

Distribute pods evenly across nodes

When deploying a new project with a lot of pods, it might happen that the scheduler decides to put most if not all the pods one node. Obviously this is not great, if that node were to crash, all the pods need to be started again on other nodes. However there is some configuration you can add to prevent this.

In the pod's spec API you can add the following:

```
spec:
  topologySpreadConstraints:
  - maxSkew: <integer>
    topologyKey: <string>
    whenUnsatisfiable: <string>
    labelSelector: <object>
```

An example would be:

```
spec:
  topologySpreadConstraints:
  - maxSkew: 1
    topologyKey: kubernetes.io/hostname
    whenUnsatisfiable: ScheduleAnyway
    labelSelector:
      matchLabels:
        app: <NAME OF APP META LABEL>
```

This should distribute them evenly.

You can read a more detailed and better explanation at the Kubernetes blog: <https://kubernetes.io/blog/2020/05/introducing-podtopologyspread/>.