Automating Cloud Forensics Lab Provisioning

Tim Ip

Security Coordination Center Incident Response

About Me - Tim Ip

Adobe Incident Response

- Work at Ground Zero
- Focus on Incident Response Automation to make Incident Responders' lives easier

Previously DevSecOps Engineer in the Life Sciences industry, Security Architect at a University, Security Big Data Consultant at Big 4

- Focus on Big Data and Automation
- Purple Teaming (Offensive Security, Detection Engineering and Big Fan of Splunk/Sysmon)
- Splunk-er for nearly 10 years

Director of Monitoring, Global Collegiate Penetration Testing Competition (https://cp.tc)

- Managing monitoring infrastructure
- Detection Engineering, Threat Hunting in competitive environment

What does a Forensics Lab look like?

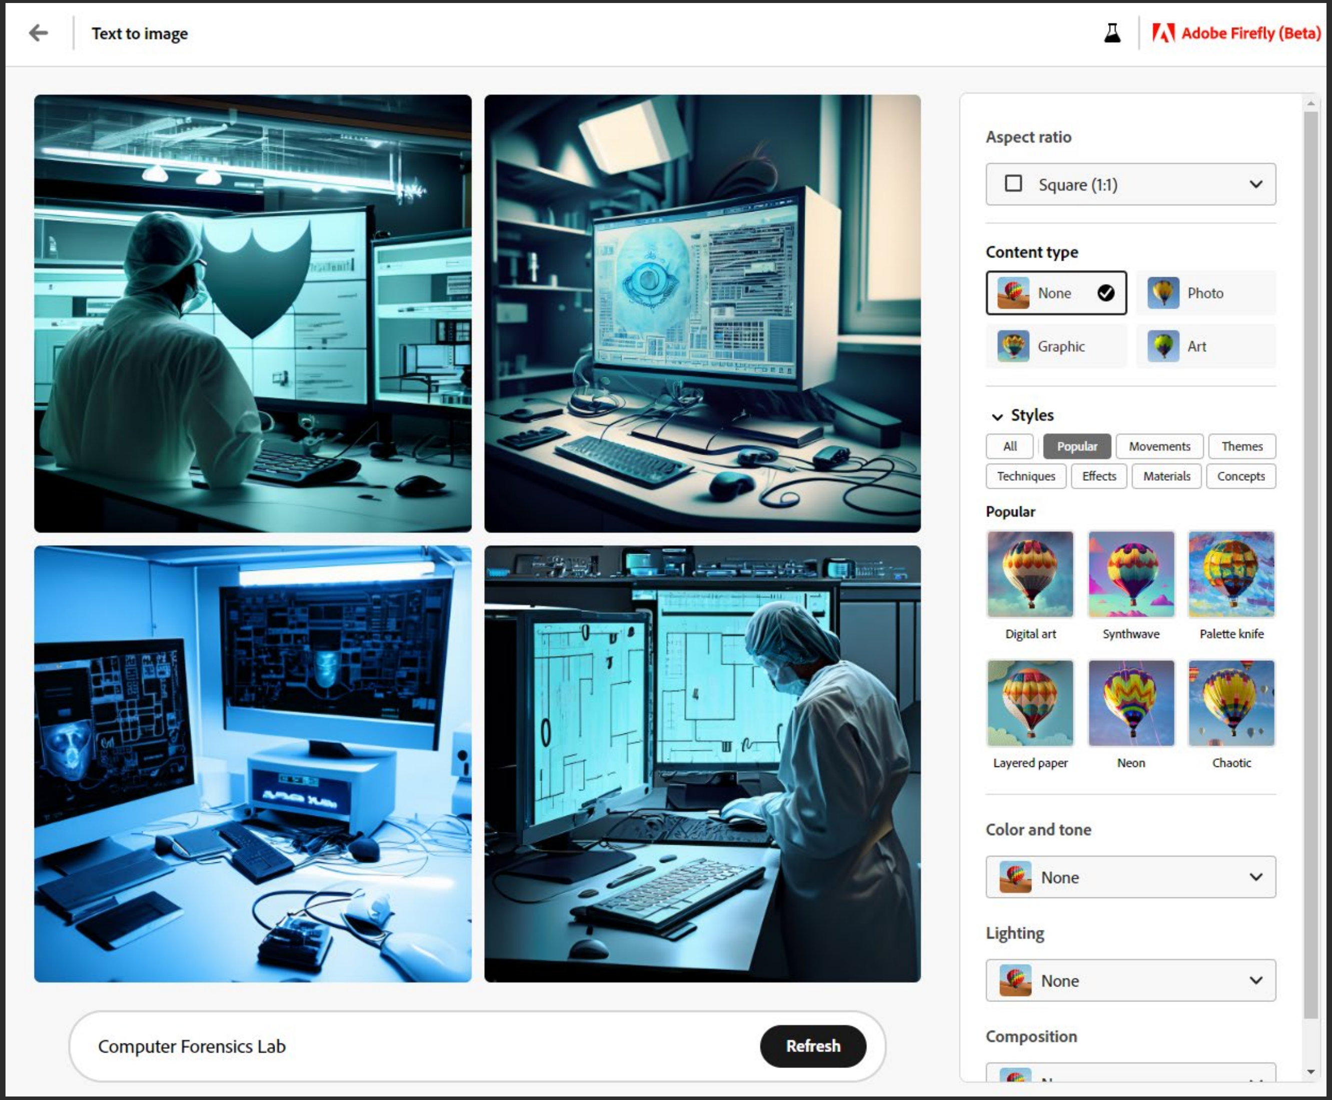

Text to image

Computer Forensics Lab

Aspect ratio: Square (1:1)

Content type: None ✓

Color and tone: None

Lighting: None

What does a Forensics Lab look like?

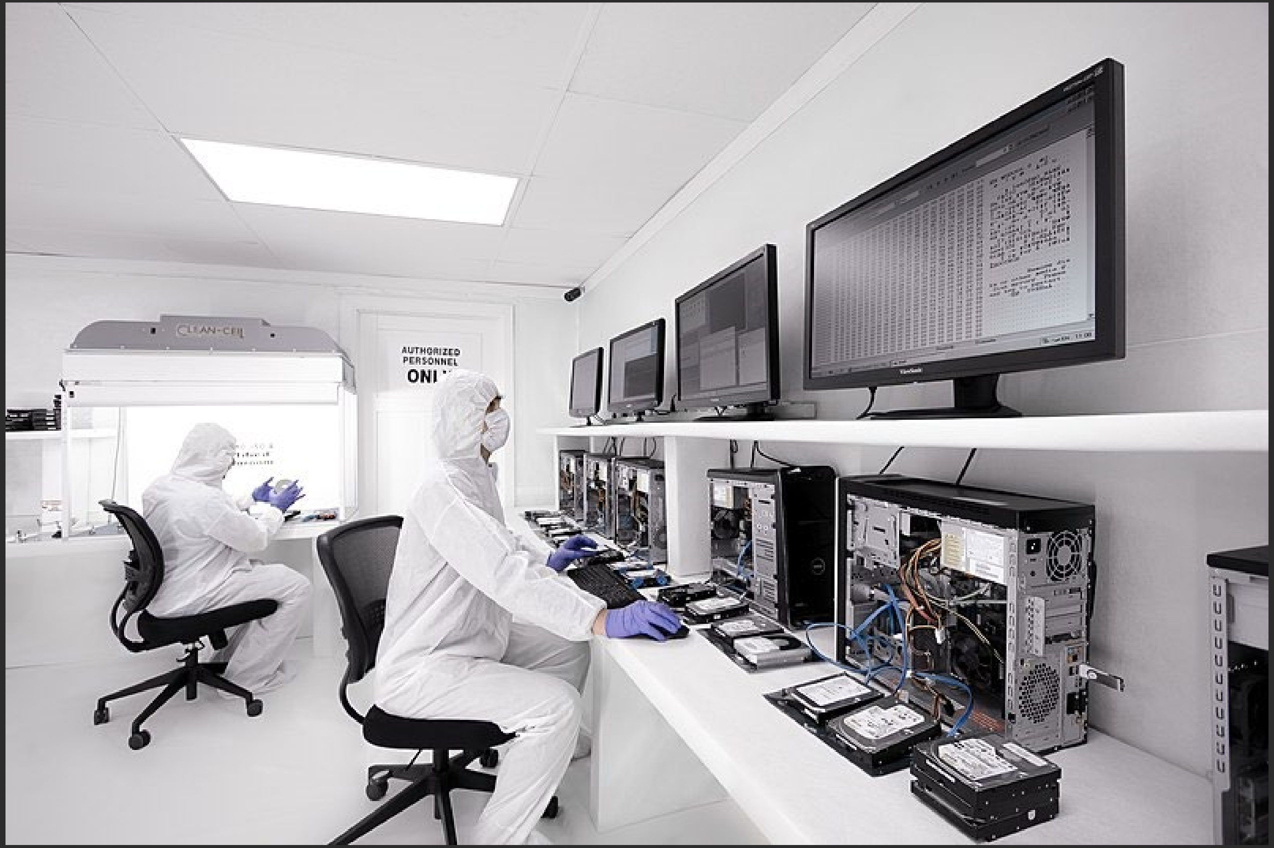

From: Wikimedia Commons (Creative Commons Attribution-Share Alike 4.0 International license)

What does a Forensics Lab look like?

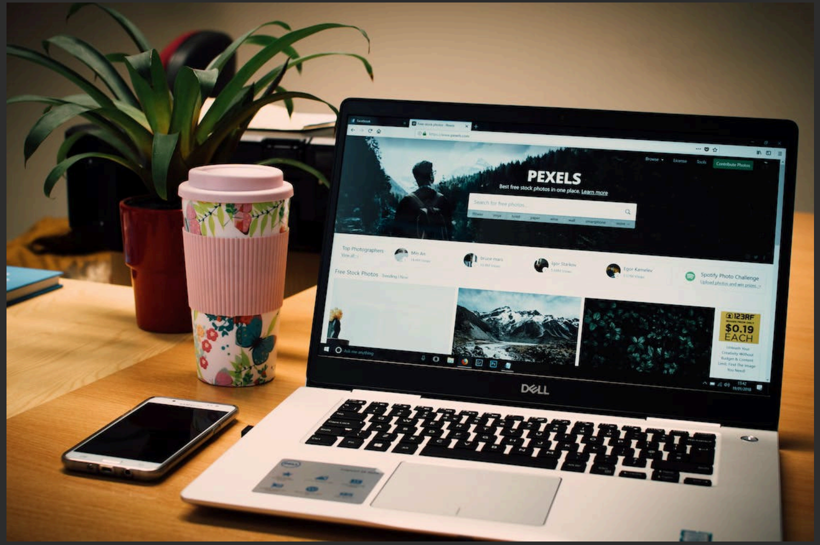

Are you able to successfully perform forensics for Cloud Compute Workloads?

What will this talk cover?

Our approach in handling forensics for Cloud Compute Workloads (Virtual Machines)

- AWS: Elastic Compute Cloud (EC2)
- Azure: Virtual Machines
- Google Cloud: Compute Engine

Forensics Lab

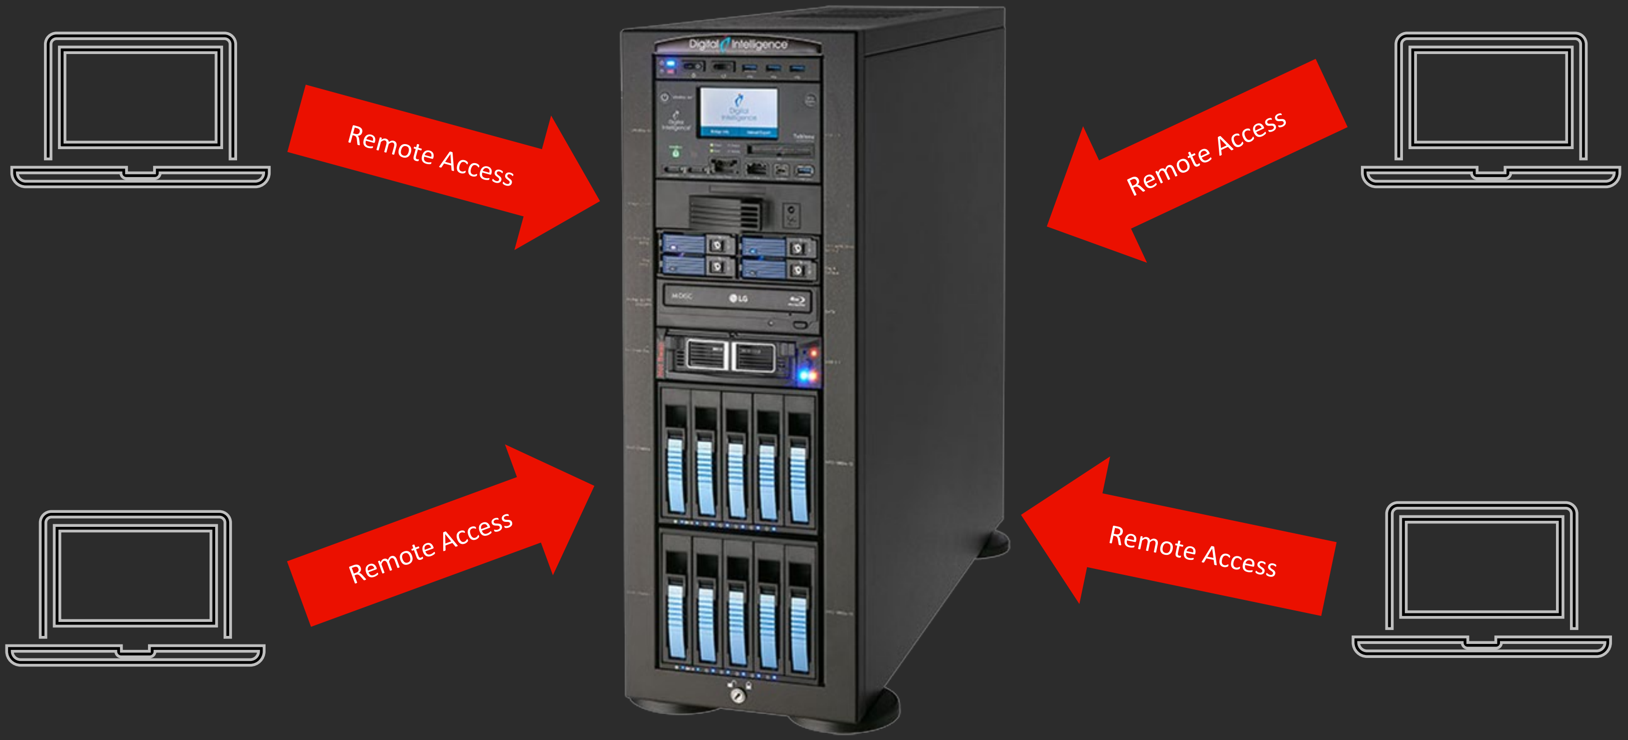

Problems with Shared Forensics Environments

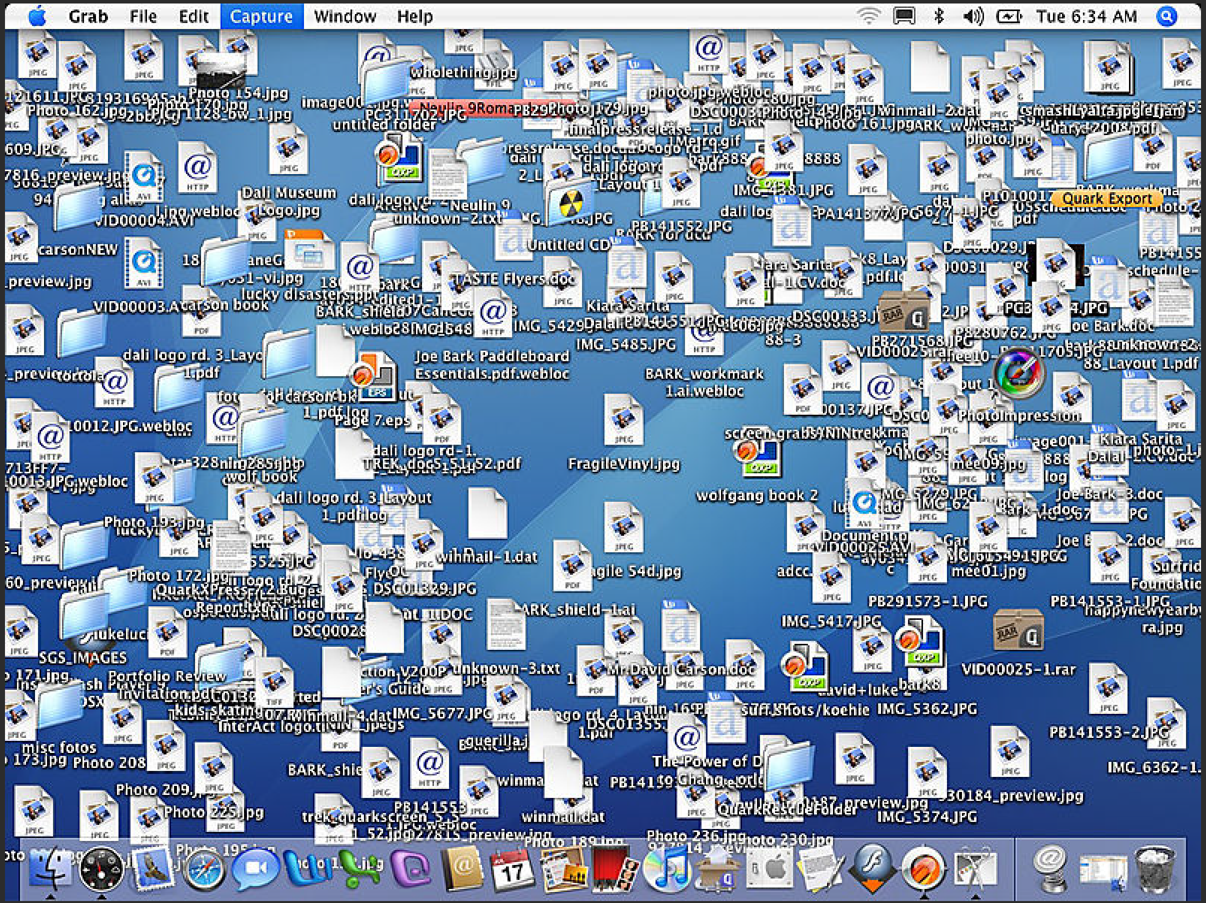

Error

Windows - System Error

Unknown Hard Error

To continue, click Next.

Adobe

Problems with Data Acquisitions

- Native Cloud Logs (CloudTrail, IAM, VPC Flow) from SIEM
- Disk Images
- Memory Dumps

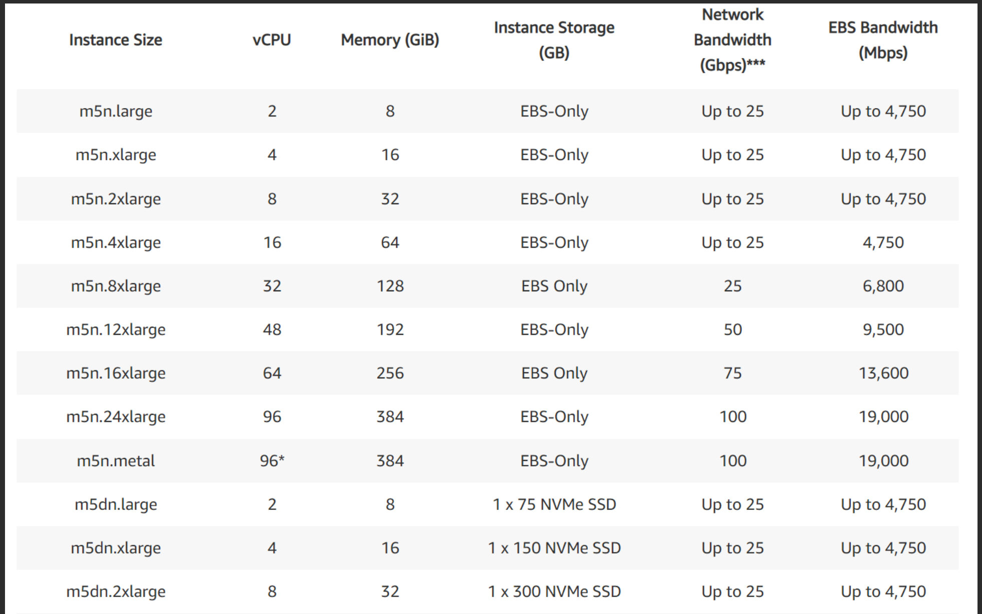

Instance Size	vCPU	Memory (GiB)	Instance Storage (GB)	Network Bandwidth (Gbps)***	EBS Bandwidth (Mbps)
m5n.large	2	8	EBS-Only	Up to 25	Up to 4,750
m5n.xlarge	4	16	EBS-Only	Up to 25	Up to 4,750
m5n.2xlarge	8	32	EBS-Only	Up to 25	Up to 4,750
m5n.4xlarge	16	64	EBS-Only	Up to 25	4,750
m5n.8xlarge	32	128	EBS Only	25	6,800
m5n.12xlarge	48	192	EBS-Only	50	9,500
m5n.16xlarge	64	256	EBS Only	75	13,600
m5n.24xlarge	96	384	EBS-Only	100	19,000
m5n.metal	96*	384	EBS-Only	100	19,000
m5dn.large	2	8	1 x 75 NVMe SSD	Up to 25	Up to 4,750
m5dn.xlarge	4	16	1 x 150 NVMe SSD	Up to 25	Up to 4,750
m5dn.2xlarge	8	32	1 x 300 NVMe SSD	Up to 25	Up to 4,750

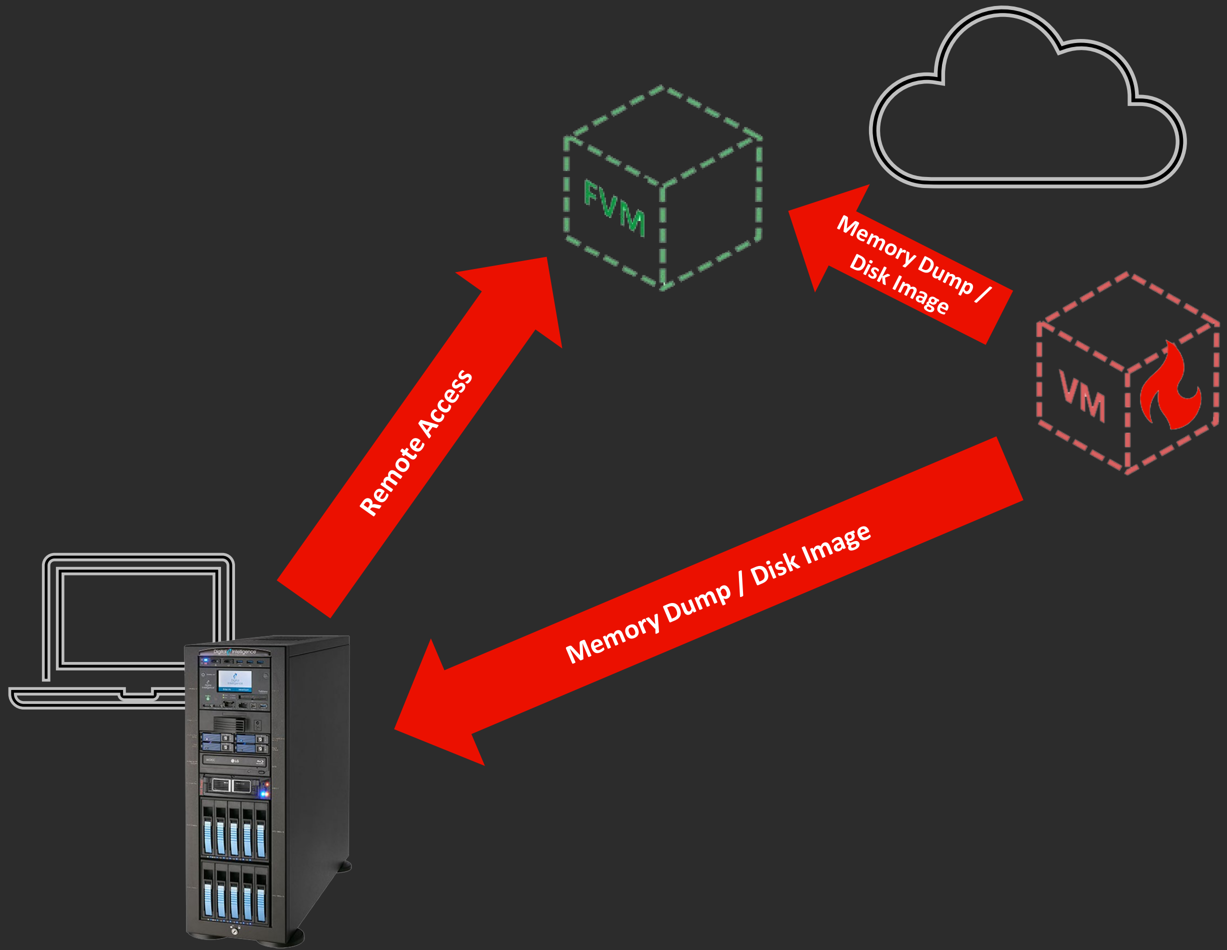

Workloads Everywhere

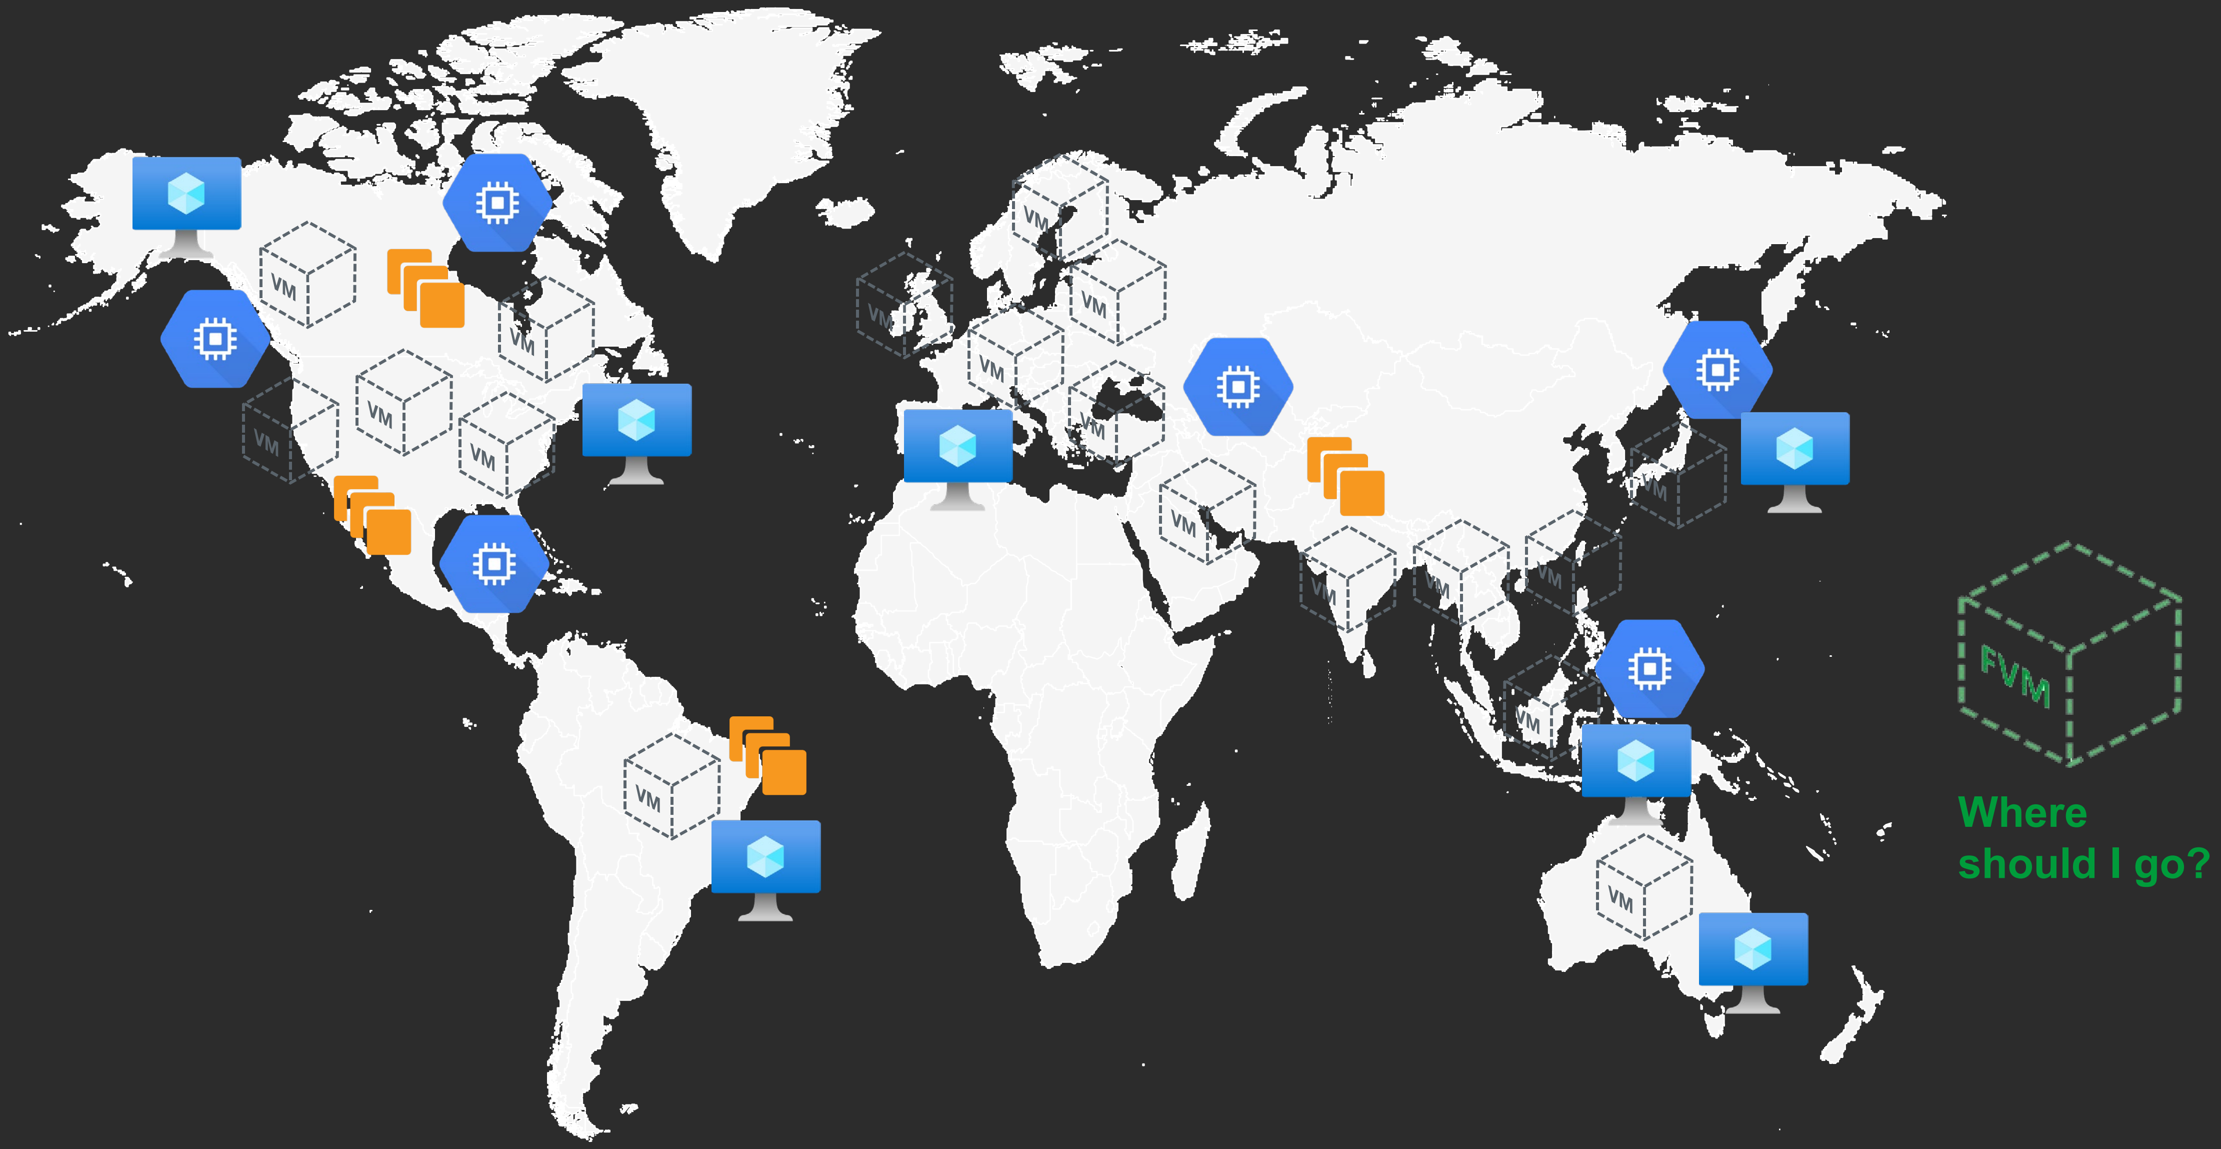

Where should I go?

Workloads Everywhere

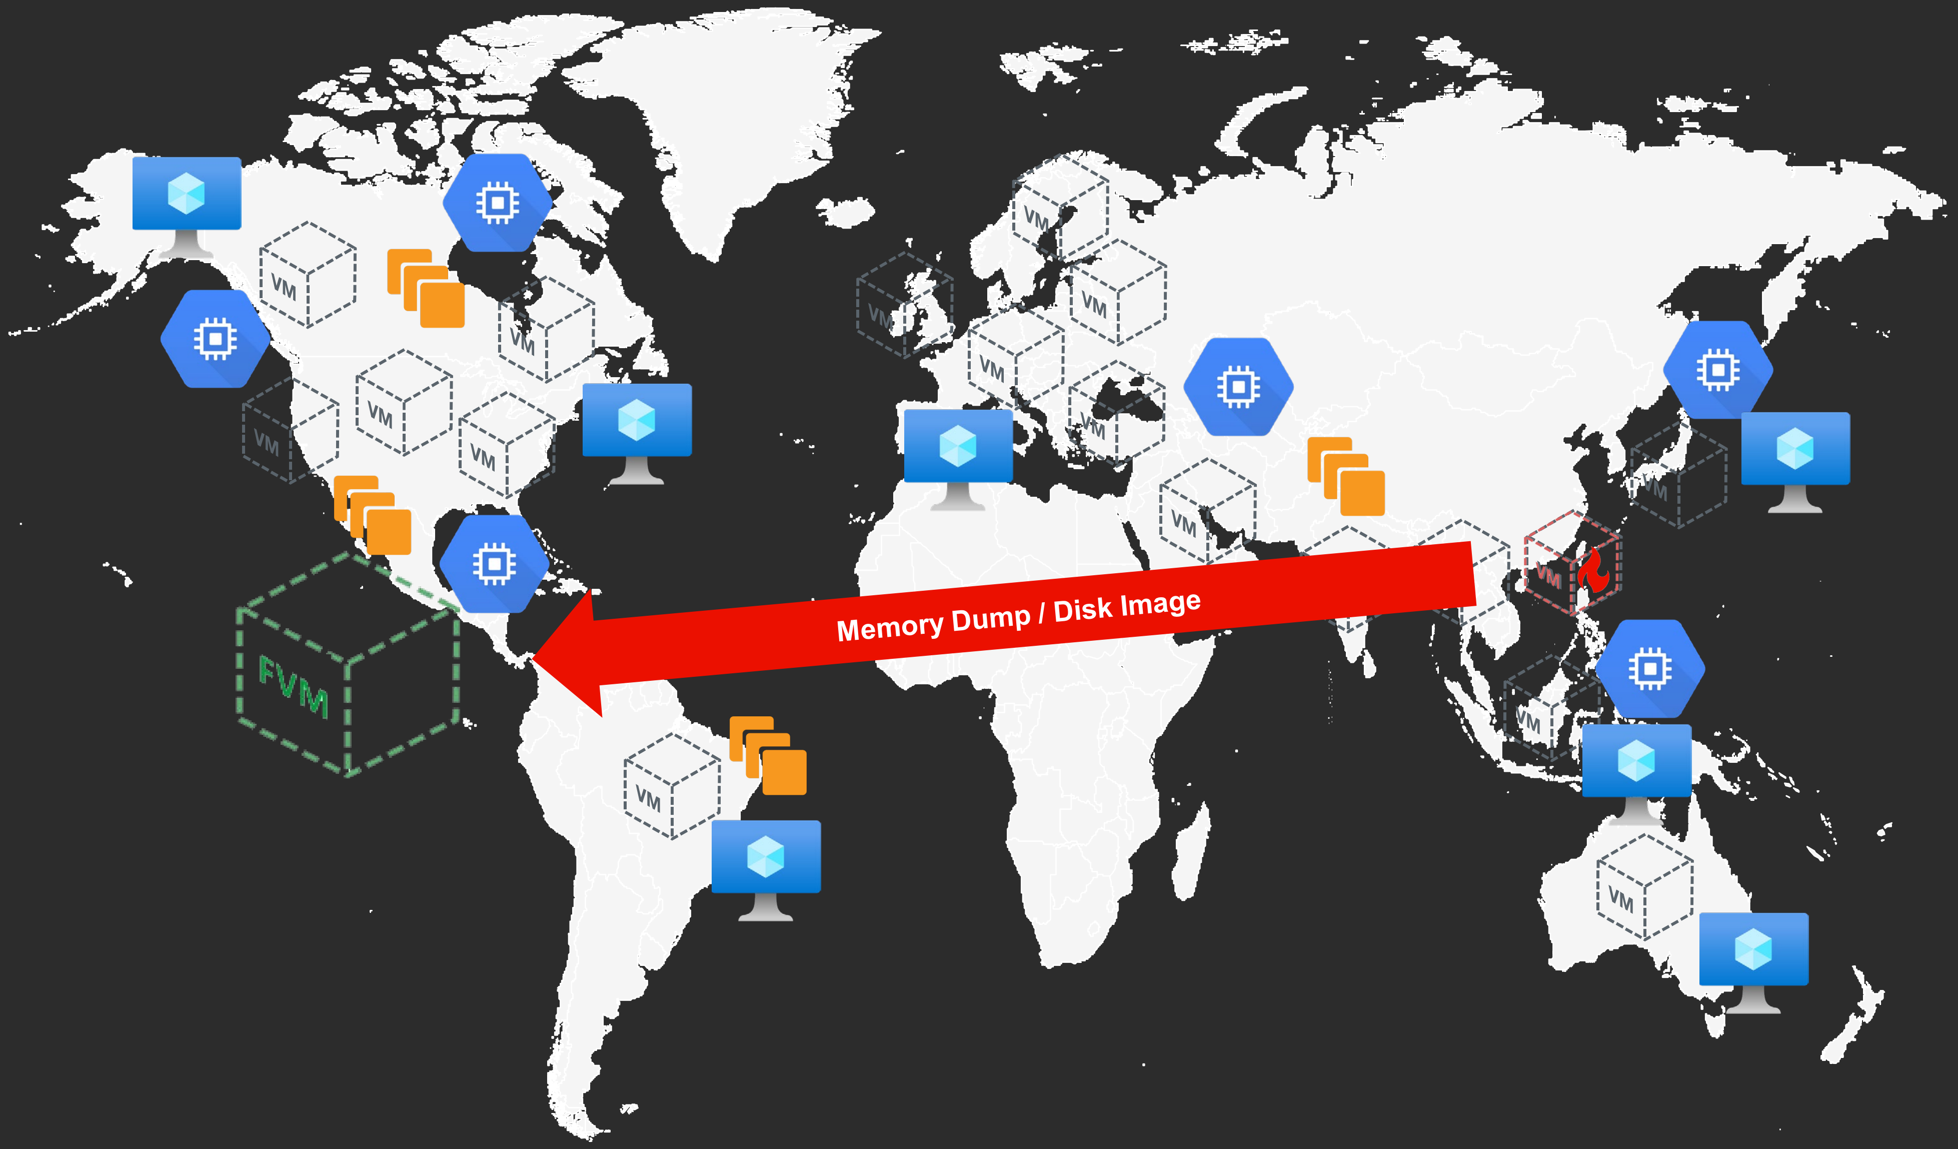

Workloads Everywhere

Requirements

- **Automate to create a forensics environment SUPER fast**
  - Using available forensics tools
  - Secure and hardened
- **Available in most regions across major cloud service providers (AWS, Azure and GCP)**
- **Make our lab environment ephemeral to save money**
  - Only up during incident
  - Create when incident starts, tear down when incident ends
- **Provide a way to easy archive forensics artifacts to permanent storage**
- **Allow collaboration for Forensics Lab development**

Solution:
Automate Forensics Lab setup using Infrastructure as Code (IaC)

HashiCorp Terraform

ANSIBLE

git

Terraform

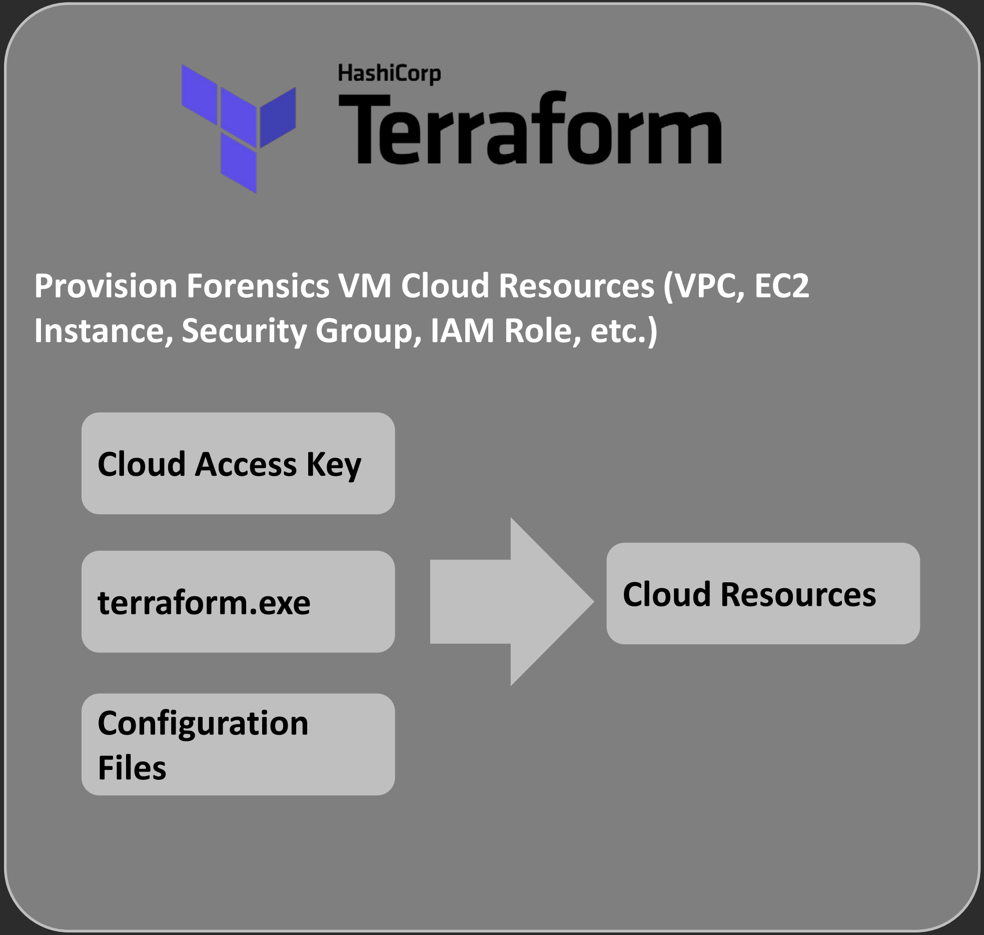

Provision Forensics VM Cloud Resources (VPC, EC2 Instance, Security Group, IAM Role, etc.)

```
module "ec2_instance" {
  source  = "terraform-aws-modules/ec2-instance/aws"

  name = "single-instance"

  instance_type          = "t2.micro"
  key_name               = "user1"
  monitoring             = true
  vpc_security_group_ids = ["sg-12345678"]
  subnet_id              = "subnet-eddcdzz4"

  tags = {
    Terraform   = "true"
    Environment = "dev"
  }
}
```

Ansible

ANSIBLE

Configure Forensics VM (Setup up directory structure, deploy forensics tools – Plaso, Volatility, Autopsy, Splunk)

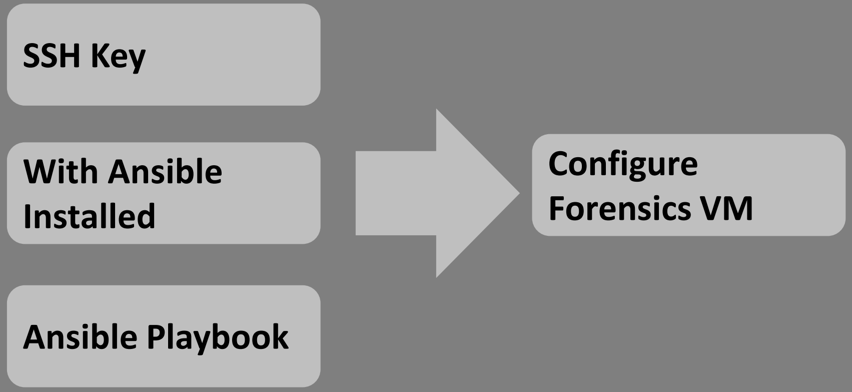

SSH Key

With Ansible Installed

Ansible Playbook

Configure Forensics VM

```
- name: Play Web - Create apache directories and username
in web servers
  hosts: webservers
  become: yes
  become_user: root
  tasks:
    - name: create username apacheadm
      user:
        name: apacheadm
        group: users,admin
        shell: /bin/bash
        home: /home/weblogic
```

Requirements

- **No need to learn Terraform/Ansible before using it**
- **Simple and fast (Single Step) to spin up/tear down Forensics Lab environment**

Solution:
A Wrapper for Terraform/Ansible

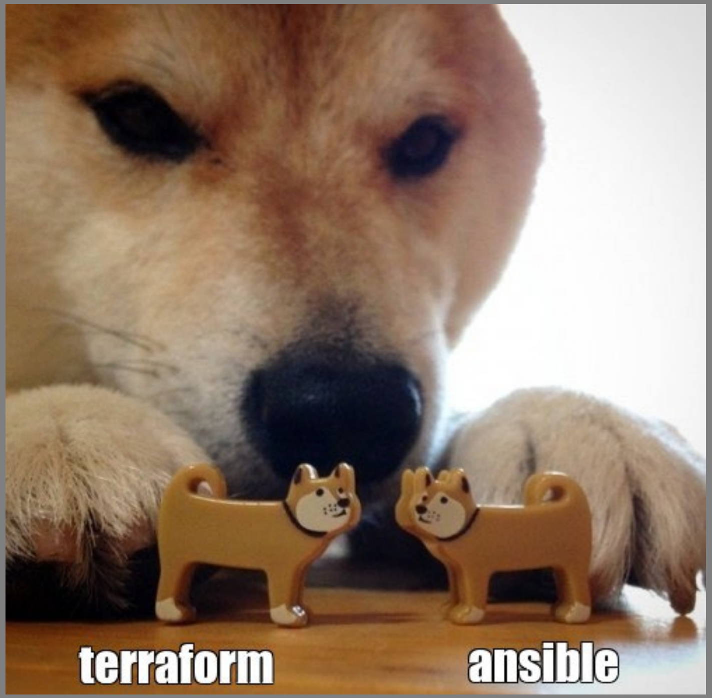

Forensics VM (FVM)

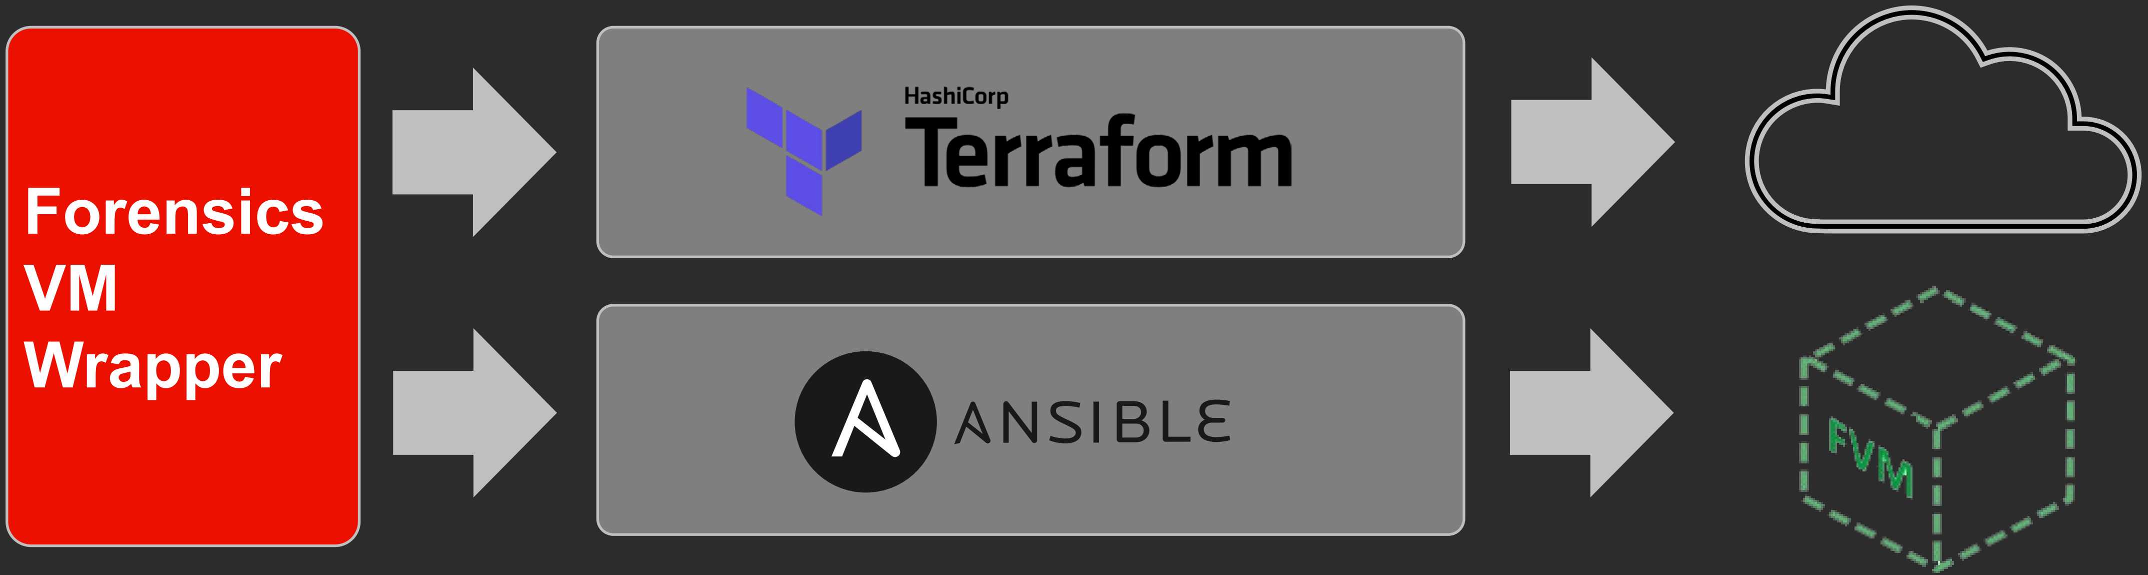

Forensics VM Wrapper

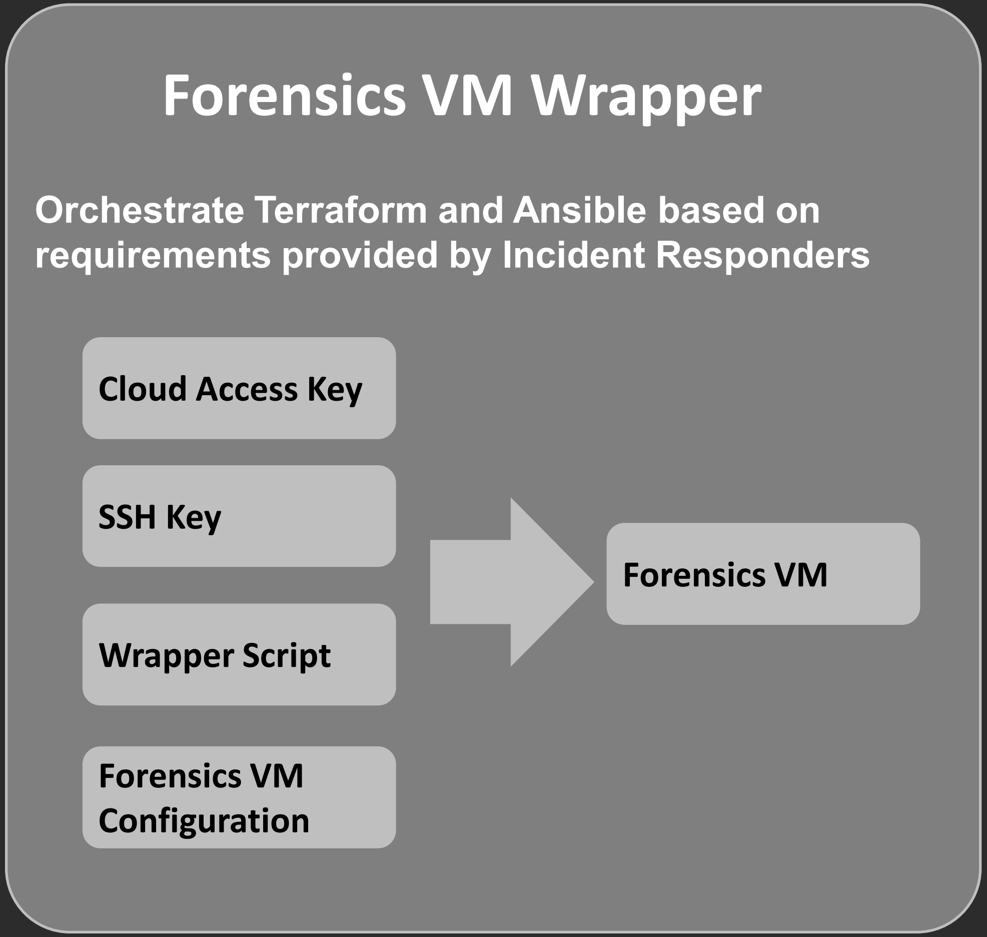

Forensics VM Wrapper

Orchestrate Terraform and Ansible based on requirements provided by Incident Responders

Cloud Access Key

SSH Key

Wrapper Script

Forensics VM Configuration

Forensics VM

```
[sample_aws]
incident_name = delawareaws
cloud_provider = aws
environment = generic
region = us-east-1
az = us-east-1b
sshkey_public_path = ~/.ssh/id_rsa.pub
disk_size_gb = 1000
instance_type = t3.2xlarge
plugins =
all_gatherfacts,all_createmountpoint,all_docker,all_forensi
cs,aws_forensics,all_tsk,all_volatility,all_tmout,all_maxlo
gins,all_addsshkeys,all_sethostname,all_falcon,all_splunk,a
ws_s3upload
```

Forensics VM Configuration (`forensicsvm.conf`)

- A configuration file to manage multiple Forensics VMs
- Stanza: A section of a configuration file. Stanzas begin with a text string enclosed in brackets and contain one or more configuration parameters defined by key/value pairs.
- Define incident name, VM location (Cloud provider/environment/region/az), disk space, plugins, etc.

```
[incident_1]
incident_name = this_is_a_template
cloud_provider = aws
environment = adobe
region = us-west-2
az = us-west-2a
sshkey_public_path = ~/.ssh/id_rsa.pub
disk_size_gb = 500
instance_type = t3.2xlarge
plugins =
all_gatherfacts,all_createmountpoint,all_docker,all_forensics,aws_forensics,all_tsk
,all_volatility,all_tmout,all_maxlogins,all_addsshkeys,all_sethostname,all_falcon,a
ll_splunk,aws_s3upload

[incident_2]
incident_name = this_is_a_template
cloud_provider = azure
environment = adobe
region = westus2
ssh_login_id = sccforensics
sshkey_public_path = ~/.ssh/id_rsa.pub
disk_size_gb = 500
instance_type = Standard_B4ms
plugins =
all_gatherfacts,all_createmountpoint,azure_mountdisk,all_docker,all_forensics,all_t
sk,all_volatility,azure_prepprofileenv,all_tmout,all_maxlogins,all_addsshkeys,all_s
ethostname,all_falcon,all_splunk,azure_allowsplunkwebfw

[incident_3]
incident_name = this_is_a_template
cloud_provider = gcp
region = us-west1
zone = us-west1-a
disk_size_gb = 500
instance_type = e2-standard-8
plugins =
all_gatherfacts,all_docker,all_forensics,all_tsk,all_volatility,all_tmout,all_maxlo
gins,all_splunk
ssh_login_id = CHANGEME_adobe_com
```

Forensics VM plugins are Ansible tags. By choosing the plugins you want in forensicsvm.conf , you can customise your FVM as well as shorten the FVM spin-up time.

Forensics VM plugins are Ansible tags. By choosing the plugins you want in forensicsvm.conf , you can customise your FVM as well as shorten the FVM spin-up time.

Plugins

- Forensics VM plugins are Ansible tags. By including plugins in configuration files, you can customize your FVM as well as shorten FVM spin-up time.
- Anyone can develop plugins to add more functionality to FVM.

Plugin Name	Usages
all_addsshkeys	Add SSH Keys
all_createmountpoint	Create mount point
all_docker	Install and configure Docker
all_falcon	Install EDR
all_forensics	Create forensics directory structure and install and configure various tools and libraries
all_gatherfacts	Default - Gather information for Ansible
all_maxlogins	Adjust maxlogins setting to allow multiple sessions for a single account
all_sethostname	Configure hostname
all_splunk	Install Splunk
all_tmout	Unlock TMOUT restriction
all_tsk	Install The Sleuth Kit
all_volatility	Install Volatility
aws_forensics	Create and install various tools and libraries specific to AWS
aws_s3upload	Confoigure AWS Role for S3 Upload
azure_allowsplunkwebfw	Configure Azure Firewall to allow SplunkWeb traffic
azure_mountdisk	Mount Forensics volume
azure_prepprofileenv	Configure Volatility profile compile environment for Azure VM

Single command to spin up/teardown FVM

```
./forensicsvm create <stanza>

./forensicsvm destroy <stanza>
```

forensicsvm.conf

```
[incident_1]
incident_name = incident_1
cloud_provider = aws
environment = adobe
region = us-west-2
az = us-west-2a
sshkey_public_path = ~/.ssh/id_rsa.pub
disk_size_gb = 500
instance_type = t3.2xlarge
plugins =
all_gatherfacts,all_createmountpoint,all_docker,all_forensics,
aws_forensics,all_tsk,all_volatility,all_tmout,all_maxlogins,a
ll_addsshkeys,all_sethostname,all_falcon,all_splunk,aws_s3uplo
ad
```

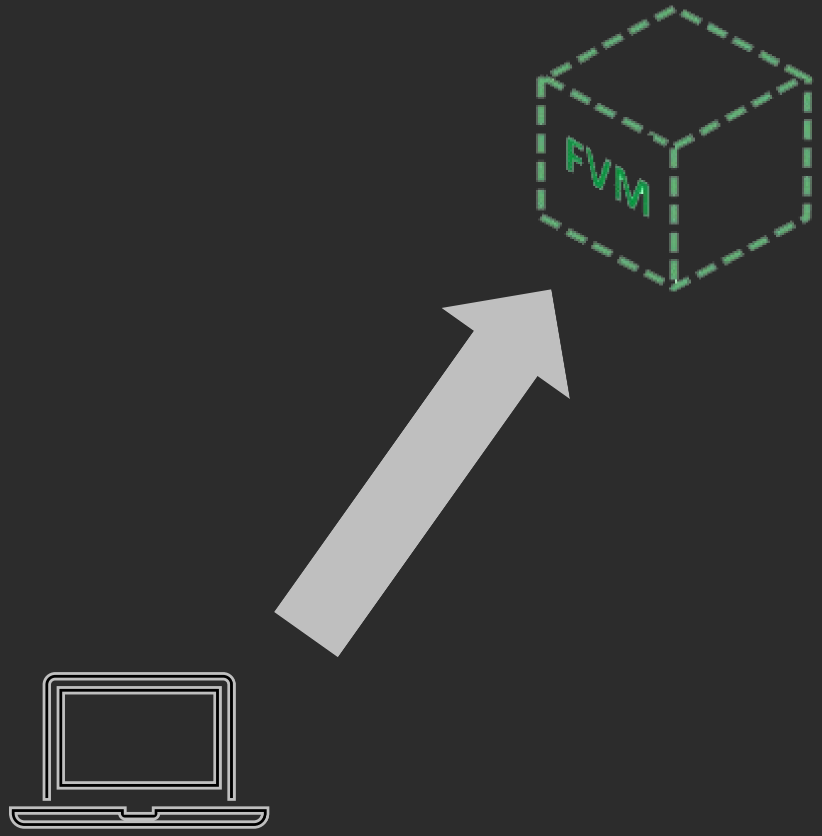

Forensics Pipeline and FVM Lifecycle

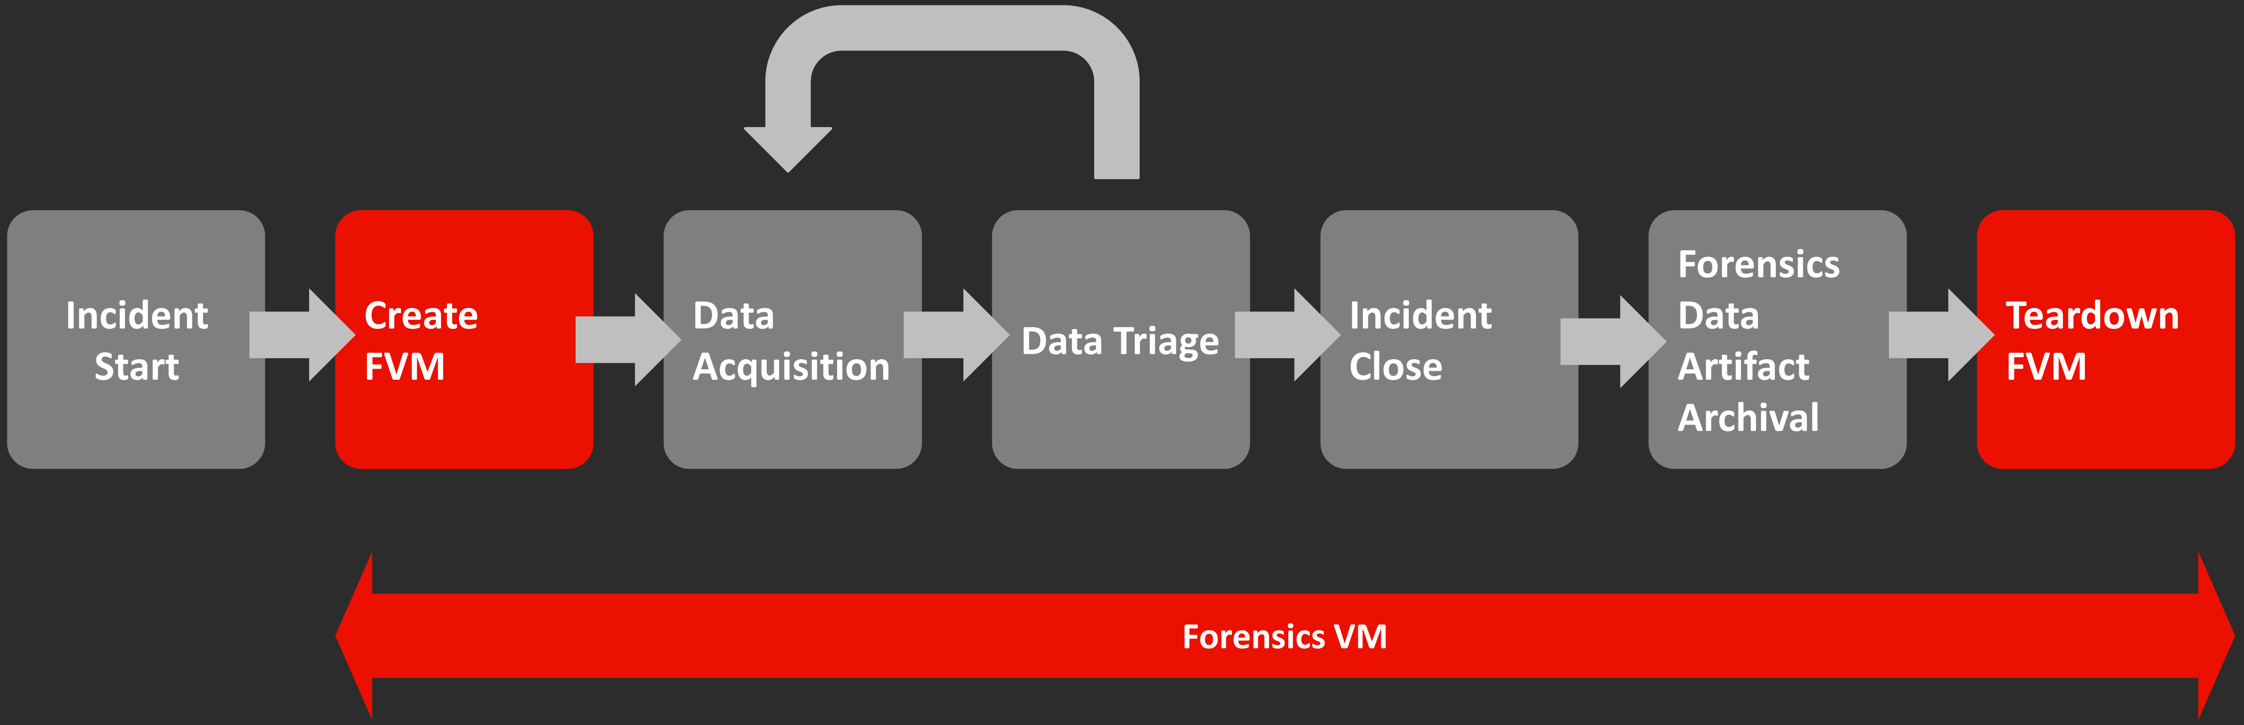

Scripts to automatic Triage and Data Archival

Automation Script for fast triage

- Volatility Triage
- Disk Image Triage / Plaso
- Application Triage

Automation Script for data archival

- S3 Archival
- Azure Storage Archival

Demo

Benefits

- Able to create Forensics Labs anywhere, anytime
- Available for popular cloud service providers
  - Analyze data locally: Avoid potential compliance issues (No need to transfer data out from jurisdictions)
- Ephemeral Lab Environment
  - Save $$$ (FVM only up during incident)
  - Fresh environment at start
- Encourage contributions and knowledge sharing
- Standardize workflow and formalize forensics pipeline

Future Development

- Expand coverage to spin-up Windows FVM
- Add more forensics triage tools to improve our triage efficiency
- Automated FVM provisioning through ticketing system

Takeaways

- Use Infrastructure as Code (IaC) to provision lab environment
- Create forensics pipeline with well-defined workflow/process
- Test the pipeline regularly
- Create a platform to encourage contributions and knowledge sharing

Adobe