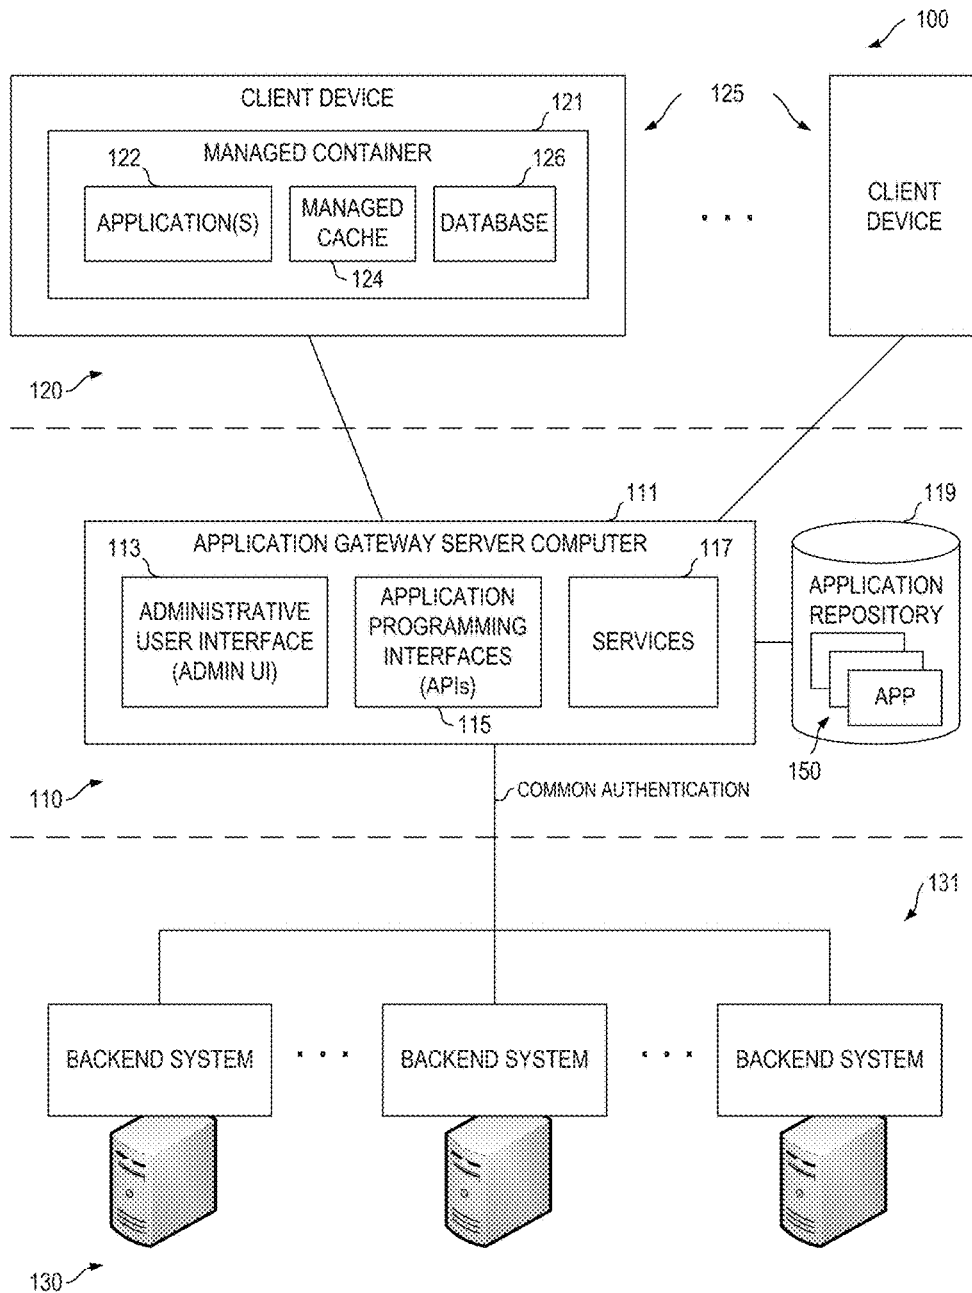


FIG. 1

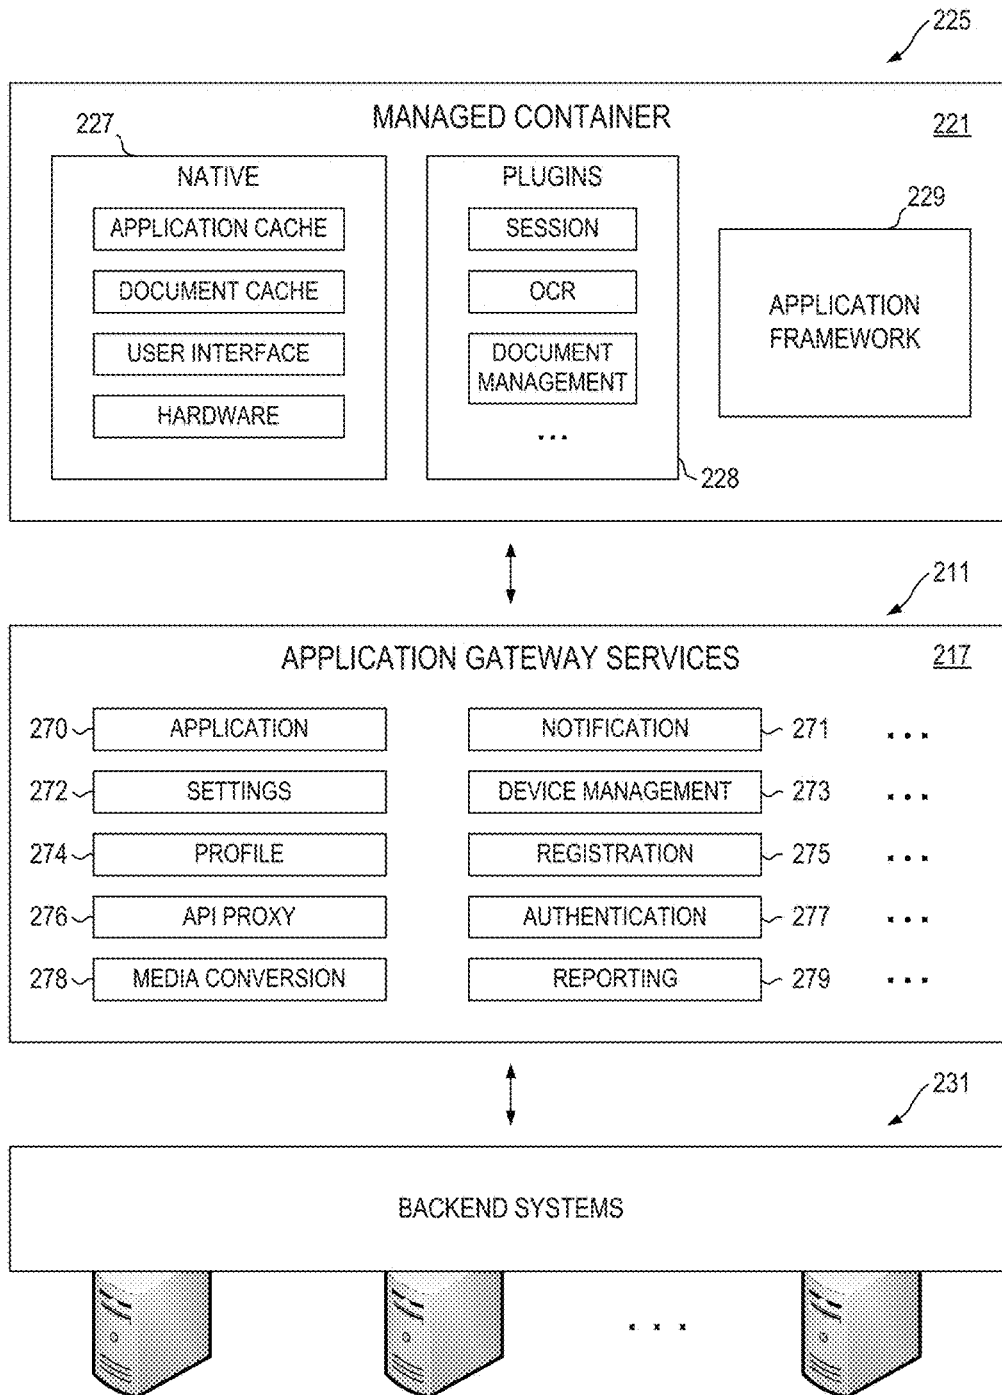




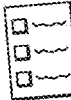



FIG. 2

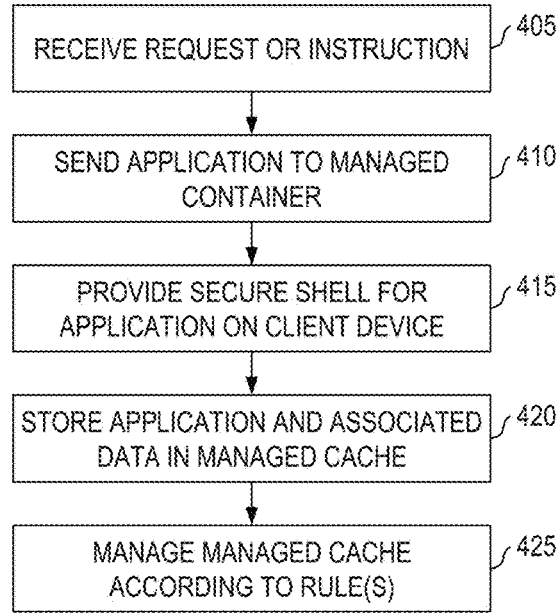
300

FIG. 3

Choose File No file chosen Install

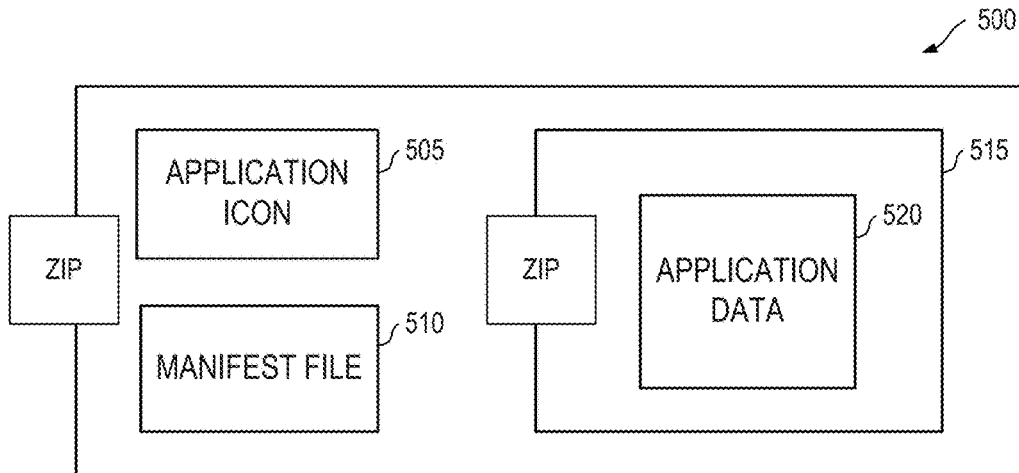
Application	Deployment	Actions
 Content Server Favorites Favorites within Content Server Version 1	Mandatory	CONFIGURE GET INFO DELETE
 Note Enterprise Note Application Version 2	Selective	CONFIGURE GET INFO DELETE
 Content Server Pulse Status and Comments within Content Server Version 1	Mandatory	CONFIGURE GET INFO DELETE
 Social Enterprise Social Application Version 4	Optional	CONFIGURE GET INFO DELETE
 Content Server Tasks Tasks within Content Server Version 1	Mandatory	CONFIGURE GET INFO DELETE
 Vender Invoice Management Vender Invoice Management Version 6	Mandatory	CONFIGURE GET INFO DELETE

330 350 370



400 ↗

FIG. 4



↘ 500

FIG. 5

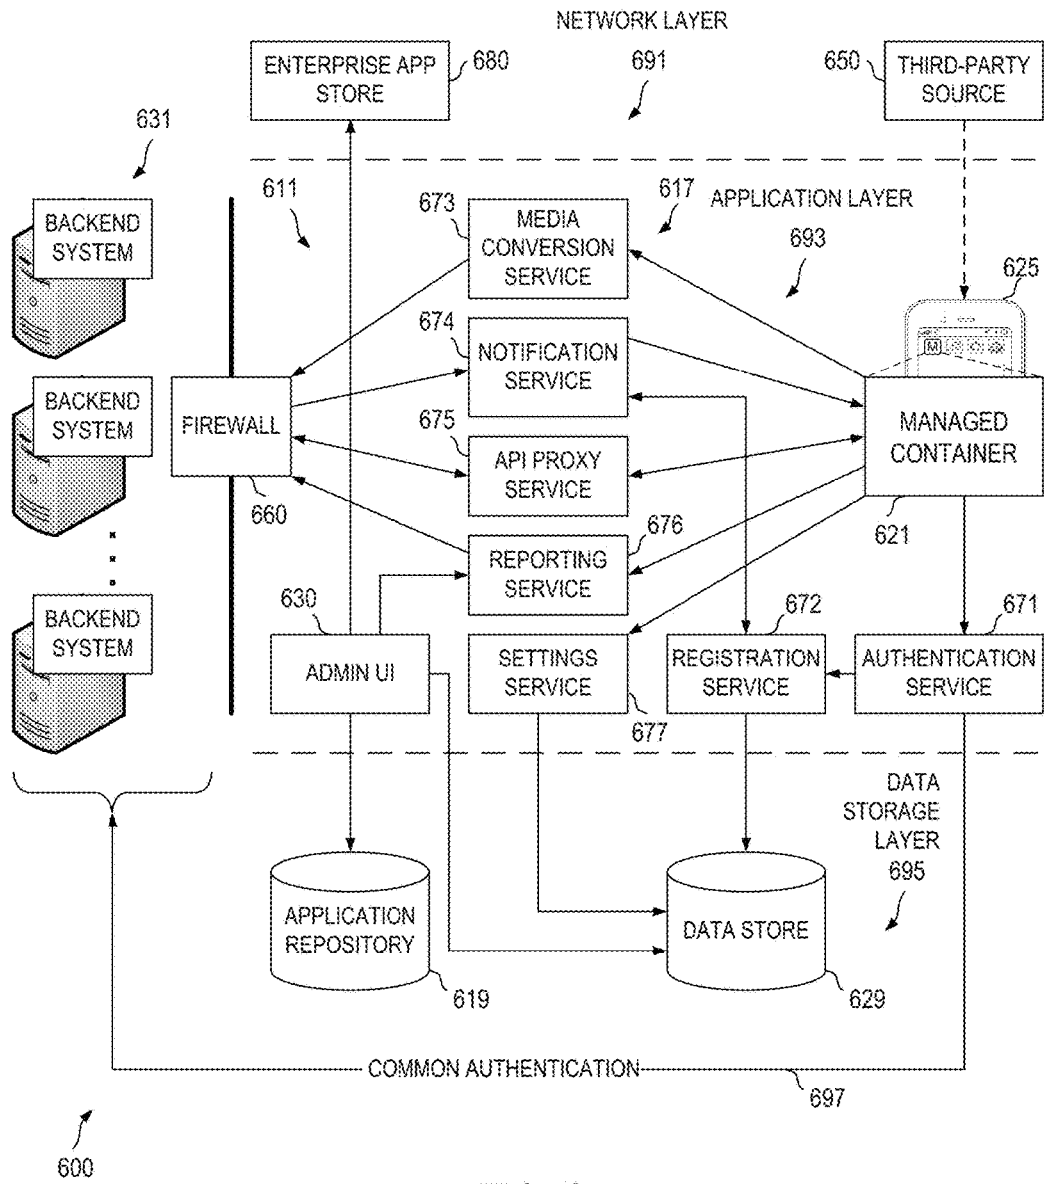


FIG. 6




FIG. 7

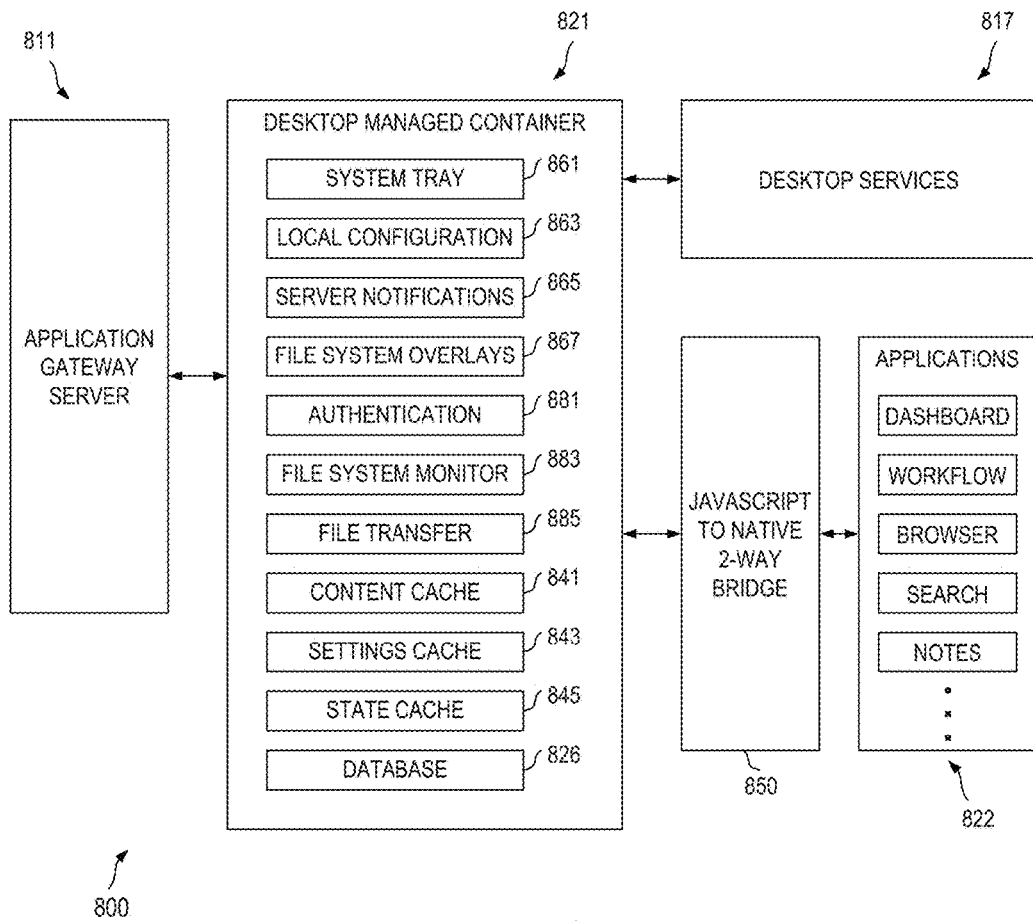


FIG. 8




FIG. 9




FIG. 10




FIG. 11A

FIG. 11B




FIG. 12A




FIG. 12B




FIG. 13A




FIG. 13B

**HOSTED APPLICATION GATEWAY
ARCHITECTURE WITH MULTI-LEVEL
SECURITY POLICY AND RULE
PROMULGATIONS**

CROSS REFERENCE TO RELATED
APPLICATIONS

[0001] This is a continuation of U.S. patent application Ser. No. 15/653,017, filed Jul. 18, 2017, entitled “HOSTED APPLICATION GATEWAY ARCHITECTURE WITH MULTI-LEVEL SECURITY POLICY AND RULE PROMULGATIONS,” which is a continuation of U.S. patent application Ser. No. 14/534,623, filed Nov. 6, 2014, entitled “HOSTED APPLICATION GATEWAY ARCHITECTURE WITH MULTI-LEVEL SECURITY POLICY AND RULE PROMULGATIONS,” issued as U.S. Pat. No. 9,747,466, which is a continuation-in-part of U.S. patent application Ser. No. 14/491,386, filed Sep. 19, 2014, entitled “APPLICATION GATEWAY ARCHITECTURE WITH MULTI-LEVEL SECURITY POLICY AND RULE PROMULGATIONS,” issued as U.S. Pat. No. 9,979,751, which claims a benefit of priority under 35 U.S.C. § 119(e) from U.S. Provisional Application No. 61/880,481, filed Sep. 20, 2013. This application relates to U.S. patent application Ser. No. 14/491,451, filed Sep. 19, 2014, which claims a benefit of priority from U.S. Provisional Application No. 61/880,502, filed Sep. 20, 2013; Ser. No. 14/491,492, filed Sep. 19, 2014, which claims a benefit of priority from U.S. Provisional Application No. 61/880,526, filed Sep. 20, 2013; and Ser. No. 14/491,483, filed Sep. 19, 2014, which claims a benefit of priority from U.S. Provisional Application No. 61/880,557, filed Sep. 20, 2013. All applications listed in this paragraph are hereby incorporated by reference as if set forth herein in their entireties, including all appendices attached thereto.

TECHNICAL FIELD

[0002] This disclosure relates generally to content management. More particularly, embodiments disclosed herein relate to a cloud based solution for controlling how backend content can be deployed and managed at client devices through managed containers operating on client devices and an application gateway connected to backend systems.

BACKGROUND OF THE RELATED ART

[0003] Conventional content control software and services are designed to control what content delivered over the Internet may be viewed or blocked from viewing. Generally, access to such controlled content can be restricted at various levels. For instance, a firewall may be used to block access to certain websites or a timer may be set up at a user’s computer to limit the amount of time that an application may be used. Additionally, filters such as e-mail filters and browser-based filters may be used.

[0004] However, such content control software and services are often inadequate to control content downloaded by users to their computers. This can be problematic for enterprises wanting to retain control over enterprise content downloaded to devices that may or may not be owned by the enterprises.

[0005] Additionally, provisioning content may require significant database resources. A procedure known as “sharding” has been used to scale databases beyond what a single

server or cluster of servers can handle. In sharding, a single large database is fragmented or sharded into multiple smaller databases that operate virtually independently. Collectively, the shards appear to form a single, very large database. However, handling relationships between objects that may be in separate shards can be problematic.

SUMMARY OF THE DISCLOSURE

[0006] An object of this disclosure is to provide an effective mechanism by which an entity can retain control over their applications and data associated therewith, even if the applications and/or data have been downloaded onto a device not owned or controlled by the entity. Another object of the disclosure is to provide a secure storage on a user device such that downloaded applications and/or data can be protected from unauthorized access. Yet another object of the disclosure is to bridge the gap between user devices and backend systems such that downloaded applications and/or data can be updated to reflect a change at the backend, for instance, a change in a data policy rule applicable to the downloaded applications and/or data.

[0007] These and other objects can be achieved through embodiments of systems, methods and computer program products disclosed herein. For example, in some embodiments, a method may comprise sending an application from an application gateway server computer to a managed container executing on a client device. Within this disclosure, a managed container refers to a special computer program that can be downloaded from a source.

[0008] The application may be hosted and/or required by a backend system such as a content server. The managed container may provide a secure shell for the application received from the application gateway server computer, store the application and data associated with the application in a managed cache, and control the managed cache in accordance with a set of rules propagated from the backend system to the managed container via the application gateway server computer. All or some of the set of rules may reside on the client device, the backend system, the application gateway server computer, or a combination thereof.

[0009] In some embodiments, the set of rules may include at least one of: a rule controlling storage of data associated with an application received from the application gateway server computer, a rule controlling access to data associated with an application received from the application gateway server computer, or a rule controlling update of data associated with an application received from the application gateway server computer.

[0010] The downloaded application—and any data associated therewith—remains under the control of the managed container regardless of whether the client device has network connectivity (i.e., regardless of whether the client device is or is not connected to application gateway server computer).

[0011] In some embodiments, the secure shell provided by the managed container includes a secure data encryption shell that encrypts the data associated with the application to limit or prevent access to the data by the client device’s own operating system and other applications residing on the client device but not received from the application gateway server computer.

[0012] In some embodiments, at least one of the set of rules propagated from the backend system may determine encryption parameters for encrypting the data stored in the

managed cache. In turn, the secure data encryption shell may encrypt the data based on the encryption parameters.

[0013] In some embodiments, the encryption parameters may be shared between the managed container and the backend system, via the application gateway server computer, to enable shared secure access to the data between and among the applications received from the application gateway server computer and the one or more backend systems.

[0014] One embodiment comprises a system comprising a processor and a non-transitory computer-readable storage medium that stores computer instructions translatable by the processor to perform a method substantially as described herein. Another embodiment comprises a computer program product having a non-transitory computer-readable storage medium that stores computer instructions translatable by a processor to perform a method substantially as described herein.

[0015] As an example, one embodiment of a system may include an application gateway server computer communicatively connected to backend systems and client devices. The backend systems as well as the client devices may operate on different platforms. The application gateway server computer may have application programming interfaces and services configured for communicating with the backend systems and managed containers operating on the client devices.

[0016] The services provided by embodiments of an application gateway server computer disclosed herein may include various types of services that may be generally categorized as core services and product services. In one embodiment, core services may refer to services necessary for building new applications. In one embodiment, product services may refer to services configured for integration of existing products. In this disclosure, these and other services are collectively referred to as “services.”

[0017] In some embodiments, a managed container may be implemented as an application (program) that is native to a client device and that can be downloaded from a source on the Internet such as a website or an app store. As disclosed herein, the managed container includes a managed cache for storing content received from the application gateway server computer, including applications. Applications received from the application gateway server computer are not downloaded from a website or third-party app store. In some embodiments, applications received from the application gateway server computer are written in a markup language for structuring and presenting content on the Internet.

[0018] A further object of this disclosure is to provide an effective mechanism by which relationships between objects in different shards may be handled. This object may be accomplished in accordance with embodiments by maintaining a function table visible to an application programming interface used to access the sharded database. In some embodiments, the function table is a split function table, with one copy stored in each shard involved in the relationship. One copy is keyed to the object in the first shard associated with the relationship and another copy is keyed to the object in the second shard associated with the relationship. Such a sharded database, and the handling of relationships therein, may be particularly advantageous in a system that includes a cloud provisioning “gateway.”

[0019] Numerous other embodiments are also possible.

[0020] These, and other, aspects of the disclosure will be better appreciated and understood when considered in con-

junction with the following description and the accompanying drawings. It should be understood, however, that the following description, while indicating various embodiments of the disclosure and numerous specific details thereof, is given by way of illustration and not of limitation. Many substitutions, modifications, additions and/or rearrangements may be made within the scope of the disclosure without departing from the spirit thereof, and the disclosure includes all such substitutions, modifications, additions and/or rearrangements.

BRIEF DESCRIPTION OF THE DRAWINGS

[0021] The drawings accompanying and forming part of this specification are included to depict certain aspects of the disclosure. It should be noted that the features illustrated in the drawings are not necessarily drawn to scale. A more complete understanding of the disclosure and the advantages thereof may be acquired by referring to the following description, taken in conjunction with the accompanying drawings in which like reference numbers indicate like features and wherein:

[0022] FIG. 1 depicts a diagrammatic representation of an example embodiment of an application gateway architecture implementing a multi-level content control mechanism;

[0023] FIG. 2 depicts a diagrammatic representation of an example embodiment of a system implementing one embodiment of application gateway architecture;

[0024] FIG. 3 depicts a diagrammatic representation of an example embodiment of an application administration interface of an application gateway server computer;

[0025] FIG. 4 is a flow chart illustrating an example of a method of operation according to some embodiments;

[0026] FIG. 5 depicts a diagrammatic representation of an example embodiment of an application according to some embodiments;

[0027] FIG. 6 depicts a diagrammatic representation of an example of an application gateway architecture in operation according to some embodiments;

[0028] FIG. 7 depicts a diagrammatic representation of an example of a managed container operating on a type of client device according to some embodiments;

[0029] FIG. 8 depicts a diagrammatic representation of an example of a managed container operating on another type of client device according to some embodiments;

[0030] FIG. 9 depicts diagrammatic representation of an example architecture that may use managed containers according to embodiments;

[0031] FIG. 10 depicts a diagrammatic representation of an example embodiment of a cloud-based environment for handling managed containers according to embodiments;

[0032] FIG. 11A and FIG. 11B depict examples of handling relationships between objects in a single database;

[0033] FIG. 12A and FIG. 12B depict examples of handling relationships across shards; and

[0034] FIG. 13A and FIG. 13B depict diagrammatic representations of example embodiments of sharding management according to embodiments;

DETAILED DESCRIPTION

[0035] The invention and the various features and advantageous details thereof are explained more fully with reference to the non-limiting embodiments that are illustrated in the accompanying drawings and detailed in the following

description. Descriptions of well-known processing techniques, components and equipment are omitted so as not to unnecessarily obscure the invention in detail. It should be understood, however, that the detailed description and the specific examples, while indicating preferred embodiments of the invention, are given by way of illustration only and not by way of limitation. Various substitutions, modifications, additions and/or rearrangements within the spirit and/or scope of the underlying inventive concept will become apparent to those skilled in the art from this disclosure. Embodiments discussed herein can be implemented in suitable computer-executable instructions that may reside on a computer readable medium (e.g., a hard disk (HD)), hardware circuitry or the like, or any combination.

[0036] Managed Container and Gateway Architecture

[0037] Embodiments disclosed herein provide a “gateway” that controls how backend content (e.g., enterprise content owned by an enterprise) is sent out of a backend system (e.g., a content server owned by the enterprise) to or downloaded by a client device. In this disclosure, this is referred to as a first layer or level of content management or server side content management layer. Embodiments also provide a second layer or level of content management at the client side. Thus, there are controls at both the server side and the client side, providing the enterprise with the ability to enforce company policy and rules on how enterprise content is managed both in and out of backend systems and at client devices.

[0038] Embodiments of an application gateway server computer disclosed herein can, on behalf of one or more backend systems connected thereto, control content distribution to managed containers operating on client devices. Within this disclosure, a managed container refers to a special computer program for reviewing, accessing, and downloading applications via an application gateway server computer. According to embodiments, a managed container can be downloaded from a source or a network site on a private or public network such as a company’s intranet or the Internet. Examples of an appropriate source may include a service or an online store (which may, in some embodiments, be referred to as an “app store”).

[0039] As described herein, a managed container can control content at a client device (e.g., how a document is stored, accessed, updated, removed, etc. . . .). For example, if a backend system (e.g., a content server, an information management system, a document repository, a business process management system, a social server, a records management (RM) system, etc.) has a policy or rule update, or a new rule on content retrieved from the backend system, the application gateway server can broadcast the update or new rule to appropriate managed containers operating on various client devices or otherwise notify the appropriate managed containers about the update or new rule (appropriate in the sense that the content, to which the update or new rule is applicable, is stored in the managed caches of such managed containers). In this way, rules can be promulgated out to appropriate client devices by the application gateway server and applied by the managed containers on those client devices to content living on the client devices without needing any help from applications associated with the content.

[0040] For example, suppose due to a policy update, access to a set of records stored in an RM system is changed to a different security level or user group. The application

gateway server can determine which managed containers store a copy of the set of records (or a portion thereof) and broadcast this change to those managed containers or notify them about the change. When a managed container receives a notice (which may be sent by the application gateway server computer using a communications channel that is different from the broadcast channel, as those skilled in the art can appreciate), the managed container may initiate a connection with the application gateway server computer to retrieve the policy update. The managed containers may then apply the policy update and change the security access to the copy of the set of records (or a portion thereof) stored in their managed caches accordingly. As this example illustrates, no user is required to open up an RM application on their device in order for the policy update from the RM system at the backend to take effect on the copy of the set of records (or a portion thereof) locally stored in the managed cache on their device.

[0041] The content control mechanism described above can be implemented in various ways. FIG. 1 depicts a diagrammatic representation of an example embodiment of an application gateway architecture implementing a multi-layer (or multi-level) content control mechanism. In the example of FIG. 1, system 100 may include application gateway server computer 110 communicatively connected to backend systems 131 and one or more client devices 125. Client device 125 shown in FIG. 1 is representative of various client devices. Those skilled in the art will appreciate that FIG. 1 shows a non-limiting example of client device 125. Backend systems 131 may comprise computer program products and/or applications developed within a company and/or by third party developers/companies. Non-limiting examples of backend systems 131 may include a content server, an information management system, a document repository, a process management system, a social server, an RM system, a database management system, an enterprise resources planning system, a collaboration and management system, a customer relationship management system, a search system, an asset management system, a case management system, etc.

[0042] In some embodiments, a first layer of content management 110 (“level 110”) can be realized in application gateway server computer 111 configured for controlling how backend content (e.g., applications that communicate with backend systems, documents created/used by such applications, etc.) can be sent out of the backend systems to client devices. A second layer of content management 120 (“level 120”) can be realized in managed containers 121 operating on client devices 125. A third layer of content management 130 (“level 130”) may include proprietary and/or third-party content management tools used by various backend systems 131.

[0043] At level 110, as shown in FIG. 1, application gateway server computer 111 may include application programming interfaces (APIs) 115 and services 117 configured for communicating with backend systems 131 and managed containers 121 operating on client devices 125. In some embodiments, applications 150 may be developed within a company and/or externally sourced and stored in application repository 119 accessible by application gateway server computer 111. Applications 150 may be associated with backend systems 131. These server-side components are explained further below.

[0044] At level 120, managed container 121 operating on client device 125 may include managed cache 124 for storing various applications 122 downloaded/pulled or received/pushed from application gateway server computer 111. All the data, documents, and files associated with applications 122 may be encrypted and stored in managed cache 124. To this end, managed cache 124 can be considered a local application repository that can provide client device 125 with offline access to cached applications 122. In some embodiments, database 126 may be used by managed container 121 to keep track of content stored in managed cache 124. Managed container 121 can be installed and run on client device 125 separate and independent of any applications that it manages. These client-side components are explained further below.

[0045] In some embodiments, managed cache 124 may store the user interface components of applications 122. However, as described below with reference to FIG. 5, not all components of applications 122 are stored in managed cache 124. In some embodiments, when an application is called (i.e., invoked by a user using client device 125 on which the application is installed), managed container 124 may obtain any data, document(s), and/or file(s) that the application needs from a backend system through application gateway server computer 111. This has the benefits of reducing the storage requirement of having applications 122 on client device 125, expediting the network transmission of applications 122, and keeping applications 122 always up-to-date.

[0046] In some embodiments, the content of managed cache 124 is managed in accordance with a set of rules. The set of rules may include rules residing at one of more backend systems 131, rules residing at application gateway server computer 111, rules residing at client device 125, or a combination thereof. In some embodiments, the set of rules may include at least one of: a data storage rule for controlling storage of the data associated with applications 122 received from application gateway server computer 111, a data policy rule for controlling access to the data associated with applications 122 received from application gateway server computer 111, an application rule for controlling at least one of applications 122 received from application gateway server computer 111, or an update rule for controlling update of the data associated with applications 122 received from application gateway server computer 111.

[0047] In some embodiments, the set of rules may be stored on client device 125. Managed container 121 may use the stored set of rules to control and/or protect the data associated with applications 122 received from application gateway server computer 111. For example, in some embodiments, when an update to one of the rules is propagated from backend system 131 to managed container 121 via application gateway server computer 111, managed container 121 may execute, based on the updated rule, an update to the data associated with applications 122 received from application gateway server computer 111. As another example, in some embodiments, managed container 121 may use the stored application rule to control application(s) 122 received from application gateway server computer 111.

[0048] In some embodiments, at least one of the set of rules may determine encryption parameters for encrypting the content of managed cache 124. Managed container 121 may encrypt the content of managed cache 124 based on the encryption parameters. In some embodiments, the encryption

parameters may be shared between managed container 121 and one or more backend systems 131, via application gateway server computer 111, to enable shared secure access to the data between and among applications 122 received from application gateway server computer 111 and one or more backend systems 131. Regardless of network connectivity of client device 125, applications 122 and the associated data stored in managed cache 124 are under control of managed container 121. In this way, unauthorized access to the data stored in managed cache 124 can be limited or prevented. Unauthorized access may include access by an operating system running on client device 125 and/or access by non-managed applications executing on client device 125 such as those downloaded onto client device 125 without going through application gateway server computer 111.

[0049] In some embodiments, users (e.g., employees of a company operating or using an application gateway server computer) do not need to or are not allowed to download (e.g., from an online app store or a website on the Internet) any application into a managed container (although they may still download and install applications on their devices as usual and such applications are outside the scope of this disclosure). Rather, an administrator may, via administrative user interface 113 (“admin UI”) load into managed containers on client devices associated with these users with select applications and/or services available on application gateway server computer 111. For example, an RM application may be needed to access an RM system, a search application may be needed to search a content repository, etc. Depending upon the role or job function of a user, one or more of these applications may be loaded into the managed container (s) on the device(s) associated with the user.

[0050] In this way, an application gateway server computer can inject new applications directly into a managed container running on a client device and remotely manage (e.g., replace, update, change, repair, remove, etc.) any of the injected applications without going through any intermediary entity such as an online app store, website, or application developer. To this end, system 100 can advantageously provide a development and integration platform for the rapid creation, administration, and distribution of applications that can be deployed and centrally managed on a variety of mobile, desktop, and web platforms. From the perspective of entities, system 100 can provide a common point of authentication where one set of credentials can provide access to various backend systems. Furthermore, system 100 can provide a secure and managed enterprise information delivery channel for client mobile and desktop platforms. From the perspective of developers, system 100 can provide a standards-based integration platform with a “write-once, run-anywhere” application development environment. Further, as explained below, system 100 can be deployed on-premises or in a cloud.

[0051] Referring to FIG. 2, which depicts an example embodiment of a system implementing one embodiment of application gateway architecture described above. In the example of FIG. 2, services 217 provided by application gateway server computer 211 may include services that are necessary for building new applications (also referred to as “core services”) and computer program product services for application developers to integrate existing products (also referred to as “product services”). In this disclosure, these and other services provided by application gateway server computer 211 are collectively referred to as “services.”

Examples of services 217 are provided below. Each of services 217 may have a corresponding API such that they can appropriately communicate with backend systems 231 and client devices 225 connected to application gateway server computer 211. As a specific example, JavaScript Object Notation (JSON) RESTful APIs may be used to communicate with backend systems 231. In some embodiments, HyperText Transfer Protocol (HTTP) APIs may be used.

[0052] Additionally, application gateway server computer 211 may receive notifications from backend systems 231 and provide web services to backend systems 231. As described below, application gateway server computer 211 may send notifications to managed containers (e.g., managed container 221) running on client devices 225.

[0053] In some embodiments, managed container 221 may be implemented as a special computer program with native managed container components 227 and managed container plugins 228 written in a programming language native to client device 225. Additionally, managed container 221 may include application framework 229 for running native components 227 and managed container plugins 228. As those skilled in the art will appreciate, application framework 229 may include an execution engine that provides a runtime environment and a set of class libraries that can be accessed at runtime. Application framework 229 may be implemented to suppose various types of client devices 225, including mobile devices, desktop computers, etc.

[0054] Managed container plugins 228 may be configured to extend the capabilities of managed container 221 to provide additional features to installed client applications. Specifically, managed container plugins 228 may include a variety of features and/or functions that leverage services 217 provided by application gateway server computer 211. Non-limiting examples of managed container plugins 228 may include a session management plugin, an optical character recognition plugin, a document management plugin, etc. To support these capabilities, native managed container components 227 may include an application cache for storing applications retrieved or received from application gateway server computer 211, a document cache for storing data associated with the applications, a user interface for providing particular user experience with managed container 221, and a hardware interface for interfacing with the hardware components of client device 225.

[0055] In some embodiments, services (e.g., services 217) provided by an application gateway server computer (e.g., application gateway server computer 211) may include one or more of the following:

[0056] an application service (e.g., application service 270) for communicating with managed containers operating on client devices and installing and managing applications on the client devices, the managing including updating, disabling, or deleting one or more of the applications;

[0057] a notification service (e.g., notification service 271) for selectively sending messages to one or more managed containers on one or more client devices, to a specific application or applications contained in the one or more of the managed containers, to one or more of the backend systems, or a combination thereof;

[0058] a settings service (e.g., settings service 272) for providing a storage mechanism for settings comprising application defaults, user preferences, and application

state information such that the settings are persisted at the application gateway server computer and consistent across the client devices;

[0059] a device management service (e.g., device management service 273) for communicating with the managed containers to enforce the set of rules independently of the application received from the application gateway server computer;

[0060] a user identity or profile service (e.g., profile service 274) for providing a common user identity (common authentication) across the backend systems connected to the application gateway server computer;

[0061] an enrollment service (e.g., registration service 275) for identifying a client device and registering the client device (for the purposes of tracking) with the application gateway server computer;

[0062] a proxy service (e.g., API proxy service 276) for communicating with one or more of the backend systems not explicitly supporting the application gateway architecture disclosed herein, or with external systems operating in another domain;

[0063] an authentication service (e.g., authentication service 277) for providing the managed container with a common authentication mechanism to the backend systems such that, once authenticated by the authentication service at the application gateway server computer, the managed container has access to the backend systems through the common authentication mechanism;

[0064] a media conversion service (e.g., media conversion service 278) for controlling content quality, size, format, watermarking, or a combination thereof such that the content is consumable by the client devices; and

[0065] a reporting service (e.g., reporting service 279) for aggregating data across backend systems and generating reports regarding same that can be viewed by an administrator or an end user.

[0066] Additional details of these services are provided below with reference to FIG. 6.

[0067] In some embodiments, an application gateway server computer may further include an application administration interface (also referred to as “admin UI”) configured for administration, deployment, and configuration of applications. A diagrammatic representation of an example embodiment of an application administration interface is depicted in FIG. 3. As illustrated in FIG. 3, through application administration interface 300, an administrator may, for instance, select an application from various applications 330, select an appropriate deployment option from a plurality of deployment options 350, and/or perform one or more associated actions 370. For example, an administrator may designate an application to be deployed to a selected group of users such as employees in a human resources department or may designate the application as mandatory for everyone in an organization. Additionally, the administrator may configure an application, obtain information on an application, delete an application, edit an application, or take other appropriate action via actions 370 (e.g., change or set permissions for one or more applications 330). In this way, application administration interface 300 may provide a centralized administration and configuration for applications 330.

[0068] In some embodiments, users may be allowed to download certain applications into managed containers. Through the application administration interface, an administrator can control which user-downloaded applications can remain in the managed container, which should be updated, and/or which should be removed. Thus, an administrator can let a particular user go directly to an app store and attempt to download certain applications. However, through the application gateway server computer operating in concert with the appropriate managed container(s) running on device(s) associated with the particular user, the administrator can still control which applications can actually be downloaded by the particular user.

[0069] An example of a method of operation according to some embodiments will now be described with reference to FIG. 4. In the example of FIG. 4, method 400 may comprise receiving, by an application gateway server computer, a request for an application from a managed container running on a client device or an instruction from an administrator using an application administration interface of the application gateway server computer to send (or “load”) an application to a managed container running on a client device (405). In response, the application gateway server computer may, for instance, obtain the application from a storage repository (e.g., application repository 119 shown in FIG. 1) and send the application to the managed container running on the client device (410). FIG. 5 shows an example of an application suitable for embodiments disclosed herein.

[0070] As shown in FIG. 5, an application sent from an application gateway server computer to a managed container may be realized via application package or file 500. In some embodiments, file 500 may be a compressed file (e.g., a zip file) and may contain application icon 505, manifest file 510, and application data 520. Application data 520 may also be compressed into file 515 to reduce file size for fast network transmission. Application icon 505 may comprise an image file containing an icon representing the particular application. Application data 520 may contain a configuration file and/or one or more code files associated with the particular application. The configuration file may include a short description of the particular application and one or more universal resource locator (URL) links to server side resources. At least one of the one or more code files may be configured for communicating with an operating system running on the client device. Specifically, in some embodiments, application data 520 may comprise user interface components of the particular application. However, other components of the particular application may not be included. In some embodiments, when the particular application is called (i.e., invoked by a user selecting application icon 505 from within a managed container running on a client device), the managed container may obtain, via its secure shell, any data, document(s), and/or file(s) that the particular application may need from a backend system through an application gateway server computer.

[0071] In some embodiments, an application may be coded or written in a markup language used for structuring and presenting content for the World Wide Web on the Internet. As a non-limiting example, the markup language may conform to the fifth revision of the HyperText Markup Language (HTML 5) standard. Those skilled in the art will appreciate that embodiments are not limited to HTML 5 content and can include any HTML, XML, text, etc. content as well. In embodiments where an application is written in

HTML/HTML 5, application data 520 may comprise associated HTML/HTML 5 application files.

[0072] Manifest file 510 may be a metadata file containing metadata about the particular application. One example of a manifest file is provided below:

```

{
  "name": "pulse",
  "displayName": "Content Server Pulse",
  "description": "Status and Comments ",
  "status": 1, "version": "8",
  "proxy_url": "https://intranet.company.com/cs/cs.dll"
  "local": true
}

```

[0073] Embodiments of client devices disclosed herein may operate on various platforms, including desktop, mobile, and web. In some embodiments, applications and components for desktop platforms are written in native binaries or HTML/HTML 5. In some embodiments, applications and components for mobile and web platforms are written in HTML/HTML 5. Accordingly, an application package such as file 500 may be implemented in various ways. For example, in some embodiments, an application package may include an icon, a metadata file, a configuration file, and at least one of a compressed file for a web platform, a compressed file for a mobile platform, or a compressed file for a desktop platform.

[0074] Referring to FIG. 4, the managed container running on the client device may provide a secure shell (SSH) for the application retrieve/received from the application gateway server computer (415). As those skilled in the art will appreciate, SSH refers to a cryptographic network protocol for secure data communication. SSH supports remote command execution, command-line login, and other secure network services. SSH enables the managed container to connect to the application gateway server computer via a secure channel over an insecure network. Once received, the managed container may extract the application (e.g., from file 500 shown in FIG. 5) and store the application and associated application data in its managed cache (420). As described above with reference to FIG. 5, at this point, the application stored in the managed cache of the managed container may comprise an icon, a manifest file, and some application data, sufficient for displaying the application within a user interface of the managed container, indicating to a user of the client device that the application is ready for use. In embodiments disclosed herein, the managed container controls the managed cache in accordance with a set of rules propagated from at least one of the backend systems to the managed container via the application gateway server computer (425). As discussed above, the set of rules or a portion thereof may reside on the client device, at the application gateway server computer, at the backend, or a combination thereof. In this way, when in use, access to the application can be controlled according to applicable rule(s) and, even when the application is not in use, the application and data stored in the managed cache can be protected by the managed container according to applicable rule(s).

[0075] Thus, embodiments of a managed container can be downloaded and installed on a client device and provides a secure runtime shell within which managed applications can be run on the client device. Specifically, the managed container can proactively retrieve or passively receive an

application (in the form of an application package, as explained below) from an application gateway server computer, extract the application locally and store the application and any data associated therewith (e.g., documents, etc.) locally in a managed cache that can be remotely managed/controlled by the application gateway server computer. Since content (which, in this example, includes the application and associated data) stored in the managed cache is encrypted/protected, the operating system running on the client device cannot open or view the managed content. In some embodiments, certain rules may be stored by the managed container on the client device, allowing the managed container to enforce the rules as needed whether the client device is online (connected to the application gateway server computer) or offline (disconnected from the application gateway server computer). One example rule may be to restrict access to certain content stored in the managed cache if the managed container is unable to communicate with the application gateway server computer.

[0076] As noted above, a user may still download and install applications on their devices as usual, allowing an application to connect with a backend system directly and bypassing embodiments of an application gateway server computer disclosed herein. However, such a client-server relationship (between an application installed on a user device and a backend system running on a server machine) has many drawbacks. For example, since such an application is not under the control of a managed container, the backend system may not retain control over the application, particularly when the user device may be offline. Furthermore, since data associated with the application is not stored in a managed cache, it is not protected by the managed container and thus may be vulnerable to unauthorized access by other applications and/or the operating system running on the user device. What is more, because the application communicates with the backend system directly and not through an application gateway server computer, it may not enjoy the many services provided by the application gateway server computer, including the authentication service. Accordingly, for each application that is not handled through the application gateway server computer, the user would need to authenticate with various backend systems separately.

[0077] An example of an application gateway architecture in operation will now be described with reference to FIG. 6. In this example, application gateway architecture 600 may include application gateway server computer 611 communicatively connected to a plurality of backend systems 631 through firewall 660 and a plurality of client devices 625. Additionally, application gateway server computer 611 may be communicatively connected to various storage devices at data storage layer 695, including application repository 619 and data store 629.

[0078] As described above, in some embodiments, a managed container may be implemented as an application (program) that is native to a client device, that can be downloaded from a source such as a website or an app store, and that can run on a client device separate and independent of any applications that it manages. In the example of FIG. 6, a user can download managed container 621 from various sources such as third-party source 650 (e.g., an online store on a public network such as the Internet) or enterprise app store 680 (e.g., a proprietary store on a private network such as a company's intranet) at network layer 691. Once installed, an icon associated with managed container 621 is

displayed on client device 625, as shown in FIG. 6. Additional details of an example of a managed container are provided below with reference to FIG. 7.

[0079] As described above, managed container 621 can extend the capabilities of client device 625 by leveraging services 617 provided by application gateway server computer 611 at application layer 693. In some embodiments, web application archive (WAR) files may be used to package/install services 617 on application gateway server computer 611. Services 617 may vary from implementation to implementation, depending upon availability of backend systems 631. Non-limiting examples of services 617 provided by application gateway server computer 611 may include authentication service 671 for providing managed container 621 with common authentication mechanism 697 across backend systems 631 and requests for services 617 provided by application gateway server computer 611. Note that services 617 provided by application gateway server computer 611 are "reusable" and "common" to applications managed by managed container 621 in that services 617 can be leveraged by any of such applications. Once a user of client device 625 is authenticated by authentication service 671 (e.g., for an application running in managed container 621), managed container 621 has access to backend systems 631 and there is no need for the user to authenticate for other applications on client device 625 to access backend systems 631.

[0080] To authenticate, managed container 621 operating on client device 625 may first identify and authenticate itself in a connection request sent to application gateway server computer 611. In response, application gateway server computer 611 (via notification service 674) may send out a notification to managed container 621 using a notification listener or a push notification channel already established on client device 625. Push notifications and acknowledgement mechanisms are known to those skilled in the art and thus are not further described herein. Managed container 621 must be able to receive the notification from notification service 674 and respond accordingly. To be able to do so, managed container 621 must be registered with application gateway server computer 611.

[0081] In some embodiments, registration or enrollment service 672 may be provided for registering and identifying (for the purposes of tracking) a client device. Specifically, service 672 may provide common registration services to track connected client devices, track and manage client devices to enable remote wipe, block authentication for lost devices, notify a backend system on connection of a new client or device, provide a broadcast point for notification services 674, etc.

[0082] Accordingly, application gateway server computer 611 may enroll various managed containers using registration service 672 and place a certificate on an enrolled managed container. A registered/enrolled client device must report all fields required by an administrator (e.g., location, jailbreak status, device ID, etc.), implement a notification listener to receive messages from notification service 674, respond to notification messages with an acknowledgement when required (e.g., delete, remote wipe, hold, permission changes, etc.), and display and run all applications deployed from application gateway server computer 611. Jailbreak refers to removing limitations on certain highly controlled devices, such as the iPhone, manufactured by Apple Computer, Inc. of Cupertino, Calif., so that root access can be

gained to allow download of additional applications, extensions, and themes not officially available on the devices. Registration or enrollment service 672 may be integrated with reporting service 676 or it may be implemented separately. This certificate or token, which is issued by application gateway server computer 611, is in addition to the certificate(s) or token(s) issued by backend system(s) 631 such as a content server that a managed container is attempting to connect. To authenticate, therefore, managed container 621 would need to provide two certificates, tokens, or the like in order to connect through application gateway server computer 611 to backend system 631.

[0083] In some embodiments, a device management service may be implemented in conjunction with registration service 672. The device management service may communicate with managed container 621 to enforce a set of rules independently of any application received from application gateway server computer 611. Specifically, the device management service may communicate with client device 621 to ensure that it is registered with application gateway server computer 611 and that it obeys the rules. The device management service allows specific instructions such as a remote wipe command to be sent to a specific client device (e.g., using the media access control address or MAC address of the client device). The device management service may perform acknowledgement verification (e.g., via an acknowledgement channel) to determine if instructions were indeed received and/or performed by a client device.

[0084] Registration data (including the identification and credentials) associated with each registered managed container may be stored in data store 629. Data store 629 may comprise a central database storing configuration data used by application gateway server computer 611. Data store 629 may be managed using admin UI 630. Admin UI 630 may implement an embodiment of application administration interface 300 described above with reference to FIG. 3. Data store 629 may further provide storage for services 617, including a server side persistent storage (e.g., a settings repository) for settings received via settings service 677. The settings repository may store and maintain client configuration and state data, provide an extensible state framework for application developers, and enable application state to be persisted across devices and clients.

[0085] To respond to the notification and to authenticate with application gateway server computer 611, managed container 621 must send an acknowledgement or an acknowledgement with additional information (e.g., meta-data that match the registration data associated with managed container 621 stored in data store 629). If managed container 621 does not acknowledge the notification, then application gateway server computer 611 will not allow managed container 621 to connect to backend system 631 using their API.

[0086] As those skilled in the art will appreciate, the above example is one of many types of notifications that may be provided by notification service 674. Indeed, using notification service 674, messages may be selectively sent from application gateway server computer 11 to appropriate managed containers operating on various client devices, to a specific application or applications contained in the managed container(s), to a particular backend system or systems, or a combination thereof. For example, notification service 674 can be used to let a managed container know that a certain application is not allowed on a particular client

device and needs to be deleted from the managed container. As another example, notification service 674 can be used to send a message to a managed container indicating that files older than ten days are to be deleted but does not specify which files are older than ten days. The managed container, in response to the message, can determine which files under its management are older than ten days and delete them from its managed cache.

[0087] Once managed container 621 is authenticated, on the one hand, application gateway server computer 611 can notify the now authenticated managed container 621 about what applications should reside on the client device, what new applications are available for download, etc., in addition to managing connections to backend systems 631. On the other hand, managed container 621 can download an application (e.g., in the form of an application package or file as described above with reference to FIG. 5) from application repository 619, extract the content (the application and associated data) into a local file system on its managed cache, encrypt the extracted content, store the encrypted content in the secure managed cache, and drop an icon so that a reference to the application shows up in a user interface of managed container 625. As explained above with reference to FIG. 3, the deployment of an application may be made mandatory, selective, or optional by an administrator. Since the application is downloaded from application repository 619 into the managed cache, an update to the application published from application repository 619 can be downloaded to managed container 621 when the update is installed (e.g., by an administrator). Furthermore, the application can be removed from the managed cache when it is deleted or otherwise removed from application repository 619.

[0088] When the icon for the application is selected or otherwise invoked from within managed container 621, the application opens up a user interface on client device 625 and makes an API call (e.g., a JSON API call) to application gateway server computer 611 (using an appropriate certificate or token issued by application gateway server computer 611) to access backend system 631 (with an appropriate certificate or token issued by backend system 631). In some cases, API calls may be handled by API proxy service 675. In some embodiments, API proxy service 675 can be configured for communicating with (e.g., making JSON and/or HTTP API calls to) any backend system that does not explicitly support application gateway 611. In some embodiments, API proxy service 675 can be configured for communicating with external systems on another domain. API proxy service 675 may maintain a list of allowed and/or disallowed third party services (e.g., in data store 629).

[0089] In some embodiments, settings service 677 may be configured for providing a storage mechanism for settings comprising application defaults, user preferences (e.g., favorite documents, color scheme, etc.), and application state information such that these settings can be persisted (e.g., in data store 629) at the server side and consistent across the client devices and/or managed applications.

[0090] In some embodiments, content (e.g., documents, video files, etc.) from backend systems 631 may be processed at application gateway server computer 611, for instance, using media conversion service 673. In some embodiments, media conversion service 673 may be configured for controlling content quality, size, format, watermarking, or a combination thereof such that the content is

consumable by particular client devices and/or per specific user preferences stored in data store 629. In some embodiments, media conversion service 673 may convert various types of content. For example, media conversion service 673 may convert a word processing document to a portable document format (PDF) to prevent changes and also watermark the PDF document. As another example, media conversion service 673 may be configured to produce only low resolution images, etc.

[0091] Note that, even if a user can open an application from within a managed container, they cannot do anything unless they have the appropriate certificate or token issued by the application gateway server computer. All the API calls that come in from client devices for connections through the application gateway server computer to the backend systems are handled by the application gateway server computer.

[0092] In an offline/disconnected scenario, the application may be opened and used (assuming such use(s) is/are allowed according to admin rules propagated from the application gateway server computer). There would still be an authentication check, but it would be based on the last set of cached credentials from the last valid authentication/connection with the application gateway server computer, due to the disconnected nature.

[0093] As described above, client devices may operate on various platforms, including desktop, mobile, and web. FIGS. 7 and 8 depict diagrammatic representations of managed containers operating on different types of client device according to some embodiments.

[0094] Specifically, FIG. 7 depicts application gateway architecture 700 including application gateway server computer 711 communicatively connected to a plurality of backend systems 731 through firewall 760 and a plurality of client devices 725. Application gateway server computer 711 may implement an embodiment of an application gateway server computer described above. Client device 725 may implement an embodiment of a managed container described above.

[0095] As shown in FIG. 7, client device may include JavaScript to Native 2-Way Bridge 750 (“bridge 750”) and device hardware 780. Device hardware 780 may include various hardware components commonly found on a mobile device such as a smart phone or a handheld or portable computing device. Bridge 750 may provide a common JavaScript API for interfacing the JavaScript side (e.g., applications 722, services provided by application gateway server computer 711, etc.) with the native (common) code in which a managed container is implemented (“managed container 721”). Bridge 750 may also interface with native device capabilities, including camera, native viewer, e-mail, contacts, GPS, etc. As explained above, applications 722 may be retrieved and/or received from application gateway server computer 711 and may not be available from commercial app stores on the Internet.

[0096] Acting as a native shell for applications 722 downloaded to client device 725, managed container 721 has knowledge (e.g., via managed file system 723) of where contents (applications 722 and data associated therewith) are stored in managed cache 724 and their corresponding settings in settings repository 729, can display a download progress bar on client device 725 via managed container user interface 730 (which includes common UI components in the native code), and can receive notifications 725 in the

background and take appropriate action accordingly. For example, if an administrator wishes to restrict access to application 722 downloaded onto client device 725, notification 725 to remove application 722 can be sent to managed container 725, as described above, and managed container 721 will respond to notification 725 and delete application 722 from managed cache 724. All related metadata and applicable cached content will be deleted as well. Correspondingly, the icon for application 722 will disappear from user interface 730 of the managed container.

[0097] In some embodiments, user interface 730 may include reusable UI components that can be leveraged by any installed applications 722. Presentation of these reusable UI components determined by managed container 721 may provide native or intuitive user experiences. Non-limiting examples of reusable UI components for user interface 730 may include a browse list with a paging capability, form input controls, a server browse function, a user profiles and lookup function, a document viewer, etc.

[0098] As described above, managed cache 724 may include an application cache and a document cache. The document cache is not limited to storing documents and may also contain other file types such as videos, photos, and so on. The application cache can be used by managed container 721 to manage applications on the client device and communicate with the application gateway server computer to access and/or update applications 722. The application cache may be configured with security 727 such that if managed container 721 is unable to communicate with the application gateway server computer, certain applications 722 stored in the application cache cannot be accessed. Another way to control access to applications 722 stored in the application cache may be done via device hardware 780. For example, managed container 721 may be configured to communicate with device hardware 780 to determine whether client device 725 has been “jail-broken” and, if so, operate to restrict access to certain applications and/or documents stored in managed cache 724.

[0099] The document cache can be used by managed container 721 to manage documents in a hierarchical manner (via managed file system 723) and control access to selected documents stored in the document cache. For example, when a user desires to open a document, an application associated with the document may call managed container 721. Managed container 721 may operate to determine, via security 727, whether the user has the proper permission to open the document and check to see if there is a version of the document already in managed cache 724. If so, managed container 721 may give that document to the application and make no calls to application gateway server computer 711. In some embodiments, managed container 721 may call application gateway server computer 711 with the appropriate certificates or tokens to connect through application gateway server computer 711 to backend system 731 to get an updated version of the document.

[0100] FIG. 8 depicts a diagrammatic representation of an example of a managed container operating on another type of client device according to some embodiments. Specifically, FIG. 8 depicts application gateway architecture 800 including application gateway server computer 811 communicatively connected to a plurality of backend systems (e.g., via a firewall such as firewall 760 shown in FIG. 7) and a plurality of desktop client devices. Examples of desktop client devices may include work stations, desktop comput-

ers, etc. Application gateway server computer **811** may implement an embodiment of an application gateway server computer described above and provide desktop services **817** to desktop clients connected thereto. Desktop services **817** may include core services and/or computer program product services similar to those described above with reference to services **117**, services **217**, and services **617**. Desktop managed container **721** may implement an embodiment of a managed container described above.

[0101] In some embodiments, each desktop client device may include library **850** that may act as an interface or bridge from the JavaScript side (applications **822**) into the native (common) code in which a desktop managed container is implemented (“desktop managed container **821**”). As explained above, applications **822** may be retrieved and/or received from application gateway server computer **811** and may not be available from commercial app stores on the Internet.

[0102] In some embodiments, desktop managed container **821** may be required for a desktop client application of a backend system running in an enterprise computing environment. Accordingly, when the desktop client application installs, it may detect whether desktop managed container **821** exists. If it does, desktop managed container **821** may be used to update and install services and components that can be used by the new desktop client application. If it does not, desktop managed container **821** may be installed first and then used to install services and components that can be used by the new desktop client application. In some embodiments, desktop client applications running in an enterprise computing environment may re-use services and components from desktop managed container **821** (e.g., via JavaScript to Native bridge **850**) and wrap/create their own experiences.

[0103] In some embodiments, desktop managed container **821** may be configured to provide a plurality of features and/or functions, including configuration push, component and service updates, application distribution and configuration, cache management and policy enforcement, state synchronization with other platforms, etc. In this way, desktop managed container **821** can provide a common integration point and a common user experience focal point in the desktop computing environment. Non-limiting example features and/or functions of desktop managed container **821** may include system tray **861**, local configuration **863**, server notifications **865**, file system overlays **867**, authentication **881**, file system monitor **883**, file transfer **885**, content cache **841**, settings cache **843**, state cache **845**, and database **826**.

[0104] System tray **861** may include a common tray icon for desktop managed container **821** that is used across all desktop client devices. System tray **861** may further include a common menu where authorized users and/or applications (including desktop client applications associated with backend systems) can place specific menu entries. Furthermore, system tray **861** may include a launch point for all the installed applications.

[0105] Local configuration **863** may include local configuration settings for desktop managed container **821** and associated users. Authorized users can extend and add additional configuration settings as needed.

[0106] Server notifications **865** may be configured to route notifications from a push notification service to appropriate application(s) (whether managed by desktop managed con-

tainer **821** or not). Additionally, server notifications **865** may be configured to display the notifications.

[0107] File system overlays **867** may be configured to provide common file system services for adding icon badges to file system objects managed by desktop managed container **821**.

[0108] Authentication **881** may include a single set of credentials and an URL for each desktop client device operating in the enterprise computing environment. Desktop managed container **821** may look up the URL based on a site name registered with application gateway server computer **811**. Application gateway server computer **811** may register a desktop client device via a registration service similar to registration service **672** described above with reference to FIG. 6.

[0109] File system monitor **883** may include a registered driver for applications managed via desktop managed container **821**. Specifically, an API may register as a watcher and may register file system locations to watch.

[0110] File transfer **885** may be configured to perform file upload/download that supports HTTP range headers. Additionally, file transfer **885** may be configured to allow for resumable transfers via application gateway server computer **811**.

[0111] Content cache **841**, settings cache **843**, and state cache **845** may collectively be referred to as a managed cache. Similar to database **126** described above with reference to FIG. 1, database **826** may be used by desktop managed container **821** to keep track of content stored in the managed cache. Additionally, similar to embodiments of a managed cache described above, content cache **841** may be configured for storing content (including desktop client applications associated with backend systems), allowing the backend systems to retain control of downloaded content and enforce applicable rules in an effective and efficient manner. To this end, settings cache **843** may store application settings and configuration and state cache **845** may provide a client side API to allow the applications to save their state and sync with application gateway server computer **811**. This allows application gateway server computer **811** to replicate the state across platforms and devices. For example, a user works on a document using a first version of an application managed by a first managed container on a first device. The state of the first version of the application is saved and sync with an application gateway server computer. The user may wish to work on the document using a second version of the application managed by a second managed container on a second device. The application gateway server computer may communicate the state information to the second managed container. When the second version of the application is opened, the second managed container may replicate the state of the application and the user can work on the document where it was left off, using the second version of the application managed by the second managed container on the second device.

[0112] Cloud-Based Architecture and Database Sharding

[0113] FIG. 9 depicts a diagrammatic representation of an example of a cloud based application gateway architecture that may employ managed containers according to some embodiments. In the example of FIG. 9, system **900** may include cloud-based gateway services **910** communicatively connected to backend systems **931** and one or more client devices **925**. Client device **925** shown in FIG. 9 is representative of various client devices. Those skilled in the art

will appreciate that FIG. 9 shows a non-limiting example of client device 925. Backend systems 931 may comprise computer program products and/or applications developed within a company and/or by third party developers/companies. Non-limiting examples of backend systems 931 may include a content server, an information management system, a document repository, a process management system, a social server, an RM system, a database management system, an enterprise resources planning system, a collaboration and management system, a customer relationship management system, a search system, an asset management system, a case management system, etc. Embodiments as shown in FIG. 9 of the cloud-based gateway may include APIs and services configured for communicating with backend systems 931 and managed containers 921 operating on client devices 925, in a manner generally similar to that discussed above.

[0114] The architecture of FIG. 9 is illustrated with more particularity with reference to FIG. 10. Specifically, architecture 1000 implements many features that can provide increased throughput and scalability. For example, architecture 1000 includes scalable cloud-based or hosted gateway server 1004 which, in some embodiments, is implemented as a Node.js application server. As those skilled in the art can appreciate, Node.js provides an event-driven architecture and can maintain a large number of connections without having to reject new incoming connections. This feature allows architecture 1000 to scale up massively by adding new “nodes” and also allows gateway server node 1004 to effortlessly handle real time applications. Some of such applications may include lightweight applications for certain functions normally provided by backend systems such as a portal or a social tool (e.g., a blogging tool) for a content server.

[0115] In some embodiments, architecture 1000 may be a multi-tenant architecture where multiple tenants can share the same application running on the same operating system on the same hardware using the same data storage mechanism. Every tenant (e.g., a group of users) is provided with a share of a software instance and/or resource. However, tenants do not share and cannot view each other’s data. In some embodiments, tenants and their contents may be stored in database shards and filtered using tenant identifiers. Embodiments of database sharding are further described below. This multi-tenant feature allows a large number of users to use the applications and/or services provided by or through gateway server node 1004 and further facilitates the scalability of architecture 1000.

[0116] In some embodiments, gateway server node 1004 may coordinate with reverse proxy server 1002, and may be in communication with a plurality of backend systems 1006-1014. In the example embodiment illustrated, such systems may include utilities 1006 such as text extraction, thumbnail generation, document conversion, and identity management; search system 1008; queue manager 1010; database manager 1012; and storage appliance 1014. Similar to the common authentication described above, a user may only need to authenticate once for all the applications delivered through architecture 1000 to the user’s managed container.

[0117] Reverse proxy server 102 may provide a plurality of functions including load balancing, web content server, and etc. In some embodiments, reverse proxy server 102 may be implemented as an Nginx server, with a focus on

high concurrency, high performance, and low memory usage. Those skilled in the art can appreciate that an Nginx reverse proxy server can be configured for HTTP, HTTPS, SMTP, POP3, and IMAP protocols and can act as a load balancer, HTTP cache, and a web server (origin server). Additionally, reverse proxy server 102 may provide an administration layer for administering new instances of nodes, further increasing the scalability of architecture 1000. [0118] Gateway server node 1004 may implement a web application framework for REST API request handlers. In addition, gateway server node 1004 may create and manage background tasks and distribute jobs to the appropriate backend engine(s) for processing. Furthermore, gateway server node 1004 may provide a framework for defining logical models and mapping them to a database’s physical model.

[0119] As noted above, embodiments provide an improved system and method for database sharding. As those skilled in the art can appreciate, database sharding is complex. In database sharding, a single large database is fragmented, or “sharded” into multiple smaller databases that operate virtually independently. Shards can be located on separate database servers or physical locations. Database sharding allows scaling near-linearly to hundreds or thousands of database clusters. Collectively, these shards form a single, very large database. In this way, database sharding can scale databases beyond what a single cluster of database servers can handle. This scalability comes with a few limitations, however. For example, cross-shard queries are not allowed, and APIs must be able to join in data that would span multiple shards (usually users).

[0120] A system according to embodiments can take away much of that complexity, making shards appear to an API as a single database. In some embodiments, this can be done by generating IDs that are unique across a cluster of shards; routing queries to the correct shards; and maintaining split relationships across shards.

[0121] To generate unique IDs, each shard is tagged with a shard ID on initialization. In some embodiments, a shard ID can range from 0-8191. This shard ID is encoded into every ID generated by a shard. Some embodiments employ 41 bits for timestamp; 13 bits for shard ID; and 10 bits for uniqueness. In some embodiments, each shard can generate object IDs completely independently.

[0122] An example of a Full ID is: FFFF FFFF FFFF FFFF

[0123] An example of the Timestamp mask is: FFFF FFFF FFE0 0000

[0124] An example of a Shard ID mask is: 0000 0000 001F FC00

[0125] An example of a Uniqueness mask is: 0000 0000 0000 03FF

[0126] An example of querying the correct shard in a multi-shard scenario follows.

[0127] First, suppose a query for an object by ID is: {where: {id: 2047} }

[0128] To perform the query, the shard ID is extracted as follows:

[0129] 0000 0000 0000 07FF (the hex value for the example full ID: 2047)

[0130] && 0000 0000 001F FC00 (the example shard ID mask)

[0131] =0000 0000 0000 0400 (compare this value with the example uniqueness mask)

[0132] 10>>0000 0000 0000 0001 (the remainder indicates the shard ID)

[0133] In this case, the object is found in shard 1 and the query is routed to that shard.

[0134] Thus, upon receiving a query for an object, the system determines which database server to query, and also handles relationships between objects that might be in completely separate shards. If a complex query is received, an API can provide a ‘context’ in which to route the query. Each repository may be routed differently, depending upon the type of objects stored therein. For example, user objects and tenant objects may be routed round-robin on creation and by ID on query. On creation, resources (e.g., docs and folders) may be routed first by their parent ID, and then by the context user’s ID. Versions may be routed based on their resource’s ID. In some embodiments, if a query is received with no ID and no context, it may be routed to a specific shard. In this way, the system can guarantee unique IDs for objects without creating a single point of failure.

[0135] As those skilled in the art can appreciate, resources in database systems can be locked using a synchronization mechanism (referred to as a resource lock) to enforce limits on access to a resource (e.g., a database record) in an environment where there are many threads of execution. Such a resource lock can enforce a mutual exclusion concurrency control policy, for instance.

[0136] FIGS. 11A and 11B illustrate how resource locking is handled in a standard (single) database system. In FIG. 11A, a data structure such as a list or a function table “ResourceLock” may include a ResourceId entry associated with a primary key (PK) and a foreign key (FK). Likewise, the locked-by party, i.e., the party for whom the resource is locked, is identified by “LockedById” and associated with a PK and a FK. In FIG. 11B, a “share” table stores an identification (Id) as a PK and also stores a plurality of FKs including an identification of the resource shared (ResourceId), an identification of the party performing the sharing (SharedById), and an identification of the party to whom the object is shared (SharedToId).

[0137] In conventional database systems, these locks are managed in memory and therefore consume memory resources. As the number of locks increase, so do the memory resources required to store and track these locks. This can significantly increase operation cost and reduce system performance.

[0138] These issues are even more challenging to address in multi-shard environments. For example, in a system that allows a user to “like,” “comment,” or “follow” a document, a first user may like a second user’s document. In such a case, the system needs to know which documents the first user likes and which user(s) like the second user’s documents. However, if this relationship (between the first user and the second user’s document) is stored with the document, a search to find what documents user 1 likes will involve searching all shards, which can be time consuming and computationally expensive. Likewise, if the relationship is stored with the user (the first user in this example), a search to find who likes this particular document will involve searching all shard.

[0139] Embodiments can handle relationships across shards in a significantly more efficient way. Specifically, embodiments of a sharding manager can generate, manage, and maintain a relationship (e.g., a resource lock) by splitting the relationship into two or more functional tables and

storing a copy of each table in each shard. Following the above example, the user item associated with the first user on one shard and the document item associated with the second user on another shard are considered to be in a relationship (the resource lock) and the relationship can be processed (e.g., searched) from either the perspective of the user item (a first split relationship) or the perspective of the document item (a second split relationship). Together these two split relationships represent one complete relationship between the user item and the document item.

[0140] This is further illustrated in FIGS. 12A and 12B. In FIG. 12A, a data structure such as a table referred to as “ResourceLock” 1202 has the same column and same data as the resource lock shown in FIG. 11A. However, instead of having the “ResourceId” and the “LockedById” be a foreign key (FK) to both the resource (e.g., a document) and the user resource on the same table as shown in FIG. 11A, ResourceLock 1202 in FIG. 12A has first split function table 1204 representing a first split relationship from a first perspective (e.g., a user item associated with the first user in the above example) and second split function table 1206 representing a second split relationship from a second perspective (a document item associated with the second user in the above example). Specifically, split table 1204 “ResourceLock_LockedBy” has “ResourceId” and “LockedById” but only the “LockedById” is the foreign key to the user table. That is, “ResourceLockLockedBy” (split table 1204) enforces the “Locked_By” half of the relationship. Split table “ResourceLock_Resource” 1206 also has “ResourceId” and “LockedById,” but its foreign key is on the “ResourceId” which is linked to the resource table. The “LockedById” is still part of the PK, but it is not the FK because the user could appear in a different shard. These two tables together represent a complete relationship between the two items in the above example. A copy of these tables is stored in each shard. The data in them represents what is local to that shard.

[0141] For example, if there is a lock on a resource, there will be an entry of a “ResourceLock_Resource” table (e.g., table 1206) on that shard. For a user obtaining that lock, there will be a “ResourceLockLockedBy” table (table 1204) in that user’s shard. If someone wants to query for the user to see all the resources that are locked by the user, the system can query for that user and all of their locks will appear in the “ResourceLockLockedBy” table on their shard. From that query result, the system can get the resource ID’s. Because those resource ID’s are using the same sharding algorithm described herein, the system can extract the shard ID (since, as described above, the shard ID mask is included in each object’s ID) and use that shard ID to query the correct database to get the resource that the user has locked. This query routing mechanism allows the system to navigate from one shard to another. As such, the generation and maintenance of these relationships can be hidden from the API layer, thereby eliminating the need for APIs to join in data that would span multiple shards.

[0142] Embodiments can handle relationships involving multiple database items in a similar manner. FIG. 12B illustrates by example how embodiments disclosed herein may handle three-database-item relationships. In particular, as shown, a share relationship is split into three parts, one associated with each shard storing an item involved in the share operation. In the embodiment illustrated, the share relationship includes a Share_SharedBy relationship in

which the SharedById is a foreign key (table 1204); a Share_SharedTo relationship, in which the SharedToId is a foreign key (table 1206); and a Share_Resource relationship in which the ResourceId is a foreign key (table 1208).

[0143] FIG. 13A schematically illustrates an example of database sharding in accordance with some embodiments. As shown, gateway server node 1300 includes or is in communication with sharding manager 1302 (which may be part of a database manager) which, in turn, communicates with database(s) 1305 via one or more APIs 1304. In the non-limiting embodiment illustrated, database 1305 includes one or more database servers 1307a, 1307b, . . . , 1307n, which maintain one or more shards 1306a, 1306b, . . . , 1306n. As discussed above, embodiments can maintain relationships across shards. For example, embodiments can store relationships between items on different shards, such as items on shards 1306a and 1306b. To do so, embodiments store split relationship 1308a, 1308b on each shard. Each shard that includes the item in the relationship maintains a corresponding portion of the relationship. In the embodiment illustrated in FIG. 13A, the resource lock relationships are shown. Thus, shard 1306a includes copy 1308a of the ResourceLock_LockedBy portion of the relationship (a first split of the relationship), while shard 1306b includes copy 1308b of the ResourceLock_Resource portion of the relationship (a second split of the relationship).

[0144] Similarly, FIG. 13B illustrates a further embodiment of a share relationship between three database items stored in different shards. A first item is stored in shard 1306, a second item is stored in shard 1306b, and a third item is stored in shard 1306c. Each item is identified by their unique ID which is generated by the respective shard and which is encoded with the respective shard ID mask as explained above. In the embodiment illustrated, the first item has a lock on a resource which is shared by the second item and which is shared to the third item. These three items, therefore, are in a relationship relative to the same resource identified by the resource ID. Accordingly, shard 1306a stores a copy 1310a of the Share_Resource portion of the relationship (a first split of the relationship), shard 1306b stores a copy 1310b of the Share_SharedBy portion of the relationship (a second split of the relationship), and shard 1306c stores a copy 1310c of the Share_SharedTo portion of the relationship (a third split of the relationship). At query time, the system can navigate from one shard to another using the query routing mechanism described above.

[0145] Embodiments disclosed here in can be implemented in various ways. For example, in some embodiments, components of an application gateway architecture described above can be deployed on premises, on premises as a virtual machine, or in a cloud computing environment (including entirely or partially hosted in the cloud). Other implementations may also be possible.

[0146] Although the invention has been described with respect to specific embodiments thereof, these embodiments are merely illustrative, and not restrictive of the invention. The description herein of illustrated embodiments of the invention, including the description in the Abstract and Summary, is not intended to be exhaustive or to limit the invention to the precise forms disclosed herein (and in particular, the inclusion of any particular embodiment, feature or function within the Abstract or Summary is not intended to limit the scope of the invention to such embodiment, feature or function). Rather, the description is

intended to describe illustrative embodiments, features and functions in order to provide a person of ordinary skill in the art context to understand the invention without limiting the invention to any particularly described embodiment, feature or function, including any such embodiment feature or function described in the Abstract or Summary. While specific embodiments of, and examples for, the invention are described herein for illustrative purposes only, various equivalent modifications are possible within the spirit and scope of the invention, as those skilled in the relevant art will recognize and appreciate. As indicated, these modifications may be made to the invention in light of the foregoing description of illustrated embodiments of the invention and are to be included within the spirit and scope of the invention. Thus, while the invention has been described herein with reference to particular embodiments thereof, a latitude of modification, various changes and substitutions are intended in the foregoing disclosures, and it will be appreciated that in some instances some features of embodiments of the invention will be employed without a corresponding use of other features without departing from the scope and spirit of the invention as set forth. Therefore, many modifications may be made to adapt a particular situation or material to the essential scope and spirit of the invention.

[0147] Reference throughout this specification to “one embodiment”, “an embodiment”, or “a specific embodiment” or similar terminology means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment and may not necessarily be present in all embodiments. Thus, respective appearances of the phrases “in one embodiment”, “in an embodiment”, or “in a specific embodiment” or similar terminology in various places throughout this specification are not necessarily referring to the same embodiment. Furthermore, the particular features, structures, or characteristics of any particular embodiment may be combined in any suitable manner with one or more other embodiments. It is to be understood that other variations and modifications of the embodiments described and illustrated herein are possible in light of the teachings herein and are to be considered as part of the spirit and scope of the invention.

[0148] In the description herein, numerous specific details are provided, such as examples of components and/or methods, to provide a thorough understanding of embodiments of the invention. One skilled in the relevant art will recognize, however, that an embodiment may be able to be practiced without one or more of the specific details, or with other apparatus, systems, assemblies, methods, components, materials, parts, and/or the like. In other instances, well-known structures, components, systems, materials, or operations are not specifically shown or described in detail to avoid obscuring aspects of embodiments of the invention. While the invention may be illustrated by using a particular embodiment, this is not and does not limit the invention to any particular embodiment and a person of ordinary skill in the art will recognize that additional embodiments are readily understandable and are a part of this invention.

[0149] Embodiments discussed herein can be implemented in a computer communicatively coupled to a network (for example, the Internet), another computer, or in a standalone computer. As is known to those skilled in the art, a suitable computer can include a central processing unit (“CPU”), at least one read-only memory (“ROM”), at least one random access memory (“RAM”), at least one hard

drive (“HD”), and one or more input/output (“I/O”) device (s). The I/O devices can include a keyboard, monitor, printer, electronic pointing device (for example, mouse, trackball, stylus, touch pad, etc.), or the like.

[0150] ROM, RAM, and HD are computer memories for storing computer-executable instructions executable by the CPU or capable of being compiled or interpreted to be executable by the CPU. Suitable computer-executable instructions may reside on a computer readable medium (e.g., ROM, RAM, and/or HD), hardware circuitry or the like, or any combination thereof. Within this disclosure, the term “computer readable medium” is not limited to ROM, RAM, and HD and can include any type of data storage medium that can be read by a processor. For example, a computer-readable medium may refer to a data cartridge, a data backup magnetic tape, a floppy diskette, a flash memory drive, an optical data storage drive, a CD-ROM, ROM, RAM, HD, or the like. The processes described herein may be implemented in suitable computer-executable instructions that may reside on a computer readable medium (for example, a disk, CD-ROM, a memory, etc.). Alternatively, the computer-executable instructions may be stored as software code components on a direct access storage device array, magnetic tape, floppy diskette, optical storage device, or other appropriate computer-readable medium or storage device.

[0151] Any suitable programming language can be used to implement the routines, methods or programs of embodiments of the invention described herein, including C, C++, Java, JavaScript, HTML, or any other programming or scripting code, etc. Other software/hardware/network architectures may be used. For example, the functions of the disclosed embodiments may be implemented on one computer or shared/distributed among two or more computers in or across a network. Communications between computers implementing embodiments can be accomplished using any electronic, optical, radio frequency signals, or other suitable methods and tools of communication in compliance with known network protocols.

[0152] Different programming techniques can be employed such as procedural or object oriented. Any particular routine can execute on a single computer processing device or multiple computer processing devices, a single computer processor or multiple computer processors. Data may be stored in a single storage medium or distributed through multiple storage mediums, and may reside in a single database or multiple databases (or other data storage techniques). Although the steps, operations, or computations may be presented in a specific order, this order may be changed in different embodiments. In some embodiments, to the extent multiple steps are shown as sequential in this specification, some combination of such steps in alternative embodiments may be performed at the same time. The sequence of operations described herein can be interrupted, suspended, or otherwise controlled by another process, such as an operating system, kernel, etc. The routines can operate in an operating system environment or as stand-alone routines. Functions, routines, methods, steps and operations described herein can be performed in hardware, software, firmware or any combination thereof.

[0153] Embodiments described herein can be implemented in the form of control logic in software or hardware or a combination of both. The control logic may be stored in an information storage medium, such as a computer-read-

able medium, as a plurality of instructions adapted to direct an information processing device to perform a set of steps disclosed in the various embodiments. Based on the disclosure and teachings provided herein, a person of ordinary skill in the art will appreciate other ways and/or methods to implement the invention.

[0154] It is also within the spirit and scope of the invention to implement in software programming or code an of the steps, operations, methods, routines or portions thereof described herein, where such software programming or code can be stored in a computer-readable medium and can be operated on by a processor to permit a computer to perform any of the steps, operations, methods, routines or portions thereof described herein. The invention may be implemented by using software programming or code in one or more digital computers, by using application specific integrated circuits, programmable logic devices, field programmable gate arrays, optical, chemical, biological, quantum or nanoengineered systems, components and mechanisms may be used. In general, the functions of the invention can be achieved by any means as is known in the art. For example, distributed, or networked systems, components and circuits can be used. In another example, communication or transfer (or otherwise moving from one place to another) of data may be wired, wireless, or by any other means.

[0155] A “computer-readable medium” may be any medium that can contain, store, communicate, propagate, or transport the program for use by or in connection with the instruction execution system, apparatus, system or device. The computer readable medium can be, by way of example only but not by limitation, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, system, device, propagation medium, or computer memory. Such computer-readable medium shall be machine readable and include software programming or code that can be human readable (e.g., source code) or machine readable (e.g., object code). Examples of non-transitory computer-readable media can include random access memories, read-only memories, hard drives, data cartridges, magnetic tapes, floppy diskettes, flash memory drives, optical data storage devices, compact-disc read-only memories, and other appropriate computer memories and data storage devices. In an illustrative embodiment, some or all of the software components may reside on a single server computer or on any combination of separate server computers. As one skilled in the art can appreciate, a computer program product implementing an embodiment disclosed herein may comprise one or more non-transitory computer readable media storing computer instructions translatable by one or more processors in a computing environment.

[0156] A “processor” includes any, hardware system, mechanism or component that processes data, signals or other information. A processor can include a system with a central processing unit, multiple processing units, dedicated circuitry for achieving functionality, or other systems. Processing need not be limited to a geographic location, or have temporal limitations. For example, a processor can perform its functions in “real-time,” “offline,” in a “batch mode,” etc. Portions of processing can be performed at different times and at different locations, by different (or the same) processing systems.

[0157] It will also be appreciated that one or more of the elements depicted in the drawings/figures can also be implemented in a more separated or integrated manner, or even

removed or rendered as inoperable in certain cases, as is useful in accordance with a particular application. Additionally, any signal arrows in the drawings/figures should be considered only as exemplary, and not limiting, unless otherwise specifically noted.

[0158] As used herein, the terms “comprises,” “comprising,” “includes,” “including,” “has,” “having,” or any other variation thereof, are intended to cover a non-exclusive inclusion. For example, a process, product, article, or apparatus that comprises a list of elements is not necessarily limited only those elements but may include other elements not expressly listed or inherent to such process, product, article, or apparatus.

[0159] Furthermore, the term “or” as used herein is generally intended to mean “and/or” unless otherwise indicated. For example, a condition A or B is satisfied by any one of the following: A is true (or present) and B is false (or not present), A is false (or not present) and B is true (or present), and both A and B are true (or present). As used herein, including the claims that follow, a term preceded by “a” or “an” (and “the” when antecedent basis is “a” or “an”) includes both singular and plural of such term, unless clearly indicated within the claim otherwise (i.e., that the reference “a” or “an” clearly indicates only the singular or only the plural). Also, as used in the description herein and throughout the claims that follow, the meaning of “in” includes “in” and “on” unless the context clearly dictates otherwise. The scope of the present disclosure should be determined by the following claims and their legal equivalents.

What is claimed is:

1. A system, comprising:

a processor;

a non-transitory computer-readable medium; and

stored instructions translatable by the processor for:

storing a first data structure representing a first split of a relationship in a first database shard, the first data structure referencing a second database item in a second database shard and keyed to a first database item in the first database shard, the first database item having a first object identifier encoded with the shard identifier of the first database shard; and

storing a second data structure representing a second split of the relationship in a second database shard, the second data structure referencing the first database item in the first database shard and keyed to the second database item in the second database shard, the second database item having a second object identifier encoded with the shard identifier of the second database shard, the second object identifier generated by the second database shard independently of the first object identifier generated by the first database shard;

wherein the second database item in the second database shard is accessible by a sharding manager using the first data structure keyed to the first database item in the first database shard, the sharding manager residing in a gateway server node or is in communication with the gateway server node.

* * * * *