



# How We Used Kubernetes to Host a CTF Competition

Liron Levin

Ariel Zelivansky

# Who we are

- Ariel Zelivansky / Security Research Leader
  - Vulnerability research on open source projects, CVEs & blog
  - Best security practices for Twistlock platform
- Liron Levin / Twistlock Chief Architect
  - Ph.D. on distributed network algorithms BGU
  - Designs and builds Twistlock platform



# Agenda

1. What is CTF + the challenge
2. Why K8S/Cloud
3. Infrastructure/Engineering
4. Securing the infrastructure
5. Results
6. Key takeouts



# What's a CTF?

- “Capture the flag” challenge
  - Jeopardy style/Attack defense/Wargames (OTW)
- Good for education, conventions



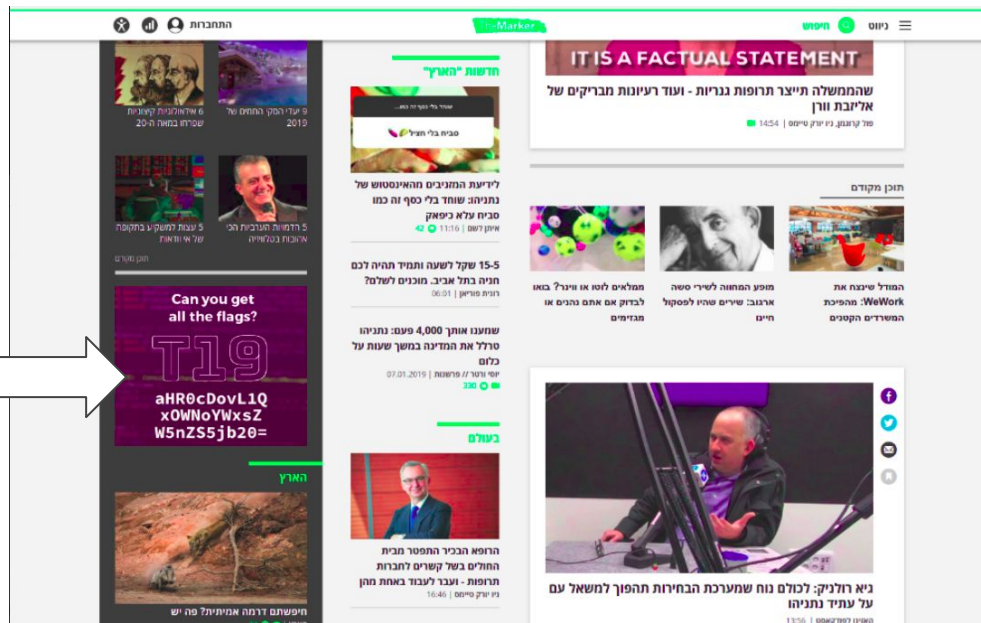
# Twistlock CTF - Why?

- Good PR among security professional
- Find good researchers
- Making challenges forces us to learn a lot
- Fun!



# Advertised!

- Reddit for CTFs (securityCTF)
- Local news sites
- Facebook/Whatsapp groups



# Making it interesting



- Wargame style
- Same machine - multiple challenges!
  - Different users, need to **escalate permissions**
  - Flags hidden as files
- Different challenge subjects - web/scripting, reverse-engineering, Linux internals, modern exploitation...

# The challenge

**THE T19  
CHALLENGE**

Welcome.

Our company developed a unique Linux binary called "cat". We recently discovered that our competitors from the Antivirus company VirusExpress are blocking our cat binary. It is now signed as a **virus**. Word is that you are a badass security researcher. We need you to **infiltrate their server** and **empty their database**.


[Download cat](#)

[Leaderboard](#) [Help](#)







# The challenge



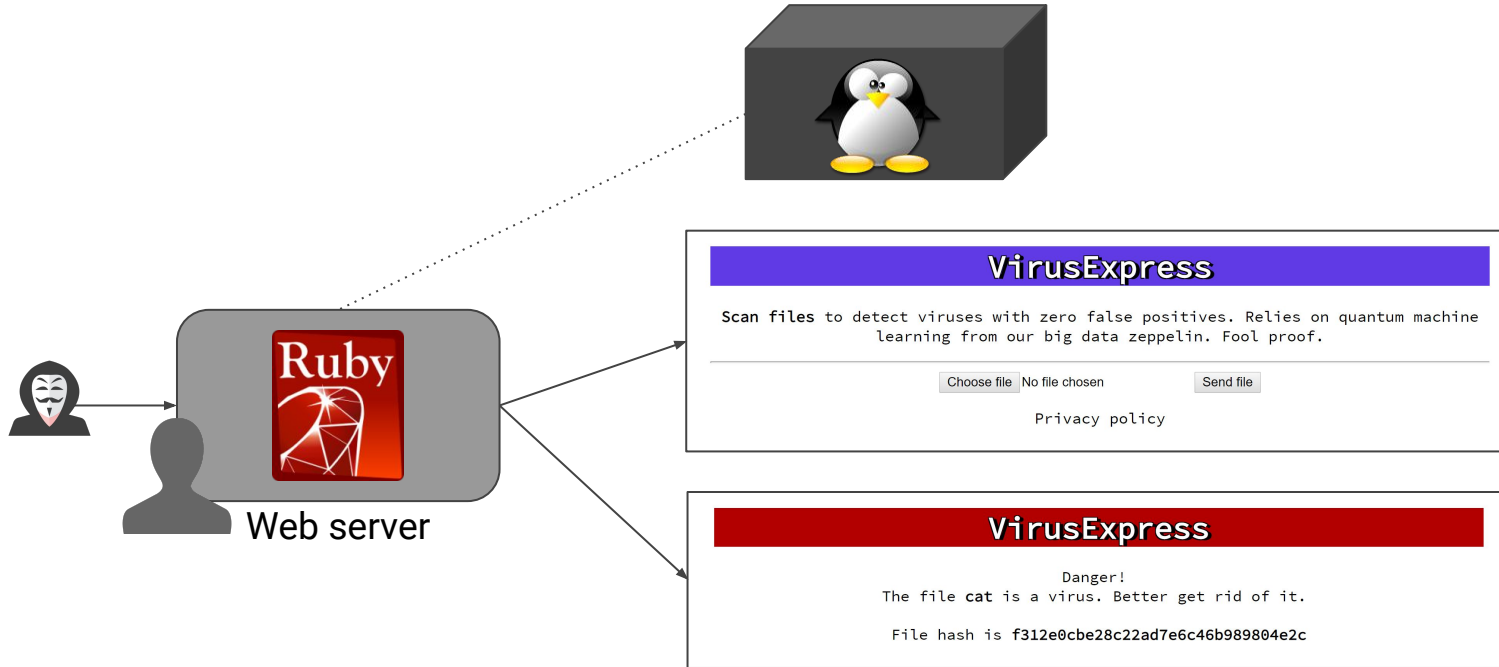
# The challenge




# The challenge



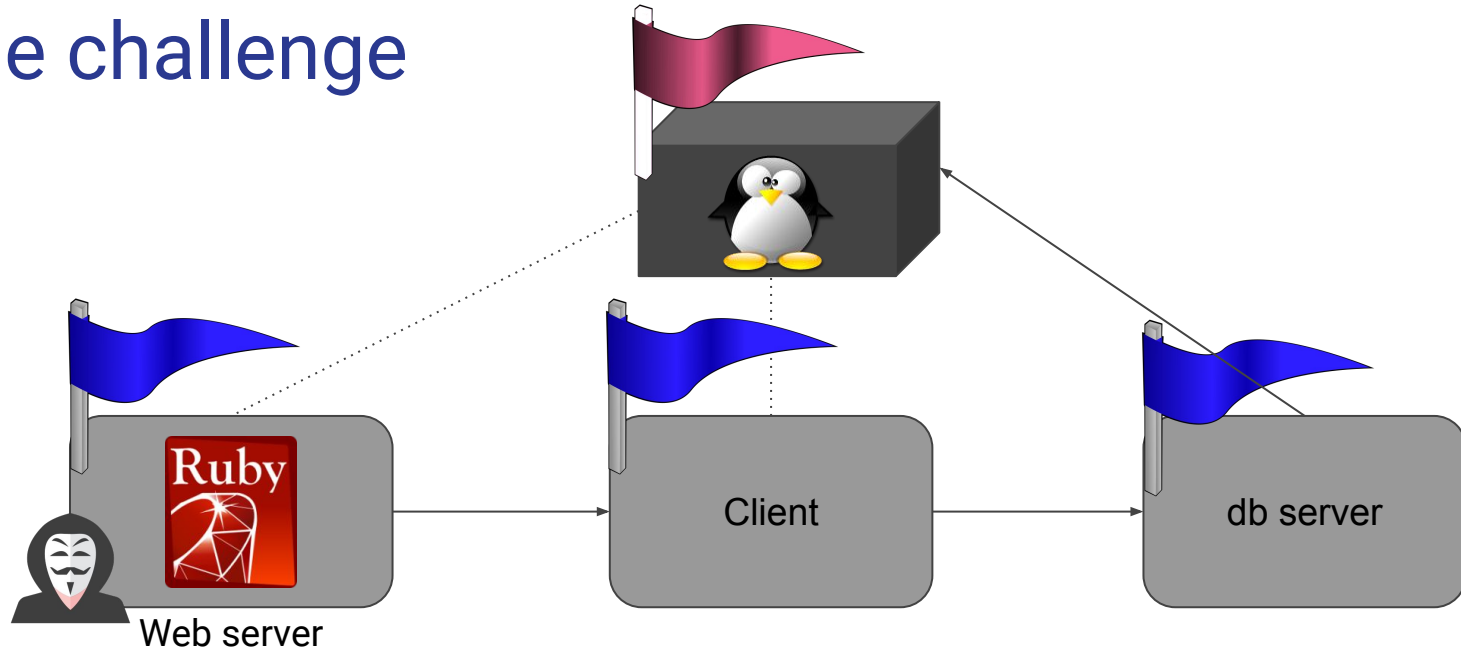
# The challenge



# The challenge



# The challenge



# Why cloud?

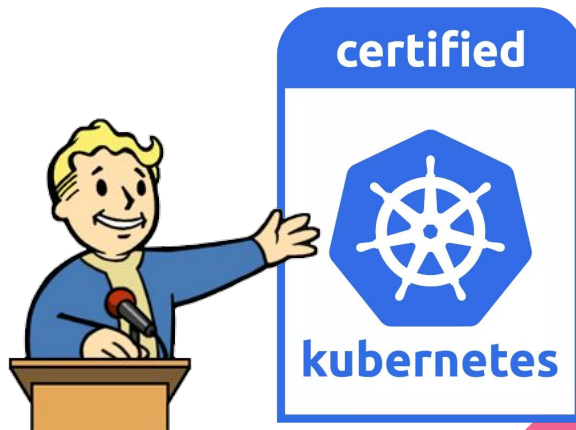
- Machines hosted on our side
  - Impossible to cheat (by reading memory/docker exec)
  - Control and monitor all instances
- Learning opportunity for cloud security





# Why Kubernetes?

- Easy to scale
- Easy to update (hotfix)
- Easy configuration management (configuration as code)
- Good baseline security



# Engineering requirements

1. Simple (but not simplistic)
2. Cheap / Cost effective (time + resources)
3. Reproducible and partially automated\*
4. Secure\* by default



# Solution


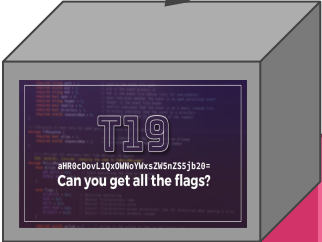
Register



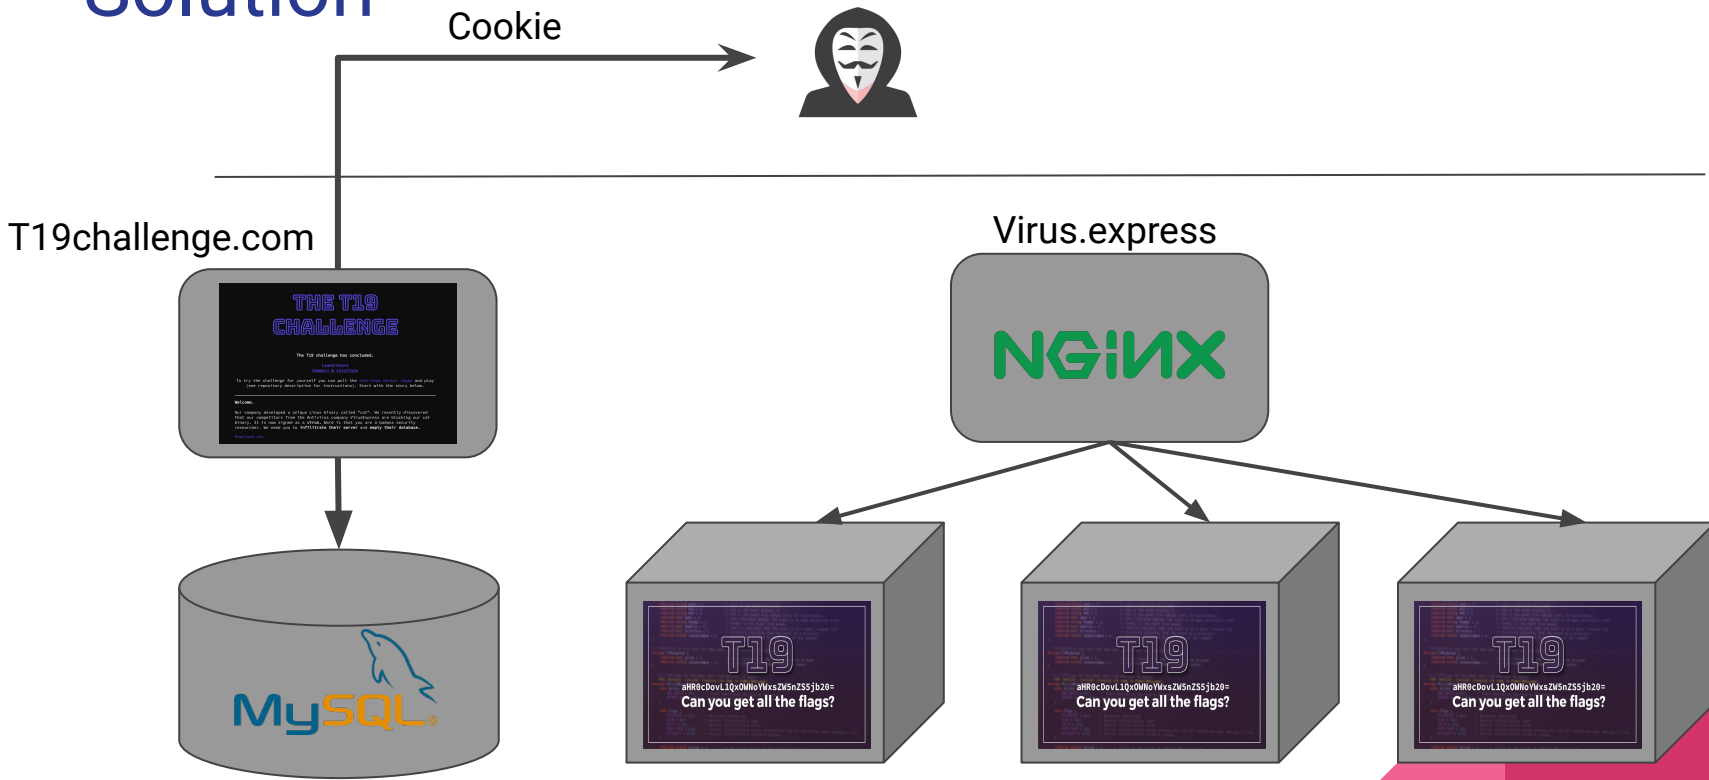
T19challenge.com



Virus.express



# Solution



# Solution



Play

T19challenge.com

Virus.express



# Possible solutions

1. Statically allocate all resources -




# Static


Register



T19challenge.com



Virus.express



# Possible solutions

1. Statically allocate all resources -  
Expensive, does not scale





# Possible solutions

1. Statically allocate all resources -  
Expensive, does not scale
2. On demand allocate pods + services -



# Dynamic


Register



T19challenge.com



Virus.express



# Dynamic


Cookie



T19challenge.com



Virus.express



# Possible solutions

1. Statically allocate all resources -  
Expensive, non-deterministic
2. On demand allocate pods + services -  
Complex, require nginx change + k8s access



# Possible solutions

1. Statically allocate all resources -  
Expensive, non-deterministic
2. On demand allocate pods + services -  
Complex, require nginx change + k8s access
3. Hybrid - statically allocate services + dynamically allocate pods



# Pre-allocated service IPs

Predefined service subnet ( --service-cidr=10.245.0.0/16)

Allocate all services (>k before) before creating pods

kind: Service

apiVersion: v1

metadata:

name: ctf-1

spec:

**clusterIP: 10.245.0.3**

selector:

app: ctf-1

ports:

- protocol: TCP

port: 13337

targetPort: 13337



# How it works



ID	Cluster-ip	
eba871ba9e58739c687e084a68f34500	10.245.0.3	...
76846a1eb5ec91e974831af1baa9e76d	10.245.0.4	...
d88ec62c1ea5b46df814f122a4641a94	10.245.0.5	...
...	10.245.0.5	
...	...	

# The load-balancer

The NGINX logo is displayed in green text on a grey rounded rectangular background.

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: nginx-config
data:
  nginx.conf: |
    http {
      limit_req_zone $binary_remote_addr zone=one:10m rate=1r/s;
      map $cookie_t19userid $backend {
        default ";

    eba871ba9e58739c687e084a68f34500 http://10.245.0.3:13337;
    76846a1eb5ec91e974831af1baa9e76d http://10.245.0.4:13337;
    d88ec62c1ea5b46df814f122a4641a94 http://10.245.0.5:13337;
```



# On demand\* pod allocation

Create pods on demand (or in batches)

```
kind: Deployment
metadata:
  name: ctf-1
  labels:
    app: ctf-1
spec:
  spec:
    containers:
      - name: ctf-1
        image: twistlock/t19
        ports:
          - containerPort: 13337
```



# Freeing unused resources

- Each CTF app takes ~20mb
- We expected ~2k registrations ~40GB RAM
- How do you detect (and shutdown) idle instance?
  - `/var/log/nginx/access.log`




# Security challenges - WHAT IF

- Too many registration (resource exhaustion) - should delete?
- One pod interfere other pods (DOS)
- Attacker breaks out of the pod (container breakout)
- Compromise network assets
- Compromised the cluster (game over)




# Too many registration (should delete?)

- Entry barriers: base64
- Doomsday solution: Captcha



# Local resource exhaustion

- CPU/memory exhaustion (deliberate or accidental)
- Resource abuse \$\$\$ (e.g. cryptomining)



# Local resource exhaustion - solution

- Pod security policy

```
apiVersion: v1
kind: Pod
metadata:
  name: ctf
spec:
  containers:
  - name: ctf-app
    image: twistlock/t19
    resources:
      requests:
        memory: "30Mi"
        cpu: "50m"
      limits:
        memory: "50Mi"
        cpu: "50m"
```



# Host compromised

- Misconfiguration (host mount/secrets)
- CVE-2019-5736 -  
Execution of malicious containers allows for container escape and access to host filesystem



# Container breakout - “solution”

- Classic container - No mounts/secrets - simple app
- Default container profile (no additional LINUX capabilities + seccomp)
- Container optimized OS - read only root partition (CVE-2019-5736 mitigation)
- [Optional] Users
- [Optional] Additional sandbox - Gvisor





# Cluster takeover

- Capturing all the flags in BSidesSF CTF by pwning our infrastructure  
<https://hackernoon.com/capturing-all-the-flags-in-bsidessf-ctf-by-pwning-our-infrastructure-3570b99b4dd0>
- SSRF in Exchange leads to ROOT access in all instances  
<https://hackerone.com/reports/341876>
- Access cloud services (\$\$) or steal sensitive data (images)



# Cluster takeover - mitigations

- Isolated environment (project)
- RBAC
- automountServiceAccountToken: false
- Metadata concealment
- Network policies



# Network policy

kind: NetworkPolicy

spec:

podSelector:

matchLabels:

app: t19

policyTypes:

- Ingress

- Egress

egress:

- to:

- ipBlock:

cidr: 0.0.0.0/0

except:

- 169.254.169.254/32


ingress:

- from:

- podSelector:

matchLabels:

app: t19-nginx



# Challenge conclusion

- 8 participants solved
  - 6 found 4th flag
- Excellent write-ups with solutions
- [Links and finalists](#)
- Challenge coins molded



# Try to solve?

- <http://t19challenge.com/>
- Follow the instructions to run
- Don't cheat and good luck!
- See you in T20?



# Key takeouts

- Good engineering == cost saving
- Good security ...
- Kubernetes is a great platform to host a live CTF
  - Little effort to deploy once built
  - Easy to monitor
  - Easy to scale
  - Hotfix on pods
- Future ideas
  - Networking CTF - more than one container in pod, need to hack via network
  - Attack/defense CTF on Kubernetes

