

ONLINE DATABASE AVAILABILITY DURING UPGRADE

BACKGROUND

[0001] Web-based online services include files that are located on web servers along
5 with data that is stored in databases. For example, there are a large number of servers
located within different networks to handle the traffic that is directed to the service.
Upgrades of databases that are associated with an application may result in the database
being unavailable for a period of time. For some services the upgrade process may result in
a significant disruption to the customers of the service.

SUMMARY

[0002] This Summary is provided to introduce a selection of concepts in a simplified
form that are further described below in the Detailed Description. This Summary is not
intended to identify key features or essential features of the claimed subject matter, nor is it
intended to be used as an aid in determining the scope of the claimed subject matter.

15 [0003] An online service includes databases that are upgraded while still processing
requests. For example, web servers continue to request operations on the database while it
is being upgraded. The schema of the database is upgraded before the web servers are
upgraded to utilize the upgraded schema. Changes that are made to the upgraded schema
are backwards compatible with the schema being used during the upgrade process.

20 Restrictions are placed on the operations performed on the database during the upgrade
process. After upgrading the schema, the web servers of the online service are upgraded to
use the upgraded schema.

BRIEF DESCRIPTION OF THE DRAWINGS

[0004] FIGURE 1 illustrates a cloud manager system for managing networks that are
25 associated with an online service, such as a content management service;

[0005] FIGURE 2 shows a cloud manager including managers and associated
databases;

[0006] FIGURE 3 shows an exemplary job record stored within a row of a database;

[0007] FIGURE 4 shows an example system for a network including front-end and
30 back-end servers for an online service;

[0008] FIGURE 5 illustrates a computer architecture for a computer; and

[0009] FIGURES 6-7 show a process for upgrading a database for an online service
while still processing requests.

DETAILED DESCRIPTION

[0010] Referring now to the drawings, in which like numerals represent like elements, various embodiment will be described.

[0011] Generally, program modules include routines, programs, components, data
5 structures, and other types of structures that perform particular tasks or implement particular abstract data types. Other computer system configurations may also be used, including hand-held devices, multiprocessor systems, microprocessor-based or programmable consumer electronics, minicomputers, mainframe computers, and the like. Distributed computing environments may also be used where tasks are performed by remote processing
10 devices that are linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote memory storage devices.

[0012] FIGURE 1 illustrates a cloud management system for managing networks that are associated with an online service. System 100 illustrates cloud manager 105 that is
15 connected to and manages different networks potentially distributed across the world. Each of the networks is configured to provide content services for one or more tenants (e.g. clients, customers). The networks may be hosted within a cloud service and/or in an on-premises data center. Cloud manager 105 is used in deploying, configuring and managing the networks. The cloud manager is configured to receive requests through an idempotent
20 and asynchronous application web service application programming interface (API) 150 that can tolerate intermittent network failures.

[0013] As illustrated, cloud manager 105 comprises work manager 110, machine manager 115, application specific manager 120, scripts 130 and a central repository, such as data store(s) 140 (e.g. databases). The functionality that is not included within one of the
25 illustrated managers may reside in some other location of the cloud manager. According to one embodiment, application manager 120 is a SharePoint tenant manager that comprises SharePoint specific logic.

[0014] Work manager 110 manages the execution of tasks and enables scheduling and retry of longer running tasks. Work manager 110 starts jobs stored in job queue 112 and
30 keeps track of running jobs. When a predetermined time has elapsed, work manager 110 may automatically cancel the task and perform some further processing relating to the task. According to one embodiment, the tasks in job queue 112 are executed by work manager 110 by invoking one or more scripts 130. For example, a scripting language such as Microsoft's PowerShell® may be used to program the tasks that are executed by work

manager 110. Each script may be run as a new process. While executing each script as a new process may have a fairly high CPU overhead, this system is scalable and helps to ensure a clean environment for each script execution plus full cleanup when the script is completed.

5 **[0015]** Machine manager 115 is configured to manage the physical machines in the networks (e.g. Network 1, Network 2, Network 3). Generally, machine manager 115 understands Networks, Physical Machines, Virtual Machines (VMs), VM Images (VHDs), and the like. The machine manager does not have a strong binding to the specific services running within the networks but keeps track of the various components in the networks in
10 terms of “roles.” For example machine manager 115 could be requested through API 150 to deploy a VM of type “Foo” with version 12.34.56.78 on Network 3. In response to a request to cloud manager 105, machine manager 115 locates a suitable Physical Machine that is located on Network 3 and configures the VM according to the VM Image associated with the VM’s Role. The physical machine is configured with a VHD of type Foo with
15 version 12.34.56.78 that is stored within a data store, such as data store 140. The images used within the network may also be stored in other locations, such as a local data share for one or more of the networks. Scripts may be run to perform the installation of the VHD on the physical machine as well as for performing any post-deployment configuration. Machine manager 115 keeps track of the configuration of the machines each network. For example,
20 machine manager 115 may keep track of a VM’s role (type of VM), state of the VM (Provisioning, Running, Stopped, Failed), version and whether the VM exists in a given farm (which implies their network).

[0016] Scripts 130 is configured to store scripts that are executed to perform work both locally for cloud manager 105 and remotely on one or more of the networks. One or more
25 of the scripts 130 may also be stored in other locations. For example, scripts to be performed on a network (e.g. Network 1, Network 2, Network 3) may be stored locally to that network. The scripts may be used for many different purposes. For example, the scripts may be used to perform configurations of machines in one or more of the networks, changing settings on previously configured machines, add a new VM, add a new database,
30 move data from one machine to another, move tenants, change schemas, and the like. According to one embodiment, the scripts are Microsoft’s PowerShell® scripts. Other programming implementations may be used. For example, a compiled and/or early-bound programming language may be used to implement the functionality. Scripting, however, is a fairly concise language to express many of the tasks that are to be performed.

Programming the equivalent in a programming language, such as C#, would often require much more verbose implementations. The scripts are also late-bound, meaning that multiple versions of underlying code-bases can be targeted without having to constantly link to different interface DLLs. Using PowerShell scripts allows a process to be started locally
5 by cloud manager 105 that may in turn start a process on a remote machine (i.e. a physical machine in one of the attached networks). Other techniques may also be used to start a process on a remote machine, such as Secure Shell (SSH) and the like.

[0017] Application specific information that cloud manager 105 is managing is performed by application manager 120. According to one embodiment, the application
10 specific information relates to Microsoft SharePoint®. As such, application manager 120 is configured to know about SharePoint Tenants, Site Collections, and the like.

[0018] Each network may be configured as a dedicated network for a tenant and/or as a multi-tenant network that services more than one client. The networks may include a changing number of physical/virtual machines with their configuration also changing after
15 deployment. Generally, a network may continue to grow as long as the networking limits (e.g. load balancer and network switches) are not exceeded. For example, a network may start out with ten servers and later expand to one hundred or more servers. The physical machines within a network may be assigned a class or type. For example, some of the machines may be compute machines (used for web front ends and app servers) and other
20 machines may be storage machines that are provisioned with more storage than compute machines. According to an embodiment, cloud manager 105 configures the machines within a network with multiple versions of the image files. According to an embodiment, farms usually have a same version of image files.

[0019] According to one embodiment, the software limits are managed by the cloud
25 manager system 100 within the network by virtualizing the machines and managing independently acting “Farms” inside the network. Each network may include one or more farms (e.g. see Network 1). According to one embodiment, a network is considered a single cluster of network load balanced machines that expose one or more VIP (Virtual IP) to the outside world and can route that traffic to any of the machines within the network. The
30 machines in the network generally are tightly coupled and have minimum latencies (i.e. < 1ms ping latency).

[0020] Farms are the basic grouping of machines used to coordinate applications that need tightly bound relationships. For example, content farms may be deployed within each of the networks for a content management application, such as Microsoft SharePoint®.

Generally, the set of machines in each of the farms provide web service and application server functions together. Typically, the machines inside the farm are running the same build of an application (i.e. SharePoint) and are sharing a common configuration database to serve specific tenants and site collections.

5 **[0021]** Farms can contain heterogeneous sets of virtual machines. Cloud manager 105 maintains a “farm goal” within data store 140 which is a target number of machines of each role for each farm. Some roles include Content Front End, Content Central Admin, Content Timer Service, Federated Central Admin, Federated App Server etc. For example, content farms are the basic SharePoint farm that handles incoming customer requests. Federated
10 Services farms contain SharePoint services that can operate cross farms such as search and the profile store. Farms may be used for hosting large capacity public internet sites. Some farms may contain a group of Active Directory servers and a Provisioning Daemon. Cloud manager 105 automatically deploys and/or decommissions virtual machines in the networks to help in meeting the defined target. These farms goals may be automatically and/or
15 manually configured. For example, the farm goals may change to respond to changes in activity and capacity needs. Network Farm- there is one network farm per Network that contains all the VM roles that scale out easily as a resource to the whole Network.

[0022] The Cloud Manager Web Service APIs 150 are designed to work in the context of a massively scalable global service. The APIs assume that any network request might fail
20 and/or hang in transit. Calls to cloud manager 105 are configured to be idempotent. In other words, the same call may be made to cloud manager 105 multiple times (as long as the parameters are identical) without changing the outcome.

[0023] Cloud manager 105 is designed to do very little processing (<10ms, <50ms) before returning a response to any given request. Cloud manager 105 maintains records to
25 keep track of current requests. For example, cloud manager 105 updates records in a local database and if necessary schedules a “job” to perform more lengthy activity later.

[0024] Cloud manager keeps track of Images (such as Virtual Disk Images) that are the templates used to deploy new machines within a network. The Image references may be stored in a database, such as database 140, and/or in some other location. The images may
30 be stored in one or more shared data stores that are local to the network(s) on which the image will be deployed. According to one embodiment, each Image includes a virtual machine (VM) role type that specifies the type of VM it can deploy, the number of processors that it should use, the amount of RAM that it will be assigned, a network ID used

to find a nearby install point (so they don't get copied repeatedly over the cross data-center links) and a share path that the deployment code can use to access the VHD.

[0025] Generally, machines in the networks being managed by cloud system 100 are not upgraded in the traditional manner by downloading data and incorporating the data into the existing software on the machine. Instead, machines are updated by replacing a VHD with an updated VHD. For example, when a new version of software is needed by a farm, a new farm is deployed that has the new version installed. When the new farm is deployed, the tenants are moved from the old farm to the new farm. In this way, downtime due to an upgrade is minimized and each machine in the farm has a same version that have been tested. When a virtual machine needs to be upgraded, the VM on the machine may be deleted and replaced with the VM that is configured to run the desired service.

[0026] While upgrades to existing software are not optimal, some servers within the networks do utilize the traditional update procedure of an in-place upgrade. For example, Active Directory Domain Controllers are upgraded by updating the current software on the server without completely replacing an image on the machine. The cloud manager may also be upgraded in place in some instances.

[0027] FIGURE 2 shows a cloud manager including managers and associated databases. As illustrated, cloud manager 200 comprises work manager 210, work database 215, machine manager 220, machine database 225, tenant manager 230, tenant database 235, secrets database 245 and web service APIs 240.

[0028] Generally, databases used within a cloud management system (e.g. system 100) are sized to enable high performance. For example, a database (such as work database 215, machine database 225, tenant database 235 and secrets database 245) may not exceed a predefined size limit (e.g. 30 GB, 50 GB, 100 GB, and the like). According to an embodiment, a database is sized such that it is small enough to fit in-memory of a physical machine. This assists in high read I/O performance. The size of the database may also be selected based on performance with an application program, such as interactions with a SQL server. The databases used in the farms may also be sized to enable high performance. For example, they may be sized to fit in-memory of the host machine and/or sized such that backup operations, move operations, copy operations, restore operations are generally performed within a predetermined period of time.

[0029] Cloud manager 200 divides the cloud manager data into four databases. The work database 215 for the work manager. The machine database 225 for the machine manager 220. The tenant database 235 for the tenant manager 230 and a secrets database

245 for storing sensitive information such as system account and password information, credentials, certificates, and the like. The databases may be on the same server and or split across servers. According to an embodiment, each database is mirrored for high availability and is a SQL database.

5 **[0030]** Cloud manager 200 is configured to interact with the databases using a reduced set of SQL features in order to assist in providing availability of the cloud manager 200 during upgrades of the databases. For example, foreign keys or stored procedures are attempted to be avoided. Foreign keys can make schema changes difficult and cause unanticipated failure conditions. Stored procedures place more of the application in the
10 database itself.

[0031] Communications with the SQL servers are attempted to be minimized since roundtrips can be expensive compared to the cost of the underlying operation. For example, it is usually much more efficient if all of the current SQL server interactions to a single database are wrapped in a single round-trip.

15 **[0032]** Constraints are rarely used within the databases (215, 225, 235). Generally, constraints are useful when it helps provide simple updates with the right kind of error handing without extra queries. For example, the fully qualified domain name (FQDN) table has a constraint placed on the “name” to assist in preventing a tenant from accidentally trying to claim the same FQDN as is already allocated to a different tenant.

20 **[0033]** Caution is used when adding indices. Indices typically improve read performance at the cost of extra I/Os for write operations. Since the data within the databases is primarily RAM resident, even full table scans are relatively fast. According to an embodiment, indices may be added once the query patterns have stabilized and a performance improvement may be determined by proposed indices. According to an
25 embodiment, if adding the index will potentially take a long time the "ONLINE=ON" option may be specified such that the table isn't locked while the index is initially built.

[0034] According to an embodiment, upgrades to databases within the cloud manager may be performed without causing downtime to the cloud manager system. In other words, even during an upgrade of the cloud manager, the cloud manager continues processing
30 received requests. As such, changes made to the schema are to be compatible with the previous schema. The SQL schema upgrade is run before the web servers used by the cloud manager are upgraded. When the web servers are upgraded they can start to use the new features enabled in the database. Database upgrades are limited such that operations involved in the upgrade are quick and efficient. For example, tables may be added and new

nullable columns may be added to existing columns. New columns may be added at the end of a table. Generally, time consuming operations to the databases are avoided. For example, adding a default value to a newly added column at creation time may be a very time consuming operation when there is a large amount of data. Adding a nullable column, however, is a very quick operation. As discussed above, adding new indices are allowed, but caution should be taken when adding a new constraint to help ensure sure that the schema upgrade won't break with the existing data. For example, when a constraint is added it may be set to a state that is not checked and avoids a costly validation of existing rows and potential errors. Old tables and unused columns are removed after a new version is being used and the cloud manager is not accessing those tables and columns.

[0035] Generally, a single row in each of the databases is used to indicate a task and/or a desired state. For example, the tenant database 235 includes a single row for each tenant. A given tenant may include a Required Version record. This record is used to help ensure that the tenant is placed on a farm running the required version. For example, for tenant 1 to stay on SharePoint 14 SP1, the required version for tenant could be set to "14.1." and any version including 14.1 would match and any other versions (e.g. 14.2.xxxx) would not match. The tenant records may include other items such as authorized number of users, quotas (e.g. allowed total data usage, per user data usage, etc.), time restrictions, and the like. Some organization might have multiple tenants that represent different geographies, organizations or capabilities. According to an embodiment, tenants are walled off from each other without explicit invitation of the users (via extranet or other features).

[0036] According to one embodiment, each tenant is locked into a specific network. Tenants are kept localized to a small set of databases. A tenant is either small (smaller than would fill one database) in which case it is in exactly one database, shared with other tenants. This implies that all the tenants sharing that database need to upgrade at the same time. When a tenant grows larger it may be moved to its own dedicated database(s) and now might have more than one, but is not sharing databases with other tenants. Maintaining a large tenant in one or more dedicated databases helps in reducing a number of databases that are needed to be upgraded simultaneously in a single upgrade.

[0037] Similarly, the work database 215 includes a single row for each job. The machine database 225 may include a row for each physical machine, VM, farm, and the like. For example, machine manager database 225 may include a version string. According to an embodiment, each VHD, Farm, and VM within a network has an associated version string.

[0038] According to one embodiment, the cloud manager includes a simple logging system that may be configured to record a log entry for each web service call. A logging system may be implemented that includes as few/many features as desired. Generally, the logging system is used for measuring usage and performance profiling.

5 **[0039]** According to an embodiment, the Web Service APIs 240 are built using SOAP with ASP.net. The various Web Methods in the APIs follow two main patterns- Gets and Updates. Generally, the update methods take a data structure as the input and return the same structure as the output. The output structure returns the current state of the underlying object in the database, potentially differing from the input object if validation or other
10 business logic changed some properties or else with additional properties filled in (for example record IDs or other values calculated by the cloud manager). The update methods are used for initial object creation as well as subsequent updates. In other words, callers to the web service APIs 240 can simply request the configuration they want and they don't need to keep track of whether the object already exists or not. In addition this means that
15 updates are idempotent in that the same update call can be made twice with the identical effect to making it only once. According to an embodiment, an update method may include a LastUpdated property. When the LastUpdated property is present, the cloud manager 200 rejects the Update if the value of LastUpdate does not match the one currently stored in the database. Some Update methods include properties that are set on the first invocation of the
20 method and are not set on other invocations of the method.

[0040] Cloud manager 200 is configured to avoid the use of callbacks. Since callbacks may be unreliable, clients interacting with cloud manager 200 may check object status using a web service API when they want to check a status of an update. According to an embodiment, a call to an update method causes cloud manager 200 to set the state of the
25 underlying object to "Provisioning" and when the updates are completed the state is set to "Active".

[0041] FIGURE 3 shows an exemplary job record stored within a row of a database. As illustrated, record 300 comprises job identifier 302, type 304, data 306, owner 308, step 310, last run 312, expire time 314, next time 316, state 318 and status 320.

30 **[0042]** Generally, for each task that is requested to be performed, the cloud manager creates a record in database 350 (e.g. work database 215 in FIGURE 2).

[0043] Job identifier 302 is used to specify a unique identifier for the requested task.

[0044] Type 304 specifies the task to perform. For example, the type may include a name of the script to be executed. For example, when the task is to run the script named

“DeployVM.ps1” then the data 306 may include the identifier (e.g. “-VMID 123”). This allows new task types to be added to the system without requiring any changes to compiled or other binary parts of the system.

5 **[0045]** Data 306 is used to store data that is associated with the task. For example, the data may be set to the tenant, machine, network, VM, etc. on which the task is to be performed. The data 306 may also store one or more values to which a value in a database is set. The process running the task may look to the job record to see what value the desired number of machines is set to. The script uses the value in the database to perform the operation.

10 **[0046]** Owner 308 specifies a process/machine that is executing the process. For example, when a cloud manager machine starts execution of a job, the machine updates the owner 308 portion of the record with an ID of the machine.

15 **[0047]** Step 310 provides an indication of a step of the current script. For example, the script may divide a task into any number of steps. As the process completes a step of the script, step 310 is updated. A process may also look at step 310 to determine what step to execute in the script and to avoid having to re-execute previously completed steps.

[0048] Last run 312 provides a time the script was last started. Each time a script is started, the last run time is updated.

20 **[0049]** Expire time 314 is a time that indicates when the process should be terminated. According to an embodiment, the expire time is a predetermined amount of time (e.g. five minutes, ten minutes...) after the process is started. The expire time may be updated by a requesting process through the web service API.

25 **[0050]** Next time 316 is a time that indicates when a task should next be executed. For example, a process may be stopped after completion of a step and be instructed to wait until the specified next time 316 to resume processing.

[0051] State 318 indicates a current state and Status 320 indicates a status of a job (e.g. Created, Suspended, Resumed, Executing, Deleted).

30 **[0052]** Duplicate rows in the database can be removed before they are performed if they have the same task type and data values. For example, multiple requests may be made to perform the same task that are stored in multiple rows of the database.

[0053] A job can have one or more locks 355 associated with it. If locks are not available then a job will not be scheduled to run until the locks are available. The locks may be configured in many different ways. For example, the locks may be based on a mutex, a semaphore, and the like. Generally, a mutex prevents code from being executed

concurrently by more than one thread and a semaphore restricts a number of simultaneous uses of a shared resource up to a maximum number. According to an embodiment, a lock is a character string that represents a resource. The resource may be any type of resource. For example, the lock may be a farm, a machine, a tenant, and the like. Generally, the locks are
5 used to defer execution of one or more tasks. Each job may specify one or more locks that it needs before running. A job may release a lock at any time during its operation. When there is a lock, the job is not scheduled. A job needing more than one lock requests all locks required at once. For example, a job already in possession of a lock may not request additional locks. Such a scheme assists in preventing possible deadlock situations caused by
10 circular lock dependencies amongst multiple jobs.

[0054] FIGURE 4 shows an example system 400 for a network including front-end and back-end servers for an online service. The example system 400 includes clients 402 and 404, network 406, load balancer 408, WFE servers 410, 412, 414 and back-end servers 416-419. Greater or fewer clients, WFEs, back-end servers, load balancers and networks
15 can be used. Additionally, some of the functionality provided by the components in system 400 may be performed by other components. For example, some load balancing may be performed in the WFEs.

[0055] In example embodiments, clients 402 and 404 are computing devices, such as desktop computers, laptop computers, terminal computers, personal data assistants, or
20 cellular telephone devices. Clients 402 and 404 can include input/output devices, a central processing unit (“CPU”), a data storage device, and a network device. In the present application, the terms client and client computer are used interchangeably.

[0056] WFEs 410, 412 and 414 are accessible to clients 402 and 404 via load balancer 408 through network 406. As discussed, the servers may be configured in farms. Back-end
25 server 416 is accessible to WFEs 410, 412 and 414. Load balancer 408 is a dedicated network device and/or one or more server computers. Load balancer 408, 420, WFEs 410, 412 and 414 and back-end server 416 can include input/output devices, a central processing unit (“CPU”), a data storage device, and a network device. In example embodiments, network 406 is the Internet and clients 402 and 404 can access WFEs 410, 412 and 414 and
30 resources connected to WFEs 410, 412 and 414 remotely.

[0057] In an example embodiment, system 400 is an online, browser-based document collaboration system. An example of an online, browser-based document collaboration system is Microsoft Sharepoint® from Microsoft Corporation of Redmond, Washington. In

system 400, one or more of the back-end servers 416-419 are SQL servers, for example SQL Server from Microsoft Corporation of Redmond, Washington.

[0058] WFEs 410, 412 and 414 provide an interface between clients 402 and 404 and back-end servers 416-419. The load balancers 408, 420 direct requests from clients 402 and 5 404 to WFEs 410, 412 and 414 and from WFEs to back-end servers 416-419. The load balancer 408 uses factors such as WFE utilization, the number of connections to a WFE and overall WFE performance to determine which WFE server receives a client request. Similarly, the load balancer 420 uses factors such as back-end server utilization, the number of connections to a server and overall performance to determine which back-end server 10 receives a request.

[0059] An example of a client request may be to access a document stored on one of the back-end servers, to edit a document stored on a back-end server (e.g. 416-419) or to store a document on back-end server. When load balancer 408 receives a client request over network 406, load balancer 408 determines which one of WFE server 410, 412 and 414 15 receives the client request. Similarly, load balancer 420 determines which one of the back-end servers 416-419 receive a request from the WFE servers. The back-end servers may be configured to store data for one or more tenants (i.e. customer).

[0060] Referring now to FIGURE 5, an illustrative computer architecture for a computer 500 utilized in the various embodiments will be described. The computer 20 architecture shown in FIGURE 5 may be configured as a server, a desktop or mobile computer and includes a central processing unit 5 ("CPU"), a system memory 7, including a random access memory 9 ("RAM") and a read-only memory ("ROM") 11, and a system bus 12 that couples the memory to the central processing unit ("CPU") 5.

[0061] A basic input/output system containing the basic routines that help to transfer 25 information between elements within the computer, such as during startup, is stored in the ROM 11. The computer 500 further includes a mass storage device 14 for storing an operating system 16, application programs 10, data store 24, files, and a cloud program 26 relating to execution of and interaction with the cloud system 100.

[0062] The mass storage device 14 is connected to the CPU 5 through a mass storage 30 controller (not shown) connected to the bus 12. The mass storage device 14 and its associated computer-readable media provide non-volatile storage for the computer 500. Although the description of computer-readable media contained herein refers to a mass storage device, such as a hard disk or CD-ROM drive, the computer-readable media can be any available media that can be accessed by the computer 100.

[0063] By way of example, and not limitation, computer-readable media may comprise computer storage media and communication media. Computer storage media includes volatile and non-volatile, removable and non-removable media implemented in any method or technology for storage of information such as computer-readable instructions, data
5 structures, program modules or other data. Computer storage media includes, but is not limited to, RAM, ROM, Erasable Programmable Read Only Memory (“EPROM”), Electrically Erasable Programmable Read Only Memory (“EEPROM”), flash memory or other solid state memory technology, CD-ROM, digital versatile disks (“DVD”), or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic
10 storage devices, or any other medium which can be used to store the desired information and which can be accessed by the computer 500.

[0064] According to various embodiments, computer 500 may operate in a networked environment using logical connections to remote computers through a network 18, such as the Internet. The computer 500 may connect to the network 18 through a network interface
15 unit 20 connected to the bus 12. The network connection may be wireless and/or wired. The network interface unit 20 may also be utilized to connect to other types of networks and remote computer systems. The computer 500 may also include an input/output controller 22 for receiving and processing input from a number of other devices, including a keyboard, mouse, or electronic stylus (not shown in FIGURE 5). Similarly, an input/output controller
20 22 may provide output to a display screen 28, a printer, or other type of output device.

[0065] As mentioned briefly above, a number of program modules and data files may be stored in the mass storage device 14 and RAM 9 of the computer 500, including an operating system 16 suitable for controlling the operation of a networked computer, such as the WINDOWS® operating systems from MICROSOFT® CORPORATION of Redmond,
25 Washington. The mass storage device 14 and RAM 9 may also store one or more program modules. In particular, the mass storage device 14 and the RAM 9 may store one or more application programs, such as cloud program 26, that perform tasks relating to the cloud system.

[0066] Referring now to FIGURES 6-7, processes for upgrading a database with an
30 online service will be described.

[0067] When reading the discussion of the routines presented herein, it should be appreciated that the logical operations of various embodiments are implemented (1) as a sequence of computer implemented acts or program modules running on a computing system and/or (2) as interconnected machine logic circuits or circuit modules within the

computing system. The implementation is a matter of choice dependent on the performance requirements of the computing system implementing the invention. Accordingly, the logical operations illustrated and making up the embodiments described herein are referred to variously as operations, structural devices, acts or modules. These operations, structural
5 devices, acts and modules may be implemented in software, in firmware, in special purpose digital logic, and any combination thereof.

[0068] Referring now to FIGURE 6, after a start operation, the process 600 flows to operation 610, where a request to upgrade a database is received. The request may be received through many different sources. For example, the request may be received through
10 an API, such as a Web Service API, a command line interface, a graphical user interface, and the like.

[0069] Moving to operation 620, the schema that is associated with the database is upgraded. According to an embodiment, upgrades made to the schema are compatible with the schema currently being used by the cloud system. Generally, upgrades to the schema
15 should attempt to avoid operations that are timely to process. For example, write operations to the database often are very time consuming, whereas read operations are generally much faster to perform.

[0070] According to an embodiment, upgrades to databases within the cloud manager may be performed without causing downtime to the cloud management system. In other
20 words, even during an upgrade of a database used by the cloud manager, the cloud manager continues receiving requests and performing operations read and write operations on the database that is being upgraded. Generally, the SQL queries from the cloud manager are written such that schema changes will not impact them.

[0071] The cloud manager system (e.g. system 100) is configured to interact with the
25 cloud manager databases using a reduced set of SQL features in order to assist in providing availability of the cloud manager during upgrades of the databases. For example, foreign keys or stored procedures are attempted to be avoided. Foreign keys can make schema changes difficult and cause unanticipated failure conditions. Stored procedures place more of the application in the database itself. As discussed above, constraints are rarely used
30 within the databases.

[0072] Moving to operation 620, the schema for the database is upgraded. The SQL schema upgrade is performed before the web front-ends (WFEs) on the web servers used by the cloud manager are upgraded. Generally, when upgrades are applied to the schema, the

underlying data in the database should remain unaffected such that the WFEs that don't have the new functionality to support the upgraded schema continue to function properly.

[0073] The upgrades to the database schema may include adding one or more columns, tables, indices, and the like. Generally, the database upgrades are limited such that

5 operations that interact with the database during the upgrade are quick and efficient. For example, new tables may be added and new nullable columns may be added at the end to an existing table. The time consuming operations to the databases are avoided. For example, adding a new column that requires the entry of a value at creation time of the column may be a very time consuming operation when there is a large amount of data. Many tables may
10 have a large number of rows. For example, adding a new column to a SharePoint tenants table may have millions of rows. Having to populate each of those rows in the column with a value could cause a significant performance hit on the database as the database writes would have to hit every page to enter a value. Adding a nullable column to an existing table, however, does not need to touch every page of the database. Instead, adding a
15 nullable column that does not require an initial value is an operation that can be performed very quickly that is generally true irrespective of the size of the table. When the schema include new indices, they should be added such that the index is computed in the background after it is added and not make it available until the background processing of the index is complete. Constraints should be added with no check such that access to the
20 database is allowed during the upgrade.

[0074] Moving to operation 630, the database being upgraded continues to receive requests for database operations from the WFEs. The operation of the cloud manager system during the upgrade process behaves as if no upgrade is being performed.

[0075] Flowing to decision block 640, a determination is made as to whether the
25 schema upgrade is complete. When the upgrade is not complete, the upgrade process continues and the system continues to receive and process received requests. When the upgrade is complete, the process moves to operation 650, where the web servers including the WFEs used for interacting with the database are upgraded to include the new functionality of the upgraded schema. The upgrade of the WFEs may be performed using
30 different methods. For example, the new WFE software may be installed on the same web server and/or placed on a different web server.

[0076] Transitioning to operation 660, the incoming requests are moved to the upgraded WFEs. The activation of the new WFEs may take place in serial and/or in parallel. For example, a portion of the WFEs may be made available at one point in time

and another portion at another point in time. According to an embodiment, ASP.net functionality is used that allows a parameter to be set to direct the requests to the upgraded software.

5 **[0077]** Following the transitioning operation 660, incoming requests to new WFEs may uncover previously undiscovered errors with the upgraded software. In this situation, a remedy is to downgrade to the previous software version operating with the upgraded schema. In one embodiment, the WFEs are quickly downgraded to the previous software version in-place while operating against the upgraded schema.

[0078] The process then moves to an end block and returns to processing other actions.

10 **[0079]** FIGURE 7 shows updating data within the database after schema upgrade has been completed.

[0080] After a start operation, the process 700 flows to operation 710 where a request to a database is received. Generally, the request comes from one of the WFEs on a web server and may be any type of request relating to a database (e.g. reads, writes, searches,
15 etc...). According to an embodiment, read operations include the explicit columns that are to be read from. In this way, even when a schema changes, the explicit columns within the data remain the same.

[0081] Moving to decision operation 720, a determination is made as to whether data has been read from the database that includes a nullable column that does not have a
20 specified value. When a nullable column has not been accessed, the process moves to decision operation 740. When a nullable column has been accessed, the process flows to operation 730.

[0082] At operation 730, a value for the nullable column may be determined. For example, a default value may be chosen or the value for the nullable column may be
25 determined from some other action (i.e performing a calculation, look up, etc.). In this way, the values for the nullable columns may be updated over time, rather than in a single time-consuming operation when they are created. When a column is accessed it may be checked to determine if the value is null and when the value is null it may be replaced with a default value if there is not already a specified value. The specific code pattern to check for
30 Null & optionally replace with a default value if there is not already a specified value. The following is an example code snippet:

```

[0083] if (!Reader.IsDBNull(5))
    {
        this.State = Reader.GetString(5);
    }
5   else
    {
        this.State = DefaultState.Active;
    }

```

[0084] After the object is written with the determined value, the null value is gone.

10 [0085] Moving to decision operation 740, a determination is made as to whether there are any unused tables/columns that should be cleaned up. This determination may occur after a predetermined period of time in which the upgraded schema has been running and/or by some other input (e.g. received command). According to an embodiment, unused tables and columns are removed after the upgraded schema is being used and the cloud manager is

15 not accessing those tables and columns. When a table or column is not to be removed, the process returns to operation 710. When a table or column is to be removed, the process moves to operation 750.

[0086] At operation 750, the unused table(s) and/or column(s) are removed.

[0087] The process then moves to an end block and returns to processing other actions.

20 [0088] The above specification, examples and data provide a complete description of the manufacture and use of the composition of the invention. Since many embodiments of the invention can be made without departing from the spirit and scope of the invention, the invention resides in the claims hereinafter appended.

What is claimed is:

1. A method for upgrading a database that is still being accessed in an online service, comprising:

starting an upgrade process of a database used in the online service that includes
5 updating a schema that is used by the database; wherein the database being upgraded continues to receive, and process, and perform requests from web servers for database operations throughout the upgrade process of the schema;

upgrading the schema of the database currently being utilized to an upgraded schema; wherein the upgraded schema is backwards compatible with the schema;

10 after upgrading the schema to the upgraded schema, upgrading web servers of the online service to use the upgraded schema; and

the upgraded web servers accessing the database using the upgraded schema.

2. The method of Claim 1, wherein upgrading the schema of the database comprises adding a nullable column to an end of a table.

15 3. The method of Claim 2, further comprising adding a value to the nullable column when a subsequent request to the database accesses data comprising the nullable column.

4. The method of Claim 1, wherein upgrading the schema of the database comprises adding a table.

20 5. The method of Claim 1, wherein upgrading the schema of the database comprises adding an index during the upgrade process and calculating the index after the upgrade process is complete.

6. The method of Claim 1, further comprising after the upgrade process of the database schema is complete and the upgraded web servers are accessing the database
25 using the upgraded schema, automatically determining when to remove unused tables and columns from the database.

7. The method of Claim 1, wherein upgrading the web servers of the online service comprises upgrading software on the web servers such that requests to the database continue to be fulfilled during the upgrading of the software on the web servers.

30 8. The method of Claim 1, wherein upgrading the web servers of the online service comprises activating a set of web servers that use the upgraded schema and directing requests to the new set of web servers.

9. The method of Claim 1, wherein the online service is a document collaboration service, and the database is a tenant database is maintained that includes a single row for each tenant.

10. A computer-readable storage medium, excluding signals, having computer-executable instructions for upgrading a database currently in use in an online service, comprising:

starting an upgrade process of a database used in the online service; wherein the database being upgraded continues to receive, process and perform requests from web servers for database operations during the upgrade process; wherein the online service comprises a document collaboration service; wherein the database is a SQL database; upgrading a schema of the database currently being utilized to an upgraded schema;

wherein the upgraded schema is backwards compatible with the schema; after upgrading the current schema has completed, upgrading the web servers of the online service to use the upgraded schema; and

accessing the database using the upgraded web servers utilizing the upgraded schema.

11. The computer-readable storage medium of Claim 10, wherein upgrading the schema of the database comprises adding a nullable column to an end of a table and after the upgrade process has completed adding a value to the nullable column when a subsequent request to the database accesses data comprising the nullable column.

12. The computer-readable storage medium of Claim 10, further comprising restricting write operations to the database that relate to the upgrade process of the schema when the schema of the database is being upgraded.

13. The computer-readable storage medium of Claim 10, wherein the database is a tenant database is maintained that includes a single row for each tenant.

14. The computer-readable storage medium of Claim 10, wherein upgrading the schema of the database comprises adding an index during the upgrade process and setting an option such that a table associated with the index remains unlocked when the index is initially built.

15. The computer-readable storage medium of Claim 10, further comprising after the upgrade process of the database schema is complete and the upgraded web servers are accessing the database using the upgraded schema, automatically determining when to remove unused tables and columns from the database.

16. The computer-readable storage medium of Claim 10, wherein upgrading the web servers of the online service comprises continuing to fulfill requests to the database during the upgrading of the software on the web servers and activating a set of web servers that use the upgraded schema and directing requests to the new set of web servers.

17. A system for upgrading a database currently in use in an online service, comprising:

a network comprising web servers, back-end servers and a database;
wherein a current schema is associated with the database;

a processor and a computer-readable storage medium excluding signals;
an operating environment stored on the computer-readable medium and
executing on the processor; and

an upgrade manager operative to:

start an upgrade process of the database; wherein the database
continues to receive, process and perform from the web servers for database
operations during the upgrade process; wherein the online service
comprises a document collaboration service;

upgrade the schema of the database currently being utilized to an
upgraded schema; wherein the upgraded schema is backwards compatible
with the schema;

after the upgrade of the schema has completed, upgrading the web
servers to use the upgraded schema; and

accessing the database using the upgraded web server utilizing the
upgraded schema.

18. The system of Claim 17, wherein upgrading the schema of the database comprises adding a nullable column to a table and after the upgrade process has completed adding a value to the nullable column when a subsequent request to the database accesses data comprising the nullable column and a check of the nullable column determines that a value of the nullable column is null.

19. The system of Claim 17, further comprising restricting write operations to the database that relate to the upgrade process of the schema when the schema of the database is being upgraded.

20. The system of Claim 17, wherein upgrading the web servers of the online service comprises continuing to fulfill requests to the database during the upgrading of the

software on the web servers and activating a set of web servers that use the upgraded schema and directing requests to the new set of web servers.