

Title: [YARN-1530] Store, manage and serve per-framework application-timeline data

By: Arun C Murthy, Billie Rinaldi, Gopal Vijayaraghavan, Hitesh Shah, Mayank Bansal, Vinod Kumar Vavilapalli, Zhijie Shen

Last modified: January 8 2014

1.0 Preamble

YARN-321 addresses the problem of storing, managing and retrieving *generic historical* information about applications. The application-history problem is solved by making ResourceManager write the events as and when they happen (application-start, container-start, app-finish etc.) to a pluggable store (default implementation based on HDFS files) and by serving the information via a generic **ApplicationHistoryServer (AHS)**. AHS deliberately avoids solving issues with storing, managing, retrieving and rendering framework-specific data - both current and historical - to manage complexity. YARN-1530 aims to address this gap as detailed in this document.

1.1 Problem statement

Today, each and every application/framework has to store, manage and serve per-framework data all by itself as YARN doesn't have a common solution. MapReduce does (and by extension any application/framework can do) this by

- Running a web-server inside the **ApplicationMaster (AM)** to serve non-generic information about the running application
- Making ApplicationMaster emit events to what are known as Job-History files, manage them on HDFS and serve them via a MR (framework) specific JobHistoryServer (daemon).
- Clients use RPCs, web-UI and web-services on JobHistoryServer to learn about finished applications.

YARN needs a common solution for handling non-generic information about running applications as well as about finished applications; with plugin points for frameworks to do their own rendering.

1.2 Proposal

YARN-1530 JIRA ticket attempts to solve the storage, management and serving of per-framework data from various applications, **both** *running* and *finished*. The aim is to change YARN to enable collection, aggregation and storage of this data in a generic manner with plugin points for frameworks to do their own thing w.r.t data-interpretation and serving. The solution is

in a way a generalization of MapReduce JobHistory events and thus involves

- **A timeline writer API/Library:** A simple REST API and/or a more involved client library that AMs and containers can use to send data across.
- **Event aggregation:** Once AMs and clients write their data to a API/library, YARN uses an aggregation system to aggregate all the events from different AMs & containers across applications; which eventually feeds into a storage system. The event-aggregation is a function of the API/library.
 - We can have REST APIs via which the clients (AMs and/or containers) directly post the events to a central location. In this case, the event-aggregation merges itself with the writer library + storage system.
 - On the other hand, event-aggregation itself can be a pluggable system that makes use of external frameworks like Kafka/Flume and potentially gives better guarantees w.r.t reliability, fault-tolerance, etc.
- **Data storage:** A pluggable storage to persist all the events into a centralized durable storage for serving and analysis. This can include
 - An embedded or local file-system based storage for simple and small deployments.
 - An implementation for larger scale deployments based on a system like HBase.
- **Serving**
 - Reusing the serving layer in Application History Server (from YARN-321) and extending it to be a generic **Application Timeline Server (ATS)**. At that point, as should be obvious, it can serve both generic history as well as per-framework current and historical data in the form of web-services (and/or read-only RPC end-points)
 - Query API: Web service functionality exposed by the Application Timeline Server either gives out the complete set of events or after filtering them based on time, event-type, etc. It *DOES NOT* do any rendering *by default*.
- **Rendering/User interfaces:**
 - To facilitate framework specific user-interfaces, *Application Timeline Server* enables administrators to optionally install rendering-libraries for *specific versions* of specific *well-known*, commonly used applications/frameworks.
 - Other than the admin-ratified rendering libraries, YARN neither encourages nor blocks users from writing their own user-interfaces making use of the serving APIs. It won't directly link to them for reasons of security.

1.3 Details

The system is centered on a (multiple-writer) framework for a given application to persist events of its lifetime in a centralized storage. The Application Timeline Server, the storage layer,

ApplicationMasters/containers are the main players in the event cycle. The AM/containers have to log the events that happen during their life-time. The storage layer and ATS are responsible for providing services to the clients for persistence, extraction and display of this information. More details on each of these individual parties follows:

1.3.1 Timeline writer API/Library

There are multiple options for writing events. Existing systems tend to write them to files or sockets directly or use some kind of an interface like a log-appender, a REST end-point or by using a client library. We've chosen the latter.

The API for application writers to use in their AMs/containers is either a simple REST API and a more involved client library that takes care of more things on behalf of the user. There are several things that may need to be done on top of a simple event-post, these things are implemented by the client library. The choice of the API has much larger implications on the event-aggregation which we inspect more in the corresponding section.

Application timeline

Writers (AMs/containers) write per-app/per-framework data using events. We call these the **Application Timeline events** as they are modeled around timestamps and specific to an application. Because of the distributed nature, the event-stream for the timeline can potentially be fed into the system by the application related entities - AM, containers - in any order and it is the responsibility of the serving layer to reorder the events to be sorted depending on the need. There are two APIs for the writers : (1) a REST API that provides minimum guarantees & enables cross-platform access and (2) an embeddable library that provides more features like buffering, failover, etc.

In the remainder of this section, we showcase the REST APIs. Corresponding APIs can be demonstrated trivially from a library standpoint too.

HTTP Request

Request: PUT http://<host>:<port>/ws/v1/timeline

Response: The status of the request

Description: Put information about a given entity of a given type.

- Does not replace / remove existing information about the entity, just adds to it.
- Can post a batch of events about different entities for efficiency.

An example of the request follows, assuming a three level entity-hierarchy: job, a number of 'vertices' in the job, a number of 'tasks' per vertex. Here, the 'relationship' between a task and a vertex is established by making the task's start-event tagged with a vertex ID. This also lets the

event be retrievable as a child of the vertex ID:

```
PUT http://<host>:<port>/ws/v1/timeline/
[
  {
    "entity" : "task_1381946403170_0003_1_00_000000",
    "entitytype" : "taskId",
    "event" : [
      {
        "ts" : 1381973272402,
        "eventtype" : "startTime",
        "info" : {
          ...
        }
      },
      {
        ...
      },
      ...
    ],
    "tag" : [
      {
        "entity" : "vertex_1381946403170_0003_1_00",
        "entitytype" : "vertexId"
      },
      ...
    ],
    "info" : {
      <name>: <value>,
      ...
    }
  }
]
```

More details about the API:

- Examples of anticipated tags include references to other entities like user, queue, rack, node, etc.
- If tags are specified as part of the call, the events are correlated to (and thus may be stored / duplicated under) each tag's child timeline; this is the only way currently provided by the API to add events to a child timeline.
- If tags are not specified, the events are only associated with the original entity, and are not correlated to any tags. This allows the application to control which child events are stored / retrievable by tag. For example, the application may want to store only the BEGIN and END events with a tag, while leaving the remaining events to be retrievable by

the original entity. Or it may want to tag all of the entity's events.

- Tags are additive. Existing tags are not replaced on event-post with new tags.
- Events are by default added into an entity's timeline, however if an identical event is put for the same entity more than once with different **info** data, the info may be overwritten.

1.3.2 Event aggregation

YARN needs to use an aggregation system to aggregate all the events from different AMs & containers across applications and then to eventually persist them into a storage layer. As mentioned before, the event-aggregation is a function of the API/library.

- Clients (AMs and/or containers) directly post the events to a central location via REST APIs. In this case, the event-aggregation merges itself with the writer library + storage system.
- On the other hand, event-aggregation itself can be a pluggable system that makes use of external frameworks like Kafka/Flume and potentially gives better guarantees w.r.t reliability, fault-tolerance, etc.

A quick description of various '*ilities*' desired with the event-aggregation sub-component:

- **1.3.2.1 Scalability:**

- A REST API based system has scalability issues with tens of thousands of AMs + containers pushing events out. The fact that events can have custom, application-specific payload too (and not just control-data) which also is routed through the web-service layer only exacerbates this. It can be scaled on the server-side by, say, using a farm of web-proxies. The good part about this is that it is a known path towards scaling web-services. The bad is that it adds to the operator burden.
- A library can let administrators choose a scalable event-aggregation framework like Kafka/Flume completely transparent to the user. The con to this approach is adding yet another dependency to YARN. The library for example could use the REST API itself for smaller installations and seamlessly change to a scalable aggregation framework.
- There is a third option of choosing a scalable framework and directly expose users to that API. But our requirement of flexibility between small and large clusters, extra vs no external dependencies pushes us to not pursue this.

- **1.3.2.2 Reliability:** If applications wish to use the aggregated timeline say as a transaction log for their recovery (this is what MapReduce ApplicationMaster does with its JobHistory), we will need reliability guarantees for the event-storage. For e.g., MR JobHistory is flushed every so often by the MR AppMaster so that it can get as many recorded events as possible when recovery happens.

- A REST based API implies that an event-write acked by a 200(OK) status code amounts to a success.

- A library based approach may involve exposing APIs to flush and/or client call-back handlers. It can explicitly implement it itself or delegate it to an underlying aggregation system if it supports such APIs.
- **1.3.2.3 Fault tolerance:** Clients writing the events may go down before the acknowledgement comes in. The aggregation-system may also go down in a similar fashion. Network failures may also happen. The Application-writers need some guarantees on the API if they are to rely on the written data for more than UI purposes.
 - A REST based API may force client to explicitly retry.
 - A library can wrap the retry functionality within itself.
- **1.3.2.4 Buffering:** This is listed as a separate requirement and not as a solution as it is one of the natural choices to address some of the scalability and the fault-tolerance problems listed above. Buffering helps scalability by queueing events locally when destination systems are loaded. Events can also be queued in the presence of a fault somewhere - client, aggregation-system or the network - and can be resent once the system recovers from the fault.
 - A REST based user-interface needs APIs that accept batches of events. On top of that, it requires each and every client to buffer events for a while before sending them across.
 - A library can wrap the buffering functionality itself. It can implement buffering itself or delegate it to the underlying aggregation system if it supports buffering.
- **1.3.2.5 Timeliness:** For serving the information while the application is running, the event-aggregation as well as the storage need to be as real-time as possible. Some lag for the end-user is acceptable but it shouldn't go beyond a few seconds at a maximum.
- **1.3.2.6 Forwarding:** Unlike other messaging/pub-sub frameworks, the destination of the events in our case is a single storage system, we do not have a need for multiple destinations. The underlying storage itself may be distributed. So we don't need to explicitly implement forwarding for the sake of load-balancing, scalability etc.

1.3.3 Data storage

To enable site-specific implementation of the storage of application data, we will use a pluggable storage to store all the events coming across from AMs & containers of various applications. Having the event-data in a centralized storage is desired both for scalable serving and for centralized analysis. There are few requirements here:

- Pluggable system to cater to small and large installations
- The primary storage that will be shipped needs to enable structured data for enabling near real-time querying (with indexes etc) **and** rendering as well as better analytics (without say running batch jobs).

For simple and small deployments, we do not want to put the burden of installing one more

system as a YARN dependency. With YARN being the lowest level of the architecture, forced addition of more systems underneath is a non-starter. So an embedded implementation that works on top of local file-system based storage will be shipped as default. For enabling a structured system that scales up, we are looking at one of the basic implementations to be based on Level-DB that can work against local-system and scale up by integrating with a nosql database like HBase.

1.3.4 Serving

Clients are served by the web-services layer in the Application Timeline Server (Previously AHS). The web-service acts as the serving layer that bridges the gap between the storage-layer and the client which is rendering the information in various ways possible. We can imagine implementing a RPC layer too (as with other things in YARN) but it isn't a priority as of now.

We now describe the APIs exposed by the serving layer.

1.3.4.1 Application Timeline REST API for retrieval

An example API follows:

HTTP Request I

```
Request: GET http://<host>:<port>/ws/v1/apptimeline/<entityType>
          [&tag=<entity>&tagType=<entityType>]
          [&limit=<limit>]
          [&startTime=<startTime>][&endTime=<endTime>]
          [&eventType=<eventType>]
```

Description: Retrieves a timeline of events for an entity type.

- With the tag & tagType options, it retrieves only events corresponding to the specified tag.
- With the startTime option, information about all events for an entity type since a start time or, if an endTime is specified, between a start time and end time are returned.
- With the limit option, one can request the most recent N events or, if an endTime is specified, the N events that occurred immediately before a given end time (inclusive).
- Events may be restricted to a particular type or set of types by specifying one or more event type query parameters.

Example

Request: GET http://<host>:<port>/ws/v1/apptimeline/appAttemptId

Response:

```
{
  "entitytype" : "appAttemptId",
  "event" : [
    {
      "entity" : "appattempt_1381946403170_0003_000001",
      "ts" : 1381973269856,
      "eventtype" : "startTime",
      "info" : {
        ...
      }
    },
    ...
  ]
}
```

HTTP Request II

Request: GET http://<host>:<port>/ws/v1/apptimeline/<entitytype>/<entityId>

Description: Retrieves information about a given entity of a given type.

- Includes the entity's timeline (if it exists), as well as any child information that has been tagged with this entity.
- The child events are grouped by entity. They are not guaranteed to include all events for the child entities; only those events that have been explicitly tagged with the queried entities. To retrieve full event information, a separate request must be made for the child entity.
- Additional information about the entity may also be stored, such as summary information, or information that does not change from event to event.

Example

Request: GET http://<host>:
<port>/ws/v1/apptimeline/vertexId/vertex_1381946403170_0003_1_00
Response:

```
{
  "entity" : "vertex_1381946403170_0003_1_00",
  "entitytype" : "vertexId",
  "event" : [
    {
      "ts" : 1381973253215,
      "eventtype" : "initRequestedTime",
      "info" : {
        ...
      }
    },
    ...
  ],
  "children" : [
    {
      "entity" : "task_1381946403170_0003_1_00_000000",
      "entitytype" : "taskId",
      "event" : [
        {
          "ts" : 1381973272402,
          "eventtype" : "startTime",
          "info" : {
            ...
          }
        },
        ...
      ]
    },
    ...
  ],
  "info" : {
    "processorName" :
    "o.a.tez.mapreduce.processor.reduce.ReduceProcessor",
    "vertexName" : "finalreduce",
    "counters" : { "hdfs_bytes_read" : 0.0, ... }
    ...
  }
}
```

1.3.5 Rendering/User interfaces

The client library can use the APIs exposed by the serving layer and build framework specific UIs. One obvious choice is to build it in JavaScript, serve them somewhere and let users access them.

To facilitate the ease of deployment of the framework specific user-interfaces, *Application Timeline Server* enables administrators to optionally install rendering-libraries for *specific versions* of specific *well-known*, commonly used applications/frameworks.

- As should be obvious, any changes in the UI at run-time by the application-developer need ratification by admins before deployment. This comes in the way of YARN's application agility requirement - other than the UI, developers are free to, and even encouraged, to evolve applications in an agile manner.
- Due to security reasons, we cannot let arbitrary code to be installed dynamically in the ATS. Because of this, at any point of time, ATS only allows specific versions for each framework to be installed by the admins. Managing lot of versions will be a burden to admins and we don't foresee more than a few versions being supported.
- ApplicationMasters can set the finalTrackingUrl as part of Application unregistration to this URL so that users are transparently redirected to the right UI on the ATS.

As mentioned previously, other than the admin-ratified & installed rendering libraries, YARN neither encourages nor blocks users from writing their own user-interfaces making use of the serving APIs, and setting up their own daemons. YARN cannot explicitly link to them for reasons of security - AMs can set finalTrackingURLs pointing to these locations, but the content may be filtered through the central web-application-proxy for security.

1.3.6 Miscellaneous

- **Time:** For the timeline to make any sense, the timestamps all need to be globally ordered and so will require us to depend on synchronized clocks.
- **Use beyond AM/Containers:** One can imagine events flowing out of YARNClient, framework clients like MR JobClient, or even upper layers like Pig, Hive, Oozie etc. We provide the API, but leave correlation across entities and their visualization to the framework specific UIs.

1.4 Merging the generic and per-framework data

It is very conceivable to use this same system with ResourceManager as the client/source of events for persisting generic historical data.

- ResourceManager can add extra fields which indicate the lifetime of a container and its allocation - e.g. container start, finish etc.
- It can add queue, user, application_type, application_name, application_id, application_state, container_id, rack and node_id etc too.

- These fields added by the RM can in addition be indexable by the storage layer for faster access of generic information.

But we are keeping the scope of the integration of generic data into this framework out of this document and for later.