

Security Breaks: TeamTNT's DockerHub Credentials Leak

 trendmicro.com/en_us/research/22/i/security-breaks-teamtnts-dockerhub-credentials-leak.html

September 12, 2022



Content added to Folio

Cloud

One of our honeypots based on exposed Docker REST APIs showed cybercriminal group TeamTNT's potential attack scenario and leak of container registry credentials for docker-abuse malware. The full version of this research will be presented at the c0c0n XV Hacking and Cyber Security Conference in September 2022.

By: Nitesh Surana September 12, 2022 Read time: (words)

We constantly deploy and study our honeypots to get a view of actively exploited vulnerabilities and misconfigurations on platforms and services that pose cloud security risks. One of these honeypots is based on exposed Docker REST API for analysis from cloud

services providers' and users' perspectives. Upon analyzing the samples, we realized and were able to understand the threat actors' use of container registry features for Docker malware and tactics, techniques, and procedures (TTPs).

Our honeypots showed threat actor TeamTNT were leaking credentials from at least two of their attacker-controlled DockerHub accounts, namely *alpineos* (with over 150,000 pulls) and *sandeep078* (with 200 pulls). We have notified Docker about these accounts.

The account *alpineos* was used in exploitation attempts on our honeypots three times, from mid-September to early October 2021, and we tracked the deployments' IP addresses to their location in Germany. The threat actors were logged in to their accounts on the DockerHub registry and probably forgot to log out. Unless a user is not logged out manually, the header "X-Registry-Auth" stores the credentials.

These DockerHub profiles were actively used to deploy malicious images containing the following:

1. Rootkits
2. Docker escape kits
3. XMRig Monero miners
4. Credential stealers
5. Kinsing malware
6. Kubernetes exploit kits

In July 2021, we published our research on TeamTNT's malicious activities and found evidence of the group infiltrating via the Docker API. As a result, we found 26 unique DockerHub accounts that are either compromised or malicious. Of the two we identified here, the most interesting account for study was the *alpineos* account, which hosted malicious container images with over 150,000 pulls.

Container registries and Docker daemon

Docker is a container services platform that helps developers follow a write-once-run-anywhere (WORA) practice. It's simple to use and is favoured by developers, as a user can write services and deploy applications at great speed. Most importantly, Docker works with any platform.

Container registries are storage and distribution platforms for container images, similar to how codes or programs are hosted on repositories like GitHub. With the right authorization context, one can simply "pull" an image, create a container based on it, and deploy applications. Many container registries such as DockerHub, Amazon Elastic Container Registry (ECR), and Alibaba Container Registry, to name a few, host container images.

When you create a container, the container daemon looks up the image from the container registry by default. In our analysis, we use DockerHub as an example.

```
ubuntu@pikachu-docker:~$ docker run hello-world
Unable to find image 'hello-world:latest' locally
latest: Pulling from library/hello-world
2db29710123e: Pull complete
Digest: sha256:37a0b92b08d4919615c3ee023f7ddb068d12b8387475d64c622ac30f45c29c51
Status: Downloaded newer image for hello-world:latest

Hello from Docker!
This message shows that your installation appears to be working correctly.

To generate this message, Docker took the following steps:
 1. The Docker client contacted the Docker daemon.
 2. The Docker daemon pulled the "hello-world" image from the Docker Hub.
    (amd64)
 3. The Docker daemon created a new container from that image which runs the
    executable that produces the output you are currently reading.
 4. The Docker daemon streamed that output to the Docker client, which sent it
    to your terminal.

To try something more ambitious, you can run an Ubuntu container with:
$ docker run -it ubuntu bash

Share images, automate workflows, and more with a free Docker ID:
https://hub.docker.com/

For more examples and ideas, visit:
https://docs.docker.com/get-started/

ubuntu@pikachu-docker:~$ █
```

Figure 1. A “Hello World”

example of the Docker daemon (dockerd) pulling the image from Docker Hub

If we don't specify the registry, DockerHub is considered by default. Docker provides a feature for developers to create containers on a remote host when the Docker daemon (on the server) is configured to listen over the TCP port, which is port 2375 by default. This makes remote development and deployment easy for developers as it provides an interface to various Docker services like images, containers, networks, and volumes using tools like curl, wget, and docker-cli.

Docker REST API for container creation

Consider a scenario where a new container with an alpine image base (Alpine Linux, a distribution based on musl libc library and BusyBox utilities) is created on remote server 172.31.42.11 via docker REST API. The remote server has the dockerd exposed over TCP port 2375.

```
ubuntu@pikachu-docker:~$ docker -H 172.31.42.11:2375 run -it alpine sh
Unable to find image 'alpine:latest' locally
latest: Pulling from library/alpine
a0d0a0d46f8b: Pull complete
Digest: sha256:e1c082e3d3c45cccac829840a25941e679c25d438cc8412c2fa221cf1a824e6a
Status: Downloaded newer image for alpine:latest
/ # exit
ubuntu@pikachu-docker:~$ █
```

Figure 2. Container created based on alpine image on a remote host

2295	9041.004789	172.31.42.11	172.25.1.2	TCP	76	34206 + 2375	[SYN] Seq=0 Win=65495 Len=0 MSS=65495 SACK_PERM=1 TSval=481606782 TSecr=0 W=128
2296	9041.004850	172.25.1.2	172.31.42.11	TCP	76	2375 + 34206	[SYN, ACK] Seq=0 Ack=1 Win=65160 Len=0 MSS=1460 SACK_PERM=1 TSval=4233902353 TSecr=481606782 W=128
2297	9041.004896	172.31.42.11	172.25.1.2	TCP	68	34206 + 2375	[ACK] Seq=1 Ack=1 Win=65536 Len=0 TSval=481606782 TSecr=4233902353
2298	9041.005297	172.31.42.11	172.25.1.2	HTTP	160	HEAD /_ping HTTP/1.1	
2299	9041.005378	172.25.1.2	172.31.42.11	TCP	68	2375 + 34206	[ACK] Seq=1 Ack=93 Win=65152 Len=0 TSval=4233902353 TSecr=481606782
2300	9041.048126	172.25.1.2	172.31.42.11	HTTP	348	HTTP/1.1 200 OK	
2301	9041.048167	172.31.42.11	172.25.1.2	TCP	68	34206 + 2375	[ACK] Seq=93 Ack=281 Win=65280 Len=0 TSval=481606825 TSecr=4233902396
2302	9041.049187	172.31.42.11	172.25.1.2	HTTP/1.1	1759	POST /v1.41/containers/create HTTP/1.1 , JavaScript Object Notation (application/json)	
2303	9041.049268	172.25.1.2	172.31.42.11	TCP	68	2375 + 34206	[ACK] Seq=281 Ack=1784 Win=64128 Len=0 TSval=4233902397 TSecr=481606826
2304	9041.452739	172.25.1.2	172.31.42.11	HTTP/1.1	320	HTTP/1.1 404 Not Found , JavaScript Object Notation (application/json)	
2305	9041.452795	172.31.42.11	172.25.1.2	TCP	68	34206 + 2375	[ACK] Seq=1784 Ack=533 Win=65408 Len=0 TSval=481607230 TSecr=4233902800
2306	9041.453833	172.31.42.11	172.25.1.2	HTTP	164	GET /v1.41/info HTTP/1.1	
2307	9041.456298	172.25.1.2	172.31.42.11	TCP	68	2375 + 34206	[ACK] Seq=533 Ack=1880 Win=64128 Len=0 TSval=4233902804 TSecr=481607231
2308	9041.559215	172.25.1.2	172.31.42.11	HTTP/1.1	2871	HTTP/1.1 200 OK , JavaScript Object Notation (application/json)	
2309	9041.559257	172.31.42.11	172.25.1.2	TCP	68	34206 + 2375	[ACK] Seq=1880 Ack=3336 Win=63616 Len=0 TSval=481607336 TSecr=4233902907
2310	9041.560110	172.31.42.11	172.25.1.2	HTTP	270	POST /v1.41/images/create?fromImage=alpine&tag=latest HTTP/1.1	
2311	9041.560189	172.25.1.2	172.31.42.11	TCP	68	2375 + 34206	[ACK] Seq=3336 Ack=2882 Win=64128 Len=0 TSval=4233902908 TSecr=481607337
2312	9044.848243	172.25.1.2	172.31.42.11	TCP	340	2375 + 34206	[PSH, ACK] Seq=3336 Ack=2882 Win=64128 Len=272 TSval=4233906196 TSecr=481607337 [TCP segment of a reassembled PDU]
2313	9044.889196	172.31.42.11	172.25.1.2	TCP	68	34206 + 2375	[ACK] Seq=2882 Ack=3608 Win=65536 Len=0 TSval=481610666 TSecr=4233906196
2314	9045.653887	172.25.1.2	172.31.42.11	TCP	145	2375 + 34206	[PSH, ACK] Seq=3608 Ack=2882 Win=64128 Len=77 TSval=4233907002 TSecr=481610666 [TCP segment of a reassembled PDU]
2315	9045.653907	172.31.42.11	172.25.1.2	TCP	68	34206 + 2375	[ACK] Seq=2882 Ack=3685 Win=65536 Len=0 TSval=481611431 TSecr=4233907002
2316	9046.550431	172.25.1.2	172.31.42.11	TCP	260	2375 + 34206	[PSH, ACK] Seq=3685 Ack=2882 Win=64128 Len=192 TSval=4233907898 TSecr=481611431 [TCP segment of a reassembled PDU]
2317	9046.550451	172.31.42.11	172.25.1.2	TCP	68	34206 + 2375	[ACK] Seq=2882 Ack=3877 Win=65408 Len=0 TSval=481612327 TSecr=4233907898

Figure 3. Network server traffic

Looking at the server network traffic log, we can see that when a new container is requested for creation on a remote server, this is the sequence that follows:

1. The client pings the target server (packet 2298) to test if the server is accessible.
2. The server responds with the status code 200 and that it is accessible (packet 2300).
3. The client requests that the server create a container from an image named “alpine” (packet 2302).
4. If the server cannot find the “alpine” image locally, it replies with the status code 404 (packet 2304).
5. The client requests for the server information from `</<version>/info endpoint>` (packet 2306).
6. The server responds to the request with system-wide information (packet 2308).

```
POST /v1.41/containers/create HTTP/1.1
Host: 172.31.42.11:2375
User-Agent: Docker-Client/20.10.7 (linux)
Content-Length: 1527
Content-Type: application/json

{"Hostname":"","Domainname":"","User":"","AttachStdin":true,"AttachStdout":true,"AttachStderr":true,"Tty":true,"OpenStdin":true,"StdinOnce":true,"Env":null,"Cmd":["sh"],"Image":"alpine","Volumes":{},"WorkingDir":"","Entrypoint":null,"OnBuild":null,"Labels":{},"HostConfig":{"Binds":null,"ContainerIDFile":"","LogConfig":{"Type":"","Config":{}},"NetworkMode":"default","PortBindings":{},"RestartPolicy":{"Name":"no","MaximumRetryCount":0},"AutoRemove":false,"VolumeDriver":"","VolumesFrom":null,"CapAdd":null,"CapDrop":null,"CgroupnsMode":"","Dns":[],"DnsOptions":[],"DnsSearch":[],"ExtraHosts":null,"GroupAdd":null,"IpcMode":"","Cgroup":"","Links":null,"OomScoreAdj":0,"PidMode":"","Privileged":false,"PublishAllPorts":false,"ReadonlyRootfs":false,"SecurityOpt":null,"UTSMode":"","UsernsMode":"","ShmSize":0,"ConsoleSize":[0,0],"Isolation":"","CpusShares":0,"Memory":0,"NanoCpus":0,"CgroupParent":"","BlkioWeight":0,"BlkioWeightDevice":[],"BlkioDeviceReadBps":null,"BlkioDeviceWriteBps":null,"BlkioDeviceReadIOps":null,"BlkioDeviceWriteIOps":null,"CpuPeriod":0,"CpuQuota":0,"CpuRealtimePeriod":0,"CpuRealtimeRuntime":0,"CpusetCpus":"","CpusetMems":"","Devices":[],"DeviceCgroupRules":null,"DeviceRequests":null,"KernelMemory":0,"KernelMemoryTCP":0,"MemoryReservation":0,"MemorySwap":0,"MemorySwappiness":-1,"OomKillDisable":false,"PidsLimit":0,"Ulimits":null,"CpuCount":0,"CpuPercent":0,"IOMaximumIOps":0,"IOMaximumBandwidth":0,"MaskedPaths":null,"ReadonlyPaths":null},"NetworkingConfig":{"EndpointsConfig":{},"Platform":null}

HTTP/1.1 404 Not Found
Api-Version: 1.41
Content-Type: application/json
Docker-Experimental: false
OSType: linux
Server: Docker/20.10.7 (linux)
Date: Sun, 17 Oct 2021 09:42:40 GMT
Content-Length: 43

{"message":"No such image: alpine:latest"}
```

Figure 4. Showing 404 status code as alpine image is locally unavailable

```

GET /v1.41/info HTTP/1.1
Host: 172.31.42.11:2375
User-Agent: Docker-Client/20.10.7 (linux)

HTTP/1.1 200 OK
Api-Version: 1.41
Content-Type: application/json
Docker-Experimental: false
Ostype: linux
Server: Docker/20.10.7 (linux)
Date: Sun, 17 Oct 2021 09:42:40 GMT
Transfer-Encoding: chunked

{"ID":"PARS:ADJ:NMPC:KFU7:PJME:IAGN:3K6J:YPET:WASJ:O2VE:MU6E:P2ND","Containers":0,"ContainersRunning":0,"ContainersPaused":0,"ContainersStopped":0,"Images":0,"Driver":"overlay2","DriverStatus":[{"Backing Filesystem":"extfs"},{"Supports d_type":"true"},{"Native Overlay Diff":"false"},{"userxattr":"true"}],"Plugins":{"Volume":["local"],"Network":{"bridge","host","ipvlan","macvlan","null","overlay"},"Authorization":null,"Log":["awslogs","fluentd","gcplogs","gelf","journald","json-file","local","logentries","splunk","syslog"],"MemoryLimit":true,"SwapLimit":true,"KernelMemory":true,"KernelMemoryTCP":true,"CpuCfsPeriod":true,"CpuCfsQuota":true,"CPUShares":true,"CPUSet":true,"PidsLimit":true,"IPV4Forwarding":true,"BridgeNfIptables":true,"BridgeNfIp6tables":true,"Debug":false,"Nfd":23,"OomKillDisable":true,"NGoroutines":35,"SystemTime":"2021-10-17T09:42:40.516990462Z","LoggingDriver":"json-file","GroupDriver":"cgroups","CgroupVersion":"1","NEventsListener":0,"KernelVersion":"5.11.0-1019-aws","OperatingSystem":"Ubuntu 20.04.2 LTS","OSVersion":"20.04","OSType":"linux","Architecture":"x86_64","IndexServerAddress":"https://index.docker.io/v1/","RegistryConfig":{"AllowNonDistributableArtifactsCIDRs":[],"AllowNonDistributableArtifactsHostnames":[],"InsecureRegistryCIDRs":["127.0.0.0/8"],"IndexConfigs":{"docker.io":{"Name":"docker.io","Mirrors":[],"Secure":true,"Official":true},"Mirrors":[],"NCPU":2,"MemTotal":834131488,"GenericResources":null,"DockerRootDir":"/var/lib/docker","HttpProxy":"","HttpsProxy":"","NoProxy":"","Name":"cde8c8bc2089","Labels":[],"ExperimentalBuild":false,"ServerVersion":"20.10.7","Runtimes":{"io.containerd.runc.v2":{"path":"runc"},"io.containerd.runtime.v1.linux":{"path":"runc"},"runc":{"path":"runc"},"DefaultRuntime":"runc"},"Swarm":{"NodeID":"","NodeAddr":"","LocalNodeState":"inactive","ControlAvailable":false,"Error":"","RemoteManagers":null},"LiveRestoreEnabled":false,"Isolation":"","InitBinary":"docker-init"},"ContainerdCommit":{"ID":"7eba5930496d9bbe375fdf71603e610ad737d2b2","Expected":"7eba5930496d9bbe375fdf71603e610ad737d2b2"},"RuncCommit":{"ID":"v1.0.0-g84113ee","Expected":"v1.0.0-g84113ee"},"InitCommit":{"ID":"de40ad0","Expected":"de40ad0"},"SecurityOptions":["name=seccomp,profile=default"],"Warnings":{"WARNING: API is accessible on http://0.0.0.0:2375 without encryption.\n      Access to the remote API is equivalent to root access on the host. Refer\n      to the 'Docker daemon attack surface' section in the documentation for\n      more information: https://docs.docker.com/go/attack-surface/"}}}

```

Figure 5. Returning system-wide information

7. The client requests the server to create a container from the alpine image. The “latest” tag is chosen when no tags are specified.

8. The server responds with the download progress of the alpine image from DockerHub.

```

POST /v1.41/images/create?fromImage=alpine&tag=latest HTTP/1.1
Host: 172.31.42.11:2375
User-Agent: Docker-Client/20.10.7 (linux)
Content-Length: 0
Content-Type: text/plain
X-Registry-Auth: e30=

HTTP/1.1 200 OK
Api-Version: 1.41
Content-Type: application/json
Docker-Experimental: false
Ostype: linux
Server: Docker/20.10.7 (linux)
Date: Sun, 17 Oct 2021 09:42:43 GMT
Transfer-Encoding: chunked

{"status":"Pulling from library/alpine","id":"latest"}
{"status":"Pulling fs layer","progressDetail":{"id":"a0d0a0d46f8b"}

```

Figure 6. Creating a container from

the alpine image

9. The client requests the server to create a container from the now-downloaded alpine image.

```

POST /v1.41/containers/create HTTP/1.1
Host: 172.31.42.11:2375
User-Agent: Docker-Client/20.10.7 (linux)
Content-Length: 1527
Content-Type: application/json

{"Hostname":"","Domainname":"","User":"","AttachStdin":true,"AttachStdout":true,"AttachStderr":true,"Tty":true,"OpenStdin":true,"StdinOnce":true,"Env":null,"Cmd":["sh","-c","image=alpine"],"Volumes":{},"WorkingDir":"","Entrypoint":null,"OnBuild":null,"Labels":{},"HostConfig":{"Binds":null,"ContainerIDFile":"","LogConfig":{"Type":"","Config":{},"NetworkMode":"default","PortBindings":{},"RestartPolicy":{"Name":"no","MaximumRetryCount":0},"AutoRemove":false,"VolumeDriver":"","VolumesFrom":null,"CapAdd":null,"CapDrop":null,"CapGroupsMode":"","Dns":{},"DnsOptions":{},"DnsSearch":{},"ExtraHosts":null,"GroupAdd":null,"IpcMode":"","Cgroup":"","Links":null,"OomScoreAdj":0,"PidMode":"","Privileged":false,"PublishAllPorts":false,"ReadOnlyRootfs":false,"SecurityOpt":null,"UTSMode":"","UsernsMode":"","ShmSize":0,"ConsoleSize":0},"Isolation":"","CpuShares":0,"Memory":0,"NanoCpus":0,"CgroupParent":"","BlkioWeight":0,"BlkioWeightDevice":{},"BlkioDeviceReadBps":null,"BlkioDeviceWriteBps":null,"BlkioDeviceReadIOps":null,"BlkioDeviceWriteIOps":null,"CpuPeriod":0,"CpuQuota":0,"CpuRealtimePeriod":0,"CpuRealtimeRuntime":0,"CpusetCpus":"","CpusetMems":"","Devices":{},"DeviceCgroupRules":null,"DeviceRequests":null,"KernelMemory":0,"KernelMemoryTCP":0,"MemoryReservation":0,"MemorySwap":0,"MemorySwappiness":-1,"OomKillDisable":false,"PidsLimit":0,"Ulimits":null,"CpuCount":0,"CpuPercent":0,"IOMaximumIOps":0,"IOMaximumBandwidth":0,"MaskedPaths":null,"ReadOnlyPaths":null},"NetworkingConfig":{"EndpointsConfig":{},"Platform":null}}
HTTP/1.1 201 Created
Api-Version: 1.41
Content-Type: application/json
Docker-Experimental: false
Ostype: linux
Server: Docker/20.10.7 (linux)
Date: Sun, 17 Oct 2021 09:42:49 GMT
Content-Length: 88

{"id":"e21f25e11b5504e639aee75f883fd370bdcc1cb2b5e26125dbbf5739c86dec","warnings":{}}

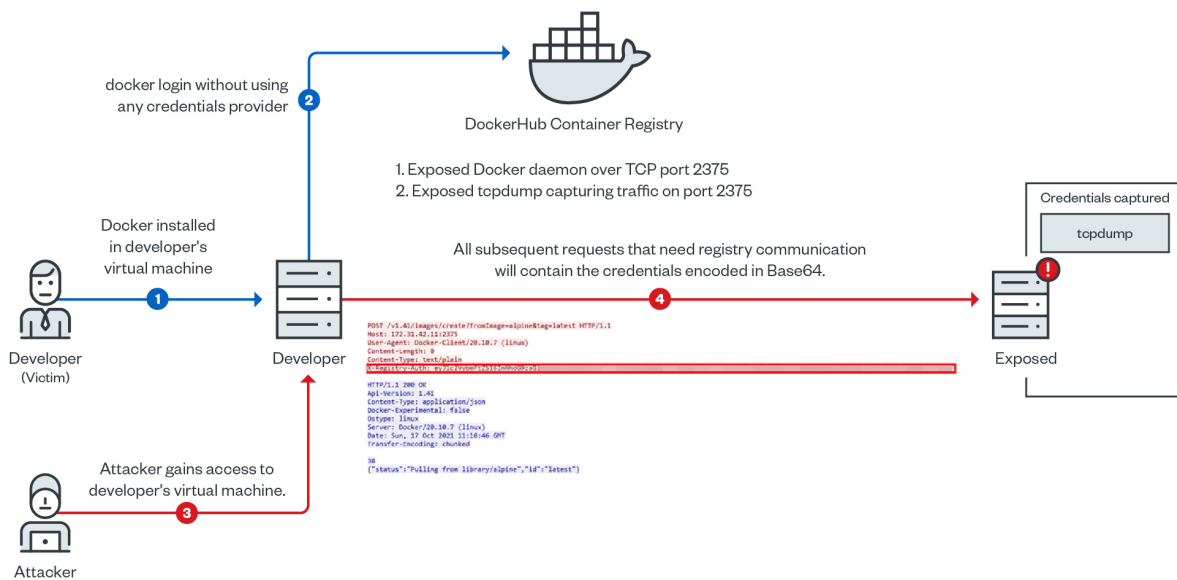
```

Figure 7. Request to create a container; the server responds with the ID of the newly created container

As a legitimate use case, a user might want to authenticate their DockerHub repository to create containers based on the images in their private repository. But if the user forgets to log out from DockerHub using “docker logout” and creates containers on untrusted hosts with dockerd exposed over TCP, the user becomes at risk since their usernames and passwords are hard-coded and non-encrypted, not to mention only encoded in Base64.

Credential leak scenarios

In the first scenario, the victim is logged in to their DockerHub registry. An attacker gains access to the victim’s virtual machine (VM) and tries to create a container on a remote server with the dockerd exposed over TCP. If the image that the container attempts for creation does not exist, the image is pulled from DockerHub and the header X-Registry-Auth is populated with the Base64-encoded credentials.



©2022 TREND MICRO

Figure 11. Scenario 1: The image pulled from DockerHub contains the Base64-encoded credentials.

Once abused, these compromised accounts can be used to view the following information and even pivot in the following ways:

1. **Associated email and reused credentials.** Cybercriminals can check for passwords being reused across different platforms and check for password leaks.
2. **Private repositories and images.** These might contain credentials like API keys and modify private images with backdoors.
3. **Access tokens.** These can be used to maintain persistent access to the account.
4. **Developers’ tools.** Pro features (like Teams, Organization, and Build pipelines) can be used to contaminate the build pipelines based on Docker and lead to supply-chain attacks.

From a different perspective, another attack scenario could involve the attacker's account itself, where the cybercriminals are logged in to their own DockerHub registry and their accounts leak credentials. While legitimate users might not be concerned about stealing threat actors' credentials, we analyzed that the procedure in creating a new container is also applicable to catching threat actors who might have left their respective accounts logged in. They might then attempt to create a container in the honeypot with an exposed dockerd over TCP. Since our honeypot has tcpdump running and captures network traffic as packet capture (PCAP), we can fetch the credentials of the attacker encoded in Base64.

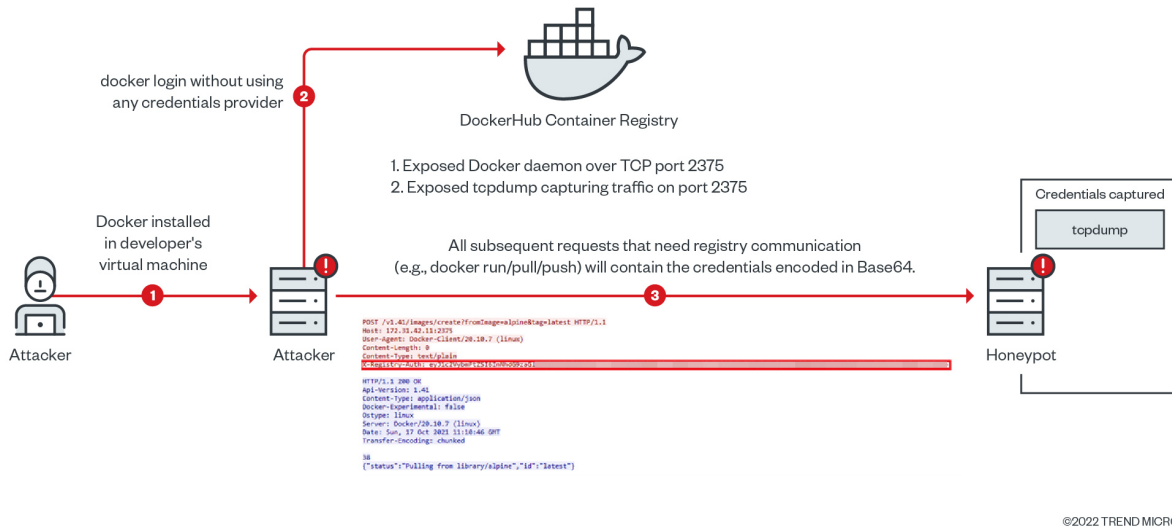


Figure 12. Scenario 2: Attackers' credentials leaked
Conclusion

Developers use containers to aid them in their development and deployments, optimize their workflows, and increase their productivity. Likewise, small gaps such as component misconfigurations have been known to be abused by cybercriminals. Malicious actors can perform damaging activities from these openings, like compromising the host for unauthorized cryptocurrency mining, exfiltrating sensitive information such as keys and credentials, or — at its worst— controlling victim servers to expand botnet malware usage, among other illicit activities. Indeed, cybercriminal groups such as TeamTNT will not stop anytime soon for as long as there are components and accounts that can be abused.

Based on our observations, we were able to identify TeamTNT's accounts because of one of the members' mistake on three occasions. There are three possible scenarios in which the user could have made this error:

1. The threat actors logged in to their DockerHub account using the credentials of alpineos.
2. The threat actors' machines were self-infected and were not using credential helpers.

3. The threat actors didn't log out from their DockerHub account while attacking exposed Docker REST API servers.

We found a total of 30 such accounts that were compromised, the credentials for which were being leaked. The registries for these were DockerHub and Alibaba Cloud Container Registry. While we have acquired this information and have access to the aforementioned credentials that might have been abused by TeamTNT, we did not access these credentials unauthorized. We have also informed Docker about these accounts and are working with them to resolve the matter.

Organizations' security teams need to be aware that developer security is critical considering this type of compromise around developer-centric tools like Docker have been observed being abused by threat actors. We advise that teams create policies for access and credential use, as well as generate threat models of their environments. Security teams can use these to educate developers about what can go wrong. Here are some mitigation practices for organizations and developers:

- While creating containers on a remote host via the Docker daemon REST API, developers should be aware that DockerHub credentials will also be shared if they are creating images from the specified container registry. They should proceed to do so only when the remote host is trusted.
- With the rising number of malicious open-source packages targeting user credentials, users should avoid storing credentials in other components such as environment variables. Instead, they must choose tools such as credential stores and [helpers](#).
- If users need to use the Docker Daemon over REST API via internet, it is recommended that they configure the exposed REST API with [TLS](#) (Transport Layer Security) protocol to avoid man-in-the-middle ([MiTM](#)) attacks sniffing for credentials.

The full details of this research will be presented at the [c0c0n XV Hacking and Cyber Security Conference](#) scheduled on Sept. 21 to 24, 2022.

Indicators of Compromise (IOCs)

For a full list of IOCs, you can visit our blog entries on previous incidents we documented [here](#) and [here](#).

sXpIBdPeKzI9PC2p0SWMpUSM2NSxWzPyXTMLibXmYa0R20xk