# UPGRADVISOR: Early Adopting Dependency Updates Using Hybrid Program Analysis and Hardware Tracing

**Yaniv David**[1], Xudong Sun[2], Raphael J Sofaer[1], Aditya Senthilnathan[3], Junfeng Yang[1], Zhiqiang Zuo[2], Guoqing Harry Xu[4], Jason Nieh[1] and Ronghui Gu[1]

[1]Columbia University [2]Nanjing University [3]IIT Delhi [4]UCLA




# Modern Software Development is fast-paced

- Facebook(Meta) updates their front-end **three times a day**, and release new iOS and Android apps **every week**

- A key enabler for new features is **pre-exiting libraries**

- Average of 12 direct dependencies 100+ transitive dependencies
  (Our survey of top-stared Python, JS, Ruby GitHub projects)

# Dependency Update Adoption Is Slow

- Dependencies' developers are releasing updates frequently, too

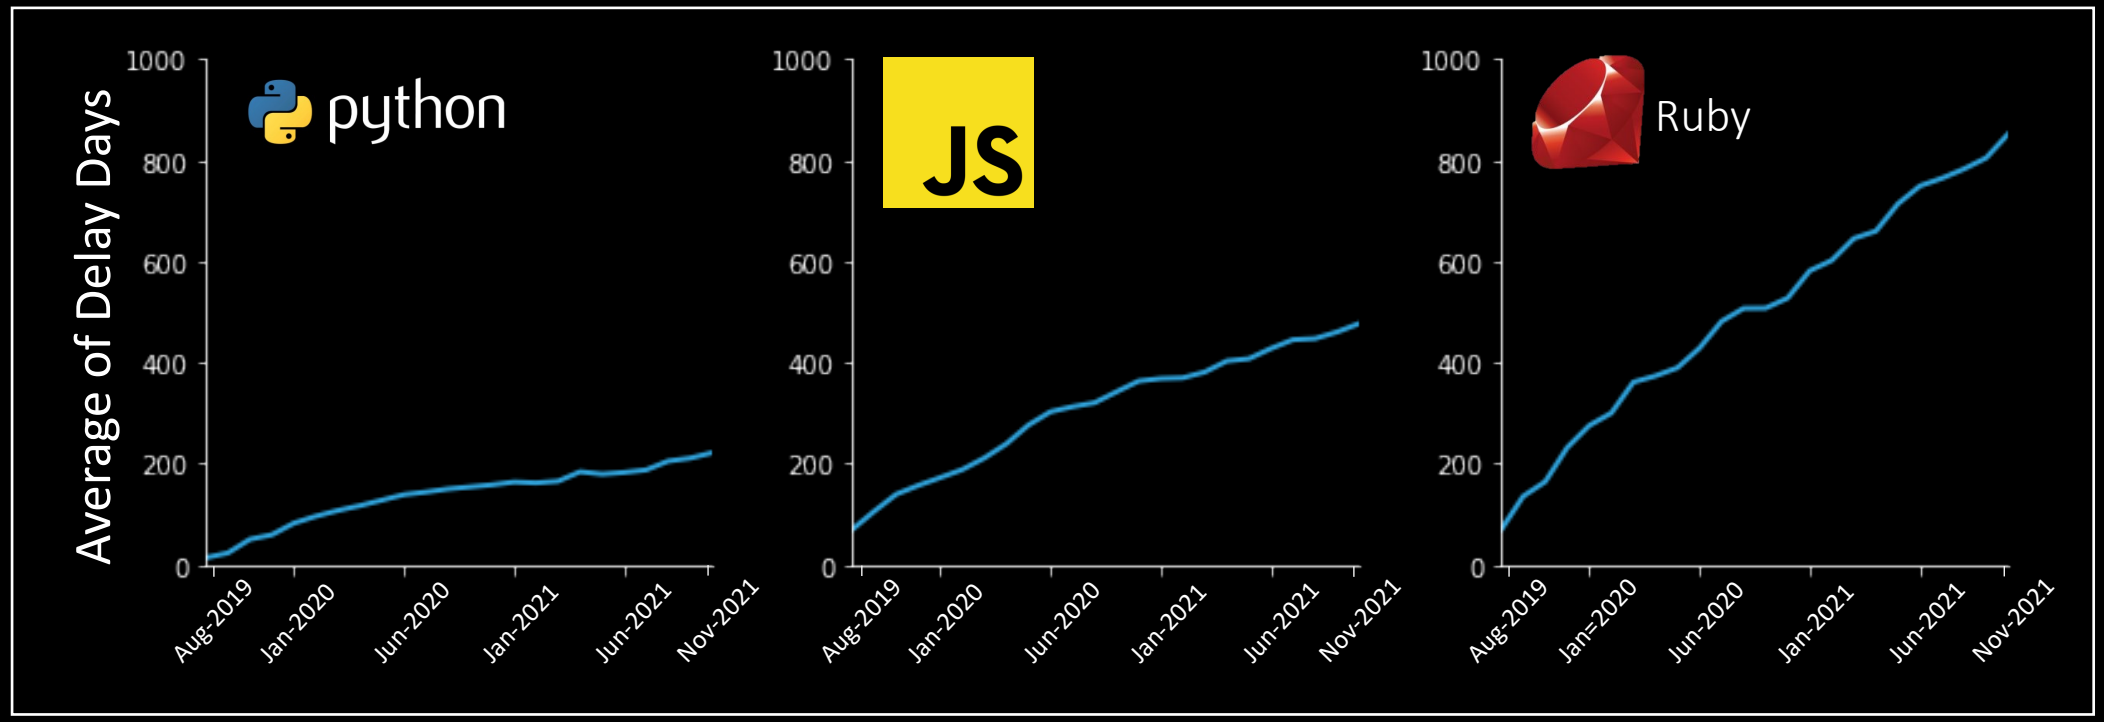


- Currently averaging **400 days** in update adaption delay

# Dependency Delays Lead to Bad Consequences

- Fixed bugs in dependencies continue to affect applications
- Closed security holes put dependent applications at risk
- Conflicts arise in transitive dependency graphs
  - Some can be resolved by using the oldest supported version
  - Other fall into a “dependency hell”

# Dependency Delays Lead to Bad Consequences

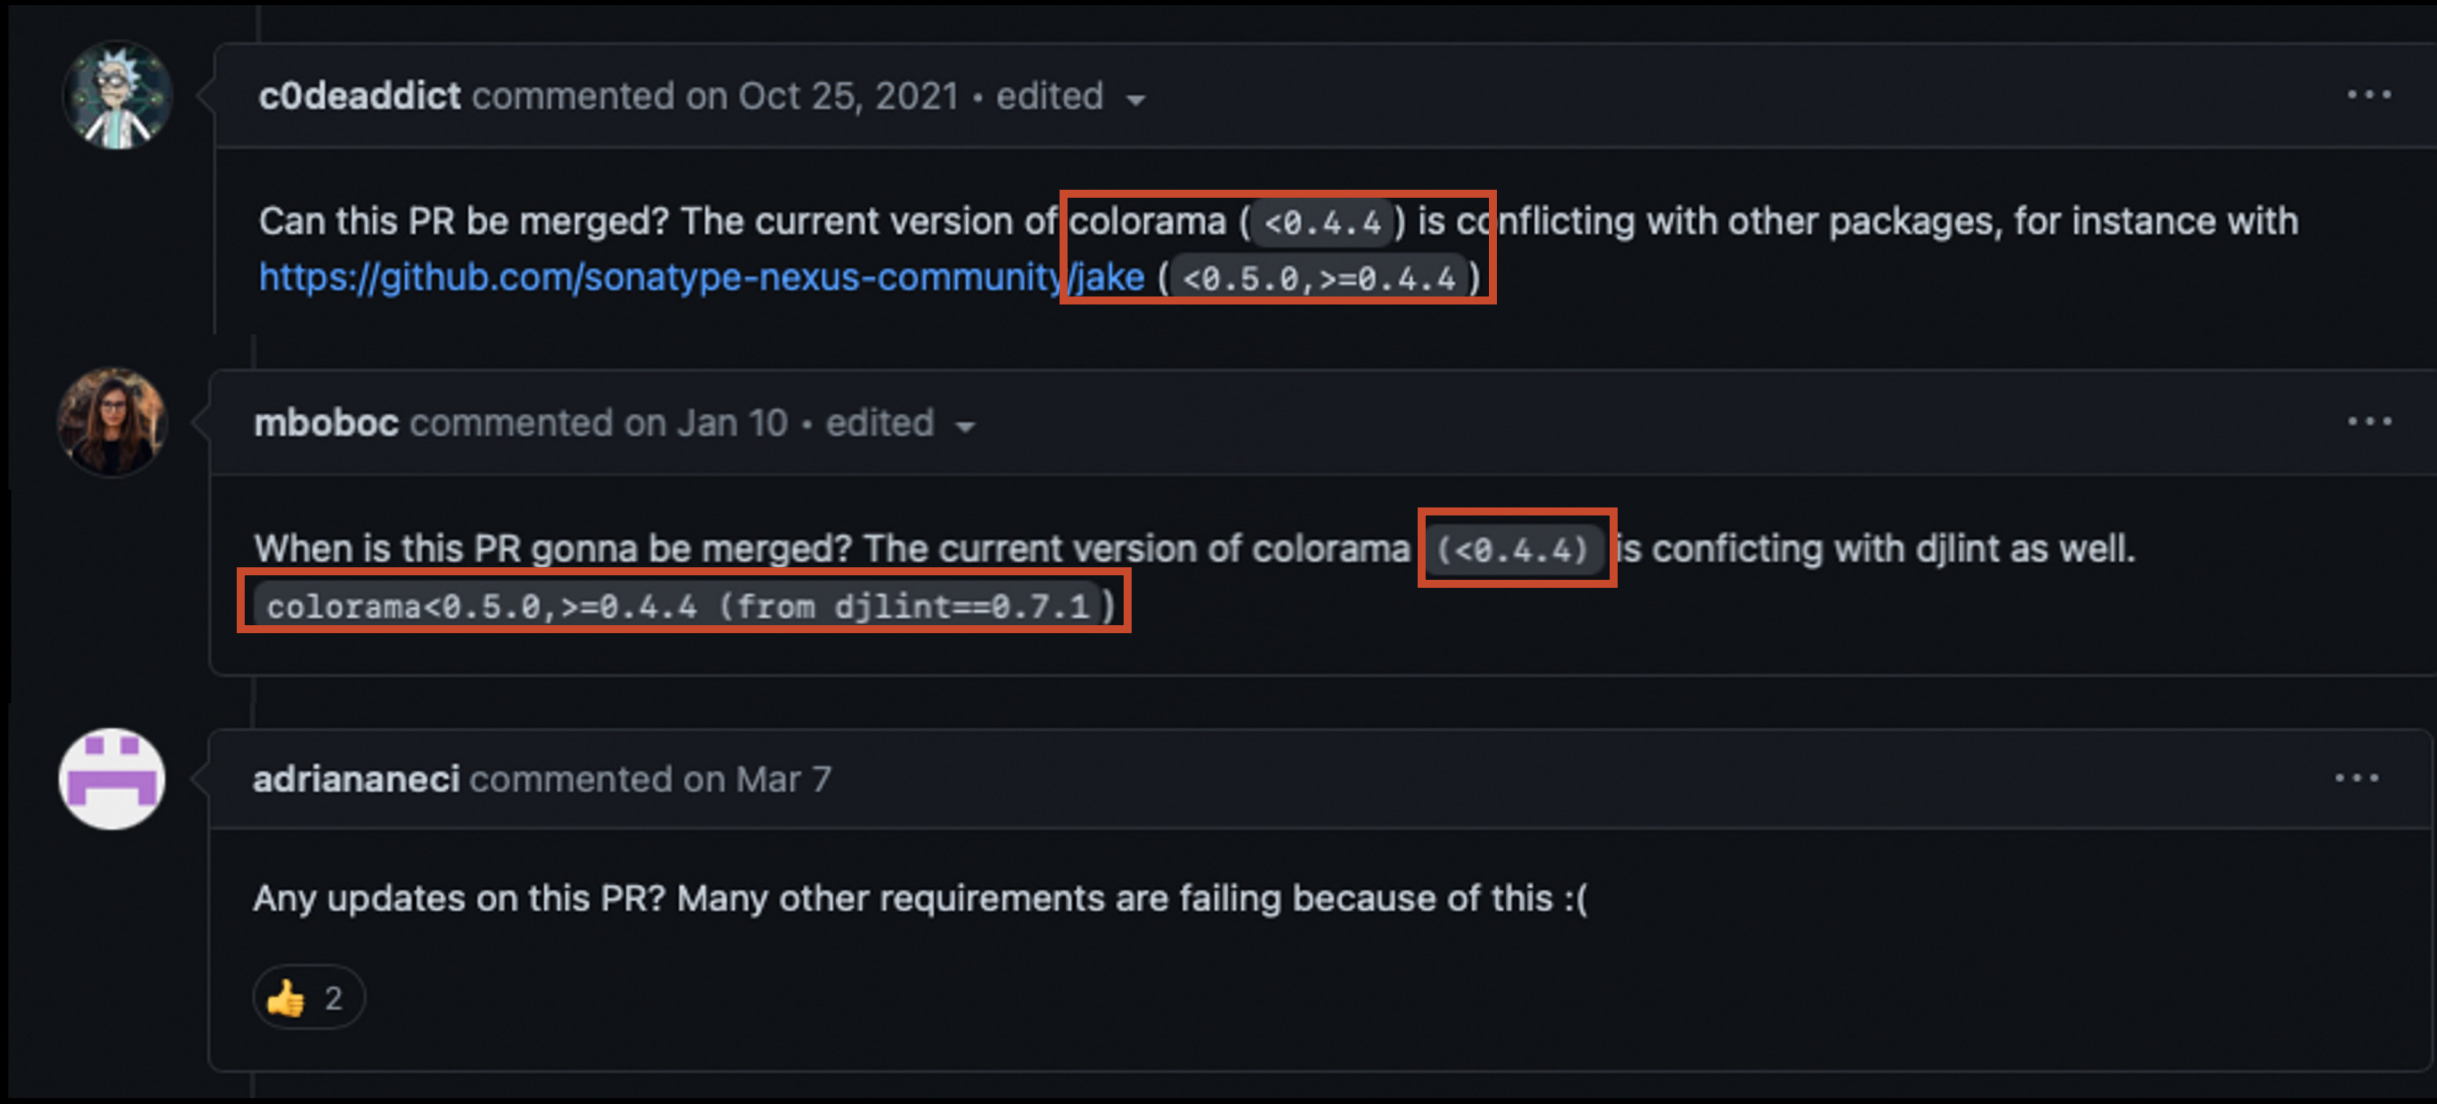

**c0deaddict** commented on Oct 25, 2021 • edited

Can this PR be merged? The current version of colorama ( `<0.4.4` ) is conflicting with other packages, for instance with https://github.com/sonatype-nexus-community/jake ( `<0.5.0,>=0.4.4` )

**mboboc** commented on Jan 10 • edited

When is this PR gonna be merged? The current version of colorama `(<0.4.4)` is conficting with djlint as well.
`colorama<0.5.0,>=0.4.4 (from djlint==0.7.1)`

**adriananeci** commented on Mar 7

Any updates on this PR? Many other requirements are failing because of this :(

👍 2

# Naïve Solutions Fall Short

- As dependency APIs change, blindly updating might fail
  - Noisy run-time crash

# Naïve Solutions Fall Short

- As dependency APIs change, blindly updating might fail
  - Noisy run-time crash

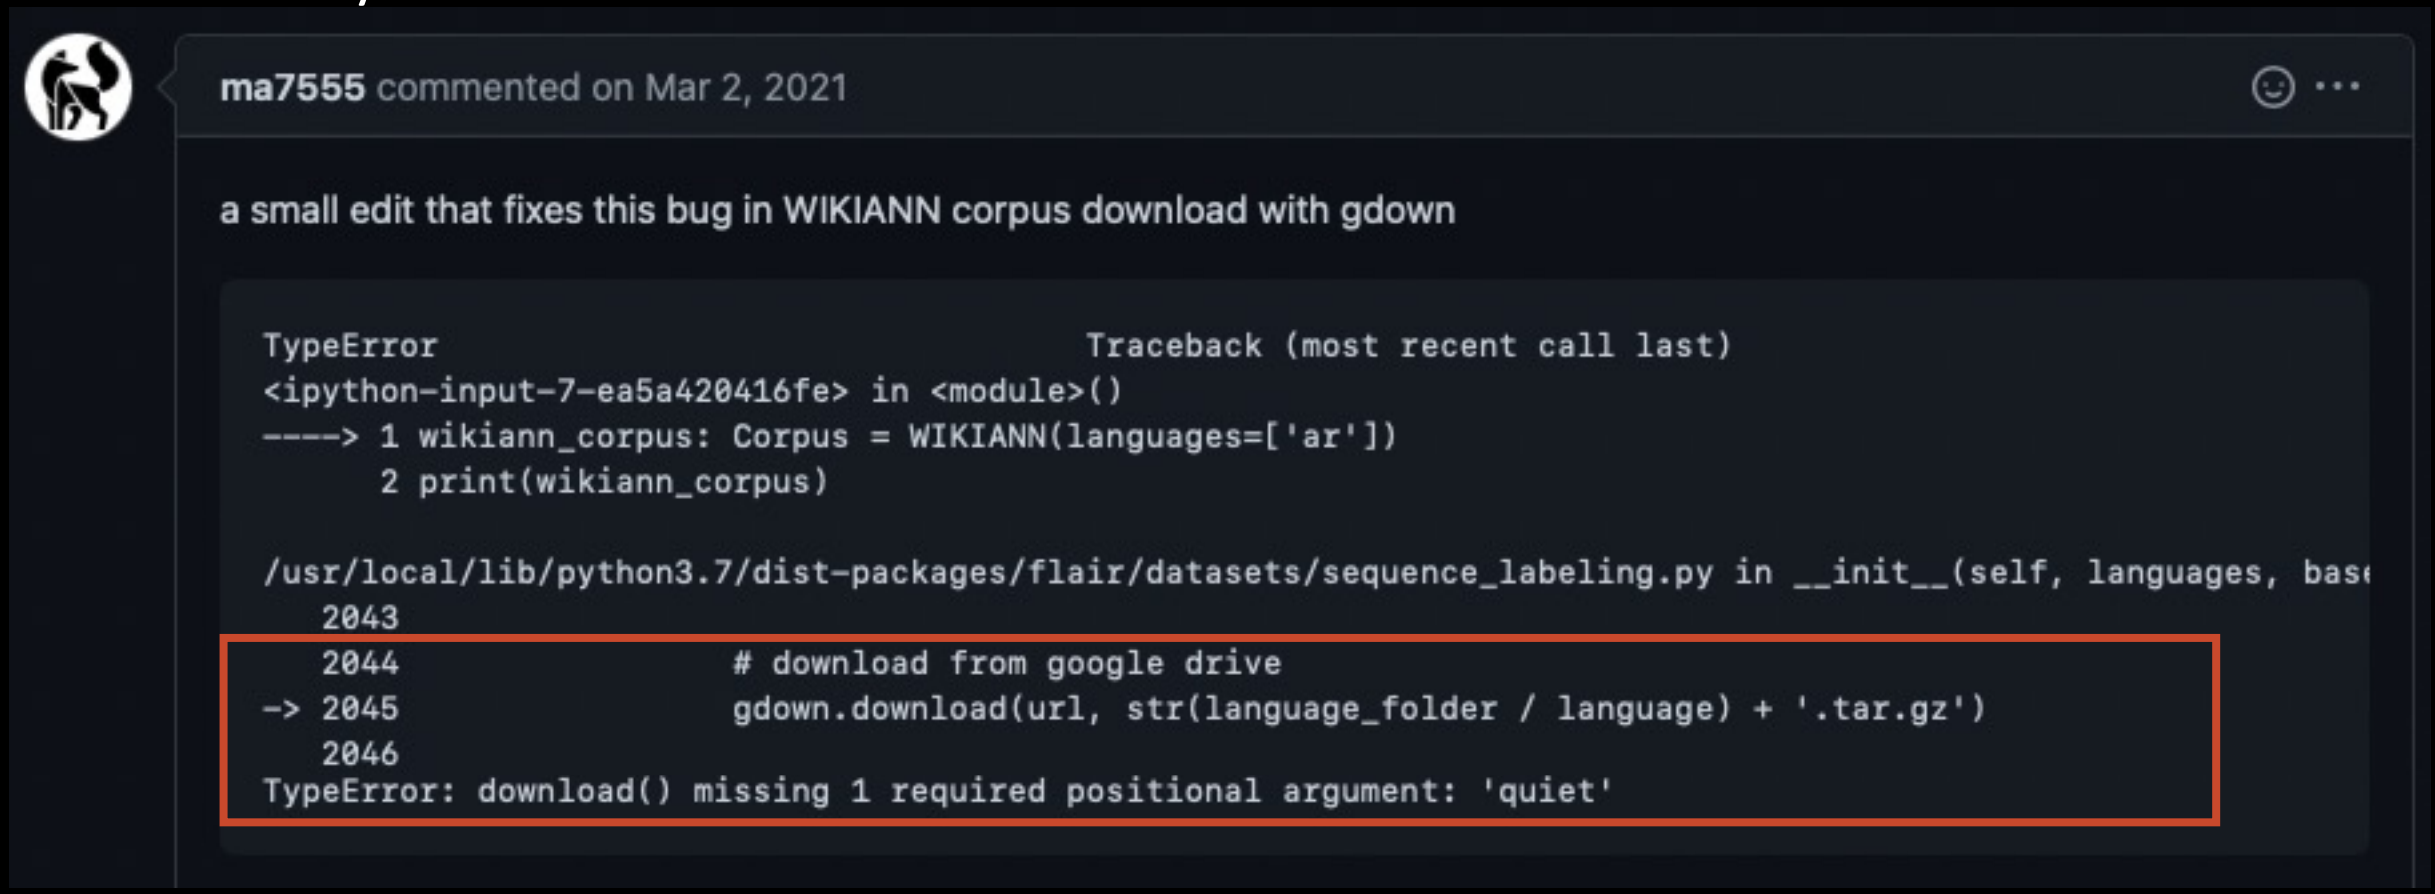

**ma7555** commented on Mar 2, 2021

a small edit that fixes this bug in WIKIANN corpus download with gdown

```
TypeError                                 Traceback (most recent call last)
<ipython-input-7-ea5a420416fe> in <module>()
----> 1 wikiann_corpus: Corpus = WIKIANN(languages=['ar'])
      2 print(wikiann_corpus)

/usr/local/lib/python3.7/dist-packages/flair/datasets/sequence_labeling.py in __init__(self, languages, base
   2043
   2044                 # download from google drive
-> 2045                 gdown.download(url, str(language_folder / language) + '.tar.gz')
   2046
TypeError: download() missing 1 required positional argument: 'quiet'
```

# Naïve Solutions Fall Short

- As dependency APIs change, blindly updating might fail
  - Noisy run-time crash
  - Or worse, fail silently

```
- async def spawn(coro, *args, loop=None, report_crash=True):
-     return spawn_sync(coro, *args, loop=loop, report_crash=report_crash)
+ async def spawn(coro, *args, loop=None, daemon=False):
+     return spawn_sync(coro, *args, loop=loop, daemon=daemon)
```

# Naïve Solutions Fall Short

- As dependency APIs change, blindly updating might fail
  - Noisy run-time crash
  - Or worse, fail silently
- Integration tests fail or don’t even try covering dependency interfaces[1]

[1]Joseph Hejderup & Georgios Gousios, “Can we trust tests to automate dependency updates? A case study of Java Projects”, Journal of Systems and Software

# Naïve Solutions Fall Short

- As dependency APIs change
  - Noisy run-time crash
  - Or worse, fail silently


- Integration tests fail or don't even try covering dependency interfaces[1]
- Manual inspection is not feasible

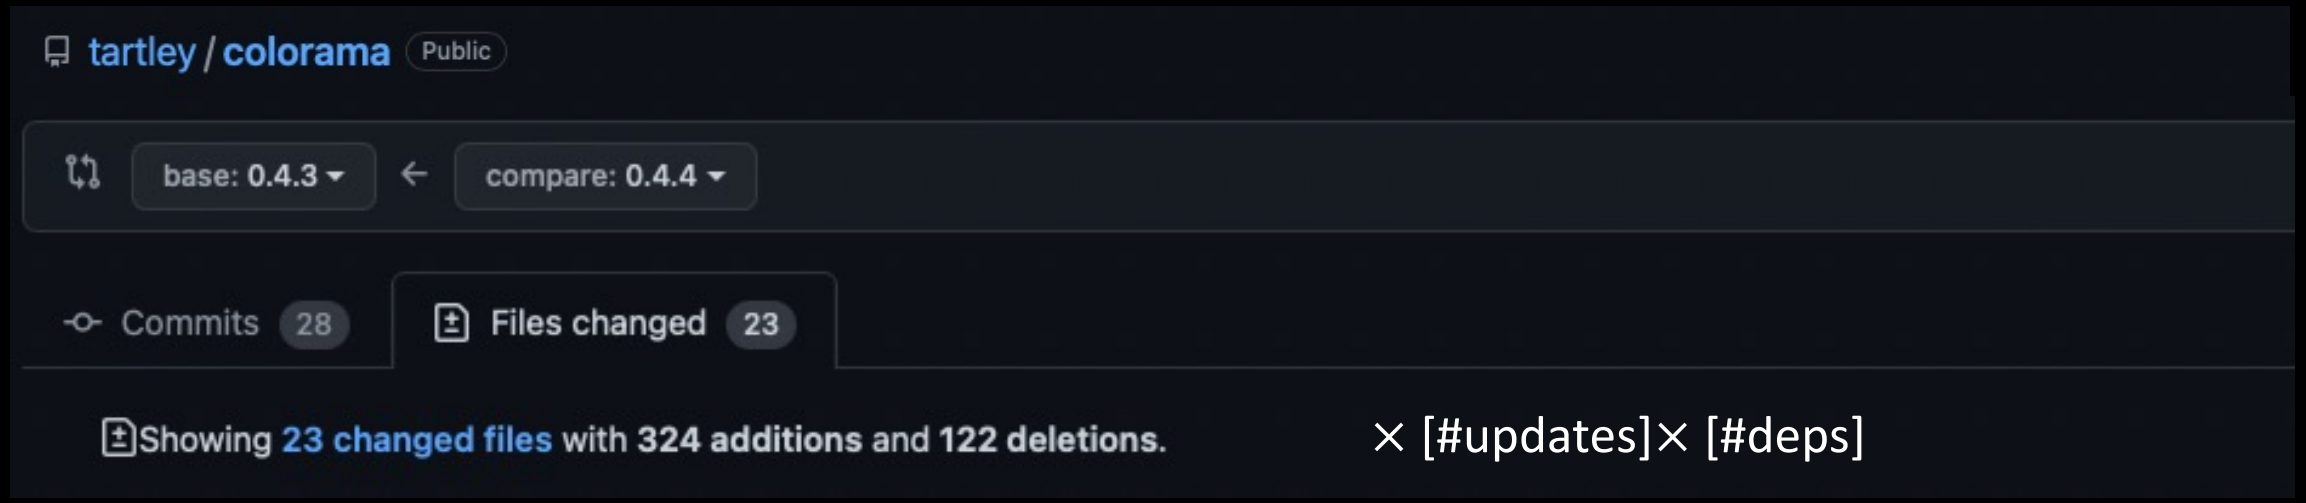


× [#updates]× [#deps]

# UPGRADVISOR: Upgrade-Advisor

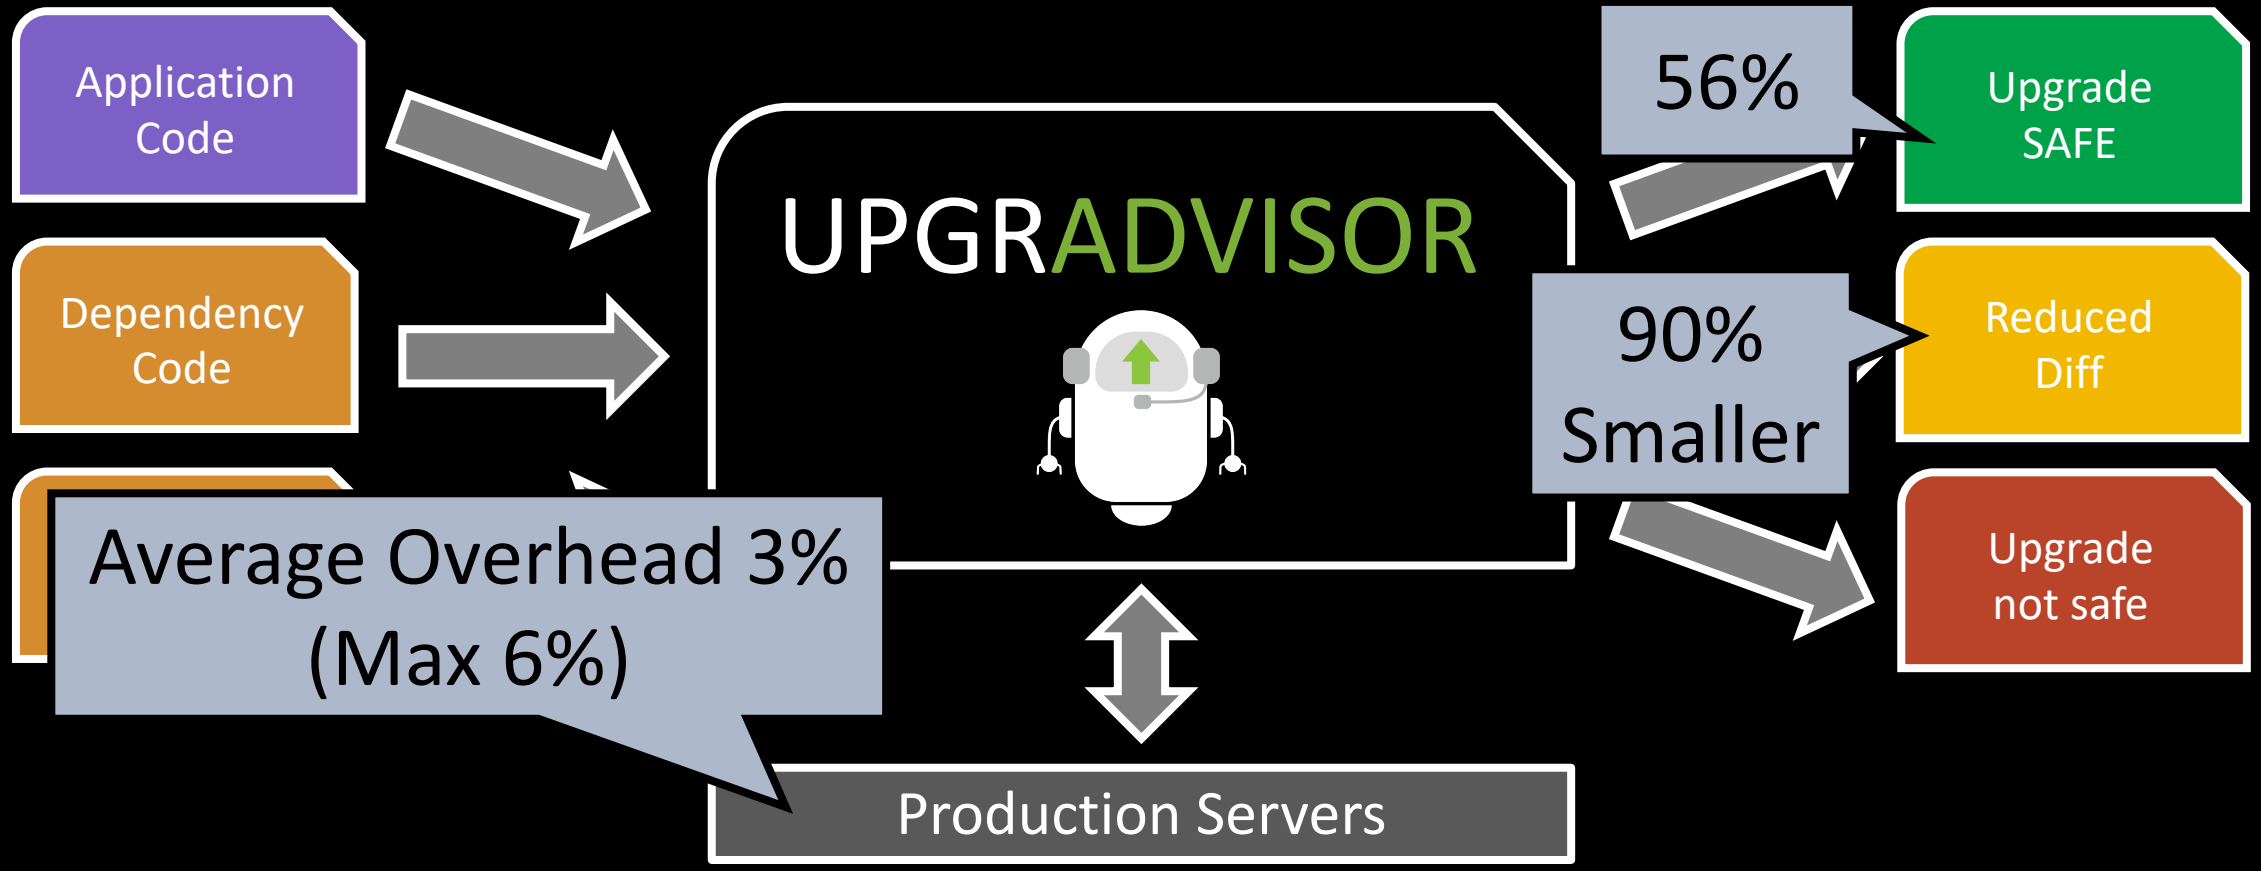


Built UPGRADVISOR-Python3 and evaluated on 172 dependency updates

# UPGRADVISOR: Upgrade-Advisor

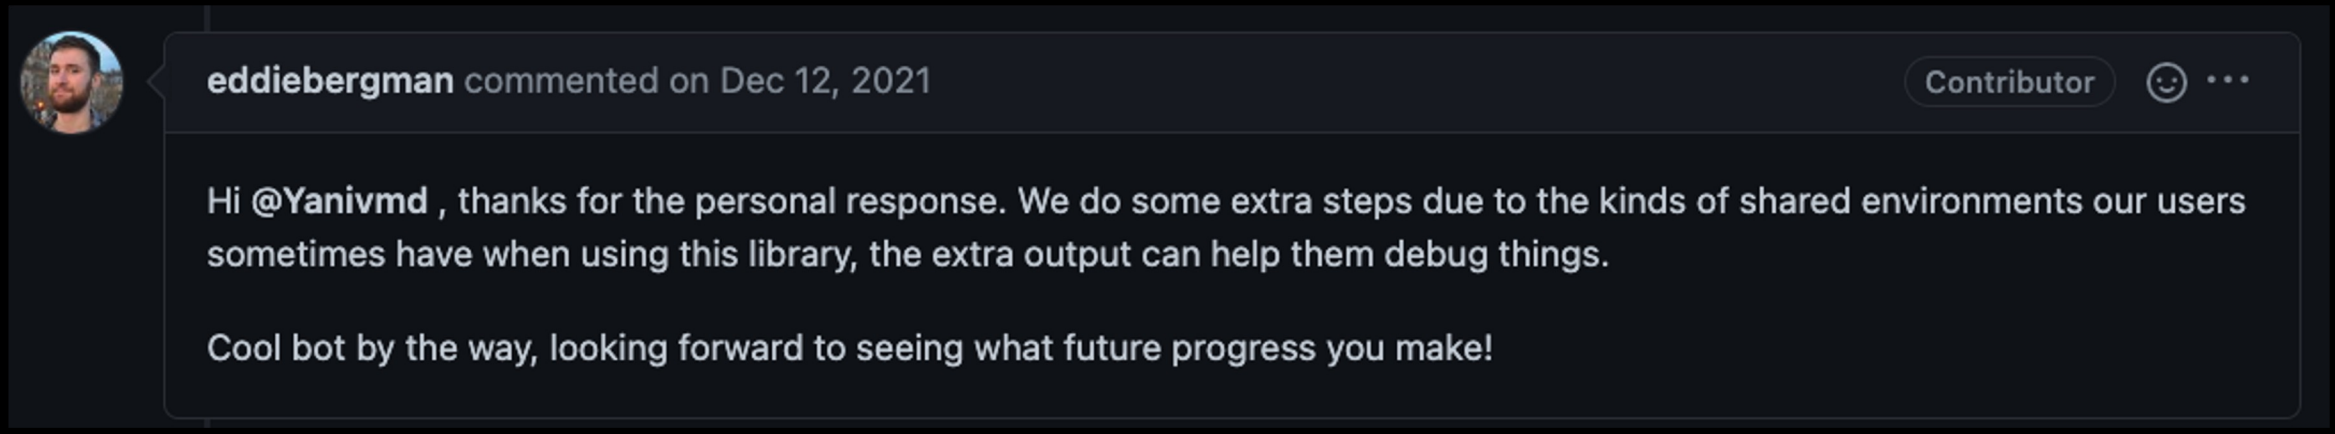


Manually submitted a sample of 9 PR
7 already merged

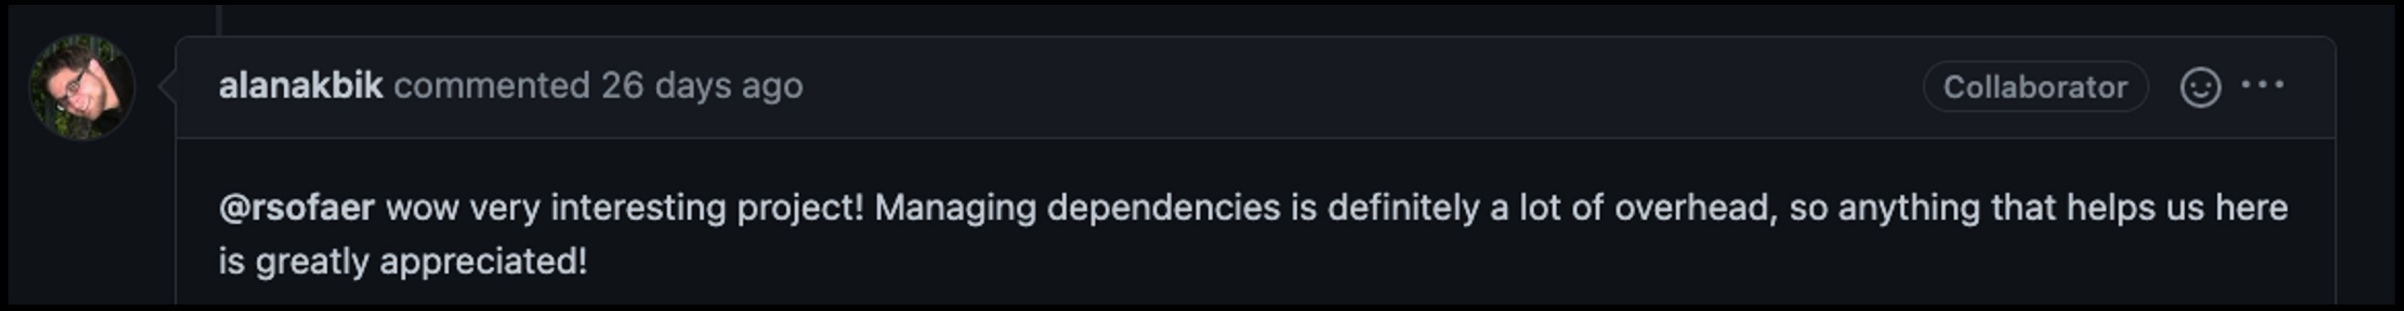

# Insight: What You Can’t Reach Won’t Hurt You

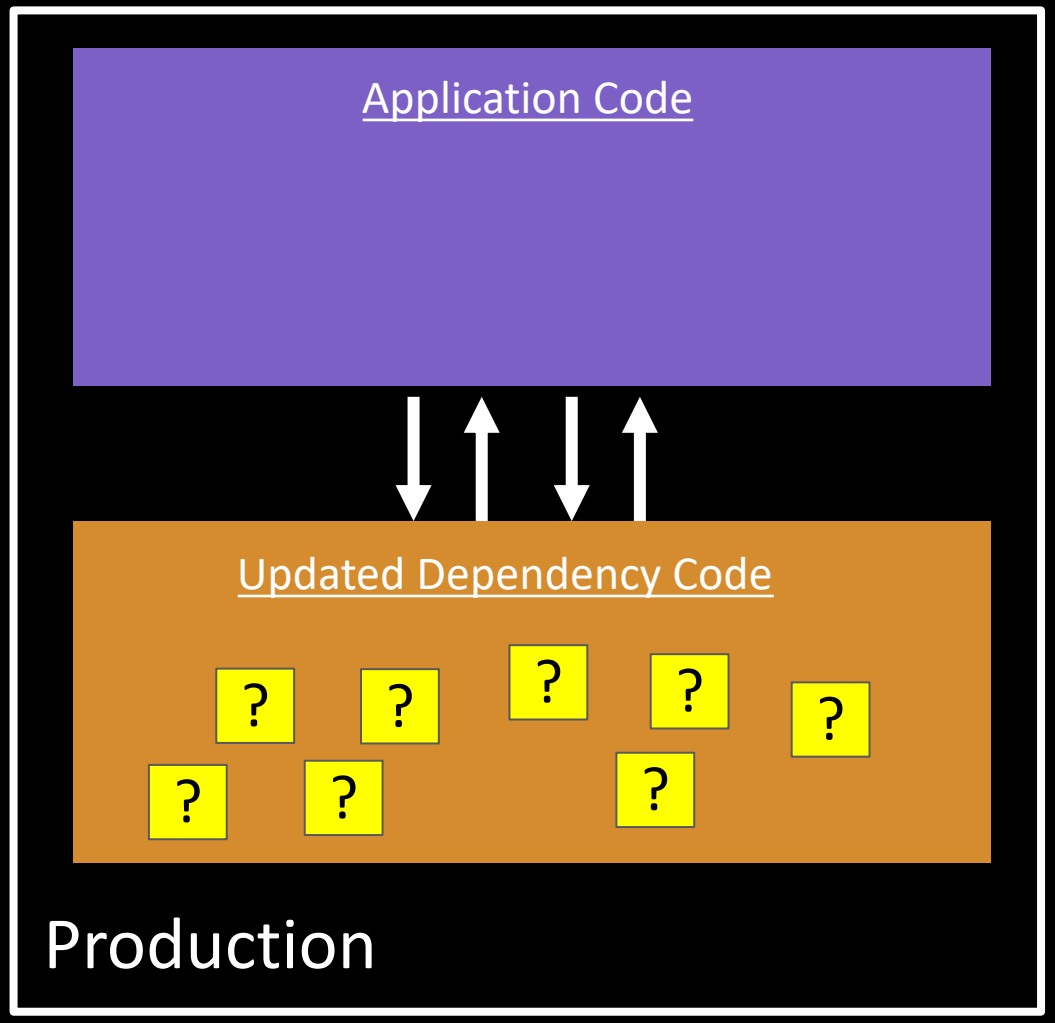

# Insight: What You Can’t Reach Won’t Hurt You

Requirements for production run:

- Study the update **without** applying it
- No interruption
- Incur low overhead

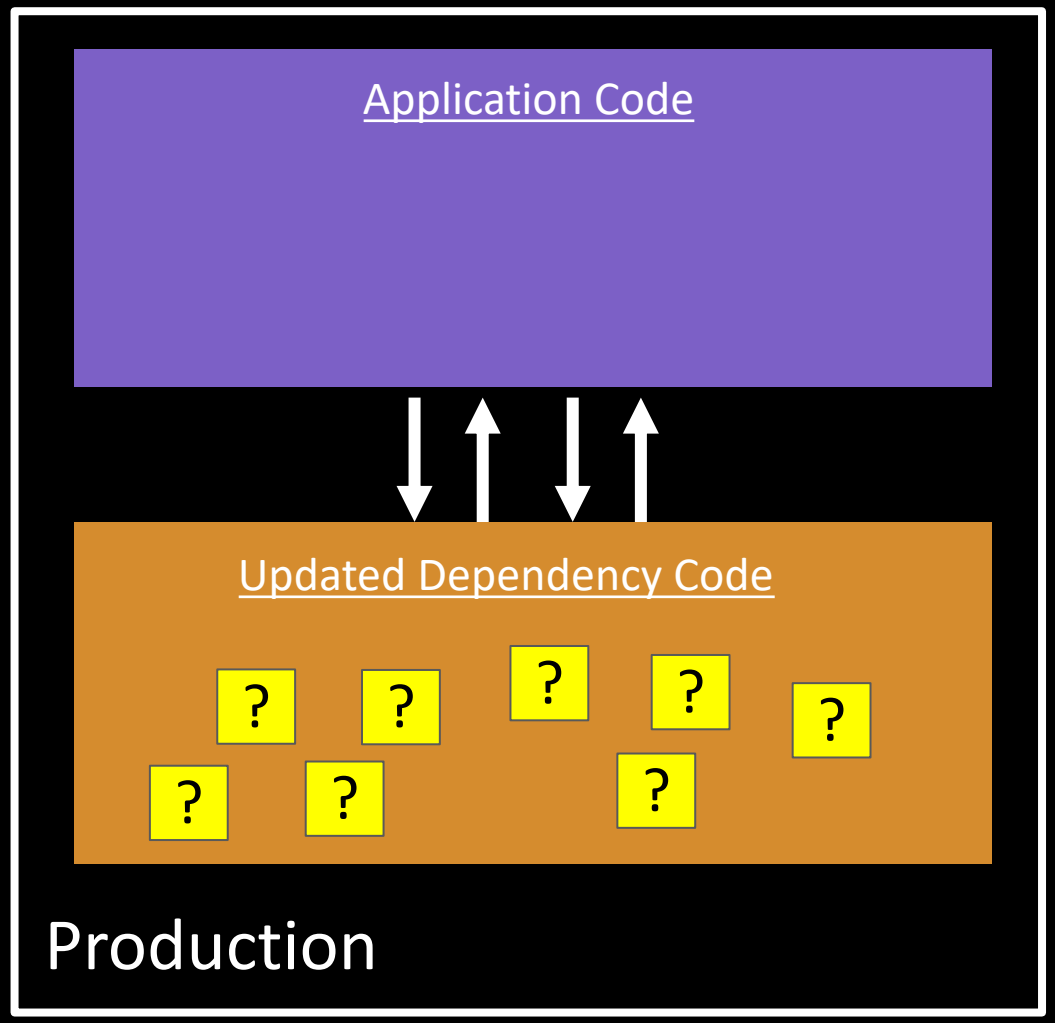

# Insight: What You Can't Reach Won't Hurt You

Requirements for production run:

- Study the update **without** applying it
- No interruption
- Incur low overhead

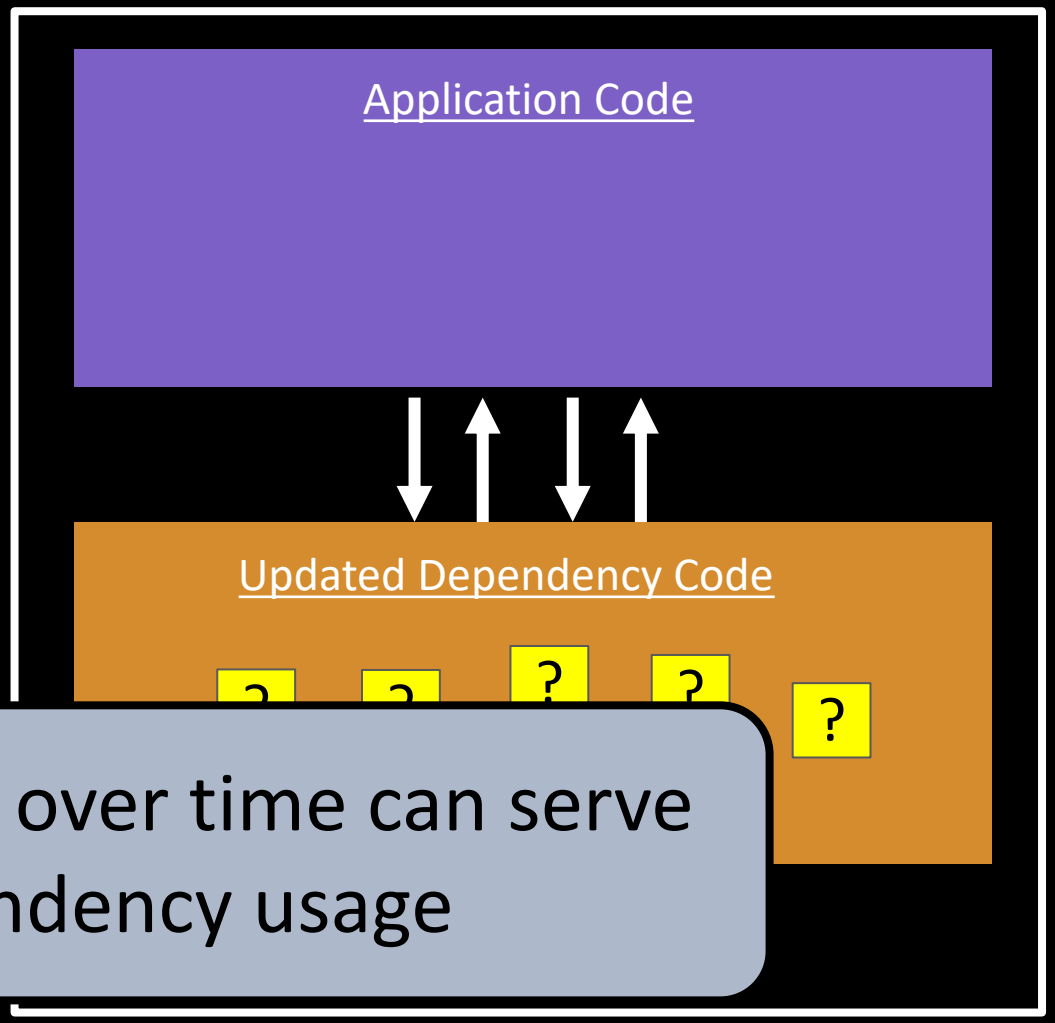


Tracing production environment over time can serve as **ground truth** for dependency usage

# Key Ideas Driving UPGRADVISOR's Design

- Safely discard non-reachable changes via hybrid program analysis
  - Static analysis to discard **never-reachable** changes
  - Dynamic analysis to test **maybe-reachable** changes
- Achieve low-overhead by employing hardware-based tracer

# Key Ideas Driving UPGRADVISOR's Design

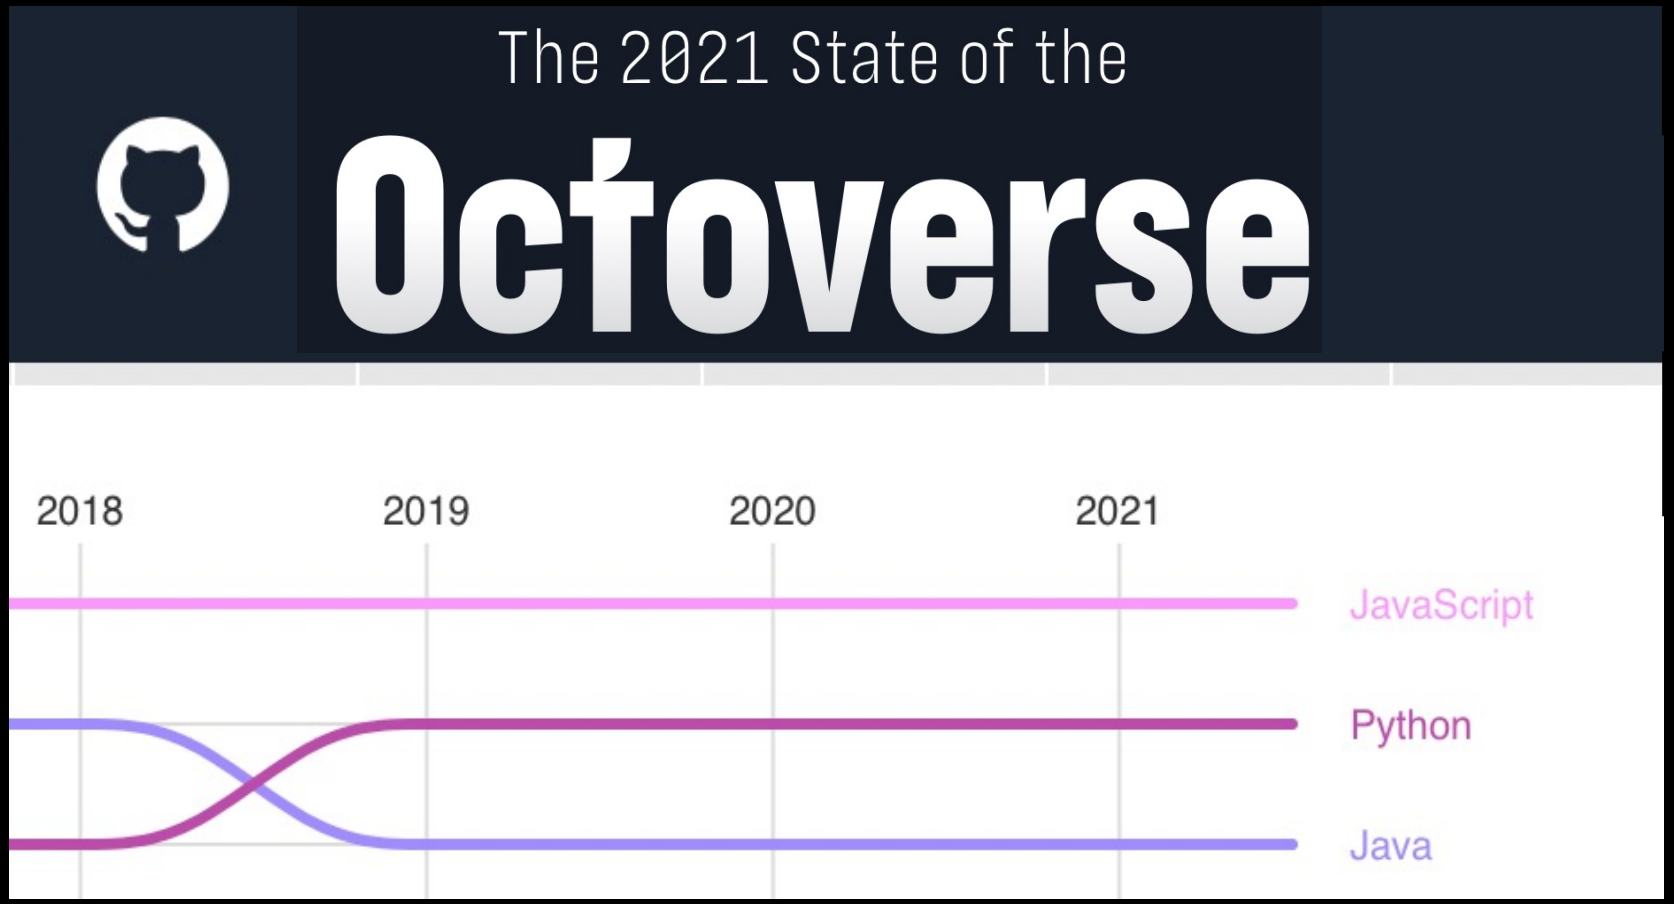


npm 2.42 M

python Package Index 455 K

RubyGems 179 K

# Key Ideas Driving UPGRADVISOR's Design

- Safely discard non-reachable changes via hybrid program analysis
  - Static analysis to discard **never-reachable** changes
  - Dynamic analysis to test **maybe-reachable** changes
- Achieve low-overhead by employing hardware-based tracer

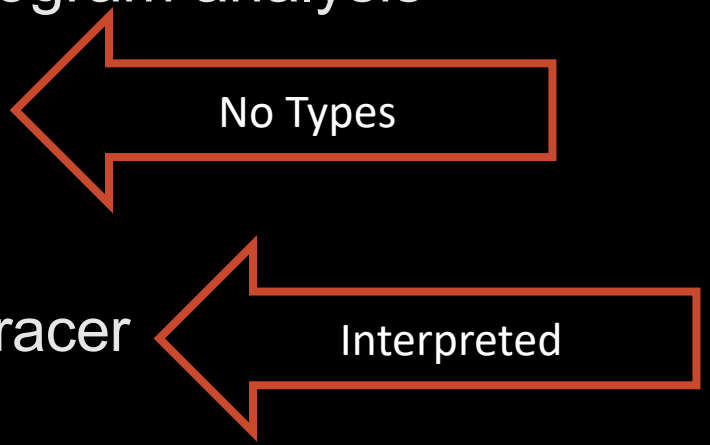


- Design for dynamic languages to maximize usability

# Analyzing Our Motivating Example

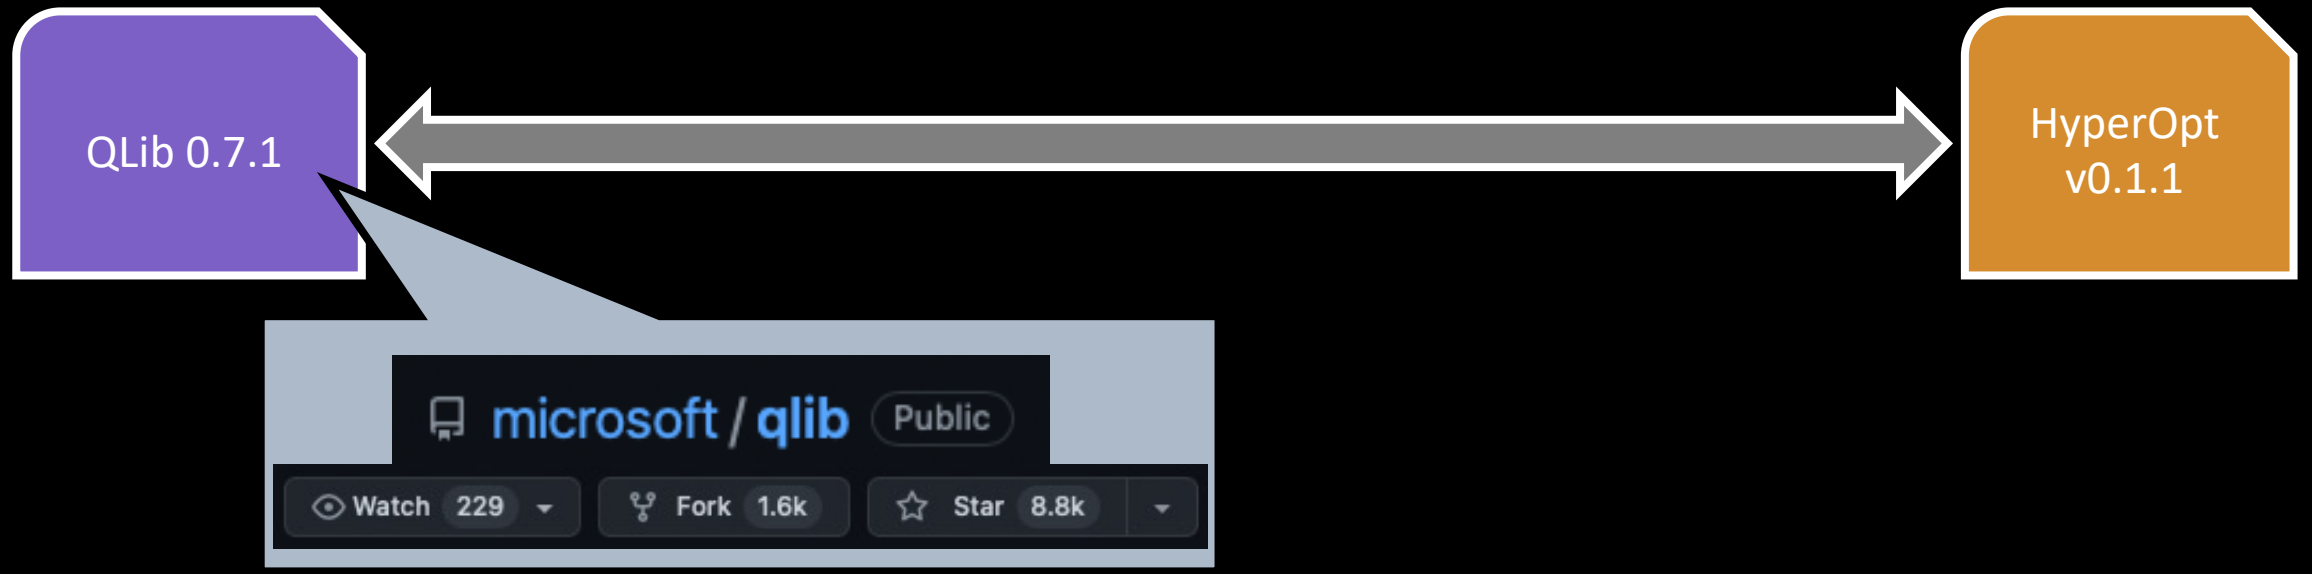

# Analyzing Our Motivating Example

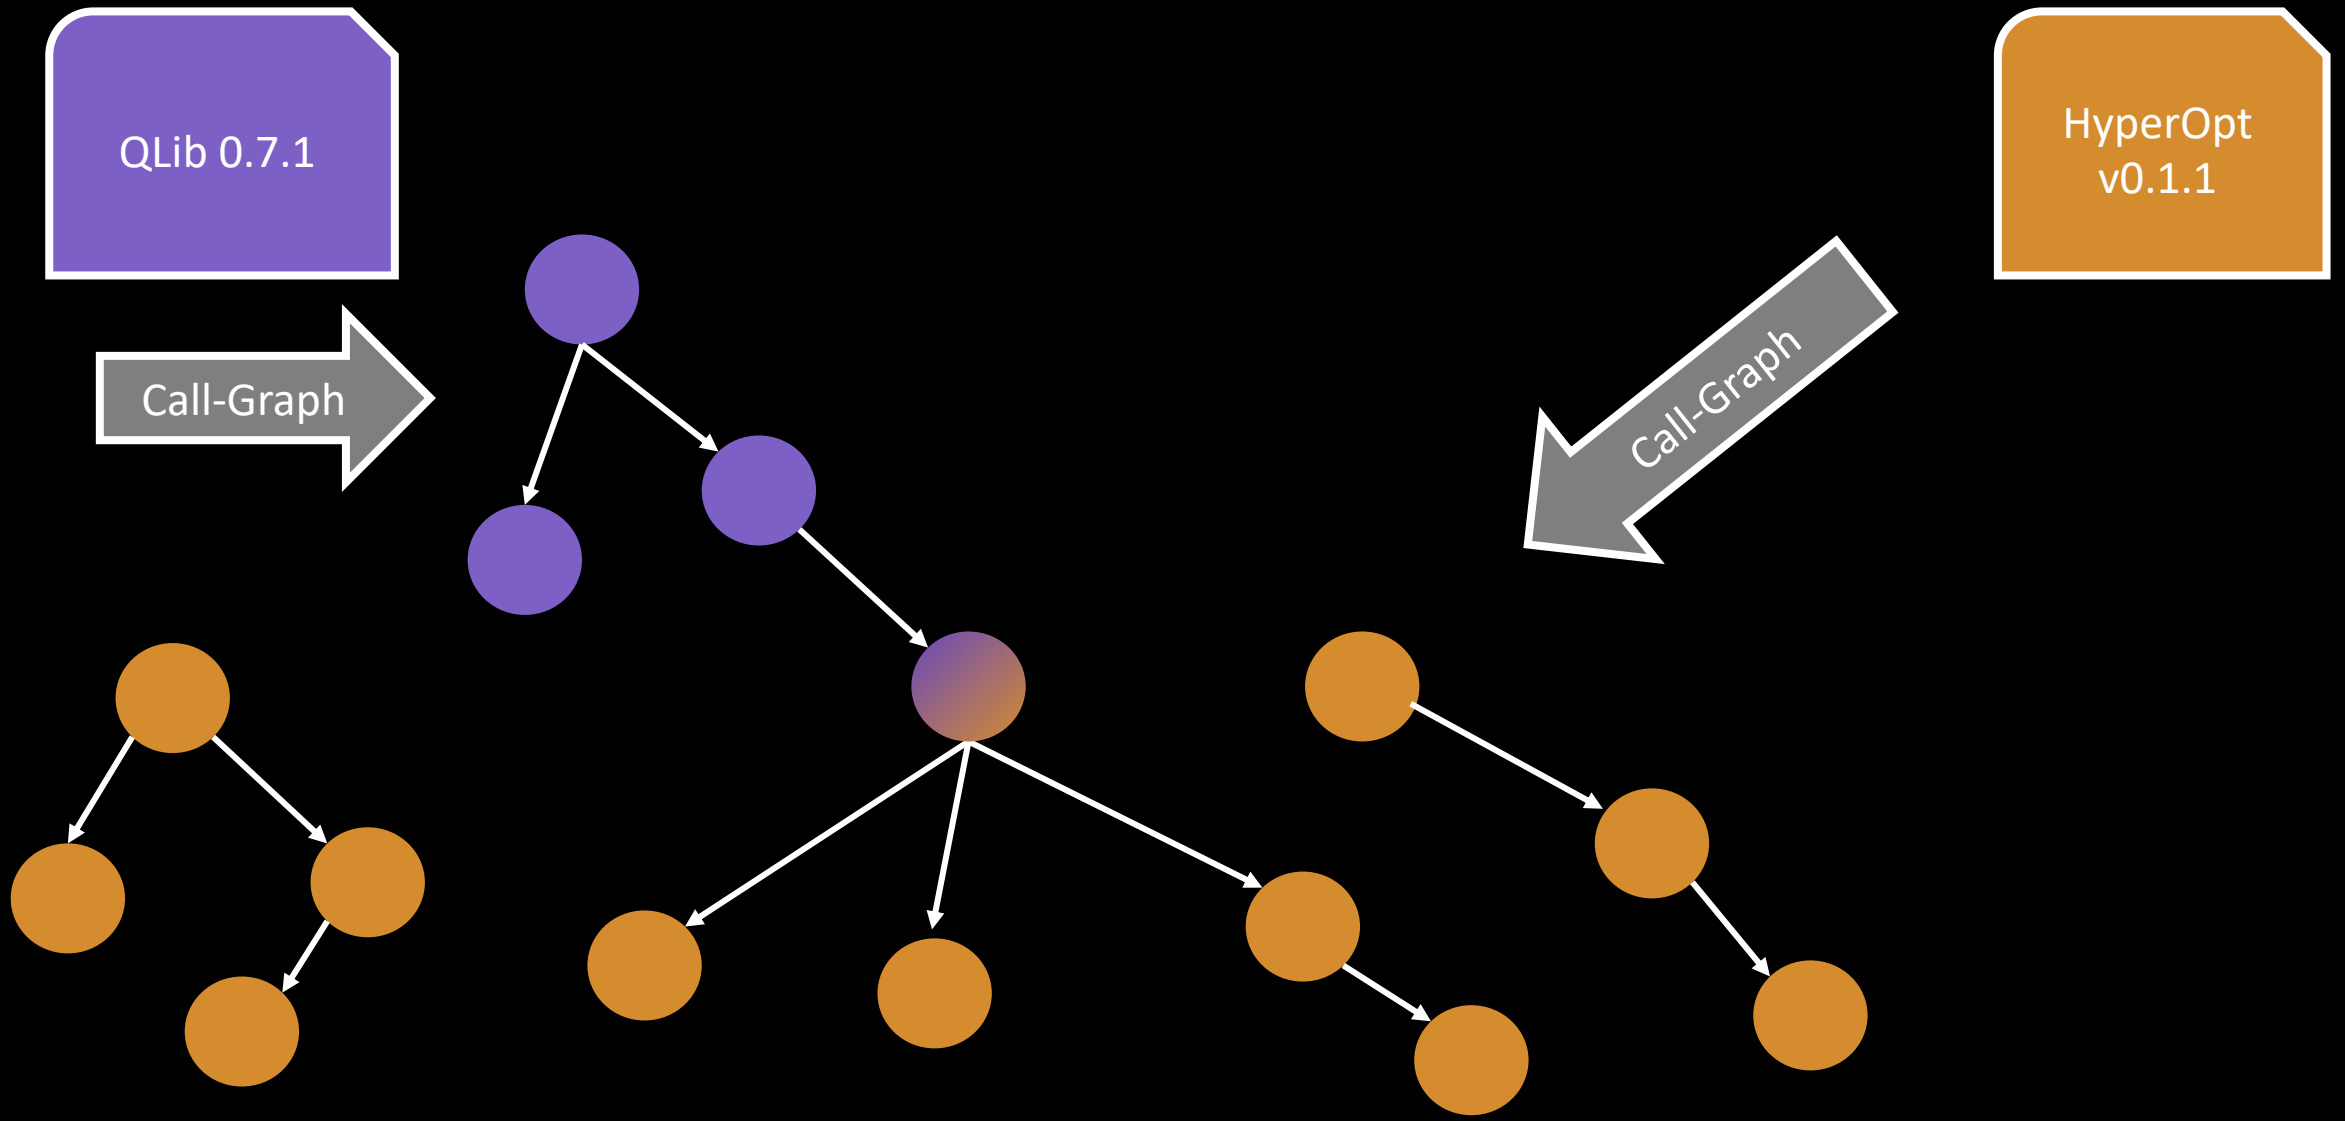

# Analyzing Our Motivating Example

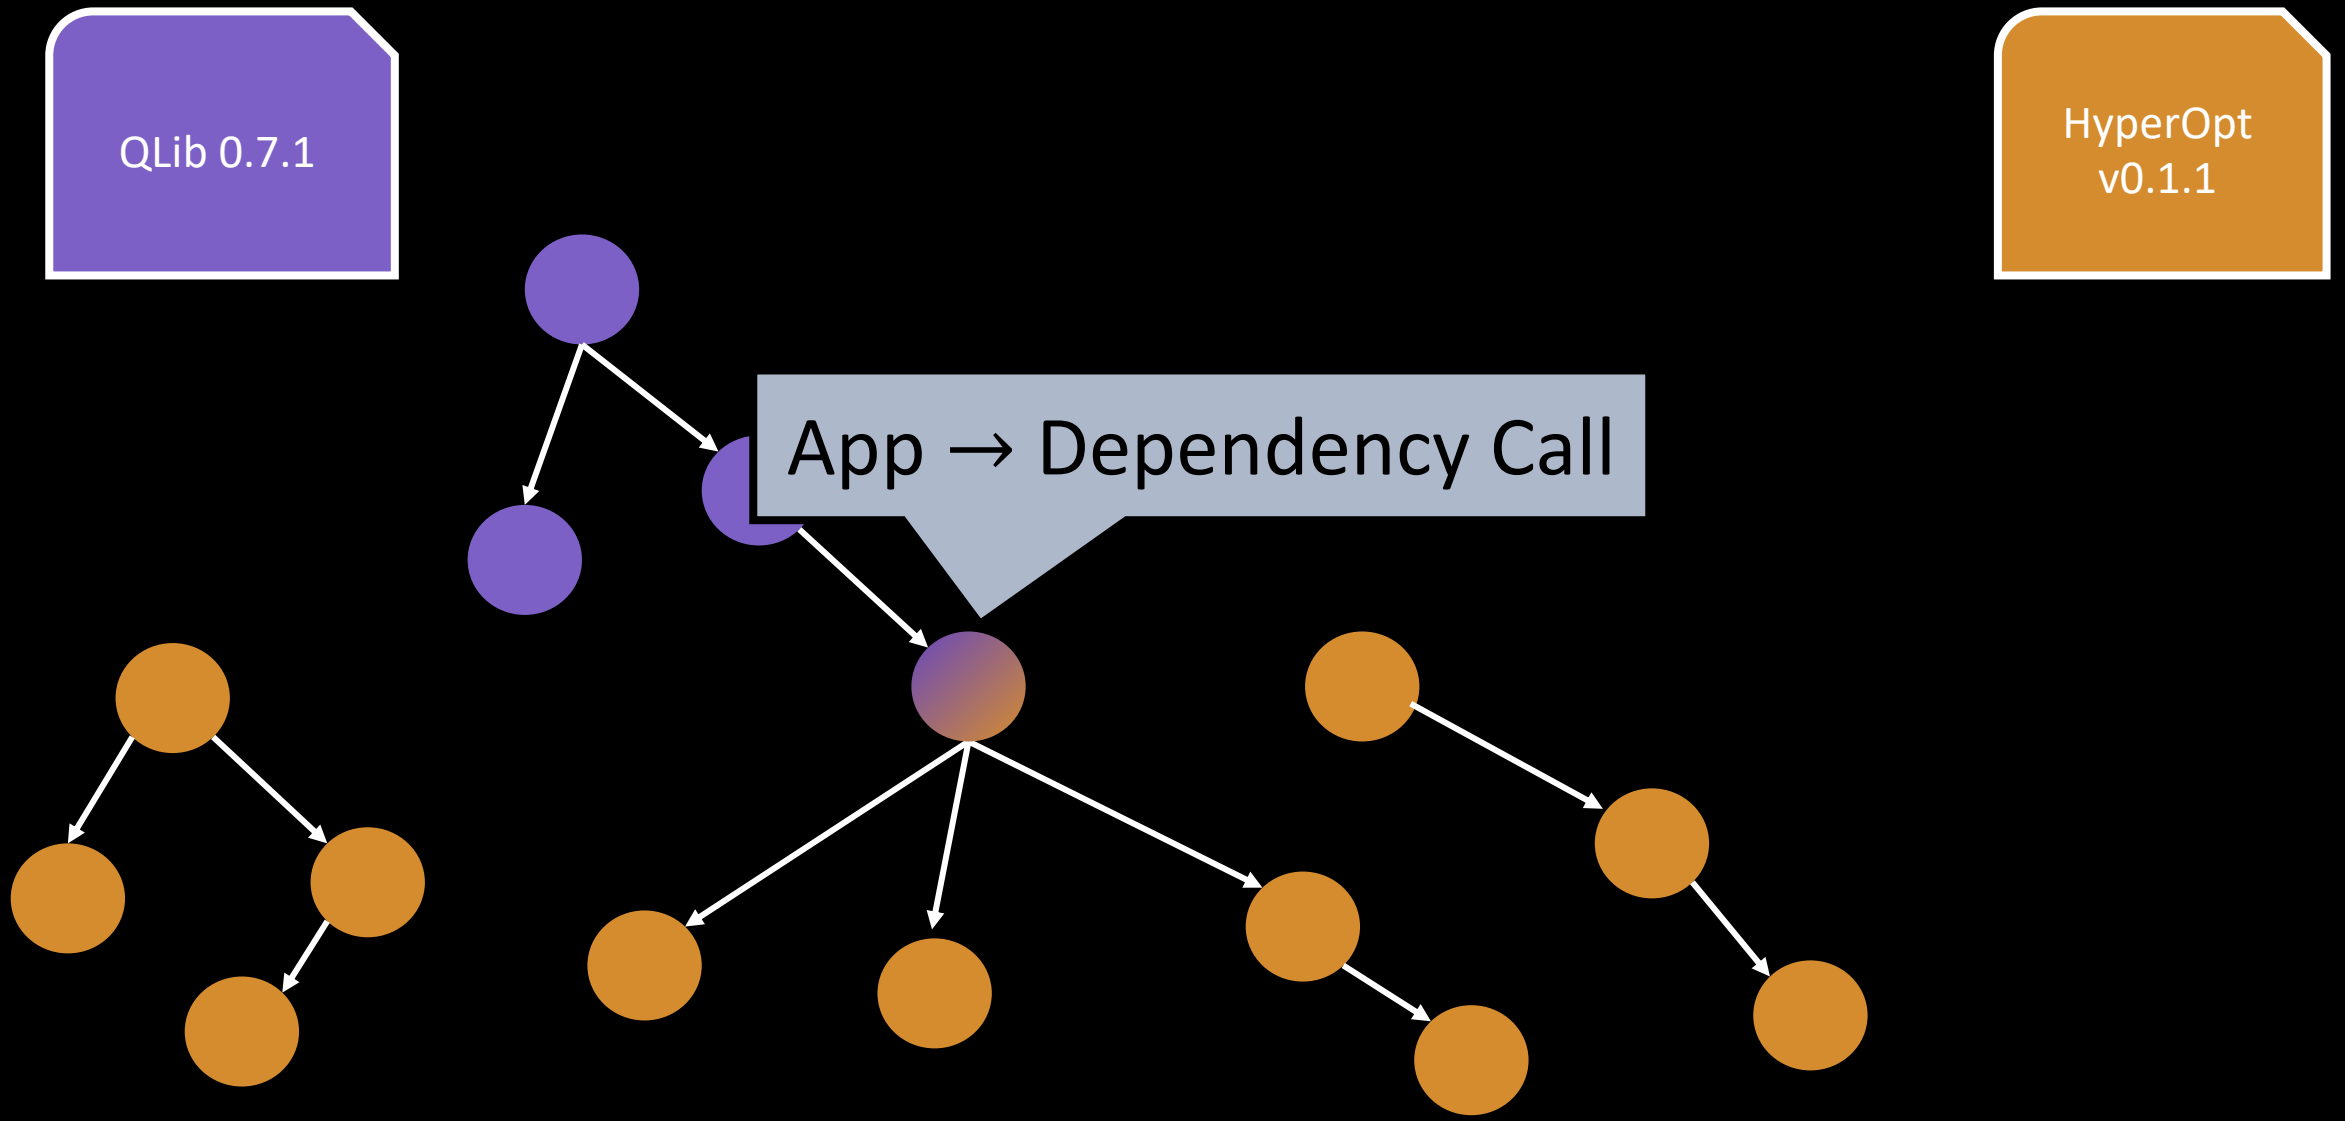

# Analyzing Our Motivating Example

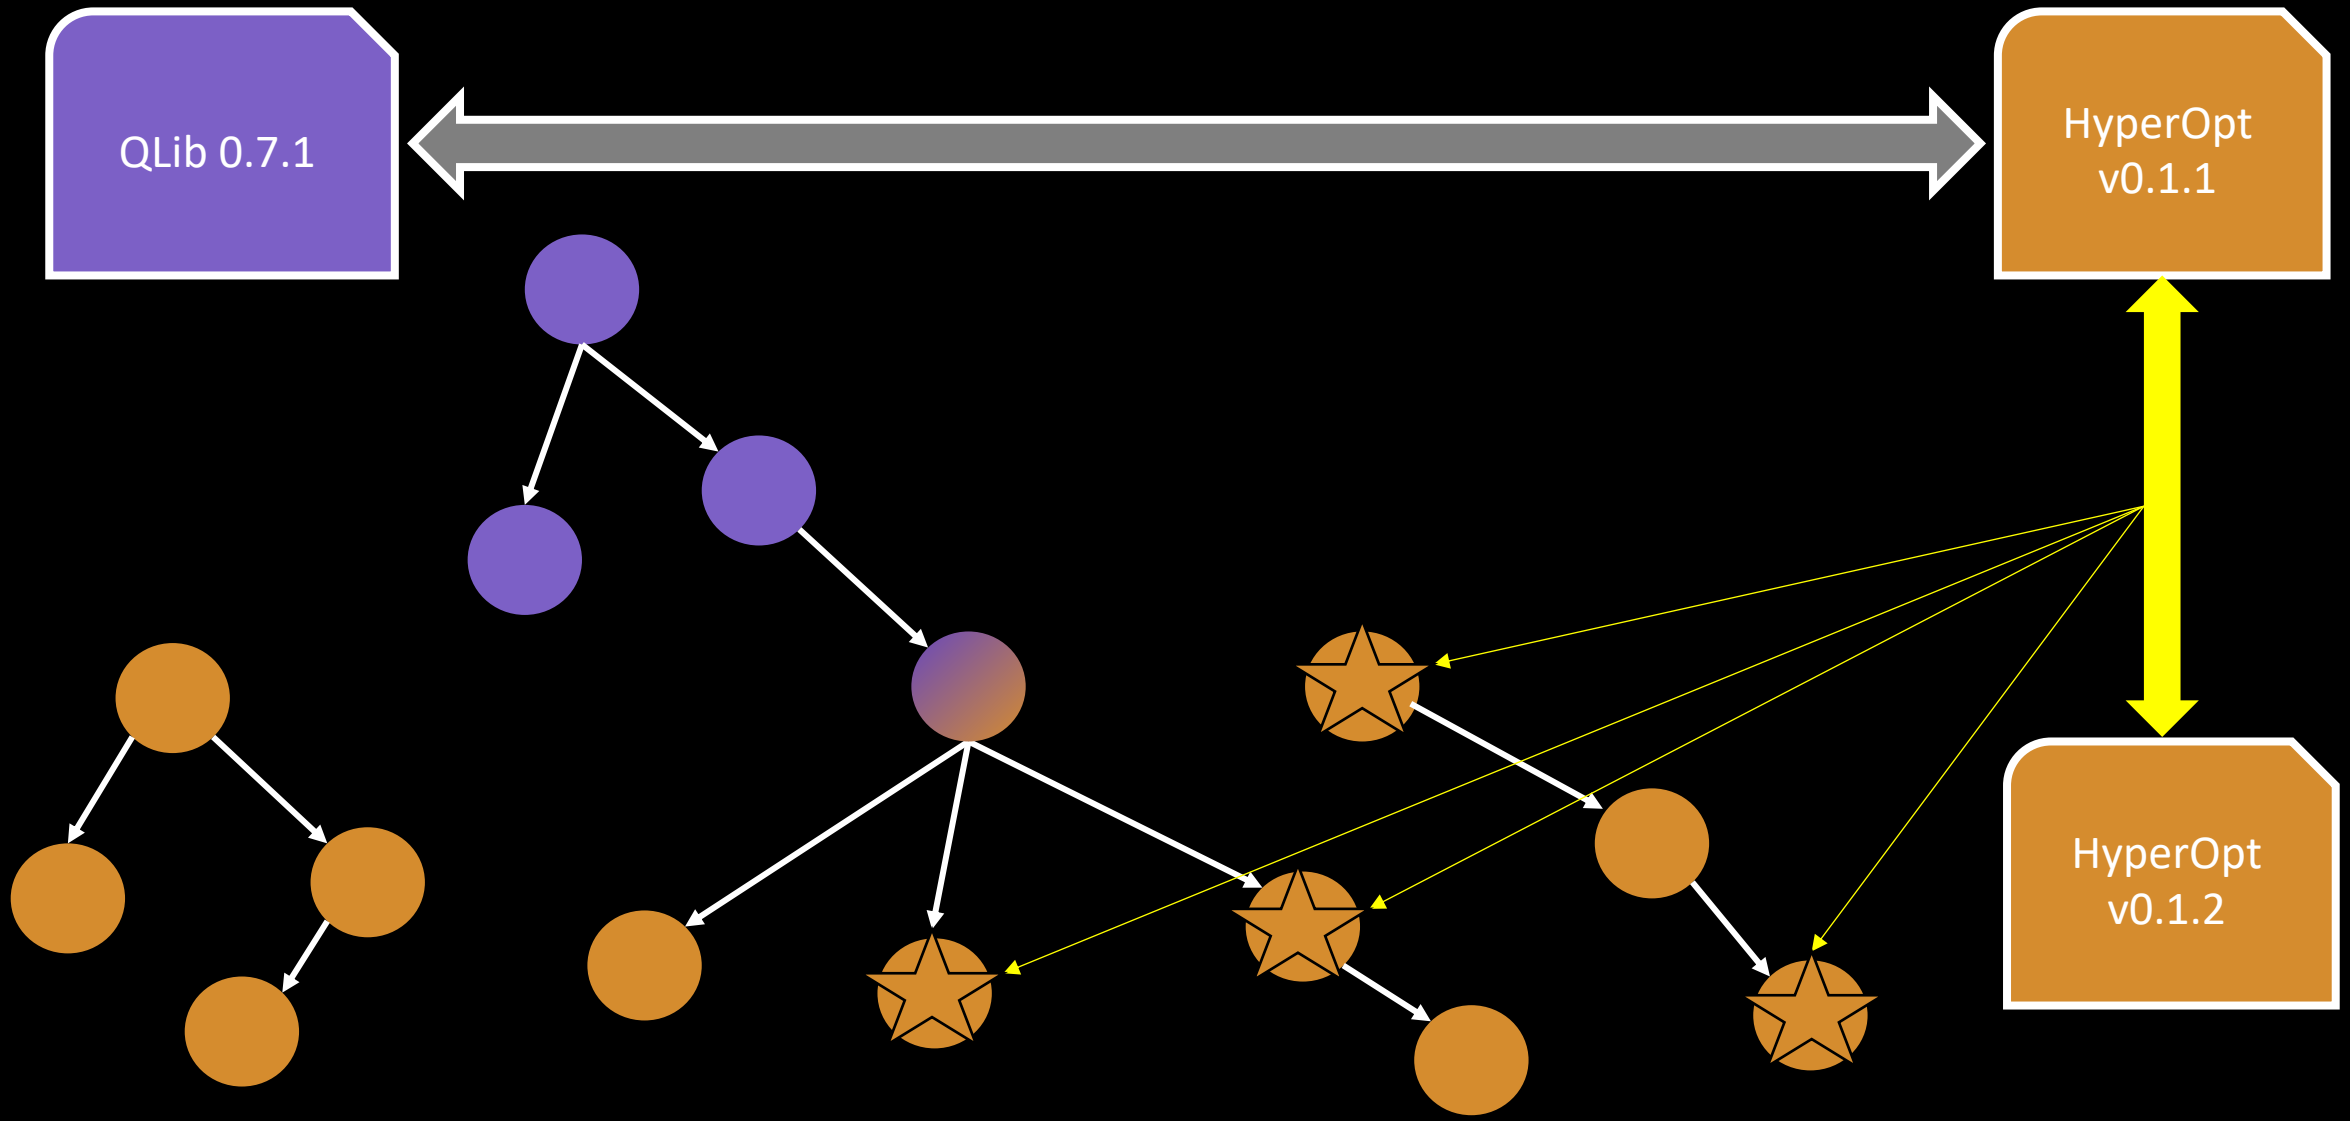

# Classify Changes In Motivating Example

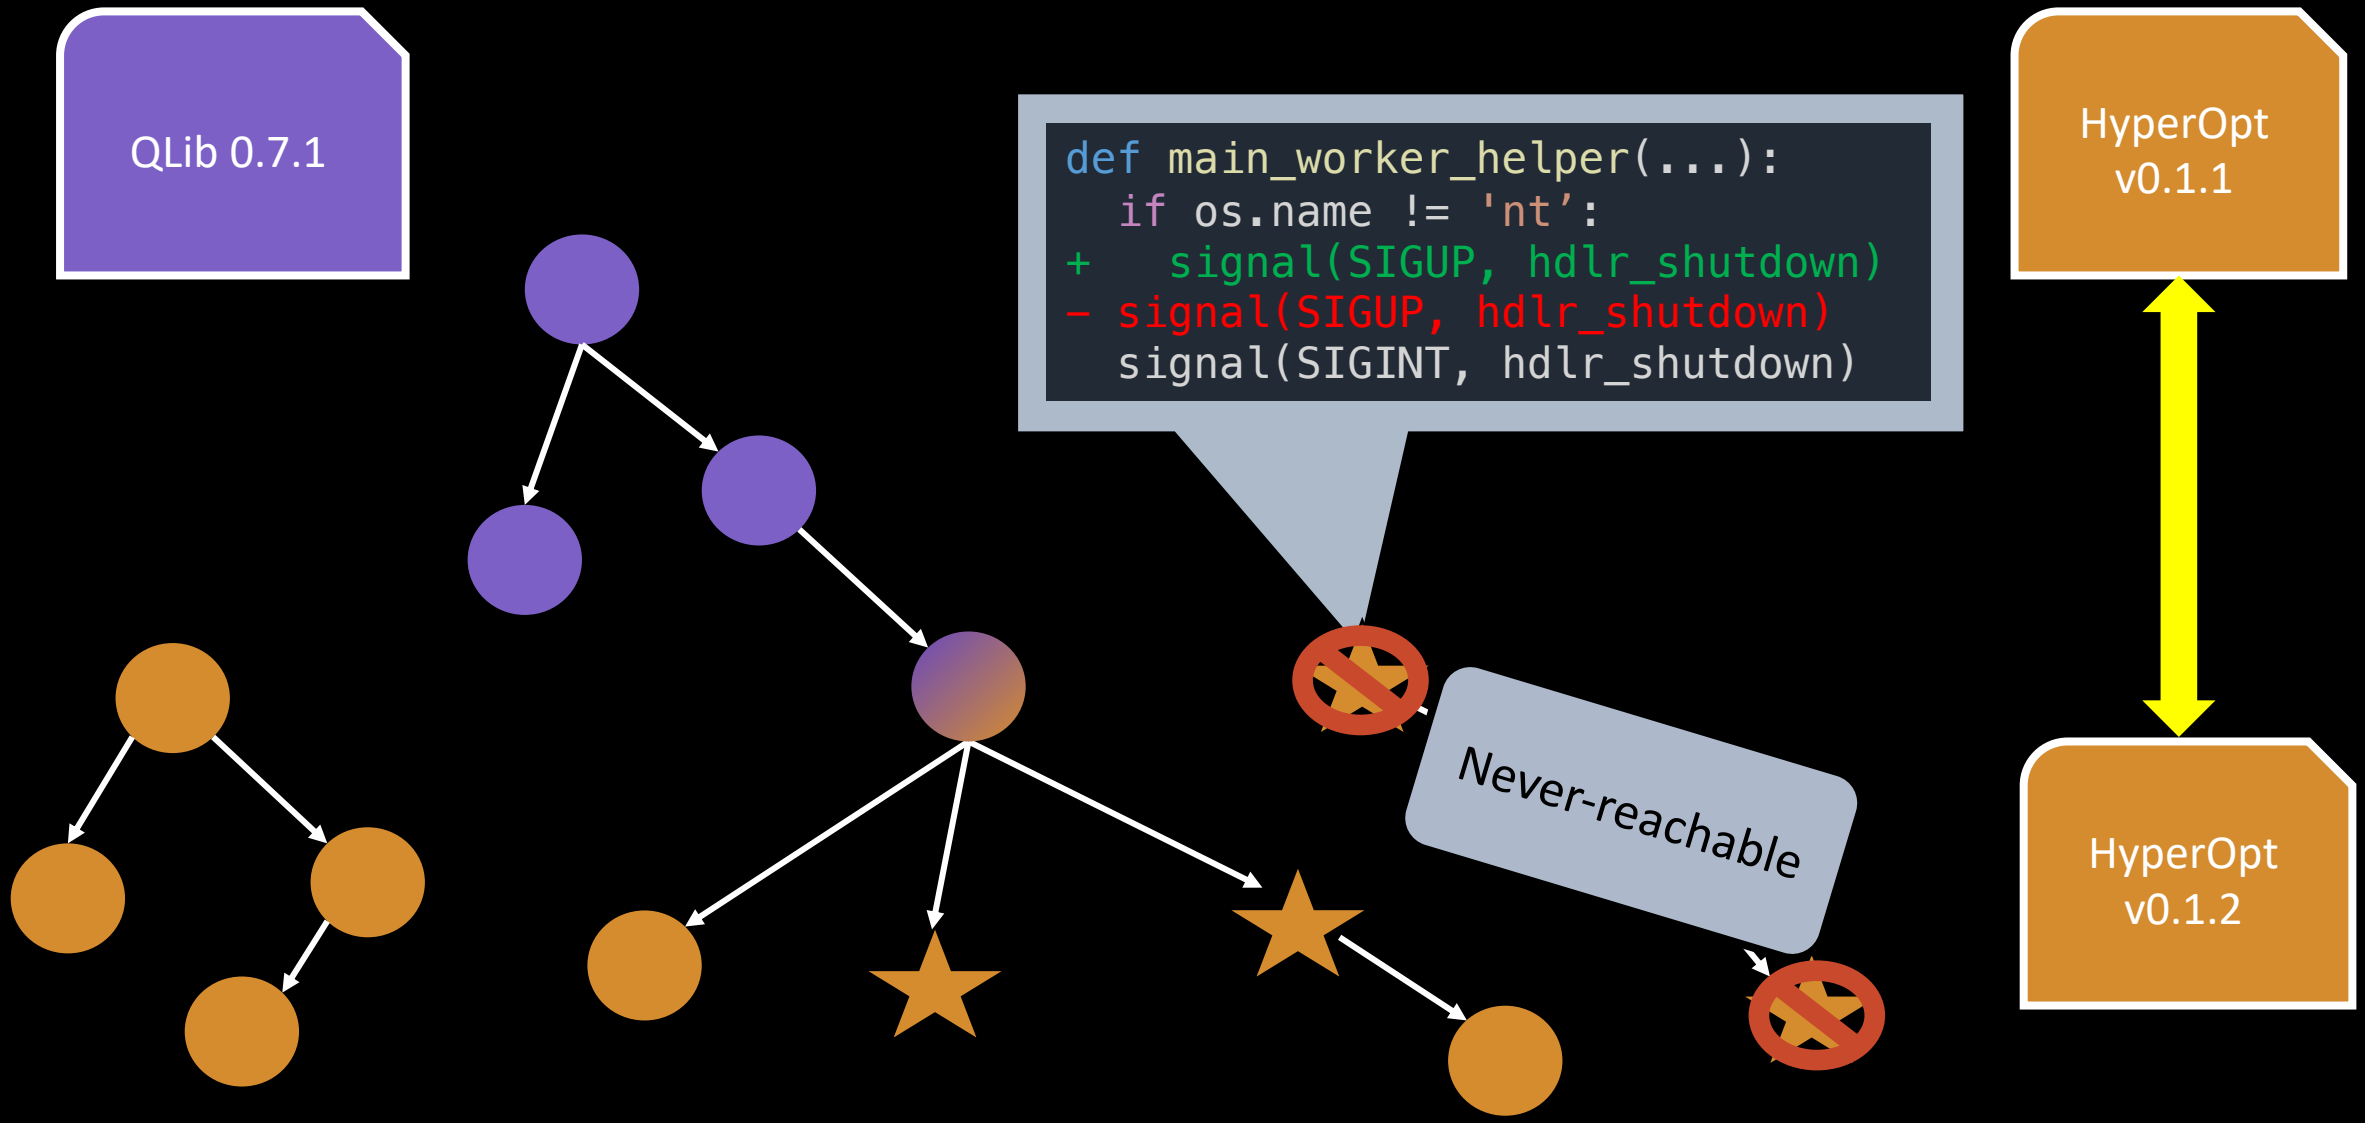

# Classify Changes In Motivating Example

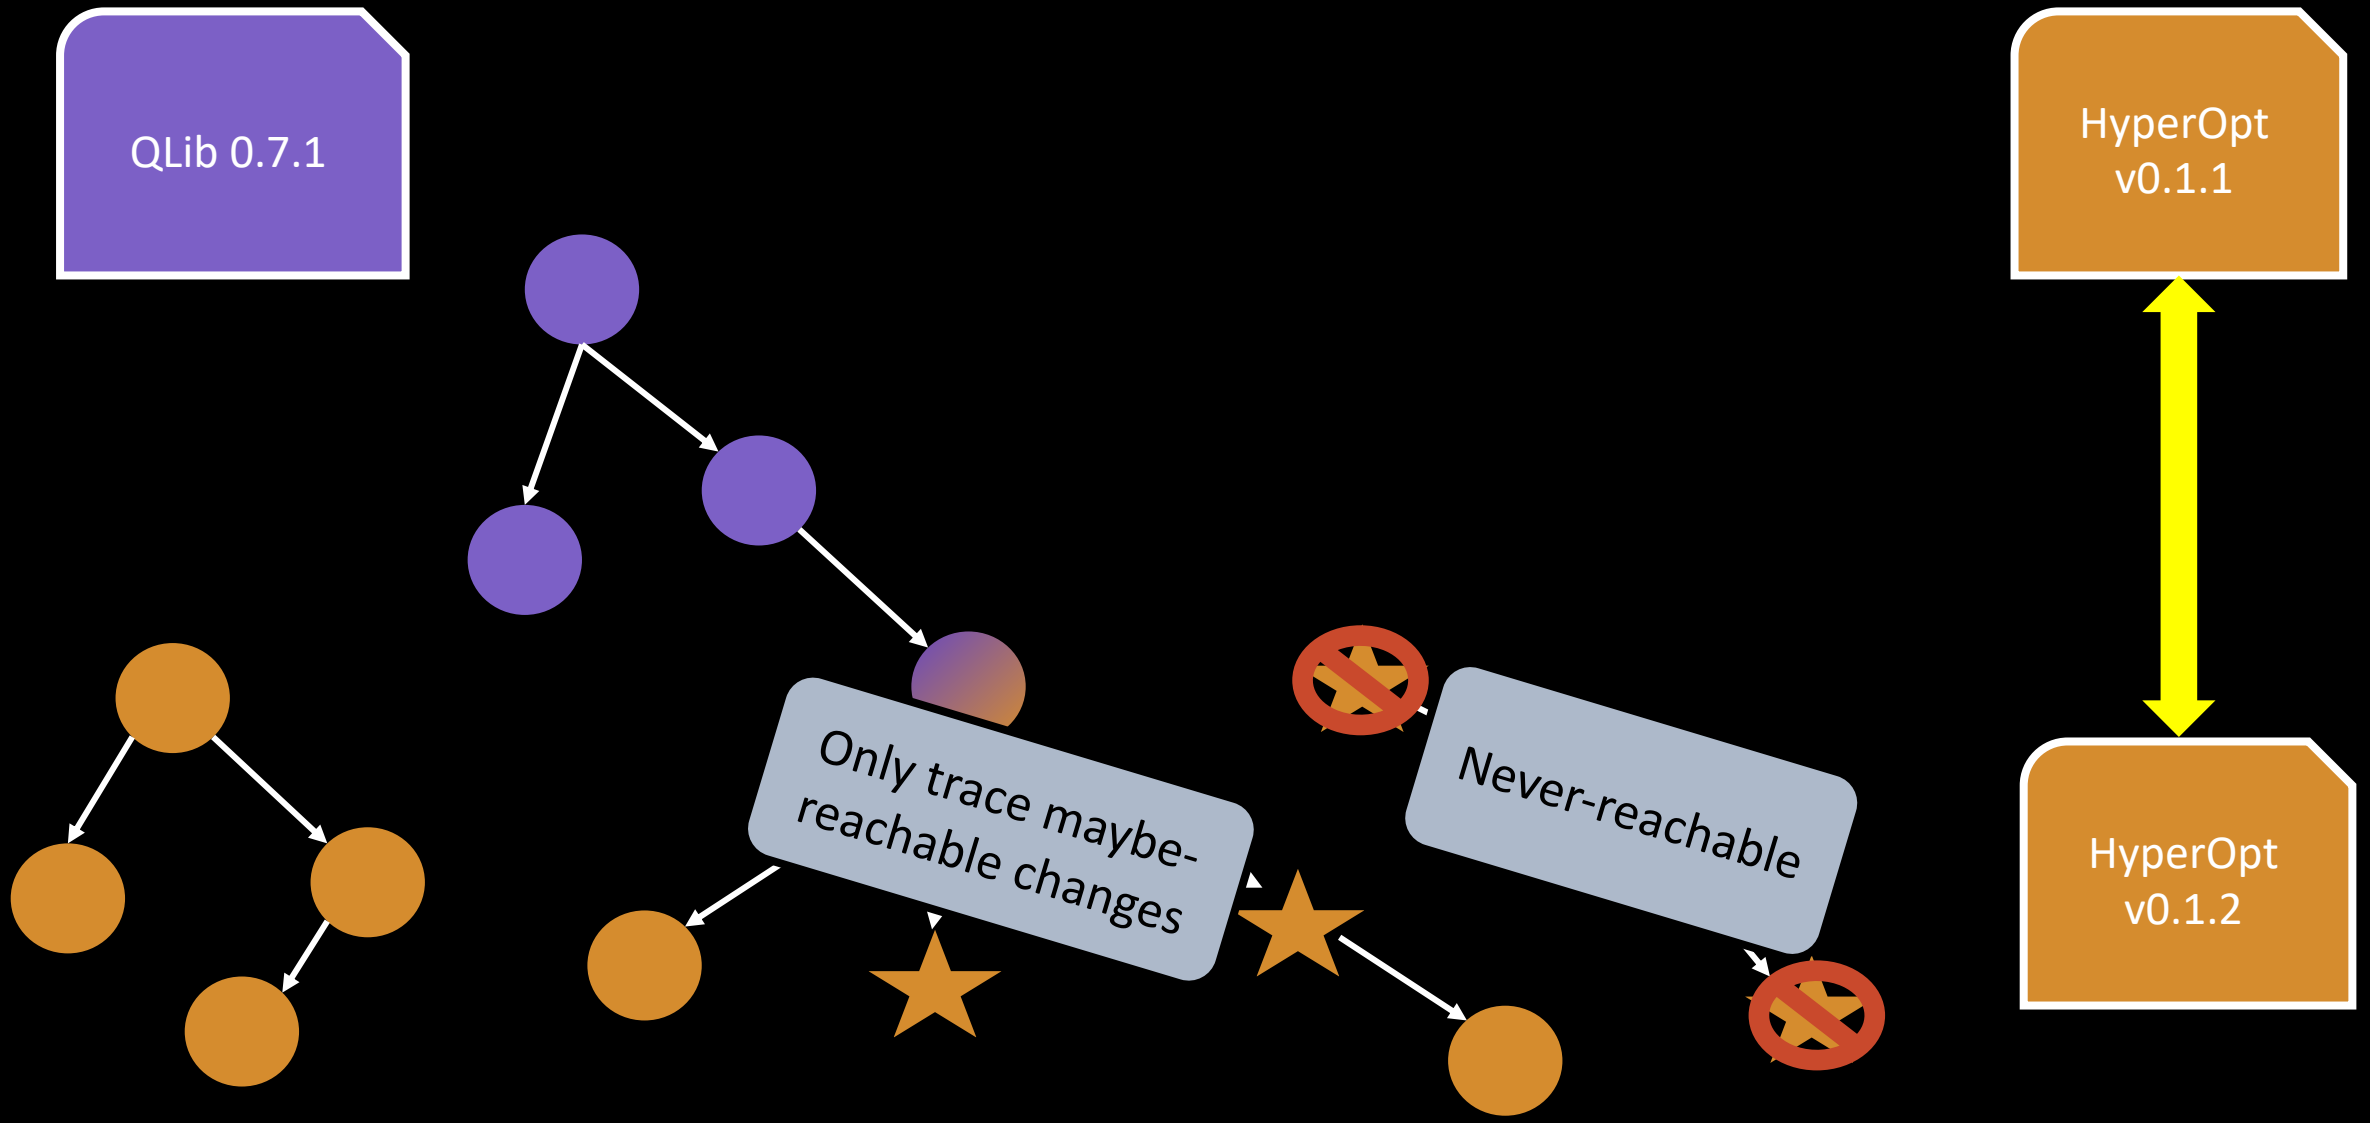

# Classify Changes In Motivating Example

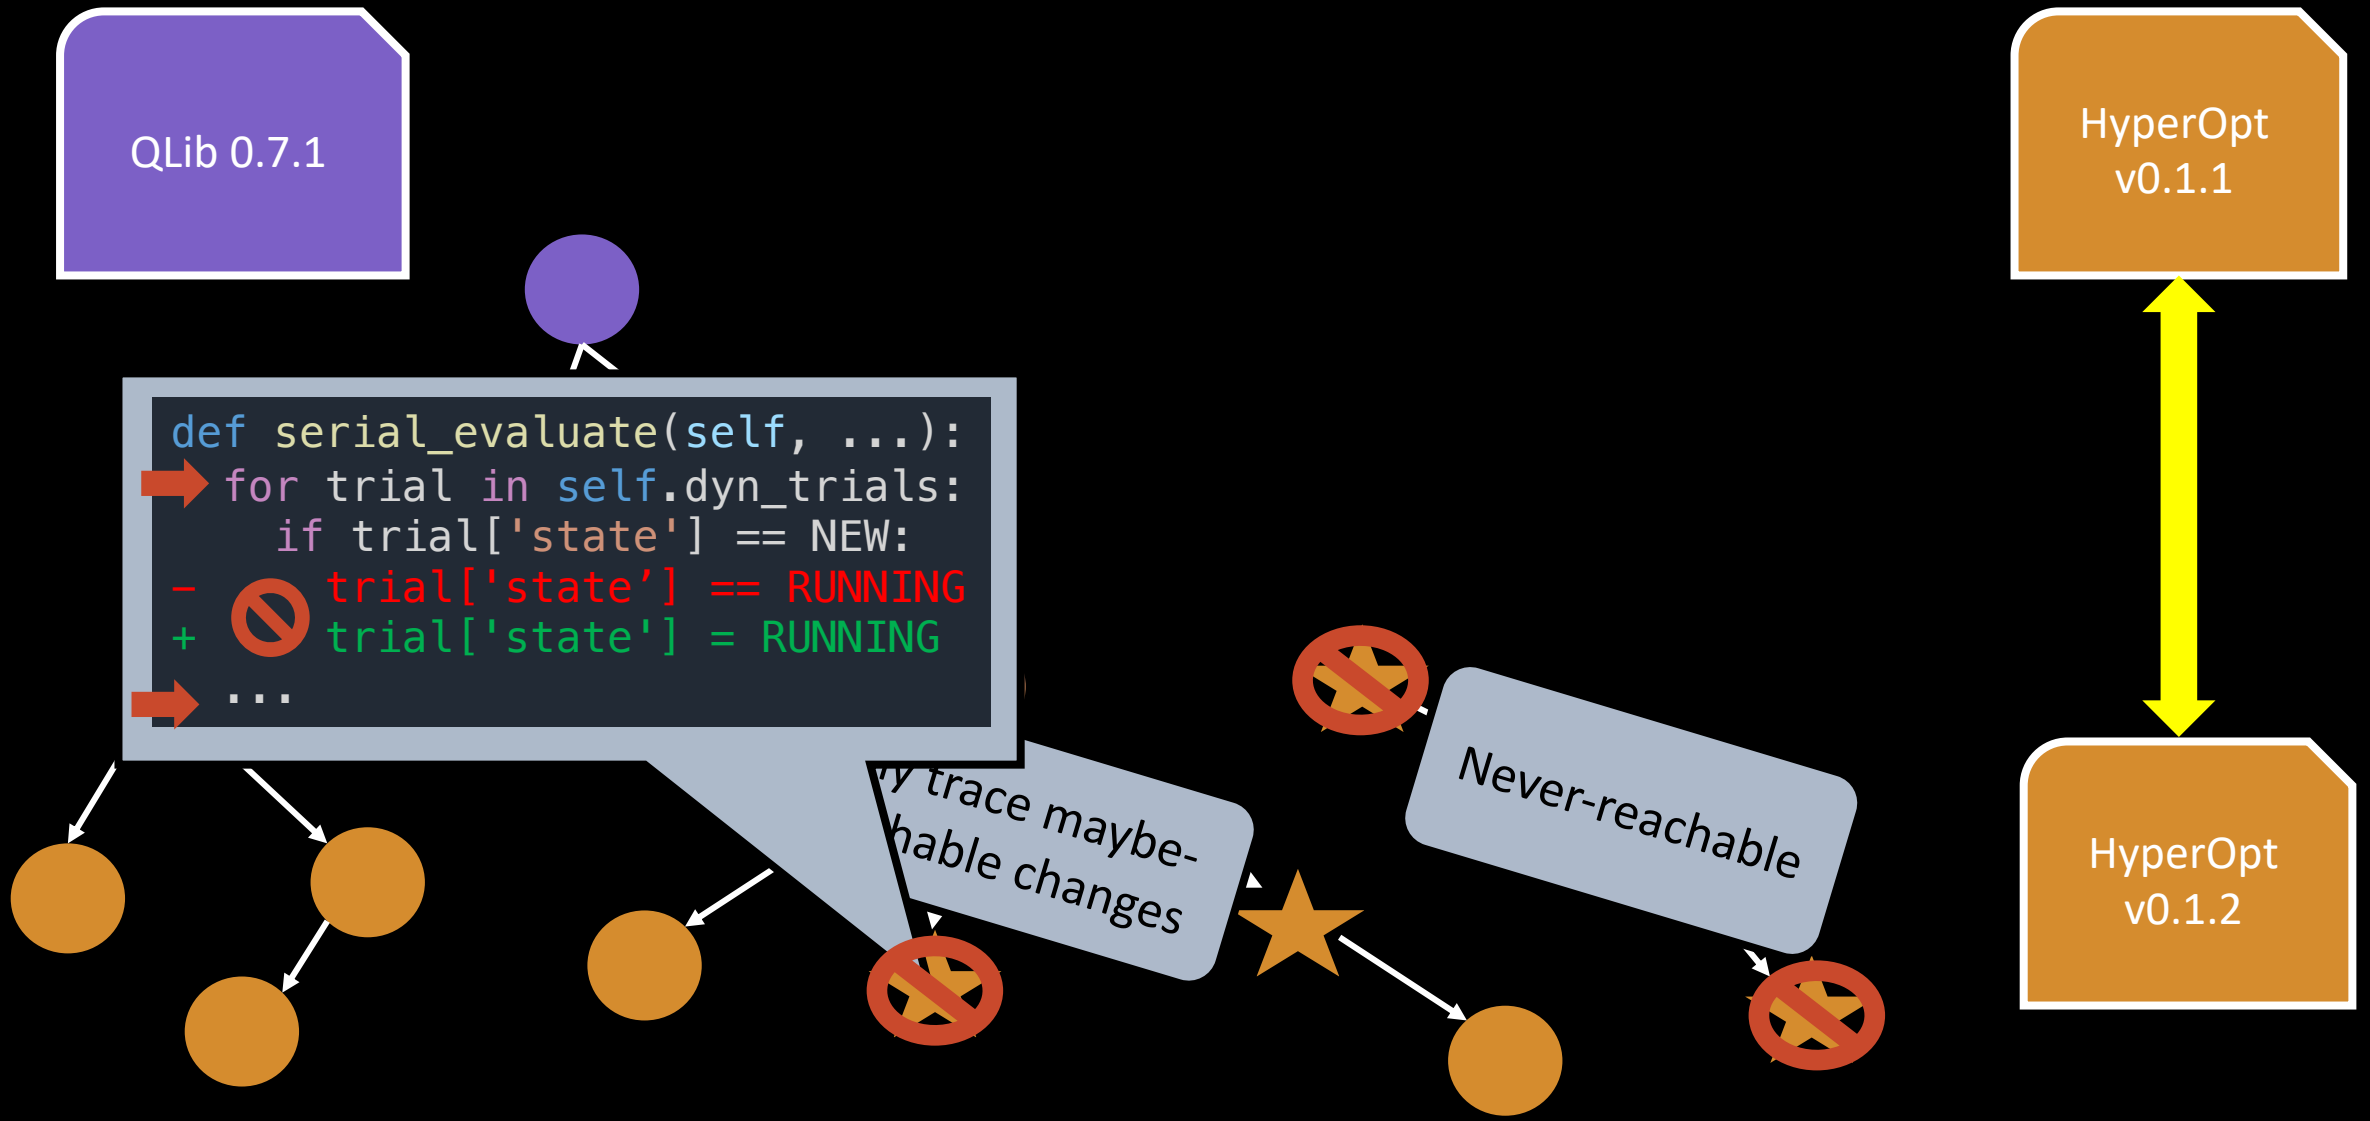

# Classify Changes In Motivating Example

QLib 0.7.1

HyperOpt v0.1.1

HyperOpt v0.1.2

| Changes | Initial Count | Discarded | | Left |
|---|---|---|---|---|
| | | Static | Dynamic | |
| Functions | 72 | 68 (94%) | 1 | 3 |

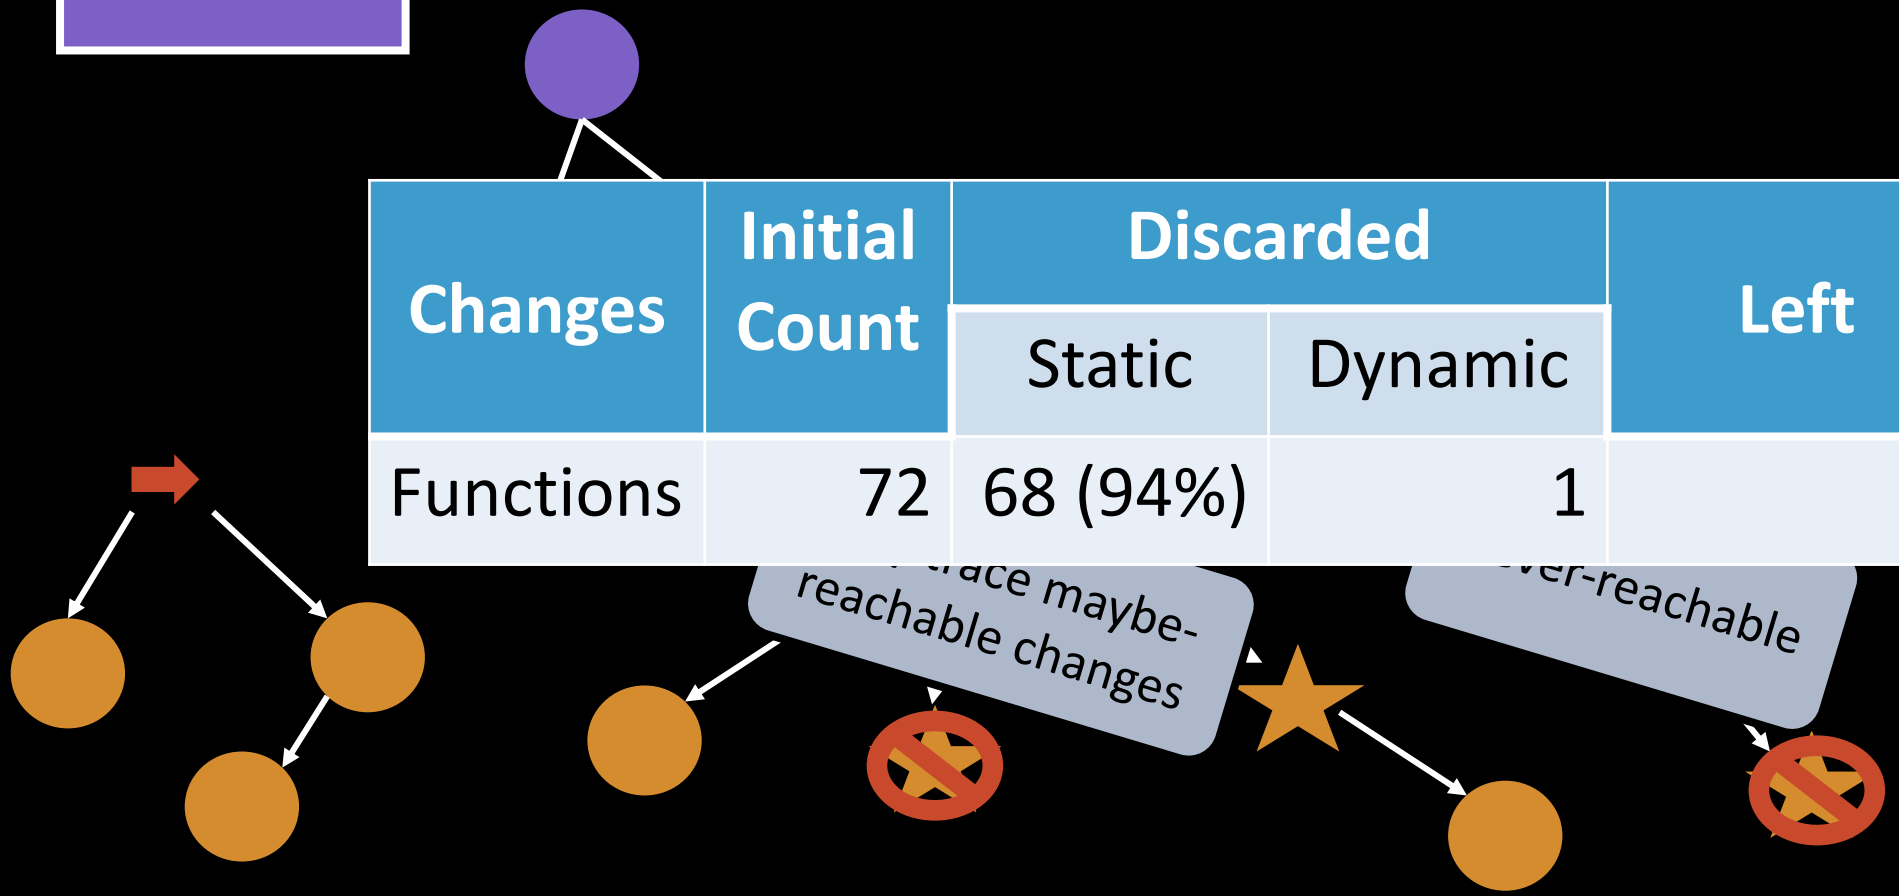

# Classify Changes In Motivating Example

## v0.8.1 🌈

github-actions released this Jan 15, 2022 · 164 commits to main since this release v0.8.1 e7954bd

### Changes

### 🌟 Features

- pylint code refine & Fix nested example **@you-n-g** (#848)
- chore: remove hard code input dimension of model pytorch_tcts **@PalanQu** (#843)
- [840] - Test case for operators. **@ChiahungTai** (#841)
- DDG-DA paper code **@you-n-g** (#743)
- Update BCELoss in MLP model **@cuicorey** (#756)
- solve VERSION.txt bug **@b4thesunrise** (#732)
- Hyperopt upgrade **@upgradvisor-bot** (#741)
- Add method parameter for volume **@you-n-g** (#734)

# Key Challenges for Designing UPGRADVISOR

- Hybrid program analysis to safely discard non-reachable changes
  - Safely discard non-reachable changes via hybrid program analysis
    - Create sound call-graphs
    - Reachable-but-non-affecting changes
  - Dynamic analysis to test maybe-reachable changes
- Achieve low-overhead using a hardware-based tracer
    - Low-overhead & selective tracing

# Key Challenges for Designing UPGRADVISOR

- Hybrid program analysis to safely discard non-reachable changes
  - Safely discard non-reachable changes via hybrid program analysis
    - Create sound call-graphs
    - Reachable-but-non-affecting changes
  - Dynamic analysis to test maybe-reachable changes
- Achieve low-overhead using a hardware-based tracer
    - Low-overhead & selective tracing

# Hardware Tracing for Native Code

```
void foo(int a){
  if (a==0){
    // something
  } else {
    // something else
  }
}
```

```
    cmp         rdi, 0
    jne         .EL
    nop
    jmp RET
.EL
    nop
.RET
    ret
```

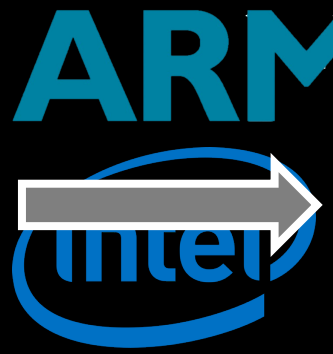


Recreate Trace Offline

Cyclic-write RAM buffer
(Usually dumped to disk)

Tracing Reco

Jump Not Taken

# Hardware Tracing for Interpreter Code

```
def foo(a):
  if a==0:
    # something
  else:
    # something else
```

```
LOAD_GLOBAL          0 (a)
LOAD_CONST           1 (0)
COMPARE_OP           2 (==)
POP_JUMP_IF_FALSE    10
```

python™

JPortal: Precise and Efficient Control-Flow Tracing for JVM Programs with Intel Processor Trace [PLDI ' 21]

```
for (ByteCode bc : allcode)
  switch (bc){
    case LOAD_GLOBAL:
      // do load global
      break;
    case LOAD_CONST:
      // do load const
      break;
    case ...:
}
```

Tracing Records

Jump to LOAD_GLOBAL

<jumps @ LOAD_GLOBAL>

Jump to LOAD_CONST

<jumps @ LOAD_CONST>

…

# Hardware Tracing for Interpreter Code

```
def foo(a):
  if a==0:
    # something
  else:
    # something else
```

```
LOAD_GLOBAL          0 (a)
LOAD_CONST           1 (0)
COMPARE_OP           2 (==)
POP_JUMP_IF_FALSE    10
```

python™

```
for (ByteCode bc : allcode)
  switch (bc){
    case LOAD_GLOBAL:
      // do load global
      break;
    case LOAD_CONST:
      // do load const
      break;
    case ...:
}
```

Recreate Bytecode Trace

Recreate Interpreter Trace

Tracing Records

Jump to LOAD_GLOBAL

<jumps @ LOAD_GLOBAL>

Jump to LOAD_CONST

<jumps @ LOAD_CONST>

…

# Hardware Tracing for Interpreter Code

```
def foo(a):
  if a==0:
    # something
  else:
    # something else
```

```
LOAD_GLOBAL          0 (a)
LOAD_CONST           1 (0)
COMPARE_OP           2 (==)
POP_JUMP_IF_FALSE    10
```


Recreate Bytecode Trace

Recreate Interpreter Trace

```
for (Byte
  switch
    case
      // do load global
      break;
    case LOAD_CONST:
      // do load const
      break;
    case ...:
}
```

High overhead
& Data loss

Tracing Records

...mp to LOAD_GLOBAL

<jumps @ LOAD_GLOBAL>

Jump to LOAD_CONST

<jumps @ LOAD_CONST>

...

# Hardware Tracing for Interpreted Code

# Hardware Tracing for Interpreted Code

```
for (ByteCode bc : all_code)
  switch (bc){
    case Op_Code_1:
      // do Op_Code_1
      break;
     case POP_JUMP_IF_F:
      jump_to_trace1()
      // do POP_JUMP_IF_F
      break;
    case : ...
}
```

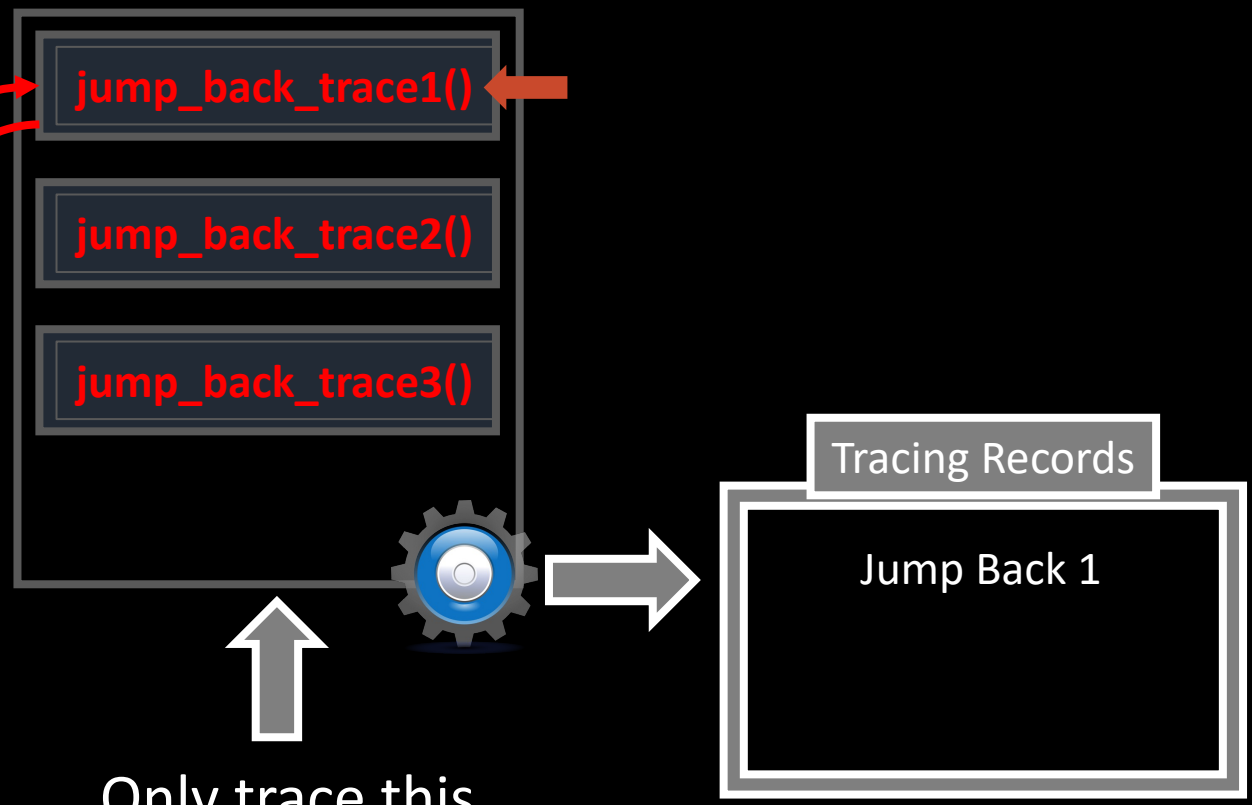


Only Selected Functions

Only trace this code segment

# Hardware Tracing for Interpreted Code

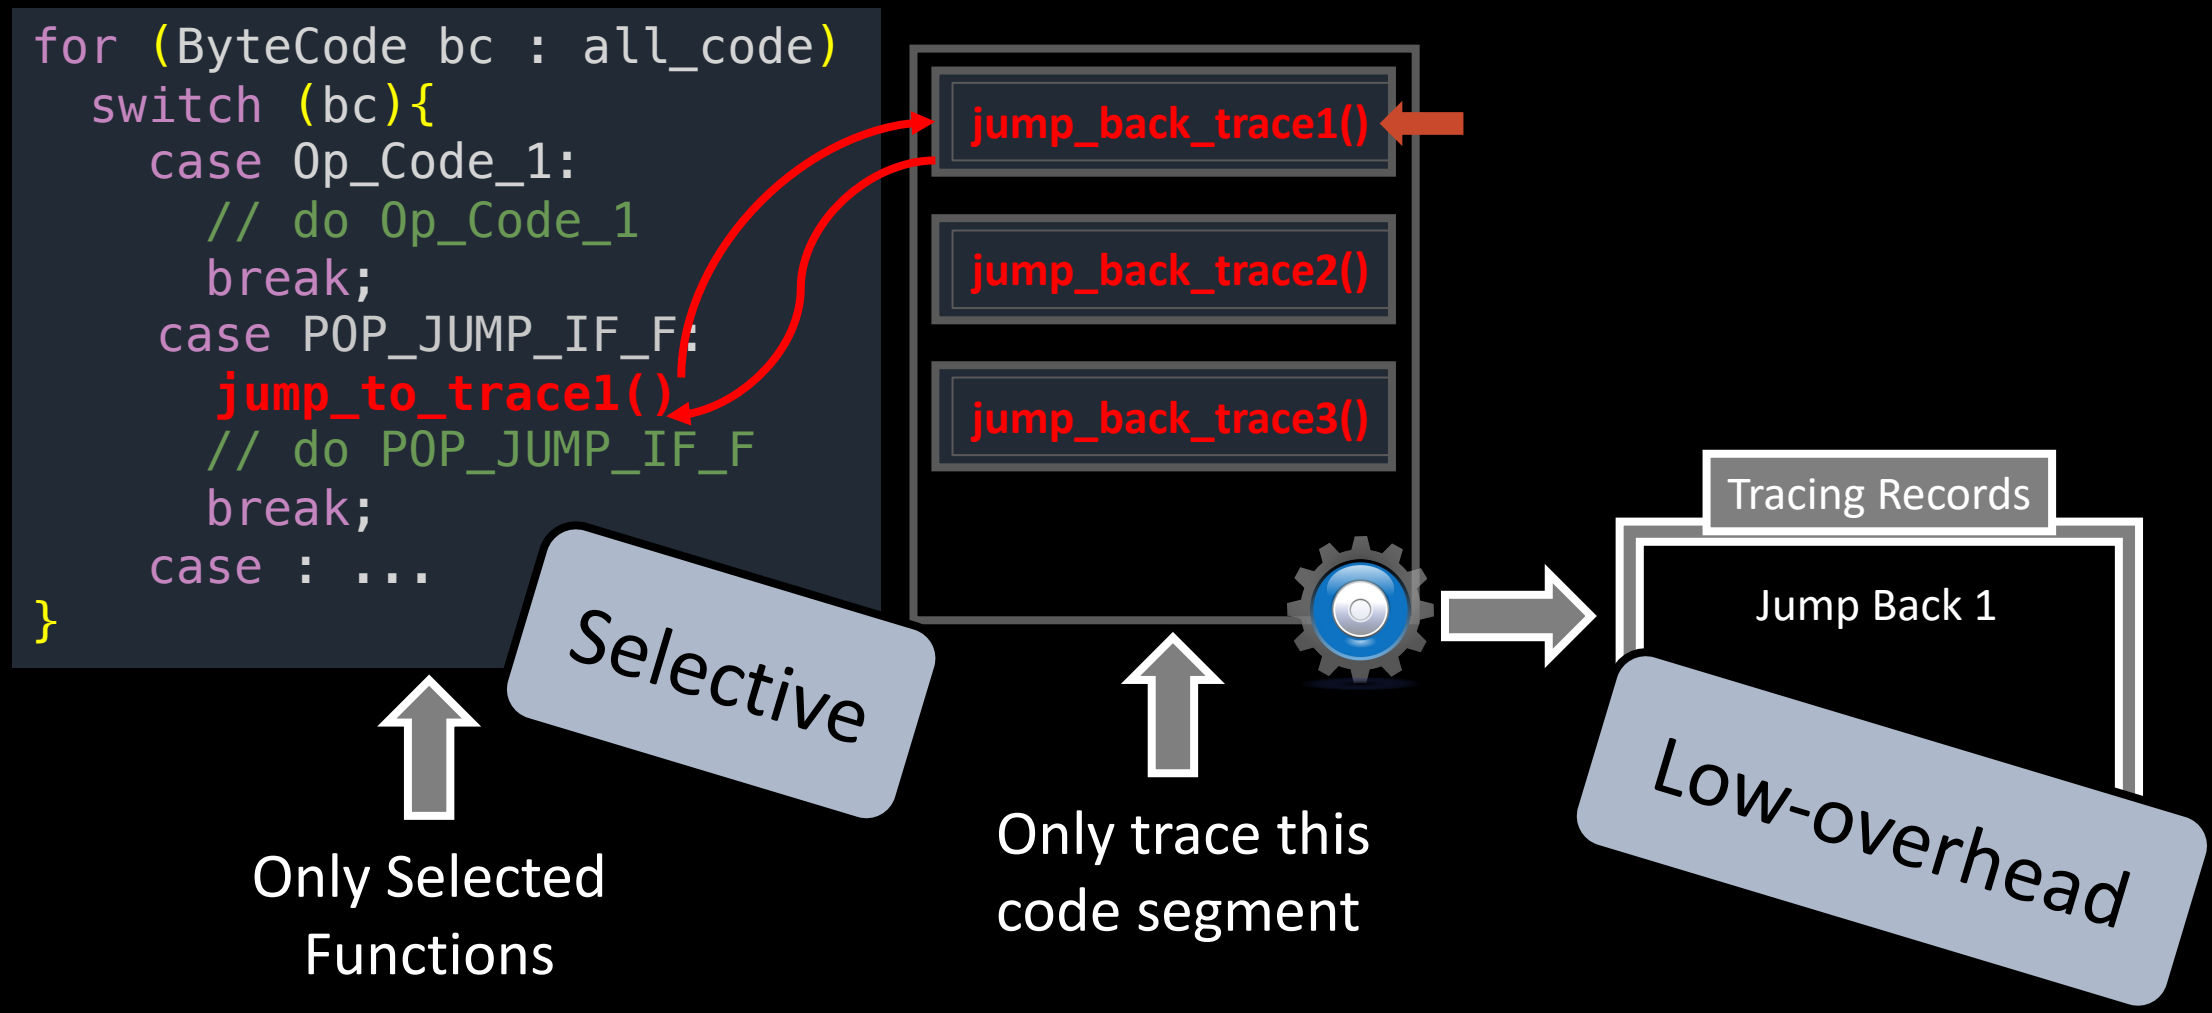

# Evaluation – Facilitating Dependency Updates

Updateable: 172

Static-Safe: 98 (56%)

Tracing Required:74 (44%)

For 5 Projects: Production-Like Tracing

# The Dynamic Tracing Contribution

| Project (Dependency) | Diff (LOC) | % Discarded | | % Left |
|---|---|---|---|---|
| | | Static | Dynamic | |
| AutoML(Distributed) | 820 | 95 | 5 | 0 |
| Electrum (qdarkstyle) | 641 | 88 | 8 | 4 |
| Flair (gdown) | 1500 | 71 | 29 | 0 |
| Qlib (Hyperopt) | 828 | 90 | 9 | 1 |
| Scylla (requests) | 449 | 90 | 8 | 2 |

# Tracer Overhead Testing Setup

- Selected Python projects with robust test-suites from our data-set
  - For Django's also running in parallel: using 1, 8, and 16 cores

# Tracer Overhead Testing Setup

- Selected Python projects with robust
  - For Django's also running in parallel
- UPGRADVISOR-Targeted

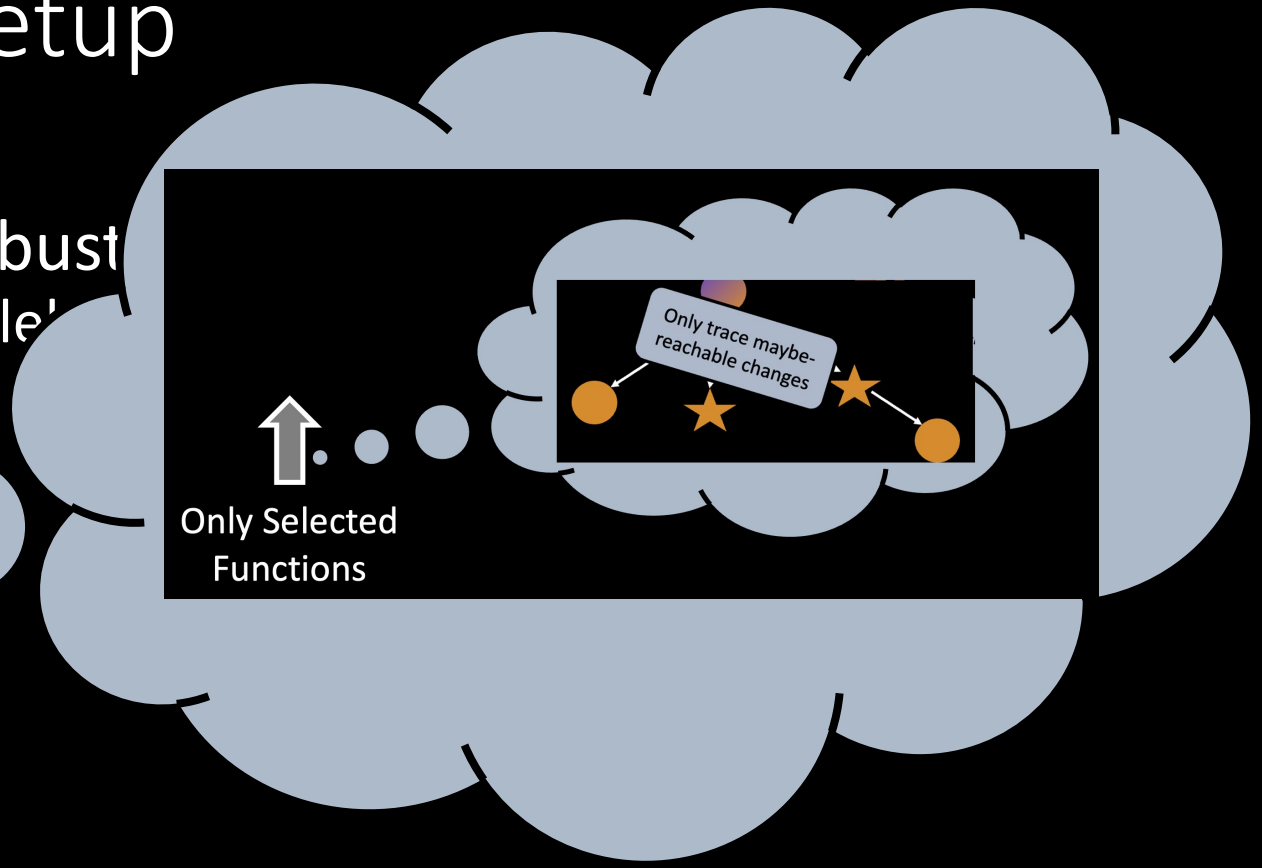

# Tracer Overhead Testing Setup

- Selected Python projects with r… our data-set
  - For Django's also running in …

- UPGRADVISOR-Targeted
- UPGRADVISOR-ALL
- Jportal4Py

# Tracer Overhead Testing Setup

- Selected Python projects with robust test-suites from our data-set
  - For Django's also running in parallel: using 1, 8, and 16 cores

- UPGRADVISOR-Targeted
- UPGRADVISOR-ALL
- Jportal4Py
- UPGRADVISOR-SW

# UPGRADVISOR's Tracer Incurs Low-Overhead

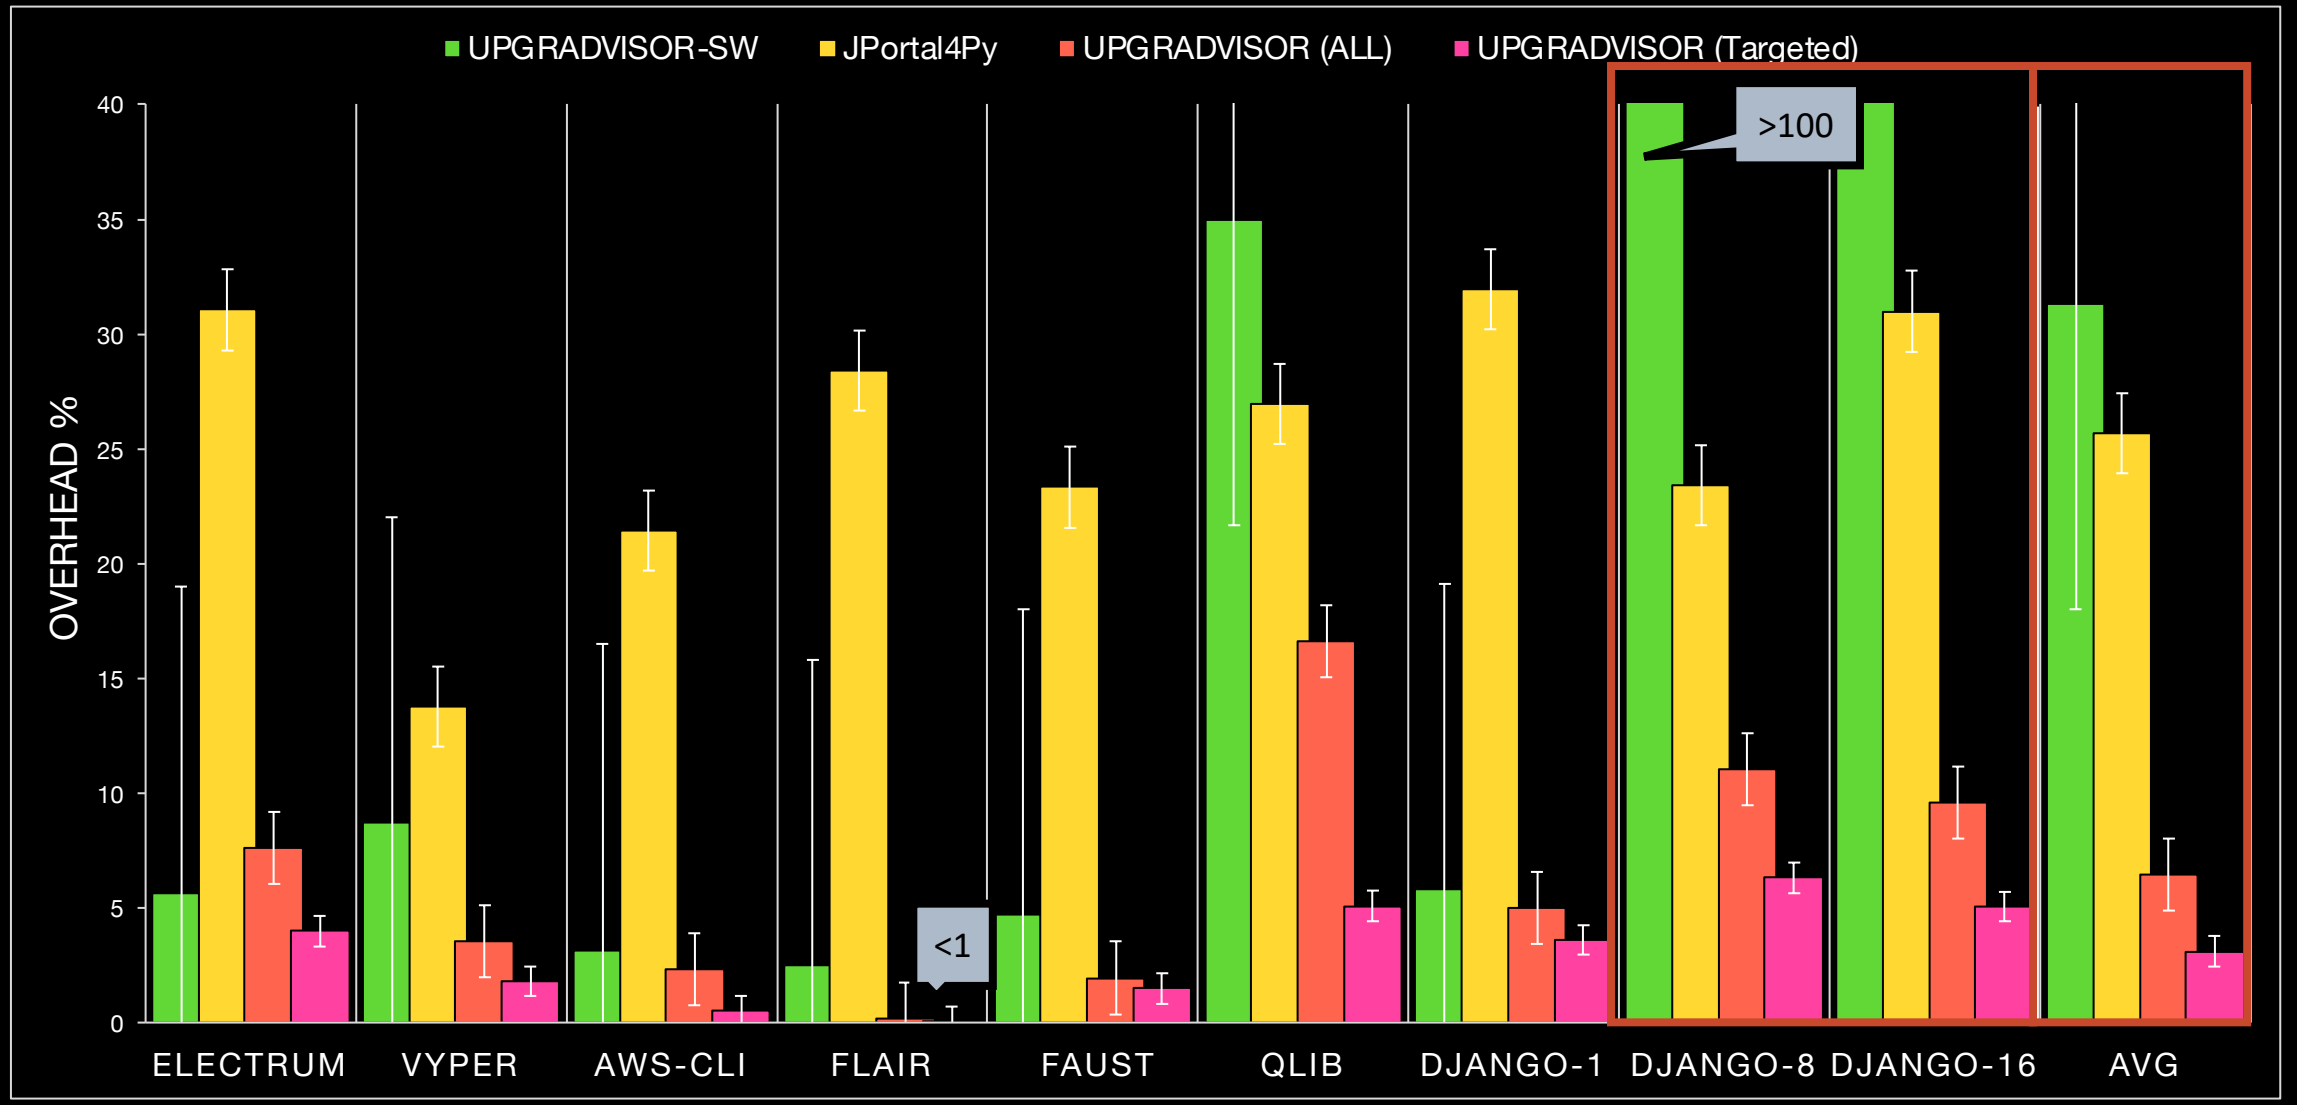

# UPGRADVISOR's Tracer Incurs Low-Overhead

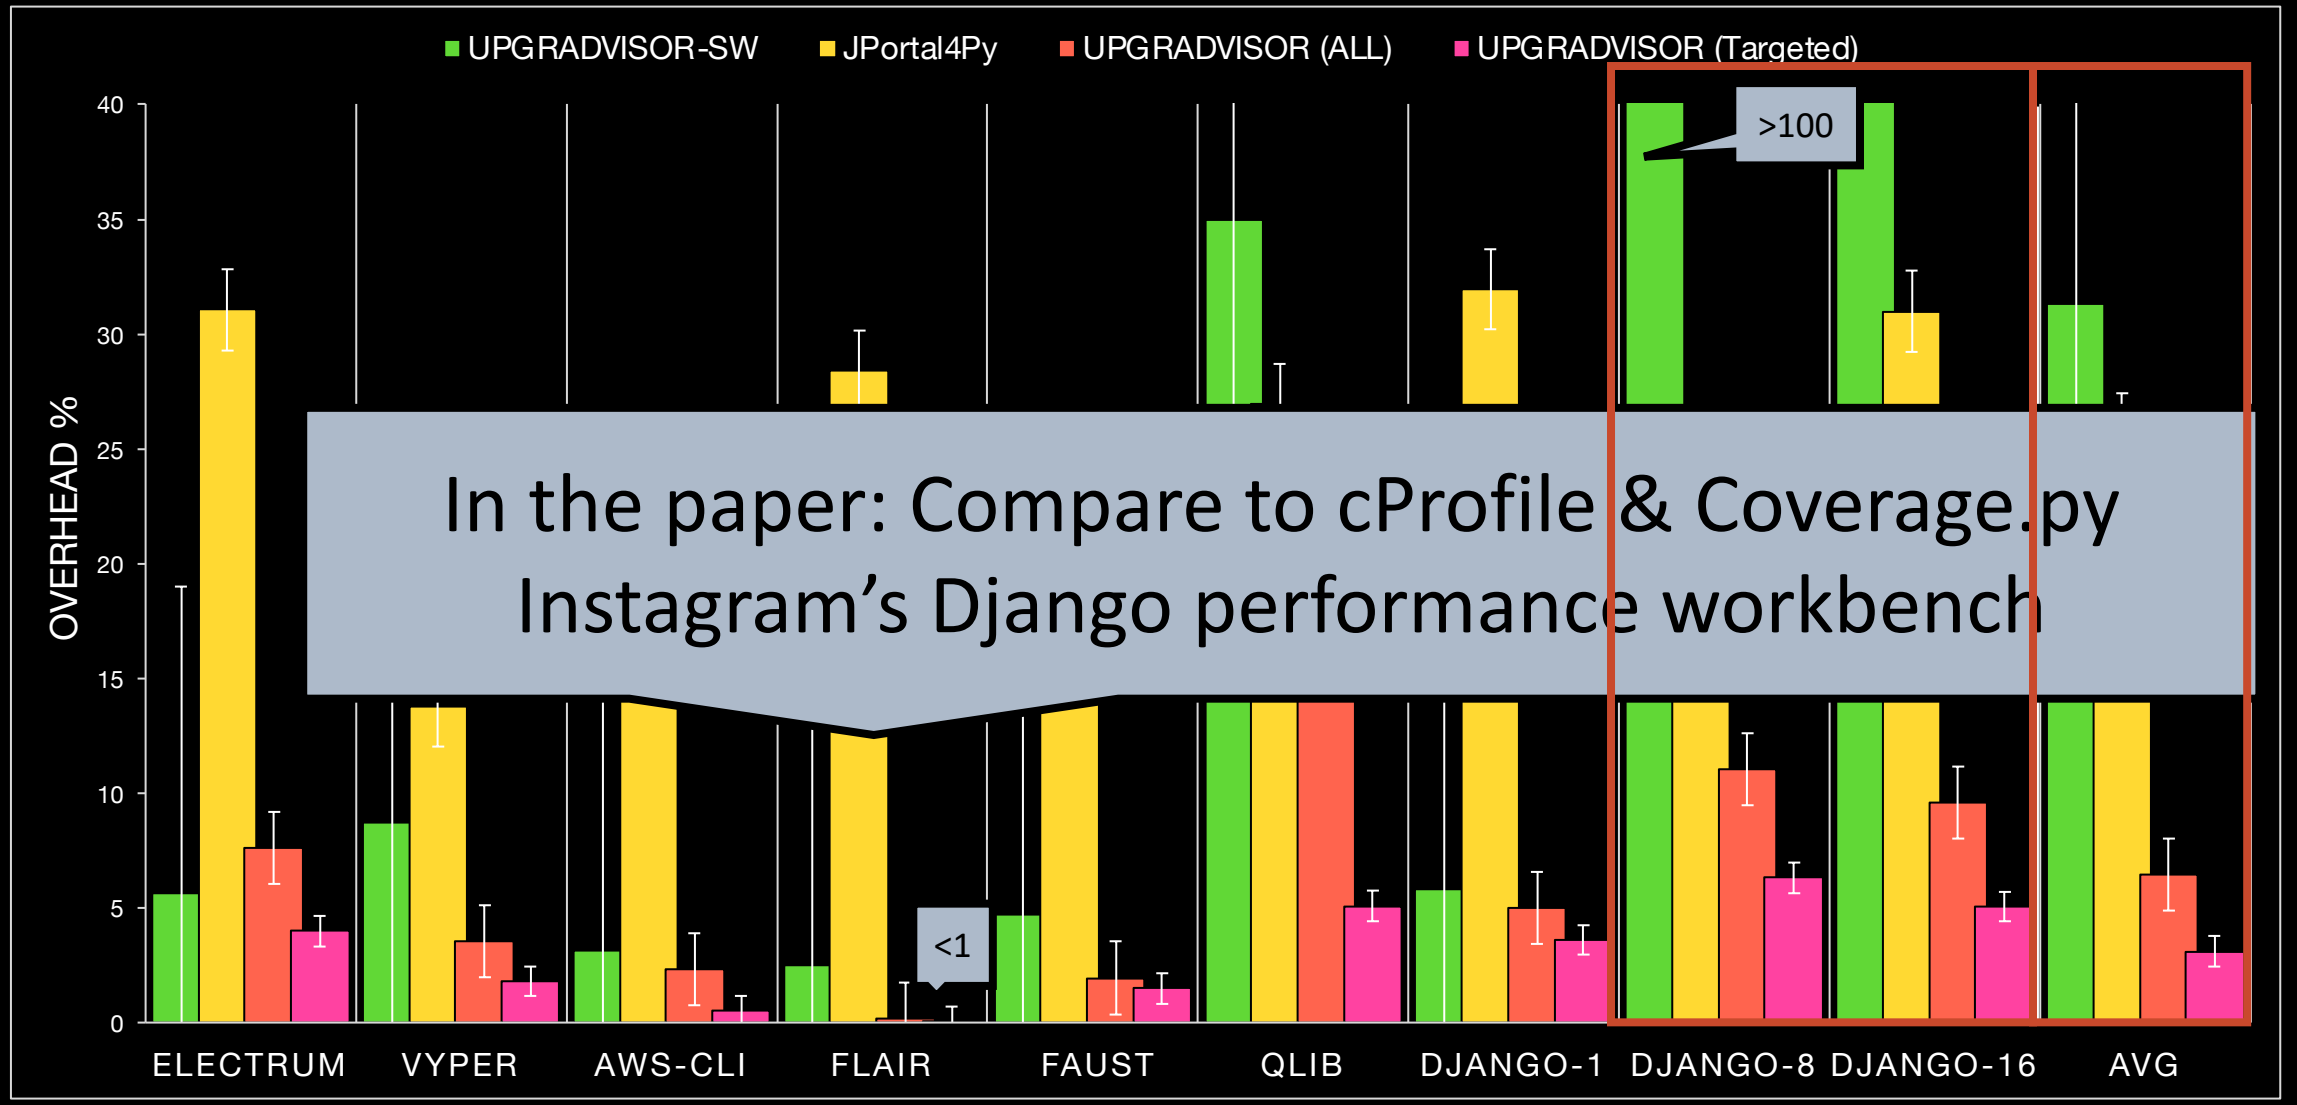

# Conclusion

- We presented UPGRADVISOR: a system for reducing developer effort and error risk in adopting dependency updates

- Want to know more? See our website!

  https://upgradvisor.github.io


- Want to use UPGRADVISOR-Python3? Install our free GitHub App

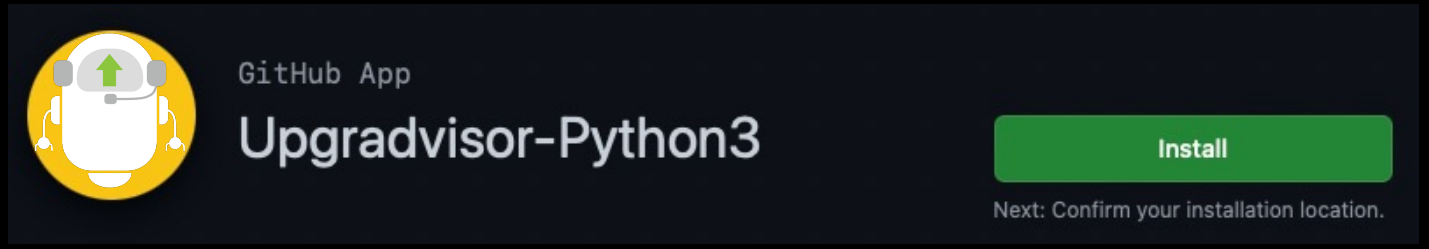

# Conclusion

- Want to use UPGRADVISOR? Install our free GitHub App

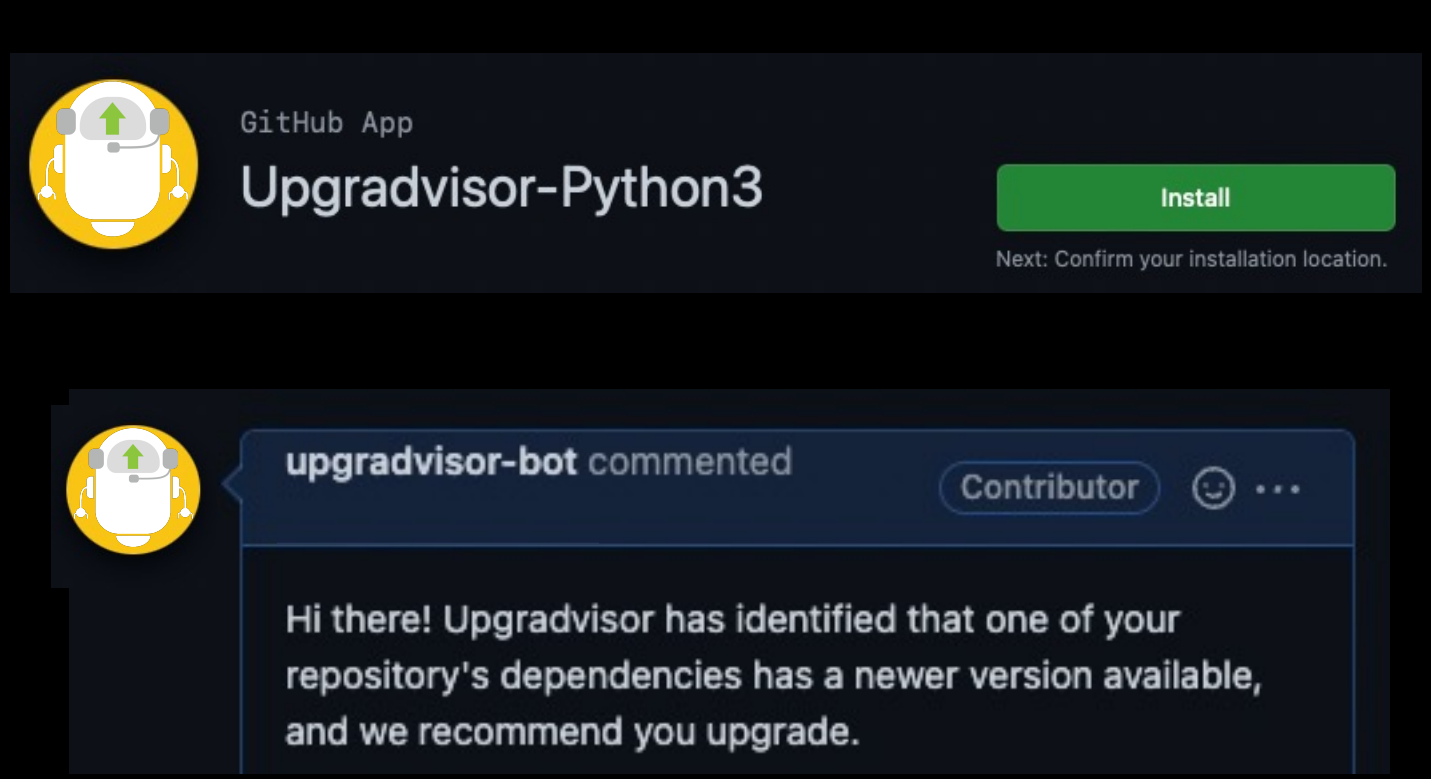

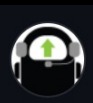

**upgradvisor-bot** commented on May 17 Contributor

Hi there! Upgradvisor has identified that one of your repository's dependencies has a newer version available, and we recommend you upgrade.

Your code currently pins `requests 2.26.0`. when requests `2.27.1` is available. Our analysis indicates that the impacts from this upgrade may fix a bug in `scylla`. The attached graph shows the dependency path of your repository relative to `requests`. Your code is shown in green (each node is a method), and your code calling `requests` is shown in orange. Changes between version `2.26.0` and `2.27.1` are shown as starred.

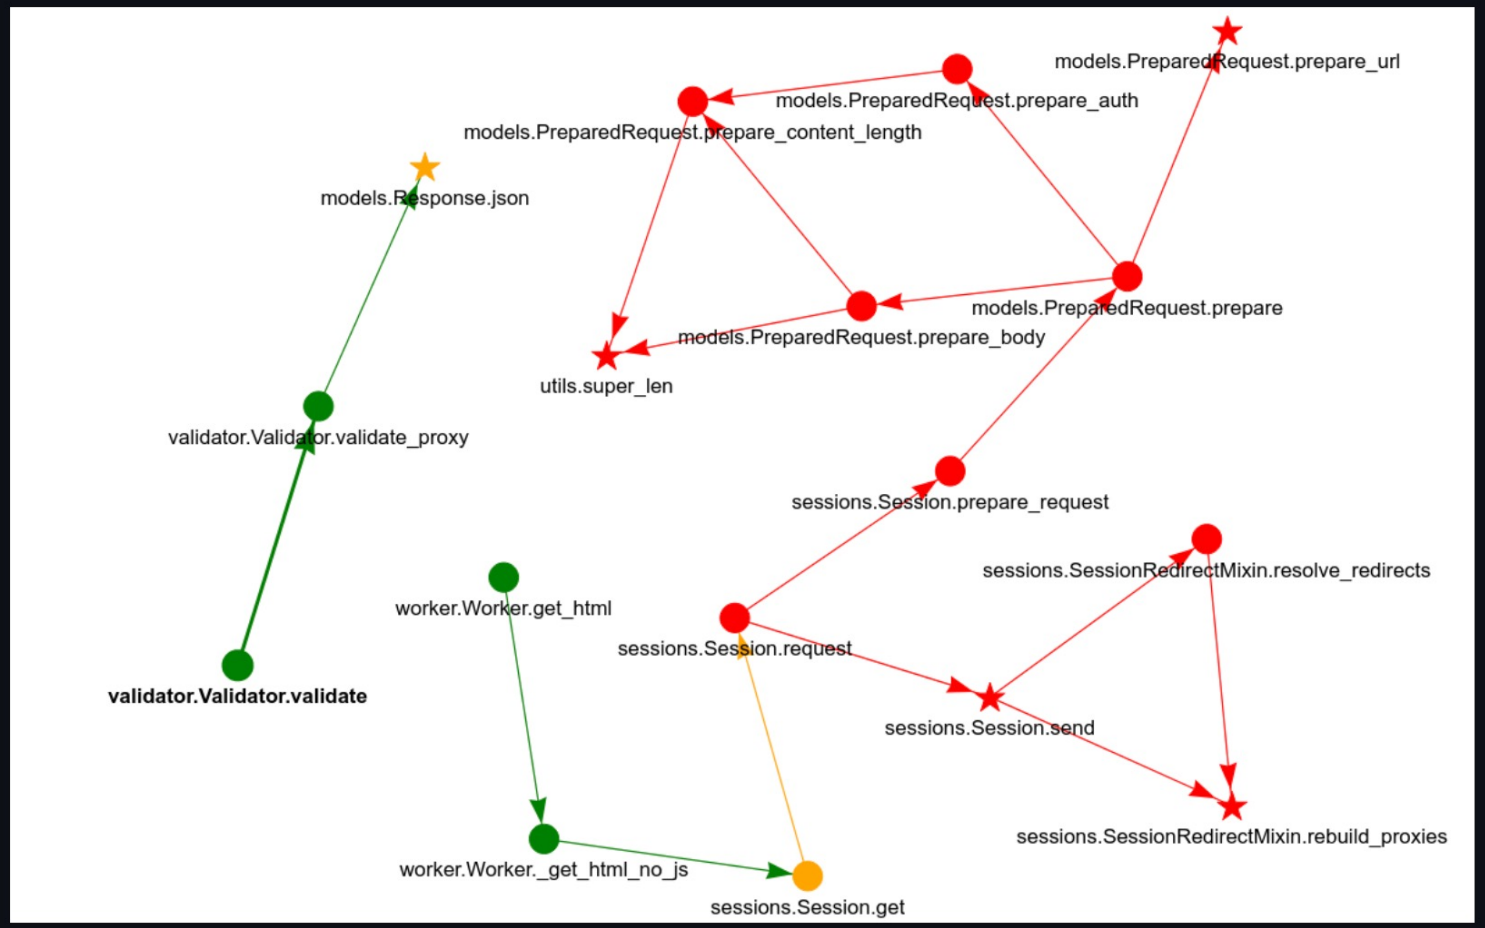

**upgradvisor-bot** commented on May 17 Contributor

Hi there! Upgradvisor has identified that one of your repository's dependencies has a newer version available, and we recommend you upgrade.

Your code currently pins `requests 2.26.0` this upgrade may fix a bug in `scylla`. The `requests`. Your code is shown in green (e Changes between version `2.26.0` and `2.`

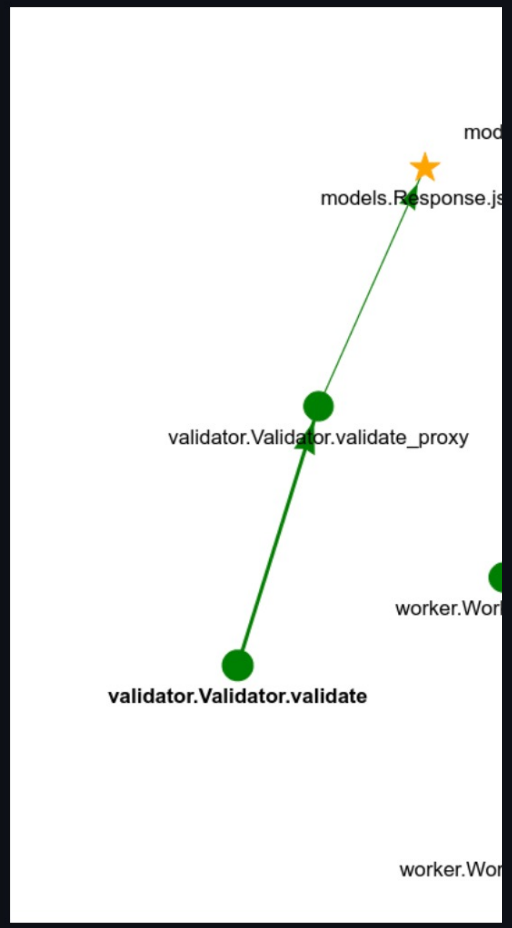


Upgradvisor is a research project from Columbia University's RCS lab. Our goal is to provide meaningful information for developers (like you) on how to upgrade their software dependencies. If you have any questions or feedback please reach out to the Upgradvisor team at yaniv.david@columbia.edu.

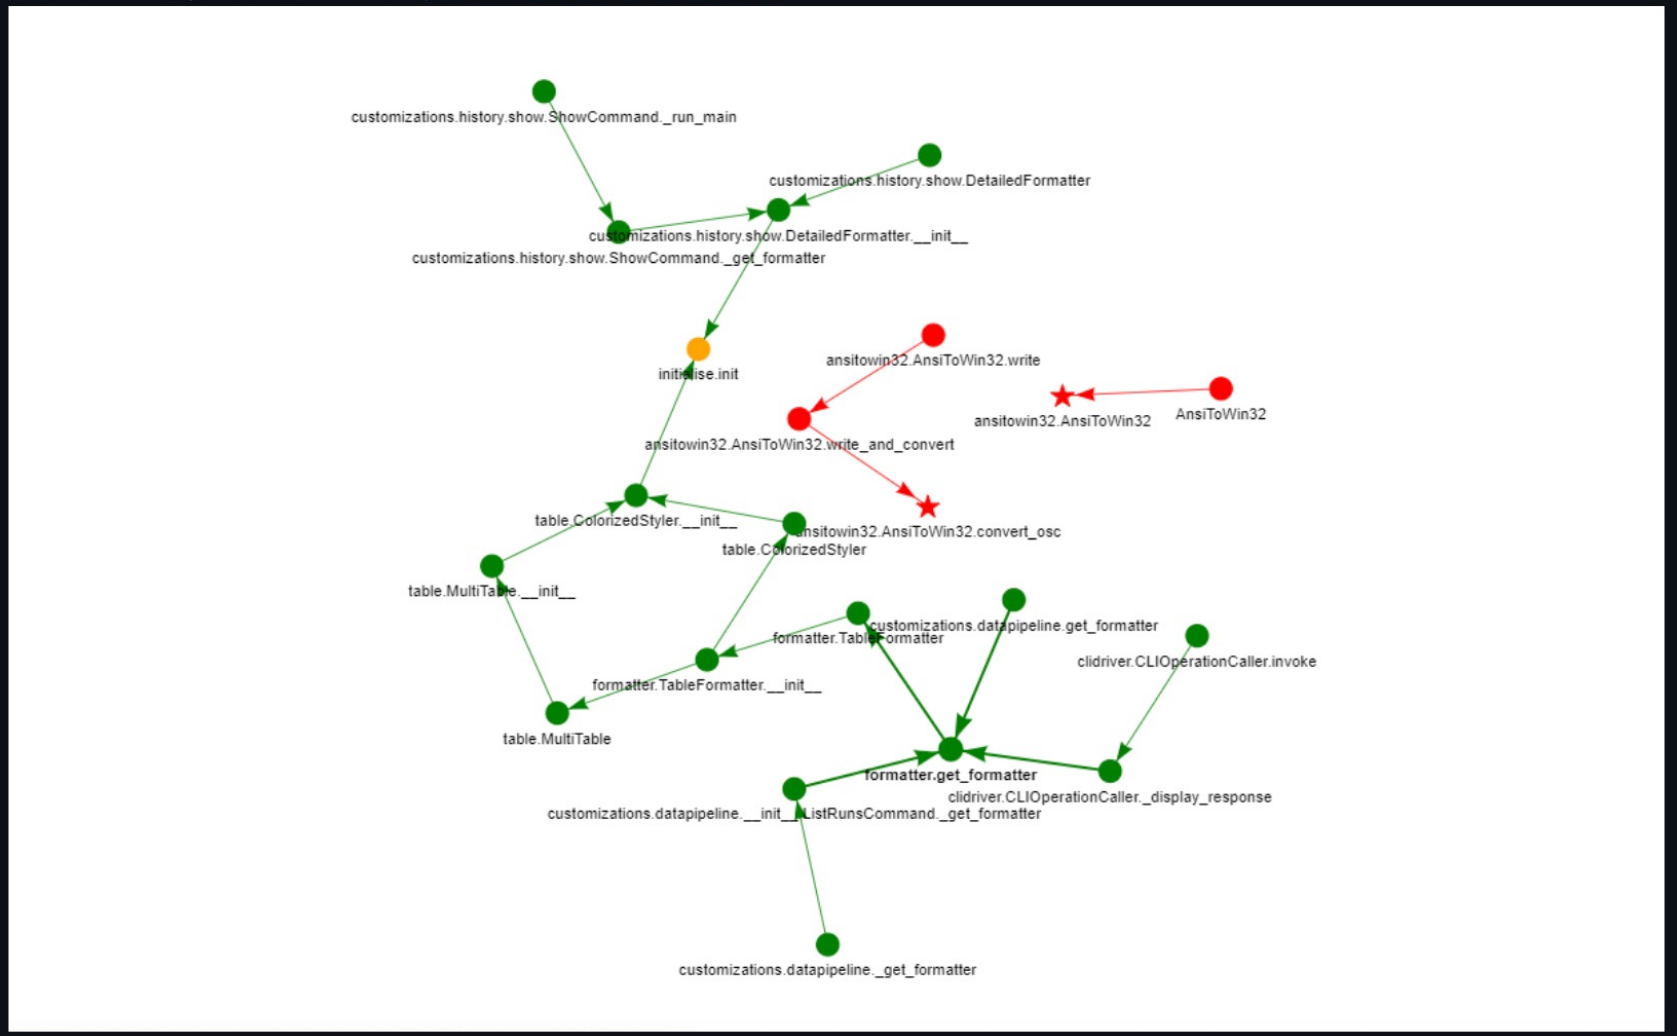

## Dependency Update Adoption Is Slow

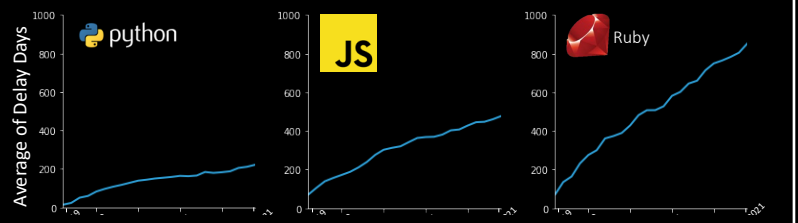


# Thank You!
# Questions?

## HW-based Tracing is Production-Ready

```
for (ByteCode bc : all_code)
  switch (bc){
    case Op_Code_1:
      // do Op_Code_1
      break;
    case POP_JUMP_IF_F:
      jump_to_trace1()
      // do POP_JUMP_IF_F
      break;
    case : ...
}
```

jump_back_trace1()
jump_back_trace2()
jump_back_trace3()

## UPGRADVISOR: a system for reducing developer effort and error risk in adopting dependency updates