

# LLVM Supply Chain Security

From Developer's Desk to User's Device

Tom Stellard  
Red Hat

---

# What is Supply Chain Security?


---

# Example: Lock Company



---

# Example: Lock Company



---

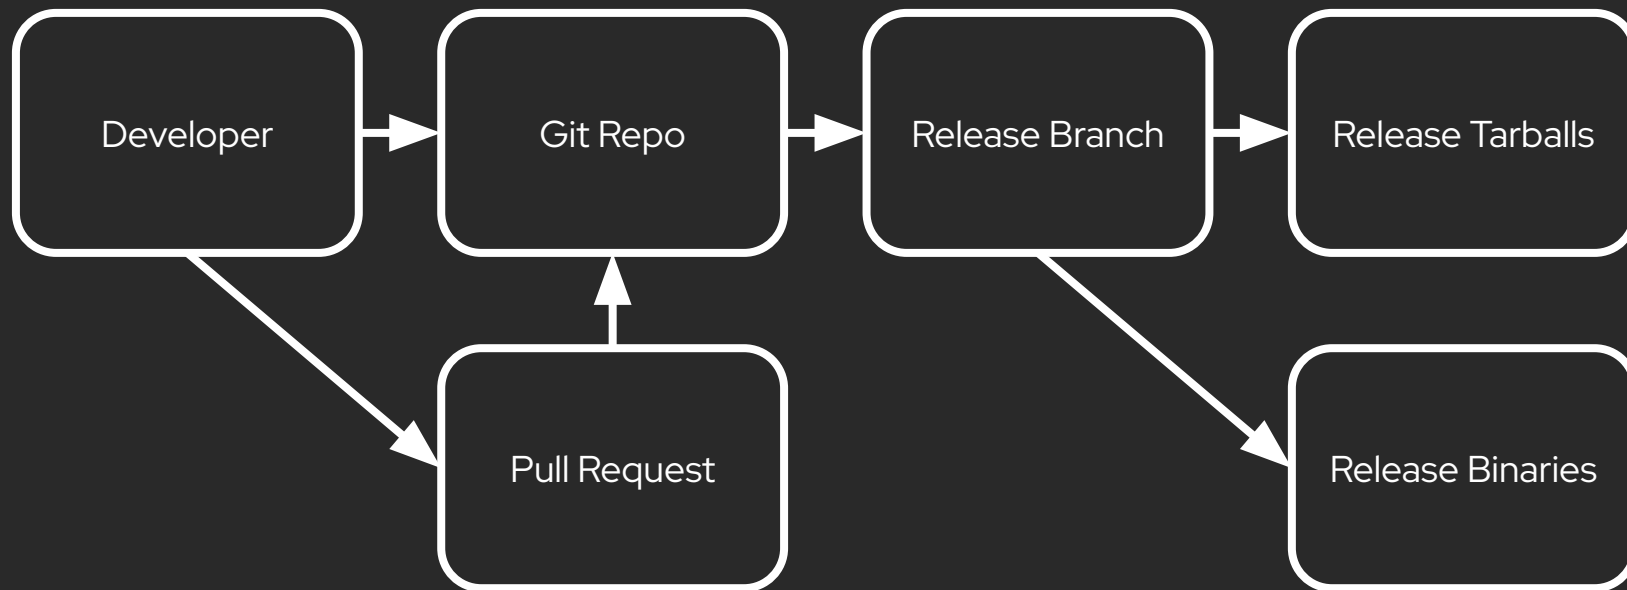
# Example: Lock Company



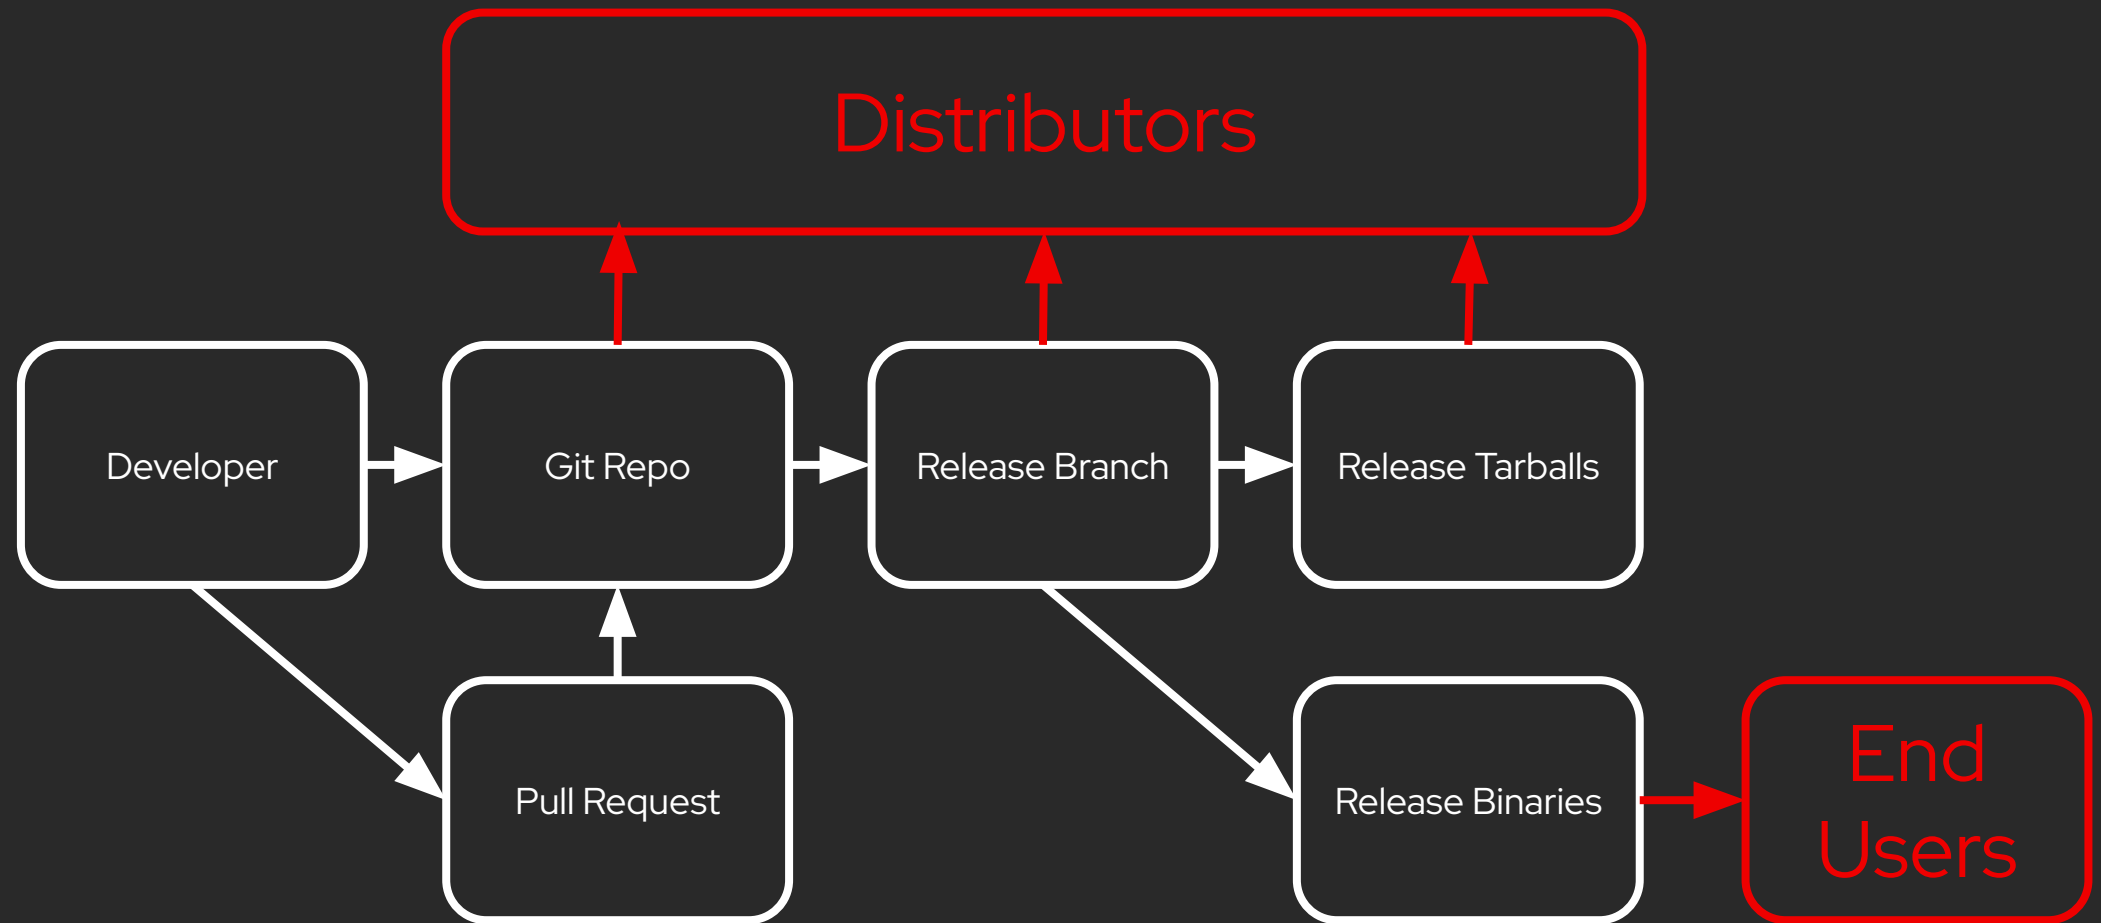
---

What does our software  
supply chain look like?

# LLVM Supply Chain

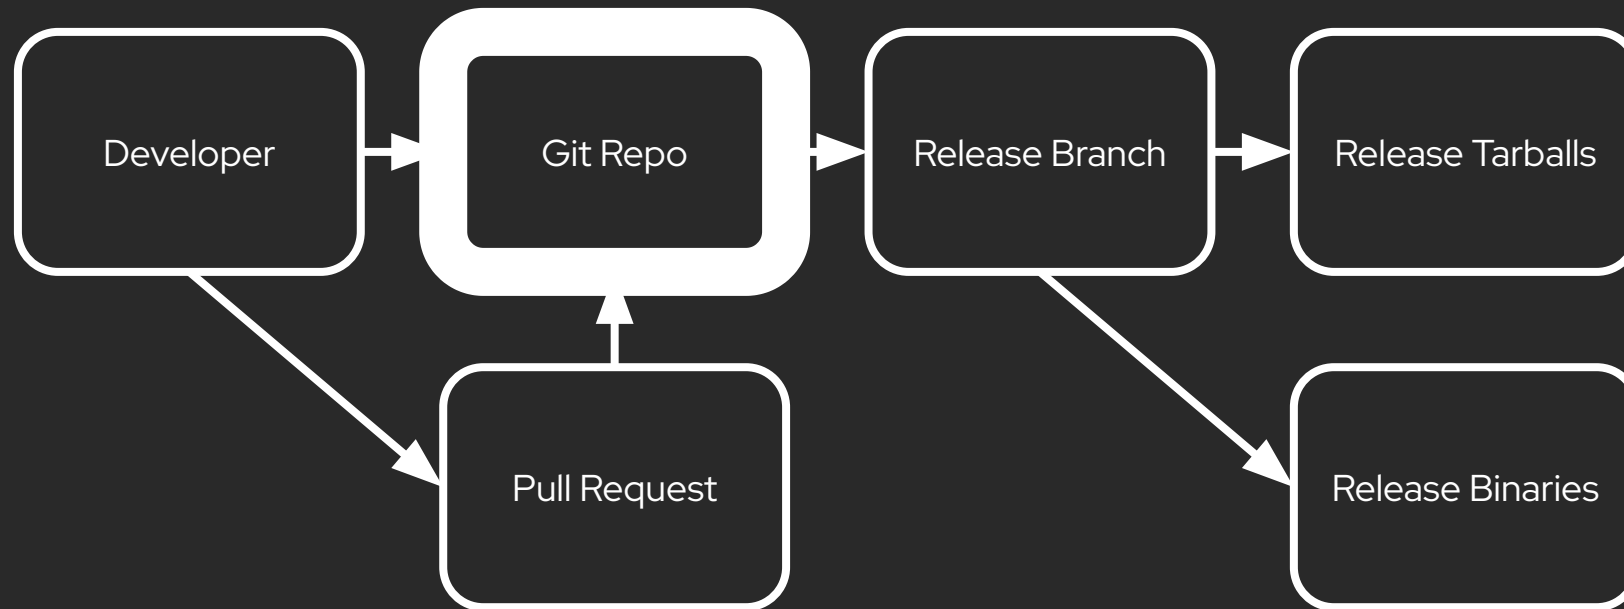


# LLVM Supply Chain





# LLVM Supply Chain



---

How do we secure the git repository?

How do we  
secure the git  
repository?

- ▶ **Restrict Commit Access.**

# What are the requirements for commit access?

- ▶ Ask!
- ▶ Give a Reason.
- ▶ No contribution requirements!

# How do we secure the git repository?

- ▶ Restrict Commit Access.
- ▶ **Commit rules.**

# What commit rules do we have?

- ▶ Major changes require an RFC.
- ▶ Pre-commit review with exceptions for 'Code Owners' and trivial changes.
- ▶ CI Tests must pass.

---

We have reactive  
enforcement.

# How do we secure the git repository?

- ▶ Restrict Commit Access.
- ▶ Commit rules.
- ▶ **Post-Commit review.**



# Post-Commit Review

- ▶ Commit lists get 1000 emails per day!
- ▶ How many people actually monitor the commit lists?

# How do we secure the git repository?

- ▶ Restrict Commit Access.
- ▶ Commit rules.
- ▶ Post-Commit review.

# Summary: Protecting the git repository.

- ▶ Anyone can get commit access!
- ▶ Anyone can commit at any time!
- ▶ We rely on people, not technology, to catch mistakes or prevent misuse.

---

Are we leaving our truck  
unlocked at night?



Who should  
be worried  
about this?

- ▶ Everyone!

# Do you build your product from main?

- ▶ What would happen if malicious code was pushed to main?
- ▶ Is your build pipeline secure against untrusted code?
- ▶ Reactive upstream policies won't protect you.

# Do you have internal CI building main?

- ▶ Even just testing main can put your systems at risk.

# What are the risks?

- ▶ Modified CMake files to execute arbitrary commands on your system.
- ▶ Backdoor inserted into compiler.
- ▶ Compiler modified to insert backdoors in everything it builds.



---

How did we get this so  
wrong?

---

We didn't actually do  
anything wrong!

---

It all comes down to  
priorities.

# What are the goals of our current policies?

- ▶ Make it easy for new contributors.
- ▶ Give experienced contributors flexibility to make their own decisions.

---

Can we update our  
policies while still meeting  
these goals?

# New Contributors.

- ▶ In the SVN days being a contributor without commit access was very frustrating!
- ▶ GitHub makes this easier: One click merge for pull requests.
- ▶ Technology has changed, but our policies haven't!

# Experienced Contributors.

- ▶ Do we need to explicitly define this group?
- ▶ Can we give extra privileges to experienced contributors?

# New policy ideas.

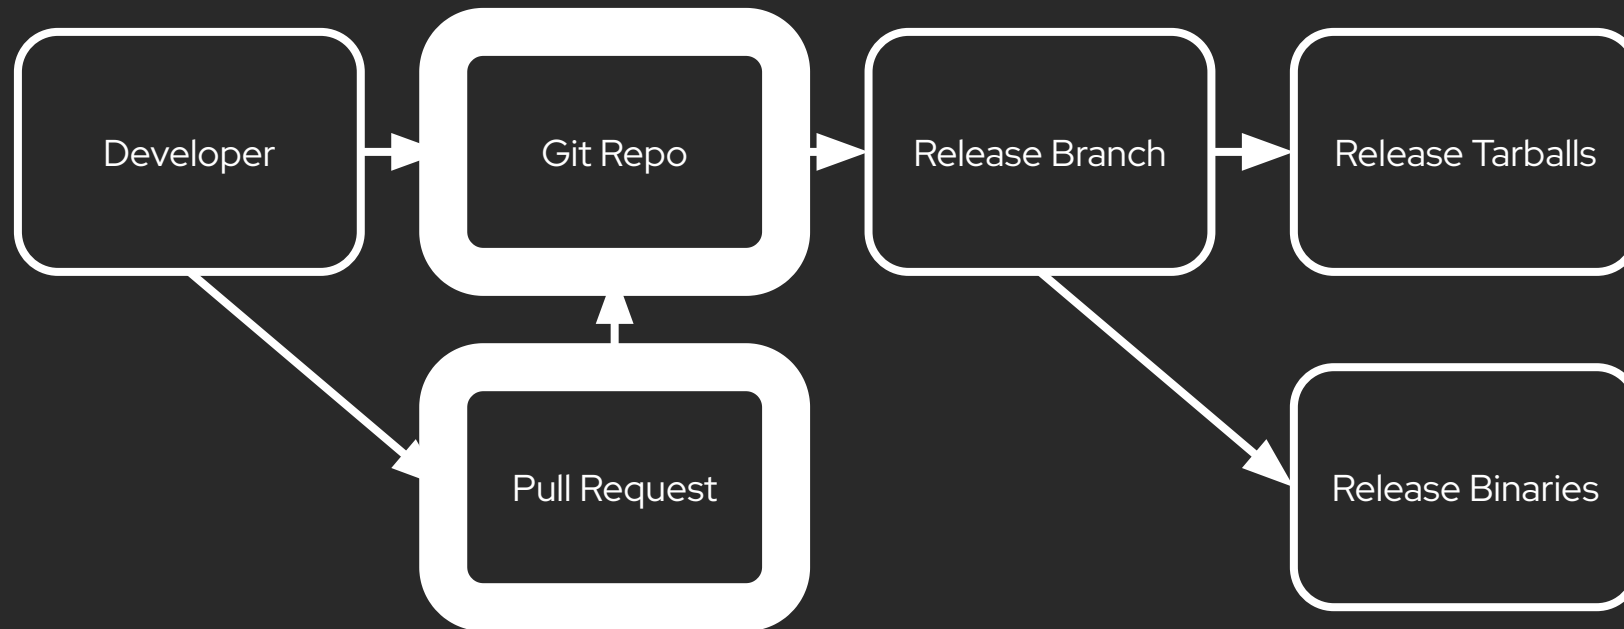
- ▶ Require pre-commit review for new contributors.
- ▶ Add requirements for obtaining commit access.
- ▶ Make CI mandatory.



# How you can help?

- ▶ Share your ideas!
- ▶ Review and comment on RFCs.
- ▶ Review your internal build/test process.

# LLVM Supply Chain



---

How do we secure our  
infrastructure?

# GitHub Actions: Overview

- ▶ Run automated jobs on GitHub infrastructure or self-hosted runners.
- ▶ Uses yaml based workflow definitions.
- ▶ One workflow per file, multiple jobs per workflow.

# GitHub Actions: Overview

- ▶ Jobs started by various events: pull request, issue comment, etc.
- ▶ Each job has its own access token:  
`GITHUB_TOKEN`
- ▶ Can add 'Secrets' for enhanced access.

# GitHub

## Actions: Risks

- ▶ Token/Secret compromise
- ▶ Repository Denial of Service attacks.
- ▶ Resource stealing.

# GitHub Actions: Tokens

- ▶ Allow access to GitHub via REST API.
- ▶ Some Examples:
  - Creating comments.
  - Adding labels.
  - Creating Pull Requests.

# GitHub Actions: Tokens

- ▶ GITHUB\_TOKEN permissions configurable for each job.
- ▶ Expires when the job finishes.
- ▶ Has permissions for current repository only.



---

Workflows will not start on  
events initiated by  
GITHUB\_TOKEN!

# workflows/commit-access-review: Use get\_collaborators() function #108313

Edit <> Code

Merged **tstelllar** merged 8 commits into `llvm:main` from `tstelllar:access-review-fix` 3 weeks ago

Conversation 3 Commits 8 Checks 9 Files changed 3

+113 -107



**tstelllar** commented last month • edited

Member ...

This gets us the full list of users with commit access to llvm/llvm-project rather than the list of people in the llvm-committers team. This will ensure we are able everyone to track everyone with commit access in case they receive access some other way.



**tstelllar** added 2 commits [last month](#)

workflows/commit-access-review: Use get\_collaborators() function 2109d17

Fix token permissions 579ca1f



**github-actions** (bot) commented last month • edited

With the latest revision this PR passed the Python code formatter.



Reviewers

**boomanaiden154**

Assignees

No one—[assign yourself](#)

Labels

**github:workflow**

Projects

None yet

Milestone

No milestone

Development

Successfully merging this pull request may close these issues.

None yet

# workflows/commit-access-review: Use get\_collaborators() function #108313

Edit <> Code

Merged **tstelllar** merged 8 commits into `llvm:main` from `tstelllar:access-review-fix` 3 weeks ago

Conversation 3 Commits 8 Checks 9 Files changed 3

+113 -107



**tstelllar** commented last month • edited

Member ...

This gets us the full list of users with commit access to llvm/llvm-project rather than the list of people in the llvm-committers team. This will ensure we are able everyone to track everyone with commit access in case they receive access some other way.



**tstelllar** added 2 commits [last month](#)

workflows/commit-access-review: Use get\_collaborators() function ... ✖ 2109d17

Fix token permissions ✖ 579ca1f



**github-actions** bot commented last month • edited

...

With the latest revision this PR passed the Python code formatter.



Reviewers

boomanaiden154 ✓

Assignees

No one—[assign yourself](#)

Labels

**github:workflow**

Projects

None yet

Milestone

No milestone

Development

Successfully merging this pull request may close these issues.

None yet

# workflow/release-binaries: Checkout sources before downloading artifacts #109349

Edit <> Code

Merged tstellar merged 1 commit into llvm:main from tstellar:release-binaries-upload-fix last week

Conversation 4 Commits 1 Checks 43 Files changed 1

+8 -8

**tstellar** commented 3 weeks ago Member

The actions/checkout step will clear the current directory, so we need to checkout the sources first so that the downloaded artifacts won't be deleted.

workflow/release-binaries: Checkout sources before downloading artifacts ✓ fde570d

**tstellar** requested a review from **tru** 3 weeks ago

**llvmbot** added the **github:workflow** label 3 weeks ago

**llvmbot** commented 3 weeks ago Member

[@llvm/pr-subscribers-github-workflow](#)

Author: Tom Stellard (tstellar)

► Changes

Reviewers

- tru** ✓
- boomanaiden154** ✓

Assignees

No one—[assign yourself](#)

Labels

- github:workflow**

Projects

- LLVM Release Status**  
Status: Needs Triage

Milestone

- LLVM 19.X Release

Development

Successfully merging this pull request may close these issues.

None yet

# workflow/release-binaries: Checkout sources before downloading artifacts #109349

Edit <> Code

Merged tstellar merged 1 commit into llvm:main from tstellar:release-binaries-upload-fix last week

Conversation 4 Commits 1 Checks 43 Files changed 1

+8 -8

tstellar commented 3 weeks ago

The actions/checkout step will clear the current directory, so we need to checkout the sources first so that the downloaded artifacts won't be deleted.

workflow/release-binaries: Checkout sources before downloading artifacts ✓ fde570d

tstellar requested a review from tru 3 weeks ago

llvmbot added the github:workflow label 3 weeks ago

llvmbot commented 3 weeks ago

@llvm/pr-subscribers-github-workflow

Author: Tom Stellard (tstellar)

► Changes

Reviewers

- tru ✓
- boomanaiden154 ✓

Assignees

No one—assign yourself

Labels

github:workflow

Projects

LLVM Release Status

Status: Needs Triage

Milestone

LLVM 19.X Release

Development

Successfully merging this pull request may close these issues.

None yet

# GitHub Actions: Secrets

- ▶ Used to allow 'chaining' of workflows.
- ▶ Can grant more permissions than `GITHUB_TOKEN`.
- ▶ Can be anything e.g. pypi token, signing keys, etc.

---

Anyone with commit  
access can view secrets.

# GitHub Actions: Real World Exploits

- ▶ [pytorch](#)
- ▶ [GitHub Runner Images](#)
- ▶ [Token leaks affecting many projects](#)



# GitHub Actions: Best Practices

- ▶ Use GITHUB\_TOKEN when possible.
- ▶ Grant GITHUB\_TOKEN minimal permissions.
- ▶ Use GitHub hosted runners.
- ▶ Limit use of secrets.
- ▶ Disable workflows for first time contributors.

## Workflow file for this run

.github/workflows/new-issues.yml at a4916d2



```
1 name: Labeling new issues
2 on:
3   issues:
4     types: ['opened']
5
6 permissions:
7   contents: read
8
9 jobs:
10  automate-issues-labels:
11    permissions:
12      issues: write
13    runs-on: ubuntu-latest
14    if: github.repository == 'llvm/llvm-project'
15    steps:
16      - uses: llvm/actions/issue-labeler@main
17        with:
18          repo-token: ${ secrets.GITHUB_TOKEN }
19          configuration-path: .github/new-issues-labeler.yml
20          include-title: 1
21          include-body: 0
22          sync-labels: 0
23          enable-versioned-regex: 0
```

## Workflow file for this run

.github/workflows/new-issues.yml at a4916d2



```
1 name: Labeling new issues
2 on:
3   issues:
4     types: ['opened']
5
6   permissions:
7     contents: read
8
9 jobs:
10  automate-issues-labels:
11    permissions:
12      issues: write
13    runs-on: ubuntu-latest
14    if: github.repository == 'llvm/llvm-project'
15    steps:
16      - uses: llvm/actions/issue-labeler@main
17        with:
18          repo-token: ${ secrets.GITHUB_TOKEN }
19          configuration-path: .github/new-issues-labeler.yml
20          include-title: 1
21          include-body: 0
22          sync-labels: 0
23          enable-versioned-regex: 0
```

## Workflow file for this run

.github/workflows/new-issues.yml at a4916d2



```
1 name: Labeling new issues
2 on:
3   issues:
4     types: ['opened']
5
6 permissions:
7   contents: read
8
9 jobs:
10  automate-issue-labeler:
11    permissions:
12      issues: write
13    runs-on: ubuntu-latest
14    if: github.repository == 'llvm/llvm-project'
15    steps:
16      - uses: llvm/actions/issue-labeler@main
17        with:
18          repo-token: ${ secrets.GITHUB_TOKEN }
19          configuration-path: .github/new-issues-labeler.yml
20          include-title: 1
21          include-body: 0
22          sync-labels: 0
23          enable-versioned-regex: 0
```

# GitHub Actions: Pull Request Events

- ▶ Two types: `pull_request` and `pull_request_target`
- ▶ `pull_request` has no secret access and read-only repo access.
- ▶ `pull_request_target` has access to secrets and write access to the repo.

---

Do not use  
`pull_request_target` when  
executing untrusted code  
(e.g. CI).

# GitHub Actions: Pull Request Events

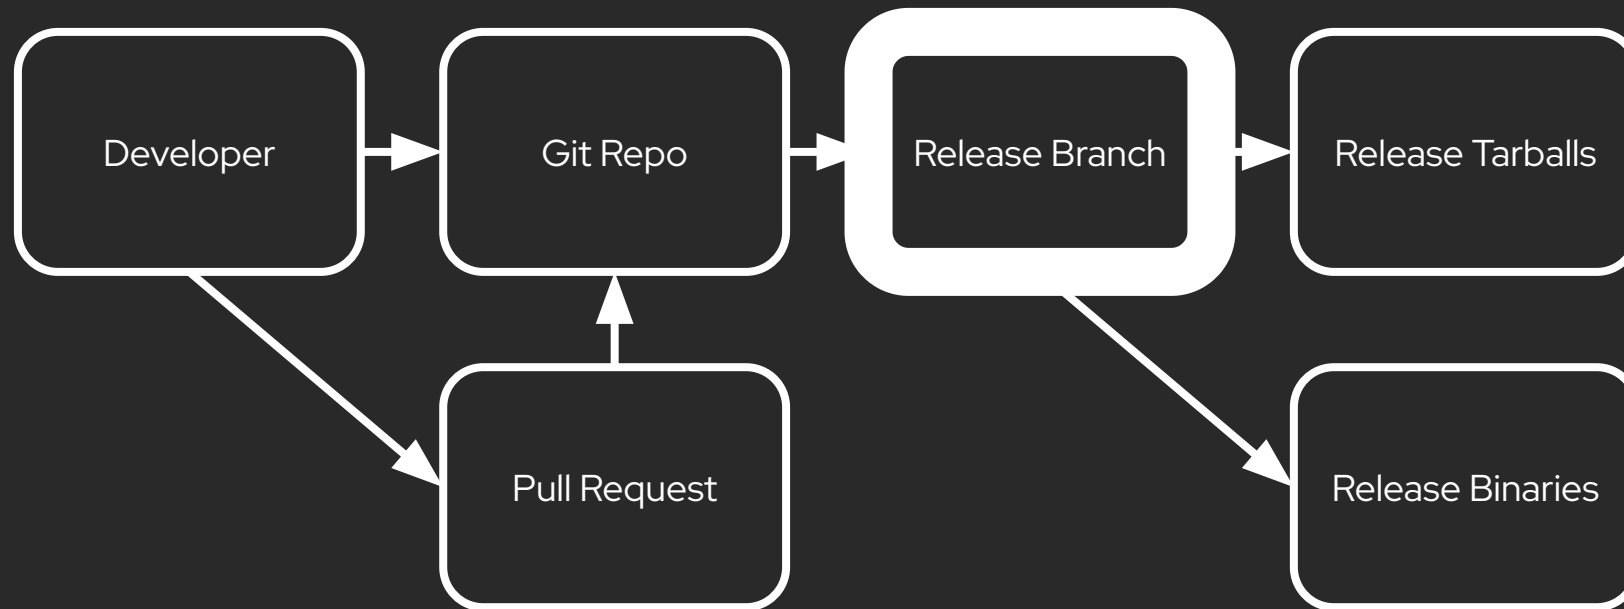
- ▶ We have three workflows using `pull_request_target`.
- ▶ Generally, considered safe since they only use code from main.
- ▶ Could be ported over to `pull_request` event.

# Other Infrastructure: Buildkite, Buildbot.

- ▶ Need to be hardened against running untrusted code.
- ▶ Using Ephemeral nodes mitigates some kinds of attacks.
- ▶ Internal CI systems carry the same risks!



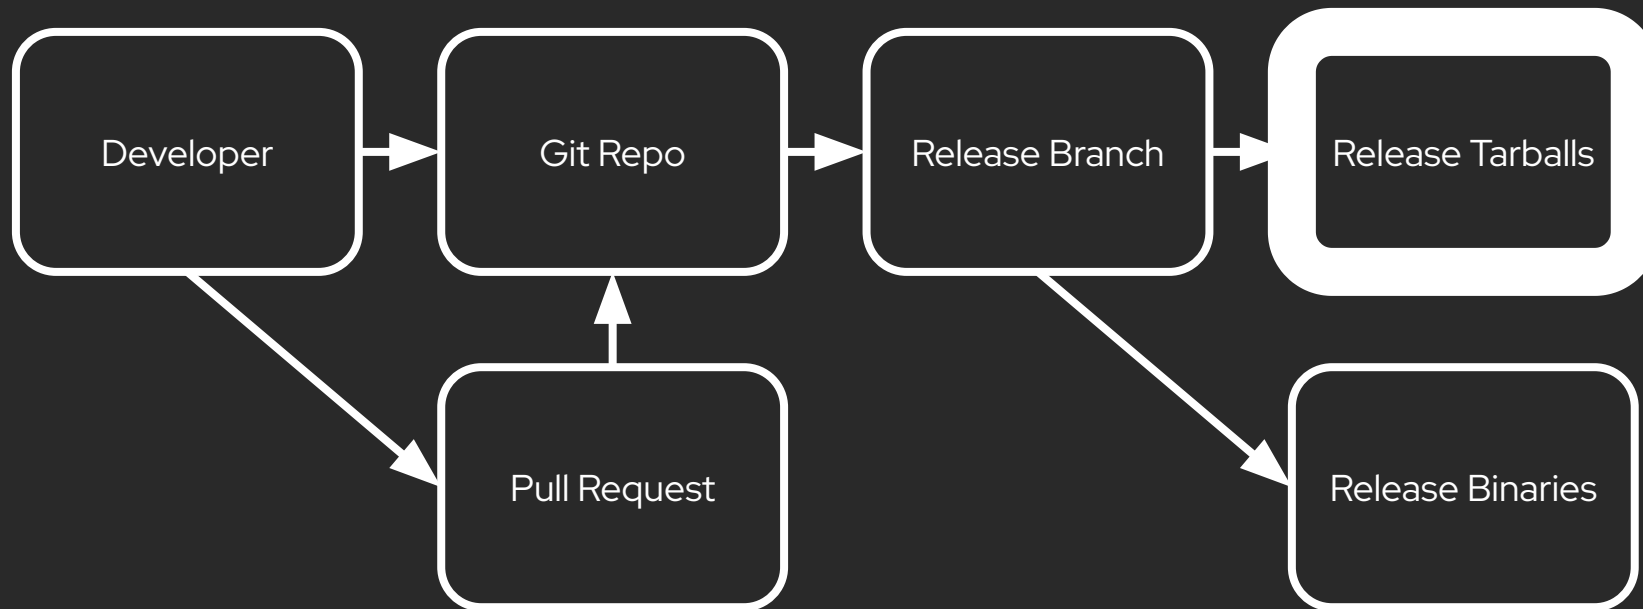
# LLVM Supply Chain



# Release Branch

- ▶ New release branch created every 6 months.
- ▶ Only release managers can commit.

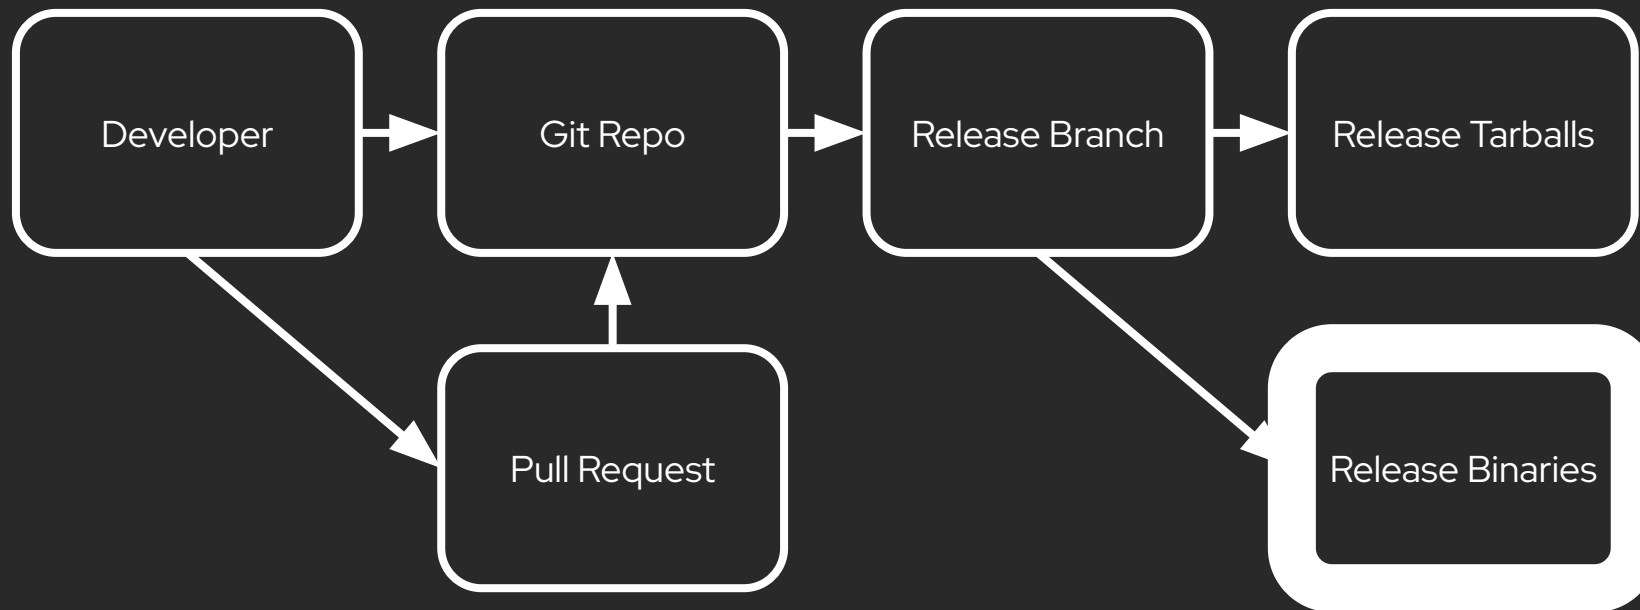
# LLVM Supply Chain



# Release Tarballs


- ▶ Generated using GitHub Actions.
- ▶ Signed by release managers.
- ▶ Provenance established using GitHub  
Artifact Attestations.

# LLVM Supply Chain



# Release Binaries

- ▶ Official binaries generated using GitHub Actions.
- ▶ Signed by release managers
- ▶ GitHub Artifact Attestations.

 LLVM-19.1.0-Linux-X64.tar.xz #2097234
[View build summary](#)[Download](#)

Created	3 weeks ago (Wed, 18 Sep 2024 04:57:41 GMT)
Commit	<a href="#">8e2dbab24276a8521d241463b4161c78bc4d39d2</a>
Subject Digest	sha256:cee77d641690466a193d9b88c89705de1c0... 
Predicate Type	<a href="https://slsa.dev/provenance/v1">https://slsa.dev/provenance/v1</a>
Workflow	<a href="#">.github/workflows/release-binaries-all.yml@refs/heads/main</a>
Verify	gh attestation verify <filename-or-url> --owner llvm --bundle <a href="#">./llvm-llvm-project-attestation-2097234.sigstore.json</a>

## Certificate Summary

Build Config Digest	8e2dbab24276a8521d241463b4161c78bc4d39d2
Build Config URI	<a href="https://github.com/llvm/llvm-project/.github/workflows/release-binaries-all.yml@refs/heads/main">https://github.com/llvm/llvm-project/.github/workflows/release-binaries-all.yml@refs/heads/main</a>
Build Signer Digest	8e2dbab24276a8521d241463b4161c78bc4d39d2
Build Signer URI	<a href="https://github.com/llvm/llvm-project/.github/workflows/release-binaries.yml@refs/heads/main">https://github.com/llvm/llvm-project/.github/workflows/release-binaries.yml@refs/heads/main</a>
Build Trigger	workflow_dispatch
Issuer	<a href="https://token.actions.githubusercontent.com">https://token.actions.githubusercontent.com</a>
Runner Invocation URI	<a href="https://github.com/llvm/llvm-project/actions/runs/10904034079/attempts/1">https://github.com/llvm/llvm-project/actions/runs/10904034079/attempts/1</a>
Runner Environment	github-hosted
Source Repository Digest	8e2dbab24276a8521d241463b4161c78bc4d39d2
Source Repository Identifier	75821432
Source Repository Owner Identifier	17149993
Source Repository Owner URI	<a href="https://github.com/llvm">https://github.com/llvm</a>
Source Repository Ref	refs/heads/main
Source Repository URI	<a href="https://github.com/llvm/llvm-project">https://github.com/llvm/llvm-project</a>
Source Repository Visibility	public

Created	3 weeks ago (Wed, 18 Sep 2024 04:57:41 GMT)
Commit	8e2dbab24276a8521d241463b4161c78bc4d39d2
Subject Digest	sha256:cee77d641690466a193d9b88c89705de1c0...
Predicate Type	<a href="https://slsa.dev/provenance/v1">https://slsa.dev/provenance/v1</a>
Workflow	<a href="#">.github/workflows/release-binaries-all.yml@refs/heads/main</a>
Verify	gh attestation verify <filename-or-url> --owner llvm --bundle <a href="#">./llvm-llvm-project-attestation-2097234.sigstore.json</a>

### Certificate Summary



Build Config Digest	8e2dbab24276a8521d241463b4161c78bc4d39d2
Build Config URI	<a href="https://github.com/llvm/llvm-project/.github/workflows/release-binaries-all.yml@refs/heads/main">https://github.com/llvm/llvm-project/.github/workflows/release-binaries-all.yml@refs/heads/main</a>
Build Signer Digest	8e2dbab24276a8521d241463b4161c78bc4d39d2
Build Signer URI	<a href="https://github.com/llvm/llvm-project/.github/workflows/release-binaries.yml@refs/heads/main">https://github.com/llvm/llvm-project/.github/workflows/release-binaries.yml@refs/heads/main</a>
Build Trigger	workflow_dispatch
Issuer	<a href="https://token.actions.githubusercontent.com">https://token.actions.githubusercontent.com</a>
Runner Invocation URI	<a href="https://github.com/llvm/llvm-project/actions/runs/10904034079/attempt/1">https://github.com/llvm/llvm-project/actions/runs/10904034079/attempt/1</a>
Runner Environment	github-hosted
Source Repository Digest	8e2dbab24276a8521d241463b4161c78bc4d39d2
Source Repository Identifier	75821432
Source Repository Owner Identifier	17149993
Source Repository Owner URI	<a href="https://github.com/llvm">https://github.com/llvm</a>
Source Repository Ref	refs/heads/main
Source Repository URI	<a href="https://github.com/llvm/llvm-project">https://github.com/llvm/llvm-project</a>
Source Repository Visibility	public



Created	3 weeks ago (Wed, 18 Sep 2024 04:57:41 GMT)
Commit	8e2dbab24276a8521d241463b4161c78bc4d39d2
Subject Digest	sha256:cee77d641690466a193d9b88c89705de1c0...
Predicate Type	<a href="https://slsa.dev/provenance/v1">https://slsa.dev/provenance/v1</a>
Workflow	<a href="#">.github/workflows/release-binaries-all.yml@refs/heads/main</a>
Verify	gh attestation verify <filename-or-url> --owner llvm --bundle <a href="#">./llvm-llvm-project-attestation-2097234.sigstore.json</a>

### Certificate Summary

Build Config Digest	8e2dbab24276a8521d241463b4161c78bc4d39d2
Build Config URI	<a href="https://github.com/llvm/llvm-project/.github/workflows/release-binaries-all.yml@refs/heads/main">https://github.com/llvm/llvm-project/.github/workflows/release-binaries-all.yml@refs/heads/main</a>
Build Signer Digest	8e2dbab24276a8521d241463b4161c78bc4d39d2
Build Signer URI	<a href="https://github.com/llvm/llvm-project/.github/workflows/release-binaries.yml@refs/heads/main">https://github.com/llvm/llvm-project/.github/workflows/release-binaries.yml@refs/heads/main</a>
Build Trigger	workflow_dispatch
Issuer	<a href="https://token.actions.githubusercontent.com">https://token.actions.githubusercontent.com</a>
Runner Invocation URI	<a href="https://github.com/llvm/llvm-project/actions/runs/10904034079/attempts/1">https://github.com/llvm/llvm-project/actions/runs/10904034079/attempts/1</a>
Runner Environment	github-hosted
Source Repository Digest	8e2dbab24276a8521d241463b4161c78bc4d39d2
Source Repository Identifier	75821432
Source Repository Owner Identifier	17149993
Source Repository Owner URI	<a href="https://github.com/llvm">https://github.com/llvm</a>
Source Repository Ref	refs/heads/main
Source Repository URI	<a href="https://github.com/llvm/llvm-project">https://github.com/llvm/llvm-project</a>
Source Repository Visibility	public

 <a href="#">LLVM-19.1.1-Linux-X64.tar.xz</a>	1.54 GB	3 days ago
 <a href="#">LLVM-19.1.1-Linux-X64.tar.xz.jsonl</a>	10.2 KB	3 days ago

```
[fedora@tstellar-uploader 19.1.1]$ gh attestation verify -R llvm/llvm-project -b LLVM-19.1.1-Linux-X64.tar.xz.jsonl LLVM-19.1.1-Linux-X64.tar.xz
Loaded digest sha256:8204de000b6a6921f0572e038336601e3225898e9a253c8aaa43b0a5fae8a4ce for file://LLVM-19.1.1-Linux-X64.tar.xz
Loaded 1 attestation from LLVM-19.1.1-Linux-X64.tar.xz.jsonl
✓ Verification succeeded!

sha256:8204de000b6a6921f0572e038336601e3225898e9a253c8aaa43b0a5fae8a4ce was attested by:
REPO          PREDICATE_TYPE          WORKFLOW
llvm/llvm-project  https://slsa.dev/provenance/v1  .github/workflows/release-binaries.yml@refs/tags/llvmorg-19.1.1
```

---

Anyone with commit  
access can upload assets  
to the release page!

# Release Assets

- ▶ No fine-grained permissions.
- ▶ We have an audit job that checks assets once per hour.
- ▶ Discussed moving assets to a different repository.

---

# Should you download this file?

<https://github.com/llvm/llvm-project/files/13166493/LLVM-19.1.1-Linux-X64.tar.xz>

---

# Should you download this file?

<https://github.com/llvm/llvm-project/files/13166493/LLVM-19.1.1-Linux-X64.tar.xz>

No!

---

# User Upload (could be malicious):

<https://github.com/llvm/llvm-project/files/13166493/LLVM-19.1.1-Linux-X64.tar.xz>

# Official Upload:

<https://github.com/llvm/llvm-project/releases/download/llvmorg-19.1.1/LLVM-19.1.1-Linux-X64.tar.xz>

---

# XZ Attack

<https://www.openwall.com/lists/oss-security/2024/03/29/4>

<https://boehs.org/node/everything-i-know-about-the-xz-backdoor>



# XZ Attack: Mechanics

- ▶ 'Trusted' user granted commit access.
- ▶ Malicious test binaries pushed to the repo.
- ▶ Release tarball hosting moved to github.
- ▶ Release tarballs modified to include script that injected malicious code into XZ.

# XZ Attack: Results

- ▶ RSA\_public\_decrypt calls redirected to malicious code.

# LLVM vs XZ: How do we compare?

01 Title of table

	XZ	LLVM
Commit Access Wait Time	8 Months [1]	Days
Test Binaries in Repo	Yes	Yes
Mandatory Code Review	No	No
Maintainers	1	Many

[1] <https://boehs.org/node/everything-i-know-about-the-xz-backdoor>

---

# What's next?

# What's Next? (my opinions)

- ▶ Having low bar for commit access is OK.
- ▶ Having minimal commit rules is OK.
- ▶ Having both at the same time is very risky.
- ▶ We should make some changes.

# Ideas

- ▶ Mandatory pull requests.
- ▶ Mandatory review for all changes.
- ▶ More strict requirements for commit access.
- ▶ We need to balance convenience and security.

# Advice for downstream.

- ▶ Know your risks.
- ▶ Contribute to upstream.
- ▶ Hire someone to work on this full time.
- ▶ Donate to the LLVM Foundation!

---

Conclusion: This is important.





---

# Acknowledgements

# Thank you

Red Hat is the world's leading provider of enterprise open source software solutions. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500.



[linkedin.com/company/red-hat](https://www.linkedin.com/company/red-hat)



[youtube.com/user/RedHatVideos](https://www.youtube.com/user/RedHatVideos)



[facebook.com/redhatinc](https://www.facebook.com/redhatinc)



[twitter.com/RedHat](https://twitter.com/RedHat)