Compute Server– New Features and Enhancements Gurobi 9.0

GUROBI
OPTIMIZATION

The World's Fastest Solver

Speaker Introduction

- **Dr. Michel Jaczynski**
  - Senior Architect at Gurobi Optimization
  - He has 20 years of experience building components, frameworks and Cloud-based services for optimization applications
  - Worked at IBM on the decision optimization Cloud services and application frameworks and also at ILOG on integrating optimization with enterprise software and a framework to optimize semiconductor Fabs operations in near real-time.
  - Dr. Jaczynski holds a Ph.D. in Computer Science from the University of Nice Sophia-Antipolis, France.

Copyright © 2019, Gurobi Optimization, LLC

Compute Server Overview

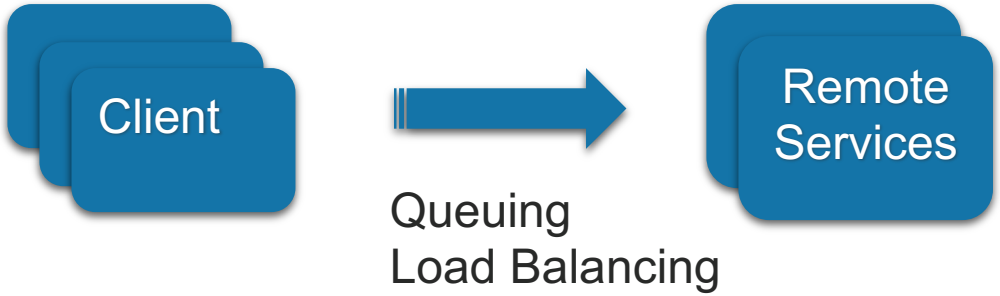

- **Seamlessly offload optimization jobs to one or more servers**
- **Common use cases:**
  - Share resources across users, teams and applications
  - Access more powerful machines
  - Ensure high availability with multiple nodes
  - Build modern applications, service-oriented architecture

Copyright © 2020, Gurobi Optimization, LLC

Gurobi Compute Server V9.0

- For IT
  - Modern and scalable architecture
  - Better security
  - New web UI to manage the cluster
- For OR Experts
  - New web UI to monitor jobs
  - New job history
  - New batch optimization

Cluster Manager

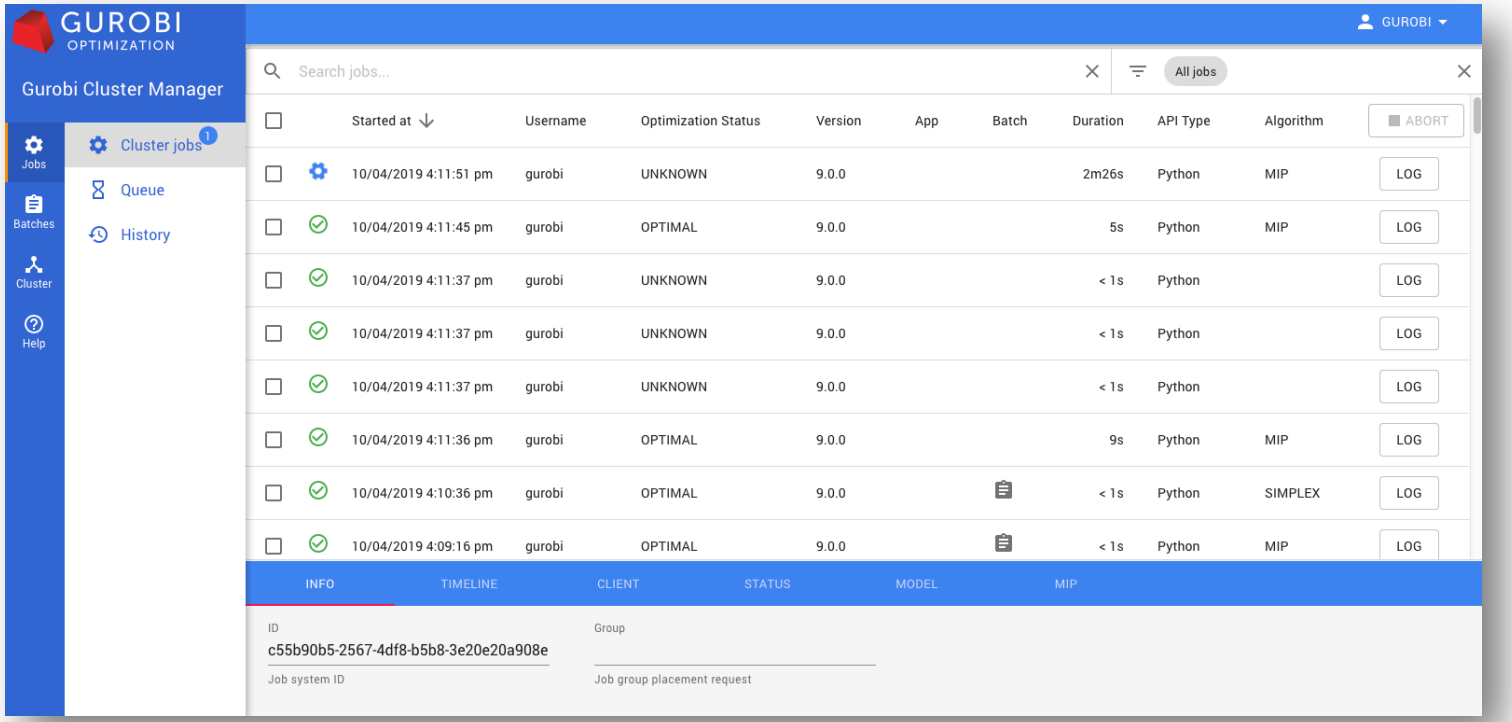

Compute Server v9 facilitates the deployment and use of optimization services on-premises or on private cloud.

Copyright © 2020, Gurobi Optimization, LLC

Architecture: Self-Managed Cluster (V8 and V9)

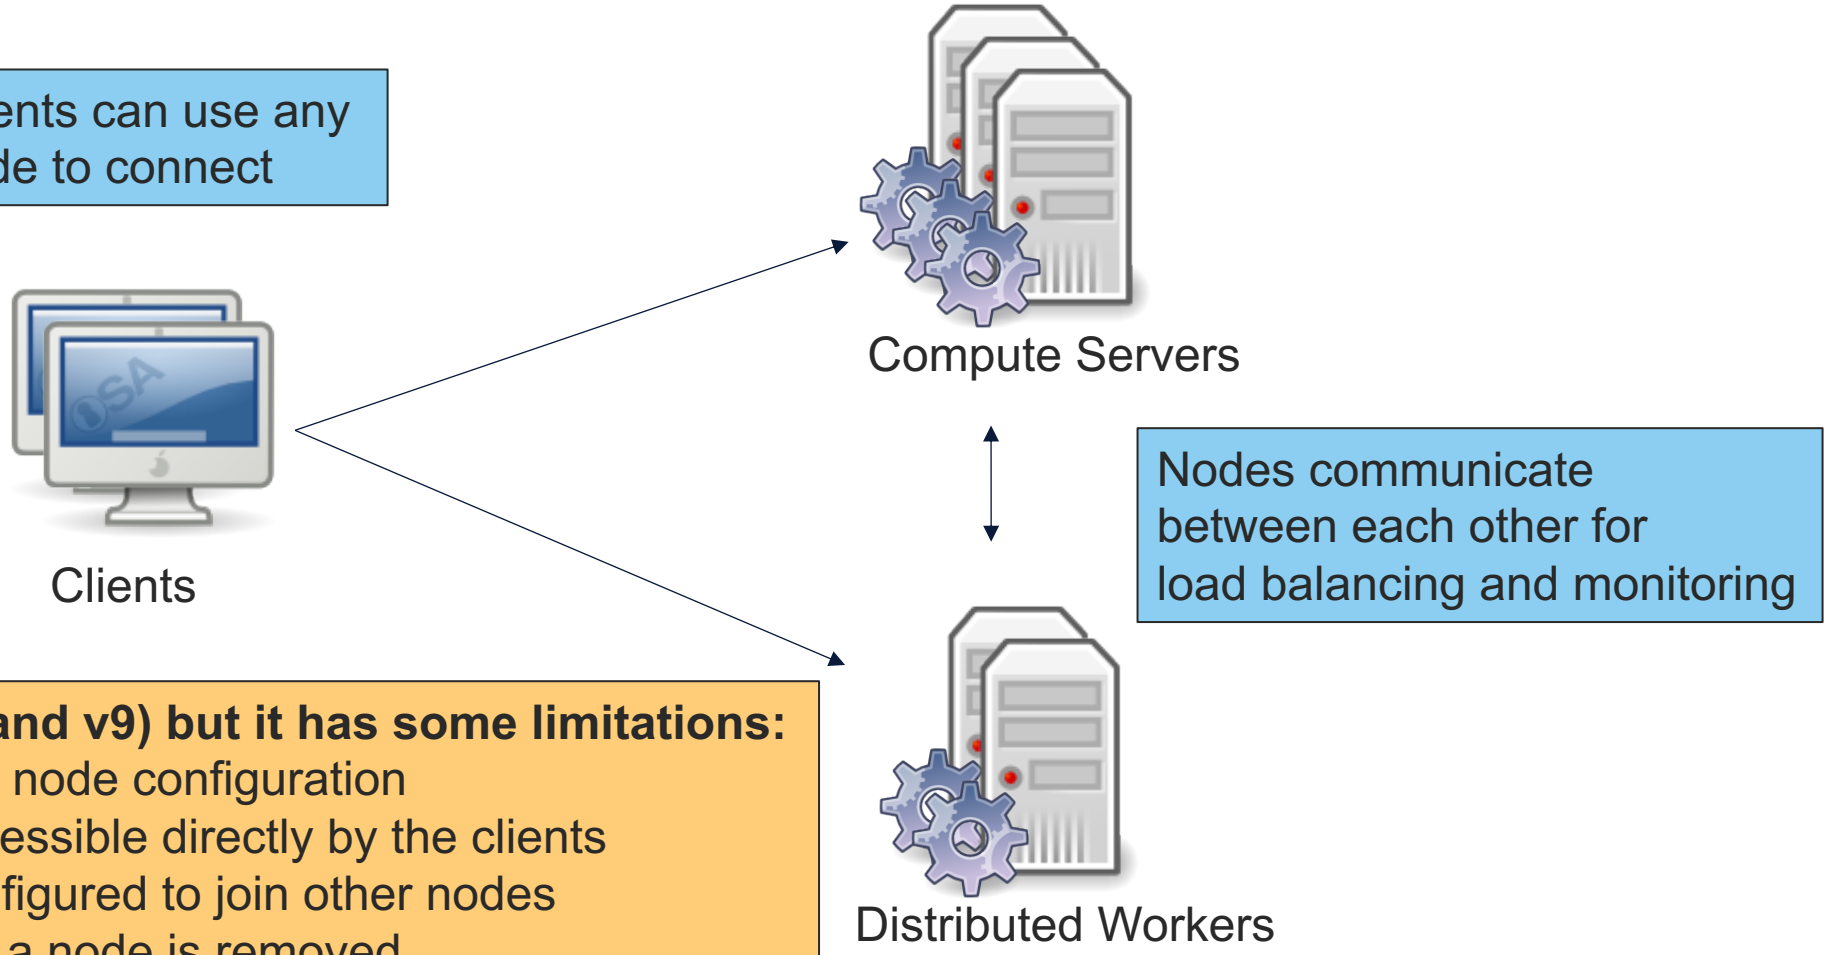

Fully supported (v8 and v9) but it has some limitations:

- Fixed passwords in node configuration
- Nodes must be accessible directly by the clients
- Nodes must be configured to join other nodes
- Job history is lost if a node is removed
- Clients must remain connected during optimization

Copyright © 2020, Gurobi Optimization, LLC

Architecture: Cluster Manager (New in V9)

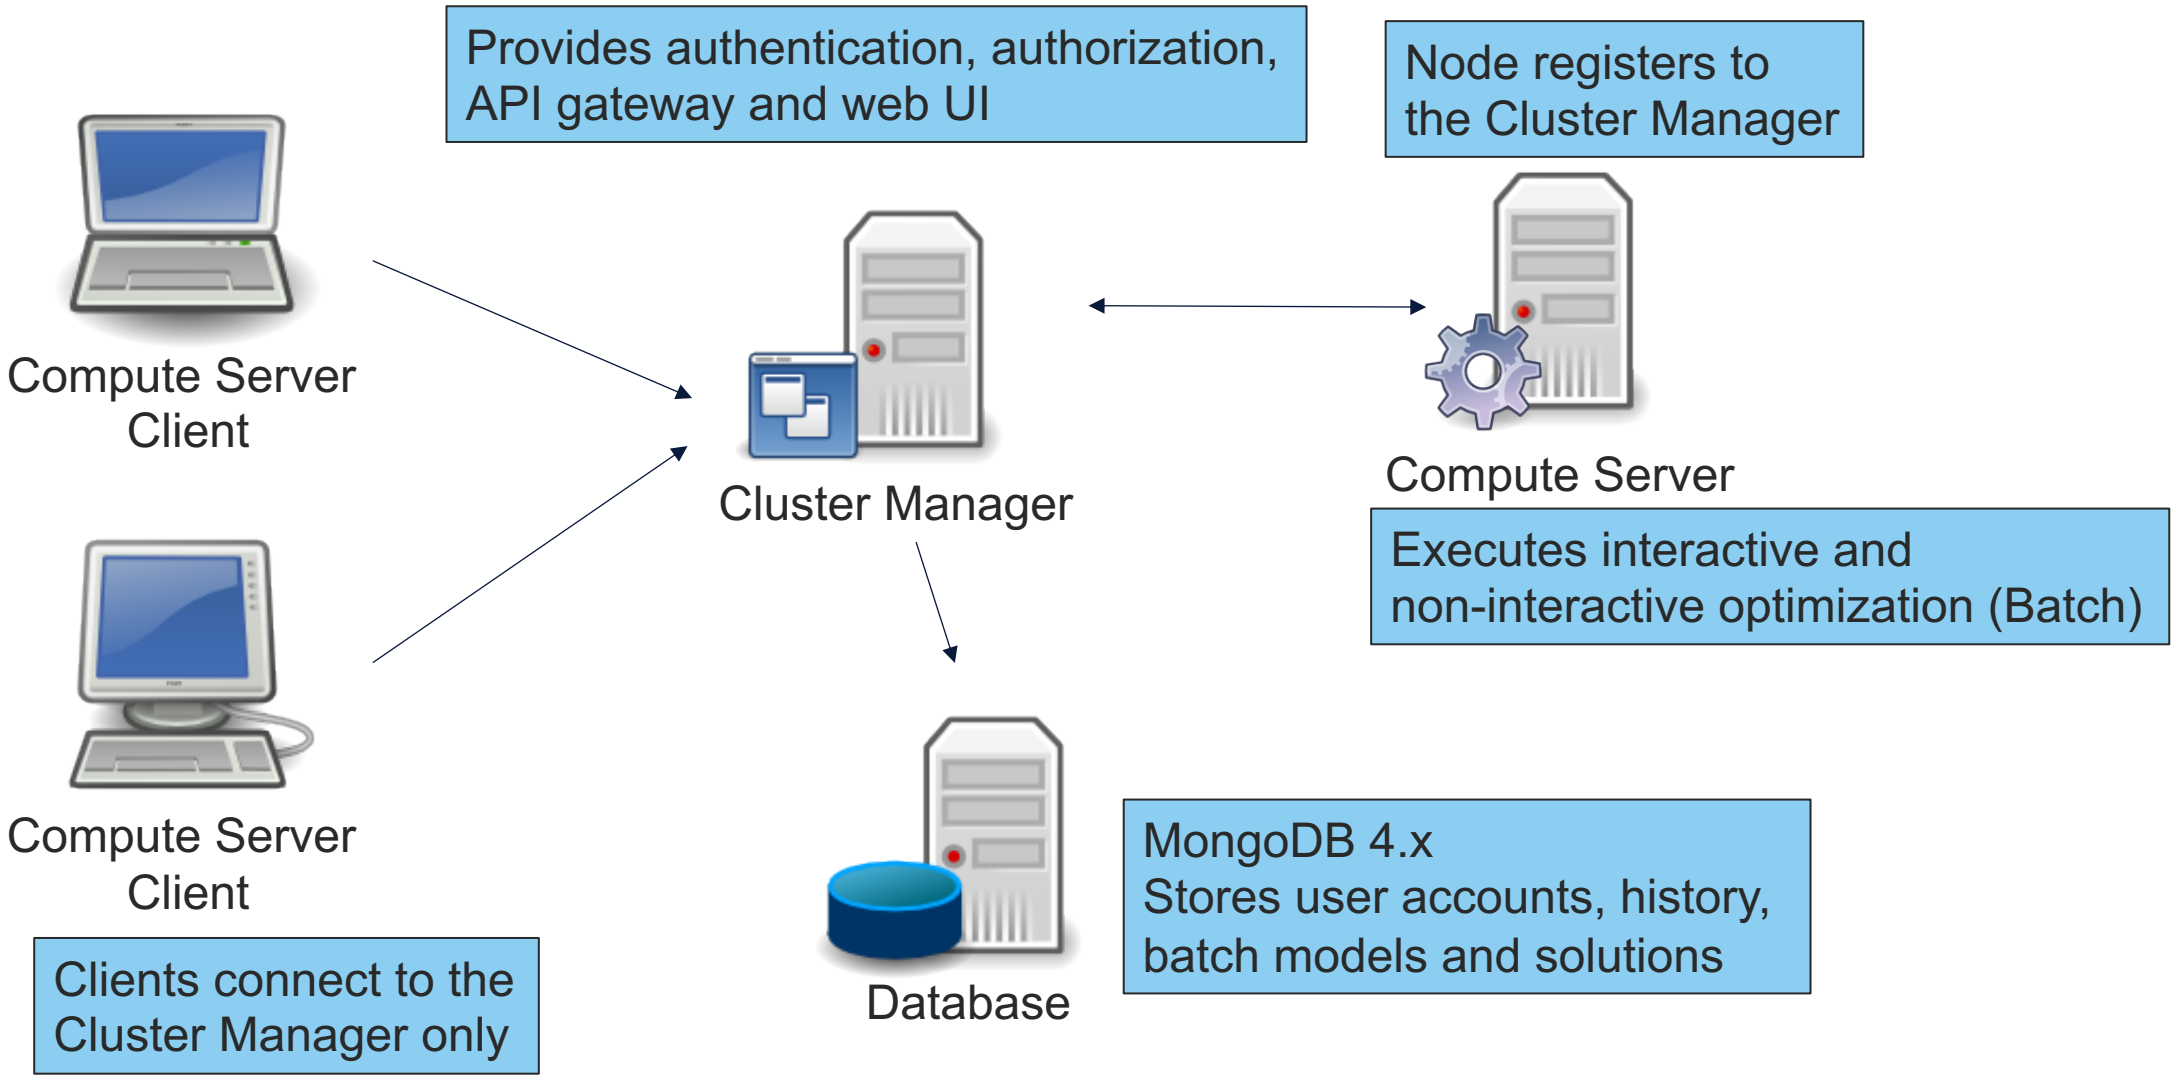

Copyright © 2020, Gurobi Optimization, LLC

Architecture: Scalable Cluster Manager (New in V9)

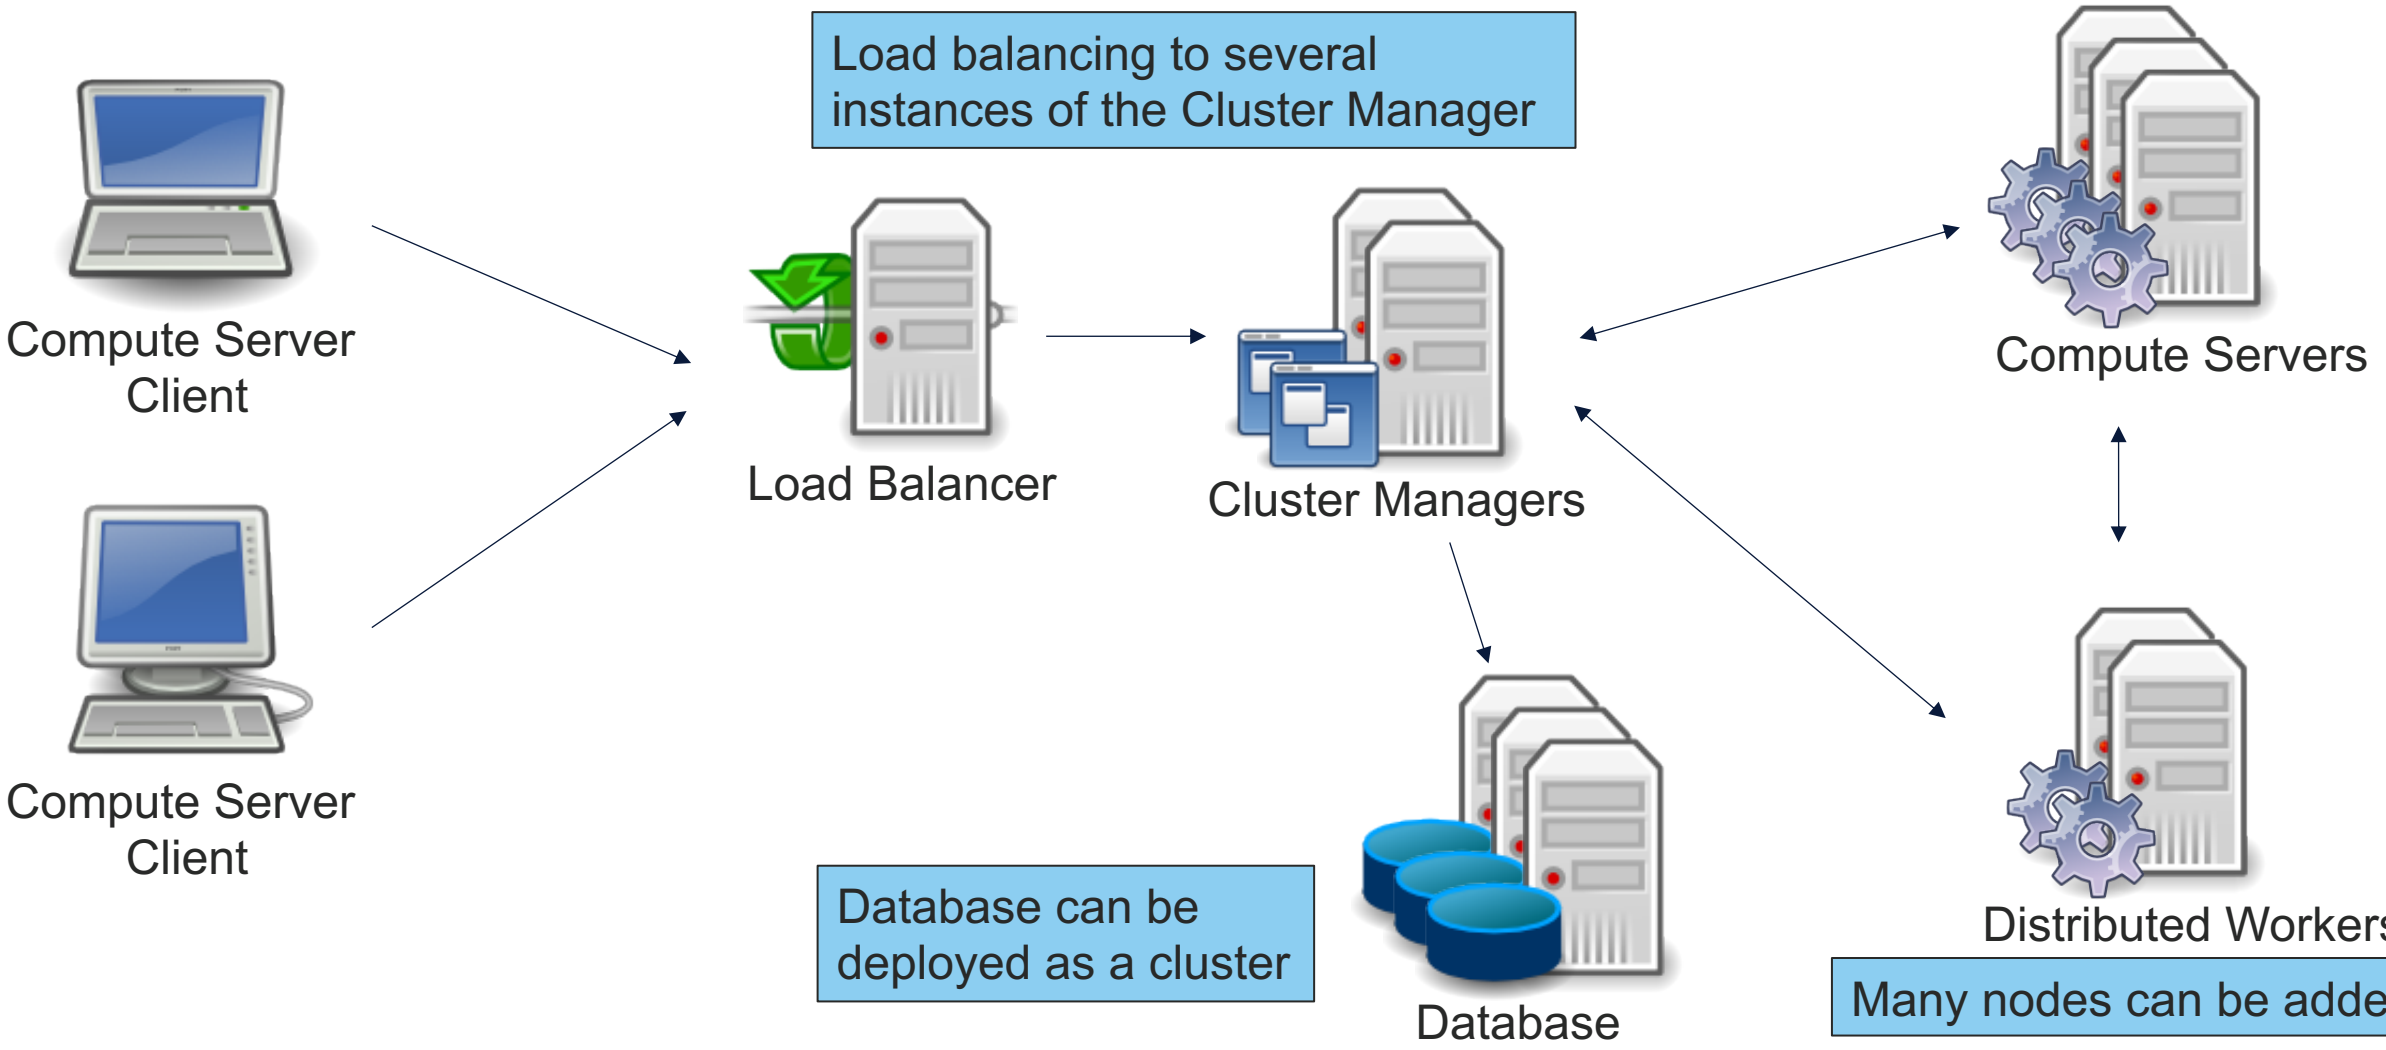

Copyright © 2020, Gurobi Optimization, LLC

Better Security: User Management

- **User Accounts**
  - Sysadmin creates and manages accounts
  - Users can provide their own passwords
  - Users must login to access the web UI
- **User Roles**
  - Standard users
  - Admin users (kill jobs, access licenses)
  - Sysadmin users
- **Default users (to delete before real deployment)**
  - gurobi/pass
  - admin/admin
  - sysadmin/cluster

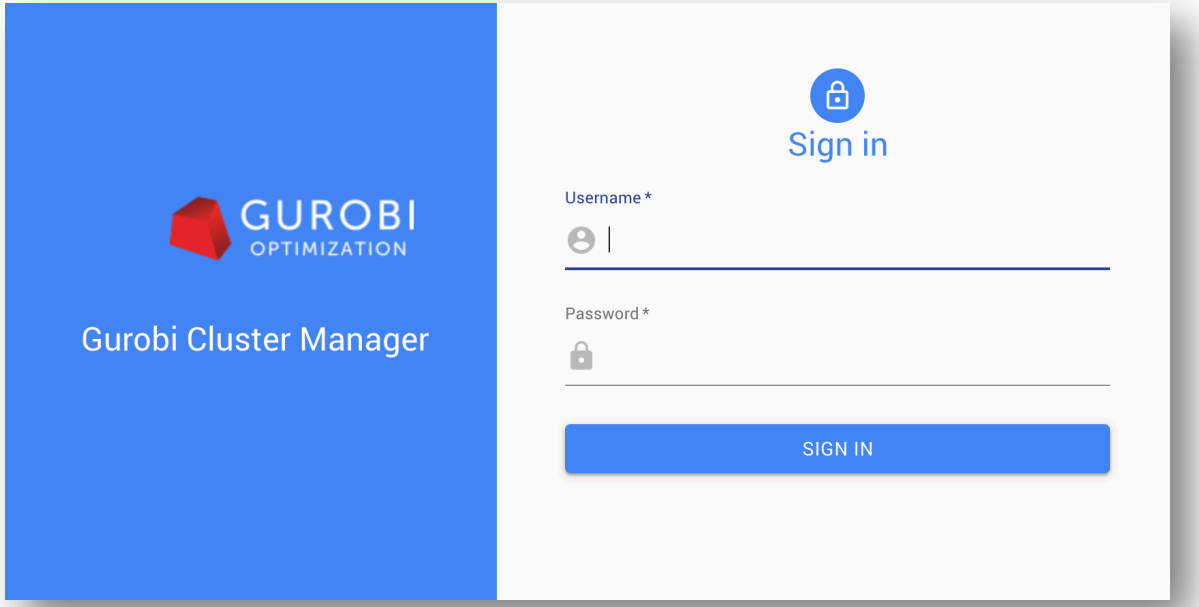

Copyright © 2020, Gurobi Optimization, LLC

Better Security: API Keys and Encryption

- **API keys for applications**
  - Composed of an authentication token and secret key
  - Each user can manage their own API keys
  - Can be embedded into applications

- **Encryption**
  - Support of HTTPS over the wire
  - End-to-end encryption or termination of TLS by the Cluster Manager

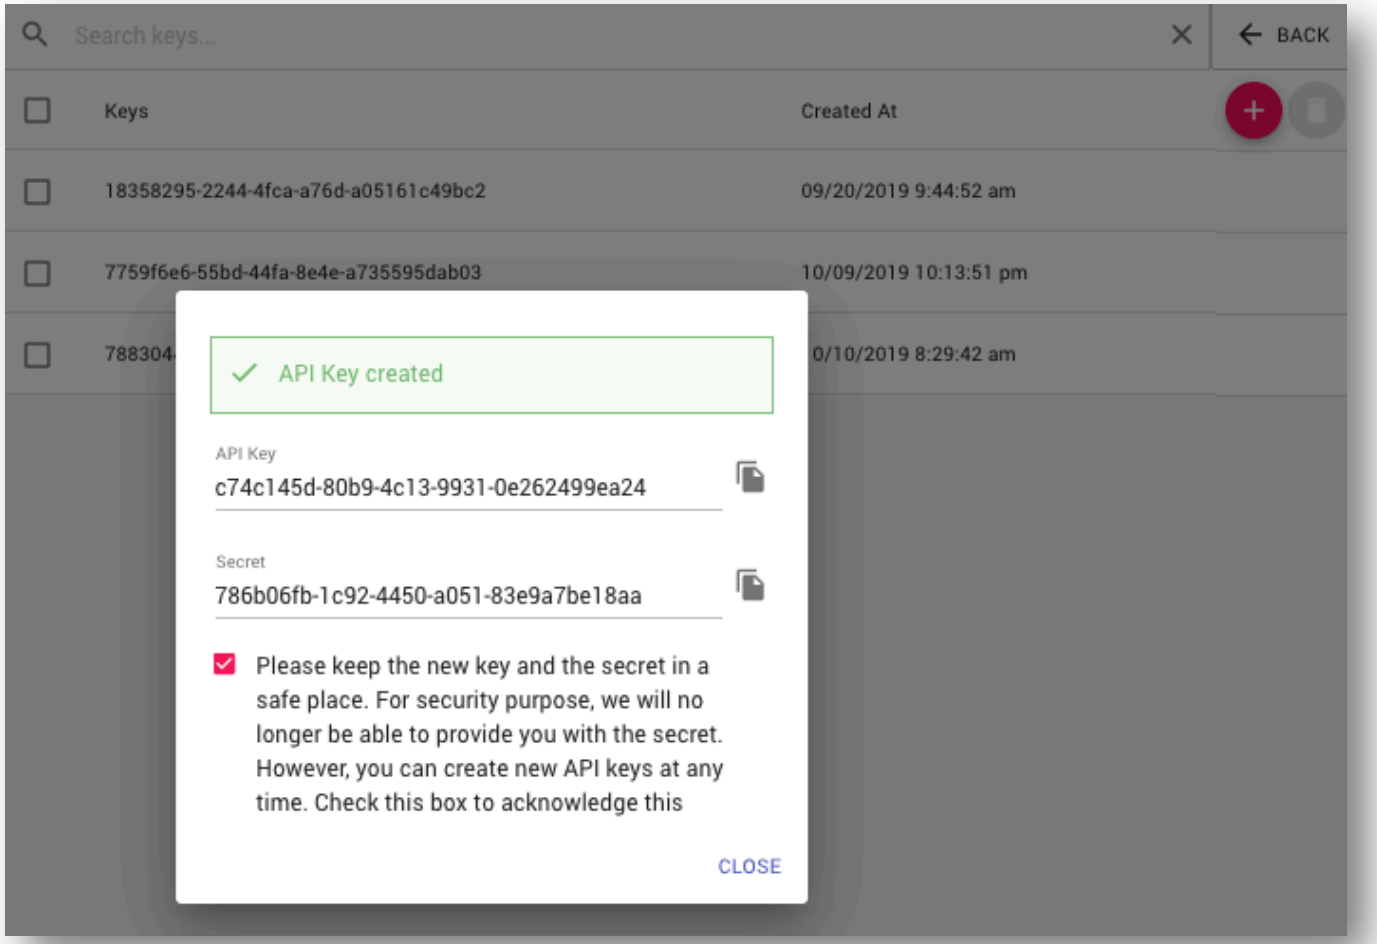

Copyright © 2020, Gurobi Optimization, LLC

Better Security: Sysadmin

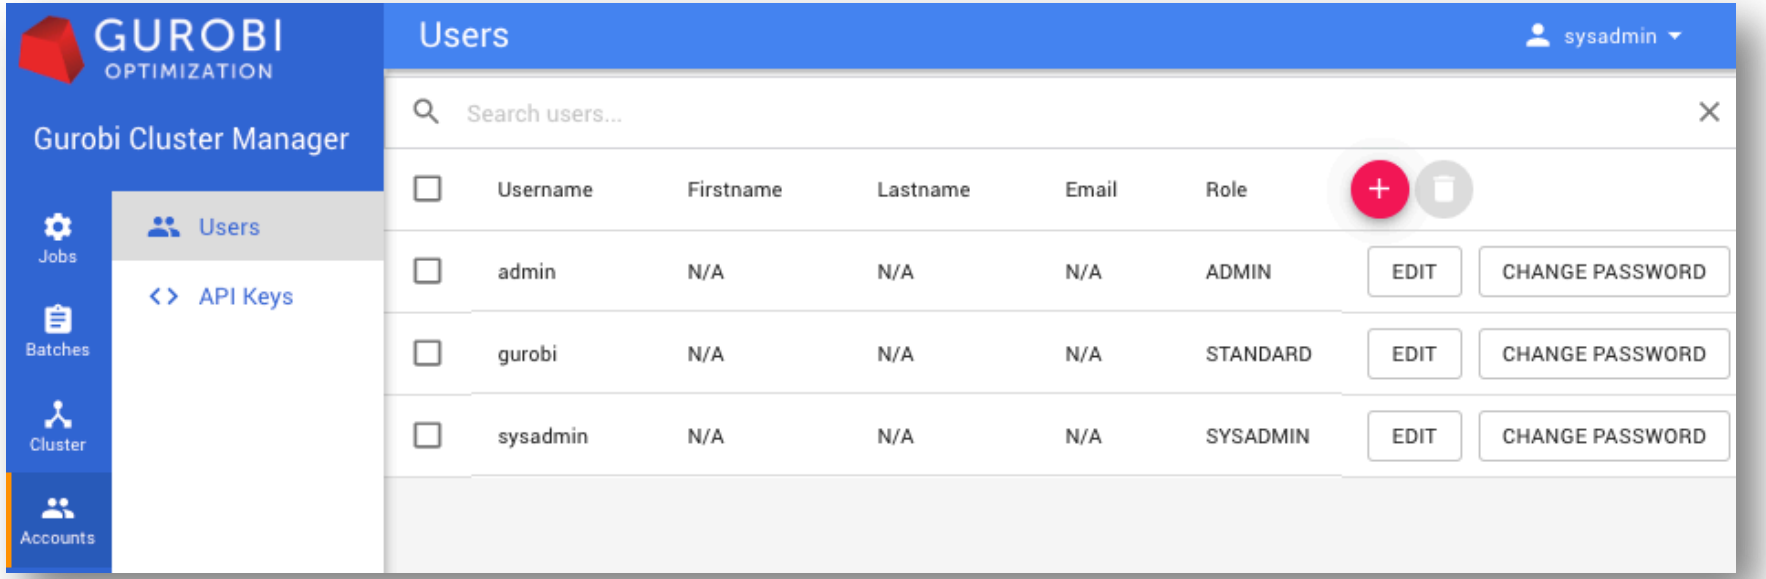

- Sysadmin can create, delete and change passwords of all users
- Sysadmin can list and delete any API key

Copyright © 2020, Gurobi Optimization, LLC

Managing the Cluster: Nodes

- List and filter nodes
- Change job limit
- Start/stop nodes

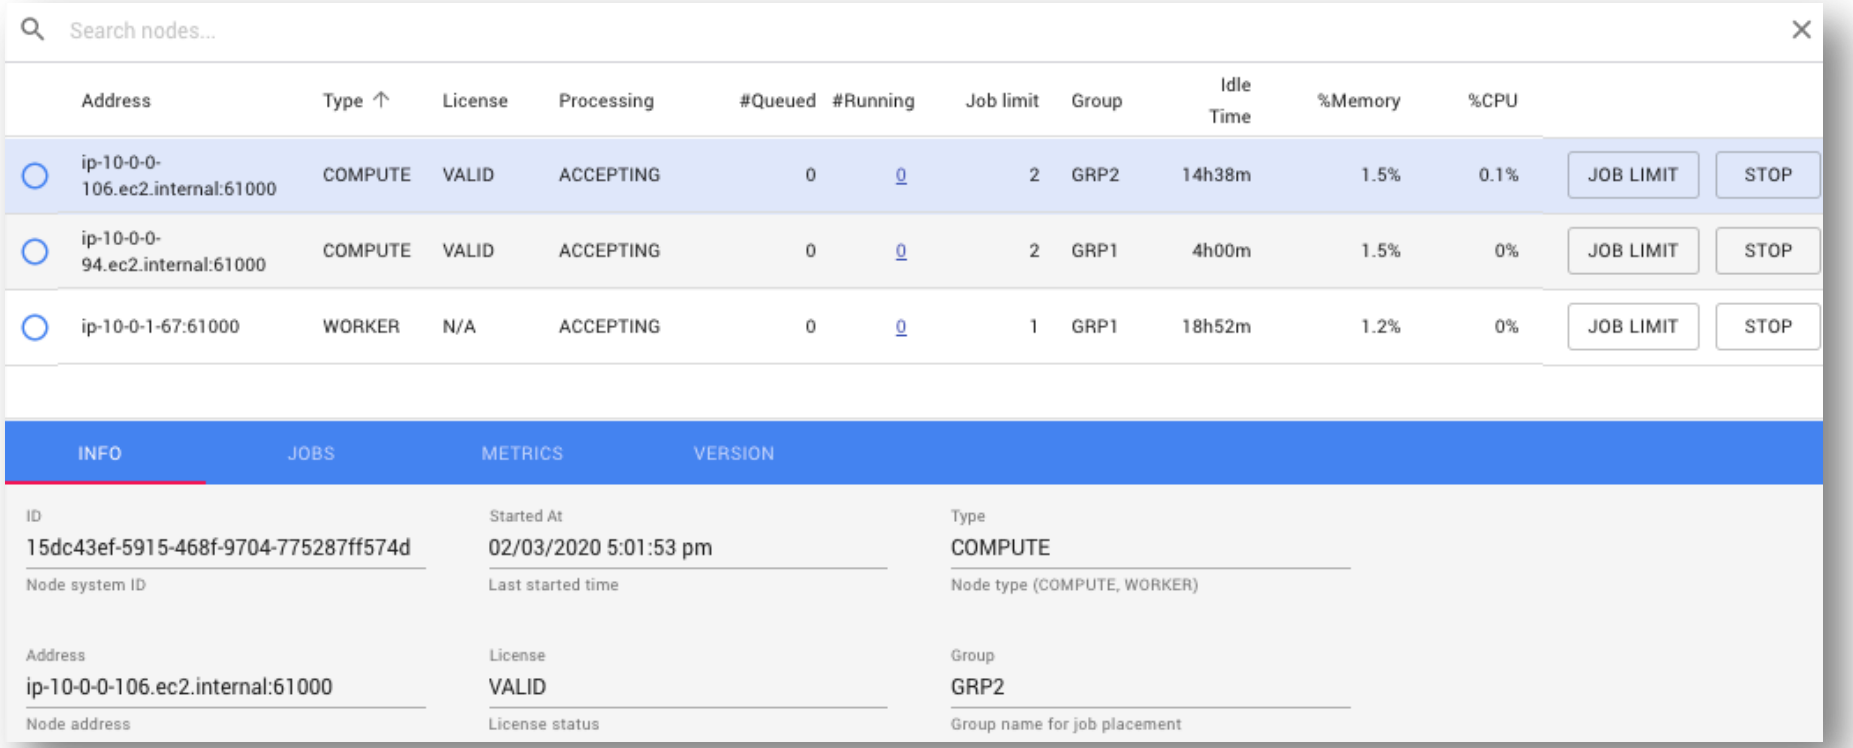

Copyright © 2020, Gurobi Optimization, LLC

Managing the Cluster: Licenses

- List node licenses
- Check expiration

- V9 adds support of token server licensing

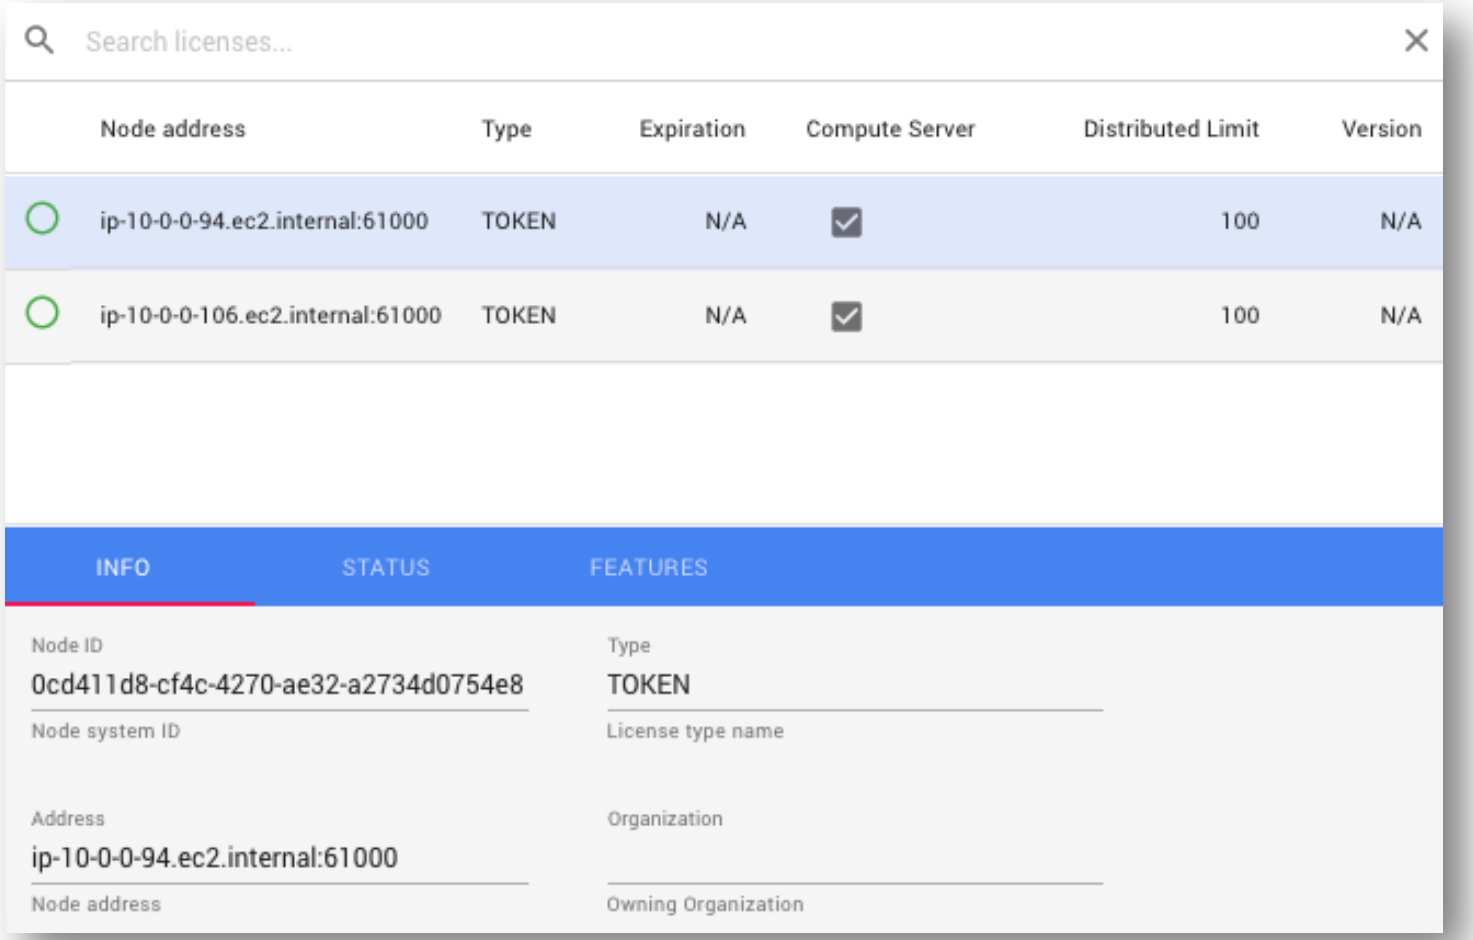

Copyright © 2020, Gurobi Optimization, LLC

Submitting Jobs: Interactive Optimization (V8 and V9)

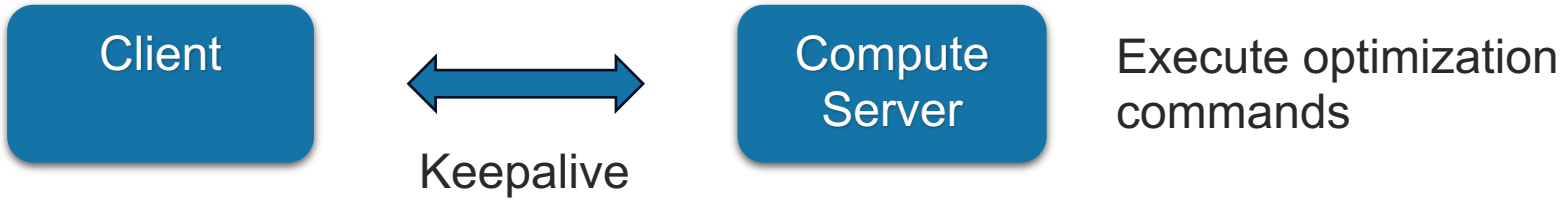

- **Interactive optimization tasks (Standard Jobs)**
  - Client always connected to server
  - Client remotely controls all optimization steps (model building, optimization)
  - Callbacks are supported

Copyright © 2020, Gurobi Optimization, LLC

Submitting Jobs: Client Configuration for Cluster Manager

GUROBI OPTIMIZATION

- Client license file has new properties
  - **CSMANAGER**: URL to the manager
  - **CSAPIACCESSID** and **CSAPISECRET**: API key and secret
  - **USERNAME** and **PASSWORD**: username/password (not recommended)
  - **CSAUTHTOKEN**: time limited token after using the login command
- New login command to update the client license file

```
grbcluster login --manager=http://localhost:61080 --username=gurobi --password=pass

gurobi_cl glass4.mps
```

- New environment parameters

```
with gp.Env(empty=True) as env:
        env.setParam('CSManager', 'http://localhost:61080')
        env.setParam('CSAPIAccessID', '0e8c35d5-ff20-4e5d-a639-10105e56b264')
        env.setParam('CSAPISecret', 'd588f010-ad47-4310-933e-1902057661c9')
        env.start()
```


Copyright © 2020, Gurobi Optimization, LLC

Submitting Jobs: Monitoring and History

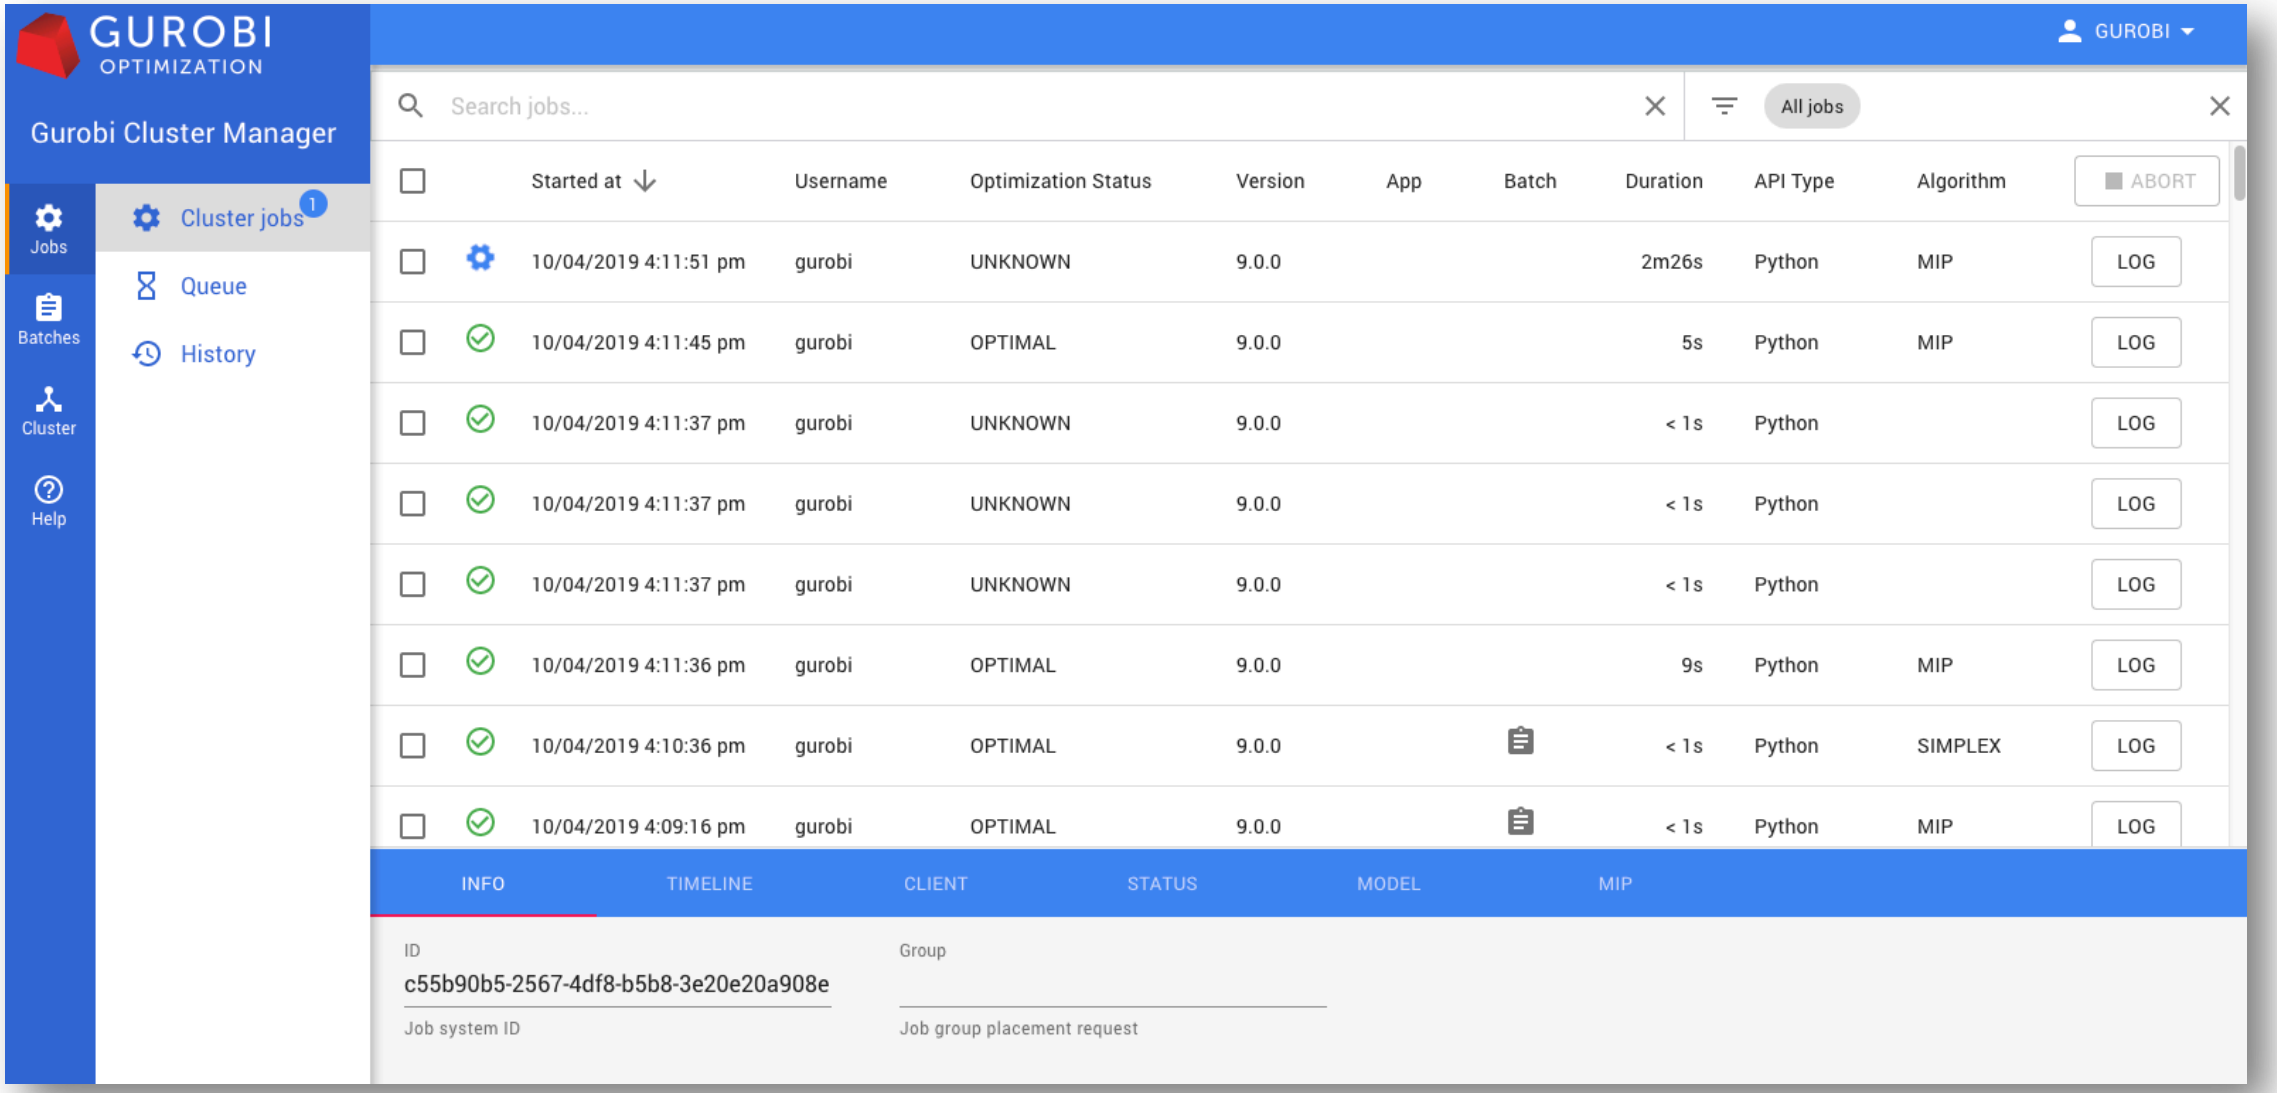

Copyright © 2020, Gurobi Optimization, LLC

Submitting Jobs: Logs

- Logs can be accessed live when the job is running
- Logs are also archived in the history

Log of job 75b3625e-0302-4e88-a610-9aea8830d0dc

DOWNLOAD

```
Compute Server job ID: 75b3625e-0302-4e88-a610-9aea8830d0dc
Optimize a model with 396 rows, 322 columns and 1815 nonzeros
Model fingerprint: 0x9be8913e
Variable types: 20 continuous, 302 integer (0 binary)
Coefficient statistics:
  Matrix range     [1e+00, 8e+06]
  Objective range  [1e+00, 1e+06]
  Bounds range     [1e+00, 8e+02]
  RHS range        [1e+00, 8e+06]
Presolve removed 4 rows and 5 columns
Presolve time: 0.01s
Presolved: 392 rows, 317 columns, 1815 nonzeros
Variable types: 19 continuous, 298 integer (298 binary)
Found heuristic solution: objective 3.133356e+09

Root relaxation: objective 8.000024e+08, 72 iterations, 0.00 seconds

    Nodes    |    Current Node    |     Objective Bounds      |     Work
 Expl Unexpl |  Obj  Depth IntInf | Incumbent    BestBd   Gap | It/Node Time

     0     0 8.0000e+08    0   72 3.1334e+09 8.0000e+08  74.5%     -    0s
H    0     0                     2.400019e+09 8.0000e+08  66.7%     -    0s
H    0     0                     2.220019e+09 8.0000e+08  64.0%     -    0s
     0     0 8.0000e+08    0   72 2.2200e+09 8.0000e+08  64.0%     -    0s
H    0     0                     2.200019e+09 8.0000e+08  63.6%     -    0s
     0     0 8.0000e+08    0   81 2.2000e+09 8.0000e+08  63.6%     -    0s
     0     0 8.0000e+08    0   77 2.2000e+09 8.0000e+08  63.6%     -    0s
     0     2 8.0000e+08    0   77 2.2000e+09 8.0000e+08  63.6%     -    0s
H  307   609                     2.066686e+09 8.0000e+08  61.3%   5.8    0s
```


Copyright © 2020, Gurobi Optimization, LLC

Batch Optimization: Non-Interactive Tasks (New in V9)

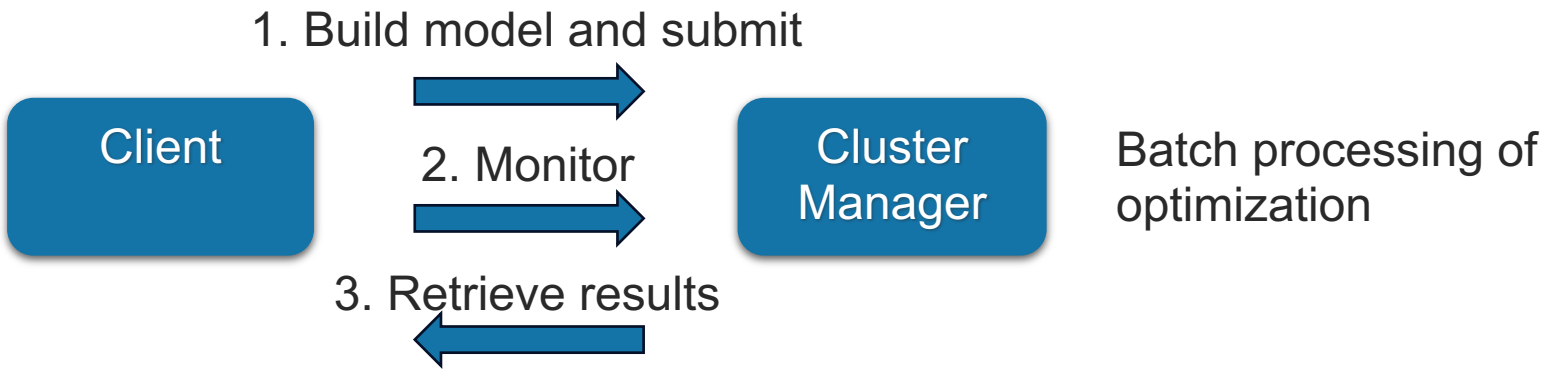

- **Three steps:**
  - Build locally and submit the model, a batch ID is returned
  - Monitor using the batch ID
  - Retrieve the results as a JSON file
- **Benefits:**
  - Clients can disconnect in between steps
  - No client resources are used once the model has been submitted
  - Results can be retrieved by a different service of the client application

Copyright © 2020, Gurobi Optimization, LLC

Batch Optimization: Architecture

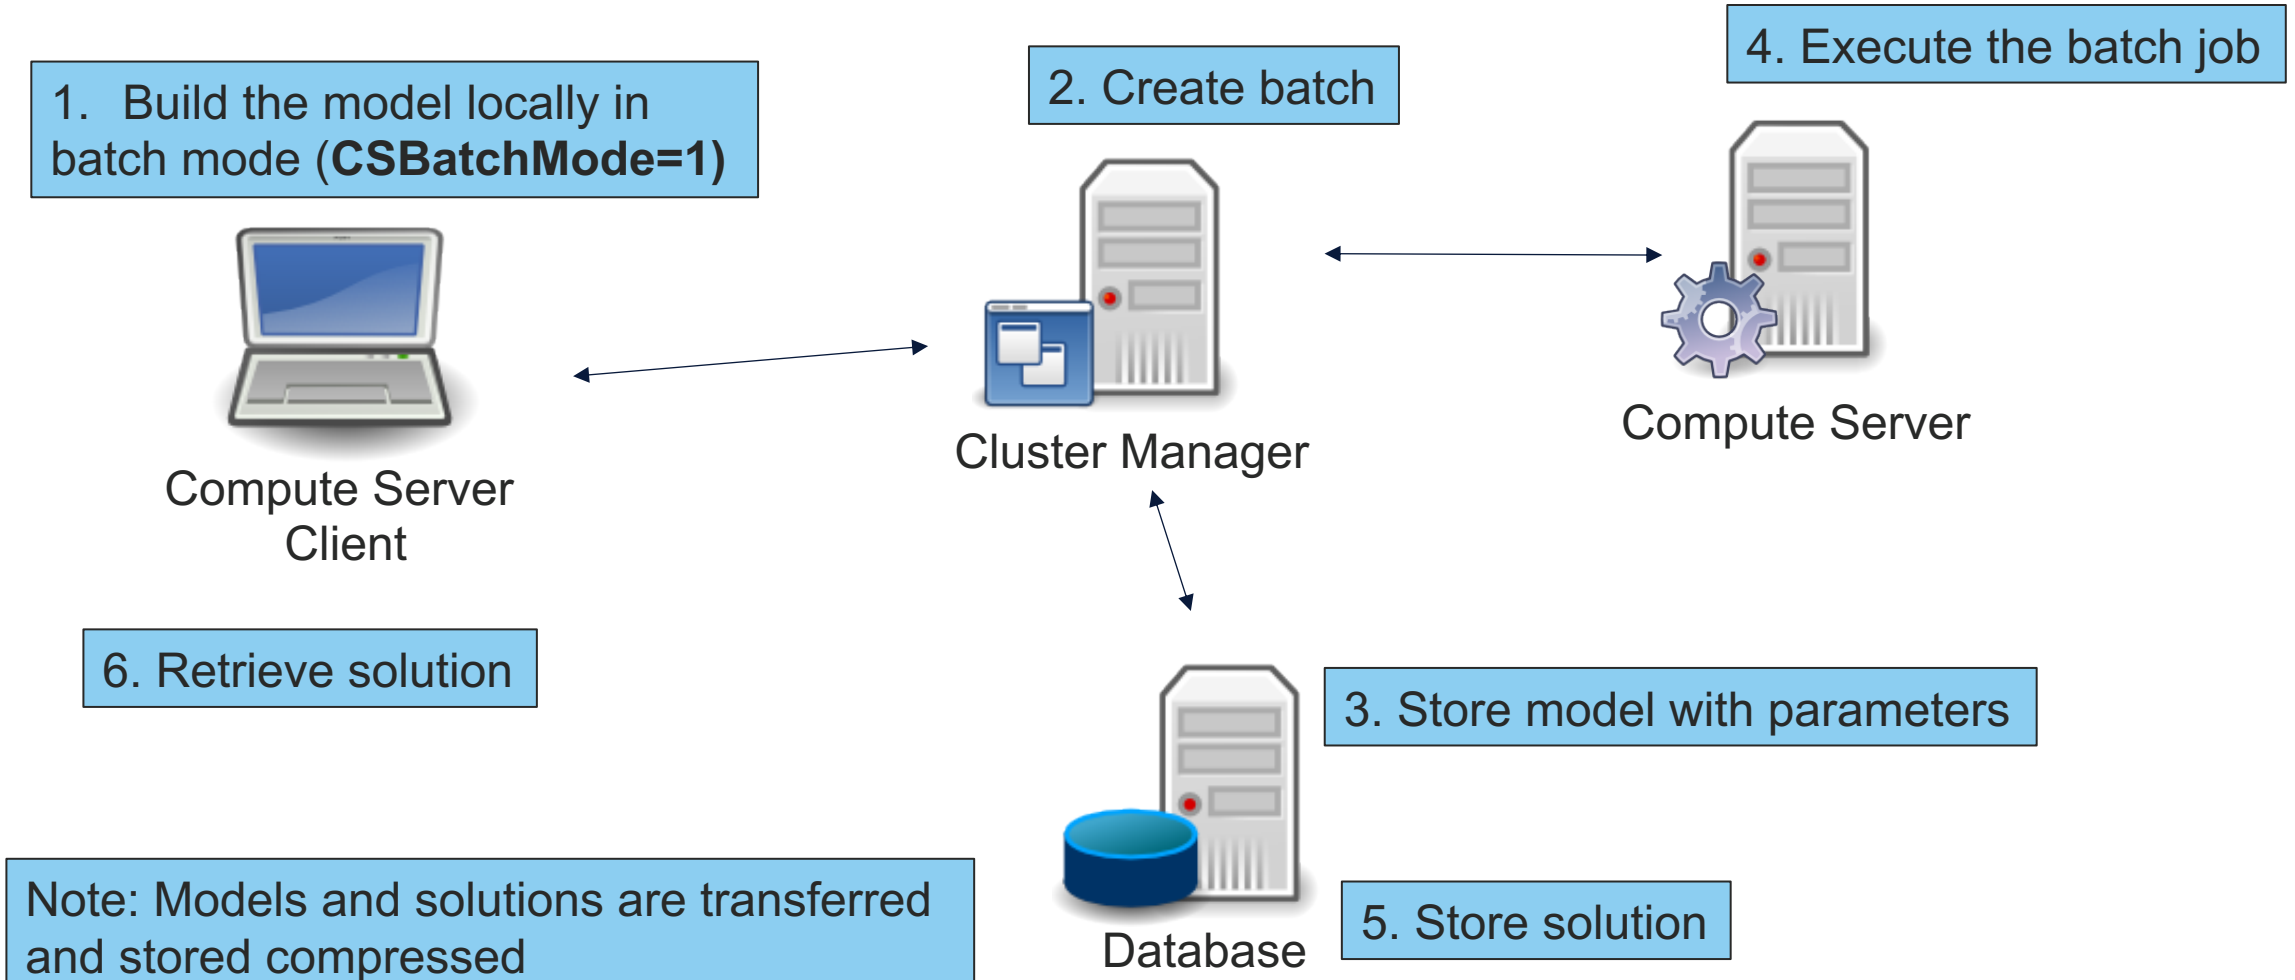

Copyright © 2020, Gurobi Optimization, LLC

Batch Optimization: Variable and Constraint Tagging

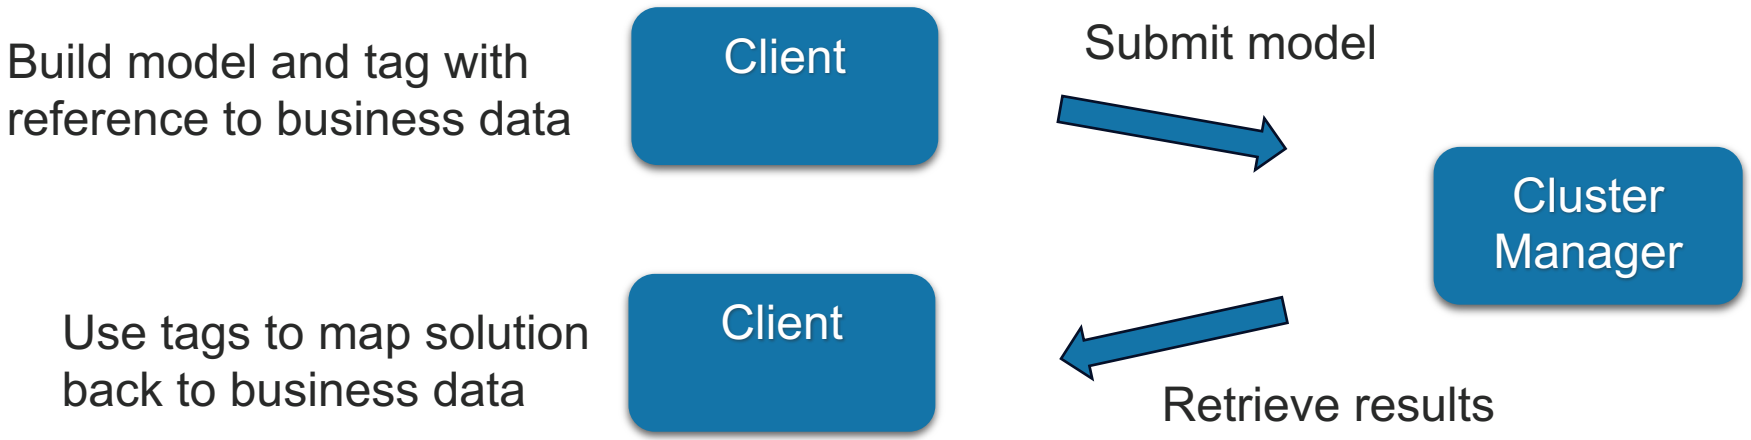

- **Benefits of tagging**
  - Solution only contains the tagged variables and constraints
  - Solution can be interpreted standalone, without the submitted model context

Copyright © 2020, Gurobi Optimization, LLC

Batch Optimization: Model Building Example

GUROBI OPTIMIZATION

- Create an environment connected to a Cluster Manager
- Set CSBatchMode=1
- Tag relevant variables and constraints with custom data
- Call model.optimizeBatch()
- Batch ID is returned

```
with gp.Model("assignment", env=env) as m:
    # Assignment variables: x[w,s] == 1 if worker w is assigned to shift s.
    x = m.addVars(availability, vtype=GRB.BINARY, name="x")

    # Set tags encoding the assignments for later retrieval of the schedule.
    # Each tag is a JSON string of the format
    #   {
    #     "Worker": "<Name of the worker>",
    #     "Shift":  "String representation of the shift"
    #   }
    #
    for k, v in x.items():
        name, timeslot = k
        d = {"Worker": name, "Shift": shiftname[timeslot]}
        v.VTag = json.dumps(d)
```


Copyright © 2020, Gurobi Optimization, LLC

Batch Optimization: Retrieving the Solution

GUROBI OPTIMIZATION

- Monitor the batch status
- When completed, get the solution file
- JSON format for easy integration
- Only tagged elements will be exported
- More details can be exported with JSONSolDetail parameter
- Extract business solution using variable value and related tags


```
{
  "SolutionInfo": {
    ...
  },
  "Vars": [
    {
      "VTag": [
        "{\"Shift\": \"Wednesday  8:00\", \"Worker\": \"Amy\"}"
      ],
      "X": "1"
    },
    {
      "VTag": [
        "{\"Shift\": \"Thursday  8:00\", \"Worker\": \"Amy\"}"
      ],
      "X": "1"
    },
...
```


Copyright © 2020, Gurobi Optimization, LLC

Batch Optimization: UI and Command Line

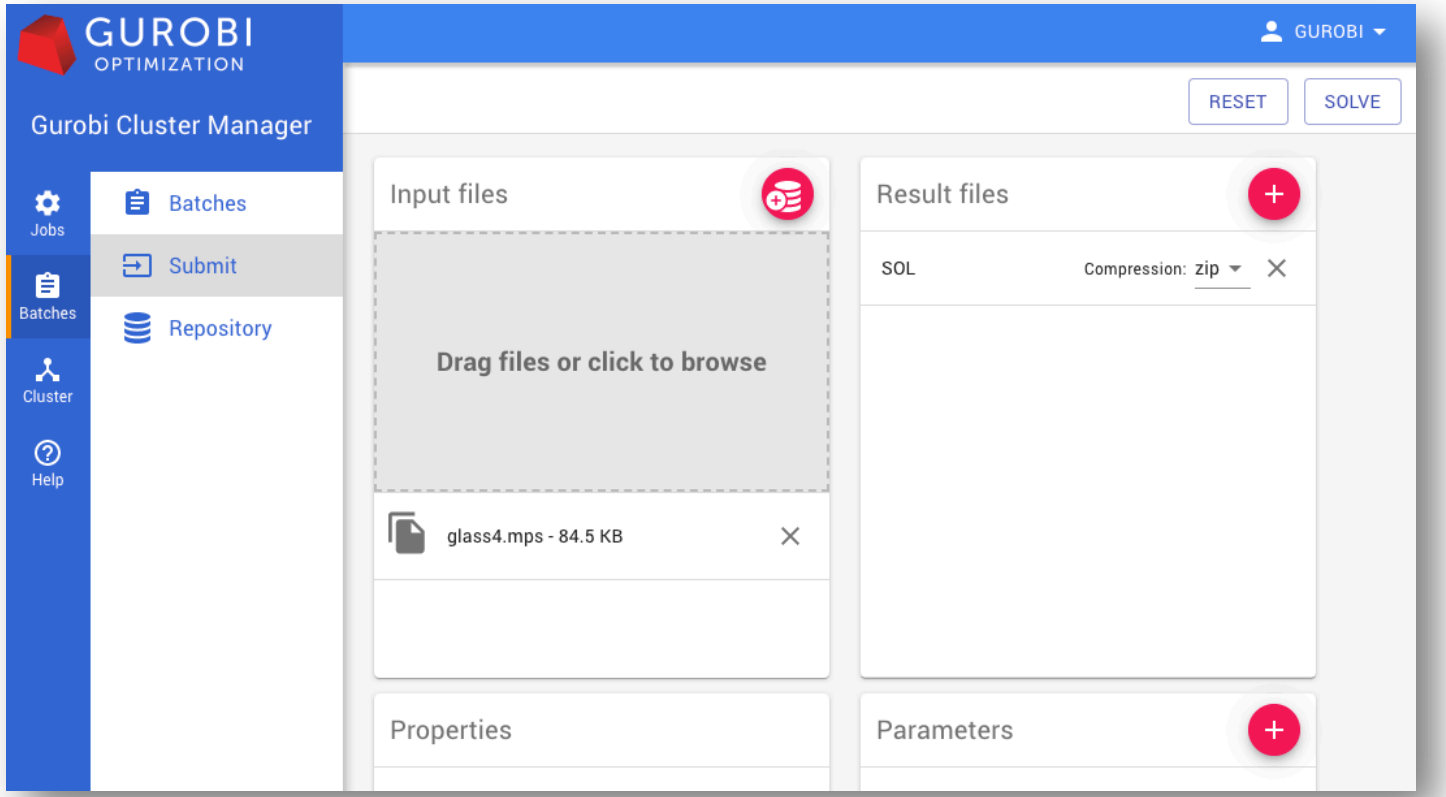


```
$ grbcluster batch solve glass4.mps ResultFile=glass4.sol
info  : Batch 10230a87-8ab9-4277-8c00-e949627516cb created
info  : Uploading glass4.mps...
info  : Batch 10230a87-8ab9-4277-8c00-e949627516cb submitted…
```


Copyright © 2020, Gurobi Optimization, LLC

Batch Optimization: Monitoring

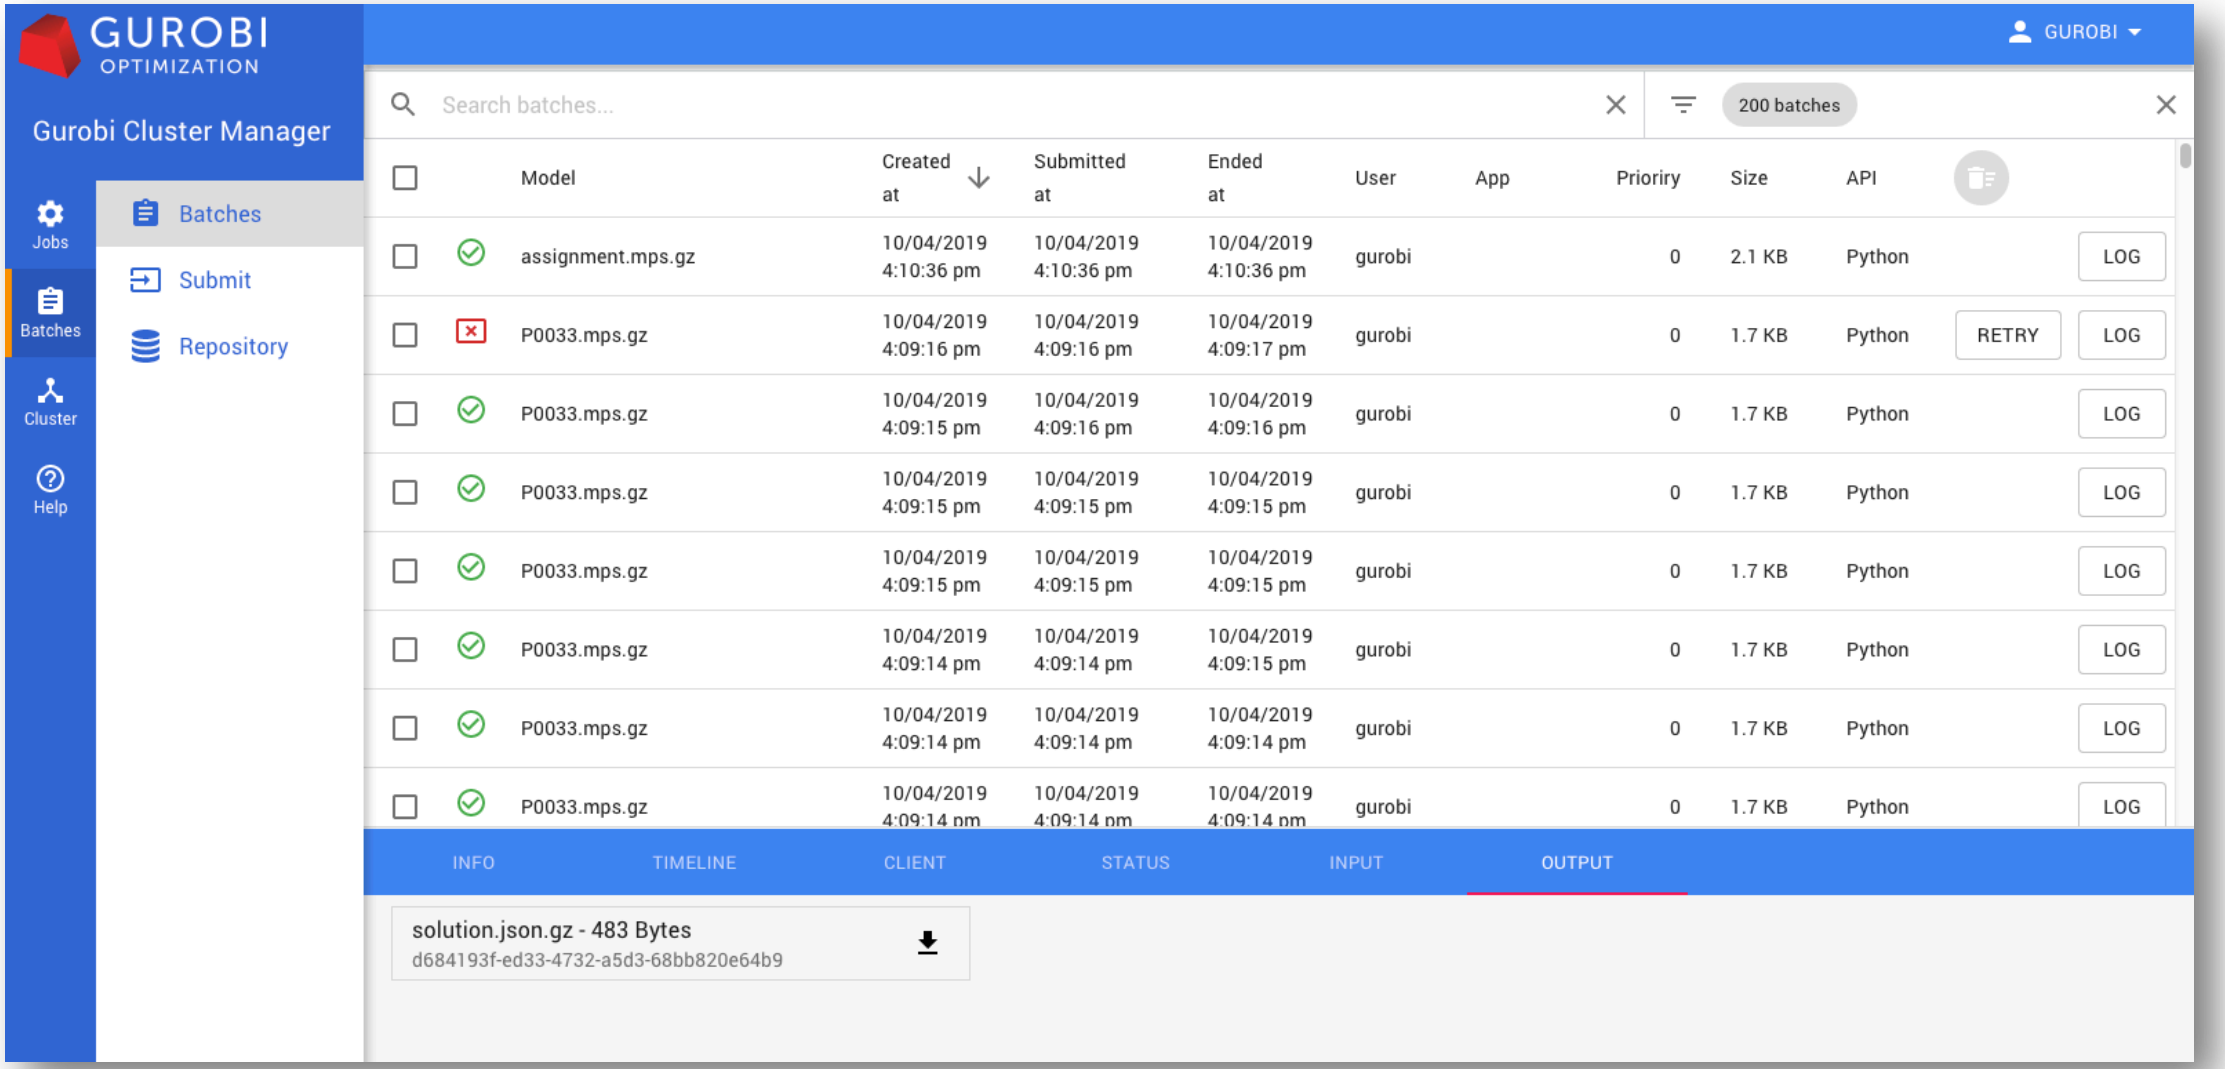

Copyright © 2020, Gurobi Optimization, LLC

Batch Optimization: File Repository

- Share models or parameters with your team
- Avoid uploading again the same files

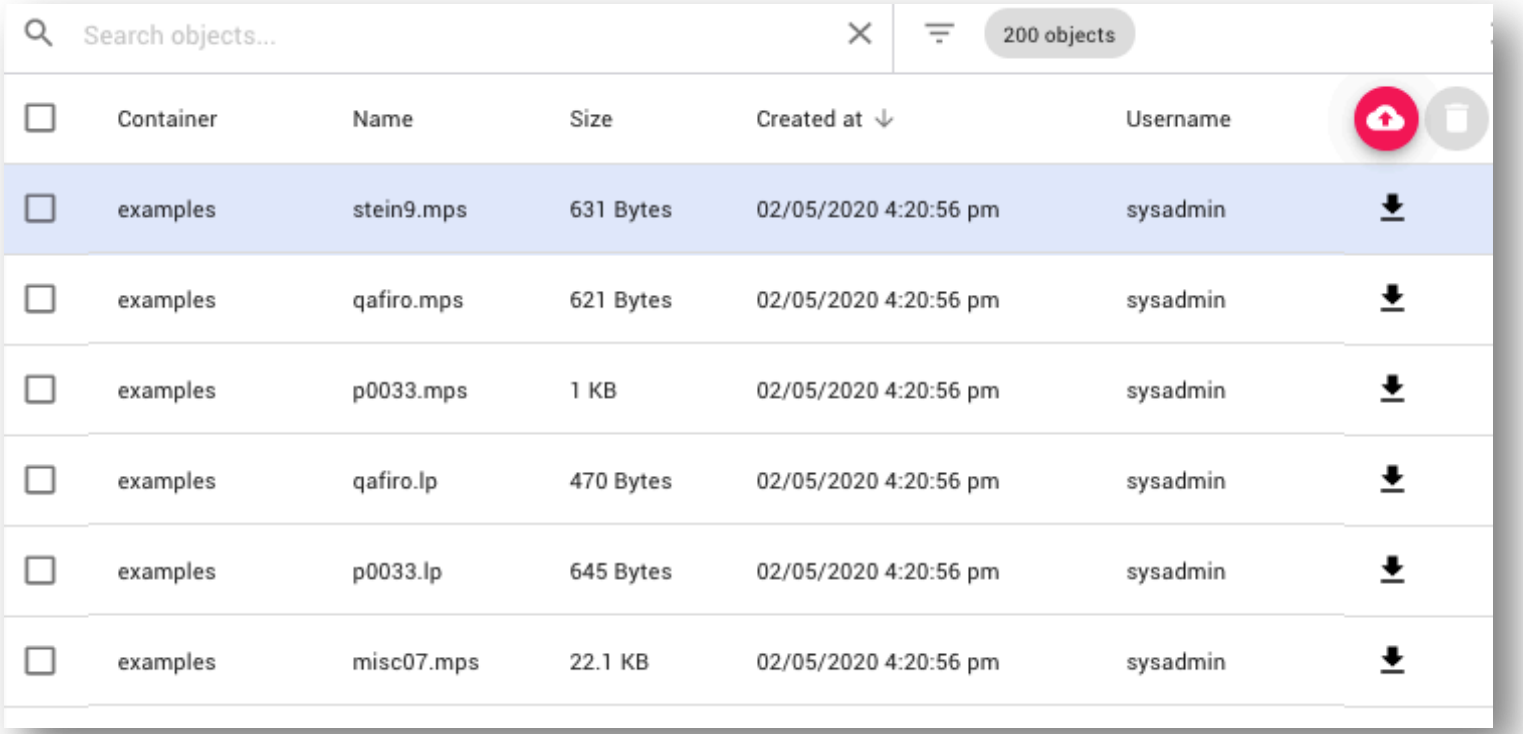

Copyright © 2020, Gurobi Optimization, LLC

REST API

- Can be used to integrate and automate cluster monitoring and management
- Covers all the entities
  - Users
  - API Keys
  - Cluster nodes
  - Jobs and history
  - Model and solution files
  - Batches

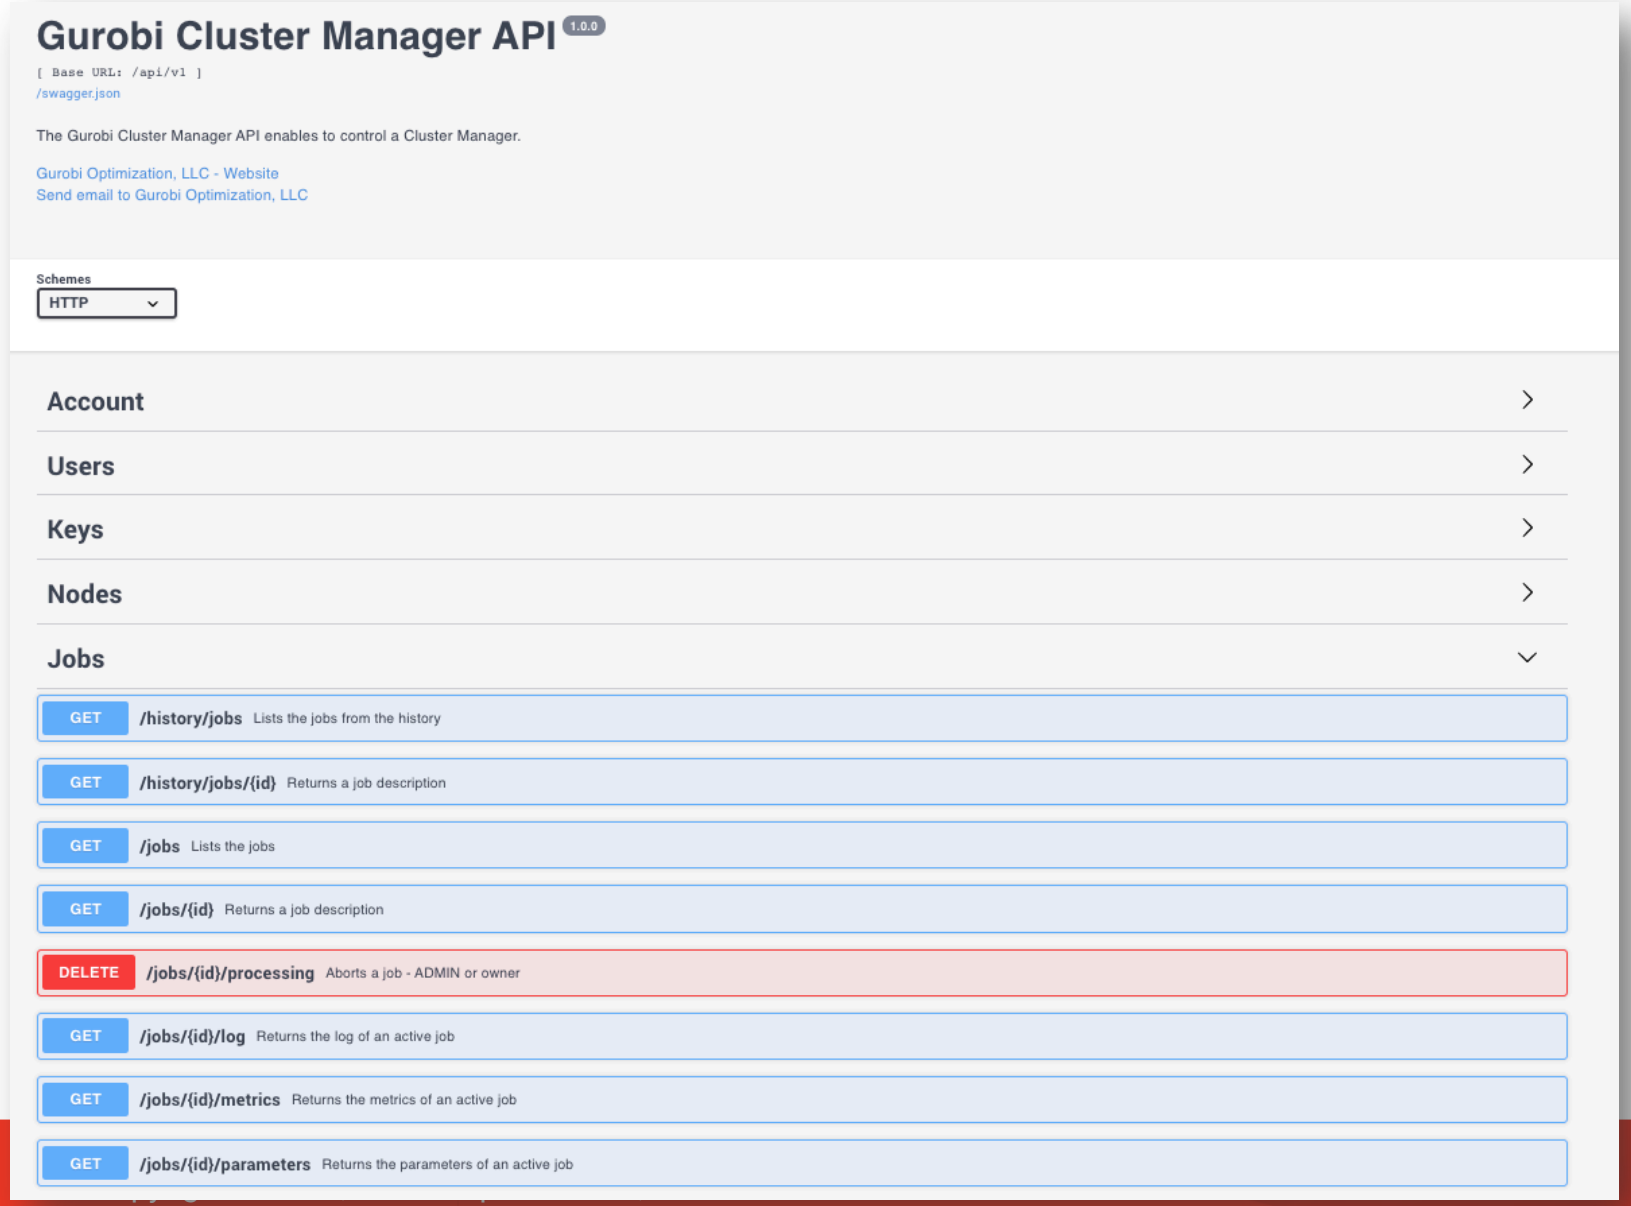

Demo

GUROBI OPTIMIZATION

- **Installation**
  - Start MongoDB and manager
  - Start two local nodes
  - Display nodes and licenses
- **Submitting Jobs**
  - Login as a user on the command line
  - Submit a job
  - Display the job info in the Cluster Manager
- **Submitting Batches**
  - Run workforce batch example
  - Submit a batch from the UI
  - Submit a batch from the command line

Copyright © 2020, Gurobi Optimization, LLC

Thank you

GUROBI
OPTIMIZATION

The World's Fastest Solver