Exercise 5.2: Explore API Calls

1. One way to view what a command does on your behalf is to use **strace**. In this case, we will look for the current endpoints, or targets of our API calls. Install the tool, if not present.

   ```
   student@cp:~$ sudo apt-get install -y strace
   ```

   ```
   student@cp:~$ kubectl get endpoints
   ```

   ```
   1  NAME          ENDPOINTS          AGE
   2  kubernetes    10.128.0.3:6443    3h
   ```

2. Run this command again, preceded by **strace**. You will get a lot of output. Near the end you will note several **openat** functions to a local directory, `/home/student/.kube/cache/discovery/k8scp_6443`. If you cannot find the lines, you may want to redirect all output to a file and `grep` for them. This information is cached, so you may see some differences should you run the command multiple times. As well your IP address may be different.

   ```
   student@cp:~$ strace kubectl get endpoints
   ```

   ```
   1  execve("/usr/bin/kubectl", ["kubectl", "get", "endpoints"], [/*....
   2  ....
   3  openat(AT_FDCWD, "/home/student/.kube/cache/discovery/k8scp_6443..
   4  <output_omitted>
   ```

3. Change to the parent directory and explore. Your endpoint IP will be different, so replace the following with one suited to your system.

   ```
   student@cp:~$ cd /home/student/.kube/cache/discovery/
   ```

   ```
   student@cp:~/.kube/cache/discovery$ ls
   ```

   ```
   1  k8scp_6443
   ```

   ```
   student@cp:~/.kube/cache/discovery$ cd k8scp_6443/
   ```

4. View the contents. You will find there are directories with various configuration information for kubernetes.

   ```
   student@cp:~/.kube/cache/discovery/k8scp_6443$ ls
   ```

   ```
   1  admissionregistration.k8s.io  certificates.k8s.io        node.k8s.io
   2  apiextensions.k8s.io          coordination.k8s.io        policy
   3  apiregistration.k8s.io        crd.projectcalico.org      rbac.authorization.k8s.io
   4  apps                          discovery.k8s.io           scheduling.k8s.io
   5  authentication.k8s.io         events.k8s.io              servergroups.json
   6  authorization.k8s.io          extensions                 storage.k8s.io
   7  autoscaling                   flowcontrol.apiserver.k8s.io  v1
   8  batch                         networking.k8s.io
   ```

5. Use the find command to list out the subfiles. The prompt has been modified to look better on this page.

   ```
   student@cp:./k8scp_6443$ find .
   ```

```
1   .
2   ./certificates.k8s.io
3   ./certificates.k8s.io/v1
4   ./certificates.k8s.io/v1/serverresources.json
5   ./certificates.k8s.io/v1beta1
6   ./certificates.k8s.io/v1beta1/serverresources.json
7   ./apiregistration.k8s.io
8   ./apiregistration.k8s.io/v1
9   ./apiregistration.k8s.io/v1/serverresources.json
10  <output_omitted>
```

6. View the objects available in version 1 of the API. For each object, or kind:, you can view the verbs or actions for that object, such as create seen in the following example. Note the prompt has been truncated for the command to fit on one line. Some are HTTP verbs, such as `GET`, others are product specific options, not standard HTTP verbs. The command may be **python**, depending on what version is installed.

```
student@cp:.$ python3 -m json.tool v1/serverresources.json
```

JSON `serverresources.json`

```
1   {
2       "apiVersion": "v1",
3       "groupVersion": "v1",
4       "kind": "APIResourceList",
5       "resources": [
6           {
7               "kind": "Binding",
8               "name": "bindings",
9               "namespaced": true,
10              "singularName": "",
11              "verbs": [
12                  "create"
13              ]
14          },
15  <output_omitted>
```

7. Some of the objects have `shortNames`, which makes using them on the command line much easier. Locate the `shortName` for endpoints.

```
student@cp:.$ python3 -m json.tool v1/serverresources.json | less
```

JSON `serverresources.json`

```
1   ....
2   {
3   "kind": "Endpoints",
4   "name": "endpoints",
5   "namespaced": true,
6   "shortNames": [
7       "ep"
8   ],
9   "singularName": "",
10  "verbs": [
11      "create",
12      "delete",
13  ....
```

8. Use the `shortName` to view the endpoints. It should match the output from the previous command.

   `student@cp:.$ kubectl get ep`

   ```
   1  NAME           ENDPOINTS         AGE
   2  kubernetes     10.128.0.3:6443   3h
   ```

9. We can see there are 37 objects in version 1 file.

   `student@cp:.$ python3 -m json.tool v1/serverresources.json | grep kind`

   ```
   1      "kind": "APIResourceList",
   2              "kind": "Binding",
   3              "kind": "ComponentStatus",
   4              "kind": "ConfigMap",
   5              "kind": "Endpoints",
   6              "kind": "Event",
   7  <output_omitted>
   ```

10. Looking at another file we find nine more.

    `student@cp:$ python3 -m json.tool apps/v1/serverresources.json | grep kind`

    ```
    1      "kind": "APIResourceList",
    2              "kind": "ControllerRevision",
    3              "kind": "DaemonSet",
    4              "kind": "DaemonSet",
    5              "kind": "Deployment",
    6  <output_omitted>
    ```

11. Delete the `curlpod` to recoup system resources.

    `student@cp:$ kubectl delete po curlpod`

    ```
    1  pod "curlpod" deleted
    ```

12. Take a look around the other files in this directory as time permits.