



(19) **United States**

(12) **Patent Application Publication**

Katti et al.

(10) **Pub. No.: US 2013/0066977 A1**

(43) **Pub. Date: Mar. 14, 2013**

(54) **MESSAGE QUEUE BEHAVIOR OPTIMIZATIONS**

(52) **U.S. Cl.**
USPC 709/206

(75) Inventors: **Jayteerth Katti**, Redmond, WA (US); **Kartik Paramasivam**, Redmond, WA (US); **Stuart John Langley**, Woollooware (AU); **Pramod Gurunath**, Sammamish, WA (US); **Muralidhar Krishnaprasad**, Redmond, WA (US); **Elvin Morales**, Kew Gardens, NY (US); **Manu Srivastava**, Redmond, WA (US)

(57) **ABSTRACT**

Network and storage calls are reduced by batching messages. Messages are collected from a client and sent to the gateway or backend application in one round trip. Alternatively, the messages are collected for different entities, and the batched messages are durably stored. Related messages, which may have the same sessionID, are grouped into a logical unit or session. This session may be locked to a single message consumer. A session may be associated with an application processing state as an atomic unit of work while other operations are performed on the messages in the session. Acknowledgements are accumulated by a broker on a message server, but the messages are not immediately deleted. Instead, multiple messages in a selected range are periodically truncated from a message store in a single operation. Expired messages for active sessions are proactively cleaned up to prevent sessions from reaching a quota or limit.

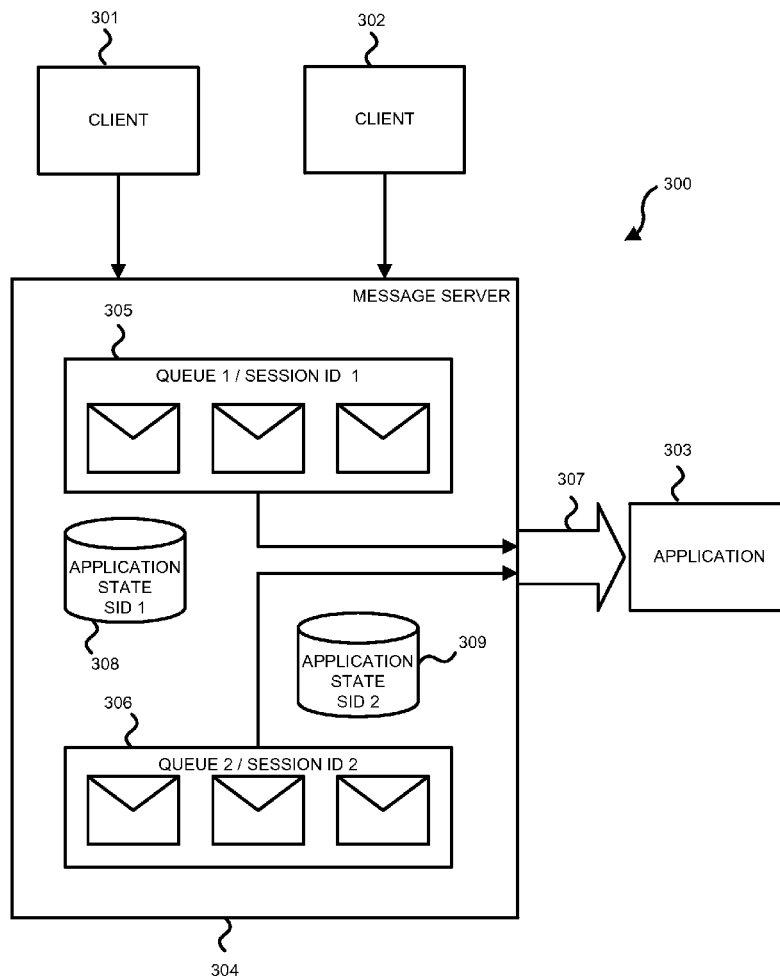
(73) Assignee: **MICROSOFT CORPORATION**, Redmond, WA (US)

(21) Appl. No.: **13/230,823**

(22) Filed: **Sep. 12, 2011**

Publication Classification

(51) **Int. Cl.**
G06F 15/16 (2006.01)






FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6

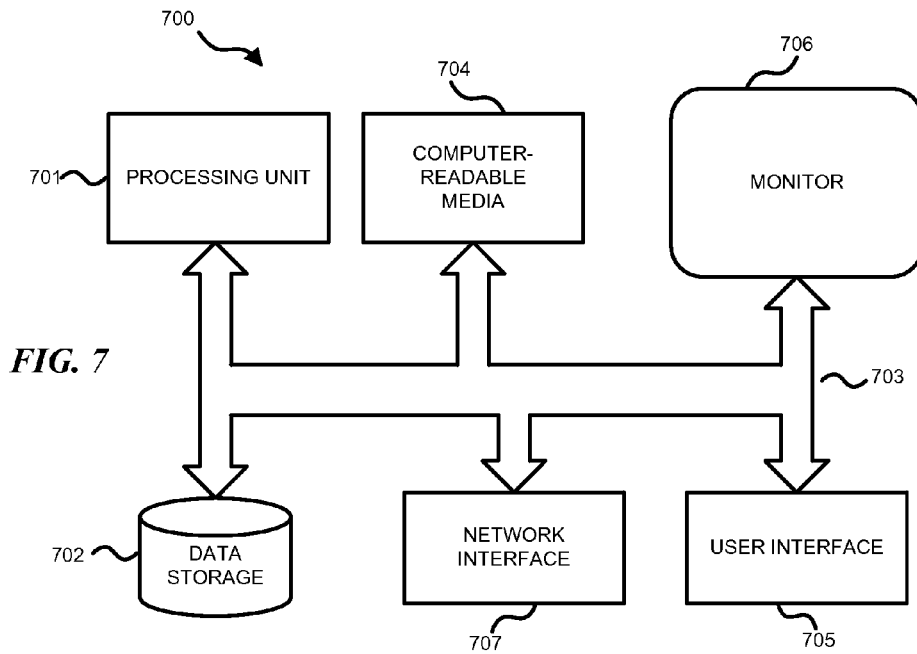


FIG. 7

MESSAGE QUEUE BEHAVIOR OPTIMIZATIONS

BACKGROUND

[0001] In a typical messaging system, many concurrent clients send and receive messages to and from the messaging system. Often, these clients send only one or just a small number of messages at a time. This interaction pattern between the clients and the messaging system increases system workload and

[0002] Independent calls to the message server by different clients increase the number of round-trips to the message server. Each call to the server typically has extra headers or other overhead that needs construction, transfer and processing. This work increases network activity, increases CPU utilization for both the client and server, and reduces the overall system throughput.

[0003] In other scenarios, when messages are persisted to a durable store, independent operations by different clients result in more round-trips to the durable store. One or more round-trips may be required per message. Durable stores such as databases and disk systems are generally optimized for small number of bulk or chunky operations. Even though each round-trip involves in just a small amount of I/O operation, an increased number of round-trips results fewer number of message operations on the durable store.

SUMMARY

[0004] 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 to limit the scope of the claimed subject matter.

[0005] Embodiments of the invention provide batching and other message handling techniques. In one form, batching of messages comprises the act of collecting messages or acknowledgements from a client and sending the collected messages to a gateway or backend application in one round trip. In another form, batching of messages comprises the act of collecting messages or acknowledgements for different entities, such as queues or topics, in a backend application or broker, and storing this information durably, such as by writing to a database.

[0006] As used in embodiments of the invention, processing messages using sessions comprises the act of grouping a bunch of related messages, which may have the same sessionID, for example, as a logical unit. This session or group of messages with a common session ID may be locked to a single consumer among all the competing consumers or backend applications. In another embodiment, processing messages using sessions comprises the ability to associate an application processing state or session state with a session. This may be accomplished as an atomic unit of work along with performing other operations on the same queue/topic.

[0007] Messages may be stored in queues. A queue can be either sessionful, wherein every message enqueued to such a queue has to have a sessionID, or sessionless, wherein the messages enqueued to the queue do not need to have a sessionID.

[0008] Completing a message does not delete the message immediately. Instead, a broker accumulates message

acknowledgements and then truncates or deletes multiple messages from the store in one operation.

[0009] In another embodiment, expired messages may be proactively cleaned up to ensure a session has sufficient space for future messages.

[0010] Embodiments of the invention reduce network and random storage communications. Additional embodiments optimize the storage operations for messages and acknowledgements, provide low-latency messaging, and manage message expiration.

[0011] Embodiments reduce the number of network and storage calls by batching groups of such calls together into one transaction. Calls from multiple messaging clients are automatically combined into one network call from a client library to a message server. This reduces the network IO and allows the message server to perform additional optimizations. Additionally, storage messaging operations from multiple clients may be combined into one store operation.

[0012] Overhead generated by message acknowledgement processing is reduced by transforming an acknowledgement list into a range of acknowledgements. Messages corresponding to the range of acknowledgements are bulk processed.

[0013] Embodiments provide optimization for low-latency scenarios when both a publisher and a consumer are connected or when the consumer is momentarily disconnected. Optimization is provided by completely bypassing or reducing store round-trips when the consumer is available to receive the message immediately.

DRAWINGS

[0014] To further clarify the above and other advantages and features of embodiments of the present invention, a more particular description of embodiments of the present invention will be rendered by reference to the appended drawings. It is appreciated that these drawings depict only typical embodiments of the invention and are therefore not to be considered limiting of its scope. The invention will be described and explained with additional specificity and detail through the use of the accompanying drawings in which:

[0015] FIG. 1 illustrates a system using batching where multiple messages between two clients are turned into single message;

[0016] FIG. 2 illustrates a system in which messages associated with the same session are grouped together;

[0017] FIG. 3 illustrates a system in which messages are grouped according to session identifiers;

[0018] FIG. 4 illustrates an embodiment in which a message server uses buffering techniques to reduce the number of store operations;

[0019] FIG. 5 illustrates a system using an optimized storage operation for message truncation according to one embodiment;

[0020] FIG. 6 illustrates a system for buffering or storing messages when a destination device is not available to receive messages; and

[0021] FIG. 7 illustrates an example of a suitable computing and networking environment on which embodiments may be implemented.

DETAILED DESCRIPTION

[0022] A message queuing server or platform may be optimized using the embodiments disclosed herein to provide durability of communications between clients, and other cli-

ents, applications, or storage. The system and methods allow an infrastructure to guarantee that messages will not be lost and will be sent in the correct order. Additionally, messaging operations may be performed atomically. If N operations need to be performed, then all N operations are processed as one unit.

[0023] Message Batching. FIG. 1 illustrates a system 100 using batching where multiple messages calls between client A 101 and a message consumer, such as back-end process 102, are turned into single call. Client A 101 sends multiple calls or messages 103 to back-end process 102. Messages 103 are processed by messaging server 104, which is responsible for routing the messages from multiple clients to the appropriate destination process. Messages 103 may be separate messages or they may be fragments or sub-parts of a larger message. In prior systems, messaging server 104 would send each message 103 to back-end process 102 individually. In embodiments of the present invention, messaging server 104 comprises a number of queues 105-107. Messaging server 104 assigns a queue 105 to client A 101, back-end process 102, or a particular interaction between client A 101 and back-end process 102. Instead of sending each message or call 103 directly to back-end process 102, messaging server 104 collects the messages 103 in queue 105. The messages 103 may be collected for a particular period of time, for example, and then they are sent as a group to back-end process 102 in a single message or call 108.

[0024] As used herein, the terms client and back-end process refer to any party, device, process, or application that sends messages to, or receives messages from, the message server. A client that sends messages may be referred to as message publisher or message publishing application. A back-end process that receives messages may be referred to as a message consumer or a message consuming application. The term server refers to a party, device, or application that provides message queuing services. A server may use a durable or non-durable store to store messages and states. Durable store refers to the storage layer of the message server. This layer persists messages to a permanent medium such as a database or disk. Non-durable store refers to the store layer, which does not store messages in persistent medium. Such a store keeps the messages and states in memory.

[0025] Messaging client 101 sends requests 103 to message server 104 using a client library. The client library exposes a shared proxy component that interacts with the message server 104 on behalf of the client 101. Calls and messages from multiple clients are buffered inside the client library.

[0026] If a second client B 109 is sending messages 110 to back-end process 102, the messages 110 are also buffered in message server 104. Messages 110 from client B 109 are buffered in queue 106, which may be the same queue used for messages 103 or a different queue. Message queues 105-107 in message server 104 may be assigned based upon a sending client identity, a receiving client identity, a session identifier, or some other parameter.

[0027] When the number of entries in the queues 105-107 reach a certain threshold or after a certain time period elapses, the client library sends all the message data in network call 108 to the destination back-end process 102. Queues 105-107 may be a buffer, a memory stack, or any other storage. The threshold used to determine when messages will be sent from queues in message server 104 to destination back-end process 102 may be based on thresholds that are set for each indi-

vidual queue 105-107 or for an overall total threshold for all the queues 105-107 combined.

[0028] In other embodiments, a user-supplied transaction may force batching of a group of messages to be sent to the message consumer in order to keep processing of the batch of messages atomic.

[0029] Message queues 105-107 may be sessionless or sessionful. In a sessionless queue, the messages are unrelated and all messages in a queue may be routed to the back-end process 102. In a sessionful queue, messages are organized by session. Each queue 105-107 may be associated with one or more sessions. As messages 103, 110 arrive at server 104, the messages are routed to the appropriate queue 105-107 based upon a session identifier in the message. All of the messages for a particular session are grouped on to the same queue 105-107. Each session is associated with a particular consumer, such as back-end process 102. Back-end process 102 may request all messages from server 104 or messages for a particular session. Message server 104 then sends all of the messages for the requested session(s) to back-end process 102. In this way, back-end process 102 only receives the messages for the sessions that it is responsible for processing, and back-end process 102 does not have to sort out messages for unrelated sessions.

[0030] In one example, back-end process may be part of an order-processing application, and clients A and B may be users submitting orders to the application. The orders may be organized into different sessions based upon client location, such as a city associated with an order. Back-end process 102 may process all messages for a particular city by requesting all messages for the associated session. The back-end process 102 may then apply the appropriate city-specific processing to those orders, such as tax rates, delivery options, etc.

[0031] Multiple back-end processes 102, 111 may receive messages from message server 104. For example, processes 102 and 111 may each be assigned one or more sessions to process. Process 102 receives only its sessions 108, and process 111 receives only its sessions 112. Processes 102 and 111 may be different applications or may be different instances of the same application. Additionally, processes 102 and 111 may be running on the same physical or virtual machine 113 or on different machines.

[0032] Using the message batching optimization as illustrated in FIG. 1, the number of message round trips required in system 100 can be reduced. Instead of sending individual messages to back-end processes 102, 111, all of the messages for a particular session may be sent as a single group 108, 112 to the appropriate destination application.

[0033] Prefetching. Another optimization provided by the system illustrated in FIG. 1 is prefetching by the destination client. For example, when client B 102 requests the next message from message server 104, all pending messages in queues 105-107 that directed to the client B 102 are batched together and delivered to client B 102 in a single message 108. In this way, even if a time or volume threshold for sending messages to a destination client has not been achieved, those messages are still available to the destination process. Accordingly, back-end process 102 can request a next message when it is ready for further processing, and back-end process 102 will receive all pending messages from queues 105-107 that are associated with its sessions.

[0034] The load for a particular application may be partitioned into a number of sessions across multiple back-end processes 102, 111. Each session is associated with a specific

back-end process, and each back-end processes may support multiple sessions. If a back-end process (e.g. 102) dies or is unavailable, then the sessions associated with that process may be reassigned to another back process (e.g. 111). The unconsumed messages in the reassigned sessions are then processed by the new back-end process (111).

[0035] Session Locking. If multiple consumers or back-end processes are consuming sessions from a queue, each session is marked as 'locked' and is assigned to a unique consumer. As long as the consumer is receiving messages from the session, the session lock is renewed for that consumer. If the consumer for a particular session dies and, therefore, stops retrieving or receiving messages for that session, then a session lock will timeout at the message server/broker, and the session will be marked as 'available' in the message server. When a new consumer requests an available session from the message server, then the message server will return this unlocked, available session to this new consumer.

[0036] Session Grouping. FIG. 2 illustrates grouping messages associated with the same session. Messages 201-203 may be related to the same session, which may be identified by a session identifier (SID), for example. As the message server receives messages, it groups messages related to the same SID into the same queue 204. When queue 204 reaches a threshold number of messages, or when a predetermined period elapses, the message server then sends all of the messages 201-203 associated with the session to the destination client in a single transaction. Alternatively, using prefetching, when a client requests a next message for a particular session, all of the messages related to that session—for example, all the messages in session queue 204—are sent to the client in a single transaction.

[0037] FIG. 3 illustrates a system 300 in which messages are grouped according to session identifiers. Clients 301, 302 interact with application 303 via message server 304. The communications between client 301 and application 303 are assigned session identifier SID 1, and communications between client 302 and application 303 are assigned session identifier SID 2. Message server 304 groups messages the messages based upon their session identifier. Messages from client 301 are grouped in queue 1 305, and messages from client 302 are grouped in queue 306 by messages server 304.

[0038] Messages from queues 305 and/or 306 may be sent to application 303 as a single transaction 307 when a threshold is reached, such when the queue holds a predetermined number of messages or at predetermined intervals. Application 303 and message server 304 may agree whether or not messages from different queues or sessions may be combined into the same transaction 307. Additionally, using prefetching, application 303 may request a next message either from any client or for a particular session. Message server 304 may send all queued messages as single transaction 307 for all sessions or for the requested session.

[0039] System 300 provides more efficient messaging between message server 304 and application 303 and allows for application to process multiple clients in parallel. Instead of requiring application 303 to receive messages from clients 301 and 302 sequentially in the order sent by the clients, application 303 can process all pending messages from one client while message server 304 is queuing messages from other clients. Message server 304 and application 303 can process each group (i.e. each SID) of client messages independently and concurrently.

[0040] Message server 304 may further comprise an application state 308, 309 for each session, which allows message server 304 to monitor the processing state for each session. Message server 304 can update the status of all messages in a queue by one update to application state 308, 309. For example, the application state may be used to set session-specific thresholds or to indicate that a particular session or client is paused, suspended, or inactive.

[0041] In an example embodiment, application 303 may be part of an order processing application that processes orders from users at clients 301 and 302. Each client order is assigned a session identifier. Order processing application 303 may process multiple orders from different clients 301, 302 in parallel. Messages associated with each client's orders are grouped into a queue based upon the session identifiers. The application state may be used to monitor the status of each order, such as tracking how far each order has been processed and whether additional applications are involved in the order. For example, a client's order may be inactive or paused while waiting for input from a separate payment or shipping application. Message server 304 may hold or forward messages for that session based upon the application state.

[0042] FIG. 4 illustrates an embodiment in which a message server stores processing state data for the sessions. Message server 401 receives messages 402 various from clients. Each message is associated with a session, which may be determined, for example, from a session identifier in the message header. Messages server 401 stores each message into the appropriate queue 403 based upon the session associated with the message. Consuming applications 404 and 405 request and process messages for different sessions in queues 403. Consuming applications 404, 405 may be a back-end process, storage medium, database, or other any application that processes client messages or commands. The system of FIG. 4 allows for cross-entity batching in which messages from multiple, independent clients 402 are grouped together into a single session 406, 407 that is processed by the consuming applications 404, 405.

[0043] From time to time, consuming applications 404, 405 may push processing state information 408, 409 back to message server 401, which stores the session-state information in durable data store or memory 410. The session-state information is associated with the relevant session in message server 401. The session-state information may indicate, for example, the current status of session processing or an acknowledgement of message processing. The session-state information is available to the original or a new consuming application, which can use the information to determine where to continue processing the session. The session-state information may be used, for example, to ensure that each message is processed only once.

[0044] In one embodiment, to keep the session information consistent, the consuming application processes messages and stores session state data in a single transaction as an atomic unit of work. The consuming application receives and processes a message for a particular session. The consuming application then notifies the message server that it has processed the message and stores updated state information for the session in the message server.

[0045] Message Truncation. FIG. 5 illustrates a system 500 using an optimized storage operation for message truncation according to one embodiment. Message server 501 processes messages and receives acknowledgements from clients 502-

504. The acknowledgements are stored to message store **505**. Message server **501** needs to durably record these acknowledgements. Typically, this is done by deleting the acknowledged messages from message store **505**. However, a high rate of operations to message store **505** reduces throughput. In some cases, messages are consumed out of order or they may be left unprocessed. This can happen, for example, when the client application is not ready to process the message. On the other hand, if the messages are deleted as soon as they are processed, the result is a high amount of random operations on the message store **505**.

[0046] These issues are addressed in one embodiment by delaying the delete operations and, instead, periodically performing bulk, sequential delete operations. As the messages are acknowledged, the acknowledgments are recorded in permanent message store **505** using a message sequence number. Periodically, the messages sequence numbers are combined together to truncate all messages up to a lowest acknowledged sequence number. In this way, all messages to be deleted can be deleted together in one operation by the message server **501**.

[0047] For the selective deletion case, message server **501** assumes that the number of unprocessed messages is small. This allows message server **501** to convert a large list of message sequence numbers into a small list of sequence number ranges. In this case, a message range entry is formed by taking only the first and last entry in the contiguous sequence of message sequence numbers. This reduces the amount of data passed to the durable store layer and allows it to perform optimizations around range deletion.

[0048] Express Messaging. In some messaging scenarios, the messages are very time critical and need to be delivered with smallest amount of latency. Such a low latency scenario might also require that the message server drop the message if the message cannot be delivered in a timely manner. At the same time, the message server should not consume too many resources, such as memory, if the messages are not consumed at the rate they are being published. In one embodiment, this problem may be addressed as illustrated in FIG. 6.

[0049] As messages **601** are published by source **602**, if the destination **603** is waiting for the messages **601**, then message server **604** delivers the messages (**604**) immediately to destination **603**. However, if destination **603** is not connected or is not actively waiting for messages **601**, then message server **604** buffers the messages **601** inside an in-memory buffer **606**. If the destination **604** starts requesting messages, then the buffered messages are delivered (**607**) by the message server **604**.

[0050] If the in-memory buffer **606** is filled to a certain threshold, the subsequent messages **601** are written directly to a durable store **608**. This prevents the message server **604** from consuming unnecessary resources by holding messages in memory.

[0051] In one embodiment, message server **604** may use a circular in-memory buffer in which it evicts the oldest messages from the in-memory buffer to make space for new messages. In this case, messages are still delivered in first-in-first-out fashion.

[0052] As the messages from in-memory buffer **606** are consumed, stored messages are loaded from the durable store **608** into buffer **606**.

[0053] In another embodiment, upon message server **604** restart or failover, any existing messages from the durable store **608** are treated as expired and are discarded.

[0054] In some cases, both publishers and consumers may be connected to the message server. In traditional messaging scenarios, the message server first puts the message inside a message store and when consumer requests for messages, it reads the message from the store and delivers them. This means that every message must go through a store write and read operation which reduces the overall throughput.

[0055] In one embodiment, this problem is addressed as follows. As messages are published, if one or more consumers are waiting for messages, then those messages are delivered to the consumers directly, without persisting them in the store. A client's publish call completes when the consumers acknowledge messages. If the consumers do not acknowledge messages within a certain time interval, the messages are put in the store and the client's publish call completes. In this manner, the message server eliminates the store round-trip in most cases without sacrificing the consistency guarantees.

[0056] Also, if consumers are temporarily disconnected as the messages are being published, the message server writes the messages to the underlying store and also caches them in-memory. This way, when consumers ask for messages later, the cached messages may be delivered. This saves trips to the store and hence improves the overall throughput.

[0057] Message Expiry. A message sending application or its administrator may want to limit the amount of time a message can stay in the message server before the message is discarded. Because there may be very many queues—on the order of thousands of queues—it is not a trivial problem to detect expired messages from a durable store without scanning through all the messages. It would require a lot of processing to read through all of the messages on the message store and determine which ones are expired. In one embodiment, the message server handles the expiring messages as described below.

[0058] First, the system determines if a queue is loaded in memory. If the queue is loaded in memory, then the system may need to clean out expired messages. If the queue is not loaded in memory, then no clients have been sending messages and there is no need to clean out expired messages.

[0059] If the queue is loaded in memory, then the system determines whether any clients have been publishing messages to the queue. If message are being published to the queue, then the system may need to clean out expired messages. If no clients have been publishing messages to the queue, then there are no messages to clean out.

[0060] If message are being published to a queue that is loaded in memory, then the system determines if any consumer has been receiving messages from the queue for some recent period of time, such as during the last five minutes or some other configurable period. If no one has been receiving message from the queue during the configurable period, then the system may need to clean out expired messages. On the other hand, if some consumer has been receiving messages from the queue, then the queue is likely to have been cleaned out while sending messages.

[0061] If no one has been receiving message from the queue during the configurable period, then the system determines if the queue is nearing its quota. If the queue is nearing its quota, then the system begins an agent job that will scan the queue for messages that are expired and then delete the expired messages. The queue may be determined to be nearing its quota, for example, if the number of stored messages for the

queue are greater than a predetermined number or are greater than a predetermined percentage of an assigned storage space.

[0062] The system will need to clean up expired messages only if messages are being sent to the queue. If no client is publishing messages to the queue, then there is no motivation to clean the queue to make room for additional messages. On the other hand, if messages are being published to the queue, the system will look for expired messages to be deleted to prevent the queue from overrunning its quota.

[0063] In other embodiments, as messages are read from the durable store, if the messages are expired, they are discarded. In this approach, the durable message store does not need to maintain information about the expiring messages. This method does not require any background processing by the message server and, therefore, provides a very efficient mechanism for discarding expired messages for the case where the percentage of expiring messages is low.

[0064] Alternatively, the message server, with the help of the durable store, tracks the list of expiring messages and removes them from the store as soon as they expire. In this case, the expired messages do not take up space in the underlying store. This process is useful for the case where a queue has a fixed quota and needs to be allocated for unexpired messages and in the case where higher percentage of messages can expire.

[0065] FIG. 7 illustrates an example of a suitable computing and networking environment **700** on which the examples of FIGS. 1-6 may be implemented. The computing system environment **700** is only one example of a suitable computing environment and is not intended to suggest any limitation as to the scope of use or functionality of the invention. The invention is operational with numerous other general purpose or special purpose computing system environments or configurations. Examples of well-known computing systems, environments, and/or configurations that may be suitable for use with the invention include, but are not limited to: personal computers, server computers, hand-held or laptop devices, tablet devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputers, mainframe computers, distributed computing environments that include any of the above systems or devices, and the like.

[0066] The invention may be described in the general context of computer-executable instructions, such as program modules, being executed by a computer. Generally, program modules include routines, programs, objects, components, data structures, and so forth, which perform particular tasks or implement particular abstract data types. The invention may also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in local and/or remote computer storage media including memory storage devices.

[0067] With reference to FIG. 7, an exemplary system for implementing various aspects of the invention may include a general purpose computing device in the form of a computer **700**. Components may include, but are not limited to, processing unit **701**, data storage **702**, such as a system memory, and system bus **703** that couples various system components including the data storage **702** to the processing unit **701**. The system bus **703** may be any of several types of bus structures including a memory bus or memory controller, a peripheral

bus, and a local bus using any of a variety of bus architectures. By way of example, and not limitation, such architectures include Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, and Peripheral Component Interconnect (PCI) bus also known as Mezzanine bus.

[0068] The computer **700** typically includes a variety of computer-readable media **704**. Computer-readable media **704** may be any available media that can be accessed by the computer **701** and includes both volatile and nonvolatile media, and removable and non-removable media. By way of example, and not limitation, computer-readable media **704** 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 structures, program modules or other data. Computer storage media includes, but is not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical disk storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by the computer **700**. Communication media typically embodies computer-readable instructions, data structures, program modules or other data in a modulated data signal such as a carrier wave or other transport mechanism and includes any information delivery media. The term "modulated data signal" means a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not limitation, communication media includes wired media such as a wired network or direct-wired connection, and wireless media such as acoustic, RF, infrared and other wireless media. Combinations of the any of the above may also be included within the scope of computer-readable media.

[0069] The data storage or system memory **702** includes computer storage media in the form of volatile and/or non-volatile memory such as read only memory (ROM) and random access memory (RAM). A basic input/output system (BIOS), containing the basic routines that help to transfer information between elements within computer **700**, such as during start-up, is typically stored in ROM. RAM typically contains data and/or program modules that are immediately accessible to and/or presently being operated on by processing unit **701**. By way of example, and not limitation, data storage **702** holds an operating system, application programs, and other program modules and program data.

[0070] Data storage **702** may also include other removable/non-removable, volatile/nonvolatile computer storage media. By way of example only, data storage **702** may be a hard disk drive that reads from or writes to non-removable, nonvolatile magnetic media, a magnetic disk drive that reads from or writes to a removable, nonvolatile magnetic disk, and an optical disk drive that reads from or writes to a removable, nonvolatile optical disk such as a CD ROM or other optical media. Other removable/non-removable, volatile/nonvolatile computer storage media that can be used in the exemplary operating environment include, but are not limited to, magnetic tape cassettes, flash memory cards, digital versatile disks, digital video tape, solid state RAM, solid state ROM, and the like. The drives and their associated computer storage

media, described above and illustrated in FIG. 7, provide storage of computer-readable instructions, data structures, program modules and other data for the computer 700.

[0071] A user may enter commands and information through a user interface 705 or other input devices such as a tablet, electronic digitizer, a microphone, keyboard, and/or pointing device, commonly referred to as mouse, trackball or touch pad. Other input devices may include a joystick, game pad, satellite dish, scanner, or the like. These and other input devices are often connected to the processing unit 701 through a user input interface 705 that is coupled to the system bus 703, but may be connected by other interface and bus structures, such as a parallel port, game port or a universal serial bus (USB). A monitor 706 or other type of display device is also connected to the system bus 703 via an interface, such as a video interface. The monitor 706 may also be integrated with a touch-screen panel or the like. Note that the monitor and/or touch screen panel can be physically coupled to a housing in which the computing device 700 is incorporated, such as in a tablet-type personal computer. In addition, computers such as the computing device 700 may also include other peripheral output devices such as speakers and printer, which may be connected through an output peripheral interface or the like.

[0072] The computer 700 may operate in a networked environment using logical connections 707 to one or more remote computers, such as a remote computer. The remote computer may be a personal computer, a server, a router, a network PC, a peer device or other common network node, and typically includes many or all of the elements described above relative to the computer 700. The logical connections depicted in FIG. 7 include one or more local area networks (LAN) and one or more wide area networks (WAN), but may also include other networks. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets and the Internet.

[0073] When used in a LAN networking environment, the computer 700 may be connected to a LAN through a network interface or adapter 707. When used in a WAN networking environment, the computer 700 typically includes a modem or other means for establishing communications over the WAN, such as the Internet. The modem, which may be internal or external, may be connected to the system bus 703 via the network interface 707 or other appropriate mechanism. A wireless networking component such as comprising an interface and antenna may be coupled through a suitable device such as an access point or peer computer to a WAN or LAN. In a networked environment, program modules depicted relative to the computer 700, or portions thereof, may be stored in the remote memory storage device. It may be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers may be used.

[0074] Although the subject matter has been described in language specific to structural features and/or methodological acts, it is to be understood that the subject matter defined in the appended claims is not necessarily limited to the specific features or acts described above. Rather, the specific features and acts described above are disclosed as example forms of implementing the claims.

What is claimed is:

1. A method, comprising:
 - receiving messages at a message server from a source; collecting the messages into one or more groups based upon properties associated with the messages; and sending all of the messages in a selected group to a destination in a single transaction.
 2. The method of claim 1, wherein each of the groups are buffered in a queue.
 3. The method of claim 1, wherein the destination is a backend application that processes the messages.
 4. The method of claim 1, wherein the destination is a gateway.
 5. The method of claim 1, wherein the destination is a durable message store.
 6. The method of claim 1, wherein the properties associated with the messages used for grouping comprise a destination address in a message header.
 7. The method of claim 1, wherein the properties associated with the messages used for grouping comprise source address in a message header.
 8. The method of claim 1, wherein the properties associated with the messages used for grouping comprise message topics.
 9. The method of claim 1, wherein the messages are acknowledgements of previously sent messages.
 10. A method, comprising:
 - receiving a plurality of messages at a message server from a source;
 - identifying related messages within the plurality of messages, the related messages having a same session identifier;
 - grouping the related messages; and
 - associating the group of related messages with a destination.
 11. The method of claim 10, wherein the related messages are grouped as a logical unit, and the method further comprising:
 - locking an association between the logical unit and the destination.
 12. The method of claim 11, further comprising:
 - updating an association between the logical unit and the destination each time the destination consumes messages from the logical unit.
 13. The method of claim 11, further comprising:
 - determining that the destination has not consumed messages from the logical unit within a threshold period; and
 - associating the logical unit to a new destination.
 14. The method of claim 10, wherein the group of related messages comprises a session, the method further comprising:
 - associating a session state with the session.
 15. The method of claim 14, wherein the destination is a message consumer that processes messages from the group of related messages, and wherein the session state represents application processing state for the destination.
 16. The method of claim 15, further comprising:
 - associating the application processing state with the session as an atomic unit of work while performing other operations on the messages in the session.
 17. A message server, comprising:
 - one or more processors adapted to process messages received from clients and to send the messages to consumers;

a system memory comprising a plurality of queues, each queue used to store groups of related messages; and
and a message broker adapted to collect acknowledgements for messages sent to consumers from the message server, the message broker periodically deleting multiple selected messages from the queues in a single operation, the selected messages corresponding to ranges of message acknowledgements received by the message broker.

18. The message server of claim 17, wherein the selected messages correspond to one or more lists of acknowledged message sequences.

19. The message server of claim 17, further comprising one or more of the processors operating to:

determine if a selected group of messages is loaded in the system memory;

determine if any clients are publishing messages to the selected group;

determine if any consumers have received messages from the selected group within a predetermined period;

if no consumer has been receiving messages within the predetermined period, then determine whether the selected group is nearing a quota; and

if the selected group is nearing the quota, identify and delete expired messages from the selected group.

* * * * *