

Python 在社会化问答网站中的应用

杨昆

y@zhihu.com

知乎 是什么

社会化问答网络

2010年12月20日上线


2011年1月26日公测

2011年9月2日 iPhone

客户端上线

The screenshot displays the Zhihu website interface. At the top, there is a navigation bar with the Zhihu logo, a search bar, and links for '提问' (Ask), '首页' (Home), '话题' (Topics), '消息' (Messages), and '管理' (Management). The main content area shows a list of questions and answers under the '最新动态' (Latest) section. The first question is '如何评价迈克尔·巴拉克的职业生涯?' (How do you evaluate Michael Barack's career?), answered by 程瀚. The second question is '为什么总有拿美国民主对比中国现状的论调?' (Why is there always a tone of comparing American democracy to China's current situation?), answered by 涛吴, 张佳玮, and 高涵琦. The third question is '本·阿弗莱克的新片《Argo》有何看点,有何不足?' (What are the highlights and shortcomings of Ben Affleck's new movie 'Argo?'), answered by Raymond Wang. The fourth question is '你看好YY语音一类产品吗?如何看待该类产品在线上教育领域的前景?' (Do you have a good view on YY voice products? How do you view the prospects of such products in the online education field?), answered by 周源. The right sidebar contains a '中国经济 迫切十问' (Ten urgent questions about China's economy) section, a '话题广场' (Topic Plaza) section, and a '你可能感兴趣的人和话题' (People and topics you may be interested in) section.

知乎 技术选型



Python



- 快速开发，第三方库丰富
- 社区活跃，用户多
- 语法简单易学，开发人员一周左右就能上手，我们自己也很熟悉

Tornado



- 简单轻量，学习成本低
- 支持异步，天然适合做实时 Comet 应用
- FriendFeed 的成熟案例，Facebook 的社区支持

SQLAlchemy



- 自带连接池
- 安全，防止 SQL 注入
- 灵活可扩展性强

知乎 技术演进



石器时代

2010/10 - 2011/03

2 machines
4k+ commits
5 developers



青铜时代

2011/4 - 2011/10

5-10 machines
1w+ commits
8 developers



蒸汽时代

2011/11 - now

15+ machines
4w+ commits
20+ developers



信息时代

now - future

?

石器时代

- Lean Startup
 - 低成本
 - 快速开发
 - 专注核心逻辑，外包周边业务
- 访问量增长带来的问题
 - 性能问题
 - 网络延迟

石器时代

- 改进
 - 硬件和网络环境改善
- 技术方案
 - MySQL master/slave
 - HAProxy 做热备
 - Tornado 解决高并发问题

Tornado Async Programming Tornado

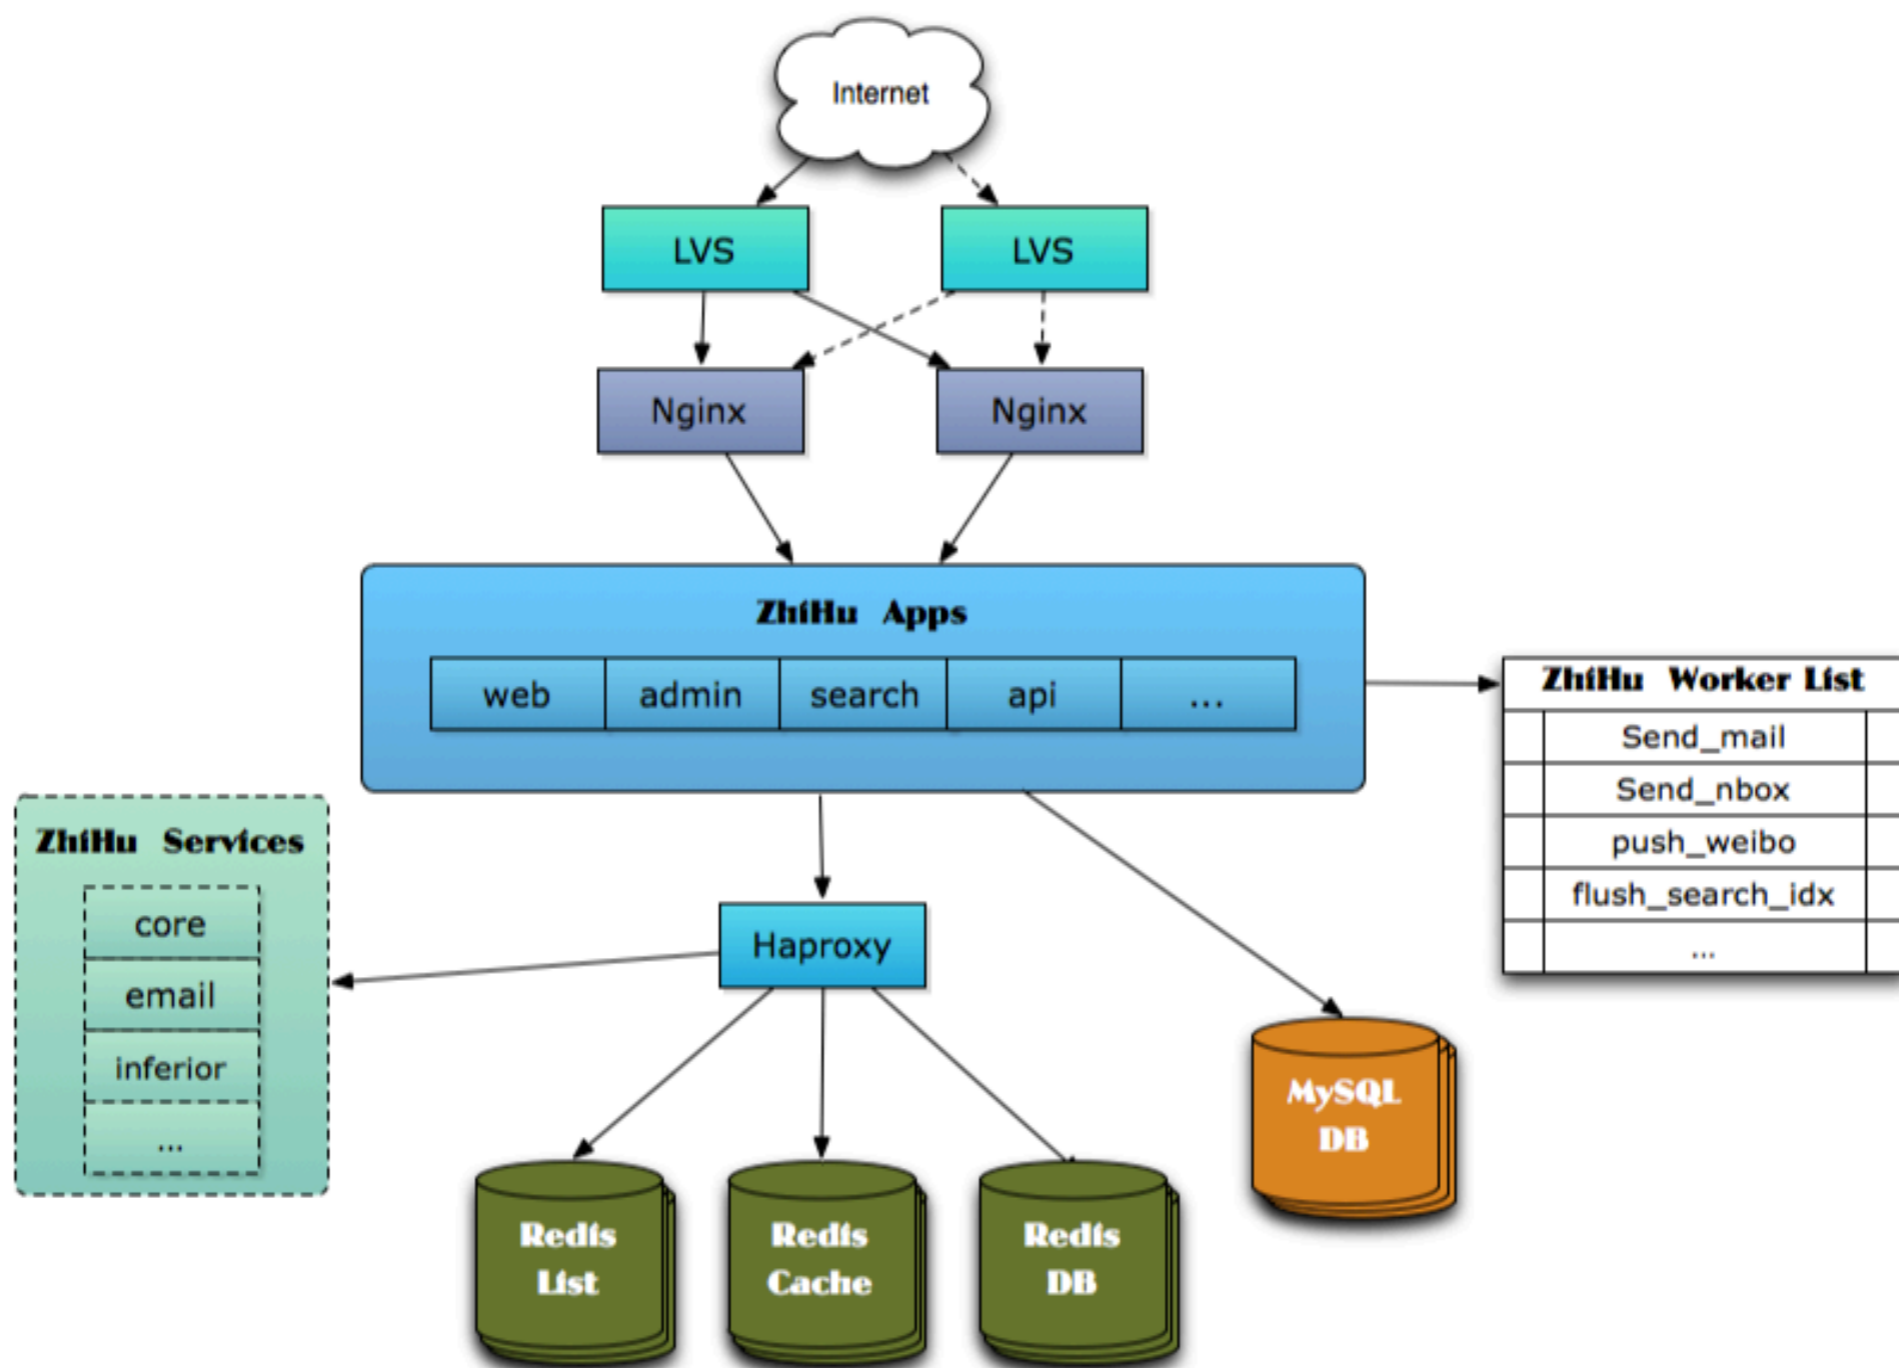
- 高性能
 - 非阻塞，使用 `epoll` 或者 `kqueue`
 - 支持同时数千个持久连接
- Async Example

```
class MyHandler(RequestHandler):  
    @asynchronous  
    @gen.engine  
    def get(self):  
        http_client = AsyncHTTPClient()  
        response = yield gen.Task(http_client.fetch,  
"http://example.com")  
        do_something_with_response(response)  
        self.render("template.html")
```

青铜时代

- 代码行数和访问量急剧膨胀
- 重构和改进
 - 架构
 - 程序逻辑
 - 存储
 - 运维和服务质量

🔨 青铜时代：架构



青铜时代：程序逻辑

- 问题
 - 急剧增长的业务需求
 - 数据库承担大量的计算任务
 - Tornado 遇到瓶颈
- 解决方案
 - Tornado
 - Celery
 - Redis

Redis 在知乎的使用



- 为什么用 Redis
 - 简单易用
 - 功能强大
 - 性能强大

Redis 在知乎的使用



- Redis 作为 Cache
 - 为什么不用 Memcached?
- Redis 作为 MQ
 - Redis 和 Celery 配合
- Redis 作为存储
 - Feed/Notification 中使用 Redis 的 sorted set 结构
 - 话题组织结构使用 Redis 的 set 结构

Redis 在知乎的使用



- Redis 数据带来的问题
 - 数据量大，单个节点放不下
 - 单点访问压力大
- Redis Sharding
 - open source @ <https://github.com/zhihu/redis-shard>
 - consistent hashing, duck-typing API

Redis 在知乎的使用



redis-py:

```
>>> client = redis.Redis()
>>> client.set('test', 1)
>>> client.zadd('testset', 'first', 1)
>>> client.zadd('testset', 'second', 2)
>>> client.zrange('testset', 0, -1)
```

redis-shard:

```
>>> client = RedisShardAPI(servers)
>>> client.set('test', 1)
>>> client.zadd('testset', 'first', 1)
>>> client.zadd('testset', 'second', 2)
>>> client.zrange('testset', 0, -1)
```

青铜时代：运维

- Cobbler 统一安装配置操作系统
- Puppet 统一管理服务器
- Nagios/Cacti 监控报警



蒸汽时代

- 问题
 - 代码量大
 - 逻辑复杂，紧耦合
 - 调试难，开发难
 - 开发周期不一致

蒸汽时代

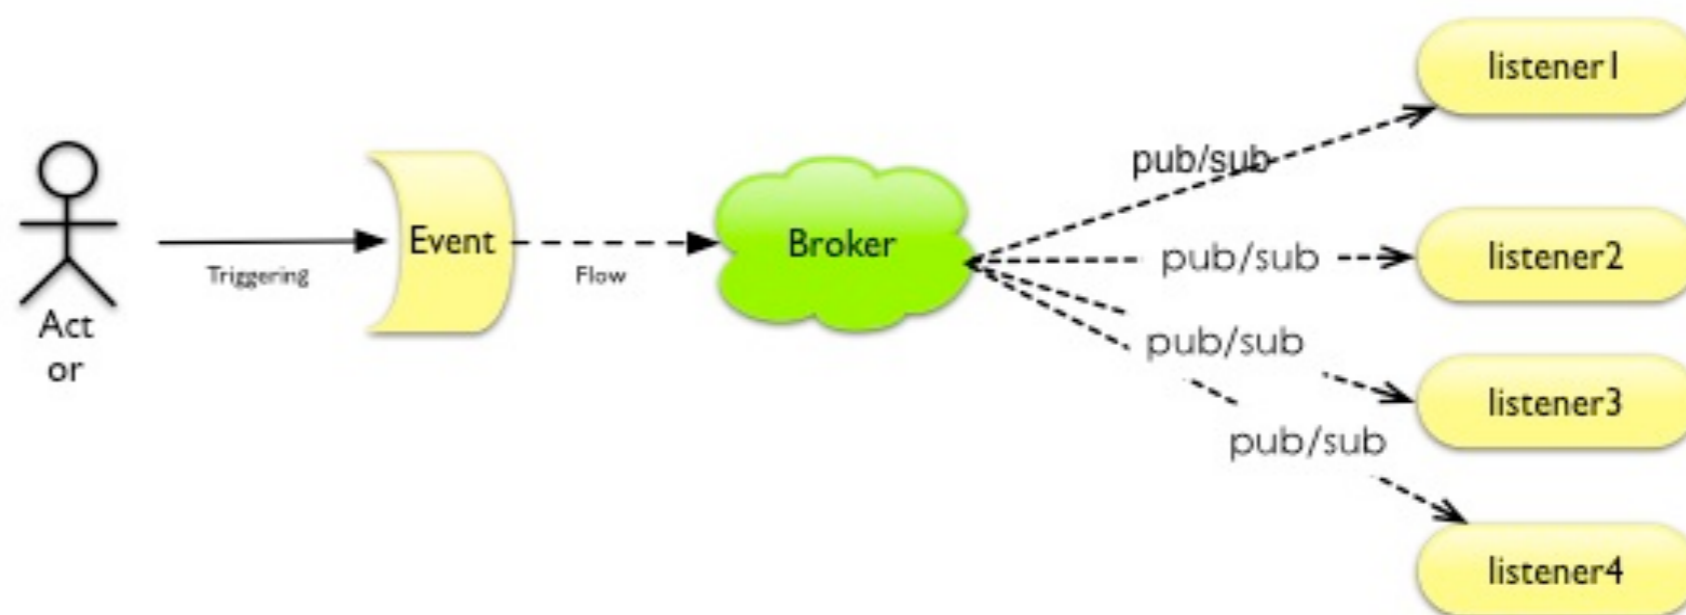
- 解决方案
 - 切分服务 \longleftrightarrow 代码量大, 开发周期不一致
 - 消息服务: Sink \longleftrightarrow 逻辑复杂, 紧耦合
 - 日志收集: Kids \longleftrightarrow 调试难
 - 虚拟开发环境: Hobox \longleftrightarrow 开发难

蒸汽时代：服务切分粒度

- 按粒度切分服务
 - Application
 - Service
 - Library (egg)

🔥 蒸汽时代: Sink

- 特性
 - Pub/Sub
 - Session



蒸汽时代: Sink

- 生产

```
connect("sink").publish("activity.{0}".format(data_dict["Action"]),  
                        json.dumps(data_dict))
```

- 消费

```
@on("activity.answer_create")  
def send_notification():  
    ...
```

蒸汽时代: Kids

- 特性
 - 树形拓扑结构
 - 支持实时订阅 (Redis 协议)
 - 集中存储
- 记录
 - logging handler



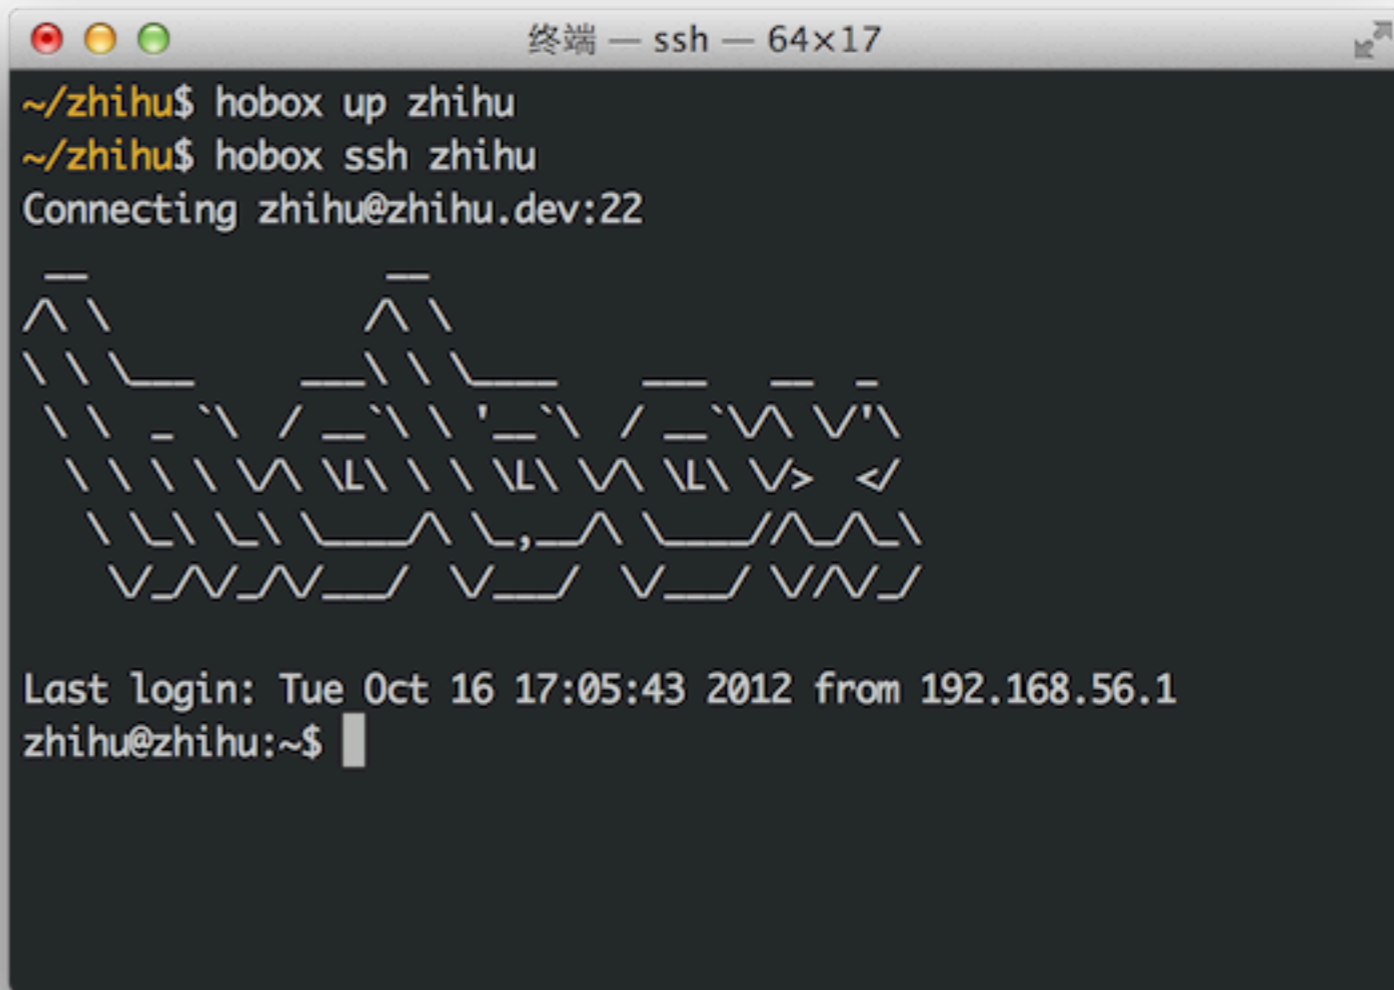
蒸汽时代: Kids

```
default zhihu.web "[I 121016 19:37:39 web:1393] 200 GET /question/20535466 (59.151.98.135) 228.64ms"
default zhihu.web "[I 121016 19:37:39 web:1393] 200 GET /question/19918637 (123.125.71.98) 74.99ms"
default zhihu.web "[I 121016 19:37:39 web:1393] 200 GET /question/20530900?nr=1&noti_id=29400937 (59.41.223.18) 321.82ms"
default zhihu.web "[I 121016 19:37:39 web:1393] 200 GET /question/20177417 (122.65.182.59) 336.22ms"
default zhihu.web "[I 121016 19:37:40 web:1393] 200 GET /register/interests (222.36.180.202) 141.99ms"
default zhihu.web "[I 121016 19:37:40 web:1393] 200 POST /answer/draft/save (123.126.50.24) 43.40ms"
default zhihu.web "[I 121016 19:37:40 web:1393] 200 GET /login?next=http%3A%2F%2Fwww.zhihu.com%2Ftopics (218.26.159.85) 15.96ms"
default zhihu.web "[I 121016 19:37:40 web:1393] 200 GET /reviewq/worker (223.202.8.69) 1249.28ms"
default zhihu.web "[I 121016 19:37:40 web:1393] 200 POST /question/uninterest (110.90.218.79) 95.77ms"
default zhihu.web "[I 121016 19:37:40 web:1393] 200 GET /topic/Google-Chrome (183.5.143.44) 394.34ms"
default zhihu.web "[I 121016 19:37:40 web:1393] 302 POST /login (117.137.253.42) 209.15ms"
warning zhihu.web "[W 121016 19:37:40 web:1393] 404 GET /question/static/img/minisite2-banner.png (198.144.104.14) 21.50ms"
default zhihu.web "[I 121016 19:37:40 web:1393] 200 GET /question/20405116 (118.165.55.249) 215.51ms"
default zhihu.web "[I 121016 19:37:41 web:1393] 200 GET /trending (202.125.231.123) 87.82ms"
default zhihu.web "[I 121016 19:37:41 web:1393] 200 GET /people/ohjmy/json (120.193.62.160) 40.69ms"
default zhihu.web "[I 121016 19:37:41 web:1393] 200 POST / (211.140.18.122) 657.20ms"
default zhihu.web "[I 121016 19:37:41 web:1393] 200 GET /question/19896725 (221.11.11.235) 624.14ms"
default zhihu.web "[I 121016 19:37:41 web:1393] 200 POST / (1.252.93.142) 535.22ms"
default zhihu.web "[I 121016 19:37:41 web:1393] 200 GET /people/ZHUH_/logs (111.172.58.123) 778.25ms"
default zhihu.web "[I 121016 19:37:41 web:1393] 200 GET / (124.78.95.126) 14.08ms"
```

Navigation: [Refresh] [Play] [Delete] [Search: /] [Filter: default error warning]

🔪 蒸汽时代: Hobox

- 特性
 - 虚拟化开发环境
 - 一键安装

A terminal window titled "终端 — ssh — 64x17" showing the execution of Hobox commands. The prompt is ~/zhihu\$. The first command is 'hobox up zhihu'. The second command is 'hobox ssh zhihu', which outputs 'Connecting zhihu@zhihu.dev:22'. This is followed by a large ASCII art graphic of a steam locomotive. Below the art, it says 'Last login: Tue Oct 16 17:05:43 2012 from 192.168.56.1' and the prompt 'zhihu@zhihu:~\$' with a cursor.

```
~/zhihu$ hobox up zhihu
~/zhihu$ hobox ssh zhihu
Connecting zhihu@zhihu.dev:22

      ^      ^      ^      ^      ^
     / \    / \    / \    / \    / \
    /   \  /   \  /   \  /   \  /   \
   /     \/     \/     \/     \/     \
  /      /      /      /      /      /
 /        /        /        /        /
/          /          /          /          /
 \        /        /        /        /
  \      /      /      /      /      /
   \     /     /     /     /     /     \
    \   /    \   /    \   /    \   /    \
     \ /      \ /      \ /      \ /      \
      ^      ^      ^      ^      ^

Last login: Tue Oct 16 17:05:43 2012 from 192.168.56.1
zhihu@zhihu:~$
```



信息时代

- 展望未来
 - ZDB: 数据中间层
 - CI: 持续集成
 -

知乎 Q&A

We are hiring **P**irate!

<http://www.zhihu.com/jobs>