



(51) International Patent Classification:
G06F 9/54 (2006.01)

(21) International Application Number:

PCT/US2018/041268

(22) International Filing Date:

09 July 2018 (09.07.2018)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

15/645,724 10 July 2017 (10.07.2017) US

(71) Applicant: **NOKIA SOLUTIONS AND NETWORKS OY** [FI/FI]; Karaportti 3, 02610 Espoo (FI).

(71) Applicant (for LC only): **NOKIA USA INC.** [US/US]; 200 South Mathilda Avenue, Sunnyvale, California 94086 (US).

(72) Inventors: **STEIN, Manuel**; Lorenzstr. 10, 70435 Stuttgart (DE). **CHEN, Ruichuan**; Lorenzstr. 10, 70435 Stuttgart (DE).


(DE). **HILT, Volker**; Lorenzstr. 10, 70435 Stuttgart (DE). **RIMAC, Ivica**; Lorenzstr. 10, 70435 Stuttgart (DE). **AKKUS, Istemi Ekin**; Lorenzstr. 10, 70435 Stuttgart (DE).

(74) Agent: **NOKIA USA INC.** et al.; Stephen J. Wyse, Attn: Docket Administrator, Room 6E-264, 600-700 Mountain Avenue, Murray Hill, New Jersey 07974-0636 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(54) Title: EVENT HANDLING IN DISTRIBUTED EVENT HANDLING SYSTEMS

FIG. 1



(57) Abstract: The present disclosure generally discloses an event handling capability configured to support handling of events. The event handling capability may be configured to support handling of events in a distributed event handling system, which may use distributed queuing of events, distributed processing of events, and so forth. The distributed event handling system may be serverless cloud system or other type of distributed event handling system. The event handling capability may be configured to support handling of events in a distributed event handling system based on use of a message bus for queuing of events and based on use of hosts for queuing and processing of events.



(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

— *as to the identity of the inventor (Rule 4.17(i))*

Published:

— *with international search report (Art. 21(3))*

EVENT HANDLING IN DISTRIBUTED EVENT HANDLING SYSTEMS

TECHNICAL FIELD

The present disclosure relates generally to network technology and, more particularly but not exclusively, to event handling in event handling systems.

BACKGROUND

Event handling systems may be used to handle various types of events that are associated with various types of services or applications. Event handling systems may be centralized or distributed. Event handling in distributed event handling systems, such as serverless cloud systems and other types of distributed event handling systems, may be complicated.

SUMMARY

The present disclosure generally discloses event handling in event handling systems.

In at least some embodiments, an apparatus is provided. The apparatus is configured to support handling of events by a message bus of an event handling system. The apparatus includes a processor and a memory communicatively connected to the processor. The processor is configured to maintain, by the message bus for an event type, a respective set of event queues for the event type, wherein the event queues for the event type are associated with respective ones of a set of hosts subscribed to the message bus for the event type. The processor is configured to receive, by the message bus for the event type, an event of the event type that is published to the message bus. The processor is configured to store, by the message bus, the event of the event type in one of the event queues for the event type based on respective sets of queue status information associated with the respective event queues. The processor is configured to send, by the message bus, the event of the event type toward one of the hosts subscribed to the message bus for the event type based on a request from the one of the hosts subscribed to the message bus for the

event type. In at least some embodiments, a non-transitory computer-readable storage medium stores instructions which, when executed by a computer, cause the computer to perform a corresponding method for supporting handling of events by a message bus of an event handling system. In at least some embodiments, a corresponding method for supporting handling of events by a message bus of an event handling system is provided.

In at least some embodiments, an apparatus is provided. The apparatus is configured to support handling of events by a host of an event handling system. The apparatus includes a processor and a memory communicatively connected to the processor. The processor is configured to subscribe, by the host, to a first event queue of a message bus, the first event queue of the message bus configured to store events of a first event type to be processed by a first event processing function of the host. The processor is configured to subscribe, by the host, to a second event queue of the message bus, the second event queue configured to store events of a second event type to be processed by a second event processing function of the host. The processor is configured to maintain, by the host, a local buffer configured to store events received from the first event queue of the message bus and events received from the second event queue of the message bus. The processor is configured to retrieve, by the host based on a local queuing policy of the host and a local event retrieval condition detected at the host, a retrieved event comprising a first event from the first event queue of the message bus or a second event from the second event queue of the message bus. In at least some embodiments, a non-transitory computer-readable storage medium stores instructions which, when executed by a computer, cause the computer to perform a corresponding method for supporting handling of events by a host of an event handling system. In at least some embodiments, a corresponding method for supporting handling of events by a host of an event handling system is provided.

BRIEF DESCRIPTION OF THE DRAWINGS

The teachings herein can be readily understood by considering the following

detailed description in conjunction with the accompanying drawings, in which:

FIG. 1 depicts an event handling system configured to support distributed handling of events;

FIG. 2 depicts an event handling system configured to support distributed
5 handling of events;

FIG. 3 depicts an embodiment of a method for use by a message bus to support event handling;

FIG. 4 depicts an embodiment of a method for use by a host to support event handling;

FIG. 5 depicts an event handling system configured to support distributed
10 handling of events based on a distributed message bus; and

FIG. 6 depicts a high-level block diagram of a computer suitable for use in performing various functions presented herein.

To facilitate understanding, identical reference numerals have been used,
15 where possible, to designate identical elements that are common to the figures.

DETAILED DESCRIPTION

The present disclosure generally discloses an event handling capability configured to support handling of events. The event handling capability may be
20 configured to support handling of events in a distributed event handling system, which may use distributed queuing of events, distributed processing of events, and so forth. The distributed event handling system may be serverless cloud system or other type of distributed event handling system. It will be appreciated that these and various other embodiments and advantages or potential advantages of the event handling capability
25 may be further understood by way of reference to the example event handling system of FIG. 1.

FIG. 1 depicts an event handling system configured to support distributed handling of events.

The event handling system 100 includes a communication system 110, a set of
30 hosts 120-1 – 120-N (collectively, hosts 120), and a management system (MS) 130.

The communication system 110 may be any type of communication system configured to support communications of the hosts 120 (e.g., communication of event processing requests entering the event processing system 100 from external sources, communication of events from ones of the hosts 120 operating as publishers of the events, communication of events to ones of the hosts 120 operating as subscribers to the events, communication of event processing responses from hosts 120 based on processing of events by the hosts 120, or the like, as well as various combinations thereof) and communications of the MS 130 (e.g., communication of events being exchanged between hosts 120, communication of management messages associated with scaling of the hosts 120, or the like, as well as various combinations thereof). The communication system 110 may be a communication network (e.g., a wireline network, a wireless network, or the like, as well as various combinations thereof), a messaging system, or the like, as well as various combinations thereof.

The hosts 120 may be any types of hosts configured to support handling of events. The hosts 120 may be configured to be publishers of events, subscribers to events, or both. For example, the hosts may be physical hosts (e.g., servers), virtual hosts (e.g., virtual machines (VMs), virtual containers (VCs), or the like), or the like, as well as various combinations thereof). The hosts 120 may be distributed in various ways (e.g., physical hosts that are geographically distributed, virtual hosts that are distributed across physical servers, or the like, as well as various combinations thereof).

The MS 130 is configured to support distributed handling of events. The MS 130 includes a message bus (MB) 131 configured to control distribution of events across the hosts 120 and a function manager (FM) 132 configured to control scaling of hosts 120 supporting processing of events. The MB 131 may be configured to control distribution of events across the hosts 120 based on monitoring of queues of the MB 131 that are used to distribute events across the hosts 120, where the dynamics of the queues of the MB 131 that are used to distribute events across the hosts 120 may be based on backpressure resulting at the MB 131 from local handling of respective queues on the respective hosts 120 (e.g., each of the hosts 120 is able to locally and

independently determine its level of participation in processing of events in a manner that enables the MB 131 to intelligently control distribution of events across the hosts 120 without the use of explicit feedback from the hosts 120). The FM 132 may be configured to control scaling of hosts 120 supporting processing of events based on
5 feedback from the MB 131, supporting scale out and scale in of hosts 120 for event types. It will be appreciated that, although primarily presented with respect to embodiments in which the MB 131 and the FM 132 are part of a particular element (namely, MS 130), the MB 131 and the FM 132 may be standalone elements, may be distributed across various elements, or the like, as well as various combinations
10 thereof.

The event handling system 100 may be a serverless cloud system. In a serverless cloud system, which also may be referred to as a function-as-a-service (FaaS) cloud resources are utilized based on demand from the applications and services with the clients being charged based on the amount of resources consumed
15 and with the management of those resources being transparent or substantially transparent to the clients. In a serverless cloud system, the applications or services generally follow an event-driven programming model in which the event processing functions of the application or service are triggered by events within the serverless cloud system (e.g., a specific web browsing action by an end user, uploading of an
20 image to a storage server, or the like). In a serverless cloud system, the handling of a request to the application or service in order to provide an associated response from the application or service may be based on a graph of event processing functions that is configured to handle respective events associated with providing the response to the request, where the graph of event processing functions may be distributed across the
25 resources in various ways. For example, the event processing functions may be distributed across hosts 120 in various ways, such that the hosts 120 may exchange events in order to support completion of the response to the request (e.g., an output based on processing of an event by an event processing function of a host 120 may be provided as an input event to be processed by that host 120 and/or one or more other
30 hosts 120). It will be appreciated that the event handling system 100 may be any other

suitable type of event handling system in which events may be processed (e.g., a traditional cloud system in which clients request virtual resources to be used to handle processing of events or the like); however, various embodiments of the event handling capability are primarily presented herein within the context of a serverless cloud system.

The event handling system 100 is configured to support distributed handling of events. The events may be events of one or more applications or services supported by the event handling system 100. The events may be external arriving at the event handling system 100 from one or more external source of the events (e.g., requests from clients, requests from network devices, or the like), internal events generated within the event handling system 100 by one or more hosts 120 of the event handling system 100 (e.g., as discussed above, outputs by hosts 120 based on processing of events by event processing functions of the hosts 120 may be inputs, or events, to be processed by hosts 120), or the like, as well as various combinations thereof. The events may include events of one or more event types (e.g., multiple event types associated with respective applications or services, multiple event types supported by a given application or service, or the like, as well as various combinations thereof). For example, where the event handling system 100 supports a database service, events of a first event type may be requests to retrieve data from the database and events of a second event type may be requests to write data in the database. For example, where the event handling system supports a web browsing service, events may include external events (e.g., requests from clients) and internal events (e.g., requests for respective pieces of information to be assembled in order to return associated responses to the clients and processing to assemble the respective pieces of information in order to return associated responses to the clients). It will be appreciated that the event handling system 100 may be configured to handle various other types of events, applications or services, or the like, as well as various combinations thereof.

It will be appreciated that the event handling system 100, although primarily presented as being organized in a particular manner, may be organized in various other

ways (e.g., using other numbers, types, or arrangements of communication systems, using other numbers, types, or arrangement of hosts, or the like, as well as various combinations thereof).

FIG. 2 depicts an event handling system configured to support distributed
5 handling of events.

The event handling system 200 is configured to support handling of events. The event handling system 200 is depicted as being associated with a single cloud node or cloud location; however, the event handling system 200 may be associated with multiple cloud nodes or cloud locations.

10 The event handling system 200 is configured to support handling of two event types (primarily referred to herein as event type A and event type B). The two event types may be associated with a common service or application supported by the event handling system 200 (e.g., the service or application may require or support handling of the two event types), may be associated with two respective services or applications
15 supported by the event handling system 200 (e.g., a first service or application requires or supports handling of events of event type A and a second service or application requires or supports handling of events of event type B), may be associated with multiple services or applications supported by the event handling system 200 (e.g., each service or application may require or support handling of events of one or more
20 of the event types), or the like, as well as various combinations thereof. It will be appreciated that the event handling system, although primarily presented with respect to handling of two event types, may be configured to handle any number of event types for any number of services or applications.

The event handling system 200 includes a message bus (MB) 210, a pair of
25 hosts 220-1 and 220-2 (collectively, hosts 220), and a function manager (FM) 230. It will be appreciated that, although only two hosts 220 are depicted, fewer or more hosts 220 may be used (and, further that the number of hosts 220 used may be controlled dynamically). The MB 210 is configured to support queueing of events for the hosts 220 while the events are awaiting processing by the hosts 220. The hosts 220 are
30 configured to subscribe to the MB 210 to receive events from the MB 210. The hosts

220 are configured to receive events from the MB 210 and process the events from the MB 210. The FM 230 is configured to control configuration and scaling of the hosts 220 for supporting processing of events.

The MB 210 is configured to manage events that are awaiting processing by the hosts 220. The MB 210 is configured to receive events that are to be processed by the hosts 220. The MB 210 may receive events based on publication of the events to the MB 210 (e.g., by hosts 220, by one or more other hosts (not presented in FIG. 2), by one or more other sources of events (not presented in FIG. 2), or the like, as well as various combinations thereof). The MB 210 is configured to provide events to the hosts 220 responsive to requests by the hosts 220 to pull the events to the host 220.

The MB 210 includes an event dispatcher (ED) 211, a set of event queues (EQs) 212 (collectively, EQs 212), and a queue monitor (QM) 213.

The ED 211 is configured to receive events and dispatch the events to the EQs 212 for storage in the EQs 212 while the events await processing by the hosts 220. The ED 211 may receive events based on publication of the events to the MB 210. The ED 211 may be configured to determine queuing of events in the EQs 212. The ED 211 may be configured to determine queuing of events in the EQs 212 based on event type information of the events that are received (e.g., an event of an event type is queued within one of the EQs 212 of that event type). The ED 211 may be configured to determine queuing of events in the EQs 212 based on queue status information of the EQs 212 that is received from the QM 213. The queue status information may include various metrics which may be determined by the QM 213 based on monitoring of the EQs 212 by the QM 213. For example, the queue status information for an EQ 212 may include one or more of a queue length (L) of the EQ 212, an arrival rate (λ) of the EQ 212, a departure rate (μ) of the EQ 212, an expected queue waiting time ($E(dt)$) for an arriving event at the EQ 212, a change of a queue length (L) of the EQ 212 over a time interval (dL/dt), or the like, as well as various combinations thereof. The ED 211 may be configured to receive an event, determine an event type of the event, and determine queuing of the event in the EQs 212 based on the queue status information of the EQs 212 configured to store events of that event type (as discussed further

below, EQs 212 correspond to combinations of event types and hosts).

The ED 211 may be configured to queue events in the EQs 212 based on various portions or combinations of the queue status information of the EQs 212 which may be determined based on monitoring of the EQs 212 by QM 213. For example, the

5 ED 211 may be configured to queue an event based on queue lengths (L) of the EQs 212 that are associated with the event type of that event (e.g., ED 211, for an event of a particular event type, may be configured to dispatch the event to one of the EQs 212, associated with that event type, having a shortest queue length (L)). For example, the

10 ED 211 may be configured to queue an event based on expected queue waiting times ($E(dt)$) of the EQs 212 that are associated with the event type of that event (e.g., ED 211, for an event of a particular event type, may be configured to dispatch the event to one of the EQs 212, associated with that event type, having a shortest expected queue waiting time ($E(dt)$)). The ED 211 may be configured to queue events in the EQs 212 based on various other portions or combinations of the queue status information of the

15 EQs 212 which may be determined based on monitoring of the EQs 212 by QM 213.

The ED 211 may be configured to queue an event based on queue scores of the EQs 212 that are associated with the event type of that event. The ED 211 may be configured to determine queue scores for the EQs 212 of an event type based on various portions or combinations of the queue status information of the EQs 212,

20 respectively. The queue scores for the EQs 212 of an event type may be configured to enable the ED 211 to select an EQ 212 having a longer queue length over an EQ 212 having a shorter queue length when the queue length of the EQ 212 having the longer queue length and the queue length of the EQ 212 having the shorter queue length are changing at different rates (e.g., where the EQ 212 having the shorter queue length is

25 draining more slowly than the EQ 212 having the longer queue length). The ED 211 may be configured to compute the queue score (S) of an EQ 212 based on the queue length (L) of the EQ 212 and the arrival rate (λ) of the EQ 212. The ED 211 may be configured to compute the queue score (S) of an EQ 212 based on the queue length (L) of the EQ 212 and the departure rate (μ) of the EQ 212. For example, the ED 211 may

30 be configured to compute the queue score (S) of an EQ 212 using exponentially

weighted moving averages of queue length (L) and departure rate (μ) as follows: $S = \text{EWMA}(\mu) / \text{EWMA}(L)$, in which case the queue score is inversely proportional to an estimated waiting time of the event in the EQ 212. It will be appreciated that, although primarily discussed with respect to embodiments in which EWMA is used for smoothing, other types of averaging functions may be used for smoothing. The ED 211 may be configured to compute the queue scores of the EQs 212 based on various other portions or combinations of the queue status information of the EQs 212 which may be determined based on monitoring of the EQs 212 by QM 213 (e.g., using one or more of a queue length (L) of the EQ 212, an arrival rate (λ) of the EQ 212, a departure rate (μ) of the EQ 212, an expected queue waiting time (E(dt)) for an arriving event at the EQ 212, a change of a queue length (L) of the EQ 212 over a time interval (dL/dt), or the like, as well as various combinations thereof).

The ED 211 may be configured to support various other functions for receiving events and dispatching the events to the EQs 212 for storage in the EQs 212 while the events await processing by the hosts 220.

The EQs 212 are configured to receive events from the ED 211 and to store the events while the events await processing by hosts 220. The EQs 212 are configured to allow for addressing based on a combination of event type and host. Namely, the set of EQs 212 includes, for each of the event types, a respective EQ 212 for each of the hosts 220 that is subscribed to handle events of that respective event type. More specifically, the set of EQs 212 includes (1) a first pair of EQs 212-A1 and 212-A2 for event type A (collectively, EQs 212-A for event type A), where the EQ 212-A1 is maintained by the MB 210 for queuing events of event type A for host 220-1 and the EQ 212-A2 is maintained by the MB 210 for queuing events of event type A for host 220-2 and (2) a second pair of EQs 212-B1 and 212-B2 for event type B (collectively, EQs 212-B for event type B), where the EQ 212-B1 is maintained by the MB 210 for queuing events of event type B for host 220-1 and the EQ 212-B2 is maintained by the MB 210 for queuing events of event type B for host 220-2. The EQs 212 are configured to provide queued events to the hosts 220 responsive to requests from the hosts 220 (e.g., host 220-1 may request events of event type A from EQ 212-A1 and

may request events of event type B from EQ 212-B1 and host 220-2 may request events of event type A from EQ 212-A2 and may request events of event type B from EQ 212-B2). It will be appreciated that, although primarily presented with respect to a situation in which both of the hosts 220 are subscribed to handle both of the event types, each of the hosts 220 may be subscribed to handle one or more event types, the hosts 220 may be subscribed to handle different sets of event types (e.g., host 220-1 may be subscribed to handle events of event type A and event type B whereas host 220-2 may be subscribed to handle events of event type B only), or the like, as well as various combinations thereof.

10 The QM 213 is configured to monitor the EQs 212 for use in supporting handling of events (e.g., for supporting ED 211 in controlling distribution of events to the EQs 212, for supporting FM 230 in controlling scaling of the hosts 220 for supporting processing of events, or the like, as well as various combinations thereof). The QM 213 may be configured to, for each of the EQs 212, determine queue status
15 information of the EQ 212 based on monitoring of the EQ 212 (e.g., monitoring the queue length (L) of the EQ 212, monitoring the manner in which events move through the EQ 212 (e.g., monitoring the arrival rate (λ) of events to the EQ 212, monitoring the departure rate (μ) of events from the EQ 212, or both), or the like, as well as various combinations thereof). The queue status information for an EQ 212 may
20 include a queue length (L) of the EQ 212, an arrival rate (λ) of events to the EQ 212, a departure rate (μ) of events from the EQ 212, an expected queue waiting time ($E(dt)$) for an arriving event at the EQ 212, a change of a queue length (L) of the EQ 212 over a time interval (dL/dt), or the like, as well as various combinations thereof. The QM
25 213 may be configured to compute an expected queue waiting time ($E(dt)$) for an arriving event at an EQ 212 by assuming that the estimated departure rate (μ) of the EQ 212 is not changing significantly before the event leaves the EQ 212: $E(dt) = L/\mu$. The QM 213, assuming or under certain conditions, may be configured to estimate a queue status parameter based on monitoring and measurement of other queue status parameters (e.g., assuming that the event arrival and departure rates change only
30 slowly so that both are approximately constant over a short time interval dt , then it is

possible to measure the change of queue length dL over the same interval dt and relate the metrics as follows: $dL/dt = \lambda - \mu$ (such that, by measuring two of the three values, the third value may be estimated)). The QM 213 may be configured to determine various other types of queue status information for the EQs 212 based on monitoring
5 of the EQs 212. The QM 213 may be configured to provide the queue status information of the EQs 212 to the ED 211 for use in controlling queueing of events in the EQs 212. The QM 213 may be configured to provide the queue status information of the EQs 212 to the FM 230 for use in determining scaling of the hosts 220. The QM 213 also may be configured to determine per-host queuing information for the
10 hosts 220 (e.g., over the sets of EQs 212 associated with the hosts 220) and provide the per-host queuing information to the FM 230 for use in determining scaling of the hosts 220.

The hosts 220 are configured to support handling of events. The hosts 220 are configured to support configuration of the hosts 220 to support event handling (e.g.,
15 obtaining event processing function capabilities for respective event types to be supported by the respective hosts 220, subscribing to event types to be supported by the respective hosts 220, configuring the hosts 220 to be able to store events to be processed at the hosts 220, or the like, as well as various combinations thereof), pull events from the MB 210 for event types supported by the respective hosts 220, queue
20 the events pulled from the MB 210 for event types supported by the respective hosts 220, and process the events pulled from the MB 210 for event types supported by the respective hosts 220.

The hosts 220 include respective sets of elements configured to enable the hosts 220 to support handling of events. The hosts 220 include respective event
25 buffers 221 (illustratively, the host 220-1 includes an event buffer 221-1 and host 220-2 includes an event buffer 221-2, which are local buffers local to the hosts 220-1 and 220-2, respectively) configured to store events retrieved by the hosts 220 for processing by the hosts 220, respectively. The hosts 220 include respective processing elements 225 (illustratively, the host 220-1 includes a processing element 225-1 and
30 the host 220-2 includes a processing element 225-2) configured to enable the hosts 220

to provide various functions discussed herein (e.g., configuration of the hosts 220 to support event handling, pulling events from the MB 210 for event types supported by the respective hosts 220, queuing the events pulled from the MB 210 for event types supported by the respective hosts 220, processing the events pulled from the MB 210 for event types supported by the respective hosts 220, or the like, as well as various combinations thereof). The hosts 220 include respective storage elements 226 (illustratively, the host 220-1 includes a storage element 226-1 and the host 220-2 includes a storage element 226-2) storing host-specific event pulling/queuing policies 227 of the hosts 220 (illustratively, the storage element 226-1 of host 220-1 stores a host-specific event pulling/queuing policy 227-1 of the host 220-1 and the storage element 226-2 of host 220-2 stores a host-specific event pulling/queuing policy 227-2 of the host 220-2), host-specific event selection policies 228 (illustratively, the storage element 226-1 of host 220-1 stores a host-specific event selection policy 228-1 of the host 220-1 and the storage element 226-2 of host 220-2 stores a host-specific event selection policy 228-2 of the host 220-2), and event processing functions 229 (illustratively, the storage element 226-1 of the host 220-1 stores an event processing function 229-1A configured to be executed by the processing element 225-1 of the host 220-1 to process events of event type A and an event processing function 229-1B configured to be executed by the processing element 225-1 of the host 220-1 to process events of event type B and the storage element 226-2 of the host 220-2 stores an event processing function 229-2A configured to be executed by the processing element 225-2 of the host 220-2 to process events of event type A and an event processing function 229-2B configured to be executed by the processing element 225-2 of the host 220-2 to process events of event type B). The operation of the hosts 220 in supporting handling of events received from MB 210 is discussed further below.

The hosts 220 are configured to support configuration of the hosts 220 to support event handling. A host 220 may be configured to support event handling for events of a particular event type by obtaining the event processing function 229 that is configured to support processing of events of the particular event type, subscribing to the MB 210 for the particular event type such that the host 220 may retrieve events of

the particular event type from MB 210 based on the host-specific event pulling/queuing policy 227 of the host 220, configuring the event buffer 221 of the host 220 to support storage of events of the particular event type at the host 220, or the like, as well as various combinations thereof. A host 220 may be configured to support event handling

5 for events of a particular event type responsive to a request by the FM 230 for the host 220 to handle events of the particular event type. It is noted that the various functions used to configure a host 220 to support event handling for events of a particular event type may be performed with various timings (e.g., serially, contemporaneously, or the like), responsive to various conditions, or the like, as well as various combinations

10 thereof. For example, a host 220 may, responsive to a request by the FM 230 for the host 220 to handle events of a particular event type, obtain the event processing function 229 that is configured to support processing of events of the particular event type and subscribe to the MB 210 for the particular event type such that the host 220 may retrieve events of the particular event type from MB 210. For example, a host 220

15 may obtain the event processing function 229 that is configured to support processing of events of the particular event type responsive to a request by the FM 230 for the host 220 to handle events of the particular event type, subscribe to receive events of the particular event type from MB 210 responsive to a determination that the host 220 has obtained the event processing function 229 that is configured to support processing of

20 events of the particular event type, and so forth. For example, a host 220 may obtain the event processing function 229 that is configured to support processing of events of the particular event type responsive to a request by the FM 230 for the host 220 to handle events of the particular event type, configure the event buffer 221 of the host 220 to support storage of events of the particular event type at the host 220 responsive

25 to a determination that the host 220 has obtained the event processing function 229 that is configured to support processing of events of the particular event type, subscribe to receive events of the particular event type from MB 210 responsive to a determination that the host 220 has configured the event buffer 221 of the host 220 to support storage of events of the particular event type at the host 220, and so forth. It will be

30 appreciated that configuration of a host 220 to support event handling may be

performed in other ways (e.g., using various other timings, responsive to various other conditions, or the like, as well as various combinations thereof).

The hosts 220 are configured to pull events from the MB 210 and queue events pulled from the MB 210. The hosts 220 may be configured to locally control

5 pulling/queuing of events from the EQs 212 of the MB 210 subscribed to by the hosts 220. The hosts 220 may be configured to locally control pulling/queuing of events from the EQs 212 of the MB 210 subscribed to by the hosts 220 based on the host-specific event pulling/queuing policies 227 of the hosts 220, respectively. The host-specific event pulling/queuing policies 227, as presented in FIG. 2, may be different

10 for the hosts 220. Namely, as presented in FIG. 2, (1) the host 220-1 uses an event buffer 221-1 that includes a pair of host event queues (HEQs) 222-A and 222-B (collectively, HEQs 222) configured to store events of event type A and events of event type B, respectively and (2) the host 220-2 uses an event buffer 221-2 that is implemented as a first-in-first-out (FIFO) buffer configured to store events of event

15 type A and events of event type B. The host-specific event pulling/queuing policies 227, although omitted from FIG. 2 for purposes of clarity, may be the same for the hosts 220. The hosts 220 may be configured to locally control pulling/queuing of events from the EQs 212 of the MB 210 subscribed to by the hosts 220 based on local event retrieval conditions detected at the hosts 220. For example, for a given host 220,

20 the local event retrieval condition(s) detected at the host 220 may include detection that a current load at the host 220 satisfies a threshold, detection of a characteristic(s) of events already pulled from the MB 210 by the host 220 or events available to be pulled from the MB 210 by the host 220, detection of one or more capabilities of the host 220, or the like, as well as various combinations thereof. The local event retrieval

25 conditions detected at the hosts 220 may be detected based on host-specific information available at the hosts 220, respectively. For example, the host-specific information available at the hosts 220 may include current load levels at the hosts 220, one or more characteristics of events already pulled from the MB 210 by the hosts 220 or events available to be pulled from the MB 220 by the hosts 220, or the like, as well

30 as various combinations thereof. The hosts 220 may be configured to locally control

pulling/queuing of events from the EQs 212 of the MB 210 subscribed to by the hosts 220 based on local event retrieval conditions detected at the hosts 220 using various combinations of local event retrieval conditions, host-specific information, or the like, as well as various combinations thereof. For example, a host 220 may decide to pull
5 memory-hungry events from MB 210 when the host 220 is CPU bottlenecked, may decide to pull CPU-hungry events from the MB 210 when the host 220 is memory bottlenecked, or the like. In this manner, decisions by the host 220 regarding pulling/queuing of events from the EQs 212 of the MB 210 remain local to the hosts 220 such that they do not need to be communicated between the hosts 220 and, thus,
10 do not require any communication or coordination between the hosts 220 for this purpose. It is noted that local decisions of the hosts 220 in pulling/queuing events from the MB 210 produce backpressures on the EQs 212 in the MB 210, which are monitored and summarized in the queue status information of the EQs 212 (e.g., the queue lengths and associated queue dynamics).

15 The hosts 220 are configured to process events queued at the hosts 220. The hosts 220 process events by retrieving events from the event buffers 221 based on the host-specific event selection policies 228 and processing the events using the event processing functions 229.

The hosts 220 are configured to retrieve events from the event buffers 221 for
20 processing by the hosts 220. The hosts 220 may be configured to locally control selection of events from the event buffers 221 of the hosts 220. The hosts 220 may be configured to locally control selection of events from the event buffers 221 of the hosts 220 based on the host-specific event selection policies 228 of the hosts 220, respectively. The host-specific event selection policies 228 may be different for the
25 hosts 220, may be the same for the hosts 220, or the like. The host-specific event selection policies 228 of the host 220 may be at least partially based on the host-specific event pulling/queuing policies 227 of the hosts (e.g., the manner in which events are selected for processing at host 220-1 may be at least partially based on the fact that the host 220-1 maintains separate HEQs 222 for queuing of different event
30 types, the manner in which events are selected for processing at host 220-2 may be at

least partially based on the fact that the host 220-2 maintains a single FIFO queue for queuing of different event types, and so forth). For example, host 220-1 may be configured to select between processing events of event type A and processing events of event type B since the host 220-1 uses the event buffer 221-1 that includes the

5 HEQs 222-A and 222-B configured to store events of event type A and events of event type B, respectively. For example, host 220-2 may be configured to process events of event type A and events of event type B based on the ordering of the events within the event buffer 221-2 since the event buffer 221-2 is implemented as a FIFO buffer. The hosts 220 may be configured to locally control selection of events from the event

10 buffers 221 of the hosts 220 based on host-specific information available at the hosts 220, such as host resources available at the hosts 220, host resources needed for processing events at the hosts 220, one or more characteristics of events awaiting processing by the hosts 220, or the like, as well as various combinations thereof. For example, a host 220 may decide to process memory-hungry events when the host 220

15 is CPU bottlenecked, may decide to process CPU-hungry events when the host 220 is memory bottlenecked, or the like. The hosts 220 may be configured to locally control selection of events from the event buffers 221 of the hosts 220 based on the host-specific event selection policies 228 and local event selection conditions detected at the hosts 220 (e.g., a current load at a host being below a threshold, detection of a

20 particular characteristic of events awaiting processing by the hosts 220, or the like, as well as various combinations thereof). In this manner, decisions by the host 220 regarding processing of events from MB 210 at the hosts 220 remain local to the hosts 220 such that they do not need to be communicated between the hosts 220 and, thus, do not require any communication or coordination between the hosts 220 for this

25 purpose.

The hosts 220 are configured to process the events, retrieved from the event buffers 221, using the event processing functions 229 available on the hosts 220. For example, the host 220-1 is configured to (1) process events of event type A, which are retrieved from the HEQ 222-A of host 220-1 based on the host-specific event selection

30 policy 228-1 of the host 220-1, based on execution of the event processing function

229-1A by the processing element 225-1 of the host 220-1 and (2) process events of event type B, which are retrieved from the HEQ 222-B of host 220-1 based on the host-specific event selection policy 228-1 of the host 220-1, based on execution of the event processing function 229-1B by the processing element 225-1 of the host 220-1.

5 Similarly, for example, the host 220-2 is configured to (1) process events of event type A, which are retrieved from the event buffer 221-2 of host 220-2 based on the host-specific event selection policy 228-2 of the host 220-2, based on execution of the event processing function 229-2A by the processing element 225-2 of the host 220-2 and (2)

10 process events of event type B, which are retrieved from the event buffer 221-2 of host 220-2 based on the host-specific event selection policy 228-2 of the host 220-2, based on execution of the event processing function 229-2B by the processing element 225-2 of the host 220-2. The processing of the events by the hosts 220 may trigger various related actions (e.g., storage or sending of state information, storage or sending of processing results, storage or sending of processing responses, or the like, as well as

15 various combinations thereof) which may vary based on various factors (e.g., the service or application being supported, event types of the events, or the like, as well as various combinations thereof. The hosts 220 are configured to control processing of events locally at the hosts 220 such that the processing of events by the hosts 220 does not require any communication or coordination between the hosts 220. The hosts 220

20 each may be configured to monitor their own resource usages locally and to apply allocation policies suitable for utilizing the local resources of the hosts 220, with the monitoring and allocation of any particular host 220 being independent of other hosts 220. In this manner, decisions by the hosts 220 regarding processing of events at the hosts 220 remain local to the hosts 220 such that they do not need to be communicated

25 between the hosts 220 and, thus, do not require any communication or coordination between the hosts 220 for this purpose.

The FM 230 is configured to control configuration and scaling of the hosts 220 for supporting processing of events. The FM 230 may be configured to control configuration of the hosts 220 for supporting processing of events (e.g., requesting or

30 instructing hosts 220 to obtain event processing functions 229 for particular event

types, to subscribe to the MB 210 for particular event types, or the like, as well as various combinations thereof). The scaling may include scale out or scale in of the hosts 220. The scaling of hosts 220 may include configuration of existing hosts 220, instantiation of new hosts 220, termination of existing hosts 220, or the like, as well as various combinations thereof. The configuration of existing hosts 220 may include configuring a host 220 to support processing of an event type (e.g., requesting or instructing the host 220 to obtain the event processing function 229 for the event type and to subscribe to receive events of that event type from the MB 210 such that the MB 210 instantiates an EQ 212 for that host 220 for that event type), configuring a host 220 to stop supporting processing of an event type (e.g., requesting or instructing the host 220 to remove the event processing function 229 for the event type and to unsubscribe from receiving events of that event type from the MB 210 such that the MB 210 removes the EQ 212 for that host 220 for that event type), or the like, as well as various combinations thereof. The FM 230 may be configured to control scaling of the hosts 220 for supporting processing of events based on information received from the QM 213 of MB 210 (e.g., queue status information of the EQs 212, per-host queuing information of the hosts 220, or the like, as well as various combinations thereof).

It will be appreciated that the event handling system 200, although primarily presented as being organized in a particular manner, may be organized in various other ways (e.g., using other numbers, types, or arrangements of hosts or host capabilities, message bus capabilities, or the like, as well as various combinations thereof).

FIG. 3 depicts an embodiment of a method for use by a message bus to support event handling. It will be appreciated that, although primarily presented herein as being performed serially, at least a portion of the functions of method 300 may be performed contemporaneously or in a different order than as presented in FIG. 3.

At block 301, method 300 begins.

At block 310, maintain, by the message bus for an event type, a respective set of event queues for the event type, wherein the event queues for the event type are associated with respective ones of a set of hosts subscribed to the message bus for the

event type.

At block 320, receive, by the message bus for the event type, an event of the event type that is published to the message bus.

At block 330, store, by the message bus, the event of the event type in one of
5 the event queues for the event type based on respective sets of queue status information associated with the respective event queues.

At block 340, send, by the message bus, the event of the event type toward one of the hosts subscribed to the message bus for the event type based on a request from the one of the hosts subscribed to the message bus for the event type.

10 At block 399, method 300 ends.

FIG. 4 depicts an embodiment of a method for use by a host to support event handling. It will be appreciated that, although primarily presented herein as being performed serially, at least a portion of the functions of method 400 may be performed contemporaneously or in a different order than as presented in FIG. 4.

15 At block 401, method 400 begins.

At block 410, subscribe, by the host, to a first event queue of a message bus, the first event queue of the message bus configured to store events of a first event type to be processed by a first event processing function of the host.

At block 420, subscribe, by the host, to a second event queue of the message
20 bus, the second event queue configured to store events of a second event type to be processed by a second event processing function of the host.

At block 430, maintain, by the host, a local buffer configured to store events received from the first event queue of the message bus and events received from the second event queue of the message bus.

25 At block 440, retrieve, by the host based on a local queuing policy of the host and a local event retrieval condition detected at the host, a retrieved event comprising a first event from the first event queue of the message bus or a second event from the second event queue of the message bus.

At block 499, method 400 ends.

30 It will be appreciated that, although primarily presented with respect to

embodiments in which the event handling system includes a centralized message bus (illustratively, MB 210 of event handling system 200), in at least some embodiments the event handling system may include a distributed message bus. An example of an event handling system using a distributed message bus is presented in FIG. 5.

5 FIG. 5 depicts an event handling system configured to support distributed handling of events based on a distributed message bus.

 The event handling system 500 is configured to support handling of events. The event handling system 500 is depicted as being associated with two cloud nodes or cloud locations; however, the event handling system 500 may be associated with a
10 single cloud node or cloud location.

 The event handling system 500 is configured to support handling of two event types (primarily referred to herein as event type A and event type B).

 The event handling system 500 includes a global message bus (GMB) 510-G, a pair of edge cloud nodes (ECNs) 520-1 and 520-2 (collectively, ECNs 520), and a
15 management system (MS) 530. It will be appreciated that, although only two ECNs 520 are depicted, fewer or more ECNs 520 may be used.

 The GMB 510-G includes a global event dispatcher (GED) 511-G, a set of global event queues (GEQs) 512 (collectively, GEQs 512), and a queue monitor (QM) 513.

20 The ECNs 520 each include a local message bus (LMB) 521 (illustratively, ECN 520-1 includes an LMB 521-1 and ECN 520-2 includes an LMB 521-2) and a respective plurality of hosts 525 (illustratively, ECN 520-1 includes a plurality of hosts 525-11 – 525-1N and ECN 520-2 includes a plurality of hosts 525-21 – 525-2M). The LMBs 521 each include a local event dispatcher (LED) 522 (illustratively, LMB 521-1
25 includes an LED 522-1 and LMB 521-2 includes an LED 522-2) and a set of local event queues (LEQs) 523 (illustratively, LMB 521-1 includes a set of LEQs 523-1 and LMB 521-2 includes a set of LEQ 523-2).

 The MS 530 includes a function manager 531. It will be appreciated that the MS 530 may be configured to include various other types of managers which may
30 provide various other types of functions.

The GED 511-G is configured to receive events and dispatch the events to the GEQs 512 for storage in the GEQs 512 while the events await transfer to the LMBs 521 of the ECNs 520 and, eventually, processing by the hosts 525 of the ECNs 520. The GED 511-G may be configured to determine queuing of events in the GEQs 512 based on event type information of the events that are received. The GED 511-G may be configured to determine queuing of events in the GEQs 512 based on queue status information of the GEQs 512 that is received from the QM 513 (e.g., queue lengths, event arrival rates, event departure rates, or the like, as well as various combinations thereof). The operation of the GED 511-G in dispatching events to the GEQs 512 may be similar to the operation of ED 211 in dispatching events to the EQs 212 (presented with respect to FIG. 2).

The GEQs 512 are configured to receive events from the GED 511-G and to store the events while the events await transfer to the LMBs 521 of the ECNs 520 and, eventually, processing by hosts 525 of the ECNs 520. The GEQs 512, as presented in FIG. 5, may be configured to allow for addressing based on event type. Namely, the set of GEQs 512 includes, for each of the event types, a respective GEQ 512 for the respective event type (illustratively, the GMB 511-G includes a GEQ 512-1 that stores events of event type A and a GEQ 512-1 that stores events of event type B). The GEQs 512, although omitted from FIG. 5, may be configured to allow for addressing based on a combination of event type and edge cloud node (e.g., the set of GEQs 512 may include, for each of the event types, a respective GEQ 512 for each of the ECNs 520 (and, thus, for each of the LMBs 521 of the ECNs 520) that is subscribed to handle events of that respective event type. The GEQs 512 are configured to provide queued events to the LMBs 521 of the ECNs 520 responsive to requests from the LMBs 521 of the ECNs 520 (e.g., LMB 521-1 of ECN 520-1 may request events of event type A from GEQ 512-1 and LMB 521-2 of ECN 520-2 may request events of event type B from GEQ 512-2). It will be appreciated that, although primarily presented with respect to a situation in which each ECN 520 is only subscribed to handle events of a single event type (namely, ECN 520-1 being subscribed to handle events of event type A and ECN 520-2 being subscribed to handle events of event type B

B), one or more of the ECNs 520 may be subscribed to handle events of a multiple event types.

The QM 513 is configured to monitor the GEQs 512 for use in supporting handling of events (e.g., for supporting GED 511-G in controlling distribution of events to the GEQs 512, for supporting FM 531 in controlling scaling of the hosts 525 for supporting processing of events, or the like, as well as various combinations thereof). The operation of the QM 513 of GMB 511-G may be similar to the operation of QM 213 of MB 210 (presented with respect to FIG. 2).

The ECNs 520 are configured to support queuing and processing of events. The events are queues by the LMBs 521 of the ECNs 520 and the events are processed by the hosts 525 of the ECNs 520.

The LMBs 521 of the ECNs 520 are configured to manage events that are awaiting processing by the hosts 525 of the ECNs 520, respectively. The LMBs 521 of the ECNs 520 are configured to receive events that are to be processed by the hosts 525 of the ECNs 520 and to provide the events to the hosts 525 of the ECNs 520 for processing by the hosts 525 of the ECNs 520.

The LMBs 521 of the ECNs 520 are configured to subscribe to the GMB 511-G in order to retrieve events from the GEQs 512 of the GMB 511-G (illustratively, the LMB 521-1 is subscribed to GMB 511-G for event type A while the LMB 521-2 is subscribed to GMB 511-G for event type B). The LMBs 521 of the ECNs 520 may be configured to subscribe to the GMB 511-G in order to retrieve events from the GEQs 512 of the GMB 511-G under the control of the FM 531 (e.g., based on a request or instruction for an ECN 520 to support handling of events of the event type, based on a request or instruction for one or more of the hosts 525 of an ECN 520 to support handling of events of the event type, or the like). The operation of the LMBs 521 of the ECNs 520 in pulling events from the GMB 511-G may be similar to the operation of the hosts 220 in pulling events from the EQs 212 of the MB 210 (e.g., based on local policies of the ECNs 520, based on local conditions at the ECNs 520, or the like, as well as various combinations thereof).

The LEDs 522 of the LMBs 521 are configured to receive events from the

GEQs 512 of the GMB 511-G and to dispatch the events to the LEQs 523 for storage in the LEQs 523 while the events await transfer to the hosts 525 of the ECNs 520 to be processed by the hosts 525 of the ECNs 520. The LEDs 522 may be configured to determine queuing of events in the LEQs 523 based on queue status information of the LEQs 523 (e.g., queue lengths, event arrival rates, event departure rates, or the like, as well as various combinations thereof), host status information associated with hosts 525 subscribed to or otherwise associated with the LEDs 522, or the like, as well as various combinations thereof. The operation of the LEDs 522 in dispatching events to the LEQs 523 may be similar to the operation of ED 211 in dispatching events to the EQs 212 (presented with respect to FIG. 2).

The LEQs 523 are configured to receive events from the LEDs 522 and to store the events while the events await transfer to the hosts 525 of the ECNs 520 to be processed by the hosts 525 of the ECNs 520. The LEQs 523, as presented in FIG. 5, may be configured to allow for addressing based on hosts 525. Namely, the set of LEQs 523 includes, for each of the hosts 525 subscribed to retrieve events from the LMB 521, a respective LEQ 523 for the respective host (illustratively, the LMB 521-1 of ECN 520-1 includes a set of LEQs 523-1 that includes an LEQ 523-1 for host 525-11 and an LEQ 523-1 for host 525-1N and the LMB 521-2 of ECN 520-2 includes a set of LEQs 523-2 that includes an LEQ 523-2 for host 525-21). The LEQs 523 are configured to provide queued events to the hosts 525 responsive to requests from the hosts 525 (e.g., host 525-11 may request events of event type A from the LEQ 523-1 to which the host 525-11 is subscribed, host 525-1N may request events of event type A from the LEQ 523-1 to which the host 525-1N is subscribed, and the host 525-21 may request events of event type B from the LEQ 523-2 to which the host 525-21 is subscribed).

The host 525 are configured to support handing of events. The hosts 525 may be configured to support configuration of the hosts 525 to support event handling (e.g., obtaining event processing functions for respective event types to be supported by the respective hosts 525, subscribing to event types to be supported by the respective hosts 525, configuring the hosts 525 to be able to store events to be processed at the hosts

525, or the like, as well as various combinations thereof), pull events from the LMBs 521 for event types supposed by the respective hosts 525, queue the events pulled from the LMBs 521 for event types supported by the respective hosts 525, and process the events pulled from the LMBs 521 for event types supported by the respective hosts 525. The details of the host 525 are omitted for purposes of clarity; however, it will be appreciated that the hosts 525 may be configured in a manner similar to the hosts 220 (presented with respect to FIG. 2). It will be appreciated that the operation of the hosts 525 may be similar to the operation of the hosts 220 (presented with respect to FIG. 2); however, rather than pulling events from a centralized message bus, the hosts 525 pull events based on a distributed message bus (e.g., pulling events from local message buses of the distributed message bus where the local message buses pull events from the global message bus of the distributed message bus).

The GMB 511-G and LMBs 521 cooperate to provide a distributed message bus for the hosts 525. In the distributed message bus, events are hierarchically passed from the GMB 511-G to the LMBs 521.

The GMB 511-G may be configured to maintain a variable number of GEQs 512 for an event type or a constant number of GEQs 512 for an event type (e.g., subsets of which may be mapped to different ECNs 520). For example, where the number of GEQs 512 is constant (e.g., assuming 10 total queues) and events of event type A are only being handled by the ECN 520-1 then, based on a decision to scale out handling of event type A to ECN 520-2, the mapping of the ten GEQs 512 may be changed such that x GEQs 512 are mapped to ECN 520-1 and 10-x GEQs 512 are mapped to ECN 520-2. It is noted that this control over scaling of the GMB 511-G may be provided by the FM 531 or other suitable element.

The LMBs 521 each may be configured to maintain a variable number of LEQs 523 for an event type or a constant number of LEQs 523 for an event type (e.g., subsets of which may be mapped to different hosts 525 or subsets of hosts 525). For example, where the number of LEQs 523-2 of ECN 520-2 is variable and events of event type B are only being handled by the host 525-21 of the ECN 520-2 then, based on a decision to scale out handling of event type B to host 525-2M of ECN 520-2, the

set of LEQs 512 may be changed such that an additional LEQ 523-2 is instantiated on LMB 521-2 of ECN 520-2 for host 525-2M. It is noted that this control over scaling of the LMB 521 may be provided locally at the ECN 520-2.

The ECNs 520 may be configured to control configuration and scaling of the
5 hosts 525, for supporting processing of events, locally within the ECNs 520,
respectively. The ECNs 520 may be configured to control configuration of the hosts
525 of the ECNs 520 for supporting processing of events (e.g., requesting or
instructing hosts 525 to obtain event processing functions for particular event types, to
subscribe to LMBs 521 for particular event types, or the like, as well as various
10 combinations thereof). The scaling may include scale out or scale in. The scaling of
hosts 525 may include configuration of existing hosts 525, instantiation of new hosts
525, termination of existing hosts 525, or the like, as well as various combinations
thereof. The localized operation of each of the ECNs 520 in controlling configuration
and scaling of the respective hosts 525 of the ECNs 520, for supporting processing of
15 events by the respective hosts 525 of the ECNs 520, may be similar to the operation of
the FM 230 in controlling configuration and scaling of the hosts 220 for supporting
processing of events (presented in FIG. 2).

The FM 531 is configured to control deployment of event processing functions
to ECNs 520. The FM 531 may be configured to control the deployment of event
20 processing functions to ECNs 520 based on information from the QM 513. The FM
531 may be configured to support configuration and scaling of the hosts 525 for
supporting processing of events (e.g., requesting or instructing hosts 525 to obtain
event processing functions for particular event types, to subscribe to LMBs 521 for
particular event types, or the like, as well as various combinations thereof) if
25 centralized control over configuration and scaling of the hosts 525 is to be used in
addition to localized control over configuration and scaling of the hosts 525 by the
respective ECNs 520. The FM 531 is configured to provide various other
management functions.

It will be appreciated that the event handling system 500, although primarily
30 presented as being organized in a particular manner, may be organized in various other

ways (e.g., using other numbers, types, or arrangements of hosts or host capabilities, message bus capabilities, or the like, as well as various combinations thereof).

Various embodiments of the event handling capability, as discussed above, may be configured to solve various problems associated with handling of events for a service or application in a distributed event handling system.

Various embodiments of the event handling capability may be configured to support event handling based on distribution of information collection and decision making between the hosts and a message bus including a global message queue that supports the hosts. The individual behaviors of the hosts are reflected in the status and the dynamics of the queues in the message bus, manifesting in different arrival and departure rates and backlogs of the individual queues. This information is available to various elements of the system and may be used for solving various problems as discussed herein.

Various embodiments of the event handling capability may be configured to solve a first problem of determining scaling (e.g., scale up and scale down) of the number of hosts participating in the event handling for the service or application in the distributed event handling system. In general, only hosts that have the software for the service or application can support the service or application and execute functions for the service or application. Installing software for the service or application generally consumes small, but non-negligible, amount of resources in the event handling system. As such, installing software for all services or applications on every host is not feasible, however, limiting the installation of software to too few hosts limits the amount of resources that are available for the service or application (which can lead to degradation of performance of the service or application (e.g., jitter in latency) and, thus, a decrease in quality-of-experience (QoE) of the end users). In view of such competing considerations, selecting the appropriate number of hosts on which software for the service or application is to be installed may be challenging. In at least some embodiments, this first problem may be solved by (1) considering the backlog and the changes in backlog of the queues in the message bus of a particular event type and (2) initiating scaling for the event type based on the queue status

information where the scaling may include scale out (e.g., if the queues for a particular event are backlogged and draining relatively slowly then this is taken as an indicator to request installation of the corresponding functions on an additional host to increase the processing capacity (which will increase the performance of the service or application for which the event type is supposed) or scale in (e.g., if the queues for a particular event type are lightly loaded and drain relatively quickly then this is taken an indicator to request that the function be uninstalled from one (or several) hosts). It will be appreciated that various other embodiments presented herein may solve or contribute to solving the first problem.

10 Various embodiments of the event handling capability may be configured to solve a second problem of determining distribution of events of the service or application to hosts that support the service or application in the distributed event handling system. In general, the distribution of events of the service or application to hosts that support the service or application should be determined such that none of
15 the hosts are overloaded and such that the hosts have the resources needed to process events efficiently. However, since hosts may support multiple services or applications, the resource availability at the hosts for event processing may depend on the workloads of the various services or applications supported by the hosts (which may vary over time). In view of such competing considerations, determining
20 distribution of events of the service or application to hosts that support the service or application may be challenging. In at least some embodiments, this second problem may be solved by (1) considering the backlog and the changes in backlog of the queues in the message bus of a particular event type by calculating a score on these two metrics for each queue where the score increases as the conditions of a queue
25 status improve (e.g., decreasing queue length and/or increasing drain rate) and (2) dispatching incoming events to the queue that has the highest score (which will minimize the latencies for the corresponding service or application for which events of the event type are being handled). It will be appreciated that various other embodiments presented herein may solve or contribute to solving the second problem.

30 Various embodiments of the event handling capability may provide various

other advantages or potential advantages.

FIG. 6 depicts a high-level block diagram of a computer suitable for use in performing various functions described herein.

The computer 600 includes a processor 602 (e.g., a central processing unit
5 (CPU), a processor having a set of one or more processor cores, or the like) and a memory 604 (e.g., a random access memory (RAM), a read only memory (ROM), or the like). The processor 602 and the memory 604 are communicatively connected.

The computer 600 also may include a cooperating element 605. The cooperating element 605 may be a hardware device. The cooperating element 605
10 may be a process that can be loaded into the memory 604 and executed by the processor 602 to implement functions as discussed herein (in which case, for example, the cooperating element 605 (including associated data structures) can be stored on a non-transitory computer-readable storage medium, such as a storage device or other storage element (e.g., a magnetic drive, an optical drive, or the like)).

15 The computer 600 also may include one or more input/output devices 606. The input/output devices 606 may include one or more of a user input device (e.g., a keyboard, a keypad, a mouse, a microphone, a camera, or the like), a user output device (e.g., a display, a speaker, or the like), one or more network communication devices or elements (e.g., an input port, an output port, a receiver, a transmitter, a
20 transceiver, or the like), one or more storage devices or elements (e.g., a tape drive, a floppy drive, a hard disk drive, a compact disk drive, or the like), or the like, as well as various combinations thereof.

It will be appreciated that computer 600 of FIG. 6 may represent a general architecture and functionality suitable for implementing functional elements described
25 herein, portions of functional elements described herein, or the like, as well as various combinations thereof. For example, computer 600 may provide a general architecture and functionality that is suitable for implementing one or more of an element of CN 110, a host 120 or a portion thereof, the MS 130 or a portion thereof, the MB 131 or a portion thereof, the FM 132 or a portion thereof, the MB 210 or a portion thereof, a
30 host 220 or a portion thereof, the FM 230 or a portion thereof, GMB 511-G or a

portion thereof, an ECN 520 or a portion thereof, an LED 521 or a portion thereof, a host 525 or a portion thereof, the MS 530 or a portion thereof, or the like, as well as various combinations thereof.

It will be appreciated that the functions depicted and described herein may be implemented in software (e.g., via implementation of software on one or more processors, for executing on a general purpose computer (e.g., via execution by one or more processors) so as to provide a special purpose computer, and the like) and/or may be implemented in hardware (e.g., using a general purpose computer, one or more application specific integrated circuits (ASIC), and/or any other hardware equivalents).

It will be appreciated that at least some of the functions discussed herein as software methods may be implemented within hardware, for example, as circuitry that cooperates with the processor to perform various functions. Portions of the functions/elements described herein may be implemented as a computer program product wherein computer instructions, when processed by a computer, adapt the operation of the computer such that the methods and/or techniques described herein are invoked or otherwise provided. Instructions for invoking the various methods may be stored in fixed or removable media (e.g., non-transitory computer-readable media), transmitted via a data stream in a broadcast or other signal bearing medium, and/or stored within a memory within a computing device operating according to the instructions.

It will be appreciated that the term "or" as used herein refers to a non-exclusive "or" unless otherwise indicated (e.g., use of "or else" or "or in the alternative").

Aspects of various embodiments are specified in the claims. Those and other aspects of various embodiments are specified in the following numbered clauses:

1. An apparatus configured to support handling of events by a message bus of an event handling system, the apparatus comprising:

a processor and a memory communicatively connected to the processor, the processor configured to:

maintain, by the message bus for an event type, a respective set of event queues for the event type, wherein the event queues for the event type

are associated with respective ones of a set of hosts subscribed to the message bus for the event type;

receive, by the message bus for the event type, an event of the event type that is published to the message bus;

5 store, by the message bus, the event of the event type in one of the event queues for the event type based on respective sets of queue status information associated with the respective event queues; and

 send, by the message bus, the event of the event type toward one of the hosts subscribed to the message bus for the event type based on a request from
10 the one of the hosts subscribed to the message bus for the event type.

2. The apparatus of clause 1, wherein the processor is configured to:
 instantiate the respective event queues for the event type based on respective requests by the respective hosts to subscribe to the message bus for the event type.

3. The apparatus of clause 1, wherein, for at least one of the event queues,
15 the respective set of queue status information comprises at least one of a queue length of the respective event queue, an event arrival rate of the respective event queue, or an event departure rate of the respective event queue.

4. The apparatus of clause 1, wherein the processor is configured to:
 compute, for each of the event queues based on the respective sets of queue
20 status information associated with the respective event queues, a respective queue score for the respective event queue.

5. The apparatus of clause 4, wherein the respective queue scores for the respective event queues are configured to enable selection of a longer one of the event queues over a shorter one of the event queues, for storing events of the event type,
25 based on an indication that a length of the longer one of the event queues and a length of the shorter one of the event queues are changing at different rates.

6. The apparatus of clause 4, wherein the respective queue score of the respective event queue is based on an exponentially weighted moving average of a queue length of the respective event queue and an exponentially weighted moving
30 average of a departure rate of the respective event queue.

7. The apparatus of clause 1, wherein the processor is configured to:
terminate the respective event queues for the event type based on respective
requests by the respective hosts to unsubscribe from the message bus for the event
type.

5 8. The apparatus of clause 1, wherein the processor is configured to:
send, toward a function manager configured to support scaling of hosts within
the event handling system, the respective queue status information of the respective
event queues.

9. The apparatus of clause 1, wherein the processor is configured to:
10 send, toward a function manager configured to support scaling of hosts within
the event handling system, per-host queue status information of the host determined
based on the respective queue status information of the respective event queue
associated with the host and based on queue status information of a second event
queue associated with the host for a second event type.

15 10. A method for handling of events by a message bus of an event
handling system, the method comprising:
maintaining, by the message bus for an event type, a respective set of event
queues for the event type, wherein the event queues for the event type are associated
with respective ones of a set of hosts subscribed to the message bus for the event type;
20 receiving, by the message bus for the event type, an event of the event type
that is published to the message bus;
storing, by the message bus, the event of the event type in one of the event
queues for the event type based on respective sets of queue status information
associated with the respective event queues; and
25 sending, by the message bus, the event of the event type toward one of the
hosts subscribed to the message bus for the event type based on a request from the one
of the hosts subscribed to the message bus for the event type.

11. An apparatus, comprising:
a processor and a memory communicatively connected to the processor, the
30 processor and the memory configured to support a message bus of an event handling

system, the message bus comprising:

5 a set of event queues for an event type supported by the event handling system, wherein the event queues for the event type are configured to store events for respective ones of a set of hosts subscribed to the message bus for the event type;

a queue monitoring element configured to monitor the event queues and to determine, for each of the event queues, respective queue status information associated with the respective event queue; and

10 an event dispatch element configured to receive an event of the event type that is published to the message bus and to dispatch the event of the event type to one of the event queues for the event type based on the respective queue status information associated with the respective event queues;

15 wherein the message bus is configured to send the event of the event type toward one of the hosts subscribed to the message bus for the event type based on a request from the one of the hosts subscribed to the message bus for the event type.

12. An apparatus configured to support handling of events by a host of an event handling system, the apparatus comprising:

20 a processor and a memory communicatively connected to the processor, the processor configured to:

subscribe, by the host, to a first event queue of a message bus, the first event queue of the message bus configured to store events of a first event type to be processed by a first event processing function of the host;

25 subscribe, by the host, to a second event queue of the message bus, the second event queue configured to store events of a second event type to be processed by a second event processing function of the host;

maintain, by the host, a local buffer configured to store events received from the first event queue of the message bus and events received from the second event queue of the message bus; and

30 retrieve, by the host based on a local queuing policy of the host and a

local event retrieval condition detected at the host, a retrieved event comprising a first event from the first event queue of the message bus or a second event from the second event queue of the message bus.

13. The apparatus of clause 12, wherein the processor is configured to
5 subscribe to the first event queue of the message bus based on a determination that the host supports the first event processing function.

14. The apparatus of clause 12, wherein the processor is configured to subscribe to the second event queue of the message bus based on a determination that the host supports the second event processing function.

10 15. The apparatus of clause 12, wherein the processor is configured receive scaling requests from a function manager, wherein the processor is configured to subscribe to the first event queue of the message bus based on a first request from the function manager and to subscribe to the second event queue of the message bus based on a second request from the function manager.

15 16. The apparatus of clause 12, wherein the local buffer of the host comprises a first host event queue configured to store events received from the first event queue of the message bus and a second host event queue configured to store events received from the second event queue of the message bus.

17. The apparatus of clause 12, wherein the local buffer of the host
20 comprises a first-in-first-out buffer configured to store events received from the first event queue of the message bus and events received from the second event queue of the message bus.

18. The apparatus of clause 12, wherein the processor is configured to:
select, by the host from the local buffer configured to store events received
25 from the first event queue of the message bus and events received from the second event queue of the message bus based on a local event selection condition detected at the host, a selected event; and
process the selected event.

19. The apparatus of clause 18, wherein the selected event is an event of
30 the first event type, wherein the selected event is processed by the host using the first

event processing function of the host.

20. The apparatus of clause 18, wherein the selected event is an event of the second event type, wherein the selected event is processed by the host using the second event processing function of the host.

5 21. The apparatus of clause 12, wherein the processor is configured to:
detect a request for the host to stop handling events of the first event type;
unsubscribe from the first event queue of the message bus based on the request
for the host to stop handling events of the first event type;
process, by the host using the first event processing function of the host, any
10 events of the local buffer of the host of the first event type; and
initiate removal of the first event processing function from the host.

22. The apparatus of clause 12, wherein the processor is configured to:
detect a request for the host to stop handling events of the second event type;
unsubscribe from the second event queue of the message bus based on the
15 request for the host to stop handling events of the second event type;
process, by the host using the second event processing function of the host,
any events of the local buffer of the host of the second event type; and
initiate removal of the second event processing function from the host.

23. A method to support handling of events by a host of an event handling
20 system, the method comprising:

subscribing, by the host, to a first event queue of a message bus, the first event
queue of the message bus configured to store events of a first event type to be
processed by a first event processing function of the host;

subscribing, by the host, to a second event queue of the message bus, the
25 second event queue configured to store events of a second event type to be processed
by a second event processing function of the host;

maintaining, by the host, a local buffer configured to store events received
from the first event queue of the message bus and events received from the second
event queue of the message bus; and

30 retrieving, by the host based on a local queuing policy of the host and a local

event retrieval condition detected at the host, a retrieved event comprising a first event from the first event queue of the message bus or a second event from the second event queue of the message bus.

5 It will be appreciated that, although various embodiments which incorporate the teachings presented herein have been shown and described in detail herein, those skilled in the art can readily devise many other varied embodiments that still incorporate these teachings.

What is claimed is:

1. An apparatus configured to support handling of events by a message bus of an event handling system, the apparatus comprising:
 - 5 a processor and a memory communicatively connected to the processor, the processor configured to:
 - maintain, by the message bus for an event type, a respective set of event queues for the event type, wherein the event queues for the event type are associated with respective ones of a set of hosts subscribed to the message
 - 10 bus for the event type;
 - receive, by the message bus for the event type, an event of the event type that is published to the message bus;
 - store, by the message bus, the event of the event type in one of the event queues for the event type based on respective sets of queue status
 - 15 information associated with the respective event queues; and
 - send, by the message bus, the event of the event type toward one of the hosts subscribed to the message bus for the event type based on a request from the one of the hosts subscribed to the message bus for the event type.
- 20 2. The apparatus of claim 1, wherein the processor is configured to:
 - instantiate the respective event queues for the event type based on respective requests by the respective hosts to subscribe to the message bus for the event type.
3. The apparatus of claim 1, wherein, for at least one of the event queues, the
- 25 respective set of queue status information comprises at least one of a queue length of the respective event queue, an event arrival rate of the respective event queue, or an event departure rate of the respective event queue.
4. The apparatus of claim 1, wherein the processor is configured to:
 - 30 compute, for each of the event queues based on the respective sets of queue

status information associated with the respective event queues, a respective queue score for the respective event queue.

5. The apparatus of claim 4, wherein the respective queue scores for the
5 respective event queues are configured to enable selection of a longer one of the event queues over a shorter one of the event queues, for storing events of the event type, based on an indication that a length of the longer one of the event queues and a length of the shorter one of the event queues are changing at different rates.
- 10 6. The apparatus of claim 4, wherein the respective queue score of the respective event queue is based on an exponentially weighted moving average of a queue length of the respective event queue and an exponentially weighted moving average of a departure rate of the respective event queue.
- 15 7. The apparatus of claim 1, wherein the processor is configured to:
terminate the respective event queues for the event type based on respective requests by the respective hosts to unsubscribe from the message bus for the event type.
- 20 8. The apparatus of claim 1, wherein the processor is configured to:
send, toward a function manager configured to support scaling of hosts within the event handling system, the respective queue status information of the respective event queues.
- 25 9. An apparatus, comprising:
a processor and a memory communicatively connected to the processor, the processor and the memory configured to support a message bus of an event handling system, the message bus comprising:
a set of event queues for an event type supported by the event handling
30 system, wherein the event queues for the event type are configured to store

events for respective ones of a set of hosts subscribed to the message bus for the event type;

a queue monitoring element configured to monitor the event queues and to determine, for each of the event queues, respective queue status information associated with the respective event queue; and

an event dispatch element configured to receive an event of the event type that is published to the message bus and to dispatch the event of the event type to one of the event queues for the event type based on the respective queue status information associated with the respective event queues;

wherein the message bus is configured to send the event of the event type toward one of the hosts subscribed to the message bus for the event type based on a request from the one of the hosts subscribed to the message bus for the event type.

10. An apparatus configured to support handling of events by a host of an event handling system, the apparatus comprising:

a processor and a memory communicatively connected to the processor, the processor configured to:

subscribe, by the host, to a first event queue of a message bus, the first event queue of the message bus configured to store events of a first event type to be processed by a first event processing function of the host;

subscribe, by the host, to a second event queue of the message bus, the second event queue configured to store events of a second event type to be processed by a second event processing function of the host;

maintain, by the host, a local buffer configured to store events received from the first event queue of the message bus and events received from the second event queue of the message bus; and

retrieve, by the host based on a local queuing policy of the host and a local event retrieval condition detected at the host, a retrieved event comprising a first event from the first event queue of the message bus or a

second event from the second event queue of the message bus.

FIG. 1






FIG. 2



3/6

FIG. 3



4/6

FIG. 4




FIG. 5

500





FIG. 6



INTERNATIONAL SEARCH REPORT

International application No
PCT/US2018/041268

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F9/54
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2011/173359 A1 (CHAKRAVARTY DIPTO [US] ET AL) 14 July 2011 (2011-07-14)	1-4,8-10
Y	figures 2-8	5-7
	paragraphs [0023] - [0098]	

Y	US 2003/007453 A1 (OGIER RICHARD G [US] ET AL) 9 January 2003 (2003-01-09)	5,6
A	figures 2-4	1-4,7-10
	paragraphs [0003] - [0016], [0041] - [0046]	

Y	US 2017/118310 A1 (GILLIARD CHRISTOPHER [US] ET AL) 27 April 2017 (2017-04-27)	7
A	figures 1-6	1-6,8-10
	paragraphs [0001], [0002], [0043] - [0045]	

	-/--	



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

3 October 2018

Date of mailing of the international search report

11/10/2018

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Freitas, Arthur

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2018/041268

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	EP 1 107 556 A2 (AVAYA TECHNOLOGY CORP [US]) 13 June 2001 (2001-06-13) figures 1-5 paragraphs [0002] - [0011] -----	1-10

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2018/041268

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2011173359	A1	14-07-2011	US 7926099 B1 12-04-2011
			US 2011173359 A1 14-07-2011
US 2003007453	A1	09-01-2003	NONE
US 2017118310	A1	27-04-2017	NONE
EP 1107556	A2	13-06-2001	CA 2326613 A1 06-06-2001
			DE 60033119 T2 08-11-2007
			EP 1107556 A2 13-06-2001
			JP 4130301 B2 06-08-2008
			JP 2001230873 A 24-08-2001
			KR 20010062178 A 07-07-2001
			US 6553114 B1 22-04-2003