Ubuntu 24.04 + Docker □□□ LiteLLM v1.87.0

□□□□□□ **Ubuntu 24.04 + Docker**□□□ LiteLLM v1.87.0 □□□□□ **Docker Compose** □□□□□□ `docker run` □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□

□□□□□□□□□□□□□□

□□□□□□

□□ Docker□

```
#Ubuntu24.04 □□Docker&Compose

sudo apt update && sudo apt install -y ca-certificates curl gnupg lsb-release git

# □□□□□□□□□□□□□□□□□□□□
curl -fsSL https://download.docker.com/linux/ubuntu/gpg | sudo gpg --dearmor -o
/usr/share/keyrings/docker-archive-keyring.gpg
# □□□□□□□
sudo chmod 644 /usr/share/keyrings/docker-archive-keyring.gpg

sudo apt update
sudo apt install -y docker-ce docker-ce-cli containerd.io docker-compose-plugin

# □□□□□
sudo systemctl enable --now docker
# □□□□□□sudo□□docker
sudo usermod -aG docker $USER
# □□□□□□□□□□□□□□□□
newgrp docker

# □□Docker□□
```

```
docker --version
docker compose version
```

□□□□ docker□□□□

```
mkdir -p /etc/docker
cat > /etc/docker/daemon.json <<EOF
{
  "registry-mirrors": [
    "https://docker.mirrors.ustc.edu.cn",
    "https://hub-mirror.c.163.com"
  ]
}
EOF
systemctl daemon-reload
systemctl restart docke
```

□□ Docker□

```
docker --version
docker compose version
```

□□□

```
Docker >= 24
Docker Compose >= 2.x
```

□□□□□□

□□□

```
mkdir -p /opt/litellm
cd /opt/litellm
```

□□□

```
/opt/litellm
├── docker-compose.yml
├── .env
└── config.yaml
```

□□□ LiteLLM □□

config.yaml

```
vi config.yaml
```

□□□

```
model_list:

  # OpenAI □□
  - model_name: gpt-5
    litellm_params:
      model: openai/gpt-5
      api_key: os.environ/OPENAI_API_KEY


  # Gemini □□
  - model_name: gemini-3.5-flash
    litellm_params:
      model: vertex_ai/gemini-3.5-flash
      api_key: os.environ/GEMINI_API_KEY


general_settings:

  master_key: os.environ/LITELLM_MASTER_KEY

  # □□□□□□□
  store_model_in_db: true
```

□□□□□□□

.env

```
vi .env
```

□□□

```
LITELLM_MASTER_KEY=sk-litellm-admin-123456

OPENAI_API_KEY=sk-xxxx

GEMINI_API_KEY=xxxx
```

□□ Docker Compose □□□□□

□□□

```
vi docker-compose.yml
```

□□□

```
services:

  litellm:
    image: ghcr.io/berriai/litellm:1.87.0

    container_name: litellm

    restart: always
```

```
  ports:
    - "4000:4000"

  environment:

    STORE_MODEL_IN_DB: "True"

    DATABASE_URL: postgresql://litellm:litellm@postgres:5432/litellm

    LITELLM_MASTER_KEY: ${LITELLM_MASTER_KEY}

    OPENAI_API_KEY: ${OPENAI_API_KEY}

    GEMINI_API_KEY: ${GEMINI_API_KEY}


  volumes:

    - ./config.yaml:/app/config.yaml


  command:
    --config /app/config.yaml


  depends_on:
    - postgres



postgres:

  image: postgres:16

  container_name: litellm-postgres

  restart: always
```

```
    environment:
      POSTGRES_DB: litellm
      POSTGRES_USER: litellm
      POSTGRES_PASSWORD: litellm

    volumes:
      - ./postgres:/var/lib/postgresql/data
```

□□□□

□□□

```
# □□□□ghcr□□□□□□□
docker pull ghcr.nju.edu.cn/berriai/litellm:1.87.0
# □□□□□□□□□□□□□compose
docker tag ghcr.nju.edu.cn/berriai/litellm:1.87.0 ghcr.io/berriai/litellm:1.87.0
docker rmi ghcr.nju.edu.cn/berriai/litellm:1.87.0

docker compose up -d
```

□□□

```
docker ps
```

□□□

```
litellm
litellm-postgres
```

□□□□□

```
docker logs -f litellm
```

□□□

```
LiteLLM Proxy Server Started
Server running on port 4000
```

□□□□ API

LiteLLM □□□□ OpenAI API□

□□□

```
curl http://127.0.0.1:4000/v1/models \
-H "Authorization: Bearer sk-litellm-admin-123456"
```

□□□

```
{
 "data":[
   {
    "id":"gpt-5"
   }
 ]
}
```

□□□□□

```
curl http://127.0.0.1:4000/v1/chat/completions \
-H "Authorization: Bearer sk-litellm-admin-123456" \
-H "Content-Type: application/json" \
-d '
{
 "model":"gpt-5",
 "messages":[
   {
    "role":"user",
    "content":"hello"
   }
```

```
  ]
}'
```

□□□□□□□□

LiteLLM v1.87 □ Dashboard□

□□□

```
http://□□□IP:4000/ui
```

□□□

```
http://192.168.4.21:4000/ui
```

□□□

```
Master Key:
sk-litellm-admin-123456
```

□□□

□□□□□□□□□□□□□□□

□□□ PostgreSQL□

```
docker run \
-d \
--name litellm \
-p 4000:4000 \
-e STORE_MODEL_IN_DB=True \
docker.litellm.ai/berriai/litellm:1.87.0
```

□□□

□□□

- □□□□
- □□□□
- □□□□□
- □ Key □□

□□□□□□

□□□□□□□□□□□ Dify / RAGFlow□

□□□□□□□□□

- Dify
- RAGFlow
- Keycloak
- PostgreSQL
- Docker Compose

LiteLLM □□□□□□□ AI Gateway□

□□□

```
                Dify
                  |
                  |
              OpenAI API
                  |
                  v
             LiteLLM 4000
                  |
    +-------------+-------------+
    |             |             |
OpenAI       Gemini       Ollama
    |
Azure OpenAI
```

Dify □□□□

□□□□□

```
OpenAI Compatible
```

Base URL:

```
http://litellm:4000/v1
```

API Key:

```
sk-litellm-admin-123456
```

□□□

□□□□□□

□□□□□□□□□□□□□□□

```
Ubuntu 24.04
        |
Docker
        |
Nginx HTTPS
        |
LiteLLM 1.87
        |
PostgreSQL 16
        |
Keycloak OAuth
```

□□□

- Redis□□□□
- Prometheus□□□□
- Grafana□ Dashboard□

□□□□□

```
LITELLM_MASTER_KEY=sk-litellm-admin-123456
```

□□ LiteLLM Dashboard □□□□□ Master Key□□

□□□□□□□

---

1. □□ .env

□□□

```
cd /opt/litellm
vi .env
```

□□□□□

```
LITELLM_MASTER_KEY=sk-litellm-admin-9f83a7d2c1e54b6a

OPENAI_API_KEY=sk-xxxx

GEMINI_API_KEY=xxxx
```

□□□□□□□

```
openssl rand -hex 32
```

□□□

```
a9d7f3e8b4c1d2e6f7a8b9c0d1e2f3456789abcd
```

□□□

```
LITELLM_MASTER_KEY=sk-a9d7f3e8b4c1d2e6f7a8b9c0d1e2f3456789abcd
```

---

2. □□ LiteLLM

□□□□□□□□□□□□□□□□

```
docker compose down
```

```
docker compose up -d
```

□□□

```
docker compose restart litellm
```

□□ `.env` □ compose □□□□□

```
docker compose up -d --force-recreate litellm
```

□□□□□□□□□□□

---

3. □□□ Key

□□□□

```
docker exec litellm env | grep LITELLM_MASTER_KEY
```

□□□□

```
LITELLM_MASTER_KEY=sk-litellm-admin-9f83a7d2c1e54b6a
```

---

4. Dashboard□□□

□□□

```
http://□□□IP:4000/ui
```

□□□

```
Master Key:
sk-litellm-admin-9f83a7d2c1e54b6a
```

□□□

注意：不要只改 `.env` 后 restart

很多人这里容易踩坑。

例如：

```
vi .env
docker compose restart litellm
```

有时不会重新加载环境变量，因为容器已经存在。

确认：

```
docker inspect litellm | grep LITELLM_MASTER_KEY
```

如果还是旧值，需要：

```
docker compose up -d --force-recreate litellm
```

生产环境建议

你当前：

```
LITELLM_MASTER_KEY=sk-litellm-admin-123456
```

属于测试密码。

建议：

```
LITELLM_MASTER_KEY=sk-prod-$(openssl rand -hex 32)
```

并且：

- Master Key 只用于管理 Dashboard
- Dify/RAGFlow 使用单独生成的 API Key
- 不要把 Master Key 放到应用配置里

后续如果接 Dify，可以在 LiteLLM Dashboard 创建：

```
Keys → Generate Key
```

例如：

```
sk-dify-prod-xxxx
```

然后 Dify 使用这个，而不是 Master Key。

你的 LiteLLM 已经部署好了，现在添加阿里百炼（DashScope / Model Studio）的 **qwen3.7-max** 很简单。

LiteLLM 对百炼支持 OpenAI Compatible 接口，配置方式和 OpenAI 类似，只需要指定：

- `model`
- `api_base`
- `api_key`

阿里云百炼 qwen3.7-max 提供 OpenAI 兼容接口，模型 ID 为：

```
qwen3.7-max
```

对应接口地址可以使用百炼 OpenAI Compatible Endpoint。(AlibabaCloud)

---

1. 修改 .env

你现在：

```
cat /opt/litellm/.env
```

类似：

```
LITELLM_MASTER_KEY=sk-litellm-admin-123456


OPENAI_API_KEY=sk-xxxx


GEMINI_API_KEY=xxxx
```

增加：

```
DASHSCOPE_API_KEY=sk-xxxxxxxxxxxxxxxx
```

□□□

```
LITELLM_MASTER_KEY=sk-litellm-admin-123456

OPENAI_API_KEY=sk-xxxx

GEMINI_API_KEY=xxxx

DASHSCOPE_API_KEY=sk-abcd123456789
```

□□ Key □□

□□□□□□□□ → API Key □□

□□□

---

2. □□ LiteLLM config.yaml

□□□□□□□

```
/opt/litellm/config.yaml
```

□□□

```
model_list:

  - model_name: gpt-5
    litellm_params:
      model: openai/gpt-5
      api_key: os.environ/OPENAI_API_KEY


  - model_name: gemini-3.5-flash
    litellm_params:
      model: gemini/gemini-3.5-flash
      api_key: os.environ/GEMINI_API_KEY


  # □□□□ qwen3.7-max
```

```
  - model_name: qwen3.7-max
    litellm_params:
      model: openai/qwen3.7-max
      api_base: https://dashscope.aliyuncs.com/compatible-mode/v1
      api_key: os.environ/DASHSCOPE_API_KEY
```

LiteLLM □ `model_name` □□□□ Dify/RAGFlow □□□□ `litellm_params.model` □□□□□□□□□ (docs.litellm.com.cn)

---

3. □□ LiteLLM

□□ `.env` □□□□

□□□ restart□

□□□

```
cd /opt/litellm

docker compose up -d --force-recreate litellm
```

□□□

```
docker compose logs -f litellm
```

---

4. □□□□□□□□

□□□

```
curl http://127.0.0.1:4000/v1/models \
-H "Authorization: Bearer sk-litellm-admin-123456"
```

□□□□□

```
{
 "data":[
   {
```

```
      "id":"gpt-5"
    },
    {
      "id":"gemini-3.5-flash"
    },
    {
      "id":"qwen3.7-max"
    }
  ]
}
```

---

5. □□ qwen3.7-max

```
curl http://127.0.0.1:4000/v1/chat/completions \
-H "Authorization: Bearer sk-litellm-admin-123456" \
-H "Content-Type: application/json" \
-d '
{
 "model":"qwen3.7-max",
 "messages":[
   {
    "role":"user",
    "content":"□□□□□□□□□"
   }
 ]
}'
```

□□□□

```
{
 "choices":[
   {
    "message":{
      "content":"..."
    }
   }
 ]
```

```
}
```

---

6. □ Dify □□

□□□ Dify □ LiteLLM□□□□□□□□□□

□□□

```
- model_name: qwen-max
  litellm_params:
    model: openai/qwen3.7-max
    api_base: https://dashscope.aliyuncs.com/compatible-mode/v1
    api_key: os.environ/DASHSCOPE_API_KEY
```

□□ Dify □□□□

```
qwen-max
```

□□□□

```
qwen3.7-max
↓
qwen4-max
↓
DeepSeek
```

Dify □□□□□

---

7. □□□□□□□□□

```
              Dify
                |
                |
          OpenAI API□□
                |
                |
```

```
            LiteLLM 1.87
                  |
      +-----------+------------+
      |           |            |
   GPT-5     Gemini       Qwen3.7-Max
   OpenAI    Google       阿里百炼
```

这正适合你之前的 Dify + RAGFlow + 企业知识库架构。

另外注意：你之前日志里 `/v1/responses` 是 OpenAI Responses API，如果后续用 Codex/Agent 类客户端调用 qwen3.7-max，需要确认该模型端点对 Responses API 的兼容性；Dify 通常走 `/v1/chat/completions`，问题会少一些。

Revision #5
Created 2026-07-21 09:36:37 UTC by Admin
Updated 2026-07-21 09:53:36 UTC by Admin