

3 июня 2025 📍 Москва, LOFT HALL#2

# БЕКОН'25

Конференция по БЕзопасности  
КОНтейнеров и контейнерных сред

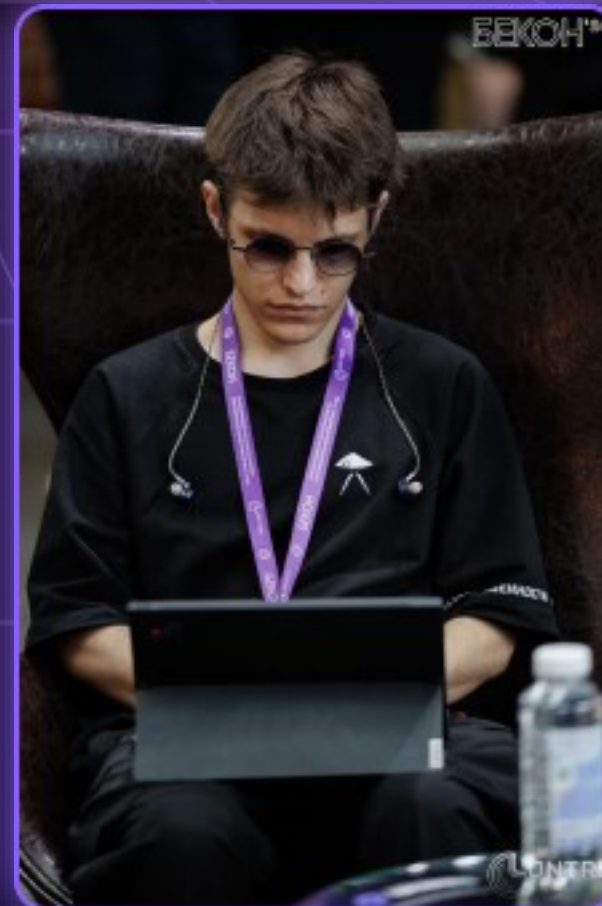
## Чем собирать контейнеры параноику?

Кожуховский Михаил

Flowmaster

# whoami

- Архитектор платформы Flowmaster
- Поддерживаю и развиваю CI/CD инфраструктуру в компании с 2020 года
- 11 лет в программировании, 7 из них в коммерческой разработке
- Интересуюсь ИБ с 2017



- Небольшое вступление
- Обзор возможных решений
  - Kaniko
  - Buildah
  - Builkit
- Небольшой эксперимент с unshare
- Альтернативный подход к сборке(Ko/Jib)
- Итоги

# Вступление



# CI Threat Matrix

Initial Access	Execution	Persistence	Privilege Escalation	Defense Evasion	Credential Access	Lateral Movement	Exfiltration	Impact
Supply Chain Compromise on CI/CD	Modify CI/CD Configuration	Compromise CI/CD Server	Get credential for Deployment(CD) on CI stage	Add Approver using Admin permission	Dumping Env Variables in CI/CD	Exploitation of Remote Services	Exfiltrate data in Production environment	Denial of Services
Valid Account of Git Repository (Personal Token, SSH key, Login password, Browser Cookie)	Inject code to IaC configuration	Implant CI/CD runner images	Privileged Escalation and compromise other CI/CD pipeline	Bypass Review	Access to Cloud Metadata	(Monorepo) Get credential of different folder's context	Clone Git Repositories	
Valid Account of CI/CD Service (Personal Token, Login password, Browser Cookie)	Inject code to source code	Modify CI/CD Configuration		Access to Secret Manager from CI/CD kicked by different repository	Read credentials file	Privileged Escalation and compromise other CI/CD pipeline		
Valid Admin account of Server hosting Git Repository	Supply Chain Compromise on CI/CD	Inject code to IaC configuration		Modify Caches of CI/CD	Get credential from CI/CD Admin Console			
	Inject bad dependency	Inject code to source code		Implant CI/CD runner images				
	SSH to CI/CD pipelines	Inject bad dependency						
	Modify the configuration of Production environment							
	Deploy modified applications or server images to production environment							

# CI Threat Matrix

Initial Access	Execution	Persistence	Privilege Escalation	Defense Evasion	Credential Access	Lateral Movement	Exfiltration	Impact
Supply Chain Compromise on CI/CD	Modify CI/CD Configuration	Compromise CI/CD Server	Get credential for Deployment(CD) on CI stage	Add Approver using Admin permission	Dumping Env Variables in CI/CD	Exploitation of Remote Services	Exfiltrate data in Production environment	Denial of Services
Valid Account of Git Repository (Personal Token, SSH key, Login password, Browser Cookie)	Inject code to IaC configuration	Implant CI/CD runner images	Privileged Escalation and compromise other CI/CD pipeline	Bypass Review	Access to Cloud Metadata	(Monorepo) Get credential of different folder's context	Clone Git Repositories	
Valid Account of CI/CD Service (Personal Token, Login password, Browser Cookie)	Inject code to source code	Modify CI/CD Configuration		Access to Secret Manager from CI/CD kicked by different repository	Read credentials file	Privileged Escalation and compromise other CI/CD pipeline		
Valid Admin account of Server hosting Git Repository	Supply Chain Compromise on CI/CD	Inject code to IaC configuration		Modify Caches of CI/CD	Get credential from CI/CD Admin Console			
	Inject bad dependency	Inject code to source code		Implant CI/CD runner images				
	SSH to CI/CD pipelines	Inject bad dependency						
	Modify the configuration of Production environment							
	Deploy modified applications or server images to production environment							

# K8s Threat Matrix

Initial Access	Execution	Persistence	Privilege Escalation	Defense Evasion	Credential Access	Discovery	Lateral Movement	Impact
Using Cloud credentials	Exec into container	Backdoor container	Privileged container	Clear container logs	List K8S secrets	Access the K8S API server	Access cloud resources	Data Destruction
Compromised images in registry	bash/cmd inside container	Writable hostPath mount	Cluster-admin binding	Delete K8S events	Mount service principal	Access Kubelet API	Container service account	Resource Hijacking
Kubeconfig file	New container	Kubernetes CronJob	hostPath mount	Pod / container name similarity	Access container service account	Network mapping	Cluster internal networking	Denial of service
Application vulnerability	Application exploit (RCE)		Access cloud resources	Connect from Proxy server	Applications credentials in configuration files	Access Kubernetes dashboard	Applications credentials in configuration files	
Exposed Dashboard	SSH server running inside container					Instance Metadata API	Writable volume mounts on the host	
							Access Kubernetes dashboard	
							Access tiller endpoint	



Initial Access	Execution	Persistence	Privilege Escalation	Defense Evasion	Credential Access	Lateral Movement	Exfiltration	Impact
Supply Chain Compromise on CI/CD	Modify CI/CD Configuration	Compromise CI/CD Server	Get credential for Deployment(CD) on CI stage	Add Approver using Admin permission	Dumping Env Variables in CI/CD	Exploitation of Remote Services	Exfiltrate data in Production environment	Denial of Services
Valid Account of Git Repository (Personal Token, SSH key, Login password, Browser Cookie)	Inject code to IaC configuration	Implant CI/CD runner images	Privileged Escalation and compromise other CI/CD pipeline	Bypass Review	Access to Cloud Metadata	(Monorepo) Get credential of different folder's context	Clone Git Repositories	
Valid Account of CI/CD Service (Personal Token, Login password, Browser Cookie)	Inject code to source code	Modify CI/CD Configuration		Access to Secret Manager from CI/CD kicked by different repository	Read credentials file	Privileged Escalation and compromise other CI/CD pipeline		
Valid Admin account of Server hosting Git Repository	Supply Chain Compromise on CI/CD	Inject code to IaC configuration		Modify Caches of CI/CD	Get credential from CI/CD Admin Console			
	Inject bad dependency	Inject code to source code		Implant CI/CD runner images				
	SSH to CI/CD pipelines	Inject bad dependency						
	Modify the configuration of Production environment							
	Deploy modified applications or server images to production environment							
Initial Access	Execution	Persistence	Privilege Escalation	Defense Evasion	Credential Access	Discovery	Lateral Movement	Impact
Using Cloud credentials	Exec into container	Backdoor container	Privileged container	Clear container logs	List K8S secrets	Access the K8S API server	Access cloud resources	Data Destruction
Compromised images in registry	bash/cmd inside container	Writable hostPath mount	Cluster-admin binding	Delete K8S events	Mount service principal	Access Kubelet API	Container service account	Resource Hijacking
Kubeconfig file	New container	Kubernetes CronJob	hostPath mount	Pod / container name similarity	Access container service account	Network mapping	Cluster internal networking	Denial of service
Application vulnerability	Application exploit (RCE)		Access cloud resources	Connect from Proxy server	Applications credentials in configuration files	Access Kubernetes dashboard	Applications credentials in configuration files	
Exposed Dashboard	SSH server running inside container					Instance Metadata API	Writable volume mounts on the host	
							Access Kubernetes dashboard	
							Access tiller endpoint	



Initial Access	Execution	Persistence	Privilege Escalation	Defense Evasion	Credential Access	Lateral Movement	Exfiltration	Impact
Supply Chain Compromise on CI/CD	Modify CI/CD Configuration	Compromise CI/CD Server	Get credential for Deployment(CD) on CI stage	Add Approver using Admin permission	Dumping Env Variables in CI/CD	Exploitation of Remote Services	Exfiltrate data in Production environment	Denial of Services
Valid Account of Git Repository (Personal Token, SSH key, Login password, Browser Cookie)	Inject code to IaC configuration	Implant CI/CD runner images	Privileged Escalation and compromise other CI/CD pipeline	Bypass Review	Access to Cloud Metadata	(Monorepo) Get credential of different folder's context	Clone Git Repositories	
Valid Account of CI/CD Service (Personal Token, Login password, Browser Cookie)	Inject code to source code	Modify CI/CD Configuration		Access to Secret Manager from CI/CD kicked by different repository	Read credentials file	Privileged Escalation and compromise other CI/CD pipeline		
Valid Admin account of Server hosting Git Repository	Supply Chain Compromise on CI/CD	Inject code to IaC configuration		Modify Caches of CI/CD	Get credential from CI/CD Admin Console			
	Inject bad dependency	Inject code to source code		Implant CI/CD runner images				
	SSH to CI/CD pipelines	Inject bad dependency						
	Modify the configuration of Production environment							
	Deploy modified applications or server images to production environment							



Initial Access	Execution	Persistence	Privilege Escalation	Defense Evasion	Credential Access	Discovery	Lateral Movement	Impact
Using Cloud credentials	Exec into container	Backdoor container	Privileged container	Clear container logs	List K8S secrets	Access the K8S API server	Access cloud resources	Data Destruction
Compromised images in registry	bash/cmd inside container	Writable hostPath mount	Cluster-admin binding	Delete K8S events	Mount service principal	Access Kubelet API	Container service account	Resource Hijacking
Kubeconfig file	New container	Kubernetes CronJob	hostPath mount	Pod / container name similarity	Access container service account	Network mapping	Cluster internal networking	Denial of service
Application vulnerability	Application exploit (RCE)		Access cloud resources	Connect from Proxy server	Applications credentials in configuration files	Access Kubernetes dashboard	Applications credentials in configuration files	
Exposed Dashboard	SSH server running inside container					Instance Metadata API	Writable volume mounts on the host	
							Access Kubernetes dashboard	
							Access tiller endpoint	

Supply chain

Побеги из контейнеров

Внутренние нарушители  
тоже бывают

А если переменные украдут?

А куда из раннера  
получится  
распространиться?

Кластер могут захватить

Разработчики жалуются  
что образы долго  
собираются

А мы точно все капы дрогнули?





# 2 самых опасных файла в репозитории

БЕКОН

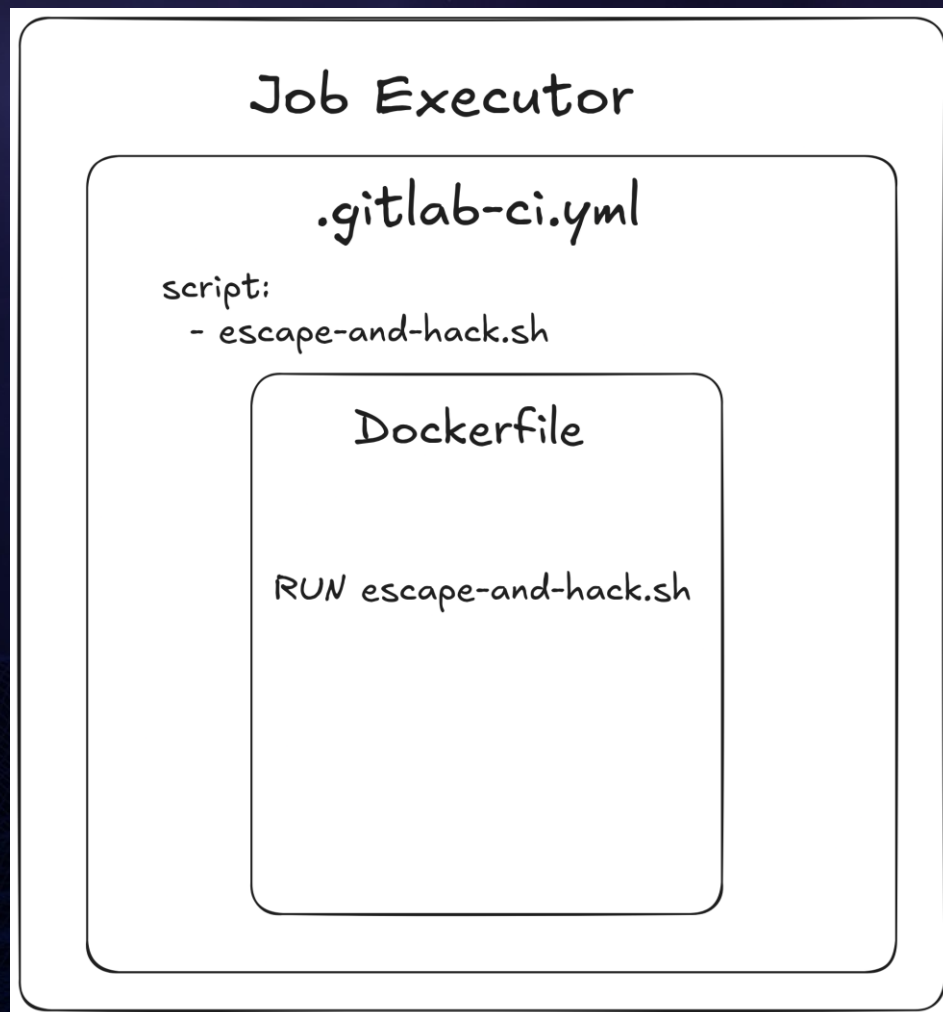




.gitlab-ci.yml



Dockerfile



.gitlab-ci.yml



Dockerfile

# Ищем решение



# Чем можно собирать образы

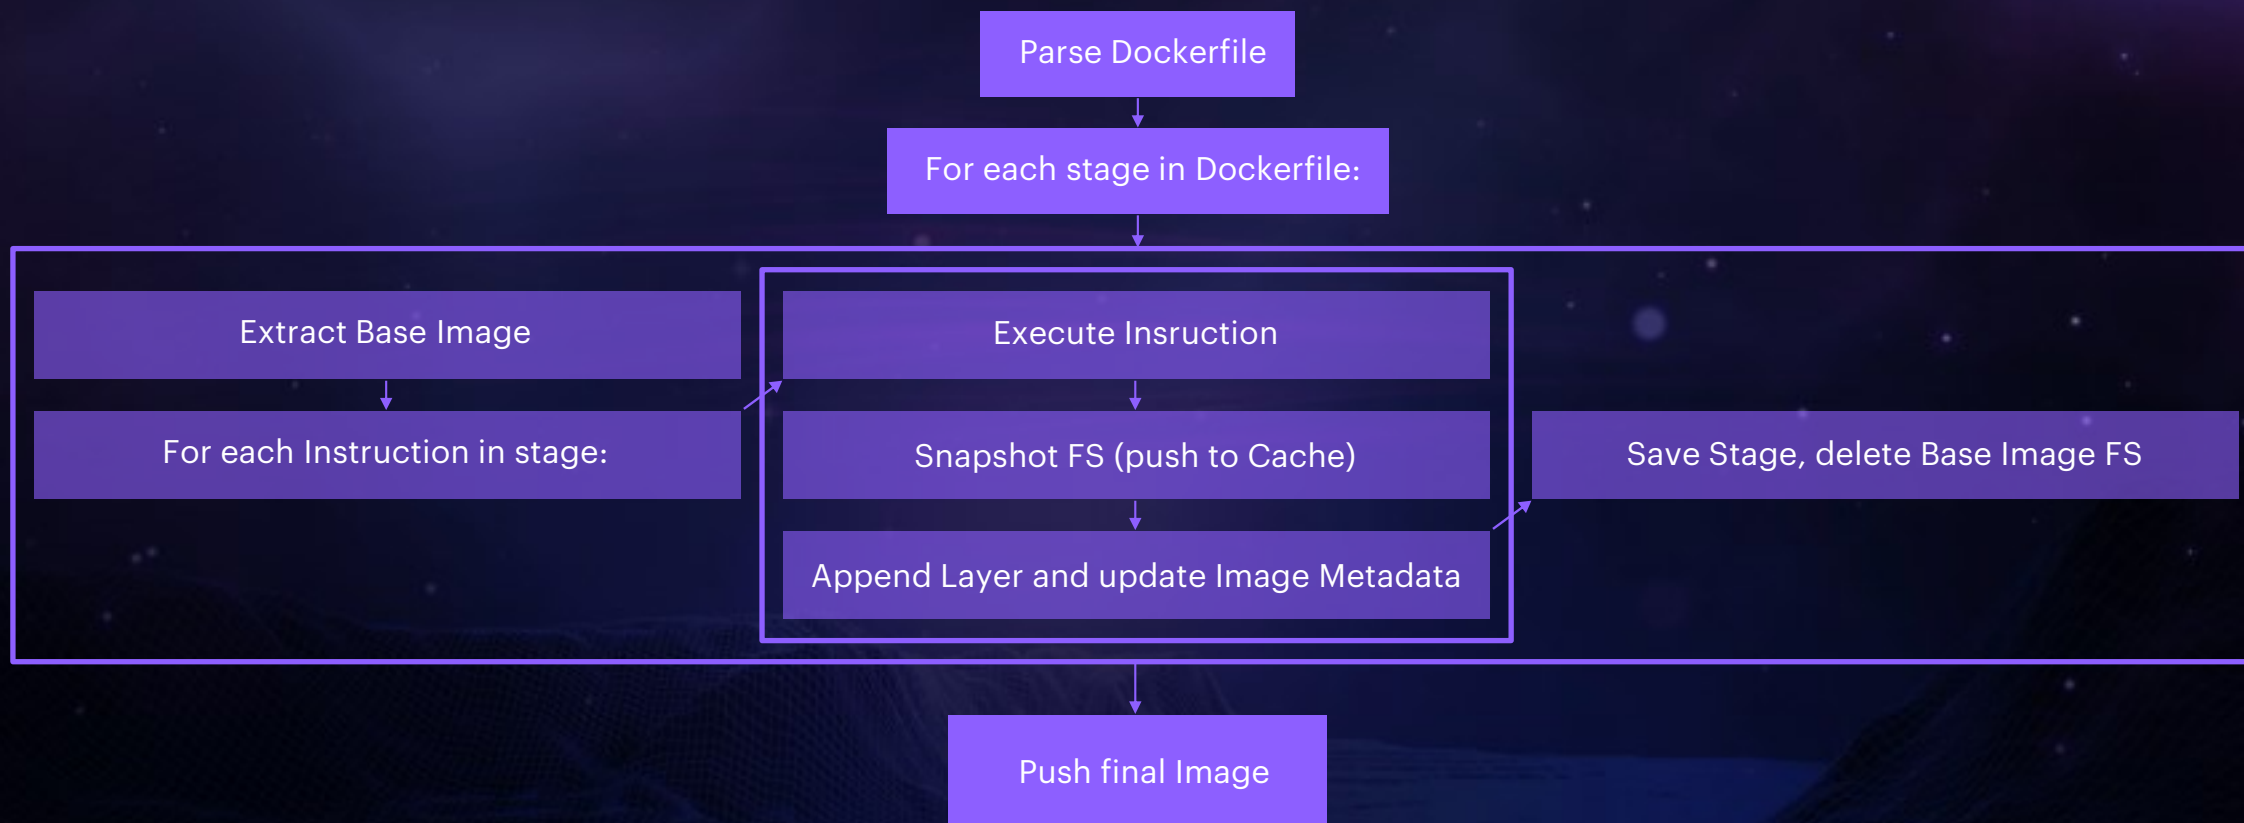
## Low level

- Kaniko
- Buildah
- Buildkit

## High level

- Buildpacks
- Nixpacks
- Werf
- Buildx
- Etc...

# Среднее устройство low level сборщика



# Kaniko



- Неофициальный проект Google
- Популярен, много сторонних гайдов
- [Мейнтейнеры уходят из проекта](#)
- [Поддерживает не весь синтаксис докерфайлов](#)
- Изоляция через chroot
- Требует root в контейнере



**Support bind Build Mounts (RUN --mount=type=bind / RUN --mount) like BuildKit and Buildah do #1568**

Open

guillaume-d opened on Feb 2, 2021

For more details see <https://github.com/moby/buildkit/blob/master/frontend/dockerfile/docs/syntax.md>

That would help avoiding contortions in the Dockerfile due to COPY's silly "partial flattening" of file hierarchies...

Also for performance some of the other types would be nice:

- `RUN --mount=type=tmpfs`
- `RUN --mount=type=cache`, but this is [Mount cache directory inside build container #969](#) already (please vote at [Mount cache directory inside build container #969 \(comment\)](#) if you are interested!) but these as a first step may be no-op implementations (just to have better compatibility with BuildKit), whereas `RUN --mount=type=bind` cannot be faked

105 5

Assignees: No one assigned

Labels: area/dockerfile-core, differs-from-docker, issue/mount, kind/feature-request, priority/p1, work-in-progress

Type: No type

Projects:

**State of Kaniko: Unmaintained? #3348**

Open

lc-guy opened on Oct 23, 2024 · edited by lc-guy

It seems that all the remaining maintainers of this project have [retired](#). Could we get some clarification regarding Kaniko's development situation right now? Is the project being abandoned?

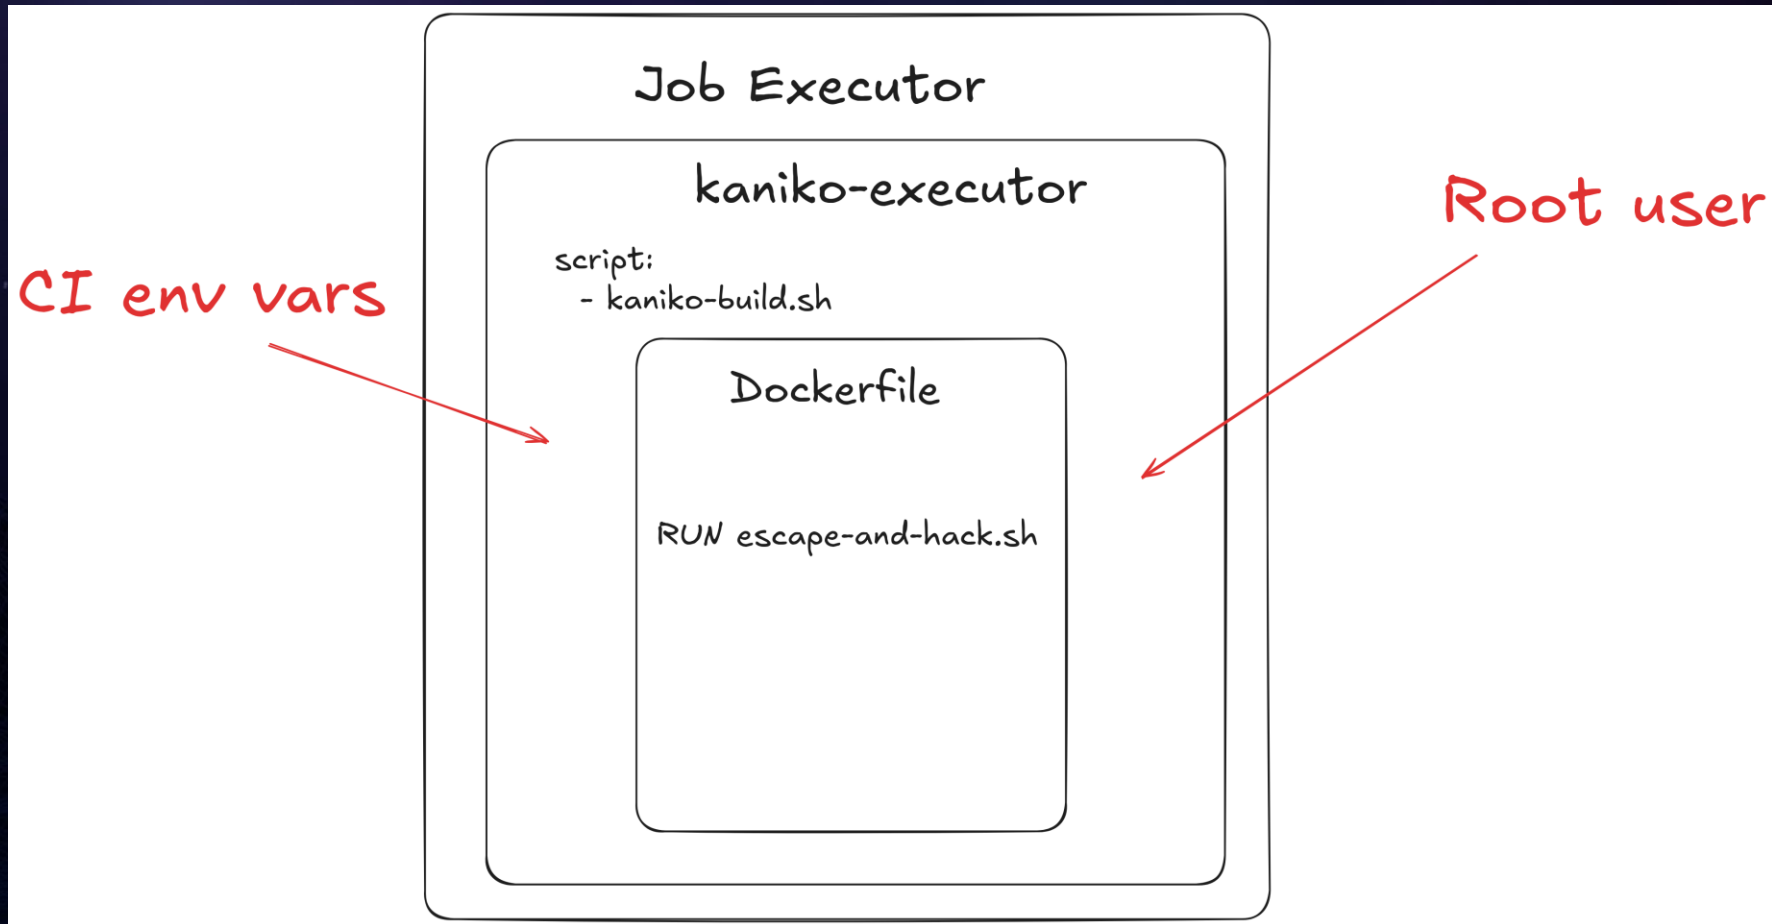
The [Gitlab CI docs](#), for example, reference Kaniko as the first choice when it comes to building container images without doing privileged Docker-in-Docker. Buildah has emerged as an alternative since, but it requires using vfs (which slows down builds quite a bit) and [seems to require additional privileges on containers now](#), which kind of ruins the point.

Thank you in advance.

35 17 72

# Kaniko problems

- Работает от рута
- При побеге из chroot доступны все переменные окружения CI джобы



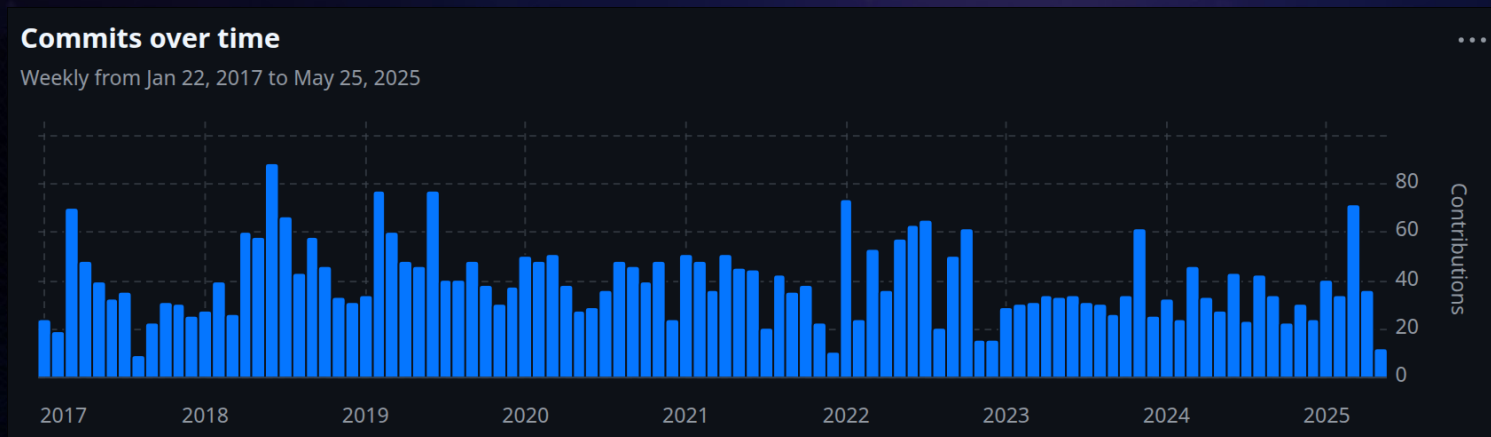
# Buildah



- В CNCF как часть Podman Container Tools
- Документация могла быть лучше, но для начала хватит
- Изоляция через chroot или вложенную контейнеризацию
- Может быть rootless(только при использовании chroot)
- Проект стабильно развивается

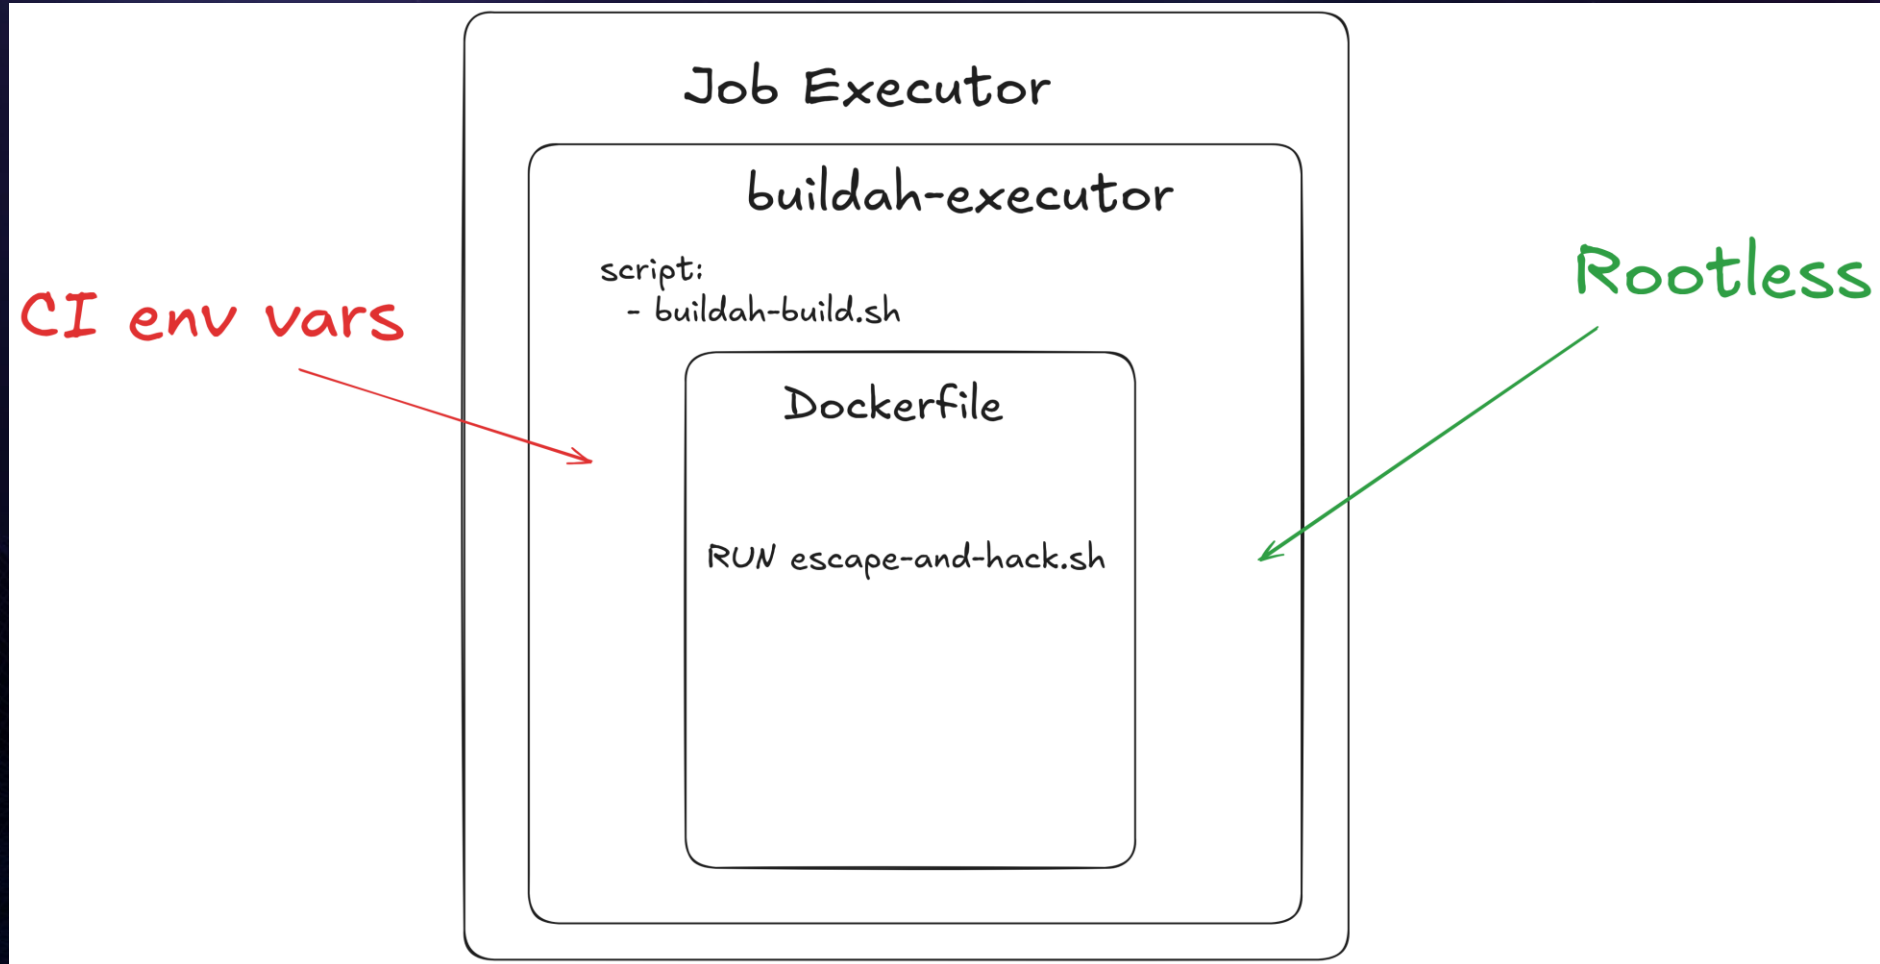


**Tutorial: Use Buildah in a rootless container with GitLab Runner Operator on OpenShift**



# Buildah problems

- При побеге из chroot доступны все переменные окружения CI джобы



# Buildkit



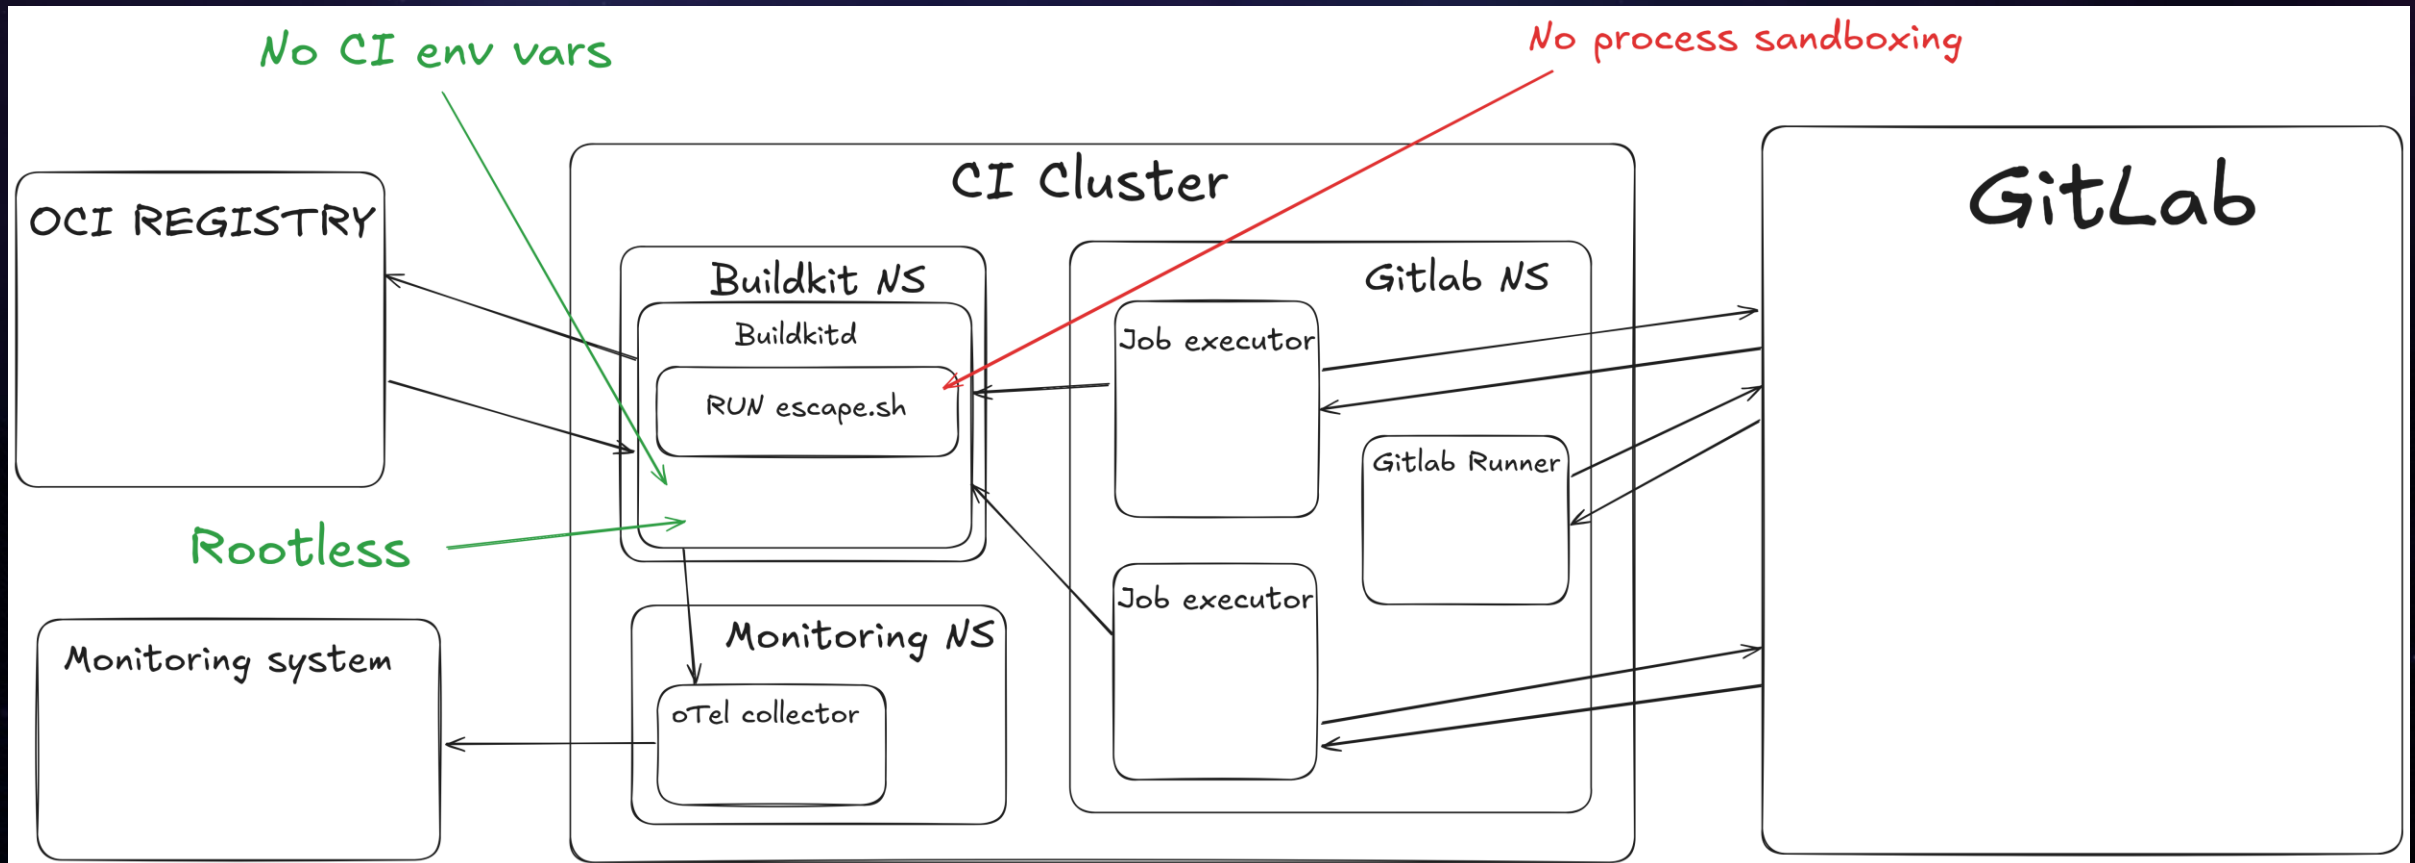
- Часть moby project
- Мало документации, порой проще почитать исходники
- [Много оптимизаций](#) вокруг скорости сборки
- Изоляция через вложенную контейнеризацию
  - Запускает с дефолтным seccomp профилем
  - Можно подложить AppArmor профиль
- Может быть rootless
- Работает как демон, можно запускать как внутри CI джобы на каждую сборку, так и держать в кластере общий демон
- Максимальная совместимость с окружением на машинах разработов
- Умеет в opentelemetry



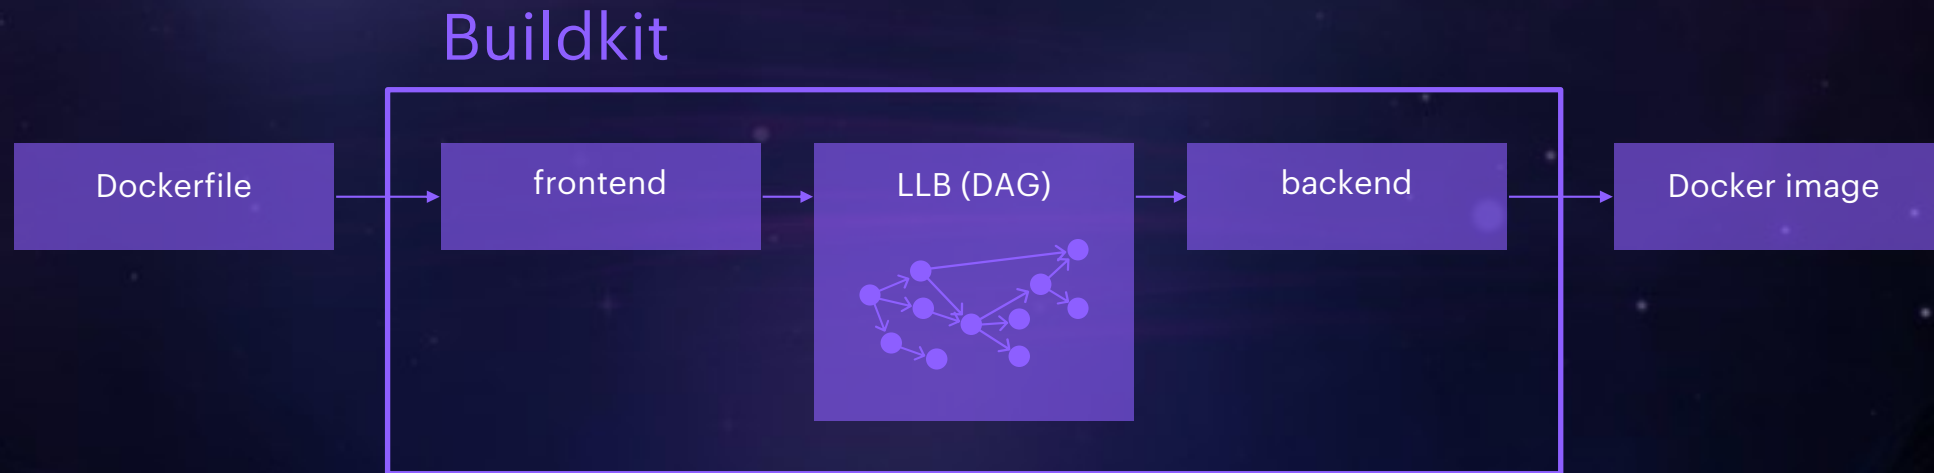


# Buildkit problems

- Архитектура на первый взгляд сложнее
- Приходится отключать process sandboxing в rootless режиме



# Buildkit Flow Diagram



**Время пробовать**





```
FROM docker.io/library/debian:bookworm
```

```
RUN useradd user
```

```
USER user
```

```
RUN whoami # user
```

```
RUN unshare -r whoami # Operation not permitted OR root
```



```
FROM docker.io/library/debian:bookworm
```

```
RUN useradd user
```

```
USER user
```

```
RUN whoami # user
```

```
RUN unshare -r whoami # Operation not permitted OR root
```

## unshare

### Failed jobs



unshare-buildah



unshare-buildkit



unshare-kaniko



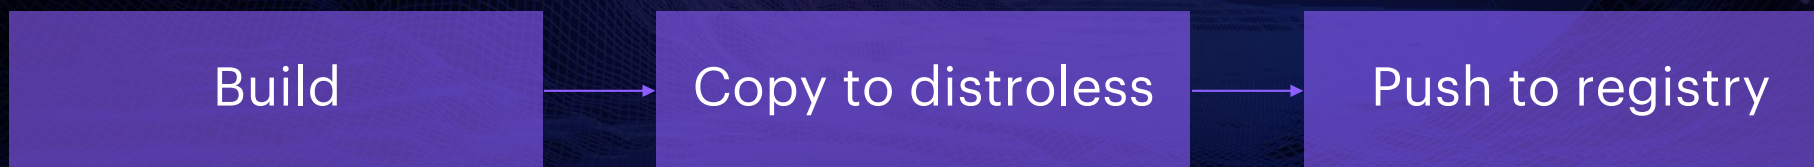
# Альтернативный способ



- Собирают бинарник/jar и копируют на последний слой distroless образа
- Ко передан в CNCF(sandbox)
- Jib поддерживается GoogleContainerTools
- Меньше компонентов – проще

## Минусы

- Ко не поддерживает CGO
- Не хватает возможностей кастомизации
- Нет поддержки кэшей вне файловой системы
- Не решают проблемы изоляции от избыточных переменных окружения





# Или свой велосипед

БЕКОН

- Лишние затраты на разработку своего решения
- Из-за потери слоев теряем в гибкости
- Distroless не всегда оправдано использовать



# Итоги



## Итого

	Kaniko	Buildah	Buildkit
Расширенный синтаксис	✗	⚠	✓
Скорость сборки	2 / 5	3 / 5	4.5 / 5
Кэш во внешних хранилищах	✓	✓	✓
Документация и сторонние гайды	4 / 5	3 / 5	2 / 5
Rootless image	✗	✓	✓
Работает с seccomp	✓	✓	⚠
Настраиваемые зеркала	✓	✓	✓
OCI spec features	Only annotations	Only annotations	in-toto attestations(SBOM, SLSA Provenance)
Дополнительный функционал		Sbom scanner	Custom dockerfile syntax, sbom scanner, daemon mode, etc

# Субъективное мнение

Kaniko

просто не надо использовать

Buildah

хорошее решение по умолчанию

Builkit

лучшее что есть если не пугает порог входа

Ко и Jib  
или свой  
аналог

подойдет не всем, но познакомиться стоит

# Этот доклад неплохо дополняют

БЕКОН

- [Побеги из контейнеров: 2024 Kubernetes Edition - Дмитрий Евдокимов и Николай Панченко](#)
- [Buildkit in depth](#)
- [BuildKit: A Modern Builder Toolkit on Top of containerd - Tonis Tiigi & Akihiro Suda](#)
- [Павел Сорокин - OPA с shared Docker executor](#)



3 июня 2025 📍 Москва, LOFT HALL#2  
Конференция по БЕзопасности  
КОНтейнеров и контейнерных сред



📍 @n0nvme

✉️ mk@apsolutions.ru

🌐 <https://flowmr.ru>