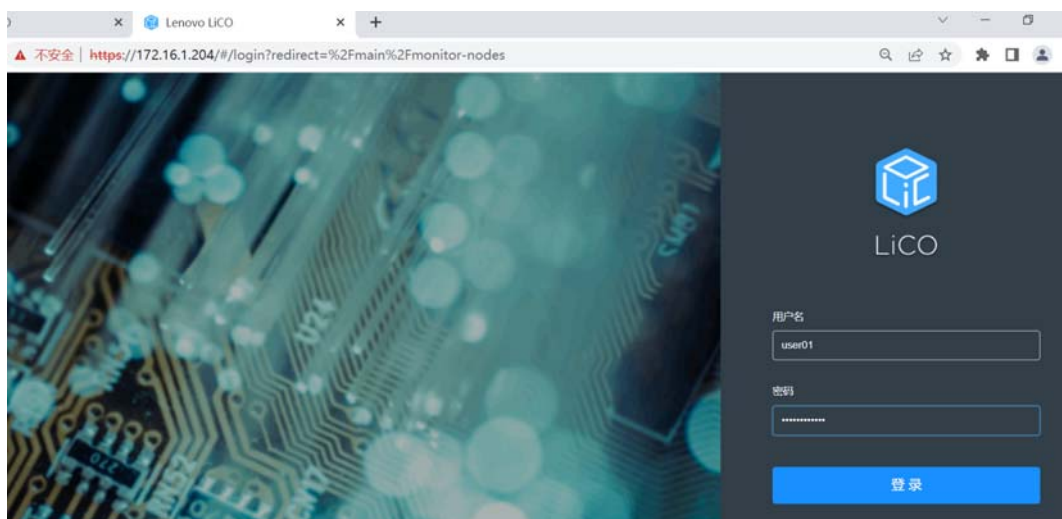


分析测试与计算中心超算平台使用手册

一. 集群登录

平台 2 台登录节点安装 portal 服务，使用账户登录集群操作界面，或使用 SSH 工具连接登录节点，进行提交作业。详细操作请参考《LiCO 6.4.0 User Guide_20220811_zh-cn》。

浏览器打开网址 <https://219.217.229.12:443> 和 <https://219.217.229.74:443> 。如下



二. Slurm 作业调度系统

[SLURM](#)（Simple Linux Utility for Resource Management）是一种可扩展的工作负载管理器，已被全世界的国家超级计算机中心广泛采用。它是免费且开源的，根据 GPL 通用公共许可证发行。下表简要列出 Slurm 常用命令，如需详细命令参数请登录 Slurm 官方网站 <https://slurm.schedmd.com/>进行查阅。

Slurm 概览

Slurm	功能
sinfo	集群状态
squeue	排队作业状态
sbatch	作业提交
scontrol	查看和修改作业参数

Slurm	功能
sacct	已完成作业报告
scancel	删除作业

Slurm 环境变量

Slurm	功能
\$SLURM_JOB_ID	作业 ID
\$SLURM_JOB_NAME	作业名
\$SLURM_JOB_PARTITION	队列的名称
\$SLURM_NTASKS	进程总数
\$SLURM_NTASKS_PER_NODE	每个节点请求的任务数
\$SLURM_JOB_NUM_NODES	节点数
\$SLURM_JOB_NODELIST	节点列表
\$SLURM_LOCALID	作业中流程的节点本地任务 ID
\$SLURM_ARRAY_TASK_ID	作业序列中的任务 ID
\$SLURM_SUBMIT_DIR	工作目录
\$SLURM_SUBMIT_HOST	提交作业的主机名

三. LAMMPS 应用

简介

LAMMPS 是大规模原子分子并行计算代码，在原子、分子及介观体系计算中均有重要应用，并行效率高，广泛应用于材料、物理、化学等模拟。

编译 LAMMPS

1、下载 LAMMPS

从以下官网链接中下载最新版本软件。

<https://www.lammps.org/download.html>

2、配置好 INTEL 编译器环境。

```
source /share/software/intel/oneapi/setvars.sh
```

```
[root@cn001 build]# which icc
```

```
/share/software/intel/oneapi/compiler/2022.0.2/linux/bin/intel64/icc
```

3、解压源文件包，并执行编译，生成 lmp 可执行程序。

```
tar xvf lammps.xxxx
```

```
cd lammps
```

```
mkdir build
```

```
cd build
```

```
cmake -DCMAKE_INSTALL_PREFIX=/share/software/lammps -DBUILD_MPI=yes -
```

```
DBUILD_OMP=yes ../cmake
```

```
make
```

```
make install
```

4、单机和多机测试

```
mpirun -np 4 lmp < in.deposit.atom
```

```
mpirun -np 4 -host cn001,cn002 -ppn 2 lmp < in.deposit.atom
```

slurm 作业文件示例

此文件位于

```
/share/software/demo/lammps/
```

需要将此文件夹和对应的输入文件复制到 home 目录后提交作业.

```
cp /share/software/demo/lammps/{in.deposit.atom,lammps.slurm} ~/
```

```
sbatch lammps.slurm
```

源文件如下：

```
#!/bin/bash

#SBATCH --output=%j.out

#SBATCH --error=%j.err

#SBATCH --partition=cpu

#SBATCH --job-name= lammmps_jobname

#SBATCH --nodes=4

#SBATCH --ntasks-per-node=56

#SBATCH --get-user-env

#SBATCH --time=0-24:00

ulimit -s unlimited

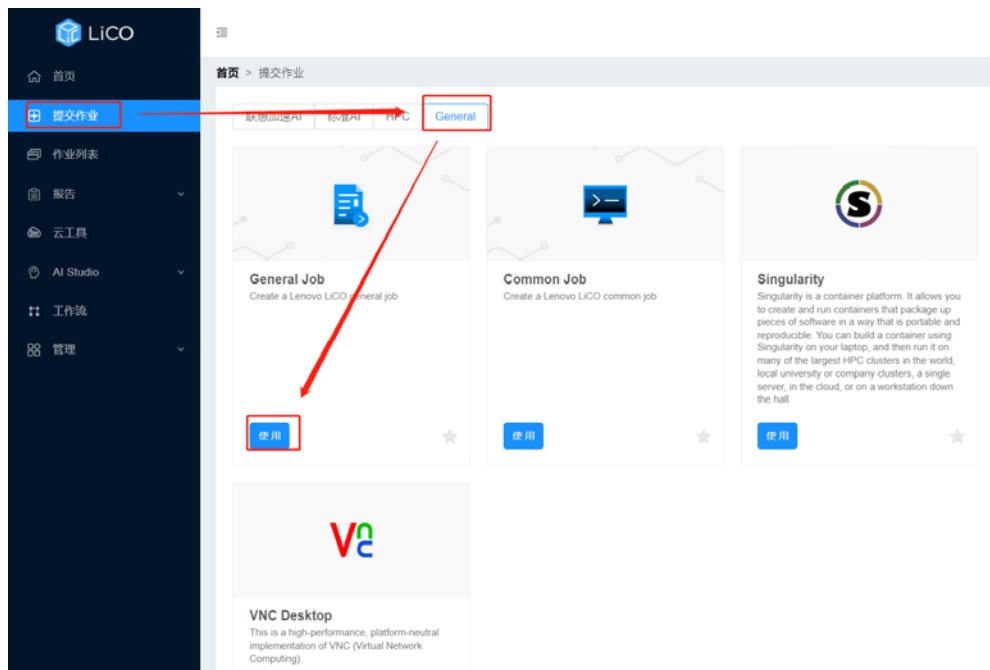
ulimit -l unlimited

source /home/software/intel/oneapi/setvars.sh

mpirun /share/software/lammps/bin/lmp -pk intel 0 omp 2 -sf intel -i in.deposit.atom
```

Lico 提交作业示例

从“提交作业” → “General” → “使用”



输入“作业名称”及作业文件后，进行提交。

首页 > 提交作业 > General Job

General Job

Create a Lenovo LiCO general job

模板信息 ▾

* 作业名称

模板参数 ▾

* 作业文件

通知选项 >

四. Intel MPI 作业提交

Slurm 作业文件示例

此脚本位于/share/software/demo/intel_mpi/下，需要将此文件复制到 home 目录下然后提交作业。

```
sbatch intel_mpi.slurm
```

源文件如下：

```
#!/bin/bash
#SBATCH --output=%j.out
#SBATCH --error=%j.err
#SBATCH --partition=cpu
#SBATCH --job-name=mpi_jobname
#SBATCH --nodes=4
```

```
#SBATCH --ntasks-per-node=56

#SBATCH --get-user-env

#SBATCH --time=0-24:00

ulimit -s unlimited

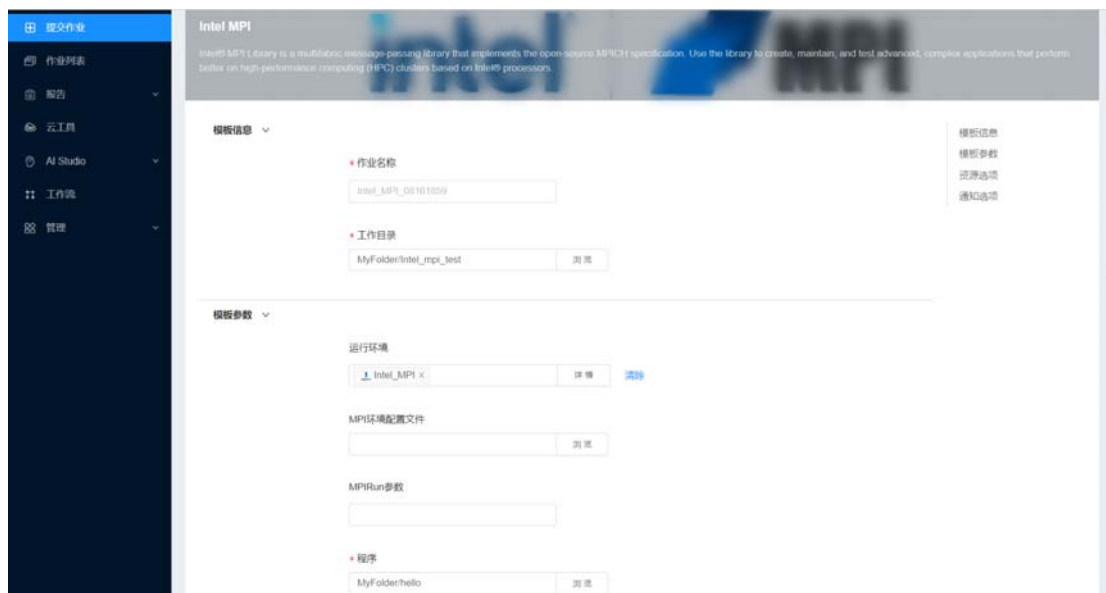
ulimit -l unlimited

source /share/software/intel/oneapi/setvars.sh

mpirun -n 64 ./hello
```

Lico 提交作业示例

选择作业相关信息



The screenshot shows the Lico web interface for submitting an Intel MPI job. The interface is divided into a sidebar and a main content area. The sidebar on the left contains navigation options: 提交作业 (Submit Job), 作业列表 (Job List), 报告 (Reports), 云工具 (Cloud Tools), AI Studio, 工作流 (Workflows), and 管理 (Management). The main content area is titled "Intel MPI" and contains a form for job configuration. The form is organized into two sections: "模板信息" (Template Information) and "模板参数" (Template Parameters). The "模板信息" section includes fields for "作业名称" (Job Name) with the value "intel_mpi_00101050" and "工作目录" (Working Directory) with the value "MyFolder/intel_mpi_test". The "模板参数" section includes fields for "运行环境" (Runtime Environment) with the value "intel_mpi_x", "MPI环境配置文件" (MPI Environment Configuration File), "MPIRun参数" (MPIRun Parameters), and "程序" (Program) with the value "MyFolder/hello". On the right side of the form, there are buttons for "模板信息", "模板参数", "环境选项", and "通知选项".

选择相应资源，如下

LICO

首页

提交作业

作业列表

报告

云工具

AI Studio

工作台

管理

首页 > 提交作业 > MPI

欢迎, zhang san

退出

MPI

Message Passing Interface (MPI) is a communication standard used in parallel computing and HPC applications. Use this template to submit your MPI application.

模板信息

- 作业名称: MPI_08101904
- 工作目录: MyFolderIntel_mpi_test

模板参数

- 运行环境: Intel_oneAPI_Base_And_HPC_To... [详情](#) [清除](#)
- MPI程序: MyFolder/hello
- MPI环境配置文件

资源选项

- 队列: cpu
- UP 72 nodes 4032 cores UNLIMITED UNLIMITED
- 节点数量: 72
- 每节点的CPU核数: 2
- 每节点的GPU数
- 内存使用(MB)
- 最大运行时间: 24h (eg 3d 4h 12m)

通知选项

[提交](#) [清除](#)

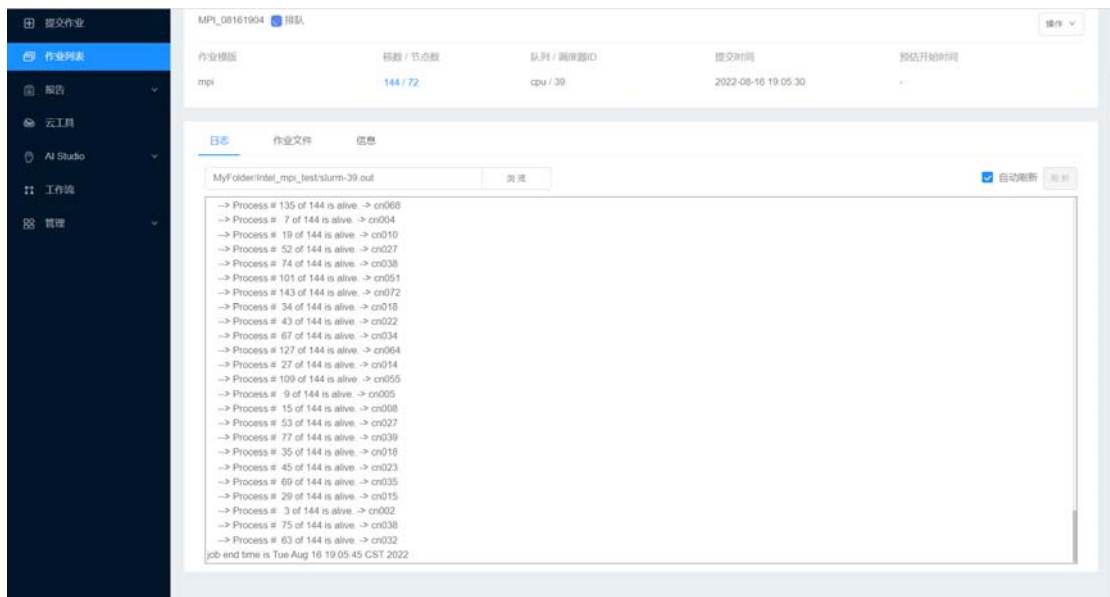
模板信息

模板参数

资源选项

通知选项

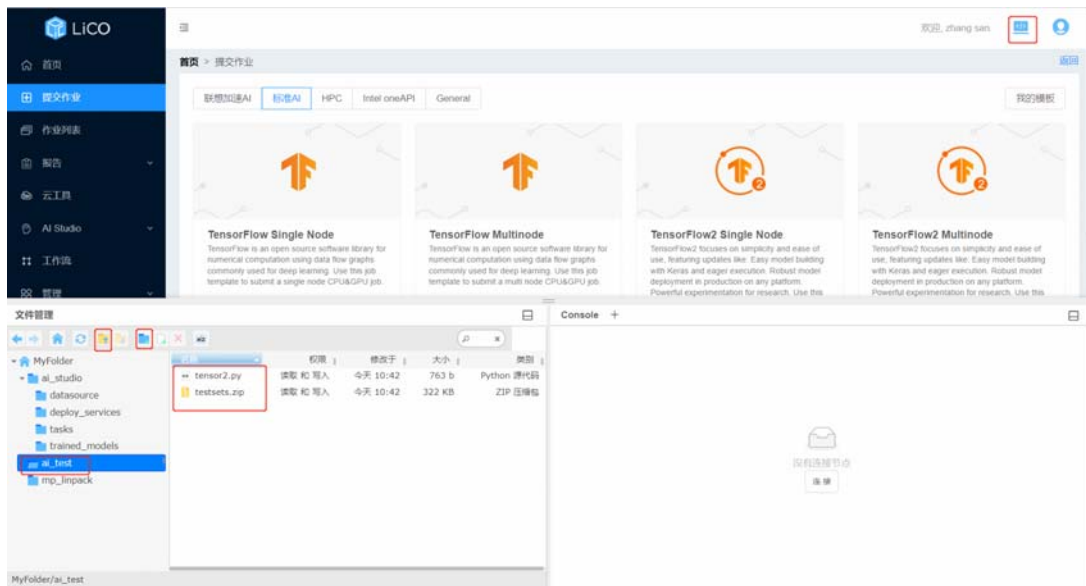
查看运行日志，如下



五. 标准 AI 作业提交（web 方式）

上传文件

点击右上角工具图标，上传数据集文件和脚本文件，如下



文件管理界面：可以进行相应的文档操作；

Console 界面：可以进行命令行操作；

tensorflow 作业提交

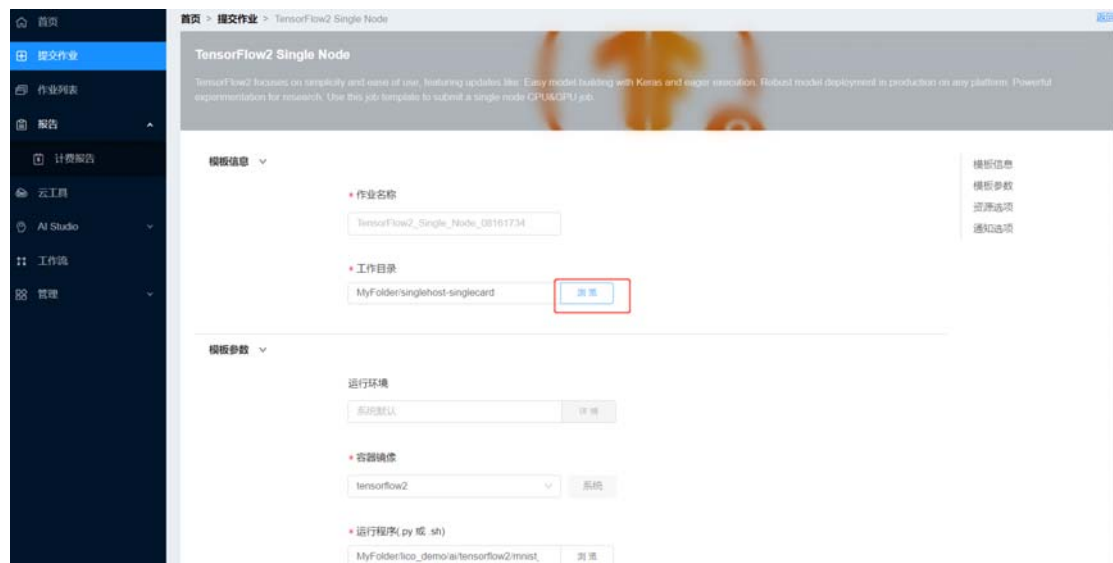
AI 标准作业提交：针对 3 种作业情况进行演示：

- ✓ 单机单卡 : mnist_single.py
- ✓ 单机多卡 : mnist_mirror_strategy.py
- ✓ 多机多 worker : mnist_multi_worker_strategy.py

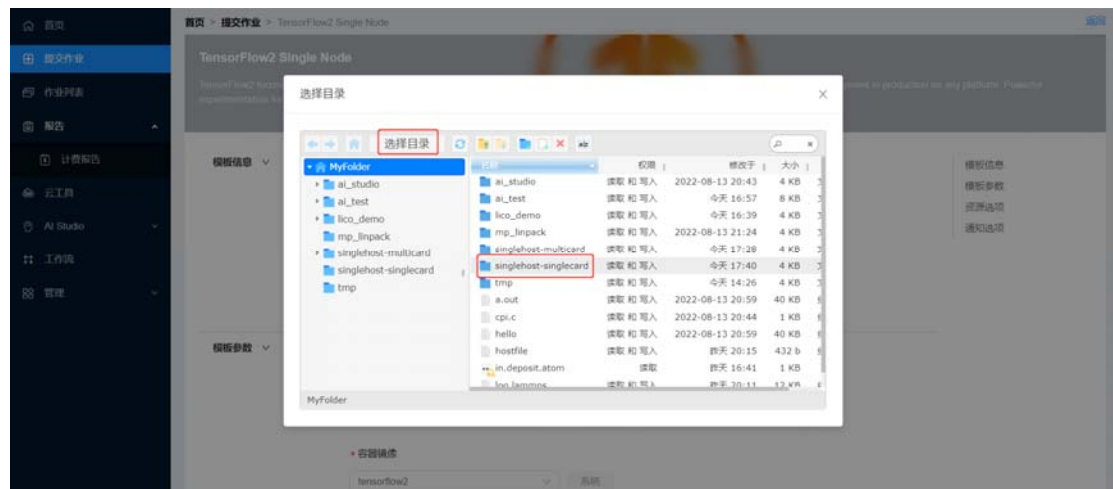
1. 单机单卡

程序文件：mnist_single.py

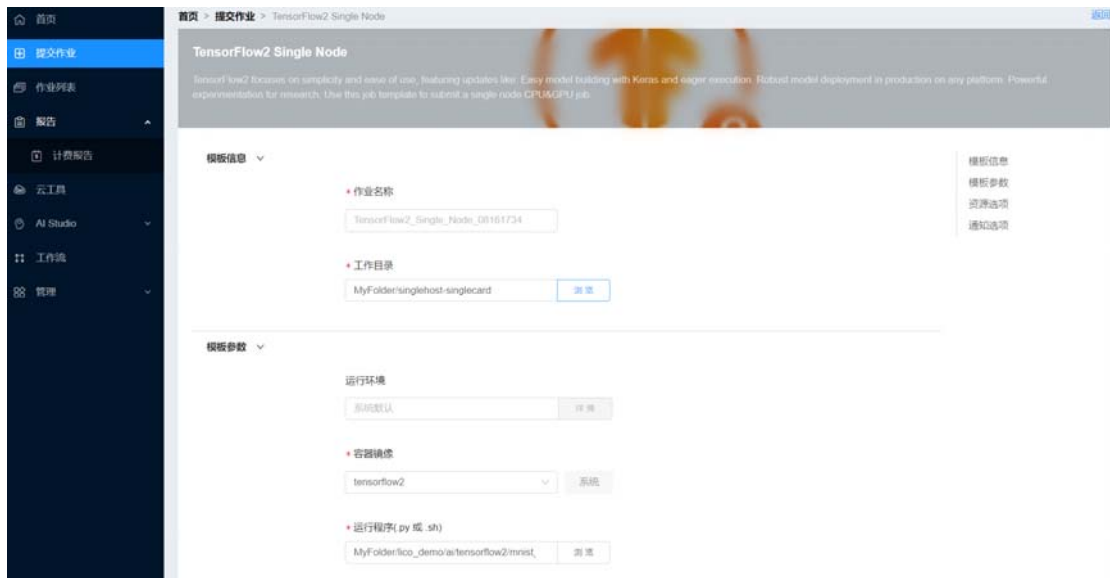
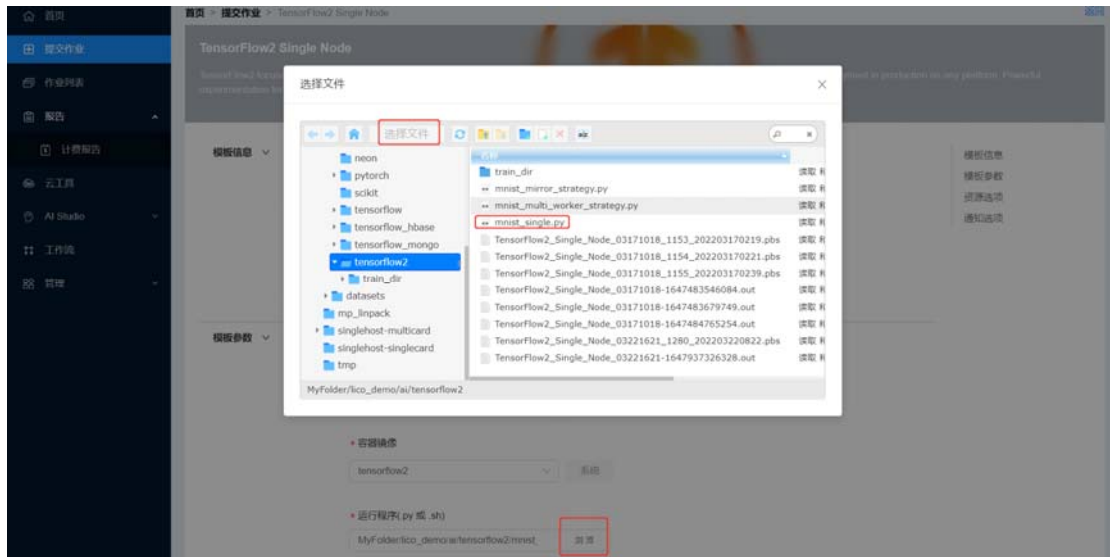
输入作业相关信息，如下



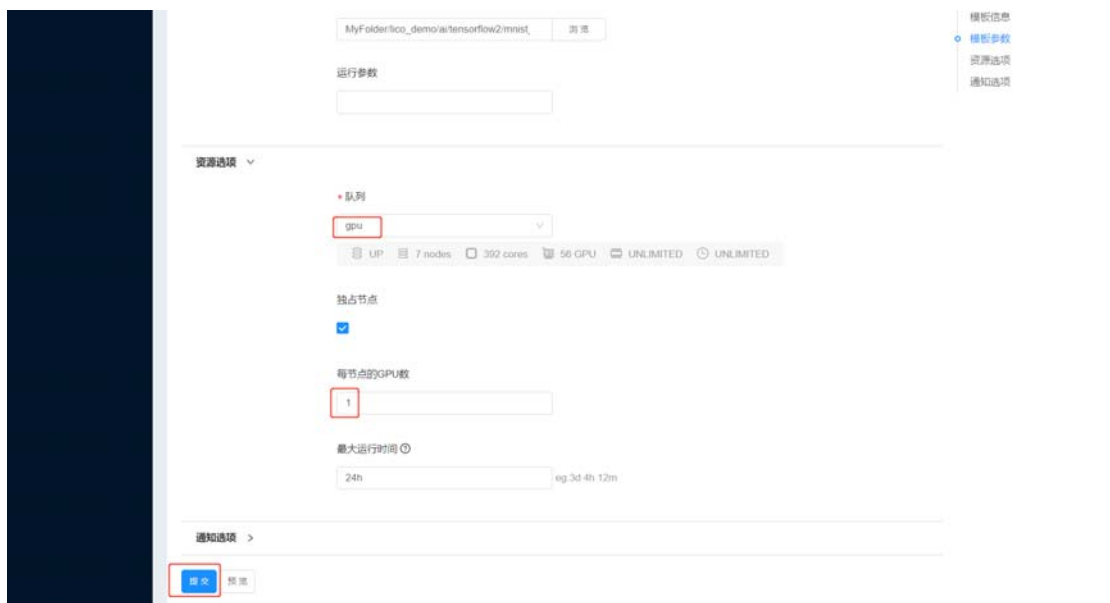
选择工作目录，如下



选择运行程序脚本，如下



选择相应作业资源后，点击提交，如下



查看作业日志的变化，如下

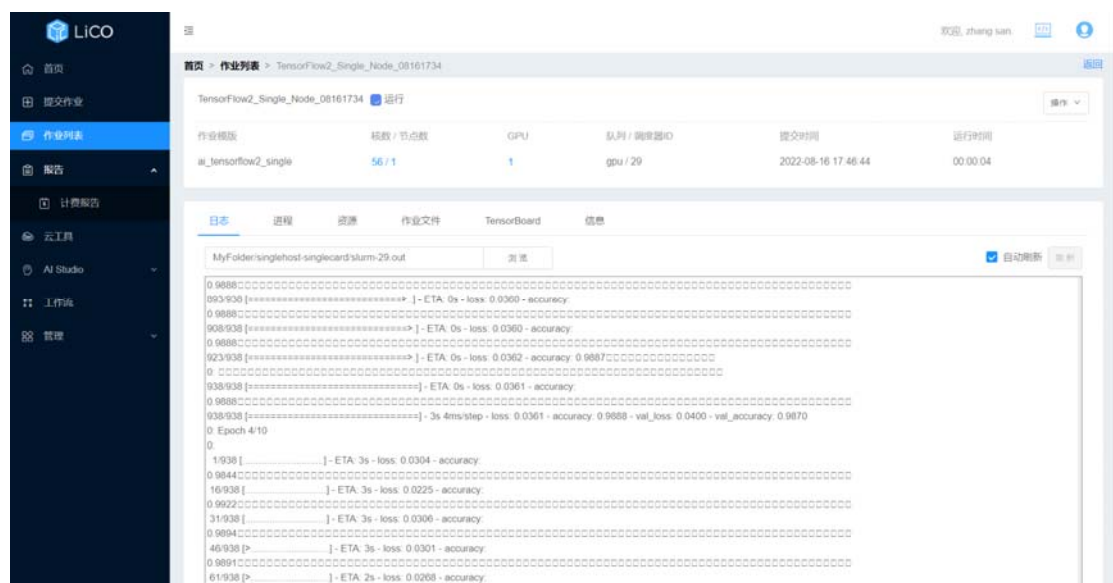


The screenshot shows the LICO dashboard with a job named 'TensorFlow2_Single_Node_08161734' in a '运行' (Running) state. The job details table is as follows:

作业模板	核数 / 节点数	GPU	队列 / 调度器ID	提交时间	运行时间
a_tensorflow2_single	56 / 1	1	gpu / 29	2022-08-16 17:46:44	00:00:04

The log content shows the start time and the command used to run the job:

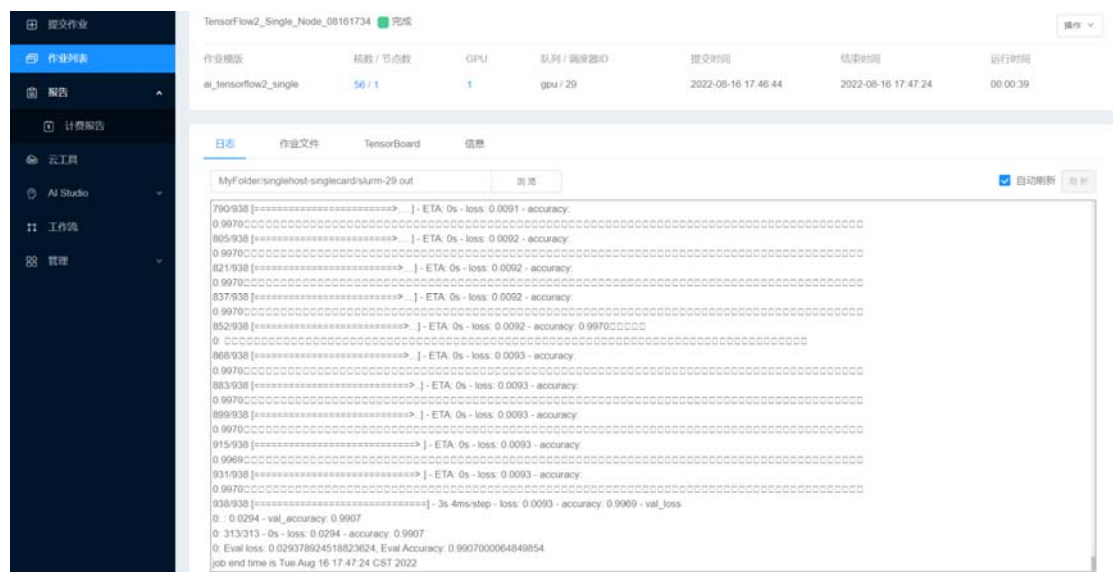
```
job start time is Tue Aug 16 17:46:45 CST 2022
[run]: "n1", "n1", "-cpus-per-task=56", "-cpu_binding=cores", "--gres=gpu:1", "--nodeslist=gpu01:", "-f", "singularity", "xrtsc", "-nv", "-B", "iSHARE/home/user01", "-pwd", "iSHARE/home/user01/singehost-singecard", "iSHARE/lco/container/tensorflow2_image", "python", "iSHARE/home/user01/lco_demo/tensorflow2/mnist_single.py"]
```



The screenshot shows the LICO dashboard with the same job in a '运行' (Running) state. The log content displays training progress over several epochs:

```
0.9888 [=====] - ETA: 0s - loss: 0.0360 - accuracy:
0.993938 [=====] - ETA: 0s - loss: 0.0360 - accuracy:
0.9888 [=====] - ETA: 0s - loss: 0.0360 - accuracy:
0.923938 [=====] - ETA: 0s - loss: 0.0362 - accuracy: 0.9887 [=====]
0. [=====]
0.938938 [=====] - ETA: 0s - loss: 0.0361 - accuracy:
0.9888 [=====] - ETA: 0s - loss: 0.0361 - accuracy: 0.9888 - val_loss: 0.0400 - val_accuracy: 0.9870
0.938938 [=====] - 3s 4ms/step - loss: 0.0361 - accuracy: 0.9888 - val_loss: 0.0400 - val_accuracy: 0.9870
0. Epoch 4/10
0. [=====]
1.938 [=====] - ETA: 3s - loss: 0.0304 - accuracy:
0.9844 [=====] - ETA: 3s - loss: 0.0225 - accuracy:
16.938 [=====] - ETA: 3s - loss: 0.0225 - accuracy:
0.9922 [=====] - ETA: 3s - loss: 0.0360 - accuracy:
31.938 [=====] - ETA: 3s - loss: 0.0360 - accuracy:
0.9894 [=====] - ETA: 3s - loss: 0.0360 - accuracy:
46.938 [p=====] - ETA: 3s - loss: 0.0301 - accuracy:
0.9891 [=====] - ETA: 3s - loss: 0.0360 - accuracy:
61.938 [p=====] - ETA: 2s - loss: 0.0268 - accuracy:
```

查看作业运行结果，如下



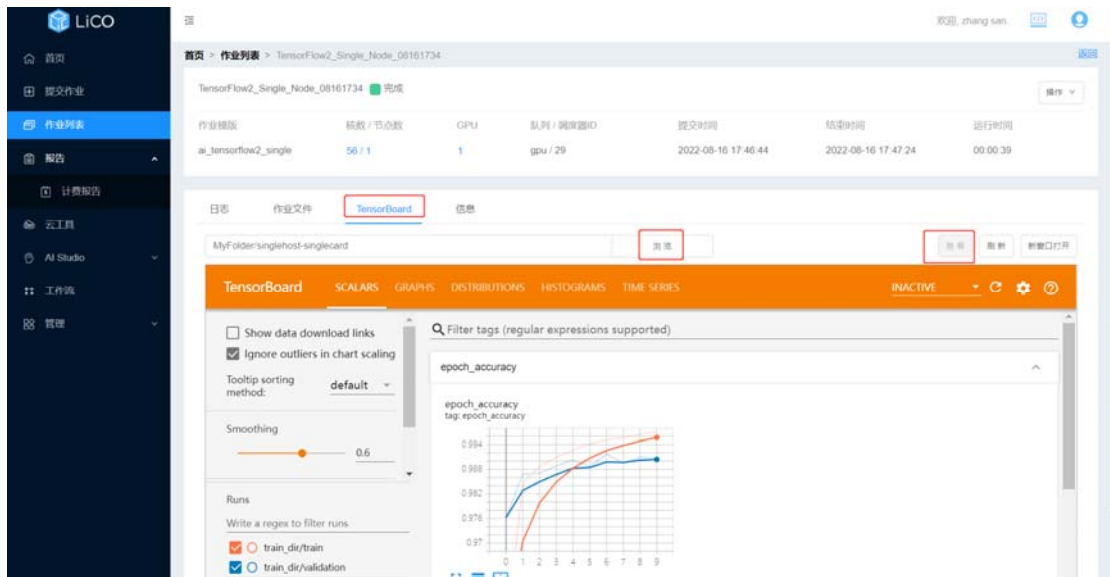
The screenshot shows the LICO dashboard with the job in a '完成' (Completed) state. The job details table is as follows:

作业模板	核数 / 节点数	GPU	队列 / 调度器ID	提交时间	结束时间	运行时间
a_tensorflow2_single	56 / 1	1	gpu / 29	2022-08-16 17:46:44	2022-08-16 17:47:24	00:00:39

The log content shows the final training results and end time:

```
790.938 [=====] - ETA: 0s - loss: 0.0091 - accuracy:
0.9970 [=====] - ETA: 0s - loss: 0.0092 - accuracy:
805.938 [=====] - ETA: 0s - loss: 0.0092 - accuracy:
0.9970 [=====] - ETA: 0s - loss: 0.0092 - accuracy:
821.938 [=====] - ETA: 0s - loss: 0.0092 - accuracy:
0.9970 [=====] - ETA: 0s - loss: 0.0092 - accuracy:
837.938 [=====] - ETA: 0s - loss: 0.0092 - accuracy:
0.9970 [=====] - ETA: 0s - loss: 0.0092 - accuracy: 0.9970 [=====]
852.938 [=====] - ETA: 0s - loss: 0.0092 - accuracy: 0.9970 [=====]
0. [=====]
868.938 [=====] - ETA: 0s - loss: 0.0093 - accuracy:
0.9970 [=====] - ETA: 0s - loss: 0.0093 - accuracy:
883.938 [=====] - ETA: 0s - loss: 0.0093 - accuracy:
0.9970 [=====] - ETA: 0s - loss: 0.0093 - accuracy:
899.938 [=====] - ETA: 0s - loss: 0.0093 - accuracy:
915.938 [=====] - ETA: 0s - loss: 0.0093 - accuracy:
0.9969 [=====] - ETA: 0s - loss: 0.0093 - accuracy:
931.938 [=====] - ETA: 0s - loss: 0.0093 - accuracy:
0.9970 [=====] - ETA: 0s - loss: 0.0093 - accuracy:
938.938 [=====] - 3s 4ms/step - loss: 0.0093 - accuracy: 0.9969 - val_loss
0. : 0.0294 - val_accuracy: 0.9907
0. 31:31:13 - 0s - loss: 0.0294 - accuracy: 0.9907
0. Eval loss: 0.029378924518823624, Eval Accuracy: 0.9907000064848854
job end time is Tue Aug 16 17:47:24 CST 2022
```

打开 TensorBoard 查看作业情况



2. 单机多卡

程序文件 mnist_mirror_strategy.py

如下输入作业相关信息

TensorFlow2 Single Node

TensorFlow2 focuses on simplicity and ease of use, featuring optimizers like Easy model building with Keras and eager execution. Rich model deployment in production on any platform. Powerful experimentation for research. Use this job template to submit a single node (CPU/GPU) job.

模板信息

- 作业名称: TensorFlow2_Single_Node_08161726
- 工作目录: MyFolder/singlehost-multicard

模板参数

- 运行环境: 系统默认
- 容器镜像: tensorflow2
- 运行程序(py或.sh): MyFolder/lico_demo/ai/tensorflow2/mnist

MyFolder/llico_demo/ai/tensorflow2/mnist 浏览

运行参数

资源选项

队列

gpu

UP 7 nodes 391 cores 56 GPU UNLIMITED UNLIMITED

独占节点

每节点的GPU数

8

最大运行时间

24h eg.3d 4h 12m

通知选项

提交 预览

作业运行日志

LiCO 首页 提交作业 作业列表 报告 云工具 AI Studio 工作流 管理

TensorFlow2_Single_Node_08161726 队列

作业模板	核数 / 节点数	GPU	队列 / 调度器ID	提交时间	预估开始时间
ai_tensorflow2_single	1 / 1	8	gpu / 27	2022-08-16 17:28:52	-

日志 作业文件 TensorBoard 信息

MyFolder/singlehost-multicard/slurm-27 out 浏览 自动刷新

```

job start time is Tue Aug 16 17:28:53 CST 2022
["srun", "NT", "-cpus-per-task=56", "-cpu_bind=cores", "-gres=gpu:8", "-nodeidsgpu01", "T", "singularity", "exec", "-nv", "0", "share/home/user01", "-pid", "share/home/user01/singlehost-multicard", "share/llico/container/tensorflow2_image", "python", "share/home/user01/llico_demo/ai/tensorflow2/mnist_mirror_strategy.py"]

```

TensorFlow2_Single_Node_08161726 运行

作业模板	核数 / 节点数	GPU	队列 / 调度器ID	提交时间	运行时间
ai_tensorflow2_single	56 / 1	8	gpu / 27	2022-08-16 17:28:52	00:00:34

日志 进程 资源 作业文件 TensorBoard 信息

MyFolder/singlehost-multicard/slurm-27 out 浏览 自动刷新

```

0 -----] - ETA: 2s - loss: 0.0126 - accuracy:
0.9958 [=====] - ETA: 2s - loss: 0.0126 - accuracy:
192/938 [=====] - ETA: 2s - loss: 0.0126 - accuracy:
0.9958 [=====] - ETA: 2s - loss: 0.0122 - accuracy:
207/938 [=====] - ETA: 2s - loss: 0.0122 - accuracy:
0.9959 [=====] - ETA: 2s - loss: 0.0131 - accuracy:
222/938 [=====] - ETA: 2s - loss: 0.0131 - accuracy:
0.9958 [=====] - ETA: 2s - loss: 0.0129 - accuracy:
237/938 [=====] - ETA: 2s - loss: 0.0129 - accuracy:
0.9958 [=====] - ETA: 2s - loss: 0.0132 - accuracy:
252/938 [=====] - ETA: 2s - loss: 0.0132 - accuracy:
0.9957 [=====] - ETA: 2s - loss: 0.0129 - accuracy:
267/938 [=====] - ETA: 2s - loss: 0.0128 - accuracy:
282/938 [=====] - ETA: 2s - loss: 0.0128 - accuracy:
0.9959 [=====] - ETA: 2s - loss: 0.0127 - accuracy:
297/938 [=====] - ETA: 2s - loss: 0.0136 - accuracy:
312/938 [=====] - ETA: 2s - loss: 0.0143 - accuracy:
0.9956 [=====] - ETA: 2s - loss: 0.0143 - accuracy:
327/938 [=====] - ETA: 2s - loss: 0.0148 - accuracy:
0.9953 [=====] - ETA: 2s - loss: 0.0148 - accuracy:
342/938 [=====] - ETA: 2s - loss: 0.0148 - accuracy:
0.9952 [=====] - ETA: 2s - loss: 0.0148 - accuracy:
357/938 [=====]

```

作业运行结果



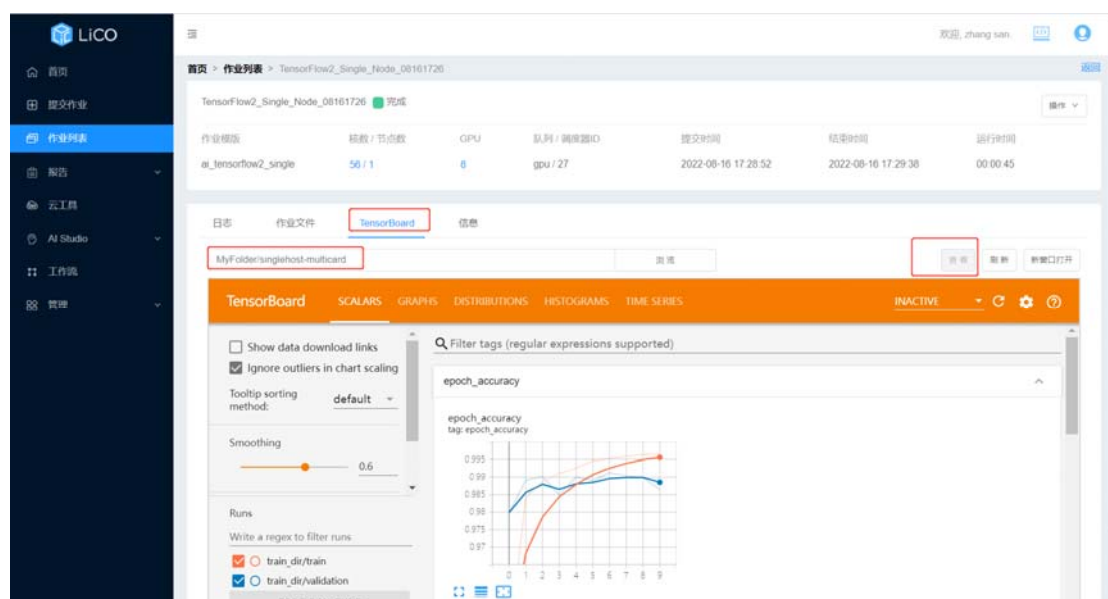
TensorFlow2_Single_Node_08161726 运行

作业描述	核数 / 节点数	GPU	队列 / 调度器ID	提交时间	运行时间
ai_tensorflow2_single	56 / 1	8	gpu / 27	2022-08-16 17:28:52	00:00:34

MyFolder/singlehost-multicard/slm-27.out

```
0: attr {
0: key "output_shapes"
0: value {
0: list {
0: shape {
0: dim {
0: size -1
0: }
0: }
0: }
0: }
0: }
0: attr {
0: key "output_types"
0: value {
0: list {
0: type: DT_INT64
0: }
0: }
0: }
0: }
0: Consider either turning off auto-sharding or switching the auto_shard_policy to DATA to shard this dataset. You can do this by creating a new 'tf.data.Options()' object then setting options.experimental_distribute.auto_shard_policy = AutoShardPolicy.DATA' before applying the options object to the dataset via 'dataset.with_options(options)'.
0: 313/313 - Ds - loss: 0.0549 - accuracy: 0.9865
0: Eval loss: 0.054916419688840485, Eval Accuracy: 0.9865000247955322
job end time is Tue Aug 16 17:29:38 CST 2022
```

查看 tensorboard



TensorFlow2_Single_Node_08161726 完成

作业描述	核数 / 节点数	GPU	队列 / 调度器ID	提交时间	结束时间	运行时间
ai_tensorflow2_single	56 / 1	8	gpu / 27	2022-08-16 17:28:52	2022-08-16 17:29:38	00:00:45

MyFolder/singlehost-multicard

TensorBoard SCALARS GRAPHS DISTRIBUTIONS HISTOGRAMS TIME SERIES INACTIVE

epoch_accuracy

epoch_accuracy
tag: epoch_accuracy

Epoch	Train Accuracy	Validation Accuracy
0	0.975	0.985
1	0.985	0.985
2	0.985	0.985
3	0.985	0.985
4	0.985	0.985
5	0.985	0.985
6	0.985	0.985
7	0.985	0.985
8	0.985	0.985
9	0.985	0.985

3. 多机多 worker

程序文件 mnist_multi_worker_strategy.py

如下输入作业相关信息

TensorFlow2 Multinode

TensorFlow2 focuses on simplicity and ease of use, featuring updates like: Easy model building with Keras and eager execution. Robust model deployment in production on any platform. Powerful experimentation for research. Use the job template to submit a multi-node CPU/GPU job.

模板信息

- 作业名称: TensorFlow2_Multinode_00101754
- 工作目录: MyFolder/multinode-multicard [浏览](#)

模板参数

- 运行环境: 系统默认 [浏览](#)
- 容器镜像: tensorflow2 [浏览](#)
- 运行程序(py或sh): MyFolder/lco_demo/tensorflow2/mnist [浏览](#)

模板信息
模板参数
资源选项
高级参数
通知选项

资源选项

队列: gpu

节点数量: 7

独占节点:

每节点的GPU数: 8

最大运行时间: 24h (eg. 3d 4h 12m)

UP 7 nodes 391 cores 56 GPU UNLIMITED UNLIMITED

作业运行日志

提交作业

作业列表

报告

云工具

AI Studio

工作台

管理

TensorFlow2_Multinode_08161754 完成

作业组名	核数 / 节点数	GPU	队列 / 调度器ID	提交时间	结束时间	运行时间
ai_tensorflow2	392 / 7	56	gpu / 31	2022-08-16 17:58:01	2022-08-16 18:02:55	00:02:23

日志 作业文件 TensorBoard 信息

MyFolder/multihost-multicard/slurm-31 out 自动刷新 刷新

```

821938 [=====] -j- ETA: 1s - loss: 0.0169 - acc:
0.0171 - accuracy: 0.9949 [=====] -j- ETA: 1s - loss: 0.0170 - accuracy:
801938 [=====] -j- ETA: 1s - loss: 0.0172 - accuracy:
0.9949 [=====] -j- ETA: 1s - loss: 0.0171 - accuracy:
805938 [=====] -j- ETA: 1s - loss: 0.0170 - accuracy:
0.9949 [=====] -j- ETA: 1s - loss: 0.0170 - accuracy:
809938 [=====] -j- ETA: 1s - loss: 0.0170 - accuracy:
0.9949 [=====] -j- ETA: 1s - loss: 0.0170 - accuracy:
813938 [=====] -j- ETA: 1s - loss: 0.0170 - accuracy:
819940 [=====] -j- ETA: 1s - loss: 0.0169 - accuracy:
817938 [=====] -j- ETA: 1s - loss: 0.0169 - accuracy:
0.9950 [=====] -j- ETA: 1s - loss: 0.0169 - accuracy:
821938 [=====] -j- ETA: 1s - loss: 0.0169 - acc:
0.9950 [=====] -j- ETA: 1s - loss: 0.0169 - accuracy:
825938 [=====] -j- ETA: 1s - loss: 0.0169 - accuracy:
0.9950 [=====] -j- ETA: 1s - loss: 0.0169 - accuracy:
829938 [=====] -j- ETA: 1s - loss: 0.0169 - accuracy:
0.9949 [=====] -j- ETA: 1s - loss: 0.0166 - accuracy:
833938 [=====] -j- ETA: 1s - loss: 0.0166 - accuracy:
837938 [=====] -j- ETA: 1s - loss: 0.0167 - accuracy:
839950 [=====] -j- ETA: 1s - loss: 0.0167 - accuracy:
841938 [=====] -j- ETA: 1s - loss: 0.0166 - accuracy:
0.9950 [=====] -j- ETA: 1s - loss: 0.0166 - accuracy: 0.9

```

作业运行结果

提交作业

作业列表

报告

云工具

AI Studio

工作台

管理

TensorFlow2_Multinode_08161754 完成

作业组名	核数 / 节点数	GPU	队列 / 调度器ID	提交时间	结束时间	运行时间
ai_tensorflow2	392 / 7	56	gpu / 31	2022-08-16 17:58:01	2022-08-16 18:02:55	00:02:23

日志 作业文件 TensorBoard 信息

MyFolder/multihost-multicard/slurm-31 out 自动刷新 刷新

```

0 terminate called without an active exception
0 terminate called without an active exception
0 terminate called without an active exception
0 terminate called without an active exception
0 /usr/bin/bash: line 1: 289646: Aborted singularity exec --nv -B /share/home/user01 --pwd /share/home/user01/multihost-multicard /share/ico/container/tensorflow2 image python
/share/home/user01/ico_demo/ai/tensorflow2/mnist_multi_worker_strategy.py --worker_hosts=gpu01.27423.gpu02.25696.gpu03.25460.gpu04.26181.gpu05.25722.gpu06.25522.gpu07.25946 --
job_name=worker --task_index=1
slurm: error: gpu02: task 0: Exited with exit code 134
0 /usr/bin/bash: line 1: 238259: Aborted singularity exec --nv -B /share/home/user01 --pwd /share/home/user01/multihost-multicard /share/ico/container/tensorflow2 image python
/share/home/user01/ico_demo/ai/tensorflow2/mnist_multi_worker_strategy.py --worker_hosts=gpu01.27423.gpu02.25696.gpu03.25460.gpu04.26181.gpu05.25722.gpu06.25522.gpu07.25946 --
job_name=worker --task_index=4
slurm: error: gpu05: task 0: Exited with exit code 134
0 /usr/bin/bash: line 1: 238788: Aborted singularity exec --nv -B /share/home/user01 --pwd /share/home/user01/multihost-multicard /share/ico/container/tensorflow2 image python
/share/home/user01/ico_demo/ai/tensorflow2/mnist_multi_worker_strategy.py --worker_hosts=gpu01.27423.gpu02.25696.gpu03.25460.gpu04.26181.gpu05.25722.gpu06.25522.gpu07.25946 --
job_name=worker --task_index=2
slurm: error: gpu03: task 0: Exited with exit code 134
0 /usr/bin/bash: line 1: 281862: Aborted singularity exec --nv -B /share/home/user01 --pwd /share/home/user01/multihost-multicard /share/ico/container/tensorflow2 image python
/share/home/user01/ico_demo/ai/tensorflow2/mnist_multi_worker_strategy.py --worker_hosts=gpu01.27423.gpu02.25696.gpu03.25460.gpu04.26181.gpu05.25722.gpu06.25522.gpu07.25946 --
job_name=worker --task_index=3
slurm: error: gpu04: task 0: Exited with exit code 134
0 /usr/bin/bash: line 1: 51709: Aborted singularity exec --nv -B /share/home/user01 --pwd /share/home/user01/multihost-multicard /share/ico/container/tensorflow2 image python
/share/home/user01/ico_demo/ai/tensorflow2/mnist_multi_worker_strategy.py --worker_hosts=gpu01.27423.gpu02.25696.gpu03.25460.gpu04.26181.gpu05.25722.gpu06.25522.gpu07.25946 --
job_name=worker --task_index=0
slurm: error: gpu01: task 0: Exited with exit code 134
job end time is Tue Aug 16 18:02:55 CST 2022

```

查看 tensorBord

The screenshot shows the LICO AI Studio interface. On the left is a dark sidebar with navigation options: 首页, 提交作业, 作业列表, 报告, 云工具, AI Studio, workflow, and 管理. The main area displays a job titled 'TensorFlow2_Multinode_08161754' with a status of '完成' (Completed). Below the job name is a table with columns: 作业模板, 核数/节点数, GPU, 队列/调度器ID, 提交时间, 结束时间, and 运行时间. The row shows 'ai_tensorflow2', '392/7', '56', 'gpu/31', '2022-08-16 17:58:01', '2022-08-16 18:02:55', and '00:02:23'. Below the table are tabs for '日志', '作业文件', 'TensorBoard', and '信息'. The 'TensorBoard' tab is active, showing a line graph of 'epoch_accuracy' over 9 epochs. The graph has two lines, one red and one blue, both showing an upward trend. The y-axis ranges from 0.965 to 0.995. The x-axis ranges from 0 to 9. The interface also includes various settings like 'Show data download links', 'Ignore outliers in chart scaling', and 'Filter tags'.

caffe 作业提交

The screenshot shows the LICO AI Studio interface for submitting a Caffe job. The left sidebar is the same as in the previous screenshot. The main area is titled 'Caffe' and contains a form for job submission. The form has two main sections: '模板信息' (Template Information) and '模板参数' (Template Parameters). Under '模板信息', there are fields for '作业名称' (Job Name) with the value 'Caffe_08161815' and '工作目录' (Working Directory) with the value 'MyFolder/caffe_test'. Under '模板参数', there are fields for '运行环境' (Runtime Environment) with the value '系统默认', '容器镜像' (Container Image) with the value 'caffe-gpu', and '运行程序(py 或 sh)' (Runtime Program) with the value 'MyFolder/lico_demo/ai/caffe/cifar10_lener'. The interface also includes a sidebar on the right with options like '模板信息', '模板参数', '资源选项', and '通知选项'.

* 队列

gpu

UP 7 nodes 391 cores 56 GPU UNLIMITED UNLIMITED

独占节点



每节点的GPU数

8

最大运行时间

24h

eg.3d 4h 12m

通知选项

提交

预览

- 提交作业
- 作业列表
- 报告
- 云工具
- AI Studio
- 数据集
- 训练任务
- 模型
- 服务
- 工作流
- 管理

Caffe_08161815 完成
操作

作业模版	核数 / 节点数	GPU	队列 / 调度器ID	提交时间	结束时间	运行时间
ai_caffe	56 / 1	8	gpu / 33	2022-08-16 18:18:53	2022-08-16 18:18:56	00:00:02

日志
作业文件
信息

MyFolder/caffe_test/slum-33.out
预览

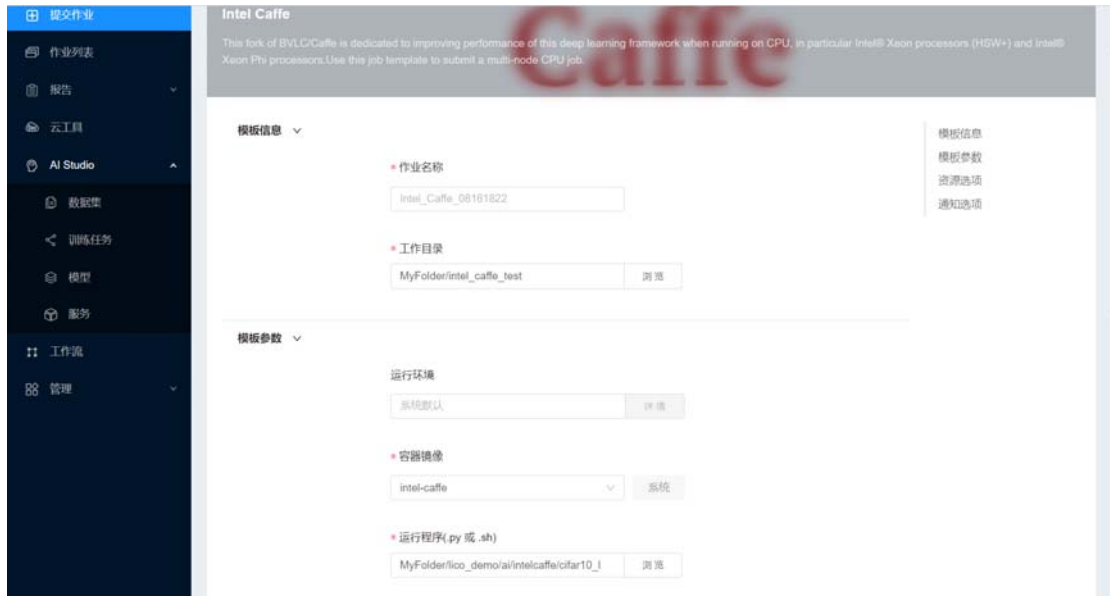
```

job start time is Tue Aug 16 18:18:54 CST 2022
[!srin!, -N1, -n1, --cpu-per-task=56, --cpu_bind=cores, --gres=gpu:8, --nodelist=gpu01, -f, 'singularity', 'exec', '-mv, -B, /share/home/user01', '-pwd',
'/share/home/user01/caffe_test', '/share/lico/container/caffe-gpu.image', 'python', '/share/home/user01/lico_demo/ai/caffe/ctar10_line1_gpu.py']
0: F0816 18:18:50.317871 53979 caffe.cpp:93] Check failed: error == cudaSuccess (30 vs. 0) unknown error
0: *** Check failure stack trace: ***
0: @ 0x7124dfe535cd google::LogMessage::Fail()
0: @ 0x7124dfe55433 google::LogMessage::SendToLog()
0: @ 0x7124dfe5315b google::LogMessage::Flush()
0: @ 0x7124dfe55e1e google::LogMessageFatal::~LogMessageFatal()
0: @ 0x40cf72 get_gpus()
0: @ 0x40e601 train()
0: @ 0x40b8c3 main
0: @ 0x7124deb3b840 __libc_start_main
0: @ 0x40c269 _start
0: @ (nil) (unknown)
0: Aborted (core dumped)
0: /share/home/user01/lico_demo/ai/caffe
job end time is Tue Aug 16 18:18:56 CST 2022

```

intel_caffe 作业提交

输入作业相关参数，如下



选择作业资源

资源选项 ▾

* 队列

cpu



72 nodes



4032 cores



UNLIMITED



UNLIMITED

* 节点数量

1

独占节点

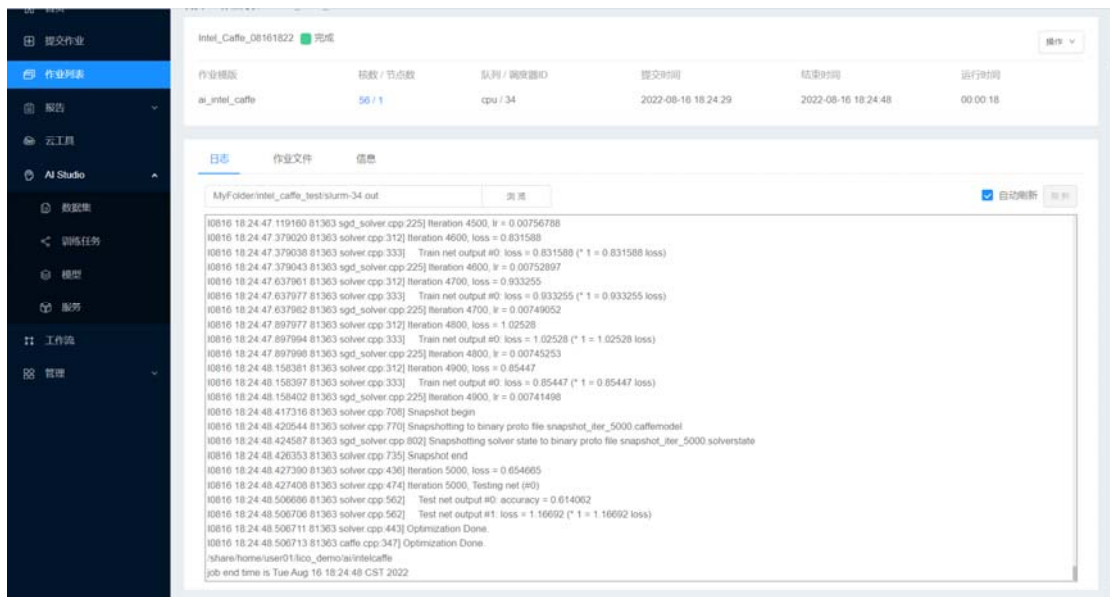


最大运行时间 ⌚

24h

eg.3d 4h 12m

运行结果，如下



六. 加速 AI 作业提交（训练/推理）

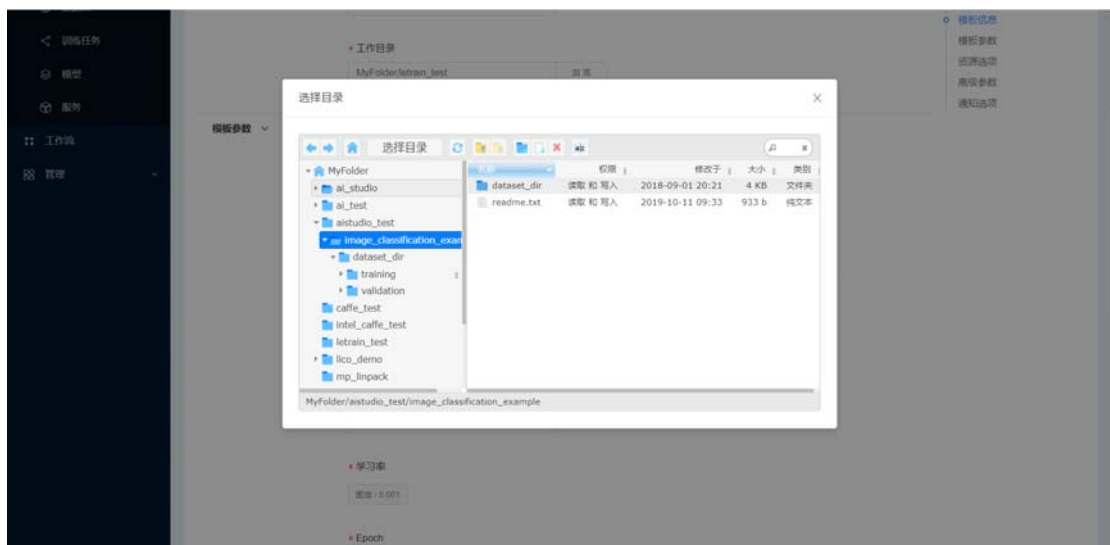
本例仅以图片识别功能，进行训练演示。

上传文件

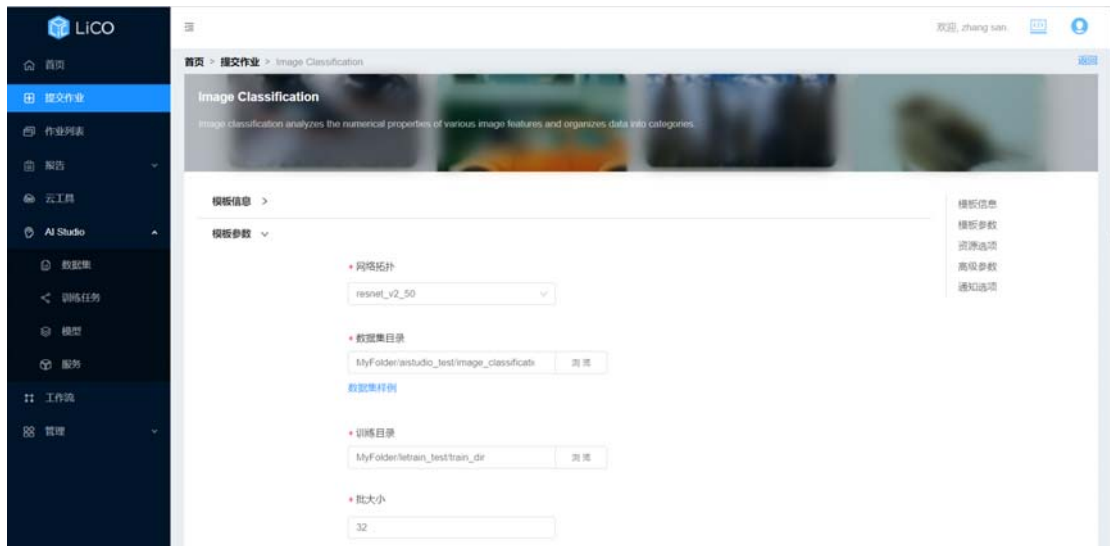
上传数据集文件（略）。

图片识别训练作业提交

选择数据集和训练目录，如下

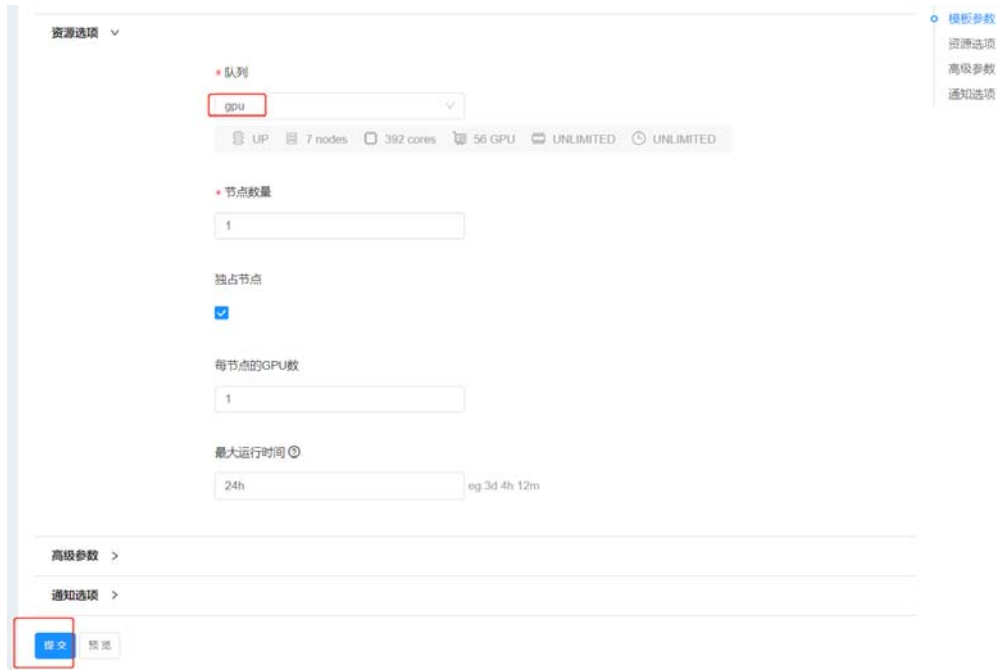


选择训练参数，如下

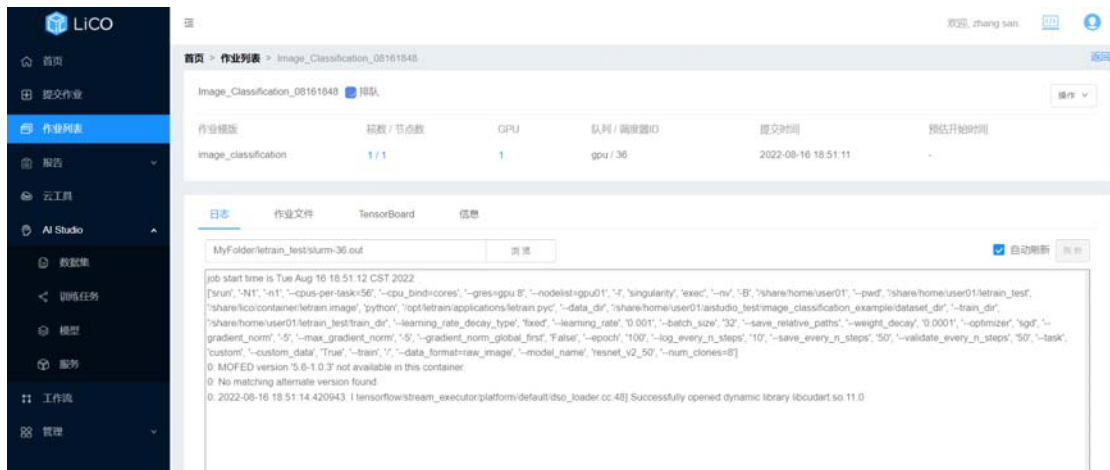


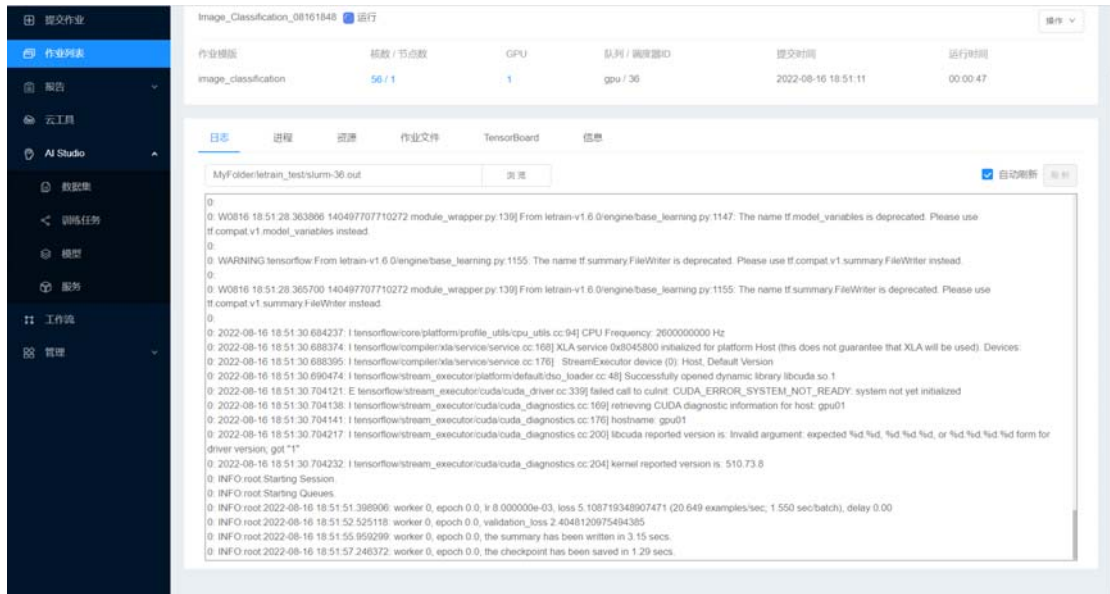
选择运行资源，点击提交，如下



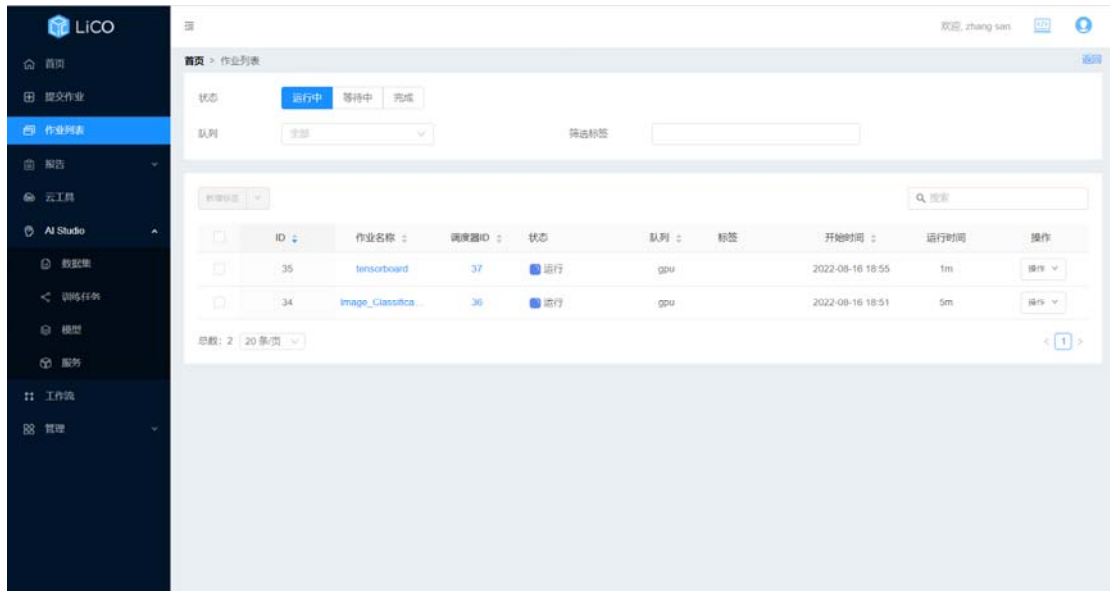


查看作业运行日志，如下





查看作业运行状态，如下



图片识别推理作业提交

(查阅 lico 用户手册)

七. AI studio

Singularity 作业提交

模板信息

作业名称: Singularity_09101907

工作目录: MyFolder/caffe_test

模板参数

容器镜像: caffe-cpu

容器内运行脚本:

```
import argparse
import subprocess
import os
import sklearn
parser = argparse.ArgumentParser(description='manual to this script')
parser.add_argument('--max_steps', type=int, default=5000)
parser.add_argument('--batch_size', type=int, default=32)
args = parser.parse_args()
max_steps=args.max_steps
batch_size=args.batch_size
path=os.getcwd()+os.path.realpath('file')
```

with open(file,"w") as f:
 f.write(file_data)
subprocess.call("caffe train --solver="+path+"/lenet_caffe_solver.prototxt",shell=True)

资源选项

队列: cpu

72 nodes 4032 cores UNLIMITED UNLIMITED

每节点的CPU核数: 1

每节点的GPU数:

内存使用(MB):

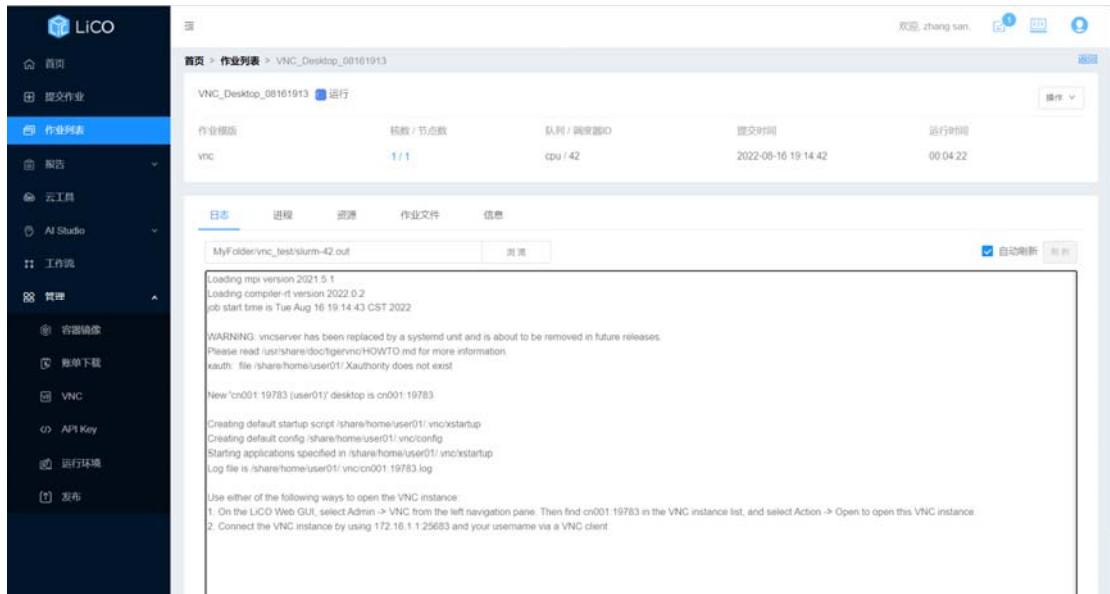
最大运行时间: 24h (eg 3d 4h 12m)

通知选项

The screenshot shows the LICO web interface with a sidebar on the left containing navigation options like '首页', '提交作业', '作业列表', '报告', '云工具', 'AI Studio', '工作流', and '管理'. The main content area displays a job summary for 'Singularity_08161907' with a status of '完成' (Completed). Below the summary is a table with columns for '作业模板', '核数 / 节点数', '队列 / 调度器ID', '提交时间', '结束时间', and '运行时间'. The job details table shows: Singularity, 1 / 1, cpu / 40, 2022-08-16 19:10:28, 2022-08-16 19:10:29, and 00:00:00. A '日志' (Log) tab is selected, showing the output of 'MyFolder\caffe_test\slurm-40 out'. The log content includes: 'job start time is Tue Aug 16 19:10:29 CST 2022', 'cn001', and five lines of 'import: command not found' errors. The job ends at 'Tue Aug 16 19:10:29 CST 2022'.

VNC 作业提交

The screenshot shows the LICO web interface with a sidebar on the left. The main content area displays a job summary for 'VNC_Desktop_08161913' with a status of '运行' (Running). Below the summary is a table with columns for '作业模板', '核数 / 节点数', '队列 / 调度器ID', '提交时间', and '运行时间'. The job details table shows: vnc, 1 / 1, cpu / 42, 2022-08-16 19:14:42, and 00:00:34. A '日志' (Log) tab is selected, showing the output of 'MyFolder\vnc_test\slurm-42 out'. The log content includes: 'Loading mpi version 2021.5.1', 'Loading compiler-rt version 2022.0.2', 'job start time is Tue Aug 16 19:14:43 CST 2022', a warning about vncserver being replaced, and instructions for opening the VNC instance. The instructions are: 1. On the LICO Web GUI, select Admin -> VNC from the left navigation pane. Then find cn001:19783 in the VNC instance list, and select Action -> Open to open this VNC instance. 2. Connect the VNC instance by using 172.16.1.125683 and your username via a VNC client.



vnc client 连接成功

