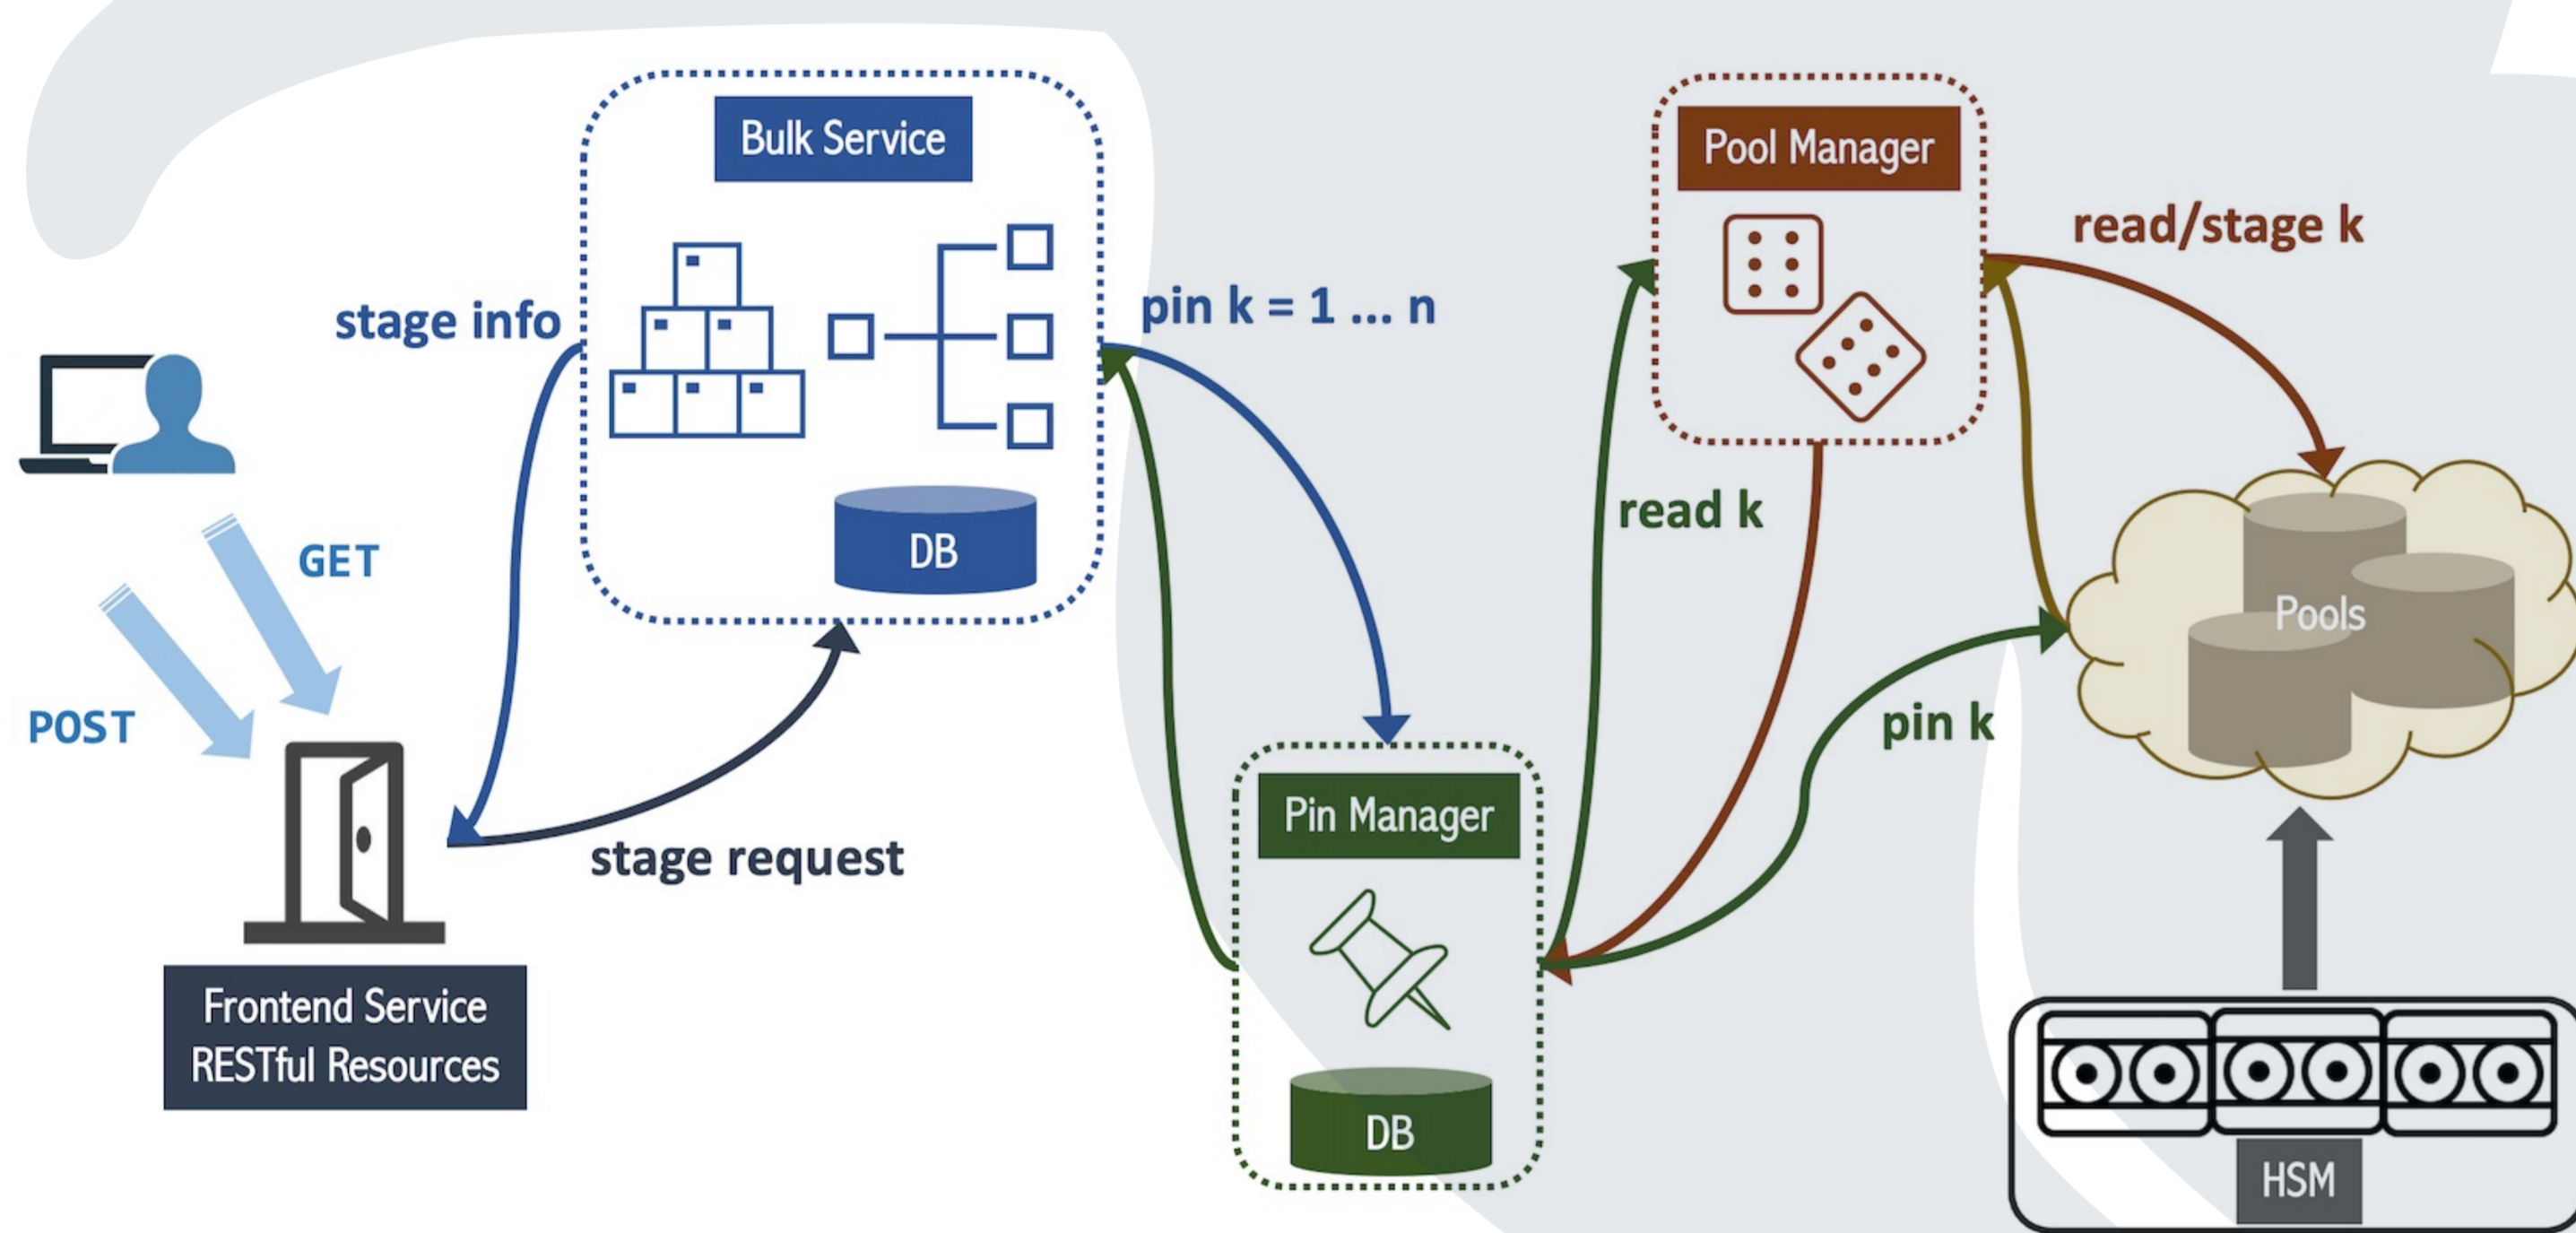


The Bulk service and WLCG TAPE API support in dCache

Authors: ALBERT ROSSI, Dmitry Litvintsev (FNAL); Svenja Meyer, Paul Millar, Tigran Mkrtchyan, Lea Morschel, Marina Sahakyan (DESY); Krishnaveni Chitrapu (NSC)

Staging via REST API and Bulk service

As part of the WLCG collaboration, sites deploying dCache will require support for the TAPE REST API allowing for massive recall of files from tape. dCache implements the specified resources in its Frontend service, relaying the requests to a general-purpose Bulk service for handling. Similarly, the status of a request is also obtained via a RESTful query to the Frontend which retrieves the information from the Bulk service. Policies affecting throughput, number of requests per user and maximum files per request are all configurable.



dCache: Staging Interactions

REST API: SWAGGER

(<https://example.org:3880/api/v1>)

tape Support for the TAPE API (bulk)	
POST	/tape/archiveinfo Return the file locality information for a list of file paths.
POST	/tape/release/{id} RELEASE files associated with a STAGE request.
POST	/tape/stage/{id}/cancel Cancel a STAGE request.
POST	/tape/stage Submit a STAGE request.
GET	/tape/stage/{id} Get the status information for an individual stage request.
DELETE	/tape/stage/{id} Clear all resources pertaining to the given stage request id.
bulk-requests	
GET	/bulk-requests/{id} Get the status information for an individual bulk request. If the number of request targets equals 10K, retry using the offset (highest returned seqNo + 1) to fetch more targets.
DELETE	/bulk-requests/{id} Clear all resources pertaining to the given bulk request id.
PATCH	/bulk-requests/{id} Take some action on a bulk request.
GET	/bulk-requests Get the summary info of current bulk operations. If the number of requests returned equals 10K, retry using the offset (highest returned seqNo + 1) to fetch more requests.
POST	/bulk-requests Submit a bulk request.

The dCache Frontend provides a SWAGGER page which describes the REST API and allows the user to try out the various resources.

WLCG TAPE: A special case of the generic Bulk request

	Pin	BULK	WLCG TAPE
POST*	Pin	api/v1/bulk-requests { "activity": "PIN", "arguments": { "lifetime": "1", "lifetimeUnit": "DAYS", "id": "<cid>", "target": { "path", ... } } }	N/A
POST*	Stage to disk	api/v1/bulk-requests { "activity": "PIN", "arguments": { "lifetime": "1", "lifetimeUnit": "DAYS", "id": "<cid>", "target": { "path", ... } } }	api/v1/tape/stage { "files": { "diskLifetime": "P1D", "path", ... } } (diskLifetime is ISO 8601)
POST*	Unpin	api/v1/bulk-requests { "activity": "UNPIN", "arguments": { "id": "<request_id>", "target": { "path", ... } } }	N/A
POST	Release	api/v1/bulk-requests { "activity": "UNPIN", "arguments": { "id": "<request_id>", "target": { "path", ... } } }	api/v1/tape/release/<request_id> { "paths": { "path", ... } }
POST*	Change file QoS	api/v1/bulk-requests { "activity": "UPDATE_QOS", "arguments": { "targetQos": "<new qos>", "target": { "path", ... } } }	N/A
POST*	Delete	api/v1/bulk-requests { "activity": "DELETE", "arguments": { "skipDirs": "true/false", "target": { "path", ... } } }	N/A

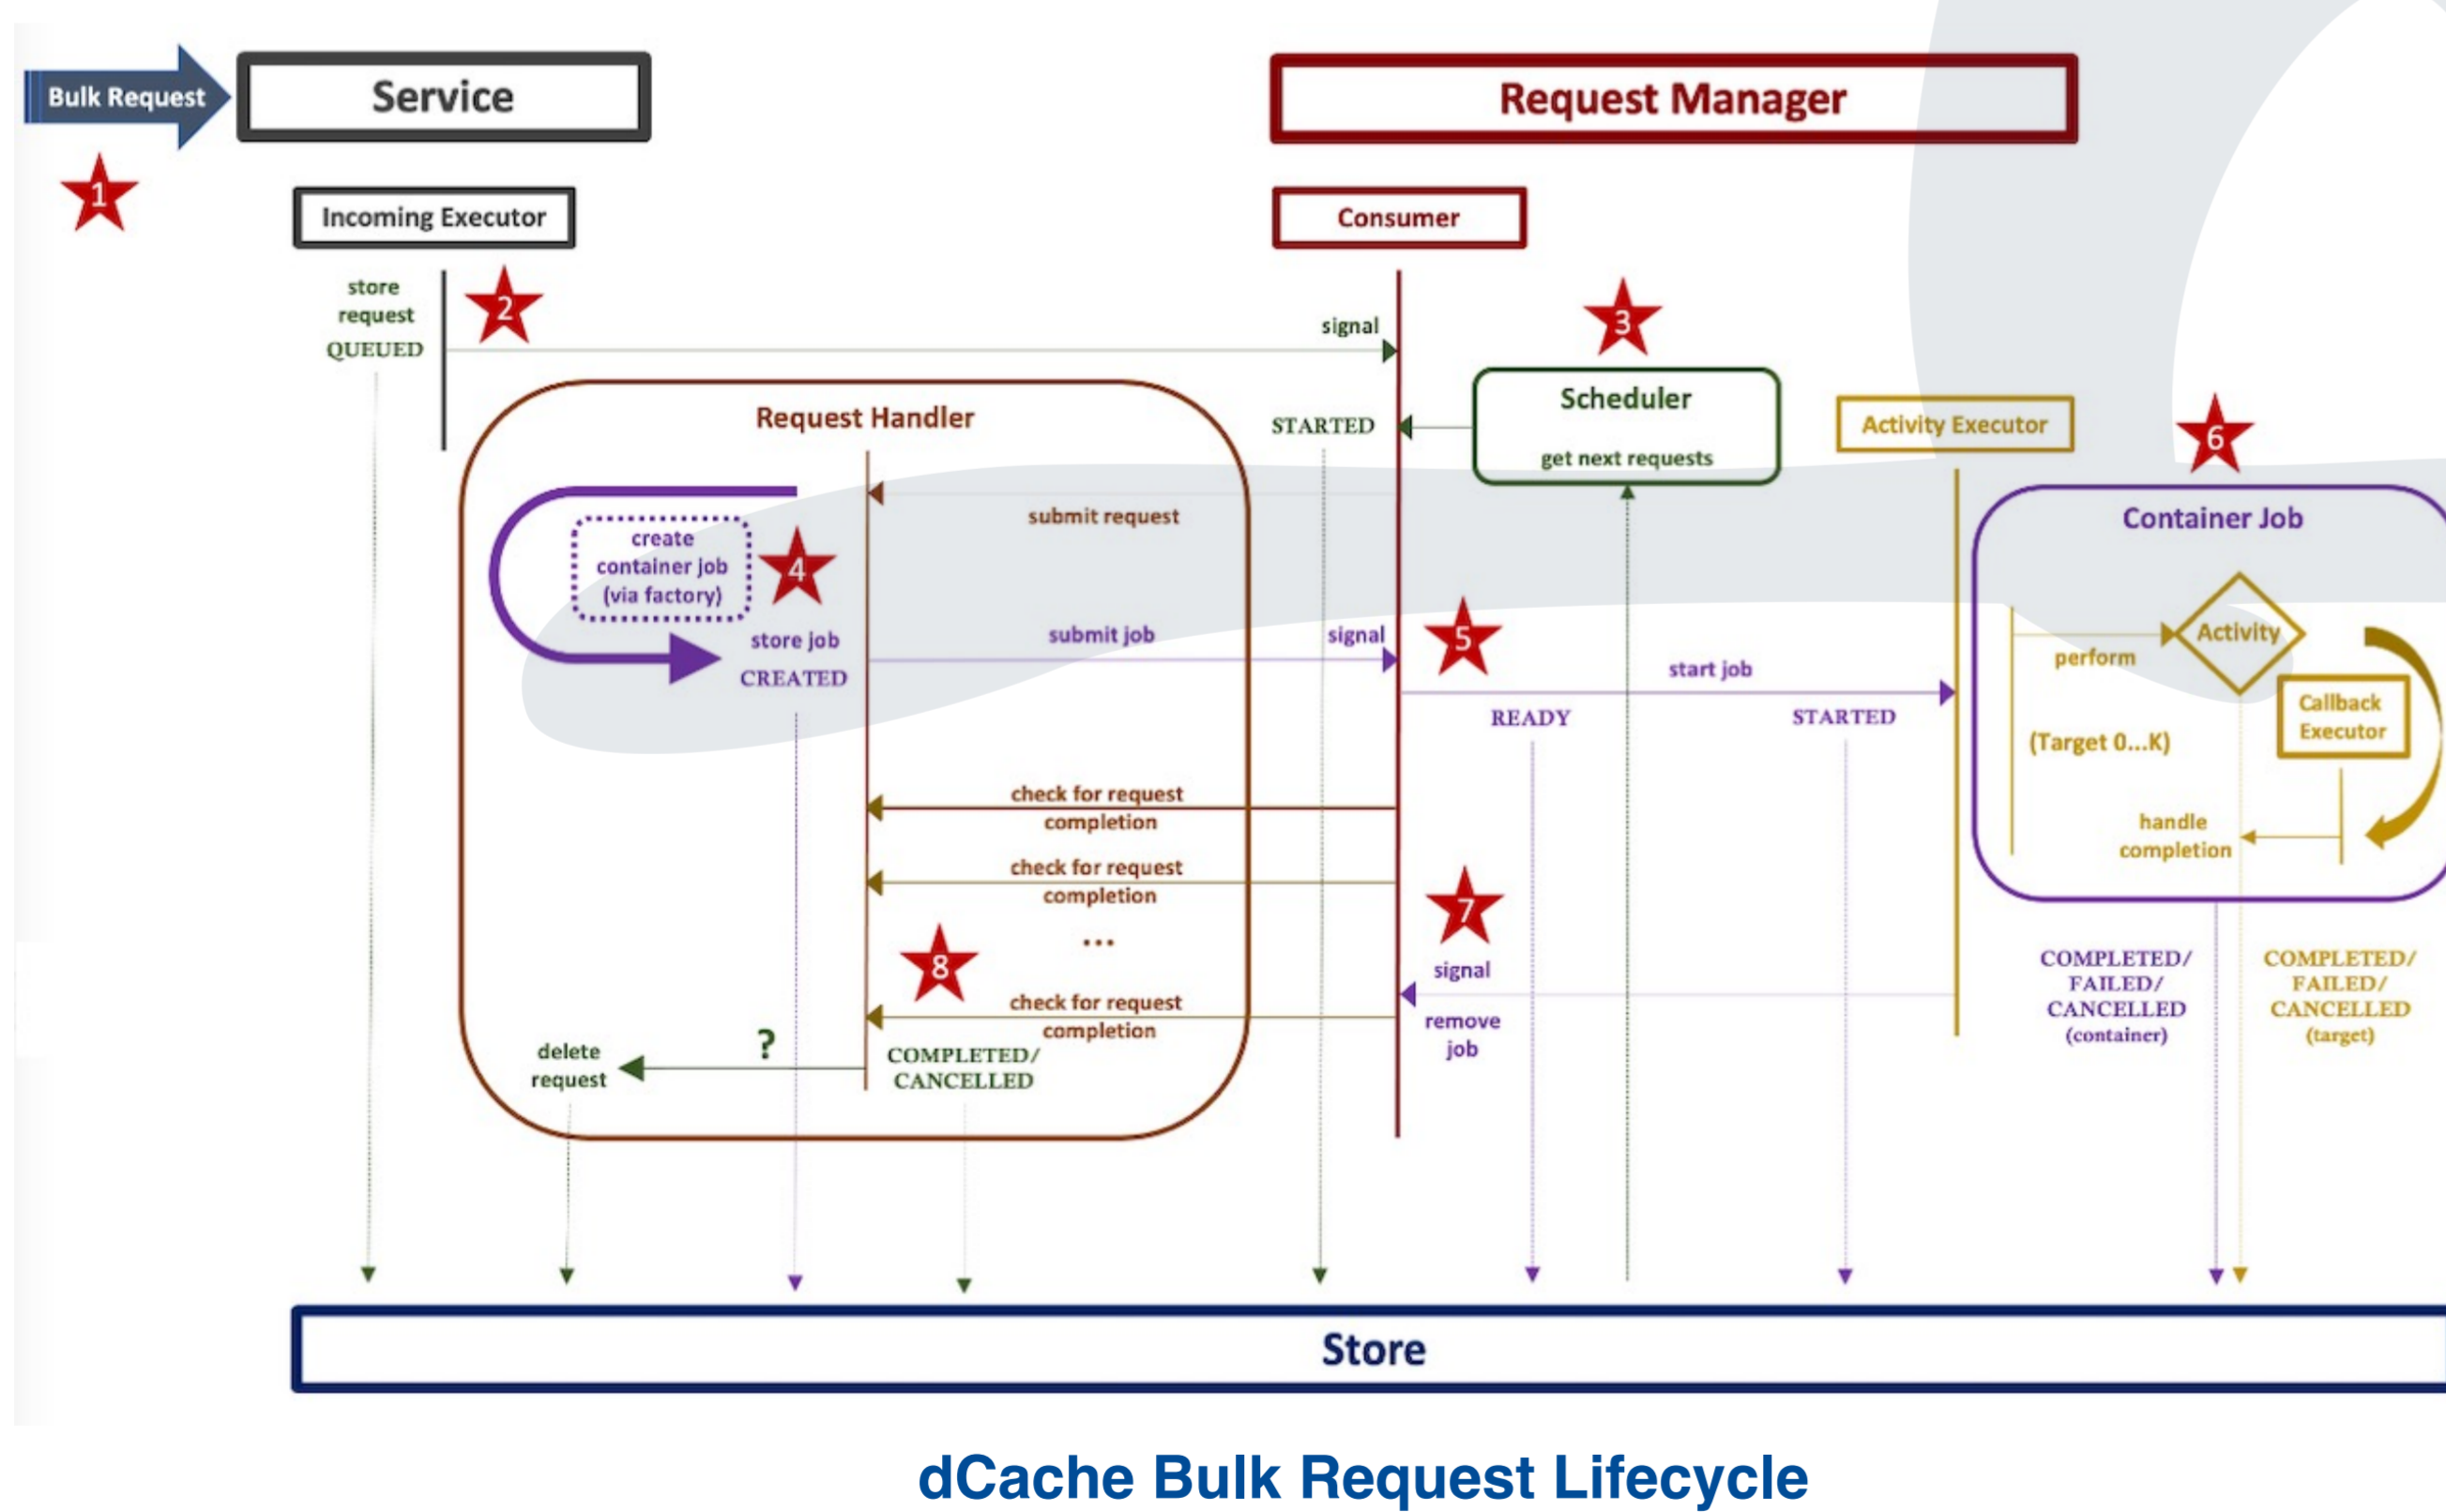
(* The request id is returned with the response.

```
{
  "nextId": 2030215,
  "uid": "0f8f4102-47f1-4b0a-aeda-e40280c2a769",
  "arrivedAt": 1652712101542,
  "startedAt": 1652712101551,
  "lastModified": 1652712369571,
  "status": "COMPLETED",
  "targets": [ {
    "target": "/pnfs/fs/usr/fermilab/users/arossi/tape/000/data-2022050912561652118976111376885",
    "state": "COMPLETED",
    "submittedAt": 1652712101670,
    "startedAt": 1652712101670,
    "finishedAt": 1652712101672,
    {
      "id": "0f8f4102-47f1-4b0a-aeda-e40280c2a769",
      "createdAt": 1652712101542,
      "startedAt": 1652712101551,
      "completedAt": 1652712369571,
      "files": [ {
        "path": "/pnfs/fs/usr/fermilab/users/arossi/tape/000/data-2022050912561652118976111376885",
        "finishedAt": 1652712101672,
        "startedAt": 1652712101670,
        "state": "COMPLETED"
      } ],
      ...
    }
  } ]
}
```

JSON GET responses, Bulk (PIN) request (left) vs WLCG TAPE stage request (right)

Lifecycle of a Bulk request

- 1) Service receives request message.
- 2) Service stores unprocessed request, signals manager.
- 3) Consumer retrieves next requests from scheduler.
- 4) Request submitted to handler, which creates and stores container job.
- 5) Job submitted to manager, which launches it using the activity's thread executor.
- 6) Job processes request targets by: expanding directories (if indicated); pre-storing discovered targets (if indicated); performing activity; processing completion as callback; storing and/or updating.
- 7) Consumer removes completed job from queue.
- 8) Consumer checks and updates state of requests. If clear is indicated, deletes request.



dCache Bulk Request Lifecycle

