

FIG. 1




FIG. 2




FIG. 3




FIG. 4

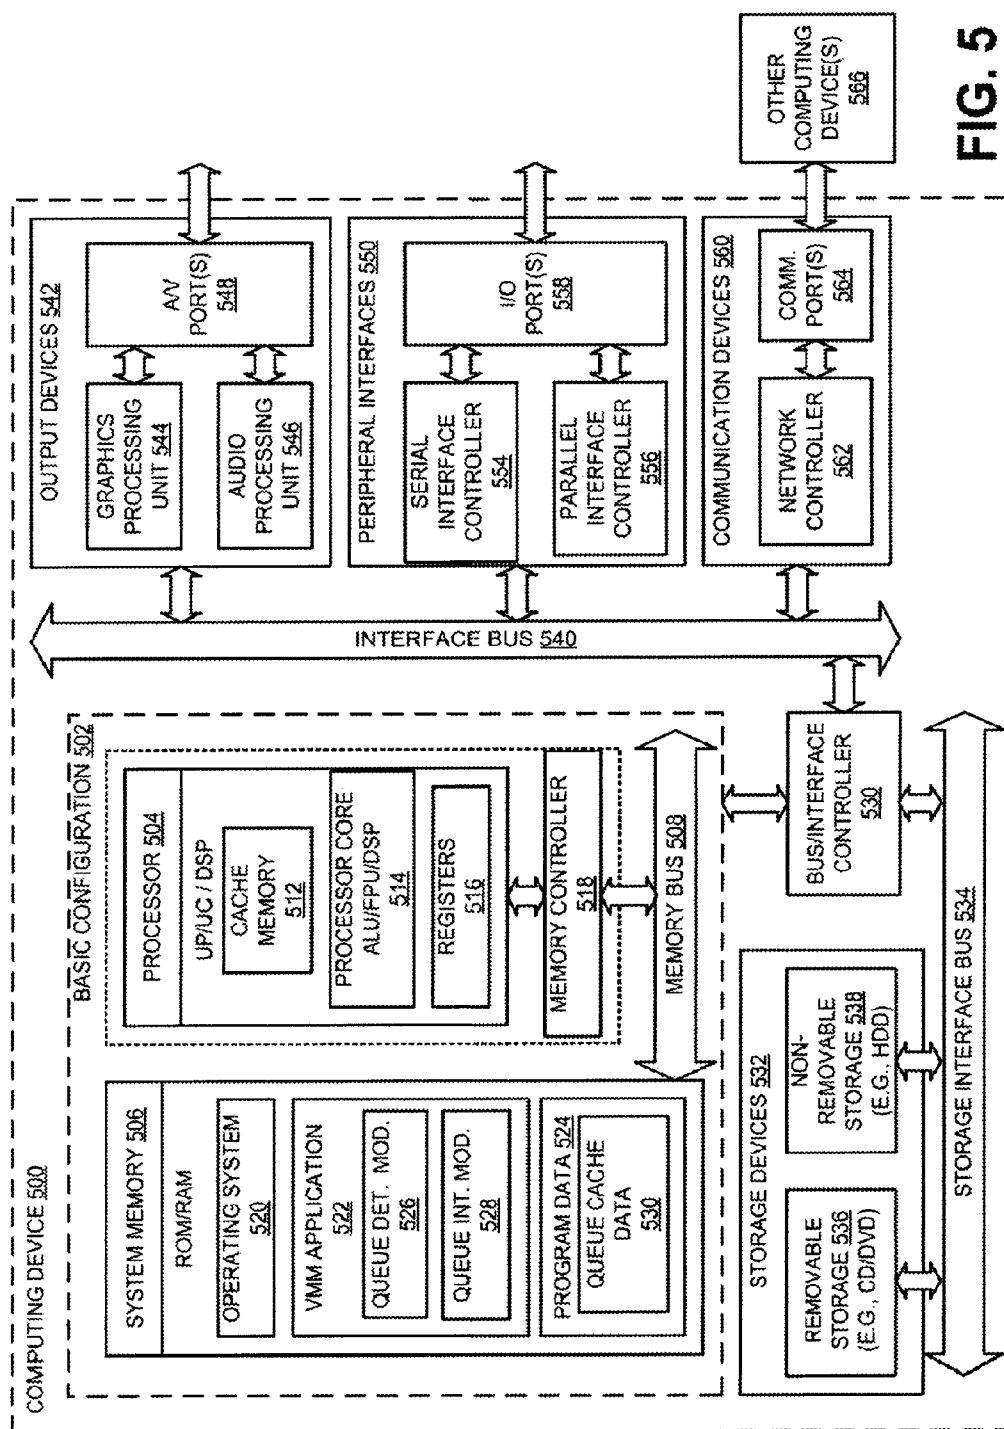
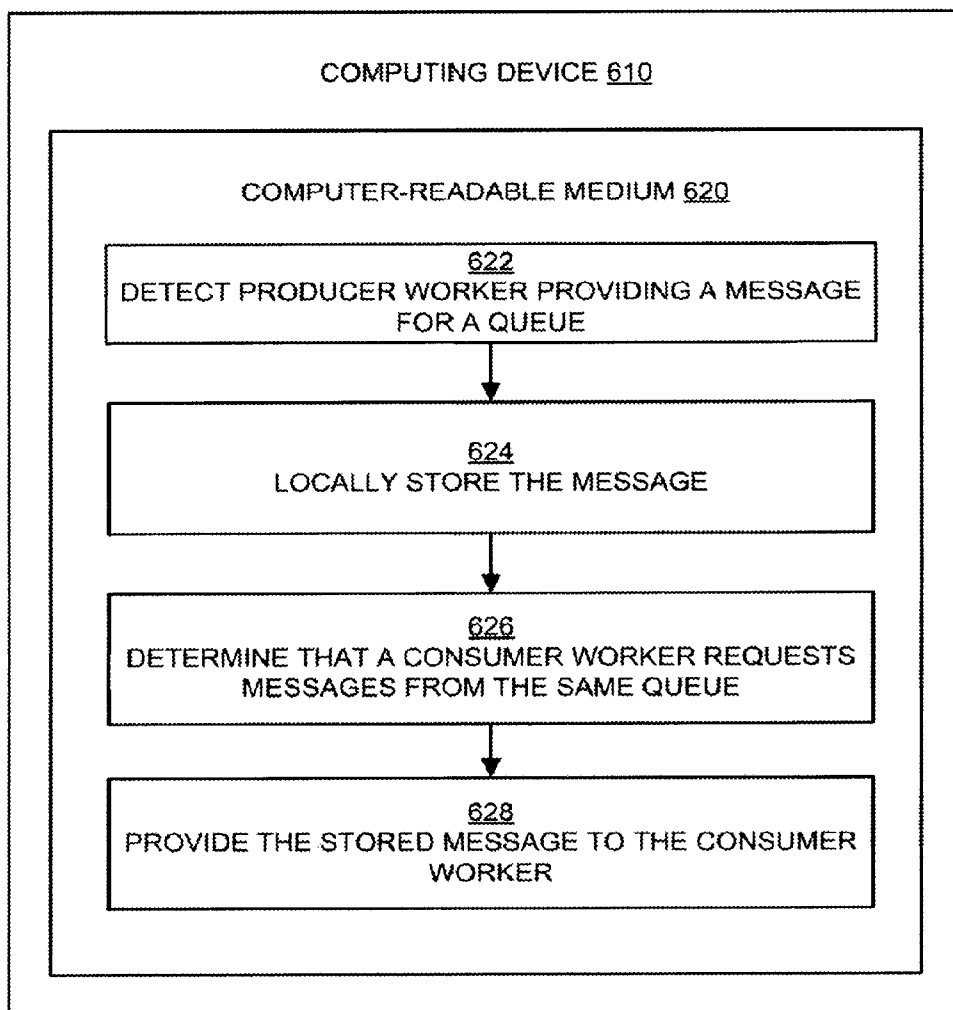
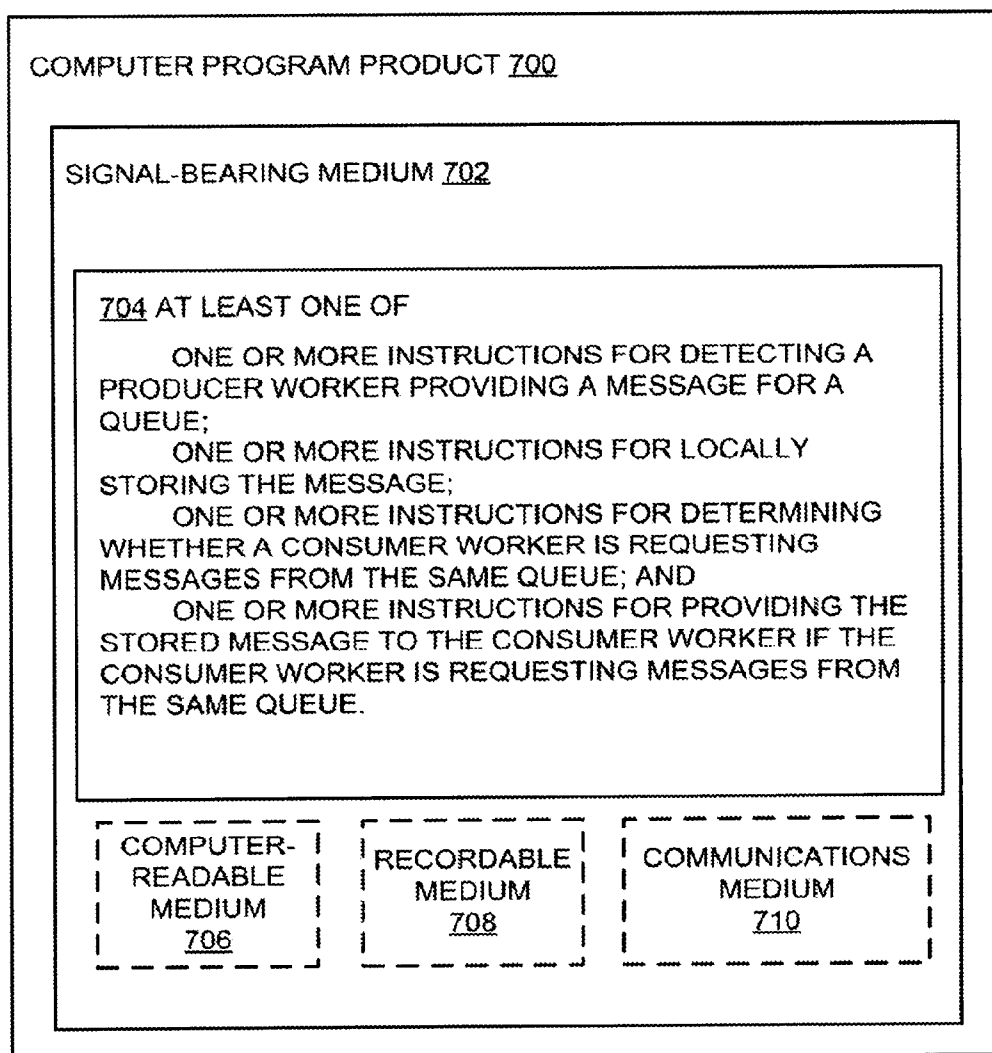


FIG. 5



**FIG. 6**

**FIG. 7**



## LOCAL MESSAGE QUEUE PROCESSING FOR CO-LOCATED WORKERS

### BACKGROUND

**[0001]** Unless otherwise indicated herein, the materials described in this section are not prior art to the claims in this application and are not admitted to be prior art by inclusion in this section.

**[0002]** Efficient management of interacting threads and workers needs fine granularity and awareness of process state and newer processors offer hardware acceleration for thread management. When interacting threads share hardware, various optimizations may be applied to hand off demand between the two. These improvements are typically available if smooth and fast worker handoff can be established between the interacting threads but they allow improvements like pipeline interleaving and fast task switching on successful speculative execution.

**[0003]** Meanwhile, datacenter multi-worker architectures are currently highly network-centric, designed around multiple minimum size worker/service instances that intercommunicate via messaging or queues so that multiple source workers can assign tasks onto a queue while multiple workers pull tasks off the queue to turn them into the next stage product output. Thus, existing multicore software can use multiple cores well. Existing web services, however, are designed to communicate via messaging capacities like message queues so even if two web services are on neighboring cores they send data between each other via the network—via a trip through at least a virtual router at the Virtual Machine Manager (VMM) level. This approach may be about 4-6 orders of magnitude slower than using intercore hardware.

### SUMMARY

**[0004]** The present disclosure generally describes techniques for locally processing queue requests from co-located workers.

**[0005]** According to some examples, a method is provided for locally processing queue requests from co-located workers. The method may include detecting a producer worker at a first server sending a first message to a datacenter queue at least partially stored at a second server, storing the first message at the first server, detecting a consumer worker at the first server sending a message request to the datacenter queue, and providing the stored first message to the consumer worker in response to the message request.

**[0006]** According to other examples, a virtual machine manager (VMM) is provided for locally processing queue requests from co-located workers. The VMM may include a queue usage detector module and a processing module. The queue usage detector module may be configured to detect a producer worker at a first server sending a first message to a datacenter queue at least partially stored at a second server and detect a consumer worker at the first server sending a message request to the datacenter queue. The processing module may be configured to store the first message at the first server and provide the stored first message to the consumer worker in response to the message request.

**[0007]** According to further examples, a cloud-based datacenter is provided for locally processing queue requests from co-located workers. The datacenter may include a first and a second virtual machine (VM) operable to be executed on one or more physical machines and a datacenter controller. The

datacenter controller may be configured to detect a producer worker executing on the first VM and sending a first message to a datacenter queue at least partially stored at a first datacenter location, store the first message at a second datacenter location different from the first, detect a consumer worker executing on the second VM and sending a message request to the datacenter queue, and provide the stored first message to the consumer worker in response to the message request, where the message is stored and provided from within a server local to the producer worker and the consumer worker.

**[0008]** According to yet further examples, a computer readable medium may store instructions for locally processing queue requests from co-located workers. The instructions may include detecting a producer worker at a first server sending a first message to a datacenter queue at least partially stored at a second server, storing the first message at the first server, detecting a consumer worker at the first server sending a message request to the datacenter queue, and providing the stored first message to the consumer worker in response to the message request.

**[0009]** The foregoing summary is illustrative only and is not intended to be in any way limiting. In addition to the illustrative aspects, embodiments, and features described above, further aspects, embodiments, and features will become apparent by reference to the drawings and the following detailed description.

### BRIEF DESCRIPTION OF THE DRAWINGS

**[0010]** The foregoing and other features of this disclosure will become more fully apparent from the following description and appended claims, taken in conjunction with the accompanying drawings. Understanding that these drawings depict only several embodiments in accordance with the disclosure and are, therefore, not to be considered limiting of its scope, the disclosure will be described with additional specificity and detail through use of the accompanying drawings, in which:

**[0011]** FIG. 1 illustrates an example datacenter-based system where local processing of queue requests from co-located workers may be implemented;

**[0012]** FIG. 2 illustrates an example datacenter-based system having multiple co-located workers;

**[0013]** FIG. 3 illustrates an example virtual machine manager implementing local processing of queue requests from co-located workers;

**[0014]** FIG. 4 illustrates the operation of the example virtual machine manager of FIG. 3;

**[0015]** FIG. 5 illustrates a general purpose computing device, which may be used to provide local processing of queue requests from co-located workers;

**[0016]** FIG. 6 is a flow diagram illustrating an example method for locally processing queue requests from co-located workers that may be performed by a computing device such as the computing device in FIG. 5; and

**[0017]** FIG. 7 illustrates a block diagram of an example computer program product,

**[0018]** all arranged in accordance with at least some embodiments described herein.

### DETAILED DESCRIPTION

**[0019]** In the following detailed description, reference is made to the accompanying drawings, which form a part hereof. In the drawings, similar symbols typically identify

similar components, unless context dictates otherwise. The illustrative embodiments described in the detailed description, drawings, and claims are not meant to be limiting. Other embodiments may be utilized, and other changes may be made, without departing from the spirit or scope of the subject matter presented herein. It will be readily understood that the aspects of the present disclosure, as generally described herein, and illustrated in the Figures, can be arranged, substituted, combined, separated, and designed in a wide variety of different configurations, all of which are explicitly contemplated herein.

**[0020]** This disclosure is generally drawn, inter alia, to methods, apparatus, systems, devices, and/or computer program products related to locally processing queue requests from co-located workers.

**[0021]** Briefly stated, technologies are generally described for locally processing queue requests from co-located workers. In some examples, information about the usage of remote datacenter queues by co-located workers may be used to determine one or more matched queues. Messages from local workers to a remote datacenter queue classified as a matched queue may be stored locally. Subsequently, local workers that request messages from matched queues may be provided with the locally-stored messages.

**[0022]** FIG. 1 illustrates an example datacenter-based system where local processing of queue requests from co-located workers may be implemented, arranged in accordance with at least some embodiments described herein.

**[0023]** As shown in a diagram 100, a physical datacenter 102 may include one or more physical servers 110, 111, and 113, each of which may be configured to provide one or more virtual machines 104. For example, the physical servers 111 and 113 may be configured to provide four virtual machines and two virtual machines, respectively. In some embodiments, one or more virtual machines may be used as a virtual datacenter. For example, the four virtual machines provided by the server 111 may be combined into a virtual datacenter 112. The virtual machines 104 and/or the virtual datacenter 112 may be configured to provide cloud-related data/computing services such as various applications, data storage, data processing, or comparable ones to a group of customers 108, such as individual users or enterprise customers, via a cloud 106.

**[0024]** FIG. 2 illustrates an example datacenter-based system having multiple co-located workers, arranged in accordance with at least some embodiments described herein.

**[0025]** As shown in a diagram 200, a datacenter (e.g., the physical datacenter 102 in FIG. 1) may include one or more layers of servers and routers, which may be arranged in a hierarchical fashion. The datacenter shown in the diagram 200 has a network of three router layers with the datacenter being a four-layer datacenter, although in other embodiments a datacenter may have more or fewer layers. The datacenter may include one or more application servers 228. In some embodiments, the application servers 228 may be configured to provide virtual machines (e.g., the virtual machines 104 described above in relation to FIG. 1) and/or execute one or more customer applications. The datacenter may also include one or more edge routers 226 that reside in the first network layer, where each edge router may couple to one or more of the application servers 228. The edge routers 226 in turn may couple to one or more aggregation routers 224, which may reside in a second layer of the network. In some embodiments, each of the aggregation routers 224 may couple to two

or more of the edge routers 226 and serve to combine or “aggregate” the network traffic from the connected edge routers. Finally, each of the aggregation routers 228 may couple to one or more core routers 222 in a third network layer.

**[0026]** As described above, the application servers 228 may be configured to provide virtual machines on which customer applications may execute. A customer application may be configured to execute as one or more virtual machine instances, also known as workers, and in some embodiments an application server may be configured to execute one or more workers. For example, an application server 230 may be configured to execute workers 232 and 234, and an application server 236 may be configured to execute workers 238 and 240. As a result, an individual application server (e.g., the application server 230 or 236) may have multiple workers co-located and executing on the same hardware. If a particular application server includes multi-core processors (processors having multiple processing cores integrated on the same die), different processing cores may execute different virtual machines and different workers on the same processor. For example, different workers may each be configured to execute on different virtual machines, where the virtual machines are configured to execute on the same physical hardware (i.e., processor or server). In some cases, the different processing cores may execute the different virtual machines/workers at approximately the same time.

**[0027]** In some embodiments, a particular worker may be responsible for executing a particular process, function, or series of functions, and multiple workers may collaborate to accomplish more complex tasks or processes. For example, one worker may generate data or information for further processing. This worker, known as a producer worker, may package the generated information into one or more messages. Another worker, known as a consumer worker, may then retrieve the message(s), extract the information in the message(s), and perform some work on the information. The consumer worker may then produce result(s) and/or generate other messages for further processing.

**[0028]** To facilitate message handling and transfer, the messages generated by producer workers may be stored at one or more datacenter queues. For example, datacenter queues may be stored at queue servers 242 and/or 244. Consumer workers may then retrieve messages from the queue(s) for processing. In some embodiments, each individual customer application may have an associated datacenter queue. Workers executing as part of a particular customer application may store and/or retrieve messages from the datacenter queue associated with that particular customer application.

**[0029]** In the diagram 200, the application servers on which workers execute (e.g., the servers 230 and 236) may be different from the application servers on which the datacenter queues reside (e.g., the queue servers 242 and 244). In this situation, a message generated by a producer worker may traverse the entire network (e.g., the edge routers 226, the aggregation routers 224, and the core routers 222) in order to be placed on a datacenter queue. For example, a message generated by a producer worker executing on virtual machine 232 may be routed through the edge routers 226, the aggregation routers 224, and the core routers 222 before reaching one of the queue servers 242 or 244. Similarly, a message retrieved by a consumer worker from a datacenter queue may also have to traverse the entire network, as would each request for such messages or lists of messages.

[0030] As discussed above, an application server having multi-core processors may be configured to execute multiple, different workers. In some embodiments, a producer worker and a consumer worker may both execute on the same application server (e.g., the workers 232 and 234 on the application server 230) and store/receive messages from the same datacenter queue. If the datacenter queue is stored at another server (e.g., the queue server 242), then produced and consumed messages may traverse the entire network, even though the producer and consumer workers may be co-located on the same server or even the same multi-core processor. As a result, message processing may be relatively slow as compared to using local message handling, and many of the advantages of multi-core systems, such as thread management hardware acceleration or multiprocessor load balancing optimizations, may be lost due to network delays.

[0031] FIG. 3 illustrates an example virtual machine manager implementing local processing of queue requests from co-located workers, arranged in accordance with at least some embodiments described herein.

[0032] According to a diagram 300, a virtual machine manager (VMM) 304 may mediate communications between virtual machines 302 and a network 306. In some embodiments, the VMM 304 may be associated with a particular application server in a datacenter, such as the application server 230 in FIG. 2. The virtual machines 302 may execute one or more producer and/or consumer workers (e.g., the workers 232 and 234 in FIG. 2), and the network 306 may correspond to the rest of the datacenter (e.g., the edge routers 226, the aggregation routers 224, and/or the core routers 222) or to an external network.

[0033] In some embodiments, producer workers executing on the virtual machines 302 may generate messages containing data for further processing. The messages may then be transmitted through the VMM 304 for forwarding to a datacenter queue stored at a different server (e.g., the queue servers 242 or 244 in FIG. 2). The VMM 304 may forward the messages through a virtualized network connection 310 to the datacenter queue, via the network 306.

[0034] While forwarding the messages to the datacenter queue, the VMM 304 may also observe network traffic through the virtualized network connection using a queue usage detector 312. The queue usage detector 312 may detect outgoing messages to queues (e.g., messages generated by producer workers executing in the virtual machines 302) and identify the producer workers and queues associated with those outgoing messages. The queue usage detector 312 may also detect message requests to queues (e.g., requests sent by consumer workers executing in the virtual machines 302) and identify the consumer workers and queues associated with those message requests. Detection may be achieved by observing network destinations associated with datacenter queues, which are often run by the datacenter. Detection may alternately be achieved by observing message headers or formats and recognizing queue commands, identities, or sessions. Detection may involve altering one or more settings within the hardware virtualization support of a processor, such as setting bits in the IOMap that will ensure the VMM receives a chance to handle messages passed to the network.

[0035] A queue user table 314 may then be constructed or updated based on the observed queue usage information. The queue user table 314 may include information about producer workers executing on the virtual machines 302, consumer workers executing on the virtual machines 302, and the

queues associated with each worker. The VMM 304 may then use the queue user table to determine matching local producer/consumer worker pairs by identifying messages (from a local producer worker) and requests (from a local consumer worker) to the same datacenter queue. This datacenter queue may be known as a “matched queue”.

[0036] Once one or more matched queues have been identified, matched queue information may be provided to a matched queue traffic intercept module 308. In some embodiments, outgoing messages and message requests to datacenter queues may be routed through the intercept module 308. The intercept module 308 may intercept outgoing messages to matched queues, store the intercepted messages in a queue cache 316, and provide stored messages from the queue cache 316 in response to message requests to the matched queues. This may increase the speed of message handling between local producer and consumer workers, because local consumer workers may receive messages from the local queue cache 316 instead of a remote datacenter queue stored on a relatively distant server. A command channel 318 may allow datacenter queues to relay certain commands to be performed on the queue cache 316.

[0037] FIG. 4 illustrates the operation of the example virtual machine manager of FIG. 3, arranged in accordance with at least some embodiments described herein.

[0038] According to a diagram 400, a producer worker executing on the virtual machines 302 may generate a message intended for a datacenter queue stored at a different server. The producer worker may transmit the message as a write command 402 through the VMM 304 for the datacenter queue. When the write command 402 reaches the VMM 304, it may first pass through the matched queue traffic intercept module 308. The intercept module 308 may forward the write command 402 on to the virtualized network connection 310 and henceforth to the network 306 and the remote datacenter queue. However, the intercept module may also determine if the write command 402 is intended for a matched queue (as described above). If so, the intercept module 308 may divert a copy of the write command 402 to the queue cache 316 in addition to forwarding the write command 402 to the virtualized network connection 310. In some cases the forwarded write command may be delayed to provide an opportunity for local consumption and deleted if the message is consumed locally within a predefined time period. The copy of the write command 402 may then be executed on the queue cache 316, writing the message generated by the producer worker into the queue cache 316.

[0039] A consumer worker executing on the virtual machines 302 may then request a message from the same datacenter queue via a read command 404. The read command 404, like the write command 402, may be transmitted to the VMM 304 for forwarding to the datacenter queue, and may pass through the intercept module 308. The intercept module 308 may forward the read command 404 on, but may also determine that the read command 404 is intended for a matched queue. The intercept module 308 may then also divert a copy of the read command 404 to the queue cache 316, as with the write command 402 described above. The copy of the read command 404 may then be executed on the queue cache 316, which may return a message 406 to the consumer worker. In some embodiments, the message 406 may be the message written by the write command 402. As a result, messages may be passed between producer workers and consumer workers both executing on the virtual machines

302 without incurring network delays from having to traverse the datacenter network, as described above in relation to FIG. 2. At the same time, the remote datacenter queue may still be updated based on the write command 402 and the read command 404 from the local workers, because the intercept module 308 forwards those commands on in addition to copying them for the local queue cache 316. In other examples, the read command may not be forwarded if it can be satisfied locally. This approach may provide efficiency if the write command has been delayed as discussed above.

[0040] The local queue cache 316 may store and provide messages from and to workers on the virtual machines 302; in effect serving as a local datacenter queue that can respond to requests from local consumer workers significantly faster than the remote datacenter queue. As described above, the intercept module 308 may direct the write command 402 and the read command 404 to both the remote datacenter queue and the local queue cache 316. In fact, in some embodiments the queue cache 316 may include a partial copy or an entire copy of a remote datacenter queue.

[0041] In order to assure that the remote datacenter queue and the local queue cache 316 remain reasonably consistent, commands from the remote datacenter queue may be provided to the VMM 304 via the command channel 318. For example, suppose that the write command 402 places a particular message on both the remote datacenter queue and the queue cache 316. Another consumer worker having access to the remote datacenter queue but not executing on the virtual machines 302 may be able to request the message. If the other consumer worker completes processing of the message, it may send a delete command to the remote datacenter queue deleting the message. In order to assure that the message copy on the queue cache 316 is also deleted, the remote datacenter queue may transmit a delete command via the command channel 318 to delete the message copy on the queue cache 316.

[0042] In some embodiments, the VMM 304 may also assist in maintaining consistency between the remote datacenter queue and the local queue cache 316. For example, the VMM 304 may be configured to forward a message request from a consumer worker to the remote datacenter queue if the remote datacenter queue is configured to hide the requested message after receiving the message request. This functionality may prevent another consumer worker from receiving the same requested message and thus duplicate processing. If the remote datacenter queue does not implement this functionality, the VMM 304 may refrain from forwarding the message request.

[0043] FIG. 5 illustrates a general purpose computing device, which may be used to provide local processing of queue requests from co-located workers, arranged in accordance with at least some embodiments described herein.

[0044] For example, the computing device 500 may be used to implement local message queue processing for co-located workers as described herein. In an example basic configuration 502, the computing device 500 may include one or more processors 504 and a system memory 506. A memory bus 508 may be used for communicating between the processor 504 and the system memory 506. The basic configuration 502 is illustrated in FIG. 5 by those components within the inner dashed line.

[0045] Depending on the desired configuration, the processor 504 may be of any type, including but not limited to a microprocessor ( $\mu$ P), a microcontroller ( $\mu$ C), a digital signal

processor (DSP), or any combination thereof. The processor 504 may include one more levels of caching, such as a level cache memory 512, a processor core 514, and registers 516. The example processor core 514 may include an arithmetic logic unit (ALU), a floating point unit (FPU), a digital signal processing core (DSP Core), or any combination thereof. An example memory controller 518 may also be used with the processor 504, or in some implementations the memory controller 518 may be an internal part of the processor 504.

[0046] Depending on the desired configuration, the system memory 506 may be of any type including but not limited to volatile memory (such as RAM), non-volatile memory (such as ROM, flash memory, etc.) or any combination thereof. The system memory 506 may include an operating system 520, a virtual machine manager (VMM) application 522, and program data 524. The VMM application 522 may include a queue detection module 526 for detecting queue usage of local workers and a queue intercept module 528 for intercepting messages and message requests from local workers to remote datacenter queues as described herein. The program data 524 may include, among other data, queue cache data 530 or the like, as described herein.

[0047] The computing device 500 may have additional features or functionality, and additional interfaces to facilitate communications between the basic configuration 502 and any desired devices and interfaces. For example, a bus/interface controller 530 may be used to facilitate communications between the basic configuration 502 and one or more data storage devices 532 via a storage interface bus 534. The data storage devices 532 may be one or more removable storage devices 536, one or more non-removable storage devices 538, or a combination thereof. Examples of the removable storage and the non-removable storage devices include magnetic disk drives such as flexible disk drives and hard-disk drives (HDD), optical disk drives such as compact disk (CD) drives or digital versatile disk (DVD) drives, solid state drives (SSD), and tape drives to name a few. Example computer storage media may include volatile and nonvolatile, 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.

[0048] The system memory 506, the removable storage devices 536 and the non-removable storage devices 538 are examples of computer storage media. Computer storage media includes, but is not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD), solid state drives, or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which may be used to store the desired information and which may be accessed by the computing device 500. Any such computer storage media may be part of the computing device 500.

[0049] The computing device 500 may also include an interface bus 540 for facilitating communication from various interface devices (e.g., one or more output devices 542, one or more peripheral interfaces 544, and one or more communication devices 566) to the basic configuration 502 via the bus/interface controller 530. Some of the example output devices 542 include a graphics processing unit 548 and an audio processing unit 550, which may be configured to communicate to various external devices such as a display or speakers via one or more A/V ports 552. One or more example

peripheral interfaces **544** may include a serial interface controller **554** or a parallel interface controller **556**, which may be configured to communicate with external devices such as input devices (e.g., keyboard, mouse, pen, voice input device, touch input device, etc.) or other peripheral devices (e.g., printer, scanner, etc.) via one or more I/O ports **558**. An example communication device **566** includes a network controller **560**, which may be arranged to facilitate communications with one or more other computing devices **562** over a network communication link via one or more communication ports **564**. The one or more other computing devices **562** may include servers at a datacenter, customer equipment, and comparable devices.

[0050] The network communication link may be one example of a communication media. Communication media may typically be embodied by 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 may include any information delivery media. A “modulated data signal” may be 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 may include wired media such as a wired network or direct-wired connection, and wireless media such as acoustic, radio frequency (RF), microwave, infrared (IR) and other wireless media. The term computer readable media as used herein may include both storage media and communication media.

[0051] The computing device **500** may be implemented as a part of a general purpose or specialized server, mainframe, or similar computer that includes any of the above functions. The computing device **500** may also be implemented as a personal computer including both laptop computer and non-laptop computer configurations.

[0052] Example embodiments may also include methods for implementing local message queue processing for co-located workers. These methods can be implemented in any number of ways, including the structures described herein. One such way may be by machine operations, of devices of the type described in the present disclosure. Another optional way may be for one or more of the individual operations of the methods to be performed in conjunction with one or more human operators performing some of the operations while other operations may be performed by machines. These human operators need not be collocated with each other, but each can be with a machine that performs a portion of the program. In other examples, the human interaction can be automated such as by pre-selected criteria that may be machine automated.

[0053] FIG. 6 is a flow diagram illustrating an example method for locally processing queue requests from co-located workers that may be performed by a computing device such as the computing device in FIG. 5, arranged in accordance with at least some embodiments described herein.

[0054] Example methods may include one or more operations, functions or actions as illustrated by one or more of blocks **622**, **624**, **626**, and/or **628**, and may in some embodiments be performed by a computing device such as the computing device **500** in FIG. 5. The operations described in the blocks **622-628** may also be stored as computer-executable instructions in a computer-readable medium such as a computer-readable medium **620** of a computing device **610**.

[0055] An example process for locally processing queue requests from co-located workers may begin with block **622**,

“DETECT PRODUCER WORKER PROVIDING A MESSAGE FOR A QUEUE”, where a queue usage detector (e.g., the queue usage detector **312**) may detect a local producer worker (e.g., a producer worker executing on the virtual machines **302**) providing messages for a remote datacenter queue. For example, the queue usage detector may detect a write command (e.g., the write command **402**) from a local worker destined for a remote datacenter queue. In some embodiments, if the local producer worker and/or the remote datacenter queue are known to be associated with a matched queue as described above, an intercept module (e.g., the intercept module **308**) may copy the message as described above in FIG. 4.

[0056] Block **622** may be followed by block **624**, “LOCALLY STORE THE MESSAGE”, where the intercept module may send the copied message to a local queue cache (e.g., the queue cache **316**) for local storage.

[0057] Block **624** may be followed by block **626**, “DETERMINE THAT A CONSUMER WORKER REQUESTS MESSAGES FROM THE SAME QUEUE”, where the intercept module may determine that a local consumer worker (e.g., a consumer worker executing on the virtual machines **302**) is requesting a message from the same remote datacenter queue. For example, the intercept module may determine that the consumer worker and/or the remote datacenter queue are associated with the same matched queue, as described above. The intercept module may then copy the message request, as described above in FIG. 4.

[0058] Finally, block **626** may be followed by block **628**, “PROVIDE THE STORED MESSAGE TO THE CONSUMER WORKER”, where the intercept module may send the copied message request to the local queue cache, which in turn may provide the message stored in block **624** to the requesting consumer worker, as described above in relation to FIG. 4.

[0059] FIG. 7 illustrates a block diagram of an example computer program product, arranged in accordance with at least some embodiments described herein.

[0060] In some examples, as shown in FIG. 7, the computer program product **700** may include a signal bearing medium **702** that may also include one or more machine readable instructions **704** that, when executed by, for example, a processor, may provide the functionality described herein. Thus, for example, referring to the processor **504** in FIG. 5, the VMM application **522** may undertake one or more of the tasks shown in FIG. 7 in response to the instructions **704** conveyed to the processor **504** by the medium **702** to perform actions associated with local message queue processing for co-located workers as described herein. Some of those instructions may include, for example, detecting a producer worker providing a message for a queue, locally storing the message, determining whether a consumer worker is requesting messages from the same queue, and/or providing the stored message to the consumer worker if the consumer worker is requesting messages from the same queue, according to some embodiments described herein.

[0061] In some implementations, the signal bearing medium **702** depicted in FIG. 7 may encompass a computer-readable medium **706**, such as, but not limited to, a hard disk drive, a solid state drive, a Compact Disc (CD), a Digital Versatile Disk (DVD), a digital tape, memory, etc. In some implementations, the signal bearing medium **702** may encompass a recordable medium **708**, such as, but not limited to, memory, read/write (R/W) CDs, R/W DVDs, etc. In some

implementations, the signal bearing medium **702** may encompass a communications medium **710**, such as, but not limited to, a digital and/or an analog communication medium (e.g., a fiber optic cable, a waveguide, a wired communications link, a wireless communication link, etc.). Thus, for example, the program product **700** may be conveyed to one or more modules of the processor **704** by an RF signal bearing medium, where the signal bearing medium **702** is conveyed by the wireless communications medium **710** (e.g., a wireless communications medium conforming with the IEEE 802.11 standard). The computer program product **700** may be delivered as a module for addition to a VMM such as a VMWare Installation Bundle (VIB).

**[0062]** According to some examples, a method for locally processing queue requests from co-located workers in a datacenter may include detecting a producer worker at a first server sending a first message to a datacenter queue at least partially stored at a second server, storing the first message at the first server, detecting a consumer worker at the first server sending a message request to the datacenter queue, and providing the stored first message to the consumer worker in response to the message request.

**[0063]** According to some embodiments, the method may further include intercepting the first message sent by the producer worker before storing the first message and/or storing the first message at the first server in a queue cache. The queue cache may include a copy or a partial copy of the datacenter queue. The method may further include building a table of queue usage based on at least one observed datacenter queue request and/or observing the at least one observed datacenter queue request.

**[0064]** According to other embodiments, the producer worker and the consumer worker may be co-located on a multi-core device at the first server and/or executed on different virtual machines, the different virtual machines configured to execute on the same physical hardware. The method may further include receiving a signal from a command channel associated with the datacenter queue and modifying the stored first message in response to receiving the signal. Modifying the stored first message may include deleting the first message.

**[0065]** According to further embodiments, the method may further include intercepting the message request from the consumer worker to the datacenter queue, forwarding the message request to the datacenter queue if a first criterion is met, and refraining from forwarding the message request to the datacenter queue if the first criterion is not met. The first criterion may include whether the datacenter queue is configured to hide a requested message upon receiving the message request from the consumer worker.

**[0066]** According to other examples, a virtual machine manager (VMM) for locally processing queue requests from co-located workers in a datacenter may include a queue usage detector module and a processing module. The queue usage detector module may be configured to detect a producer worker at a first server sending a first message to a datacenter queue at least partially stored at a second server and detect a consumer worker at the first server sending a message request to the datacenter queue. The processing module may be configured to store the first message at the first server and provide the stored first message to the consumer worker in response to the message request.

**[0067]** According to some embodiments, the processing module may be further configured to intercept the first mes-

sage sent by the producer worker before storing the first message. The VMM may further include a queue cache configured to store the first message. The queue cache may include a copy or a partial copy of the datacenter queue. The processing module may be further configured to build a table of queue usage based on at least one observed datacenter queue request, and the queue usage detector module may be configured to observe the at least one observed datacenter queue request.

**[0068]** According to other embodiments, the producer worker and the consumer worker may be co-located on a multi-core device at the first server and/or executed on different virtual machines, the different virtual machines configured to execute on the same physical hardware. The processing module may be further configured to receive a signal from a command channel associated with the datacenter queue and modify the stored first message in response to receiving the signal. Modifying the stored first message may include deleting the first message.

**[0069]** According to further embodiments, the processing module may be further configured to intercept the message request from the consumer worker to the datacenter queue, forward the message request to the datacenter queue if a first criterion is met, and refrain from forwarding the message request to the datacenter queue if the first criterion is not met. The first criterion may include whether the datacenter queue is configured to hide a requested message upon receiving the message request from the consumer worker.

**[0070]** According to further examples, a cloud-based datacenter is provided for locally processing queue requests from co-located workers. The datacenter may include a first and a second virtual machine (VM) operable to be executed on one or more physical machines and a datacenter controller. The datacenter controller may be configured to detect a producer worker executing on the first VM and sending a first message to a datacenter queue at least partially stored at a first datacenter location, store the first message at a second datacenter location different from the first, detect a consumer worker executing on the second VM and sending a message request to the datacenter queue, and provide the stored first message to the consumer worker in response to the message request, where the message is stored and provided from within a server local to the producer worker and the consumer worker.

**[0071]** According to some embodiments, the controller may be further configured to intercept the first message sent by the producer worker before storing the first message and/or store the first message in a queue cache at the second datacenter location. The queue cache may include a copy or a partial copy of the datacenter queue. The controller may be further configured to build a table of queue usage based on at least one observed datacenter queue request and/or observe the at least one observed datacenter queue request.

**[0072]** According to other embodiments, the producer worker and the consumer worker may be co-located on a multi-core device at the first datacenter location and/or executed on different virtual machines, the different virtual machines configured to execute on the same physical hardware. The controller may be further configured to receive a signal from a command channel associated with the datacenter queue and modify the stored first message in response to receiving the signal. Modifying the stored first message may include deleting the first message.

**[0073]** According to further embodiments, the controller may be further configured to intercept the message request

from the consumer worker to the datacenter queue, forward the message request to the datacenter queue if a first criterion is met, and refrain from forwarding the message request to the datacenter queue if the first criterion is not met. The first criterion may include whether the datacenter queue is configured to hide a requested message upon receiving the message request from the consumer worker.

**[0074]** According to yet further examples, a computer-readable storage medium may store instructions for locally processing queue requests from co-located workers in a datacenter. The instructions may include detecting a producer worker at a first server sending a first message to a datacenter queue at least partially stored at a second server, storing the first message at the first server, detecting a consumer worker at the first server sending a message request to the datacenter queue, and providing the stored first message to the consumer worker in response to the message request.

**[0075]** According to some embodiments, the instructions may further include intercepting the first message sent by the producer worker before storing the first message and/or storing the first message at the first server in a queue cache. The queue cache may include a copy or a partial copy of the datacenter queue. The instructions may further include building a table of queue usage based on at least one observed datacenter queue request and/or observing the at least one observed datacenter queue request.

**[0076]** According to other embodiments, the producer worker and the consumer worker may be co-located on a multi-core device at the first server and/or executed on different virtual machines, the different virtual machines configured to execute on the same physical hardware. The instructions may further include receiving a signal from a command channel associated with the datacenter queue and modifying the stored first message in response to receiving the signal. Modifying the stored first message may include deleting the first message.

**[0077]** According to further embodiments, the instructions may further include intercepting the message request from the consumer worker to the datacenter queue, forwarding the message request to the datacenter queue if a first criterion is met, and refraining from forwarding the message request to the datacenter queue if the first criterion is not met. The first criterion may include whether the datacenter queue is configured to hide a requested message upon receiving the message request from the consumer worker.

**[0078]** There is little distinction left between hardware and software implementations of aspects of systems; the use of hardware or software is generally (but not always, in that in certain contexts the choice between hardware and software may become significant) a design choice representing cost vs. efficiency tradeoffs. There are various vehicles by which processes and/or systems and/or other technologies described herein may be effected (e.g., hardware, software, and/or firmware), and that the preferred vehicle will vary with the context in which the processes and/or systems and/or other technologies are deployed. For example, if an implementer determines that speed and accuracy are paramount, the implementer may opt for a mainly hardware and/or firmware vehicle; if flexibility is paramount, the implementer may opt for a mainly software implementation; or, yet again alternatively, the implementer may opt for some combination of hardware, software, and/or firmware.

**[0079]** The foregoing detailed description has set forth various embodiments of the devices and/or processes via the use

of block diagrams, flowcharts, and/or examples. Insofar as such block diagrams, flowcharts, and/or examples contain one or more functions and/or operations, it will be understood by those within the art that each function and/or operation within such block diagrams, flowcharts, or examples may be implemented, individually and/or collectively, by a wide range of hardware, software, firmware, or virtually any combination thereof. In one embodiment, several portions of the subject matter described herein may be implemented via Application Specific Integrated Circuits (ASICs), Field Programmable Gate Arrays (FPGAs), digital signal processors (DSPs), or other integrated formats. However, those skilled in the art will recognize that some aspects of the embodiments disclosed herein, in whole or in part, may be equivalently implemented in integrated circuits, as one or more computer programs running on one or more computers (e.g., as one or more programs running on one or more computer systems), as one or more programs running on one or more processors (e.g., as one or more programs running on one or more microprocessors), as firmware, or as virtually any combination thereof, and that designing the circuitry and/or writing the code for the software and/or firmware would be well within the skill of one of skill in the art in light of this disclosure.

**[0080]** The present disclosure is not to be limited in terms of the particular embodiments described in this application, which are intended as illustrations of various aspects. Many modifications and variations can be made without departing from its spirit and scope, as will be apparent to those skilled in the art. Functionally equivalent methods and apparatuses within the scope of the disclosure, in addition to those enumerated herein, will be apparent to those skilled in the art from the foregoing descriptions. Such modifications and variations are intended to fall within the scope of the appended claims. The present disclosure is to be limited only by the terms of the appended claims, along with the full scope of equivalents to which such claims are entitled. It is to be understood that this disclosure is not limited to particular methods, reagents, compounds compositions or biological systems, which can, of course, vary. It is also to be understood that the terminology used herein is for the purpose of describing particular embodiments only, and is not intended to be limiting.

**[0081]** In addition, those skilled in the art will appreciate that the mechanisms of the subject matter described herein are capable of being distributed as a program product in a variety of forms, and that an illustrative embodiment of the subject matter described herein applies regardless of the particular type of signal bearing medium used to actually carry out the distribution. Examples of a signal bearing medium include, but are not limited to, the following: a recordable type medium such as a floppy disk, a hard disk drive, a Compact Disc (CD), a Digital Versatile Disk (DVD), a digital tape, a computer memory, a solid state drive, etc.; and a transmission type medium such as a digital and/or an analog communication medium (e.g., a fiber optic cable, a waveguide, a wired communications link, a wireless communication link, etc.).

**[0082]** Those skilled in the art will recognize that it is common within the art to describe devices and/or processes in the fashion set forth herein, and thereafter use engineering practices to integrate such described devices and/or processes into data processing systems. That is, at least a portion of the devices and/or processes described herein may be integrated into a data processing system via a reasonable amount of



experimentation. Those having skill in the art will recognize that a typical data processing system generally includes one or more of a system unit housing, a video display device, a memory such as volatile and non-volatile memory, processors such as microprocessors and digital signal processors, computational entities such as operating systems, drivers, graphical user interfaces, and applications programs, one or more interaction devices, such as a touch pad or screen, and/or control systems including feedback loops and control motors (e.g., feedback for sensing position and/or velocity of gantry systems; control motors for moving and/or adjusting components and/or quantities).

**[0083]** A typical data processing system may be implemented utilizing any suitable commercially available components, such as those typically found in data computing/communication and/or network computing/communication systems. The herein described subject matter sometimes illustrates different components contained within, or connected with, different other components. It is to be understood that such depicted architectures are merely exemplary, and that in fact many other architectures may be implemented which achieve the same functionality. In a conceptual sense, any arrangement of components to achieve the same functionality is effectively “associated” such that the desired functionality is achieved. Hence, any two components herein combined to achieve a particular functionality may be seen as “associated with” each other such that the desired functionality is achieved, irrespective of architectures or intermediate components. Likewise, any two components so associated may also be viewed as being “operably connected”, or “operably coupled”, to each other to achieve the desired functionality, and any two components capable of being so associated may also be viewed as being “operably couplable”, to each other to achieve the desired functionality. Specific examples of operably couplable include but are not limited to physically connectable and/or physically interacting components and/or wirelessly interactable and/or wirelessly interacting components and/or logically interacting and/or logically interactable components.

**[0084]** With respect to the use of substantially any plural and/or singular terms herein, those having skill in the art can translate from the plural to the singular and/or from the singular to the plural as is appropriate to the context and/or application. The various singular/plural permutations may be expressly set forth herein for sake of clarity.

**[0085]** It will be understood by those within the art that, in general, terms used herein, and especially in the appended claims (e.g., bodies of the appended claims) are generally intended as “open” terms (e.g., the term “including” should be interpreted as “including but not limited to,” the term “having” should be interpreted as “having at least,” the term “includes” should be interpreted as “includes but is not limited to,” etc.). It will be further understood by those within the art that if a specific number of an introduced claim recitation is intended, such an intent will be explicitly recited in the claim, and in the absence of such recitation no such intent is present. For example, as an aid to understanding, the following appended claims may contain usage of the introductory phrases “at least one” and “one or more” to introduce claim recitations. However, the use of such phrases should not be construed to imply that the introduction of a claim recitation by the indefinite articles “a” or “an” limits any particular claim containing such introduced claim recitation to embodiments containing only one such recitation, even when the

same claim includes the introductory phrases “one or more” or “at least one” and indefinite articles such as “a” or “an” (e.g., “a” and/or “an” should be interpreted to mean “at least one” or “one or more”); the same holds true for the use of definite articles used to introduce claim recitations. In addition, even if a specific number of an introduced claim recitation is explicitly recited, those skilled in the art will recognize that such recitation should be interpreted to mean at least the recited number (e.g., the bare recitation of “two recitations,” without other modifiers, means at least two recitations, or two or more recitations).

**[0086]** Furthermore, in those instances where a convention analogous to “at least one of A, B, and C, etc.” is used, in general such a construction is intended in the sense one having skill in the art would understand the convention (e.g., “a system having at least one of A, B, and C” would include but not be limited to systems that have A alone, B alone, C alone, A and B together, A and C together, B and C together, and/or A, B, and C together, etc.). It will be further understood by those within the art that virtually any disjunctive word and/or phrase presenting two or more alternative terms, whether in the description, claims, or drawings, should be understood to contemplate the possibilities of including one of the terms, either of the terms, or both terms. For example, the phrase “A or B” will be understood to include the possibilities of “A” or “B” or “A and B.”

**[0087]** In addition, where features or aspects of the disclosure are described in terms of Markush groups, those skilled in the art will recognize that the disclosure is also thereby described in terms of any individual member or subgroup of members of the Markush group.

**[0088]** As will be understood by one skilled in the art, for any and all purposes, such as in terms of providing a written description, all ranges disclosed herein also encompass any and all possible subranges and combinations of subranges thereof. Any listed range can be easily recognized as sufficiently describing and enabling the same range being broken down into at least equal halves, thirds, quarters, fifths, tenths, etc. As a non-limiting example, each range discussed herein can be readily broken down into a lower third, middle third and upper third, etc. As will also be understood by one skilled in the art all language such as “up to,” “at least,” “greater than,” “less than,” and the like include the number recited and refer to ranges which can be subsequently broken down into subranges as discussed above. Finally, as will be understood by one skilled in the art, a range includes each individual member. Thus, for example, a group having 1-3 cells refers to groups having 1, 2, or 3 cells. Similarly, a group having 1-5 cells refers to groups having 1, 2, 3, 4, or 5 cells, and so forth.

**[0089]** While various aspects and embodiments have been disclosed herein, other aspects and embodiments will be apparent to those skilled in the art. The various aspects and embodiments disclosed herein are for purposes of illustration and are not intended to be limiting, with the true scope and spirit being indicated by the following claims.

1. A method to locally process queue requests from co-located workers in a datacenter, the method comprising:

detecting a producer worker at a first server sending a first message to a datacenter queue at least partially stored at a second server;

storing the first message in a queue cache at the first server, wherein the queue cache includes one of a copy and a partial copy of the datacenter queue;



detecting a consumer worker at the first server sending a message request to the datacenter queue;  
 providing the stored first message to the consumer worker in response to the message request;  
 receiving a signal from a command channel associated with the datacenter queue; and  
 modifying the stored first message in response to receiving the signal.

2. The method of claim 1, further comprising intercepting the first message sent by the producer worker before storing the first message.

3.-6. (canceled)

7. The method of claim 1, wherein the producer worker and the consumer worker are co-located on a multi-core device at the first server.

8. The method of claim 1, wherein the producer worker and the consumer worker are executed on different virtual machines, the different virtual machines configured to execute on the same physical hardware.

9. (canceled)

10. The method of claim 1, wherein modifying the stored first message includes deleting the first message.

11. The method of claim 1, further comprising:

intercepting the message request from the consumer worker to the datacenter queue;

forwarding the message request to the datacenter queue if a first criterion is met; and

refraining from forwarding the message request to the datacenter queue if the first criterion is not met.

12. The method of claim 11, wherein the first criterion includes whether the datacenter queue is configured to hide a requested message upon receiving the message request from the consumer worker.

13. A virtual machine manager (VMM) to locally process queue requests from co-located workers in a datacenter, the VMM comprising:

a queue usage detector module configured to:

detect a producer worker at a first server, wherein the producer worker sends a first message to a datacenter queue at least partially stored at a second server; and

detect a consumer worker at the first server, wherein the consumer worker sends a message request to the datacenter queue, and wherein the producer worker and the consumer worker are co-located on a multi-core device at the first server; and

a processing module configured to:

intercept the first message sent by the producer worker;

store the first message at the first server;

provide the stored first message to the consumer worker in response to the message request;

receive a signal from a command channel associated with the datacenter queue; and

modify the stored first message in response to receiving the signal.

14.-16. (canceled)

17. The VMM of claim 13, wherein the processing module is further configured to build a table of queue usage based on at least one observed datacenter queue request.

18. The VMM of claim 17, wherein the queue usage detector module is further configured to observe the at least one observed datacenter queue request.

19. (canceled)

20. The VMM of claim 13, wherein the producer worker and the consumer worker are executed on different virtual

machines, the different virtual machines configured to execute on the same physical hardware.

21. (canceled)

22. (canceled)

23. The VMM of claim 13, wherein the processing module is further configured to:

intercept the message request from the consumer worker to the datacenter queue;

forward the message request to the datacenter queue if a first criterion is met; and

refrain from forwarding the message request to the datacenter queue if the first criterion is not met.

24. The VMM of claim 23, wherein the first criterion includes whether the datacenter queue is configured to hide the requested message upon receiving the message request from the consumer worker.

25. A cloud-based datacenter configured to locally process queue requests from co-located workers in the datacenter, the datacenter comprising:

a first and a second virtual machine (VM) operable to be executed on one or more physical machines; and

a datacenter controller configured to:

detect a producer worker that is executed on a first VM and sends a first message to a datacenter queue at least partially stored at a first datacenter location;

intercept the first message sent by the producer worker before storing the first message;

store the first message in a queue cache at a second datacenter location different from the first;

detect a consumer worker that is executed on a second VM and sends a message request to the datacenter queue;

provide the stored first message to the consumer worker in response to the message request, wherein the first message is stored and provided from within a server to the producer worker and the consumer worker;

receive a signal from a command channel associated with the datacenter queue; and

modify the stored first message in response to receiving the signal.

26. The datacenter of claim 25, wherein the controller is further configured to intercept the first message sent by the producer worker before storing the first message.

27. (canceled)

28. The datacenter of claim 25, wherein the queue cache includes one of a copy and a partial copy of the datacenter queue.

29. The datacenter of claim 25, wherein the controller is further configured to build a table of queue usage based on at least one observed datacenter queue request.

30. The datacenter of claim 29, wherein the controller is further configured to observe the at least one observed datacenter queue request.

31. The datacenter of claim 25, wherein the producer worker and the consumer worker are co-located on a multi-core device at the first datacenter location.

32. The datacenter of claim 25, wherein the first and second VMs are configured to execute on the same physical machine.

33. (canceled)

34. The datacenter of claim 25, wherein the controller is further configured to modify the stored first message by deleting the first message.

35. The datacenter of claim 25, wherein the controller is further configured to:

intercept the message request from the consumer worker to the datacenter queue;  
forward the message request to the datacenter queue if a first criterion is met; and  
refrain from forwarding the message request to the datacenter queue if the first criterion is not met.

**36.** The datacenter of claim **35**, wherein the first criterion includes whether the datacenter queue is configured to hide the requested message upon receiving the message request from the consumer worker.

**37.** (canceled)

\* \* \* \* \*