CernVM File System

Status and Roadmap

CernVM Users Workshop 2015
René Meusel

CERN - 5th of March 2015

Agenda

- **Introduction**
  - Developing and Testing CernVM-FS
- **CernVM-FS Server Migration at CERN**
- **New Features in CernVM-FS**
  - Configuration Consolidation
  - History-less Repositories
  - **Planned Developments**
    - Multiple Union File Systems
    - REST API for Stratum0
    - Smart Stratum1 Servers
    - Consolidating Parrot

Introduction

What is CernVM-FS?

- **Scalable software distribution system**
  - Infrequent atomic updates in a central location
  - Read-only access on the clients
  - Repository signed by a trusted release manager
- **HTTP based global data transfer**
  - Minimal protocol requirements
  - Aggressive hierarchical cache strategy
    - Assumption: Coherent working set on physically close nodes (cf. software vs. data distribution)
- **Accessible through a mounted file system (POSIX)**
  - FUSE module, NFS exported FUSE volume or Parrot

Usage Statistics of CERN Repositories

Repository	Files	Objects	Logical Volume	avg. File Size	
cms.cern.ch	40 M	6.4 M	1.2 TB	31.4 kB	Software
alice.cern.ch	9 M	0.3 M	0.8 TB	94.6 kB	
lhcb.cern.ch	14 M	4.8 M	0.6 TB	44.5 kB	
atlas.cern.ch	40 M	4.2 M	2.5 TB	68.5 kB	
ams.cern.ch	4 M	2.5 M	2.1 TB	0.6 MB	Cond. Data
alice-ocdb.cern.ch	0.8 M	0.8 M	0.1 TB	0.2 MB	
atlas-condb.cern.ch	8 k	9 k	0.5 TB	62 MB	
lhcbdev.cern.ch	126 M	21.1 M	6.5 TB	56.9 kB	Nightly
cms-ib.cern.ch	9 M	1.6 M	0.3 TB	42.2 kB	

- *Files* and *Volume* as stored in the CernVM-FS catalogs
- Actual number of *Referenced Objects* is compressed and de-duplicated
- Based on latest revision - no history involved

(Effective: February 2015)

Developing and Testing CernVM-FS

Developing the CernVM-File System

- Active Development Team
  - Jakob and René
  - José recently joined the team
  - Seppo and Dave provide regular contributions

- Development workflow based on **GitHub Pull Requests**
- **Code Review** before any changes are merged
- Regular **static code analysis** using Coverity
- Regular **automated testing** is crucial
  - New features should to come with tests
  - Idea: *Untested code is broken!*

git-flow: http://nvie.com/posts/a-successful-git-branching-model/

CernVM-FS Source Code Figures

- SLOC of CernVM-File System (no blank lines and comments)

- CernVM-FS **Production Code**
  - 41’840 C/C++ and Header Files
  - 5’008 Shell Scripts and Perl

- CernVM-FS **Unit and Integration Tests**
  - 15’789 C++
  - 16’007 Shell Scripts

- Overall: 475 files and ~115’000 LOC (*including* comments and blanks)
- Measured in *ecce387494ab26758a38943179239accaf03c917*

CernVM-FS 2.1.x Server Migration

Repository Server Migration to 2.1.x

- CernVM-FS 2.1.x comes with a different server architecture
  - Based on a **union file system** instead of shadow directory
  - Schema changes in the catalog data structure

- Preconditions for the server migration
  - All clients on CernVM-FS 2.1.x ✓
  - Stratum 1 replication servers on CernVM-FS 2.1.x ✓
  - Automatic repository migration available in CernVM-FS ✓
    (First appeared in version 2.1.15 - fully stabilised in 2.1.20)

- **Successful migration campaign at CERN in 2014**
  for all hosted CernVM-FS repositories

Outcome and Experience

DONE

- End of Life of CVMFS 2.0.x at CERN on Sept. 2nd 2014
- Overall **smooth transition** with only minor issues
  - Sporadic **outages on two Tier 3 sites** and individual users (still running CernVM-FS 2.0.x clients)
  - **Test4Theory outage** after migrating grid.cern.ch (LHC@Home 2.0) (a central job server was running CernVM-FS 2.0.x)
  - Some **minor bugs** in the server tools (CernVM-FS 2.1.19)
    - Release manager machine locks up repository after reboot (Hotfix: http://cernvm.cern.ch/portal/cvmfs/fix-failed-remount)
    - Possible file descriptor leakage during a publish operation (Hotfix: ulimit -n 100000)
    - All fixed with CernVM-FS 2.1.20

Configuration Consolidation

Consolidated Configuration of CernVM-FS

- Goals:
  - Disentangle CernVM-FS from CERN-specific configuration
  - Simplify CernVM-FS client configuration
  - Allow for 3rd party configuration packages
  - Facilitate support for non-HEP VOs

- New configuration methods in **CernVM-FS 2.1.20**
  - Introduction of *cvmfs-config-**** packages
  - Ability to use *configuration repositories*
  - Automatic location aware *ordering of Stratum1 servers*

CernVM-FS Configuration Changes

- Until **CernVM-FS 2.1.19** - Client package depended on:
  - **cvmfs-keys** package
    - contains CERN public keys
    - later also keys and configuration for EGI and OSG
  - (**cvmfs-init-scripts** package)
    - optional configuration for some CERN hosted repositories

- As of **CernVM-FS 2.1.20** - Client package requires *cvmfs-config*
  - Abstract meta-package to be fulfilled by *something*
  - We provide **cvmfs-config-default** and **cvmfs-config-none**
    - replaces legacy *cvmfs-keys* and *cvmfs-init-scripts*

CernVM-FS Configuration Repository

- Bootstrap repository for CernVM-FS clients
- Maintains configuration baseline and keys centrally

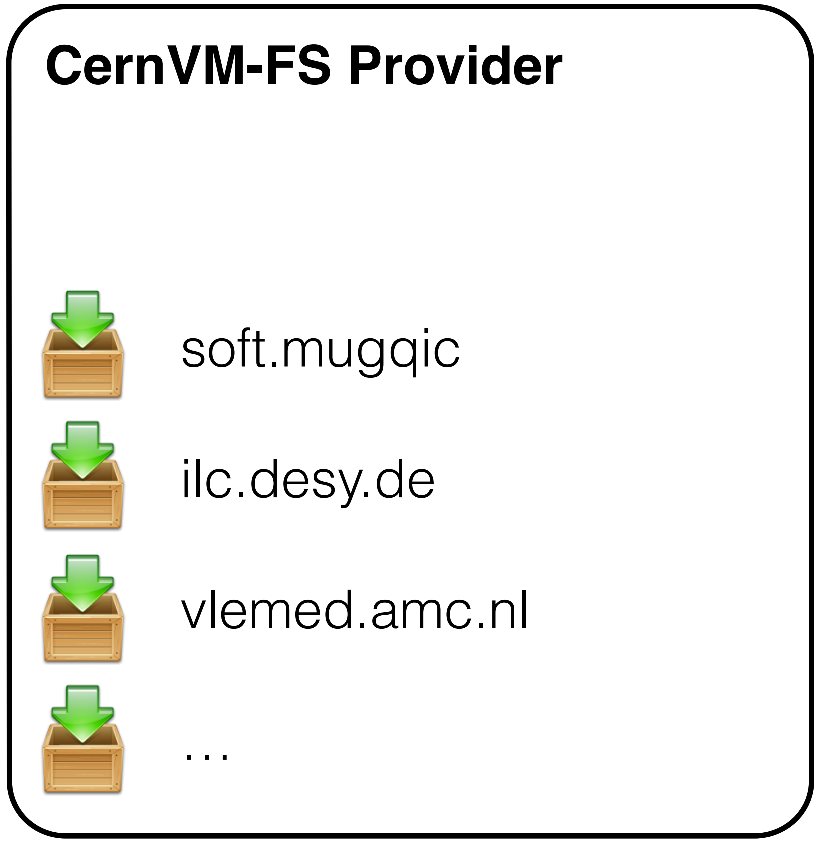
CernVM-FS Provider

soft.mugqic

ilc.desy.de

vlemed.amc.nl

…

CernVM-FS Configuration Repository

- Bootstrap repository for CernVM-FS clients
- Maintains configuration baseline and keys centrally

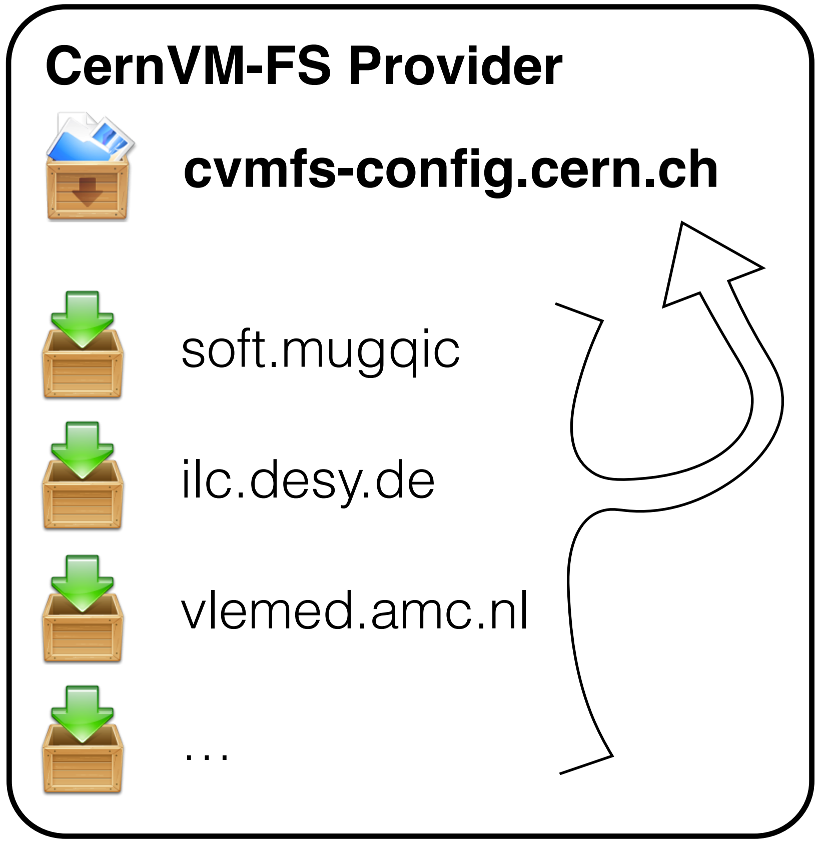

CernVM-FS Provider

cvmfs-config.cern.ch

soft.mugqic

ilc.desy.de

vlemed.amc.nl

…

CernVM-FS Configuration Repository

- Bootstrap repository for CernVM-FS clients
- Maintains configuration baseline and keys centrally

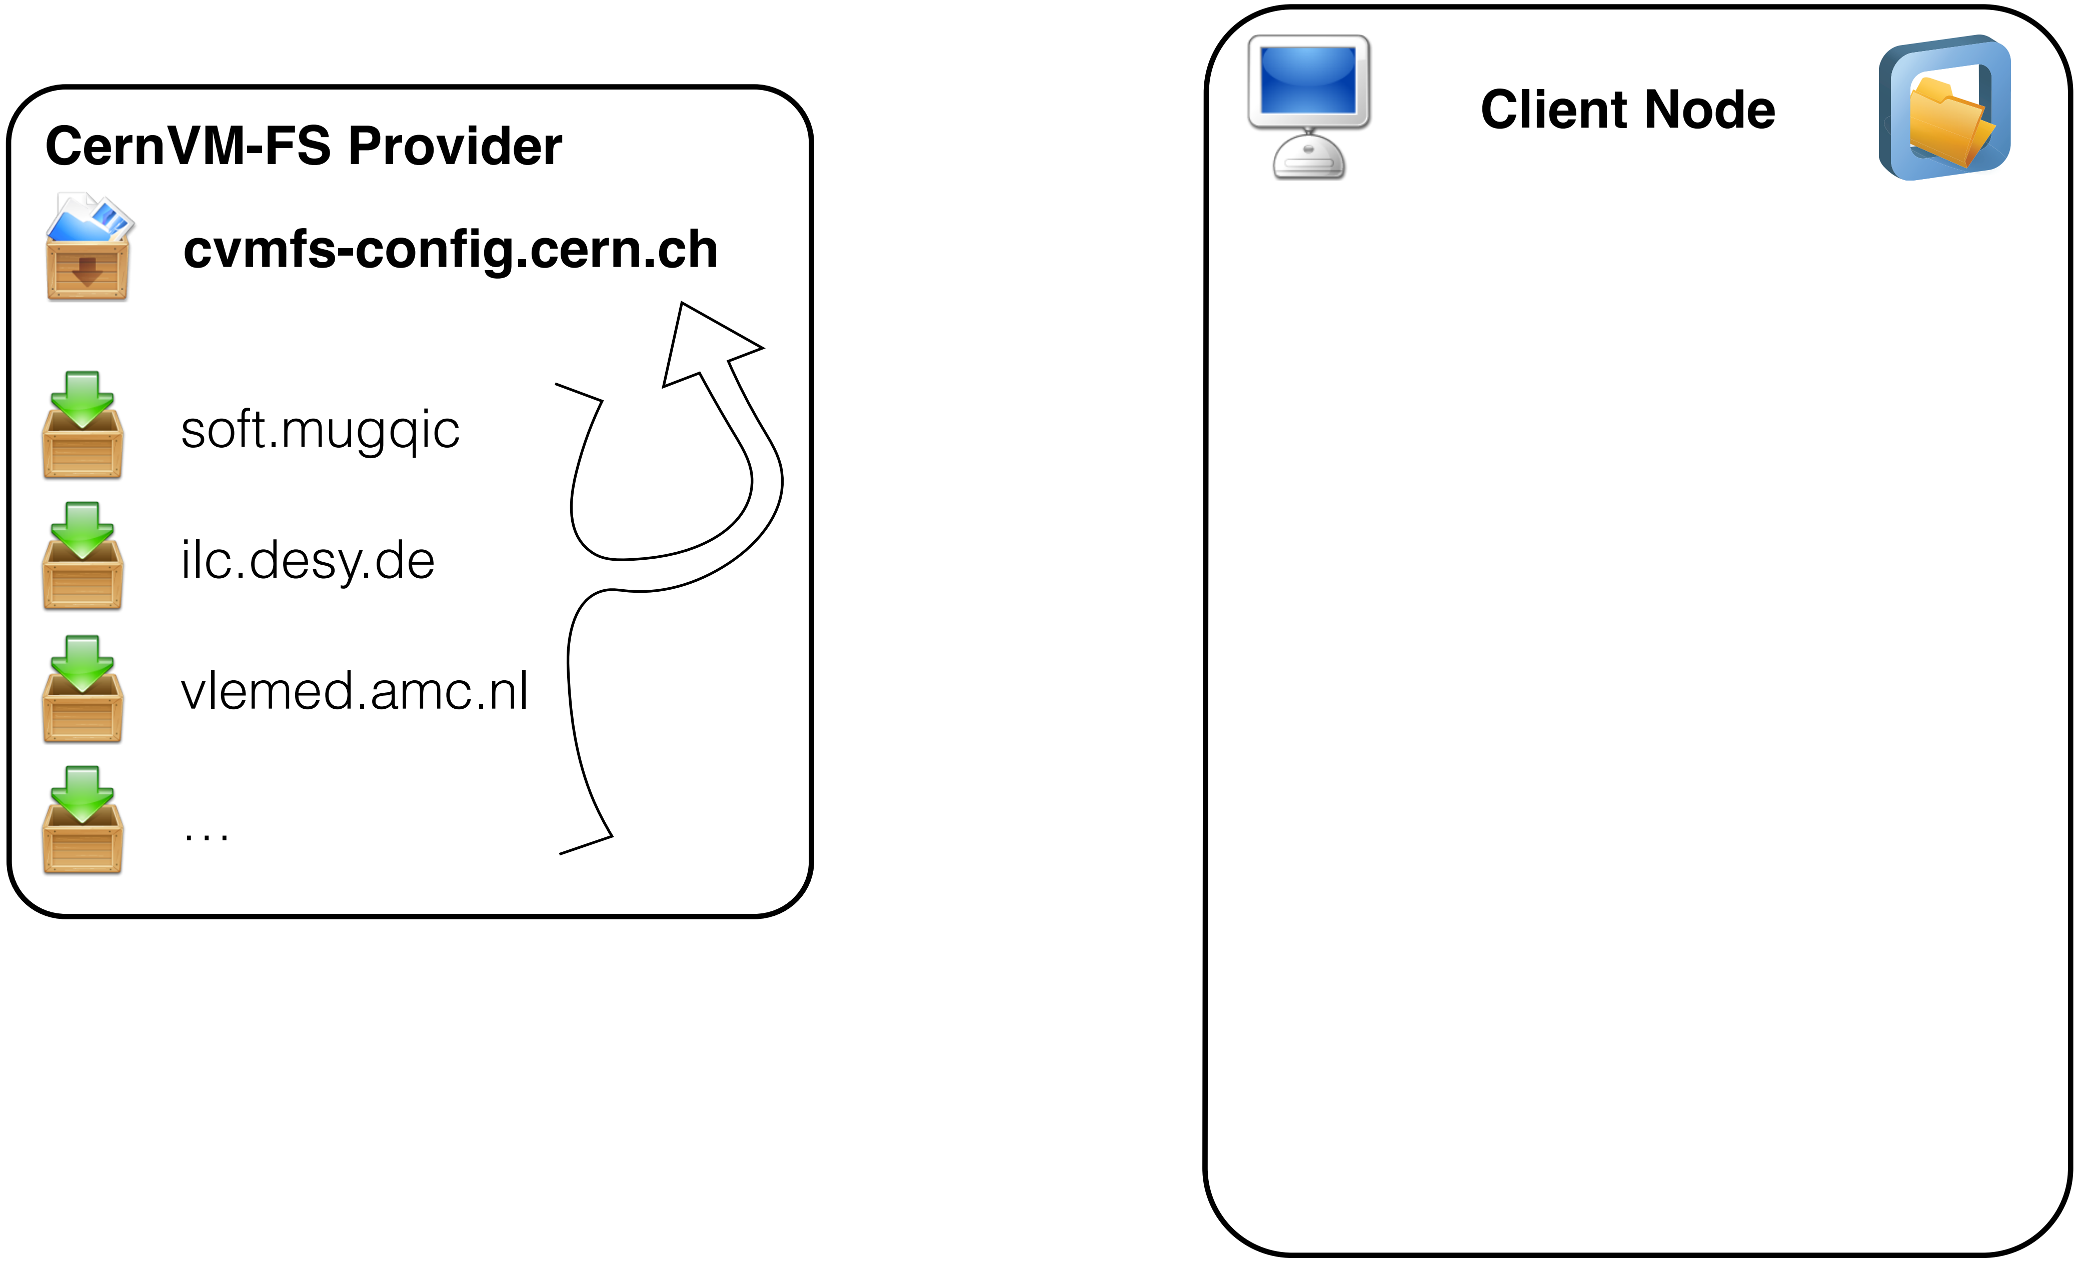

CernVM-FS Provider

cvmfs-config.cern.ch

soft.mugqic

ilc.desy.de

vlemed.amc.nl

…

Client Node

CernVM-FS Configuration Repository

- Bootstrap repository for CernVM-FS clients
- Maintains configuration baseline and keys centrally

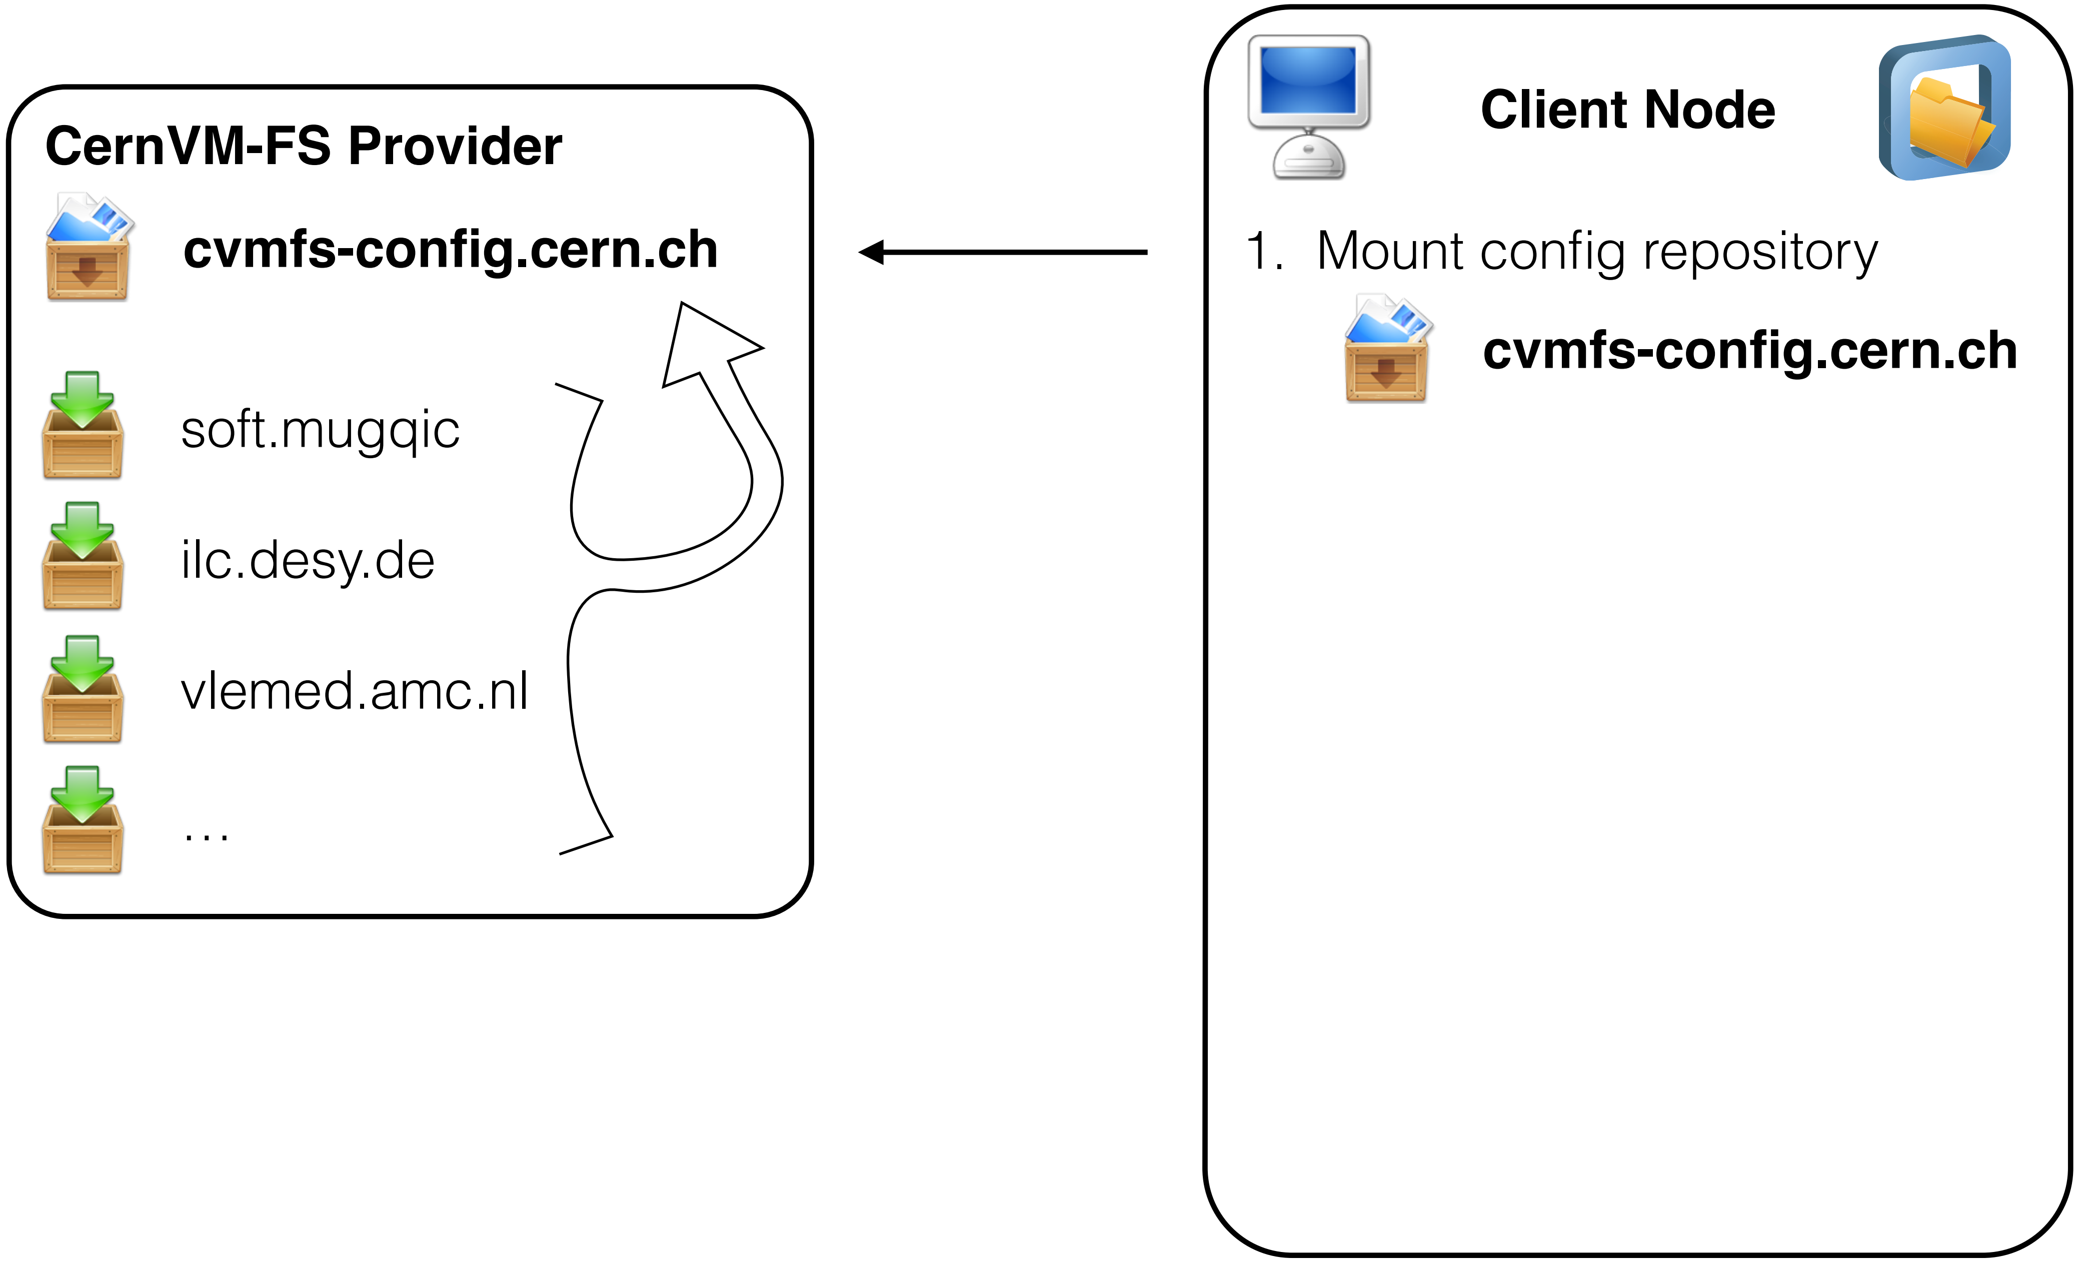

CernVM-FS Provider

cvmfs-config.cern.ch

soft.mugqic

ilc.desy.de

vlemed.amc.nl

…

Client Node

1. Mount config repository

cvmfs-config.cern.ch

CernVM-FS Configuration Repository

- Bootstrap repository for CernVM-FS clients
- Maintains configuration baseline and keys centrally

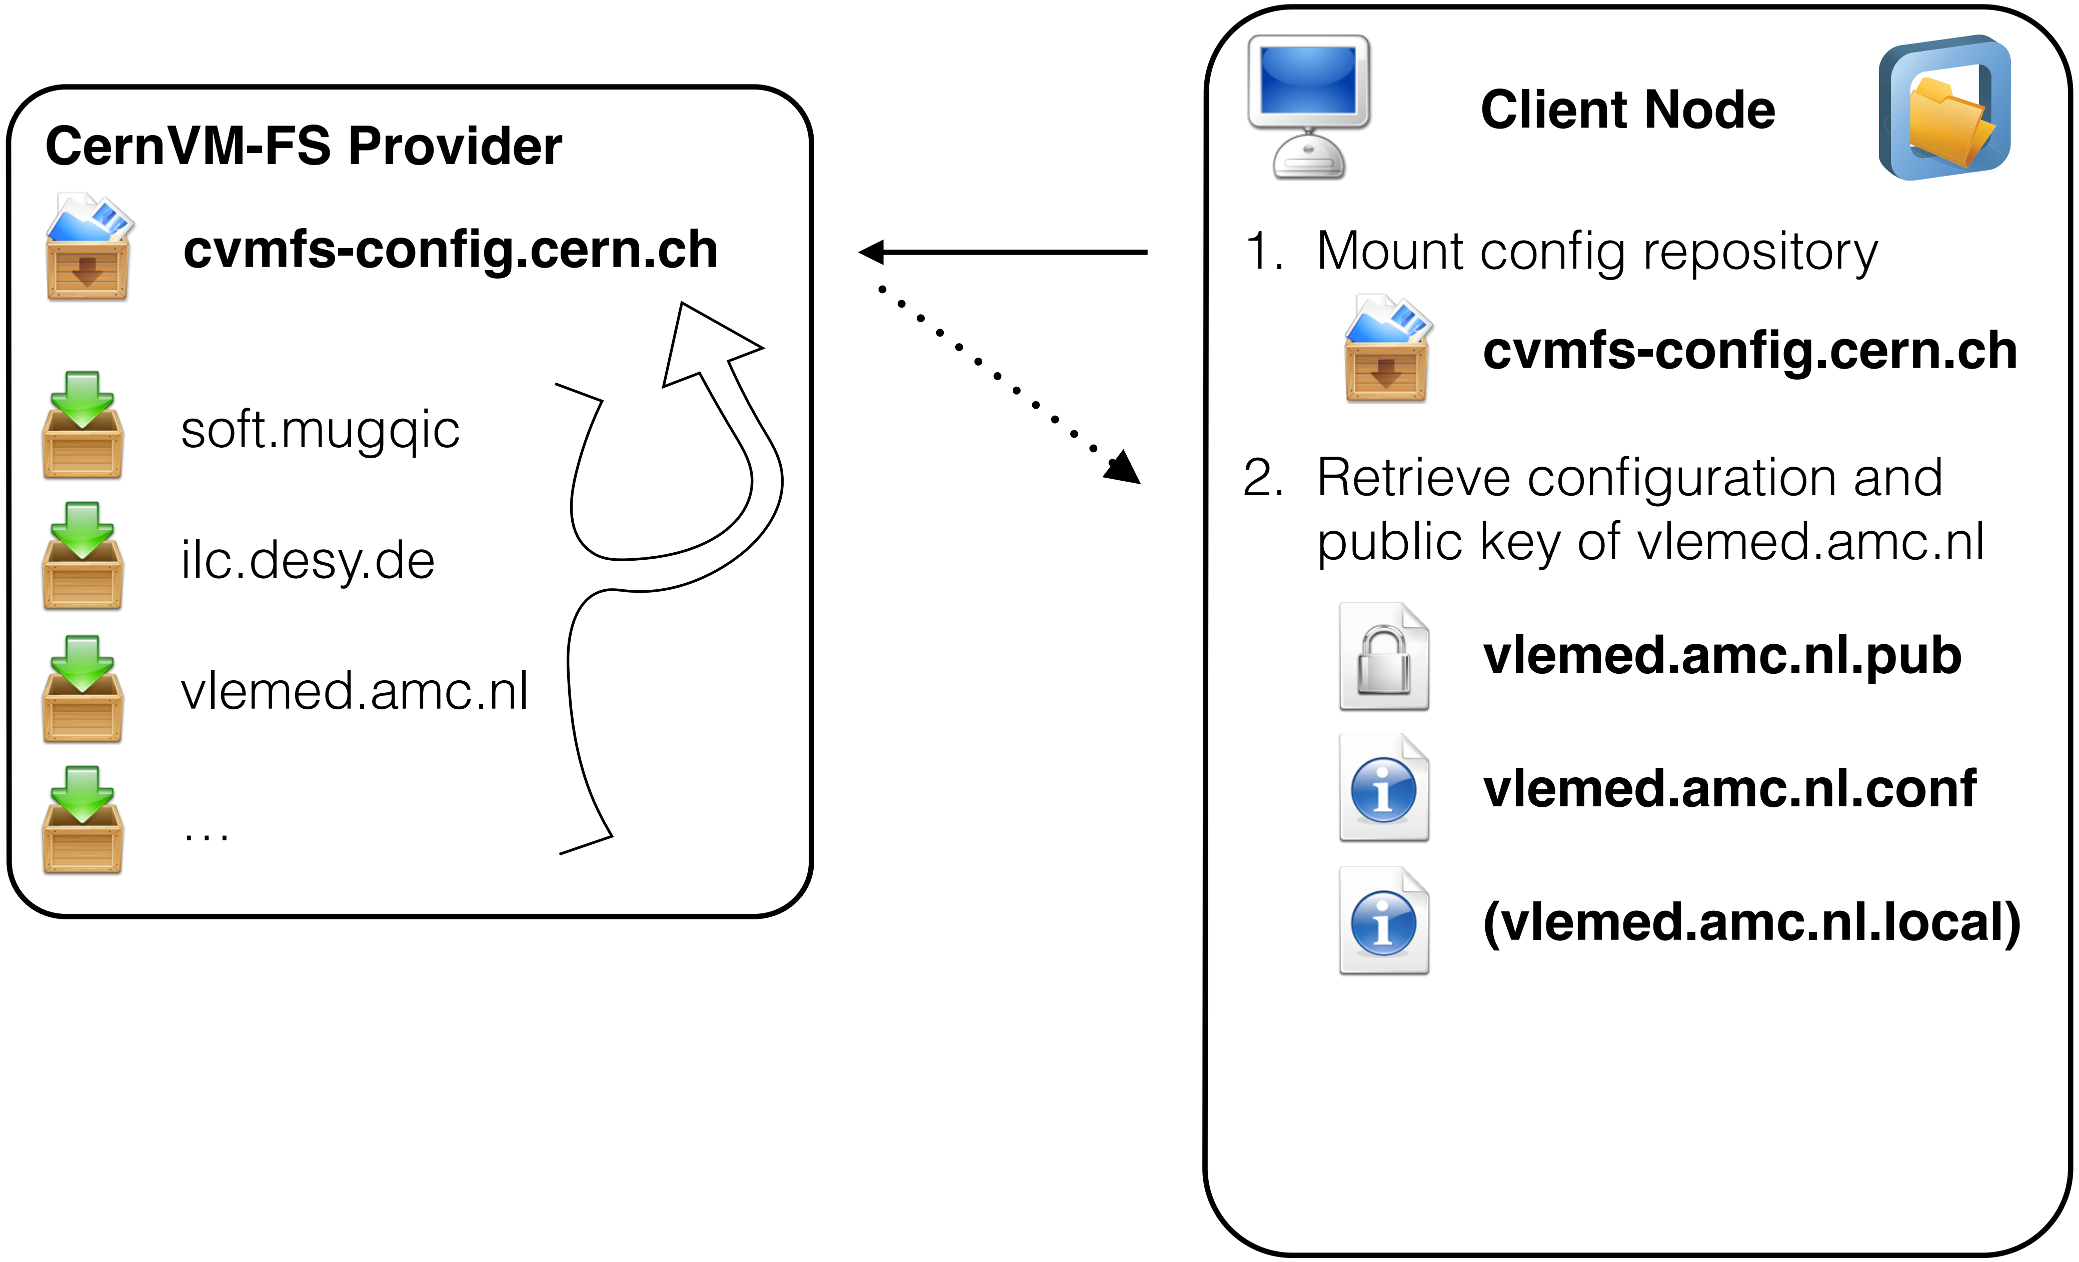

CernVM-FS Provider

cvmfs-config.cern.ch

soft.mugqic

ilc.desy.de

vlemed.amc.nl

…

Client Node

1. Mount config repository

   cvmfs-config.cern.ch

2. Retrieve configuration and public key of vlemed.amc.nl

   vlemed.amc.nl.pub

   vlemed.amc.nl.conf

   (vlemed.amc.nl.local)

CernVM-FS Configuration Repository

- Bootstrap repository for CernVM-FS clients
- Maintains configuration baseline and keys centrally

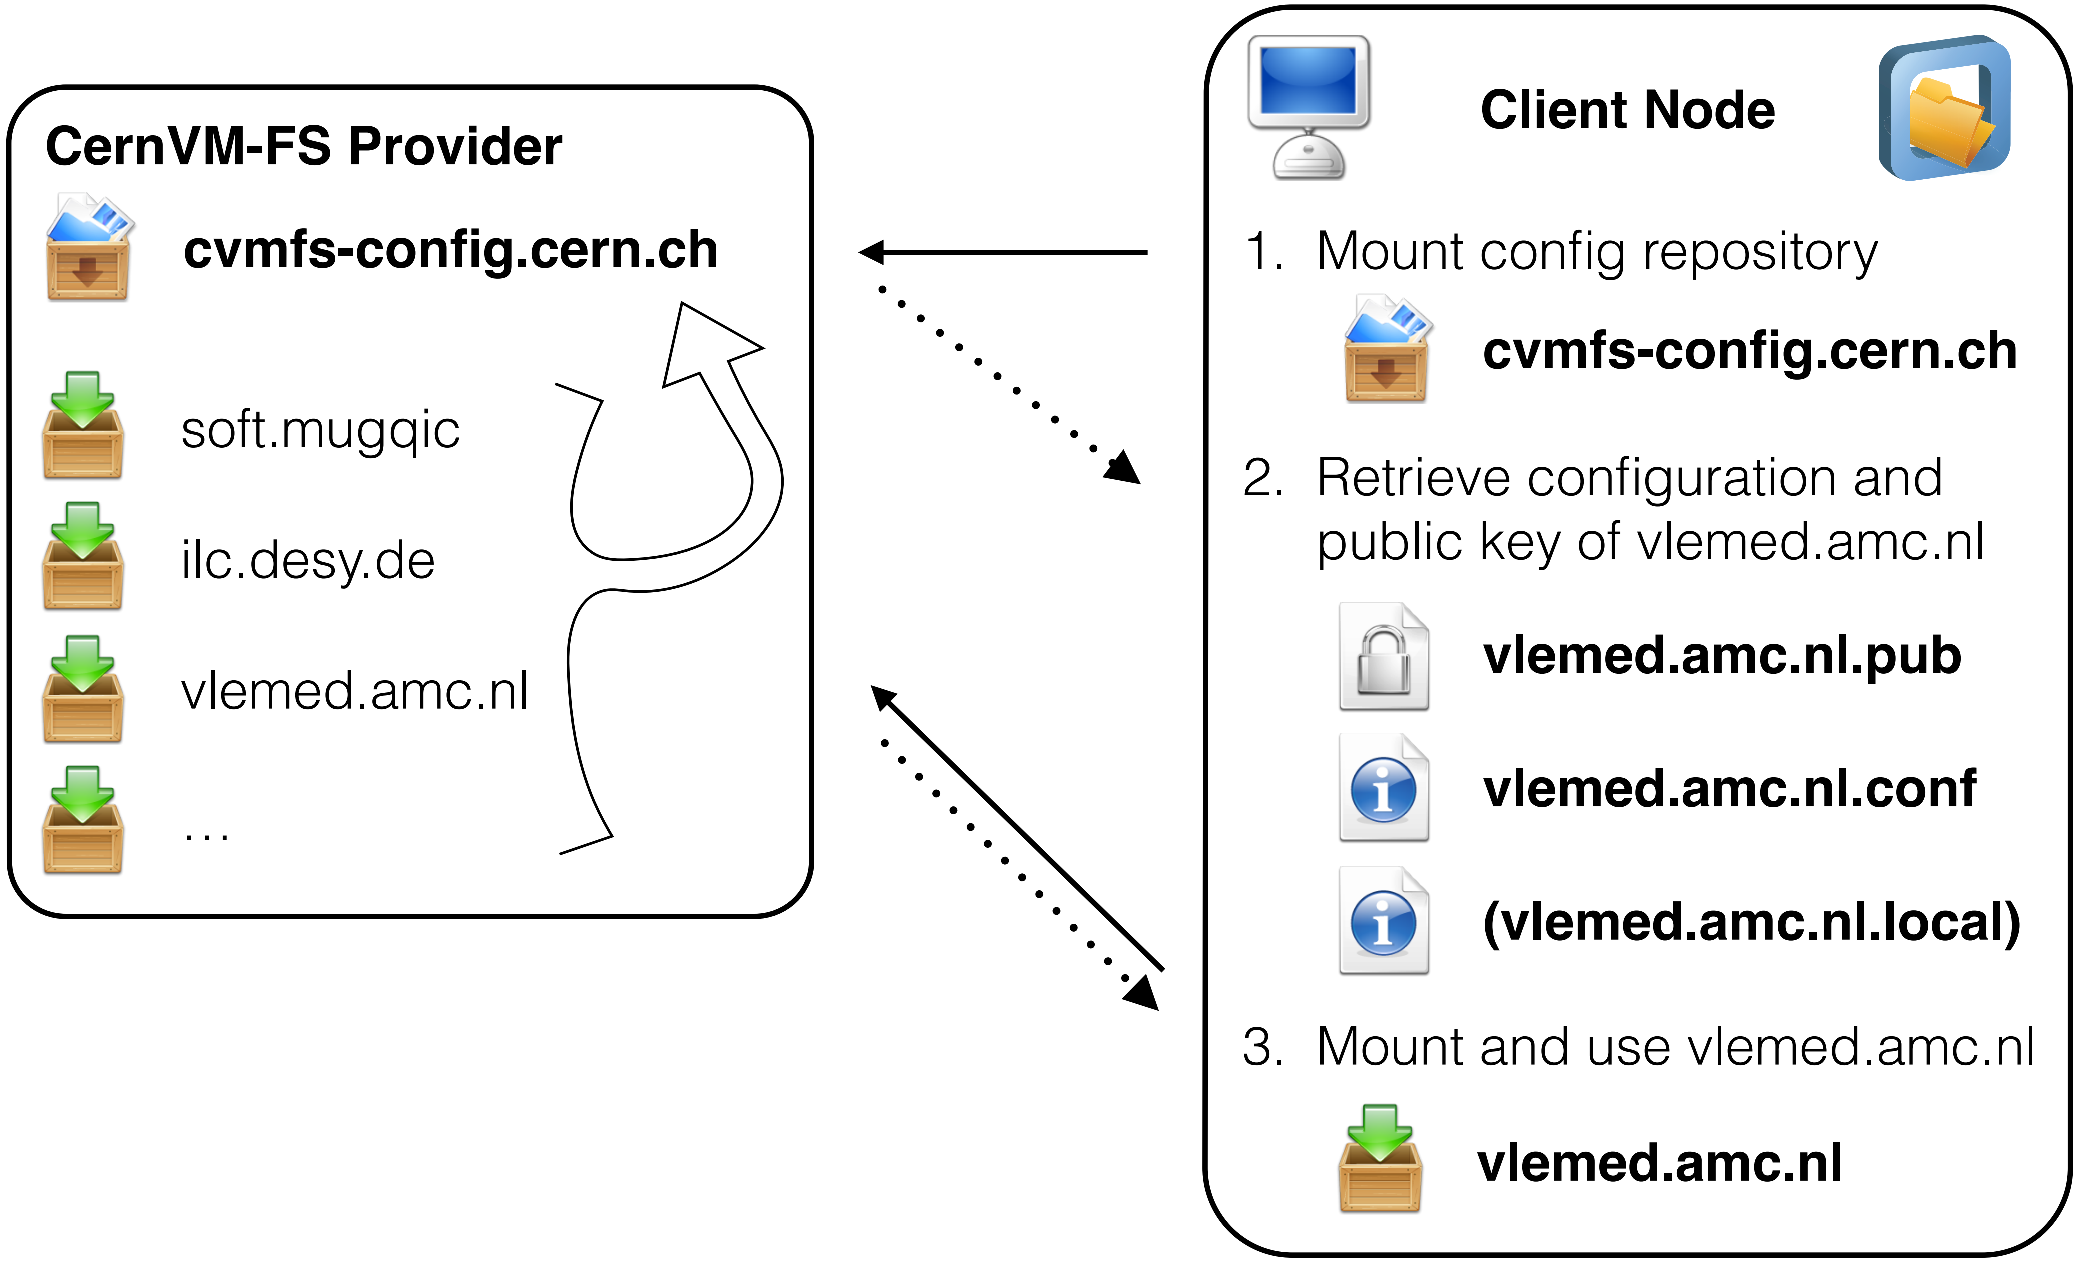

CernVM-FS Provider

cvmfs-config.cern.ch

soft.mugqic

ilc.desy.de

vlemed.amc.nl

…

Client Node

1. Mount config repository

   cvmfs-config.cern.ch

2. Retrieve configuration and public key of vlemed.amc.nl

   vlemed.amc.nl.pub

   vlemed.amc.nl.conf

   (vlemed.amc.nl.local)

3. Mount and use vlemed.amc.nl

   vlemed.amc.nl

GeoIP-based Stratum1 Ordering

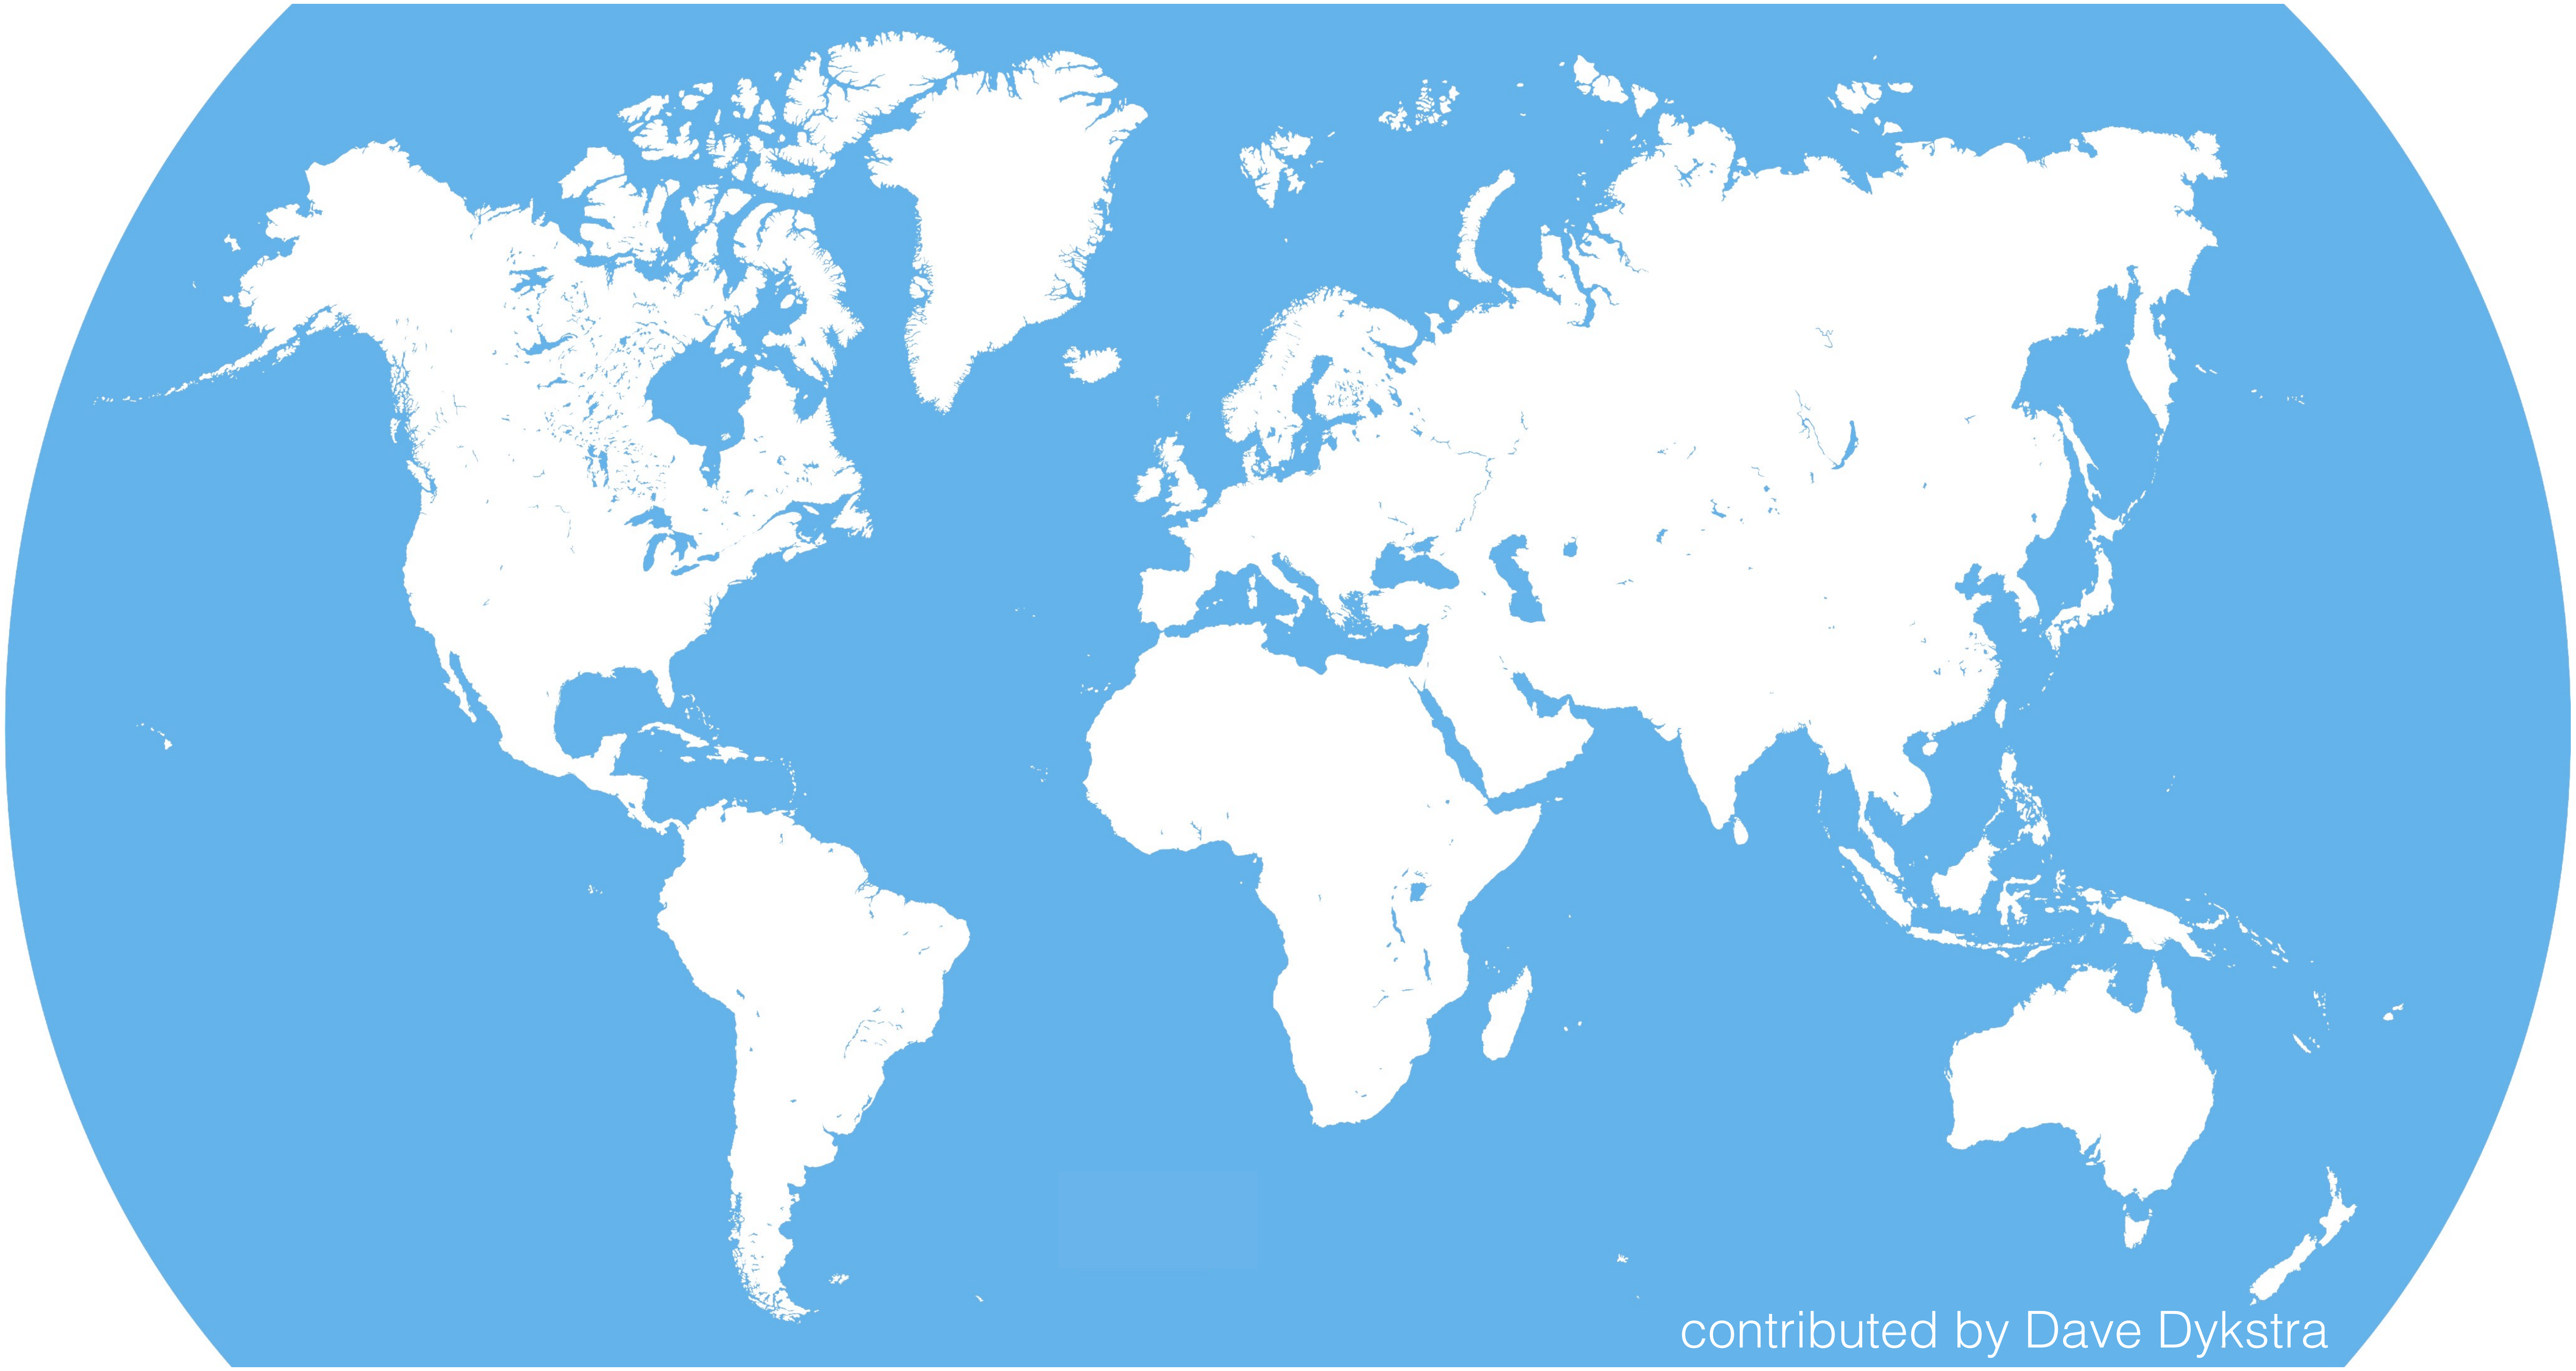

contributed by Dave Dykstra

GeoIP-based Stratum1 Ordering

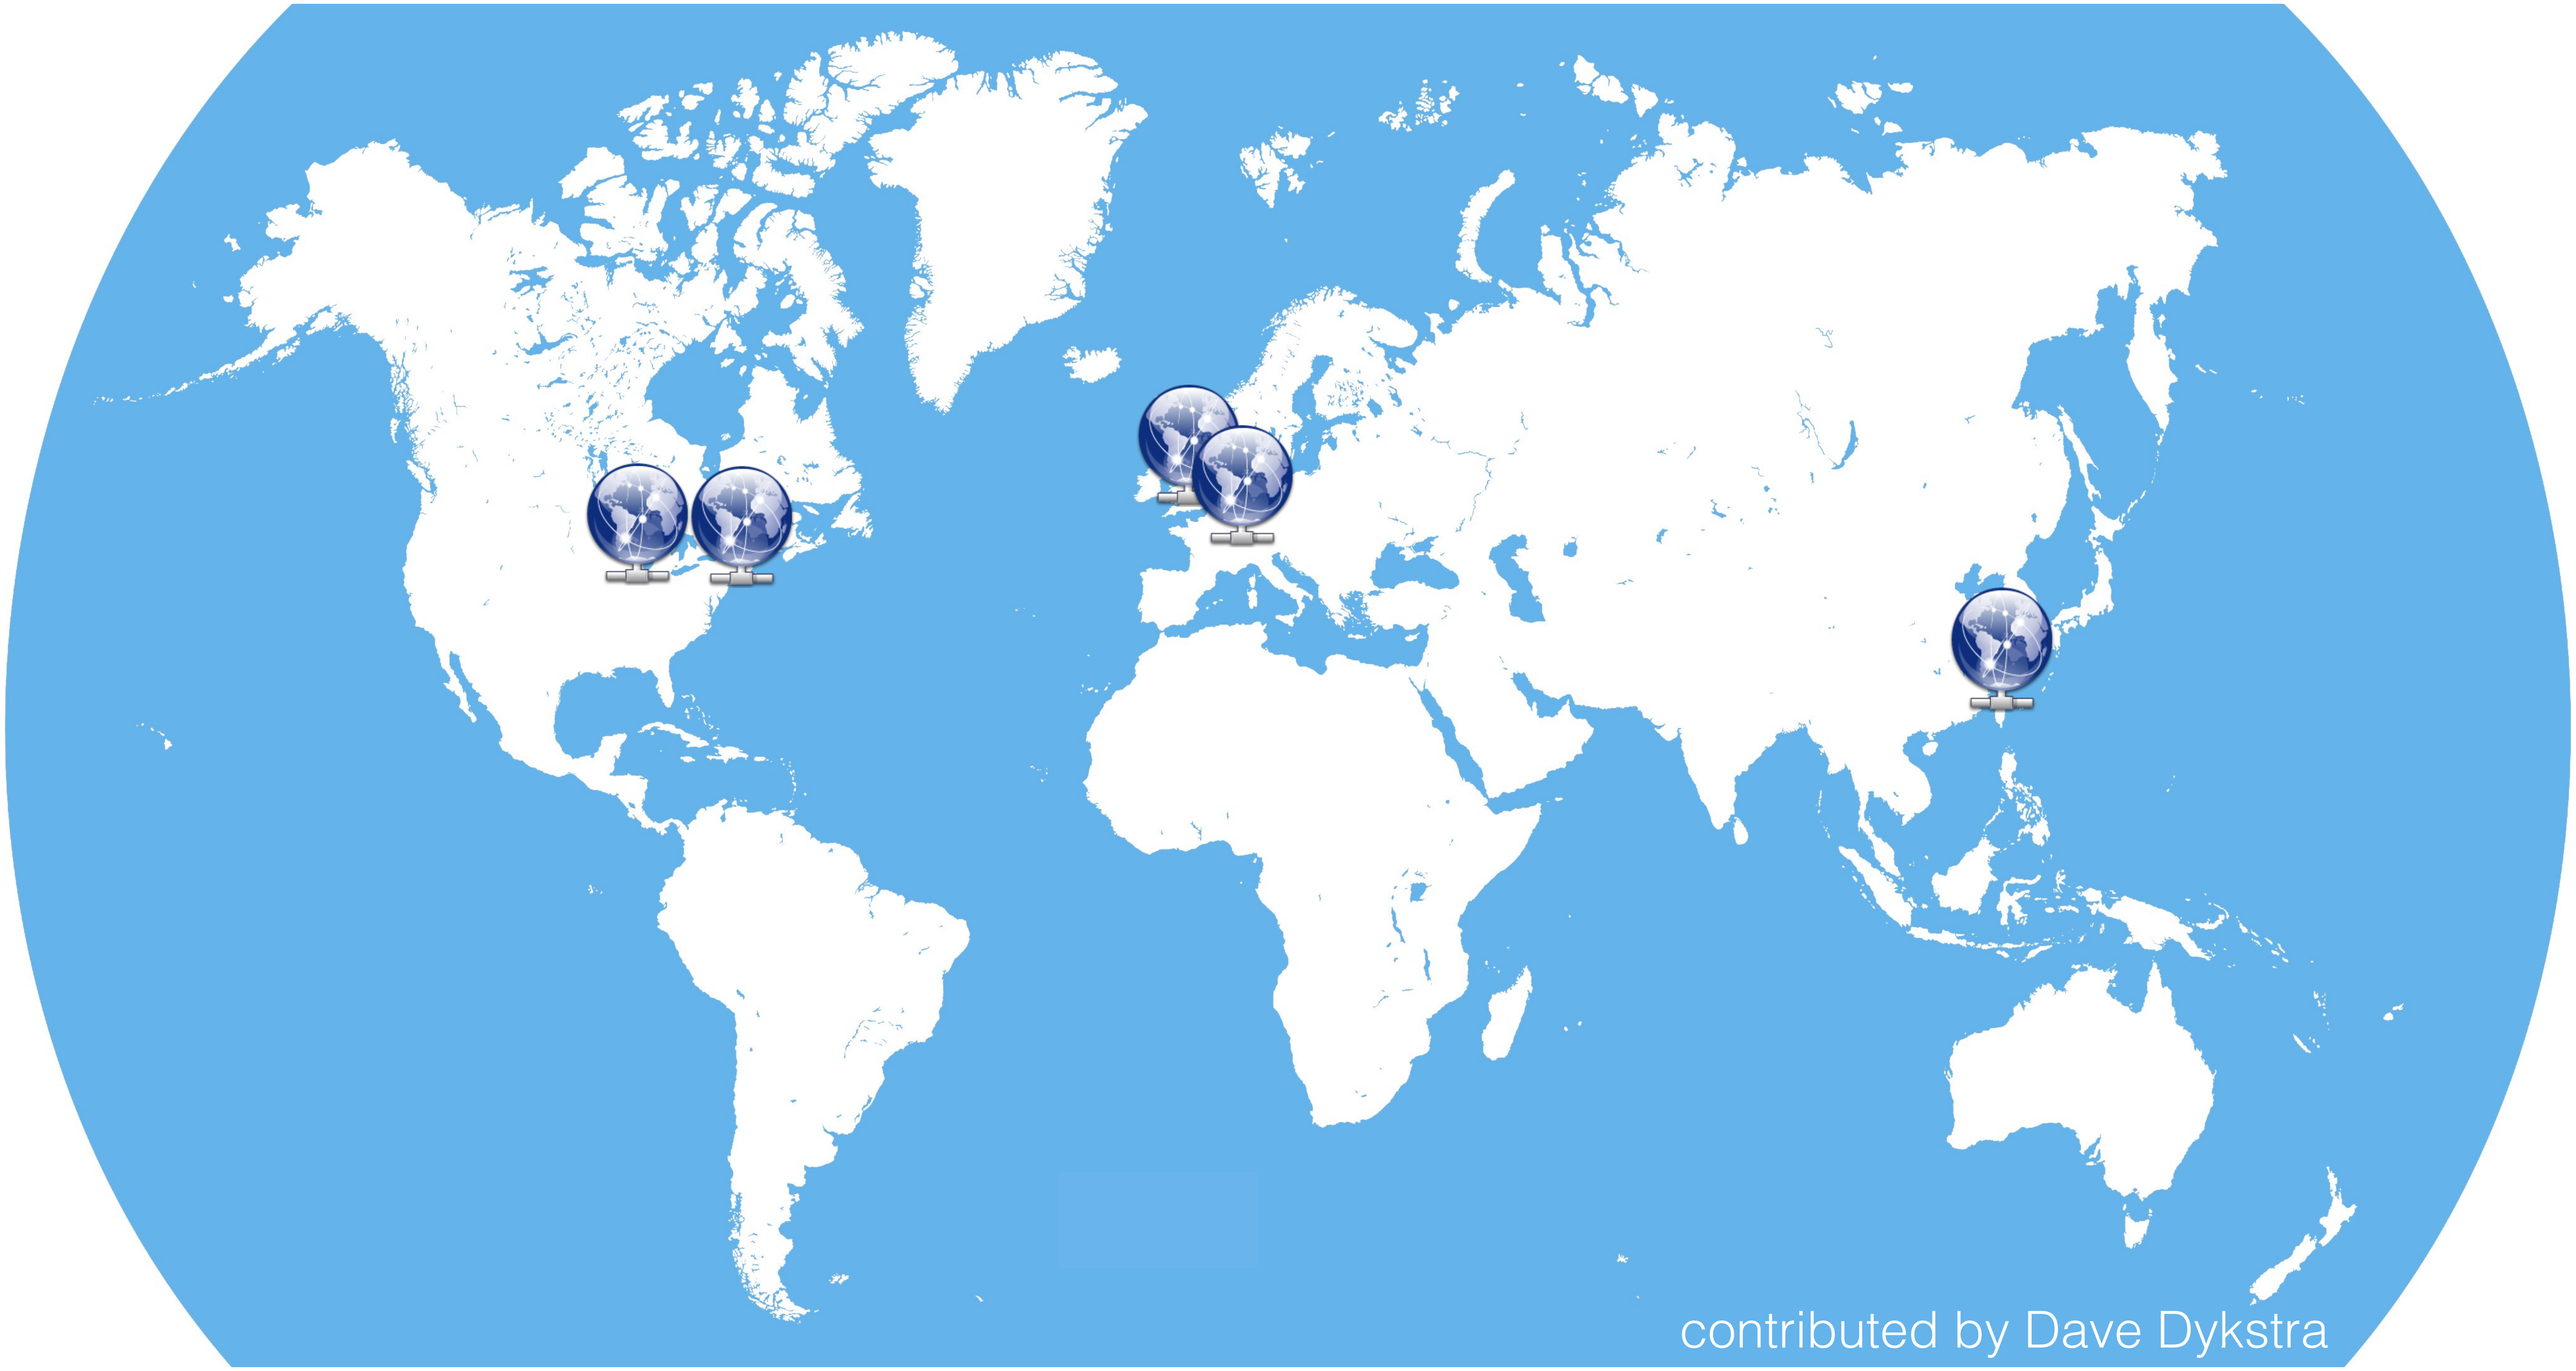

contributed by Dave Dykstra

GeoIP-based Stratum1 Ordering

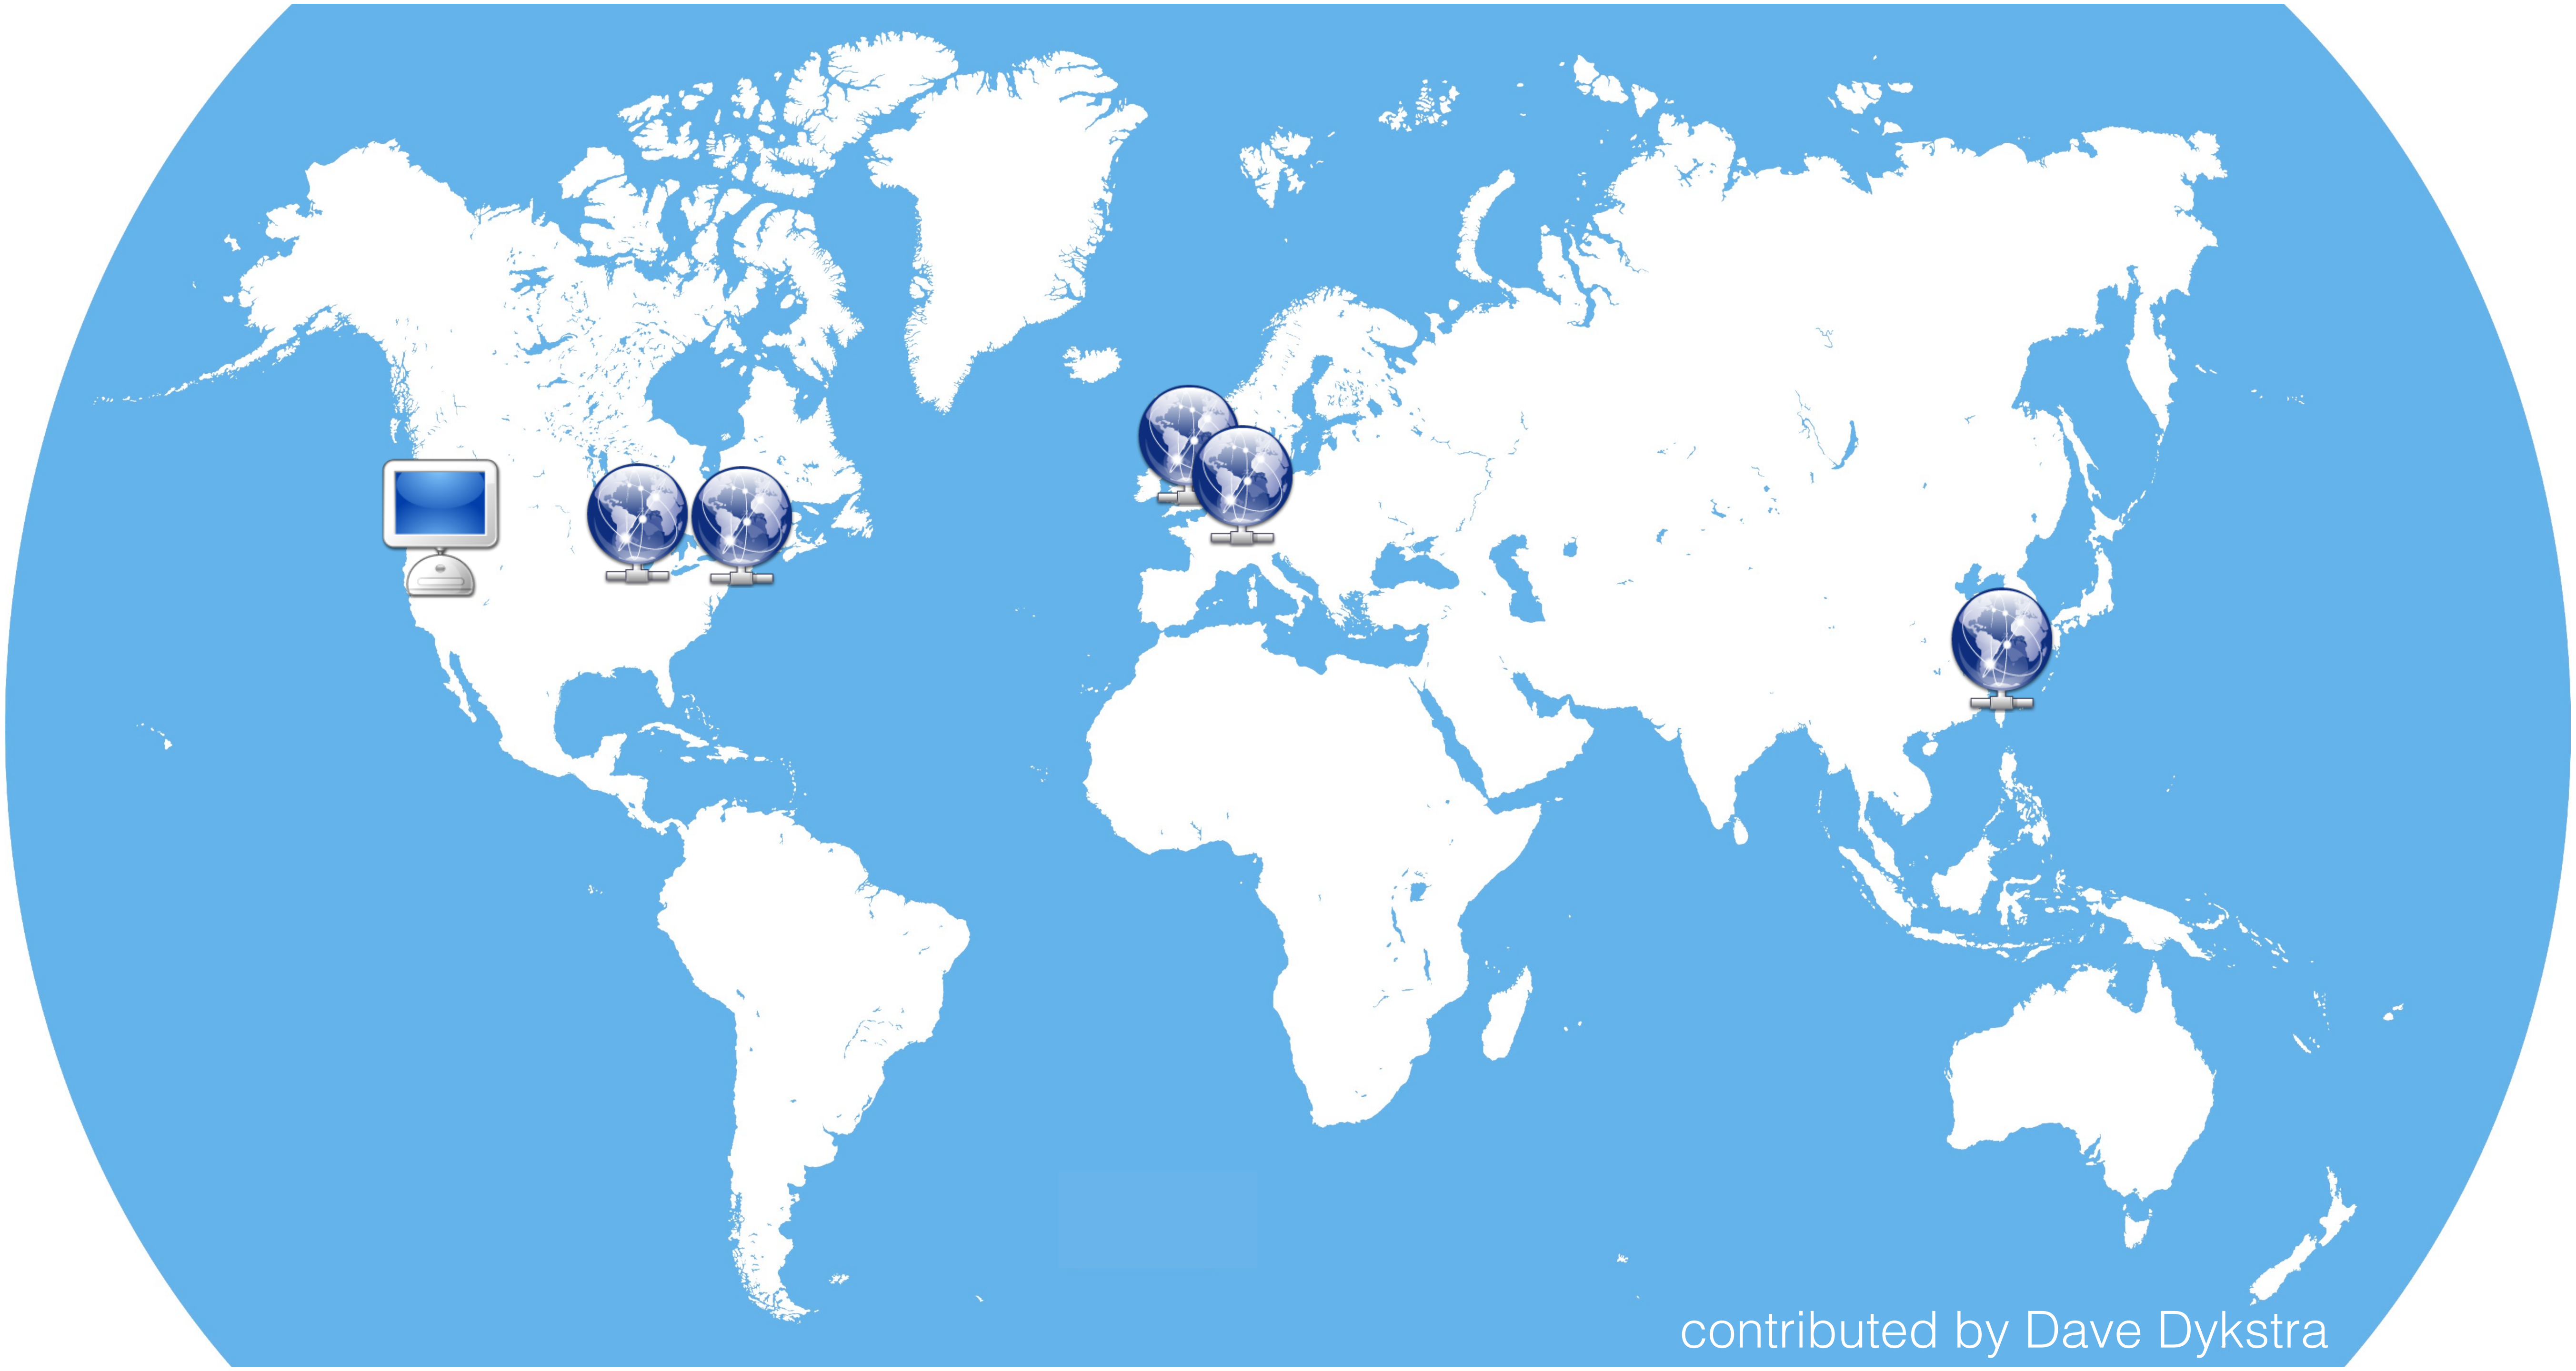

contributed by Dave Dykstra

GeoIP-based Stratum1 Ordering

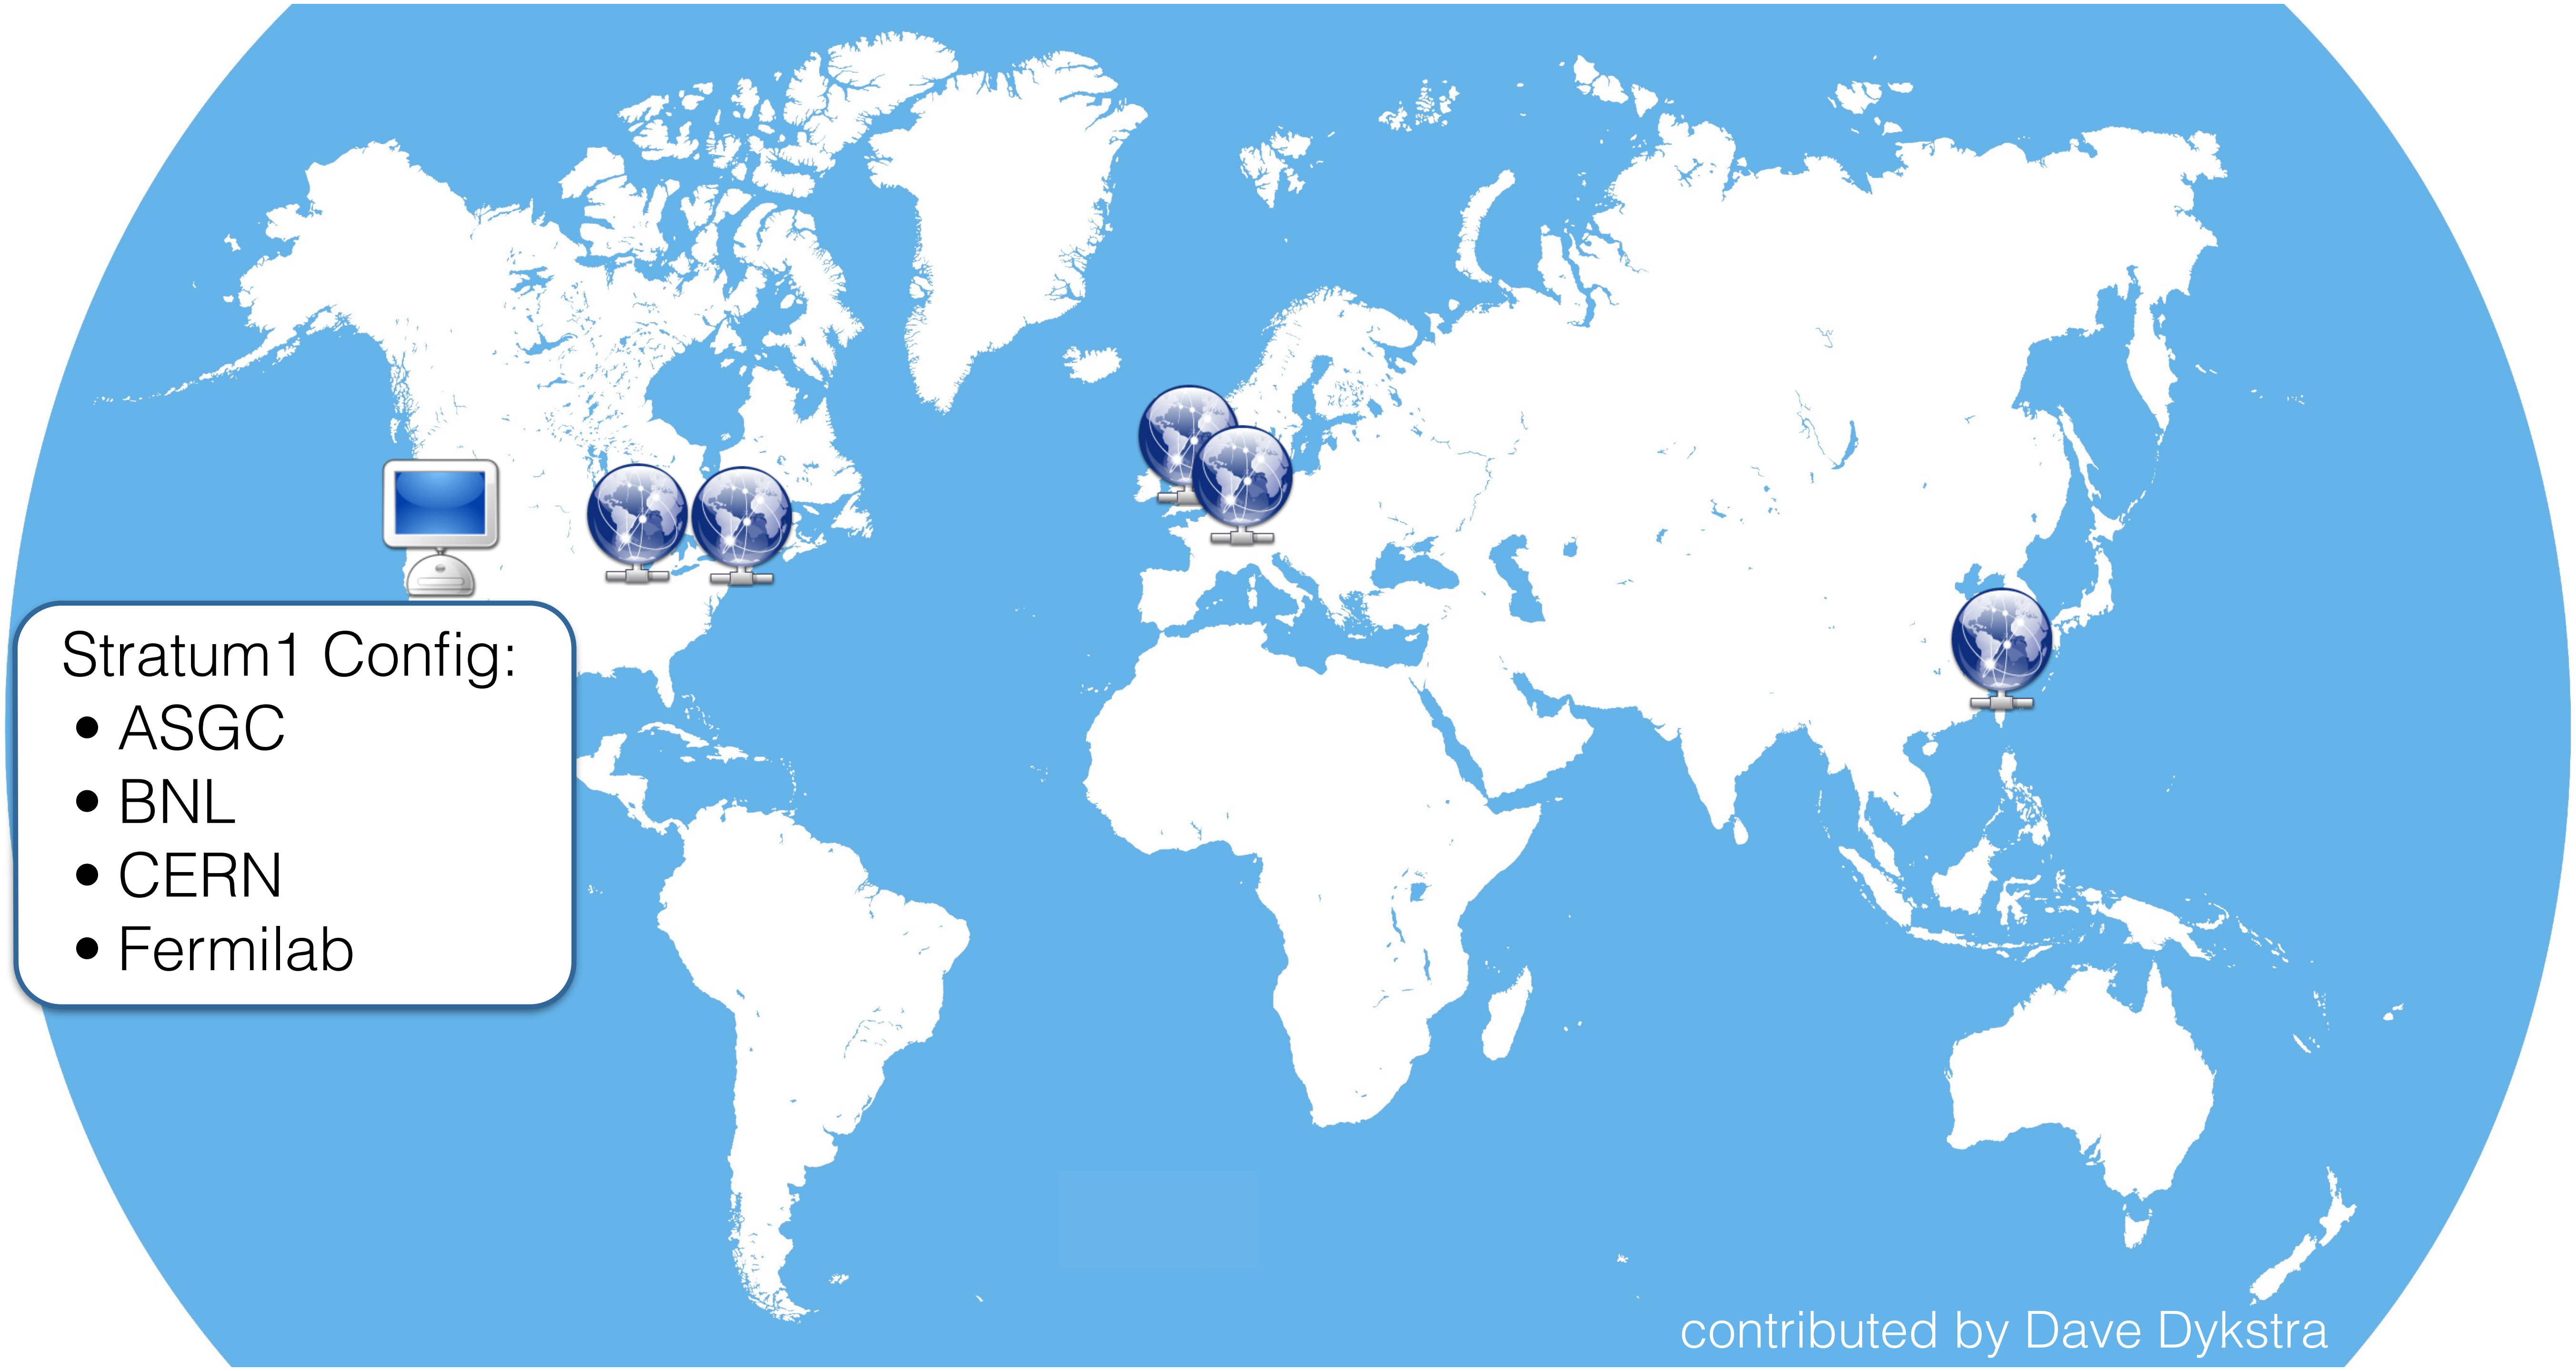

Stratum1 Config:

- ASGC
- BNL
- CERN
- Fermilab

contributed by Dave Dykstra

GeoIP-based Stratum1 Ordering

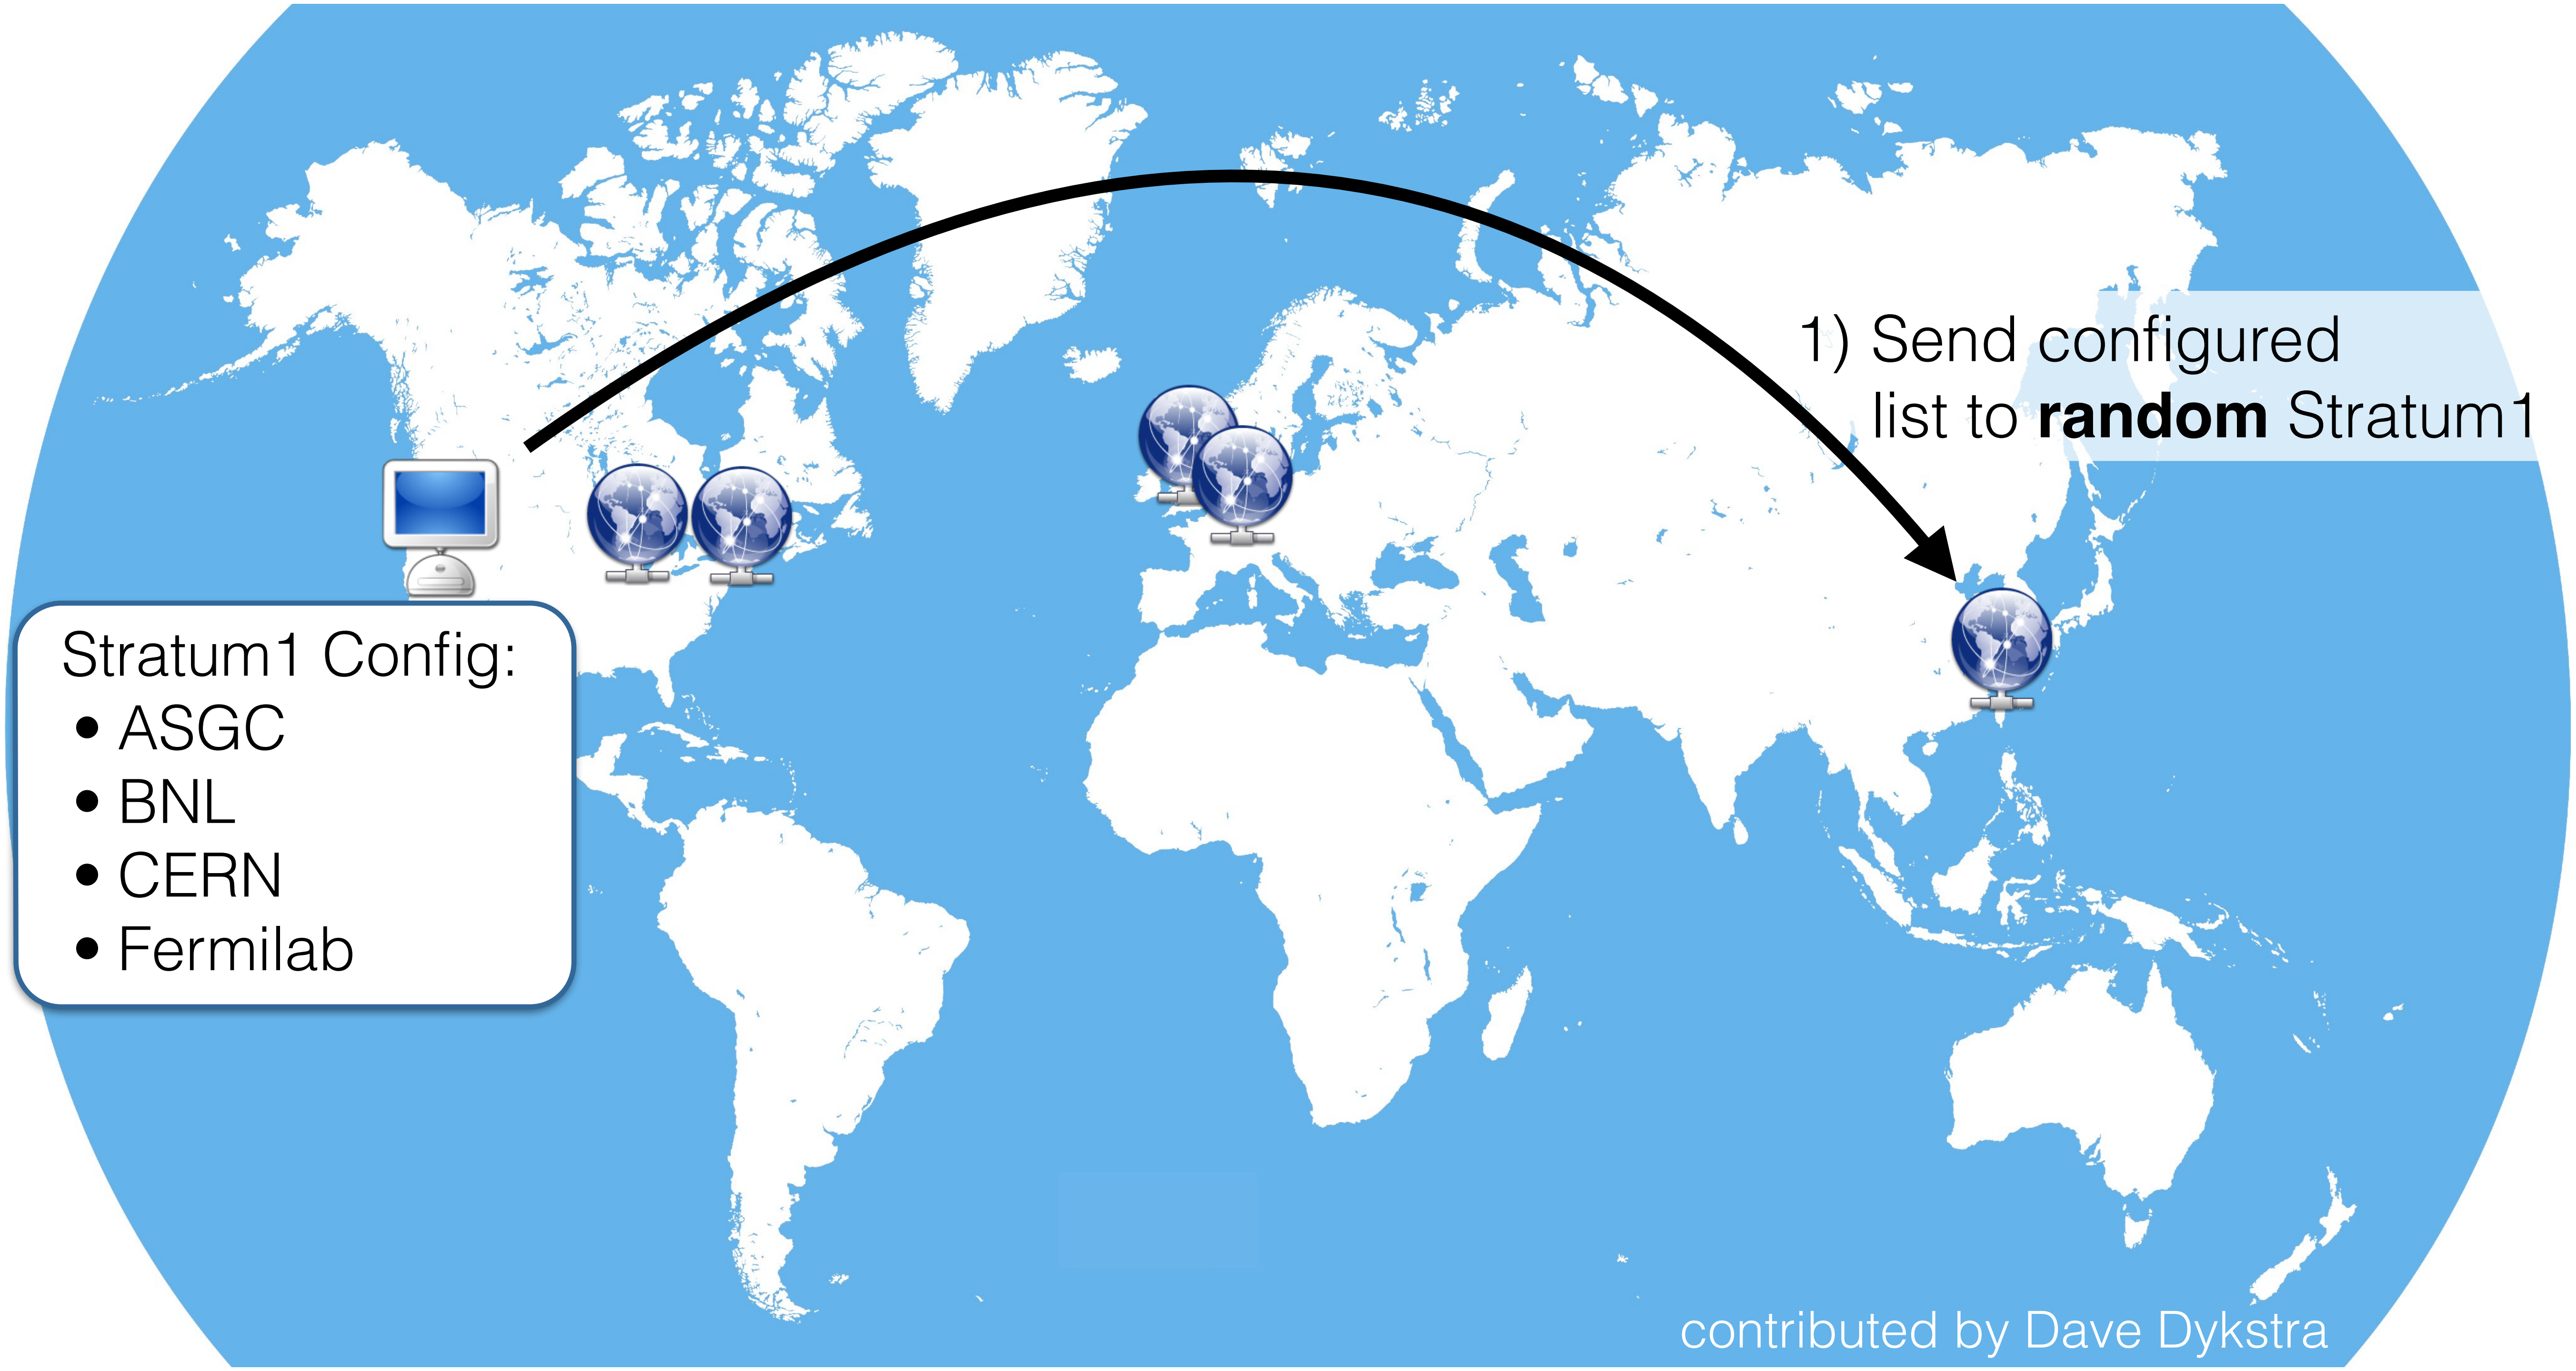

Stratum1 Config:

- ASGC
- BNL
- CERN
- Fermilab

1) Send configured list to **random** Stratum1

contributed by Dave Dykstra

GeoIP-based Stratum1 Ordering

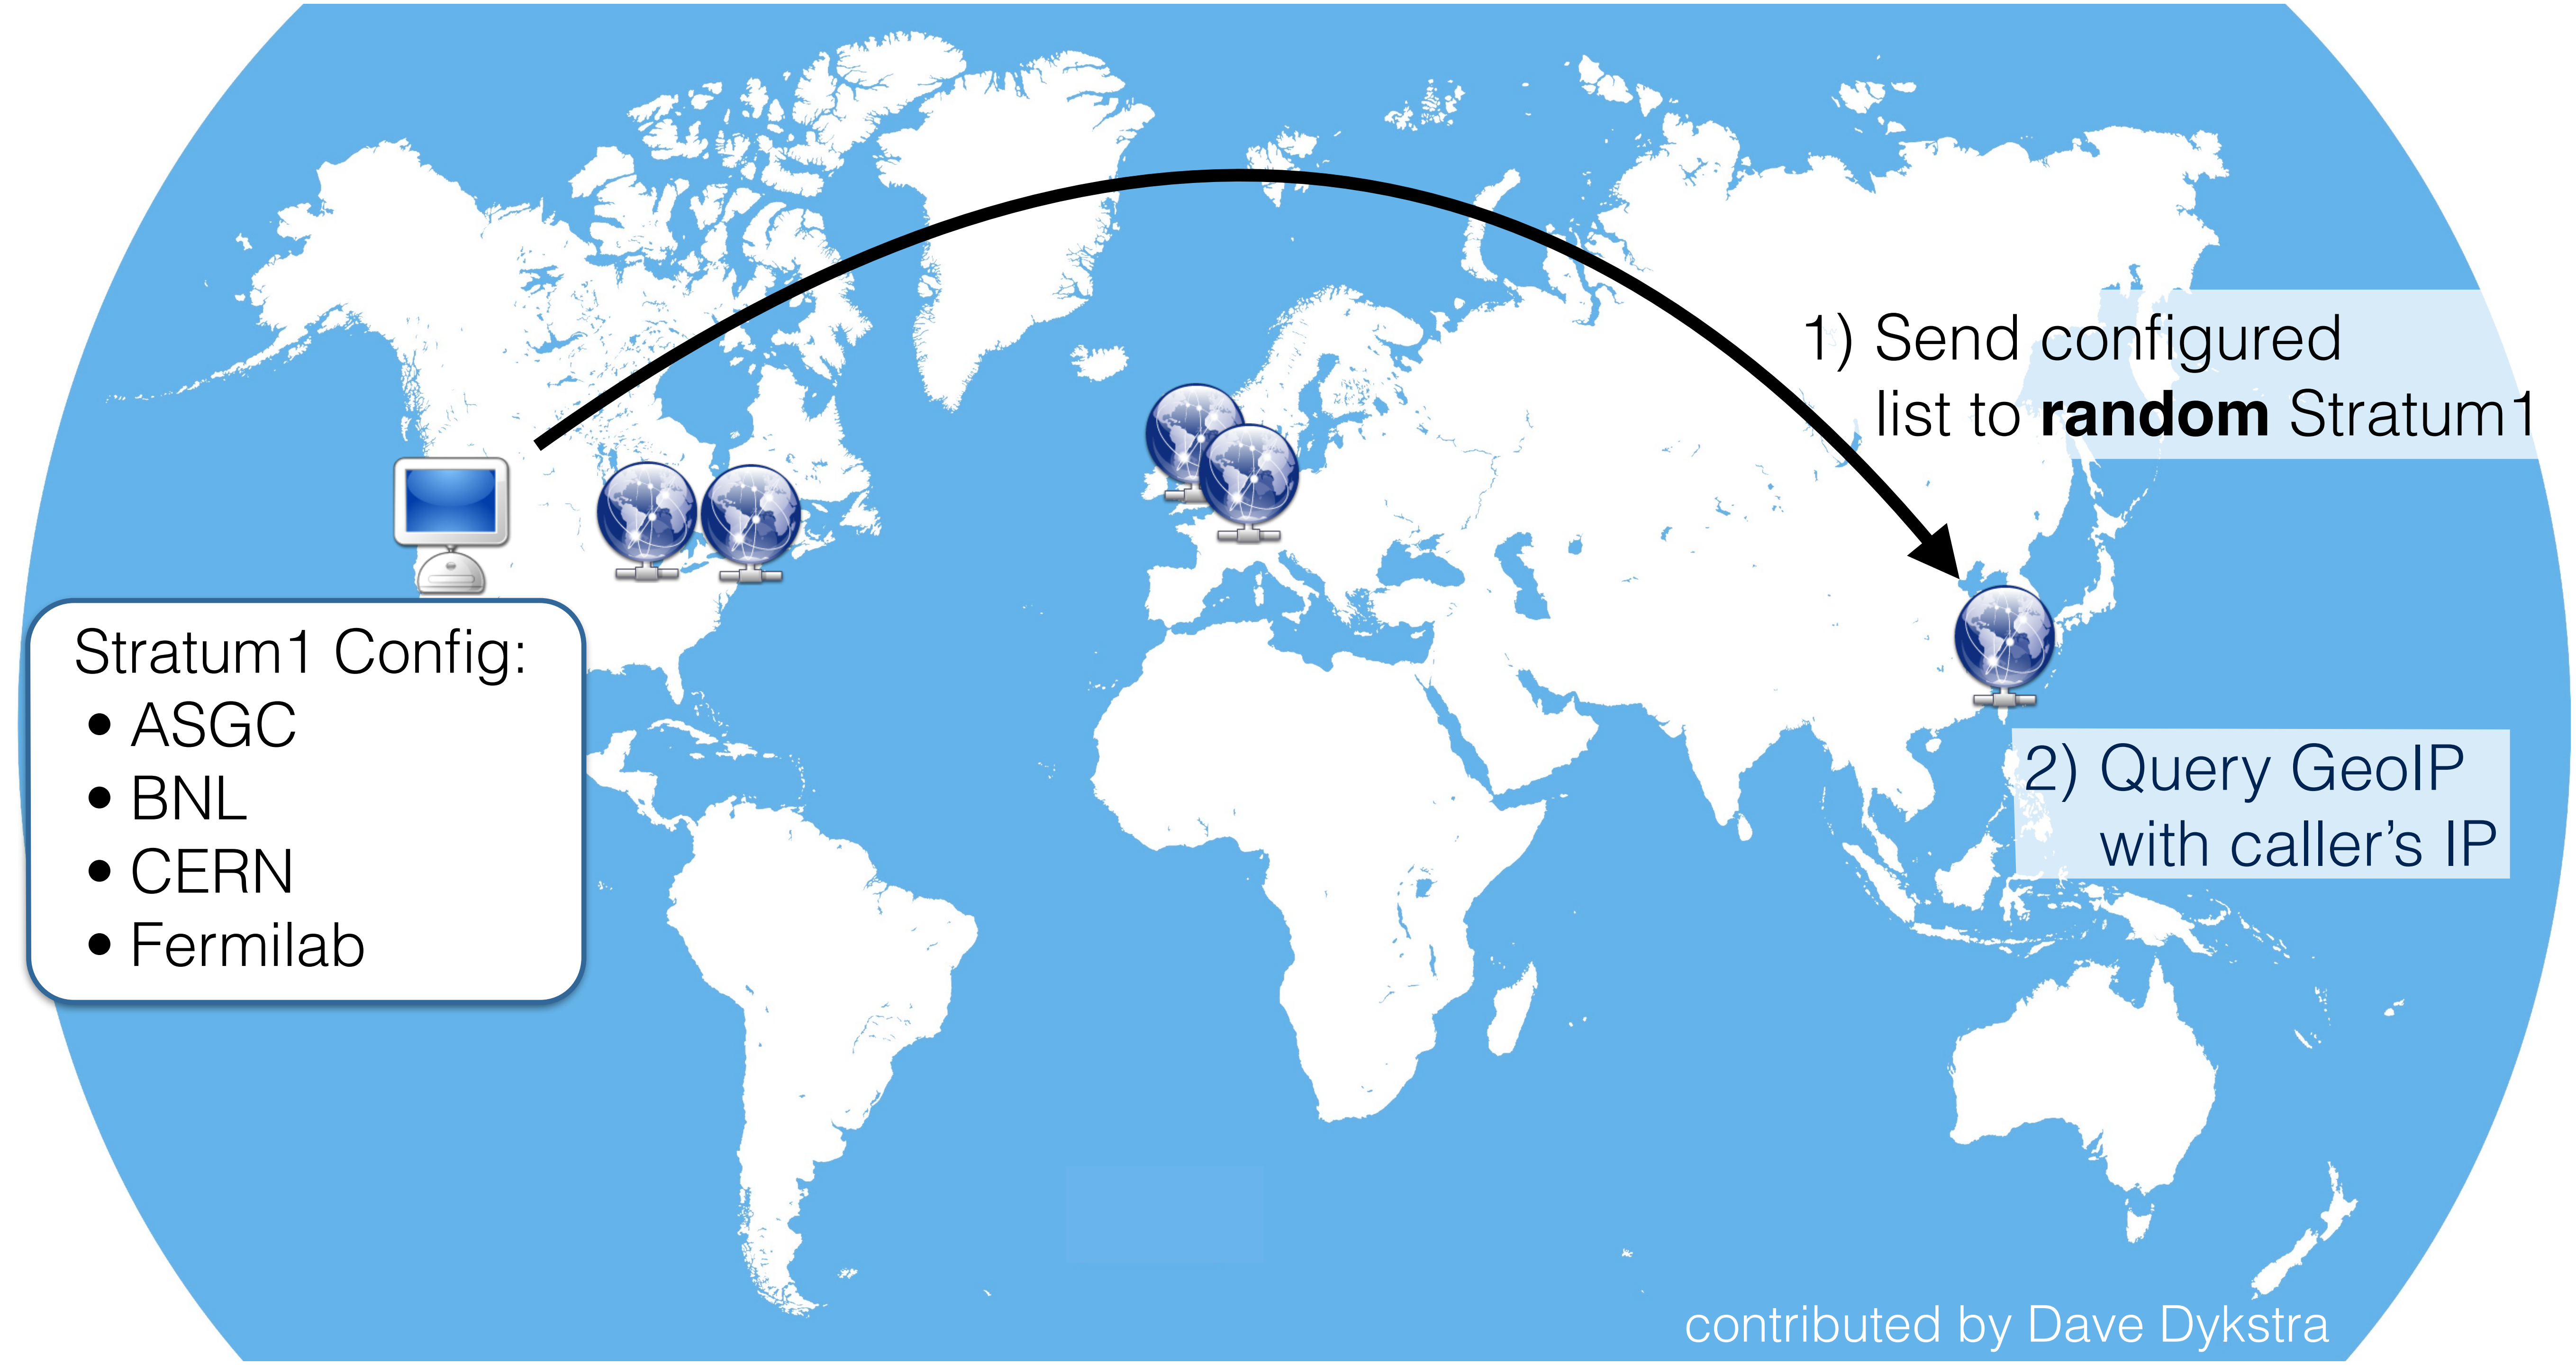

Stratum1 Config:

- ASGC
- BNL
- CERN
- Fermilab

1) Send configured list to **random** Stratum1

2) Query GeoIP with caller's IP

contributed by Dave Dykstra

GeoIP-based Stratum1 Ordering

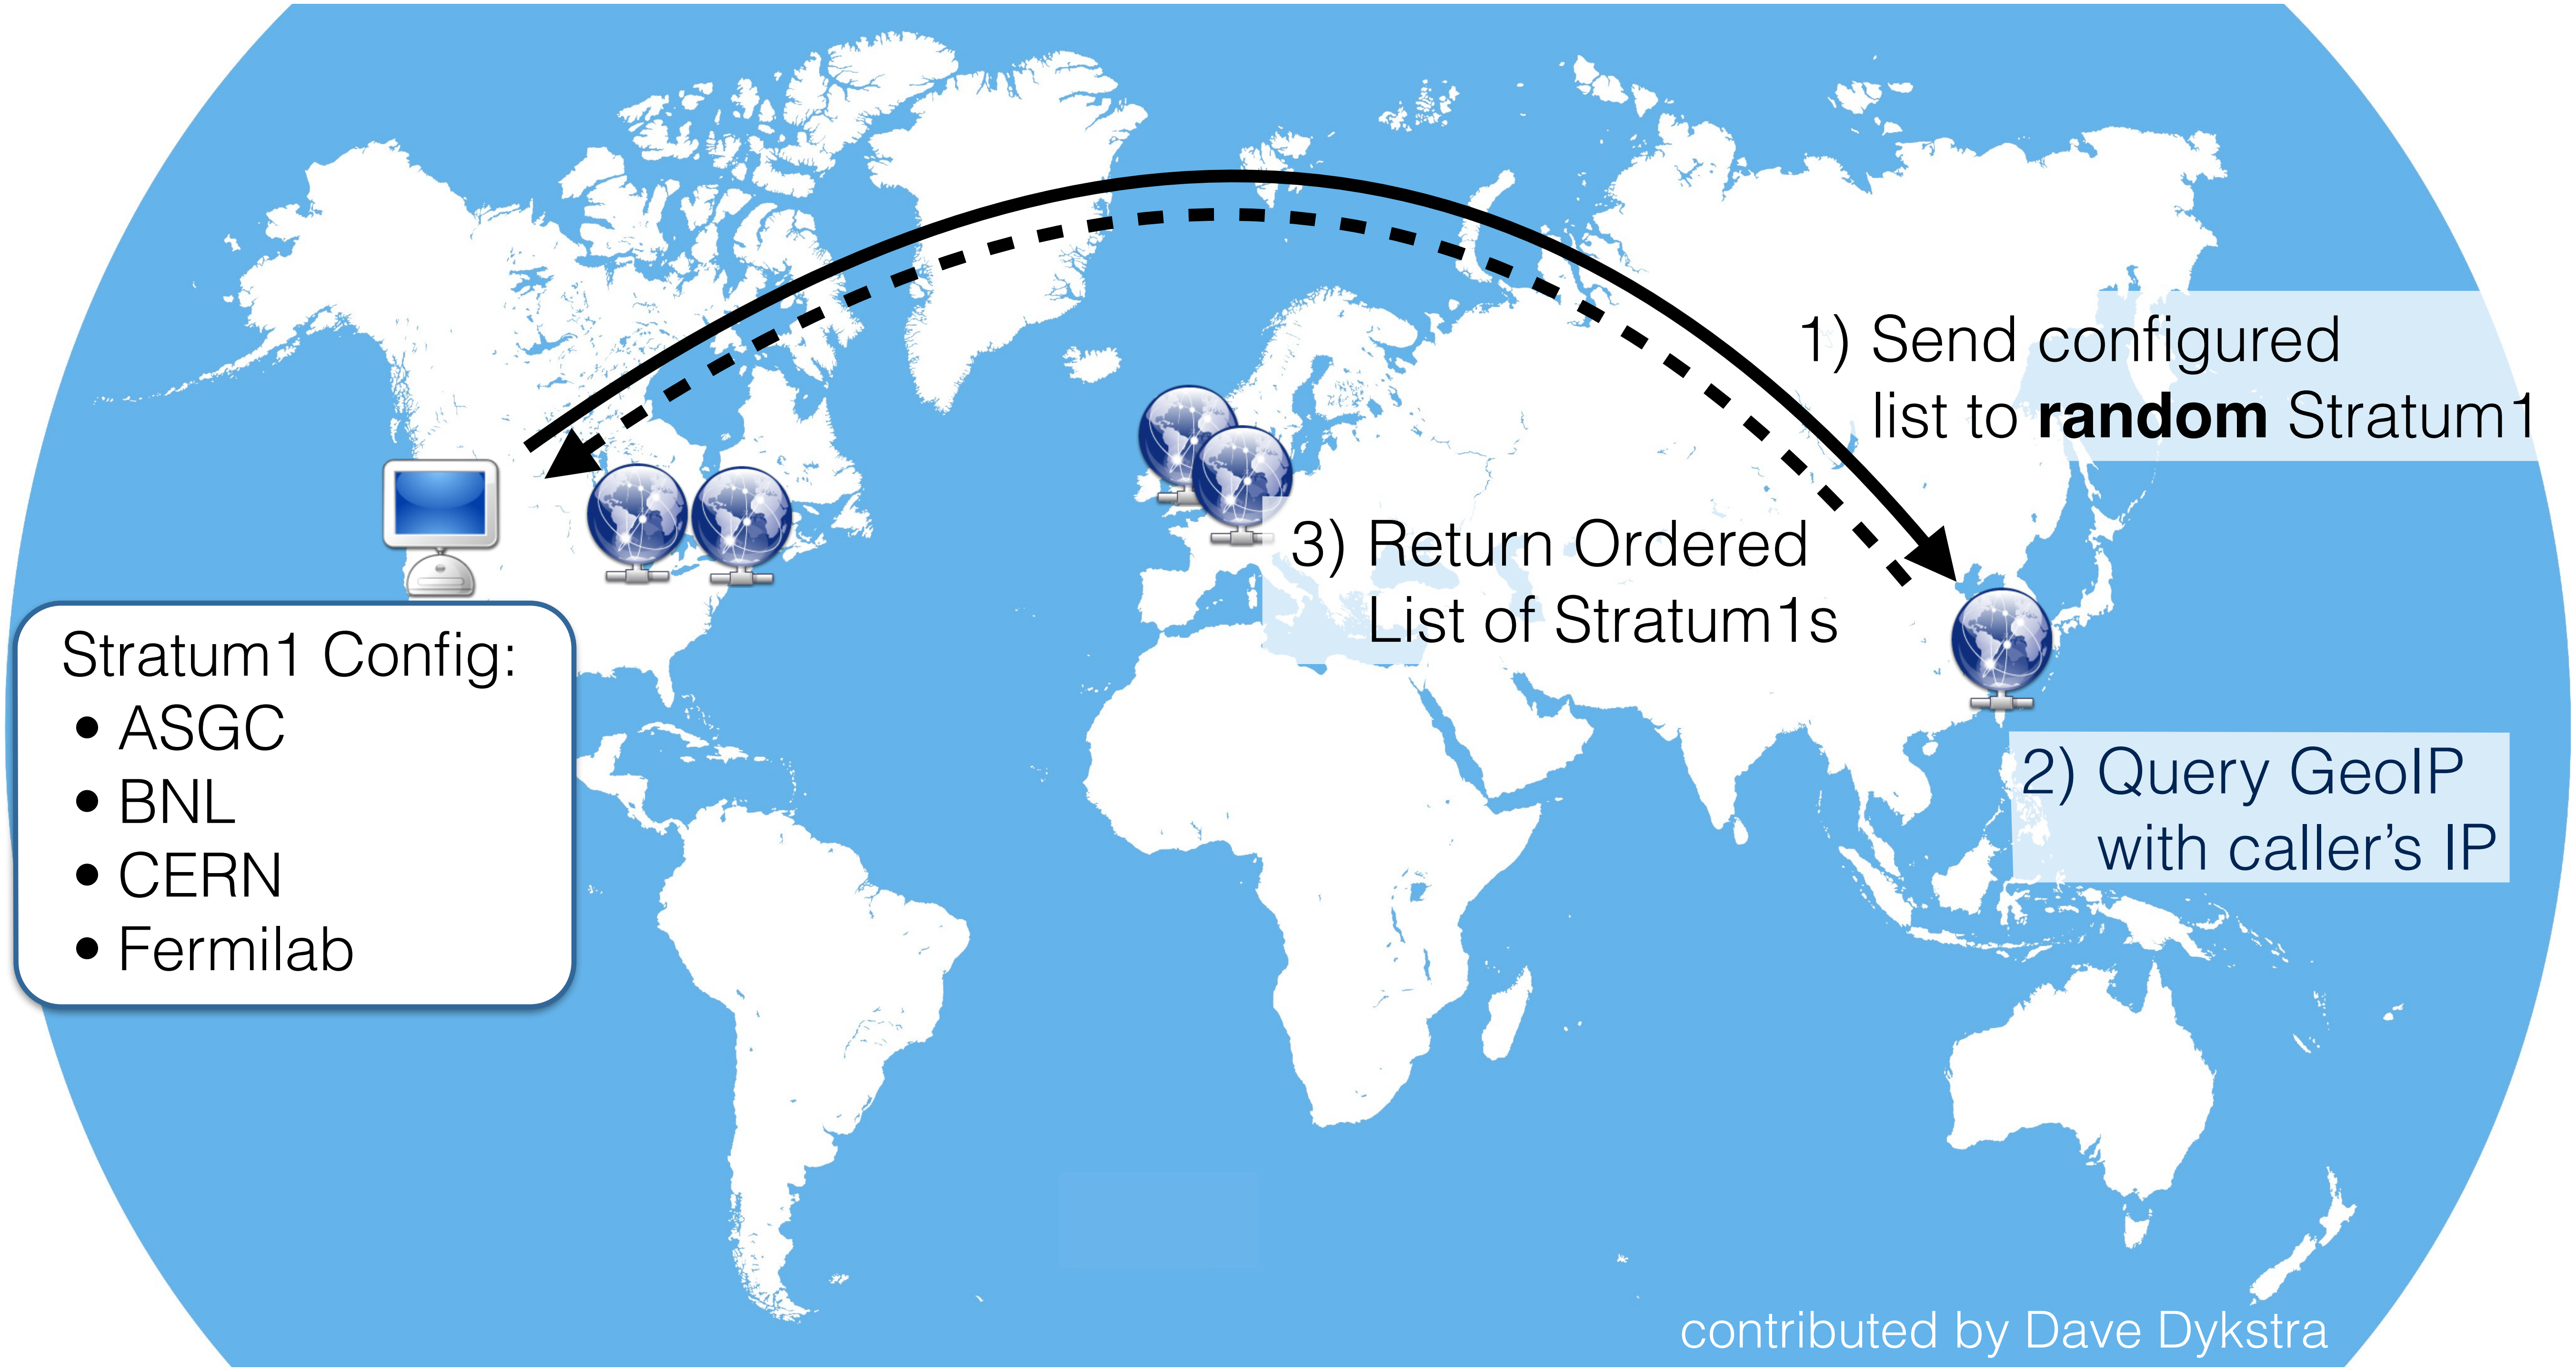

Stratum1 Config:

- ASGC
- BNL
- CERN
- Fermilab

1) Send configured list to **random** Stratum1

2) Query GeoIP with caller's IP

3) Return Ordered List of Stratum1s

contributed by Dave Dykstra

GeoIP-based Stratum1 Ordering

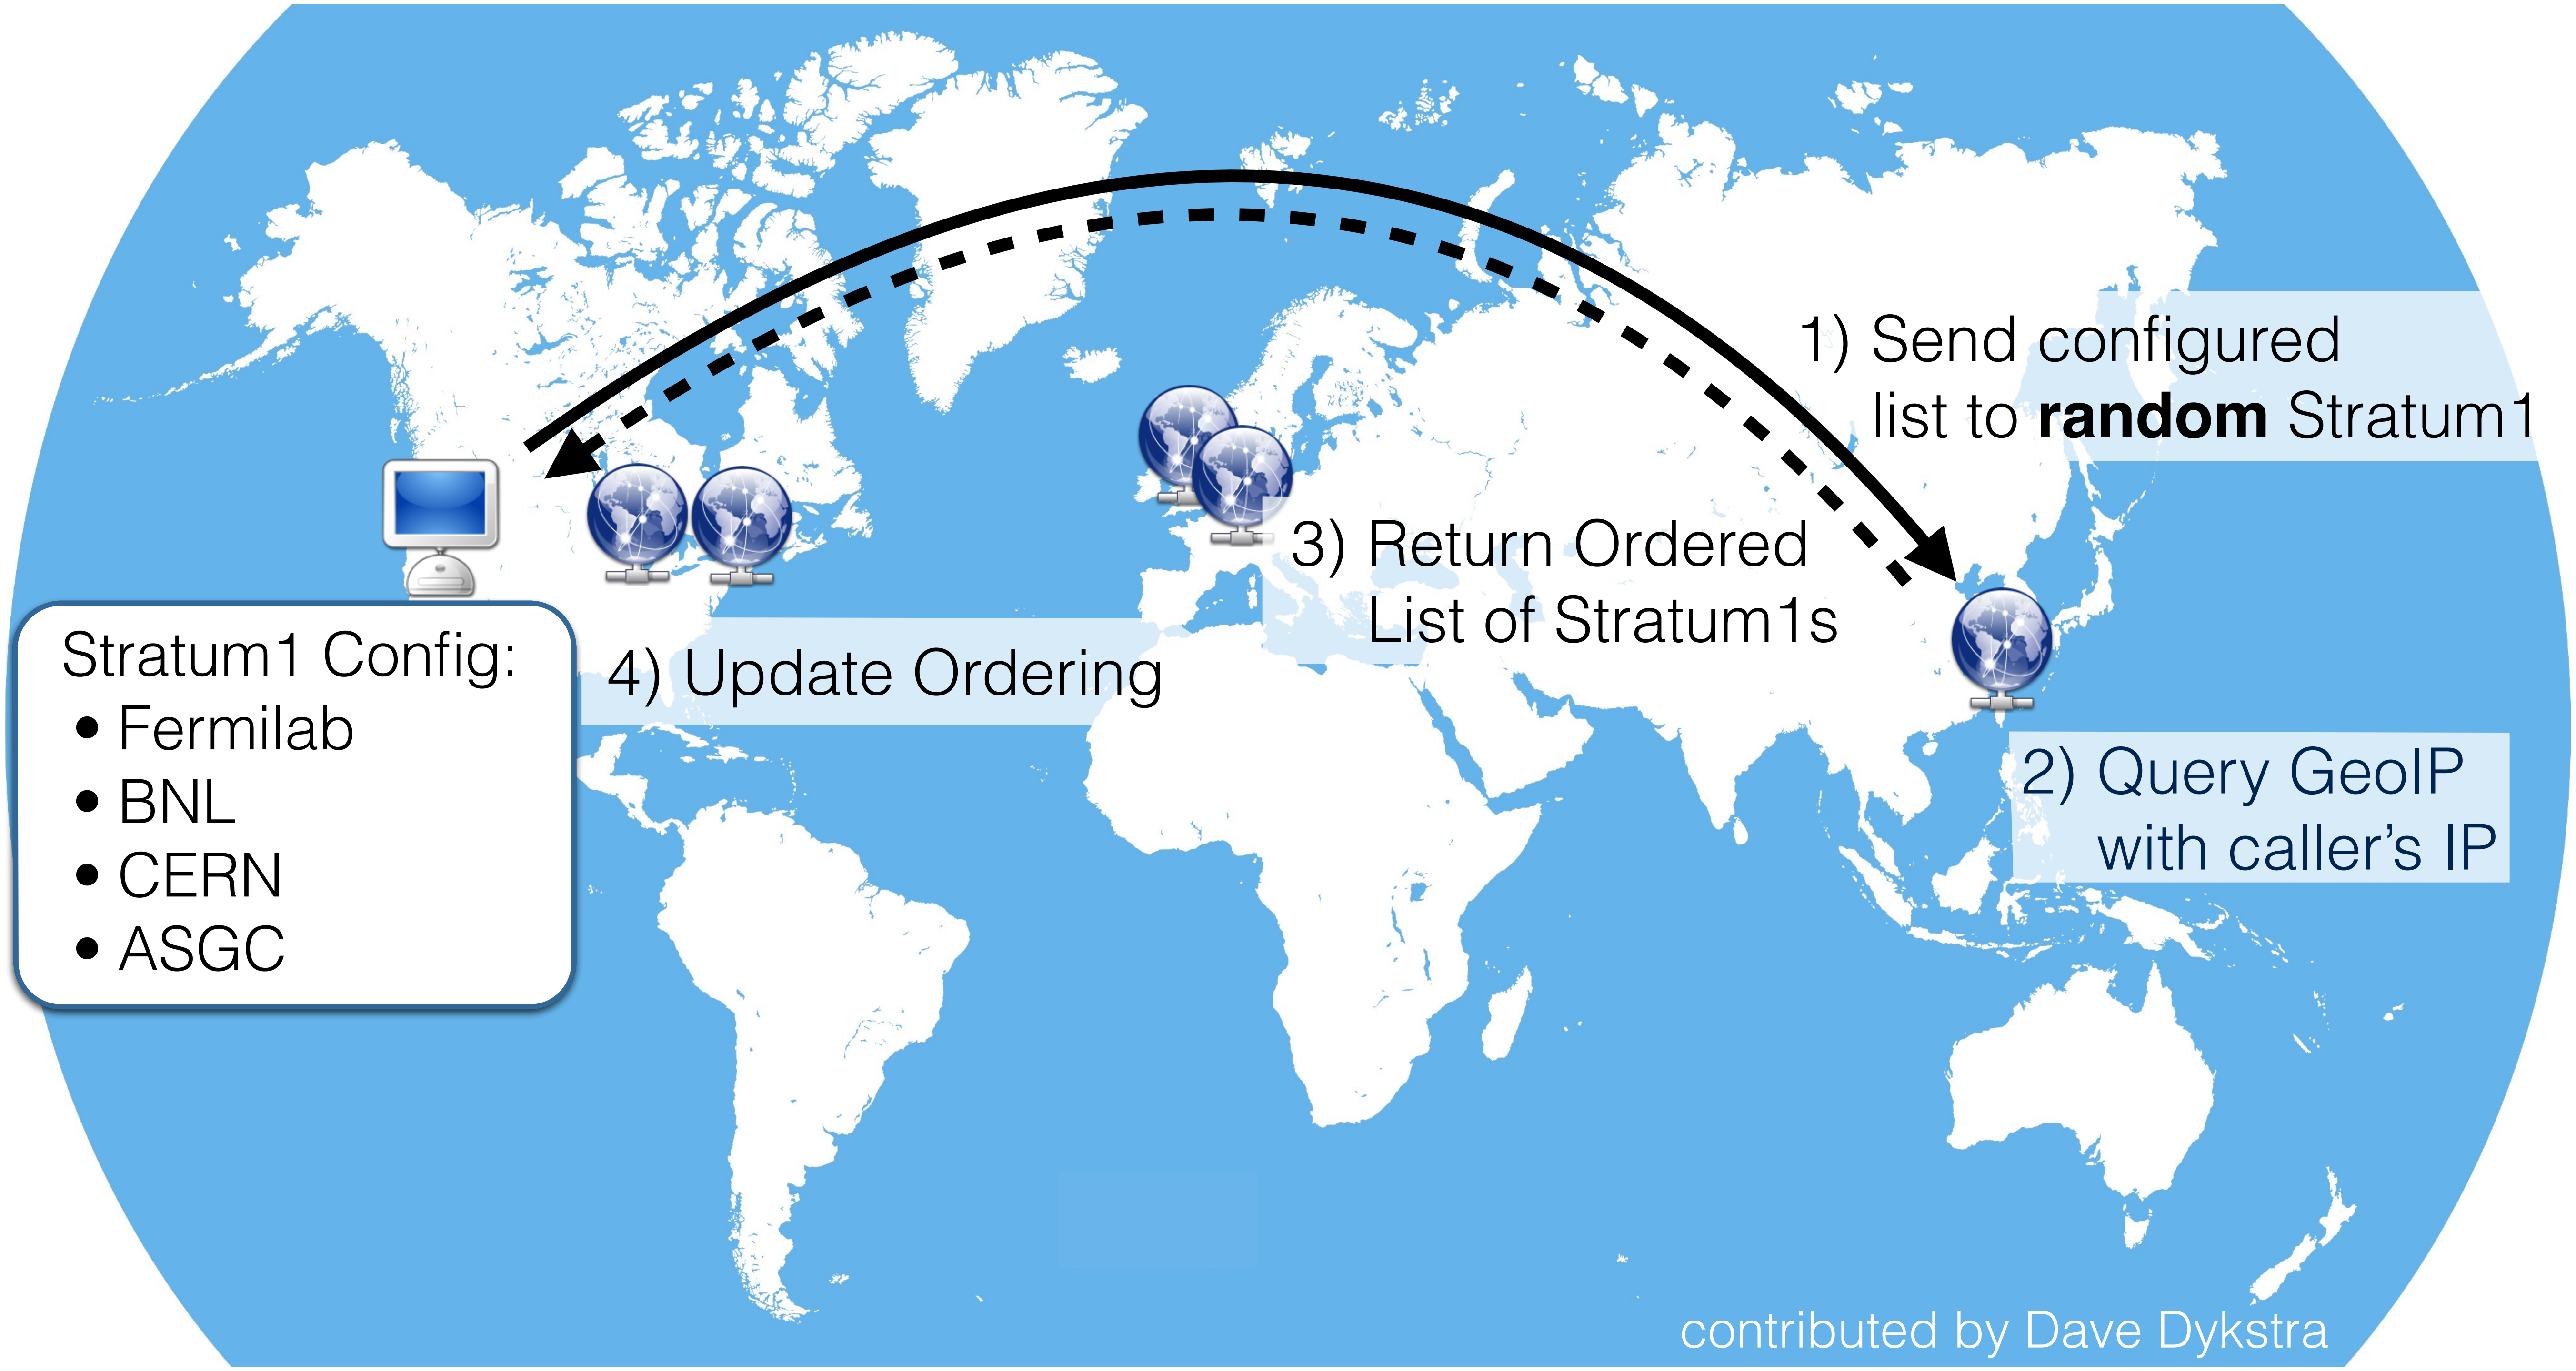

1) Send configured list to **random** Stratum1

2) Query GeoIP with caller's IP

3) Return Ordered List of Stratum1s

4) Update Ordering

Stratum1 Config:

- Fermilab
- BNL
- CERN
- ASGC

contributed by Dave Dykstra

GeoIP-based Stratum1 Ordering

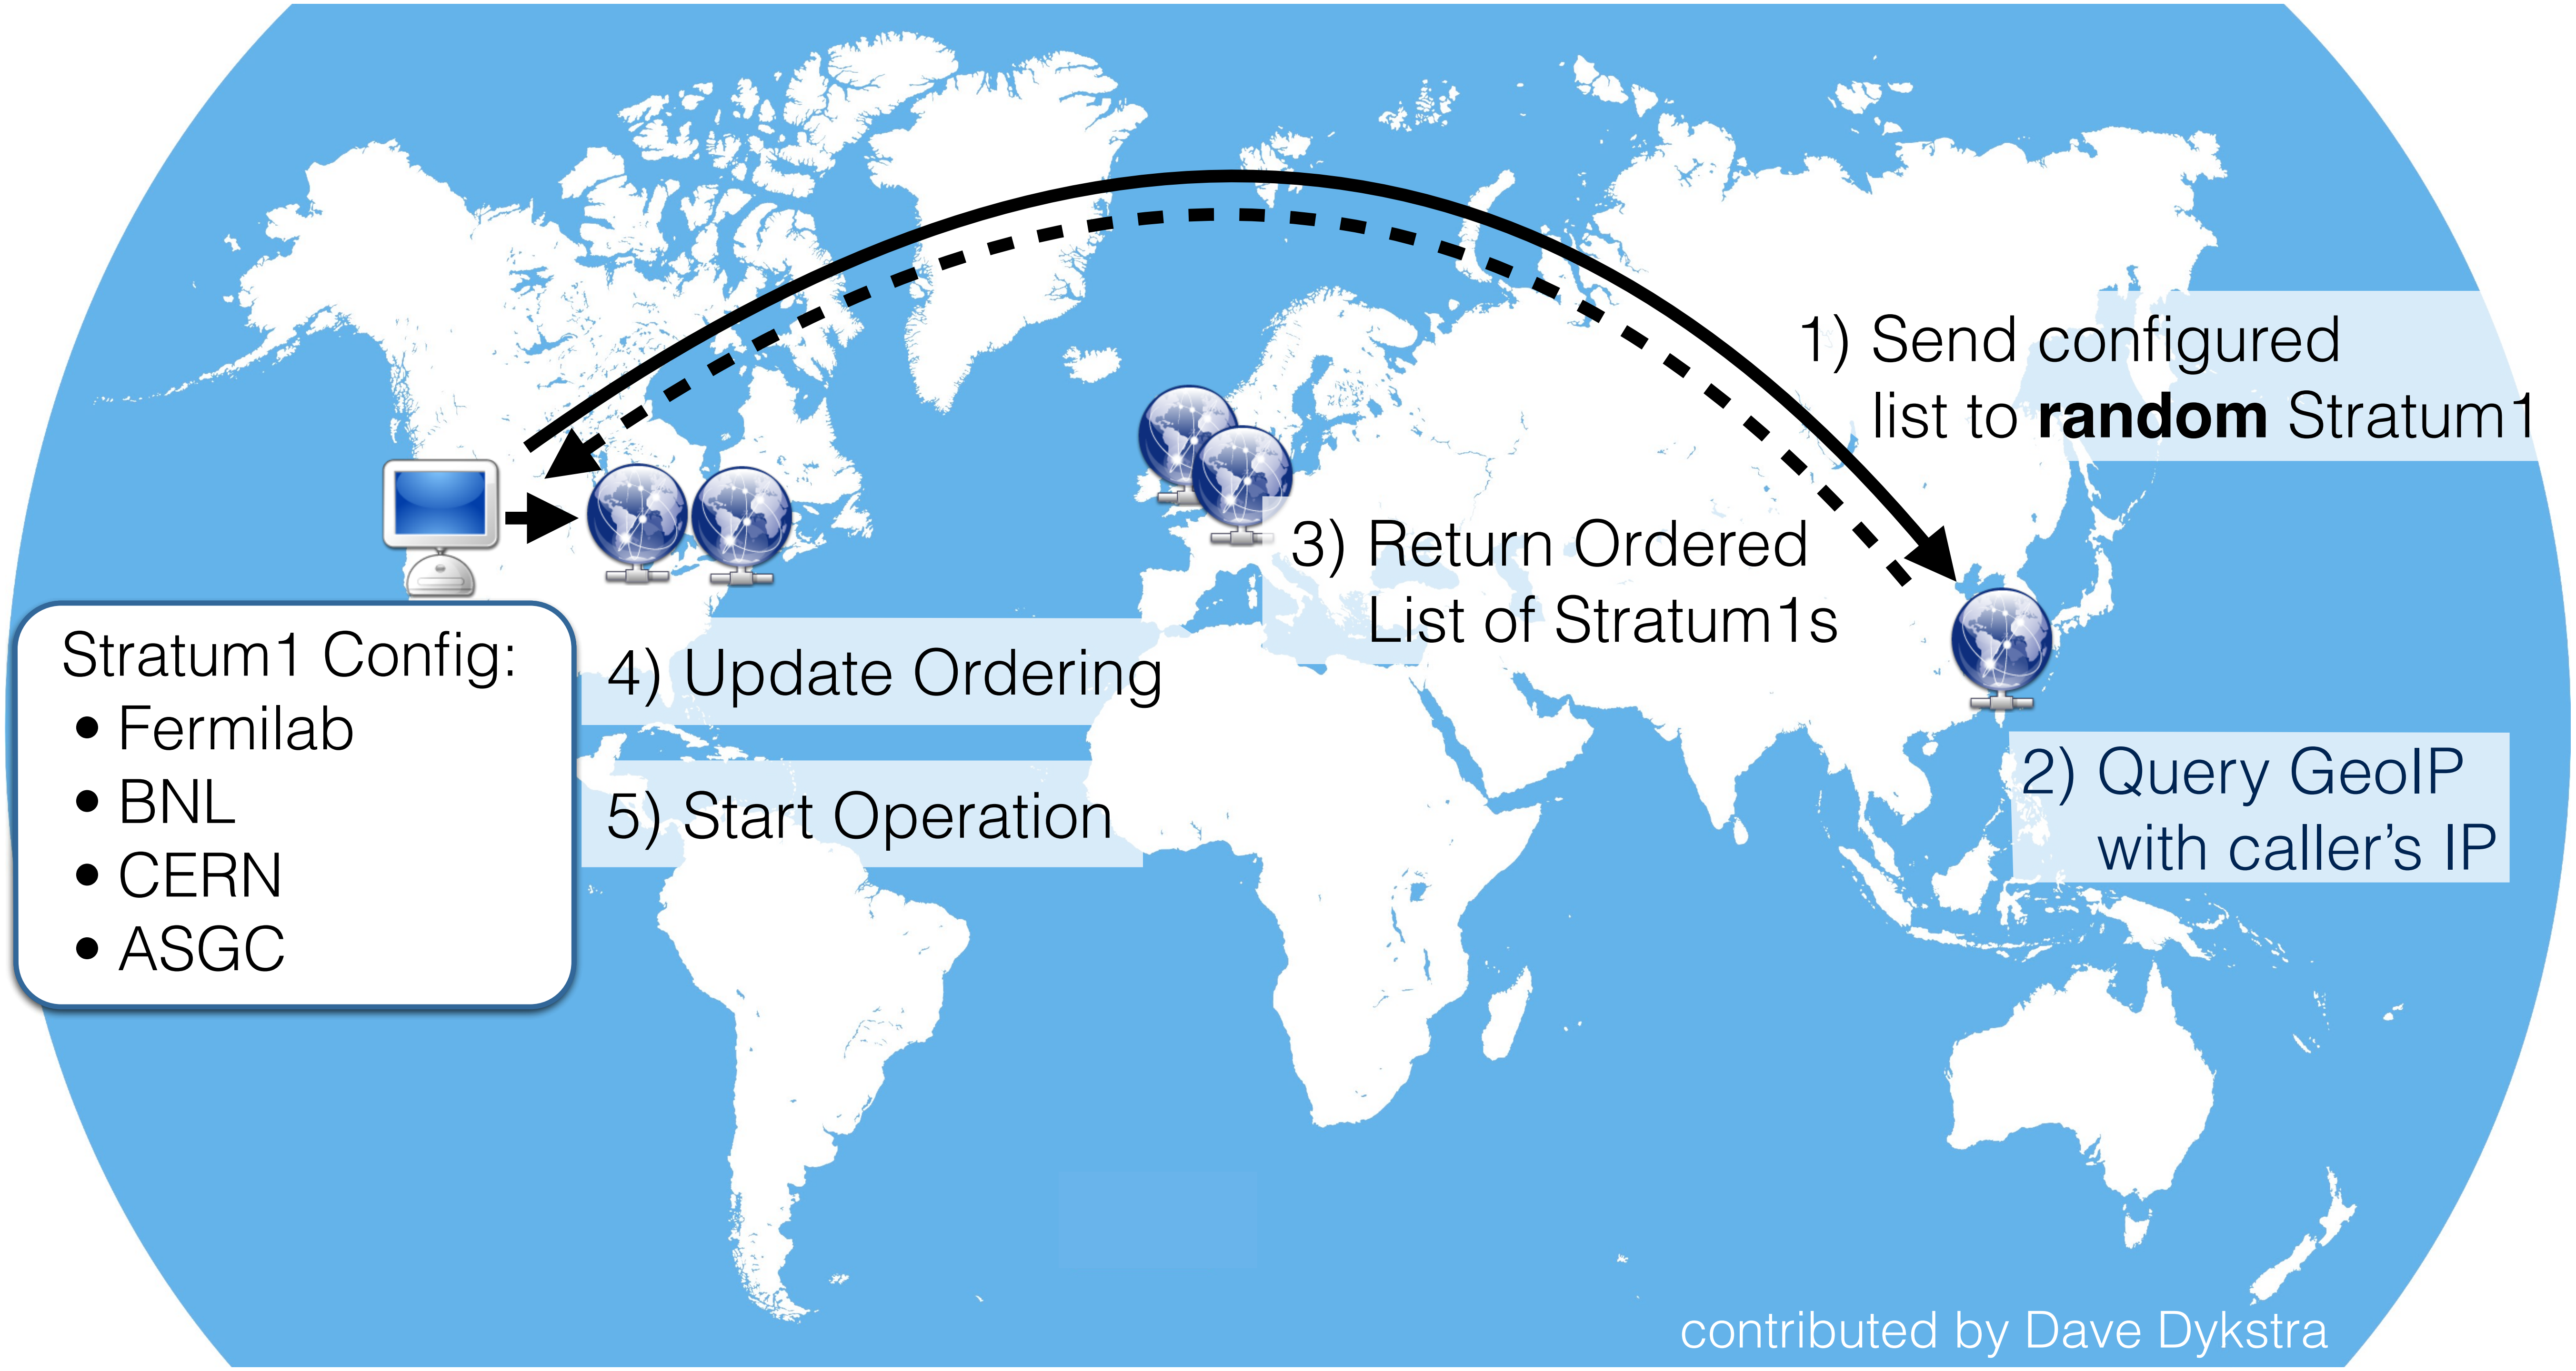

1) Send configured list to **random** Stratum1

2) Query GeoIP with caller's IP

3) Return Ordered List of Stratum1s

4) Update Ordering

5) Start Operation

Stratum1 Config:

- Fermilab
- BNL
- CERN
- ASGC

contributed by Dave Dykstra

History-less Repositories

History-less Repositories

- CernVM-FS backend initially designed as **insert-only**
- **New use-case:** LHC experiment’s nightly integration builds
  - High update rate (up to twice a day)
  - Large volume of newly staged files (10-100GiB per day)
  - Short lived installations (maximal two weeks)

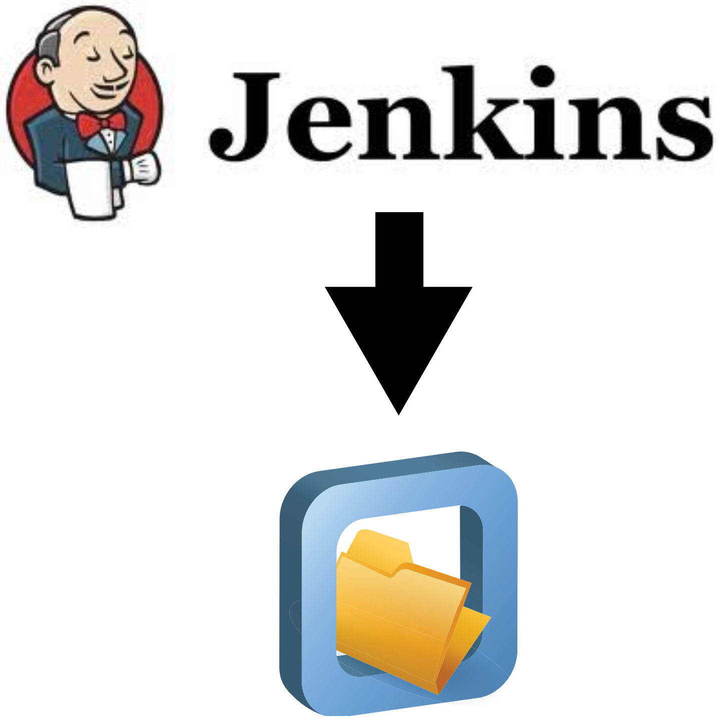

- Insert-only quickly fills up backend storage!
- History preservation is not needed
- Transparent backend cleanup required

➡ Automatic **garbage collection** mechanism

Repository Growth without GC

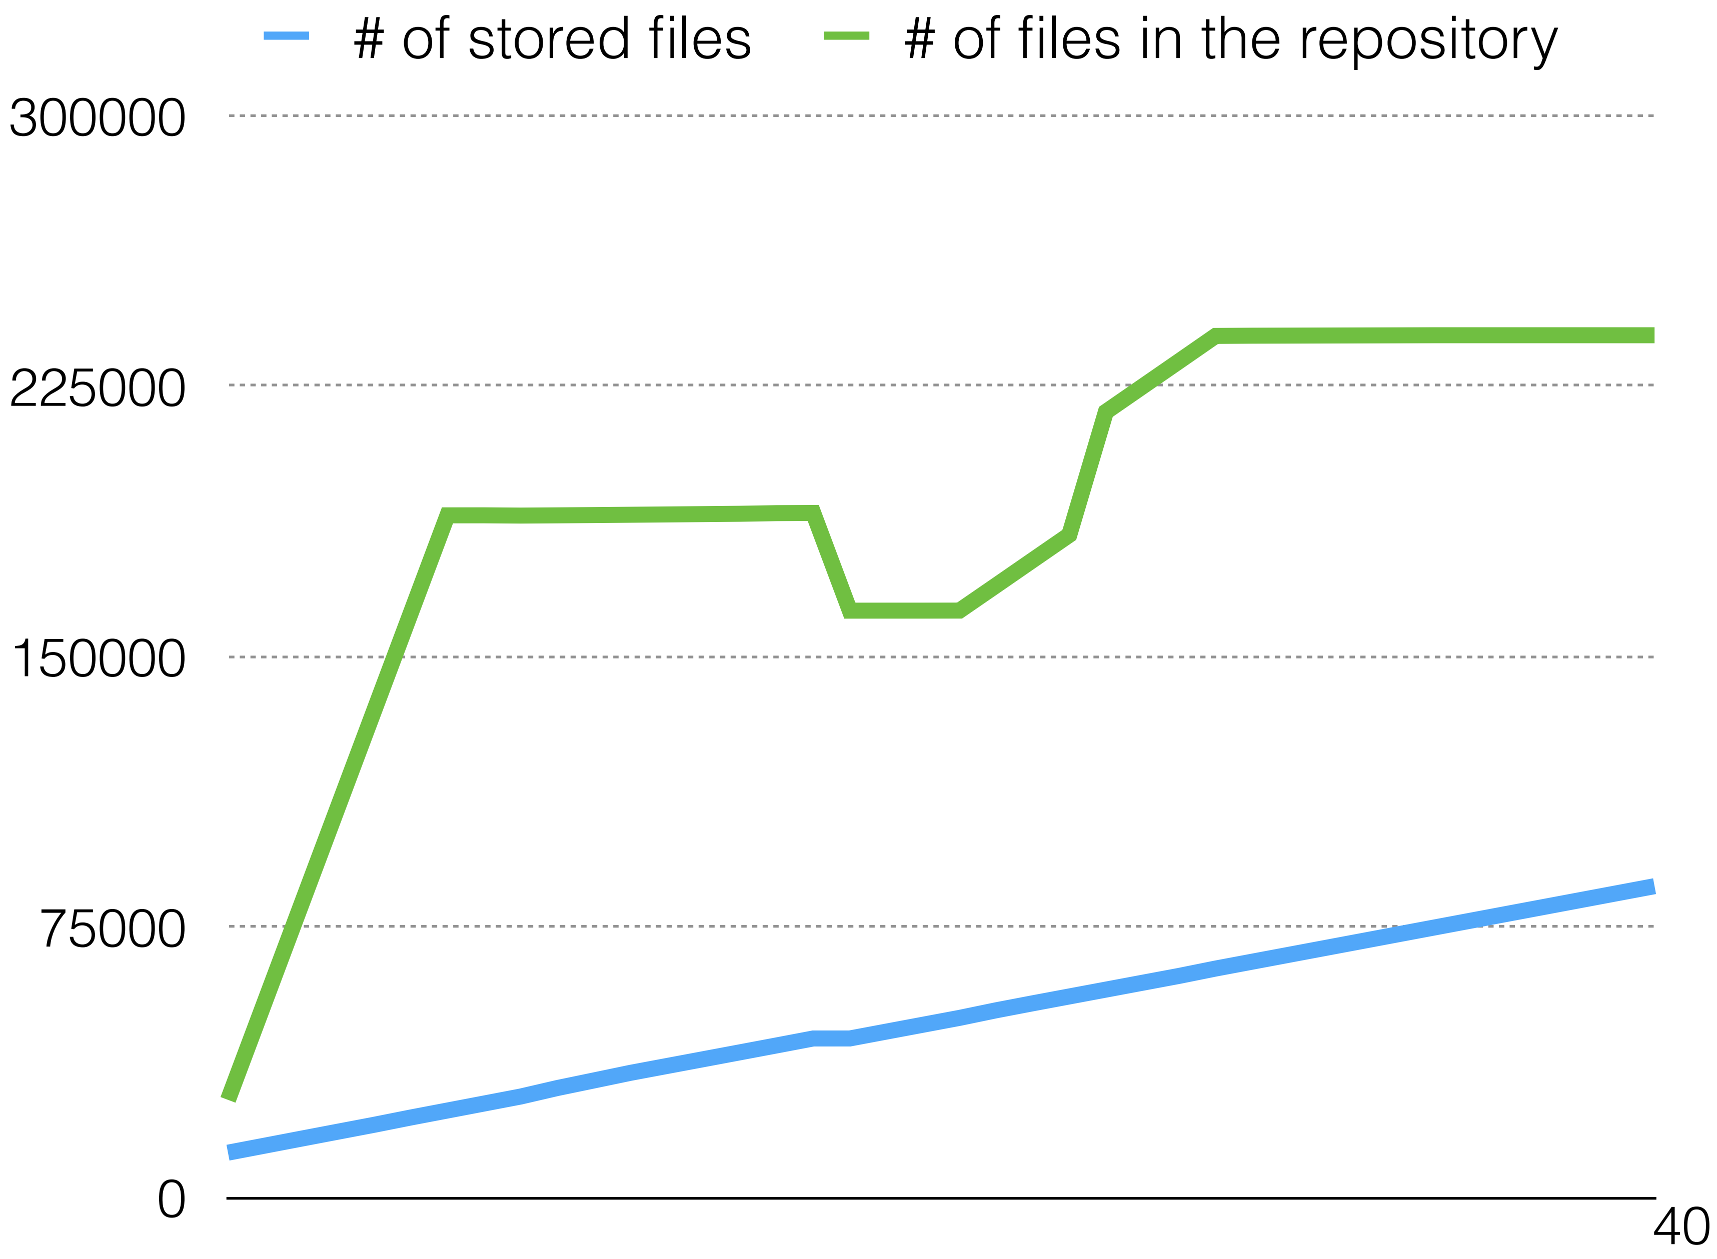

Repository Content:

- Installation of nightly integration build results
- Installed once per day
- Deleted after 7 days

Observation:

- Steady growth in the backend storage
- Quick accumulation of garbage in storage
- De-duplication: 1/8

Repository Growth without GC

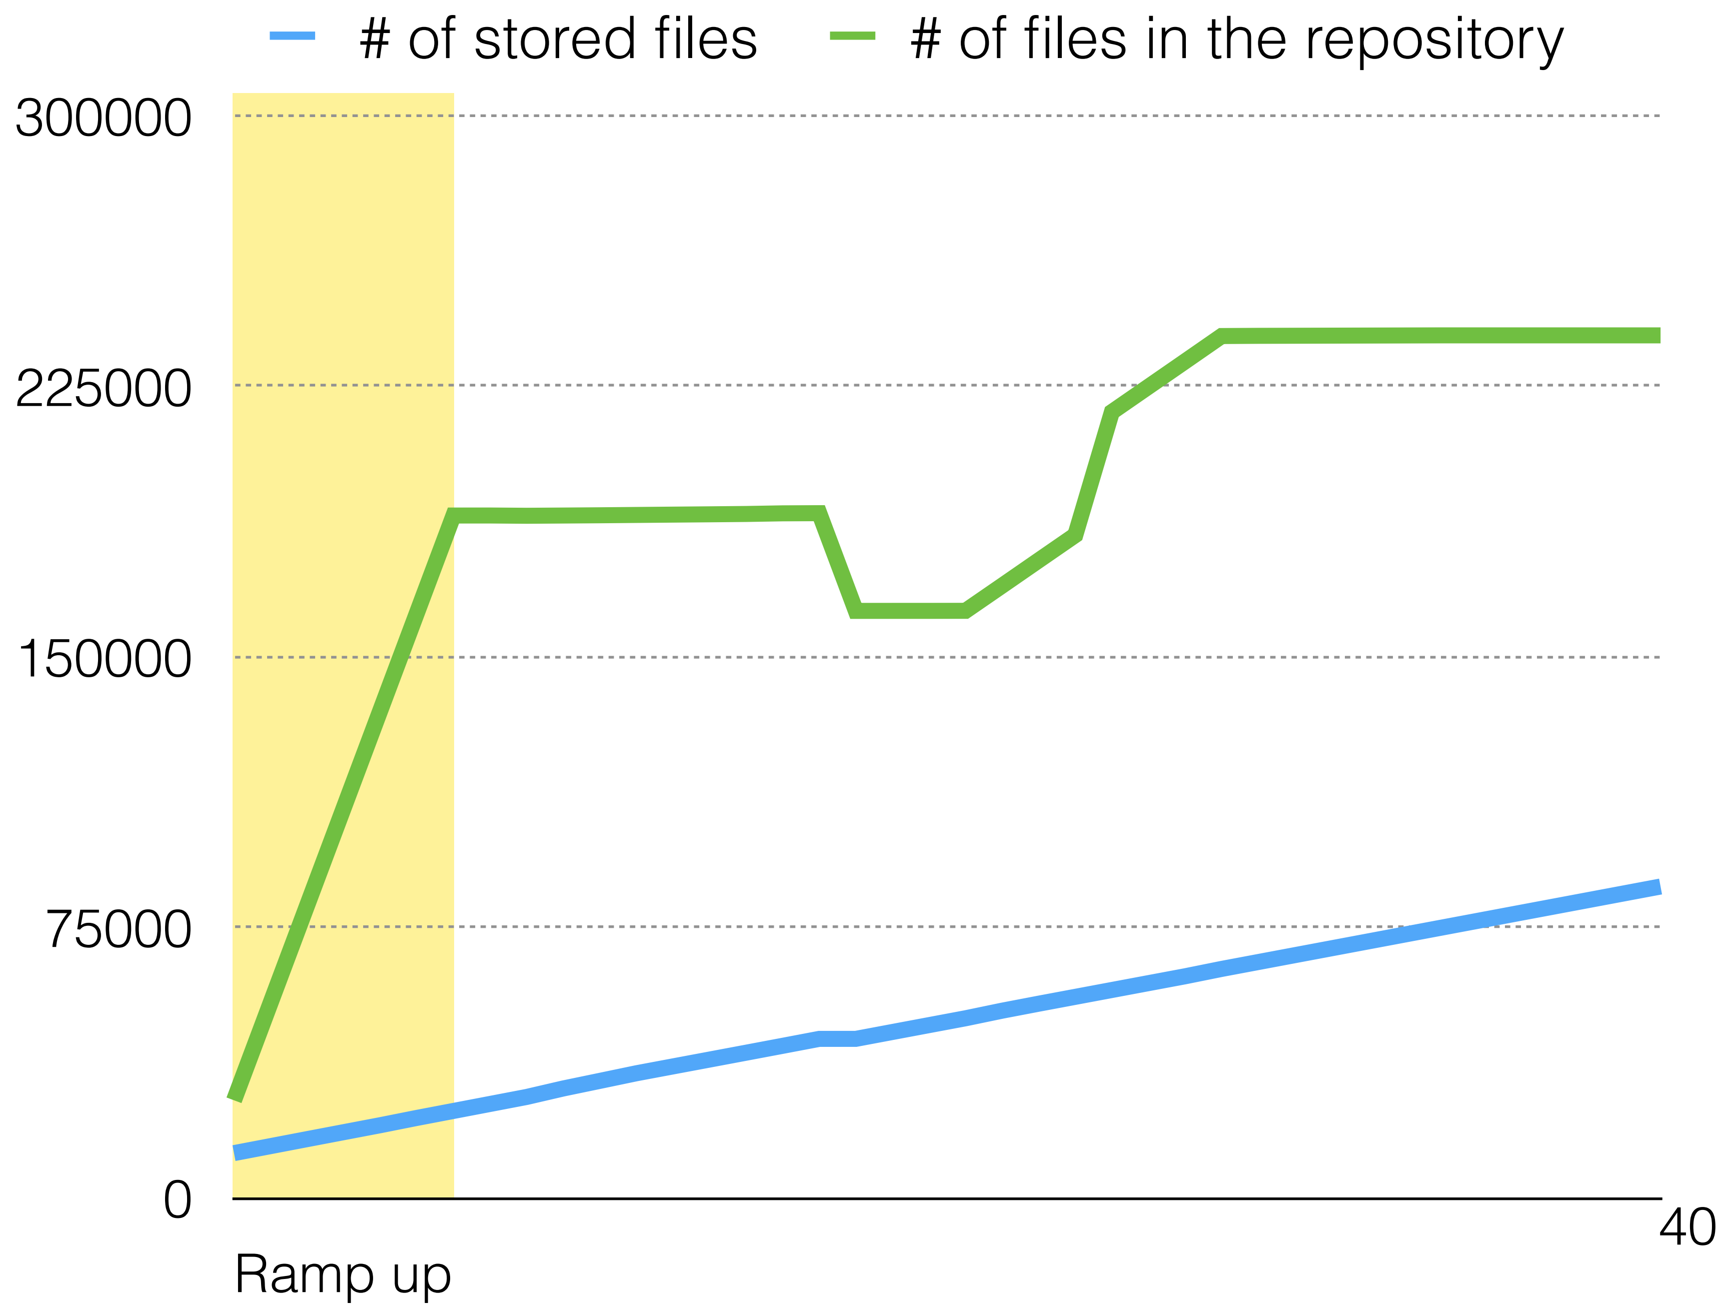

Repository Content:

- Installation of nightly integration build results
- Installed once per day
- Deleted after 7 days

Observation:

- Steady growth in the backend storage
- Quick accumulation of garbage in storage
- De-duplication: 1/8

Repository Growth without GC

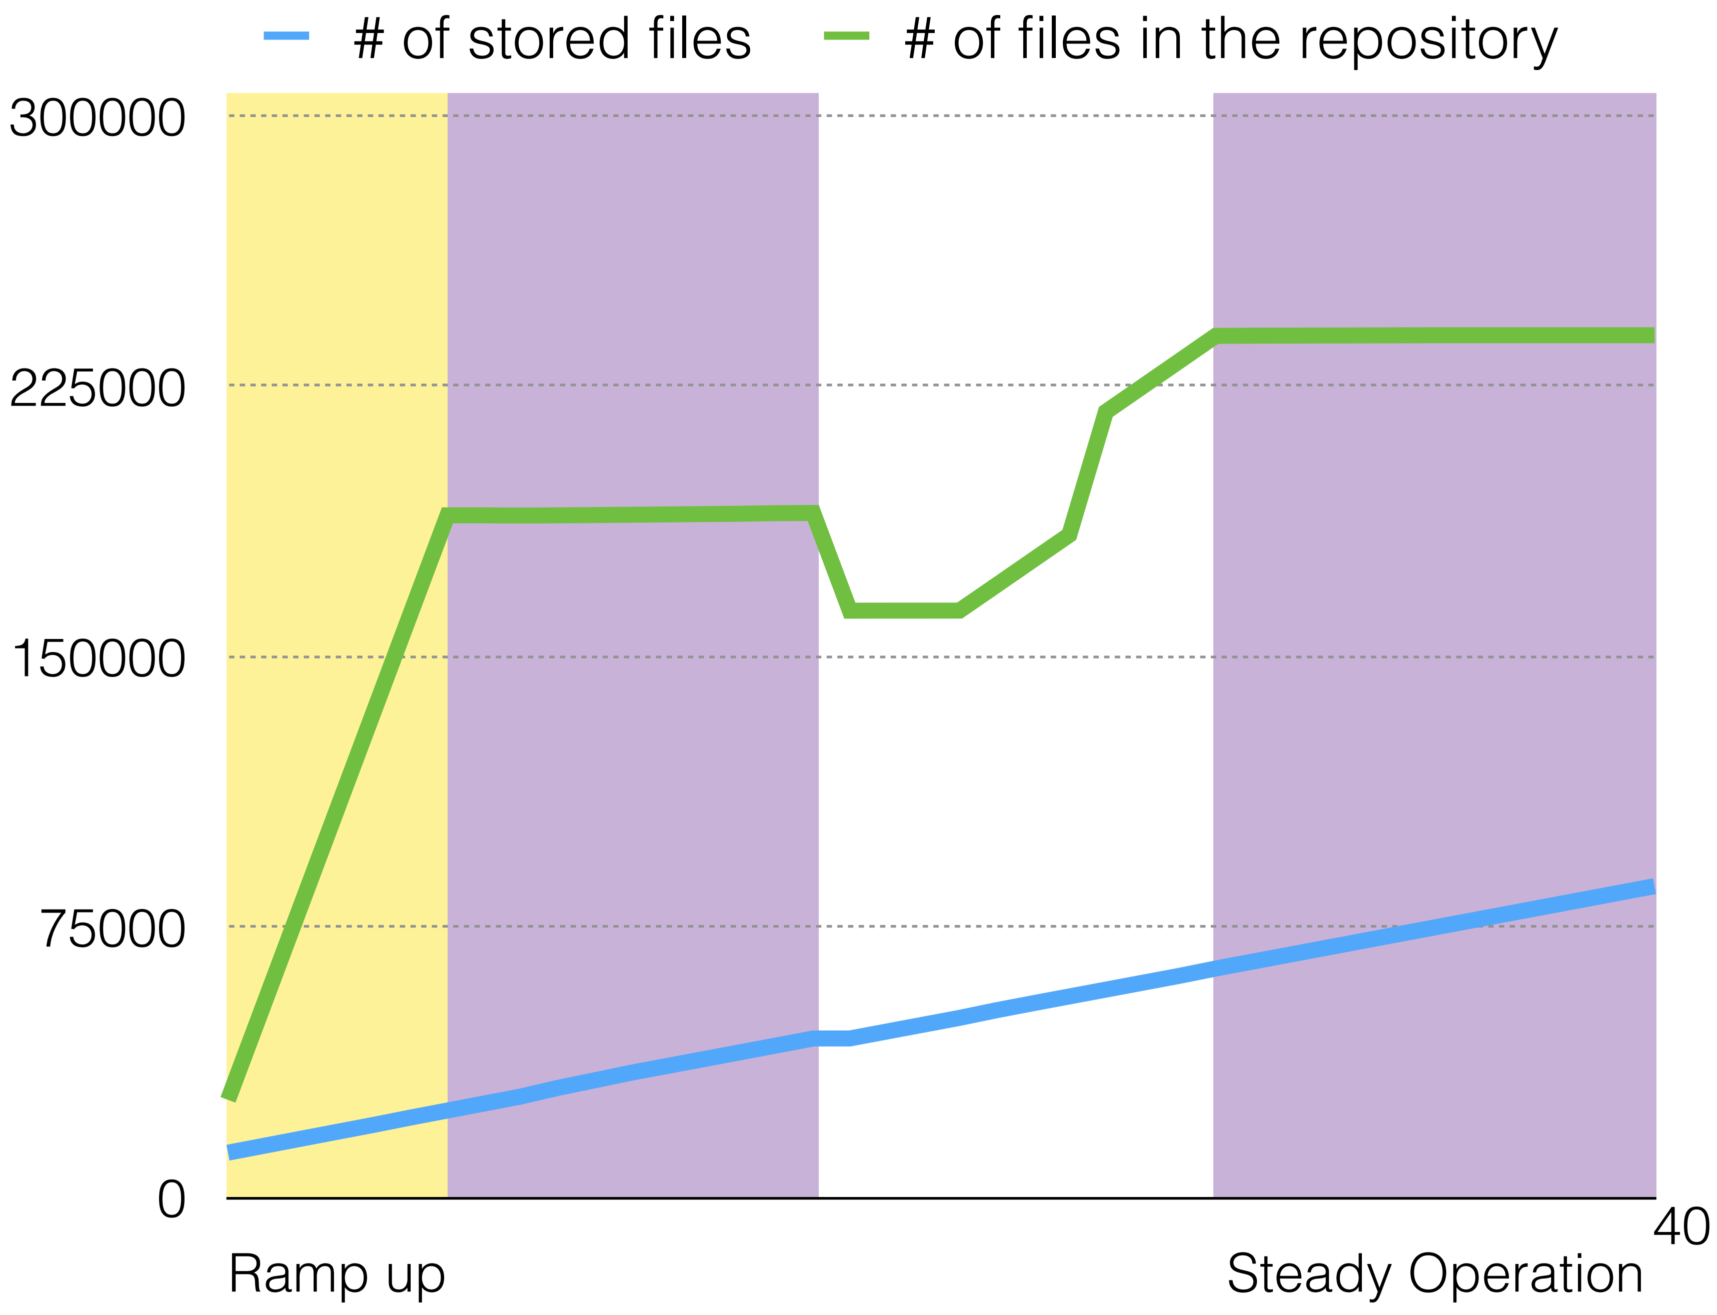

Repository Content:

- Installation of nightly integration build results
- Installed once per day
- Deleted after 7 days

Observation:

- Steady growth in the backend storage
- Quick accumulation of garbage in storage
- De-duplication: 1/8

Repository Growth without GC

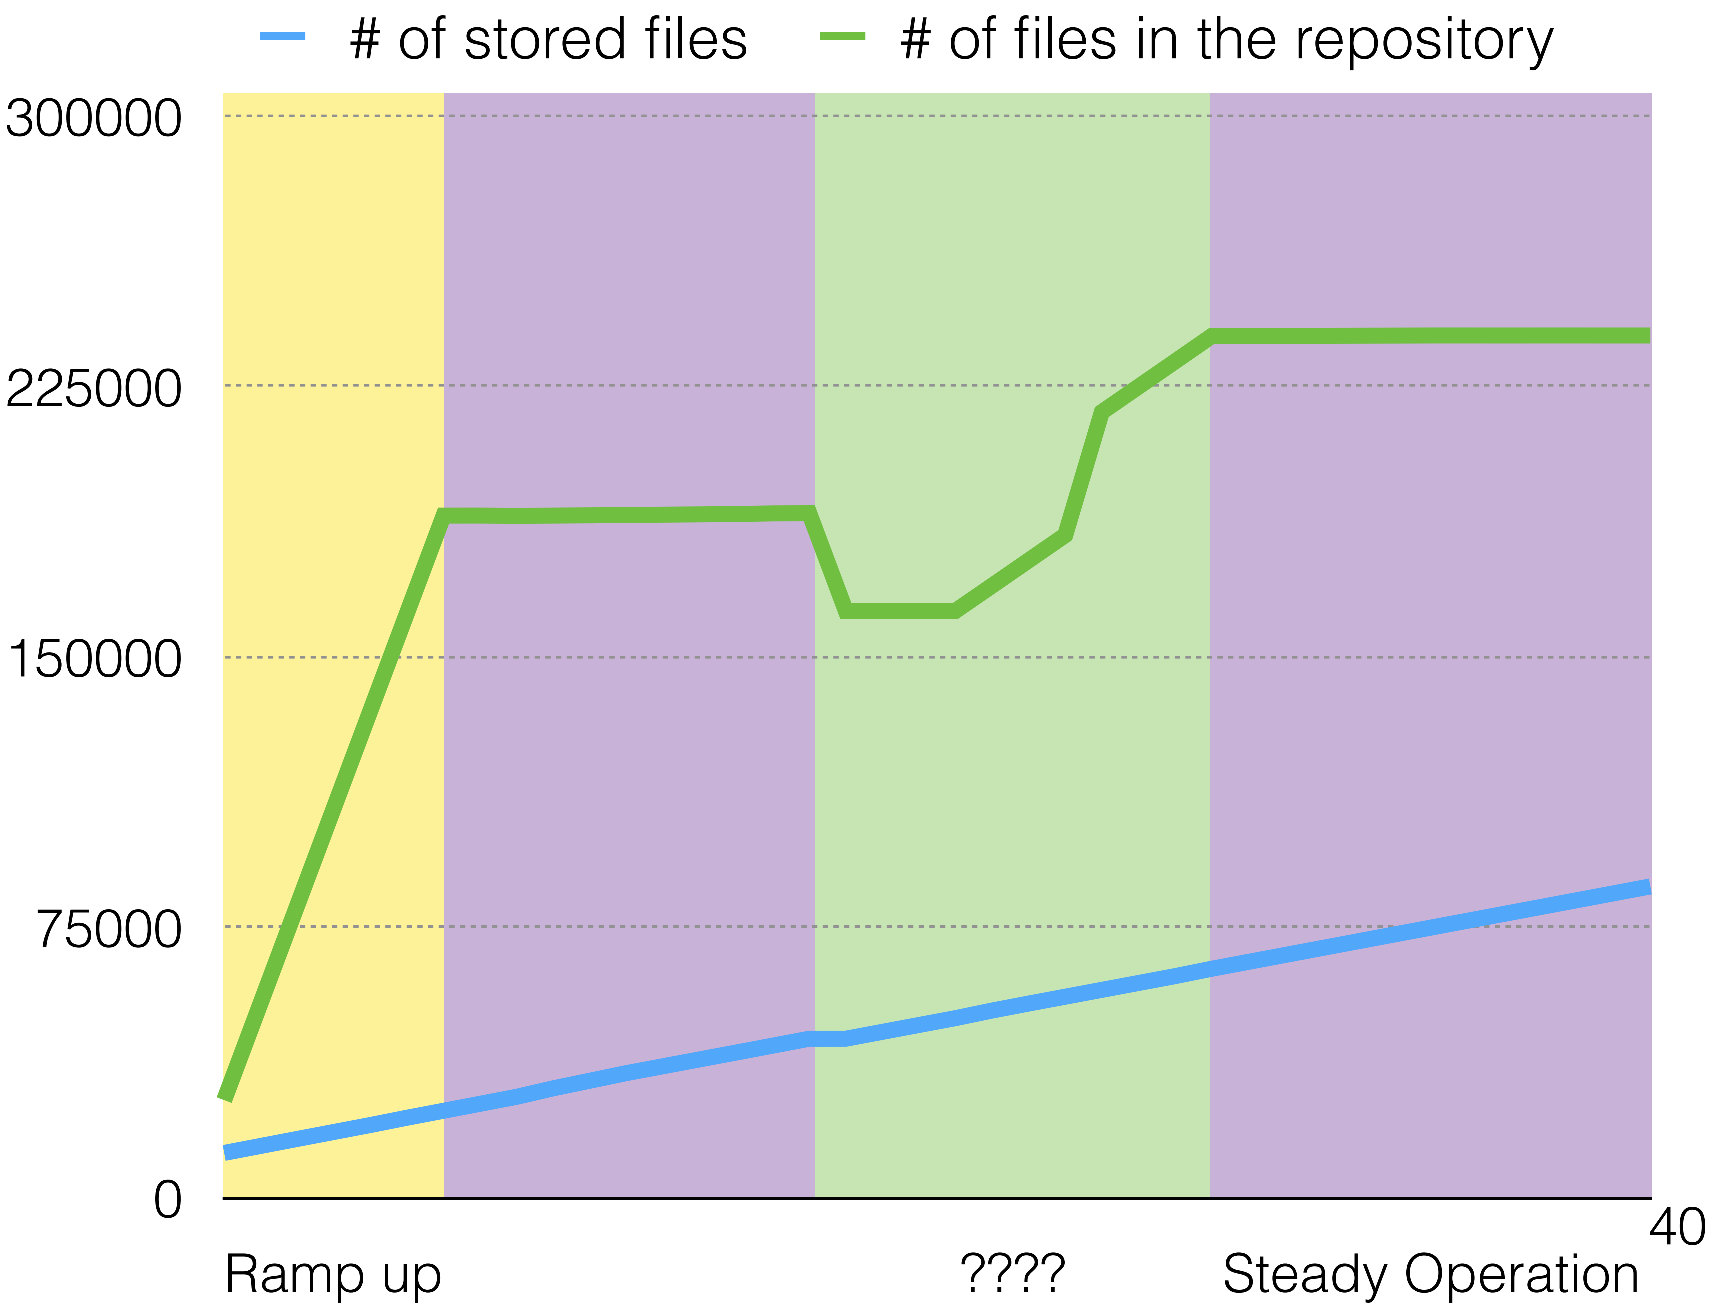

Repository Content:

- Installation of nightly integration build results
- Installed once per day
- Deleted after 7 days

Observation:

- Steady growth in the backend storage
- Quick accumulation of garbage in storage
- De-duplication: 1/8

Repository Growth with GC

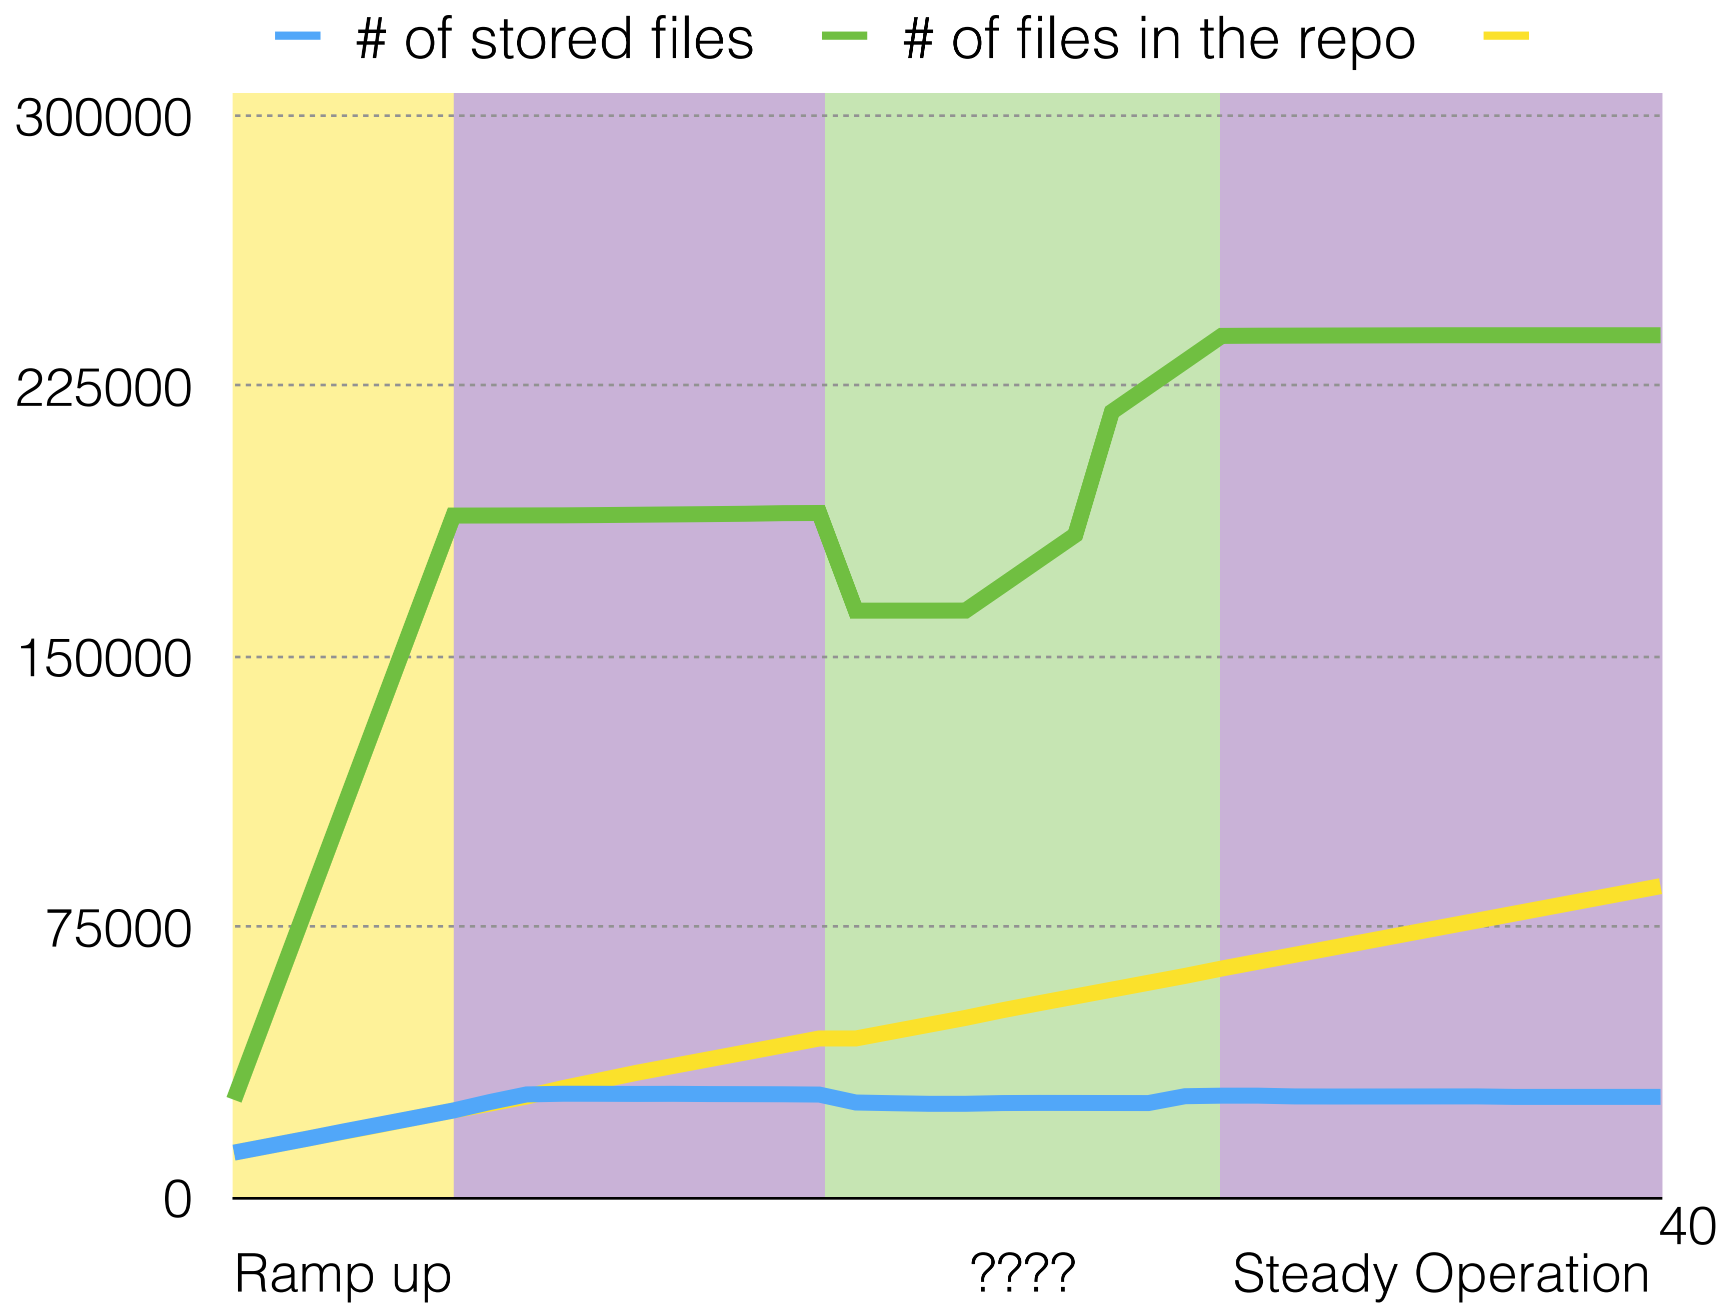

Repository Content:

- Installation of nightly integration build results
- Installed once per day
- Deleted after 7 days

Observation:

- Backend volume size remains steady as well
- Garbage outweighed referenced backend-data rather quickly

Repository Growth with GC

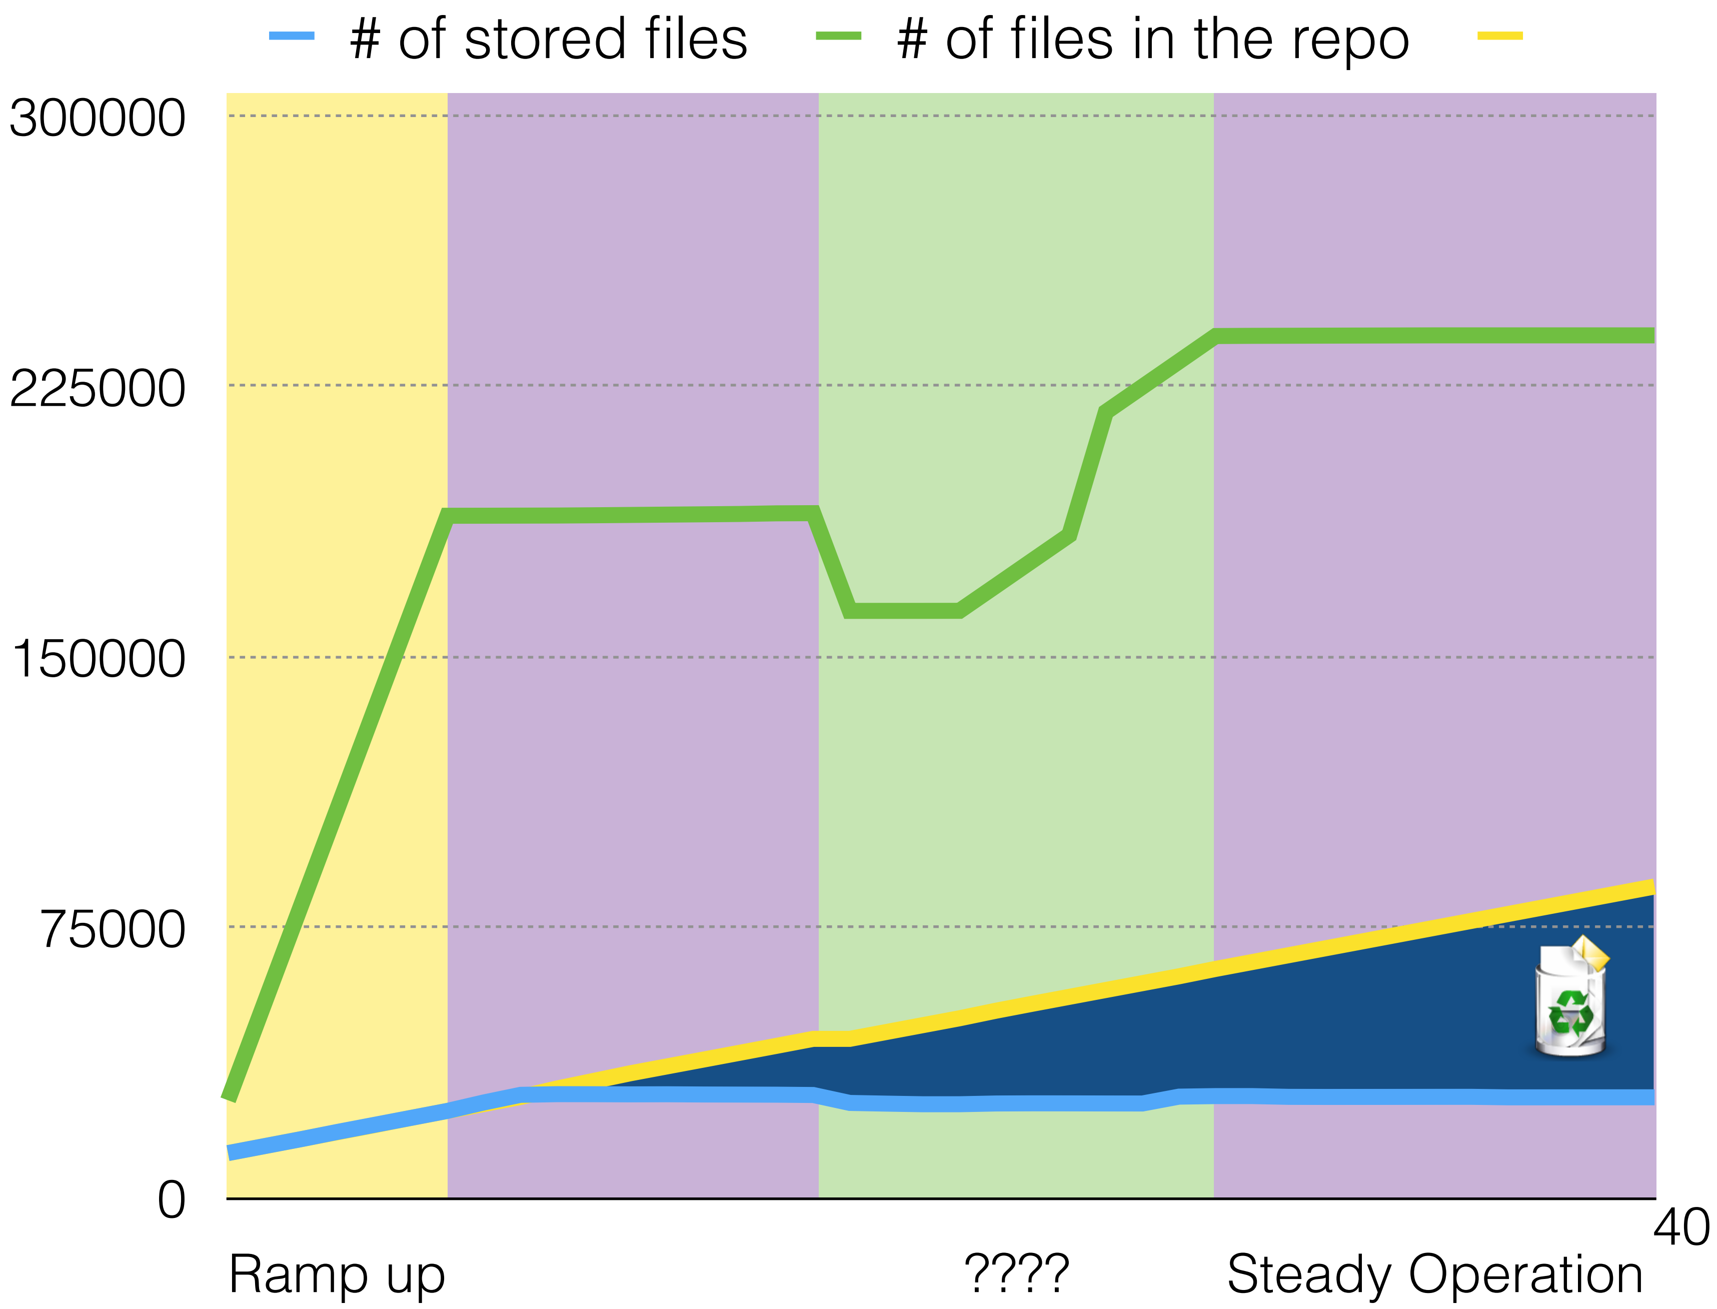

Repository Content:

- Installation of nightly integration build results
- Installed once per day
- Deleted after 7 days

Observation:

- Backend volume size remains steady as well
- Garbage outweighed referenced backend-data rather quickly

`Garbage Collection Details

- To come in **CernVM-FS 2.1.20**
- Configurable history preservation grace period (default: 3 days)
- Automatic cleanup of older revisions after each publish/snapshot
- Cleanup on both the stratum 0 and the stratum 1
- Garbage Collection based on *“Mark-and-Sweep”*

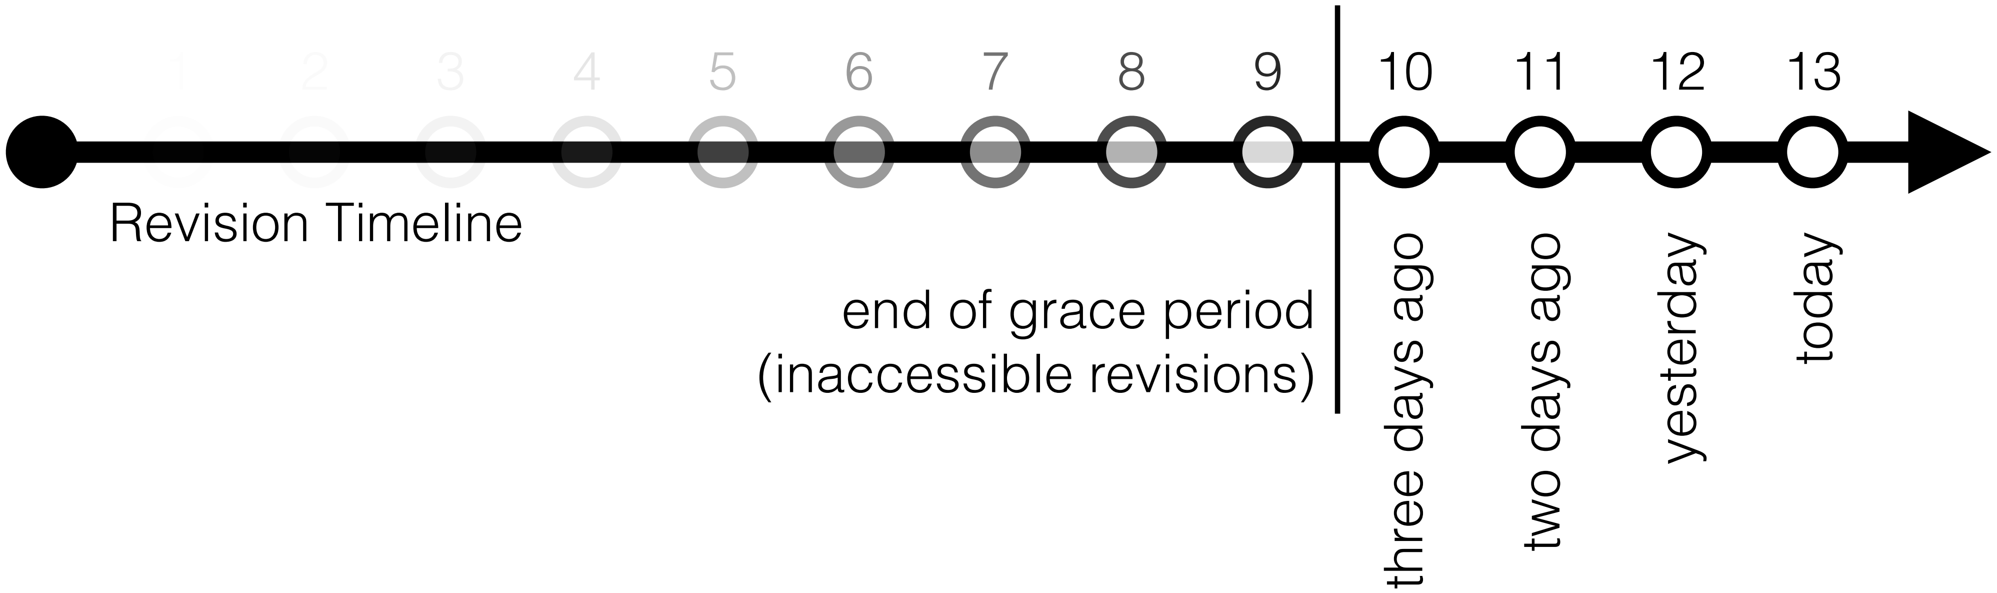

Near Future Developments

Support for Multiple Union File Systems

Stratum0
(backend storage)

Support for Multiple Union File Systems

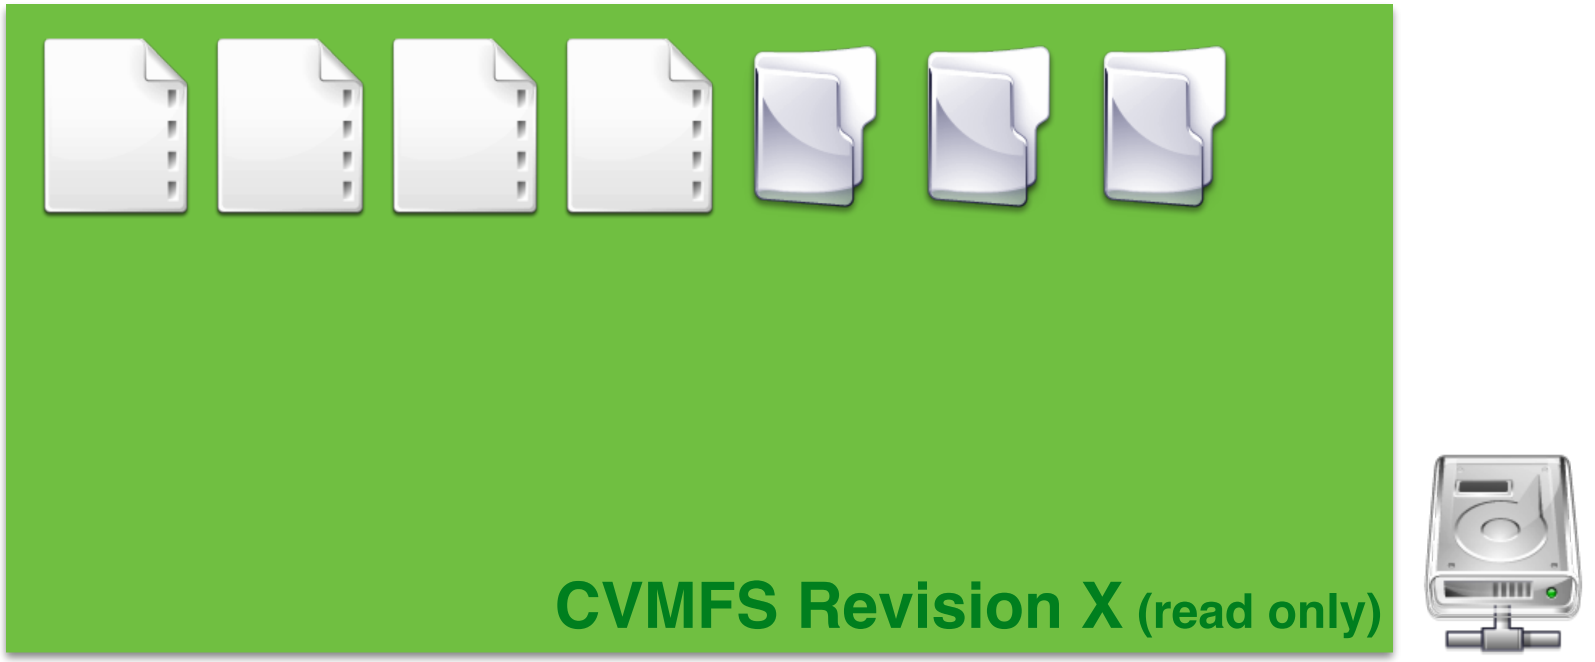

Stratum0
(backend storage)

Support for Multiple Union File Systems

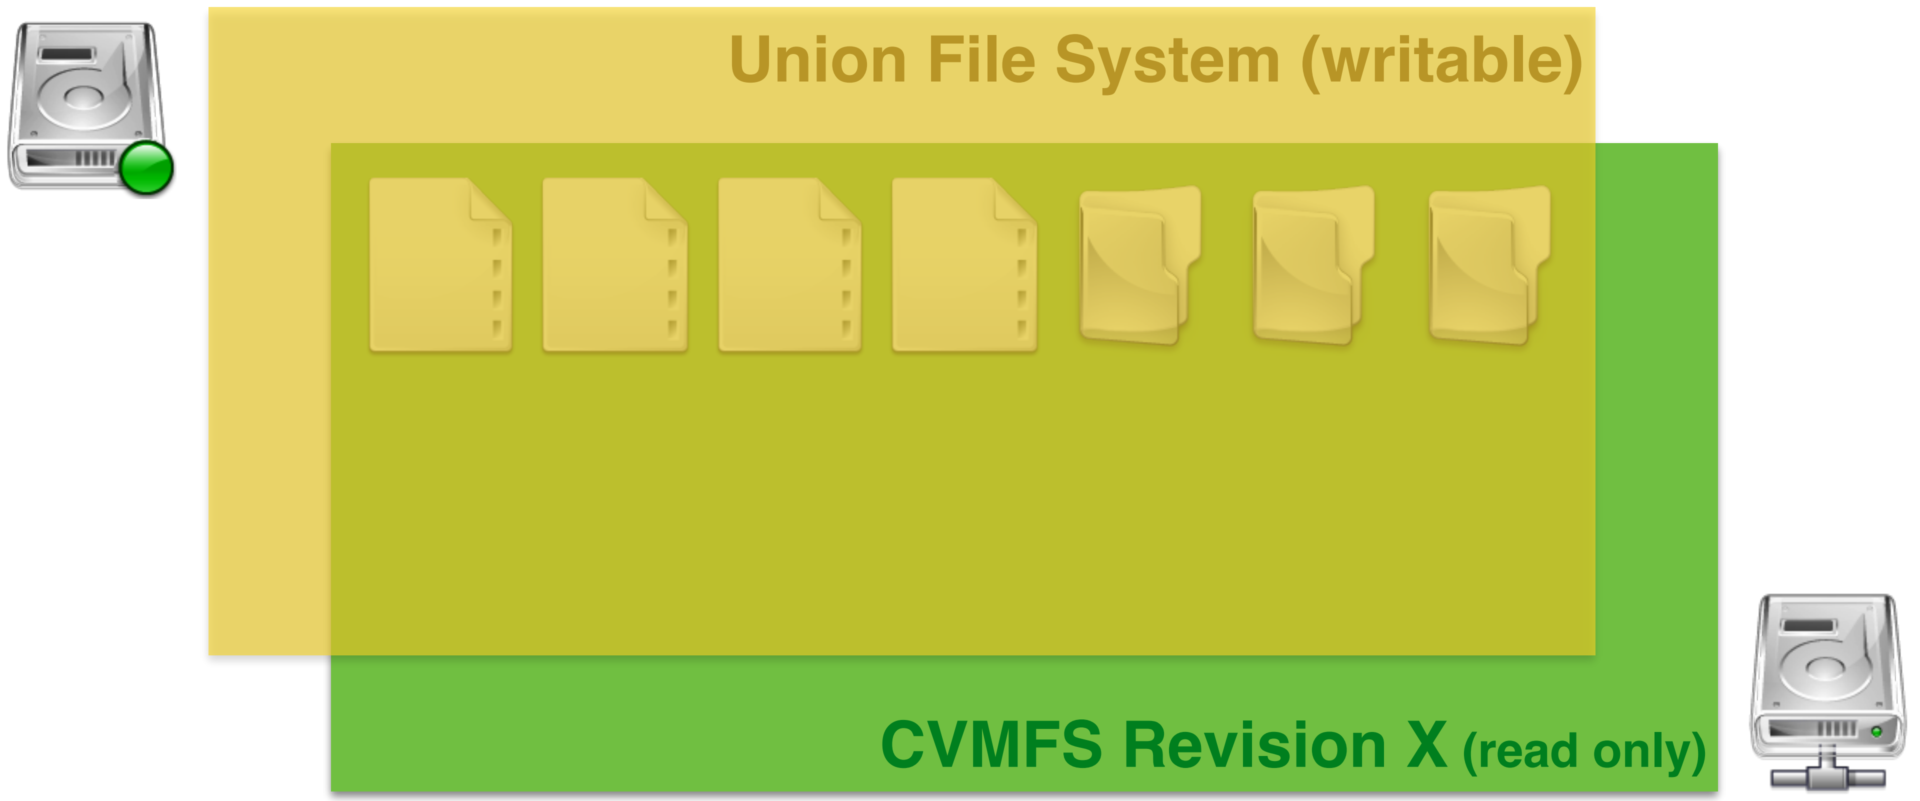

Stratum0
(backend storage)

Support for Multiple Union File Systems

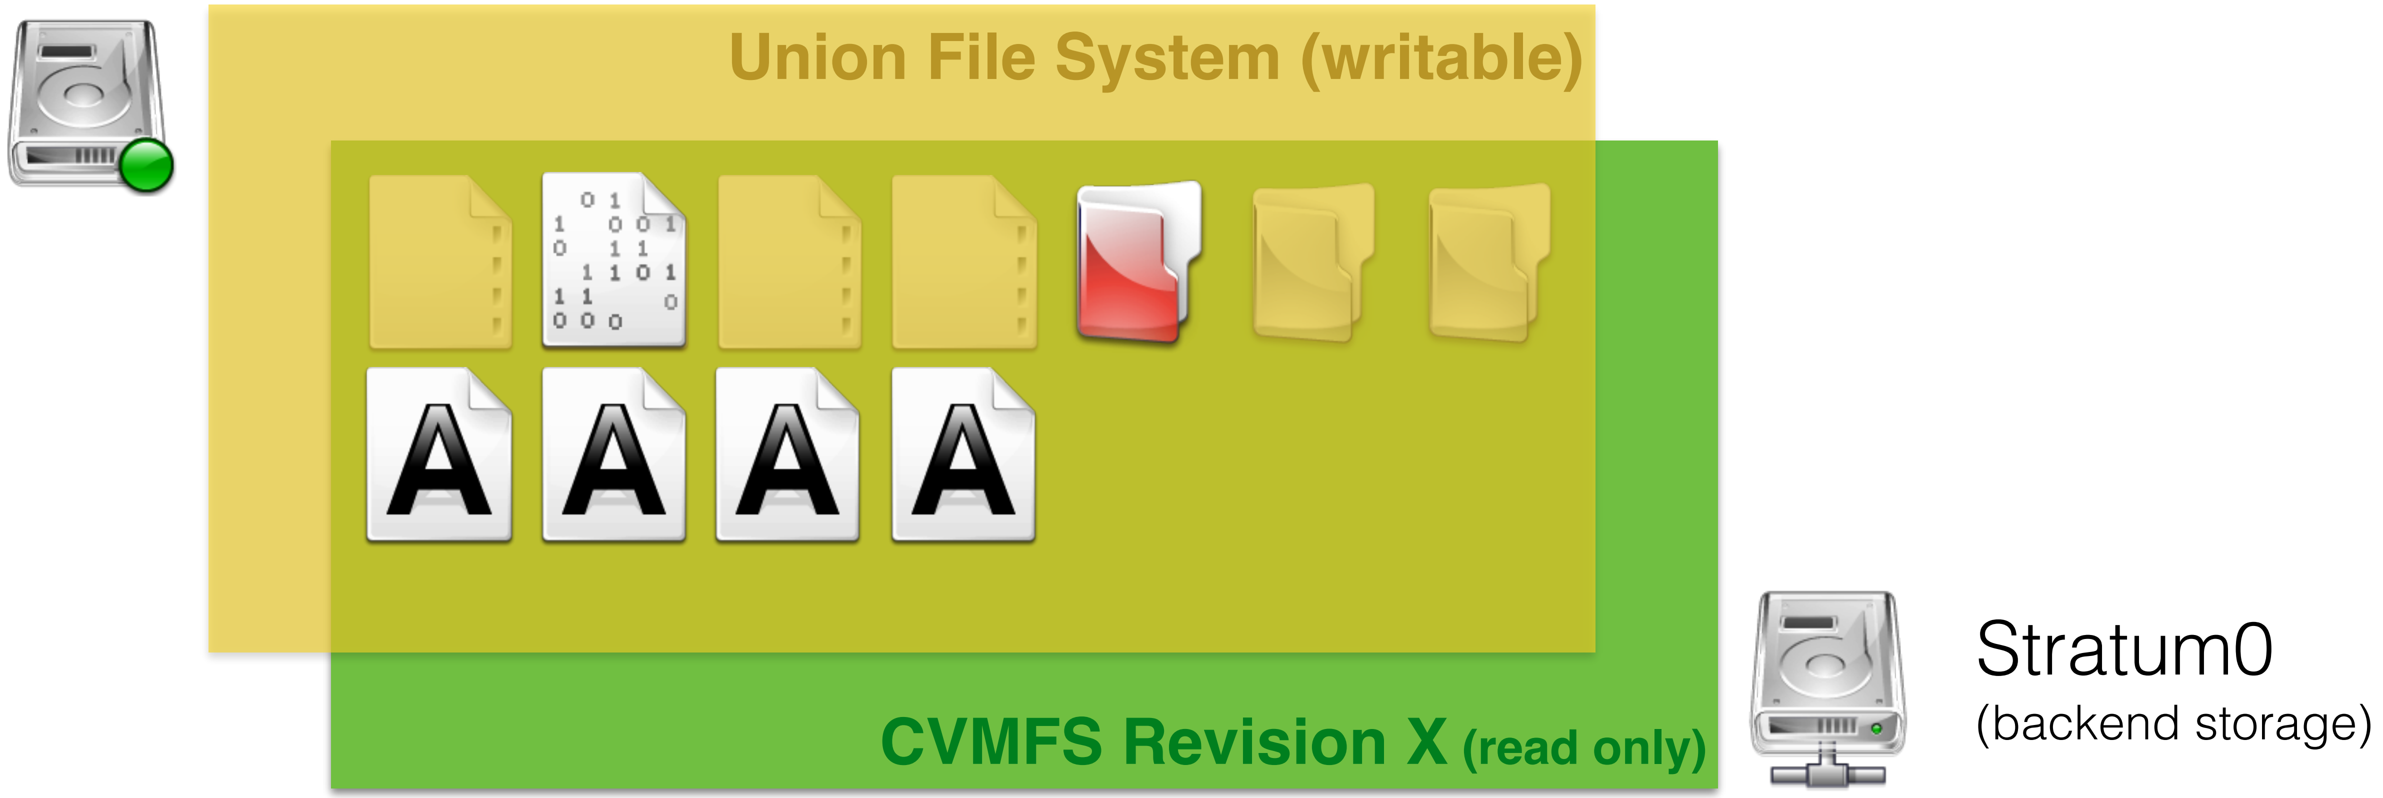

Support for Multiple Union File Systems

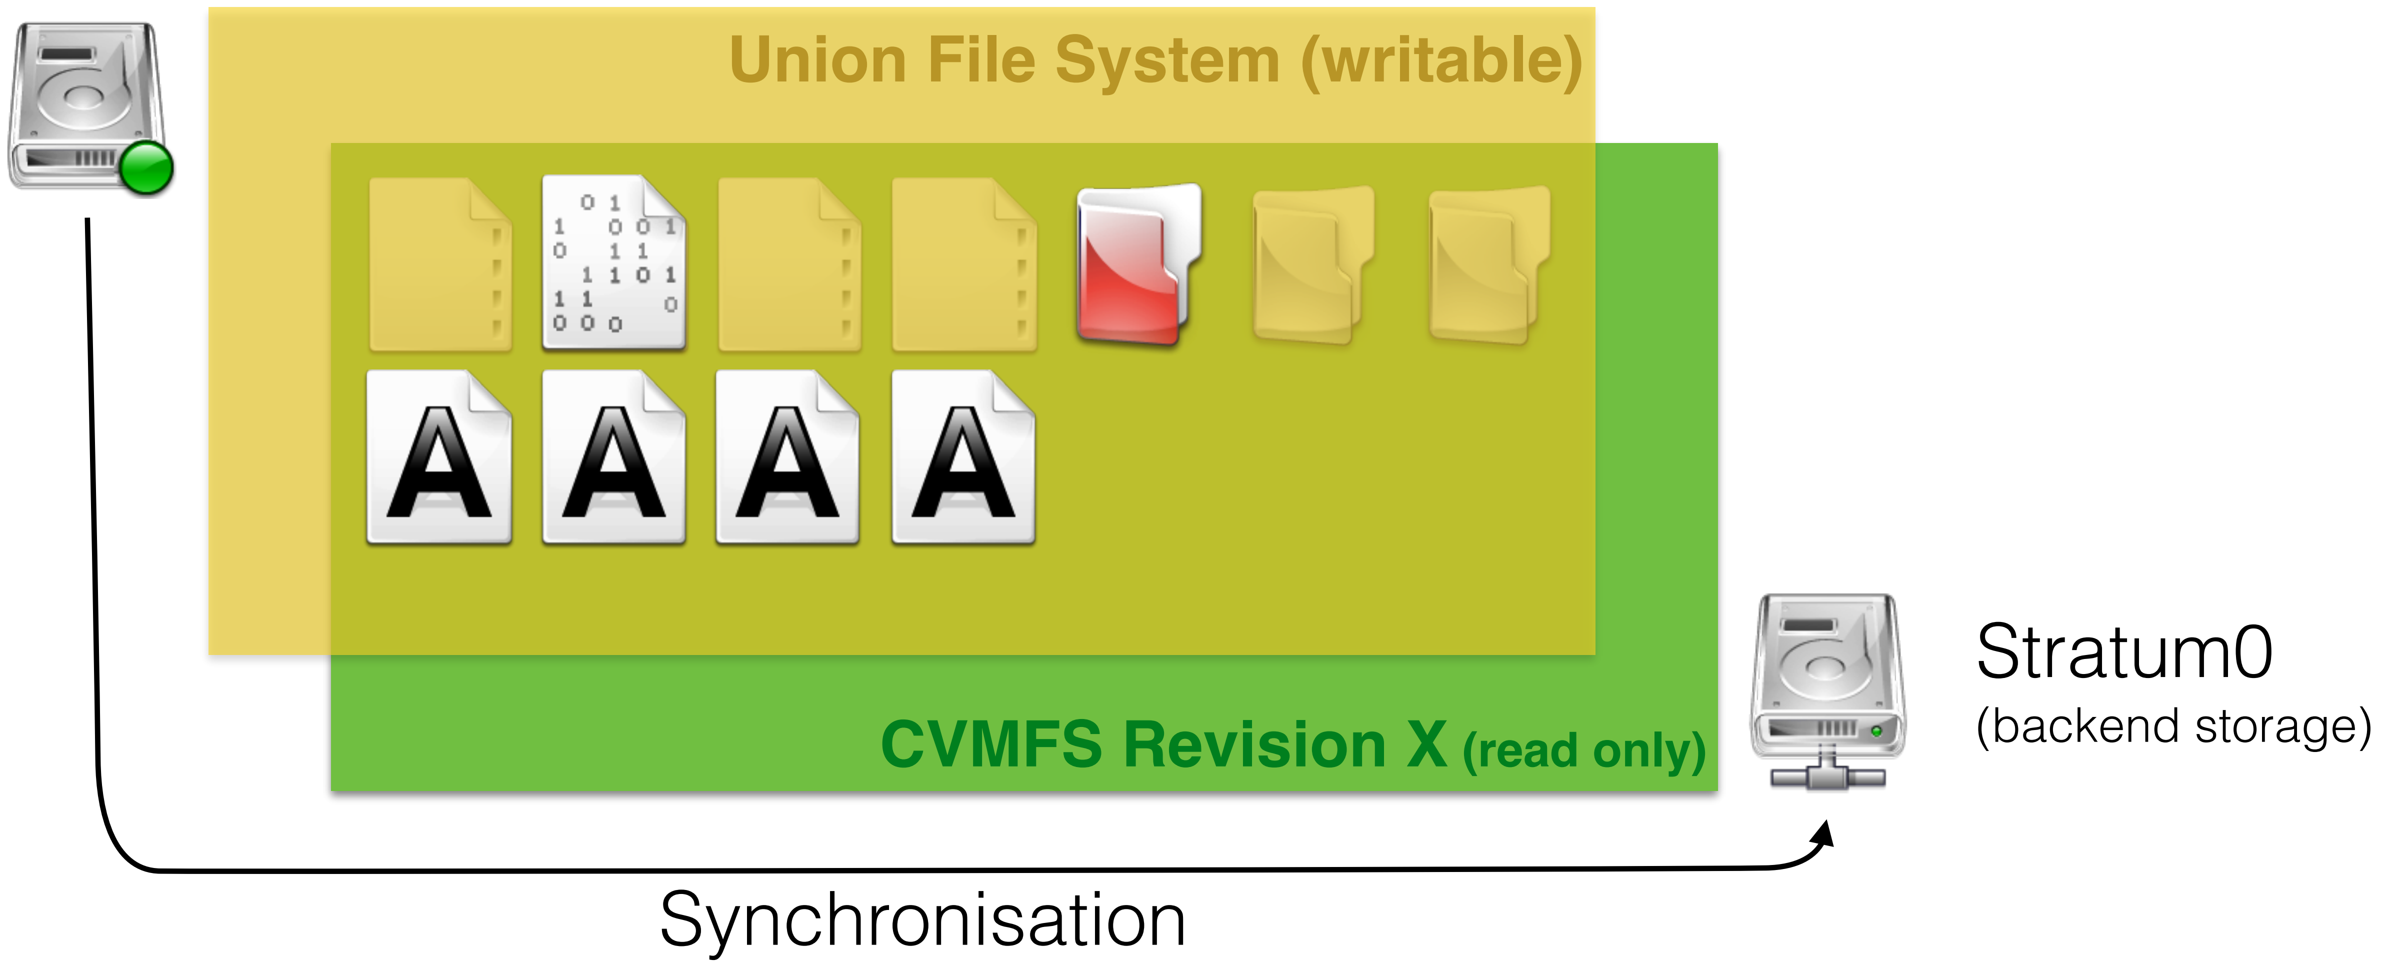

Support for Multiple Union File Systems

- **AUFS** is used at the moment
  - Comes closest to POSIX file system semantics
  - Officially supported by RedHat 5 and Ubuntu
  - Unfortunately not part of RedHat 6 and 7
    - We provide a patched kernel following upstream releases

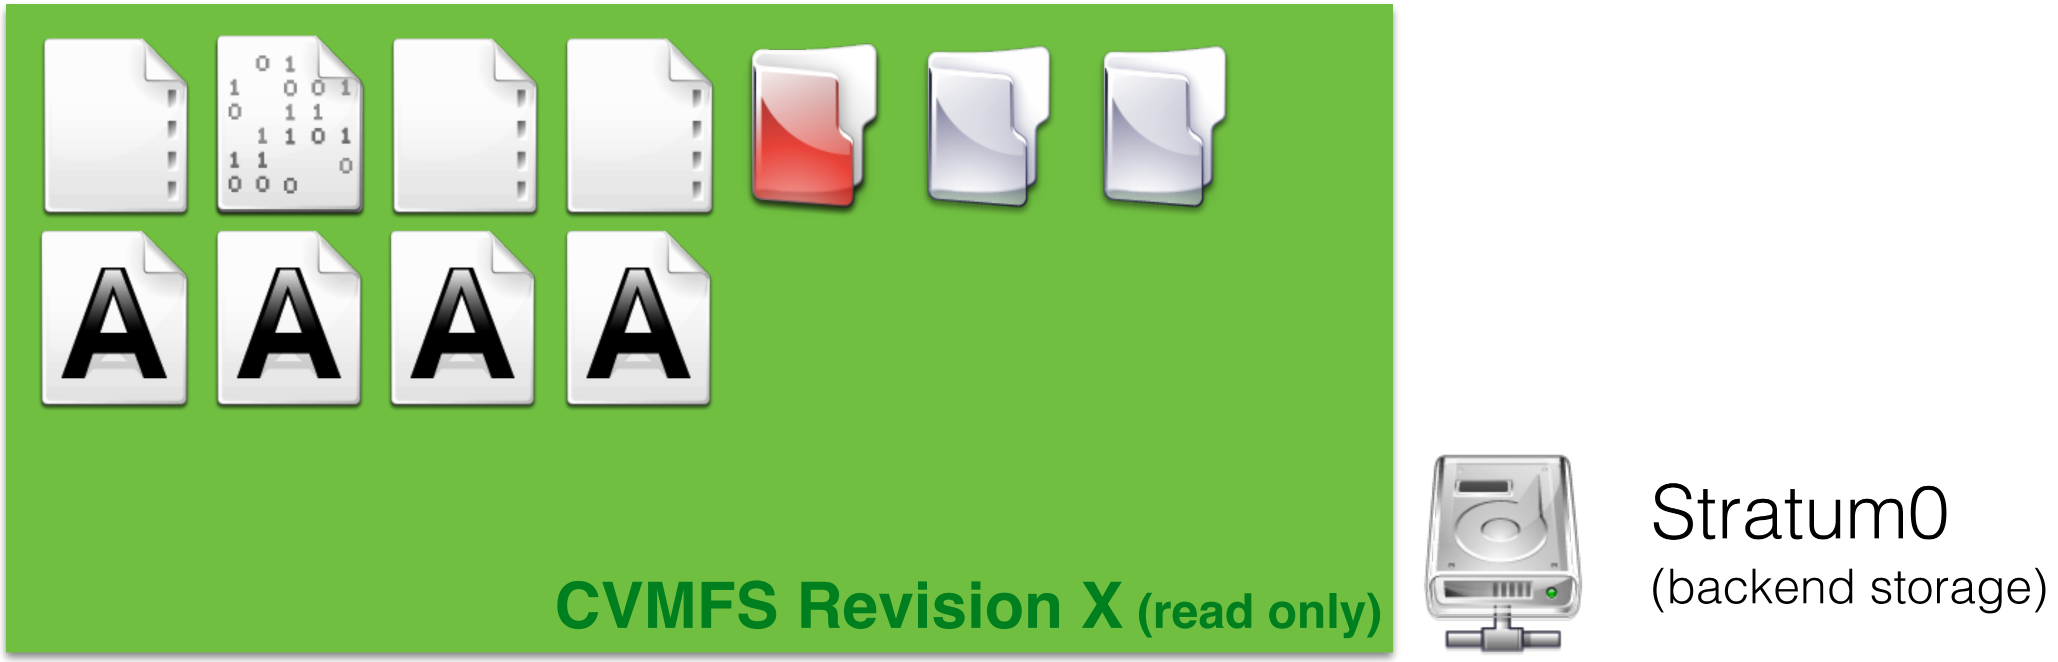

Support for Multiple Union File Systems

- **AUFS** is used at the moment
  - Comes closest to POSIX file system semantics
  - Officially supported by RedHat 5 and Ubuntu
  - Unfortunately not part of RedHat 6 and 7
    - We provide a patched kernel following upstream releases

- **OverlayFS** was merged into the Linux kernel 3.18
  - CernVM-FS has (currently untested) support for OverlayFS (contributed by Joshua C. Randall in 2013)
  - We will work on supporting OverlayFS for the server
  - Unfortunately it breaks POSIX compliance (hard links)

Customisable Release Manager Tools

- Right now: `cvmfs_server` is used by repository maintainers
- Planned: **REST API support** to remote control a Stratum 0
- Use Cases: **CernVM-FS as a Service** (i.e. OASIS)

- Rough Requirements:
  - Scripted installation and publish
  - Pre-publish transaction inspection
  - Instrumented publishing process
  - Auto-create nested catalogs

- Details are to be discussed
  - Feedback appreciated!

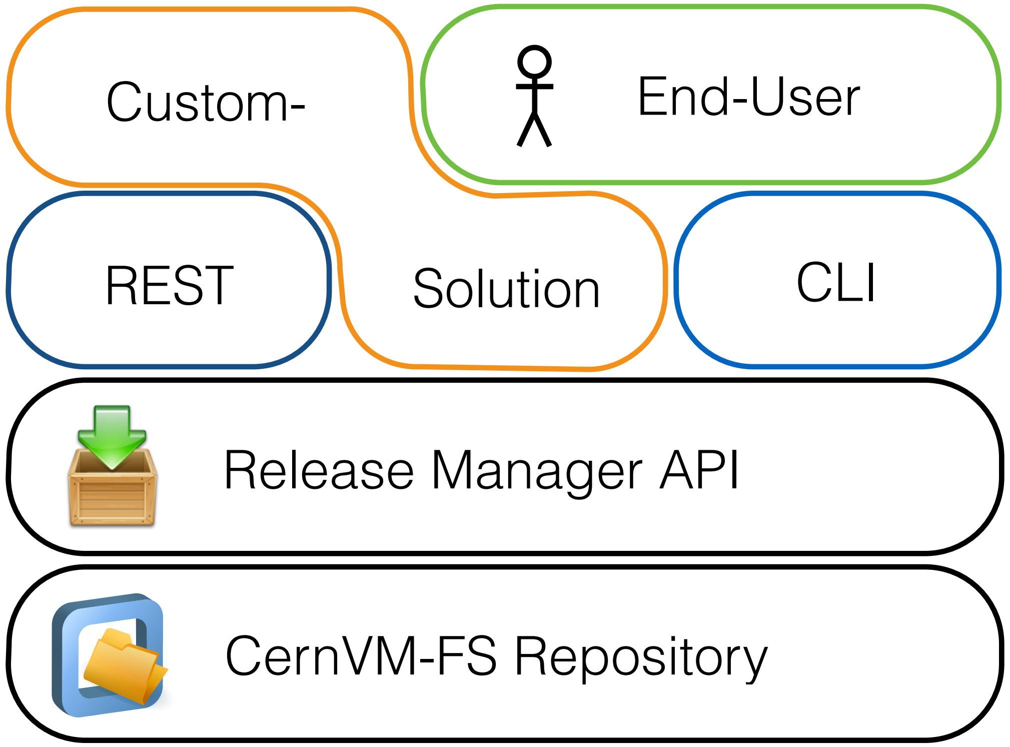

Smart Stratum1 Servers

- Carefully **enhance Stratum1 servers** with a REST API
- CernVM-FS 2.1.20's **Stratum1 GeoIP ordering** is the first step
- Potential Use Cases
  - **Push Replication**
    - Stratum0 notifies the Stratum1s about published revisions
    - Improves update propagation latency
    - Stratum0 needs a register of Stratum1s
  - Monitoring Status Information
  - Stratum1 Peer-to-Peer Updates?

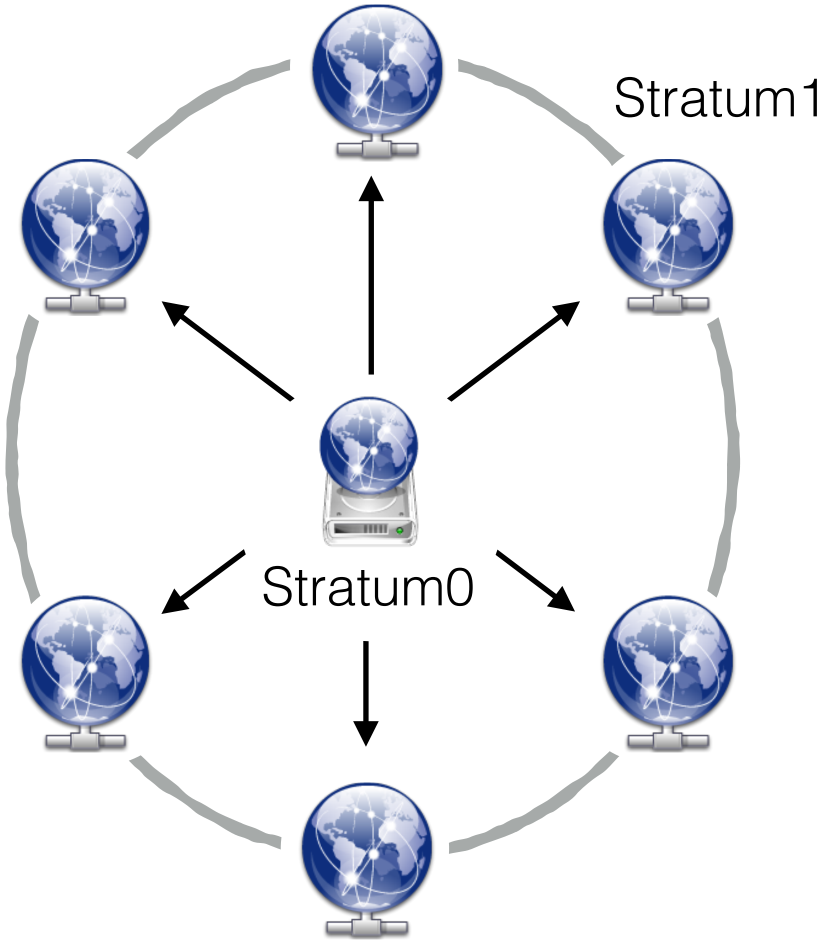

Consolidating Parrot (libcvmfs)

- Mount CernVM-FS *entirely* in user space *without* FUSE

- Main Use-Case: **Opportunistic Resources**
- Successfully used in production by CMS
- *However:* Current implementation has issues
  - No support for chunked files
  - No local cache management
  - Not covered by our test suite

- Consolidation, bug fixes and performance improvements planned for CernVM-FS 2.1.21

Release Plan

(Future) Release Plan

CernVM-FS **2.1.19** *28th of May 2014*

- Still current stable release
- Was mainly a consolidation release
- No known major issues

CernVM-FS **2.1.20** *March 2015 (published for testing*)*

- Feature release
  - Separation of software and configuration
  - Automatic GeoIP ordering of Stratum 1
  - History-less repositories
  - Alternative Storage Backends (i.e. S3)

CernVM-FS **2.1.21** *End 2015*

- Push Replication
- Parrot Improvements

* http://cernvm.cern.ch/portal/filesystem/release-procedure

Thanks!

Backup Slides

New Feature in CernVM-FS 2.1.x

Repository Import

No Need for a Shadow Directory

Stratum 1 Aliases

Revision History Database

File System Snapshotting

pkcs#7 signed whitelists

Parallel File Processing

Transactional Repository Updates

Alien Cache

Compatibility to ARM

RipeMD 160 Hash Support

S3 Compatible Server Backend

Filesystem Snapshot Rollbacks

Shared Cache Manager

Multiple Repositories on one Release Manager Machine

Proxy Auto-Config Support

Large File Chunking

Aggregated Catalog Statistics

GeoIP-Based Stratum1 Ordering

Post-Mortem Stack Generation

Pluggable Storage Backends

Client Hot-Patching

Repository Garbage Collection

NFS Exports

Named File System Snapshots

Micro-Syslog

Read-only Cache Mode

Configuration Meta-Repository

Volatile Repositories

Testing of CernVM-File System

- Testing is part of **continuous integration**
  - Pull Requests are built and tested before being merged
  - Nightly integration build and test
  - Regular full integration test runs (~ 16 hour runtime)
- Multiple Testing Layers
  - Unittests
  - Integration and Regression Tests
  - Stress and Performance Testing (currently not automated)
  - Gradual roll-out of new CernVM-FS releases

googletest
Google C++ Testing Framework

Jenkins

openstack™

Continuous Multi-Platform Build and Test

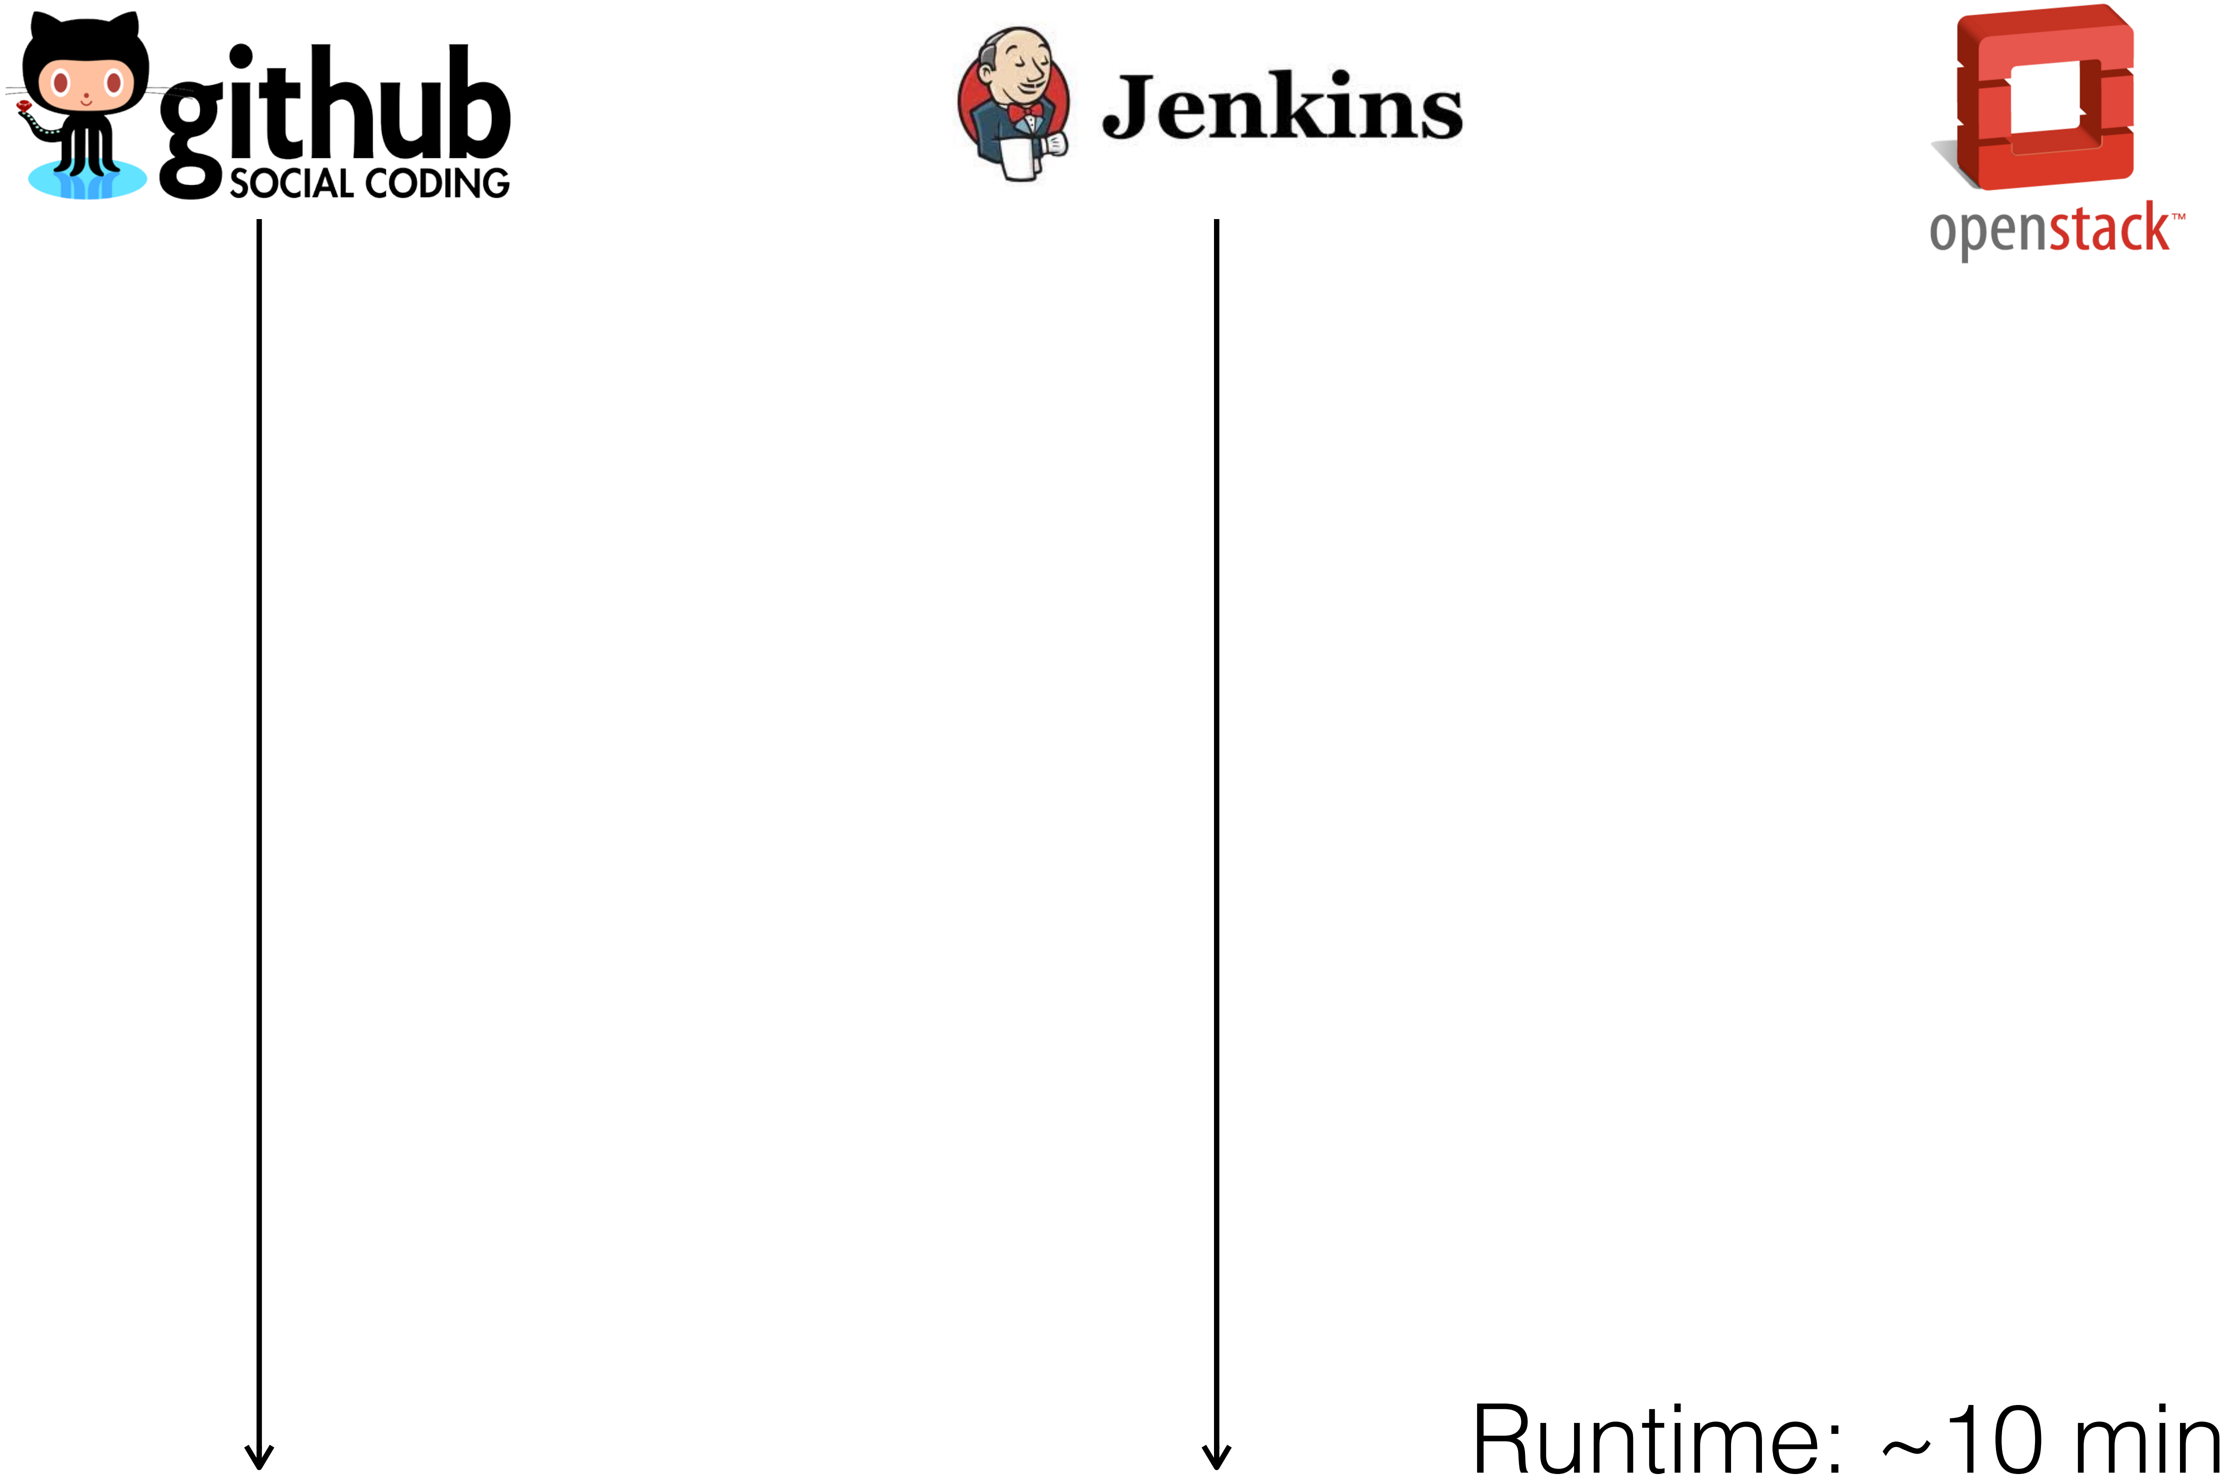

Runtime: ~10 minutes

Continuous Multi-Platform Build and Test

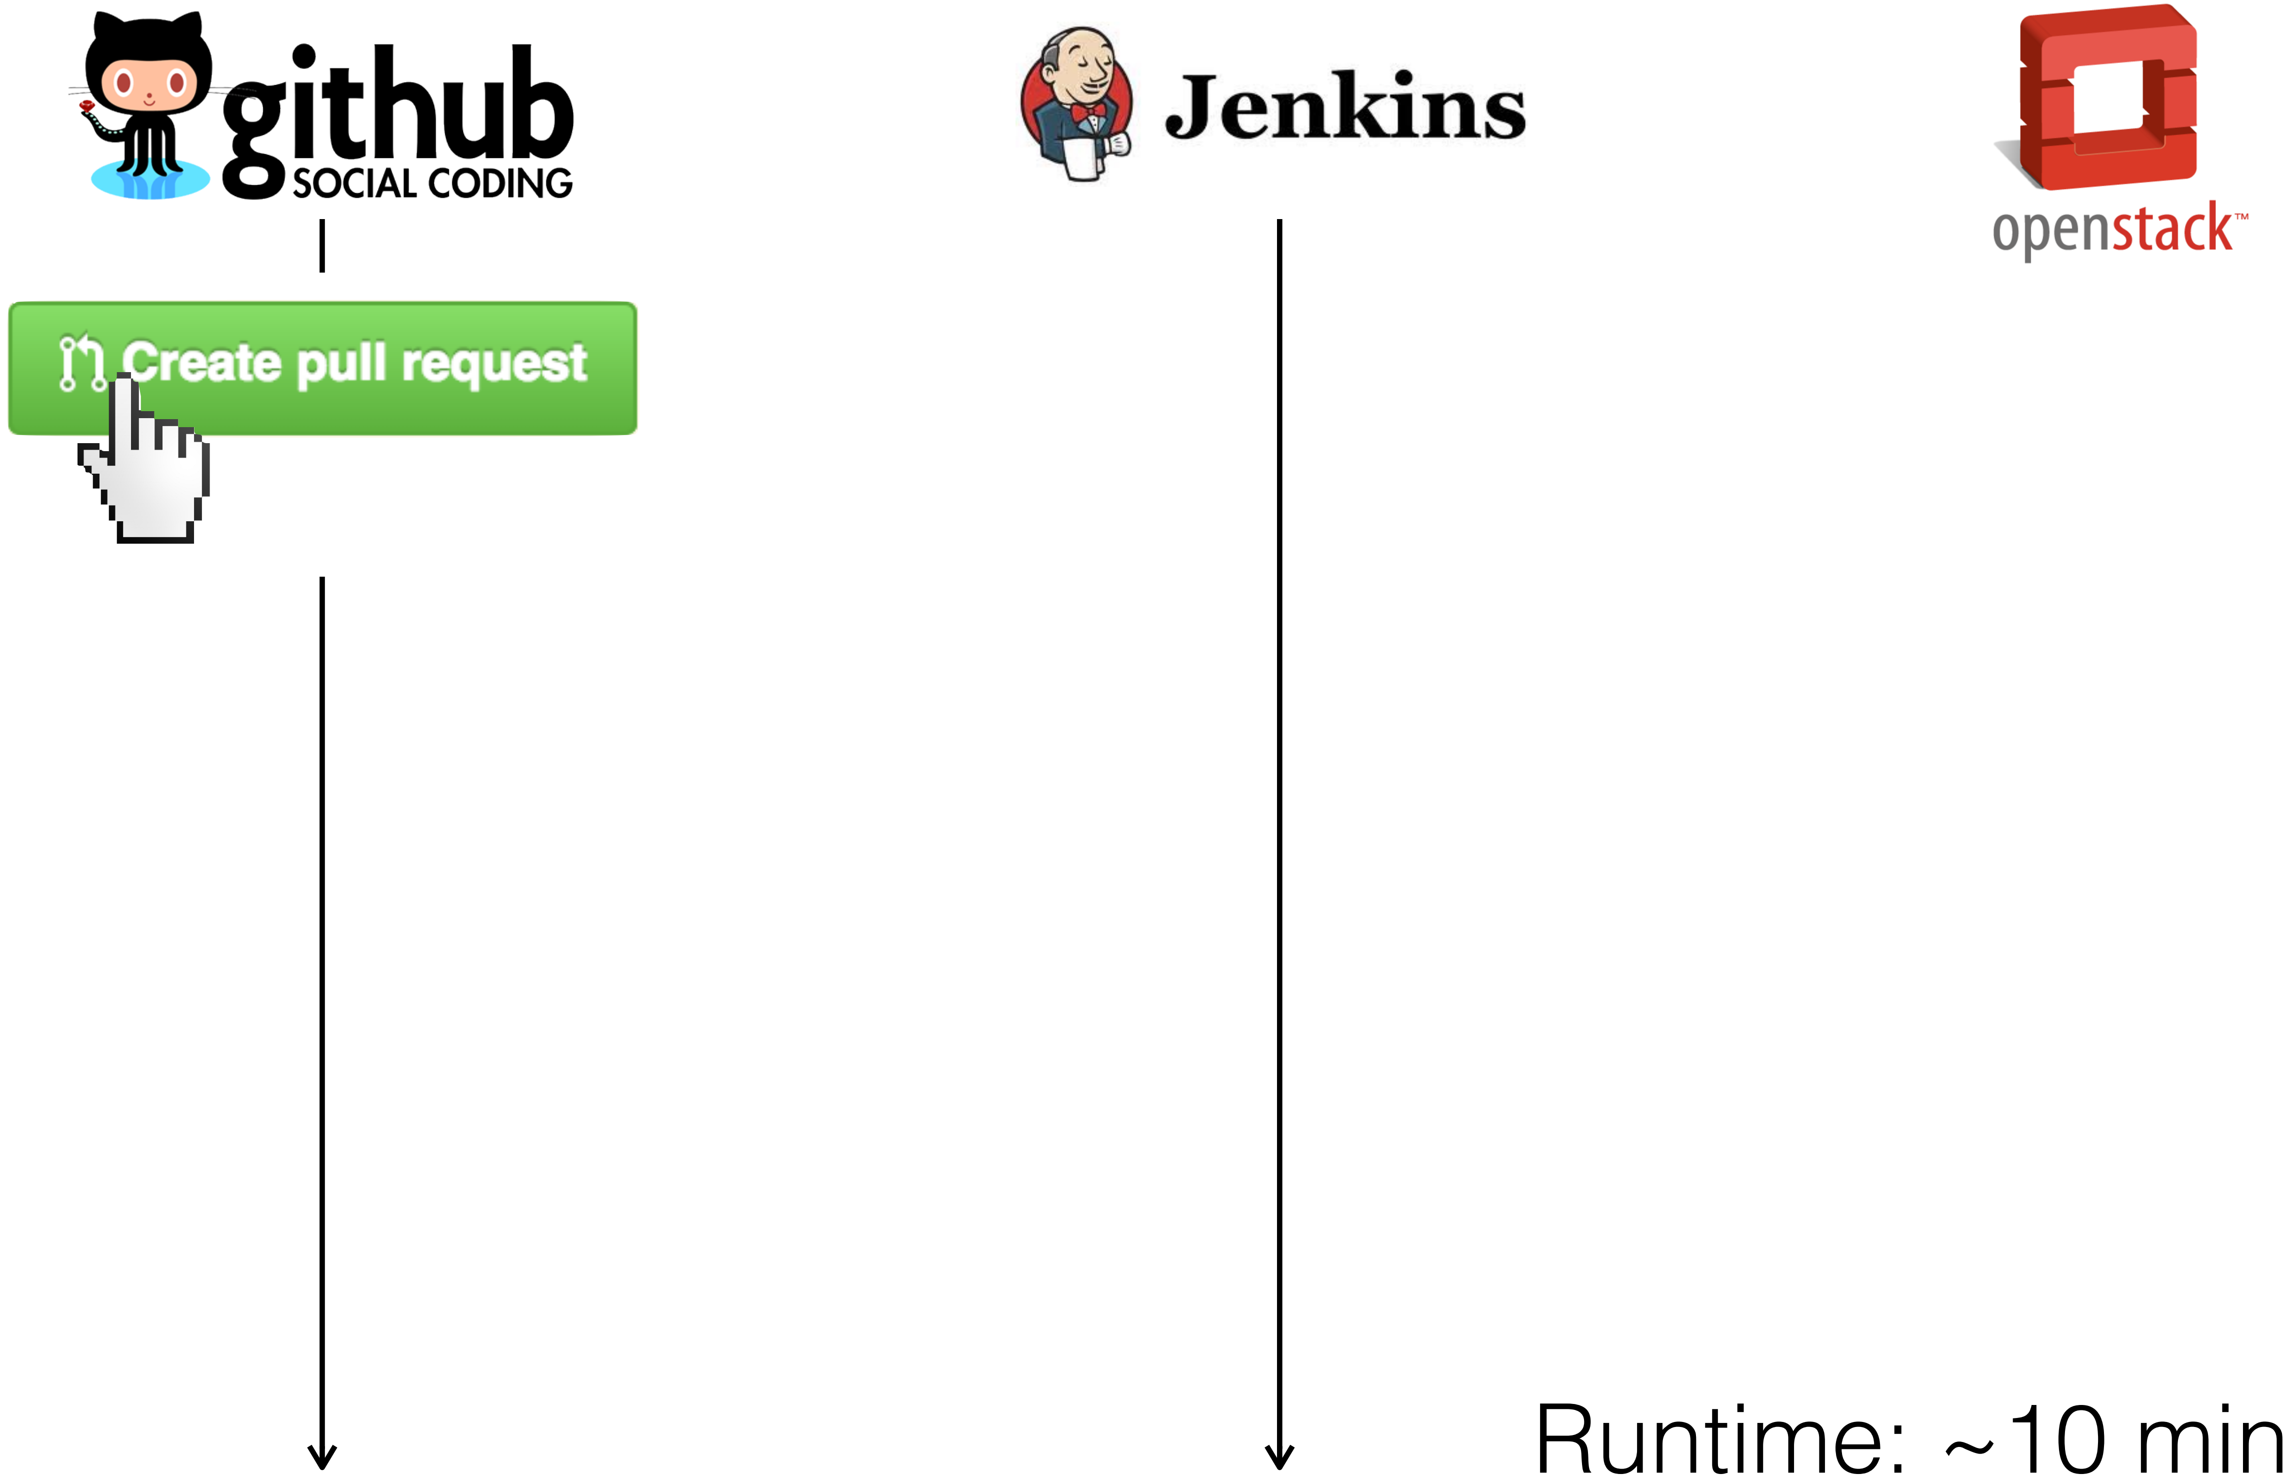

Runtime: ~10 minutes

Continuous Multi-Platform Build and Test

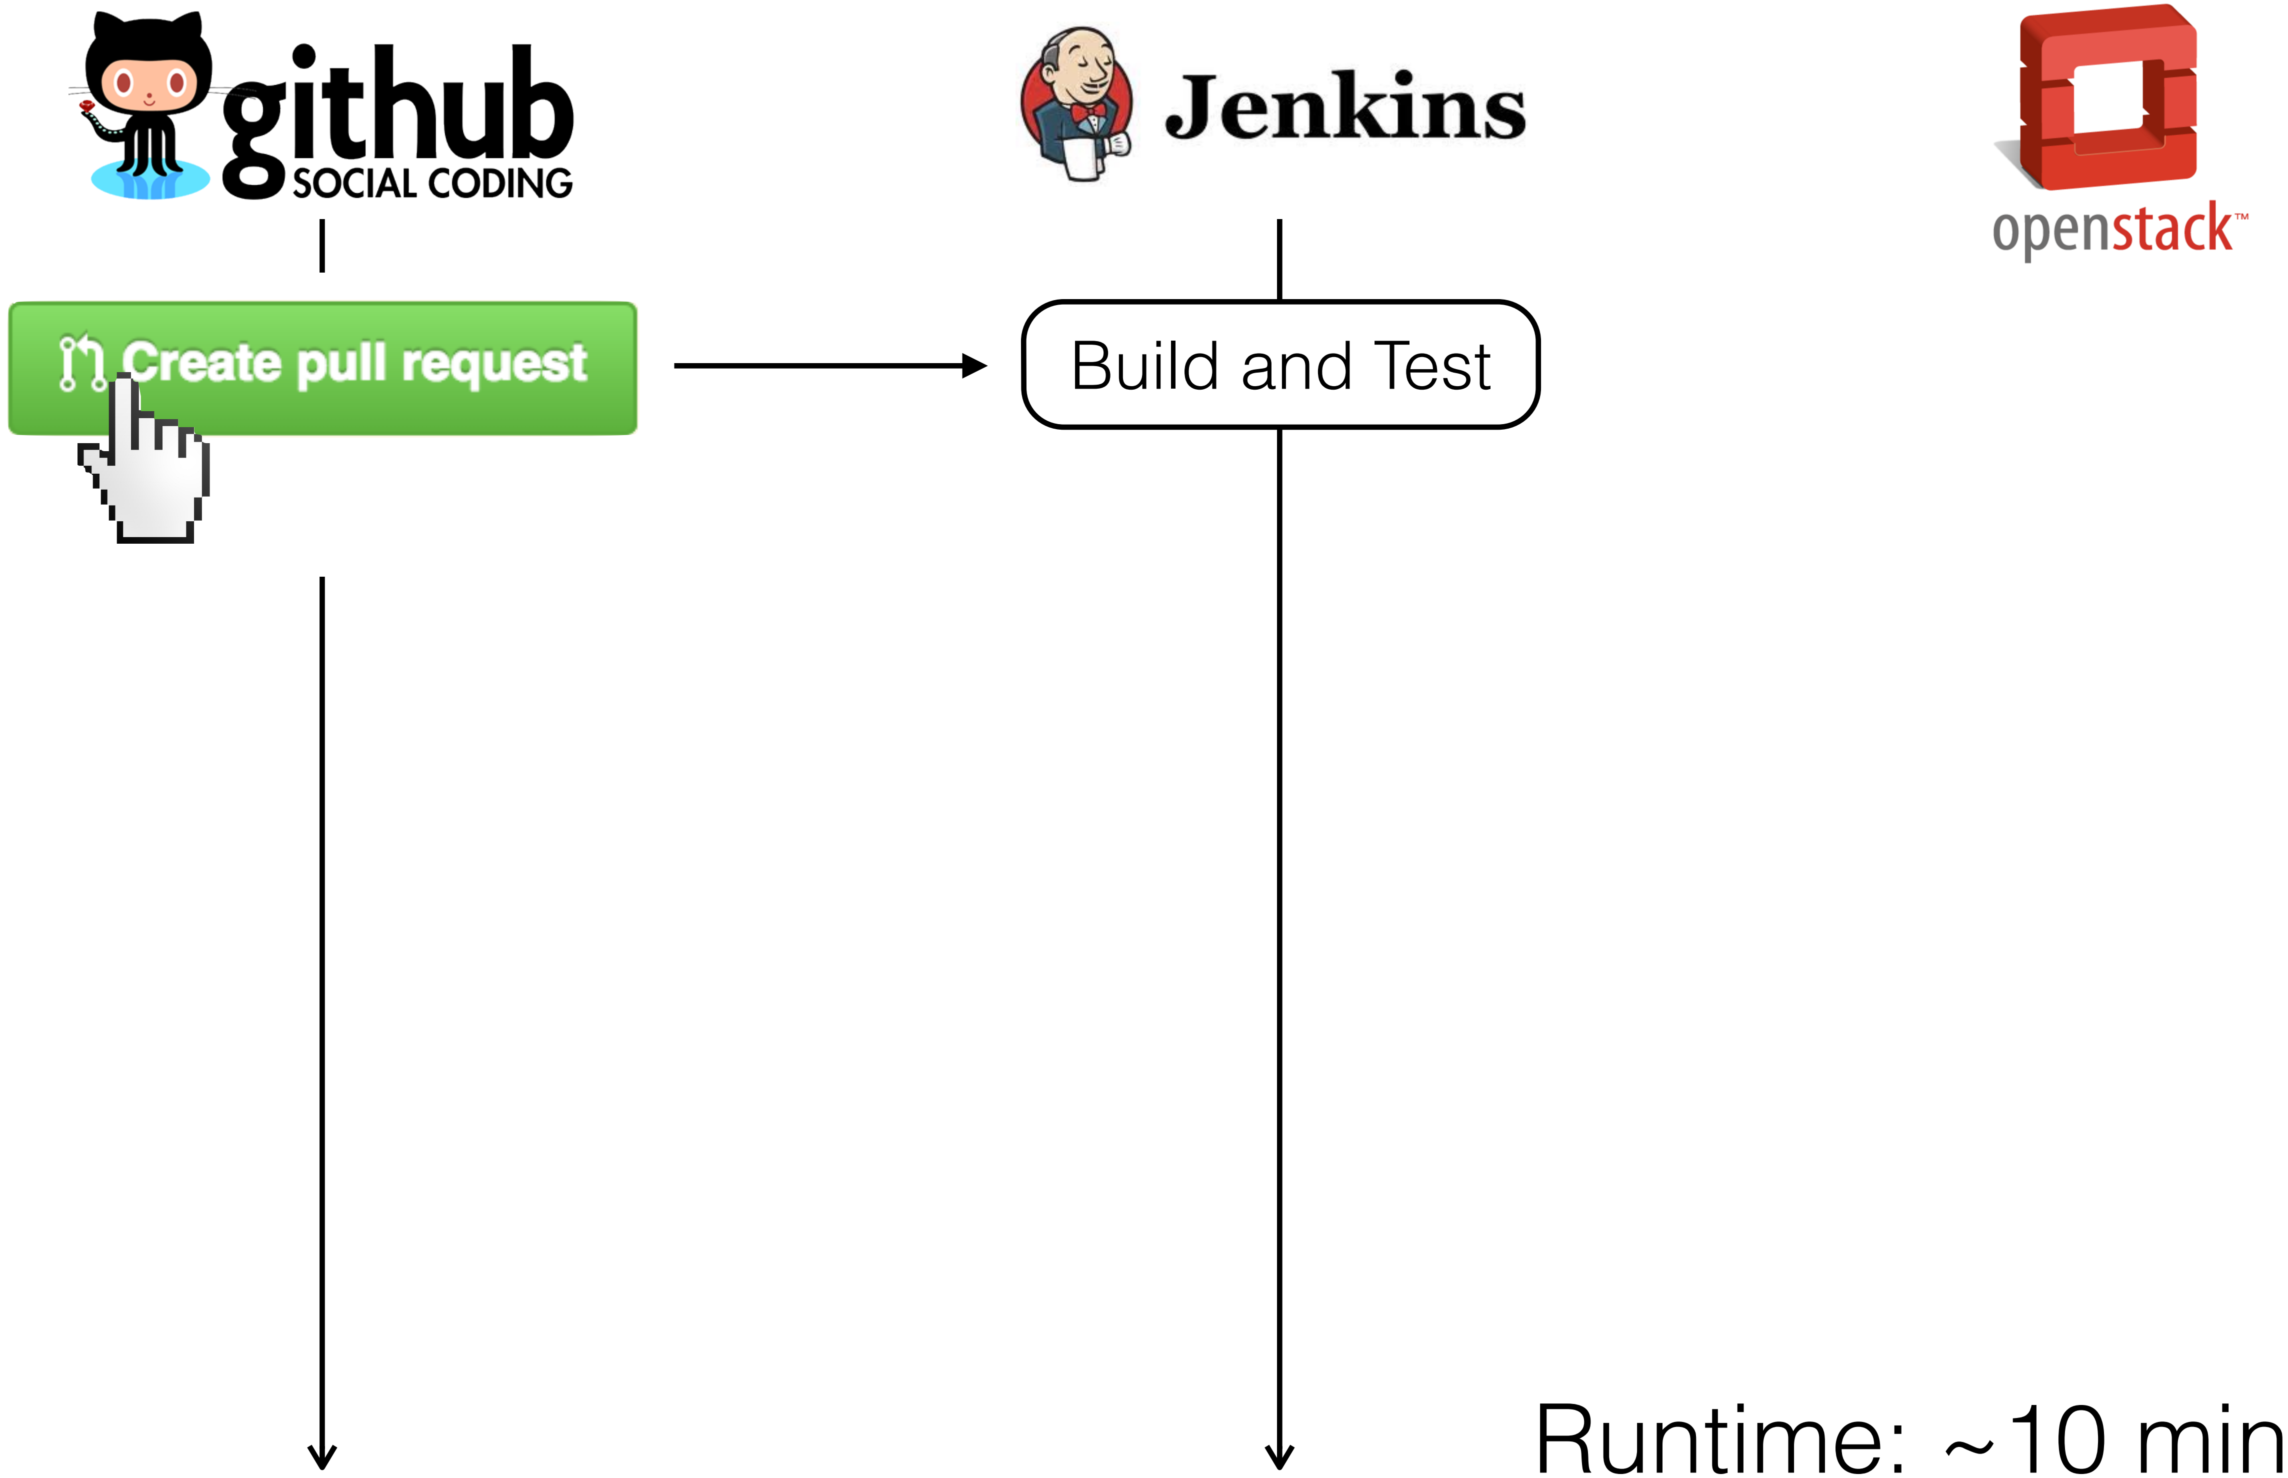

Create pull request

Build and Test

Runtime: ~10 minutes

Continuous Multi-Platform Build and Test

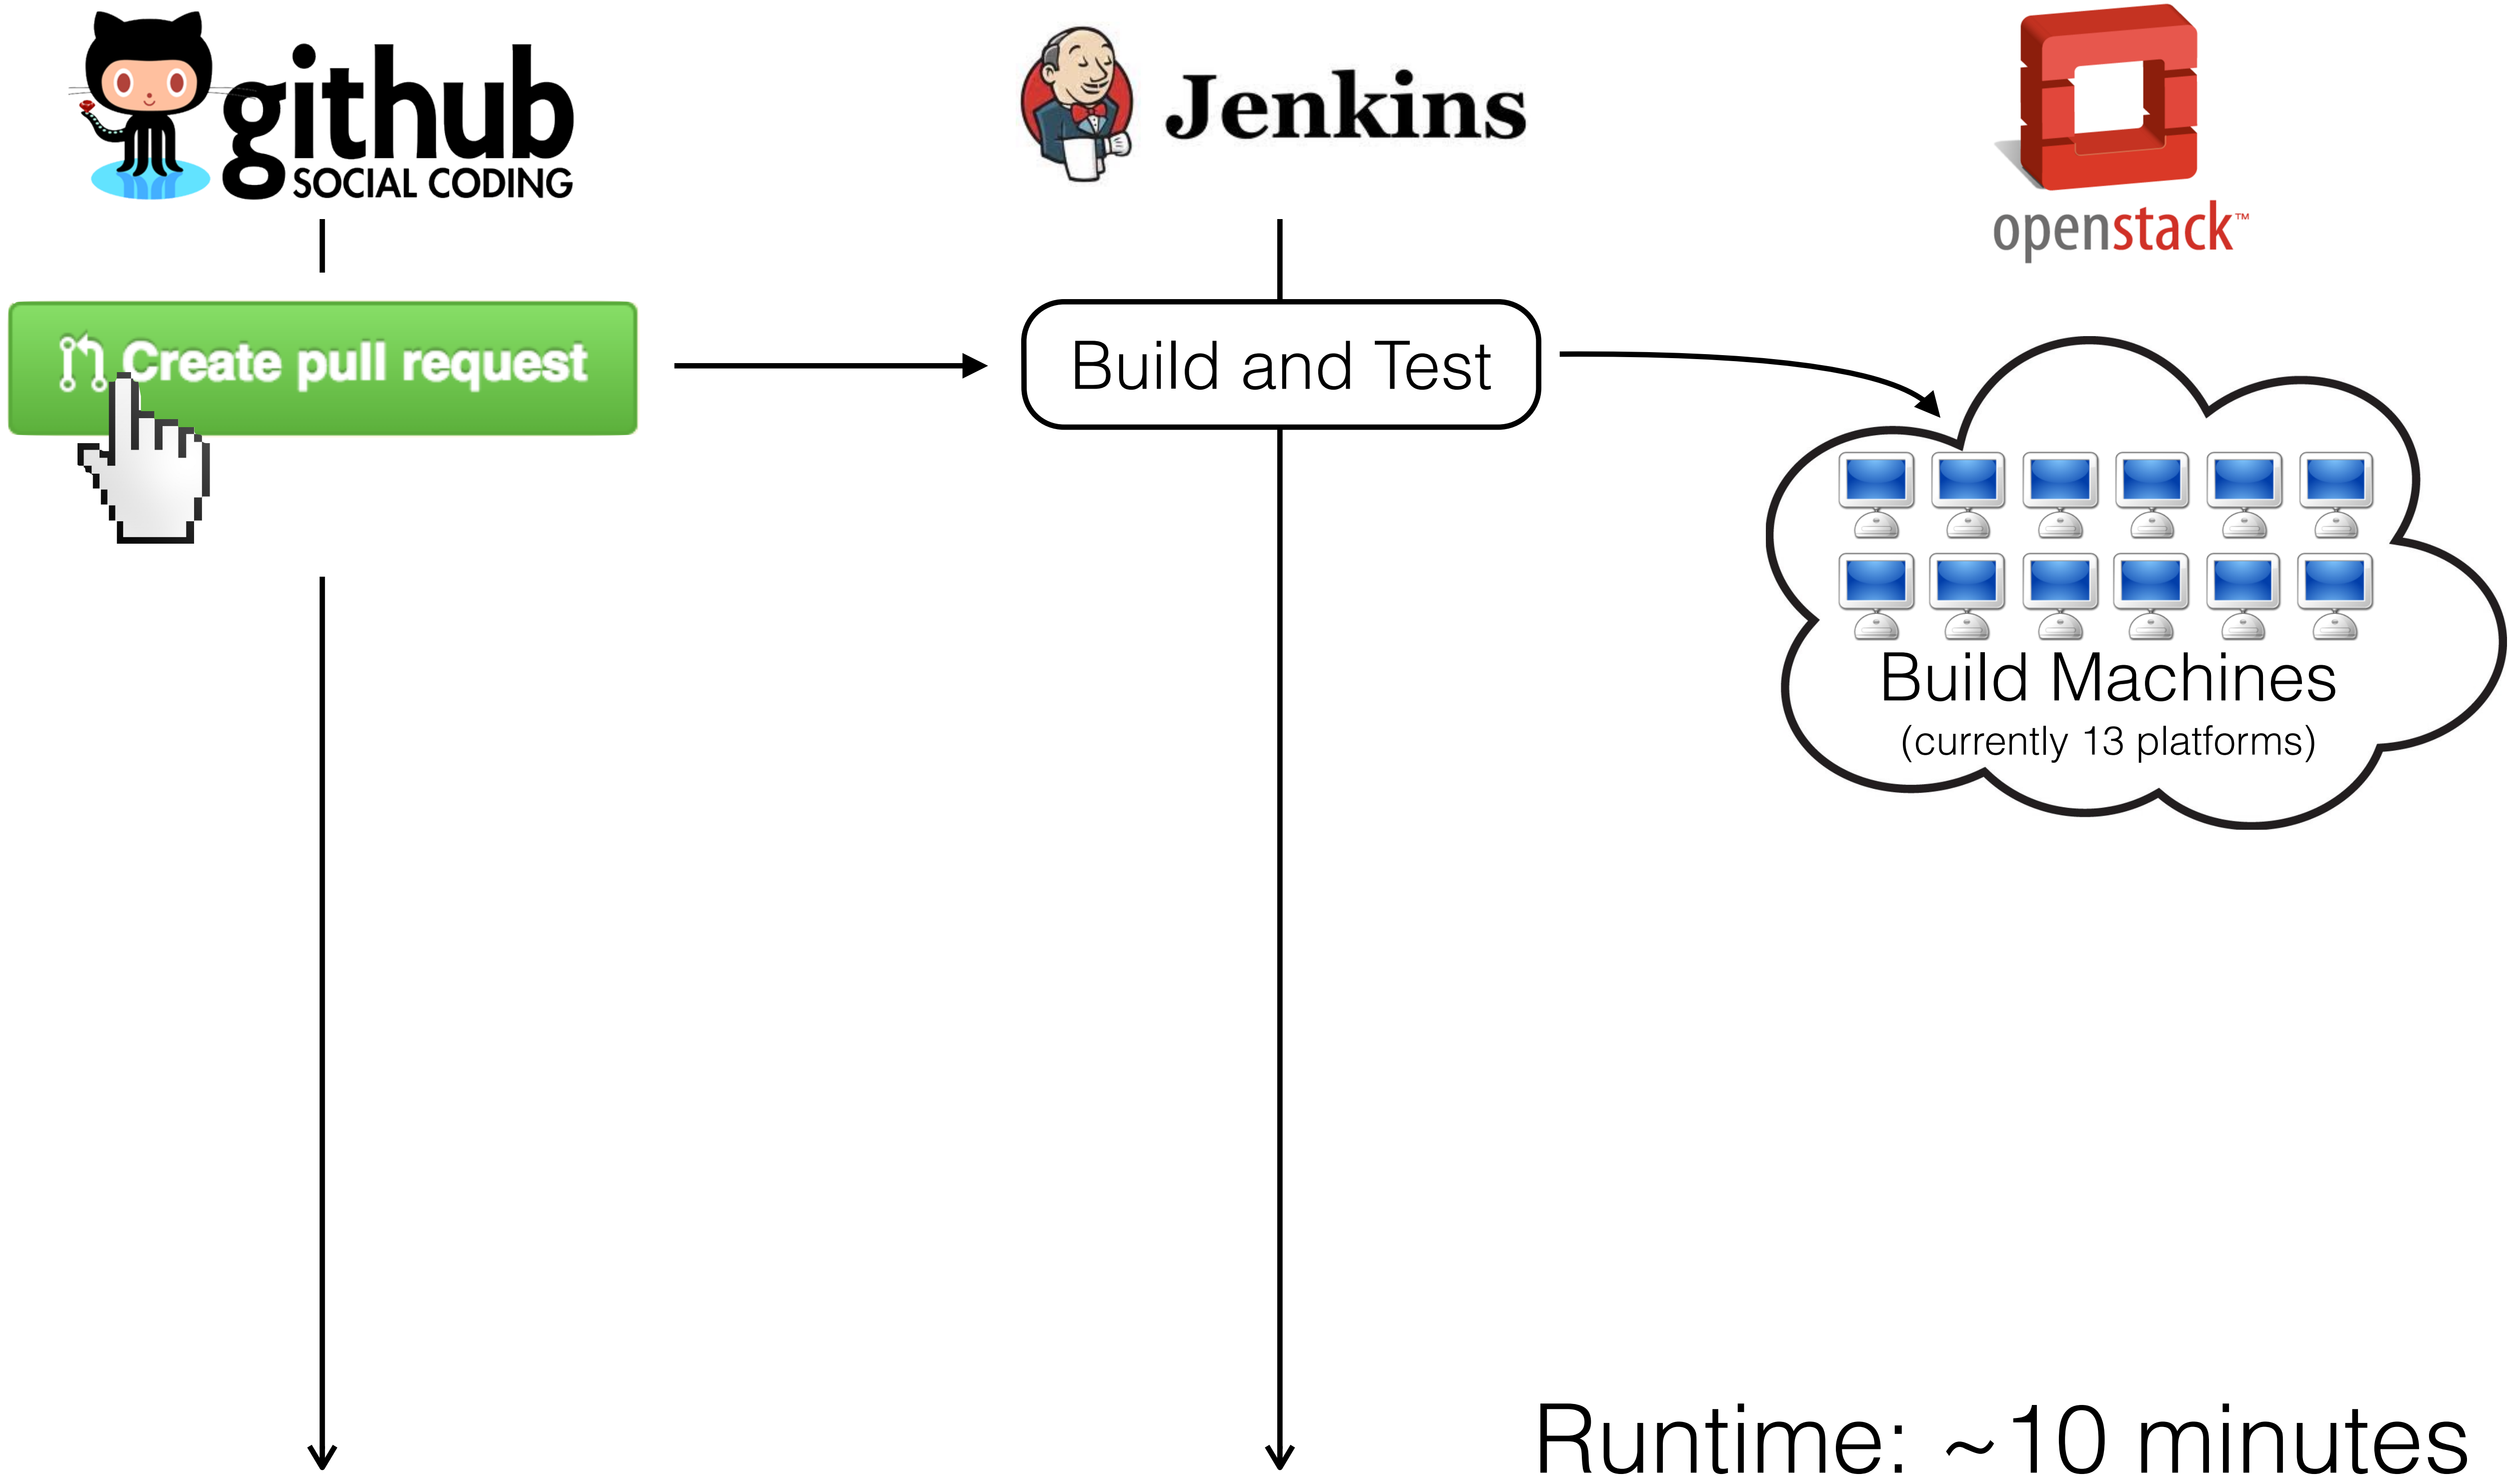

Build and Test

Build Machines
(currently 13 platforms)

Runtime: ~10 minutes

Continuous Multi-Platform Build and Test

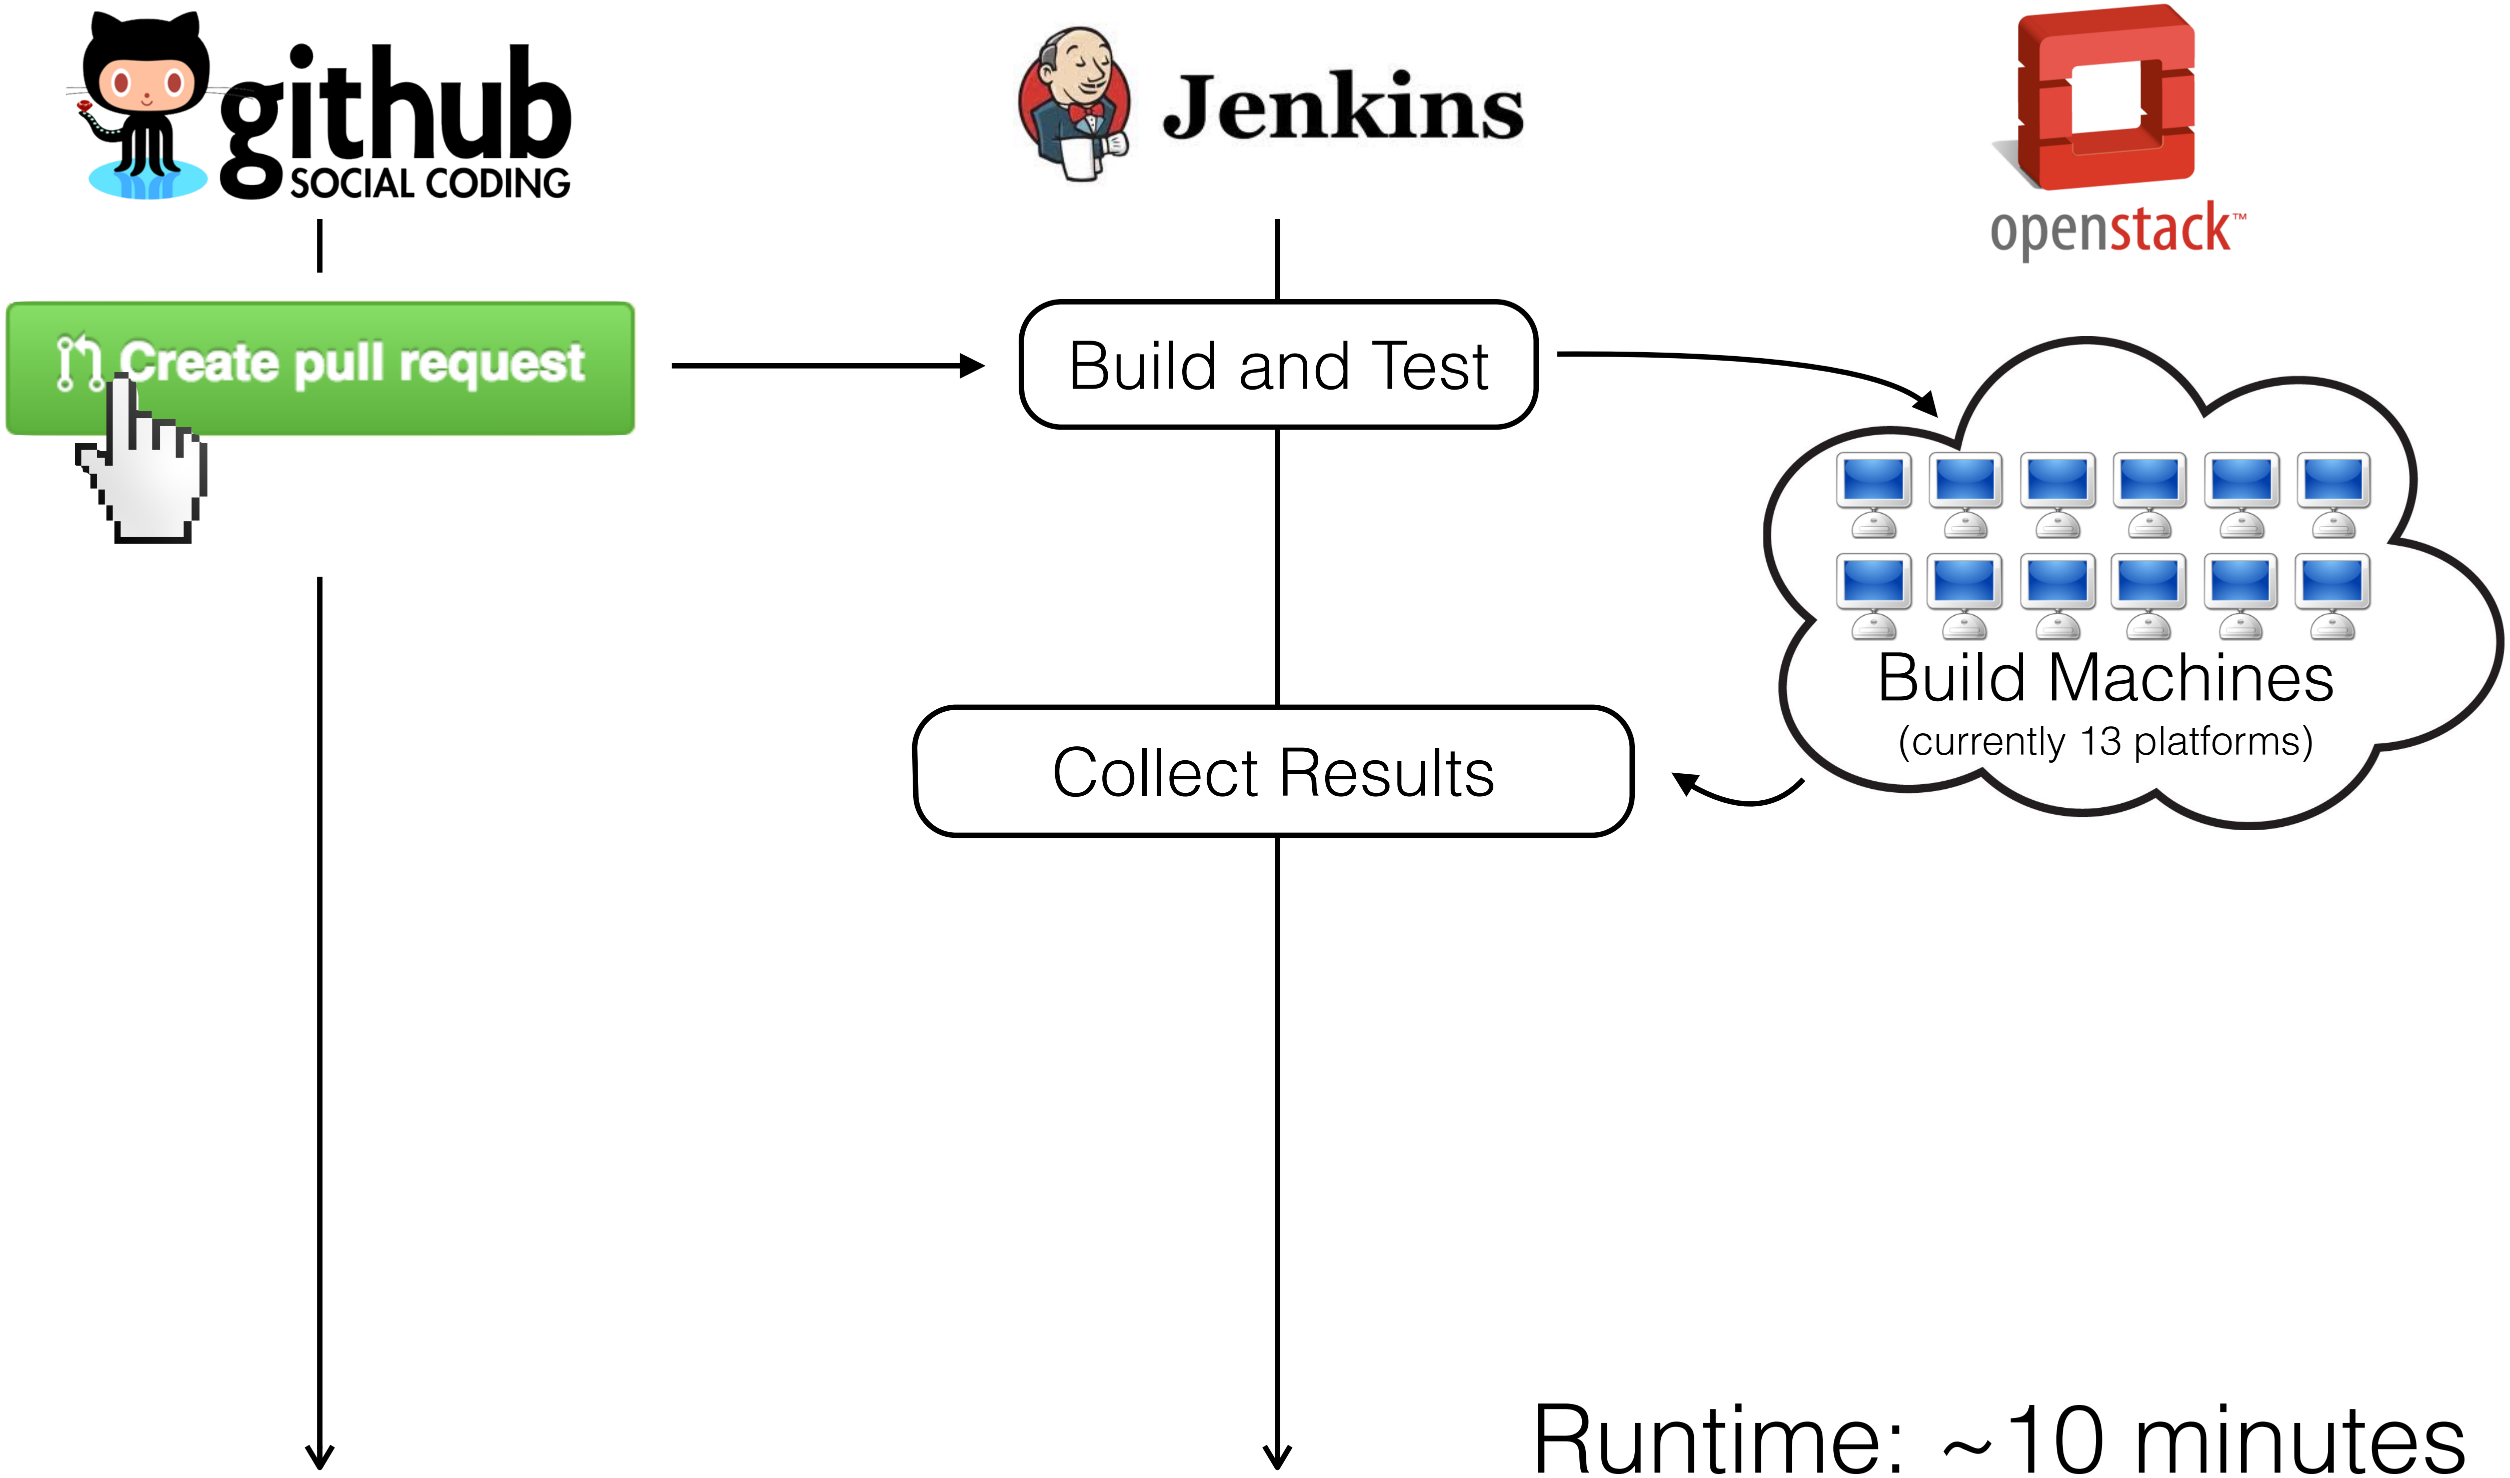

Build and Test

Build Machines
(currently 13 platforms)

Collect Results

Runtime: ~10 minutes

Continuous Multi-Platform Build and Test

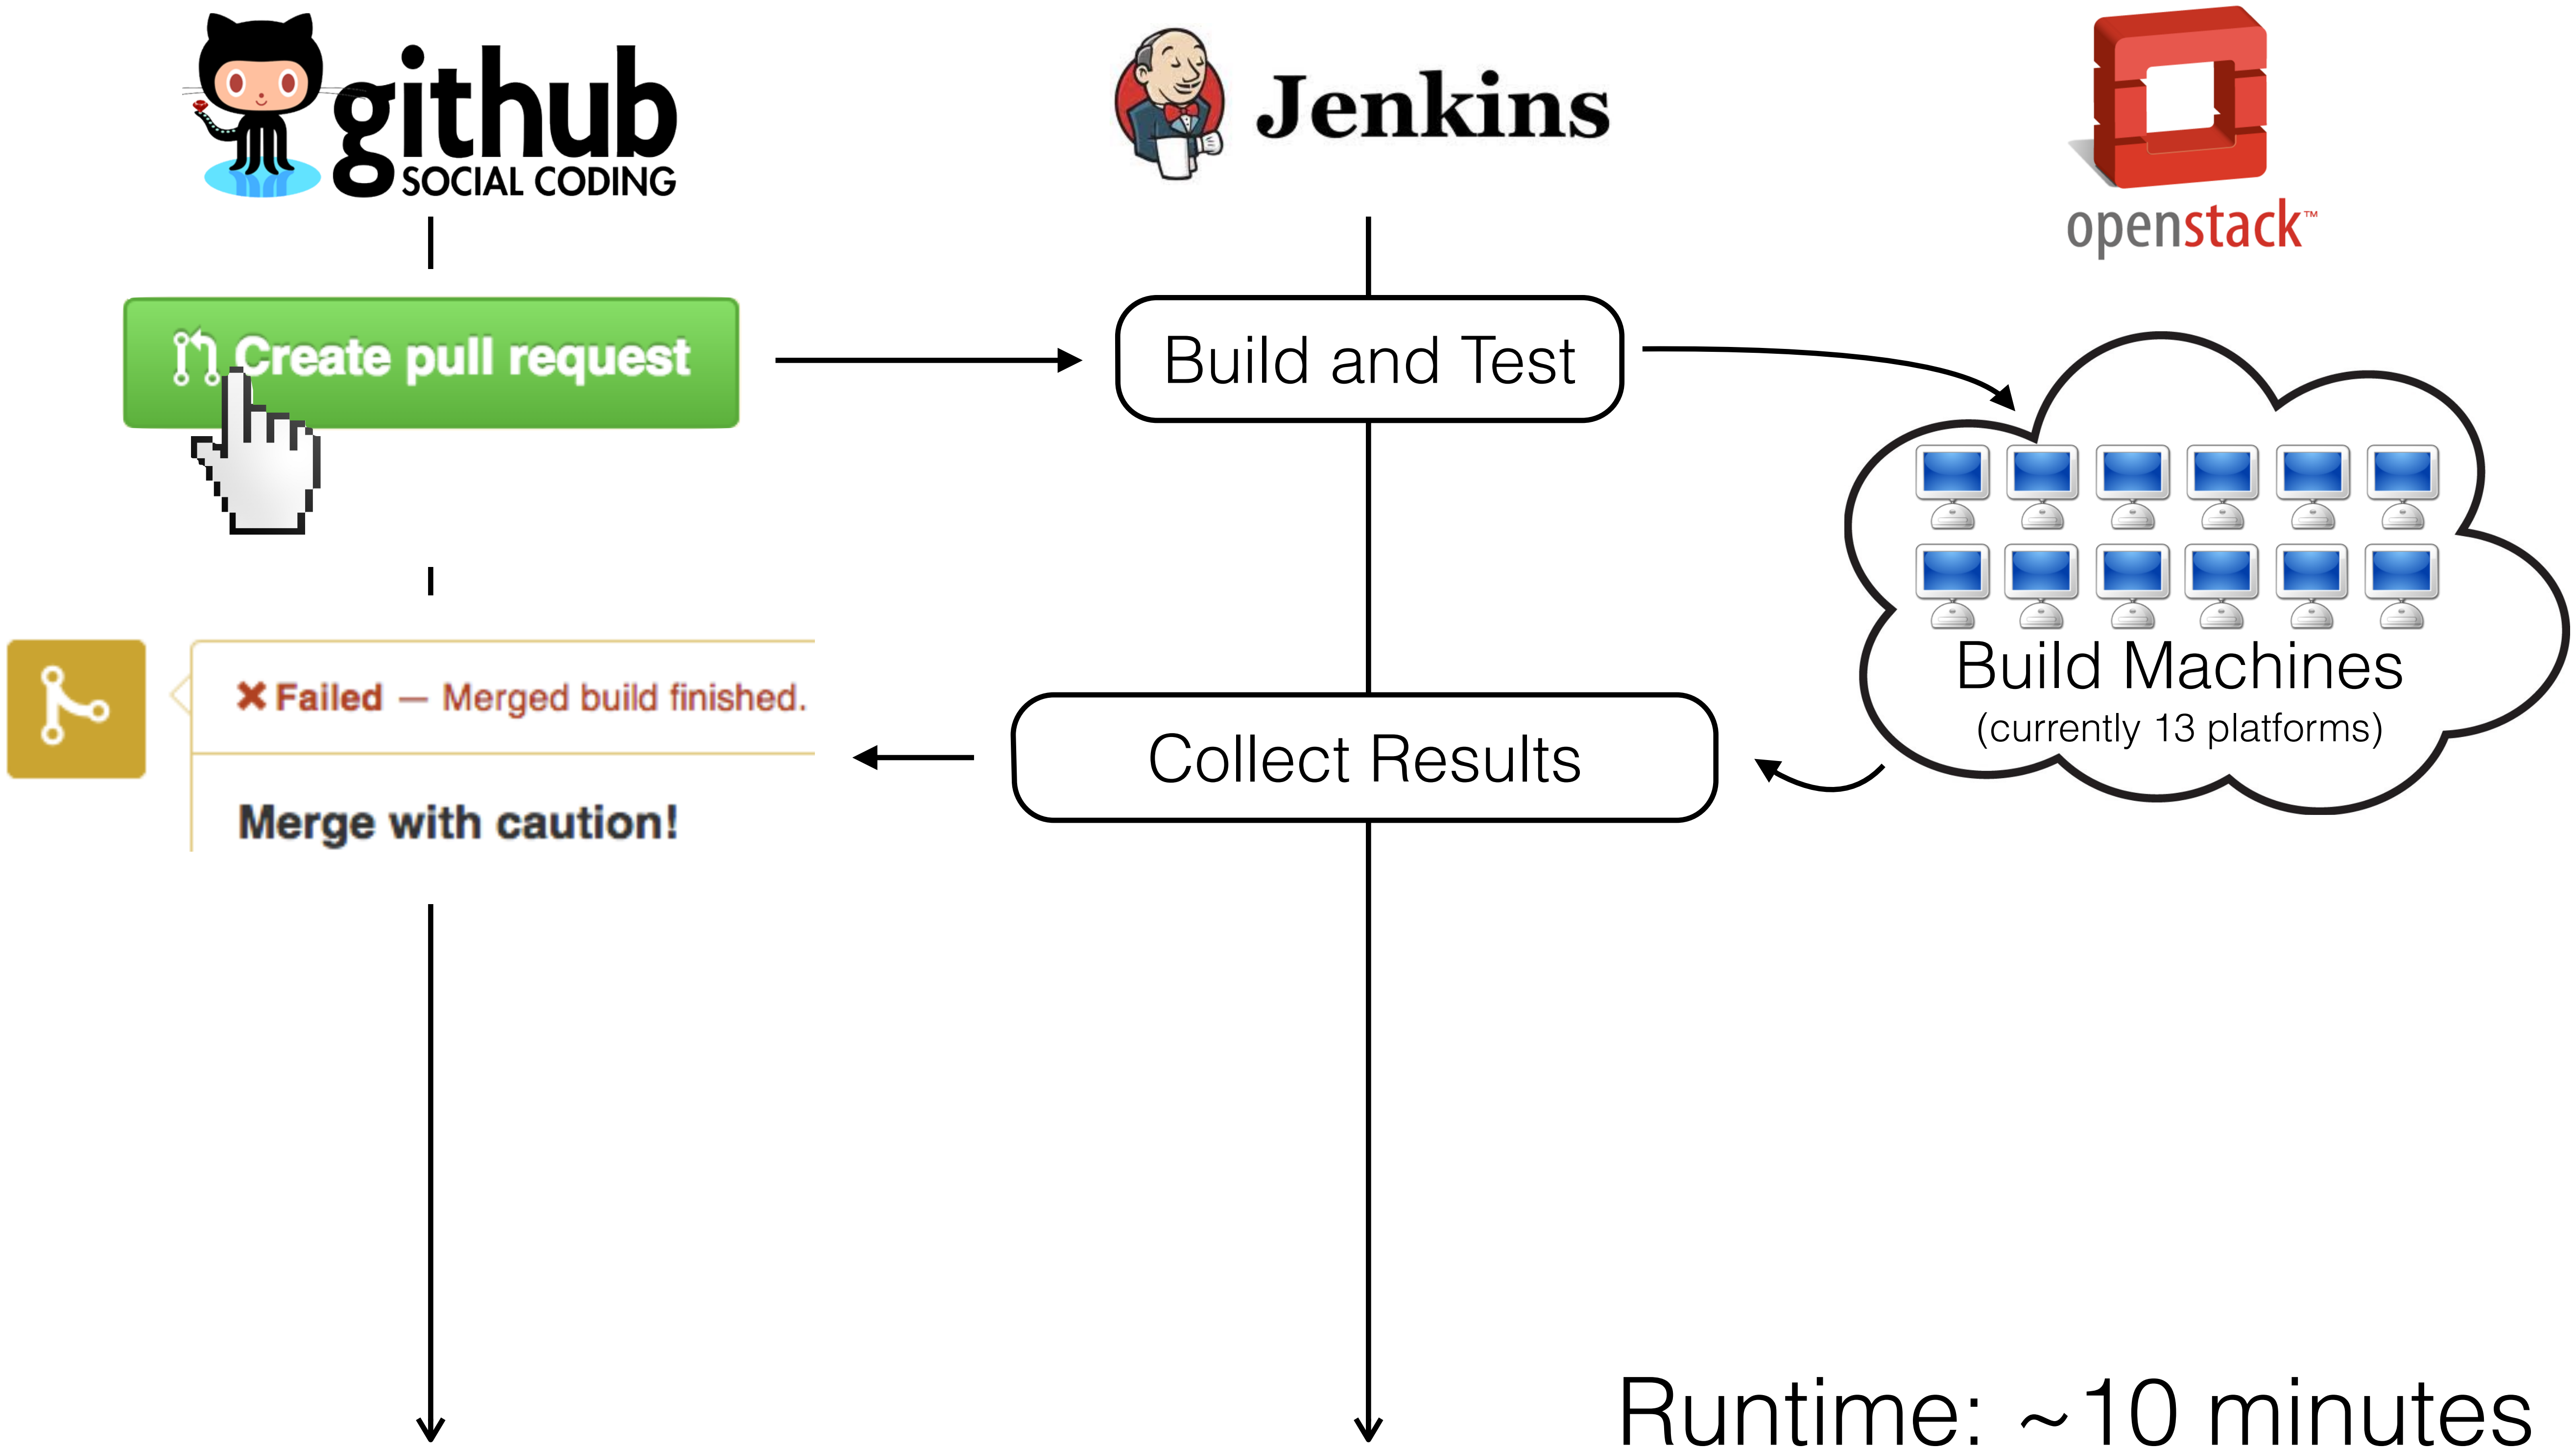

Build and Test

Build Machines
(currently 13 platforms)

Collect Results

✕ Failed — Merged build finished.

Merge with caution!

Runtime: ~10 minutes

Continuous Multi-Platform Build and Test

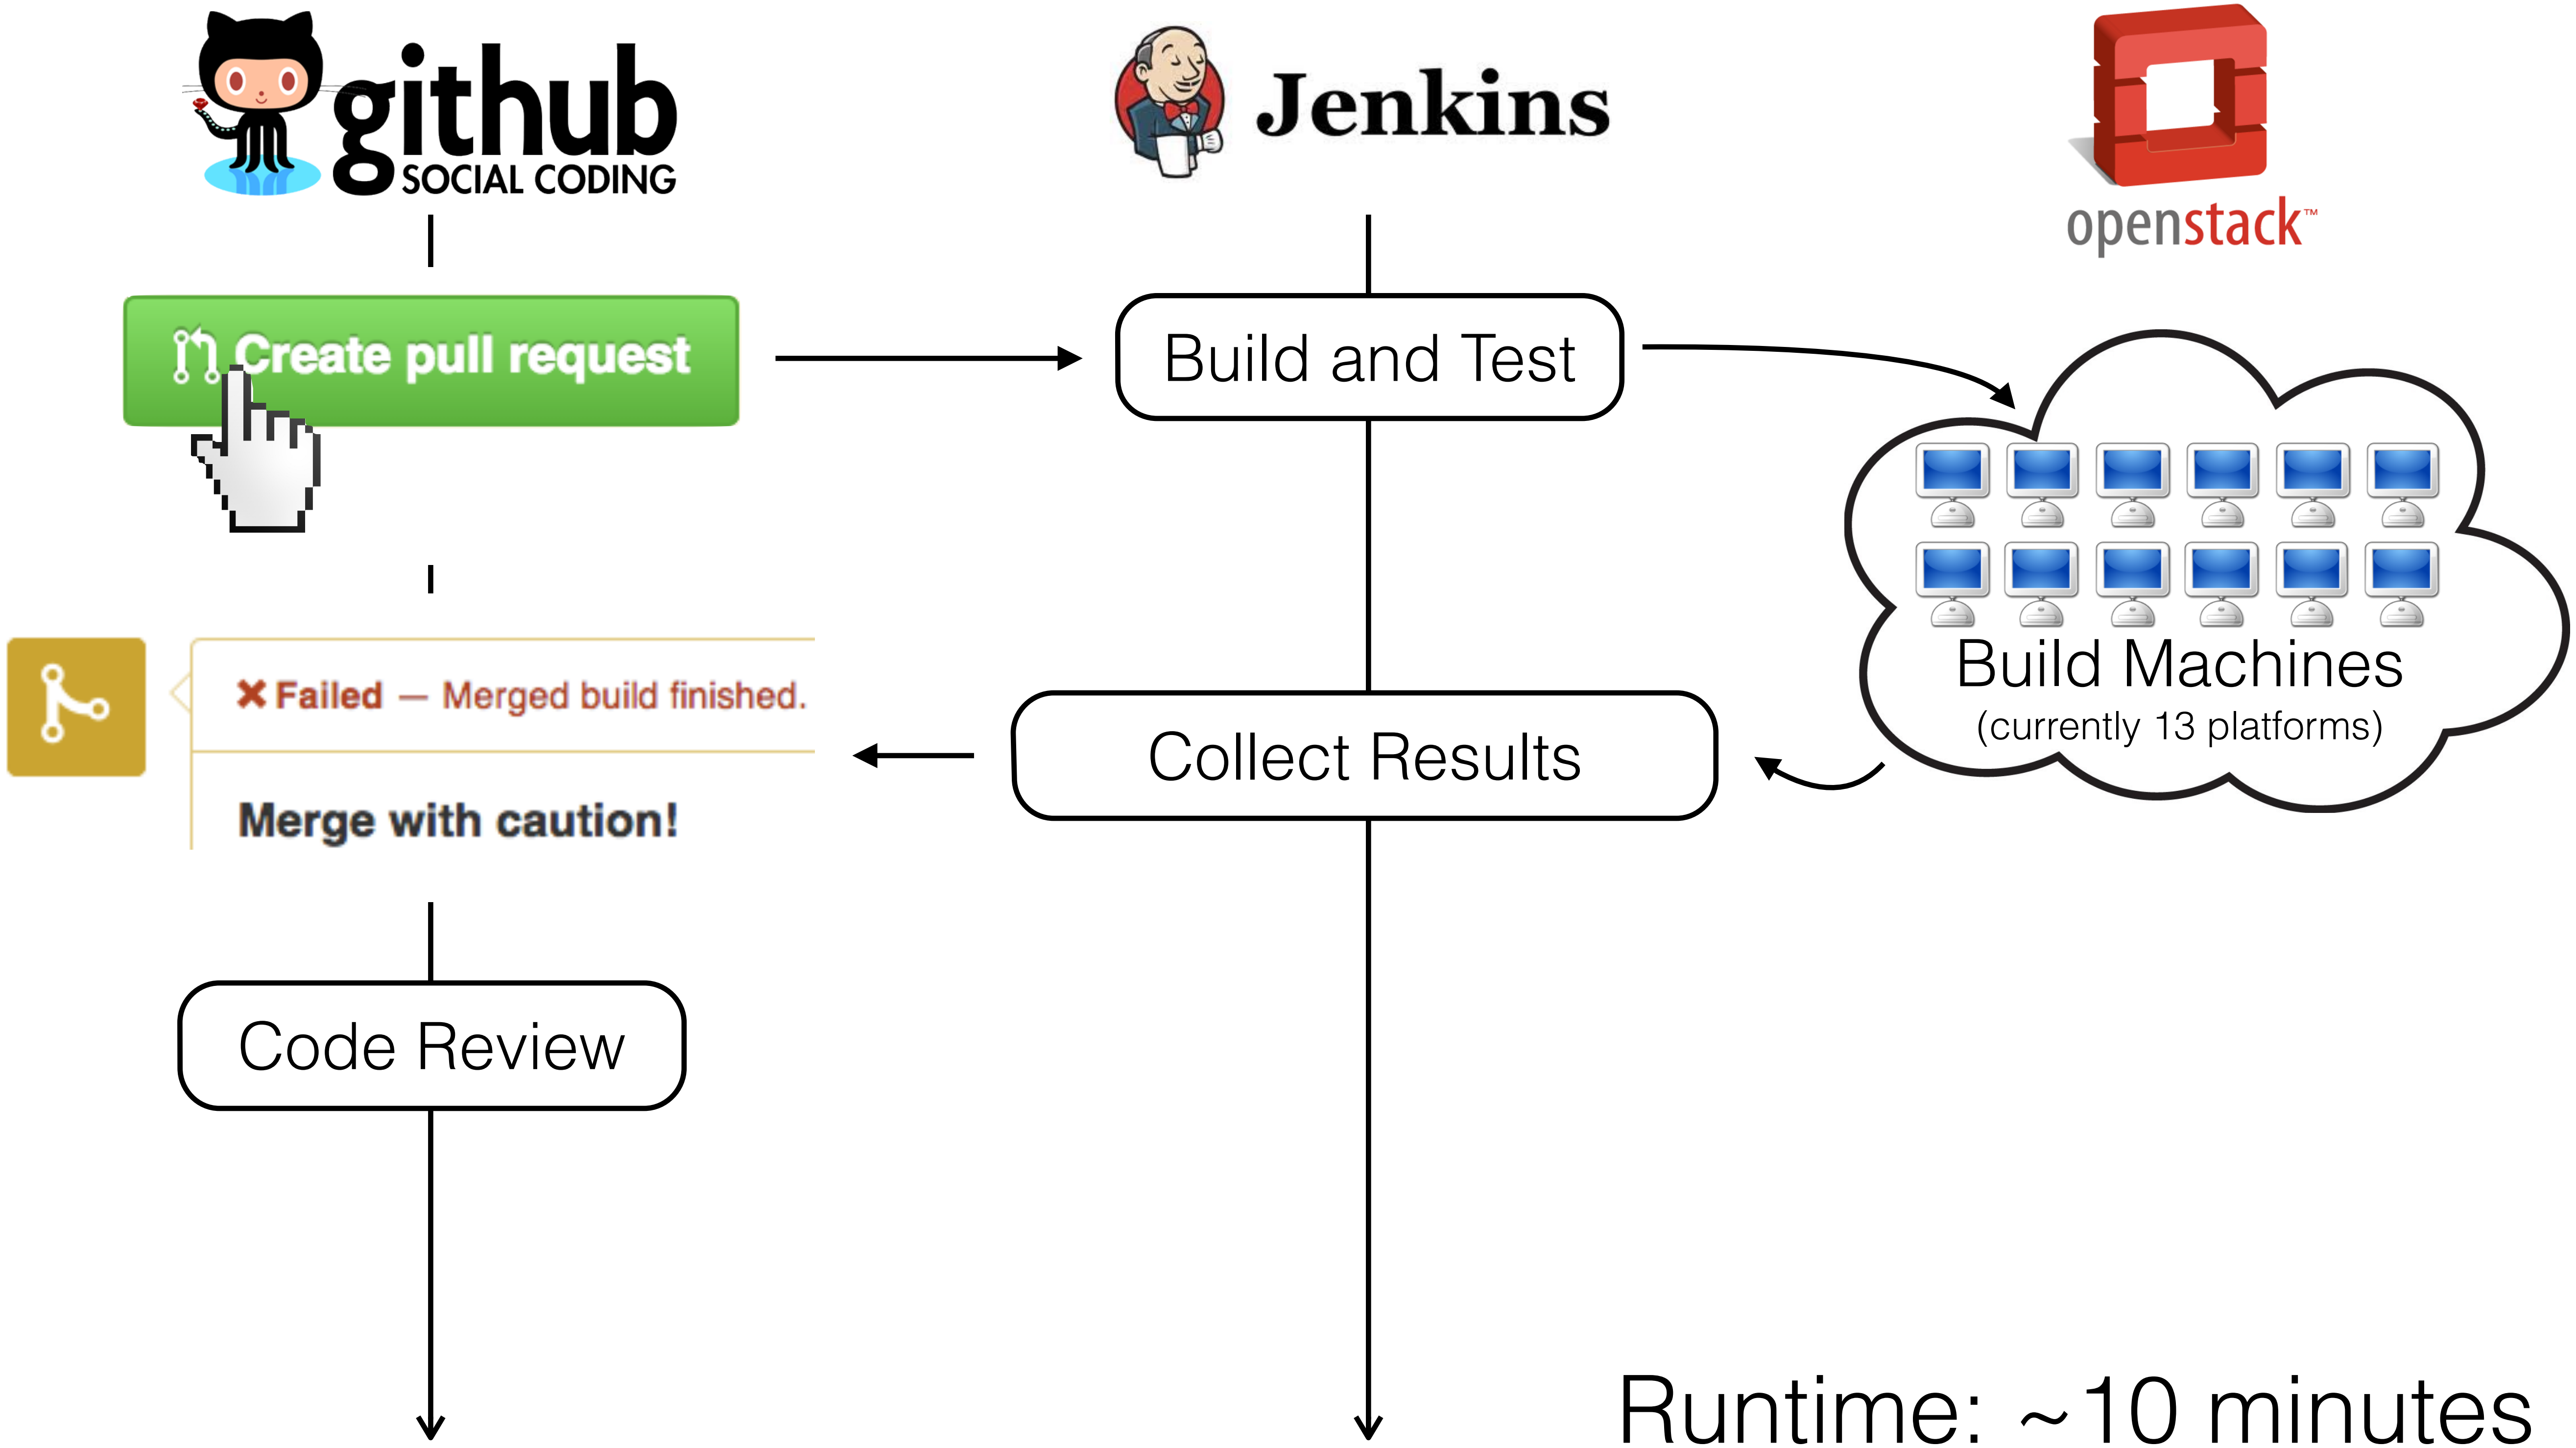

Build and Test

Build Machines
(currently 13 platforms)

Collect Results

✖ Failed — Merged build finished.

Merge with caution!

Code Review

Runtime: ~10 minutes

Continuous Multi-Platform Build and Test

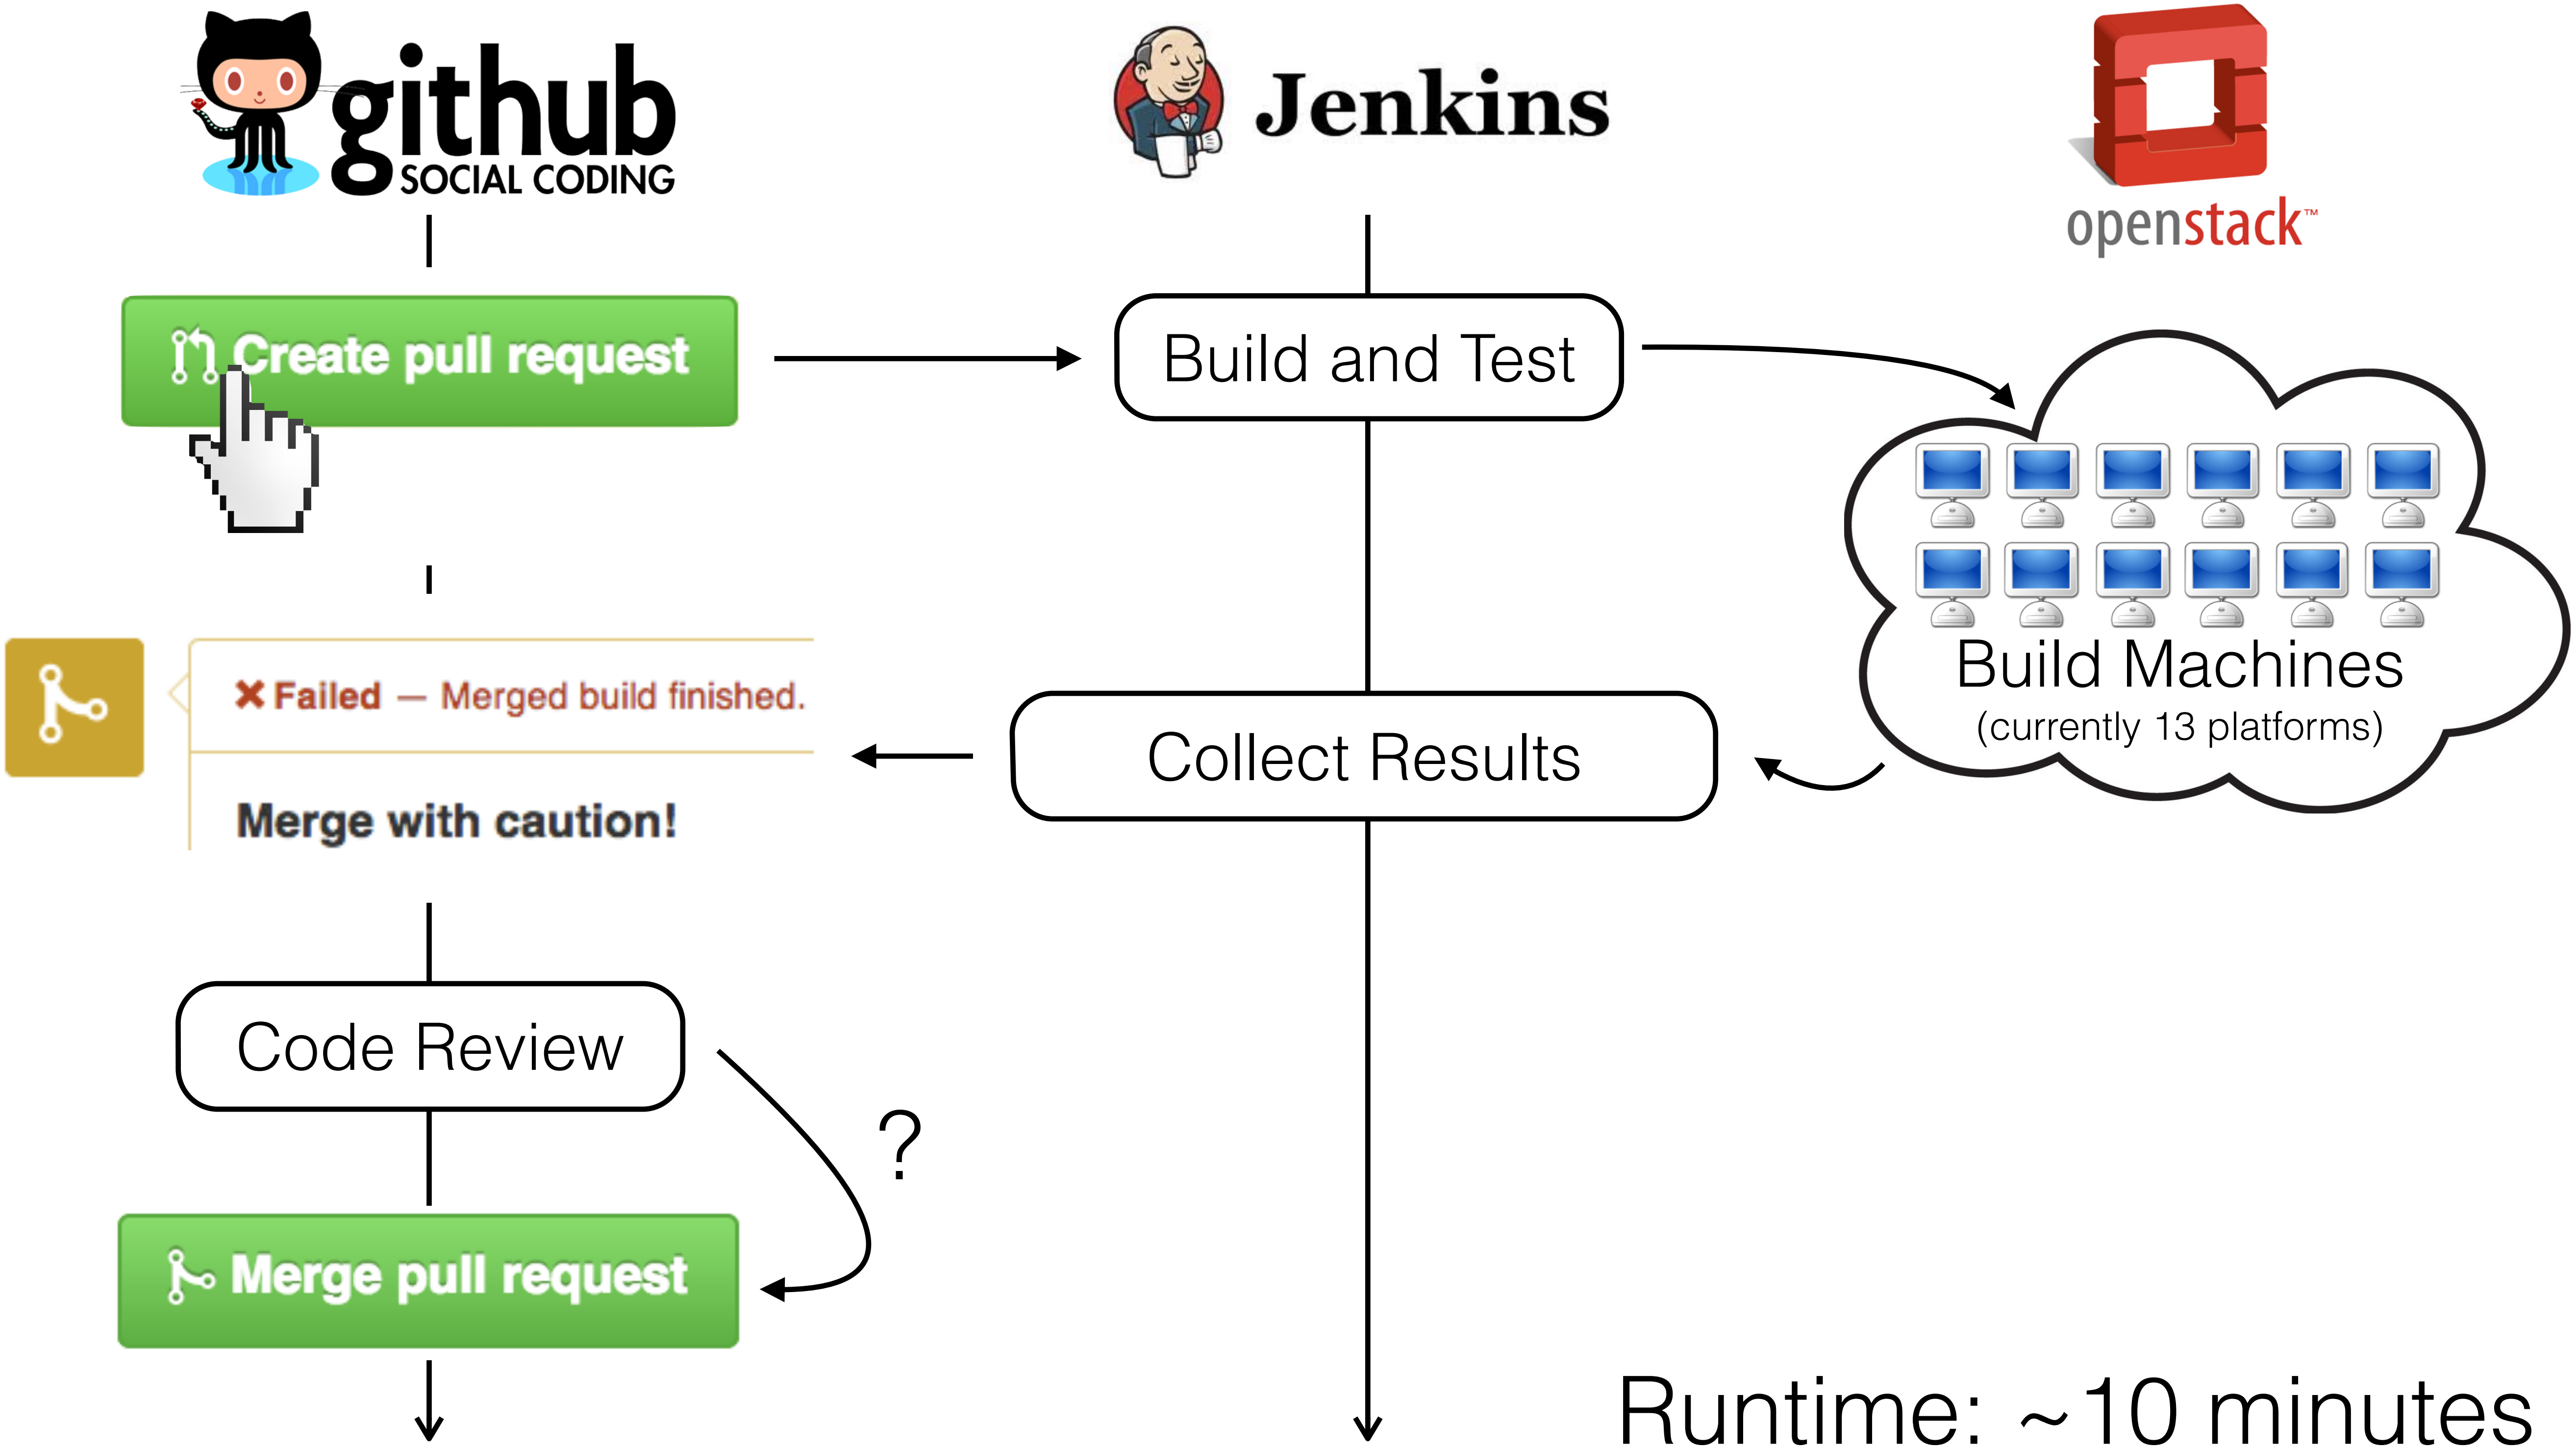

github SOCIAL CODING

Jenkins

openstack™

Create pull request

Build and Test

Build Machines

(currently 13 platforms)

Collect Results

✕ Failed — Merged build finished.

Merge with caution!

Code Review

?

Merge pull request

Runtime: ~10 minutes

Full Integration Test Execution

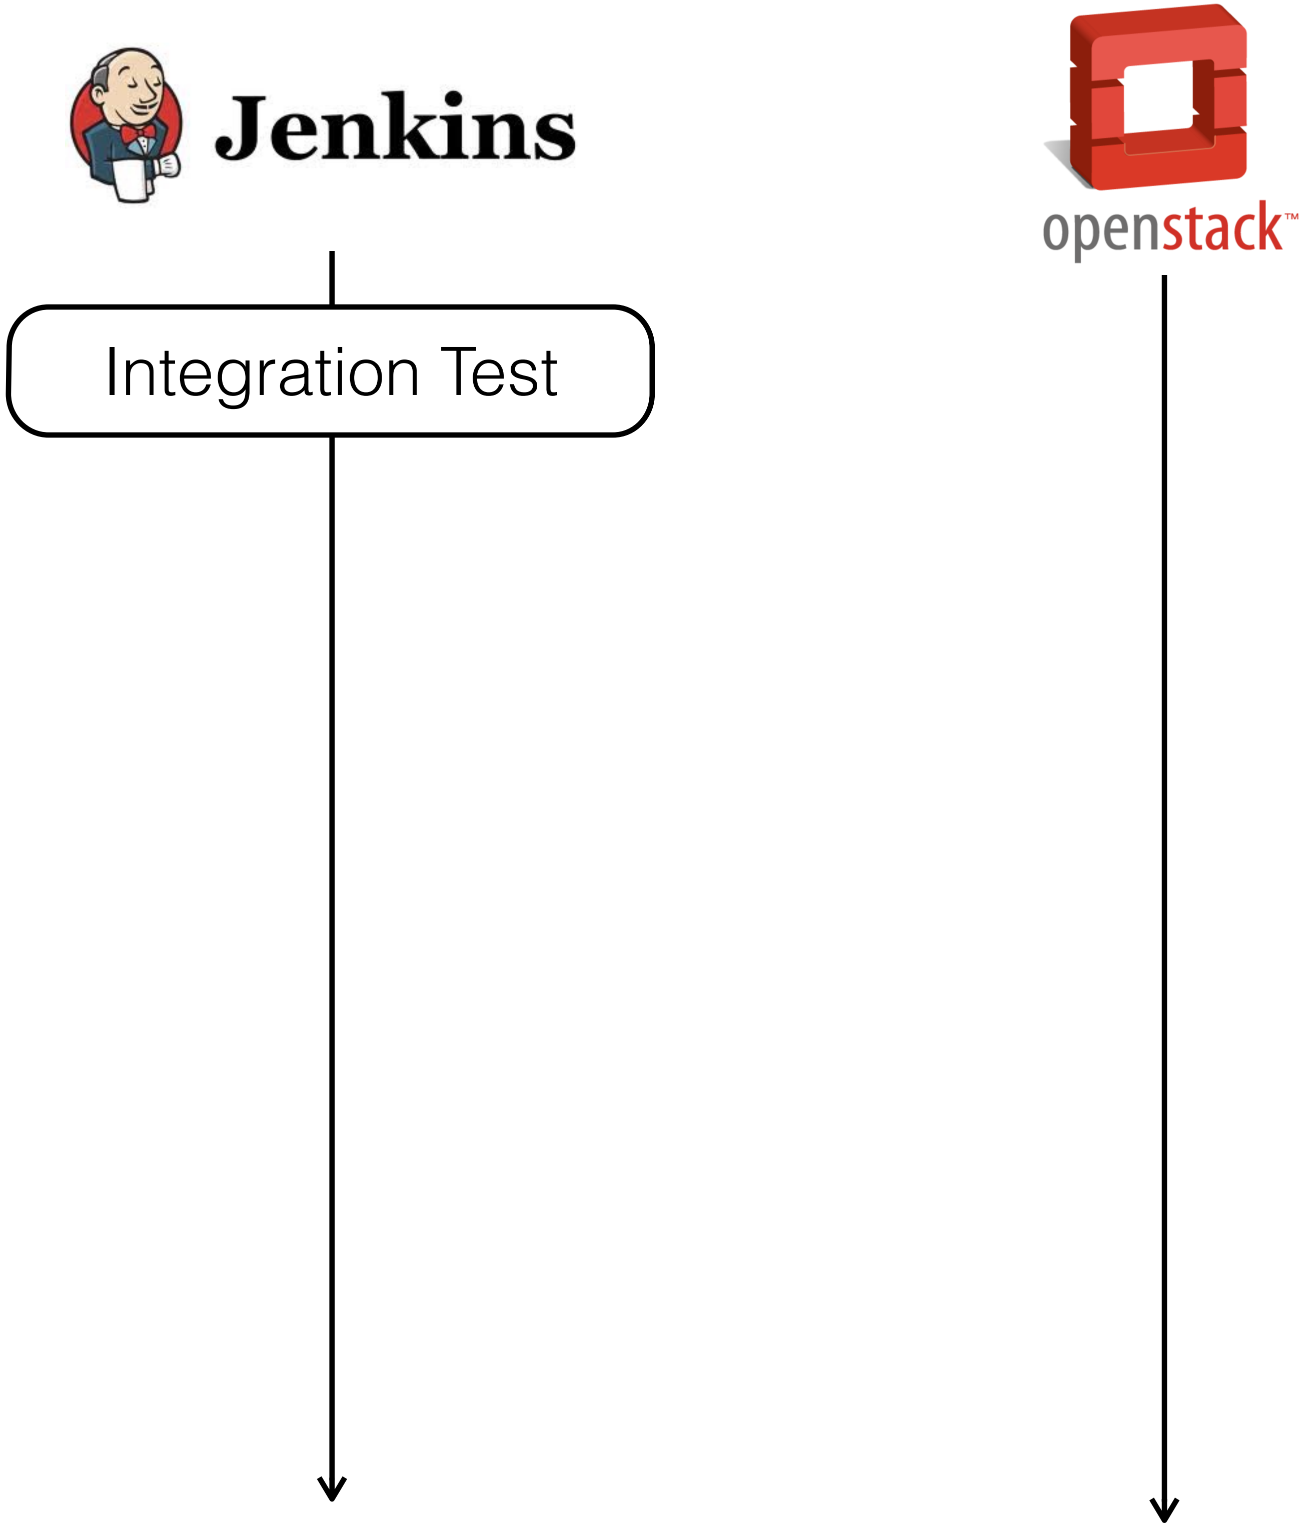

Integration Test

Runtime: ~16 hours

Full Integration Test Execution

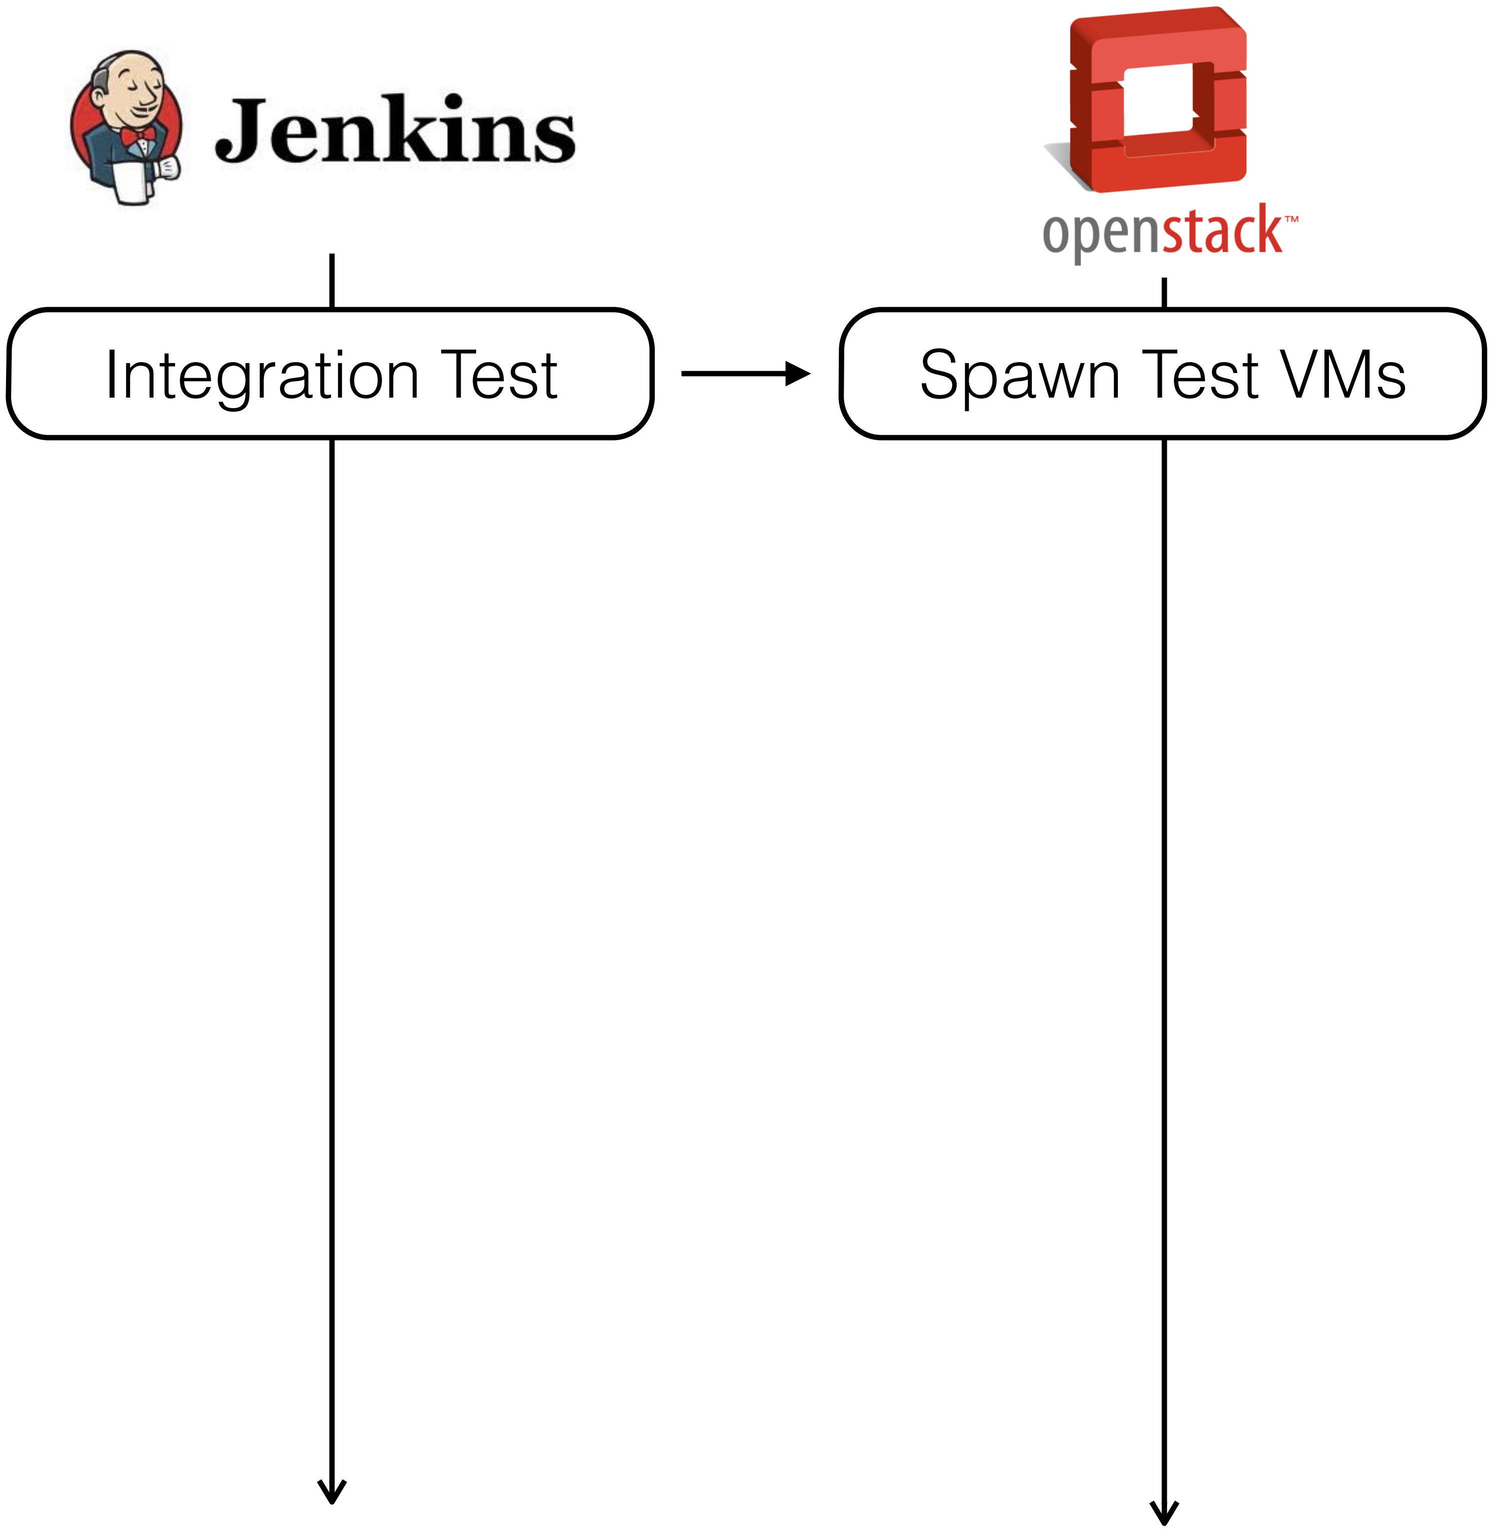

Integration Test

Spawn Test VMs

Runtime: ~16 hours

Full Integration Test Execution

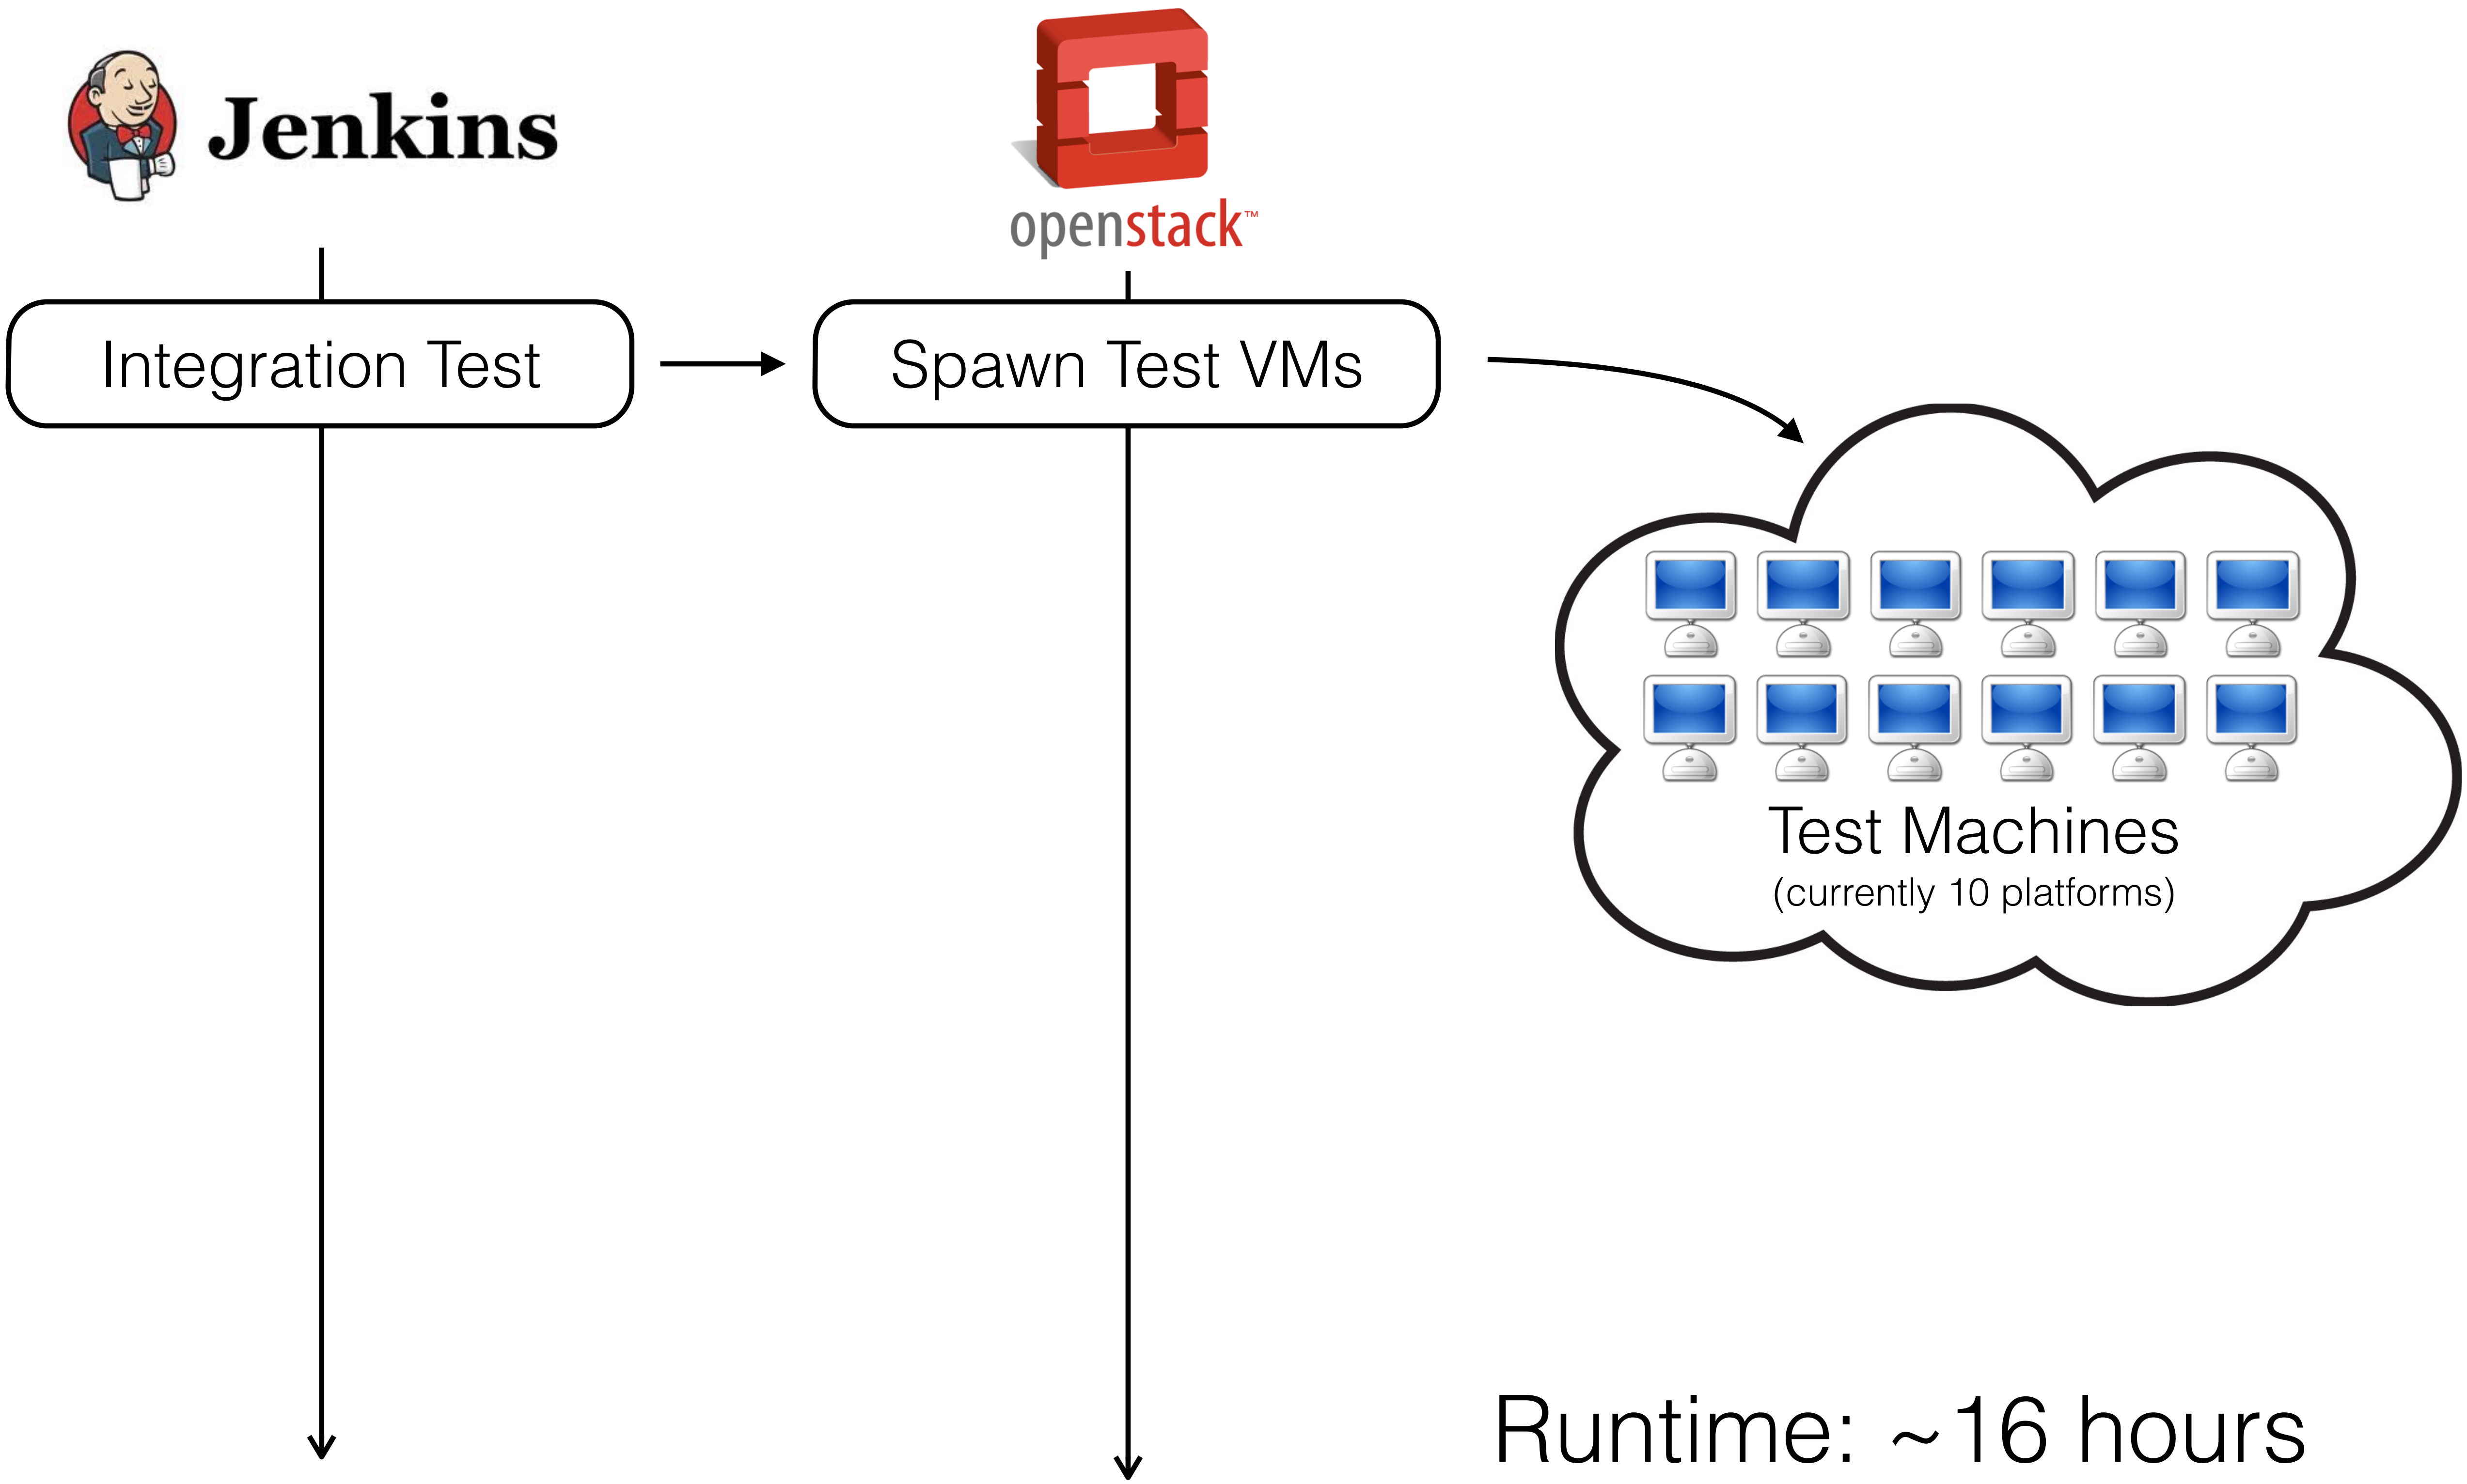

Jenkins

openstack™

Integration Test

Spawn Test VMs

Test Machines

(currently 10 platforms)

Runtime: ~16 hours

Full Integration Test Execution

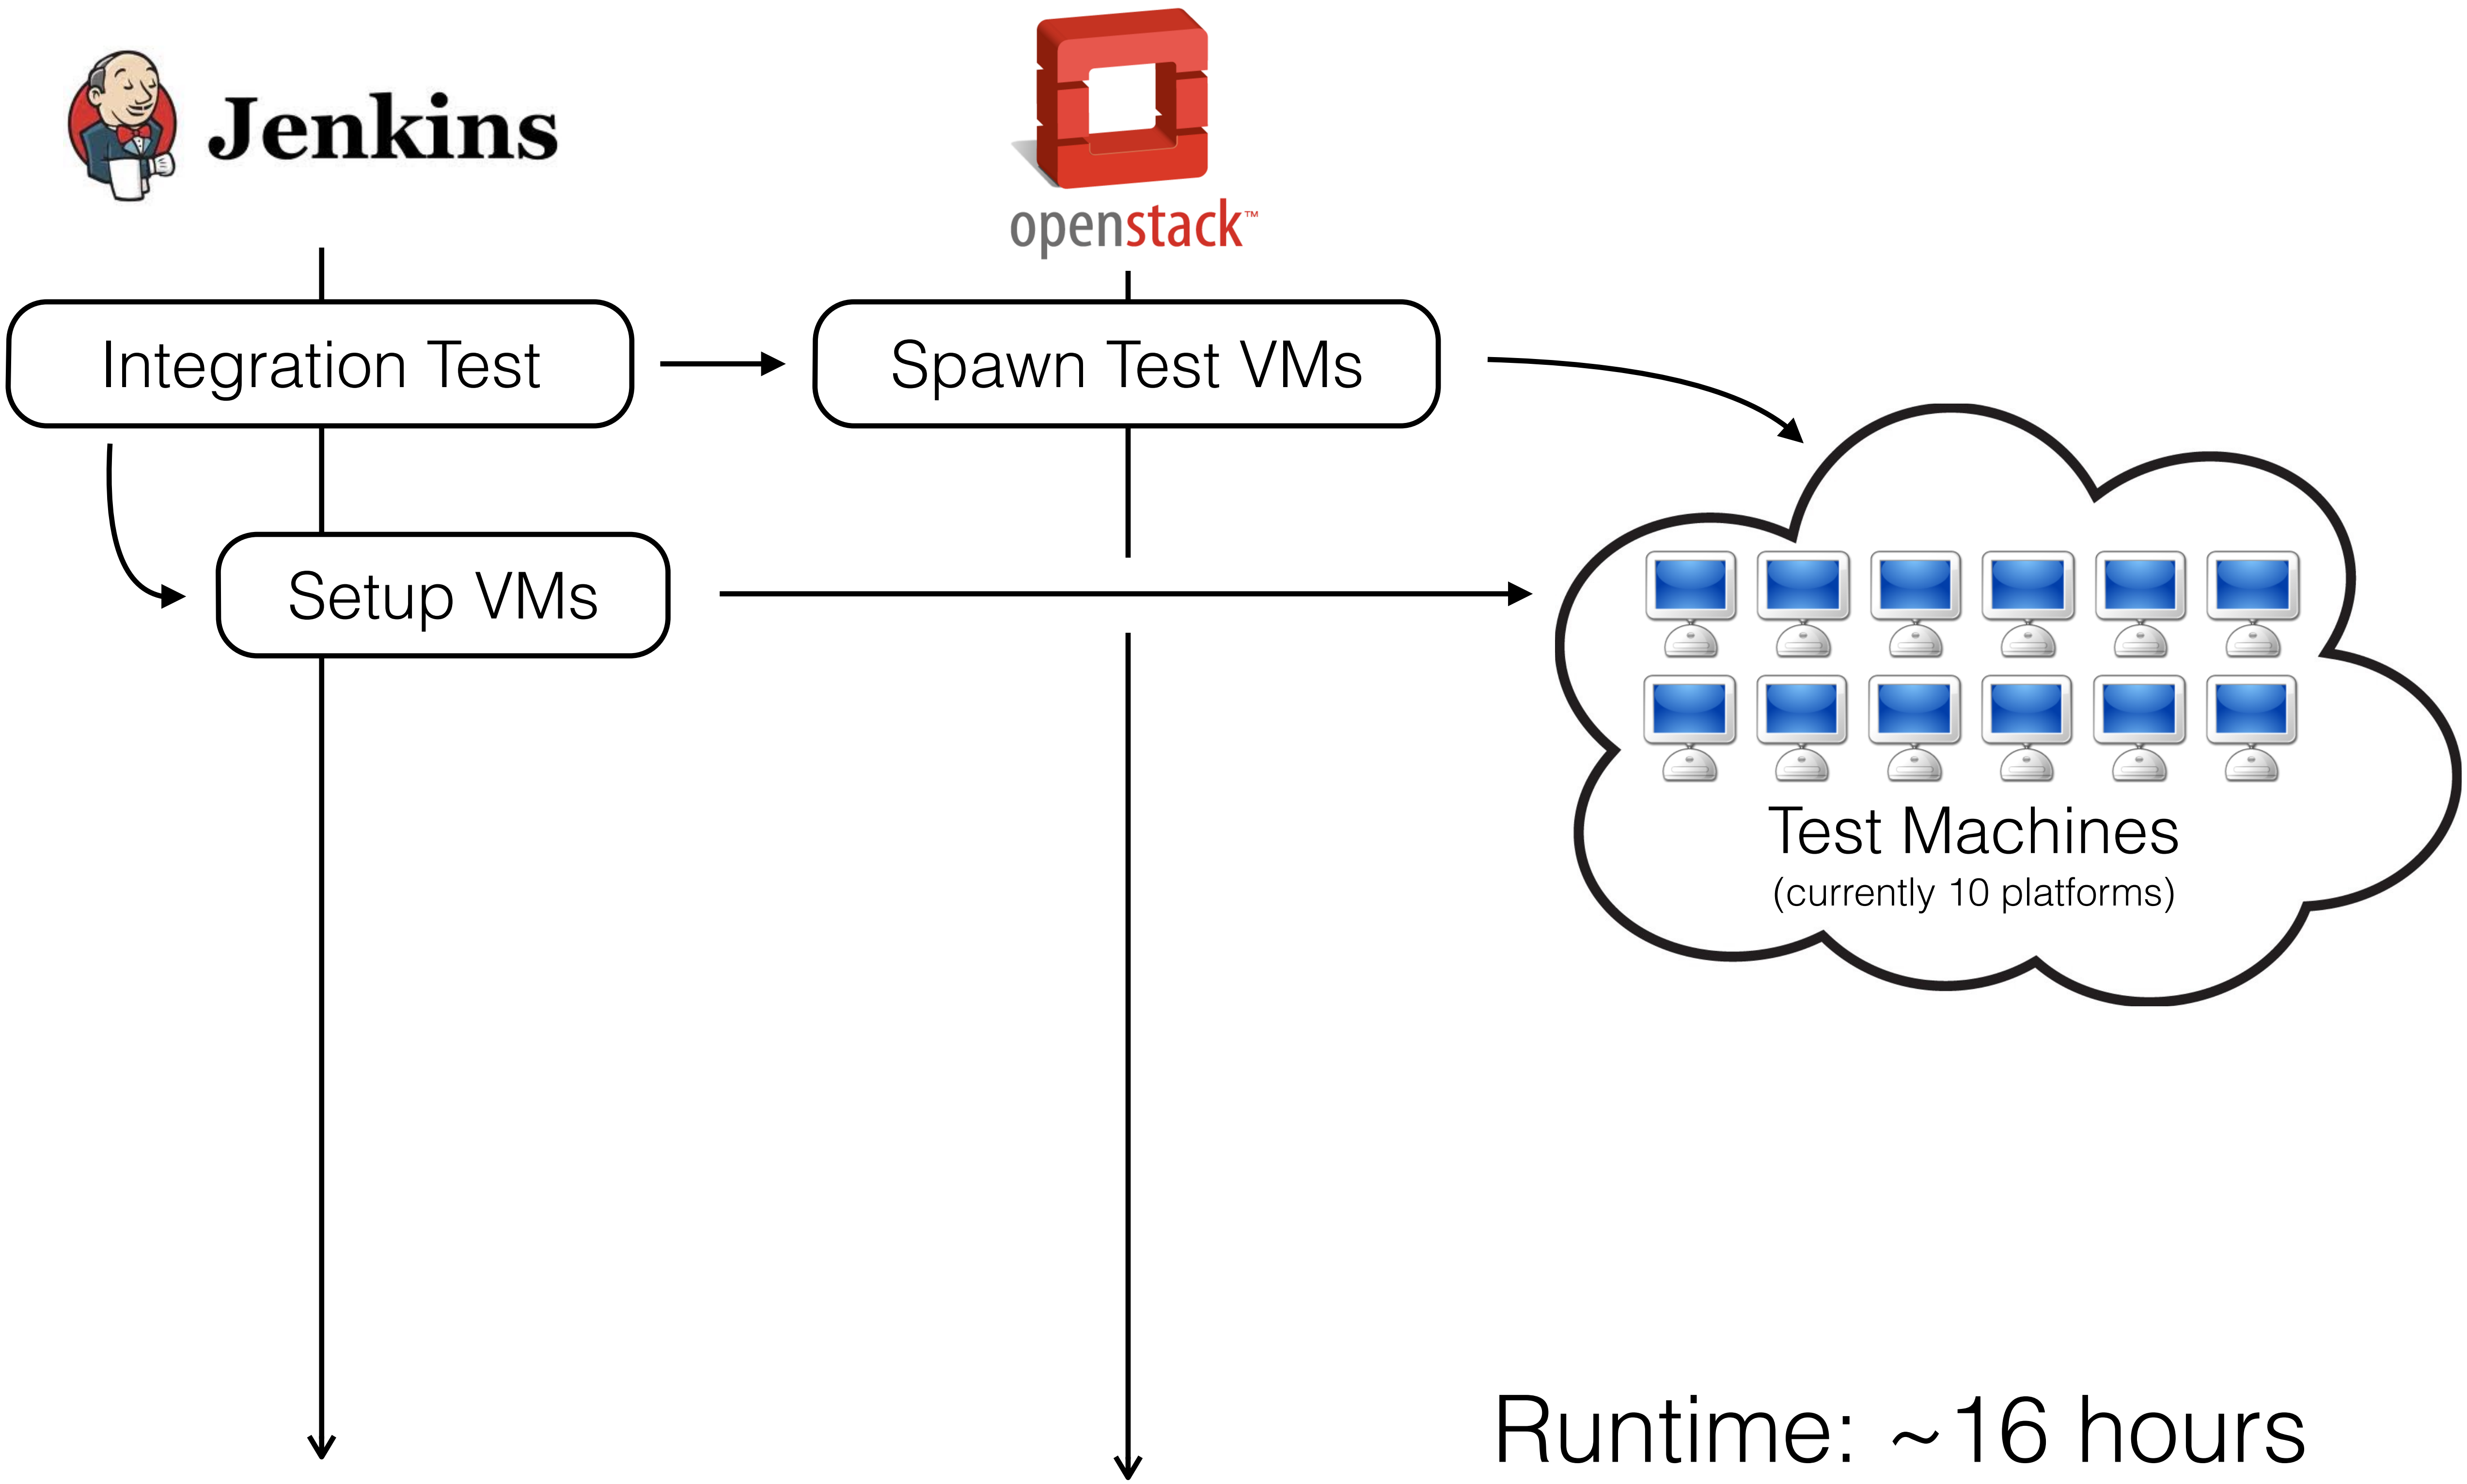

Runtime: ~16 hours

Full Integration Test Execution

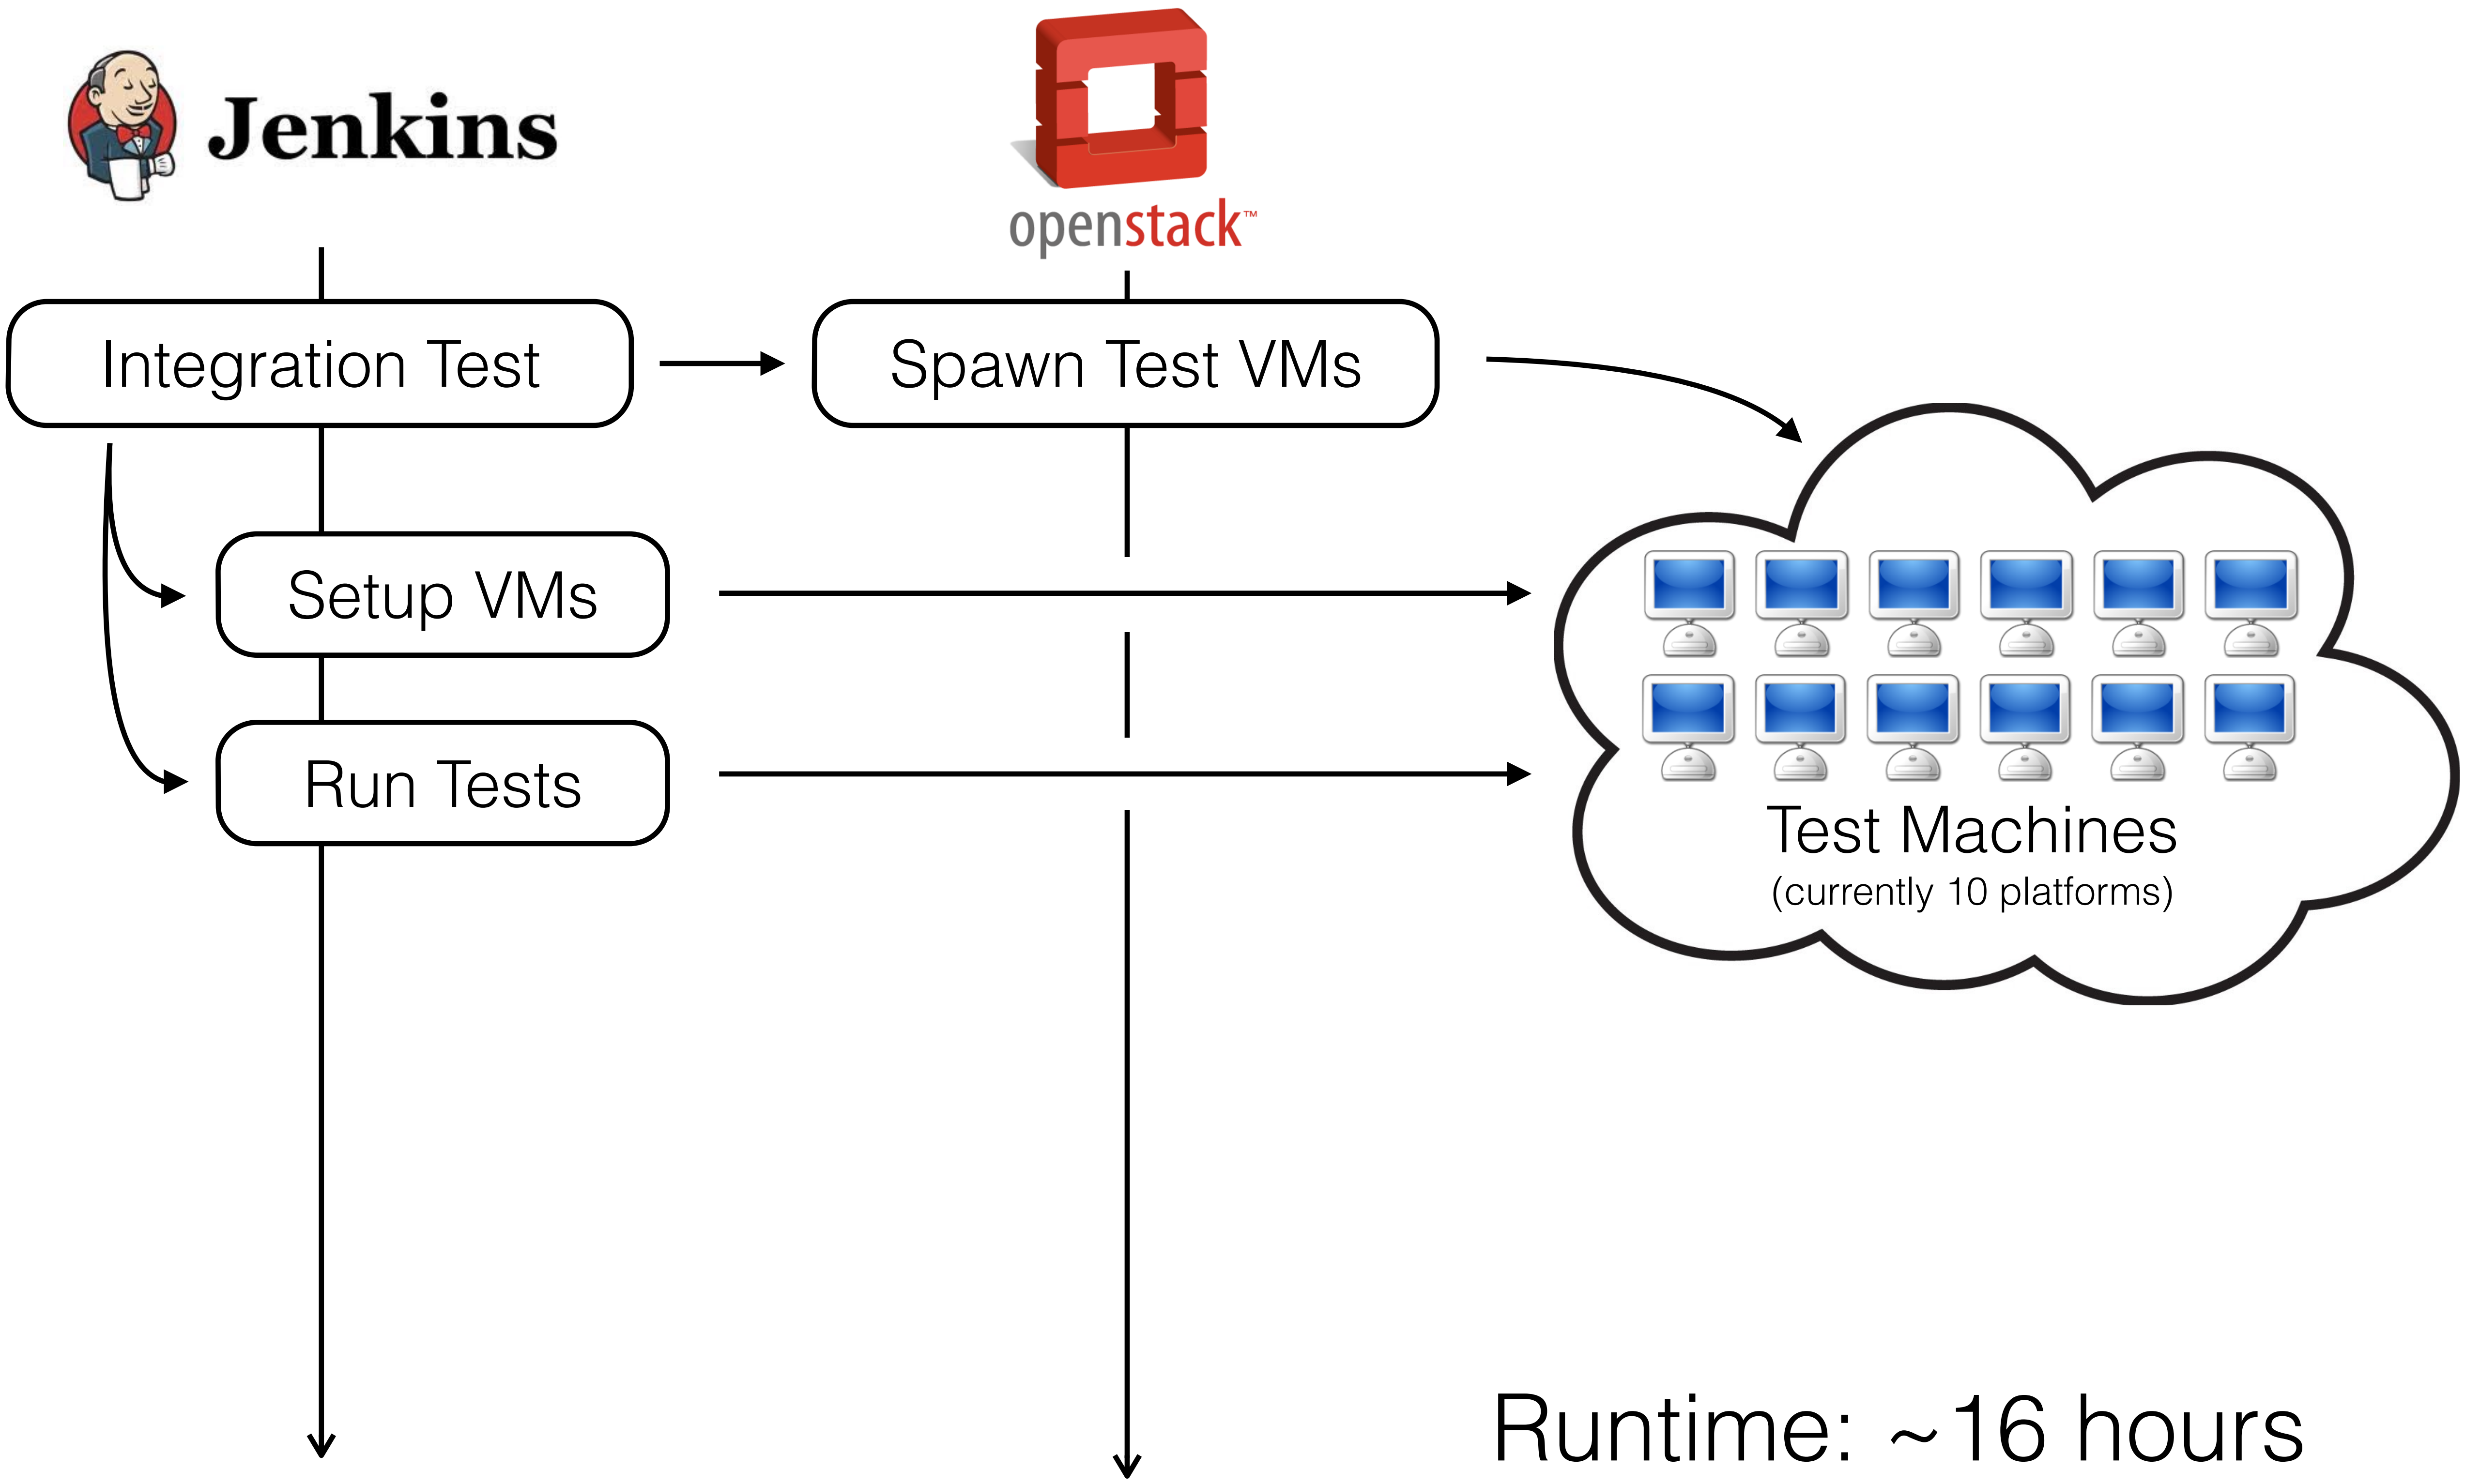

Runtime: ~16 hours

Full Integration Test Execution

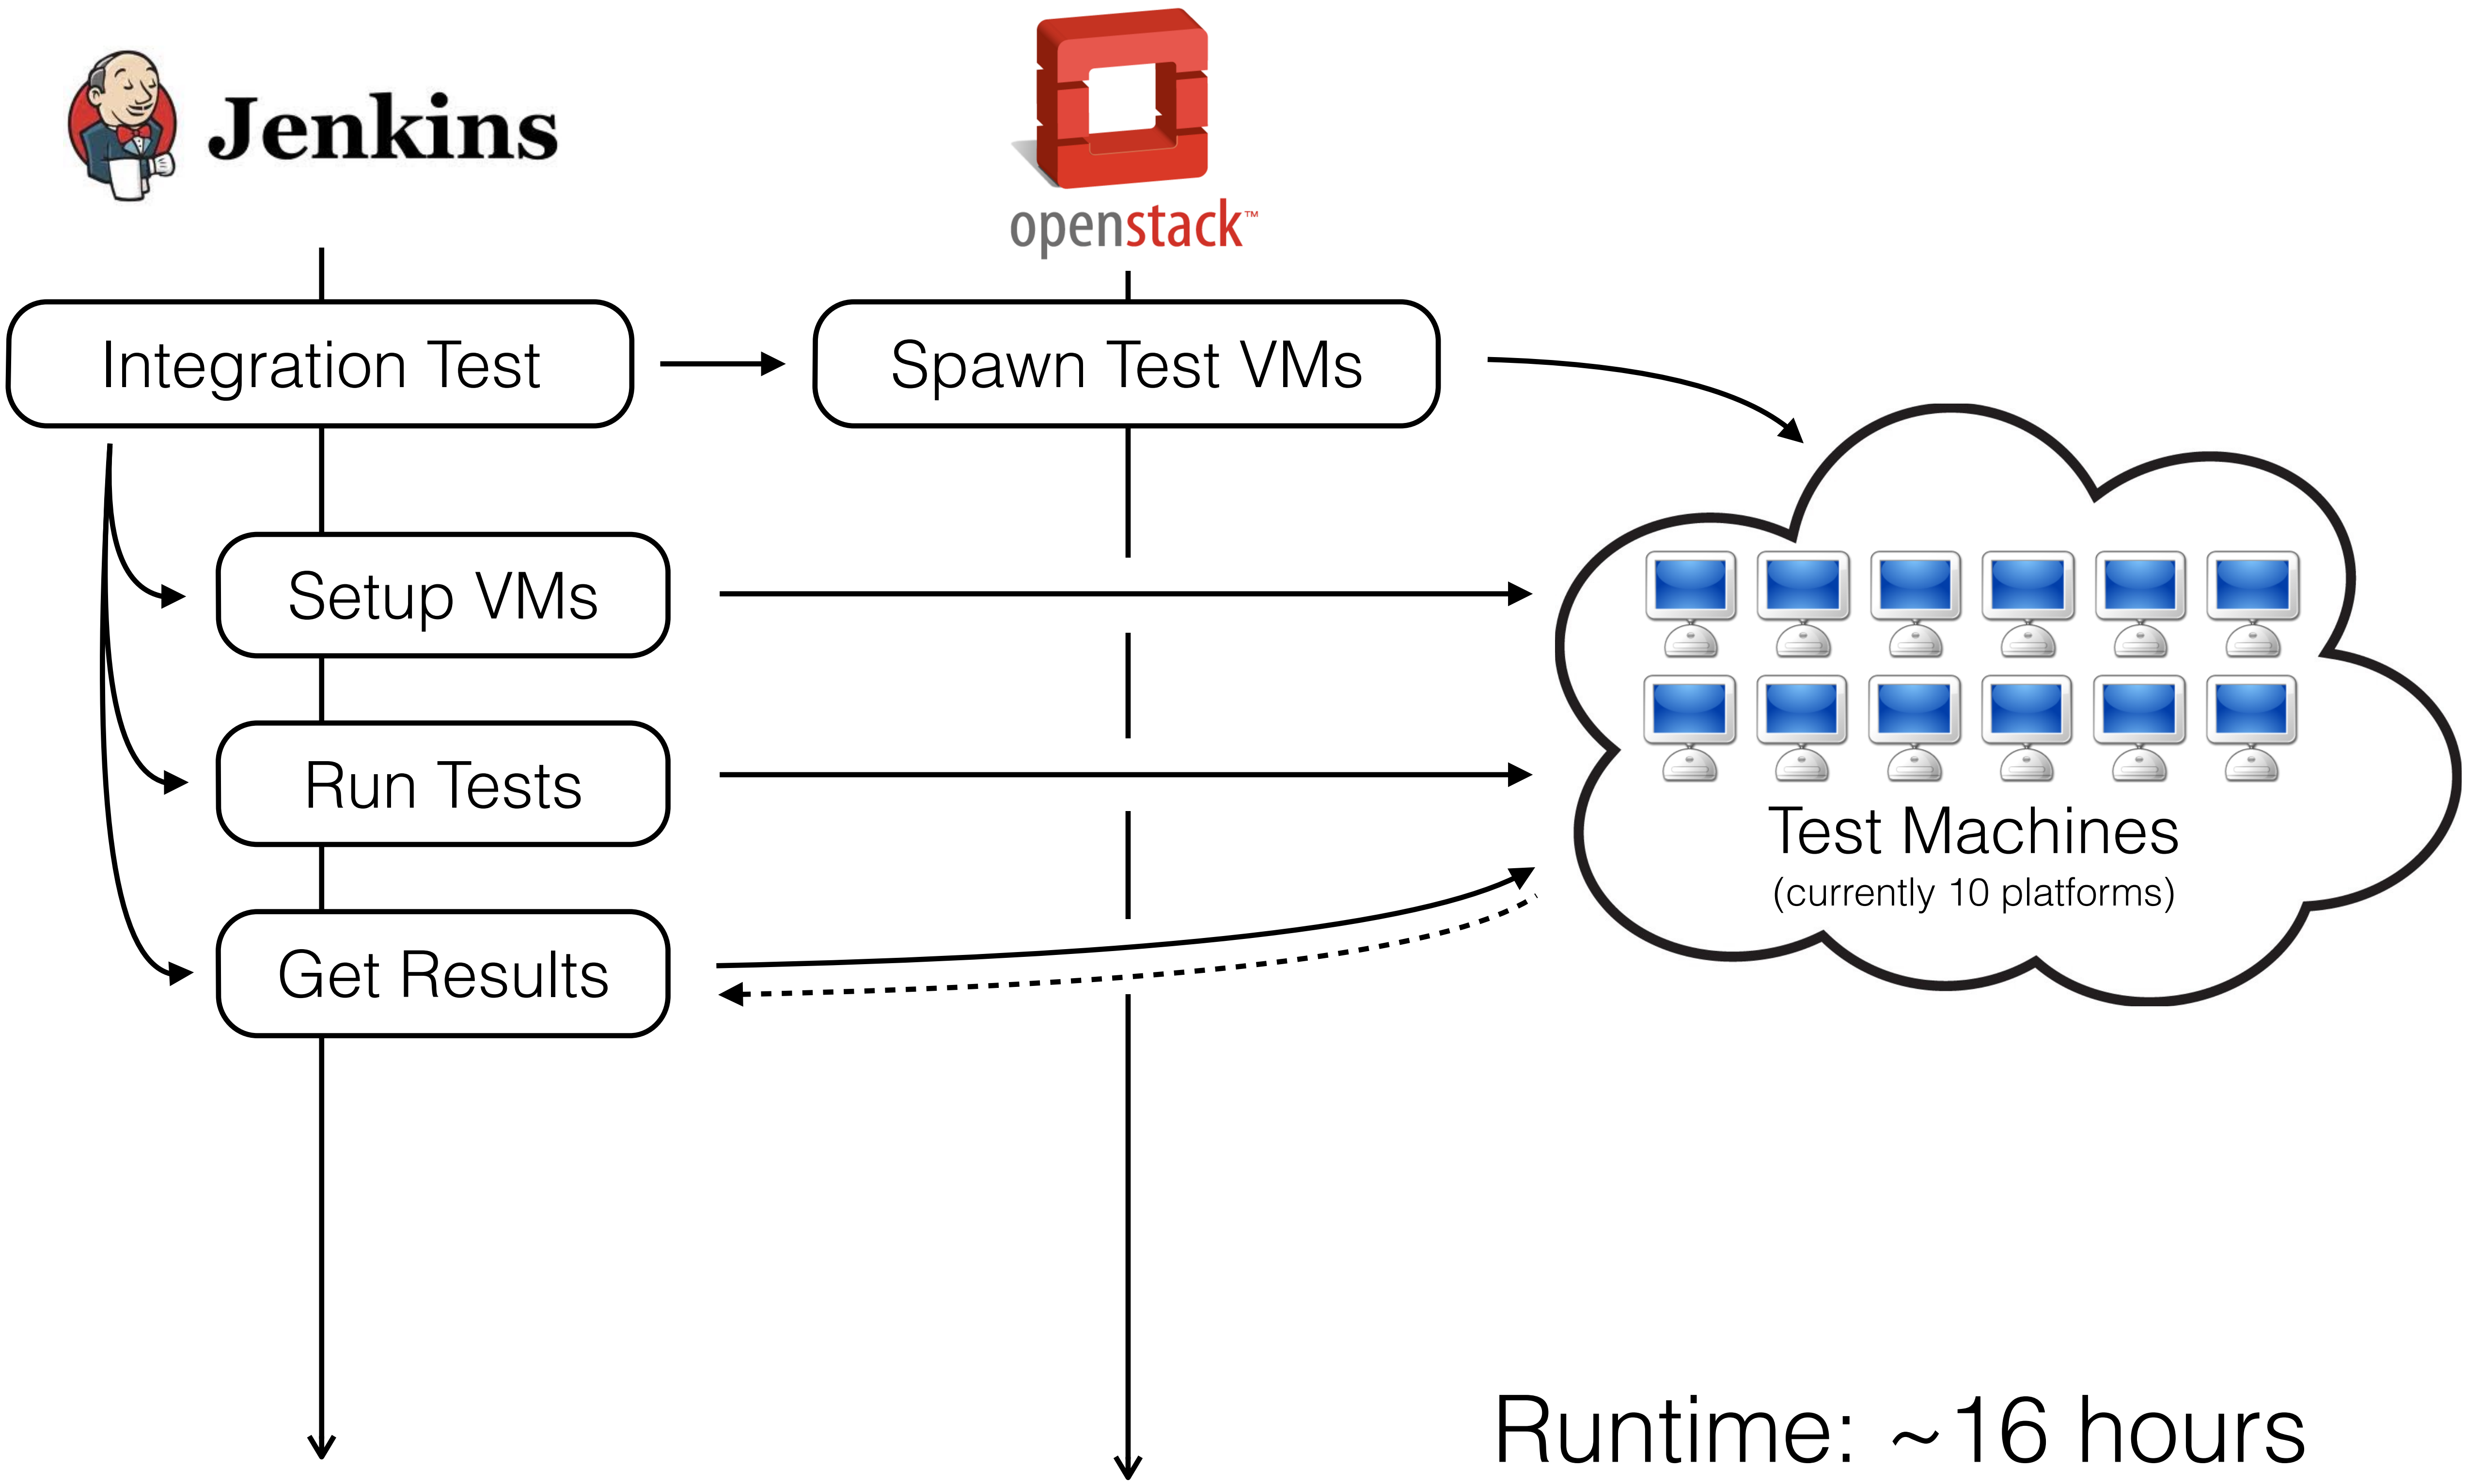

Runtime: ~16 hours

Full Integration Test Execution

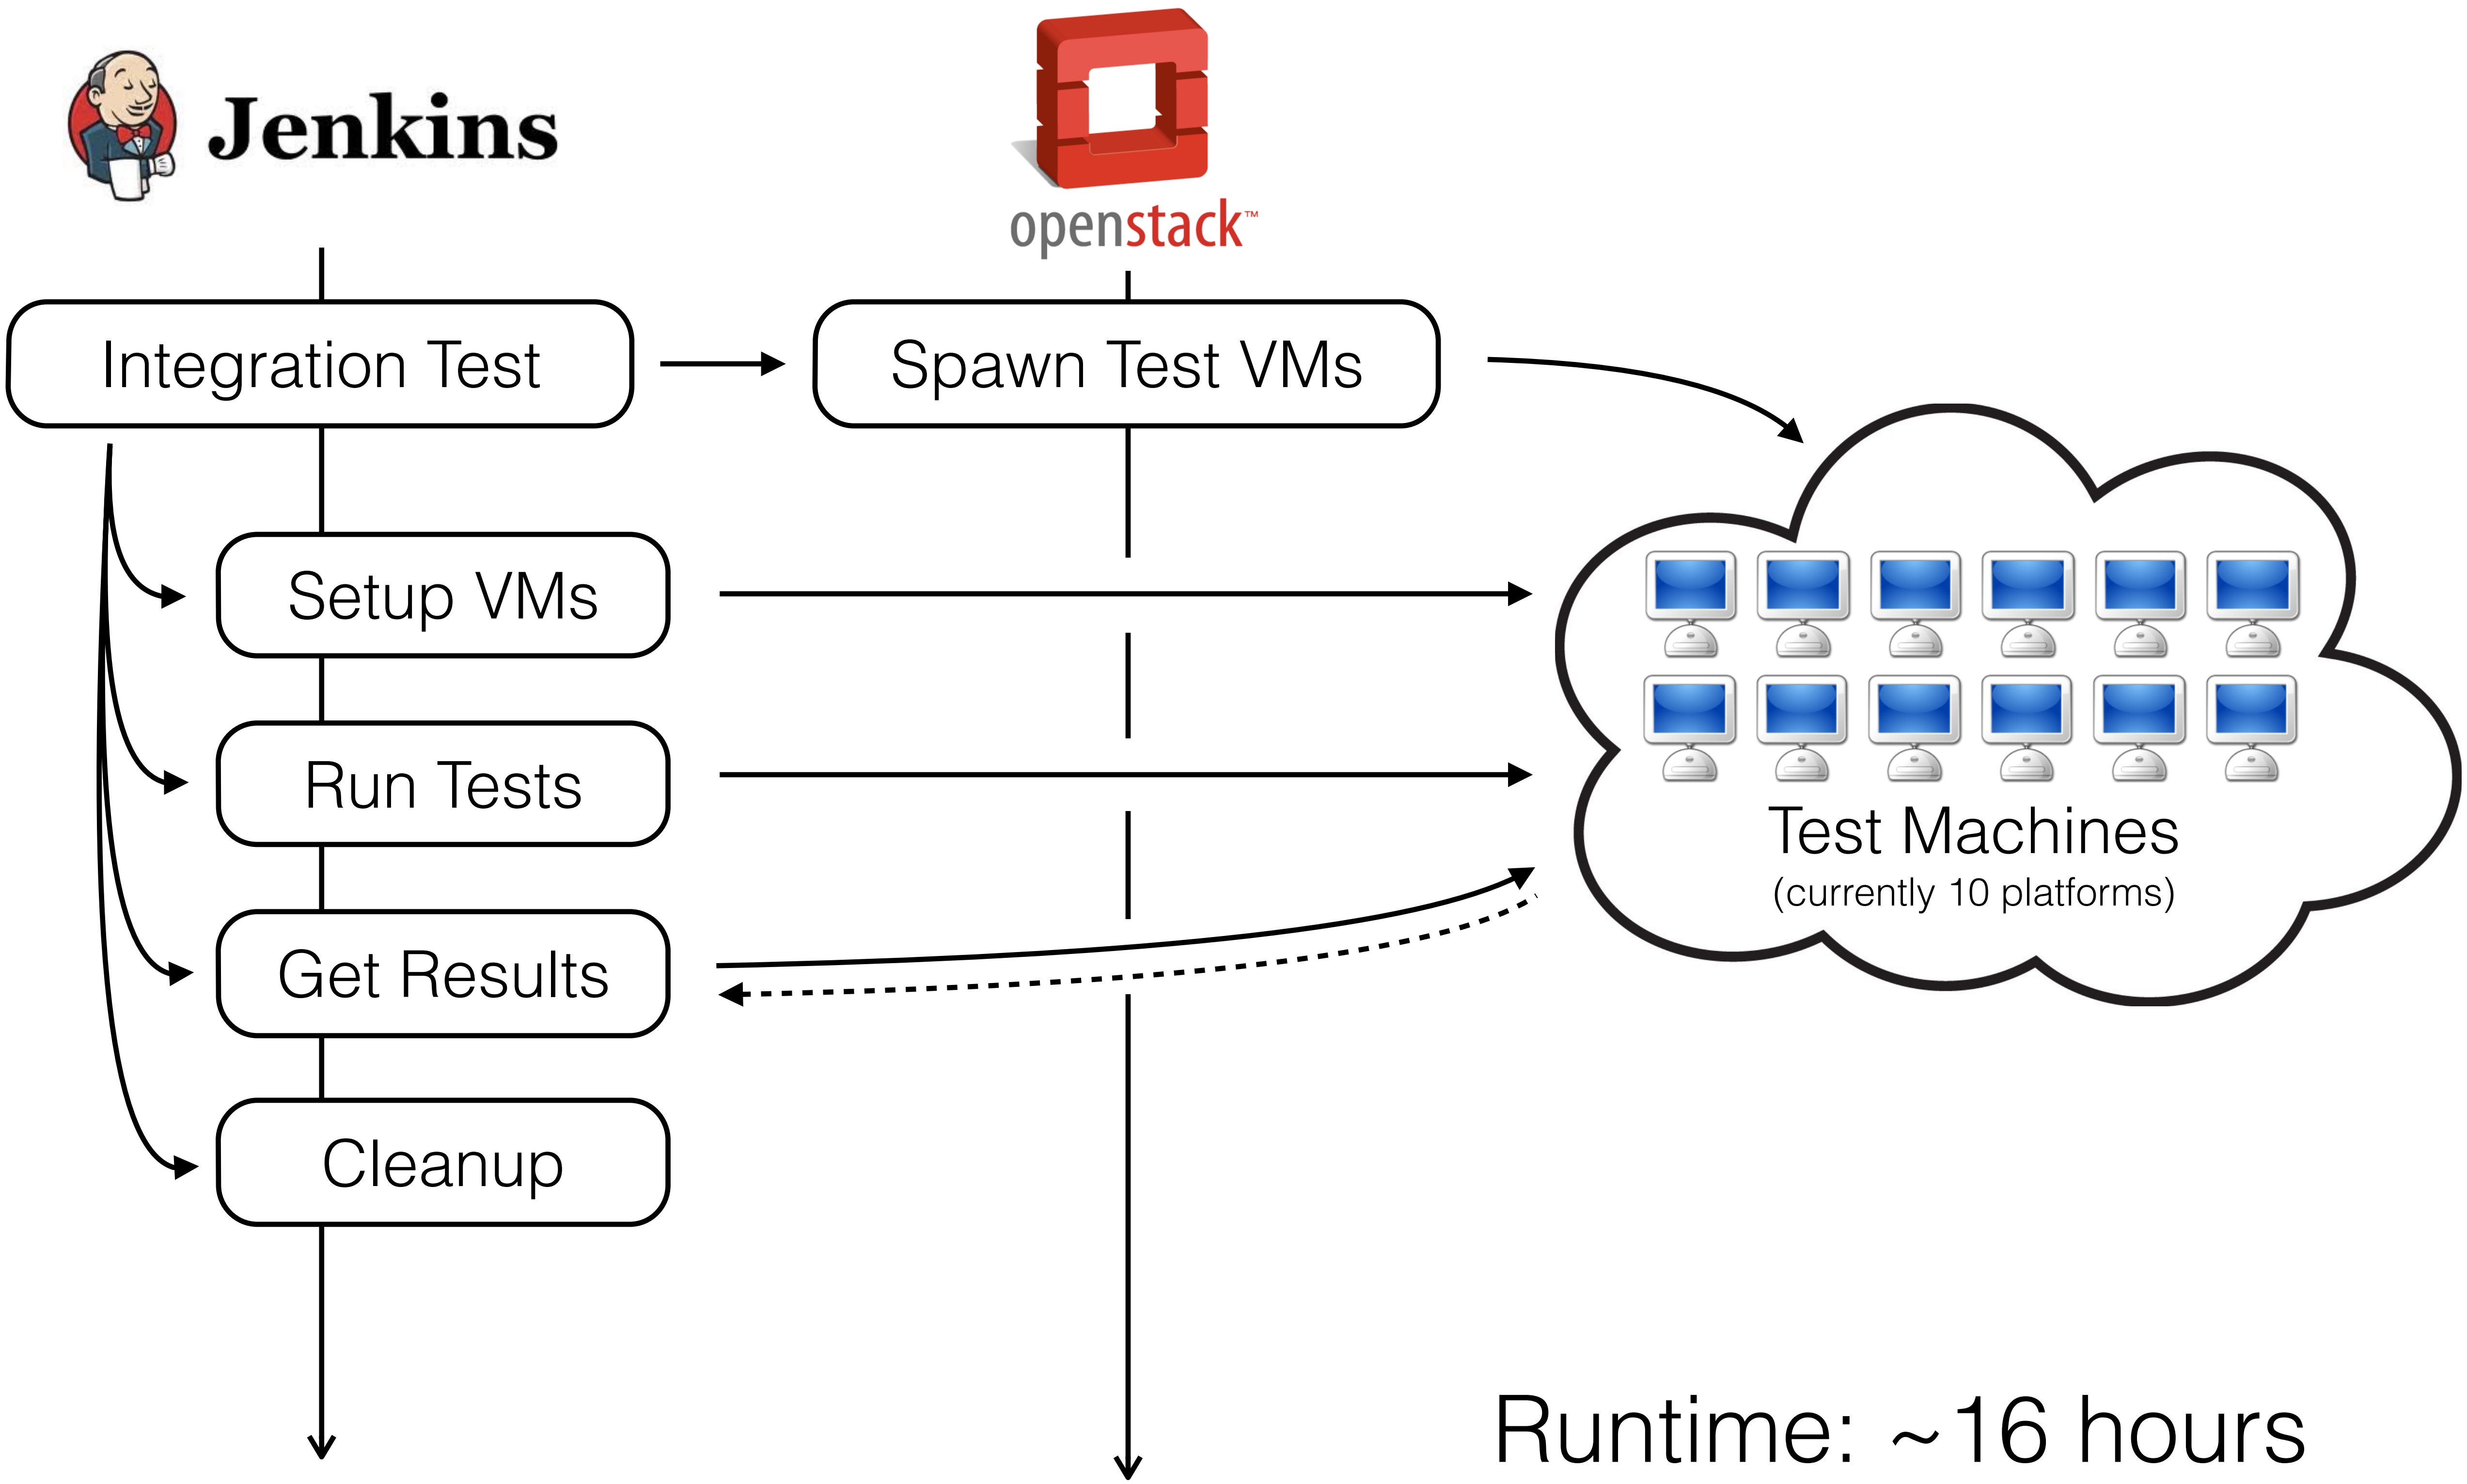

Runtime: ~16 hours

Full Integration Test Execution

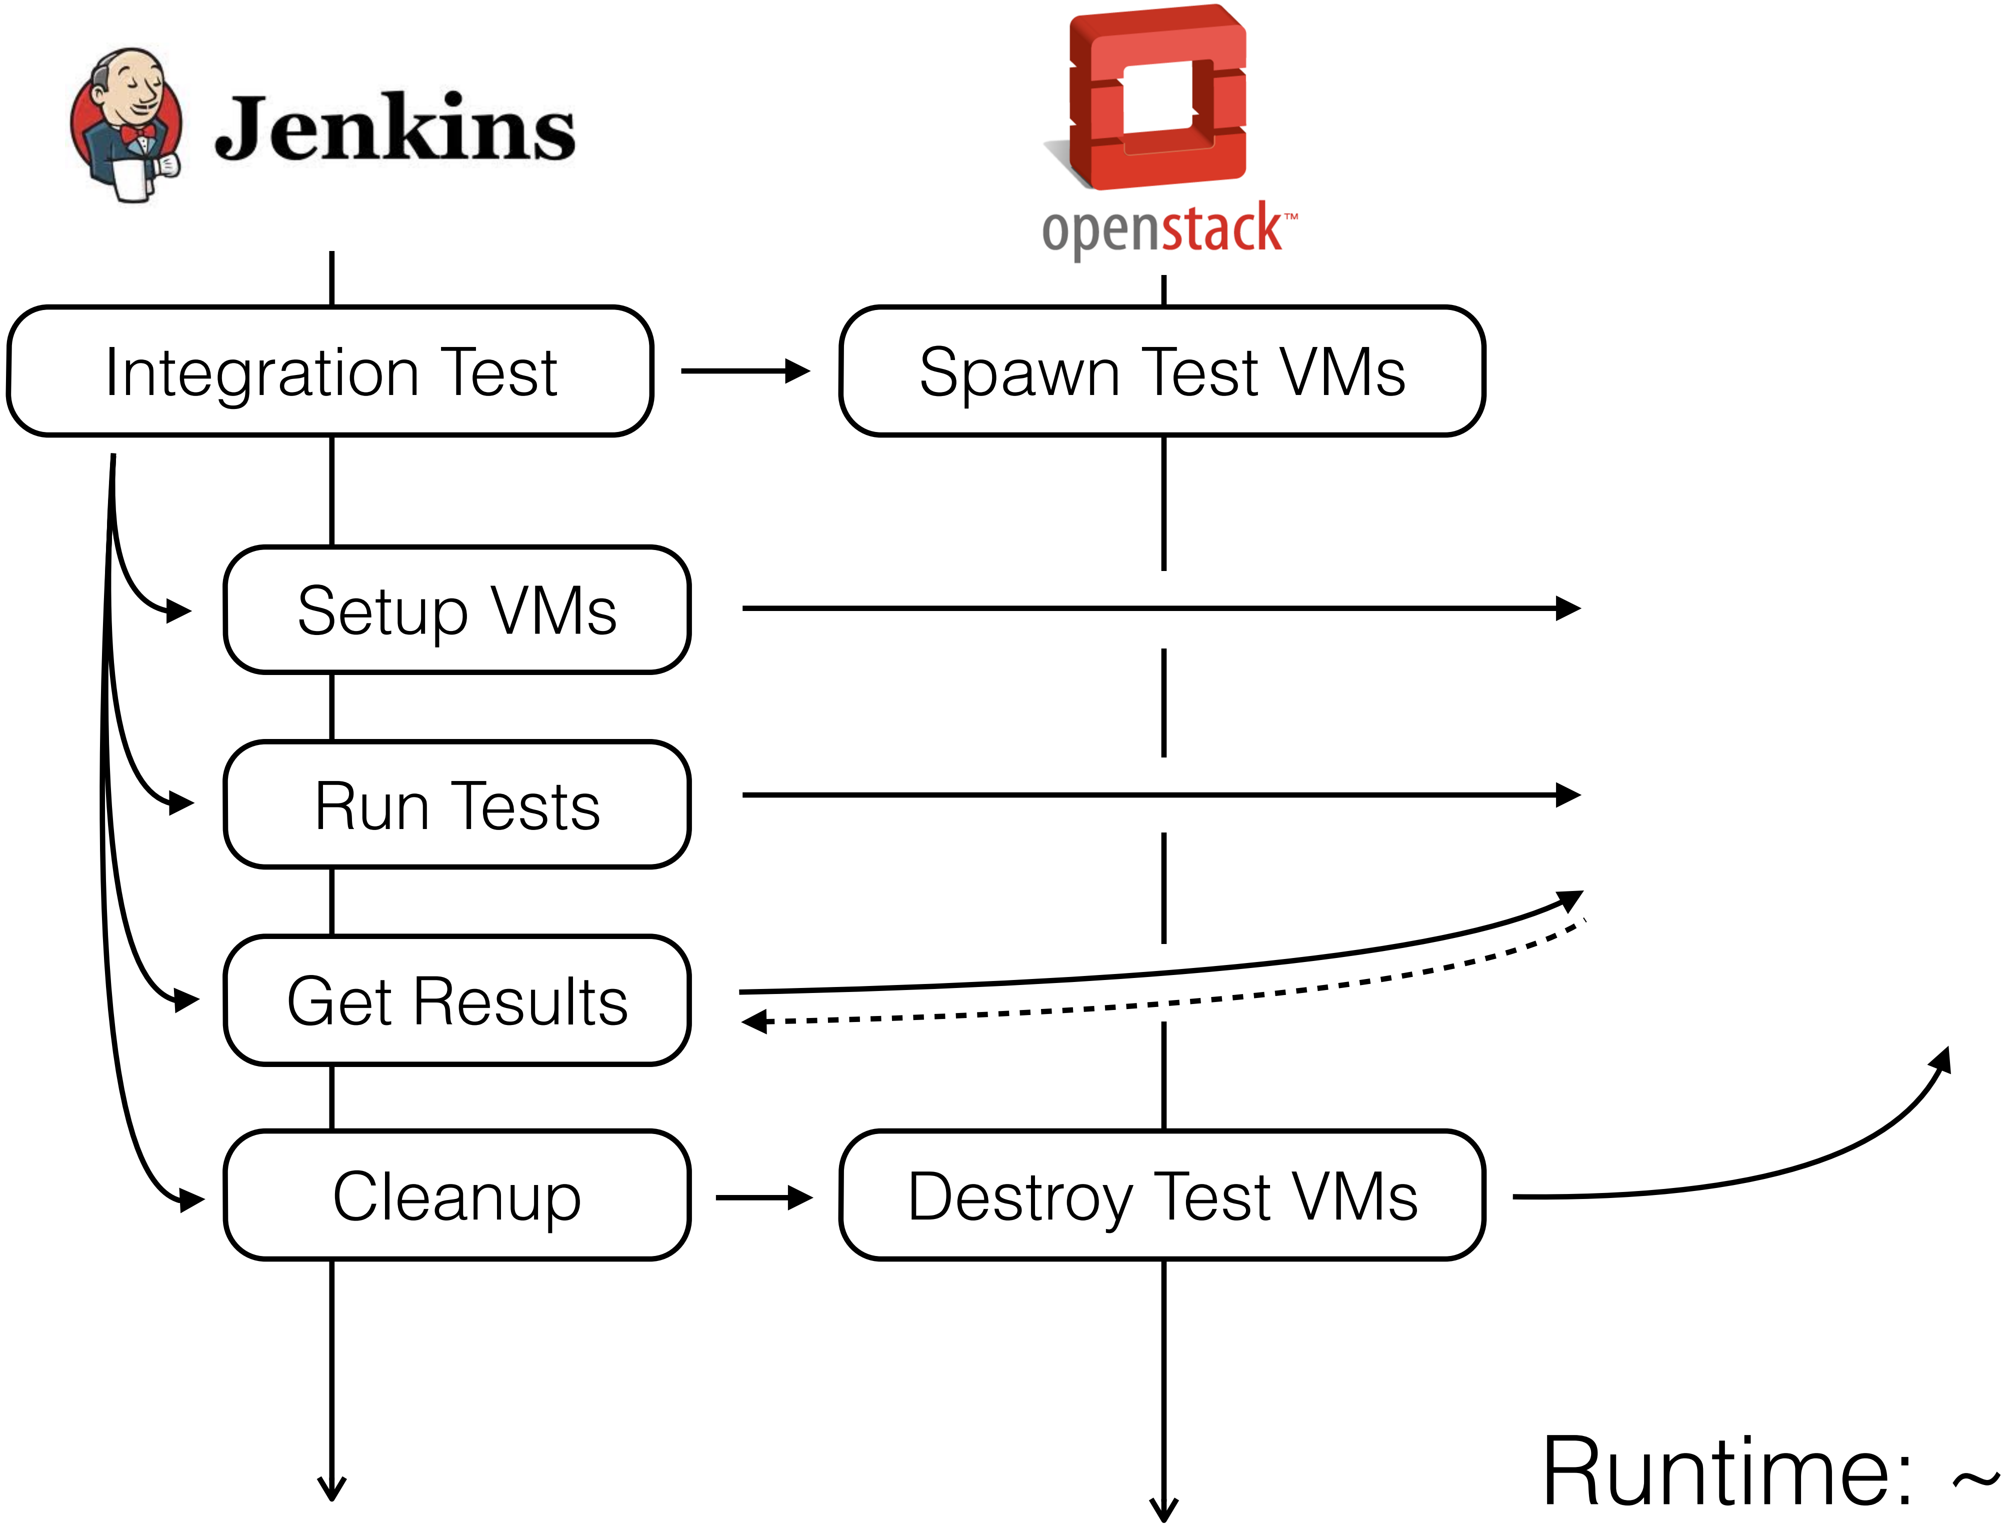

Runtime: ~16 hours

Full Integration Test Execution

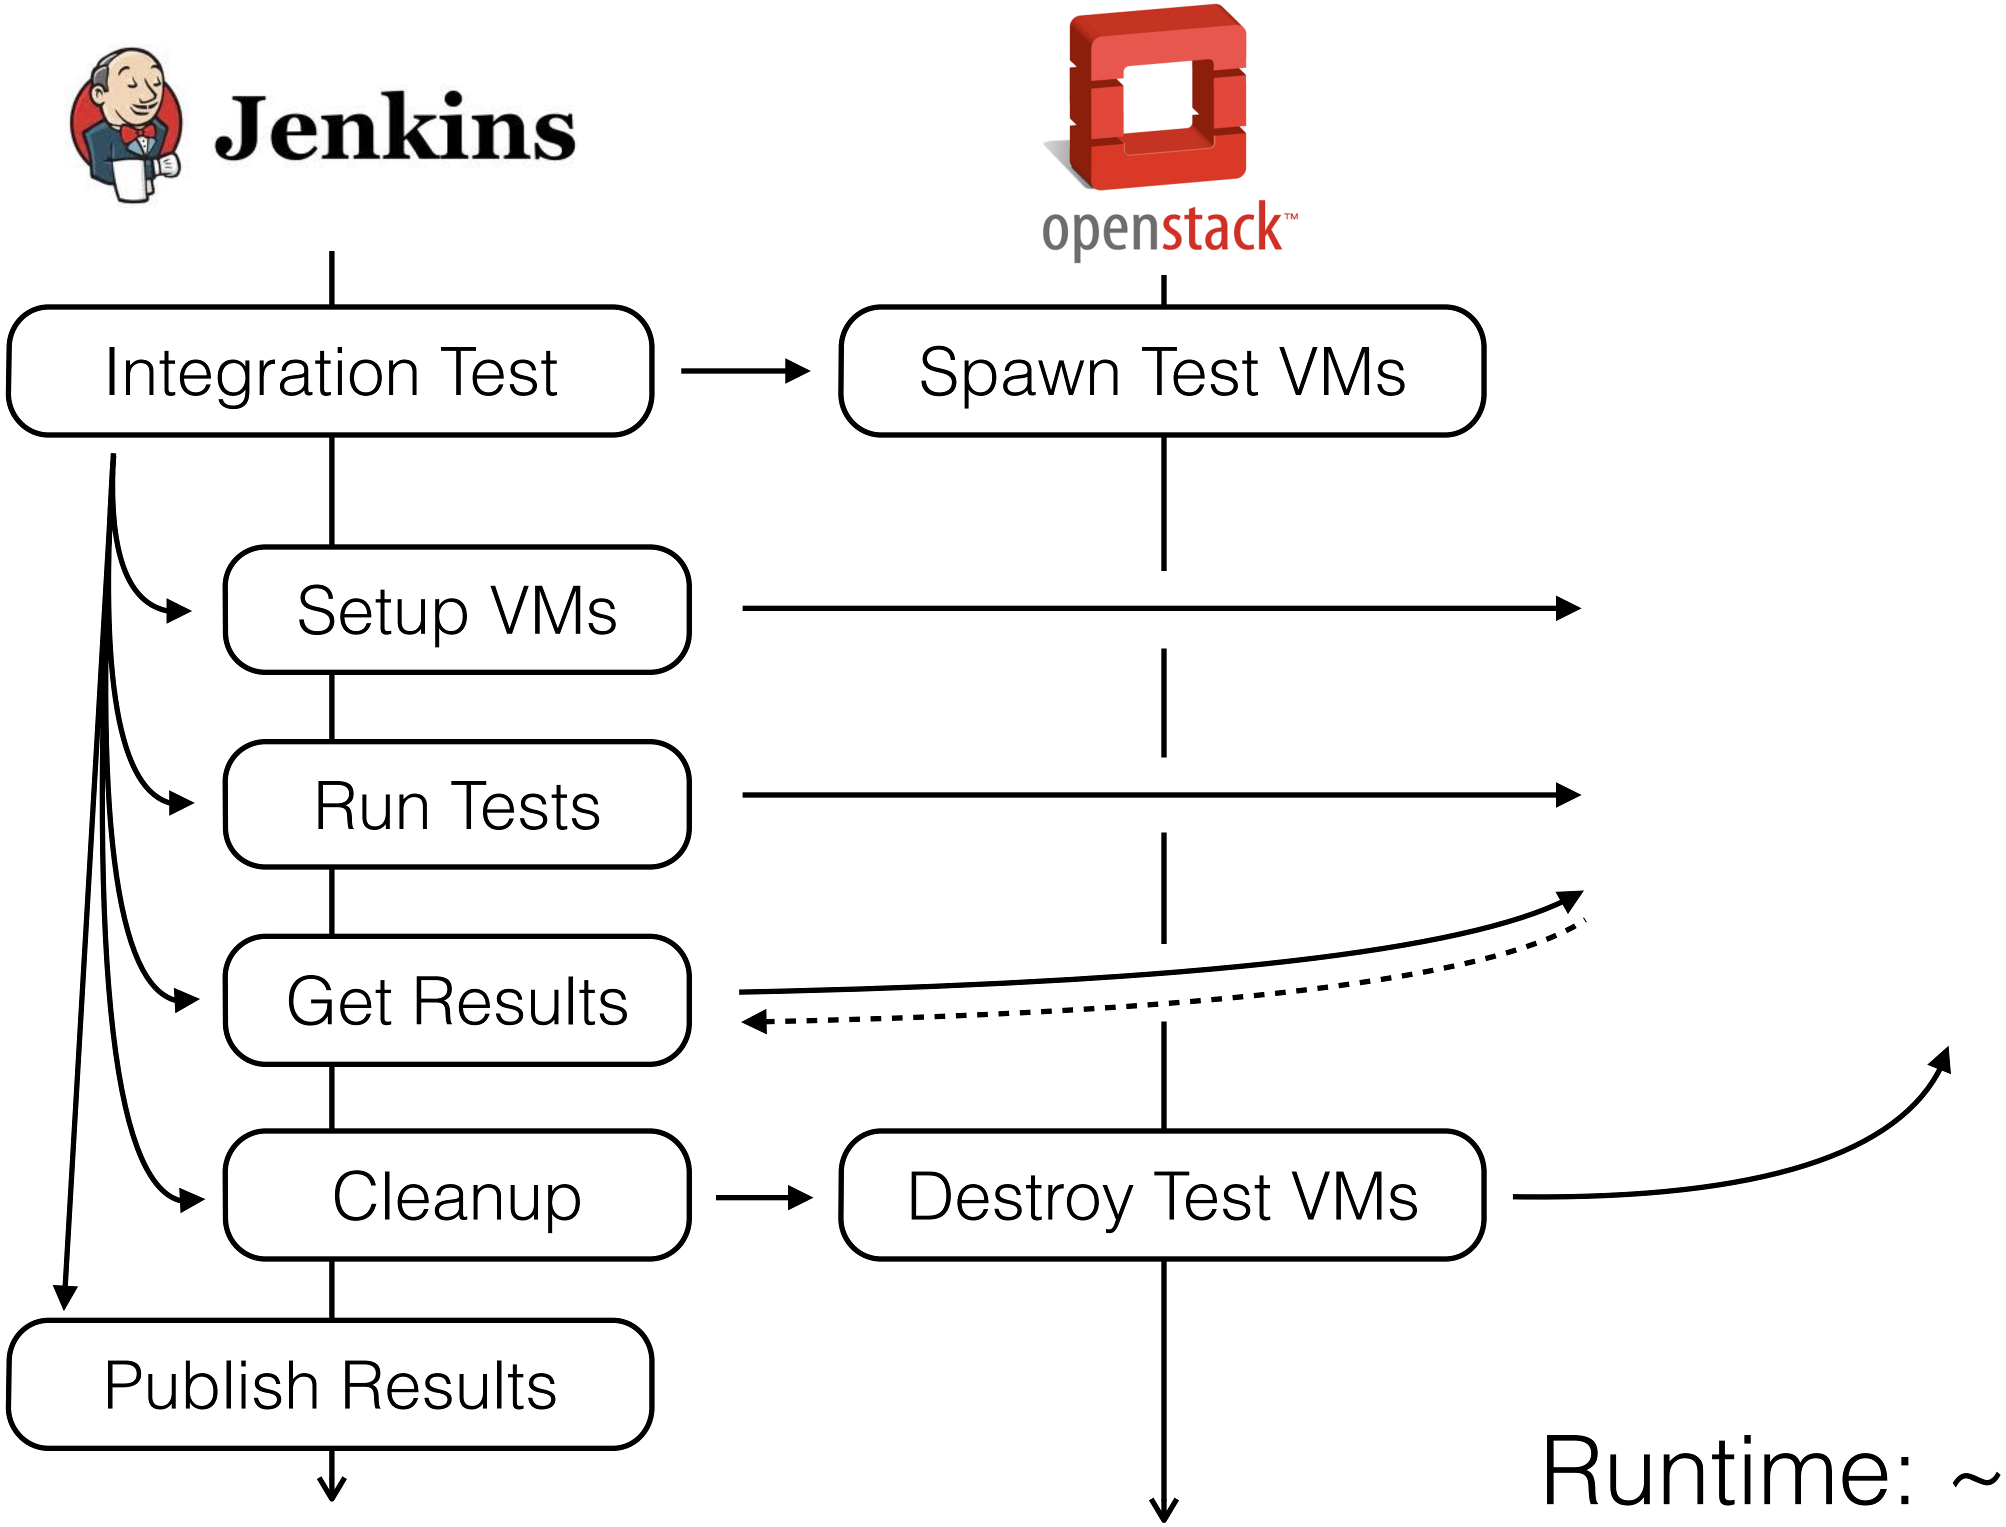

Runtime: ~16 hours

Server Migration as Done at CERN (2014)

- **Gradual roll-out** from April to September **2014**

- April 11th First migrated repository: **geant4**.cern.ch

  ↳ A couple of minor fixes in CernVM-FS 2.1.19

- April/May Migrated all “small” repositories

  boss.cern.ch, **belle**.cern.ch, **grid**.cern.ch, **na49**.cern.ch, **na61**.cern.ch

- August 5th CernVM-FS 2.1.19 installed on all sites

  (decided in: WLCG Ops Meeting - June 5th [1])

- August/ September Migrated “large” repositories

  sft.cern.ch, **ams**.cern.ch, **atlas**.cern.ch, **alice**.cern.ch, **atlas-condb**.cern.ch, **cms**.cern.ch, **lhcb**.cern.ch

- Sept. 2nd End of Life of CernVM-FS 2.0.x at CERN

[1] https://twiki.cern.ch/twiki/bin/view/LCG/WLCGOpsMinutes140605#Migration_of_the_CVMFS_servers

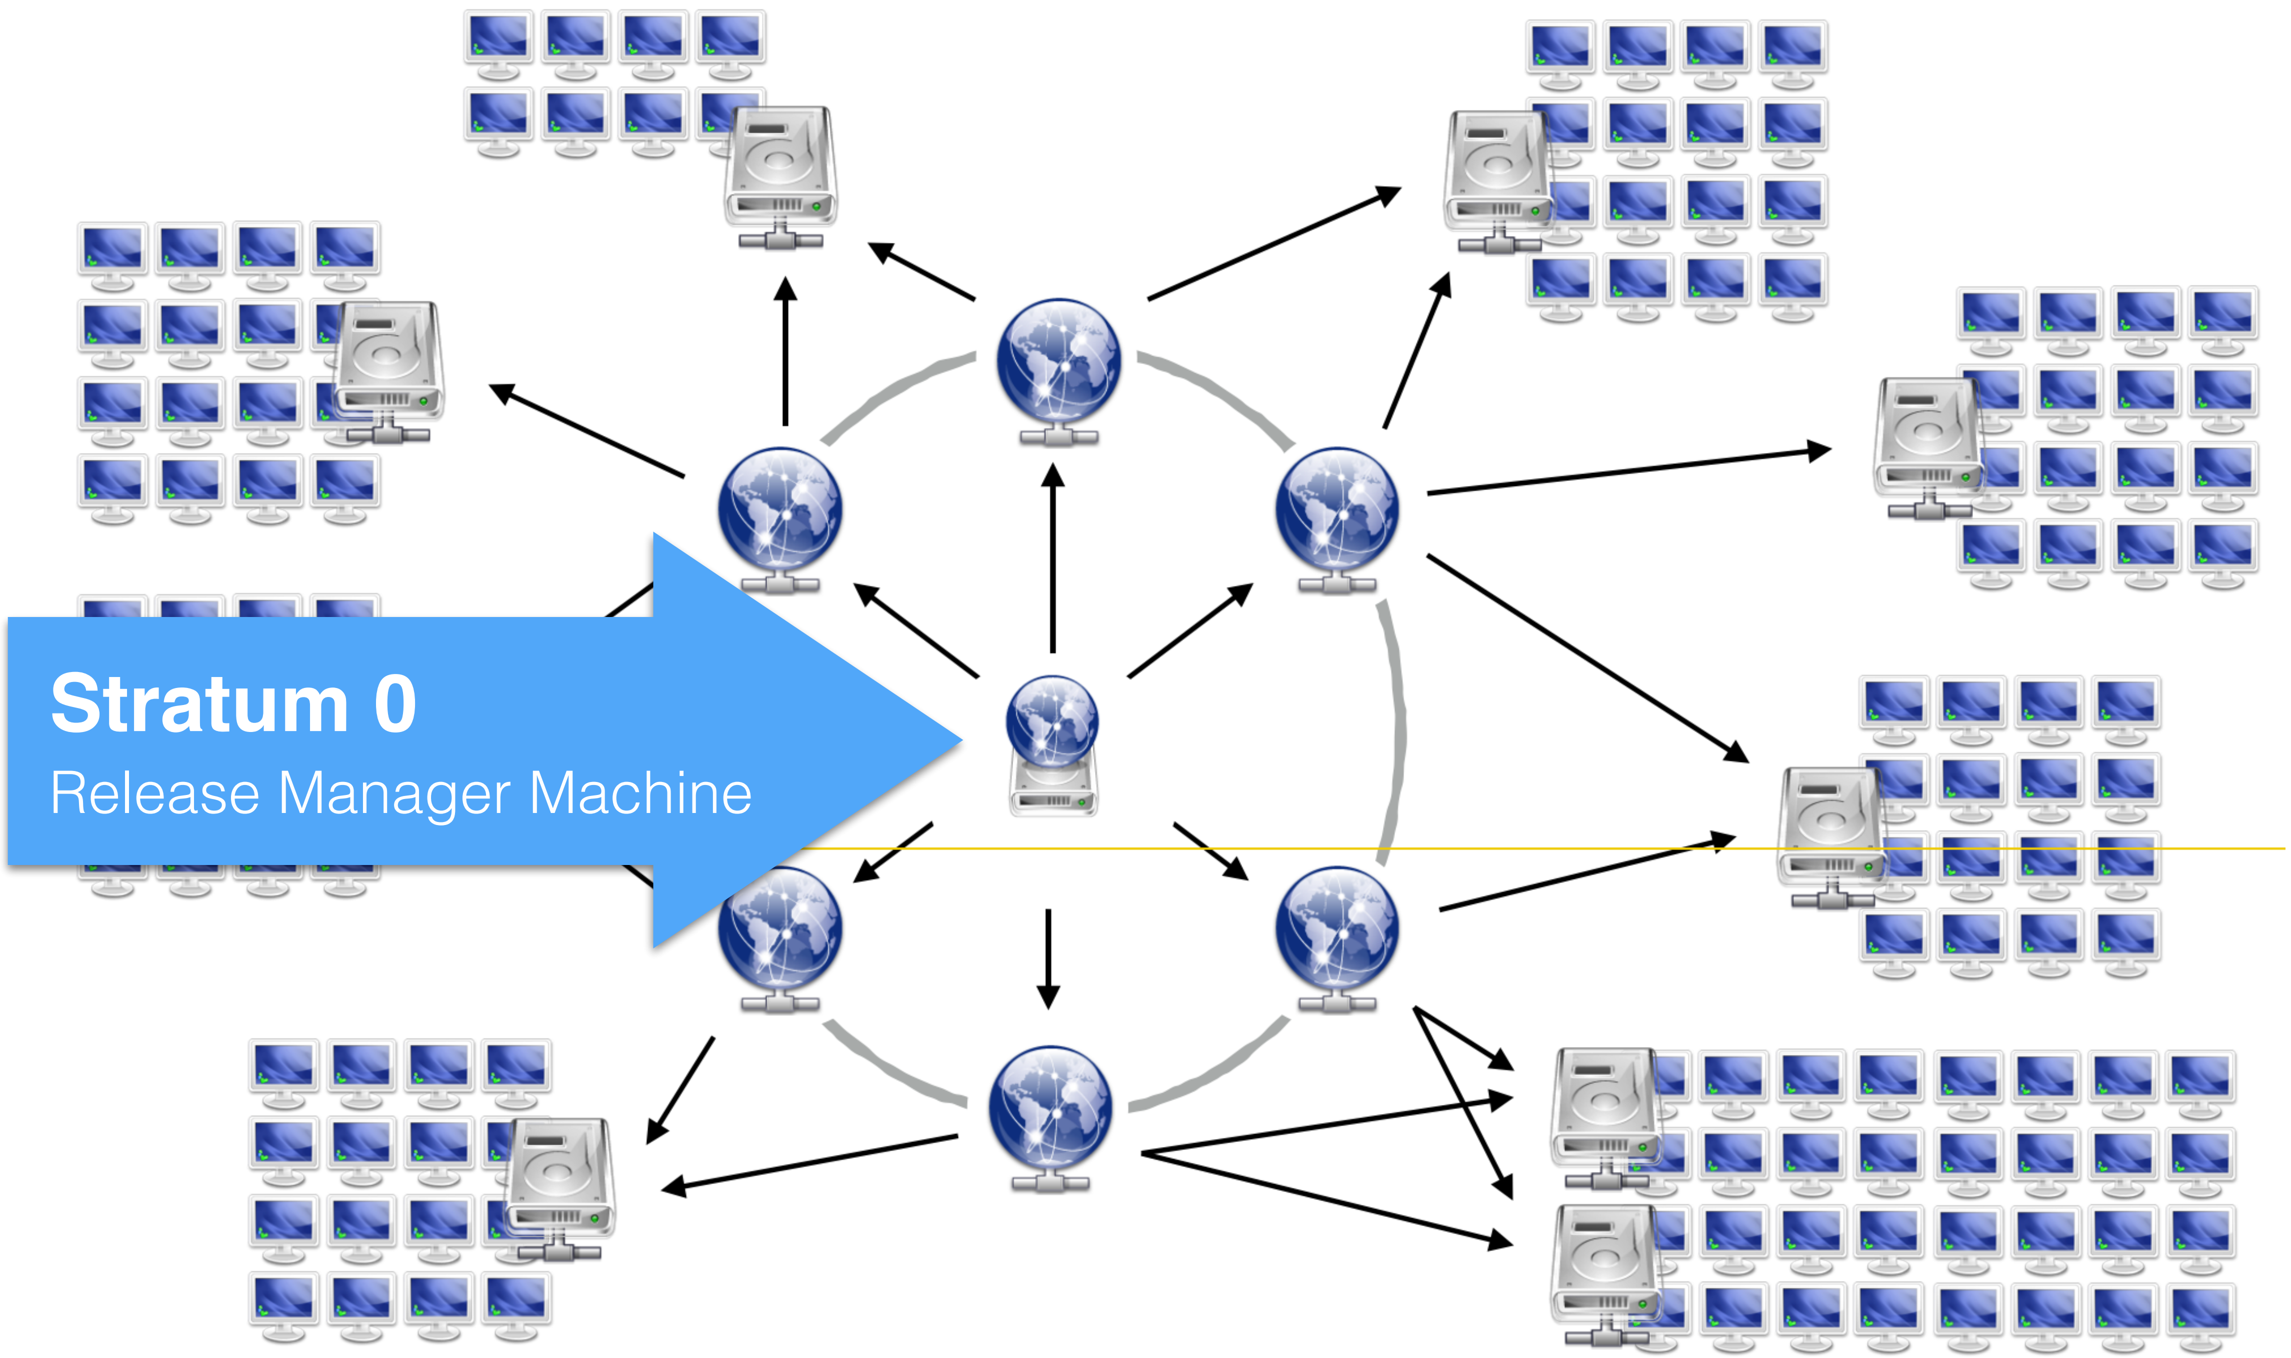

Stratum 0

Release Manager Machine

What is CernVM-FS?

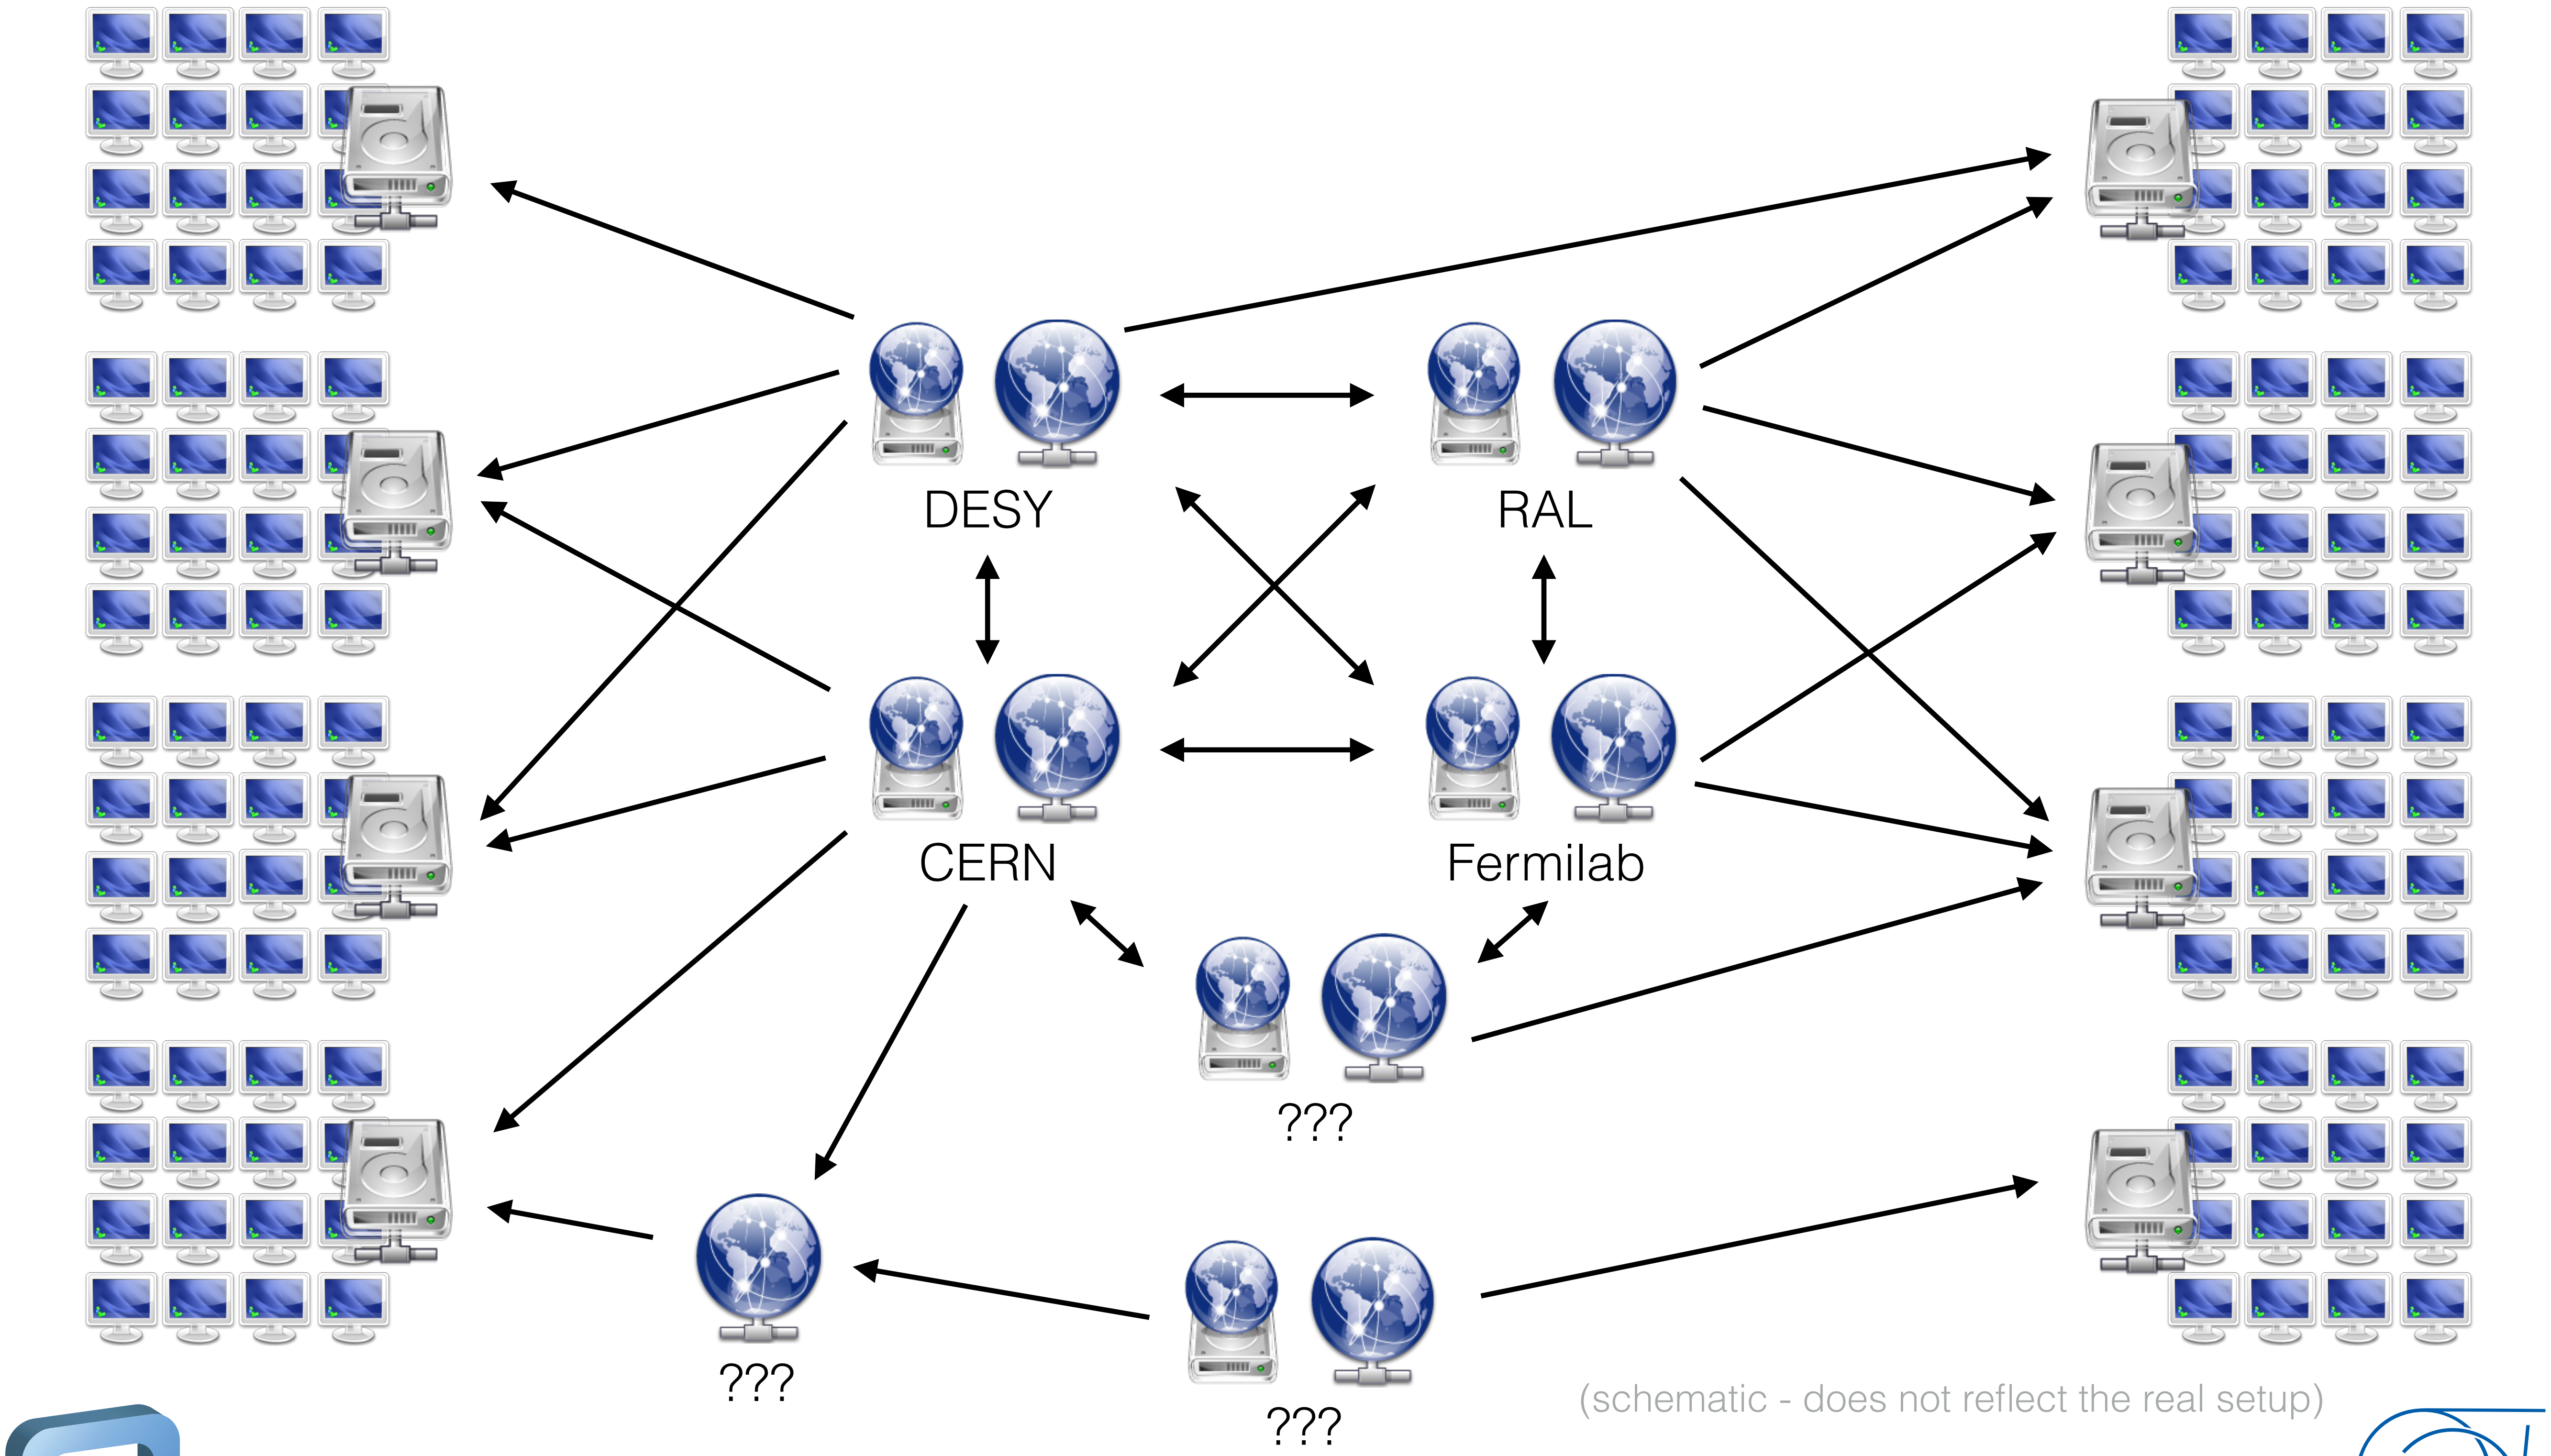

(schematic - does not reflect the real setup)

CernVM-FS Client - Control Flow

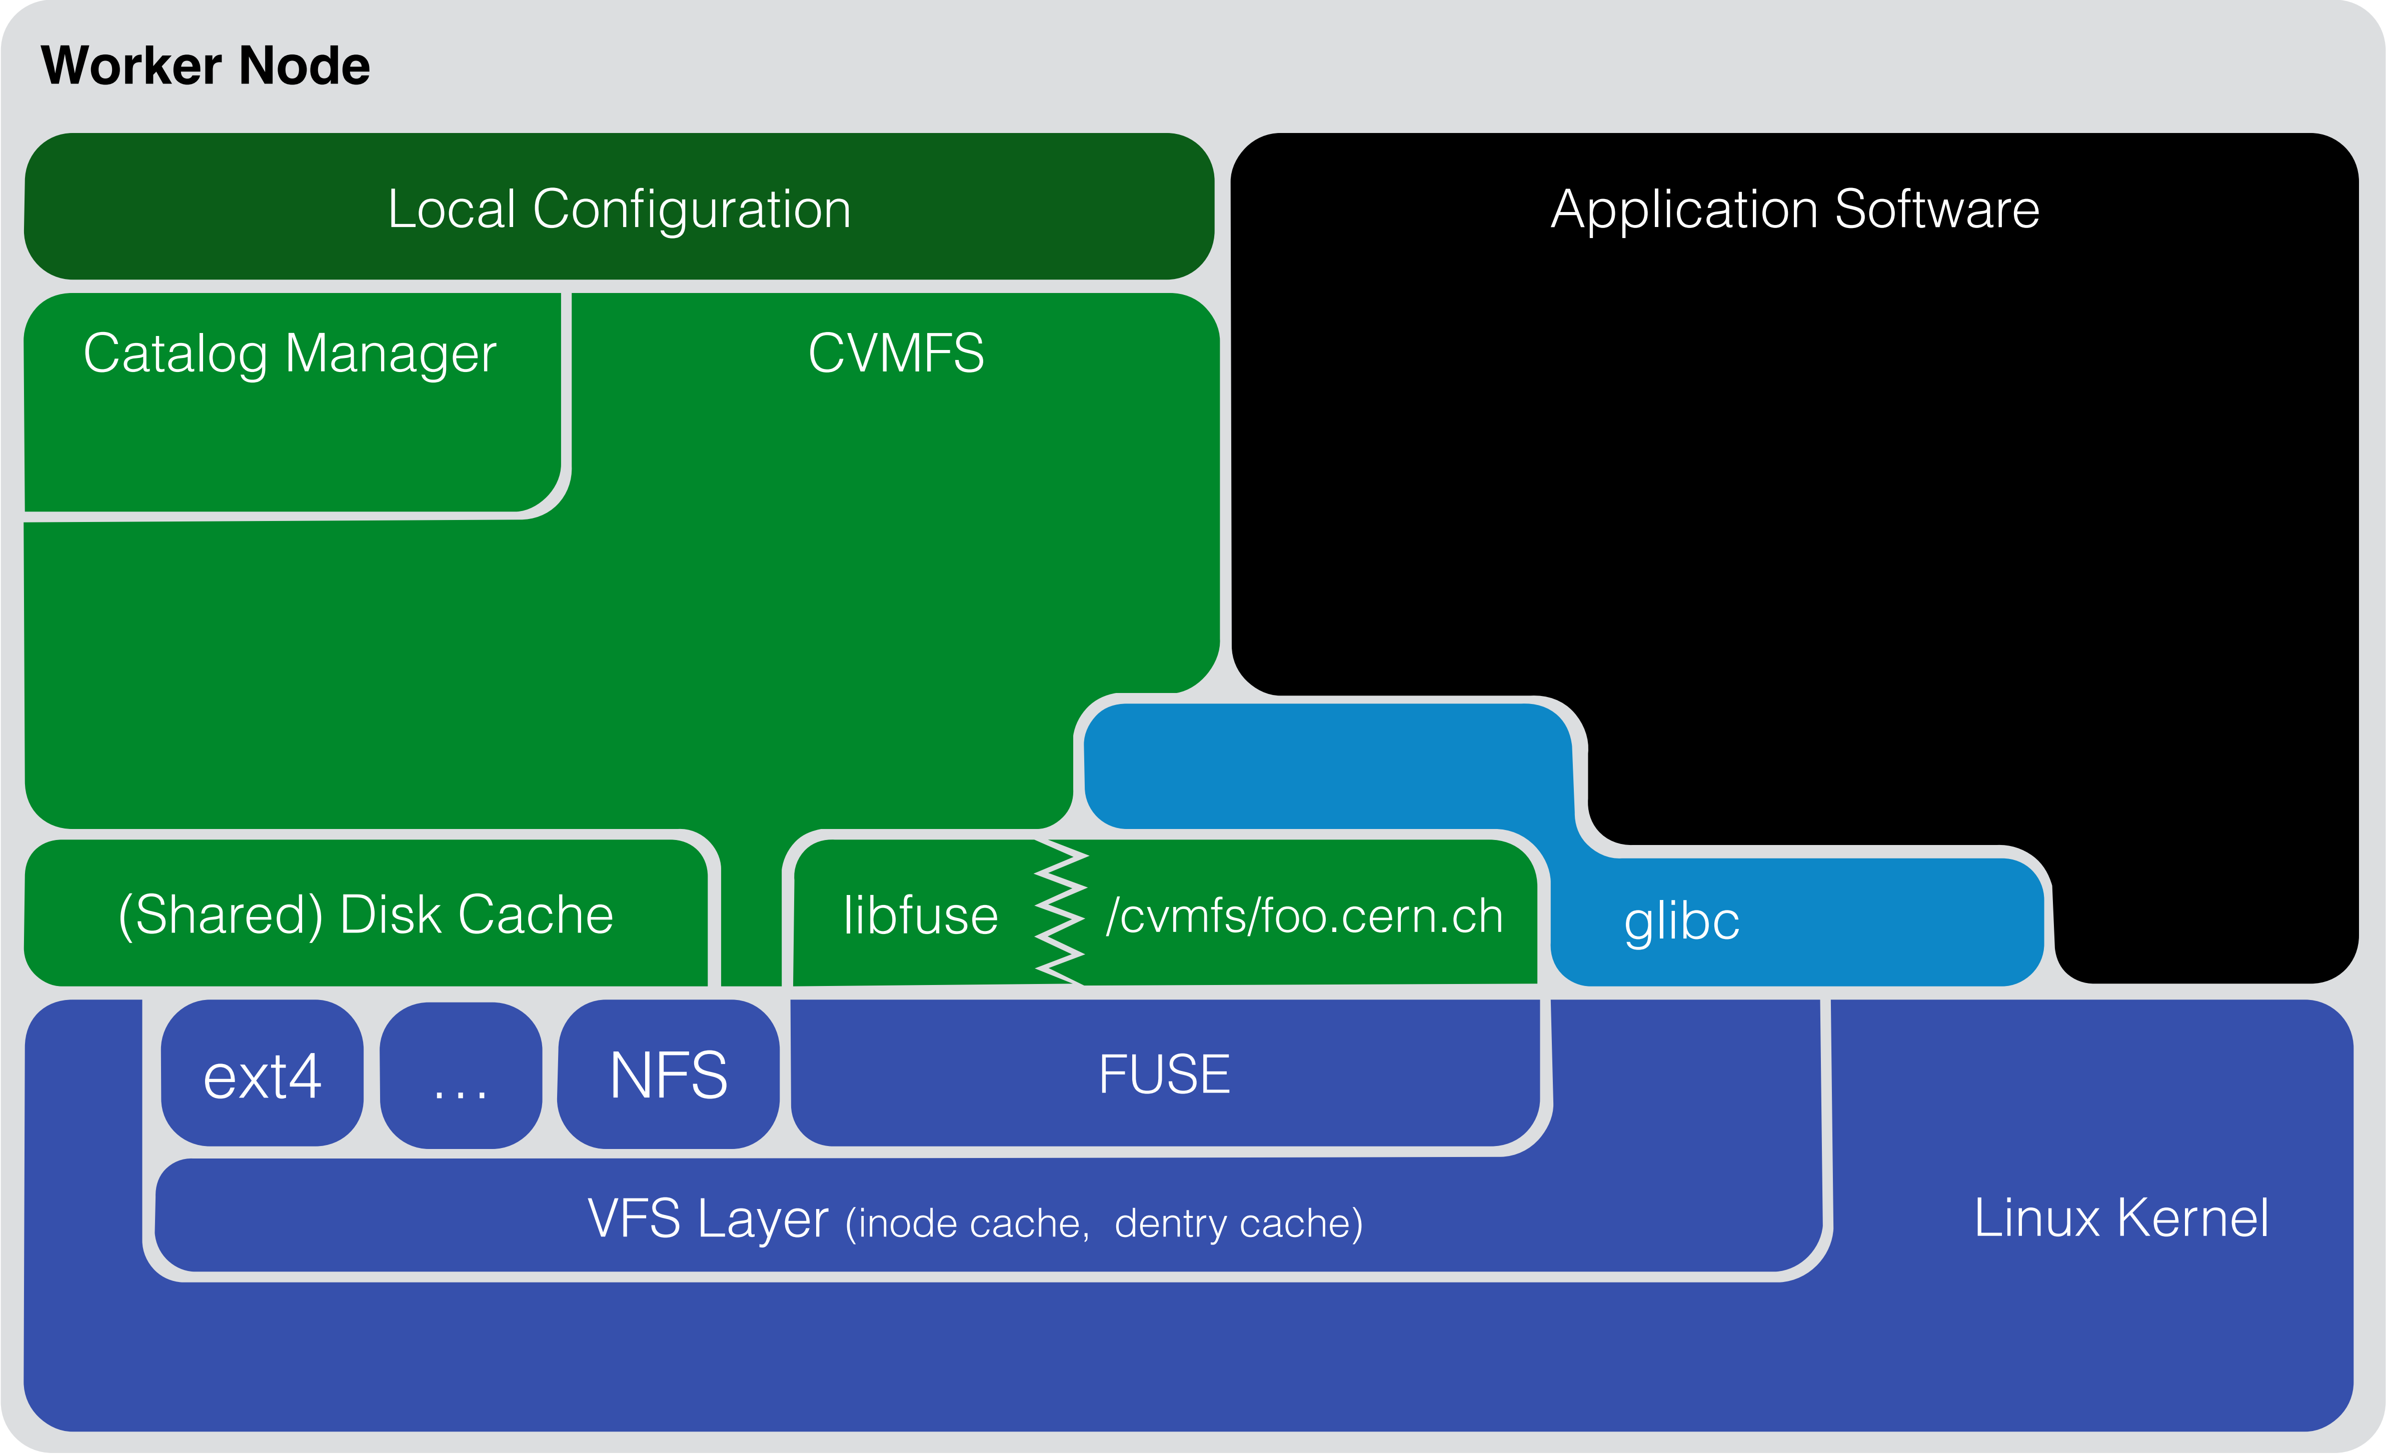

CernVM-FS Client - Control Flow

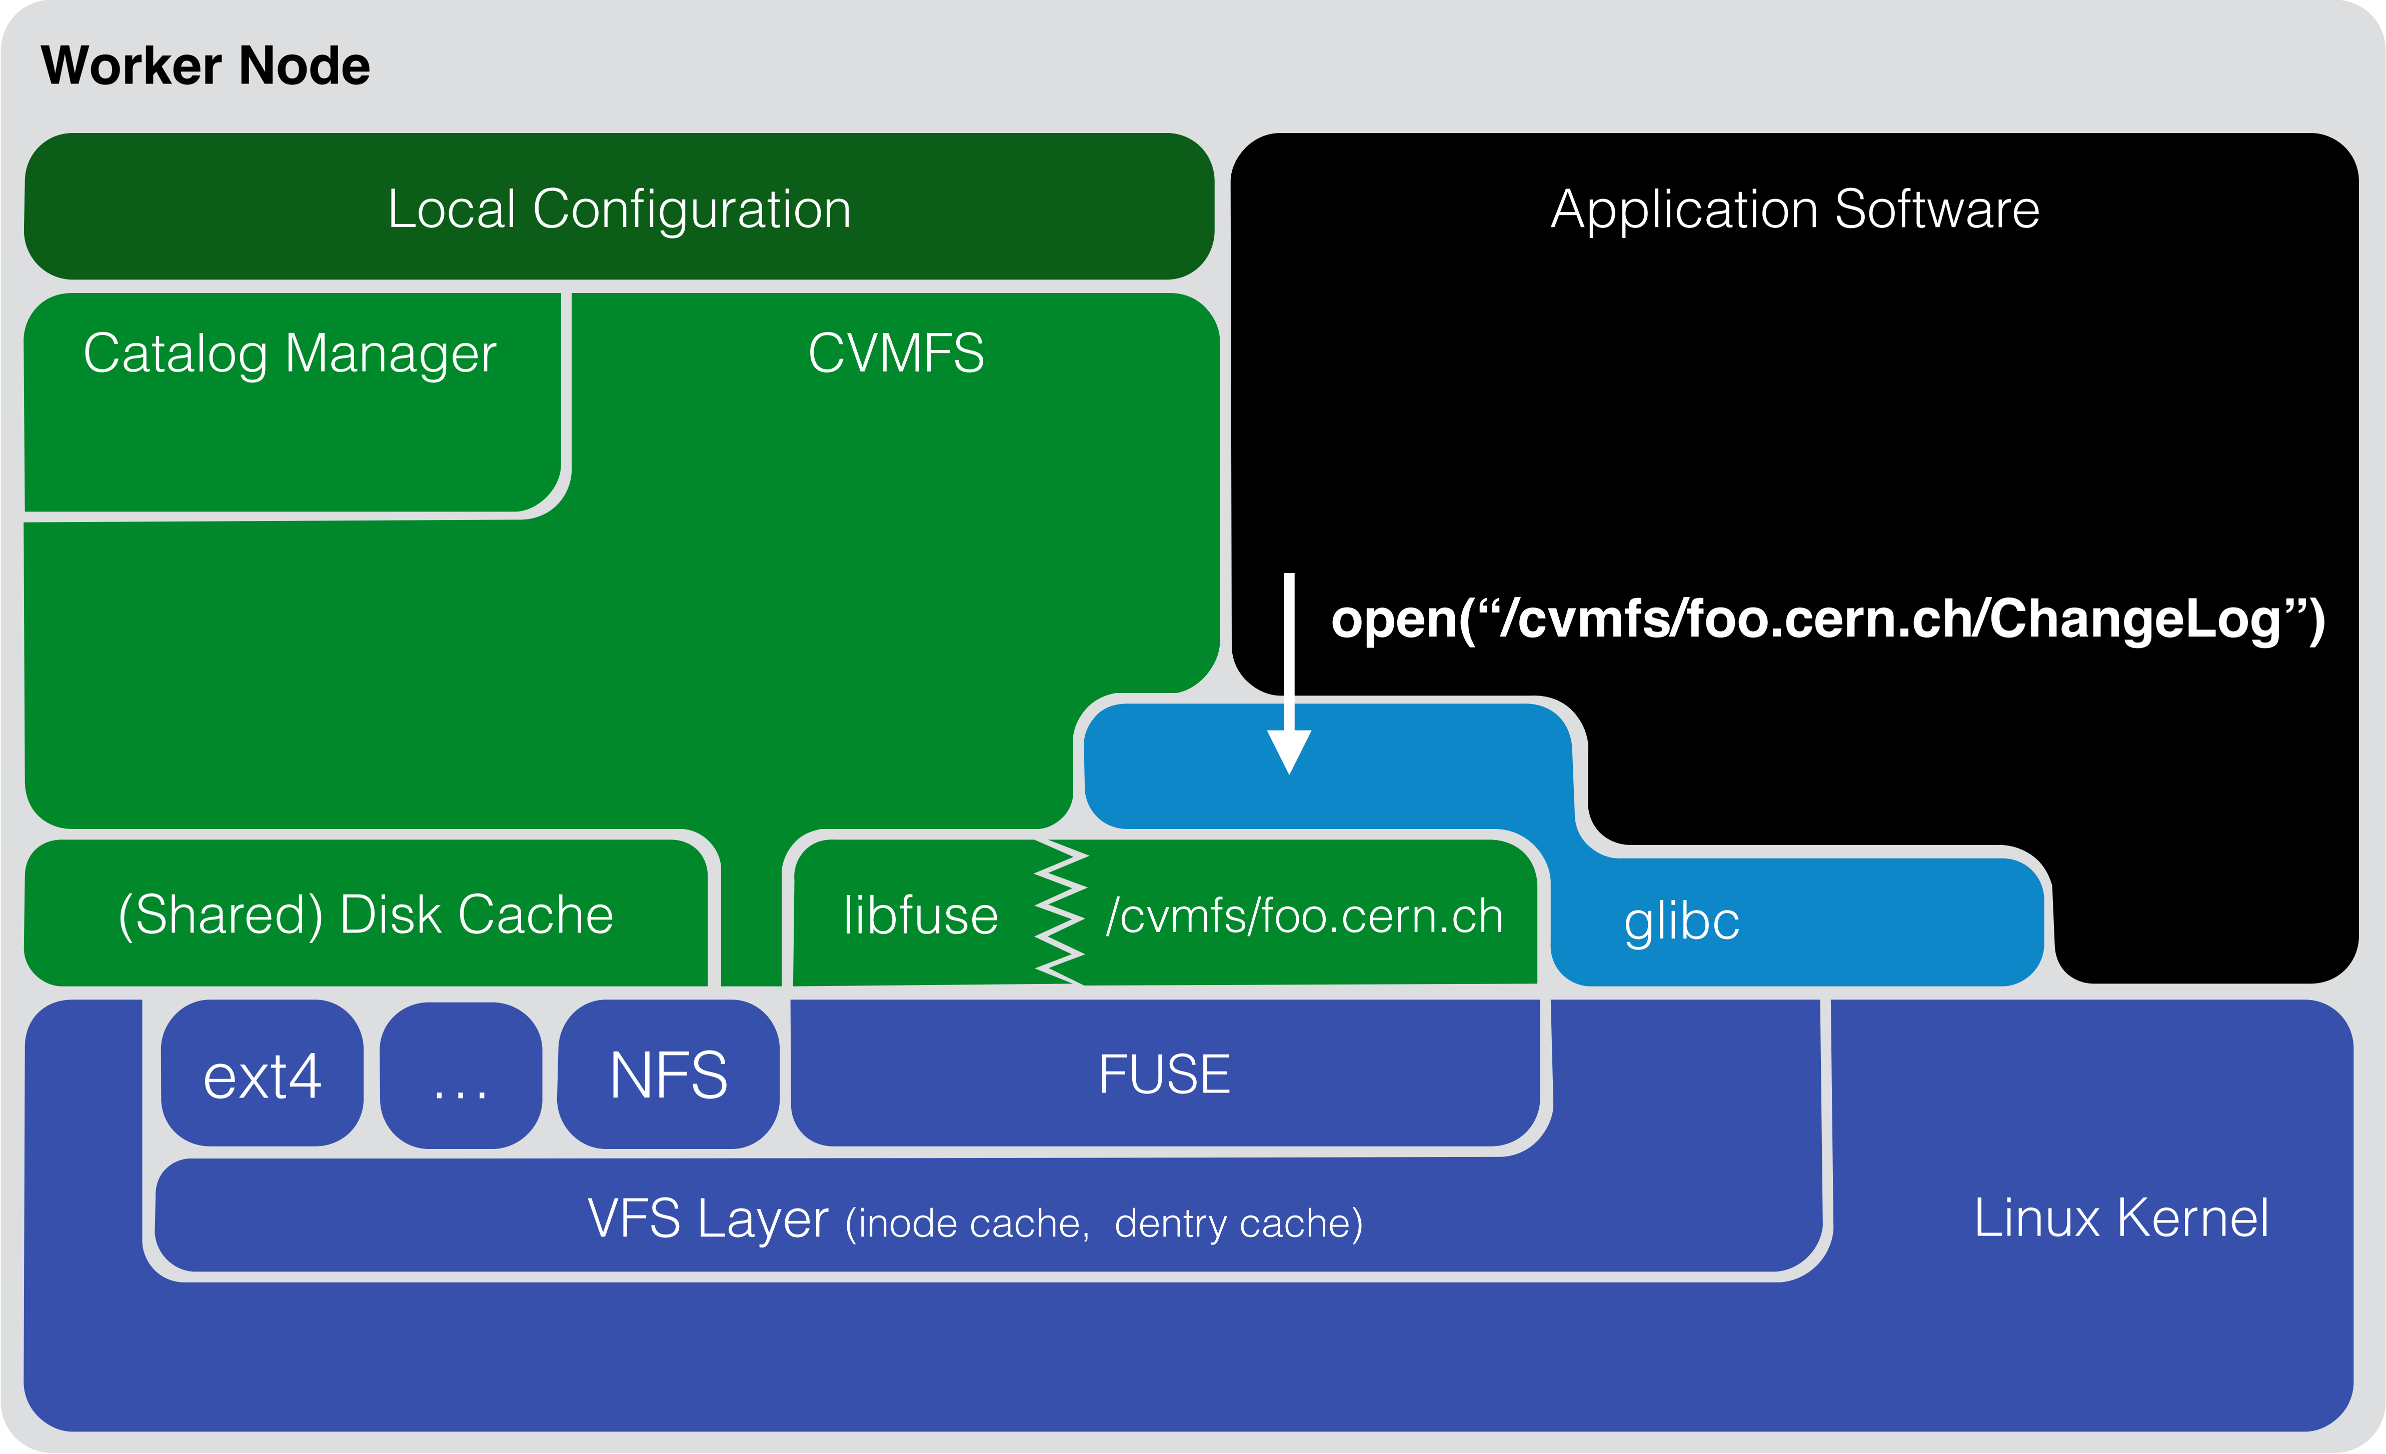

CernVM-FS Client - Control Flow

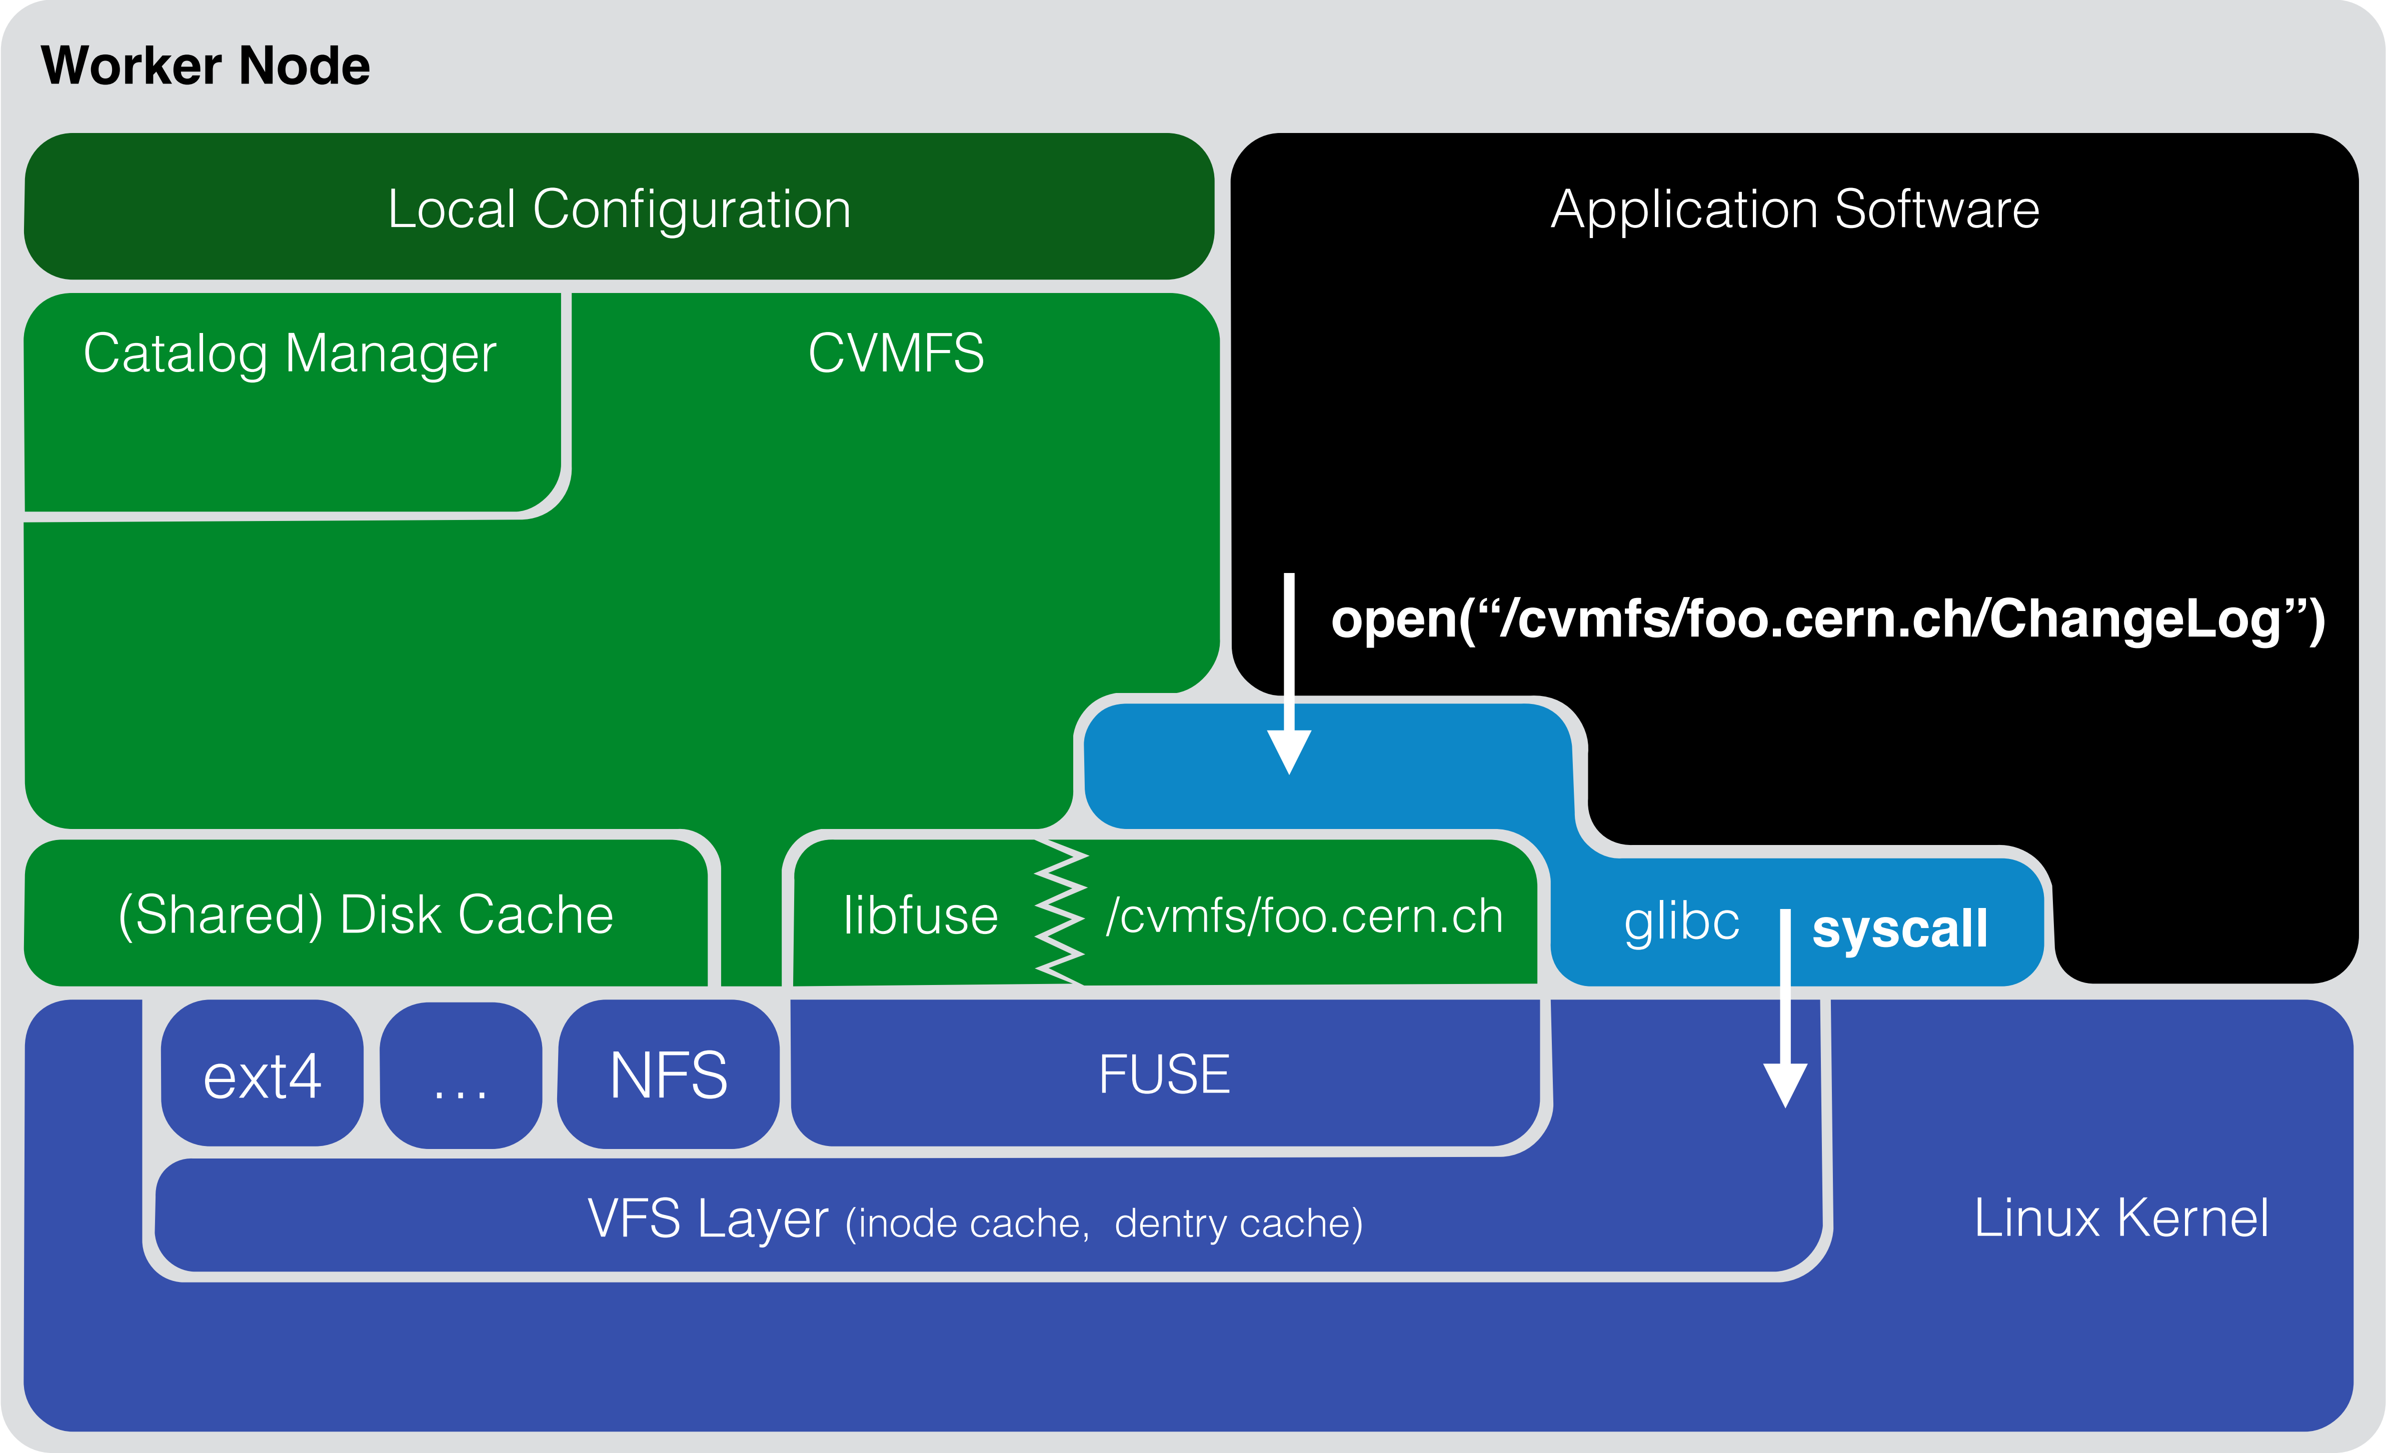

CernVM-FS Client - Control Flow

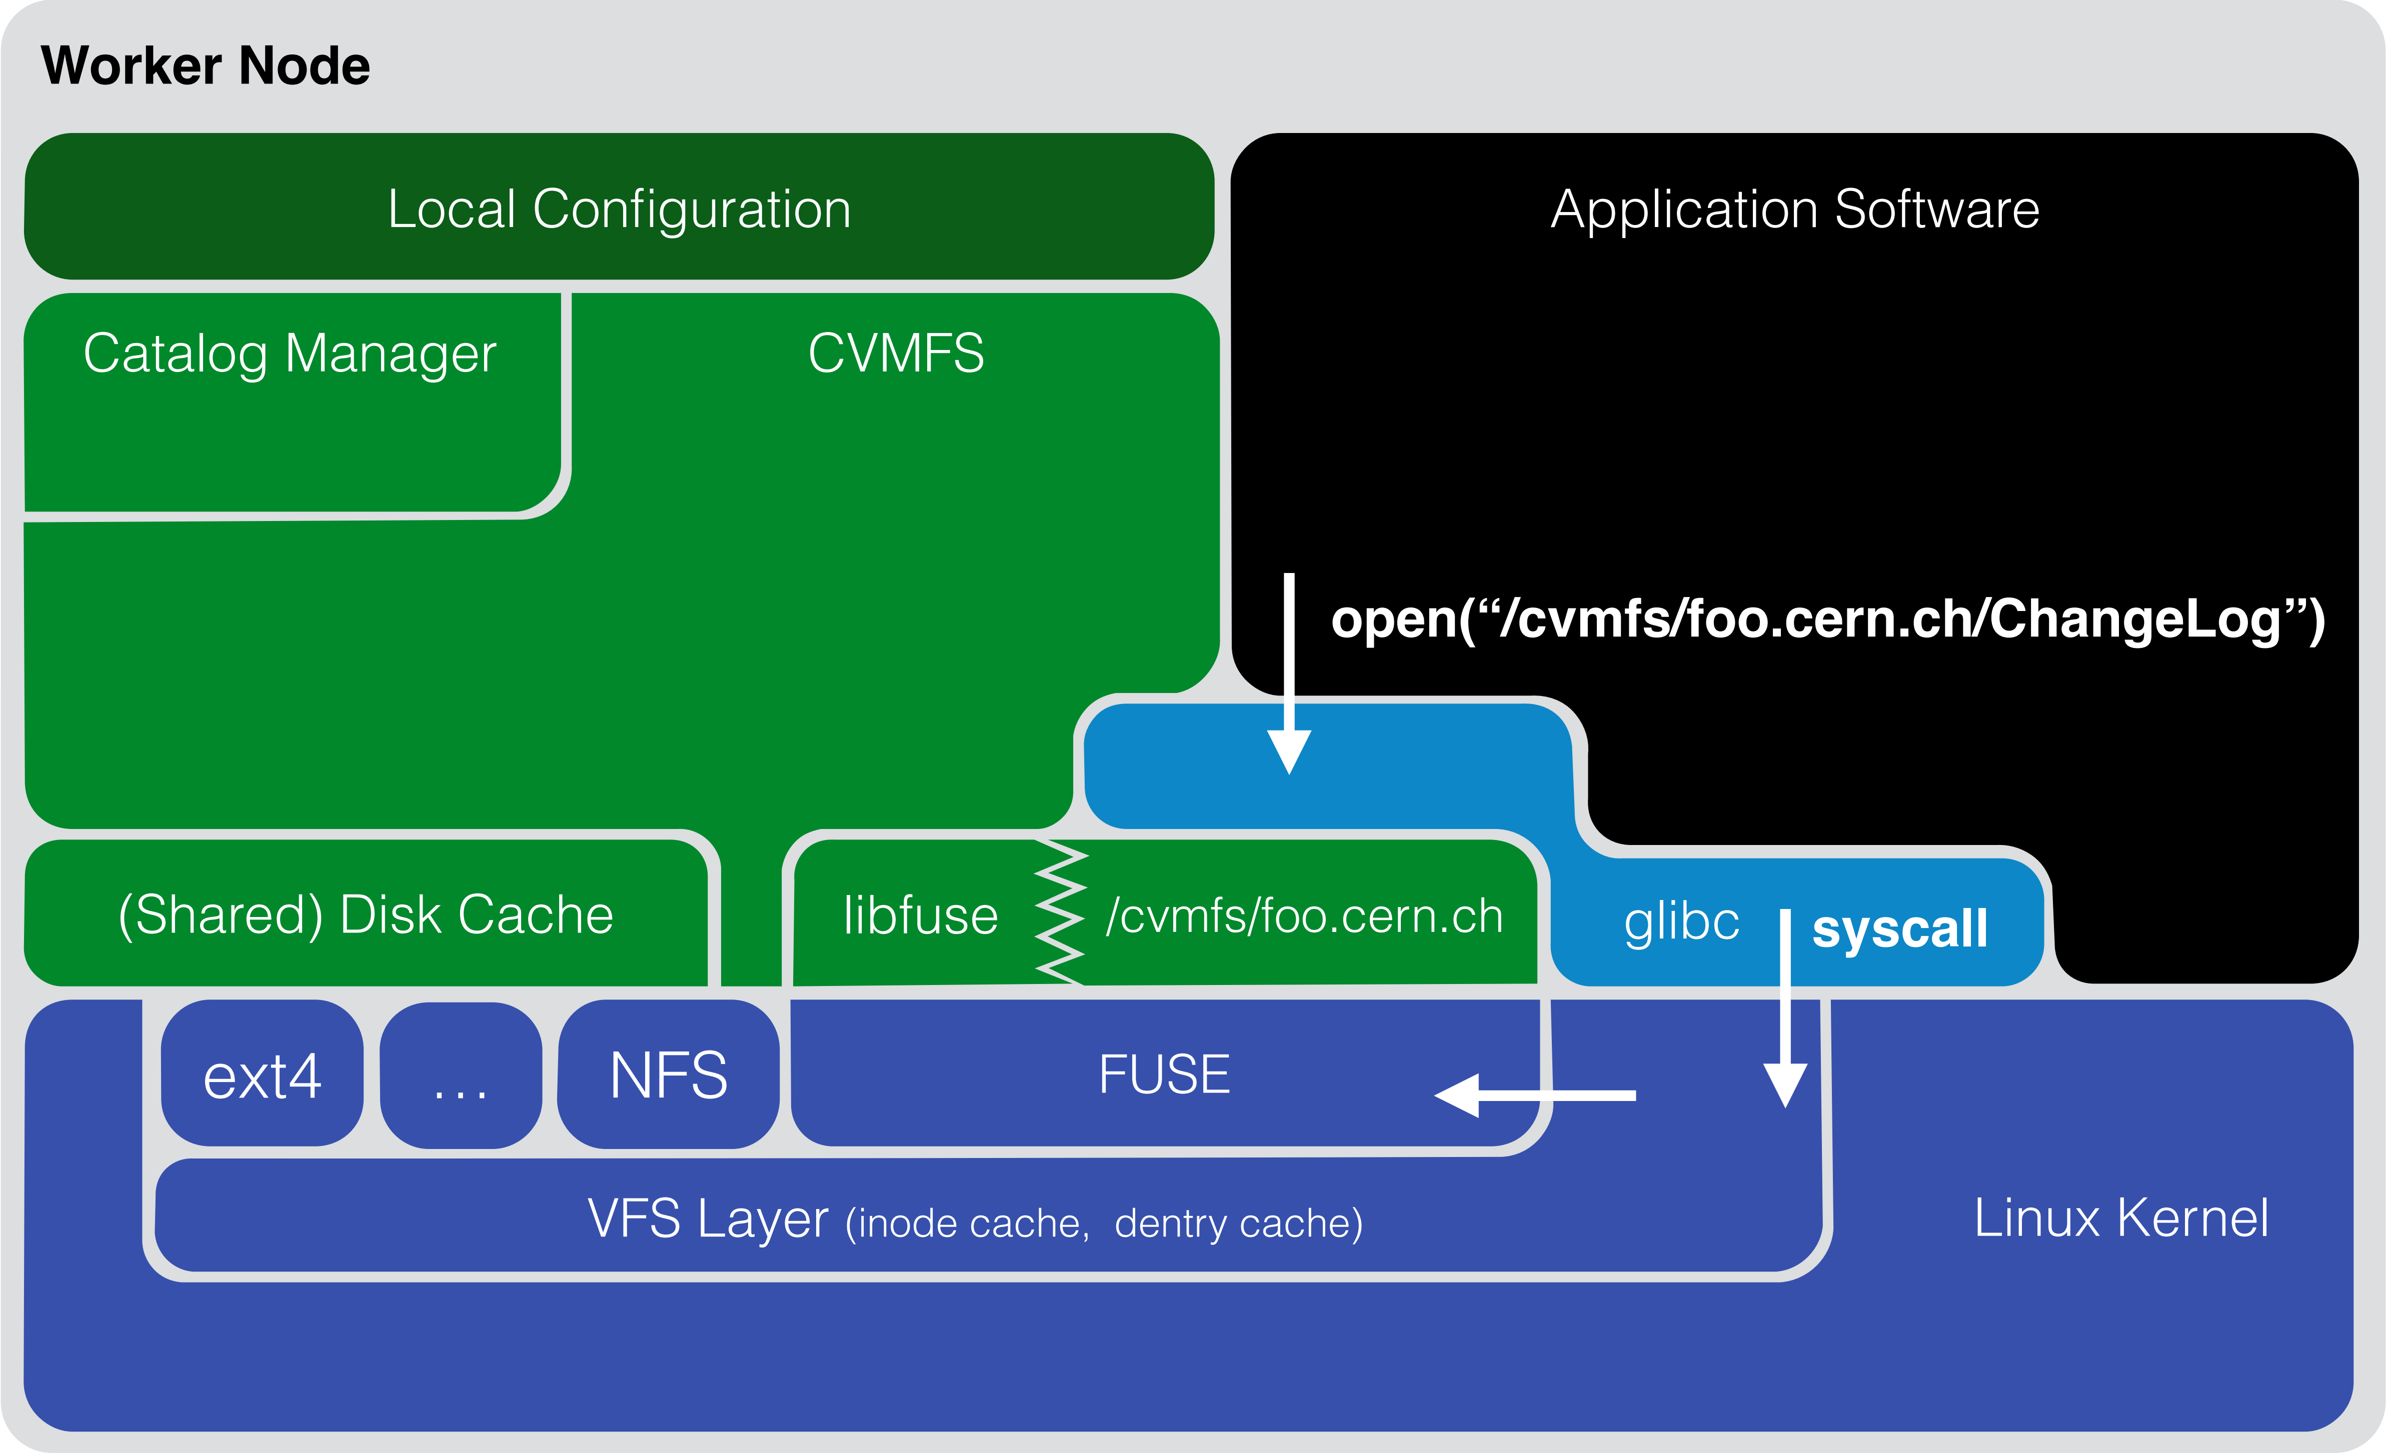

CernVM-FS Client - Control Flow

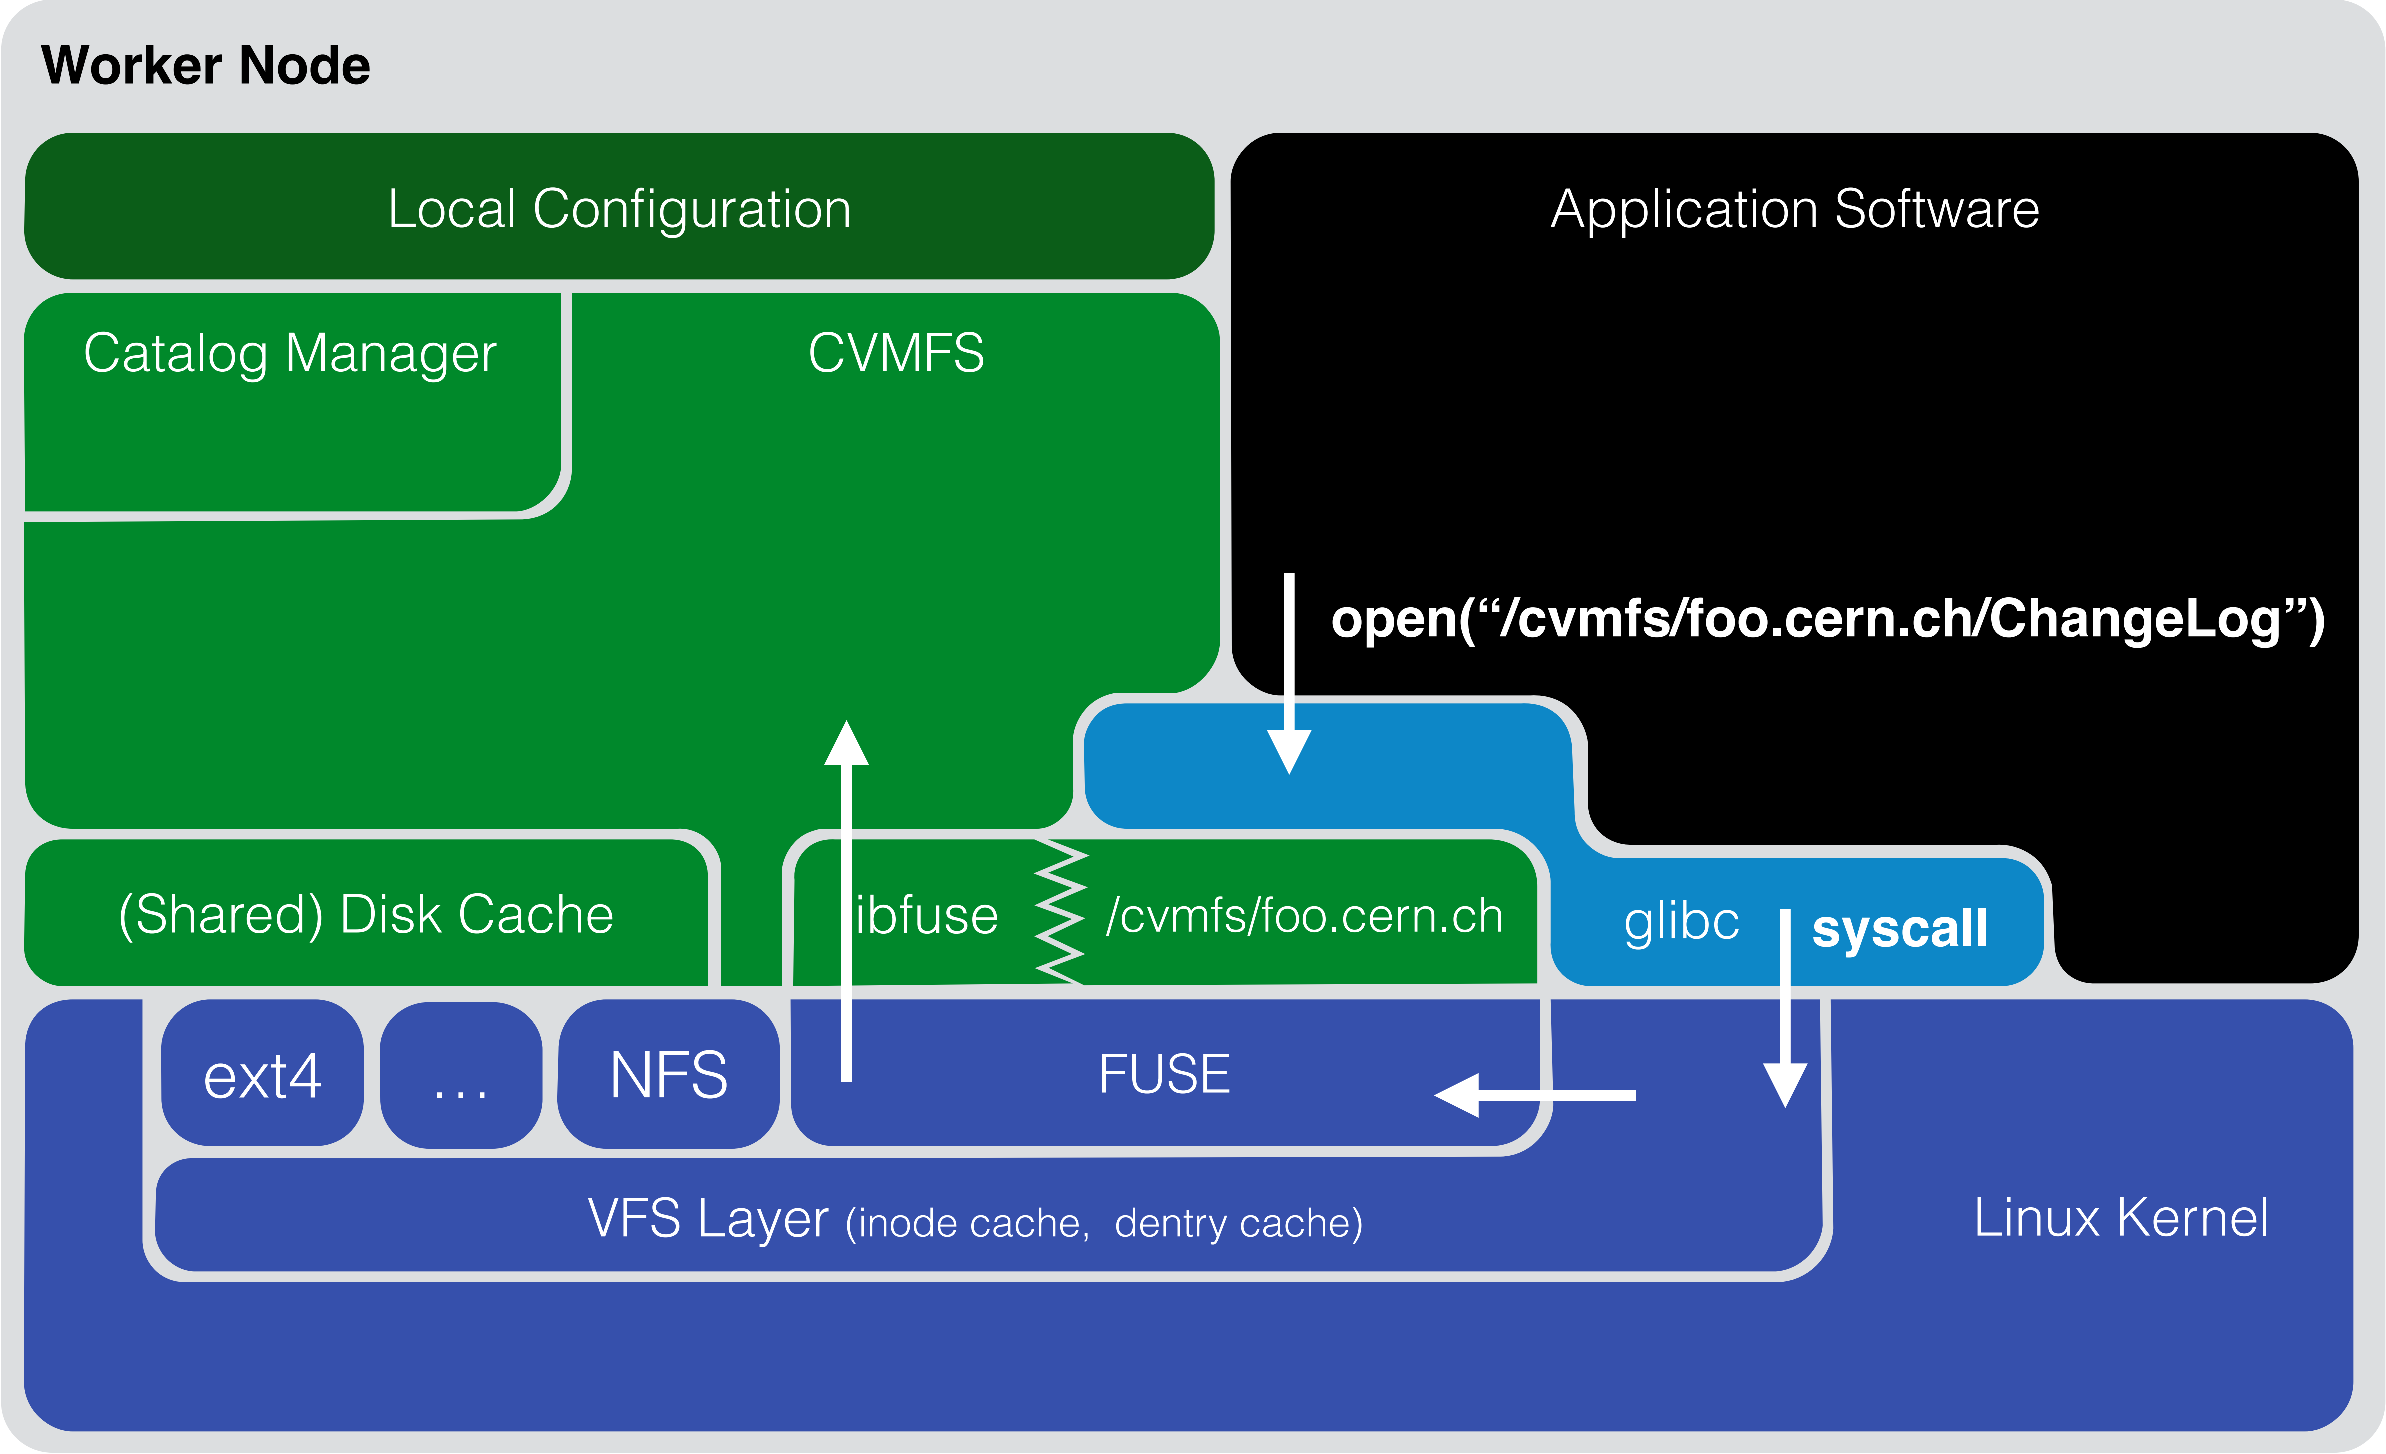

CernVM-FS Client - Control Flow

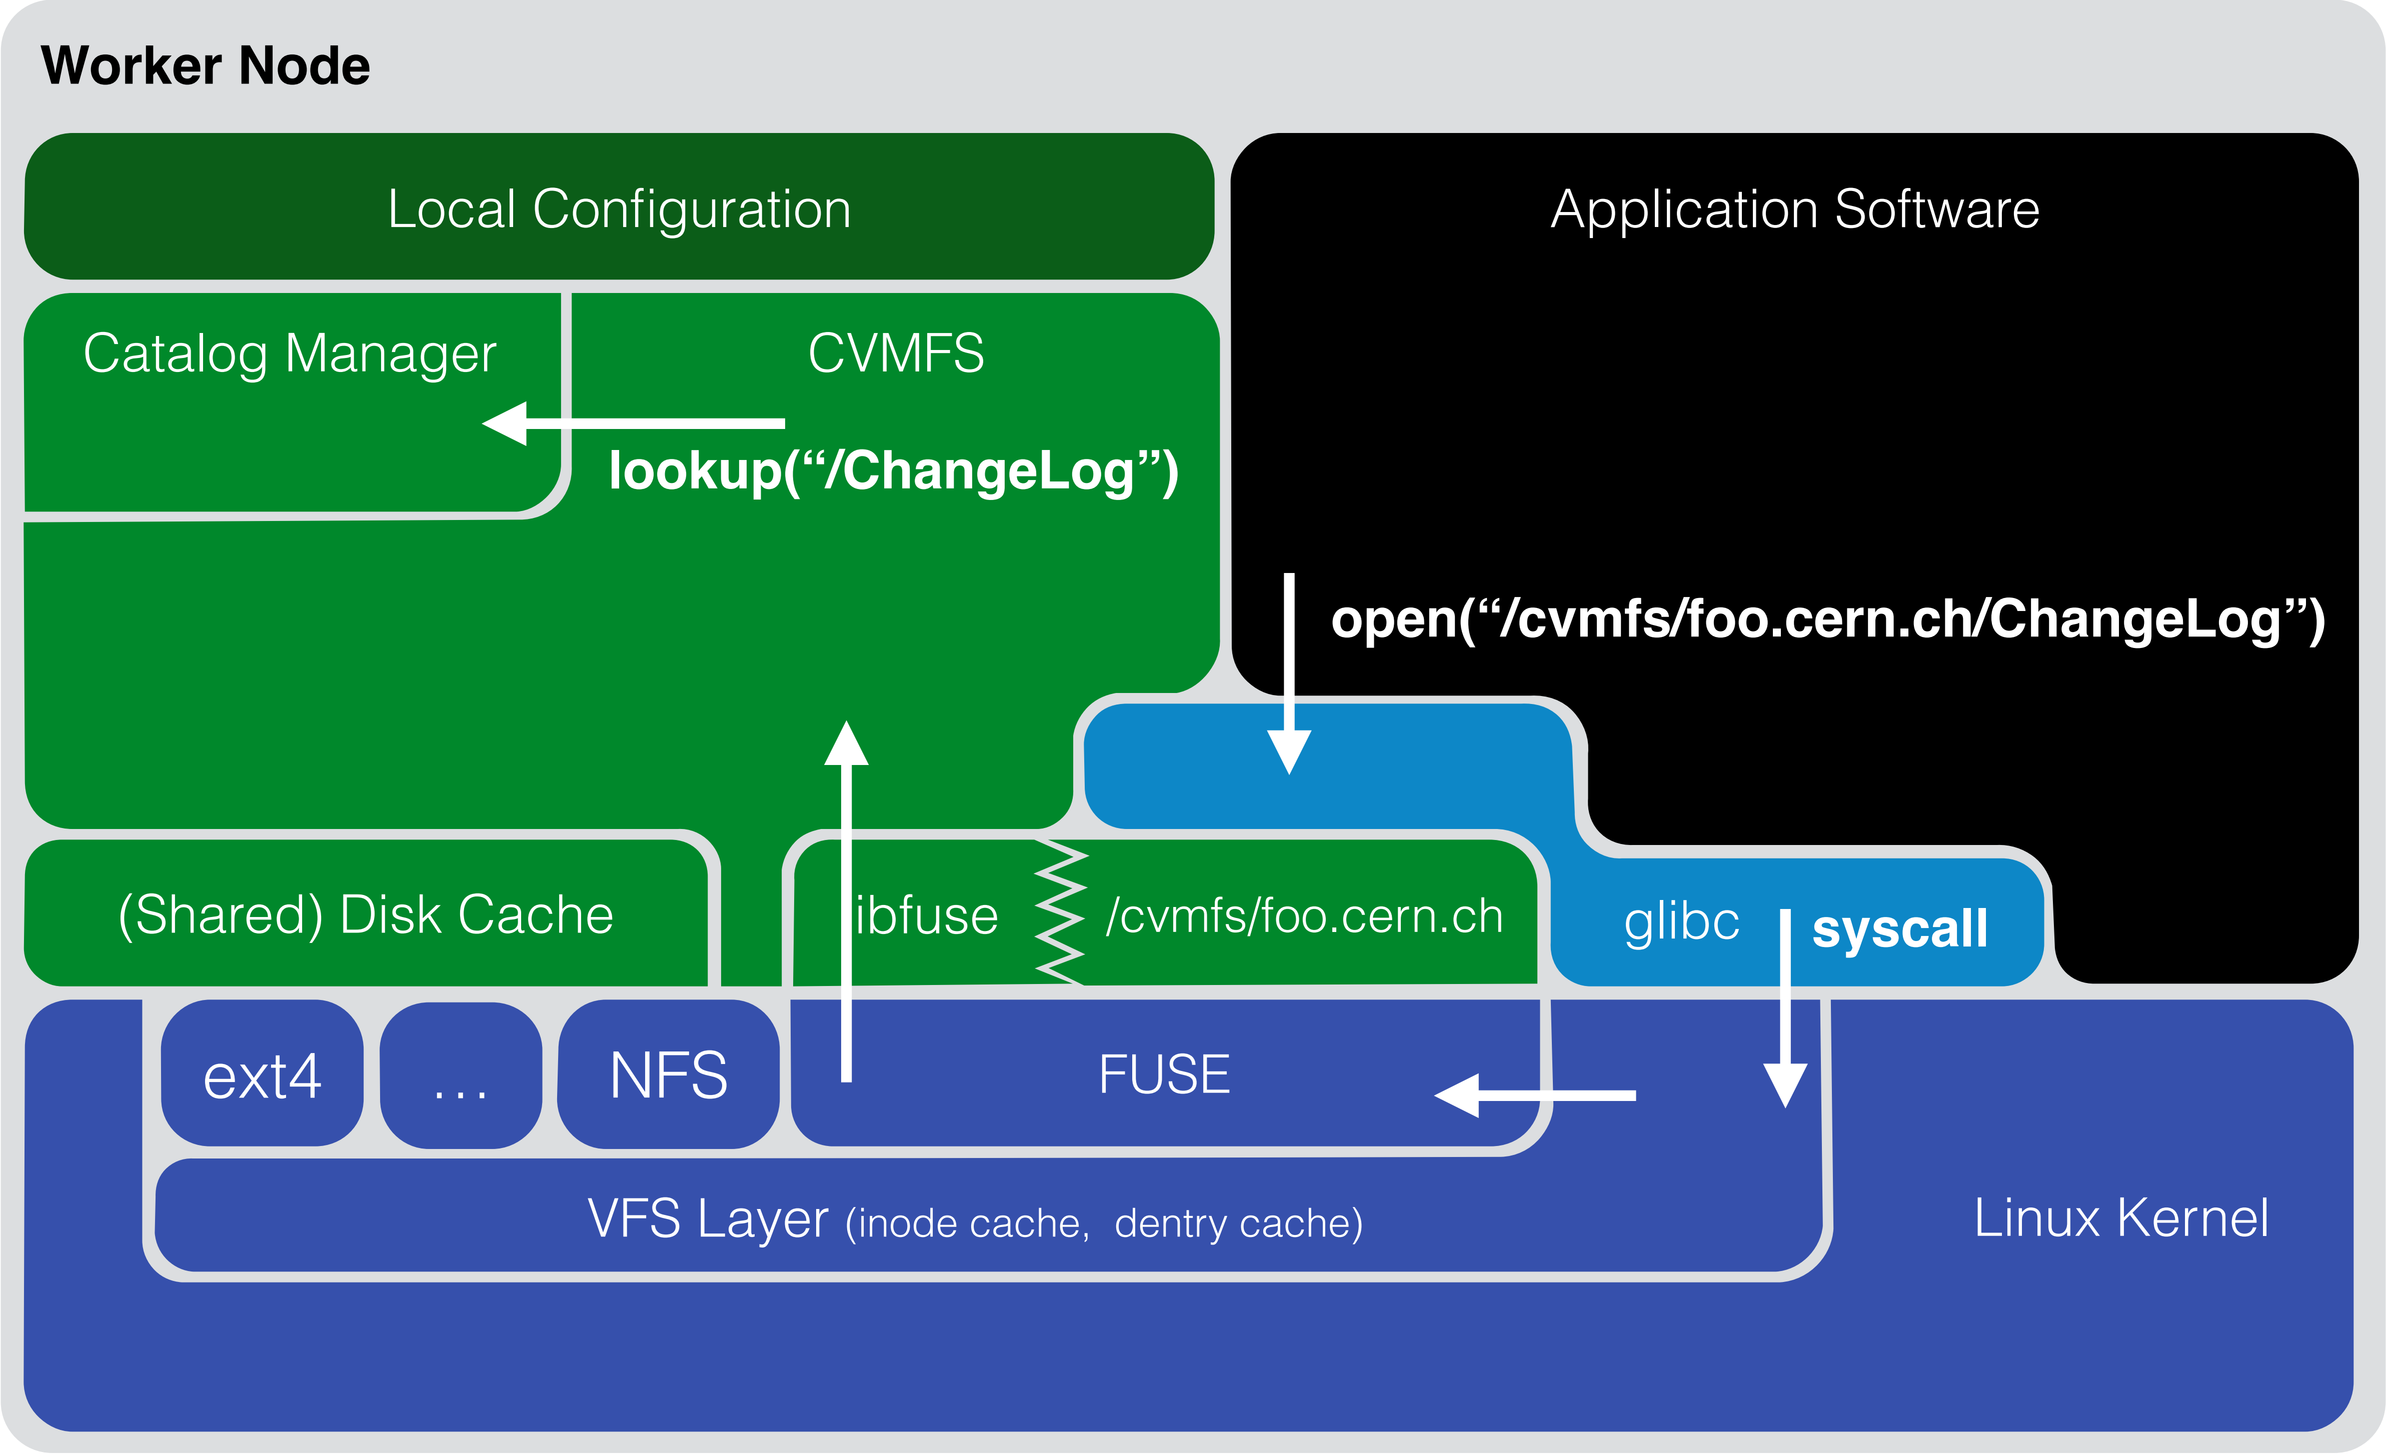

CernVM-FS Client - Control Flow

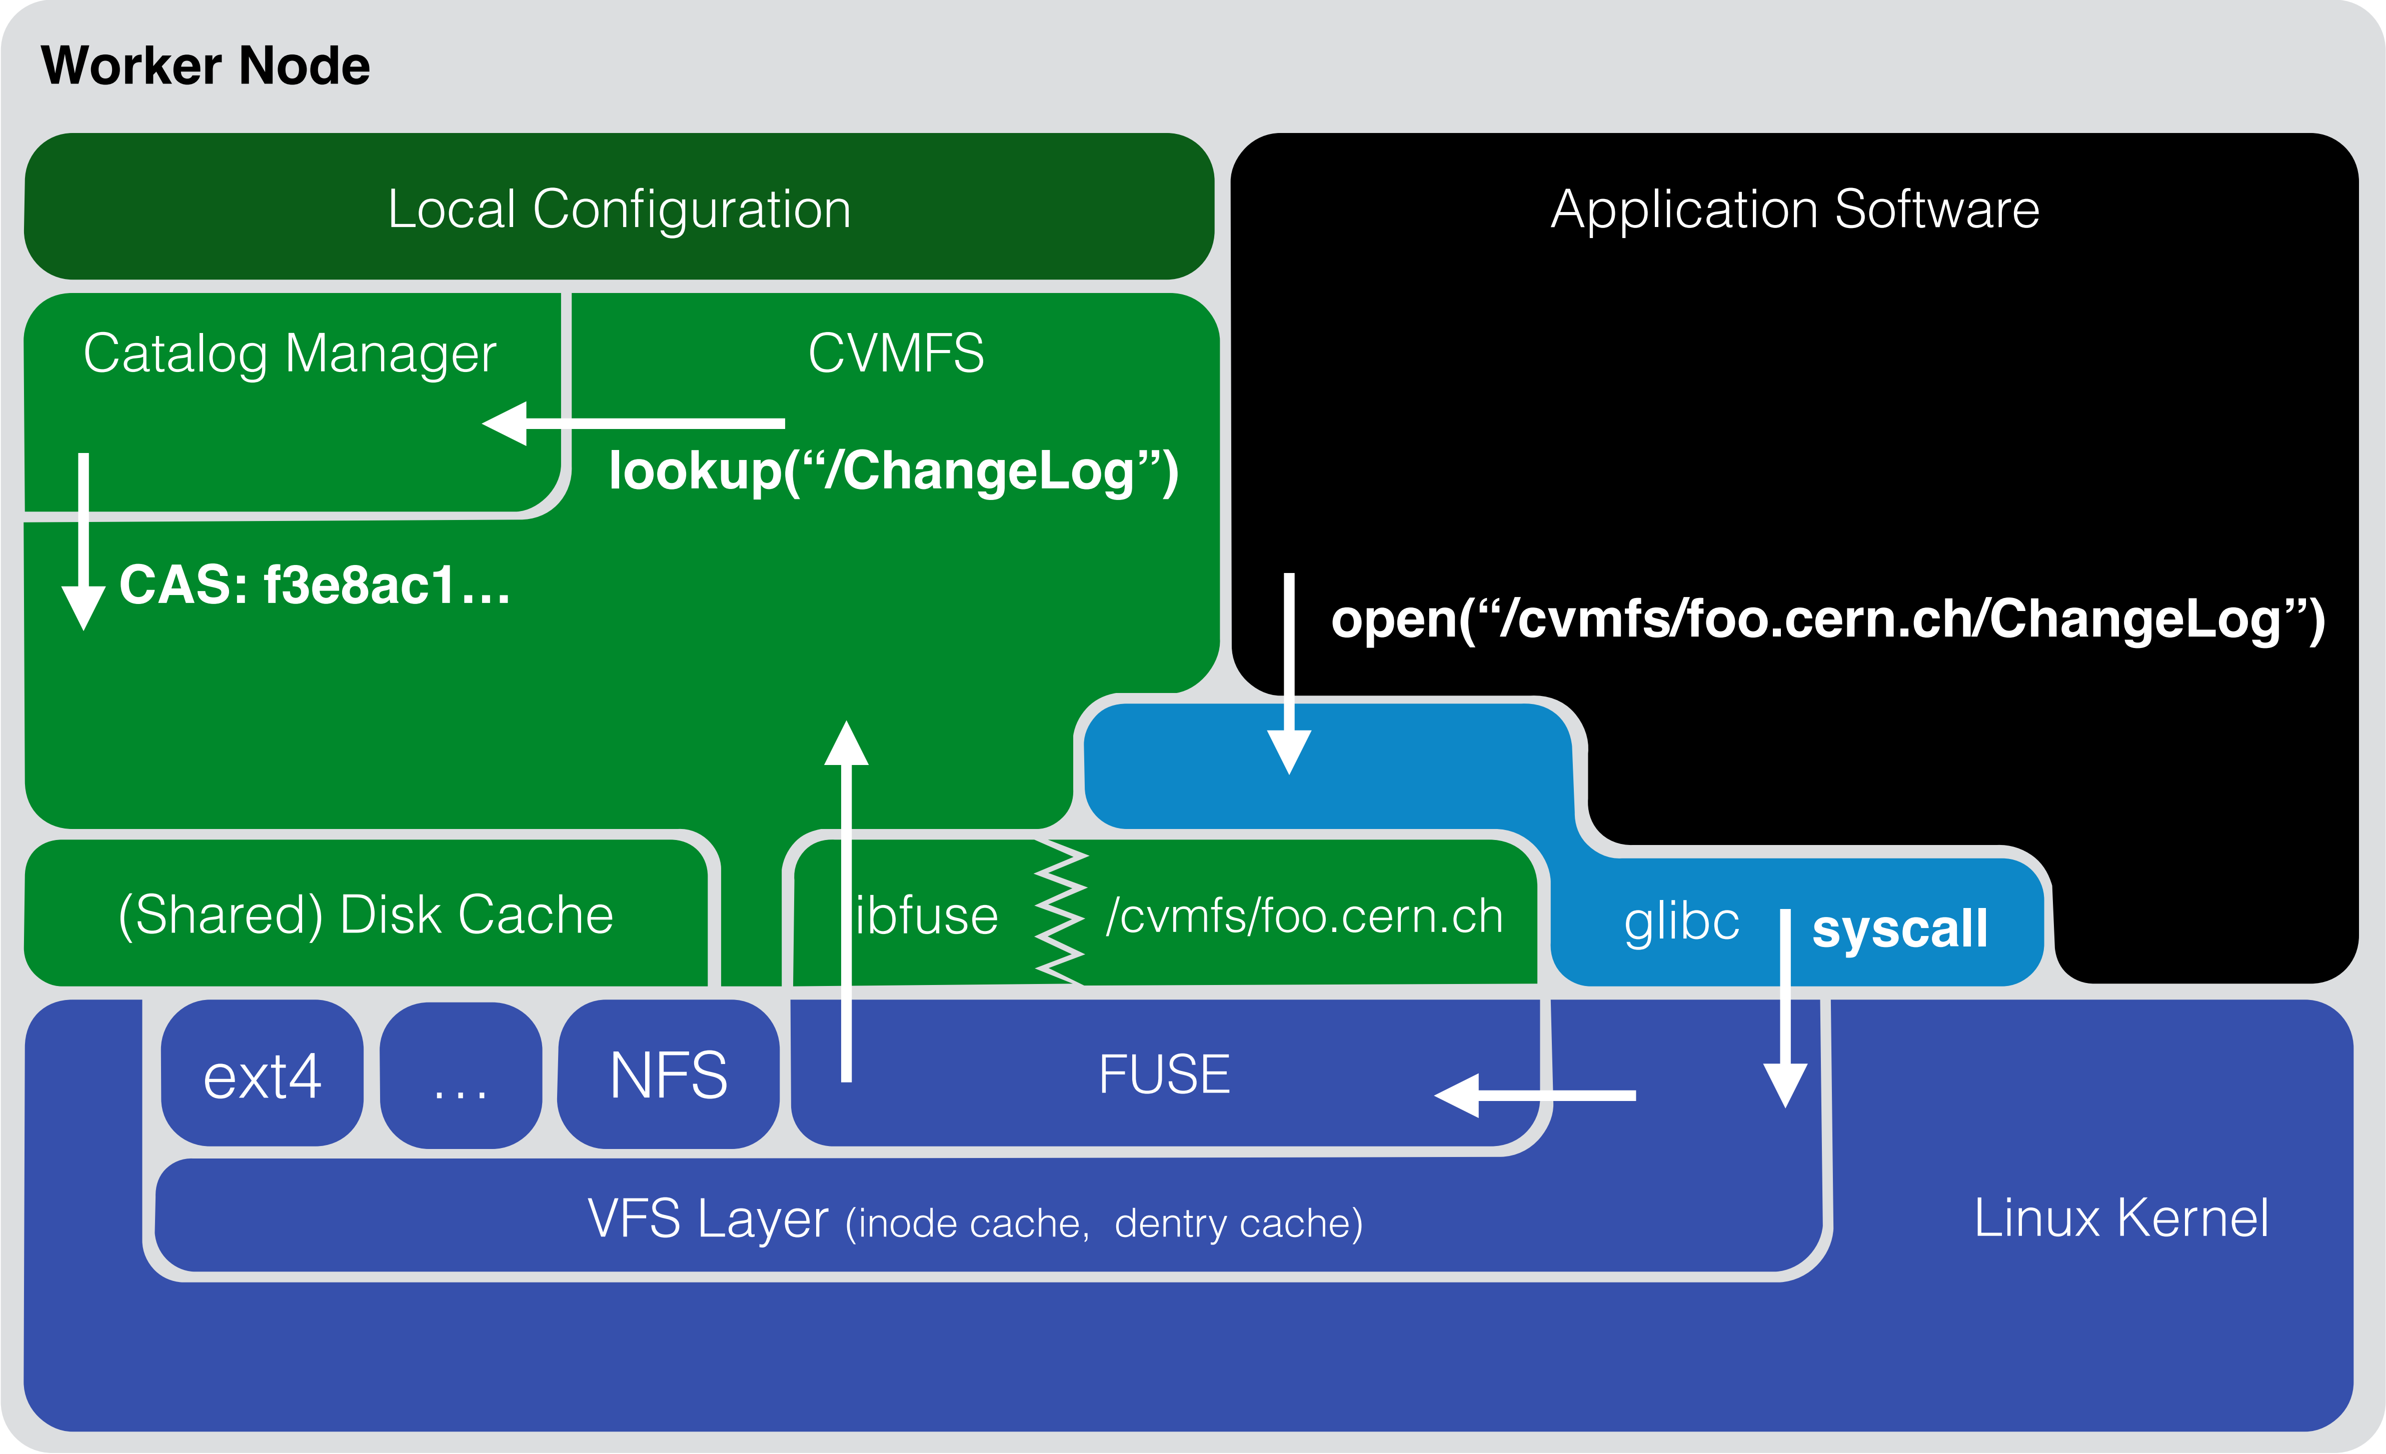

CernVM-FS Client - Control Flow

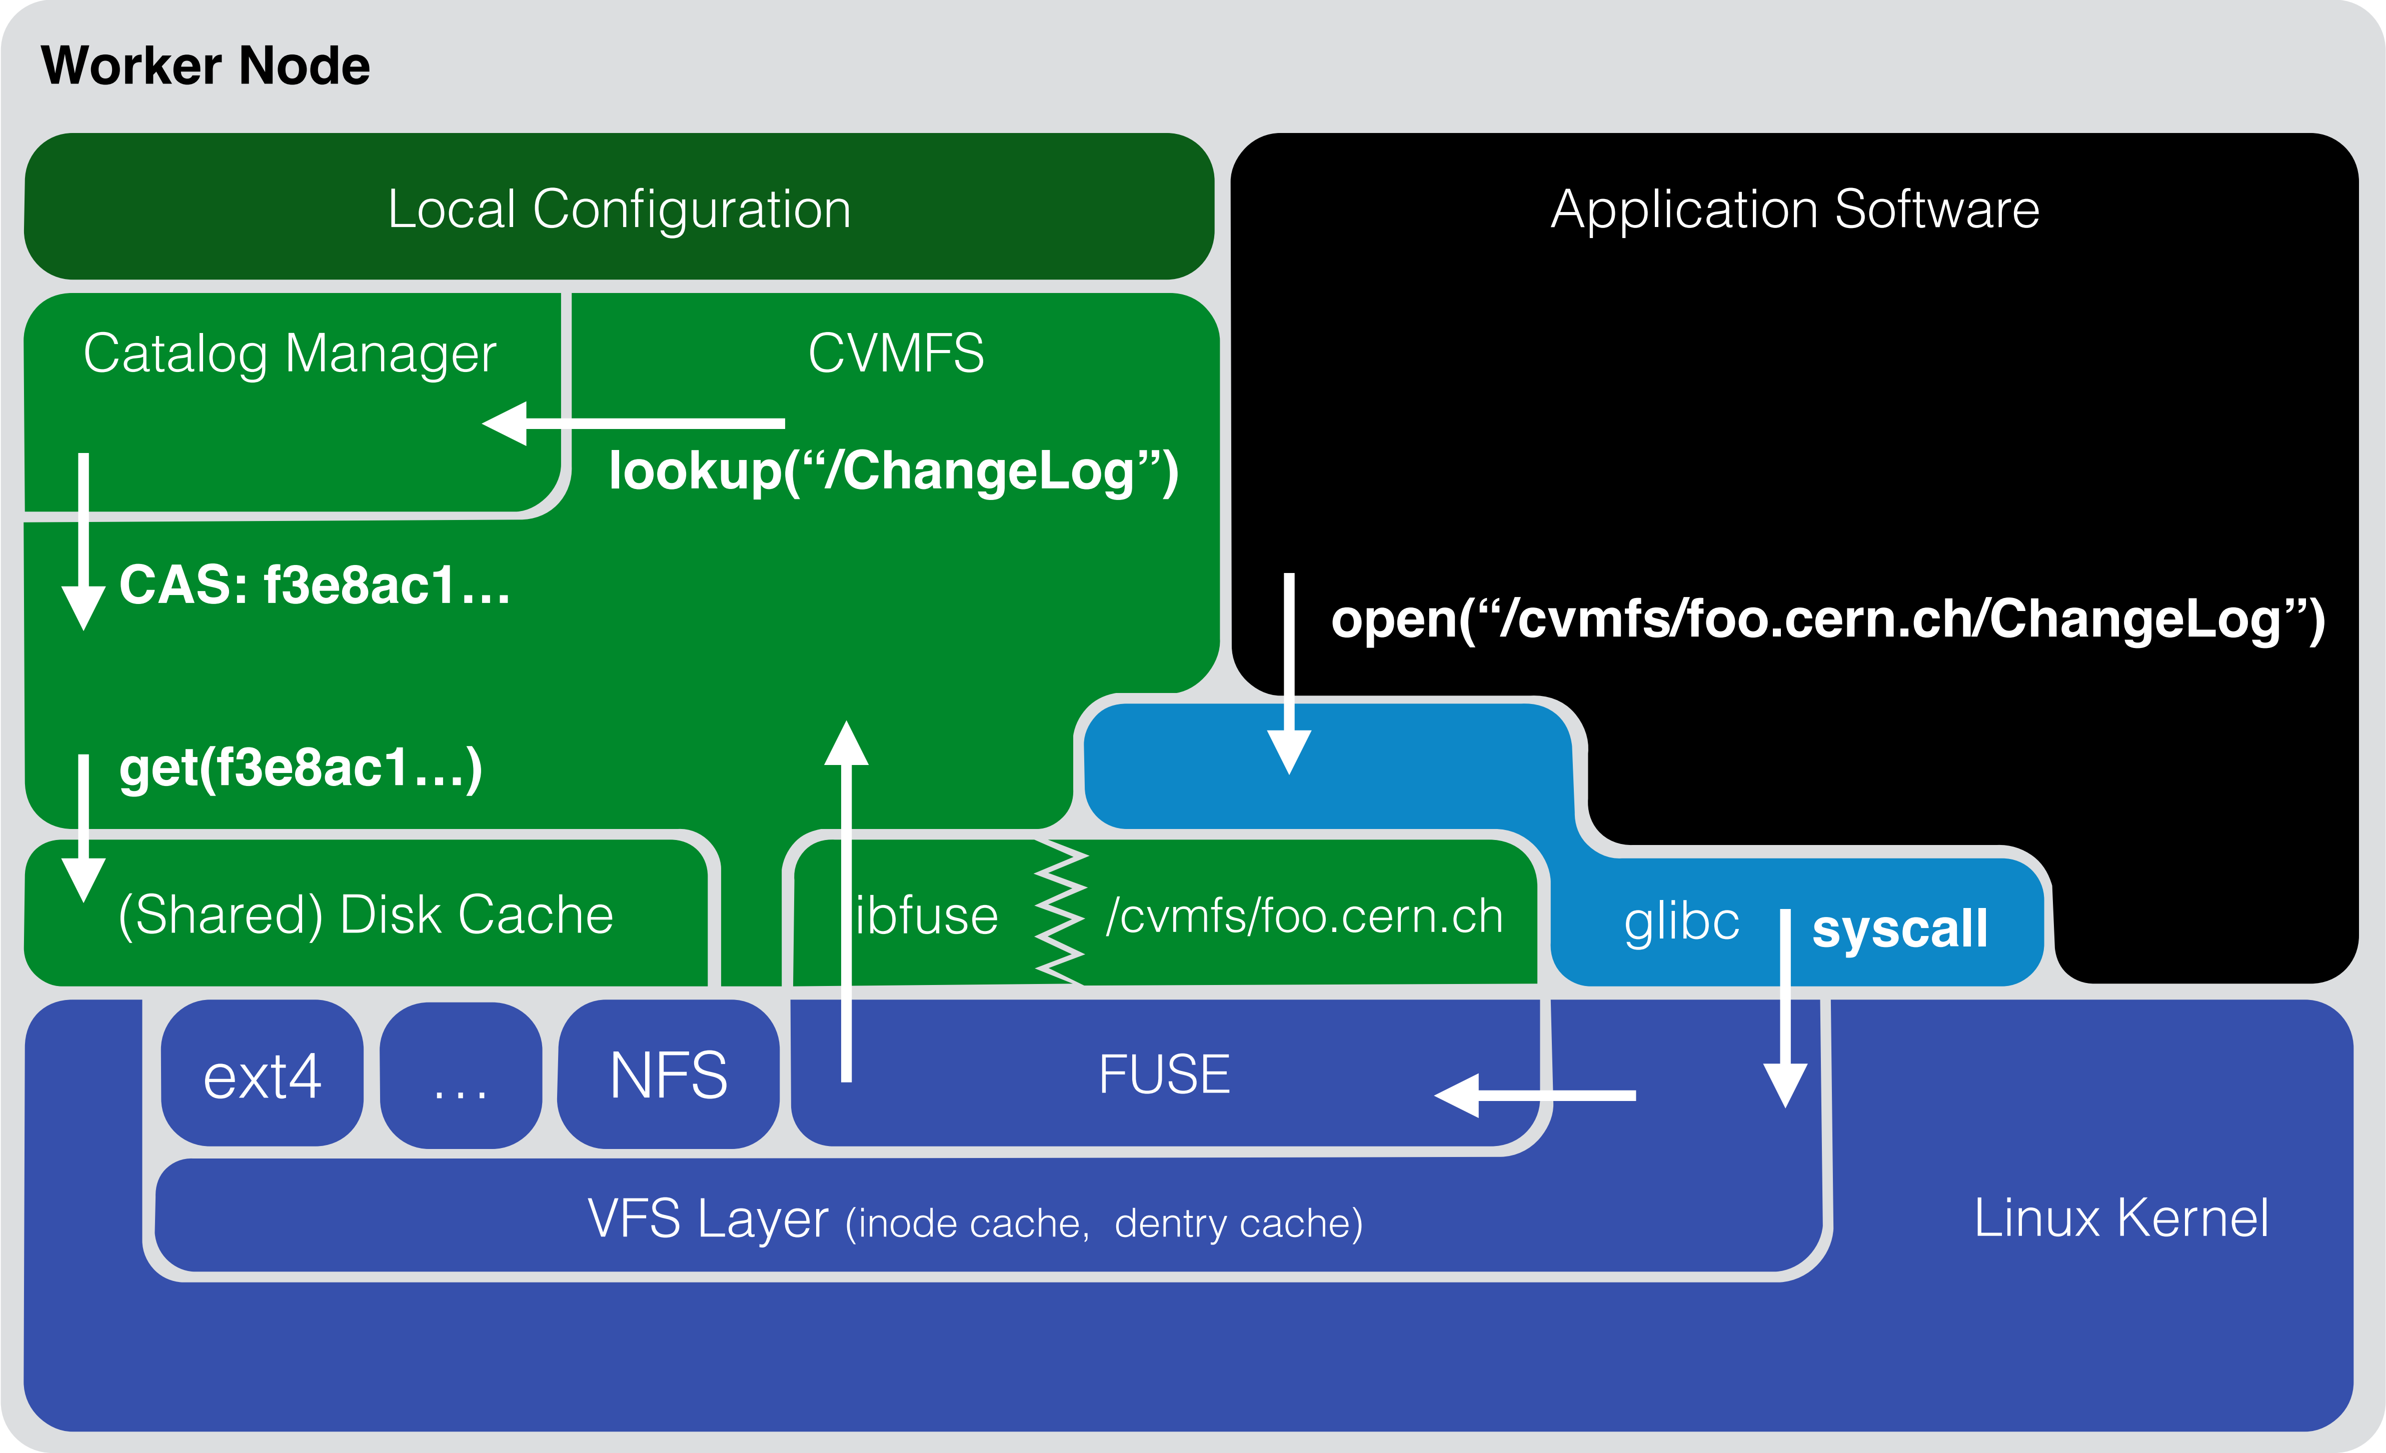

CernVM-FS Client - Control Flow

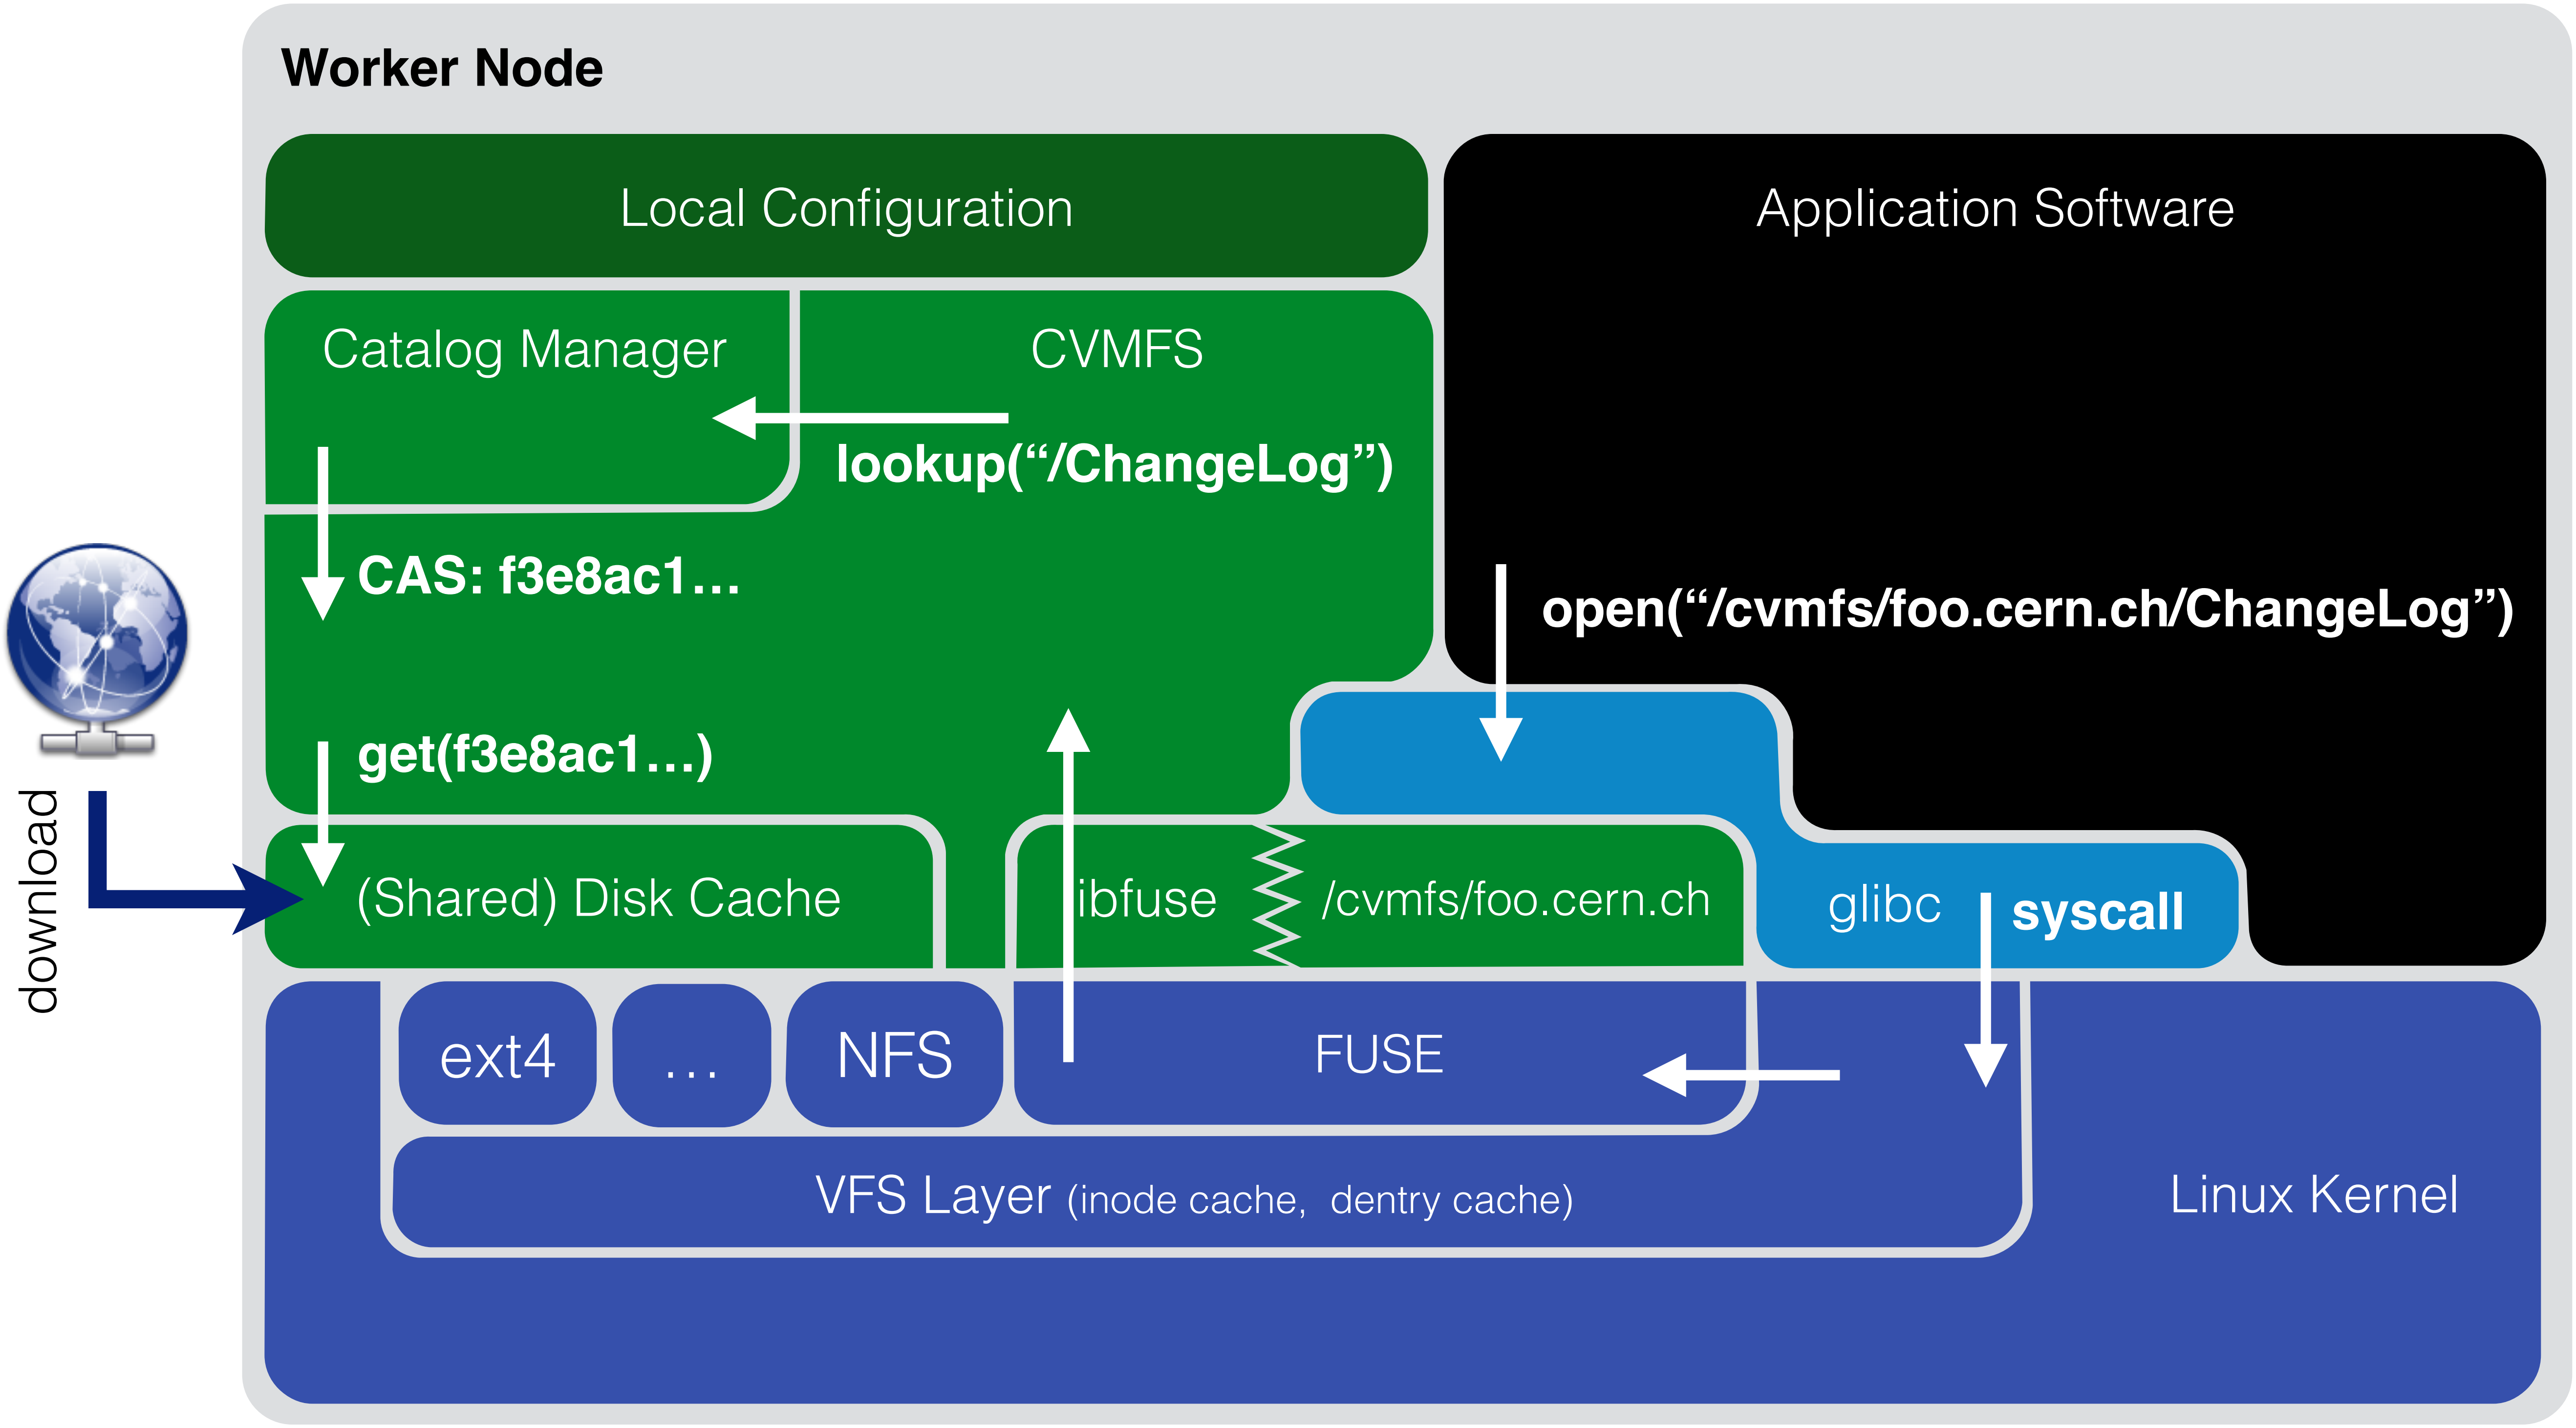

CernVM-FS Client - Control Flow

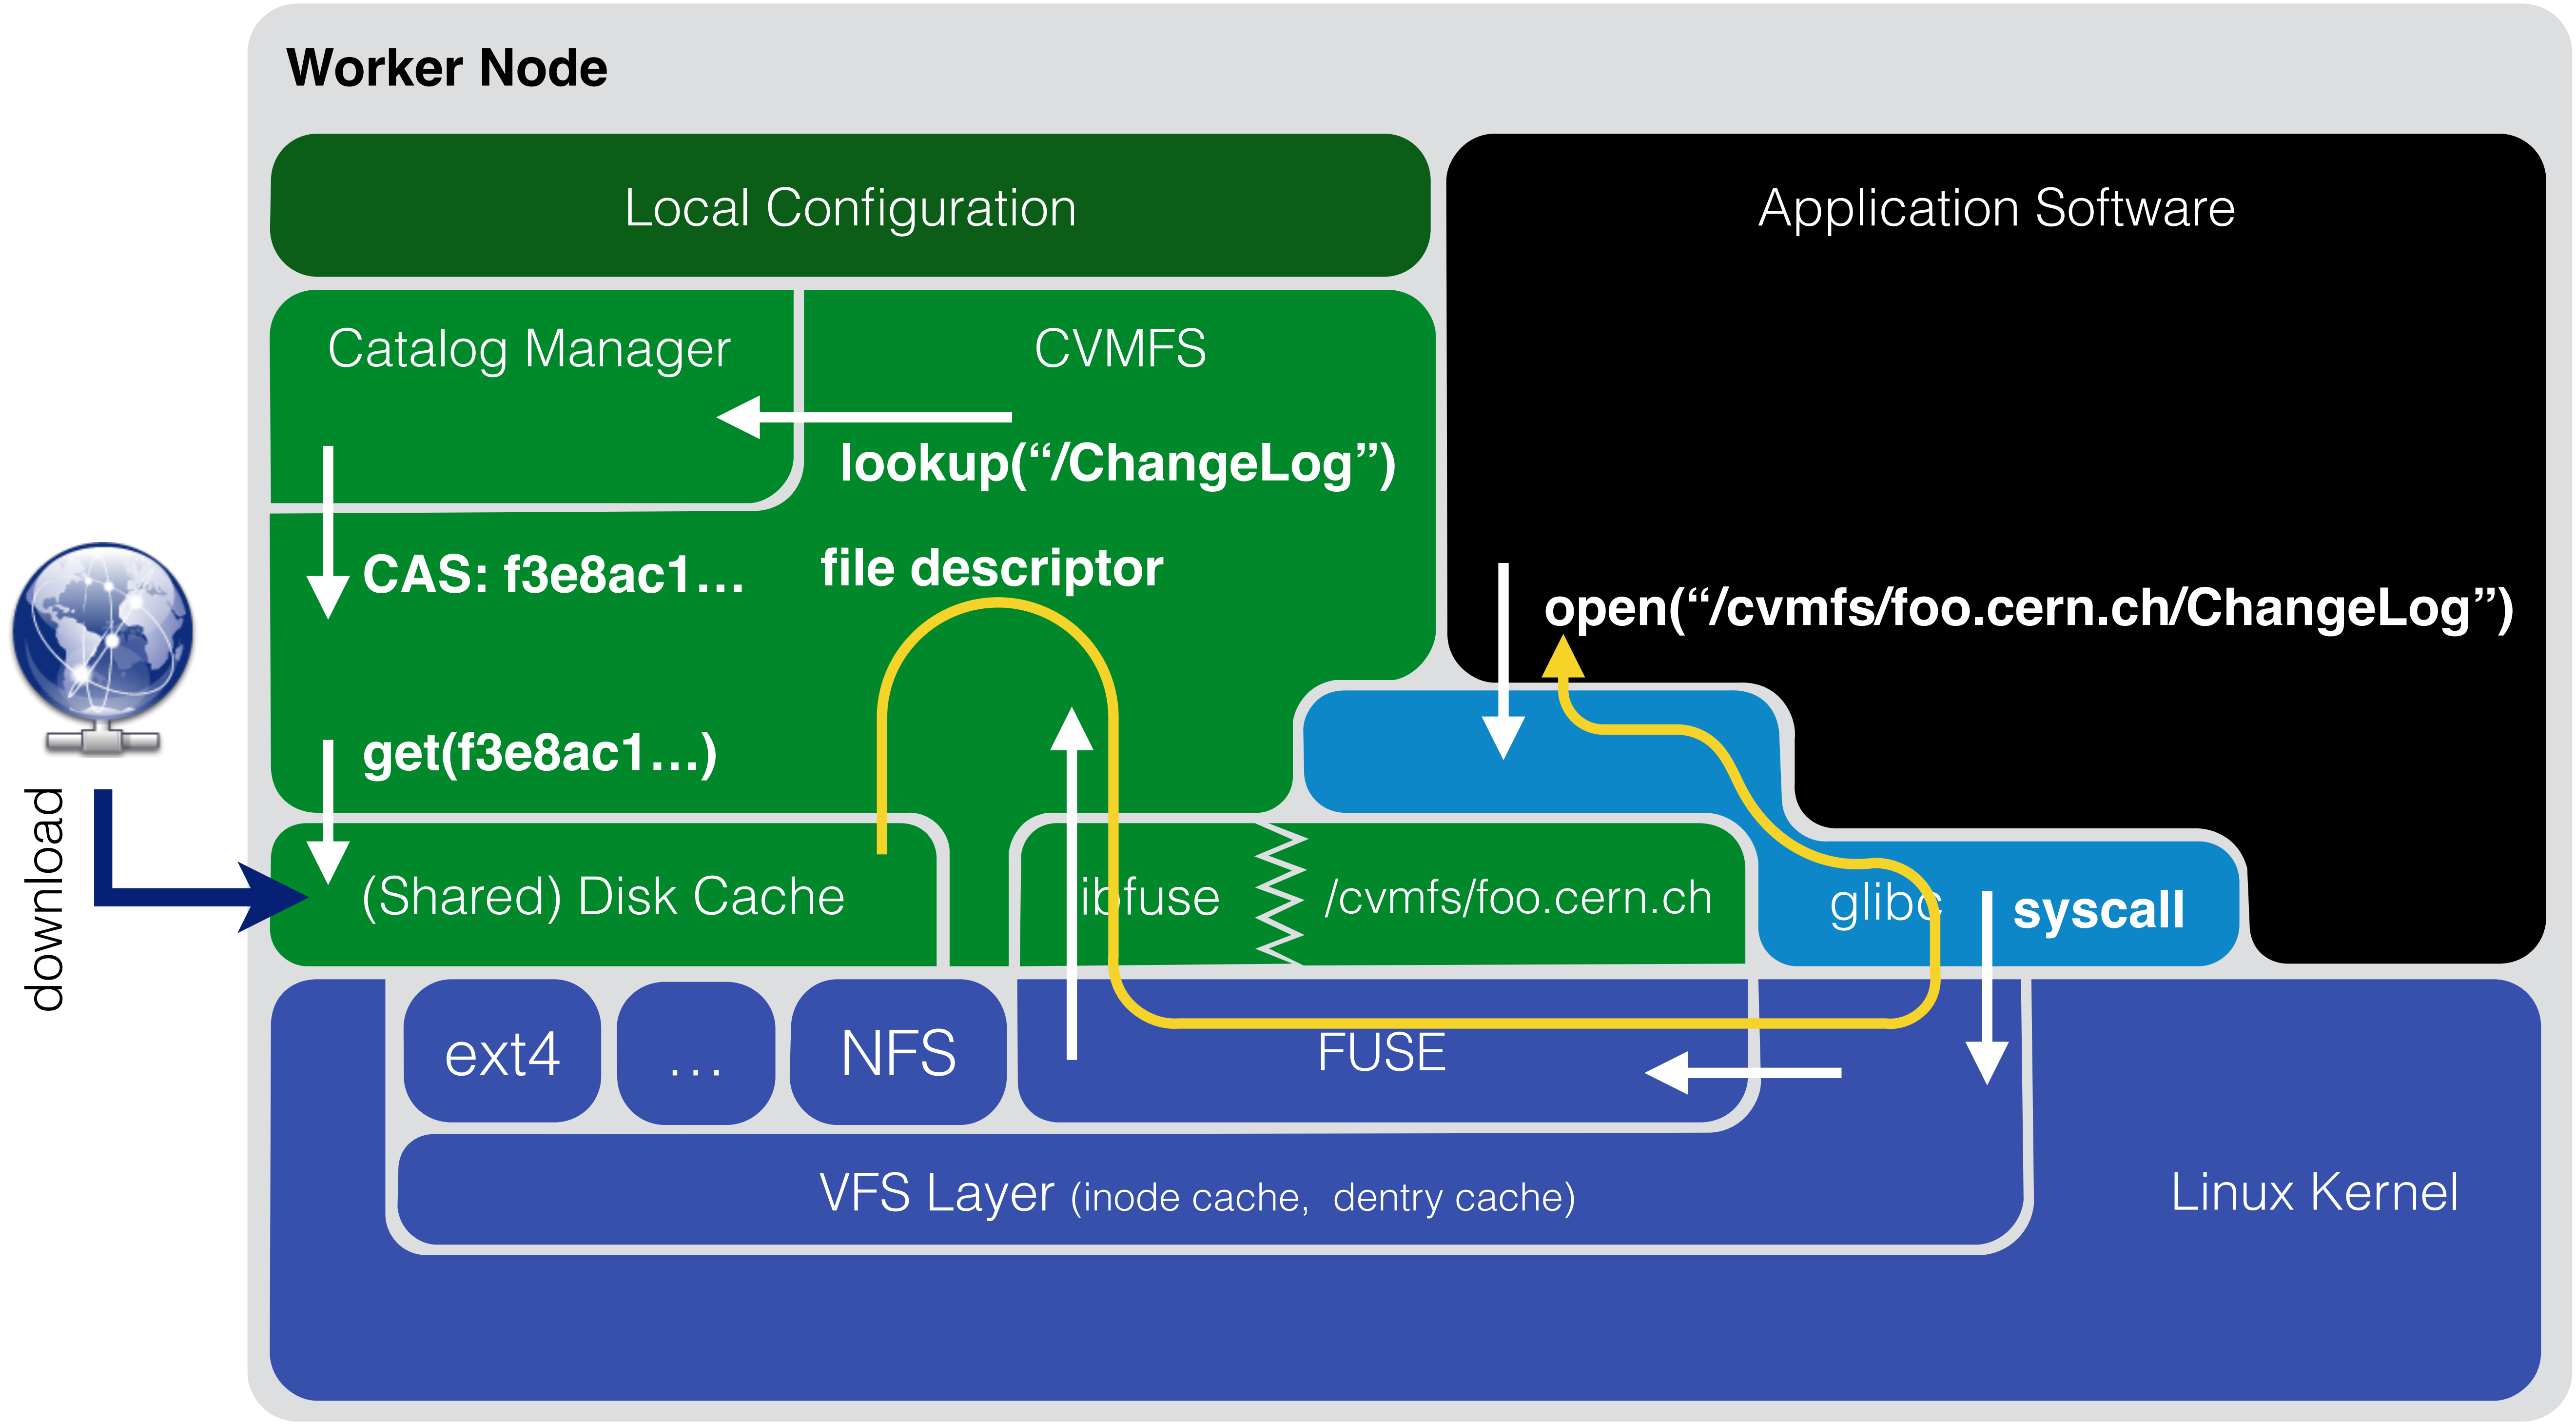

CernVM-FS Client - Hotpatching/Reloading

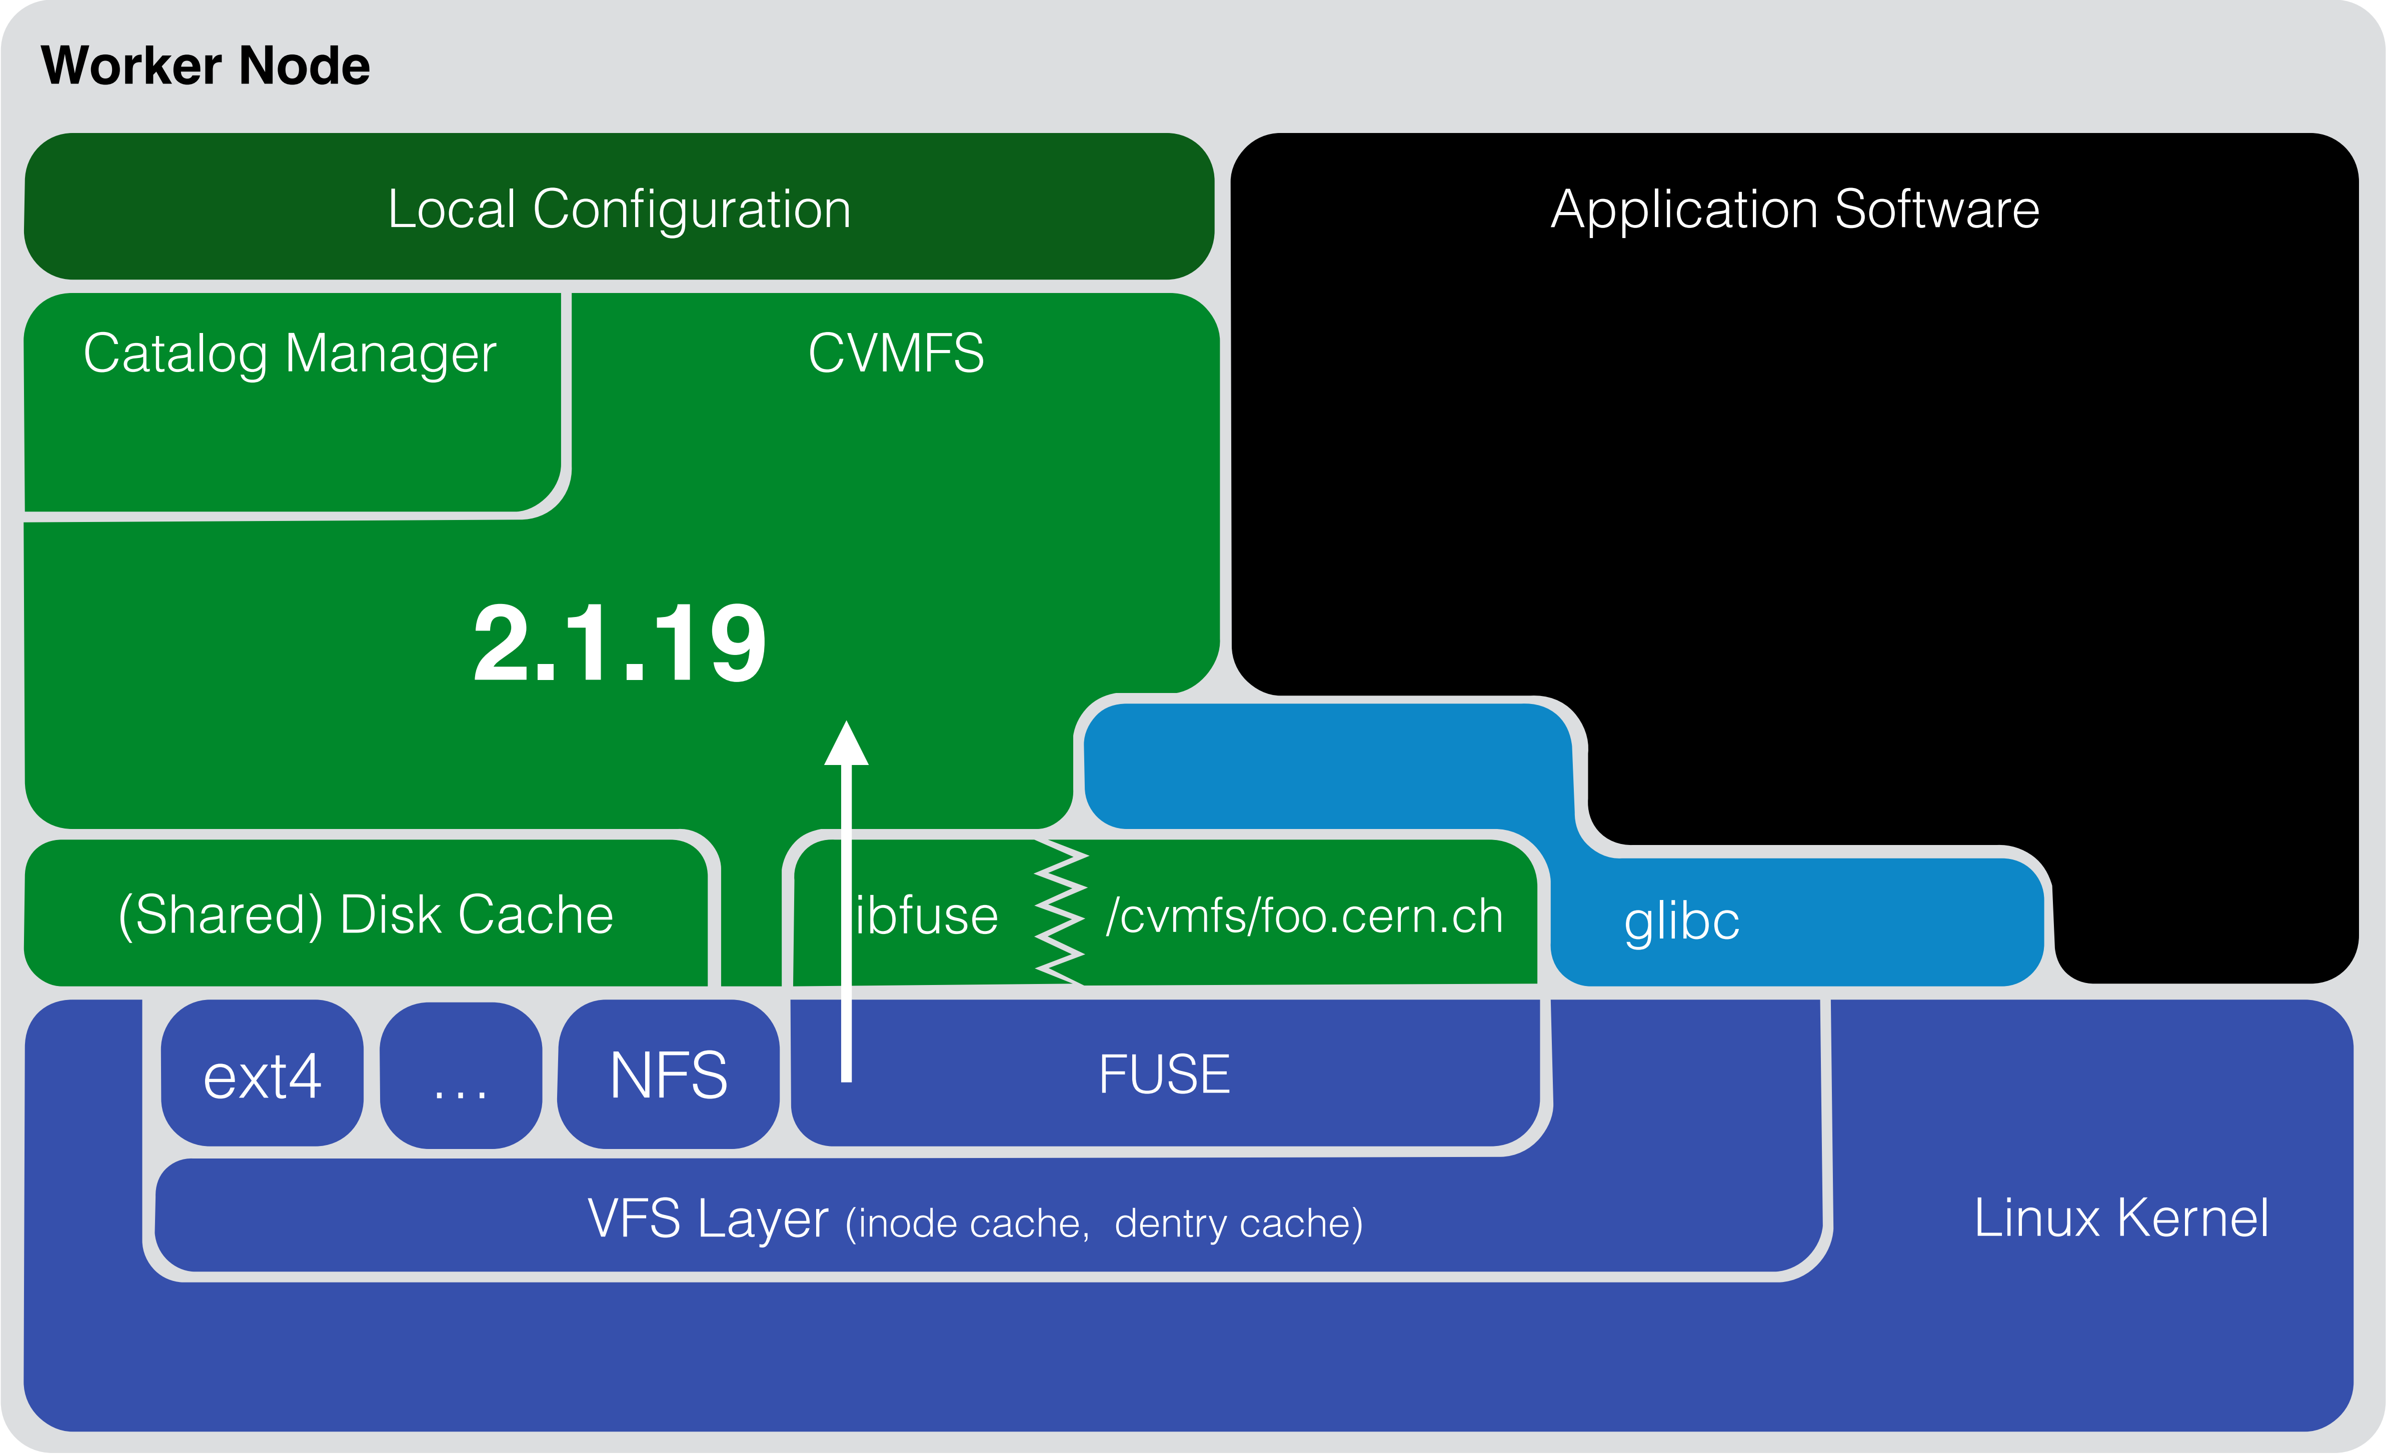

CernVM-FS Client - Hotpatching/Reloading

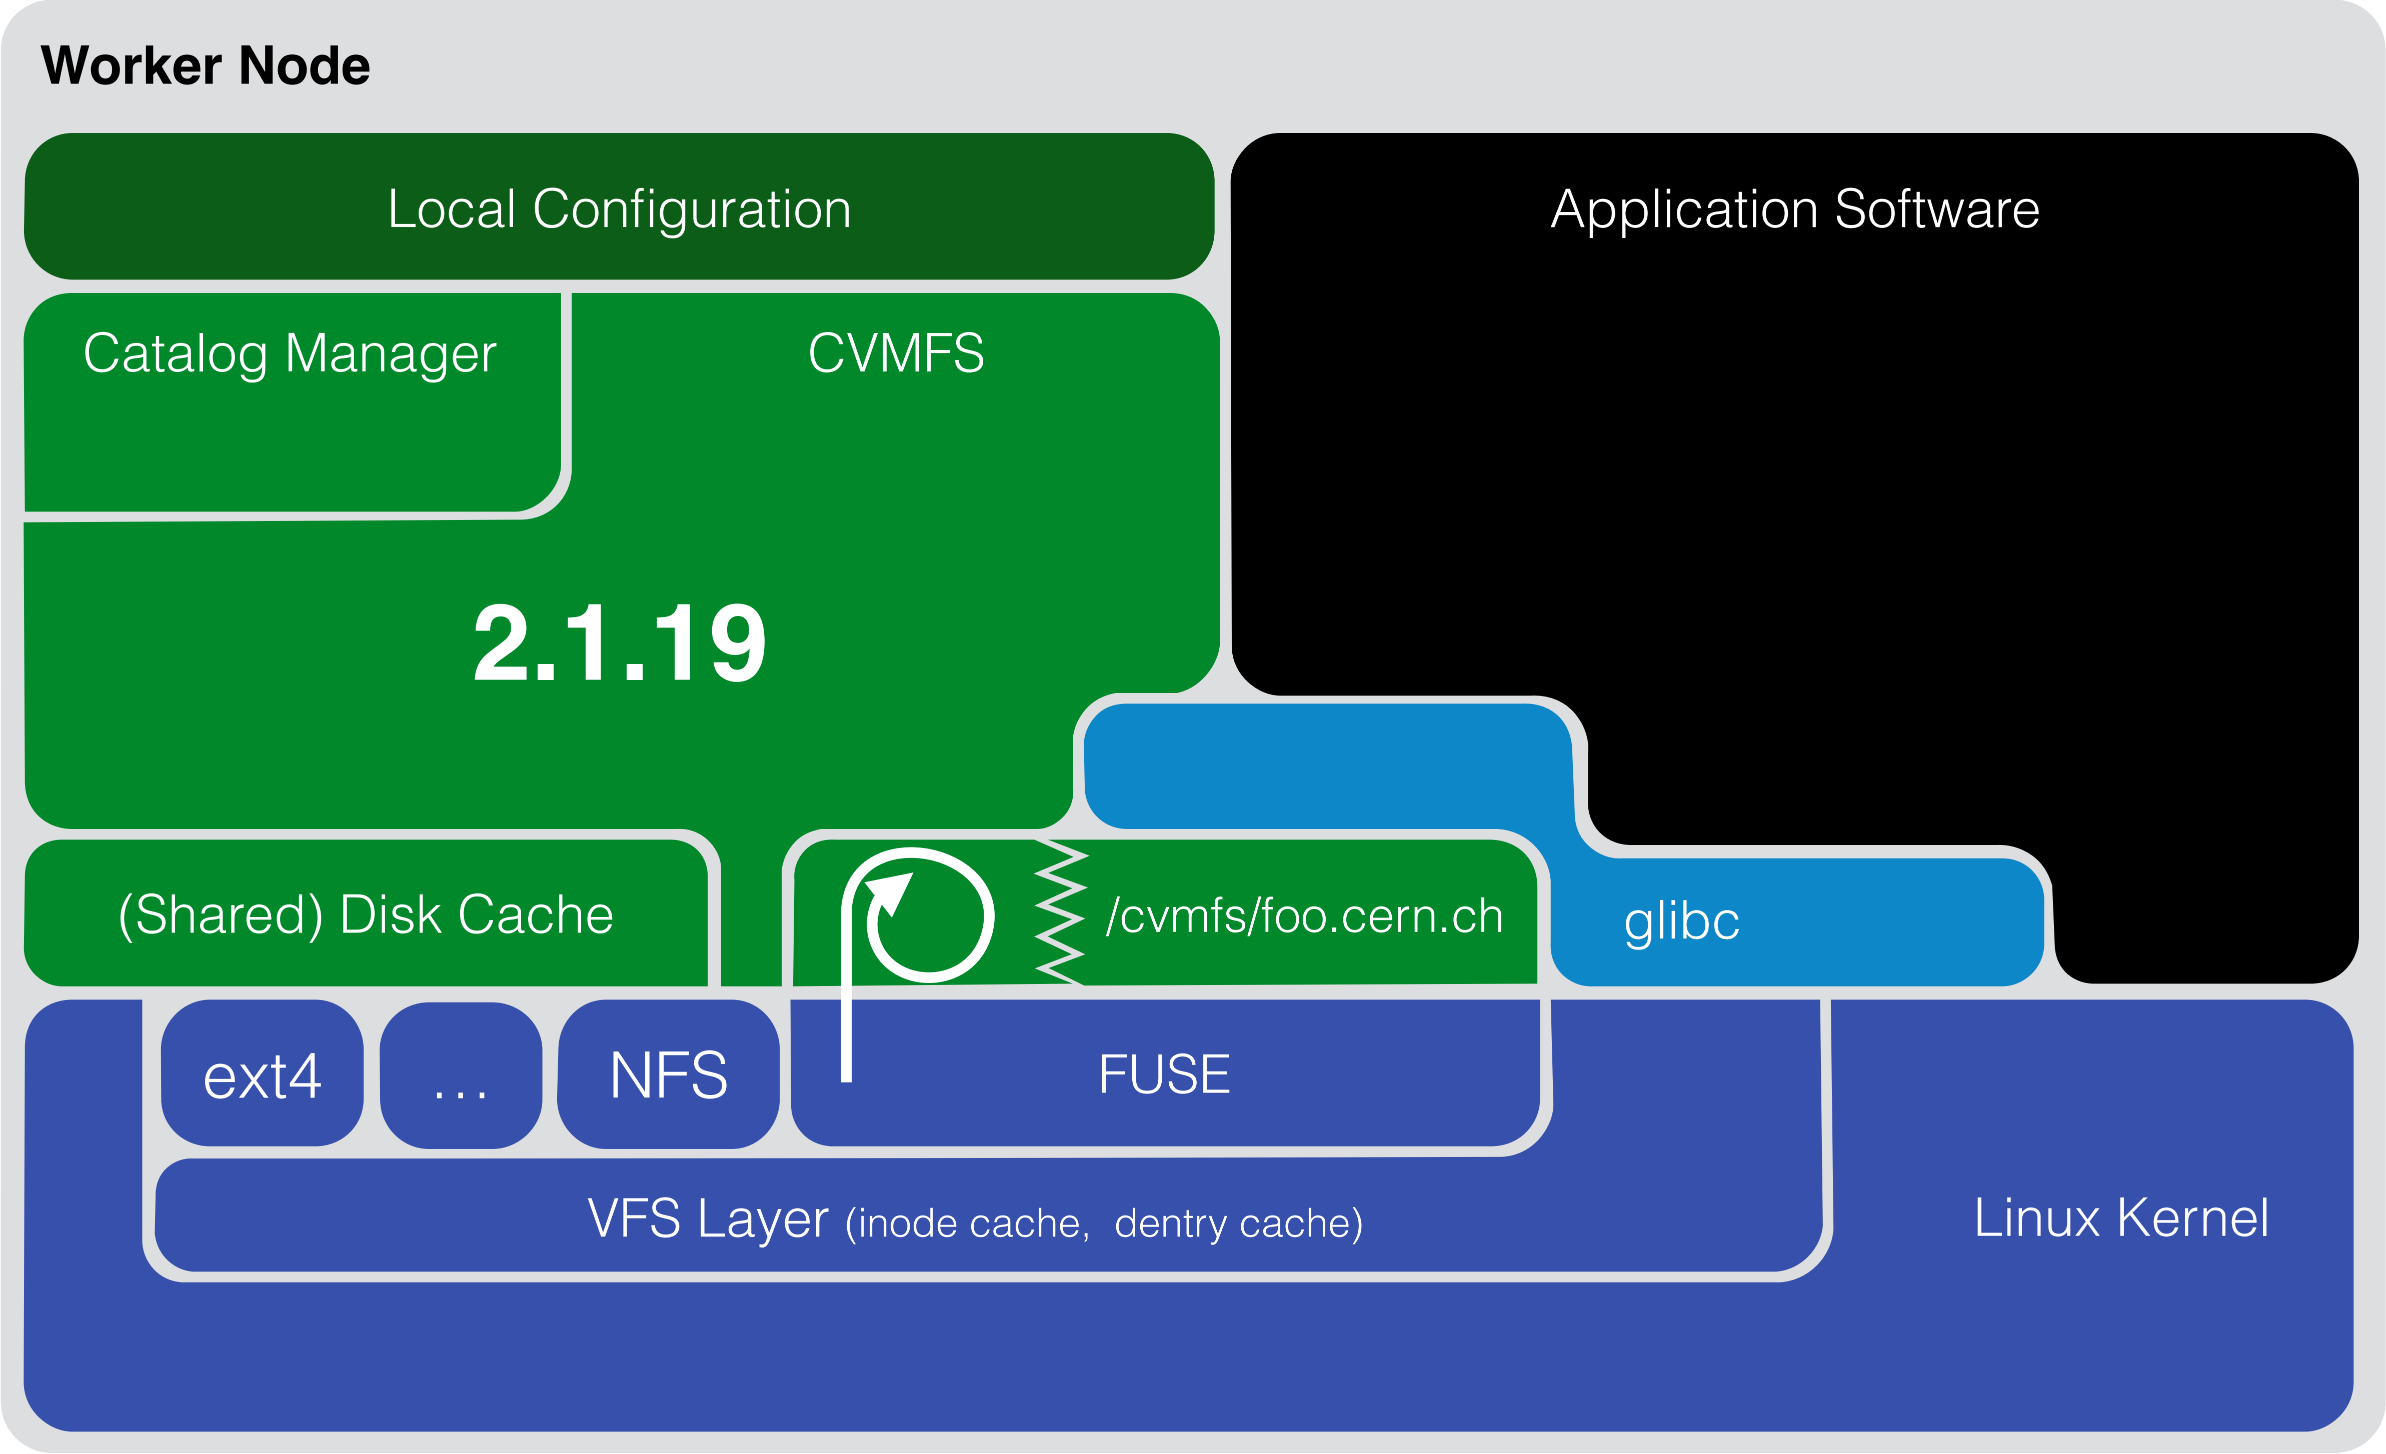

CernVM-FS Client - Hotpatching/Reloading

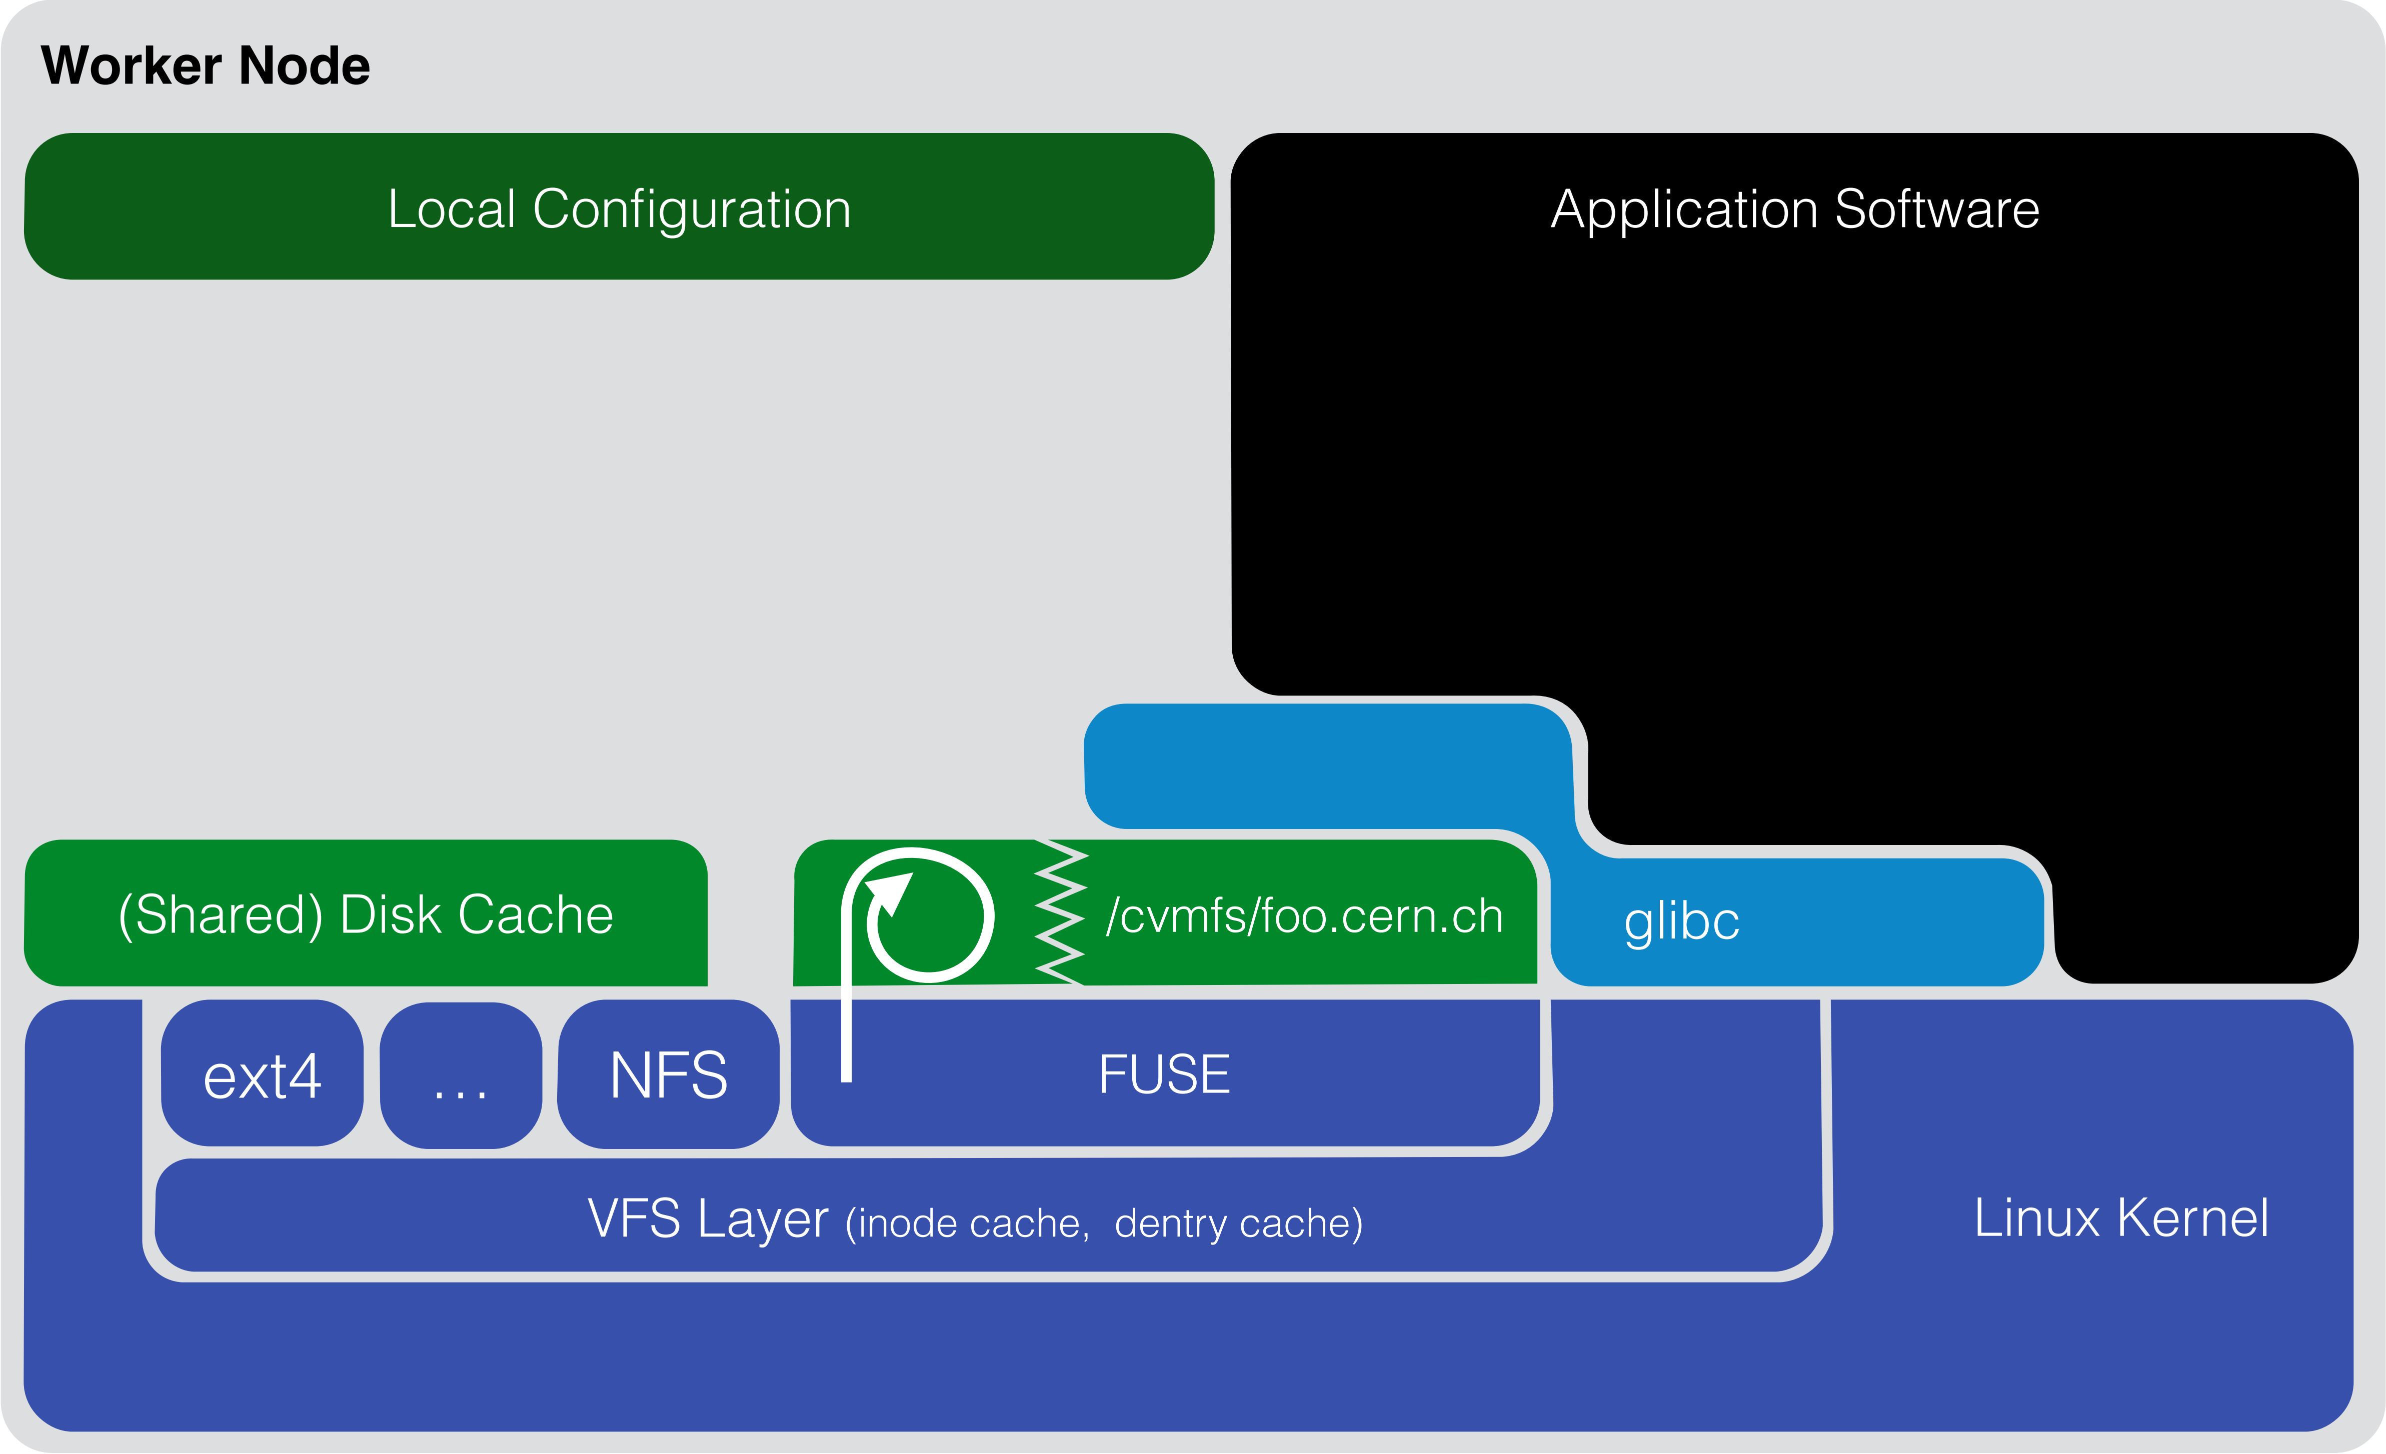

CernVM-FS Client - Hotpatching/Reloading

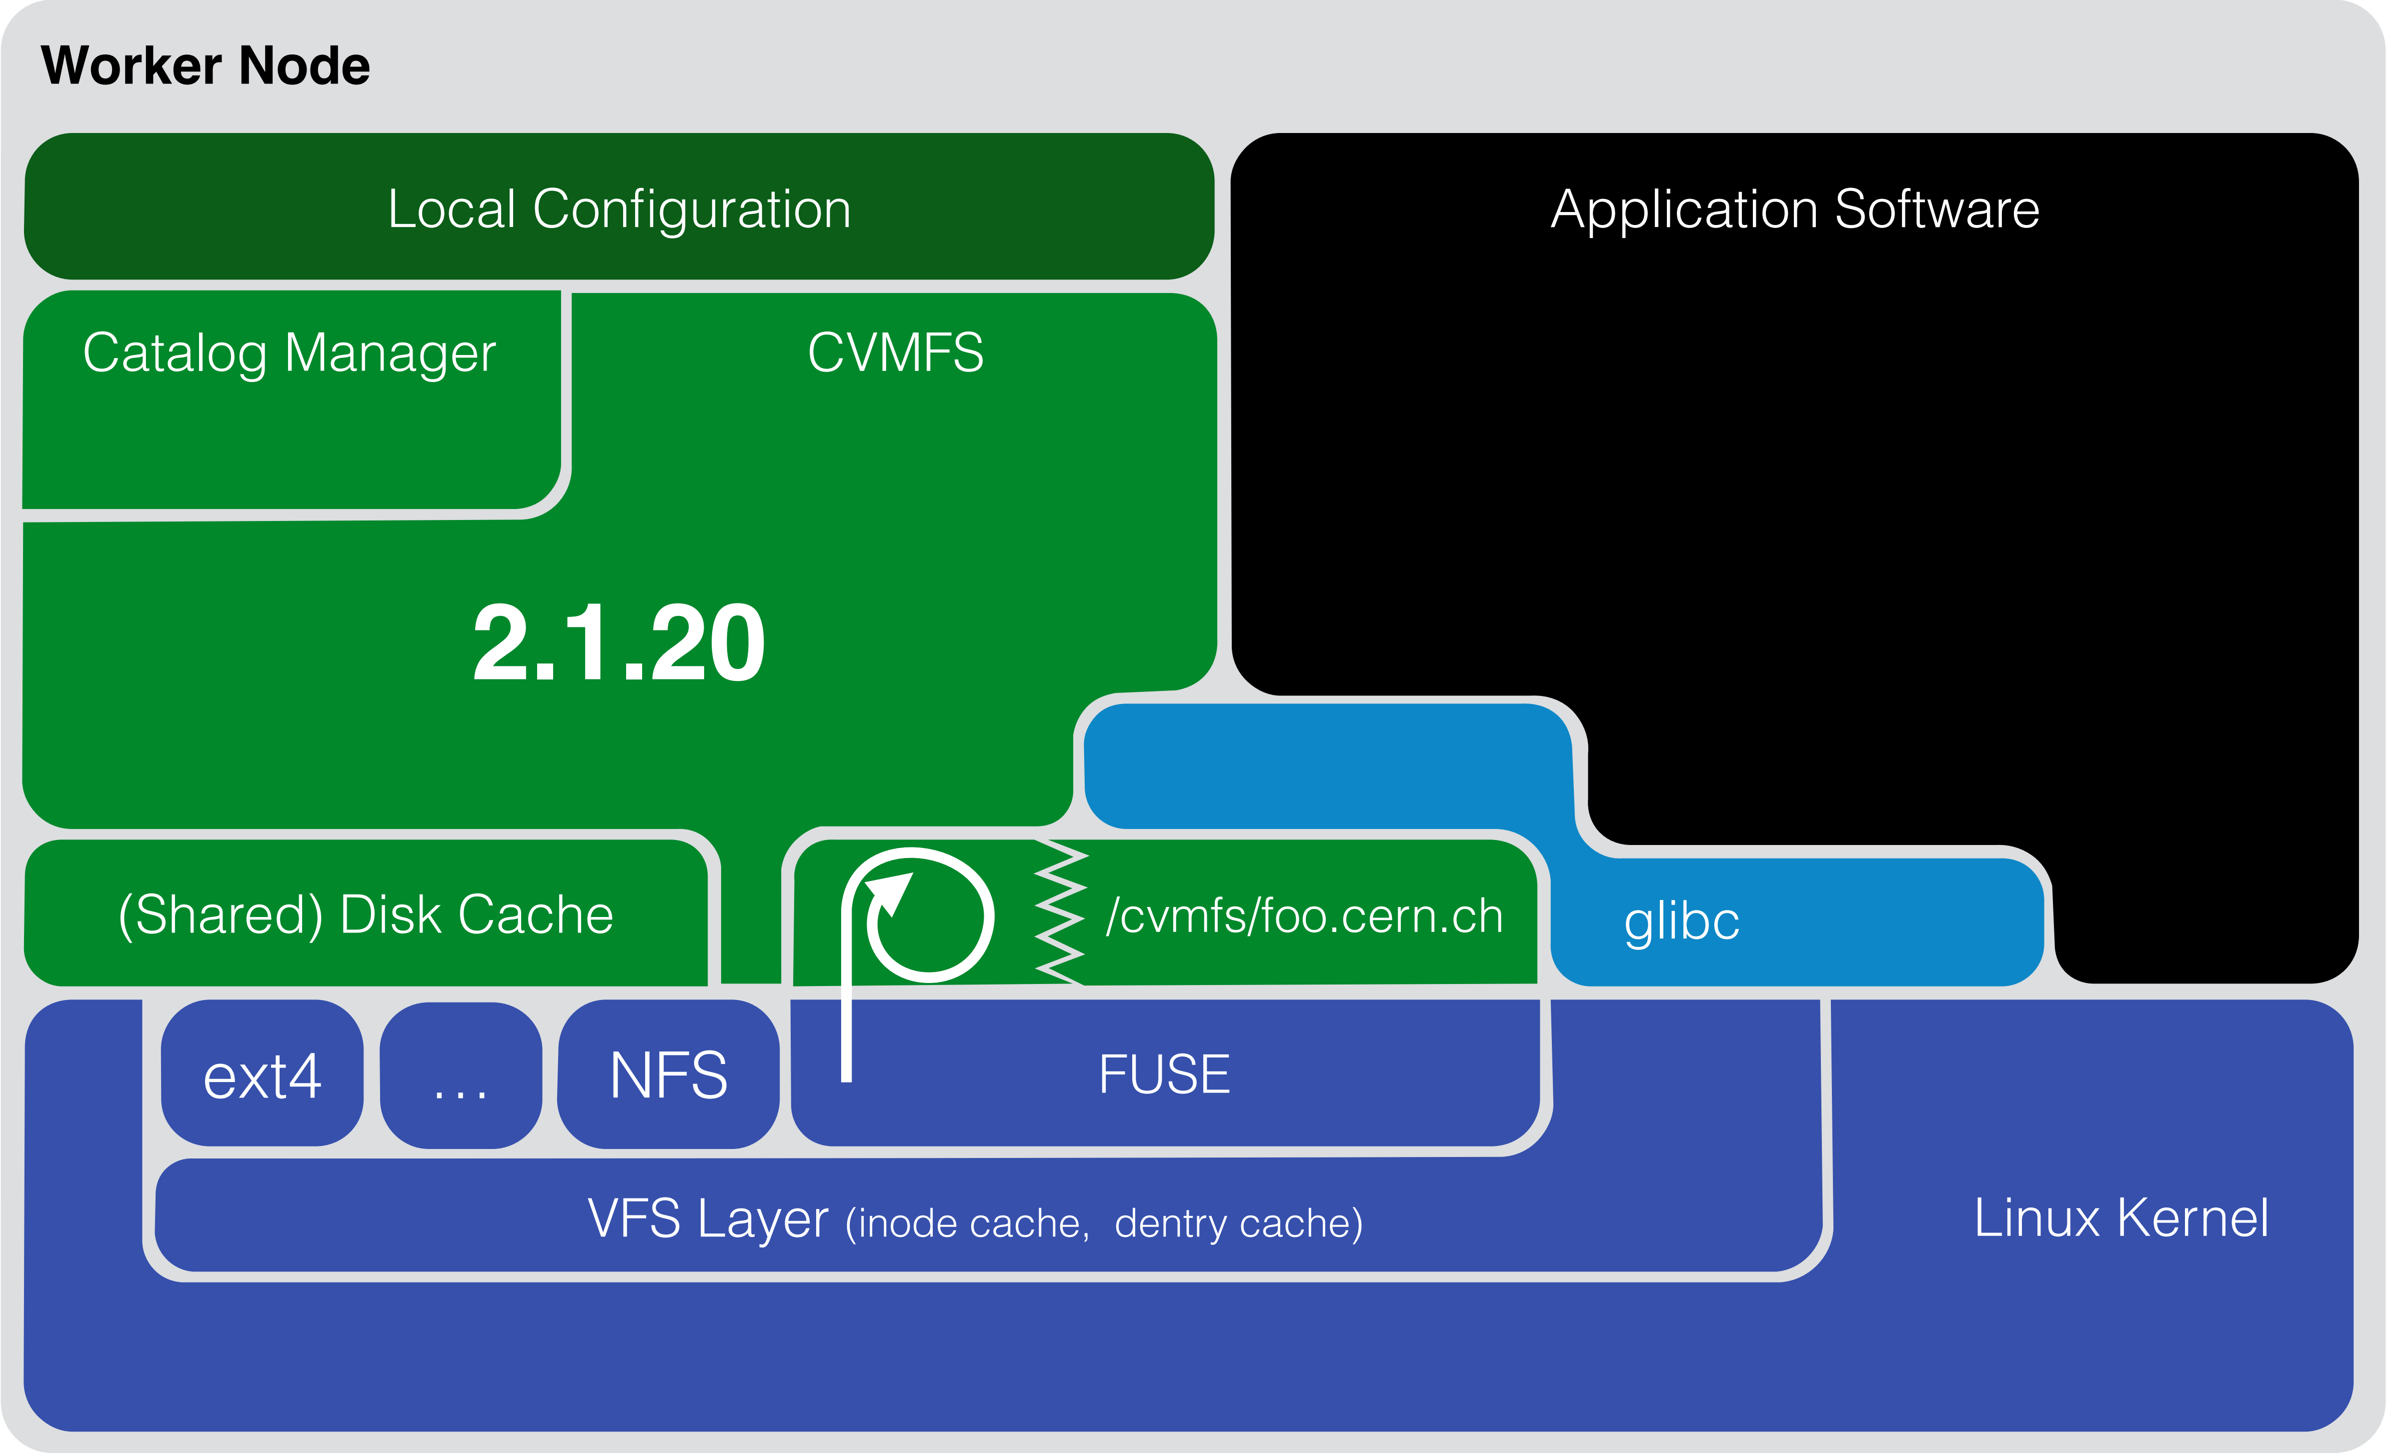

CernVM-FS Client - Hotpatching/Reloading

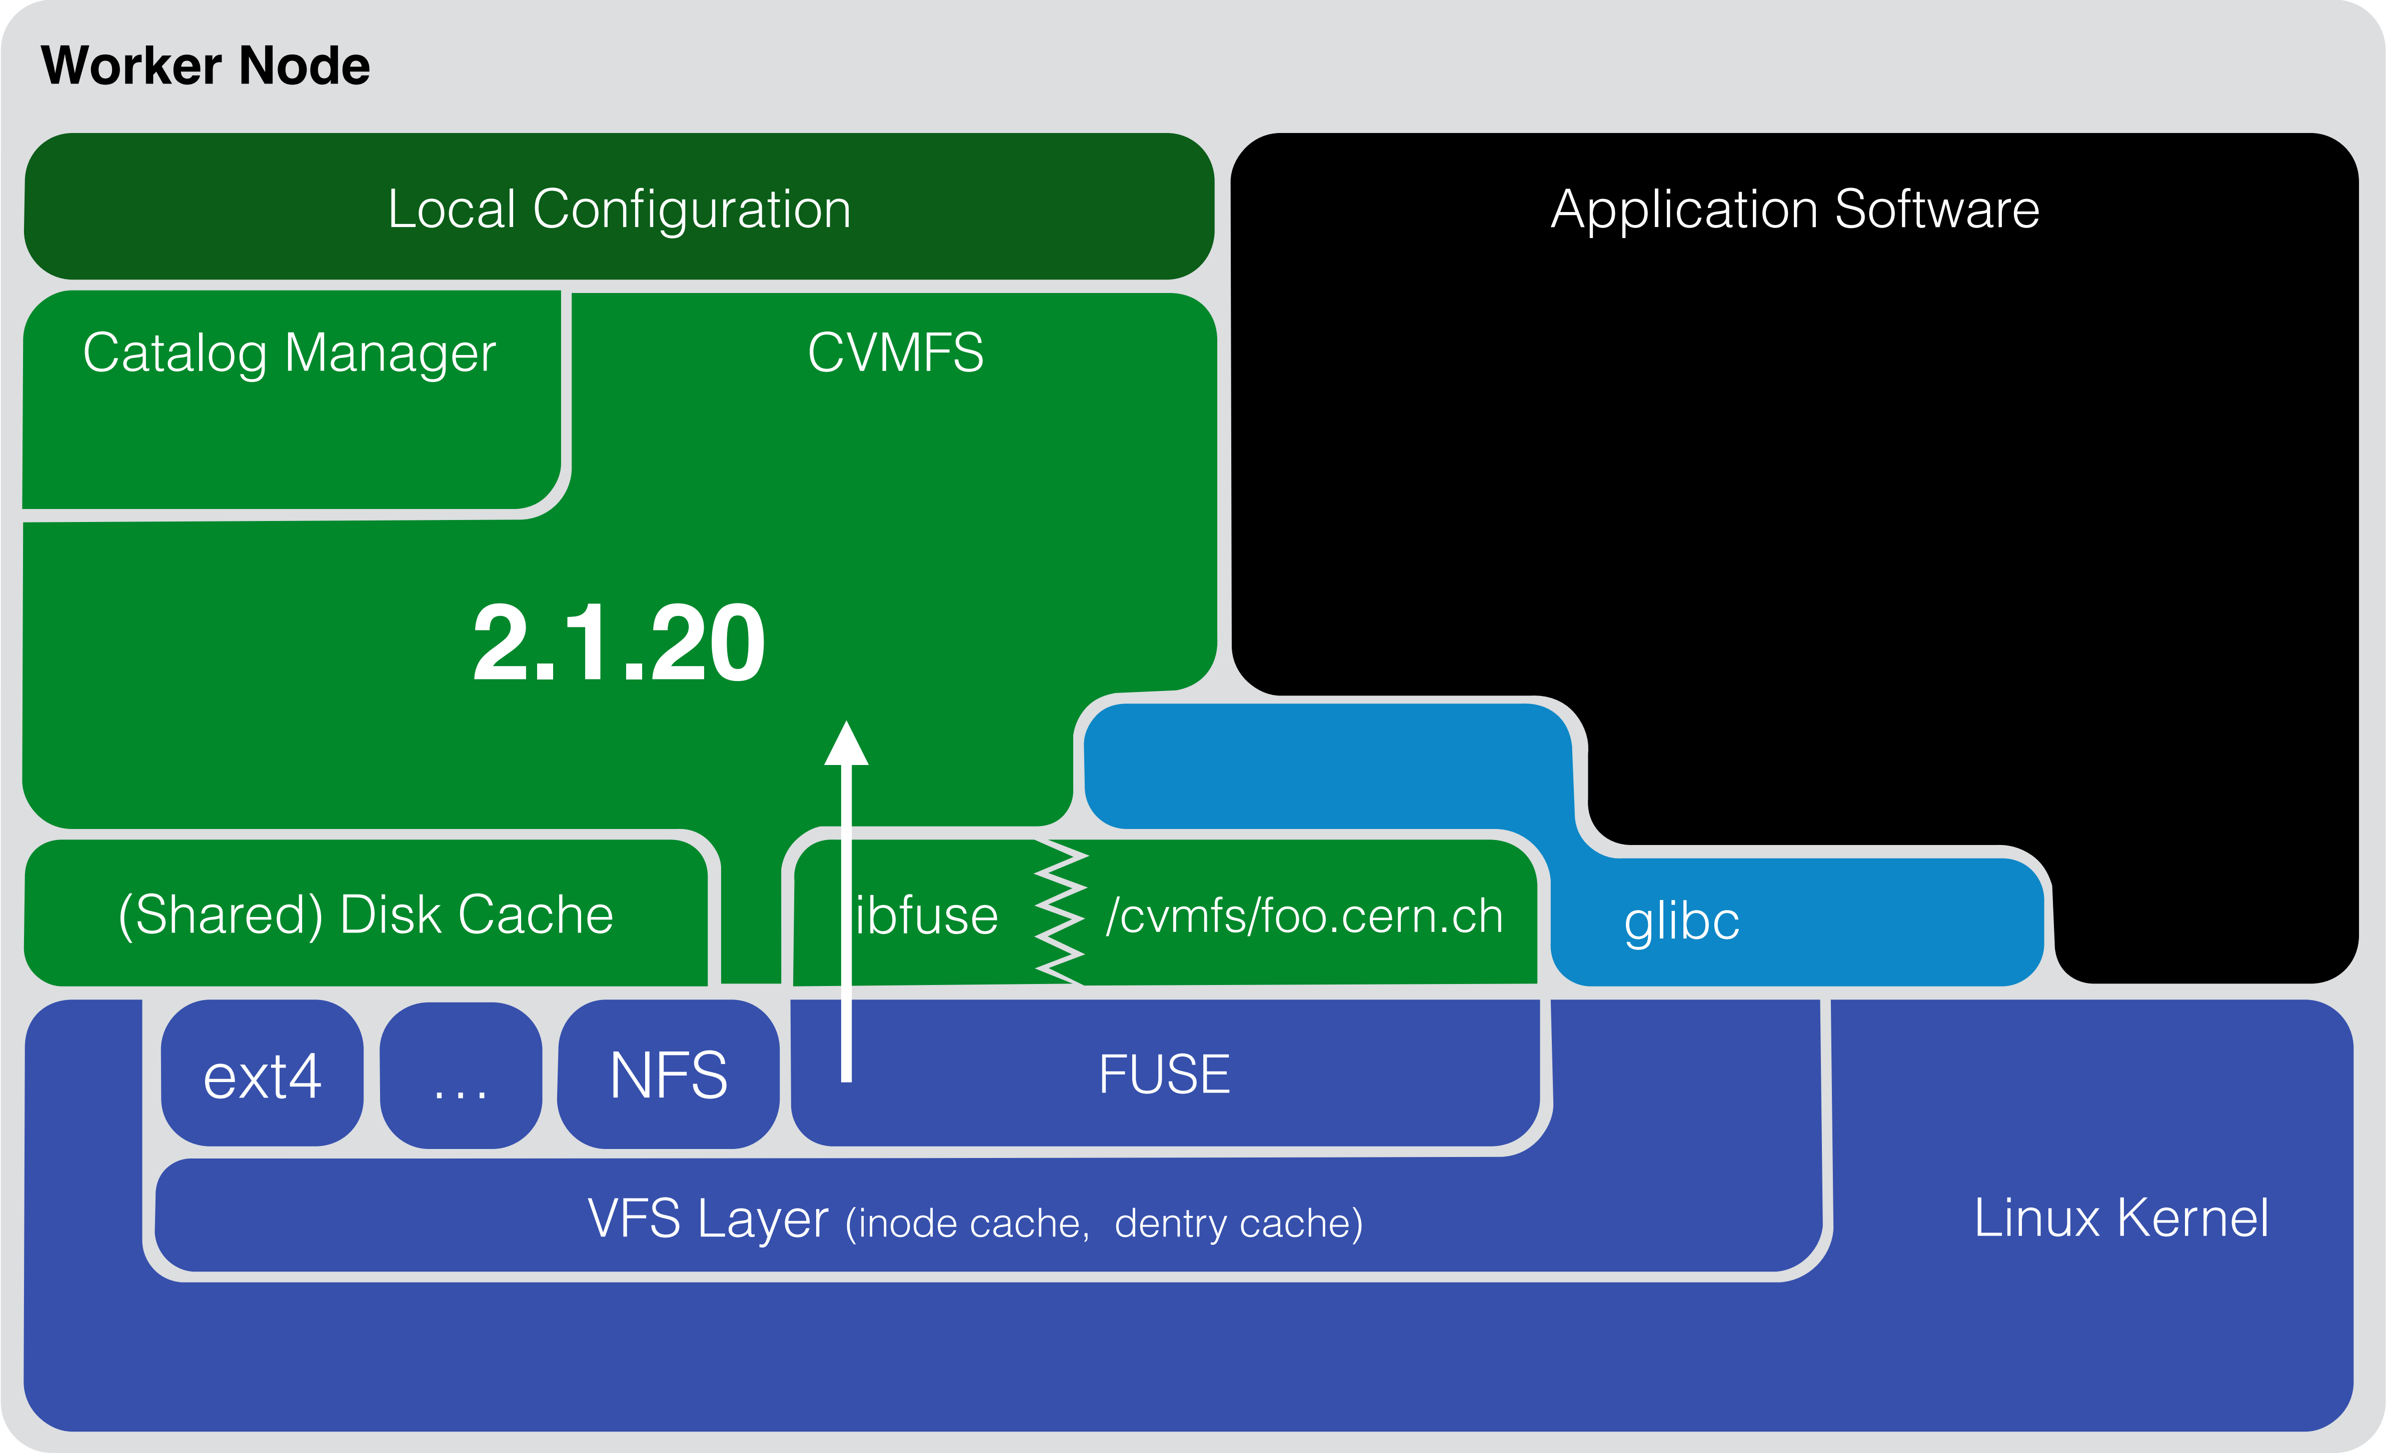

Alternative Storage Backends

- “Plug-in” Architecture since CernVM-FS Server 2.1.17
  - Potential for adding alternative storage drivers (S3, Ceph, Basho Riak, OpenStack Swift, …)

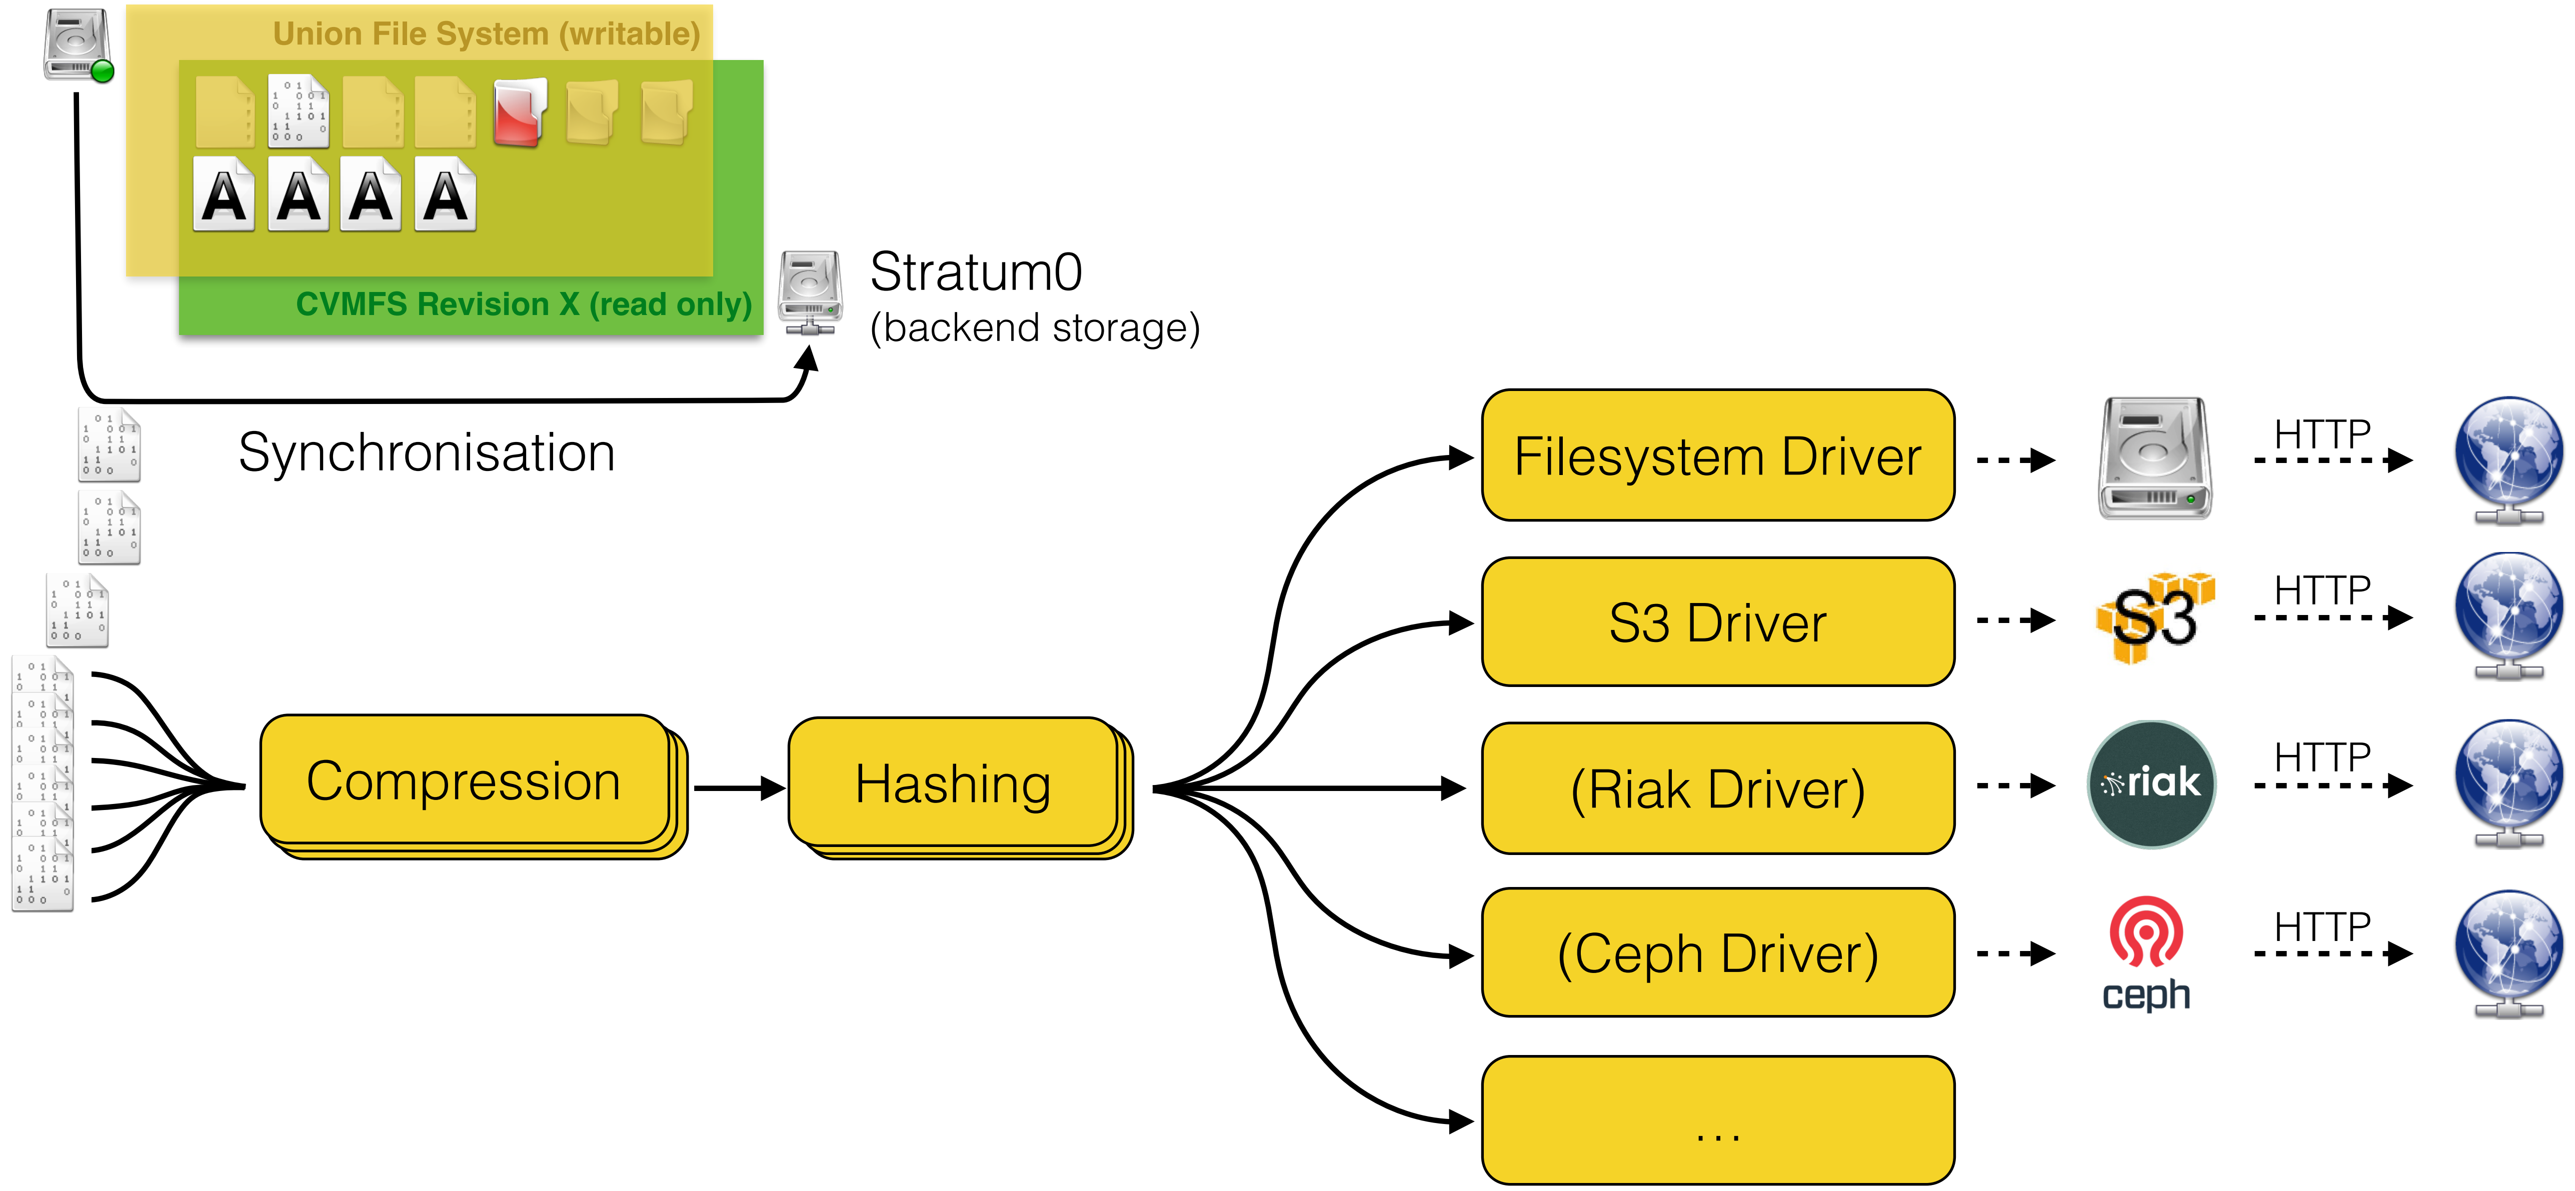

From POSIX to Blob-Objects

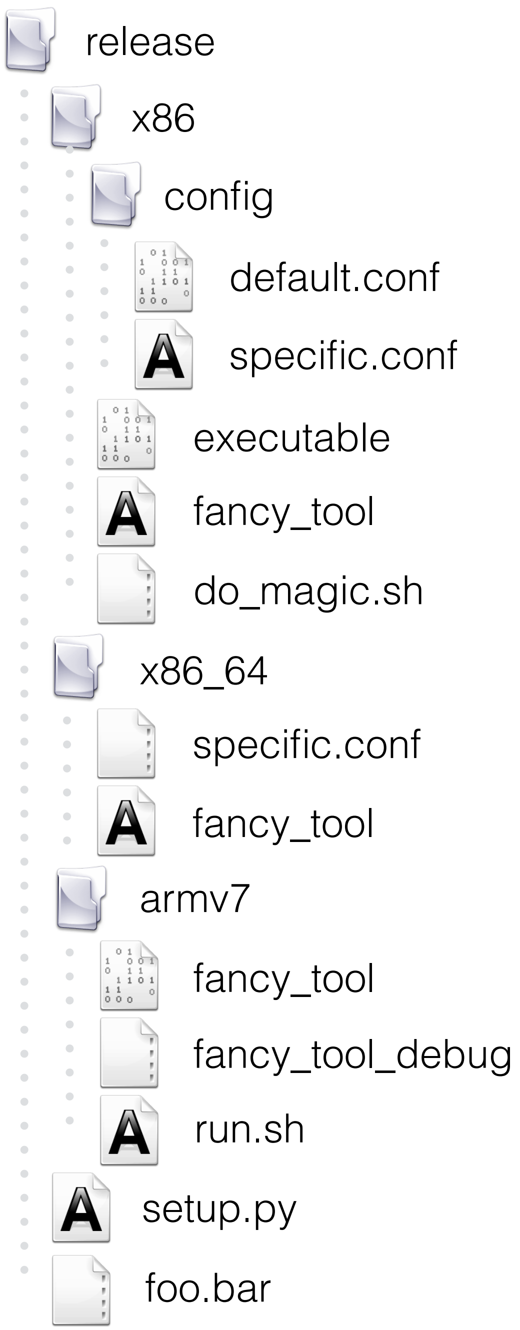

- release
  - x86
    - config
      - default.conf
      - specific.conf
    - executable
    - fancy_tool
    - do_magic.sh
  - x86_64
    - specific.conf
    - fancy_tool
  - armv7
    - fancy_tool
    - fancy_tool_debug
    - run.sh
  - setup.py
  - foo.bar

From POSIX to Blob-Objects

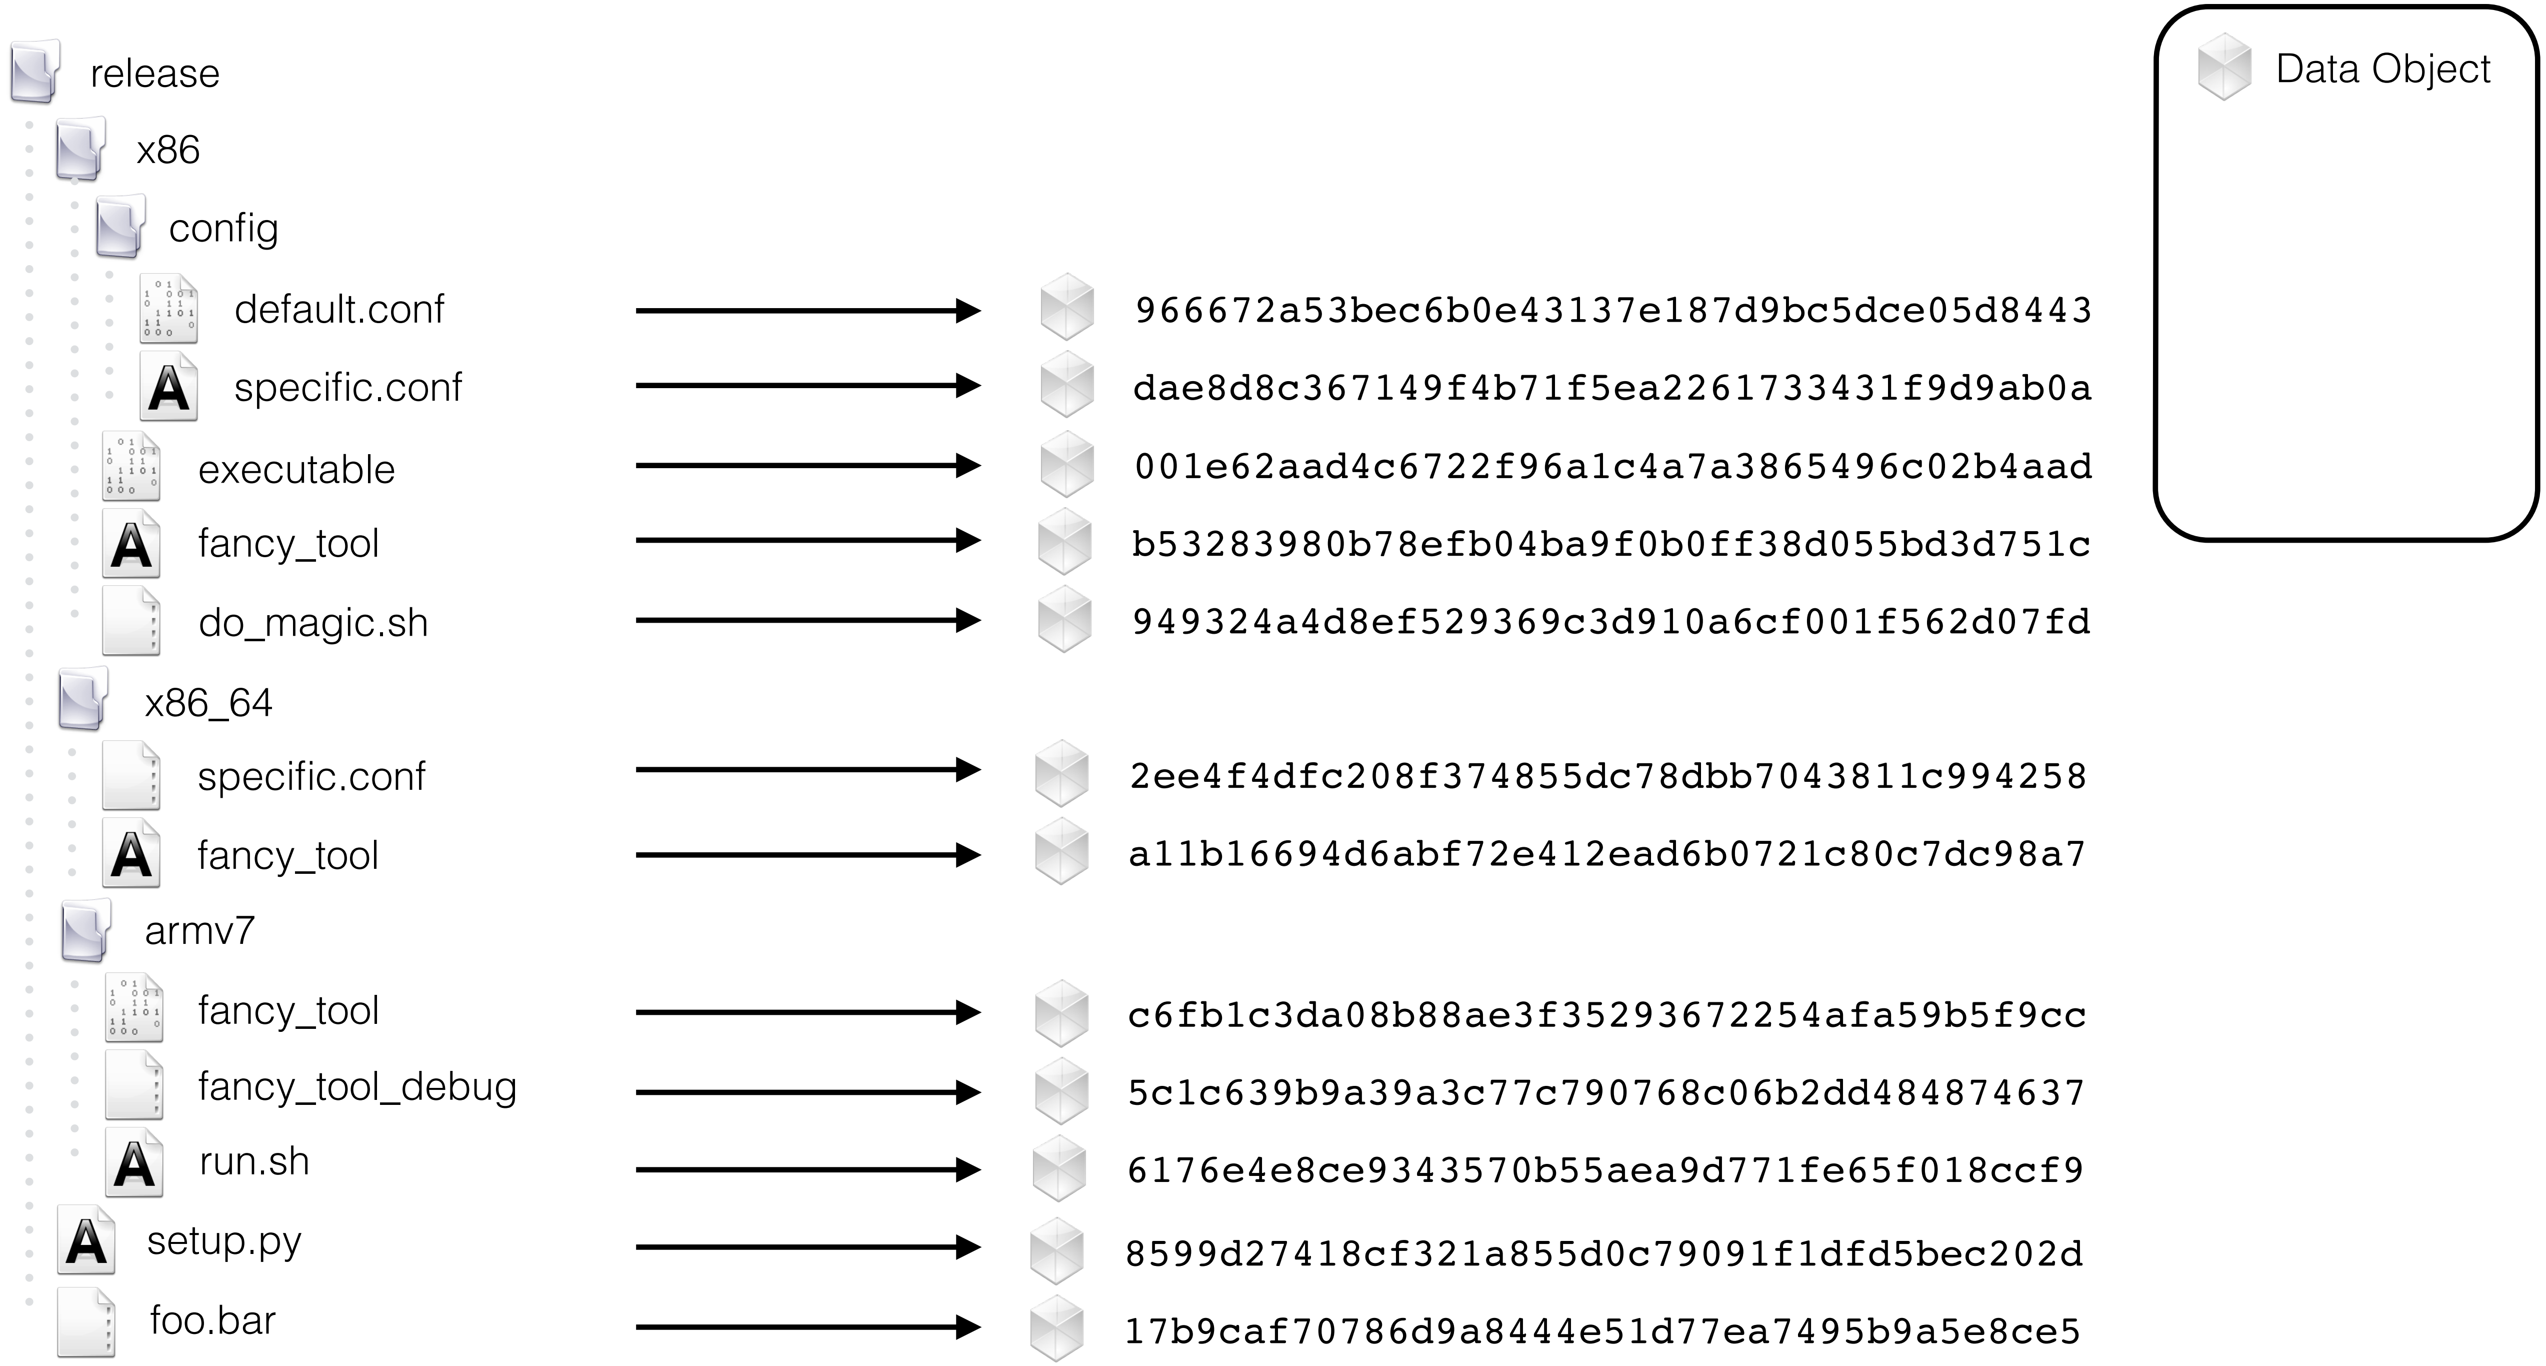

Data Object

release

x86

config

default.conf	966672a53bec6b0e43137e187d9bc5dce05d8443
specific.conf	dae8d8c367149f4b71f5ea2261733431f9d9ab0a
executable	001e62aad4c6722f96a1c4a7a3865496c02b4aad
fancy_tool	b53283980b78efb04ba9f0b0ff38d055bd3d751c
do_magic.sh	949324a4d8ef529369c3d910a6cf001f562d07fd

x86_64

specific.conf	2ee4f4dfc208f374855dc78dbb7043811c994258
fancy_tool	a11b16694d6abf72e412ead6b0721c80c7dc98a7

armv7

fancy_tool	c6fb1c3da08b88ae3f35293672254afa59b5f9cc
fancy_tool_debug	5c1c639b9a39a3c77c790768c06b2dd484874637
run.sh	6176e4e8ce9343570b55aea9d771fe65f018ccf9
setup.py	8599d27418cf321a855d0c79091f1dfd5bec202d
foo.bar	17b9caf70786d9a8444e51d77ea7495b9a5e8ce5

From POSIX to Blob-Objects

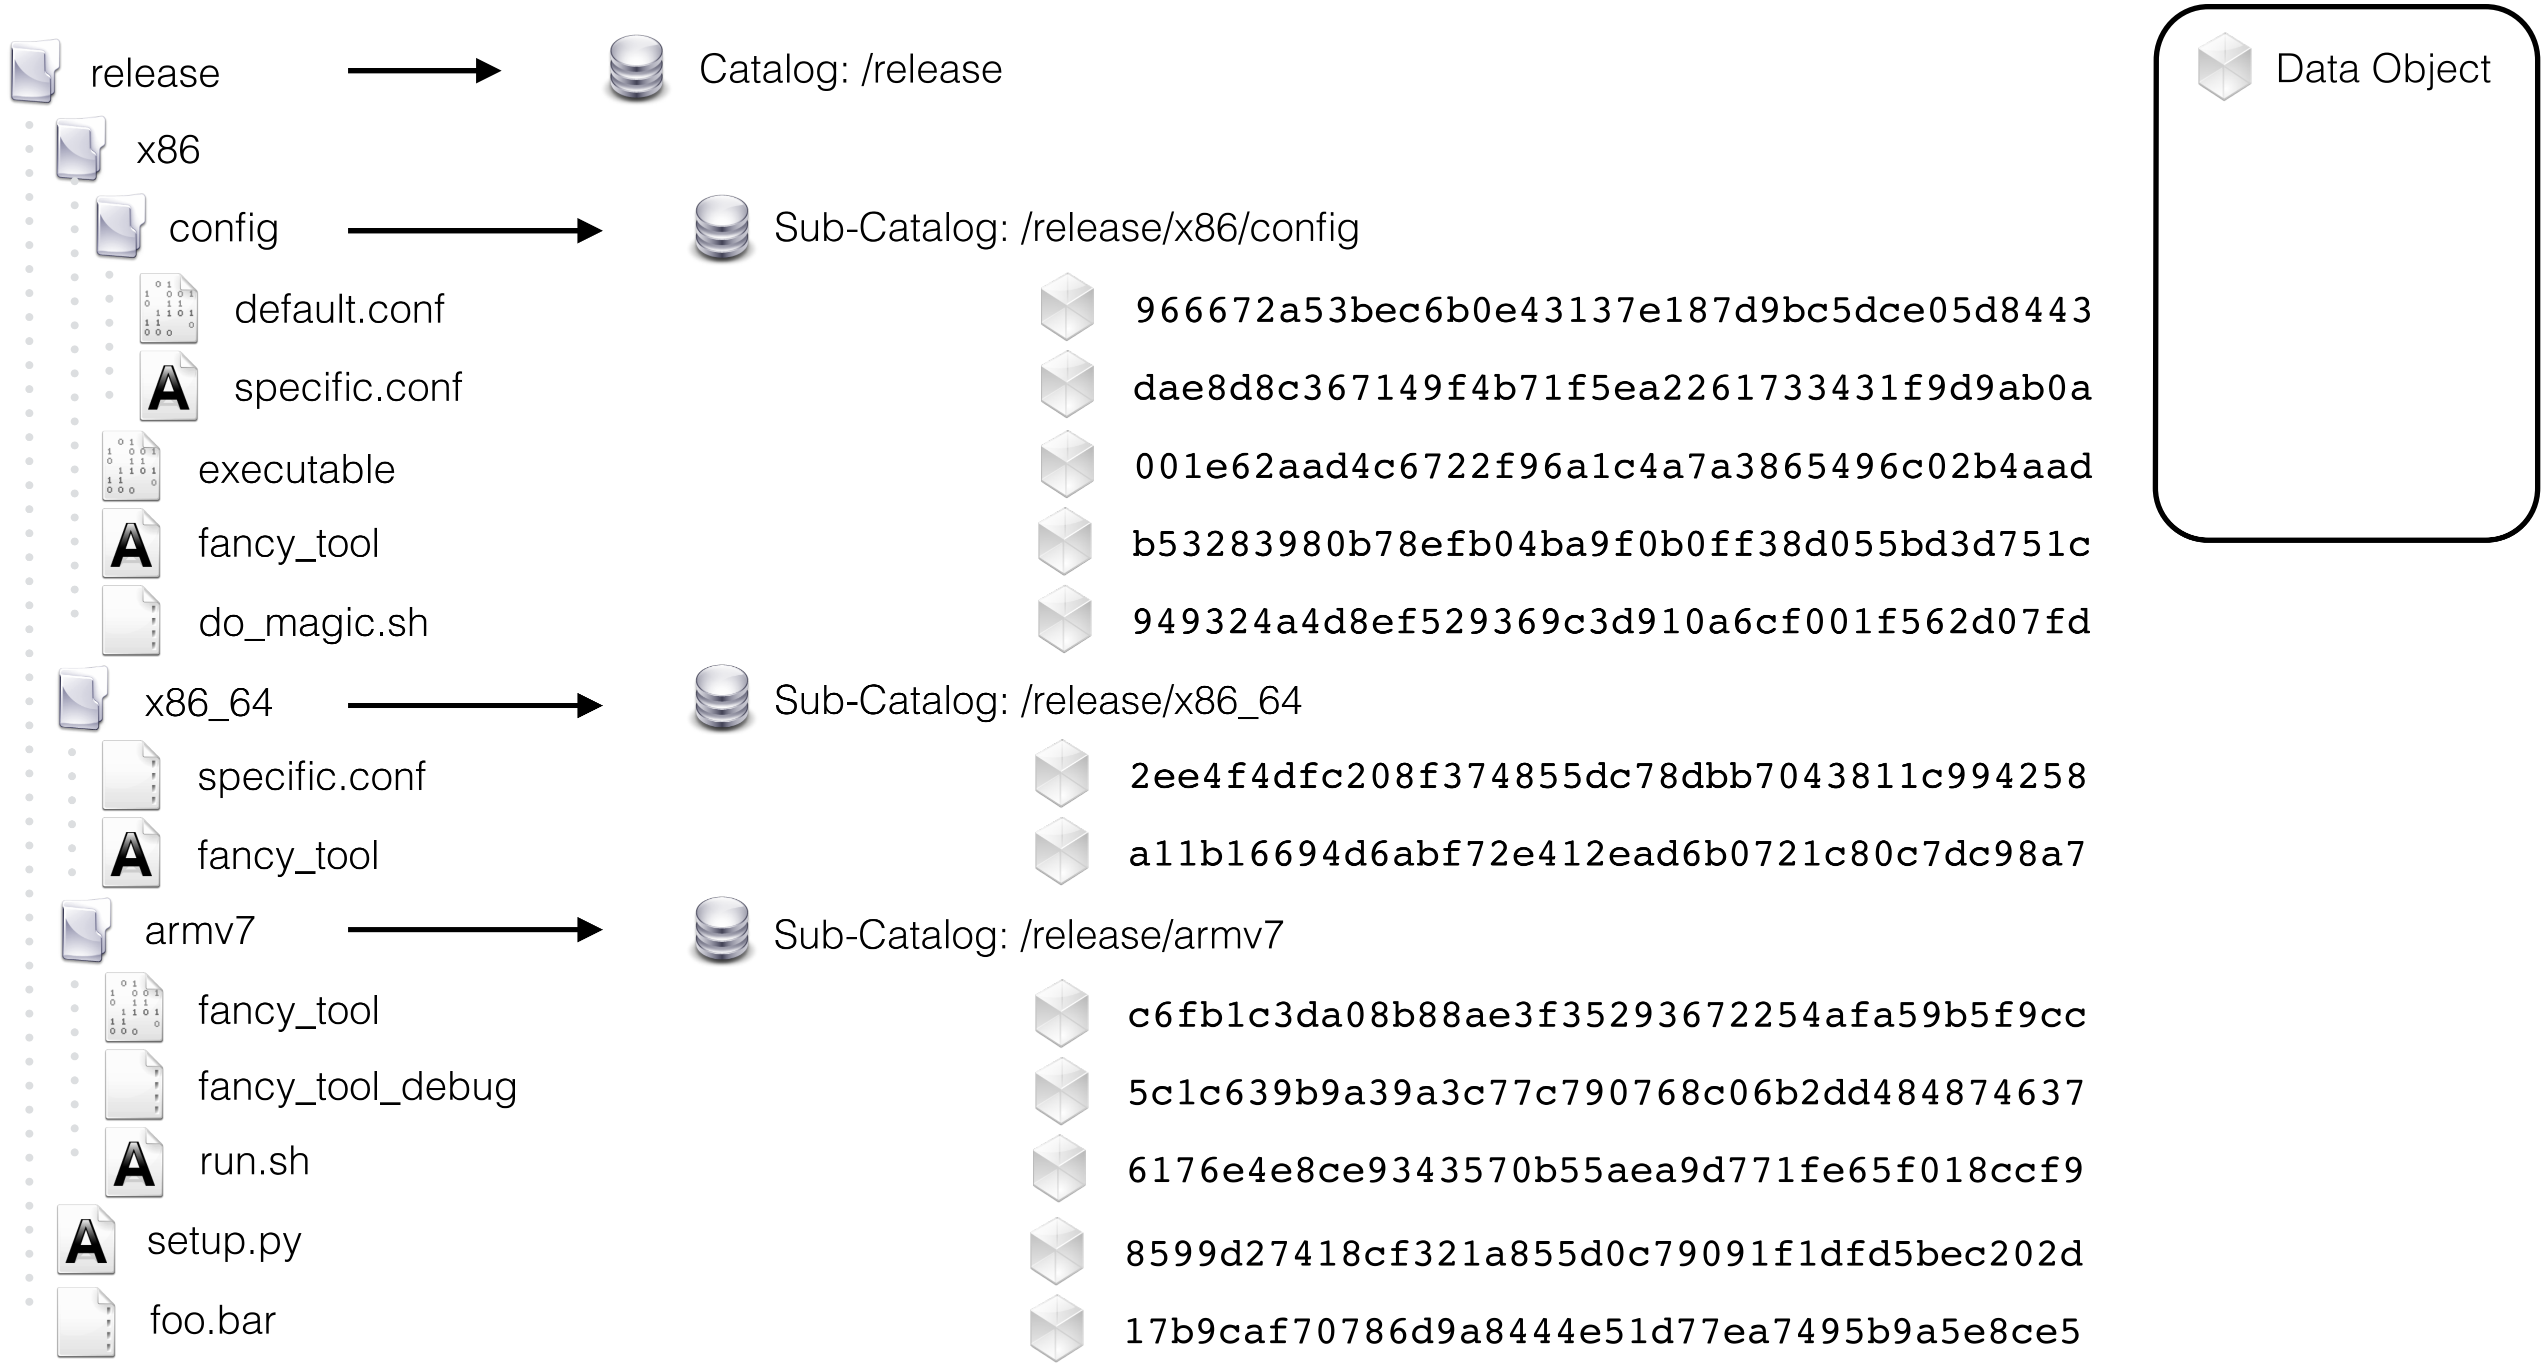

release → Catalog: /release

Data Object

x86

config → Sub-Catalog: /release/x86/config

default.conf 966672a53bec6b0e43137e187d9bc5dce05d8443

specific.conf dae8d8c367149f4b71f5ea2261733431f9d9ab0a

executable 001e62aad4c6722f96a1c4a7a3865496c02b4aad

fancy_tool b53283980b78efb04ba9f0b0ff38d055bd3d751c

do_magic.sh 949324a4d8ef529369c3d910a6cf001f562d07fd

x86_64 → Sub-Catalog: /release/x86_64

specific.conf 2ee4f4dfc208f374855dc78dbb7043811c994258

fancy_tool a11b16694d6abf72e412ead6b0721c80c7dc98a7

armv7 → Sub-Catalog: /release/armv7

fancy_tool c6fb1c3da08b88ae3f35293672254afa59b5f9cc

fancy_tool_debug 5c1c639b9a39a3c77c790768c06b2dd484874637

run.sh 6176e4e8ce9343570b55aea9d771fe65f018ccf9

setup.py 8599d27418cf321a855d0c79091f1dfd5bec202d

foo.bar 17b9caf70786d9a8444e51d77ea7495b9a5e8ce5

From POSIX to Blob-Objects

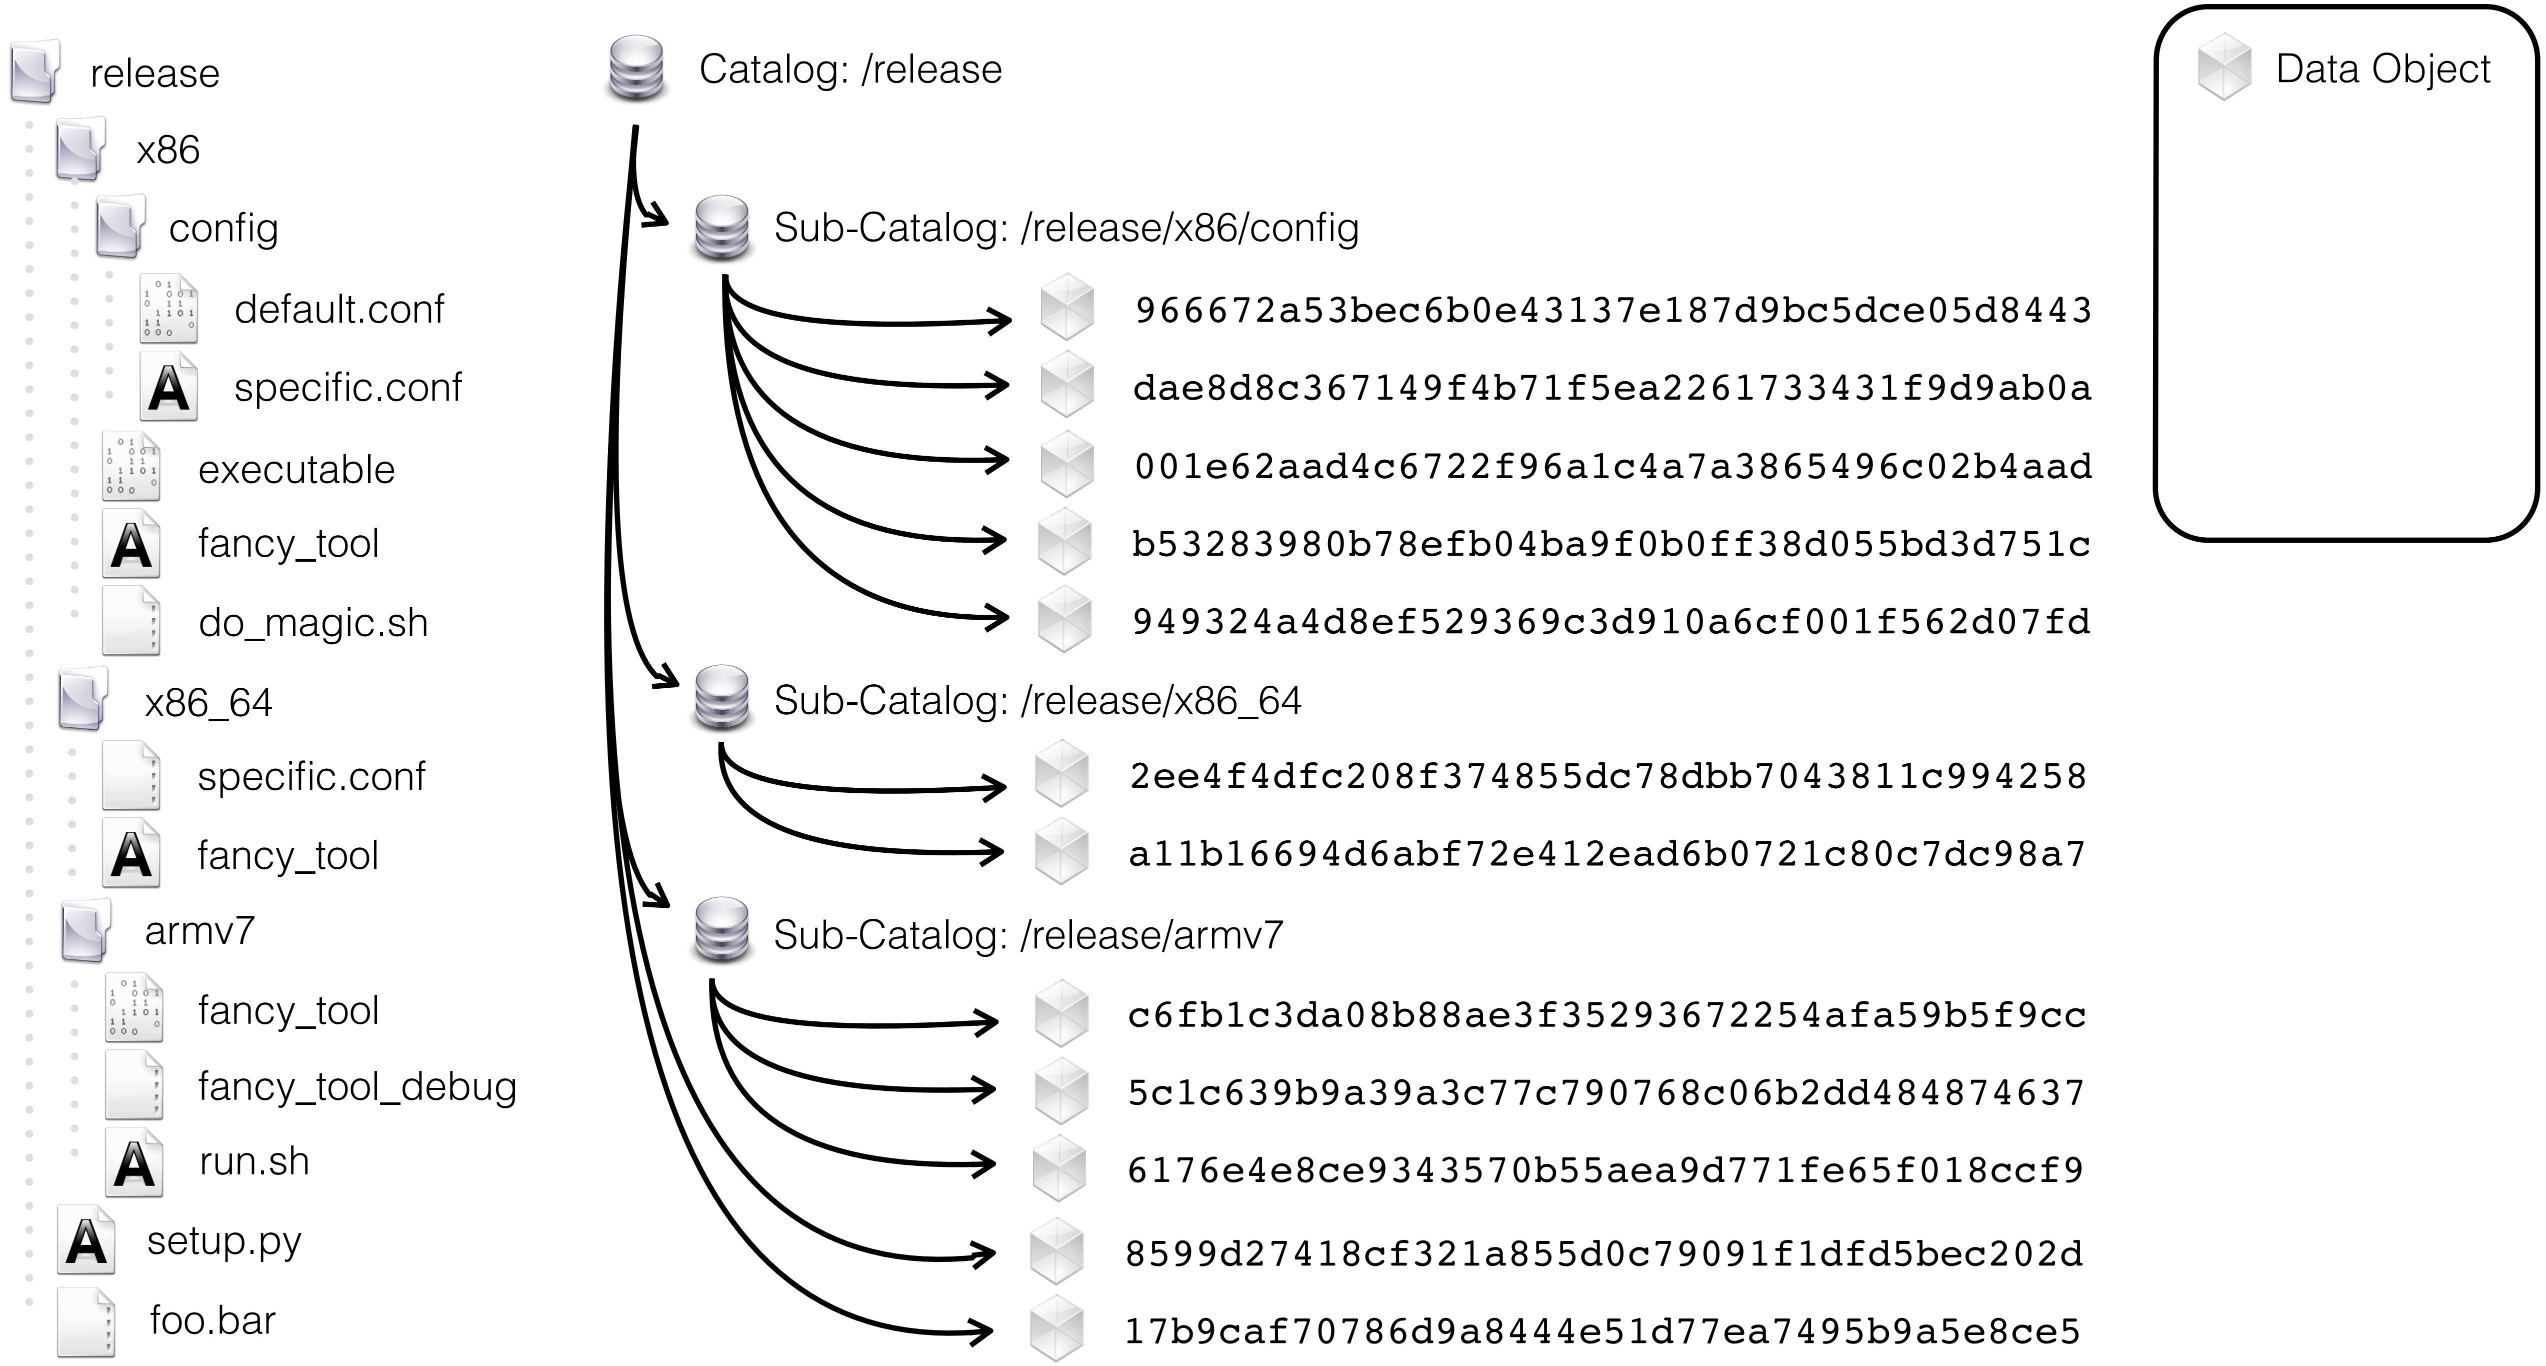

release
x86
config
default.conf
specific.conf
executable
fancy_tool
do_magic.sh
x86_64
specific.conf
fancy_tool
armv7
fancy_tool
fancy_tool_debug
run.sh
setup.py
foo.bar

Catalog: /release

Sub-Catalog: /release/x86/config

966672a53bec6b0e43137e187d9bc5dce05d8443
dae8d8c367149f4b71f5ea2261733431f9d9ab0a
001e62aad4c6722f96a1c4a7a3865496c02b4aad
b53283980b78efb04ba9f0b0ff38d055bd3d751c
949324a4d8ef529369c3d910a6cf001f562d07fd

Sub-Catalog: /release/x86_64

2ee4f4dfc208f374855dc78dbb7043811c994258
a11b16694d6abf72e412ead6b0721c80c7dc98a7

Sub-Catalog: /release/armv7

c6fb1c3da08b88ae3f35293672254afa59b5f9cc
5c1c639b9a39a3c77c790768c06b2dd484874637
6176e4e8ce9343570b55aea9d771fe65f018ccf9
8599d27418cf321a855d0c79091f1dfd5bec202d
17b9caf70786d9a8444e51d77ea7495b9a5e8ce5

Data Object

From POSIX to Blob-Objects

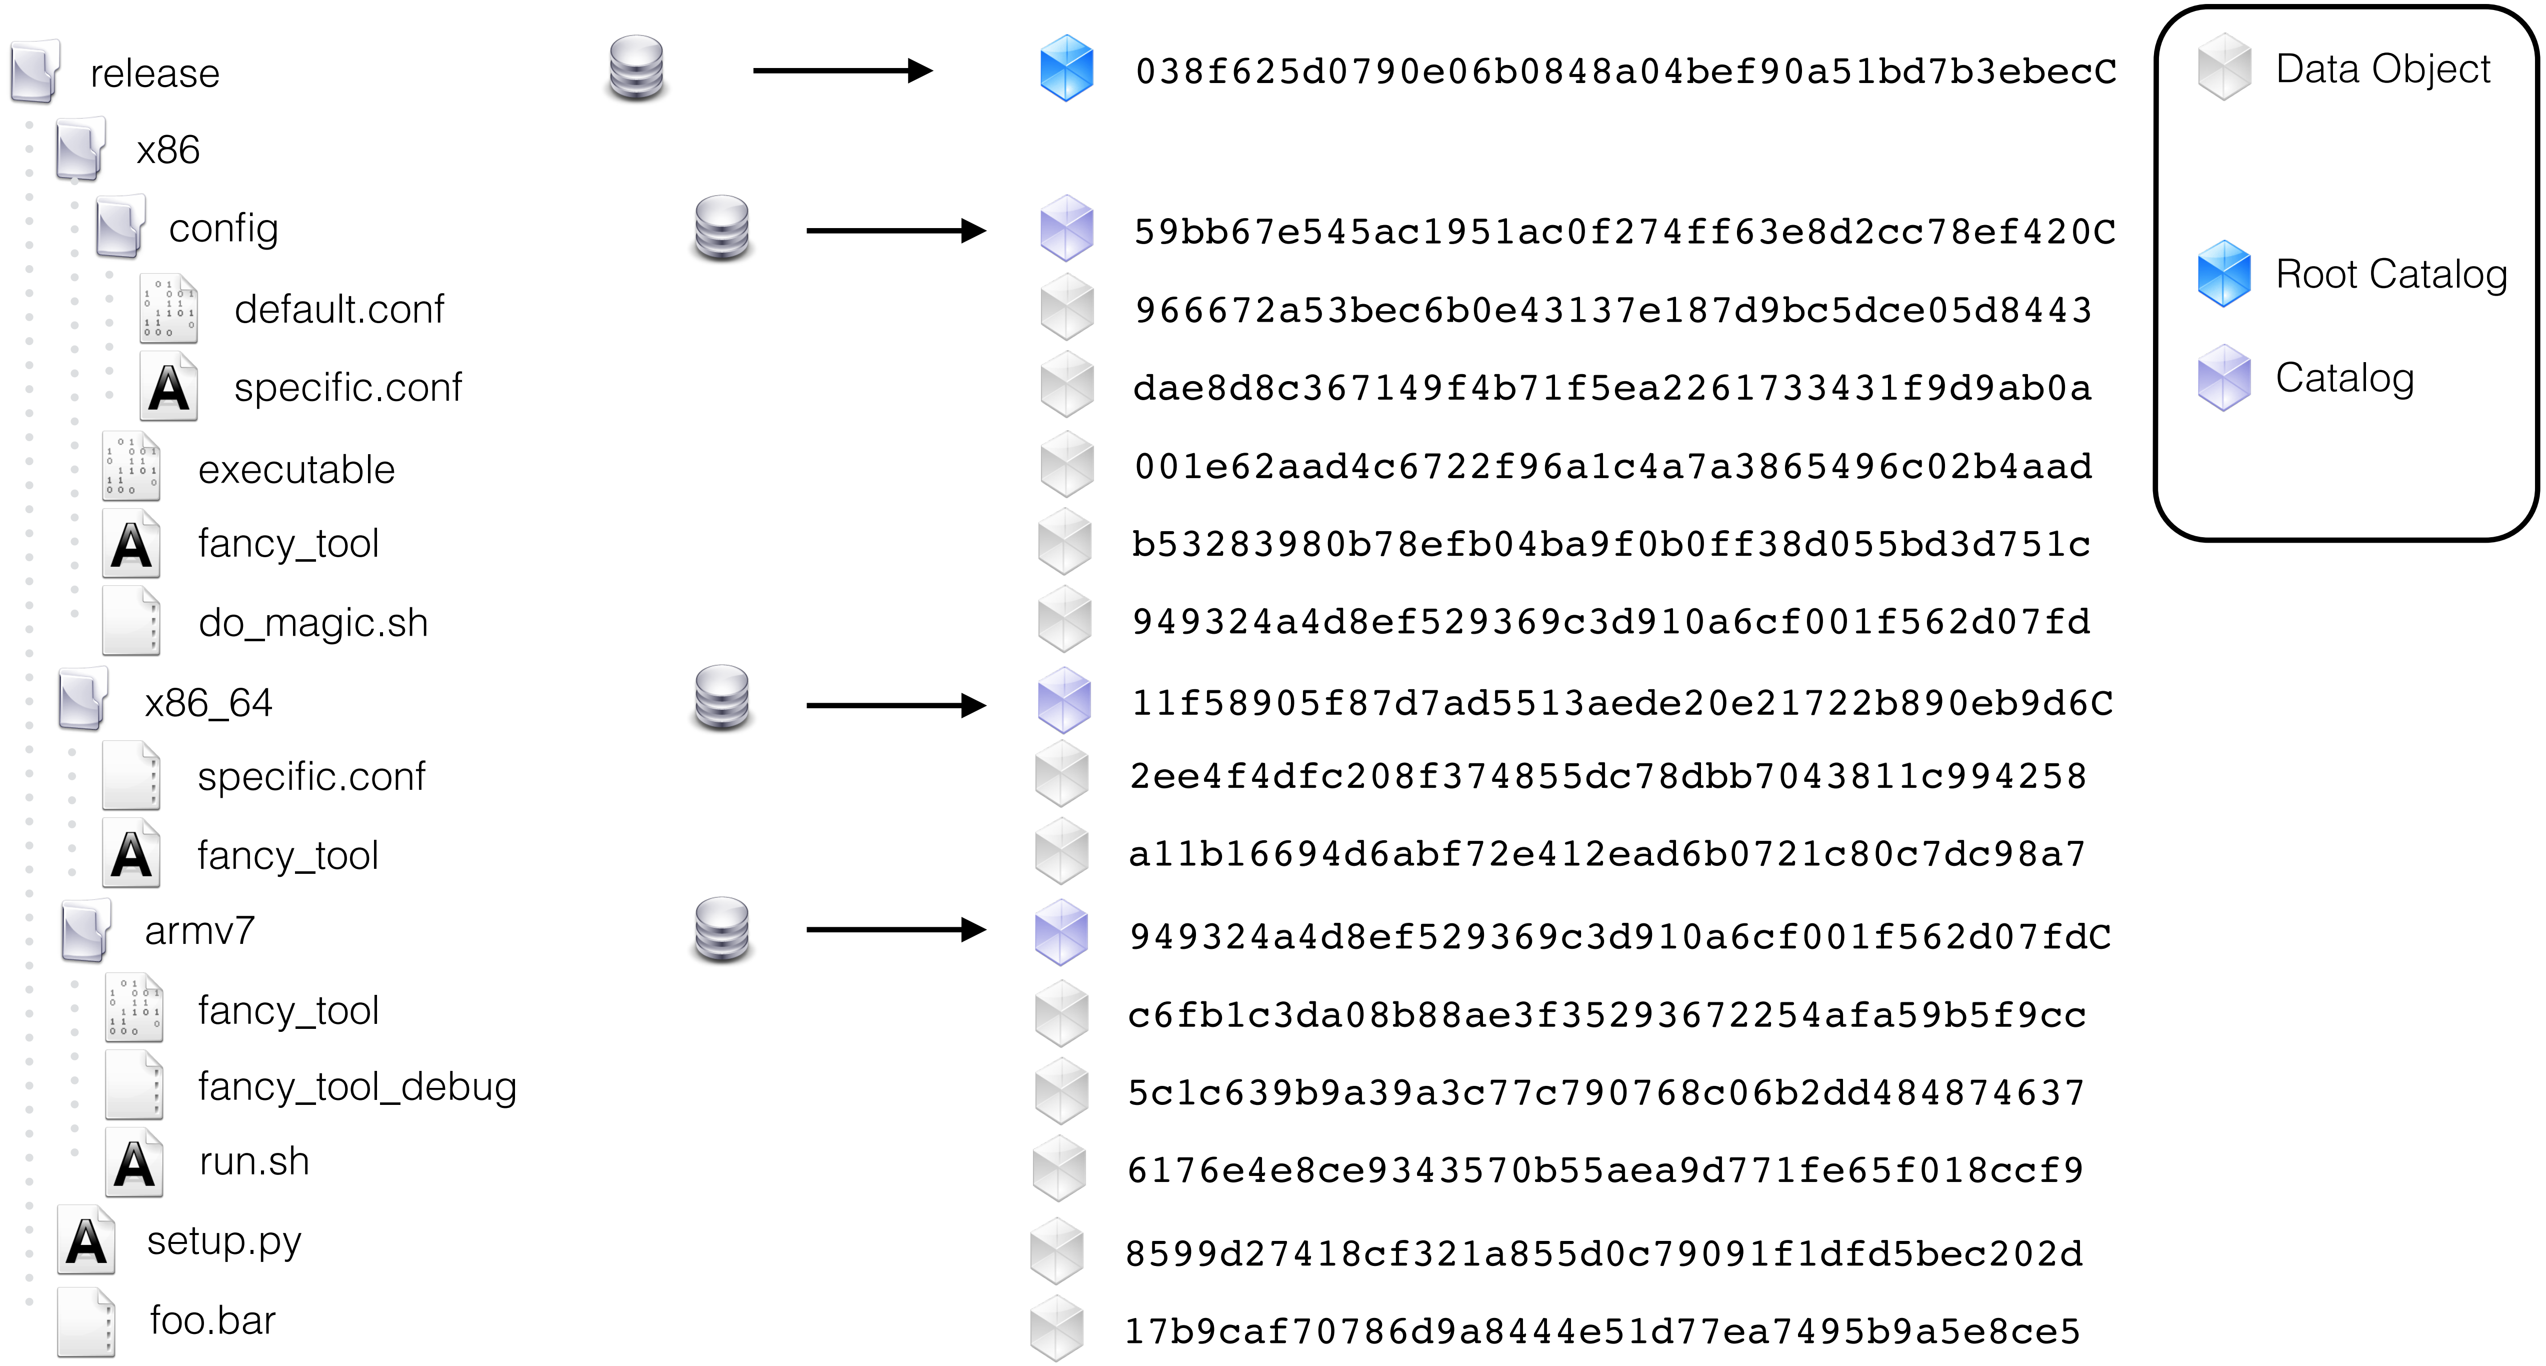

release
x86
config
default.conf
specific.conf
executable
fancy_tool
do_magic.sh
x86_64
specific.conf
fancy_tool
armv7
fancy_tool
fancy_tool_debug
run.sh
setup.py
foo.bar

038f625d0790e06b0848a04bef90a51bd7b3ebecC
59bb67e545ac1951ac0f274ff63e8d2cc78ef420C
966672a53bec6b0e43137e187d9bc5dce05d8443
dae8d8c367149f4b71f5ea2261733431f9d9ab0a
001e62aad4c6722f96a1c4a7a3865496c02b4aad
b53283980b78efb04ba9f0b0ff38d055bd3d751c
949324a4d8ef529369c3d910a6cf001f562d07fd
11f58905f87d7ad5513aede20e21722b890eb9d6C
2ee4f4dfc208f374855dc78dbb7043811c994258
a11b16694d6abf72e412ead6b0721c80c7dc98a7
949324a4d8ef529369c3d910a6cf001f562d07fdC
c6fb1c3da08b88ae3f35293672254afa59b5f9cc
5c1c639b9a39a3c77c790768c06b2dd484874637
6176e4e8ce9343570b55aea9d771fe65f018ccf9
8599d27418cf321a855d0c79091f1dfd5bec202d
17b9caf70786d9a8444e51d77ea7495b9a5e8ce5

Data Object
Root Catalog
Catalog

From POSIX to Blob-Objects

🟦 038f625d0790e06b0848a04bef90a51bd7b3ebecC

🟪 59bb67e545ac1951ac0f274ff63e8d2cc78ef420C

⬜ 966672a53bec6b0e43137e187d9bc5dce05d8443

⬜ dae8d8c367149f4b71f5ea2261733431f9d9ab0a

⬜ 001e62aad4c6722f96a1c4a7a3865496c02b4aad

⬜ b53283980b78efb04ba9f0b0ff38d055bd3d751c

⬜ 949324a4d8ef529369c3d910a6cf001f562d07fd

🟪 11f58905f87d7ad5513aede20e21722b890eb9d6C

⬜ 2ee4f4dfc208f374855dc78dbb7043811c994258

⬜ a11b16694d6abf72e412ead6b0721c80c7dc98a7

🟪 949324a4d8ef529369c3d910a6cf001f562d07fdC

⬜ c6fb1c3da08b88ae3f35293672254afa59b5f9cc

⬜ 5c1c639b9a39a3c77c790768c06b2dd484874637

⬜ 6176e4e8ce9343570b55aea9d771fe65f018ccf9

⬜ 8599d27418cf321a855d0c79091f1dfd5bec202d

⬜ 17b9caf70786d9a8444e51d77ea7495b9a5e8ce5

⬜ Data Object

🟦 Root Catalog

🟪 Catalog

From POSIX to Blob-Objects

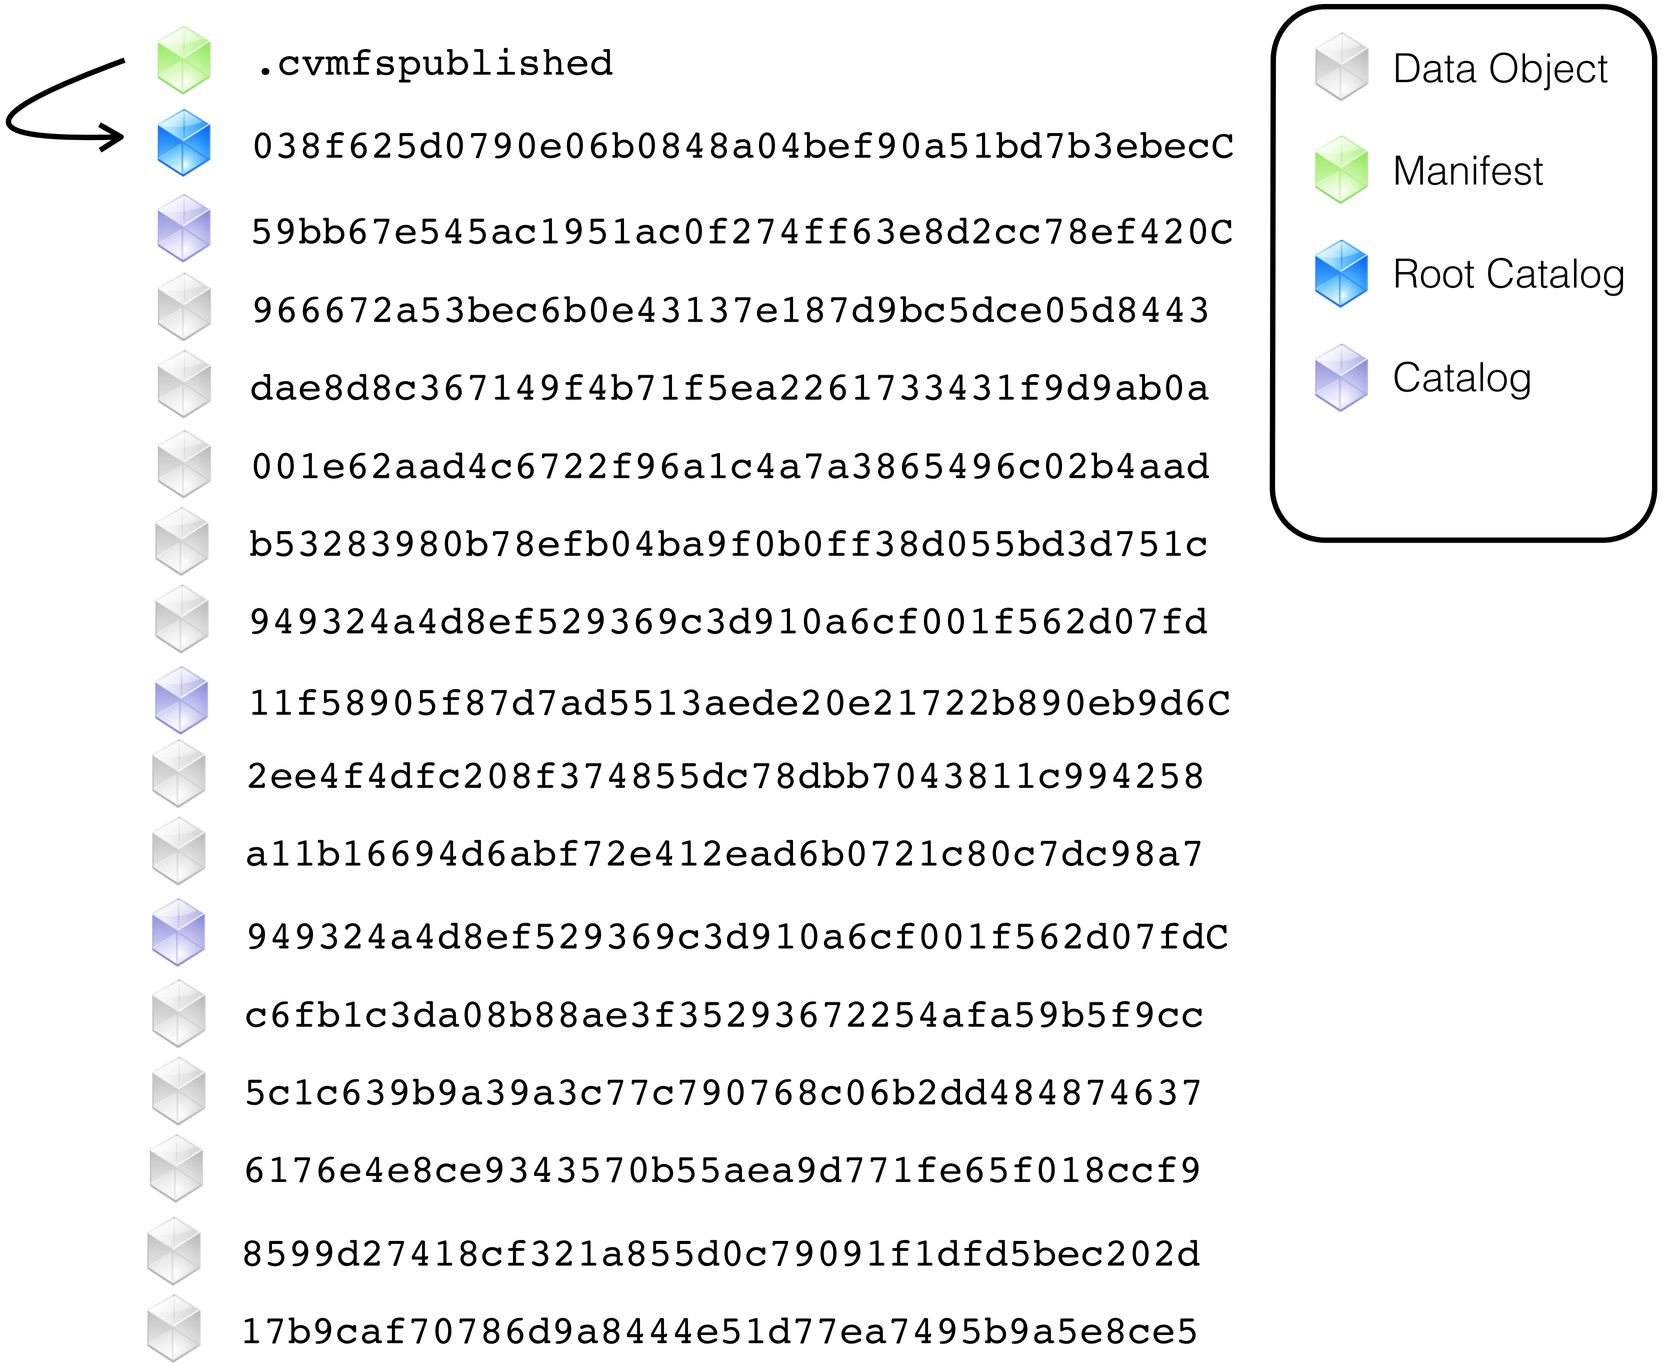

```
.cvmfspublished
038f625d0790e06b0848a04bef90a51bd7b3ebecC
59bb67e545ac1951ac0f274ff63e8d2cc78ef420C
966672a53bec6b0e43137e187d9bc5dce05d8443
dae8d8c367149f4b71f5ea2261733431f9d9ab0a
001e62aad4c6722f96a1c4a7a3865496c02b4aad
b53283980b78efb04ba9f0b0ff38d055bd3d751c
949324a4d8ef529369c3d910a6cf001f562d07fd
11f58905f87d7ad5513aede20e21722b890eb9d6C
2ee4f4dfc208f374855dc78dbb7043811c994258
a11b16694d6abf72e412ead6b0721c80c7dc98a7
949324a4d8ef529369c3d910a6cf001f562d07fdC
c6fb1c3da08b88ae3f35293672254afa59b5f9cc
5c1c639b9a39a3c77c790768c06b2dd484874637
6176e4e8ce9343570b55aea9d771fe65f018ccf9
8599d27418cf321a855d0c79091f1dfd5bec202d
17b9caf70786d9a8444e51d77ea7495b9a5e8ce5
```

From POSIX to Blob-Objects

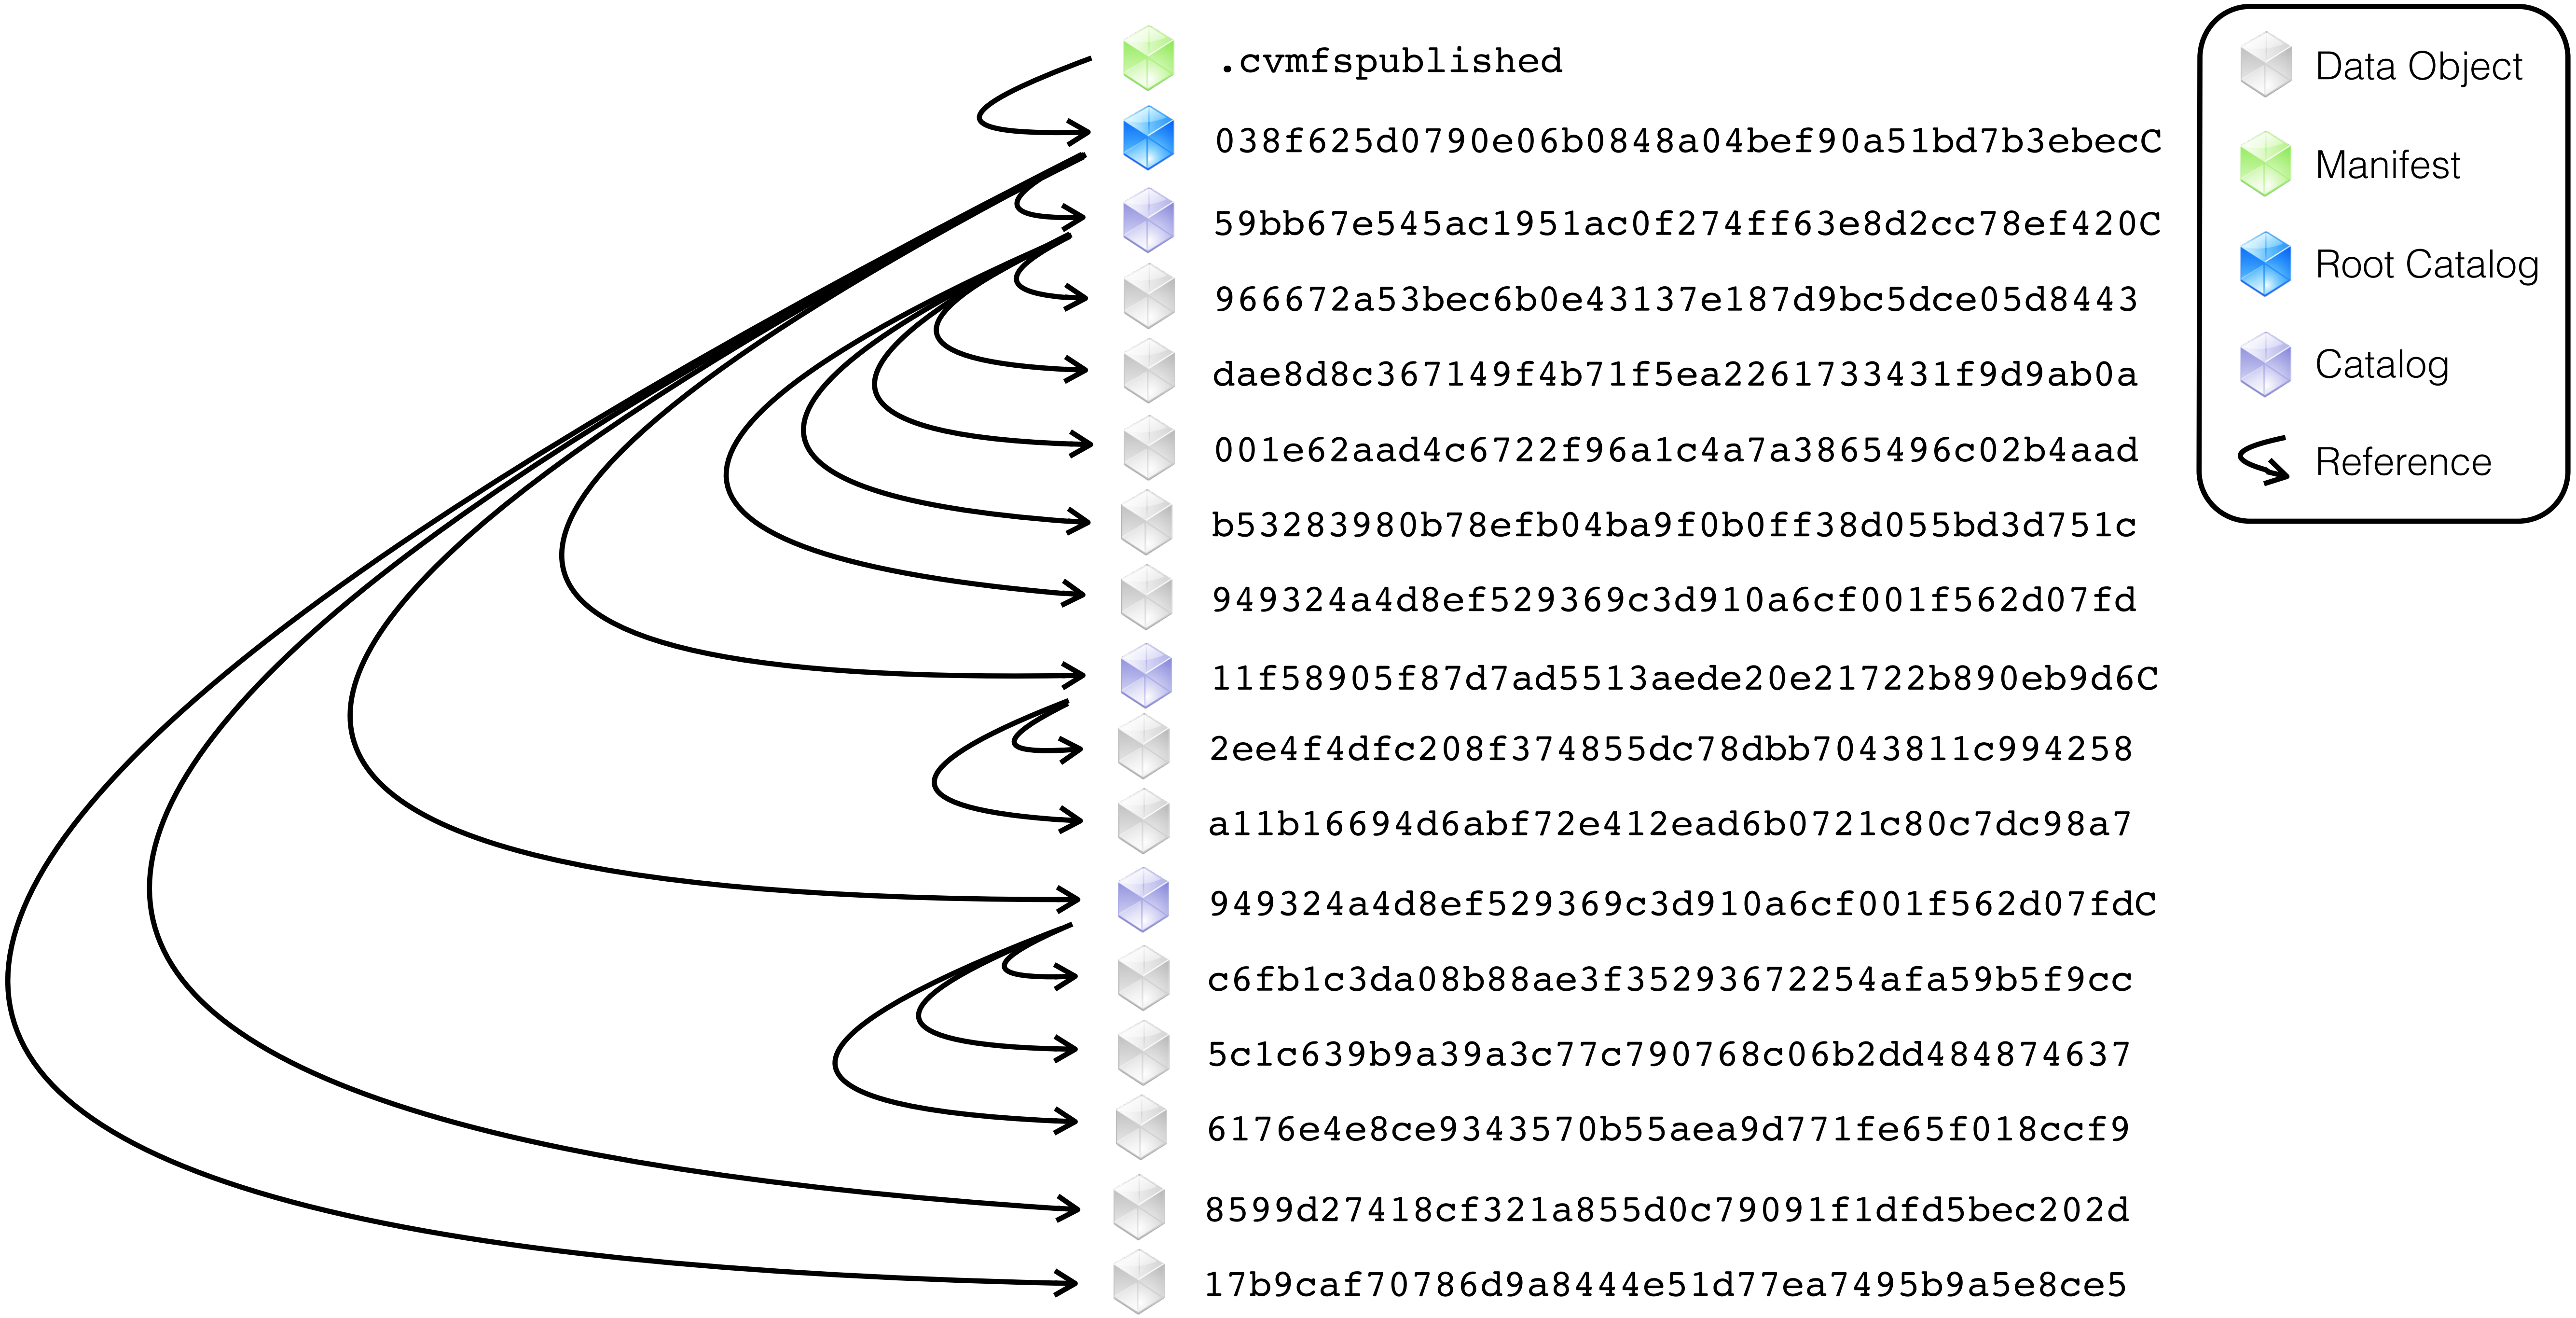

.cvmfspublished

038f625d0790e06b0848a04bef90a51bd7b3ebecC

59bb67e545ac1951ac0f274ff63e8d2cc78ef420C

966672a53bec6b0e43137e187d9bc5dce05d8443

dae8d8c367149f4b71f5ea2261733431f9d9ab0a

001e62aad4c6722f96a1c4a7a3865496c02b4aad

b53283980b78efb04ba9f0b0ff38d055bd3d751c

949324a4d8ef529369c3d910a6cf001f562d07fd

11f58905f87d7ad5513aede20e21722b890eb9d6C

2ee4f4dfc208f374855dc78dbb7043811c994258

a11b16694d6abf72e412ead6b0721c80c7dc98a7

949324a4d8ef529369c3d910a6cf001f562d07fdC

c6fb1c3da08b88ae3f35293672254afa59b5f9cc

5c1c639b9a39a3c77c790768c06b2dd484874637

6176e4e8ce9343570b55aea9d771fe65f018ccf9

8599d27418cf321a855d0c79091f1dfd5bec202d

17b9caf70786d9a8444e51d77ea7495b9a5e8ce5

Data Object

Manifest

Root Catalog

Catalog

Reference

From POSIX to Blob-Objects

- **Merkle tree**
  - only .cvmfspublished needs to be signed
- **Content-Addressable Storage**
  - File de-duplication
  - Trivial file integrity checks
- **Flat Namespace**
  - Perfect for HTTP caching
  - Minimal storage API requirements (PUT, GET, [DELETE])

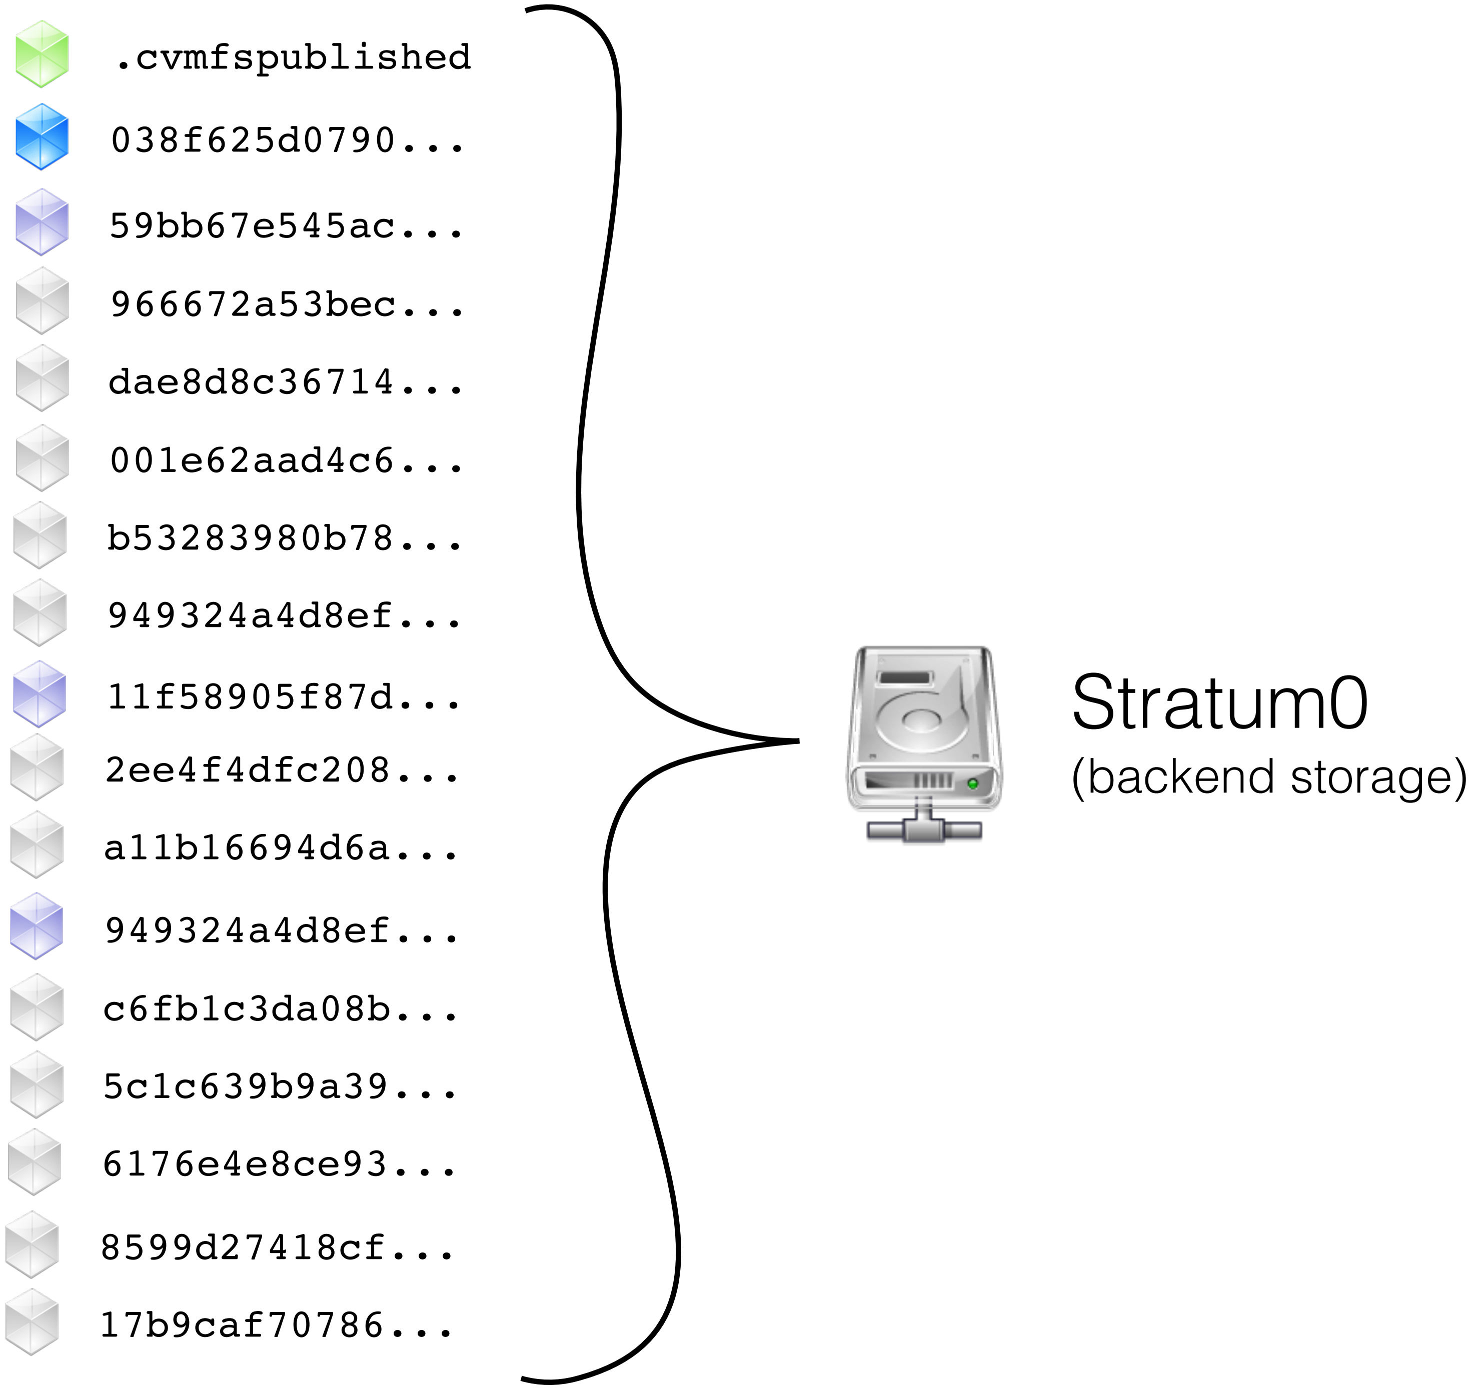

.cvmfspublished
038f625d0790...
59bb67e545ac...
966672a53bec...
dae8d8c36714...
001e62aad4c6...
b53283980b78...
949324a4d8ef...
11f58905f87d...
2ee4f4dfc208...
a11b16694d6a...
949324a4d8ef...
c6fb1c3da08b...
5c1c639b9a39...
6176e4e8ce93...
8599d27418cf...
17b9caf70786...

Stratum0
(backend storage)

From POSIX to Blob-Objects

- **Merkle tree**
  - only .cvmfspublished needs to be signed
- **Content-Addressable Storage**
  - File de-duplication
  - Trivial file integrity checks
- **Flat Namespace**
  - Perfect for HTTP caching
  - Minimal storage API requirements (PUT, GET, [DELETE])

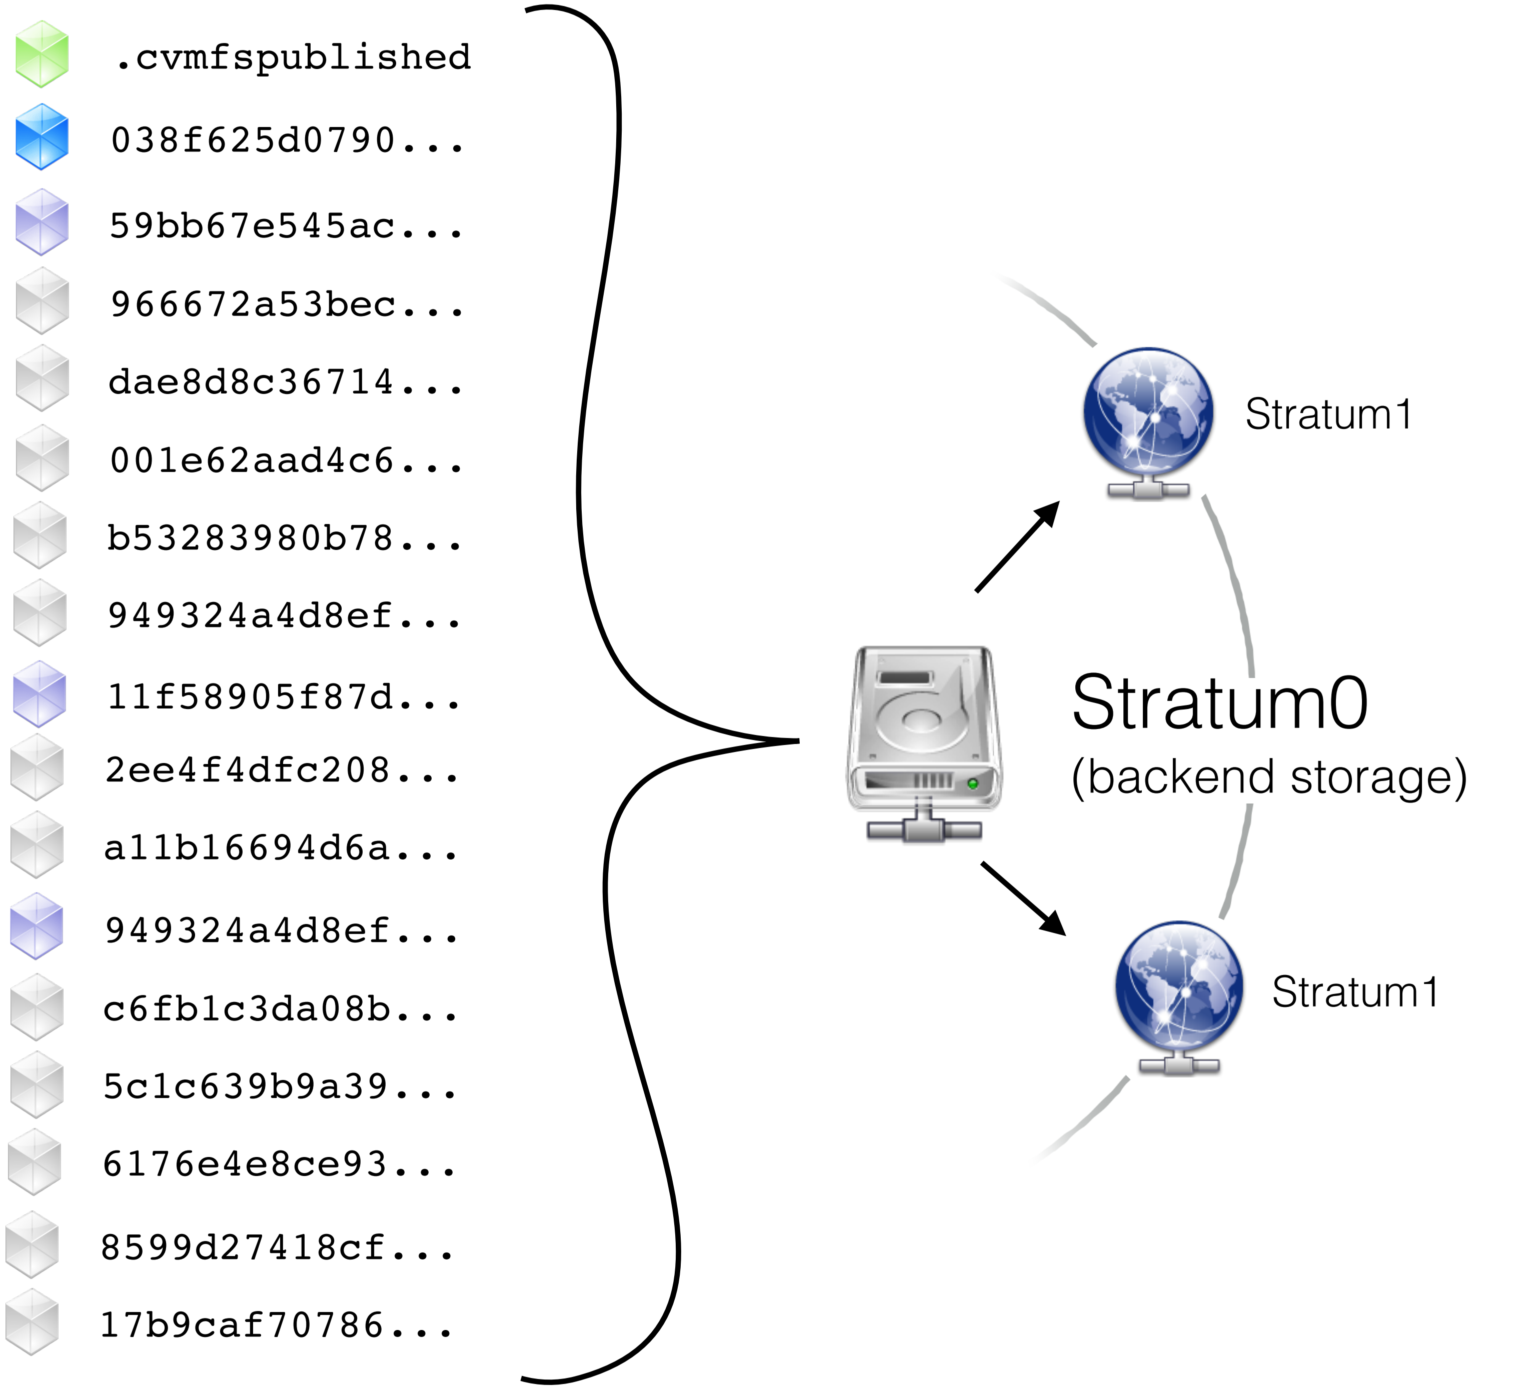

.cvmfspublished
038f625d0790...
59bb67e545ac...
966672a53bec...
dae8d8c36714...
001e62aad4c6...
b53283980b78...
949324a4d8ef...
11f58905f87d...
2ee4f4dfc208...
a11b16694d6a...
949324a4d8ef...
c6fb1c3da08b...
5c1c639b9a39...
6176e4e8ce93...
8599d27418cf...
17b9caf70786...

Stratum0
(backend storage)

Stratum1

Stratum1

From POSIX to Blob-Objects

- **Merkle tree**
  - only .cvmfspublished needs to be signed
- **Content-Addressable Storage**
  - File de-duplication
  - Trivial file integrity checks
- **Flat Namespace**
  - Perfect for HTTP caching
  - Minimal storage API requirements (PUT, GET, [DELETE])

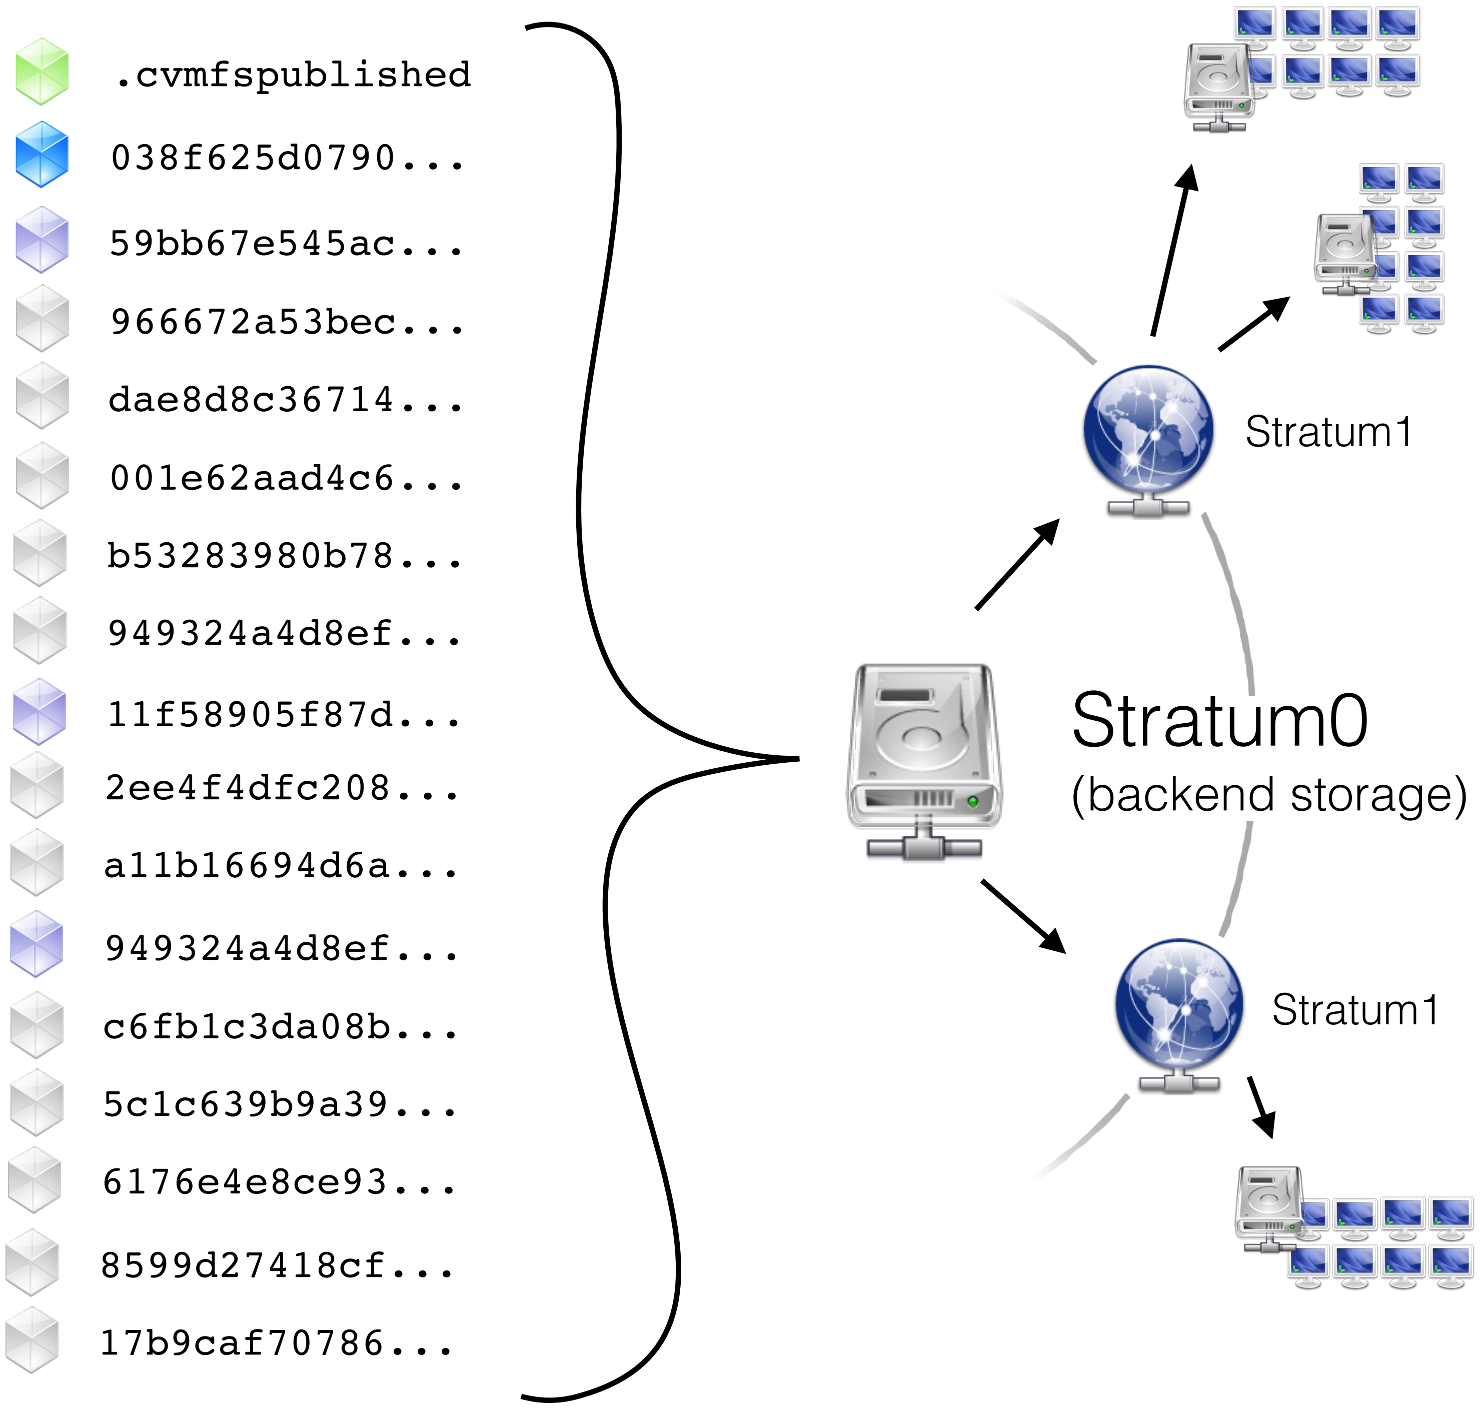

Garbage Collection

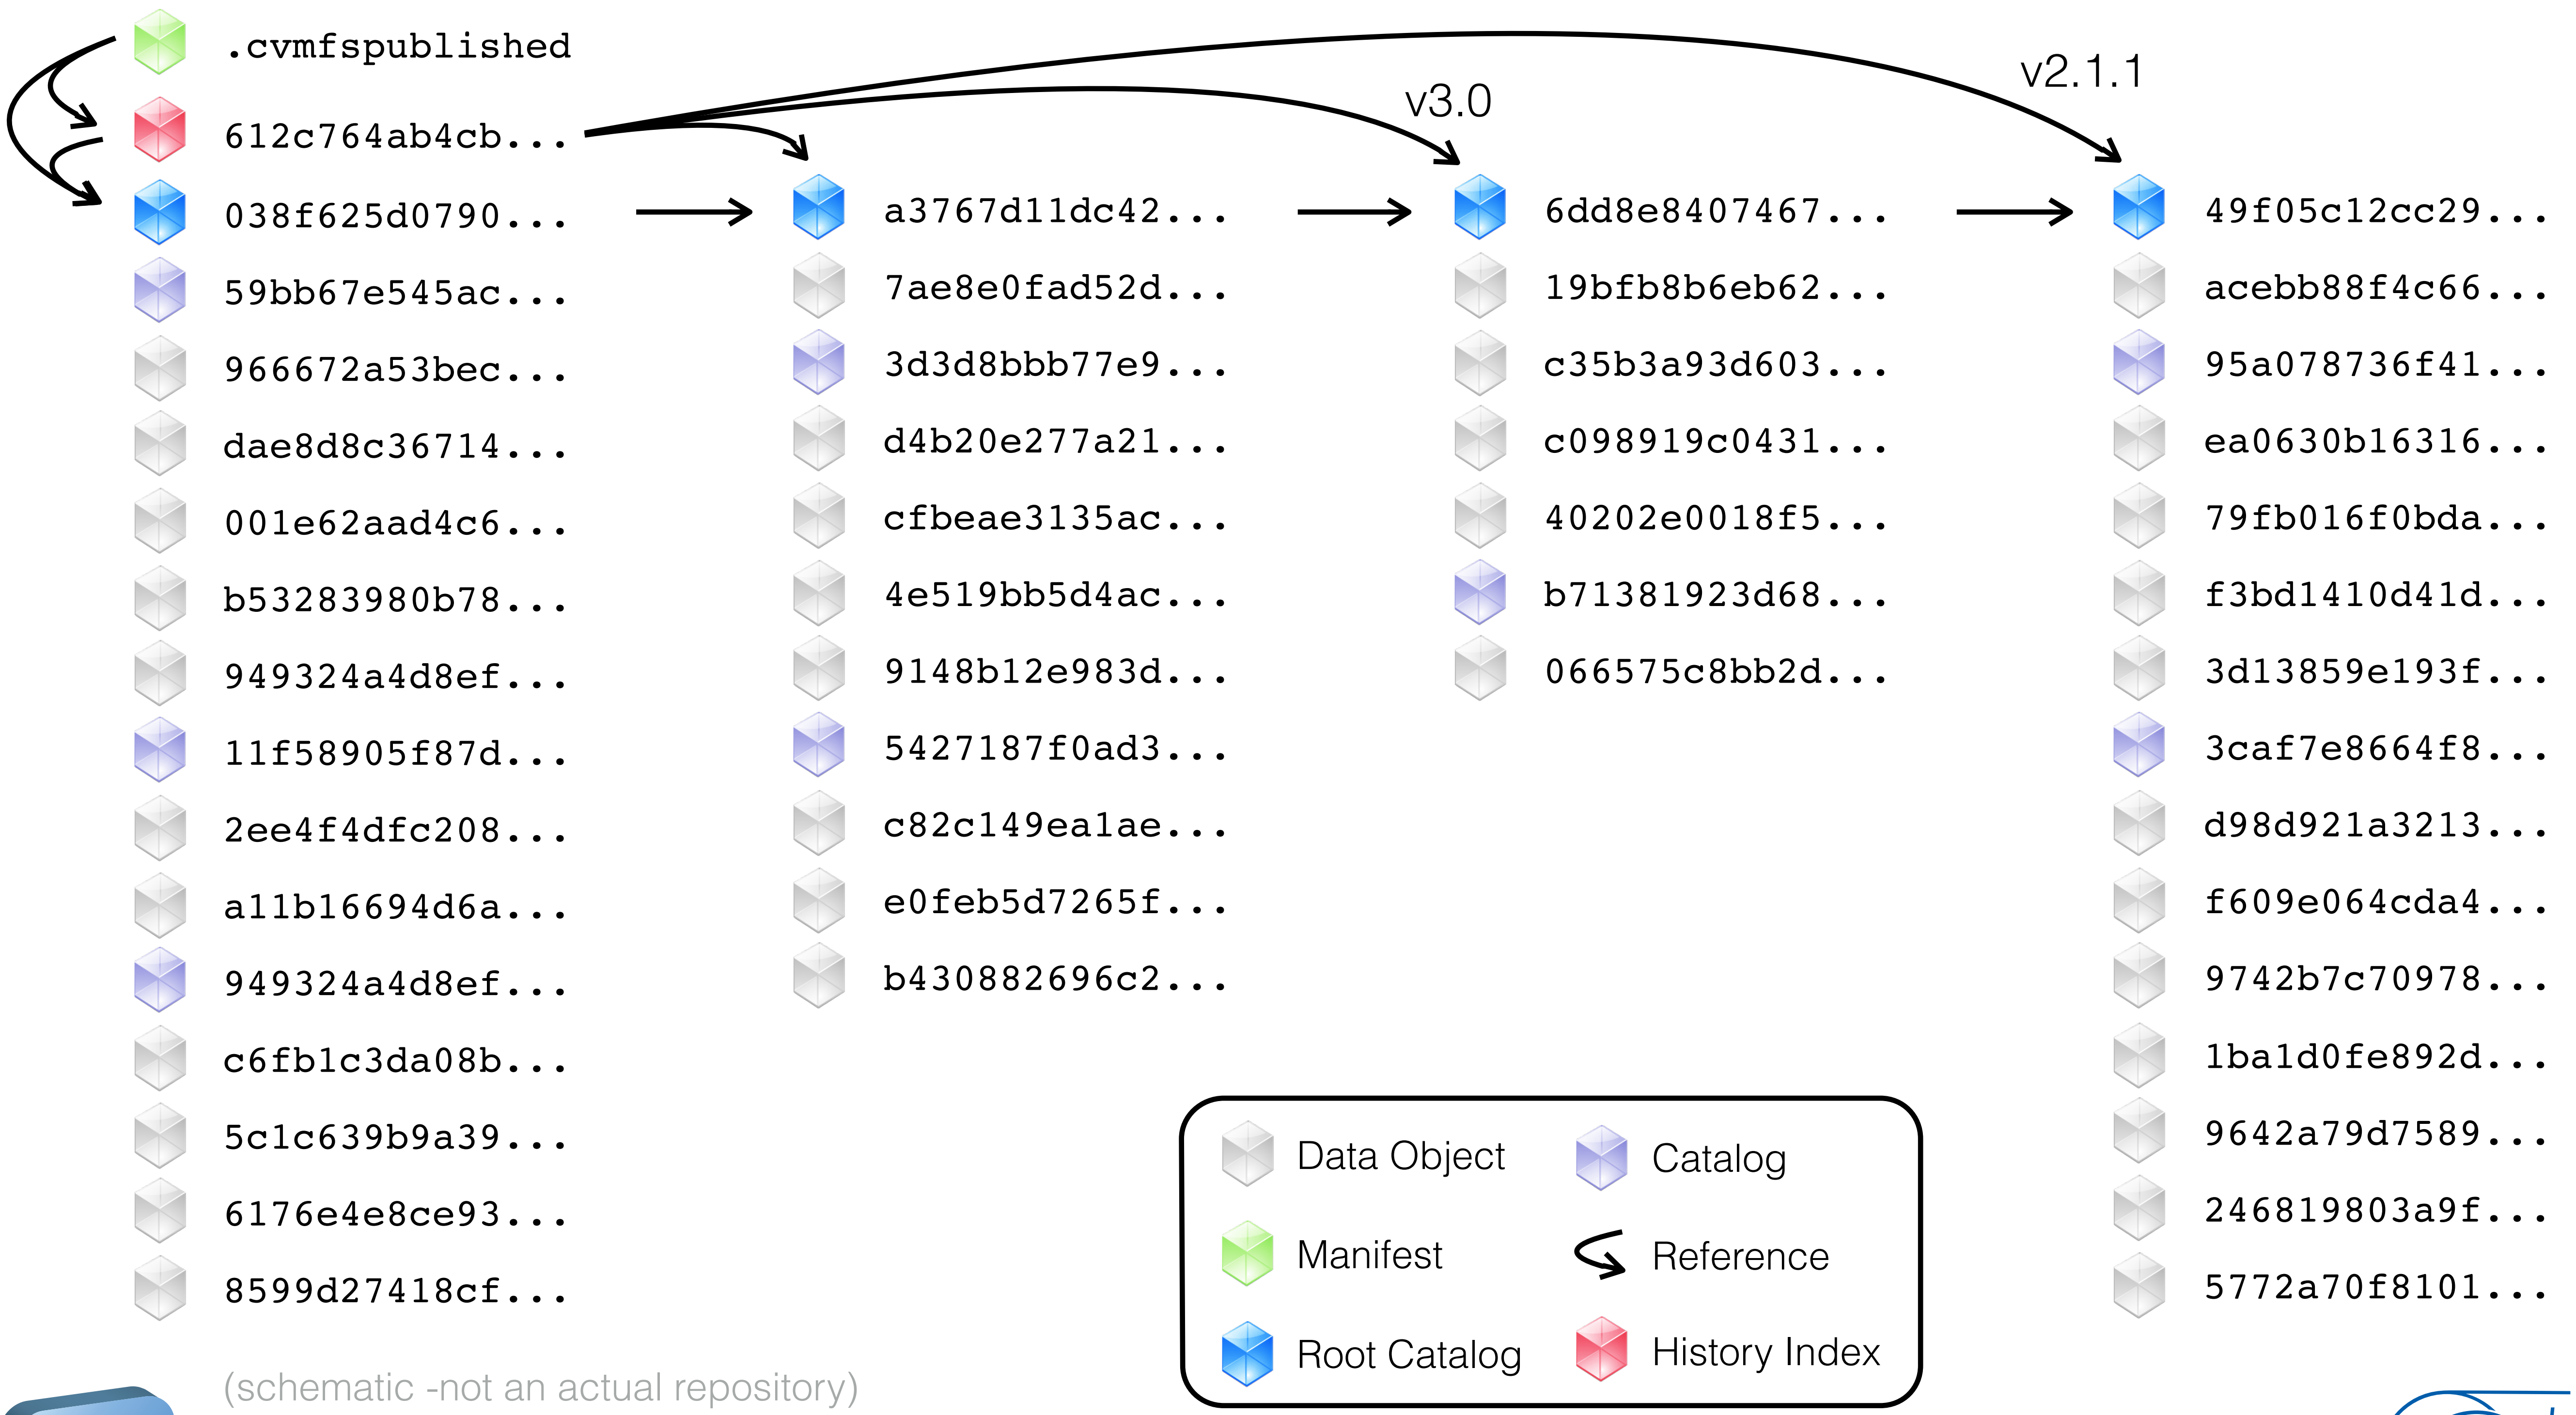

.cvmfspublished
612c764ab4cb...
038f625d0790...
59bb67e545ac...
966672a53bec...
dae8d8c36714...
001e62aad4c6...
b53283980b78...
949324a4d8ef...
11f58905f87d...
2ee4f4dfc208...
a11b16694d6a...
949324a4d8ef...
c6fb1c3da08b...
5c1c639b9a39...
6176e4e8ce93...
8599d27418cf...

a3767d11dc42...
7ae8e0fad52d...
3d3d8bbb77e9...
d4b20e277a21...
cfbeae3135ac...
4e519bb5d4ac...
9148b12e983d...
5427187f0ad3...
c82c149ea1ae...
e0feb5d7265f...
b430882696c2...

v3.0

6dd8e8407467...
19bfb8b6eb62...
c35b3a93d603...
c098919c0431...
40202e0018f5...
b71381923d68...
066575c8bb2d...

v2.1.1

49f05c12cc29...
acebb88f4c66...
95a078736f41...
ea0630b16316...
79fb016f0bda...
f3bd1410d41d...
3d13859e193f...
3caf7e8664f8...
d98d921a3213...
f609e064cda4...
9742b7c70978...
1ba1d0fe892d...
9642a79d7589...
246819803a9f...
5772a70f8101...

Data Object
Catalog
Manifest
Reference
Root Catalog
History Index

(schematic -not an actual repository)

Garbage Collection

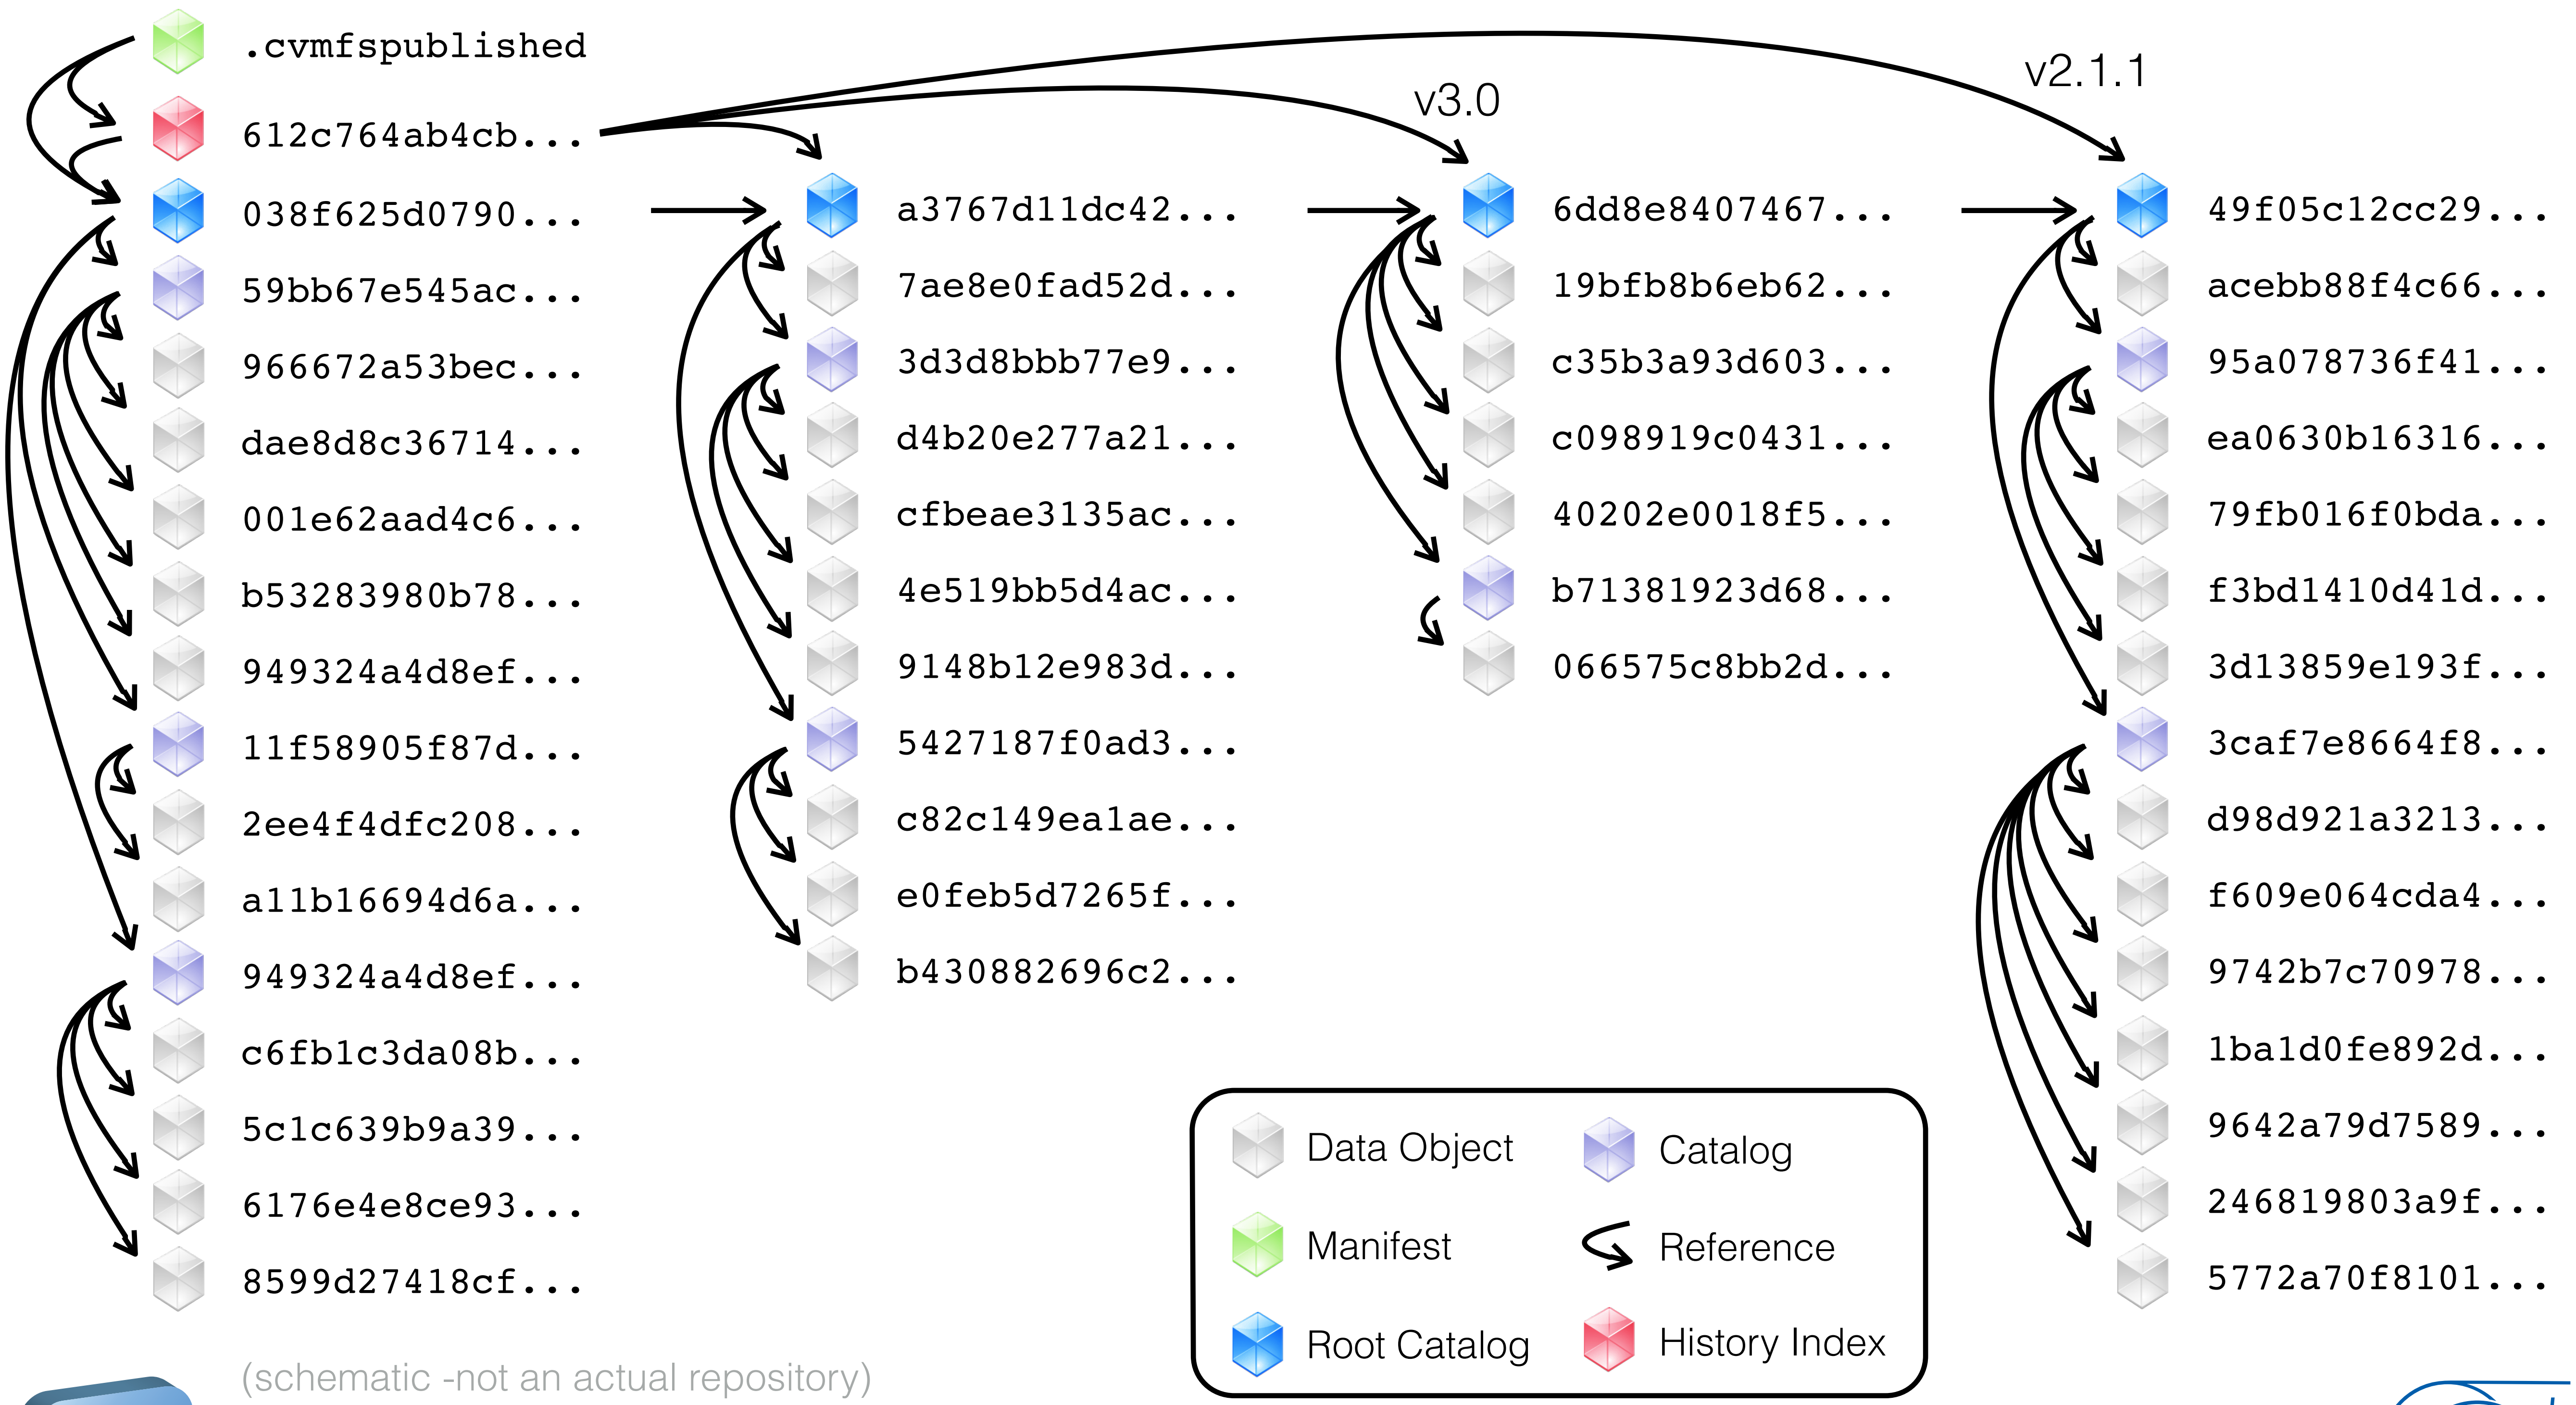

.cvmfspublished
612c764ab4cb...
038f625d0790...
59bb67e545ac...
966672a53bec...
dae8d8c36714...
001e62aad4c6...
b53283980b78...
949324a4d8ef...
11f58905f87d...
2ee4f4dfc208...
a11b16694d6a...
949324a4d8ef...
c6fb1c3da08b...
5c1c639b9a39...
6176e4e8ce93...
8599d27418cf...

a3767d11dc42...
7ae8e0fad52d...
3d3d8bbb77e9...
d4b20e277a21...
cfbeae3135ac...
4e519bb5d4ac...
9148b12e983d...
5427187f0ad3...
c82c149ea1ae...
e0feb5d7265f...
b430882696c2...

v3.0

6dd8e8407467...
19bfb8b6eb62...
c35b3a93d603...
c098919c0431...
40202e0018f5...
b71381923d68...
066575c8bb2d...

v2.1.1

49f05c12cc29...
acebb88f4c66...
95a078736f41...
ea0630b16316...
79fb016f0bda...
f3bd1410d41d...
3d13859e193f...
3caf7e8664f8...
d98d921a3213...
f609e064cda4...
9742b7c70978...
1ba1d0fe892d...
9642a79d7589...
246819803a9f...
5772a70f8101...

(schematic -not an actual repository)

Garbage Collection

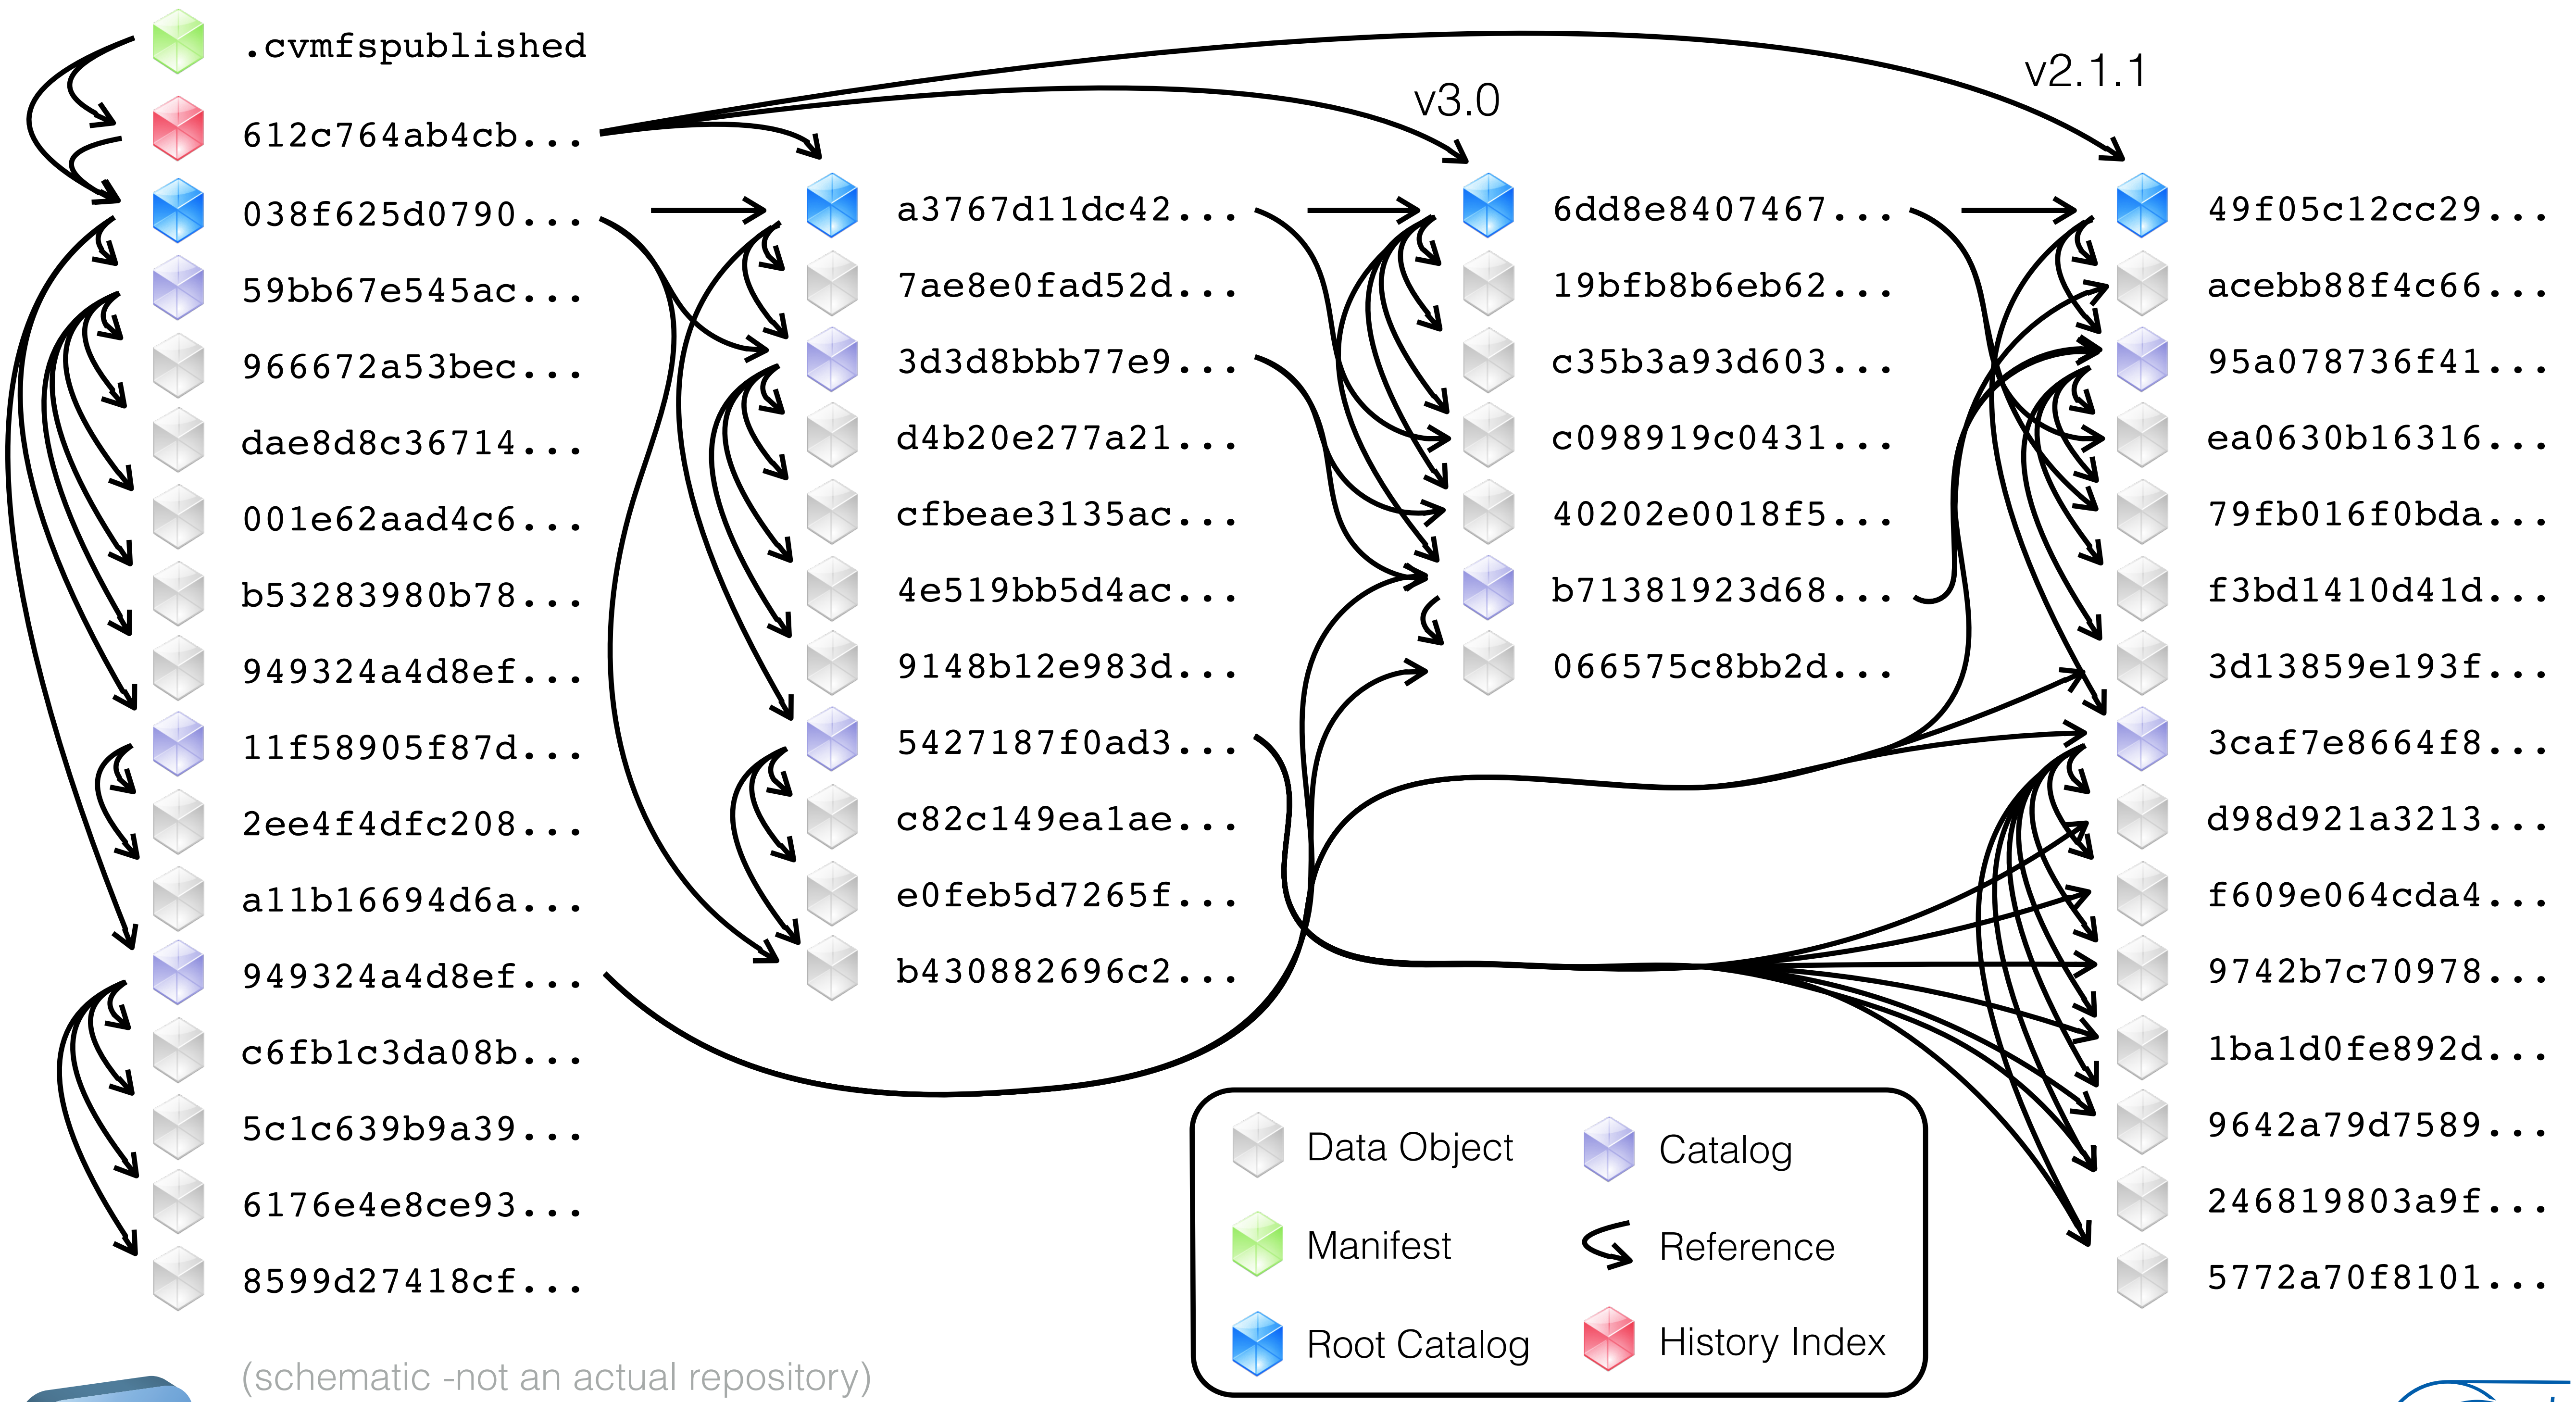

.cvmfspublished
612c764ab4cb...
038f625d0790...
59bb67e545ac...
966672a53bec...
dae8d8c36714...
001e62aad4c6...
b53283980b78...
949324a4d8ef...
11f58905f87d...
2ee4f4dfc208...
a11b16694d6a...
949324a4d8ef...
c6fb1c3da08b...
5c1c639b9a39...
6176e4e8ce93...
8599d27418cf...

a3767d11dc42...
7ae8e0fad52d...
3d3d8bbb77e9...
d4b20e277a21...
cfbeae3135ac...
4e519bb5d4ac...
9148b12e983d...
5427187f0ad3...
c82c149ea1ae...
e0feb5d7265f...
b430882696c2...

v3.0

6dd8e8407467...
19bfb8b6eb62...
c35b3a93d603...
c098919c0431...
40202e0018f5...
b71381923d68...
066575c8bb2d...

v2.1.1

49f05c12cc29...
acebb88f4c66...
95a078736f41...
ea0630b16316...
79fb016f0bda...
f3bd1410d41d...
3d13859e193f...
3caf7e8664f8...
d98d921a3213...
f609e064cda4...
9742b7c70978...
1ba1d0fe892d...
9642a79d7589...
246819803a9f...
5772a70f8101...

Data Object

Catalog

Manifest

Reference

Root Catalog

History Index

(schematic -not an actual repository)

Garbage Collection

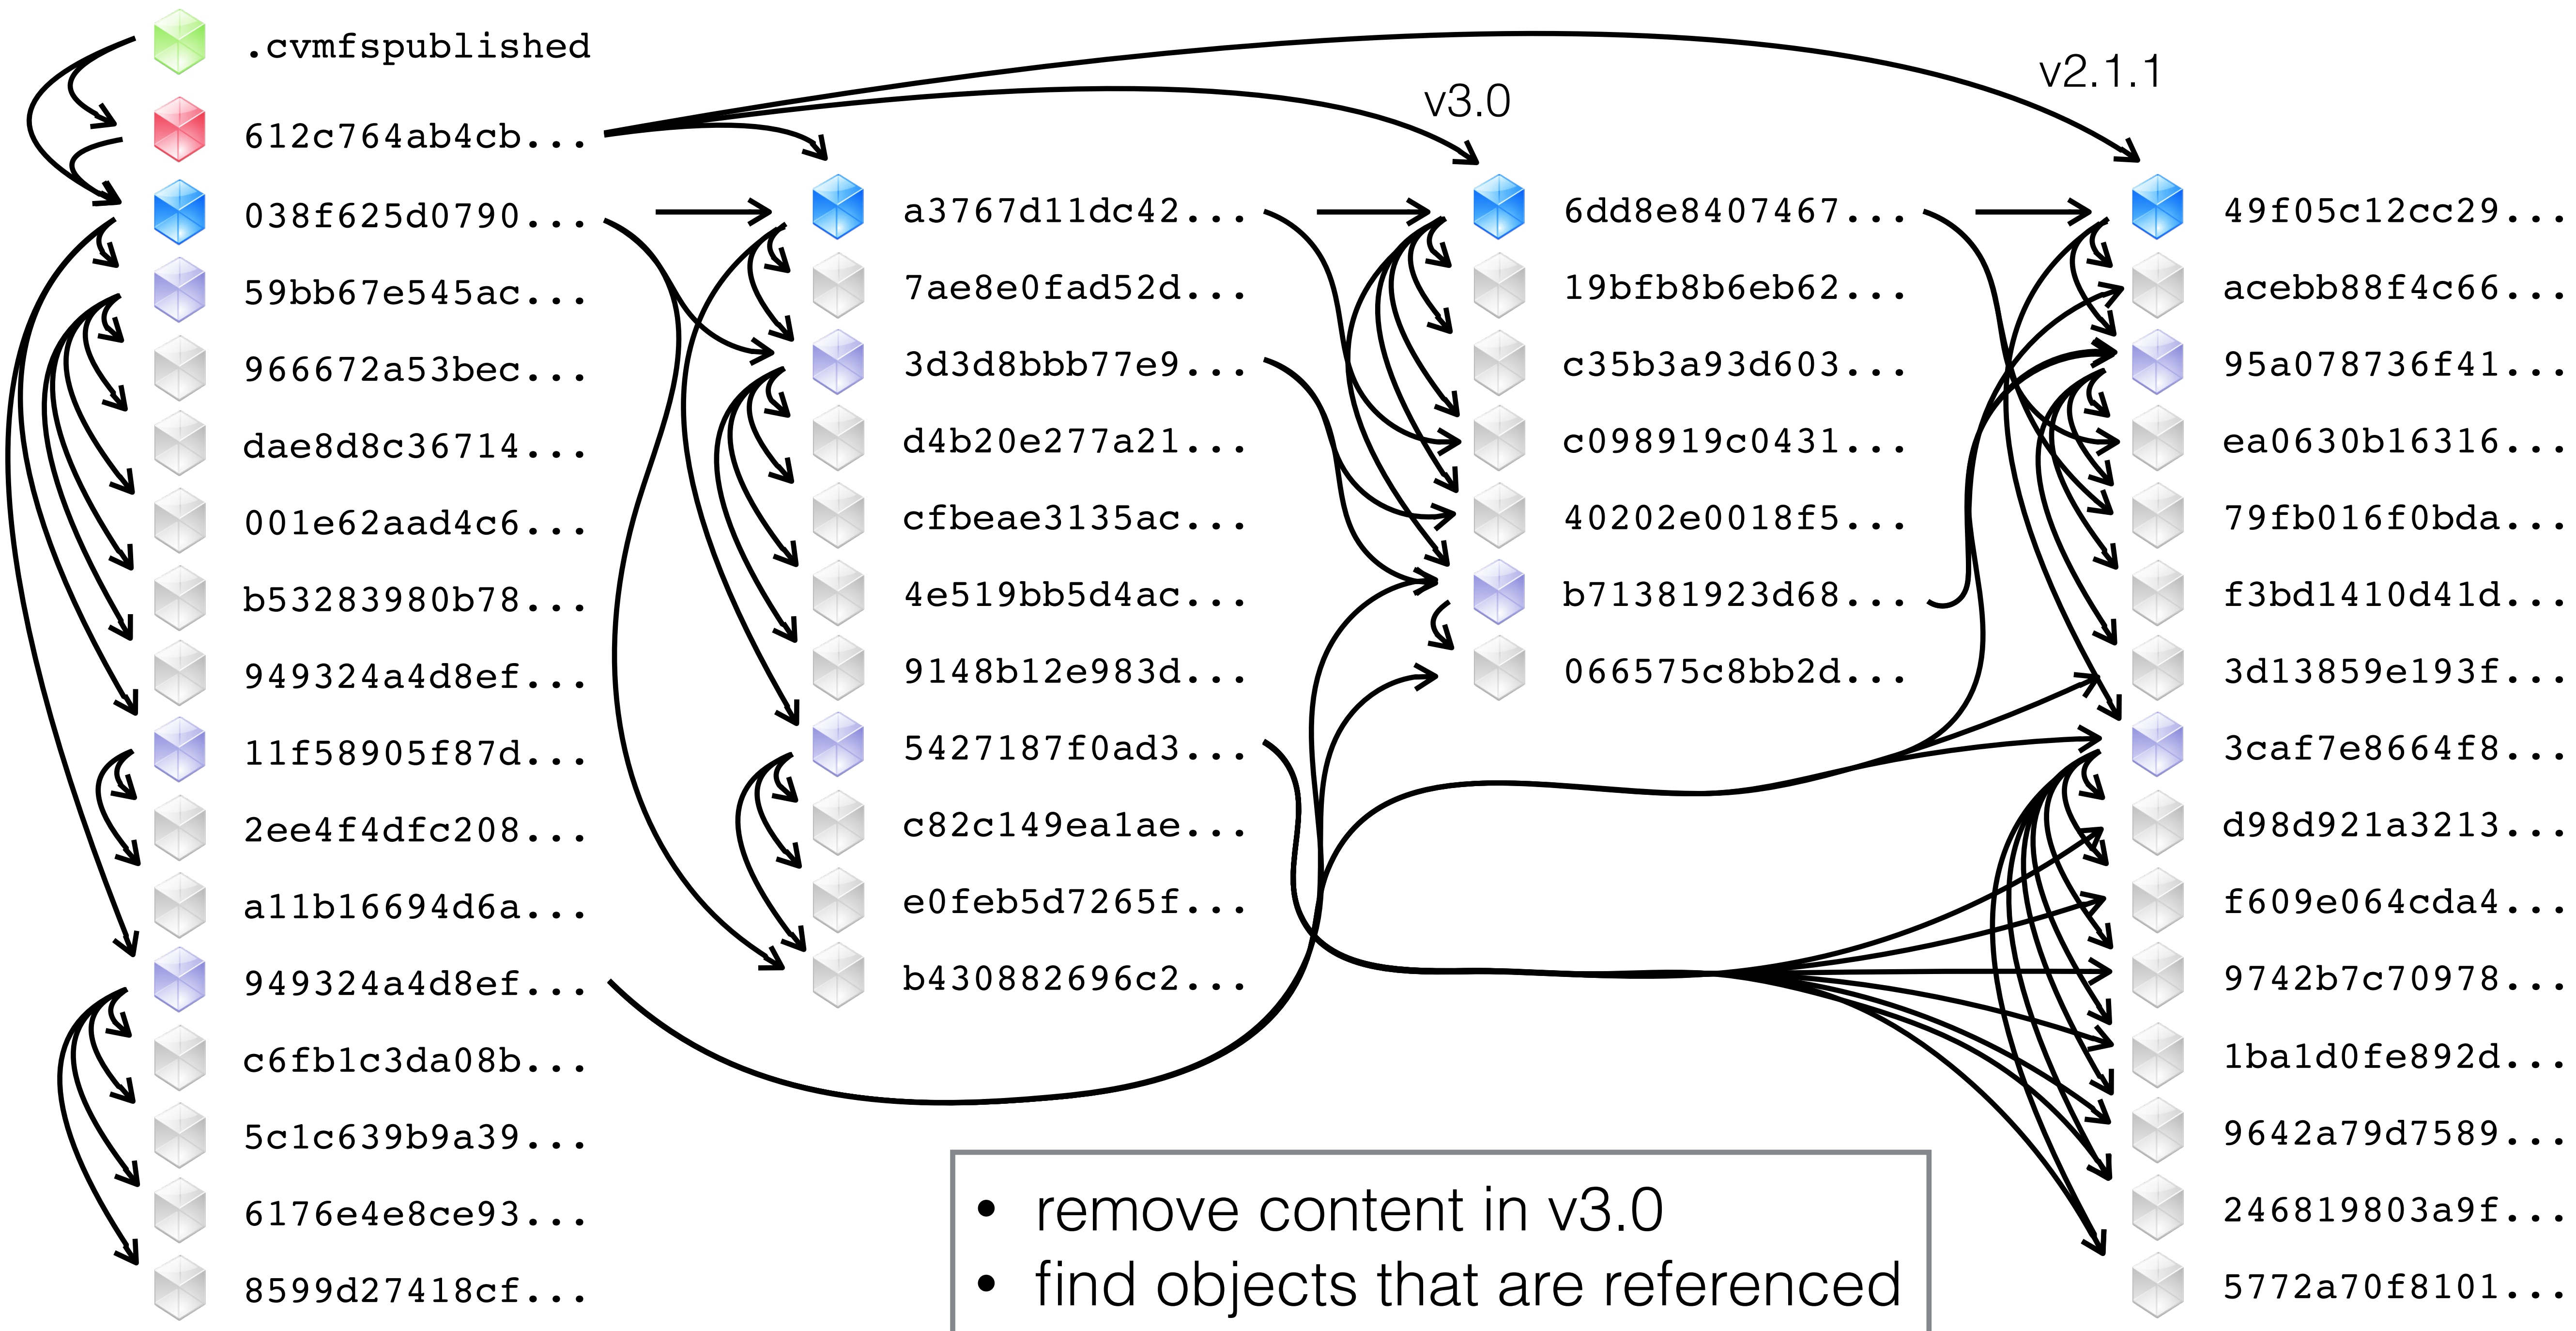

.cvmfspublished
612c764ab4cb...
038f625d0790...
59bb67e545ac...
966672a53bec...
dae8d8c36714...
001e62aad4c6...
b53283980b78...
949324a4d8ef...
11f58905f87d...
2ee4f4dfc208...
a11b16694d6a...
949324a4d8ef...
c6fb1c3da08b...
5c1c639b9a39...
6176e4e8ce93...
8599d27418cf...

a3767d11dc42...
7ae8e0fad52d...
3d3d8bbb77e9...
d4b20e277a21...
cfbeae3135ac...
4e519bb5d4ac...
9148b12e983d...
5427187f0ad3...
c82c149ea1ae...
e0feb5d7265f...
b430882696c2...

v3.0

6dd8e8407467...
19bfb8b6eb62...
c35b3a93d603...
c098919c0431...
40202e0018f5...
b71381923d68...
066575c8bb2d...

v2.1.1

49f05c12cc29...
acebb88f4c66...
95a078736f41...
ea0630b16316...
79fb016f0bda...
f3bd1410d41d...
3d13859e193f...
3caf7e8664f8...
d98d921a3213...
f609e064cda4...
9742b7c70978...
1ba1d0fe892d...
9642a79d7589...
246819803a9f...
5772a70f8101...

- remove content in v3.0
- find objects that are referenced nowhere else

(schematic -not an actual repository)

Garbage Collection

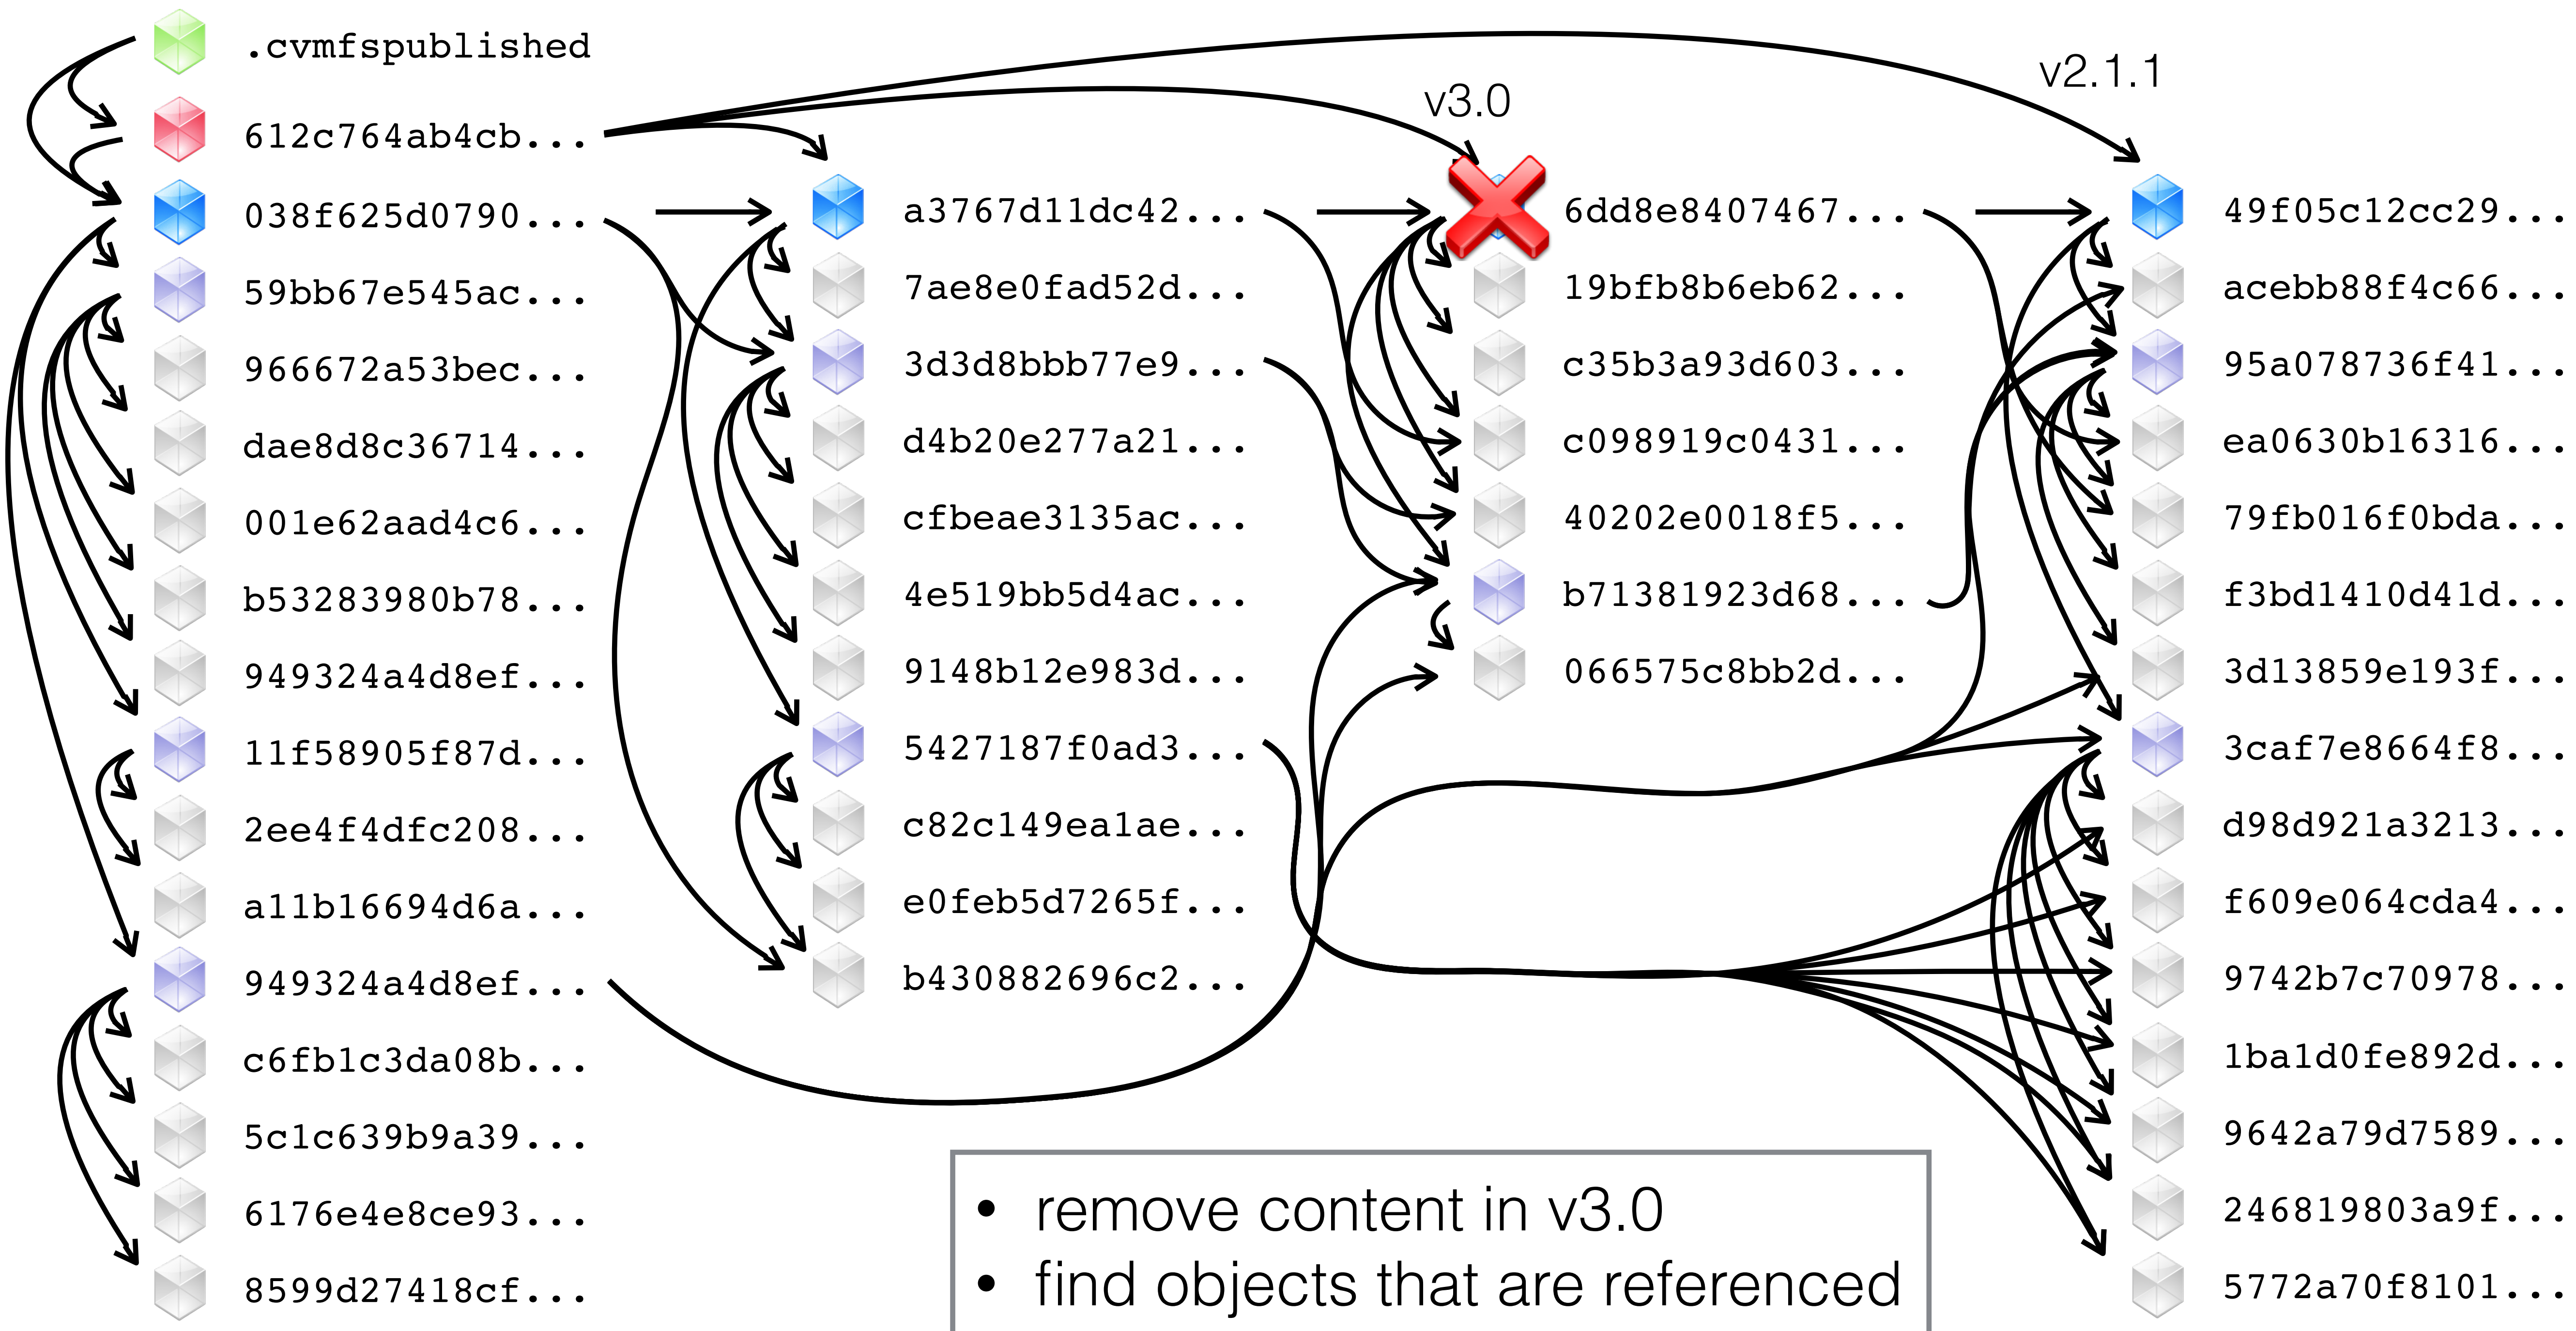

- remove content in v3.0
- find objects that are referenced nowhere else

(schematic -not an actual repository)

Garbage Collection

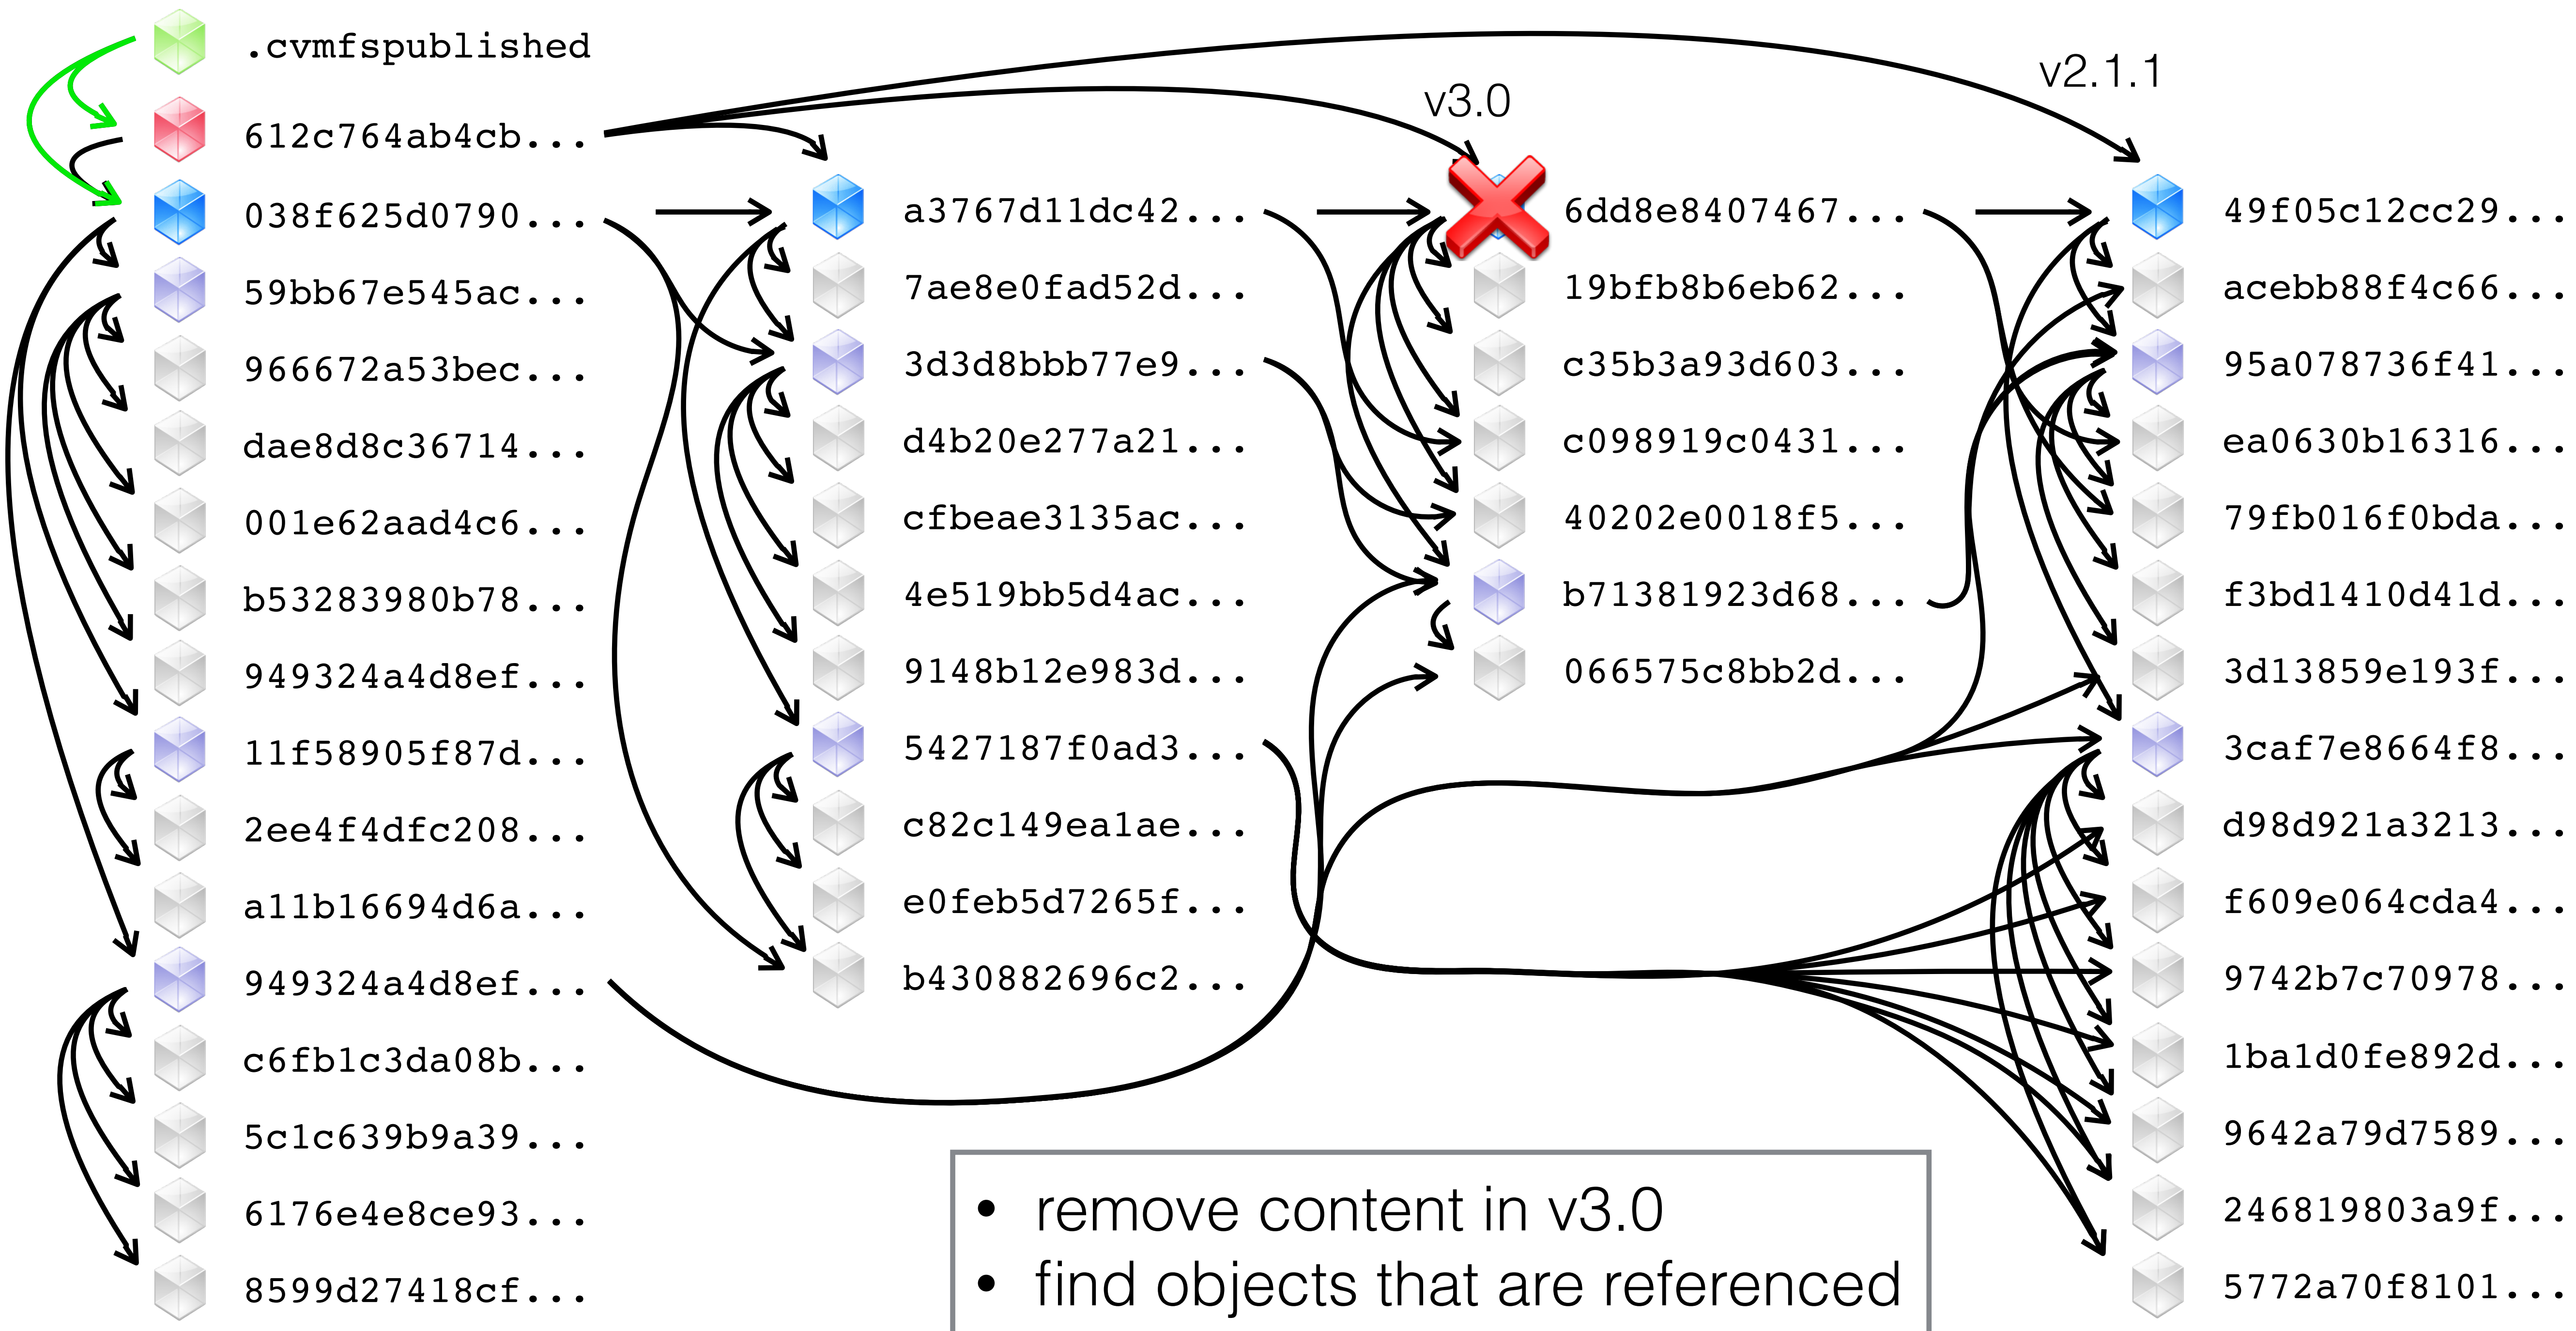

.cvmfspublished
612c764ab4cb...
038f625d0790...
59bb67e545ac...
966672a53bec...
dae8d8c36714...
001e62aad4c6...
b53283980b78...
949324a4d8ef...
11f58905f87d...
2ee4f4dfc208...
a11b16694d6a...
949324a4d8ef...
c6fb1c3da08b...
5c1c639b9a39...
6176e4e8ce93...
8599d27418cf...

a3767d11dc42...
7ae8e0fad52d...
3d3d8bbb77e9...
d4b20e277a21...
cfbeae3135ac...
4e519bb5d4ac...
9148b12e983d...
5427187f0ad3...
c82c149ea1ae...
e0feb5d7265f...
b430882696c2...

v3.0
6dd8e8407467...
19bfb8b6eb62...
c35b3a93d603...
c098919c0431...
40202e0018f5...
b71381923d68...
066575c8bb2d...

v2.1.1
49f05c12cc29...
acebb88f4c66...
95a078736f41...
ea0630b16316...
79fb016f0bda...
f3bd1410d41d...
3d13859e193f...
3caf7e8664f8...
d98d921a3213...
f609e064cda4...
9742b7c70978...
1ba1d0fe892d...
9642a79d7589...
246819803a9f...
5772a70f8101...

- remove content in v3.0
- find objects that are referenced nowhere else

(schematic -not an actual repository)

Garbage Collection

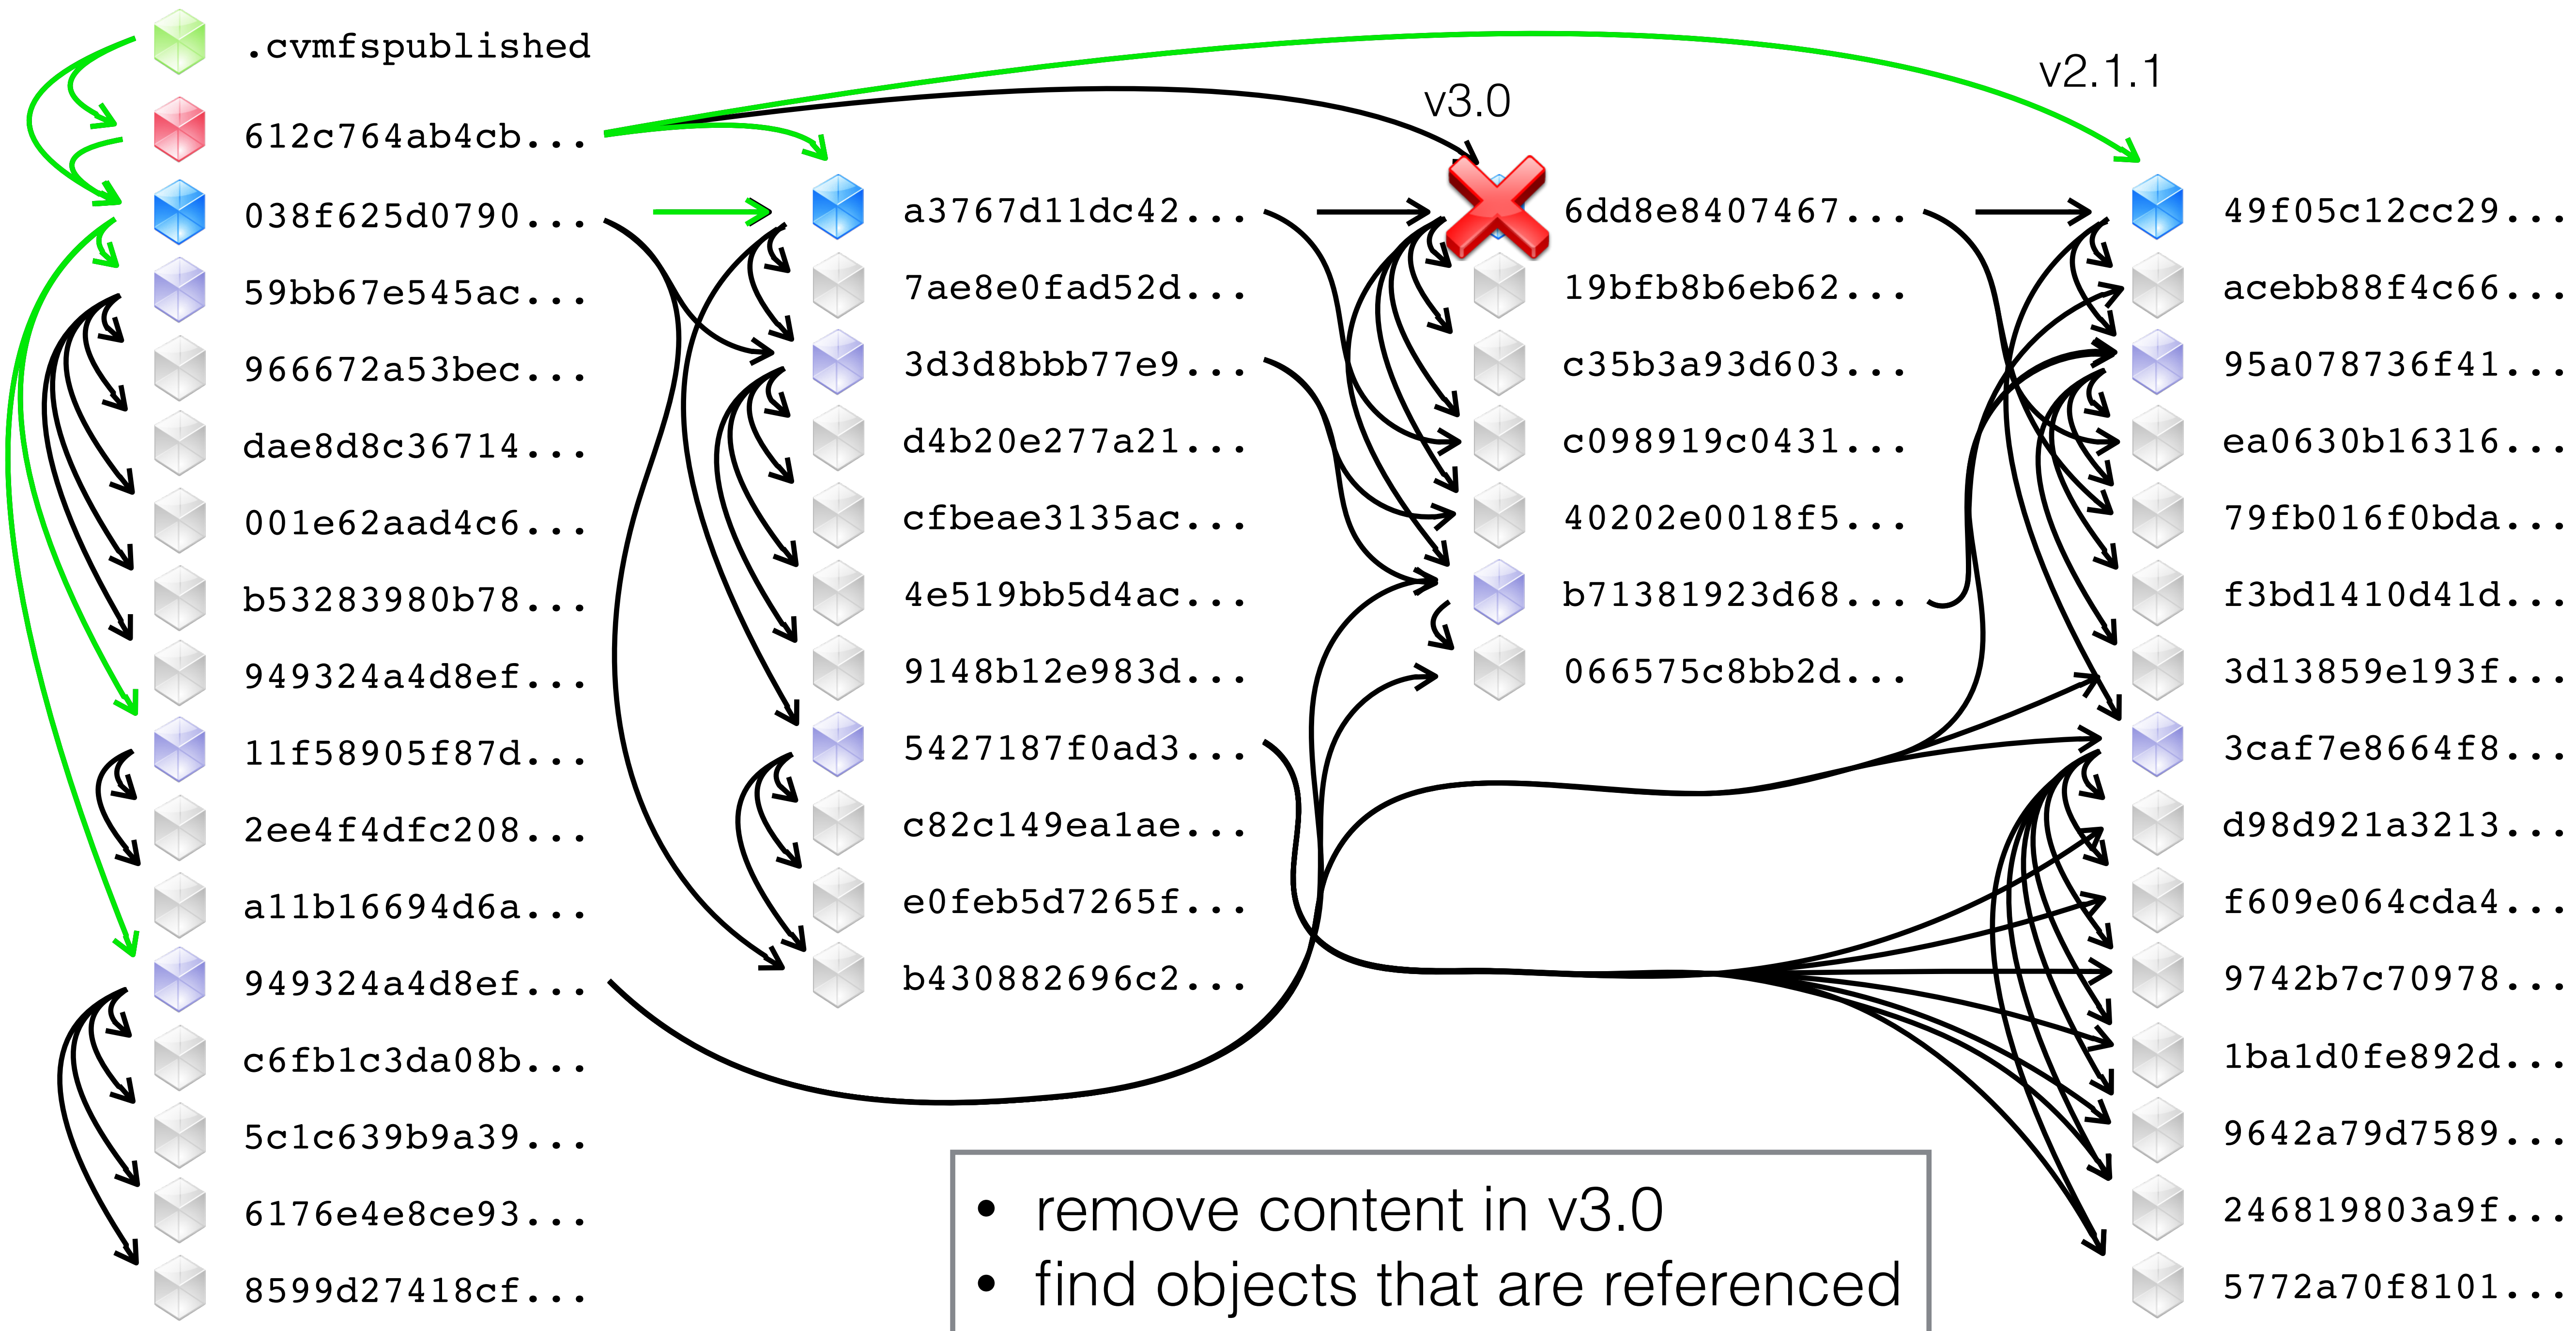

.cvmfspublished
612c764ab4cb...
038f625d0790...
59bb67e545ac...
966672a53bec...
dae8d8c36714...
001e62aad4c6...
b53283980b78...
949324a4d8ef...
11f58905f87d...
2ee4f4dfc208...
a11b16694d6a...
949324a4d8ef...
c6fb1c3da08b...
5c1c639b9a39...
6176e4e8ce93...
8599d27418cf...

a3767d11dc42...
7ae8e0fad52d...
3d3d8bbb77e9...
d4b20e277a21...
cfbeae3135ac...
4e519bb5d4ac...
9148b12e983d...
5427187f0ad3...
c82c149ea1ae...
e0feb5d7265f...
b430882696c2...

v3.0

6dd8e8407467...
19bfb8b6eb62...
c35b3a93d603...
c098919c0431...
40202e0018f5...
b71381923d68...
066575c8bb2d...

v2.1.1

49f05c12cc29...
acebb88f4c66...
95a078736f41...
ea0630b16316...
79fb016f0bda...
f3bd1410d41d...
3d13859e193f...
3caf7e8664f8...
d98d921a3213...
f609e064cda4...
9742b7c70978...
1ba1d0fe892d...
9642a79d7589...
246819803a9f...
5772a70f8101...

- remove content in v3.0
- find objects that are referenced nowhere else

(schematic -not an actual repository)

Garbage Collection

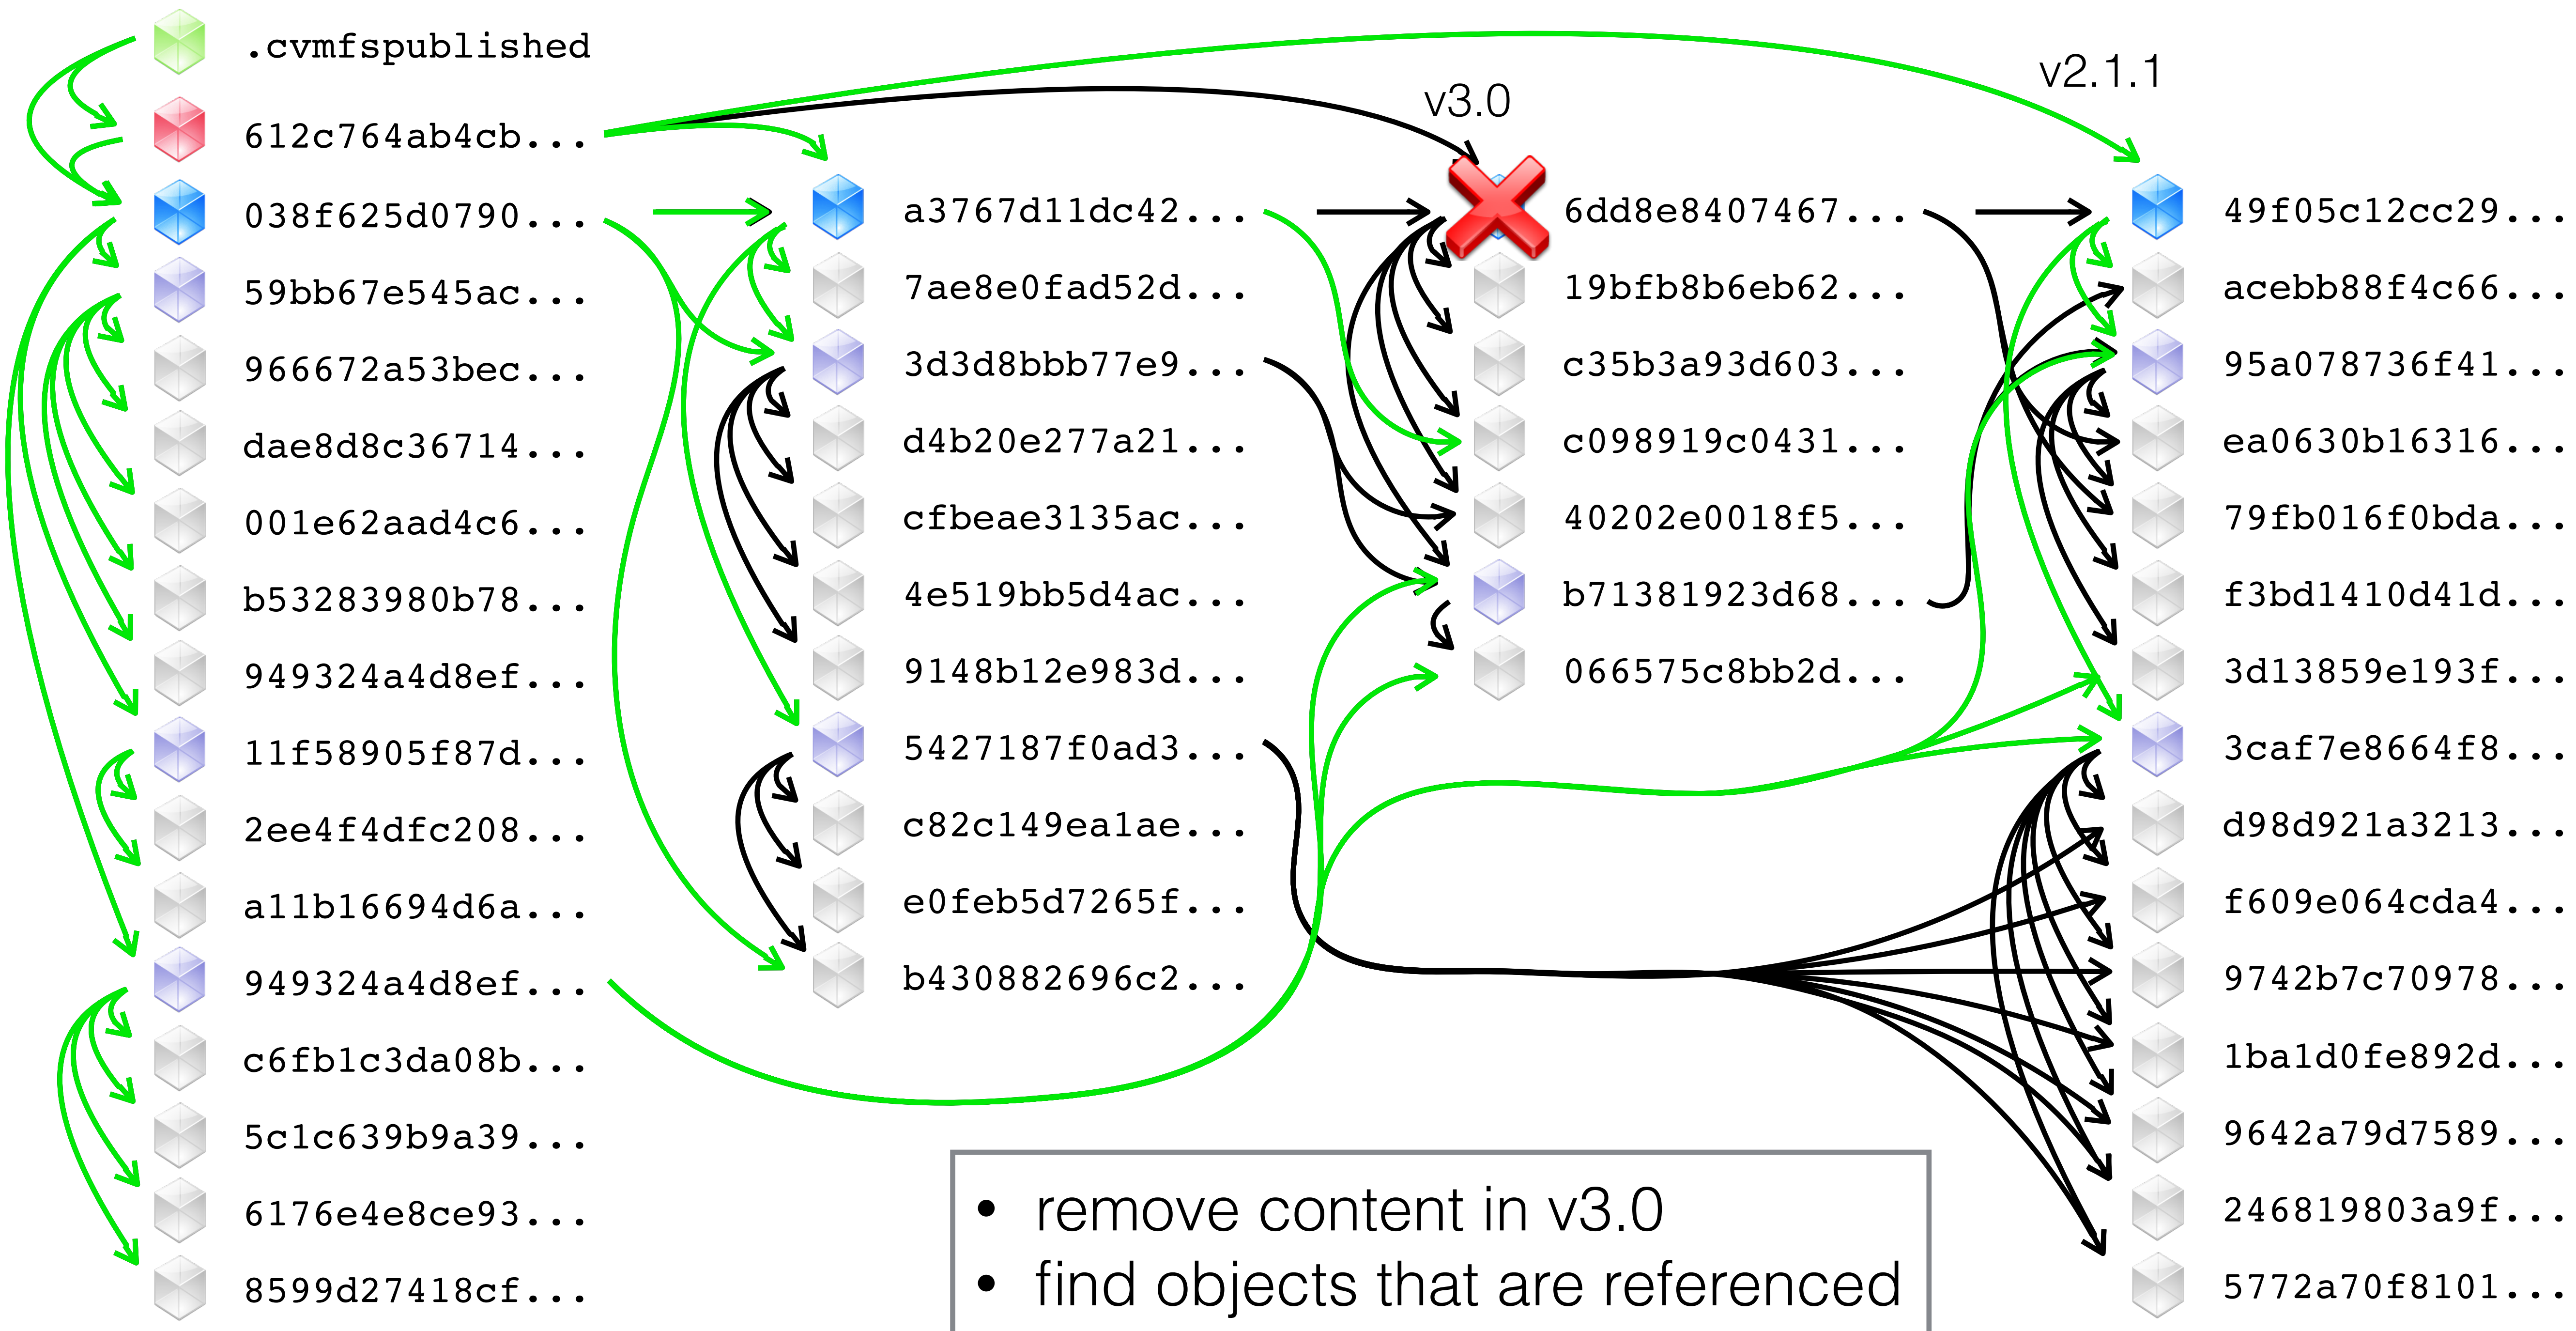

.cvmfspublished
612c764ab4cb...
038f625d0790...
59bb67e545ac...
966672a53bec...
dae8d8c36714...
001e62aad4c6...
b53283980b78...
949324a4d8ef...
11f58905f87d...
2ee4f4dfc208...
a11b16694d6a...
949324a4d8ef...
c6fb1c3da08b...
5c1c639b9a39...
6176e4e8ce93...
8599d27418cf...

a3767d11dc42...
7ae8e0fad52d...
3d3d8bbb77e9...
d4b20e277a21...
cfbeae3135ac...
4e519bb5d4ac...
9148b12e983d...
5427187f0ad3...
c82c149ea1ae...
e0feb5d7265f...
b430882696c2...

v3.0
6dd8e8407467...
19bfb8b6eb62...
c35b3a93d603...
c098919c0431...
40202e0018f5...
b71381923d68...
066575c8bb2d...

v2.1.1
49f05c12cc29...
acebb88f4c66...
95a078736f41...
ea0630b16316...
79fb016f0bda...
f3bd1410d41d...
3d13859e193f...
3caf7e8664f8...
d98d921a3213...
f609e064cda4...
9742b7c70978...
1ba1d0fe892d...
9642a79d7589...
246819803a9f...
5772a70f8101...

- remove content in v3.0
- find objects that are referenced nowhere else

(schematic -not an actual repository)

Garbage Collection

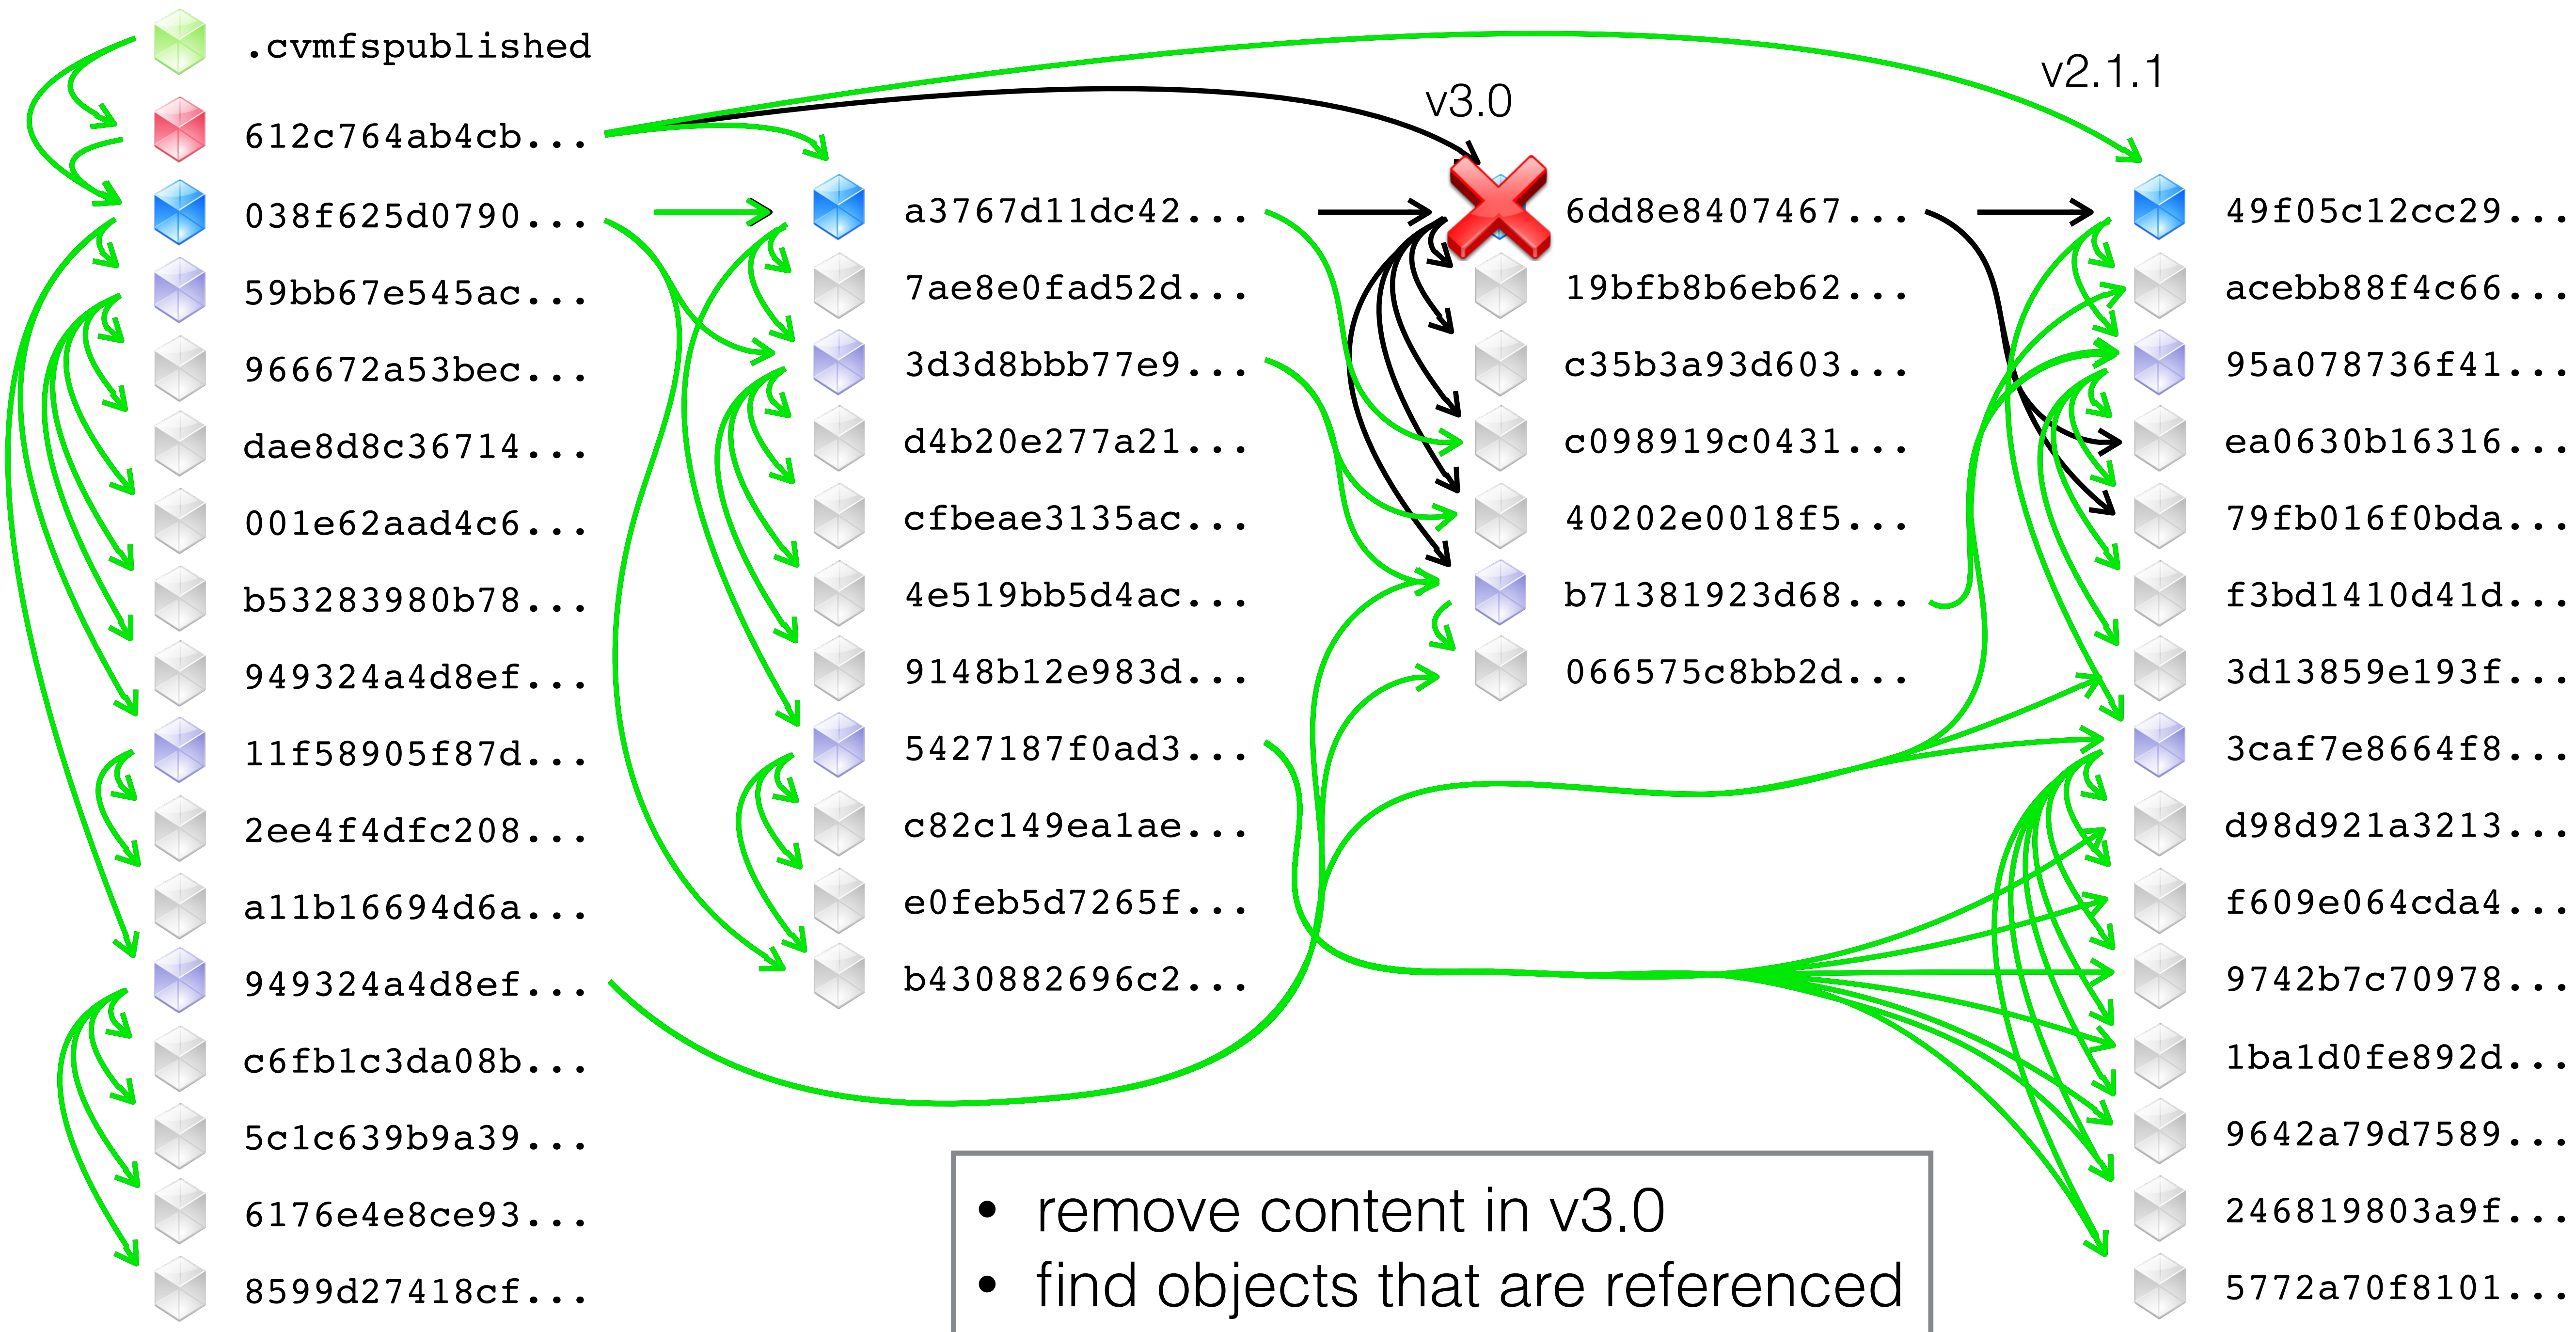

.cvmfspublished
612c764ab4cb...
038f625d0790...
59bb67e545ac...
966672a53bec...
dae8d8c36714...
001e62aad4c6...
b53283980b78...
949324a4d8ef...
11f58905f87d...
2ee4f4dfc208...
a11b16694d6a...
949324a4d8ef...
c6fb1c3da08b...
5c1c639b9a39...
6176e4e8ce93...
8599d27418cf...

a3767d11dc42...
7ae8e0fad52d...
3d3d8bbb77e9...
d4b20e277a21...
cfbeae3135ac...
4e519bb5d4ac...
9148b12e983d...
5427187f0ad3...
c82c149ea1ae...
e0feb5d7265f...
b430882696c2...

v3.0

6dd8e8407467...
19bfb8b6eb62...
c35b3a93d603...
c098919c0431...
40202e0018f5...
b71381923d68...
066575c8bb2d...

v2.1.1

49f05c12cc29...
acebb88f4c66...
95a078736f41...
ea0630b16316...
79fb016f0bda...
f3bd1410d41d...
3d13859e193f...
3caf7e8664f8...
d98d921a3213...
f609e064cda4...
9742b7c70978...
1ba1d0fe892d...
9642a79d7589...
246819803a9f...
5772a70f8101...

- remove content in v3.0
- find objects that are referenced nowhere else

(schematic -not an actual repository)

Garbage Collection

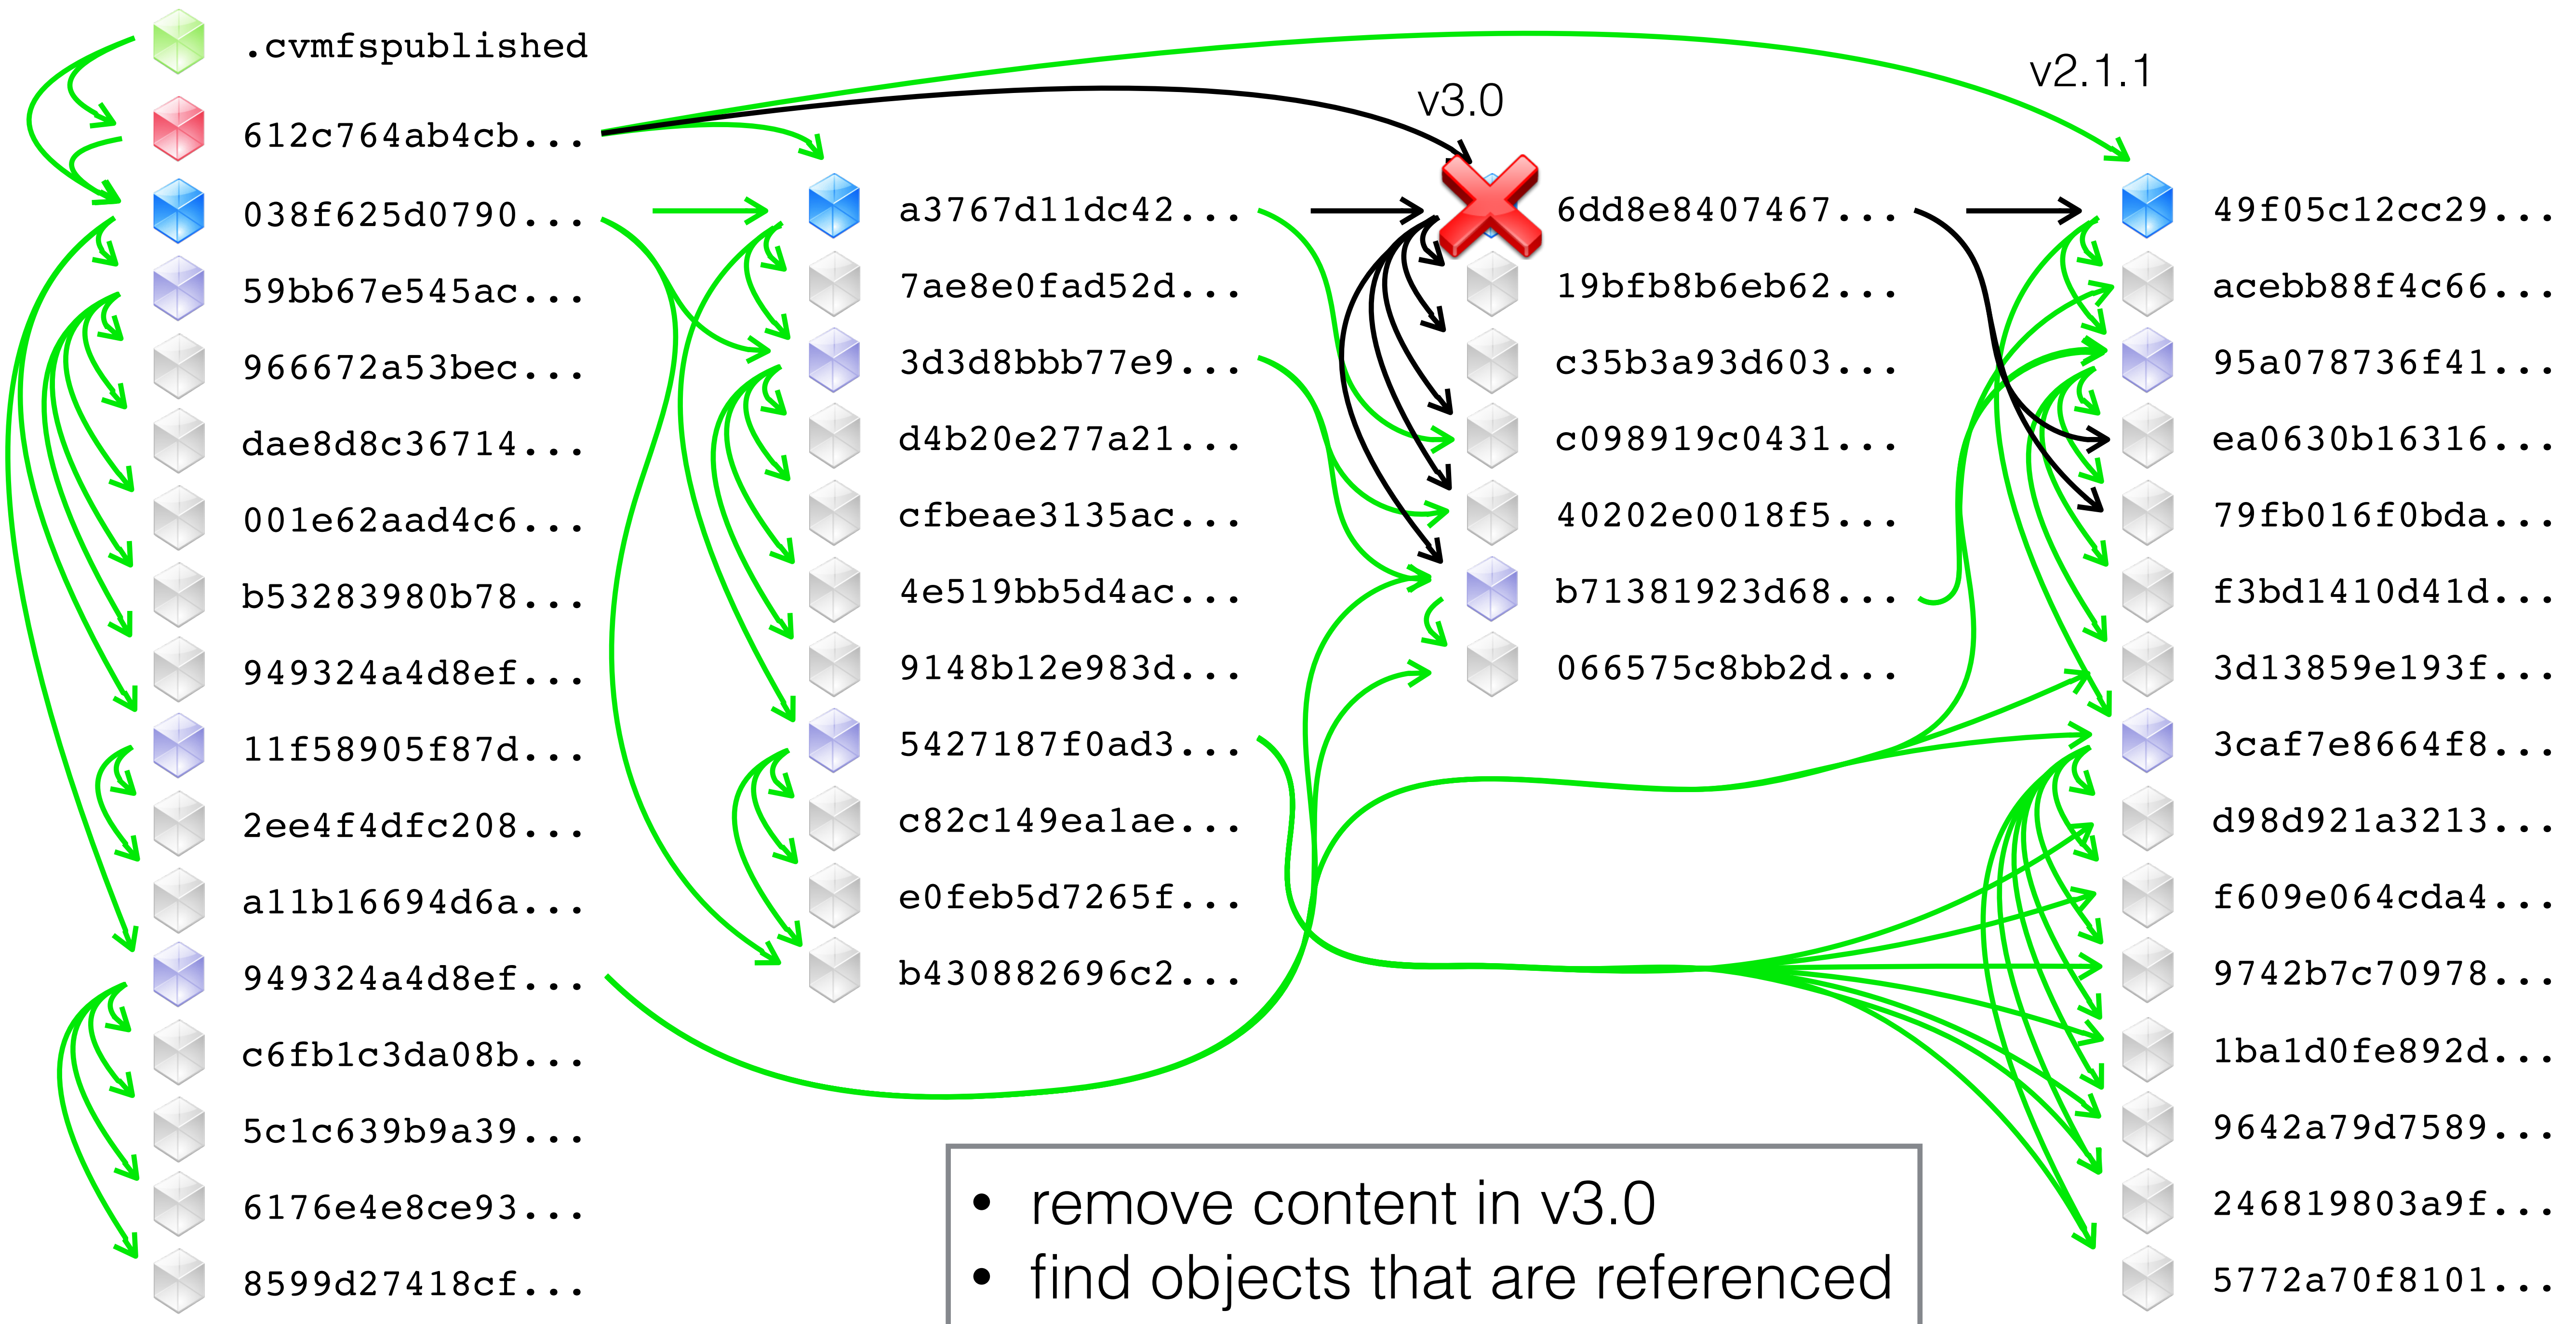

.cvmfspublished
612c764ab4cb...
038f625d0790...
59bb67e545ac...
966672a53bec...
dae8d8c36714...
001e62aad4c6...
b53283980b78...
949324a4d8ef...
11f58905f87d...
2ee4f4dfc208...
a11b16694d6a...
949324a4d8ef...
c6fb1c3da08b...
5c1c639b9a39...
6176e4e8ce93...
8599d27418cf...

a3767d11dc42...
7ae8e0fad52d...
3d3d8bbb77e9...
d4b20e277a21...
cfbeae3135ac...
4e519bb5d4ac...
9148b12e983d...
5427187f0ad3...
c82c149ea1ae...
e0feb5d7265f...
b430882696c2...

v3.0

6dd8e8407467...
19bfb8b6eb62...
c35b3a93d603...
c098919c0431...
40202e0018f5...
b71381923d68...
066575c8bb2d...

v2.1.1

49f05c12cc29...
acebb88f4c66...
95a078736f41...
ea0630b16316...
79fb016f0bda...
f3bd1410d41d...
3d13859e193f...
3caf7e8664f8...
d98d921a3213...
f609e064cda4...
9742b7c70978...
1ba1d0fe892d...
9642a79d7589...
246819803a9f...
5772a70f8101...

- remove content in v3.0
- find objects that are referenced nowhere else

(schematic -not an actual repository)

Garbage Collection

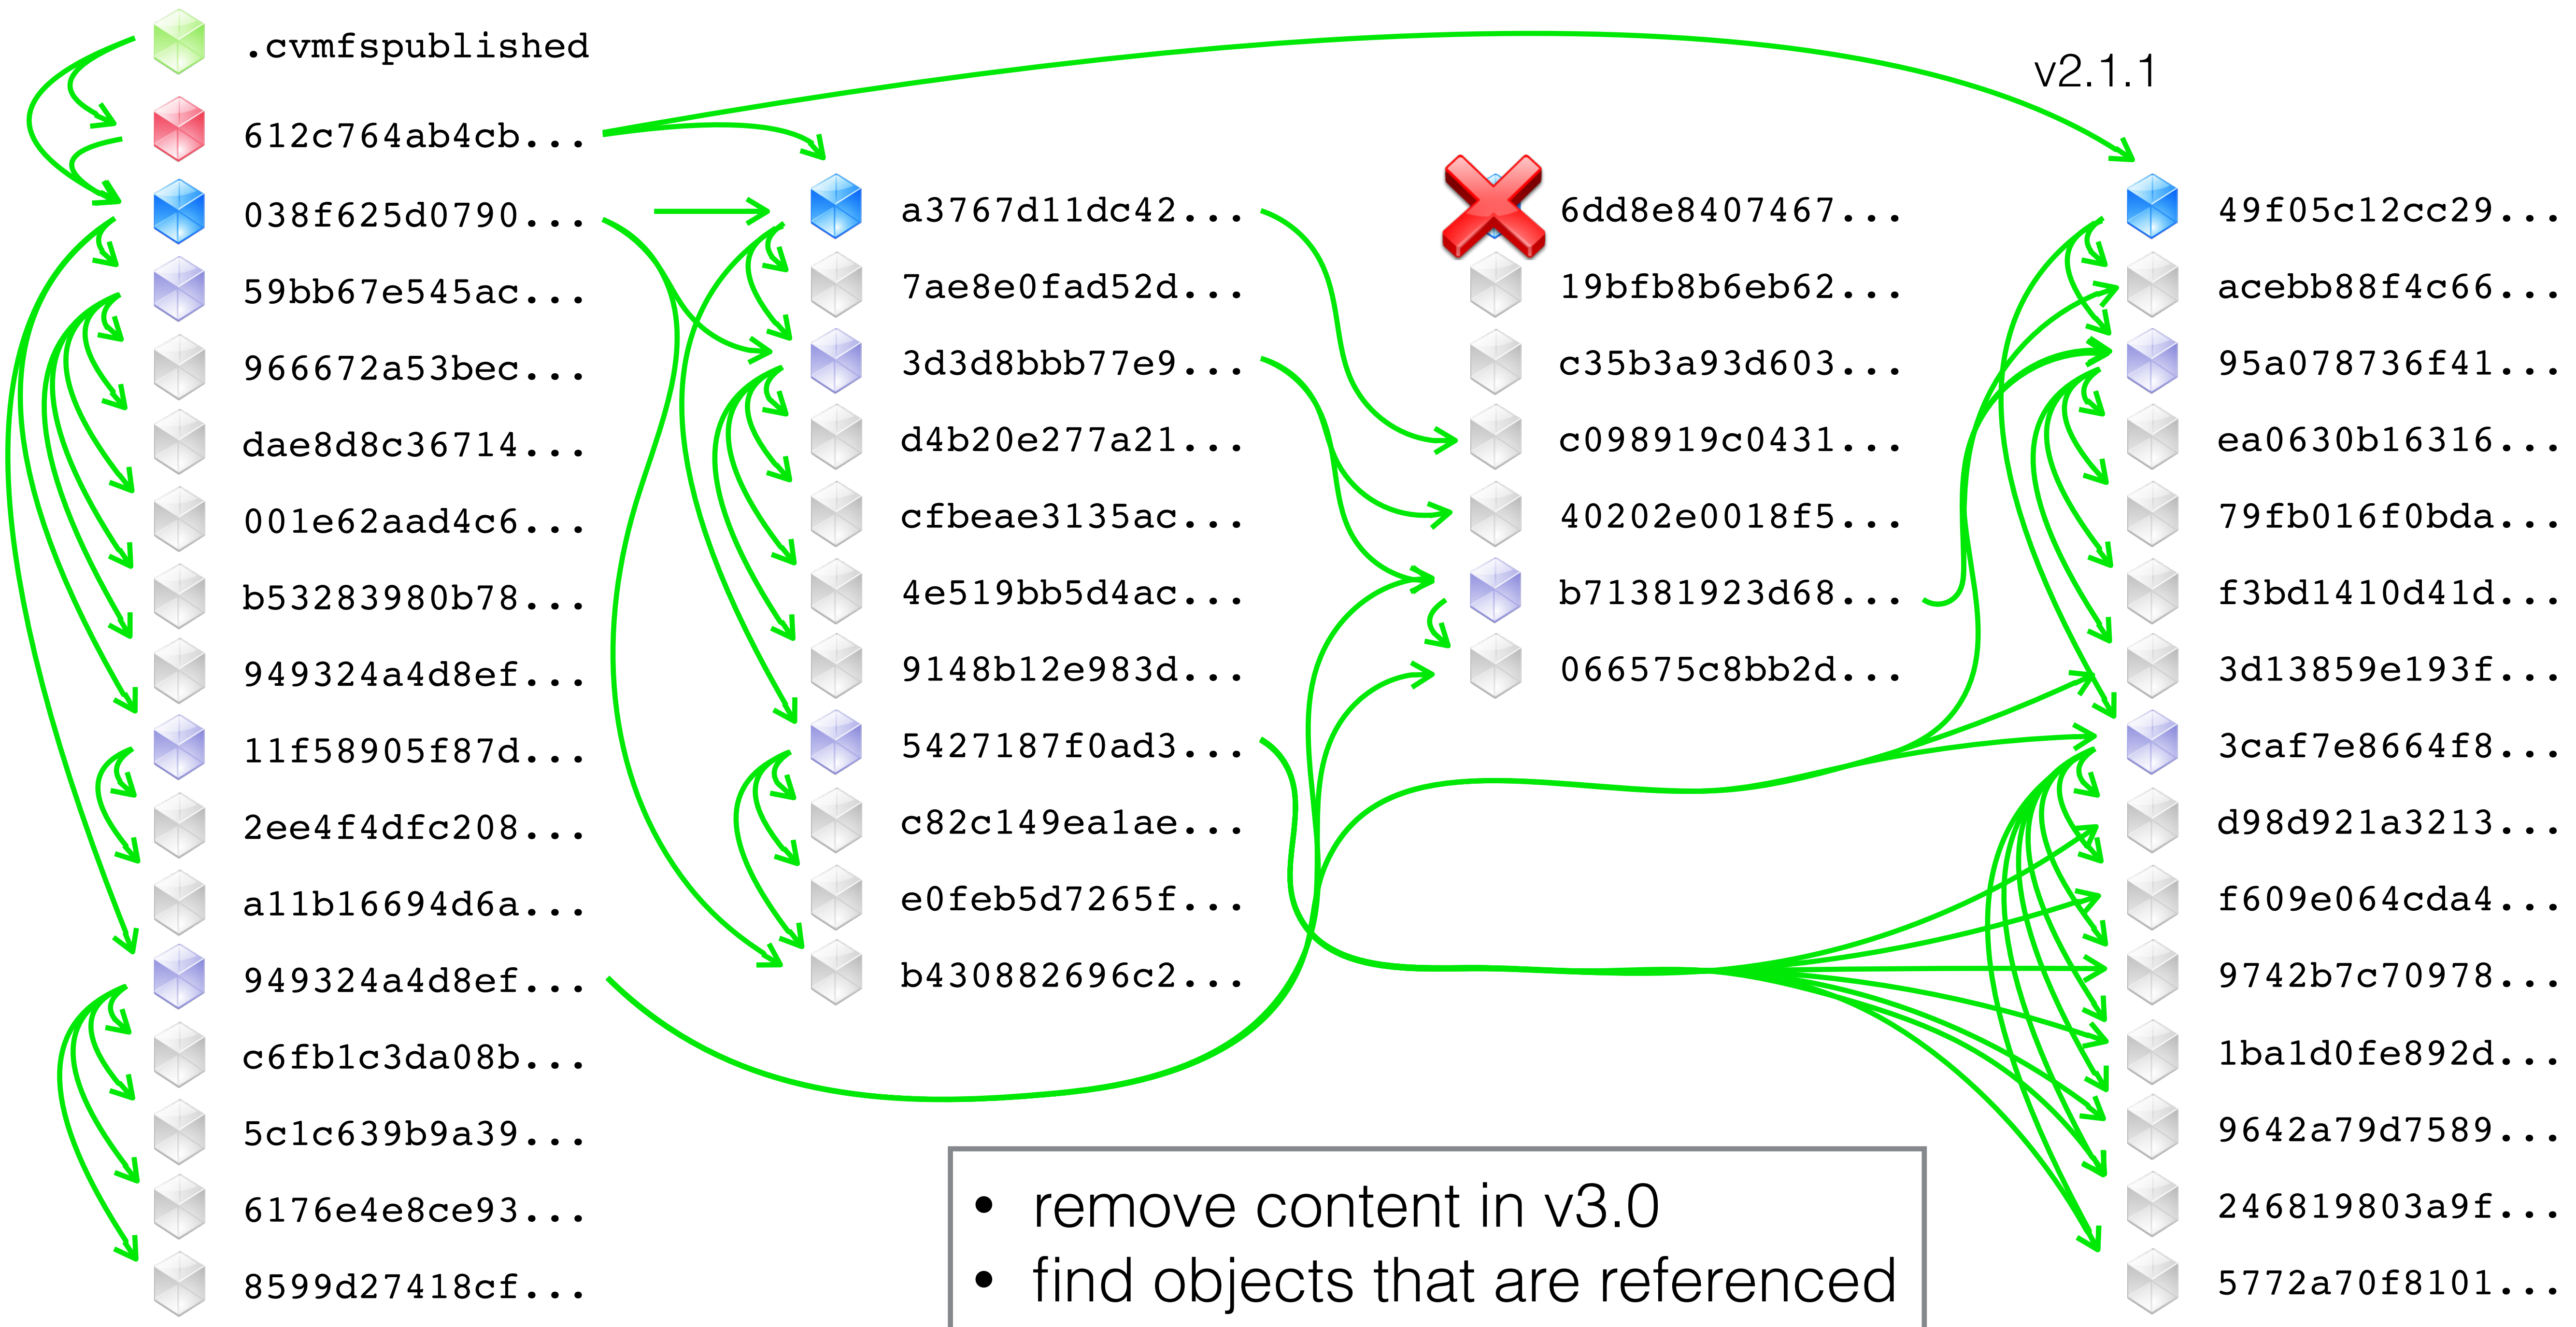

.cvmfspublished
612c764ab4cb...
038f625d0790...
59bb67e545ac...
966672a53bec...
dae8d8c36714...
001e62aad4c6...
b53283980b78...
949324a4d8ef...
11f58905f87d...
2ee4f4dfc208...
a11b16694d6a...
949324a4d8ef...
c6fb1c3da08b...
5c1c639b9a39...
6176e4e8ce93...
8599d27418cf...

a3767d11dc42...
7ae8e0fad52d...
3d3d8bbb77e9...
d4b20e277a21...
cfbeae3135ac...
4e519bb5d4ac...
9148b12e983d...
5427187f0ad3...
c82c149ea1ae...
e0feb5d7265f...
b430882696c2...

6dd8e8407467...
19bfb8b6eb62...
c35b3a93d603...
c098919c0431...
40202e0018f5...
b71381923d68...
066575c8bb2d...

v2.1.1

49f05c12cc29...
acebb88f4c66...
95a078736f41...
ea0630b16316...
79fb016f0bda...
f3bd1410d41d...
3d13859e193f...
3caf7e8664f8...
d98d921a3213...
f609e064cda4...
9742b7c70978...
1ba1d0fe892d...
9642a79d7589...
246819803a9f...
5772a70f8101...

- remove content in v3.0
- find objects that are referenced nowhere else

(schematic -not an actual repository)

Garbage Collection

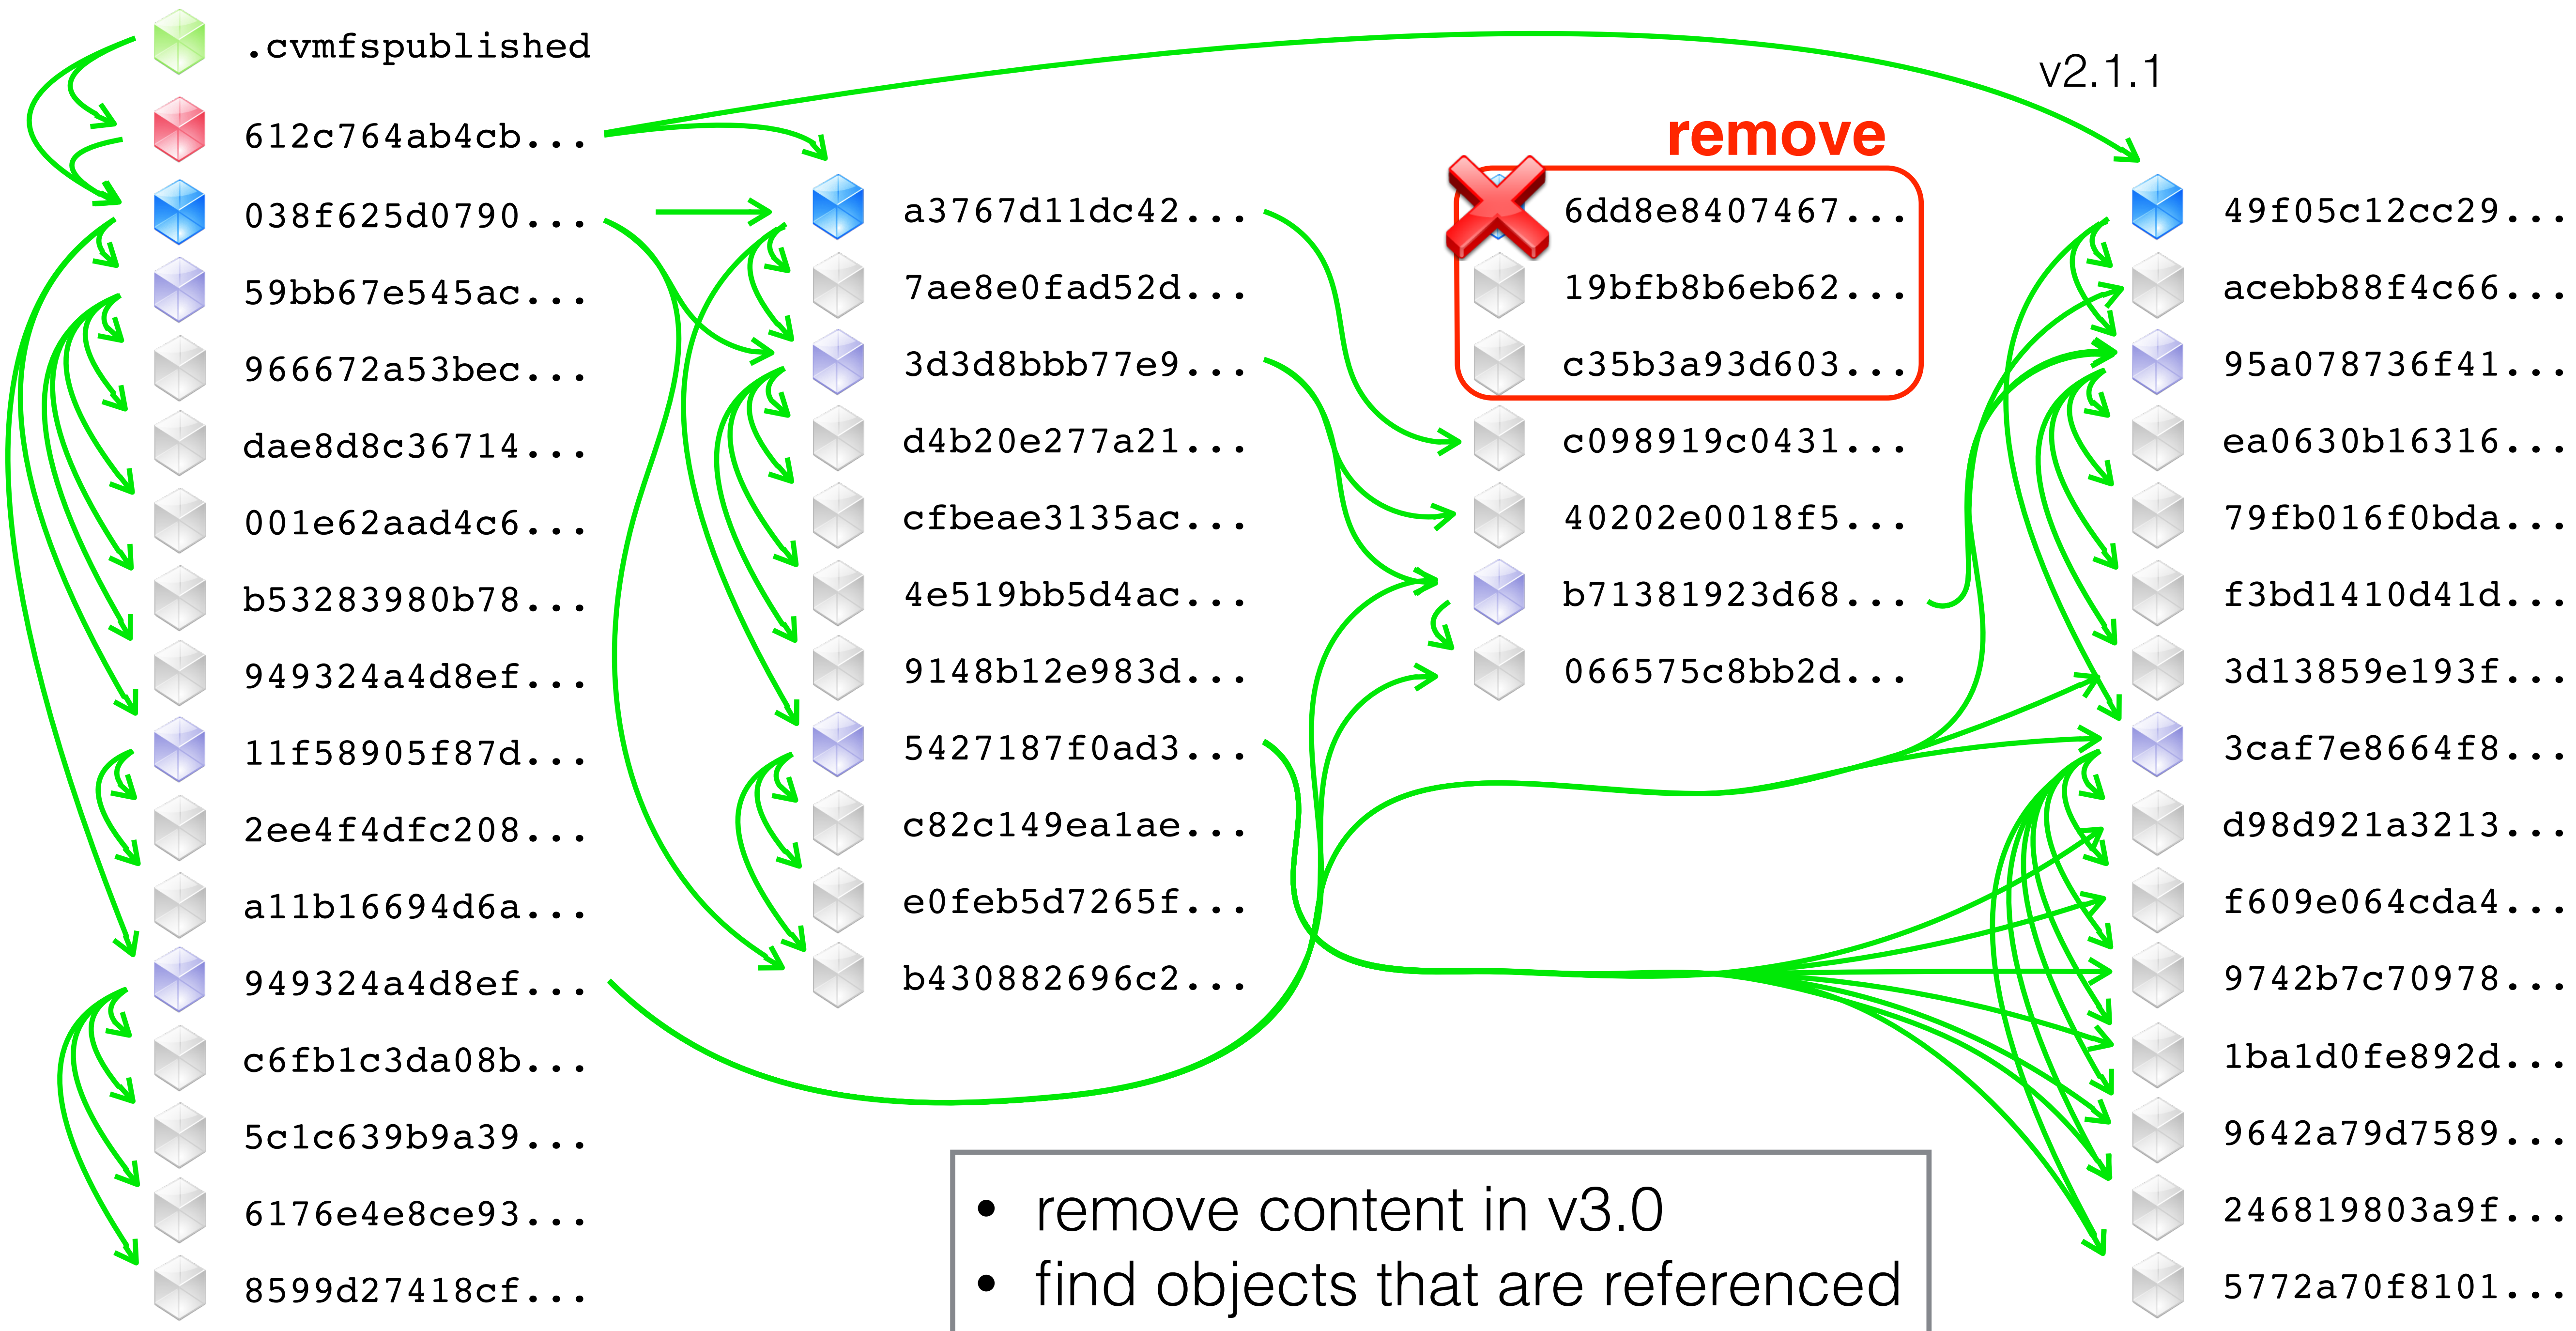

.cvmfspublished
612c764ab4cb...
038f625d0790...
59bb67e545ac...
966672a53bec...
dae8d8c36714...
001e62aad4c6...
b53283980b78...
949324a4d8ef...
11f58905f87d...
2ee4f4dfc208...
a11b16694d6a...
949324a4d8ef...
c6fb1c3da08b...
5c1c639b9a39...
6176e4e8ce93...
8599d27418cf...

a3767d11dc42...
7ae8e0fad52d...
3d3d8bbb77e9...
d4b20e277a21...
cfbeae3135ac...
4e519bb5d4ac...
9148b12e983d...
5427187f0ad3...
c82c149ea1ae...
e0feb5d7265f...
b430882696c2...

remove

6dd8e8407467...
19bfb8b6eb62...
c35b3a93d603...
c098919c0431...
40202e0018f5...
b71381923d68...
066575c8bb2d...

v2.1.1

49f05c12cc29...
acebb88f4c66...
95a078736f41...
ea0630b16316...
79fb016f0bda...
f3bd1410d41d...
3d13859e193f...
3caf7e8664f8...
d98d921a3213...
f609e064cda4...
9742b7c70978...
1ba1d0fe892d...
9642a79d7589...
246819803a9f...
5772a70f8101...

- remove content in v3.0
- find objects that are referenced nowhere else

(schematic -not an actual repository)

CERN