



- (51) **International Patent Classification:**
H04L 29/08 (2006.01) **G06F 11/30** (2006.01)
G06F 11/34 (2006.01)
- (21) **International Application Number:**
 PCT/US2016/039371
- (22) **International Filing Date:**
 24 June 2016 (24.06.2016)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
- | | | |
|------------|---------------------------|----|
| 14/752,754 | 26 June 2015 (26.06.2015) | US |
| 14/752,752 | 26 June 2015 (26.06.2015) | US |
| 14/752,759 | 26 June 2015 (26.06.2015) | US |
| 14/752,760 | 26 June 2015 (26.06.2015) | US |
| 14/752,756 | 26 June 2015 (26.06.2015) | US |
- (71) **Applicant:** **AMAZON TECHNOLOGIES, INC.**
 [US/US]; P.O. Box 81226, Seattle, Washington 98108-1226 (US).
- (72) **Inventors:** **TAYLOR, Gary**; 410 Terry Avenue North, Seattle, Washington 98109-5210 (US). **DENNEHY, Brian**; 410 Terry Avenue North, Seattle, Washington 98109-5210 (US). **RODRIGUEZ, Sebastien Alexandre Roland**; 410 Terry Avenue North, Seattle, Washington 98109-5210

(US). **MANGAS, Julio Delgado**; 410 Terry Avenue North, Seattle, Washington 98109-5210 (US).

- (74) **Agents:** **ADAMS, Scott S.** et al.; c/o IP Docketing Dept., Davis Wright Tremaine LLP, 1201 Third Avenue, Suite 2200, Seattle, Washington 98101-3045 (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, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, 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.

(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).

[Continued on next page]

- (54) **Title:** DATASTORE FOR AGGREGATED MEASUREMENTS FOR METRICS




FIG. 1

(57) **Abstract:** A computing resource monitoring service receives a request to store a measurement for a metric associated with a computing resource. The request includes the measurement itself and metadata for the measurement, which specifies attributes of the measurement. Based at least in part on the metadata, the computing resource monitoring service generates a fully qualified metric identifier and, using the identifier, selects a logical partition for placement of the measurement. From the logical partition, the computing resource monitoring service transmits the measurement to an aggregator sub-system comprising one or more in-memory datastores. The computing resource monitoring service stores the measurement in an in-memory datastore within the aggregator sub-system.

**Declarations under Rule 4.17:**

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

- *with international search report (Art. 21(3))*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

DATASTORE FOR AGGREGATED MEASUREMENTS FOR METRICS

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] This application incorporates by reference for all purposes the full disclosure of co-
5 pending U. S. Patent Application Nos. 14/752,754, filed June 26, 2016, entitled
“DATASTORE FOR AGGREGATED MEASUREMENTS FOR METRICS” (Attorney
Docket No. 0097749-515US0), U. S. Patent Application No. 14/752,752, filed June 26, 2016,
entitled “AGGREGATION OF METRICS DATA WITH FINE GRANULARITY” (Attorney
Docket No. 0097749-492US0), U. S. Patent Application No. 14/752,759, filed June 26, 2016,
10 entitled “ARCHITECTURE FOR METRICS AGGREGATION WITHOUT SERVICE
PARTITIONING” (Attorney Docket No. 0097749-547US0), U. S. Patent Application No.
14/752,760, filed June 26, 2016, entitled “AUTOMATIC SCALING OF COMPUTING
RESOURCES USING AGGREGATED METRICS” (Attorney Docket No. 0097749-
548US0), and U. S. Patent Application No. 14/752,756, filed June 26, 2016, entitled
15 “RETRIEVAL OF AUTHORITATIVE MEASUREMENT DATA FROM IN-MEMORY
DATASTORES” (Attorney Docket No. 0097749-549US0).

BACKGROUND

[0002] Customers and administrators of a computing resource service provider, as well as
20 other users of computing resources of the computing resource service provider often utilize
monitoring services to measure, diagnose, and improve how they operate their computing
resources. For instance, through these monitoring services, customers, administrators, and
other users can obtain data for their computing resources and use this data to determine
whether their computing resources are functioning properly. If their computing resources are
25 not functioning properly, the data can be used to identify and enable customers,
administrators, and other users to troubleshoot any issues that may be present. However,
during high impact operational events, time is often of the essence. Monitoring computing
resources in a distributed system is very complex, and the complexity grows with the size of
the distributed system. As a result, conventional techniques for monitoring computing
30 resources often result in significant latencies between when data indicative of an issue is
obtained and when that data is available for consumption. Some customers, administrators,

and other users may need to acquire data rapidly to track their computing resources and address any issues as soon as possible.

BRIEF DESCRIPTION OF THE DRAWINGS

- 5 [0003] Various embodiments in accordance with the present disclosure will be described with reference to the drawings, in which:
- [0004] FIG. 1 shows an illustrative example of an environment in which various embodiments can be implemented;
- 10 [0005] FIG. 2 shows an illustrative example of an environment in which a front-end server of a computing resource monitoring service processes one or more application programming interface calls to store and retrieve measurements from the service in accordance with at least one embodiment;
- [0006] FIG. 3 shows an illustrative example of an environment in which a partitioner sub-system of a computing resource monitoring service partitions measurements into various
15 logical partitions and delivers measurements in each logical partition to an aggregator sub-system in accordance with at least one embodiment;
- [0007] FIG. 4 shows an illustrative example of an environment in which an aggregator sub-system performs aggregation of measurements and provides read access for measurements stored within one or more datastores in accordance with at least one embodiment;
- 20 [0008] FIG. 5 shows an illustrative example of an environment in which a customer computer system partitions measurements into various logical partitions and delivers the measurements from each logical partition to one or more aggregator sub-systems of a computing resource monitoring service in accordance with at least one embodiment;
- [0009] FIG. 6 shows an illustrative example of an environment in which metadata is
25 provided with a measurement for a metric to enable future storage of measurements for the metric without requiring additional metadata transmissions in accordance with at least one embodiment;
- [0010] FIG. 7 shows an illustrative example of an environment in which ingestion of measurements for metrics are used to determine whether automatic scaling of one or more
30 resources is to be performed in accordance with at least one embodiment;

[0011] FIG. 8 shows an illustrative example of a process for partitioning measurements for a metric for delivery to one or more aggregator sub-systems in accordance with at least one embodiment;

[0012] FIG. 9 shows an illustrative example of a process for aggregating measurements from one or more partitioner sub-systems with measurements from one or more datastores of a computing resource monitoring service in accordance with at least one embodiment;

[0013] FIG. 10 shows an illustrative example of a process for retrieving one or more measurements from one or more aggregator datastores in response to a GET application programming interface call in accordance with at least one embodiment;

[0014] FIG. 11 shows an illustrative example of a process for partitioning measurements to be transmitted to a computing resource monitoring service for publishing of the measurements in accordance with at least one embodiment;

[0015] FIG. 12 shows an illustrative example of a process for storing measurements in one or more aggregator datastores based at least in part on a metadata hash in accordance with at least one embodiment;

[0016] FIG. 13 shows an illustrative example of a process for retrieving measurements from one or more aggregator datastores based at least in part on metadata included within a request to obtain the measurements in accordance with at least one embodiment; and

[0017] FIG. 14 shows an illustrative example of an environment in which various embodiments can be implemented.

DETAILED DESCRIPTION

[0018] In the following description, various embodiments will be described. For purposes of explanation, specific configurations and details are set forth in order to provide a thorough understanding of the embodiments. However, it will also be apparent to one skilled in the art that the embodiments may be practiced without the specific details. Furthermore, well-known features may be omitted or simplified in order not to obscure the embodiment being described.

[0019] Techniques described and suggested herein relate to a computing resource monitoring service configured to aggregate measurements for metrics and store the aggregated measurements in memory during a retention period specified by a service provider. The computing resource monitoring service may include a front-end server, which

may be configured to obtain measurements from a variety of different sources, including customers of the computing resource service provider, various other services of the computing resource service provider, and computing resources made available to the customers through the various services. This front-end server may transform these

5 measurements into a binary serialization format that can be utilized by the various components of the computing resource monitoring service. Once the transformation of the measurements is completed, the front-end service may publish these measurements to various partitioner sub-systems to divide the measurements based on a fully qualified metric identifier (FQMI) for each measurement or metric and the timestamp for the measurement for
10 distribution to various aggregator datastores.

[0020] A partitioner sub-system may store the measurements in various queues based on the FQMI. Each queue may be associated with a respective aggregator sub-system for storage of the measurements into the various aggregator datastores utilized to support application programming interface (API) calls to the computing resource monitoring service to retrieve
15 measurements. The partitioner sub-system may further partition the measurements within each queue based on the timestamp for each measurement. This enables the partitioner sub-system to prioritize delivery of the measurements to the aggregator sub-systems based on the timestamps, delivering the most recent measurements first. Further, the partitioner sub-system may purge measurements from the various queues if the timestamp for this data is earlier than
20 the oldest retention period. For measurements with timestamps in the future (e.g., beyond the latest retention period), the partitioner sub-system may queue these measurements in a separate unpartitioned queue, where a measurement will remain until the timestamp for the measurement is covered by the present retention period at a later time. Once the retention period has been updated, measurements with timestamps within this retention period may be
25 transferred from the separate unpartitioned queue and processed.

[0021] The aggregator sub-system may obtain the measurements from various queues of the partitioner sub-systems and aggregate these measurements for a retention period. When an aggregator sub-system obtains the measurements from various partitioner sub-systems, the aggregator sub-system may obtain existing measurements and determine whether any of the
30 existing measurements correspond to a time period earlier than the oldest retention period. If so, the aggregator sub-system may purge these existing measurements. The aggregator sub-system may store the measurements obtained from the partitioner sub-systems within various file-based queues for asynchronous processing by the aggregator datastores of the aggregator sub-system. This aggregation of the newly obtained measurements with existing

measurements may be performed through deserialization of the existing measurements and the newly obtained measurements and aggregating the newly obtained measurements with the remaining measurements into a serialized format. This serialized data may then be stored within an aggregator datastore.

5 [0022] When a customer of the computing resource service provider submits an API call to the front-end server to retrieve measurements from the computing resource monitoring service (e.g., GET, etc.), the front-end server may transmit a request to the aggregator sub-systems to obtain the measurements. This may cause the various aggregator sub-systems to each access a metric mapping registry to determine where the requested measurements are
10 stored. The aggregator sub-systems may then access their respective aggregator datastores to obtain the requested measurements and provide these measurements to a metric consolidation engine. This metric consolidation engine may compile the measurements from the aggregator sub-systems and provide these measurements to the front-end server. The front-end server may fulfill the customer's request by providing the compiled measurements from the metric
15 consolidation engine.

[0023] In this manner, the computing resource monitoring service may ingest measurements for various metrics from a variety of sources, compute real time analytics and provide measurements to customers, administrators, and other entities to enable rapid
20 response to any inherent computing resource issues and to allow metrics evaluation in a short amount of time. In addition, the techniques described and suggested herein facilitate additional technical advantages. For example, because measurements ingested by the computing resource monitoring service are divided and placed into logical partitions for distribution to the various partitioner sub-systems, any failure of a partitioner sub-system may result in minimal impact to the overall aggregation of measurements by the computing
25 resource monitoring service. This may ensure that customers of the computing resource service provider may store measurements for the retention period quickly with minimal risk of significant data loss.

[0024] FIG. 1 shows an illustrative example of an environment 100 in which various embodiments can be implemented. In the environment 100, one or more customers 102 (e.g.,
30 individuals, organizations, automated processes, computing resource agents, etc.) of a computing resource service provider may submit a request to a computing resource monitoring service 104 to store a measurement of a particular metric for a computing resource. The request to store the measurement may be provided in the form of an API call to

the computing resource monitoring service 104, such as a PUT call. In an embodiment, the initial request to store the measurement includes metadata that can be used to describe the measurement that is being provided to the computing resource monitoring service 104. For instance, the metadata may specify the name of the originating computing resource that has generated the data, an auto-scaling group to which the originating computing resource belongs to, a virtual machine image identifier for the originating computing resource, and the like. Upon receiving the measurement and the metadata through the API call to store the measurement, the computing resource monitoring service 104 may generate a hash of the metadata, which may be used to index the measurement upon storage. A measurement may be a numerical value from a scale of possible values, a Boolean value, or an alphanumeric or other indicator of a state of a set of possible states. The measurement may further serve as an indication of at least one aspect of the operation of an associated computing resource, such as virtual machine instances, customer computer systems, object-based datastores and/or other computing resources. For instance, measurements may include processor usage over time, memory usage over time, a count for errors detected within a computing resource over a period of time, and the like.

[0025] In an embodiment, the computing resource monitoring service 104 includes one or more partitioner sub-systems 108 and one or more aggregator sub-systems 110 within each datacenter 106 of the computing resource monitoring service 104 for storage of the measurement. Each datacenter 106 may correspond to a physical location where computer systems of the computing resource monitoring service 104, as well as other computer systems for services of the computing resource service provider, may be located. When the computing resource monitoring service 104 obtains the measurement from a customer 102, the computing resource monitoring service 104 may utilize the provided metadata to identify the computing resource that has generated the data and determine the datacenter 106 where this computing resource is located. Based at least in part on this determination, the computing resource monitoring service 104 may select a corresponding datacenter 106 for storage of the data. In some embodiments, the computing resource monitoring service 104 replicates the measurement for storage in a subset of a plurality of datacenters 106 to provide redundancy for the measurement. This may enable the computing resource monitoring service 104 to generate one or more conflict resolution rules in order to select a particular measurement if the subset of the plurality of datacenters 106 is unable to provide a consistent measurement answer in response to a request. For instance, the computing resource monitoring service 104 may require a quorum where the computing resource monitoring service 104 may select the

measurement that is provided by the most datacenters 106 so long as the number of datacenters 106 providing the same measurement is greater than or equal to a predetermined number of datacenters 106.

[0026] The metadata for the measurement provided to the computing resource monitoring service 104 may uniquely identify the particular metric for the measurement. For instance, the metadata may uniquely identify the measurement by a combination of a customer account number, the namespace (e.g., associated service for the computing resource), the dimensions for the measurement (e.g., key/value pairs), the name of the measurement itself, and the like. In an embodiment, when the computing resource monitoring service 104 receives the measurement and metadata from the customer 102, the computing resource monitoring service 104 utilizes a hash function to generate a hash of the metadata, which can be used as an FQMI for the measurement that is to be stored. Further, the computing resource monitoring service 104 may utilize a binary serialization format to compress the measurement to be stored. In this binary serialization format, the serialized measurement may include the FQMI generated by hashing the metadata, the timestamp for the data, the unit of measurement for the data, and the measurement itself. In an alternative embodiment, the customer 102, through use of a customer computer system or other computing resource, calculates the FQMI based at least in part on the metadata for the measurement. This may enable the customer 102 to provide the measurement with the FQMI for the measurement without requiring the computing resource monitoring service 104 to generate the FQMI itself.

[0027] In an embodiment, once the computing resource monitoring service 104 has hashed the metadata to generate the FQMI for the measurement to be stored, the customer 102 is no longer required to repeat the metadata for future submissions of the measurements for the particular metric, so long as the metadata has not expired or the customer 102 has not initiated a new session with the computing resource monitoring service 104. For instance, when a customer 102 supplies additional measurements to the computing resource monitoring service 104, the customer 102 may provide the FQMI for the measurements corresponding to the same metric instead of the metadata. The computing resource monitoring service 104 may serialize the measurements and the FQMI in anticipation of storage of the various measurements within the datastores of the computing resource monitoring service 104. The computing resource monitoring service 104 may provide the customer 102 with the FQMI for the measurements in response to the customer's request to store a first measurement that includes metadata for the measurement. Once the computing resource monitoring service 104 generates the FQMI based at least in part on the metadata

provided by the customer 102, the computing resource monitoring service 104 may transmit the FQMI to the customer 102 for future use, include PUT and GET requests.

[0028] The API call that may be made to the computing resource monitoring service 104 for storage of the measurement (e.g., an observation for a metric) may further include one or more parameters that may be used to specify that the data provided for a particular time series for the metric is authoritative for the time series. For instance, as measurements for a metric are aggregated for a particular retention period, the customer 102 may specify that for the time series, the provided one or more measurements are authoritative for the time series such that the computing resource monitoring service 104 may indicate, to an entity requesting the measurements, that the provided measurement for the time series is in fact authoritative and steady. In an embodiment, when the computing resource monitoring service 104 receives an API call including the one or more parameters usable to specify that the measurement provided is authoritative for the time series, the computing resource monitoring service 104 will make the various measurements for the time series available once the received measurement has been aggregated with other measurements for the time series and stored. This measurement, as well as other measurements for the time series, may be available to the customer 102 or other entities such that the computing resource monitoring service 104 may indicate that these measurements are authoritative for the time series.

[0029] Once the computing resource monitoring service 104 has transformed the measurement into a serialized format, the computing resource monitoring service 104 may transmit the measurement to one or more partitioner sub-systems 108 within each datacenter 106. In an embodiment, each partitioner sub-system 108 determines, based at least in part on the FQMI and the timestamp for the measurement, which logical partition of a plurality of logical partitions will be used for storing the measurement. For various measurements received, the partitioner sub-system 108 may split these various measurements into various logical partitions further based at least in part on a number of active aggregator sub-systems 110 within the datacenter 108 for one or more retention periods. In order to determine the number of active aggregator sub-systems 110, the partitioner sub-system 108 may access a metric mapping registry, which may provide a mapping of logical partitions to the aggregator sub-systems 110 active for the one or more retention periods. The partitioner sub-system 108 may store the measurement within a file-based queue associated with the selected logical partition, bound to a single aggregator sub-system 110 within the datacenter 106.

[0030] The file-based queue may include various measurements, which may be batched and delivered asynchronously to the corresponding aggregator sub-system 110 after a period of time has passed (e.g., every few milliseconds, etc.). The measurements within these file-based queues may be delivered to only one aggregator sub-system 110 per datacenter 106

5 during each retention period, as determined through the metric mapping registry. The measurements within the file-based queues may be further sub-partitioned based at least in part on the observation time for each measurement. This may enable the partitioner sub-system to deliver the measurements to the aggregator sub-systems 110 in a manner that prioritizes delivery of the most recent measurement first. Measurements with timestamps in
10 the future (e.g., beyond the latest time series of the latest retention period) may be queued in a separate unpartitioned queue that may be processed only when the time series becomes current. Additionally, the partitioner sub-system 108 may purge measurements with timestamps that are earlier than the oldest retention period.

[0031] An aggregator sub-system 110 is an in-memory key-value storage system that may
15 be optimized for aggregating measurements and serving time series measurement data for various metrics. In an embodiment, the in-memory key-value storage system relies on main memory of one or more computer systems for storage of aggregated measurements. For instance, the in-memory key-value storage system may be directly or indirectly connected to a central processing unit (CPU) of a computer system via a memory bus. This may enable the
20 CPU to obtain any data stored within the in-memory key-value storage system in response to GET requests or to store any data within the in-memory key-value storage system in response to GET requests. The main memory may also be directly accessible to the CPU. The in-memory key-value storage system may be formed using volatile storage systems (e.g., Random Access Memory (RAM), etc.) or non-volatile storage systems. In some
25 embodiments, the measurements may be stored in solid state storage, such as solid state drives (SSDs).

[0032] The aggregator sub-system 110 may obtain a plurality of measurements from various partitioner sub-systems 108 within a datacenter 106 and may place the measurements into file-based queues for asynchronous processing. If it is the first time that a particular
30 metric has been observed for an aggregation period, the aggregation sub-system 110 may store the measurement for this particular metric in its serialized format within an aggregator datastore. Alternatively, if the particular metric has been previously observed, the aggregator sub-system 110 may obtain the stored measurements for the metric, de-serialize the stored measurements and the newly obtained measurements, and aggregate the newly obtained

measurements with the existing measurements in the serialized format. These aggregated measurements may then be stored within the aggregator datastore for fulfillment of query requests. In an embodiment, the aggregator sub-system 110 will purge any expired measurements from the aggregator datastores once the oldest retention period has been replaced with a later retention period for storage of the measurements.

[0033] At any time, a customer 102 of the computing resource service provider may transmit a request to the computing resource monitoring service 104 to retrieve any number of measurements from the various aggregator datastores. The request to retrieve the measurements may be provided in the form of an API call to the computing resource monitoring service 104, such as a GET call. In order for the computing resource monitoring service 104 to obtain the measurements necessary to fulfill the request, the customer 102 may be required to provide the metadata for the measurements that are to be obtained. This may enable the computing resource monitoring service 104 to hash the metadata to obtain the one or more FQMI for the measurements. The computing resource monitoring service 104 may utilize the one or more FQMI to identify, from the metric mapping registry, the one or more locations of the measurements for the one or more current retention periods.

[0034] In an embodiment, the API call from the customer 102 to the computing resource monitoring service 104 to retrieve measurement data (e.g., one or more measurements for a metric) from the various aggregator datastores can include a parameter, which may be used to indicate that the customer 102 is requesting authoritative results. For instance, if the API call includes the parameter for requesting authoritative results for a particular time range and metric, the computing resource monitoring service 104 may retrieve the measurements for the metric within the specified time range and determine whether any of these measurements are authoritative. For instance, as noted above, a customer 102, through use of the PUT API call, may specify that a particular measurement or series of measurements are authoritative for a time series. The computing resource monitoring service 104 may utilize this information to identify the authoritative measurements and provide these to the customer 102 or other requesting entity. In some embodiments, the customer 102 can submit a GET API call that includes a parameter for indication that the customer 102 is agnostic to obtaining the authoritative measurements for a particular time range. This may cause the computing resource monitoring service 104 to provide all measurements for the specified time range regardless of whether the measurements are authoritative or not.

[0035] Upon determining the locations where the requested measurements are stored, the computing resource monitoring service 104 may utilize a metric consolidation engine 112 to transmit requests to the appropriate aggregator sub-systems 110 from the various datacenters 106 to obtain the measurements necessary to fulfill the GET request. The metric

consolidation engine 112 may be a computer system module of the computing resource monitoring service 104, which may include a client library to enable the module to read measurements from the various aggregator sub-systems 110. Once the metric consolidation engine 112 has obtained the measurements from the various aggregator sub-systems 110, the measurements from each datacenter 106 may be compiled to generate a datacenter 106 measurement response. The metric consolidation engine 112 may utilize one or more conflict resolution rules to determine the appropriate response to the GET request if the generated responses are in conflict. For instance, the metric consolidation engine 112 may select the response with the highest sample count from the various datacenters 106. The metric consolidation engine 112 may provide a response to the GET request in the form of the measurements in a de-serialized format that may be used by the customer 102 for its own purposes.

[0036] In an embodiment, the metric consolidation engine 112 is configured to provide these measurements to an auto-scale group comprising one or more server computers and computing resources. The measurements may be obtained by the metric consolidation engine 112 in response to GET requests from one or more computing resource managers within the auto-scale group for metrics associated with the computing resources within the group. These one or more computing resource managers may utilize the obtained measurements for the particular metrics to determine whether scaling of existing computing resources is necessary. If so, the computing resource managers within each server computer in the auto-scale group may provision additional computing resources or remove computing resources from the auto-scale group as needed.

[0037] It should be noted that in some embodiments, the partitioner sub-systems 108 are utilized within one or more client-side computing systems of the customer 102. For instance, the client-side computing systems may obtain the measurements from the various computing resources and may hash the metadata for the measurements directly. Further, the customer client device may access the metric mapping registry provided by the computing resource monitoring service 104 to determine which aggregator sub-systems 110 are to receive the measurements stored within the logical partitions of the partitioner sub-systems 108 within the customer client device. The customer client device may transmit a PUT request to the

corresponding aggregator sub-systems 110 of the computing resource monitoring service 104 to store the measurements. The measurements may be transmitted using one or more communications protocols that may enable the customer client device to determine whether each measurement has been transmitted successfully to the corresponding aggregator sub-systems 110. If the transmission of the measurements was unsuccessful, the customer client device may refresh the file-based queues of the partitioner sub-systems 108 and transmit a new PUT request to the aggregator sub-systems 110 for each failed request. Otherwise, the customer client device may receive acknowledgment from the aggregator sub-systems 110 that delivery of the measurements has been performed successfully.

[0038] As noted above, a computing resource monitoring service may obtain measurements from customers, administrators, computing resources, and other entities and store these measurements within one or more aggregator data stores over a particular retention period as defined by the service. The measurements obtained by the computing resource monitoring service may be aggregated as it arrives in real time. This may enable the computing resource monitoring service to respond to requests to obtain measurements from the aggregator datastores in an expeditious manner, enabling the customers and other entities to utilize the measurements for their own needs. Accordingly, FIG. 2 shows an illustrative example of an environment 200 in which a front-end server 204 of a computing resource monitoring service 202 processes one or more API calls to store and retrieve measurements from the service 202 in accordance with at least one embodiment.

[0039] In the environment 200, a front-end server 204 of the computing resource monitoring service 202 may receive a PUT request from a customer computer system to store a measurement within an aggregator datastore. The customer computer system may be a computer system hosted by a computing resource service provider but managed by a customer. Alternatively, the customer computer system may include computer systems and resources hosted by a customer in its own datacenters. The PUT request may include a measurement for a particular metric at a given timestamp. Additionally, the PUT request may further include metadata for the provided measurement. This metadata may include one or more metric attributes that may uniquely identify the associated measurement and metric within the computing resource monitoring service 202. For instance, the metadata may specify, individually or in combination, the account number of a customer associated with the customer computer system submitting the PUT request or with another computing resource responsible for generation of the measurement, the name of the particular service responsible for managing the computing resource (e.g., virtual computer system service, database service,

etc.), the dimensions for the measurement (e.g., key/value pairs), and an identifier for the computing resource producing the measurement.

[0040] When the front-end server 204 receives the PUT request from the customer computer system, the front-end server 204 may obtain the metadata from the PUT request and utilize a hash function to generate an FQMI for the measurement that is to be stored within the aggregator datastore. Additionally, the front-end server 204 may transform the measurement into a serialized format for publication to a partitioner sub-system 208 within the one or more datacenters 206. The serialized format for a particular measurement obtained from the customer computer system may include the FQMI, the timestamp for the measurement, the unit of measurement for the measurement, and the measurement itself. The front-end server 204 may spool unpartitioned measurements from various PUT requests for asynchronous delivery to the one or more partitioner sub-systems 208 of the one or more datacenters 206. For instance, the front-end server 204 may include one or more file-based queues that may each be configured to transmit the measurements to at least one partitioner sub-system 208 of the partitioner sub-systems 208 maintained within each datacenter 206.

[0041] In an embodiment, if a customer computer system submits a later PUT request to store a measurement for the same metric, the customer computer system is only required to provide the FQMI for the metric instead of the complete metadata from the measurement, as long as the metadata has not expired. Alternatively, the PUT request may lack at least some of the metadata that was received within the initial PUT request to store another measurement for the same metric. As will be described in greater detail below, the one or more partitioner sub-systems 208 may utilize the provided FQMI to store the measurement within an aggregator datastore such that information responsive to a request to retrieve measurements and that specifies the metadata includes at least both measurements.

[0042] The PUT request may further provide the customer computer system with a variety of options for storage of the provided measurement. For instance, in an embodiment, the PUT API may enable the customer computer system to specify whether the provided measurement is the last measurement to be provided for a particular metric over a period of time for aggregation of measurements by the computing resource monitoring service 202. For instance, through the PUT API, if the front-end server 204 is configured to batch measurements for a particular time interval (e.g., one second, one minute, one hour, etc.) prior to delivery to the partitioner sub-systems 208, a customer computer system may specify that the provided measurement is the final measurement for the particular time interval. This

may cause the front-end server 208 to prepare the measurements for delivery to the partitioner sub-systems 208 even through the current time interval has not elapsed. This may enable faster storage and availability of measurements, as the front-end server 204 may not be required to wait for additional measurements from the customer computer system. In another embodiment, the front-end server 204 will utilize run length encoding to compress the obtained measurements for one or more sequential measurements having a same value over a time interval for the measurements. This may reduce the storage space required for the serialized measurements.

[0043] In an embodiment, the front end server 204 sends the serialized measurement generated based at least in part on the received PUT request to each datacenter 206 for redundant storage of the measurement. Each datacenter 206 may include a partitioner load balancer (not shown) that may be configured to transmit the serialized measurement to a partitioner sub-system 208 within the particular datacenter 206. The partitioner sub-system may include one or more computer systems collectively configured to partition measurements for various metrics based at least in part on the FQMI and timestamp of each measurement for delivery to the one or more aggregator sub-systems 210 within the datacenter 206. The partitioner sub-system 208, upon obtaining the measurement from the front-end server 204, may access a metric mapping registry 212 to obtain a mapping of logical partitions to aggregator sub-systems 210 within the datacenter 206 over time. The metric mapping registry may include a database for the mapping and one or more computer systems for creating partition schedules based at least in part on available aggregator sub-systems 210 within each datacenter. For instance, the metric mapping registry 212 may evaluate the various aggregator sub-systems 210 within the various datacenters 206 at predetermined (e.g., defined by the computing resource monitoring service 202 or the computing resource service provider) time intervals. This evaluation may be used to generate the partition schedule for a time in the future. In some embodiments, the partition schedule can be updated manually in order handle emergent issues with aggregator sub-system 210 availability, required throughput for storage of the measurements, and other issues. The mapping may correspond to a particular retention period for which it is active. Once the retention period has elapsed, the metric mapping registry 212 may make the next mapping current while keeping a history of the recent changes to the partition schedule to enable GET requests to be routed to the correct aggregator sub-systems 210.

[0044] The partitioner sub-system 208 may utilize the mapping from the metric mapping registry 212 to determine, for each logical partition, a destination aggregator sub-system 210

for the partitioned measurements. The partitioner sub-system 208 may select, based at least in part on the FQMI and timestamp of the measurement obtained through the PUT request, a logical partition from the plurality of logical partitions for storage of the measurement. As noted above, the FQMI may be created using a hash function to hash the metric attributes provided in the metadata. For instance, the FQMI may include a multiple-byte hash, wherein a portion of the hash is used to determine the logical partition for each measurement metric, while another portion of the hash is used for unique identification of the measurement within the logical partition. This logical partition may be mapped to a particular aggregator sub-system 210 by the metric mapping registry 212.

[0045] The partitioner sub-system 208 may store the partitioned measurements within file-based queues each bound to a single aggregator sub-system 210 within the datacenter 206. The measurements within each file-based queue may be batched and delivered asynchronously to each aggregator sub-system 210 at certain short time intervals (e.g., every few milliseconds, etc.). This may enable the partitioner sub-system 208 to deliver measurements from each logical partition to only one aggregator sub-system 210 per datacenter 206 during each retention period, as determined by the metric mapping registry 212. Each logical partition may be further sub-partitioned by the measurement timestamp. The partitioner sub-system 208 may prioritize delivery of the most recent measurements to the aggregator sub-systems 210 first. If at any point a file-based queue is filled to capacity, the partitioner sub-system 208 may purge the oldest measurements from the file-based queue. Further, any measurements for timestamps earlier than the oldest retention period may also be purged. The partitioner sub-system 208 may maintain a separate unpartitioned queue for measurements whose timestamps correspond to a time in the future (e.g., the timestamp specifies a time beyond the latest retention period time). These future observations may be processed only when the timestamp for these measurements becomes current and is within the retention period.

[0046] As noted above, the aggregator sub-system 210 may be an in-memory key-value storage system optimized for aggregating measurements for various metrics in real time and serving time series measurements for the various metrics. The aggregator sub-system 210 may obtain the measurements from the logical partitions of the partitioner sub-systems 208 within the particular datacenter 206 and place these measurements into one or more file-based queues for asynchronous processing by dedicated aggregator datastores of the aggregator sub-system 210. The aggregator datastore may obtain a measurement for a metric from the file-based queue and determine whether there are existing measurements for the

particular metric over an aggregation period. If there are no existing measurements for the particular metric, the aggregator datastore may store the serialized measurement and the aggregator sub-system 210 may update the metric mapping registry 212 to indicate that the serialized measurement has been successfully stored within the datastore. Otherwise, the

5 aggregator sub-system 210 may retrieve the existing serialized measurements from the aggregator datastore and de-serialize these measurements and the newly obtained measurements from the queue. The aggregator sub-system 210 may aggregate the two de-serialized sets of measurements. These aggregated measurements may then be serialized and stored within the aggregator datastore. The aggregator sub-system may update the metric
10 mapping registry 212 to indicate that storage of the measurements from the logical partition was performed successfully.

[0047] At any time, the front-end server 204 may receive a GET request from a customer computer system, computing resource, administrator, or other authorized entity to retrieve measurements for a particular metric from the one or more aggregator datastores in the one or
15 more aggregator sub-systems 210. The GET request may include metadata for the measurements to be obtained from the computing resource monitoring service 202, which may be similar to metadata provided at an earlier time to store a measurement for the metric. The front-end server 204 may hash the received metadata to generate an FQMI that can be used to locate the measurements being sought. The front-end server 204 may transmit a
20 query, along with the FQMI, to the metric consolidation engine 214 in order to obtain the measurements required to fulfill the GET request. In response to the query, the metric consolidation engine 214 may access the metric mapping registry 212 to determine where the measurements needed to fulfill the GET request are stored.

[0048] In order to determine the aggregator sub-systems 210 that have the measurements
25 necessary to fulfill the GET request, the metric consolidation engine 214 may provide the FQMI to the metric mapping registry 212. The metric mapping registry 212 may analyze the current mapping and past mappings for a number of retention periods to determine where the measurements are stored. Once the metric consolidation engine 214 determines, based at least in part on the information provided by the metric mapping registry 212, the aggregator
30 datastores from the aggregator sub-systems 210 where the measurements are stored, the metric consolidation engine 214 may access the aggregator datastores within each datacenter 206 to obtain the requested measurements.

[0049] In an embodiment, the metric consolidation engine 214 compiles the measurements obtained from each of the datacenters 206. If the compiled measurements for the datacenters 206 are inconsistent, the metric consolidation engine 214 may utilize one or more conflict resolution rules to determine which measurements are to be provided in response to the GET request. For instance, the metric consolidation engine 214 may select the response with the highest sample count from the various datacenters 206. Once the metric consolidation engine 214 has resolved any conflicts among the datacenters, if any, the metric consolidation engine 214 may provide the measurements to the front-end server 204. This may enable the front-end server 204 to fulfill the GET request by providing the measurements to the requesting entity.

[0050] The front-end server 204, in some embodiments, is further configured to provide measurements for various metrics to one or more computing resource managers of an auto-scale group once the measurements have been aggregated. The metric mapping registry 212 may maintain one or more entries for computing resources of the auto-scale group, which may specify, for each retention period, the aggregator datastores within the various datacenters 206 that include measurements for these computing resources. Additionally, the metric consolidation engine 214 may be configured to periodically access the metric mapping registry 212 to identify these aggregator datastores that include these measurements. This may enable the metric consolidation engine 212 to obtain the measurements for these computing resources and provide the measurements to the front-end server 204. The front-end server 204 may transmit the measurements to the one or more computing resource managers of the auto-scale group to enable the computing resource managers to scale the computing resources accordingly. In some instances, the one or more computing resource managers may instead transmit GET requests to the front-end server 204, which may fulfill the GET request by performing the process described above to obtain the requested measurements for the one or more computing resources.

[0051] As noted above, the computing resource monitoring service may include, within one or more datacenters, one or more partitioner sub-systems. A partitioner sub-system may include one or more computer systems that are collectively configured to publish measurements from customer computer systems to aggregator sub-systems for storage. Accordingly, FIG. 3 shows an illustrative example of an environment 300 in which a partitioner sub-system 302 of a computing resource monitoring service partitions measurements from various customer computer systems into various logical partitions and delivers measurements from each logical partition to one or more aggregator sub-systems

in accordance with at least one embodiment. The partitioner sub-system 302 may include a metric ingestion engine 304, which may be a module of the one or more computer systems of the partitioner sub-system 302 that is configured to obtain serialized measurements from the front-end server.

5 [0052] The metric ingestion engine 304 may access the metric mapping registry 310 to determine a method for partitioning the obtained measurements from the front-end server into one or more logical partitions for storage in one or more aggregator sub-systems 312. For instance, the metric ingestion engine 304 may utilize the metric mapping registry 310 to obtain a mapping of logical partitions to aggregator sub-systems 312 for a particular retention
10 period. The metric ingestion engine 304 may partition the received measurements into one or more logical partitions based at least in part on the received FQMI for each measurement, the timestamp of each measurement, and the number of aggregator sub-systems 312 as defined by the metric mapping registry 310 through the mapping.

[0053] The FQMI, as noted above, may be a hash of the metadata for a measurement. A
15 portion of the FQMI may be used to determine the logical partition for the associated metric, while another portion of the FQMI may be used for identification of the measurement within a logical partition. The partitioning of the measurements may be performed with consistent hashing to minimize the amount of measurements that are transmitted between aggregator sub-systems 312 as an aggregator sub-system 312 is added or removed from the particular
20 datacenter. The partitioning of the measurements may also be based at least in part on the measurement timestamp for each measurement to eliminate any impact due to clock skew.

[0054] Once the metric ingestion engine 304 has partitioned the obtained measurements into one or more logical partitions based at least in part on the mapping from the metric mapping registry 310, the metric ingestion engine 304 may place the measurements from the
25 one or more logical partitions into one or more file-based queues 308. Each file-based queue 308 may be bound to a single aggregator sub-system 312, as illustrated in FIG. 3. This binding of the file-based queue 308 to a particular aggregator sub-system 312 may be based at least in part on the logical partition associated with the file-based queue 308, as each logical partition may be mapped to a particular aggregator sub-system 312 based at least in
30 part on the mapping. The measurements in each file-based queue 308 may be batched and delivered asynchronously to each aggregator sub-system 312 at regular time intervals (e.g., every few milliseconds, every few seconds, etc.) based at least in part on the time series of the measurements.

[0055] Within each file-based queue 308, the measurements may be further sub-partitioned based at least in part on the timestamp for each measurement. This may enable the partitioner sub-system 302 to prioritize delivery of the measurements such that the most recent measurements are delivered to the aggregator sub-systems 312 first. In an embodiment, if a file-based queue 308 reaches capacity, the partitioner sub-system 302 identifies the measurements with the oldest timestamp and purges these measurements from the queue 308. Further, any measurements with timestamps earlier than the oldest retention period may also be purged from the partitioner sub-system 302. This may enable the metric ingestion engine 304 to continuously populate the file-based queues 308 with new measurements obtained from the front-end server, which may have obtained the measurements from the customer computer systems.

[0056] The partitioner sub-system 302 may further include a future observation queue 306, which may be a separate unpartitioned queue not associated with any aggregated sub-systems 312. The metric ingestion engine 304 may transfer any measurements with timestamps in the future (e.g., measurements with timestamps later than the newest retention period) to the future observation queue 306 as the metric mapping registry 310 may not include a mapping of logical partitions to the aggregator sub-systems 312 for a retention period that would include these measurements. Thus, the metric ingestion engine 304 may only process measurements in the future observation queue 306 if the timestamps for these measurements become current and a mapping is available for storing the measurements in a logical partition and transmitting the measurements from the logical partition to the corresponding aggregator sub-system 312.

[0057] As noted above, the computing resource monitoring service may include, within each datacenter, one or more aggregator sub-systems configured to aggregate metric measurements in real time and to serve time series measurements for a variety of metrics. An aggregator sub-system may include one or more computer systems collectively configured to perform the aforementioned tasks, storing measurements within one or more aggregator datastores and making these measurements available for fulfillment of GET requests from customer computer systems and other entities. Accordingly, FIG. 4 shows an illustrative example of an environment 400 in which an aggregator sub-system 402 performs aggregation of measurements and provides read access for measurements stored within one or more datastores 408 in accordance with at least one embodiment. In the environment 400, the aggregator sub-system 402 may include a metric processing engine 404, which may comprise a module operating within the one or more computer systems of the aggregator sub-system

402. The metric processing engine 404 may receive the measurements from the logical partitions of the various partitioner sub-systems for storage in one or more datastores 408. In response to the PUT requests, the metric processing engine 404 may place the received one or more measurements into file-based queues 406 for asynchronous processing by a dedicated aggregator datastore 408.

[0058] Each aggregator datastore 408 may comprise one or more storage devices and at least one module configured to obtain measurements from a file-based queue 406 to store the measurements within the one or more storage devices. In an embodiment, the aggregator datastore 408 obtains a measurement from the file-based queue 406 and determines whether measurements for the same metric over the same aggregation period are stored within the aggregator datastore 408. If so, the aggregator datastore 408 may de-serialize the stored measurements and the measurement from the queue 404 in order to aggregate the measurements for the particular metric. The aggregator datastore 408 may obtain the results of the aggregation of the two sets of measurements and serialize the aggregated measurements using the binary serialization format described above. Once the aggregated measurements have been serialized, the aggregator datastore 408 may store these measurements within the one or more storage devices of the datastore 408 and transmit a notification to the metric processing engine 404 that the latest measurement has been stored successfully. Alternatively, if it is the first time that a measurement has been observed in an aggregation period, the measurement may simply be stored within the aggregator datastore 408 without need for de-serialization, as the measurement may already be in the binary serialization format. Once the metric processing engine 404 receives confirmation that a measurement has been stored successfully within an aggregator datastore 408, the metric processing engine 404 may access the metric mapping registry 410 to modify an entry within the mapping to specify that storage of the measurement within the aggregator sub-system 402 has been completed successfully.

[0059] The metric processing engine 404 may be configured to also process GET requests from a metric consolidation engine 410 of the computing resource monitoring service. As noted above, the metric mapping registry may include a mapping of logical partitions and measurements to aggregator sub-systems 402 within each datacenter. The metric consolidation engine may utilize this mapping to identify the aggregator sub-systems 402 that have the requisite measurements to fulfill the GET request from a customer computer system. Based at least in part on the mapping, the metric consolidation engine may transmit a request

to the metric processing engine 404 of each of these identified one or more aggregator sub-systems 402 to obtain the needed measurements.

[0060] Once the metric processing engine 404 receives the GET request from the metric consolidation engine, the metric processing engine 404 may determine which aggregator datastores 408 have the necessary measurements to fulfill the request. The metric processing engine 404 may transmit a query including the FQMI and time range to each aggregator datastore 408. Each aggregator datastore 408 may utilize the provided information to identify any measurements that may be used to fulfill the request. If one or more measurements are identified, the aggregator datastore 408 may provide these measurements to the metric processing engine 404. The metric processing engine 404 may compile the measurements that may be used to fulfill the GET request and provide the compiled measurements to the metric consolidation engine 410 for further processing.

[0061] In an embodiment, the aggregator datastores 408 are configured to purge expired measurements from the one or more storage devices once the oldest retention period has been replaced by a later retention period. For instance, the computing resource monitoring service may serve GET requests for one or more consecutive retention periods. As a new retention period begins, the computing resource monitoring service may maintain a number of consecutive retention periods for fulfillment of GET requests from customer computer systems or other computing resources. The computing resource monitoring service may also maintain measurements for an older retention period that has become eligible for expiry but may not yet have been purged. When measurements for this older retention period have become eligible for expiry, the aggregator datastores 408 may purge any measurements with a timestamp that is within this older retention period.

[0062] As noted above, the customer computer system may internally perform the partitioning of the measurements prior to aggregation and storage through use of one or more aggregator sub-systems within the computing resource monitoring service. This architecture may obviate the need for the computing resource monitoring service to maintain the one or more partitioner sub-systems described above in connection with FIGS. 2 and 3, as the computing resource monitoring service may obtain the serialized measurements directly from the logical partitions of the customer computer system. Accordingly, FIG. 5 shows an illustrative example of an environment 500 in which a customer computer system 502 partitions measurements into various logical partitions and delivers these measurements to

one or more aggregator sub-systems 512 of a computing resource monitoring service in accordance with at least one embodiment.

[0063] In the environment 500, the customer computer system 502 may include a metric ingestion engine 504, which may be a module of the customer computer system 502 that is configured to obtain measurements from computing resources associated with the customer computer system 502 or from the customer computer system 502 itself. The metric ingestion engine 504 may perform similar operations as the metric ingestion engine described above in connection with FIG. 3. For instance, the metric ingestion engine 504 may access the metric mapping registry 510 within the computing resource monitoring service to determine a method for partitioning the obtained measurements into one or more logical partitions for storage in one or more aggregator sub-systems 512. The metric ingestion engine 504 may partition the received measurements into one or more logical partitions based at least in part on the received FQMI for each measurement, the timestamp of each measurement, and the number of aggregator sub-systems 512 as defined by the metric mapping registry 510 through the mapping. As opposed to the process described above in connection with FIG. 2, the customer computer system 502 may obtain the metadata for each measurement and utilize a hash function to generate an FQMI for the measurement. Additionally, the customer computer system 502 may transform the measurements into a serialized format for processing by the metric ingestion engine 504. The serialized format for a particular measurement obtained from the customer computer system 502 may include the FQMI, the timestamp for the measurement, the unit of measurement for the observation, and the measurement itself.

[0064] Once the metric ingestion engine 504 has partitioned the measurements into one or more logical partitions based at least in part on the mapping from the metric mapping registry 510, the metric ingestion engine 504 may place the one or more measurements from each logical partition into a corresponding file-based queue 508. Similar to the file-based queues illustrated in FIG. 3, each file-based queue 508 may be bound to a single aggregator sub-system 512. The measurements in each file-based queue 508 may be batched and delivered asynchronously to each aggregator sub-system 512 at regular time intervals (e.g., every few milliseconds, every few seconds, etc.) based at least in part on the time series of the measurements. For instance, at each time interval, the customer computer system 502 may transmit a PUT request that includes the measurements to one or more aggregator sub-systems 512 selected based at least in part on the mapping. The PUT request may be transmitted to these aggregator sub-systems 512 using a communications protocol compatible for both the customer computer system 502 and the aggregator sub-systems 512. If a PUT

request is unsuccessful, the communications protocol may cause a notification to be generated specifying that the transmission of the measurements within the logical partition has failed. Further, the communications protocol may cause the customer computer system 502 to refresh the communications channel between the customer computer system 502 and the aggregator sub-systems 512. This may enable the customer computer system 502 to again attempt transmission of the PUT request to the aggregator sub-systems 512.

[0065] Within each file-based queue 508, the measurements may be further sub-partitioned based at least in part on the measurement timestamps. This may enable the customer computer system 502 to prioritize delivery of the measurements such that the most recent measurements are delivered to the aggregator sub-systems 512 first. Any measurements with timestamps earlier than the oldest retention period may be purged from the customer computer system 502 or delivered to an alternative metric monitoring service. This may enable the metric ingestion engine 504 to continuously populate the file-based queues 508 with new measurements generated or obtained by the customer computer system 502.

[0066] Similar to the partitioner sub-system illustrated in FIG. 3, the customer computer system 502 may further include a future observation queue 506, which may be a separate unpartitioned queue not associated with any aggregated sub-systems 512. The metric ingestion engine 504 may transfer any measurements with timestamps in the future (e.g., measurements with timestamps later than the newest retention period) to the future observation queue 506 as the metric mapping registry 510 may not include a mapping of logical partitions to the aggregator sub-systems 512 for a retention period that would include measurements with these timestamps. Thus, the metric ingestion engine 504 may only process measurements in the future observation queue 506 if the timestamps for these measurements become current and a mapping is available for generating a logical partition for these measurements and transmitting the measurements from the logical partition to the corresponding aggregator sub-system 512.

[0067] In an alternative embodiment, the metric ingestion engine 504 can provide unpartitioned measurements to the any number of aggregator sub-systems 512 of the computing resource monitoring service. The one or more aggregator sub-systems 512, upon obtaining the unpartitioned measurements, may access the metric mapping registry 510 to determine, based at least in part on the FQMI and timestamps for the measurements, which aggregator sub-systems 512 are to be used for storing the unpartitioned measurements. If an aggregator sub-system 512 determines, based at least in part on one or more mappings from

the metric mapping registry 510, that unpartitioned measurements are to be stored within an in-memory datastore of another aggregator sub-system, the aggregator sub-system may transmit these unpartitioned measurements to the other aggregator sub-system, which may then store the measurements within its one or more in-memory datastores.

- 5 [0068] FIG. 6 shows an illustrative example of an environment 600 in which metadata is provided with a measurement to enable future storage of measurements without requiring additional metadata transmissions in accordance with at least one embodiment. In the environment 600, a customer computer system 602 may transmit a PUT request to a computing resource monitoring service to store a measurement for a particular metric within
10 aggregator sub-systems 610 of one or more datacenters 606. The PUT request may include, in addition to the measurement that is to be stored, metadata 604 for the measurement. This metadata 604 may include one or more metric attributes that may uniquely identify the associated measurement within the computing resource monitoring service. For instance, the metadata 604 may specify, individually or in combination, the account number of a customer
15 associated with the customer computer system 602 submitting the PUT request or with another computing resource responsible for generation of the measurement, the name of the particular service responsible for managing the computing resource (e.g., virtual computer system service, database service, etc.), the dimensions for the measurement (e.g., key/value pairs), and an identifier for the computing resource producing the measurement.
- 20 [0069] When the computing resource monitoring service receives the measurement and metadata 604 from the customer computer system 602, the computing resource monitoring service may utilize a hash function to generate a hash of the metadata 604, which can be used as an FQMI for the measurement that is to be stored. Further, the computing resource monitoring service may utilize a binary serialization format to compress the measurement to
25 be stored. In this binary serialization format, the serialized measurement may include the FQMI generated by hashing the metadata 604, the timestamp for the measurement, the unit of measurement for the measurement, and the measurement itself. Once the computing resource monitoring service has transformed the measurement into a serialized format, the computing resource monitoring service may transmit the measurement to a partitioner sub-system 608
30 within each datacenter 606. The partitioner sub-system 608 may determine a logical partition from a plurality of logical partitions for storage of the measurement based at least in part on the FQMI and timestamp of the received measurement, as well as the number of active aggregator sub-systems 610 within the particular datacenter 606 for one or more retention periods, as specified in the mapping of logical partitions to aggregator sub-systems 610. The

partitioner sub-system 608 may transmit the measurement from a queue associated with the logical partition to the aggregator sub-systems 610 within the datacenter 606 according to a mapping of the logical partitions to the aggregator sub-systems 610 obtained from a metric mapping registry.

5 [0070] In an embodiment, once the computing resource monitoring service has hashed the metadata 604 to generate the FQMI for the measurement to be stored, the customer computer system 602 is no longer required to repeat the metadata 604 for future submissions of measurements for the metric, so long as the metadata 604 has not expired or the customer computer system 602 has not initiated a new session with the computing resource monitoring service. For instance, if the customer computer system 602 submits an additional PUT request
10 at a later time to the computing resource monitoring service, the customer computer service 602 may provide the FQMI (e.g., metadata hash 612) for the measurement corresponding to the same metric instead of the metadata 604. The computing resource monitoring service may serialize the measurement and the FQMI in anticipation of storage of the measurement within the datastores of the computing resource monitoring service. Alternatively, if the customer
15 computer system 602 does not provide the FQMI with the measurement, the computing resource monitoring service may identify, based at least in part on the measurement and the metric, the FQMI that is to be utilized based at least in part on metric entries within the metric mapping registry. For instance, if the PUT request includes a continuation of measurements for a particular time series, the computing resource monitoring service may
20 determine the FQMI based at least in part on previously obtained measurements for the time series. These previously obtained measurements may have been obtained through an earlier PUT request to the computing resource monitoring service to store the measurements. This earlier PUT request may have included the metadata 604, from which the FQMI was
25 generated.

[0071] At any time, the customer computer system 602 may submit a GET request to retrieve measurements from one or more aggregator sub-systems 610 of the computing resource monitoring service. The GET request may include the metadata 604 for the measurements that are to be obtained. The computing resource monitoring service may hash
30 the received metadata 604 to generate an FQMI that can be used to locate the measurements being sought. The computing resource monitoring service may transmit a query, along with the FQMI, to the metric consolidation engine 614 in order to obtain the measurements required to fulfill the GET request. In response to the query, the metric consolidation engine 614 may access the metric mapping registry to determine where the measurements needed to

fulfill the GET request are stored. The metric consolidation engine 614 may access the metric mapping registry and utilize the FQMI and timestamps for the measurements to determine which aggregator sub-systems 610 have the data necessary to fulfill the GET request. Once the metric consolidation engine 614 determines, based at least in part on the information
5 provided by the metric mapping registry, the aggregator datastores from the aggregator sub-systems 610 where the measurements are stored, the metric consolidation engine 614 may access the aggregator datastores within each datacenter 606 to obtain the requested measurements.

[0072] As noted above, the computing resource monitoring service may make available
10 measurements to one or more computing resource managers of an auto-scale group to enable these computing resource managers to quickly determine whether additional computing resources are necessary within the auto-scale group based at least in part on the obtained measurements. Accordingly, FIG. 7 shows an illustrative example of an environment 700 in which ingestion of measurements is used to determine whether automatic scaling of one or
15 more resources is to be performed in accordance with at least one embodiment. In this example, the networked environment 700 includes a computing resource service provider 704 in data communication with a client device 706 and server computers 742 over a network 706. In one embodiment, the server computers 742 comprise one or more computer hardware devices that are used to implement instances 720 (e.g., computing resources). For example,
20 the server computers 742 may include hardware for implementing types of computing resources, such as storage devices, virtualized storage devices, networking devices, and the like. Additionally, the implemented computing resources may be programmatically and remotely managed by a customer of the distributed computing resource service provider 704.

[0073] The server computers 742 include a plurality of computer system devices that are
25 each capable of executing one or more instances 720 created by the computing resource service provider 704. In one embodiment, each of the server computers 742 includes a processor, a data store, an input/output bus, and/or any other component known in the art for executing instances 720. Additionally, the instances 720 may be virtual machine instances. As known in the art, a virtual machine instance is an instance of a software implementation on a
30 machine (i.e., a computer) that executes programs like a physical machine. For example, each of the server computers 742 may be configured to execute an instance manager 718 capable of implementing the instances 720. For example, the instance manager 718 may be a hypervisor, virtualization layer, or another type of program configured to enable the execution of multiple instances 720 on a single server computer 742. As discussed above,

each of the instances 720 may be configured to execute all or a portion of an application. The networked environment 700 may span one or more datacenters, where each datacenter may be geographically distinct from each other. Additionally, the networked environment 700 shown in FIG. 7 may be one of several embodiments employed by the computing resource service provider 704.

[0074] In one embodiment, the computing resource service provider 704 includes a load balancer database 710, an instance service 712, a placement service 726, an auto-scale service 730, a maintenance service 732, a computing resource monitoring service 734, a load balancing service 736, and/or other components. The load balancer database 710 may include load balancer data 742. For example, the load balancer database 710 may include one or more records of load balancers 740 associated with the auto-scale group 702. Each one of the records of the load balancer data 746 may correspond to a load balancer 740 of the networked environment 700.

[0075] The instance service 712 instantiates instances 720 based at least in part on a set of preferences provided by the customer. In one embodiment, the instance service 712 receives, from the customer on the client device 706, a request 716 to create one or more instances 732 and optionally assign the created instances 720 to an auto-scale group 702. Additionally, the request 716 received from the customer on the client device 706 may also indicate a time to start execution of the requested instances 720. In response to receiving the request, the instance service 712 instantiates instances 720. In various embodiments, the auto-scale service 730 receives the request and transmits a command to the instance service 712 to instantiate the instances 720 such that the instances are associated with the auto-scale group, for example, by associating auto-scale group 702 metadata with the instances 720. In one embodiment, the instance service 712 may place instances in standby or detach instances from the auto-scale group in response to a request from the client device 706 and/or auto-scale service 730. For example, the auto-scale service 730 may transmit a request to the instance service 712 to remove the auto-scale group 702 metadata associated with the instances 720 being detached from the auto-scale group 702 according to the request 716. Furthermore, the auto-scale service 730 may de-assign a detached instance from the load balancer 740.

[0076] The customer may interact with the computing resource service provider 704 (via appropriately configured and authenticated API calls) to provision, operate, and manage instances 720 associated with the auto-scale group 702 that is instantiated on server

computers 742 and operated by the computing resource service provider 704. Additionally, the customer may create one or more auto-scale groups 702, and the auto-scale groups 702 may be a logical collection of instances 720. Furthermore, the instances 720 may be assigned to the auto-scale group 702 or may be members of the auto-scale group 702. The auto-scale service 730 may allow customers to interact with and manage various auto-scale groups 702. For example, the customer may, through the auto-scale service 730, set a maximum or minimum capacity for an auto-scale group 702. The auto-scale group 702 may then manage the instances 720 assigned to the auto-scale group in order to maintain the settings provided by the customer. In various embodiments, the customer can create and manage auto-scale groups 702 through a management console, as described above, provided by the computing resource service provider 704.

[0077] The management console may be exposed to the customers as a webpage; by interacting with the webpage (e.g., through a browser application) the customer may cause API calls to be generated. The generated API calls may cause the computing resource service provider 704 or component thereof to perform various operations indicated by the customer. Once the customer has created the auto-scale group 702, the customer may assign one or more load balancers to the auto-scale group 702 by submitting requests 716. The requests 716, in this case, may be processed by the auto-scale service 730 or other component of the computing resource service provider 304. The instances 720 of the auto-scale group 702 may be used for various purposes, such as to operate as servers supporting a website, to operate business applications, or generally, to serve as computing power for the customer. Other applications for the instances 720 may be to support database applications, electronic commerce applications, business applications and/or other applications. Additionally, load balancers 740 may distribute traffic to various instances 720 of the auto-scale group 702 to enable operation of the instances for the various purposes described above and prevent the instances 720 of the auto-scale group 702 from being overloaded. Although the instance service 712 is shown in FIG. 7, any other computer system or computer system service may be utilized by the computing resource service provider 702, such as a computer system or computer system service that does not employ virtualization or instantiation and instead provisions computing resources on dedicated or shared computers/servers and/or other physical devices.

[0078] The placement service 726 provisions the instances 720 to one or more of the server computers 742. In one embodiment, the placement service 726 determines the server computers 742 to provision the new instances 720 based at least in part on the indicated auto-

scale group 702 of the new instances 720. For example, the placement service 726 may identify one or more server computers 742 with the appropriate capacity to execute the instances 720. To this end, the placement service 726 determines the capacity of each server computer 742 from the resource data 710 stored in the data store and accordingly provisions the instances 730, as will be described. The auto-scale service 730 automatically scales the capacity of a collection of previously requested instances 720 up or down based at least in part on circumstances defined by the customer that requested the instances 720. For example, the auto-scale service 730 may decrease the number of instances 720 allocated to the customer during demand lulls and increase the number of instances 720 allocated to the customer during demand peaks.

[0079] In one embodiment, the auto-scale service 730 sheds a subset of the requested instances 720 during a period of low usage and/or idle time. For example, the auto-scale service 730 may determine that the amount of instances 720 requested by the customer is redundant and/or excessive. In response, the auto-scale service 730 may terminate a certain number of instances 320 allocated to the customer such that the remaining number of instances 720 allocated to the customer is not redundant and/or excessive. In another embodiment, the auto-scale service 730 may shed the subset of the requested instances 730 if the usage rate does not exceed a predetermined threshold. Similarly, the auto-scale service 730 increases the amount of instances 320 during a period of high usage. In one embodiment, the auto-scale service 730 may increase the amount of instances 720 if the usage rate exceeds a predetermined threshold.

[0080] The maintenance service 732 schedules maintenance, software updates, and/or firmware updates for the server computers 742. In one embodiment, the maintenance service 732 schedules the maintenance and software updates at an appropriate time based at least in part on the available capacity of the server computers 742. For example, the maintenance service 732 may schedule the maintenance and software updates at a time when the respective server computer 742 has a projected availability. In one embodiment, the maintenance service 732 may patch and restart the server computers 742 when the maintenance service 732 determines that the server computer 742 is not hosting any instances 720. Additionally, the maintenance service 732 may patch virtual machines associated with the instance 742 if necessary prior to instantiating new images that are associated with the respective virtual machines. For example, the maintenance service 732 may schedule a patch of the machine image based at least in part on the health status of the instances 720. In one

embodiment, no additional instances may be provisioned on the server computer 742 until the scheduled maintenance is completed.

[0081] The maintenance service 732 may also periodically or aperiodically check the health status of the instances 720, including instances assigned to the auto-scale group 702 and/or load balancers 740. The health check may include determining the load, utilization, and operation of various components of the instances 720 such as the central processing unit, memory, networking interface, operating system, application, and other components of the instances 720. In various embodiments, when the maintenance service 732 determines that an instance 720 is unhealthy, based at least in part on the health check, the maintenance service 732 or other component of the service provider 704, such as the auto-scale service 730, may initiate a workflow to remove the unhealthy instances from the auto-scale group 702.

Additionally, if the maintenance service 732, determines that a previously unhealthy instance 720 has returned to a healthy status the maintenance service 732 or other component of the service provider 704, such as the auto-scale service 730, may move the instances 720 into service or attach the instances 720 to the auto-scale group 702. Furthermore, if an instance 722 assigned to a load balancer 740 returns a healthy status the auto-scale group 702 may update the status of the load balancer 740 to in-service, as described above.

[0082] The computing resource monitoring service 734 may be responsible for collecting resource data corresponding to the instances 720. The resource data obtained by the computing resource monitoring service 734 may indicate the utilization of various components of the instances 720 such as the central processing unit, memory, networking interface, operating system, applications, and other components of the instances 720. This information may be used for a variety of different purposes, for example, determining whether to allocate or deallocate resources to the auto-scale group 702. Additionally, the information may be used by the maintenance service 732 to determine the health of an instance 720 and/or a server computer 742. The computing resource monitoring service 734 may obtain and aggregate utilization information for all of the instances 720 assigned to the auto-scale group 702.

[0083] In an embodiment, the computing resource monitoring service 734 obtains, from the one or more load balancers 740 for the instances 720 of the auto-scale group 702, one or more PUT requests to store measurements for the one or more instances 720 within one or more in-memory datastores of the computing resource monitoring service 734. Each PUT request may include metadata for each measurement, specifying the metric attributes for the

measurement. The computing resource monitoring service 734, through a front-end server, may obtain the measurement and the metadata from a load balancer 740. The front-end server may utilize the metadata and a hash function to generate an FQMI for the measurement that is to be stored within the aggregator datastore of the computing resource monitoring service 734. Additionally, the front-end server may transform the measurement into a serialized format for publication to a partitioner sub-system within the one or more datacenters of the service 734. The front-end server may transmit the serialized measurement from the instance 720 to a partitioner sub-system.

[0084] The partitioner sub-system may determine, based at least in part on the FQMI and the timestamp for the measurement, which logical partition of a plurality of logical partitions will be used for storing the measurement. The partitioner sub-system may access a metric mapping registry, which may provide a mapping of logical partitions to the aggregator sub-systems active for one or more retention periods. This may enable the partitioner sub-system to determine which aggregator sub-system of the computing resource monitoring service 734 the measurement will be stored in for the processing of GET requests. The partitioner sub-system may store the measurement within a file-based queue associated with the selected logical partition, bound to a single aggregator sub-system.

[0085] Once the aggregator sub-system of the computing resource monitoring service 734 obtains the measurement from the partitioner sub-system, the aggregator sub-system may determine whether measurements for the same metric over the same aggregation period are stored within an aggregator datastore. If so, the aggregator datastore may de-serialize the stored measurements and the measurement from the partitioner sub-system in order to aggregate the measurements for the particular metric. The aggregator datastore may obtain the results of the aggregation of the two sets of measurements and serialize the aggregated measurements using the binary serialization format described above. Once the aggregated measurements have been serialized, the aggregator datastore may store these measurements within the one or more storage devices of the datastore and transmit a notification to the metric processing engine 404 that the latest measurement has been stored successfully. Alternatively, if it is the first time that a measurement has been observed in an aggregation period, the measurement may simply be stored within the aggregator datastore without need for de-serialization.

[0086] In an embodiment, the front-end server of the computing resource monitoring service 734 is further configured to provide measurements for various metrics to the one or

more instance managers 718 of the auto-scale group 702 once the measurements have been aggregated, after a particular time period has passed, and/or in response to a triggering event, such as a spike in the measurements for a particular metric. The metric mapping registry may maintain one or more entries for the instances 720 of the auto-scale group 702, which may specify, for each retention period, the aggregator datastores within the various datacenters of the computing resource monitoring service 734 that include measurements for these instances 720. Additionally, the metric consolidation engine of the computing resource monitoring service 734 may be configured to periodically access the metric mapping registry to identify these aggregator datastores that include these measurements. This may enable the metric consolidation engine to obtain the measurements for these instances 720 and provide the measurements to the front-end server. The front-end server may transmit the measurements to the one or instance managers 718 of the auto-scale group 702 to enable the instance managers 718 to automatically scale the instances 720 accordingly. In some instances, the one or more instance managers 718 may instead transmit GET requests to the computing resource monitoring service 734, which may fulfill the GET request by performing the process described above to obtain the requested measurements for the one or more instances 720.

[0087] In an embodiment, when the one or more instance managers 718 of the auto-scale group 702 obtains the measurements from the front-end server, the one or more instance managers 718 can determine whether to perform modification of any of the one or more instances 720 of the auto-scale group 702. For instance, the one or more instance managers 718, based at least in part on the obtained measurements for a particular instance 720, may change the instance type to address any issues highlighted by the obtained measurements. For example, if the one or more instance managers 718 determine that based at least in part on the measurements a particular instance 720 requires additional processing power and memory, the one or more instance managers 718 may provision an instance that provides greater processing power and memory, thereby enabling the one or more instance managers 718 to transfer the instance assets from the original instance to this newly provisioned instance.

[0088] The one or more instance managers 718 may perform modification of the one or more instances 720 after evaluation of measurements over a particular period of time has passed. For instance, if the obtained measurements for a particular retention period are indicative of a need to automatically scale the auto-scale group 702 (e.g., add instances 720, remove instances 720, modify existing instances 720, etc.), the one or more instance managers 718 may await modification of the auto-scale group 702 until it obtains and evaluates measurements for additional retention periods. For example, if the measurements

for subsequent retention periods indicate that there it is not necessary to perform any modifications of the auto-scale group 702, the one or more instance managers 718 may not perform these modifications, as the measurements indicative of a need to automatically scale the auto-scale group 702 may be deemed an aberration.

5 [0089] A load balancing service 736 may be offered to customers of a computing resource service provider 704 in order to facilitate request processing by instances 720 of the customer. In various embodiments, the instances 720 may be assigned to the auto-scale group 702 and the load-balancer service 736 may distribute traffic to the instances 722 assigned to the auto-scale group 702. For example, the customer may operate a website using instances
10 720 assigned to the auto-scale group 702 using the resources of computing resource service provider 704. Additionally, the website may receive requests from multiple other customers over the network 706. The computing resource service provider 704 may configure a load balancer of the load balancing service 736 to direct the requests to the instances 720 of the auto-scale group 702 executing the website, in such a way that the load generated by
15 processing the requests is distributed among the instances 720 of the auto-scale group 702 executing the website. The load balancing service 736 may be a computer system or virtual computer system configured to distribute the request to the instances 720 assigned to the load balancer in order to optimize resource utilization and/or avoid overloading a particular server computer 742. For example, the load balancer may include physical hardware connected to a
20 server rack or otherwise included in a data center. In another example, the load balancer may include one or more virtual machines supported by the server computer 742.

[0090] As noted above, a computing resource monitoring service may include one or more partitioner sub-systems, which may serve as entry points for the publication of measurements for various metrics. The partitioner sub-systems may obtain measurements from a front-end
25 server and split these measurements into logical partitions based at least in part an FQMI and the timestamp for each measurement and a number of aggregator sub-systems as provided by the metric mapping registry. The partitioner sub-systems may provide these measurements to the various aggregator sub-systems based at least in part on the mapping from the metric mapping registry. Accordingly, FIG. 8 shows an illustrative example of a process 800 for
30 partitioning measurements for a metric for delivery to one or more aggregator sub-systems in accordance with at least one embodiment. The process 800 may be performed by a front-end server of the computing resource monitoring service, which may obtain measurements from various computing resources and/or customer computer systems and performed operations prior to providing the measurements to the aforementioned partitioner sub-system of the

computing resource monitoring service, which may perform various operations of the process 800.

[0091] At any point, the front-end server may receive 802, from a customer computer system or other computing resource, a PUT API call to publish a measurement within an aggregator sub-system for in-memory storage of the measurement within an aggregator datastore. The PUT API call may include the measurement to be stored as well as metadata for the particular measurement. The metadata may specify one or more metric attributes for the measurement. For instance, the metadata may uniquely identify the measurement by a combination of a customer account number, the namespace (e.g., associated service for the computing resource), the dimensions for the measurement (e.g., key/value pairs), the name of the measurement itself, and the like. Thus, the front-end server may obtain 804, from the PUT API call, the measurement that is to be stored in-memory within an aggregator datastore of the computing resource monitoring service.

[0092] Once the front-end server has obtained the measurement from the PUT API call, the front-end server may utilize a hash function and the provided metadata to generate an FQMI for the measurement. This FQMI may include a multiple-byte hash, wherein a portion of the hash may be used to determine a logical partition for each measurement metric based at least in part on a mapping within a metric mapping registry, while another portion of the hash may be used for unique identification of the measurement once placed within the corresponding logical partition. The front-end server may further transform the measurement into a binary serialization format. The measurement, in this format, may include the FQMI, timestamp for the measurement, measurement unit for the measurement, and the measurement itself. The front-end server may transmit this serialized measurement to a partitioner sub-system within each datacenter of the computing resource monitoring service for redundant storage of the measurement. The partitioner sub-system within each datacenter may be selected through use of a partitioner load-balancer.

[0093] When the partitioner sub-system receives the serialized measurement from the front-end server, the partitioner sub-system may determine 806 the FQMI and timestamp for the measurement. The partitioner sub-system may access 808 the metric mapping registry to obtain a mapping of logical partitions to various aggregator sub-systems to determine which logical partition will be used for placement of the measurement and for delivery to an aggregator sub-system. For instance, the partitioner sub-system may utilize the FQMI and

timestamp from the serialized measurement to identify, based at least in part on the mapping, which logical partition to place the measurement in.

[0094] In an embodiment, the partitioner sub-system utilizes the timestamp and the mapping to determine 810 whether the measurement is for a time period beyond the latest retention period, as specified in the mapping. If the measurement includes a future timestamp (e.g., timestamp is for a time beyond the latest retention period), the partitioner sub-system may queue 812 the measurement in a separate unpartitioned queue, where the measurement may remain until the timestamp for the measurement is covered by the present retention period at a later time. As the retention period is updated, the partitioner sub-system may obtain the measurement from this queue and determine whether the timestamp for the measurement is covered by the present retention period. If not, the measurement may remain in the queue.

[0095] If the timestamp for the measurement is for a current retention period, the partitioner sub-system may place the measurement within a logical partition of a plurality of logical partitions of the partitioner sub-system, as determined based at least in part on the timestamp, FQMI, and the mapping. Each logical partition may be associated with a particular queue, which may be used to transmit measurements from the partitioner sub-system to the corresponding aggregator sub-system, as determined through use of the mapping. Thus, when the measurement is placed in a logical partition, the partitioner sub-system may move 814 the measurement into an associated file-based queue bound to a corresponding aggregator sub-system. The measurements within this queue may be partitioned based at least in part on the timestamp for each measurement such that delivery of the most recent measurements may be completed first. Thus, based on this partitioning of the measurements, the partitioner sub-system may deliver 816 the measurement asynchronously from the queue to the corresponding aggregator sub-system for in-memory storage of the measurement.

[0096] As noted above, the partitioner sub-system may transmit serialized measurements to one or more aggregator sub-systems configured to aggregate metric measurements in real time and to serve time series measurements for a variety of metrics. An aggregator sub-system may include one or more computer systems collectively configured to perform the aforementioned tasks, storing measurements within one or more aggregator datastores and making these measurements available for fulfillment of GET requests from customer computer systems and other entities. Accordingly, FIG. 9 shows an illustrative example of a

process 900 for aggregating measurements from one or more partitioner sub-systems with measurements from one or more datastores of a computing resource monitoring service in accordance with at least one embodiment. The process 900 may be performed by any aggregator sub-system of the computing resource monitoring service, which may be
5 configured to process measurements from the partitioner sub-systems and store the measurements within one or more in-memory datastores.

[0097] At any point during a retention period, the aggregator sub-system may receive 902 a new serialized measurement from a partitioner sub-system. As described above in connection with FIG. 8, the partitioner sub-system may include one or more file-based queues, which
10 may be configured to asynchronously transmit data to the one or more aggregator sub-systems within a datacenter. Each queue may be associated with a single aggregator sub-system based at least in part on a mapping for the given retention period included within the metrics mapping registry. Based at least in part on the FQMI and timestamp of the received measurement, the aggregator sub-system may identify a particular in-memory datastore for
15 storage of the received measurement.

[0098] Once the aggregator sub-system has identified the in-memory datastore for storage of the obtained measurement, the aggregator sub-system may determine 904 whether it is the first time that the aggregator sub-system has encountered a measurement for the particular metric for an aggregation period. If this is the first time that the aggregator sub-system has
20 observed a measurement for this metric during the aggregation period, the aggregator sub-system may store 912 the serialized measurement within the identified in-memory datastore. However, if the aggregator sub-system determines that other measurements for the metric over the aggregation period are stored within the identified in-memory datastore, the aggregator sub-system may obtain 906 these other serialized measurements from the in-
25 memory datastore.

[0099] The aggregator sub-system may de-serialize 908 the newly obtained serialized measurement from the partitioner sub-system and the other serialized measurements from the in-memory datastore for aggregation of the measurements. Once the aggregator sub-system has de-serialized the newly obtained serialized measurement and the measurements from the
30 in-memory datastore, the aggregator sub-system may aggregate 910 the newly obtained measurement from the partitioner sub-system with the other measurements for the metric previously stored within the in-memory datastore. This aggregation may include serializing the aggregated measurements into the binary serialization format described above. Once the

aggregator sub-system has aggregated and serialized the measurements, the aggregator sub-system may store 912 the serialized measurements within the in-memory datastore.

[0100] As noted above, the front-end server of a computing resource monitoring service, in conjunction with the metric consolidation engine of the service, may process and fulfill GET requests from customer computer systems and/or other computing resources. For instance, in response to a GET request, the metric consolidation engine may obtain measurements necessary to fulfill the request from one or more aggregator datastores as specified within the metric mapping registry and compile these measurements for delivery to the requesting entity. Accordingly, FIG. 10 shows an illustrative example of a process 1000 for retrieving one or more measurements from one or more aggregator datastores in response to a GET application programming interface call in accordance with at least one embodiment. The process 1000 may be performed by the front-end server, which may obtain the GET requests from various entities, and a metric consolidation engine, which may obtain the measurements and compile these measurements for delivery to the requesting entities.

[0101] At any point, a front-end server of the computing resource monitoring service may receive 1002 a GET API call to obtain one or more measurements from the computing resource monitoring service. The GET API call may include metadata for the measurements to be obtained. The front-end server may utilize a hash function and the metadata to generate one or more FQMI s for measurements to be obtained from one or more aggregator datastores of the computing resource monitoring service. The front-end server may access the metric mapping registry and utilize the one or more FQMI s to determine 1004 the location for each of the requested measurements.

[0102] Once the front-end server has obtained the location for each of the requested measurements, the front-end server may determine 1006 whether any of the requested measurements are no longer stored within the one or more aggregator datastores of the computing resource monitoring service due to expiration of the measurements. If it is determined that one or more measurements have expired, the front-end server may indicate 1008, to the requesting entity, that these one or more measurements are no longer available from the one or more aggregator datastores. If these expired measurements have been moved to an alternative datastore, such as an archival datastore, the front-end server may transmit a request to a computer system of the archival datastore to retrieve the expired measurements.

[0103] Once the front-end server has determined the location of any unexpired measurements within the one or more aggregator datastores, the front-end server may cause a

metric consolidation engine to access 1010 the one or more aggregator datastores specified within the metric mapping registry to obtain 1012 the one or more measurements from the aggregator datastores, the measurements being necessary to fulfill the GET request. The metric consolidation engine may aggregate 1014 and compile the measurements from the one or more aggregator datastores and de-serialize the measurements in order to fulfill the GET request. In an embodiment, if the measurements are redundantly stored within more than one datacenter of the computing resource monitoring service, the metric consolidation obtains multiple responses to the GET request.

[0104] If the compiled measurements for the datacenters are inconsistent, the metric consolidation engine may utilize one or more conflict resolution rules to determine which measurements are to be provided in response to the GET request. For instance, the metric consolidation engine may select the response with the highest sample count from the various datacenters. Once the metric consolidation engine has resolved any conflicts among the datacenters, if any, the metric consolidation engine may provide the measurements to the front-end server. This may enable the front-end server to fulfill the GET request by providing the measurements to the requesting entity.

[0105] As noted above, the customer computer system or other computing resource providing measurements to the computing resource monitoring service may internally perform the partitioning of measurements prior to aggregation and storage through use of one or more aggregator sub-systems of the computing resource monitoring service. This architecture may obviate the need for the computing resource monitoring service to maintain the one or more partitioner sub-systems or load balancers for the partitioner sub-systems described above in connection with FIGS. 2 and 3, as the computing resource monitoring service may obtain the serialized measurements directly from the logical partitions of the customer computer system or other computing resource. Accordingly, FIG. 11 shows an illustrative example of a process 1100 for partitioning measurements to be transmitted to a computing resource monitoring service for publishing of the measurements in accordance with at least one embodiment. The process 1100 may be performed by a customer computer system or other computing resource, which may be configured to internally obtain and partition measurements for delivery to aggregator sub-systems of the computing resource monitoring service for in-memory storage.

[0106] The customer computer system or other computing resource may include one or more monitoring agents configured to monitor the health and other metrics for the customer

computer system or computing resource. These monitoring agents may record various measurements for various metrics over time and provide these measurements to a front-end module of the customer computer system or computing resource. Thus, the customer computer system may obtain 1102 a measurement from a monitoring agent for storage within an in-memory datastore. Additionally, the customer computer system may obtain metadata associated with the obtained measurement. Similar to the process 800 described above, the customer computer system server may utilize a hash function and the provided metadata to determine 1104 the FQMI for the measurement. This FQMI may include a multiple-byte hash, wherein a portion of the hash may be used to determine a logical partition for each measurement metric based at least in part on a mapping within a metric mapping registry of the computing resource monitoring service, while another portion of the hash may be used for unique identification of the measurement once placed within the corresponding logical partition. The customer computer system may further transform the measurement into a binary serialization format.

[0107] The customer computer system may access 1106 the metric mapping registry of the computing resource monitoring service to obtain a mapping of logical partitions to aggregator sub-systems within the computing resource monitoring service. The customer computer system may utilize the mapping and the FQMI and timestamp of the measurement to identify active aggregator sub-systems for one or more retention periods. This may enable the customer computer system to associate each logical partition to a corresponding aggregator sub-system within the computing resource monitoring service.

[0108] In an embodiment, the customer computer system utilizes the timestamp and the mapping to determine 1108 whether the measurement is for a time period beyond the latest retention period, as specified in the mapping. If the measurement includes a future timestamp (e.g., timestamp is for a time beyond the latest retention period), the customer computer system may queue 1110 the measurement in a separate unpartitioned queue, where the measurement may remain until the timestamp for the measurement is covered by the present retention period at a later time. As the retention period is updated, the customer computer system may obtain the measurement from this queue and determine whether the timestamp for the measurement is covered by the present retention period. If not, the measurement may remain in the queue.

[0109] If the timestamp for the measurement is for a current retention period, the customer computer system may place the measurement within a logical partition of a plurality of

logical partitions of the partitioner sub-system, as determined based at least in part on the timestamp, FQMI, and the mapping. Each logical partition may be associated with a particular queue, as described above. Thus, when the measurement is placed in a logical partition, the customer computer system may queue 1112 the measurement within an

5 associated file-based queue bound to a corresponding aggregator sub-system. The measurements within this queue may be partitioned based at least in part on the timestamp for each measurement such that delivery of the most recent measurements may be completed first. Thus, based on this partitioning of the measurements, the customer computer system may transmit 1114 one or more PUT API calls including the measurement to the
10 corresponding aggregator sub-system within the computing resource monitoring service.

[0110] In an embodiment, the PUT API calls to the aggregator sub-system are transmitted through use of a communications protocol which may indicate whether delivery of the PUT API call was completed successfully or not. Thus, through use of this protocol, the customer computer system may determine 1106 whether delivery of the PUT API calls and, hence, the
15 measurements was successful. If delivery of the measurements is unsuccessful, the customer computer system may refresh the communications channel between the customer computer system and the aggregator sub-system and transmit 1114 the one or more PUT API calls including the measurement to the corresponding aggregator sub-system within the computing resource monitoring service. Otherwise, if delivery was successful, the customer computer
20 system may receive 1116 acknowledgement of successful delivery and aggregation of the measurements within the aggregator sub-system.

[0111] As noted above, once the computing resource monitoring service has hashed the metadata to generate the FQMI for the measurement to be stored, the customer computer system may no longer required to repeat the metadata for future submissions of the
25 measurements for the particular metric, so long as the metadata has not expired or the customer computer system has not initiated a new session with the computing resource monitoring service. For instance, when customer computer system supplies additional measurements to the computing resource monitoring service, the customer computer system may provide the FQMI for the measurements corresponding to the same metric instead of the
30 metadata. The computing resource monitoring service may serialize the measurements and the FQMI in anticipation of storage of the various measurements within the in-memory datastores of the computing resource monitoring service. Accordingly, FIG. 12 shows an illustrative example of a process 1200 for storing measurements in one or more aggregator datastores based at least in part on a metadata hash in accordance with at least one

embodiment. The process 1200 may be performed by the aforementioned computing resource monitoring service, which may obtain requests to store measurements that may include either the metadata for the measurements or the FQMI for each measurement.

[0112] At any point, the computing resource monitoring service may receive 1202 a PUT API call to publish a measurement within an aggregator datastore. This PUT API call may include the measurement to be stored, as well as metadata for the measurement or an FQMI for the measurement. Thus, through the PUT API call, the computing resource monitoring service may obtain 1204 the measurement that is to be stored within an aggregator datastore. The computing resource monitoring service may further determine 1206 whether the PUT API call includes metadata for the measurement to be stored. For instance, the metadata may include one or more metric attributes that may uniquely identify the associated measurement and metric within the computing resource monitoring service. For example, the metadata may specify, individually or in combination, the account number of a customer associated with the customer computer system submitting the PUT request or with another computing resource responsible for generation of the measurement, the name of the particular service responsible for managing the computing resource (e.g., virtual computer system service, database service, etc.), the dimensions for the measurement (e.g., key/value pairs), and an identifier for the computing resource producing the measurement.

[0113] If the PUT API call includes the metadata for the measurement, the computing resource monitoring service may utilize a hash function and the provided metadata to generate 1208 a metadata hash, which may include FQMI for the measurement. This FQMI may include a multiple-byte hash, wherein a portion of the hash may be used to determine a logical partition for each measurement metric based at least in part on a mapping within a metric mapping registry of the computing resource monitoring service, while another portion of the hash may be used for unique identification of the measurement once placed within the corresponding logical partition. If the PUT API call does not include the metadata for the measurement, the computing resource monitoring service may determine 1210 whether the PUT API call includes the metadata hash (e.g., FQMI) for the measurement. If the PUT API call does not include either the metadata or the FQMI for the measurement, the computing resource monitoring service may deny 1212 the request to store the measurement.

Alternatively, the computing resource monitoring service may still process the measurement but the measurement may not be obtainable by the customer computer system or other computing resource, as the measurement may not be properly indexed due to the missing metadata and FQMI.

[0114] Once the computing resource monitoring service has obtained the FQMI for the measurement, the computing resource monitoring service may utilize 1214 the FQMI to identify an in-memory storage destination for the measurement. For instance, the computing resource monitoring service may access the metric mapping registry to obtain a mapping of logical partitions to various aggregator sub-systems to determine which logical partition will be used for placement of the measurement and for delivery to an aggregator sub-system. The computing resource monitoring service may utilize the FQMI and timestamp from the serialized measurement to identify, based at least in part on the mapping, which logical partition to place the measurement in. This may enable the computing resource monitoring service to store 1216 the measurement in the appropriate storage destination, which may include an aggregator sub-system associated with the logical partition in which the measurement is placed in based at least in part on the FQMI and the timestamp of the measurement.

[0115] As noted above, the computing resource monitoring service may obtain various GET requests from one or more entities to retrieve measurements from the in-memory datastores. These GET requests may include metadata for the measurements to be obtained and timestamps for these measurements. This may enable the computing resource monitoring service to determine the storage locations for the measurements and provide these measurements to the requesting entity in response to a GET request. Accordingly, FIG. 13 shows an illustrative example of a process 1300 for retrieving measurements from one or more aggregator datastores based at least in part on metadata included within a request to obtain the measurements in accordance with at least one embodiment. The process 1300 may be performed by the aforementioned computing resource monitoring service.

[0116] At any point, the computing resource monitoring service may receive 1302 a GET API call to obtain one or more measurements from the computing resource monitoring service. Once the computing resource monitoring service receives the GET API call, the computing resource monitoring service may determine 1304 whether the GET API call includes metadata for the measurements to be retrieved. If the GET API call does not include metadata for the measurements to be retrieved, the computing resource monitoring service may deny 1306 the request. However, if the GET API call does include metadata for the measurements to be obtained, the computing resource monitoring service may utilize a hash function and the metadata to generate 1308 one or more FQMI (e.g., metadata hashes) for measurements to be obtained from one or more aggregator datastores of the computing resource monitoring service.

[0117] Once the computing resource monitoring service generates the one or more FQMI for the measurements to be obtained, the computing resource monitoring service may access a metric mapping registry and utilize 1310 the generated one or more FQMIs to identify any in-memory storage locations (e.g., aggregator datastores) where the measurements may be stored. For instance, the computing resource monitoring service may obtain a mapping of measurements to aggregated datastores from the metrics mapping registry and utilize the FQMIs to determine 1312 the location of the aggregated datastores that may include the requested measurements. If the computing resource monitoring service is unable to identify any in-memory storage locations for the requested measurements (e.g., the measurements have expired, etc.), the computing resource monitoring service may deny 1306 the request.

[0118] If the computing resource monitoring service is able to identify one or more in-memory storage locations for the requested measurements, the computing resource monitoring service may retrieve 1314 the requested measurements from the identified locations and compile 1316 the measurements for delivery to the requesting entity. For instance, the computing resource monitoring service may obtain the measurements from the various aggregator sub-systems from each datacenter where the measurements may have been redundantly stored. If the obtained measurements from the one or more datacenters are not identical, the computing resource monitoring service may utilize one or more conflict resolution rules to determine the appropriate response to the GET API call. For instance, the computing resource monitoring service may select the response with the highest sample count from the various datacenters. The computing resource monitoring service may provide 1318 a response to the GET request in the form of the compiled measurements in a de-serialized format that may be used by the requesting entity for its own purposes.

[0119] FIG. 14 illustrates aspects of an example environment 1400 for implementing aspects in accordance with various embodiments. As will be appreciated, although a web-based environment is used for purposes of explanation, different environments may be used, as appropriate, to implement various embodiments. The environment includes an electronic client device 1402, which can include any appropriate device operable to send and/or receive requests, messages, or information over an appropriate network 1404 and, in some embodiments, convey information back to a user of the device. Examples of such client devices include personal computers, cell phones, handheld messaging devices, laptop computers, tablet computers, set-top boxes, personal data assistants, embedded computer systems, electronic book readers, and the like. The network can include any appropriate network, including an intranet, the Internet, a cellular network, a local area network, a

satellite network or any other such network and/or combination thereof. Components used for such a system can depend at least in part upon the type of network and/or environment selected. Protocols and components for communicating via such a network are well known and will not be discussed herein in detail. Communication over the network can be enabled
5 by wired or wireless connections and combinations thereof. In this example, the network includes the Internet, as the environment includes a web server 1406 for receiving requests and serving content in response thereto, although for other networks an alternative device serving a similar purpose could be used as would be apparent to one of ordinary skill in the art.

10 [0120] The illustrative environment includes at least one application server 1408 and a data store 1410. It should be understood that there can be several application servers, layers or other elements, processes or components, which may be chained or otherwise configured, which can interact to perform tasks such as obtaining data from an appropriate data store. Servers, as used herein, may be implemented in various ways, such as hardware devices or
15 virtual computer systems. In some contexts, servers may refer to a programming module being executed on a computer system. As used herein, unless otherwise stated or clear from context, the term “data store” refers to any device or combination of devices capable of storing, accessing and retrieving data, which may include any combination and number of data servers, databases, data storage devices and data storage media, in any standard,
20 distributed, virtual or clustered environment. The application server can include any appropriate hardware, software and firmware for integrating with the data store as needed to execute aspects of one or more applications for the client device, handling some or all of the data access and business logic for an application. The application server may provide access control services in cooperation with the data store and is able to generate content including,
25 but not limited to, text, graphics, audio, video and/or other content usable to be provided to the user, which may be served to the user by the web server in the form of HyperText Markup Language (“HTML”), Extensible Markup Language (“XML”), JavaScript, Cascading Style Sheets (“CSS”) or another appropriate client-side structured language. Content transferred to a client device may be processed by the client device to provide the content in one or more
30 forms including, but not limited to, forms that are perceptible to the user audibly, visually and/or through other senses including touch, taste, and/or smell. The handling of all requests and responses, as well as the delivery of content between the client device 1402 and the application server 1408, can be handled by the web server using PHP: Hypertext
Preprocessor (“PHP”), Python, Ruby, Perl, Java, HTML, XML or another appropriate server-

side structured language in this example. It should be understood that the web and application servers are not required and are merely example components, as structured code discussed herein can be executed on any appropriate device or host machine as discussed elsewhere herein. Further, operations described herein as being performed by a single device may, unless otherwise clear from context, be performed collectively by multiple devices, which may form a distributed and/or virtual system.

[0121] The data store 1410 can include several separate data tables, databases, data documents, dynamic data storage schemes and/or other data storage mechanisms and media for storing data relating to a particular aspect of the present disclosure. For example, the data store illustrated may include mechanisms for storing production data 1412 and user information 1416, which can be used to serve content for the production side. The data store also is shown to include a mechanism for storing log data 1414, which can be used for reporting, analysis or other such purposes. It should be understood that there can be many other aspects that may need to be stored in the data store, such as page image information and access rights information, which can be stored in any of the above listed mechanisms as appropriate or in additional mechanisms in the data store 1410. The data store 1410 is operable, through logic associated therewith, to receive instructions from the application server 1408 and obtain, update or otherwise process data in response thereto. The application server 1408 may provide static, dynamic, or a combination of static and dynamic data in response to the received instructions. Dynamic data, such as data used in web logs (blogs), shopping applications, news services and other such applications may be generated by server-side structured languages as described herein or may be provided by a content management system ("CMS") operating on, or under the control of, the application server. In one example, a user, through a device operated by the user, might submit a search request for a certain type of item. In this case, the data store might access the user information to verify the identity of the user and can access the catalog detail information to obtain information about items of that type. The information then can be returned to the user, such as in a results listing on a web page that the user is able to view via a browser on the user device 1402. Information for a particular item of interest can be viewed in a dedicated page or window of the browser. It should be noted, however, that embodiments of the present disclosure are not necessarily limited to the context of web pages, but may be more generally applicable to processing requests in general, where the requests are not necessarily requests for content.

[0122] Each server typically will include an operating system that provides executable program instructions for the general administration and operation of that server and typically

will include a computer-readable storage medium (e.g., a hard disk, random access memory, read only memory, etc.) storing instructions that, when executed by a processor of the server, allow the server to perform its intended functions. Suitable implementations for the operating system and general functionality of the servers are known or commercially available and are readily implemented by persons having ordinary skill in the art, particularly in light of the disclosure herein.

[0123] The environment, in one embodiment, is a distributed and/or virtual computing environment utilizing several computer systems and components that are interconnected via communication links, using one or more computer networks or direct connections. However, it will be appreciated by those of ordinary skill in the art that such a system could operate equally well in a system having fewer or a greater number of components than are illustrated in FIG. 14. Thus, the depiction of the system 1400 in FIG. 14 should be taken as being illustrative in nature and not limiting to the scope of the disclosure.

[0124] Additionally, embodiments of the present disclosure can be described in view of the following clauses:

1. A computer-implemented method, comprising:
under the control of one or more computer systems configured with executable instructions, receiving, from a customer computer system, one or more requests resulting in receipt of a plurality of measurements, the one or more requests including measurements and metadata specifying attributes of the measurements;
utilizing a hash function and the metadata to generate a fully qualified metric identifier;
partitioning the measurements based at least in part on the fully qualified metric identifier for the measurements into a plurality of logical partitions;
transmitting the measurements from the plurality of logical partitions to corresponding aggregator sub-systems comprising one or more in-memory datastores for aggregation and storage of the measurements with other measurements associated with the fully qualified metric identifier stored within the one or more in-memory datastores;
storing the measurements in the one or more in-memory datastores;
receiving a request to retrieve one or more measurements, the request including metadata associated with the one or more measurements;
utilizing the hash function and the metadata associated with the one or more measurements to generate the fully qualified metric identifier;

utilizing the fully qualified metric identifier to identify the one or more in-memory datastores wherein the one or more measurements and the other measurements are stored;
obtaining the one or more measurements and the other measurements from the identified one or more in-memory datastores; and

5 providing the one or more measurements and the other measurements in response to the request.

2. The computer-implemented method of clause 1, further comprising aggregating the measurements with the other measurements associated with the fully qualified identifier for storage in the one or more in-memory datastores.

10 3. The computer-implemented method of clauses 1 or 2, further comprising partitioning the measurements of the plurality of logical partitions based at least in part on a timestamp for the measurements to enable transmission of most recent measurements first.

4. The computer-implemented method of any of clauses 1-3, further comprising:
obtaining, from a metric mapping registry, a mapping of the plurality of logical partitions to a
15 plurality of aggregator sub-systems;
identifying, based at least in part on the mapping and a timestamp for each measurement of the measurements, the corresponding aggregator sub-systems; and
associating the plurality of logical partitions with the identified corresponding aggregator sub-systems to enable transmission of the measurements to the identified corresponding
20 aggregator sub-systems.

5. A system, comprising at least one computing device configured to implement one or more services, wherein the one or more services are configured to:
receive a request to store a measurement, the request including the measurement and metadata specifying attributes of the measurement;
25 generate, based at least in part on the metadata, a metric identifier for the measurement;
select, based at least in part on the metric identifier, a partition of a plurality of partitions at a plurality of datacenters for placement of the measurement;
transmit the measurement from the partition to a corresponding aggregator system comprising one or more in-memory datastores;
30 aggregate the measurement with one or more measurements previously stored within an in-memory datastore of the one or more in-memory datastores; and

store the measurement and the one or more measurements in the in-memory datastore of the one or more in-memory datastores.

6. The system of clause 5, wherein the one or more services are further configured to: receive a request to retrieve one or more measurements, the request including metadata

- 5 associated with the one or more measurements;
generate, based at least in part on the metadata associated with the one or more measurements, the metric identifier;
use the metric identifier to identify the one or more in-memory datastores within the plurality of datacenters wherein the one or more measurements are stored;
10 obtain the one or more measurements from the identified one or more in-memory datastores of the plurality of datacenters;
utilize one or more conflict resolution rules to determine the one or more measurements to be transmitted in response to the request if the one or more measurements are different from the plurality of datacenters; and
15 transmit the one or more measurements in response to the request.

7. The system of clauses 5 or 6, wherein:

at least one of the attributes of the measurement includes a timestamp for the measurement;
and
the selection of the partition is further based at least in part on the timestamp.

- 20 8. The system of any of clauses 5-7, wherein the one or more services are further configured to remove, from the in-memory datastore and based at least in part on expiration of a retention period, measurements having a timestamp exceeding the retention period.

9. The system of any of clauses 5-8, wherein the one or more services are further configured to serialize the measurement in a binary serialization format such that the
25 measurement in the binary serialization format includes the metric identifier, a timestamp for the measurement, a unit of measurement for the measurement, and the measurement itself.

10. The system of any of clause 5-9, wherein the one or more services are further configured to:

- obtain a mapping of the plurality of partitions to a plurality of aggregator systems;
30 identify, based at least in part on the mapping and a timestamp of the measurement, the corresponding aggregator system for the selected partition; and

associate the selected partition to the corresponding aggregator system to enable transmission of the measurement from the selected partition to the corresponding aggregator system.

11. The system of any of clauses 5-10, wherein the one or more services are further configured to:

- 5 receive a request to store a second measurement, the request including the second measurement and metadata specifying attributes of the second measurement;
generate, based at least in part on the metadata, a metric identifier for the second measurement;
determine, based at least in part on a timestamp for the second measurement and the metric
10 identifier, that no partitions are available for placement of the second measurement as a result of the timestamp for second measurement being outside a retention period; and
store the second measurement in a queue until a partition is available based at least in part on the timestamp and the metric identifier.

12. The system of any of clauses 5-11, wherein the one or more services are further
15 configured to, if the selected partition does not have capacity for placement of the measurement, purge one or more other measurements having timestamps earlier than a timestamp for the measurement.

13. A non-transitory computer-readable storage medium having stored thereon executable
instructions that, when executed by one or more processors of a computer system, cause the
20 computer system to at least:
receive, from a partition of a plurality of partitions, a plurality of measurements for a metric;
obtain, from an in-memory datastore, other measurements for the metric;
aggregate the received plurality of measurements with the other measurements; and
store the aggregated measurements in the in-memory datastore.

- 25 14. The non-transitory computer-readable storage medium of clause 13, wherein the instructions further cause the computer system to:
receive a request to retrieve one or more measurements from one or more in-memory
datastores, the request including a metric identifier for the one or more measurements;
use the metric identifier to identify the one or more in-memory datastores wherein the one or
30 more measurements are stored;
obtain the one or more measurements from the identified one or more in-memory datastores;
and

transmit the one or more measurements.

15. The non-transitory computer-readable storage medium of clauses 13 or 14, wherein the plurality of measurements are received in a binary serialization format, wherein the plurality of measurements in the binary serialization format include a metric identifier, a

5 timestamp for the measurement, a unit of measurement for the measurement, and the plurality of measurements themselves.

16. The non-transitory computer-readable storage medium of clause 15, wherein the instructions further cause the computer system to:

de-serialize the plurality of measurements and the other measurements from the in-memory
10 datastore;

aggregate the de-serialized plurality of measurements with the de-serialized measurements from the in-memory datastore; and

serialize the aggregated measurements prior to storing the aggregated measurements in the in-memory datastore.

15 17. The non-transitory computer-readable storage medium of any of clauses 13-16, wherein the instructions further cause the computer system to remove, based at least in part on expiration for a retention period, measurements from the in-memory datastore having a timestamp exceeding the retention period.

18. The non-transitory computer-readable storage medium of any of clauses 13-17,
20 wherein the partition is associated with the computer system based at least in part on a mapping of the plurality of partitions to a plurality of aggregator systems obtained from a registry.

19. The non-transitory computer-readable storage medium of any of clauses 13-18,
25 wherein the partition is provided with the plurality of measurements based at least in part on a partitioning of measurements based at least in part on a metric identifier and a timestamp for the plurality of measurements.

20. The non-transitory computer-readable storage medium of clause 19, wherein the metric identifier for the measurements is generated using a hash function and metadata for the measurement, the metadata specifying attributes of the measurement.

30 21. A computer-implemented method, comprising:
under the control of one or more computer systems configured with executable instructions,

receiving, from a customer computer system, a request to store measurements for a metric obtained within a time period for the metric, the request including the measurements and metadata specifying attributes of the measurements;

utilizing, for the measurements, a hash function and the metadata to generate a fully qualified metric identifier for the measurements;

determining, based at least in part on the fully qualified metric identifier and a timestamp for the measurements, one or more datastores for the measurements;

storing the measurements within the determined first one or more datastores;

receiving, from the customer computer system, a second request to store a second plurality of measurements specifying measurements obtained within the time period for the metric, the request including the second plurality of measurements and the fully qualified metric identifier but lacking at least a portion of the metadata;

utilizing the fully qualified metric identifier and timestamps for the second plurality of measurements to make a determination where to store the second plurality of measurements;

and

storing the second plurality of measurements according to the determination.

22. The computer-implemented method of clause 21, further comprising:

receiving, from the customer computer system, a third request to retrieve the measurements, the request including the fully qualified metric identifier; and

using the fully qualified metric identifier to obtain the measurements from the one or more datastores for the measurements.

23. The computer-implemented method of clauses 21 or 22, further comprising utilizing run length encoding to compress the measurements and the second plurality of measurements for one or more sequential measurements having a same value.

24. The computer-implemented method of any of clauses 21-23, further comprising aggregating the measurements with the second plurality of measurements if the determination is to store the second plurality of measurements in the same one or more datastores as the measurements.

25. A system, comprising at least one computing device configured to implement one or more services, wherein the one or more services are configured to:

receive a request to store a plurality of measurements for a metric, the request not having the same metadata that was received with an earlier request to store other measurements for the metric;

determine a metric identifier for the plurality of measurements; and

- 5 use the metric identifier to store the plurality of measurements such that information responsive to a request to retrieve measurements for the metric that specifies the metadata includes at least both the plurality of measurements and the other measurements.

26. The system of clause 25, wherein the request to store the plurality of measurements includes the metric identifier.

- 10 27. The system of clauses 25 or 26, wherein the one or more services are further configured to aggregate the plurality of measurements with the other measurements based at least in part on the determined metric identifier.

28. The system of any of clauses 25-27, wherein use of the determined fully qualified metric identifier to store the metrics data includes:

- 15 utilizing the metric identifier and a timestamp of the measurements of the plurality of measurements to partition the plurality of measurements into one or more partitions; transmitting the plurality of measurements from the one or more partitions to one or more aggregator sub-system of a plurality of aggregator sub-systems for storage of the plurality of measurements; and
- 20 storing the plurality of measurements within one or more aggregator datastores of the one or more aggregator sub-systems.

29. The system of clause 28, wherein the one or more services are further configured to: obtain, from a metric mapping registry and based at least in part on timestamps for the measurements of the plurality of measurements, one or more mappings between the one or
- 25 more partitions and the plurality of aggregator sub-systems; and select the one or more aggregator sub-systems from the plurality of aggregator sub-systems based at least in part on the one or more mappings.

30. The system of any of clauses 25-29, wherein the one or more services are further configured to:

- 30 receive a second request to retrieve the plurality of measurements, the request specifying the metadata; determine the metric identifier based at least in part on the metadata; and

retrieve, through use of the metric identifier and one or more time periods specified in the metadata, at least both the plurality of measurements and the other measurements.

31. The system of any of clauses 25-30, wherein the plurality of measurements are stored in a binary serialization format, the binary serialization format including the metric identifier, a timestamp for the metrics data, a unit of measurement for the metrics data, and a measurement for a metric represented by the metrics data.

32. The system of any of clauses 25-31, wherein the metric identifier is generated through use of a hash function and one or more metric attributes specified within the metadata.

33. A non-transitory computer-readable storage medium having stored thereon executable instructions that, when executed by one or more processors of a computer system, cause the computer system to at least:

receive a first application programming interface call to store one or more measurements for a metric, the first application programming interface call specifying the one or more measurements and a metric identifier for the one or more measurements, the metric identifier generated based at least in part on metadata included in an earlier request to store other measurements for the metric;

utilize the metric identifier to identify a datastore for storing the one or more measurements; and

store the one or more measurements in the identified datastore such that the one or more measurements and the other measurements can be obtained in response to a request that includes the metadata or the fully qualified metric identifier.

34. The non-transitory computer-readable storage medium of clause 33, wherein the instructions further cause the computer system to:

receive a second application programming interface call to retrieve a plurality of

measurements for the metric, the second application programming interface call including the metadata;

generate, using the included metadata, the metric identifier;

identify, based at least in part on the generated metric identifier and one or more time periods specified in the metadata, one or more datastores that include the one or more measurements and the other measurements; and

provide the one or more measurements and the other measurements from the identified one or more datastores.

35. The non-transitory computer-readable storage medium of clauses 33 or 34, wherein the instructions further cause the computer system to transform the one or more measurements into a binary serialization format, the one or more measurements in the binary serialization format including the metric identifier and the measurement.

5 36. The non-transitory computer-readable storage medium of clause 35, wherein storing the one or more measurements in the identified datastore includes:
retrieving, from the datastore, serialized measurements for the metric, the serialized
measurements including the other measurements in the binary serialization format;
de-serializing the serialized measurements and the one or more measurements in the binary
10 serialization format;
aggregating the one or more measurements with the other measurements; and
transforming the one or more measurements and the other measurements into the binary
serialization format.

37. The non-transitory computer-readable storage medium of clauses 35 or 36, wherein
15 the instructions further cause the computer system to utilize run length encoding to compress the one or more measurements for one or more sequential measurements having a same value in the binary serialization format.

38. The non-transitory computer-readable storage medium of any of clauses 33-37,
wherein the instructions further cause the computer system to:
20 generate, based at least in part on the metric identifier and a timestamp of the one or more measurements, a partition;
place the one or more measurements in the logical partition into a queue for distribution to an aggregator sub-system, the aggregator sub-system including the identified datastore; and
transmit, from the queue, the one or more measurements to the aggregator sub-system for
25 storage of the one or more measurements in the identified datastore.

39. The non-transitory computer-readable storage medium of clause 38, wherein the instructions further cause the computer system to:
obtain, from a metric mapping registry, a mapping between partitions and aggregator sub-systems over time; and
30 utilize the metric identifier and the timestamp of the metrics data to determine, from the mapping, the aggregator sub-system for storage of the one or more measurements into the identified datastore.

40. The non-transitory computer-readable storage medium of any of clauses 33-40, wherein the metric identifier is generated through use of a hash function to hash one or more metric attributes specified in the metadata included with the request to store the other measurements.

5 41. A computer-implemented method, comprising:

processing incoming metrics at a web server computer system by at least, at the web server computer system:

obtaining, from a computing resource associated with a customer computer system, a

plurality of measurements for a metric and including, for each measurement of the plurality

10 of measurements, metadata specifying attributes of the measurement;

utilizing a hash function and the metadata to generate a fully qualified metric identifier for each measurement of the plurality of measurements;

obtaining, from a metric mapping registry of a computing resource monitoring service and

based at least in part on a timestamp for each measurement of the plurality of measurements,

15 a mapping of logical partitions to aggregator sub-systems of the computing resource

monitoring service, the one or more aggregator sub-systems for storing the plurality of measurements;

partitioning the plurality of measurements, based at least in part on the fully qualified metric identifier for each measurement and the timestamp for each measurement of the plurality of

20 measurements, into a plurality of logical partitions;

sub-partitioning the plurality of measurements in the plurality of logical partitions based at least in part on the timestamp for each measurement of the plurality of measurements;

prioritizing the plurality of measurements in the plurality of logical partitions based at least in part on the timestamp for each measurement of the plurality of measurements; and

25 transmitting one or more application programming interface calls to the one or more

aggregator sub-systems of the computing resource monitoring service to store the plurality of measurements of the plurality of logical partitions within one or more datastores of the one or

more aggregator sub-systems, the one or more application programming interface calls including the plurality of measurements.

30 42. The computer-implemented method of clause 41, further comprising:

placing the plurality of measurements from the plurality of logical partitions into one or more

file-based queues, the one or more file-based queues associated with the one or more

aggregator sub-systems; and

delivering, from the file-based queues, the plurality of measurements to the one or more aggregator sub-systems.

43. The computer-implemented method of clauses 41 or 42, further comprising selecting the at least one measurement from the plurality of measurements based at least in part on the at least one measurement including a latest timestamp.

44. The computer-implemented method of any of clauses 41-43, further comprising: transmitting, from the web server computer system, a request to retrieve the plurality of measurements from the one or more datastores; determining, at a metric consolidation engine and based at least in part on the mapping of the plurality of logical partitions, the one or more aggregator sub-systems including the one or more datastores; obtaining, from the one or more datastores, the plurality of measurements; and providing the plurality of measurements to the web server computer system.

45. A system, comprising: one or more processors; and memory including instructions that, when executed by the one or more processors, cause the system to: receive a plurality of measurements for a metric, measurements of the plurality of measurements including metadata for the measurements; determine, for the measurements and based at least in part on the metadata, a metric identifier; partition the measurements into a plurality of partitions; and transmit, from the system, a message to one or more aggregator systems of a computing resource monitoring service to store the measurements.

46. The system of clause 45, wherein the instructions further cause the system to: transmit a request to retrieve one or more measurements from the one or more datastores, the request including the metric identifier and specifying a time period associated with the one or more measurements; the metric identifier and the specified time period used to identify one or more datastores wherein the one or more measurements are stored; and obtain from the computing resource monitoring service the one or more measurements.

47. The system of clauses 45 or 46, wherein the instructions further cause the system to serialize the measurements in a binary serialization format such that the measurements in the

binary serialization format include the metric identifier, a timestamp for the measurements, a unit of measurement for the measurements, and the measurements themselves.

48. The system of any of clauses 45-47, wherein the instructions further cause the system to:

- 5 obtain, based at least in part on a timestamp for the measurements, a mapping of partitions to a plurality of aggregator sub-systems;
identify, based at least in part on the mapping, the one or more aggregator sub-systems for the plurality of partitions; and
associate the plurality of partitions to the one or more aggregator sub-system to enable
10 selection of the measurements from the plurality of logical partitions for creation of the message.

49. The system of any of clauses 45-48, wherein the instructions further cause the system to:

- 15 determine if any partition of the plurality of partitions lacks capacity for at least one measurement of the plurality of measurements; and
purge one or more measurements having timestamps earlier than a timestamp for the at least one measurement to create capacity for placement of the at least one measurement.

50. The system of any of clauses 45-49, wherein the instructions further cause the system to:

- 20 determine, based at least in part on the metric identifier and a timestamp being outside a retention period, that no logical partitions are available for placement of a measurement; and
place the measurement in a queue until a logical partition is available for placement of the measurement.

51. The system of any of clauses 45-50, wherein the message is transmitted to the one or
25 more aggregator sub-systems using a communications protocol configured to provide acknowledgment of successful delivery of the measurements.

52. The system of any of clauses 45-51, wherein the instructions further cause the system to:

- 30 place the measurements from the plurality of partitions into one or more queues, the one or more queues associated with the one or more aggregator sub-systems; and
deliver, from the queues, the measurements to the one or more aggregator sub-systems.

53. A non-transitory computer-readable storage medium having stored thereon executable instructions that, when executed by one or more processors of a computer system, cause the computer system to at least:

obtain a measurement for a metric, the measurement including metadata specifying attributes
5 of the measurement;

determine, based at least in part on the metadata, a metric identifier for the measurement;

partition a plurality of measurements that includes the measurement into a plurality of
partitions; and

distribute the plurality of measurements from the plurality of partitions among a plurality of

10 aggregator systems of a computing resource monitoring service.

54. The non-transitory computer-readable storage medium of clause 53, wherein the instructions further cause the computer system to:

obtain a mapping of the plurality of partitions to the plurality of aggregator systems of the
computing resource monitoring service; and

15 associate the partitions to corresponding aggregator systems to determine a distribution of the plurality of measurements from the partitions to the corresponding aggregator systems.

55. The non-transitory computer-readable storage medium of clauses 53 or 54, wherein the instructions further cause the computer system to:

place the plurality of measurements into a plurality of queues, the one or more queues

20 associated with the plurality of aggregator systems; and

distribute, from the plurality of queues, the plurality of measurements among the plurality of
aggregator systems.

56. The non-transitory computer-readable storage medium of any of clauses 53-55,

wherein the instructions further cause the computer system to distribute the plurality of

25 measurements among the plurality of aggregator systems using a communications protocol configured to provide acknowledgment of successful distribution of the plurality of measurements upon successful storage of the plurality of measurements.

57. The non-transitory computer-readable storage medium of any of clauses 53-56, wherein the instructions further cause the computer system to:

30 transmit a message to the plurality of aggregator systems to retrieve one or more

measurements, the request including the metric identifier and specifying a time period for the one or more measurements to enable identification of the plurality of aggregator systems; and

obtain from the plurality of aggregator systems the one or more measurements.

58. The non-transitory computer-readable storage medium of any of clauses 53-57, wherein the instructions further cause the computer system to:

obtain a second measurement for the metric, the second measurement including metadata

5 specifying attributes of the second measurement;

determine, based at least in part on the metadata of the second measurement, a metric identifier for the second measurement;

determine, based at least in part on the metric identifier for the second measurement and a timestamp for the second measurement being later than a latest retention period, that no

10 logical partitions are available for the second measurement; and

place the second measurement in a queue until a partition is available for the second measurement.

59. The non-transitory computer-readable storage medium of any of clauses 53-58,

wherein the instructions further cause the computer system to transform the measurement into

15 a binary serialization format such that the measurement in the binary serialization format further includes at least the metric identifier.

60. The non-transitory computer-readable storage medium of any of clauses 53-59,

wherein the instructions further cause the computer system to, for any partition that lacks

capacity for at least one measurement, remove one or more measurements having timestamps

20 earlier than a timestamp for the at least one measurement from the partition.

61. A computer-implemented method, comprising:

under the control of one or more computer systems configured with executable instructions,

receiving, from one or more load balancers for one or more virtual machine instances of an

auto-scale group, a plurality of measurements for a metric including metadata specifying

25 attributes of measurements of the plurality of measurements;

utilizing a hash function and the metadata to generate a fully qualified metric identifier for the measurements;

partitioning the measurements based at least in part on the fully qualified metric identifier and a timestamp for the measurements, resulting in a plurality of logical partitions;

30 transmitting the plurality of measurements from the plurality of logical partitions to one or more aggregator sub-systems for storage within one or more datastores of the one or more aggregator sub-systems;

obtaining, at one or more instance managers for the one or more virtual machine instances, the measurements from the one or more datastores; and
performing automatic scaling of the auto-scale group based at least in part on the obtained measurements by at least changing a number of virtual machine instances in the auto-scale
5 group.

62. The computer-implemented method of clause 61, further comprising:
obtaining, from a metric mapping registry, a mapping of the plurality of logical partitions to a plurality of aggregator sub-systems;
identifying, for each logical partition of the plurality of logical partitions and based at least in
10 part on the mapping, the corresponding aggregator sub-system; and
associating each logical partition with the identified corresponding aggregator sub-system to enable transmission of the measurements of the logical partition to the identified corresponding aggregator sub-system.

63. The computer-implemented method of clauses 61 or 62, further comprising:
15 receiving, from the one or more instance managers, a request to retrieve one or more measurements, the request including the fully qualified metric identifier;
utilizing the fully qualified metric identifier to identify the one or more datastores wherein the one or more measurements are stored; and
providing the one or more measurements to the one or more instance managers in response to
20 the request.

64. The computer-implemented method of any of clauses 61-63, further comprising aggregating, within one or more aggregator sub-systems, the measurements with existing measurements within the one or more in-memory datastores to enable storage of the plurality of measurements.

25 65. A system, comprising at least one computing device configured to implement one or more services, wherein the one or more services are configured to:
receive a plurality of measurements for a metric associated with a group of computing resources, measurements of the plurality of measurements associated with metadata for the measurements;
30 determine, for the measurements and based at least in part on the metadata, a metric identifier;
partition the measurements into a plurality of partitions;

transmit the measurements from the plurality of partitions to one or more in-memory datastores for storage; and
provide the measurements to one or more computing resource managers for the group to enable modification of one or more computing resources of the group.

- 5 66. The system of clause 65, wherein the one or more services are configured to:
receive a request to retrieve one or more measurements, the request including the metric
identifier and specifying a time period;
use the metric identifier and the specified time period to identify the one or more in-memory
datastores wherein the one or more measurements are stored;
10 obtain the one or more measurements; and
transmit the one or more measurements to the one or more computing resource managers to
enable the modification of the one or more computing resources of the group.

67. The system of clauses 65 or 66, wherein the one or more services are further
configured to aggregate the measurements with existing measurements of the one or more in-
15 memory datastores for storage of the measurements in the one or more in-memory datastores.

68. The system of any of clauses 65-67, wherein the one or more services are further
configured to serialize the measurements in a binary serialization format such that the
measurements in the binary serialization format includes the metric identifier, a timestamp
for the measurements, a unit of measurement for the measurements, and the measurements
20 themselves.

69. The system of any of clauses 65-68, wherein the one or more services are further
configured to, if one or more partitions of the plurality of partitions do not have capacity for
the measurements, purge one or more measurements having timestamps earlier than an
earliest timestamp for the measurements.

- 25 70. The system of any of clauses 65-69, wherein the one or more services are configured
to:
obtain a mapping of the plurality of partitions to a plurality of aggregator systems, the
plurality of aggregator systems including the one or more in-memory datastores;
identify, based at least in part on the mapping, corresponding aggregator systems for
30 partitions of the plurality of partitions; and
associate the partitions to the corresponding aggregator systems to enable transmission of the
measurements from the partitions to the one or more in-memory datastores.

71. The system of any of clauses 65-70, wherein the one or more services are configured to remove, from the one or more in-memory datastores, expired measurements as a result of an old retention period for the expired measurements being replaced with a newer retention period.

5 72. The system of any of clauses 65-71, wherein the one or more computing resource managers are configured to provision additional computing resources for the auto-scale group based at least in part on the provided measurements.

73. A non-transitory computer-readable storage medium having stored thereon executable instructions that, when executed by one or more processors of a computer system, cause the
10 computer system to at least:
receive a measurement for a metric associated with a group of computing resources, the measurement associated with metadata for the measurement;
determine, based at least in part on the metadata, a metric identifier;
select, based at least in part on the metric identifier for the measurement, a partition from a
15 plurality of partitions;
transmit the measurement from the selected partition to an in-memory datastore for storage;
and
provide the measurement to one or more computing resource managers for the group of
computing resources to enable one or more operations to be performed on the group of
20 computing resources.

74. The non-transitory computer-readable storage medium of clause 73, wherein the one or more computing resource managers are configured to terminate one or more computing resources from the auto-scale group based at least in part on the provided measurement.

75. The non-transitory computer-readable storage medium of clauses 73 or 74, wherein
25 the instructions further cause the computer system to:
determine, based at least in part on a mapping of partitions to aggregator systems, a corresponding aggregator system for the selected partition, the corresponding aggregator system including the in-memory datastore; and
associate the partition to the corresponding aggregator system to enable transmission of the
30 measurement from the selected partition to the in-memory datastore.

76. The non-transitory computer-readable storage medium of any of clauses 73-75, wherein the instructions further cause the computer system to:

determine if the selected partition does not have capacity for the measurement; and
purge one or more measurements having timestamps earlier than a timestamp for the
measurement.

77. The non-transitory computer-readable storage medium of any of clauses 73-76,
5 wherein the instructions further cause the computer system to:
determine that an old retention period has been replaced with a newer retention period for
measurements stored within the in-memory datastore; and
remove, from in-memory datastores, measurements having timestamps within the old
retention period.

10 78. The non-transitory computer-readable storage medium of any of clauses 73-78,
wherein the instructions further cause the computer system to:
receive a message to retrieve one or more measurements for a metric, the request including
the metric identifier associated with the one or more measurements and specifying a time
range;
15 use the metric identifier and the specified time range to identify one or more in-memory
datastores;
obtain the one or more measurements from the identified one or more in-memory datastores;
and
transmit the one or more measurements to the one or more computing resource managers to
20 enable performance of the one or more operations on the group.

79. The non-transitory computer-readable storage medium of any of clauses 73-78,
wherein the instructions further cause the computer system to serialize the measurement
using a binary serialization format such that the measurement in the binary serialization
format include the metric identifier, a timestamp for the measurement, a unit of measurement
25 for the measurement, and the measurement itself.

80. The non-transitory computer-readable storage medium of clause 79, wherein the
instructions further cause the computer system to:
obtain, from the in-memory datastore, a plurality of measurements, the plurality of
measurements being in the binary serialization format;
30 de-serialize the plurality of measurements and the received measurement for the metric; and
aggregate the plurality of measurements with the received measurement for the metric for
storage within the in-memory datastore.

81. A computer-implemented method, comprising:

under control of one or more computer systems of a computing resource monitoring service provided by a computing resource service provider,

receiving, from a customer computer system of a customer of the computing resource service

5 provider, a first application programming interface call to retrieve a first set of measurements for a metric, the first application programming interface call including a fully qualified metric identifier, and specifying a time range for the first set of measurements and a parameter indicating that the first set of measurements are to be authoritative;

determining, based at least in part on the fully qualified metric identifier and the specified

10 time range for the first set of measurements, one or more datastores where the first set of measurements are stored;

obtaining the first set of measurements from the determined one or more datastores where the first set of measurements are stored;

determining if the first set of measurements obtained from the determined one or more

15 datastores are authoritative;

as a result of the first set of measurements being authoritative, providing the first set of measurements in response to the first application programming interface call;

receiving, from the customer computer system, a second application programming interface call to retrieve a second set of measurements for the metric, the second application

20 programming interface call including the fully qualified metric identifier, and specifying a time range for the second set of measurements and a parameter indicating that the second set of measurements are to be obtained regardless of whether a time period for storage of measurements for the metric has elapsed;

determining, based at least in part on the fully qualified metric identifier and the specified

25 time range for the second set of measurements, one or more datastores where the second set of measurements are stored;

obtaining the second set of measurements from the determined one or more datastores where the second set of measurements are stored; and

30 providing the second set of measurements in response to the second application programming interface call.

82. The computer-implemented method of clause 81, further comprising:

receiving, from the customer computer system, a third application programming interface call to store a measurement for a second metric obtained within a time period for the second

metric, the third application programming interface call specifying a parameter indicating that the measurement to be stored is authoritative for the time period;

determining a fully qualified metric identifier for the measurement for the second metric;

determining, based at least in part on the fully qualified metric identifier for the measurement

5 for the second metric, a datastore for the measurement for the second metric;

storing the measurement for the second metric within the determined datastore; and

as a result of the third application programming interface call specifying the parameter

indicating that the measurements to be stored is authoritative for the time period, making the measurement for the second metric available for retrieval and indicating that the

10 measurement for the second metric is authoritative for the time period.

83. The computer-implemented method of clause 82, further comprising:

selecting, based at least in part on the fully qualified metric identifier for the measurement for the second metric, a logical partition of a plurality of logical partitions for placement of the measurement;

15 obtaining, from a metric mapping registry and based at least in part on a timestamp for the measurement for the second metric, a mapping of the plurality of logical partitions to a plurality of in-memory datastores; and

transmitting the measurement for the second metric from the logical partition to a

corresponding aggregator sub-system comprising the in-memory datastore based at least in

20 part on the mapping.

84. The computer-implemented method of clauses 82 or 83, further comprising updating a metric mapping registry to indicate that, as a result of the measurement being stored within the determined one or more datastores, the measurement is authoritative for the time period.

85. A system, comprising at least one computing device configured to implement one or
25 more services, wherein the one or more services are configured to:

receive a request to retrieve measurement data for a metric for an interval of time;

determine, based at least in part on information in the request, one or more datastores where the measurement data is stored;

generate a response to the request such that:

30 if the request indicates that authoritative data is to be provided in the generated response, the generated response limits the measurement data to authoritative measurement data; and

if the request does not indicate that authoritative data is to be provided in the generated response, the generated response is not limited to authoritative measurement data.

86. The system of clause 85, wherein the information in the request includes information indicating that the measurement data to be retrieved is authoritative.

87. The system of clauses 85-86, wherein the information in the request includes information indicating that measurement data to be retrieved is not limited to the authoritative measurement data.

88. The system of any of clauses 85-87, wherein:

the information in the request includes a metric identifier and time information for the measurement data; and

the one or more services are further configured to:

obtain, for at least the interval of time, a mapping indicating the one or more datastores where the measurement data is stored; and

utilize the metric identifier to identify the one or more datastores where the measurement data is stored.

89. The system of any of clauses 85-88, wherein the one or more services are further configured to:

receive a second request to store a measurement for a second metric;

determine, based at least in part on information in the second request, that measurement data for the second metric is authoritative for a time interval, the measurement data for the second metric including the measurement for the second metric; and

store the measurement for the second metric in a second in-memory datastore as authoritative for the time interval.

90. The system of clause 89, wherein:

the second request to store the measurement for the second metric includes the measurement for the second metric and metadata specifying attributes of the measurement for the second metric; and

the one or more services are further configured to:

generate, based at least in part on the metadata, a metric identifier for the measurement for the second metric;

select, based at least in part on the metric identifier for the measurement for the second

metric, a partition of a plurality of logical partitions for placement of the measurement for the second metric; and

transmit the measurement from the partition to a corresponding aggregator system comprising the second in-memory datastore.

91. The system of clause 90, wherein the one or more services are further configured to: remove, from the second in-memory datastore, expired measurements for the second metric;

5 and

aggregate the measurement for the second metric with unexpired measurements for the second metric for storage of the measurement for the second metric.

92. The system of clauses 90 or 91, wherein the one or more services are further configured to serialize the measurement for the second metric in a binary serialization format such that the measurement in the binary serialization format includes the metric identifier, a
10 timestamp for the measurement for the second metric, a unit of measurement for the measurement for the second metric, and the measurement for the second metric itself.

93. A non-transitory computer-readable storage medium having stored thereon executable instructions that, when executed by one or more processors of a computer system, cause the
15 computer system to at least:

receive a request to retrieve measurement data for a metric;

identify, based at least in part on information in the request, an in-memory datastore where the measurement data is stored;

provide measurement data that is authoritative if the information in the request indicates that
20 authoritative measurements are to be provided; and

provide the measurement data regardless of whether the measurement data is authoritative if the information does not indicate that authoritative measurements are to be provided .

94. The non-transitory computer-readable storage medium of clause 93, wherein the instructions further cause the computer system to:

25 receive a second request to store a measurement for a second metric, the second request specifying the measurement for the second metric and without information indicating that measurements for the second metric collected before a time interval has elapsed are authoritative for the time interval; and

store the measurement for the second metric in a second in-memory datastore as a result of
30 the request lacking the information indicating that the measurements for the second metric are authoritative for the time interval, as not authoritative for the time interval.

95. The non-transitory computer-readable storage medium of clauses 93 or 94, wherein the information in the request specifies a parameter indicating that the measurement data for the metric to be retrieved is to be the authoritative measurements.

96. The non-transitory computer-readable storage medium of any of clauses 93-96,

5 wherein the instructions further cause the computer system to:

obtain, based at least in part on the information in the request, a mapping indicating the in-memory datastore where the measurement data is stored; and

utilize the information in the request to identify the in-memory datastore where the measurement data is stored from the mapping.

10 97. The non-transitory computer-readable storage medium of clause 96, wherein the information in the request includes:

time information usable to obtain the mapping; and

a metric identifier usable to identify, from the mapping, the in-memory datastore where the measurement data is stored.

15 98. The non-transitory computer-readable storage medium of any of clauses 93-97, wherein the instructions further cause the computer system to:

receive a second request to store a measurement for a second metric, the second request including information indicating that measurement data for the second metric is authoritative for a time interval; and

20 store the measurement for the second metric in a second in-memory datastore as authoritative, such that they are returned in response to requests to obtain authoritative measurements for the second metric.

99. The non-transitory computer-readable storage medium of clause 98, wherein the instructions further cause the computer system to:

25 generate, based at least in part on the information included in the second request, a metric identifier for the measurement for the second metric;

utilize the metric identifier for the measurement for the second metric to identify a partition of a plurality of partitions for placement of the measurement for the second metric; and transmit the measurement from the partition to the second in-memory datastore.

30 100. The non-transitory computer-readable storage medium of clause 99, wherein: the information included in the second request specifies a timestamp for the measurement for the second metric; and

the identification of the partition of the plurality of partitions is further based at least in part on the timestamp for the measurement for the second metric.

[0125] The various embodiments further can be implemented in a wide variety of operating environments, which in some cases can include one or more user computers, computing devices or processing devices which can be used to operate any of a number of applications. User or client devices can include any of a number of general purpose personal computers, such as desktop, laptop or tablet computers running a standard operating system, as well as cellular, wireless and handheld devices running mobile software and capable of supporting a number of networking and messaging protocols. Such a system also can include a number of workstations running any of a variety of commercially-available operating systems and other known applications for purposes such as development and database management. These devices also can include other electronic devices, such as dummy terminals, thin-clients, gaming systems and other devices capable of communicating via a network. These devices also can include virtual devices such as virtual machines, hypervisors and other virtual devices capable of communicating via a network.

[0126] Various embodiments of the present disclosure utilize at least one network that would be familiar to those skilled in the art for supporting communications using any of a variety of commercially-available protocols, such as Transmission Control Protocol/Internet Protocol ("TCP/IP"), User Datagram Protocol ("UDP"), protocols operating in various layers of the Open System Interconnection ("OSI") model, File Transfer Protocol ("FTP"), Universal Plug and Play ("UpnP"), Network File System ("NFS"), Common Internet File System ("CIFS") and AppleTalk. The network can be, for example, a local area network, a wide-area network, a virtual private network, the Internet, an intranet, an extranet, a public switched telephone network, an infrared network, a wireless network, a satellite network, and any combination thereof.

[0127] In embodiments utilizing a web server, the web server can run any of a variety of server or mid-tier applications, including Hypertext Transfer Protocol ("HTTP") servers, FTP servers, Common Gateway Interface ("CGI") servers, data servers, Java servers, Apache servers, and business application servers. The server(s) also may be capable of executing programs or scripts in response to requests from user devices, such as by executing one or more web applications that may be implemented as one or more scripts or programs written in any programming language, such as Java[®], C, C# or C++, or any scripting language, such as Ruby, PHP, Perl, Python or TCL, as well as combinations thereof. The server(s) may also

include database servers, including without limitation those commercially available from Oracle[®], Microsoft[®], Sybase[®], and IBM[®] as well as open-source servers such as MySQL, Postgres, SQLite, MongoDB, and any other server capable of storing, retrieving, and accessing structured or unstructured data. Database servers may include table-based servers, document-based servers, unstructured servers, relational servers, non-relational servers or combinations of these and/or other database servers.

[0128] The environment can include a variety of data stores and other memory and storage media as discussed above. These can reside in a variety of locations, such as on a storage medium local to (and/or resident in) one or more of the computers or remote from any or all of the computers across the network. In a particular set of embodiments, the information may reside in a storage-area network (“SAN”) familiar to those skilled in the art. Similarly, any necessary files for performing the functions attributed to the computers, servers or other network devices may be stored locally and/or remotely, as appropriate. Where a system includes computerized devices, each such device can include hardware elements that may be electrically coupled via a bus, the elements including, for example, at least one central processing unit (“CPU” or “processor”), at least one input device (e.g., a mouse, keyboard, controller, touch screen or keypad) and at least one output device (e.g., a display device, printer or speaker). Such a system may also include one or more storage devices, such as disk drives, optical storage devices and solid-state storage devices such as random access memory (“RAM”) or read-only memory (“ROM”), as well as removable media devices, memory cards, flash cards, etc.

[0129] Such devices also can include a computer-readable storage media reader, a communications device (e.g., a modem, a network card (wireless or wired), an infrared communication device, etc.), and working memory as described above. The computer-readable storage media reader can be connected with, or configured to receive, a computer-readable storage medium, representing remote, local, fixed, and/or removable storage devices as well as storage media for temporarily and/or more permanently containing, storing, transmitting, and retrieving computer-readable information. The system and various devices also typically will include a number of software applications, modules, services or other elements located within at least one working memory device, including an operating system and application programs, such as a client application or web browser. It should be appreciated that alternate embodiments may have numerous variations from that described above. For example, customized hardware might also be used and/or particular elements might be implemented in hardware, software (including portable software, such as applets) or

both. Further, connection to other computing devices such as network input/output devices may be employed.

[0130] Storage media and computer readable media for containing code, or portions of code, can include any appropriate media known or used in the art, including storage media and communication media, such as, but not limited to, volatile and non-volatile, removable and non-removable media implemented in any method or technology for storage and/or transmission of information such as computer readable instructions, data structures, program modules or other data, including RAM, ROM, Electrically Erasable Programmable Read-Only Memory (“EEPROM”), flash memory or other memory technology, Compact Disc Read-Only Memory (“CD-ROM”), digital versatile disk (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices or any other medium which can be used to store the desired information and which can be accessed by the system device. Based on the disclosure and teachings provided herein, a person of ordinary skill in the art will appreciate other ways and/or methods to implement the various embodiments.

[0131] The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense. It will, however, be evident that various modifications and changes may be made thereunto without departing from the broader spirit and scope of the invention as set forth in the claims.

[0132] Other variations are within the spirit of the present disclosure. Thus, while the disclosed techniques are susceptible to various modifications and alternative constructions, certain illustrated embodiments thereof are shown in the drawings and have been described above in detail. It should be understood, however, that there is no intention to limit the invention to the specific form or forms disclosed, but on the contrary, the intention is to cover all modifications, alternative constructions and equivalents falling within the spirit and scope of the invention, as defined in the appended claims.

[0133] The use of the terms “a” and “an” and “the” and similar referents in the context of describing the disclosed embodiments (especially in the context of the following claims) are to be construed to cover both the singular and the plural, unless otherwise indicated herein or clearly contradicted by context. The terms “comprising,” “having,” “including” and “containing” are to be construed as open-ended terms (i.e., meaning “including, but not limited to,”) unless otherwise noted. The term “connected,” when unmodified and referring to physical connections, is to be construed as partly or wholly contained within, attached to or

joined together, even if there is something intervening. Recitation of ranges of values herein are merely intended to serve as a shorthand method of referring individually to each separate value falling within the range, unless otherwise indicated herein and each separate value is incorporated into the specification as if it were individually recited herein. The use of the term “set” (e.g., “a set of items”) or “subset” unless otherwise noted or contradicted by context, is to be construed as a nonempty collection comprising one or more members. Further, unless otherwise noted or contradicted by context, the term “subset” of a corresponding set does not necessarily denote a proper subset of the corresponding set, but the subset and the corresponding set may be equal.

[0134] Conjunctive language, such as phrases of the form “at least one of A, B, and C,” or “at least one of A, B and C,” unless specifically stated otherwise or otherwise clearly contradicted by context, is otherwise understood with the context as used in general to present that an item, term, etc., may be either A or B or C, or any nonempty subset of the set of A and B and C. For instance, in the illustrative example of a set having three members, the conjunctive phrases “at least one of A, B, and C” and “at least one of A, B and C” refer to any of the following sets: {A}, {B}, {C}, {A, B}, {A, C}, {B, C}, {A, B, C}. Thus, such conjunctive language is not generally intended to imply that certain embodiments require at least one of A, at least one of B and at least one of C each to be present.

[0135] Operations of processes described herein can be performed in any suitable order unless otherwise indicated herein or otherwise clearly contradicted by context. Processes described herein (or variations and/or combinations thereof) may be performed under the control of one or more computer systems configured with executable instructions and may be implemented as code (e.g., executable instructions, one or more computer programs or one or more applications) executing collectively on one or more processors, by hardware or combinations thereof. The code may be stored on a computer-readable storage medium, for example, in the form of a computer program comprising a plurality of instructions executable by one or more processors. The computer-readable storage medium may be non-transitory.

[0136] The use of any and all examples, or exemplary language (e.g., “such as”) provided herein, is intended merely to better illuminate embodiments of the invention and does not pose a limitation on the scope of the invention unless otherwise claimed. No language in the specification should be construed as indicating any non-claimed element as essential to the practice of the invention.

[0137] Embodiments of this disclosure are described herein, including the best mode

known to the inventors for carrying out the invention. Variations of those embodiments may become apparent to those of ordinary skill in the art upon reading the foregoing description. The inventors expect skilled artisans to employ such variations as appropriate and the inventors intend for embodiments of the present disclosure to be practiced otherwise than as specifically described herein. Accordingly, the scope of the present disclosure includes all
5 modifications and equivalents of the subject matter recited in the claims appended hereto as permitted by applicable law. Moreover, any combination of the above-described elements in all possible variations thereof is encompassed by the scope of the present disclosure unless otherwise indicated herein or otherwise clearly contradicted by context.

10 [0138] All references, including publications, patent applications, and patents, cited herein are hereby incorporated by reference to the same extent as if each reference were individually and specifically indicated to be incorporated by reference and were set forth in its entirety herein.

CLAIMS

WHAT IS CLAIMED IS:

1. A system, comprising at least one computing device configured to implement one or more services, wherein the one or more services are configured to:
 - receive a request to store a plurality of measurements for a metric, the request not having the same metadata that was received with an earlier request to store other
 - 5 measurements for the metric;
 - determine a metric identifier for the plurality of measurements; and
 - use the metric identifier to store the plurality of measurements such that information responsive to a request to retrieve measurements for the metric that specifies the metadata includes at least both the plurality of measurements and the other measurements.
- 10 2. The system of claim 1, wherein the request to store the plurality of measurements includes the metric identifier.
3. The system of claim 1 or 2, wherein the one or more services are further configured to aggregate the plurality of measurements with the other measurements based at least in part on the determined metric identifier.
- 15 4. The system of any of claims 1-4, wherein use of the determined fully qualified metric identifier to store the metrics data includes:
 - utilizing the metric identifier and a timestamp of the measurements of the plurality of measurements to partition the plurality of measurements into one or more
 - partitions;
 - 20 transmitting the plurality of measurements from the one or more partitions to one or more aggregator sub-system of a plurality of aggregator sub-systems for storage of the plurality of measurements; and
 - storing the plurality of measurements within one or more aggregator datastores of the one or more aggregator sub-systems.

5. The system of claim 8, wherein the one or more services are further configured to:

obtain, from a metric mapping registry and based at least in part on timestamps for the measurements of the plurality of measurements, one or more mappings
5 between the one or more partitions and the plurality of aggregator sub-systems; and
select the one or more aggregator sub-systems from the plurality of aggregator sub-systems based at least in part on the one or more mappings.

6. The system of any of claims 1-5, wherein the one or more services are further configured to:

10 receive a second request to retrieve the plurality of measurements, the request specifying the metadata;

determine the metric identifier based at least in part on the metadata; and
retrieve, through use of the metric identifier and one or more time periods specified in the metadata, at least both the plurality of measurements and the other
15 measurements.

7. The system of any of claims 1-6, wherein the plurality of measurements are stored in a binary serialization format, the binary serialization format including the metric identifier, a timestamp for the metrics data, a unit of measurement for the metrics data, and a measurement for a metric represented by the metrics data.

20 8. The system of any of claims 1-7, wherein the metric identifier is generated through use of a hash function and one or more metric attributes specified within the metadata.

9. A non-transitory computer-readable storage medium having stored thereon executable instructions that, when executed by one or more processors of a computer system, cause the computer system to at least:

receive a first application programming interface call to store one or more measurements for a metric, the first application programming interface call specifying the one or more measurements and a metric identifier for the one or more measurements, the metric identifier generated based at least in part on metadata included in an earlier request to store other measurements for the metric;

utilize the metric identifier to identify a datastore for storing the one or more measurements; and

store the one or more measurements in the identified datastore such that the one or more measurements and the other measurements can be obtained in response to a request that includes the metadata or the fully qualified metric identifier.

10. The non-transitory computer-readable storage medium of claim 9, wherein the instructions further cause the computer system to:

receive a second application programming interface call to retrieve a plurality of measurements for the metric, the second application programming interface call including the metadata;

generate, using the included metadata, the metric identifier;

identify, based at least in part on the generated metric identifier and one or more time periods specified in the metadata, one or more datastores that include the one or more measurements and the other measurements; and

provide the one or more measurements and the other measurements from the identified one or more datastores.

11. The non-transitory computer-readable storage medium of claims 9 or 10, wherein the instructions further cause the computer system to transform the one or more measurements into a binary serialization format, the one or more measurements in the binary serialization format including the metric identifier and the measurement.

12. The non-transitory computer-readable storage medium of any of claims 9-11, wherein storing the one or more measurements in the identified datastore includes:

retrieving, from the datastore, serialized measurements for the metric, the serialized measurements including the other measurements in the binary serialization format;

5 de-serializing the serialized measurements and the one or more measurements in the binary serialization format;

aggregating the one or more measurements with the other measurements; and

transforming the one or more measurements and the other measurements into the binary serialization format.

10 13. The non-transitory computer-readable storage medium of any of claims 9-12, wherein the instructions further cause the computer system to utilize run length encoding to compress the one or more measurements for one or more sequential measurements having a same value in the binary serialization format.

14. The non-transitory computer-readable storage medium of any of claims 15 9-13, wherein the instructions further cause the computer system to:

generate, based at least in part on the metric identifier and a timestamp of the one or more measurements, a partition;

place the one or more measurements in the logical partition into a queue for distribution to an aggregator sub-system, the aggregator sub-system including the identified 20 datastore; and

transmit, from the queue, the one or more measurements to the aggregator sub-system for storage of the one or more measurements in the identified datastore.

15. The non-transitory computer-readable storage medium of any of claims 9-14, wherein the instructions further cause the computer system to:

25 obtain, from a metric mapping registry, a mapping between partitions and aggregator sub-systems over time; and

utilize the metric identifier and the timestamp of the metrics data to determine, from the mapping, the aggregator sub-system for storage of the one or more measurements into the identified datastore.




FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6








FIG. 7


**FIG. 8**


**FIG. 9**

10/14


**FIG. 10**

**FIG. 11**

**FIG. 12**

**FIG. 13**

14/14

**FIG. 14**

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2016/039371

A. CLASSIFICATION OF SUBJECT MATTER

INV. H04L29/08 G06F11/34 G06F11/30
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)

H04L 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, INSPEC, IBM-TDB

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	Amazon: "Amazon CloudWatch Developer Guide", 10 November 2014 (2014-11-10), pages 1-236, XP055308341, Retrieved from the Internet: URL:https://awsdocs.s3.amazonaws.com/AmazonCloudWatch/latest/acw-dg.pdf [retrieved on 2016-10-06] page 24, step 2 page 3, paragraph 6 page 25, paragraph 1 page 39 - page 40 ----- -/-	1-15



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

6 October 2016

Date of mailing of the international search report

17/10/2016

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

Hoareau, Samuel

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2016/039371

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	Anonymous: "Distributed hash table - Wikipedia, the free encyclopedia", 9 June 2015 (2015-06-09), XP055308352, Retrieved from the Internet: URL:https://web.archive.org/web/20150609172655/https://en.wikipedia.org/wiki/Distributed_hash_table [retrieved on 2016-10-06] page 2, paragraphs "Structure" and "Properties" -----	1-15
A	WO 2013/184935 A1 (RACKSPACE US INC [US]) 12 December 2013 (2013-12-12) page 7 - page 10 -----	1-15
A	TRIHINAS DEMETRIS ET AL: "JCatascopia: Monitoring Elastically Adaptive Applications in the Cloud", 2014 14TH IEEE/ACM INTERNATIONAL SYMPOSIUM ON CLUSTER, CLOUD AND GRID COMPUTING, IEEE, 26 May 2014 (2014-05-26), pages 226-235, XP032614504, DOI: 10.1109/CCGRID.2014.41 [retrieved on 2014-06-30] page 6; figure 5 -----	1-15
A	Eben Hewitt: "Cassandra- The definitive guide" In: "Cassandra- The definitive guide", 1 November 2010 (2010-11-01), O'reilly, XP055308367, ISBN: 978-1-4493-9041-9 pages 1-330, page 51, last paragraph page 269, paragraph 1 page 50, paragraph 4 -----	1-15

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2016/039371

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 2013184935 A1	12-12-2013	AU 2013271538 A1	22-01-2015
		EP 2859473 A1	15-04-2015
		HK 1208745 A1	11-03-2016
		US 2013332484 A1	12-12-2013
		US 2015169650 A1	18-06-2015
		WO 2013184935 A1	12-12-2013
