

Federated Cluster Support



Dominik Bartkiewicz, Brian Christiansen
SchedMD LLC

Slurm User Group Meeting 2016

Background



- Slurm has long had limited support for federated clusters
 - Most commands support a “--cluster (-M)” option to route requests to different clusters
- Submitted jobs are routed to one cluster
- Each cluster operates independently
 - Job IDs on each cluster are independent (two jobs can have same ID)
 - No cross-cluster job dependencies
 - No job migration between clusters
 - No unified view of system state, each cluster largely independent

Related Work



- There was a community project that was done to address these shortcomings in functionality, but was lacking in scalability and was never integrated into the official release
- The major problem was the use of a single daemon to manage the mapping of job ID to cluster
- The slurmdbd maintained a table identifying which job IDs were on each clusters
- Slurmdbd used for job dependency testing
- Placed very heavy load on slurmdbd to locate jobs

New Capabilities



- **Job Distribution**

- Jobs distributed across federation (packed or spread)
- Unique job IDs

- **Unified Views**

- Appear as one cluster

- **Easy Administration**

- Add/remove clusters to/from the federation with database commands

Design Goals



- **Performance:** Little to no reduction in throughput of each cluster, performance scales with cluster count
- **Scalability:** No reduction in scalability of individual clusters, able to support many federated clusters
- **Ease of use:** Unified enterprise-wide view, minimize change in user interface
- **Stability:** No change in behavior for clusters not explicitly placed into a federation

Eliminating the Bottlenecks



- Need mechanism to identify the cluster associated with a job ID without using slurmdbd lookup
- Make use of 32-bit job ID
 - Embed cluster ID within the job ID
 - Bits 26-31: Cluster ID (1 to 63)
 - Bits 0-25: Local Job ID (1 to 67,108,863)
- Unique job ID across all clusters
 - Large but unique: 134296327 (ClusterID:2 + JobID:78,599)

Configuration

- A cluster can only be part of one federation at a time
- `sacctmgr add federation <fedname> [flags=LLC][clusters=<list>]`
 - `sacctmgr mod federation <fedname> flags[+|-]=LLC`
 - `sacctmgr mod federation <fedname> clusters[+|-]=<list>`
- `sacctmgr mod cluster <cluster> ...`
 - `set federation=<federation>`
 - `set weight=#`
 - Used to prioritize clusters that can start job the soonest
 - `set features=<features>`
 - Features at a cluster level that can be requested by jobs

Configuration

- `sacctmgr mod cluster <cluster> ...`
 - `set fedstate=<state>`
 - **ACTIVE:** Cluster will actively accept and schedule federated jobs.
 - **INACTIVE:** Cluster will not schedule or accept any jobs.
 - **DRAIN:** Cluster will not accept any new jobs and will let existing federated jobs complete.
 - **DRAIN+REMOVE:** Cluster will not accept any new jobs and will remove itself from the federation once all federated jobs have completed. When removed from the federation, the cluster will accept jobs as a non-federated cluster.

Configuration



- A cluster can only be part of one federation at a time
- Jobs can't span multiple clusters

Configuration

```
$ sacctmgr show fed tree
```

Federation	Flags	Cluster	ID	Weight	FedState

fed					
		fed1	1	1	ACTIVE
		fed2	2	1	ACTIVE
		fed3	3	1	ACTIVE

```
$ scontrol show fed
```

```
Federation: fed
```

```
Self: fed1:172.19.0.3:30002 ID:1 FedState:ACTIVE Weight:1
```

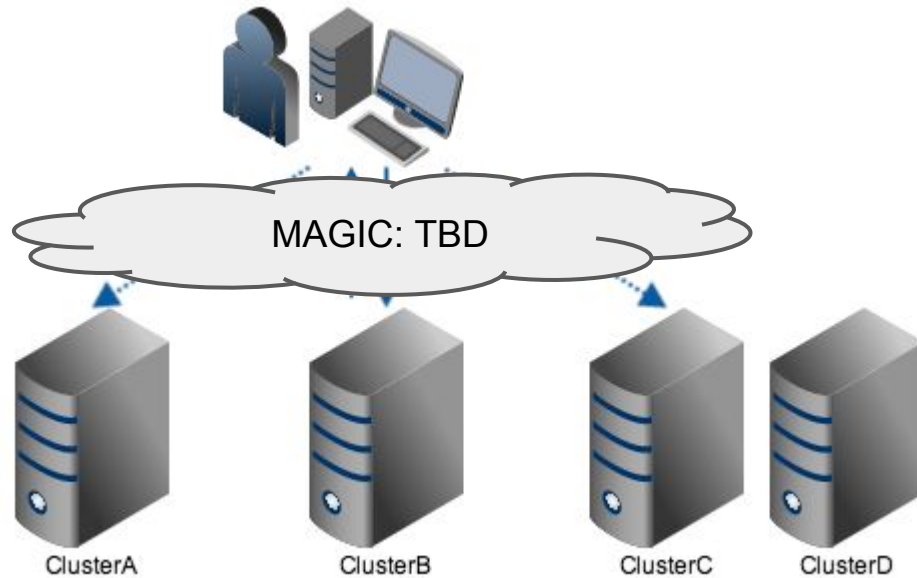
```
Sibling: fed2:172.19.0.4:30002 ID:2 FedState:ACTIVE Weight:1 PersistConnSend/Recv:Yes/Yes
```

```
Sibling: fed3:172.19.0.5:30002 ID:3 FedState:ACTIVE Weight:1 PersistConnSend/Recv:Yes/Yes
```

Persistent Connections

- Clusters talk to each other over persistent connections
 - Reduces communication overhead -- only authenticate once
 - Broken connections detected immediately and established when needed
 - Controller and SlurmDBD use the same code

Job Submission 2015



Job Submission

- sbatch, salloc, srun supported
- Jobs submitted to local cluster
- Local cluster attempts to find a cluster that can start the job now
 - By default packs jobs onto busiest clusters to reduce fragmentation
 - Can spread to least loaded clusters with LLC federation flag.
 - Cluster weights influence which clusters get used first.
 - Similar to node weights.
 - Try to find clusters with weight 1 first, then weight 2, ...

Job Submission



- If a cluster can start the job now, a “sibling” job will be submitted to that cluster only
- If no cluster can start the job now, “sibling” jobs are submitted to each cluster in the federation
- A tracker job stays on the local cluster to coordinate and route requests to siblings jobs
 - Job starts, updates, cancellations

Job Submission

- Interactive Jobs
 - Origin cluster will tell waiting salloc/srun which cluster to talk to
 - srun will talk directly to nodes on selected cluster
 - All nodes need to be accessible by each cluster

Scheduling

- Federated jobs contain the locations of all “sibling” jobs
- Each cluster independently schedules each sibling job
- Coordinates with “origin” cluster to start job
 - The origin cluster is determined from the job id
 - Prevents multiple jobs from being started at the same time
 - Policies in place to handle if origin cluster fails
- Once sibling job is started, origin cluster cancels remaining siblings jobs
- Batch jobs can be requeued to federation

Unified Views

- Provide unified view of federated clusters by default
 - `squeue`, `sinfo`, `sacct`, `sreport`, etc.
- Subset of federated clusters can be requested using `-M` option
 - Existing functionality

```
$squeue
JOBID      FED_ORIGIN_RAW FED_ORIGIN  FED_SIBLINGS_RAW  FED_SIBLINGS  TIME ST  NODES  NODELIST
201330292      3      fed3      1,2,3      fed1,fed2,fed3  0:00 PD   5
134219096      2      fed2      1,2,3      fed1,fed2,fed3  0:00 PD   5
67116596       1      fed1      1,2,3      fed1,fed2,fed3  0:00 PD   5
134219088      2      fed2              1      fed1      0:31 R   5 fed1_[6-10]
67116588       1      fed1              1      fed1      0:34 R   5 fed1_[1-5]
67116592       1      fed1              2      fed2      0:22 R   5 fed2_[6-10]
201330284      3      fed3              2      fed2      0:27 R   5 fed2_[1-5]
201330288      3      fed3              3      fed3      0:17 R   5 fed3_[6-10]
134219092      2      fed2              3      fed3      0:20 R   5 fed3_[1-5]
```

Schedule

- Scheduled to be completed May 2017 (Slurm 17.11)

Questions?

