Use of Java / JVM at Twitter

@TonyPrintezis | @TwitterBoston

tprintezis@twitter.com

#JCP EC Twitter reps

Tony Printezis

- VM Team | Infrastructure Org | Twitter
- ex-HotSpot (Sun / Oracle, 6+ years), ex-SunLabs (3+ years)

Ramki Ramakrishna

- VM Team | Infrastructure Org | Twitter
- ex-HotSpot (Sun / Oracle, 10+ years)

#Twitter

#Twitter : the short version

#Platform

#RealTime

#Twitter : services

- A #huge distributed system
  - huge number of machines running a huge number of JVMs
  - multiple data centers
  - services communicate with other services via RPC
  - typically, many JVMs per service
- Most services run on
  - CentOS Linux distribution (x64)
  - OpenJDK distribution (TwitterJDK)
  - Mesos (data center scheduling)
  - Finagle (async RPC for Scala / Java)

#Twitter : programming languages

server-side

- Scala (most used by far)
- Java
- Ruby (moving away from it…)
- C/C++ (a bit)
- Python (a bit)

#Twitter : Java libraries

- We rely heavily on
  - NIO
  - collections
  - concurrent collections
  - concurrent utilities
  - etc.

#TwitterJDK

#TwitterJDK : vm team

- Infrastructure Org
- Responsible for releasing, maintaining, customizing, improving
  - TwitterJDK
- Developer support
  - consultancy
  - troubleshooting
  - education

#TwitterJDK : why?

Why deploy our own JDK?

- Commercial support wouldn't work for us
  - too expensive
  - too slow
- OpenJDK development is slow
  - dreadfully slow process
  - change turnaround feels like eternity
    - current development repo is JDK 9, months / years away
    - backports to JDK 8: slow turnaround
- We believe we can support ourselves best

#TwitterJDK : pushing the limits

- We push the JVM's limits more than most
    - large heaps
    - huge scale
    - strict latency requirements
- Optimize for our stack
    - Scala, Finagle, CentOS, x64, etc.
- Optimize for our workloads
    - quite different vs. most other JVM workloads I’m familiar with

#TwitterJDK : release

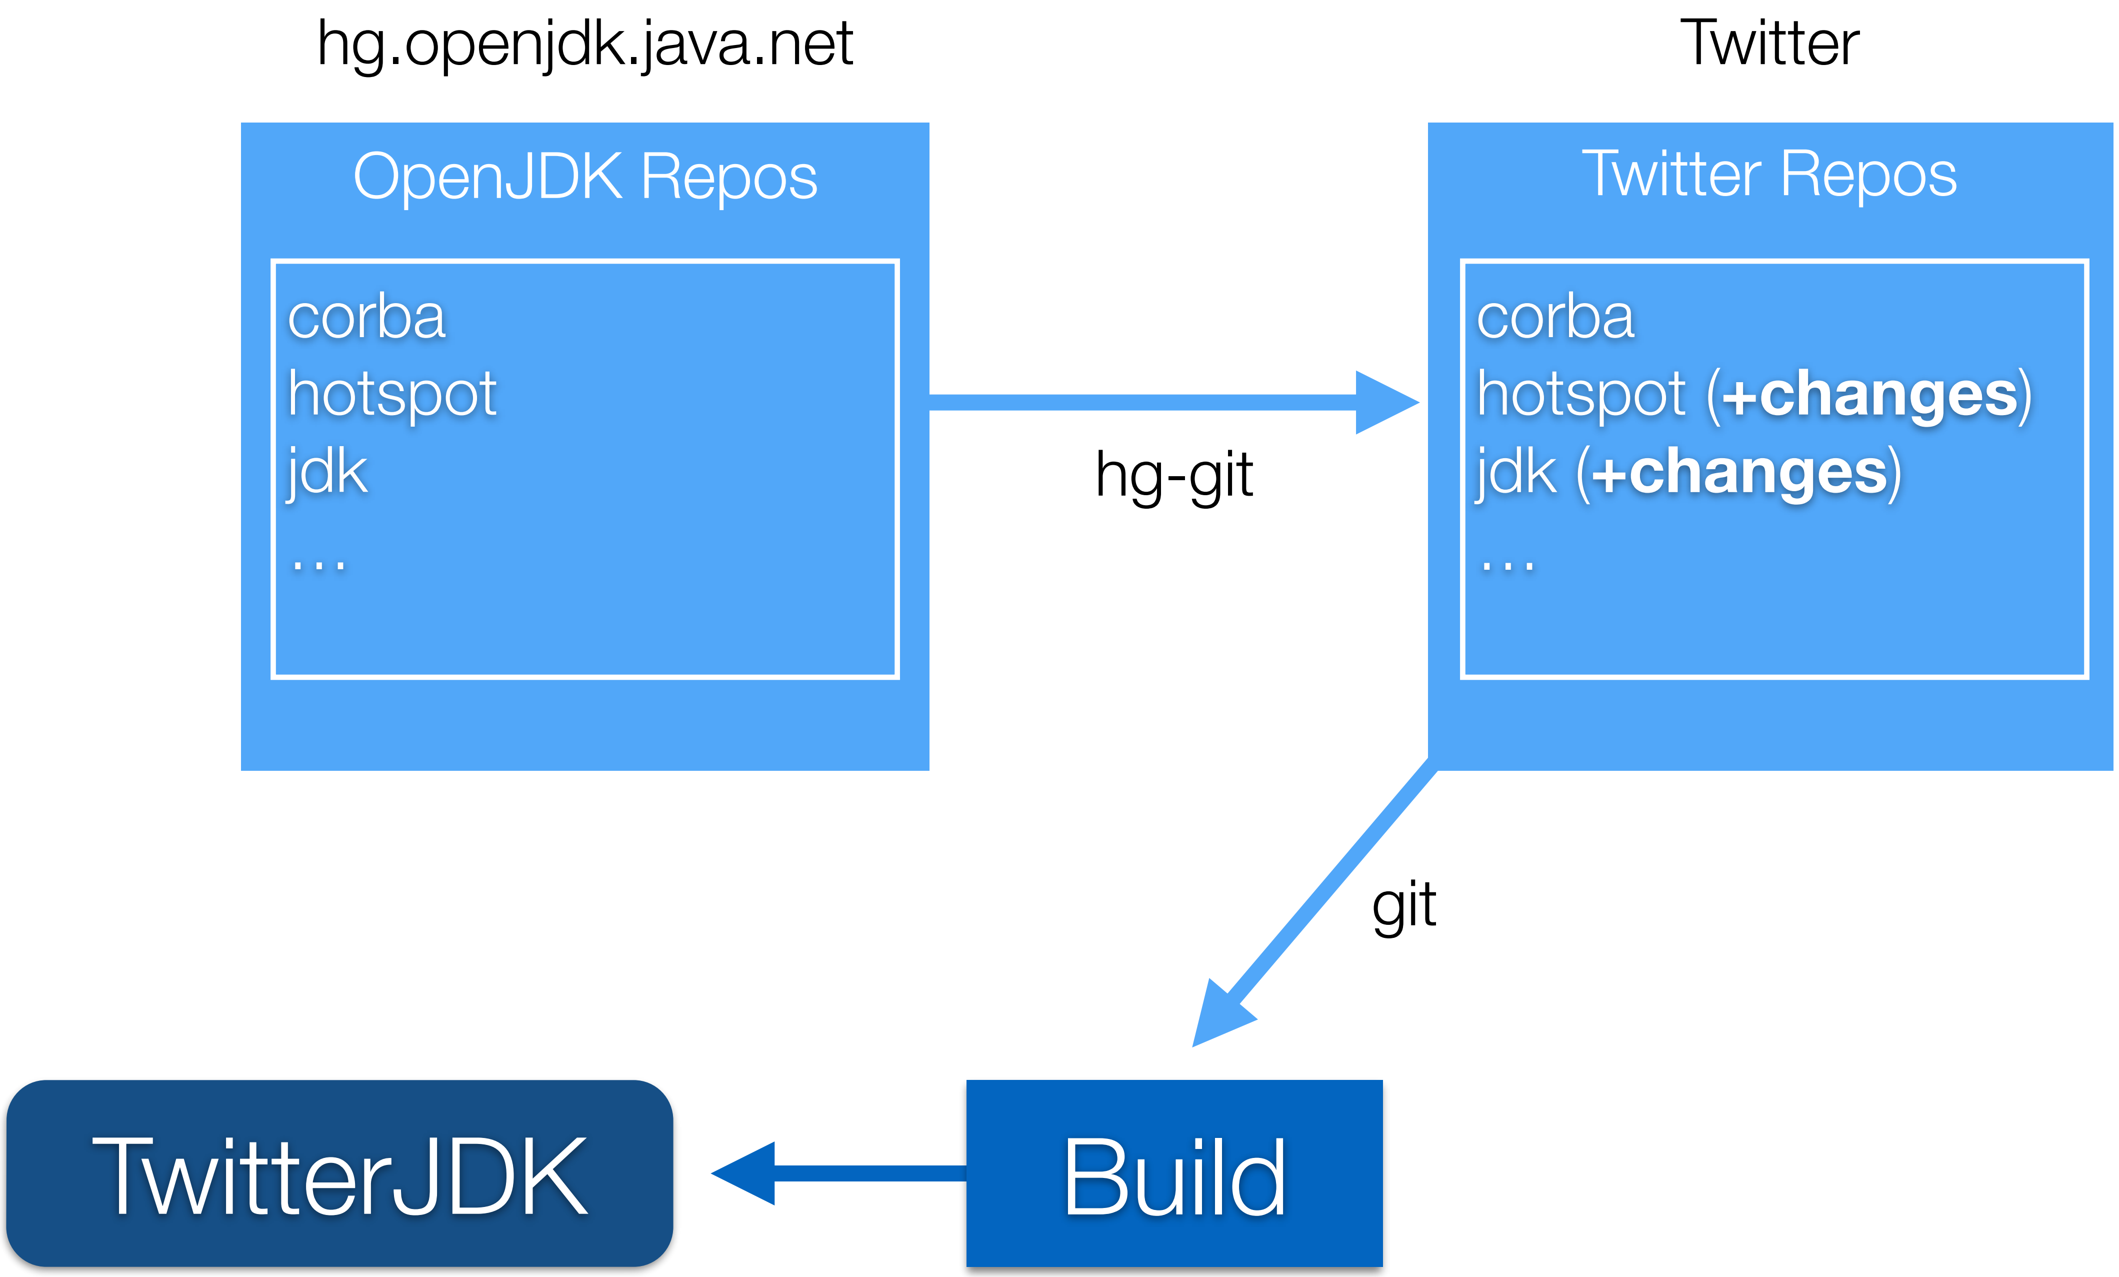

hg.openjdk.java.net

Twitter

OpenJDK Repos

corba
hotspot
jdk
…

hg-git

Twitter Repos

corba
hotspot (**+changes**)
jdk (**+changes**)
…

git

TwitterJDK

Build

#TwitterJDK : release

- Based on OpenJDK
  - Synced up to the latest update release
    - leverage all OpenJDK testing
  - Plus our own changes
    - which we'll be happy to contribute back BTW…
  - Plus (small number of) additional patches (e.g., security / critical fixes)
- Monthly releases
- Current main release
  - JDK 8 / Tiered Compilation / 64-bit / CentOS

#TwitterJDK : enhancements

- Heap Profiling
- Binary Logging Framework (JVM + Java tracepoints)
- Intermediate Generation(s) for G1 GC
- Misc Bug / Performance Fixes
- Mostly in HotSpot + a few small patches to the JDK libraries

#Questions