


How SAP is using Python to test its database SAP HANA

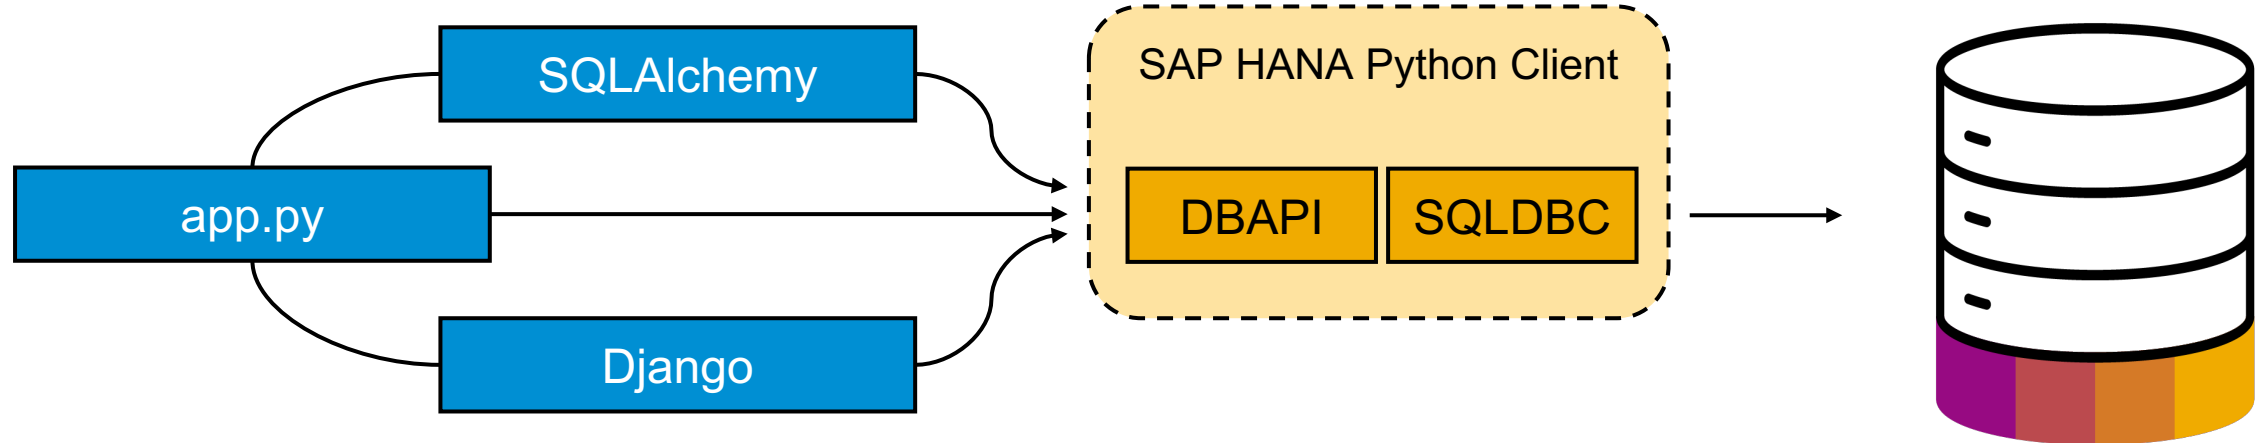
Christoph Heer
EuroPython 2017 - July 11

SAP HANA


- In-Memory database
- Column & Row store
- OLAP + OLTP
- Single host or distributed
- SQL
- C++ & Python



Python + SAP HANA



Testing a database




End-to-end tests
Python

Integration tests
Python

Unit/Component tests
C++






Gerrit

Pending changes

Central
Repository

Pre-merge Quality Assurance




Build

Test

Static code analysis

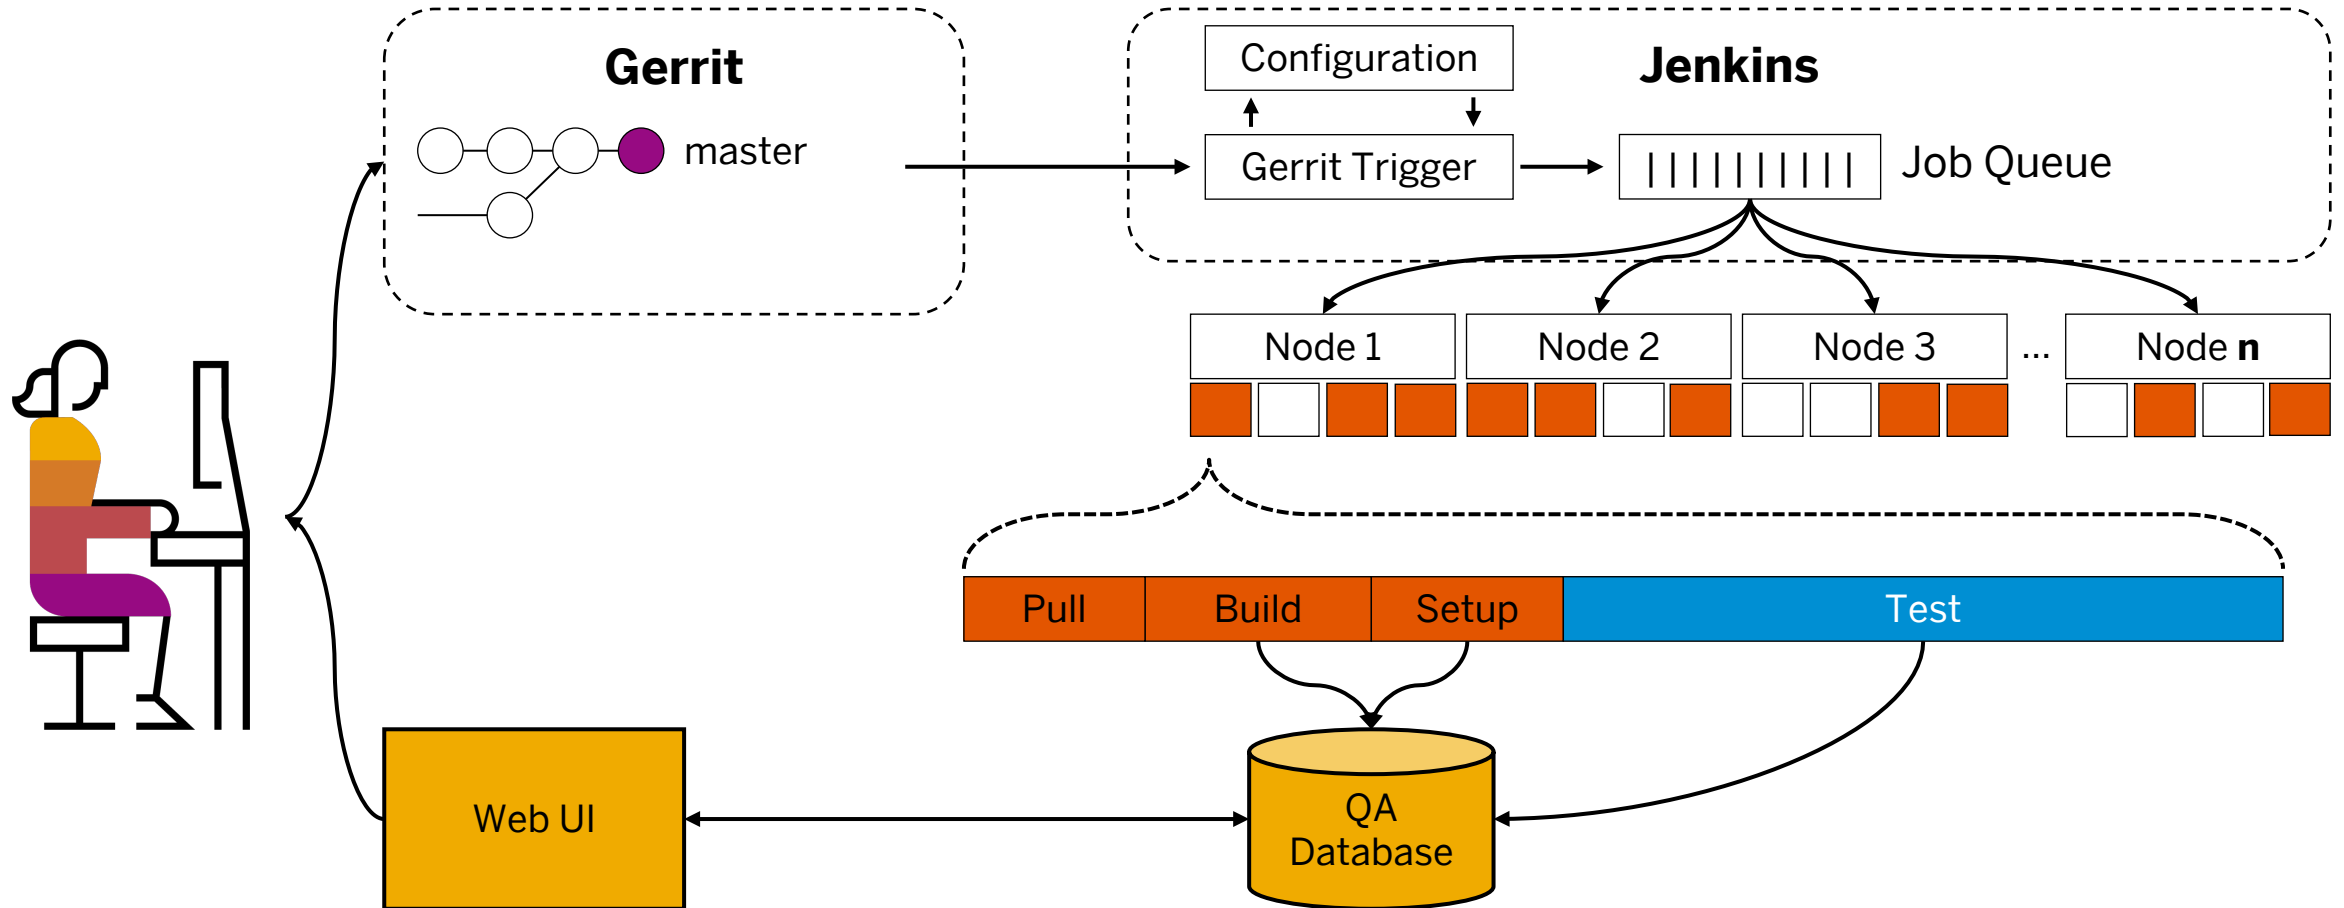
Style checks

C++ sanitizers



Review

Implementation - 2010






600

Developers



700

Commits per day




13 mil.

lines Python tests



36,000h

Test runtime per day



1,300

Jenkins nodes




408 TB

Memory


Scaling domains



Test Runtime



Test Scheduling

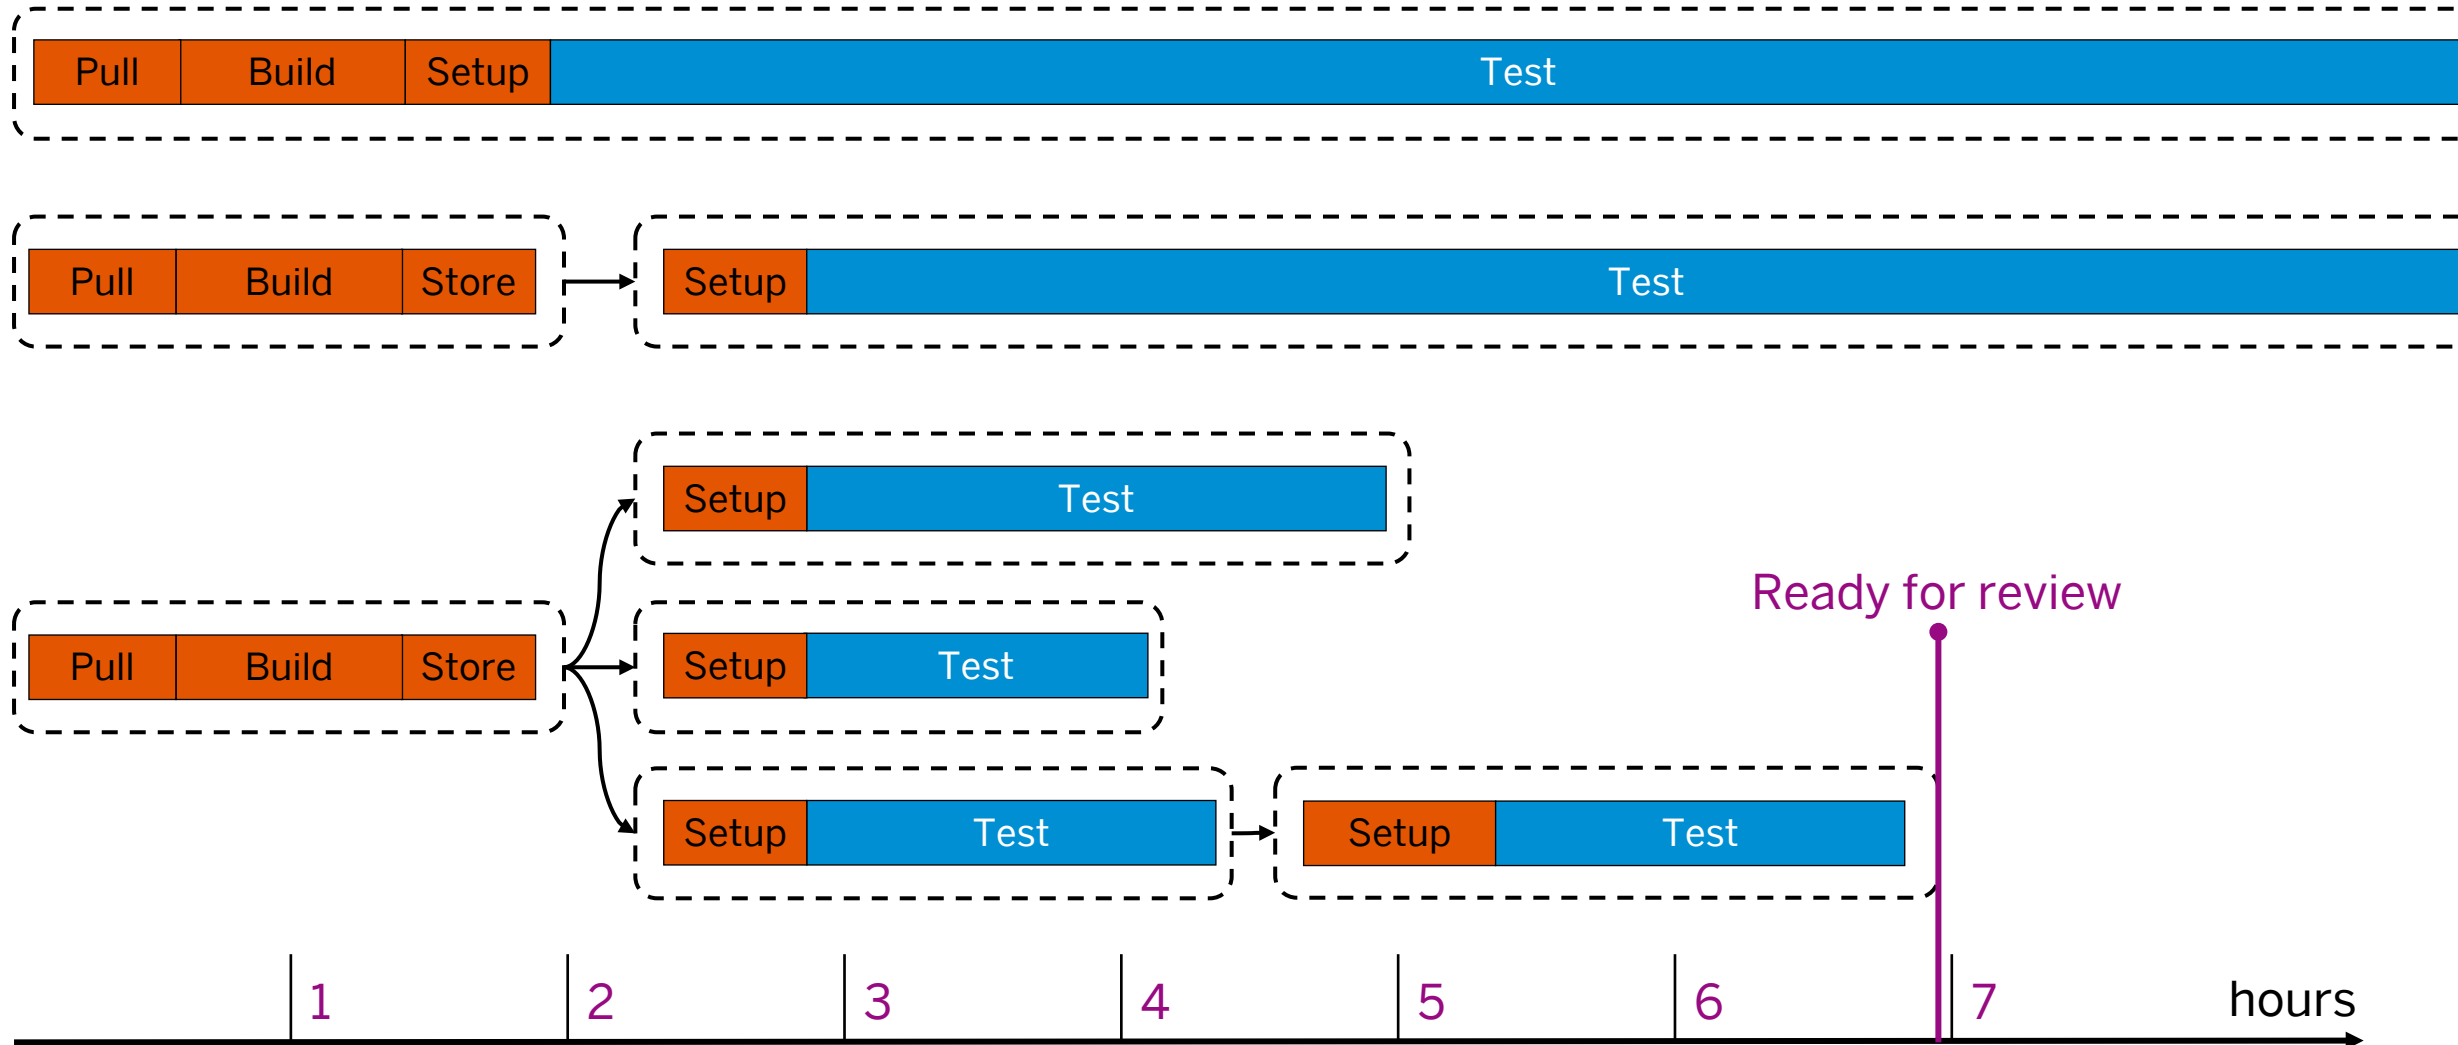


Artifacts



Healthy Test Environment

Test Runtime: Divide and conquer test



Test Scheduling: Handle test failures



git 003adae	TestBaseFunctionality	FAILED
git 003adae	TestBaseFunctionality_rerun	OK

- Regression, sporadic failure or infrastructure problem
- Rerun test to
 - confirm stable failure => regression
 - exclude infrastructure problem
- Who restarts failed tests?

Test Scheduling

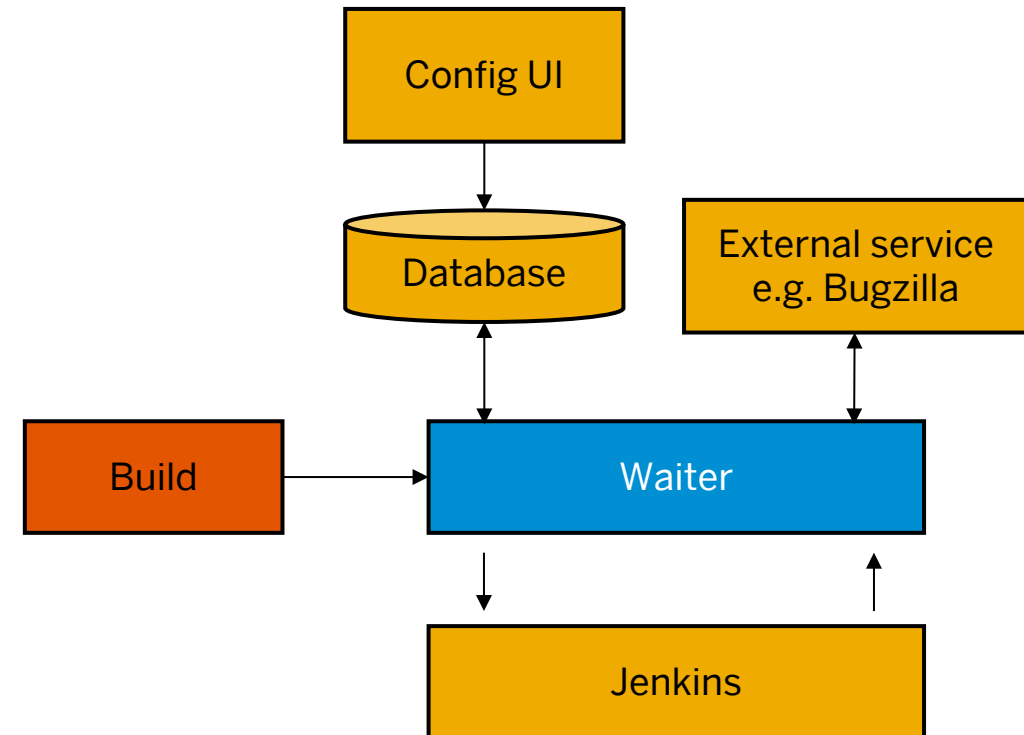


Configuration:

- What tests should run?
- Branch specific configuration
- Layered testing
- Exclude current broken tests

Observe and react:

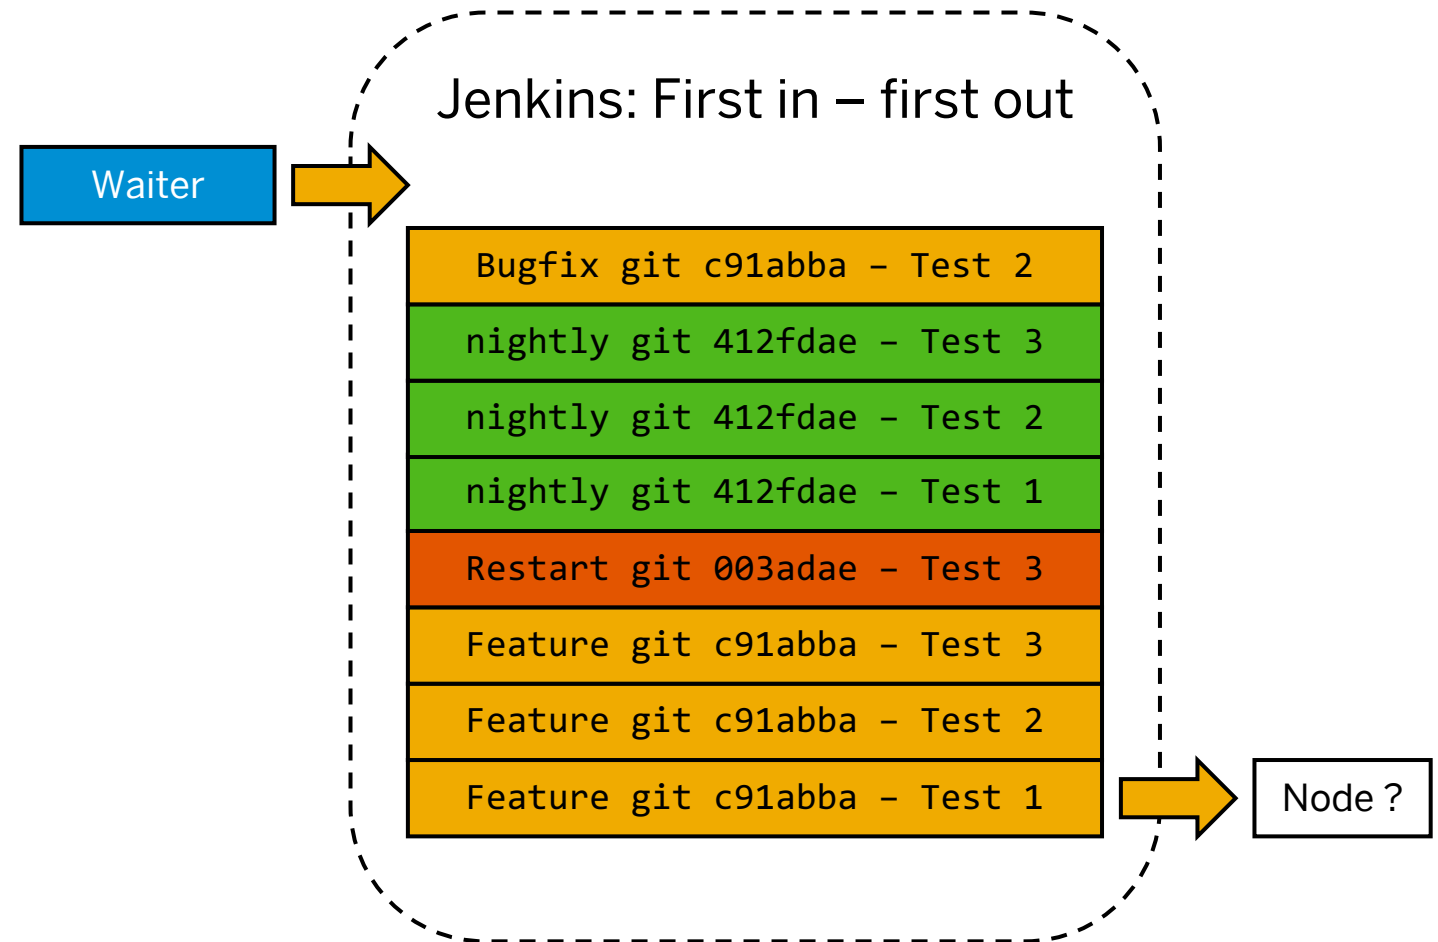
- Re-schedule failed tests
- Automated review
- Notify about completion



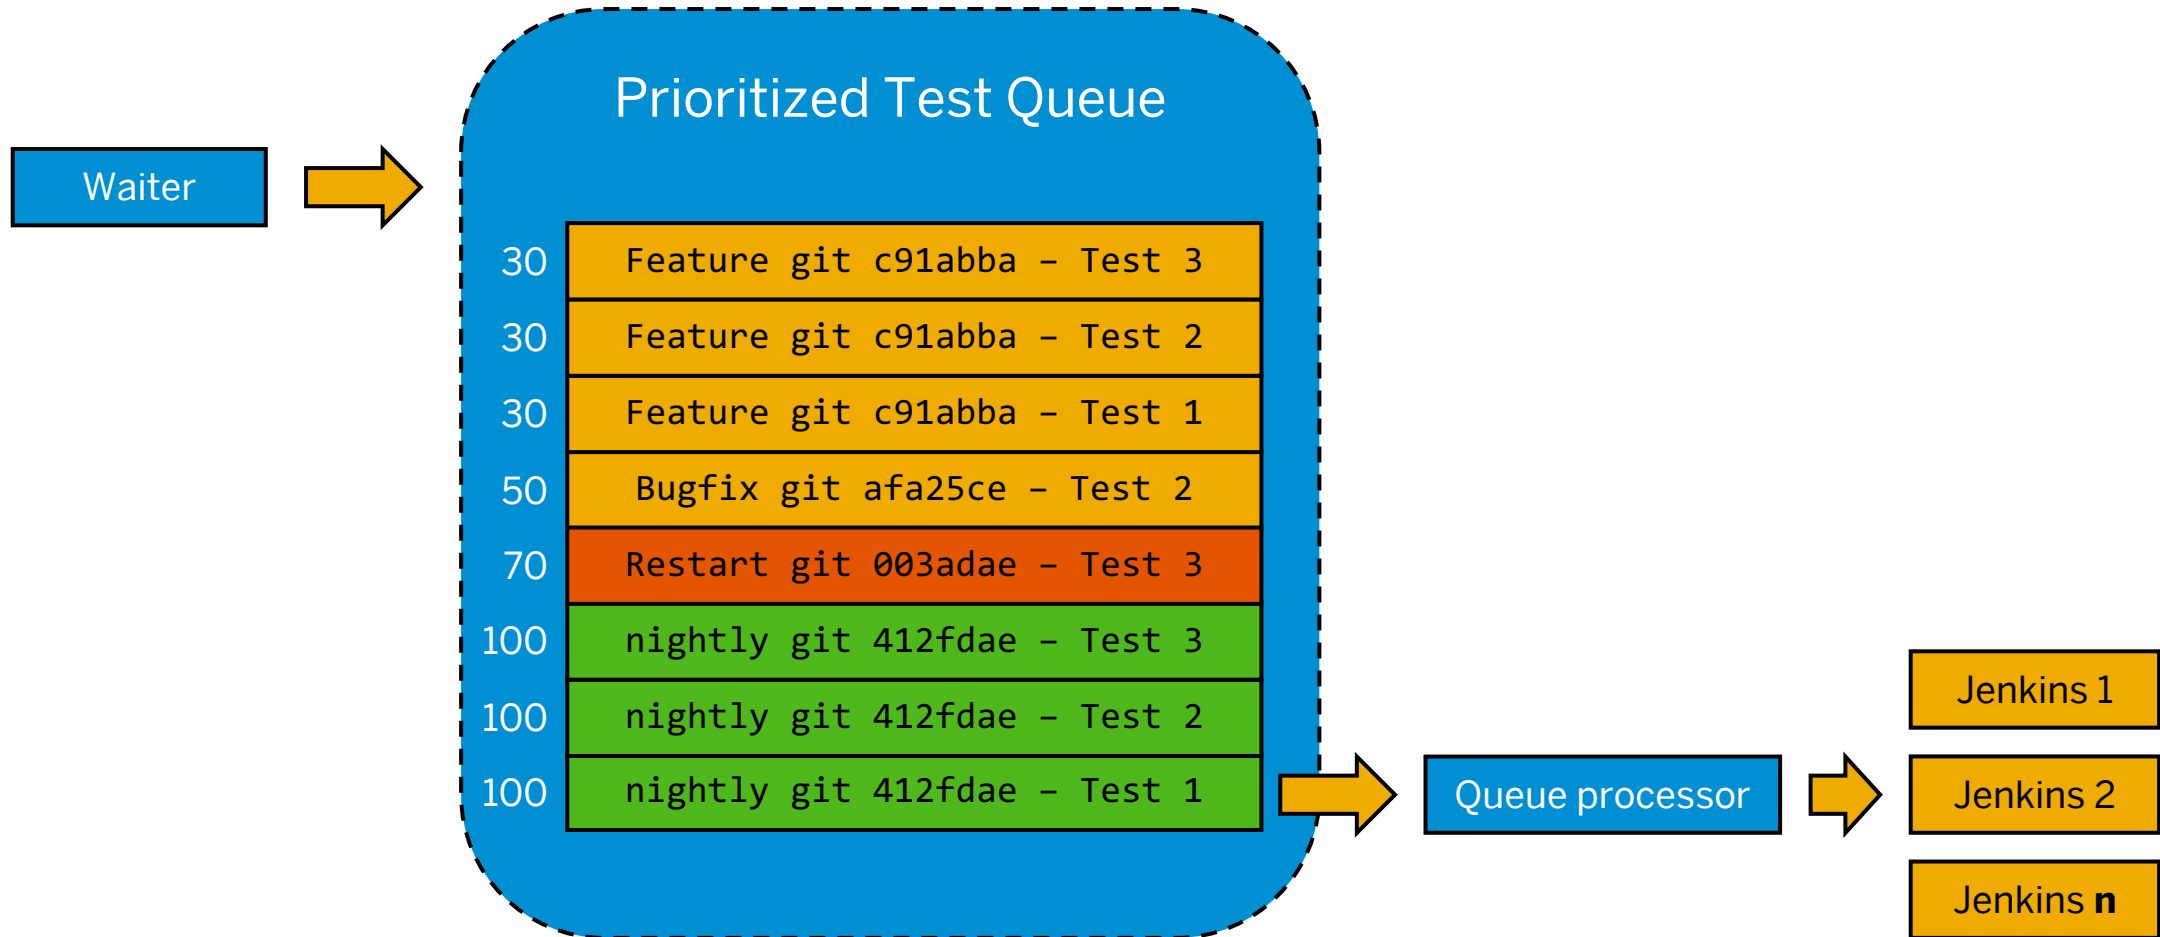
Test Scheduling: Queuing



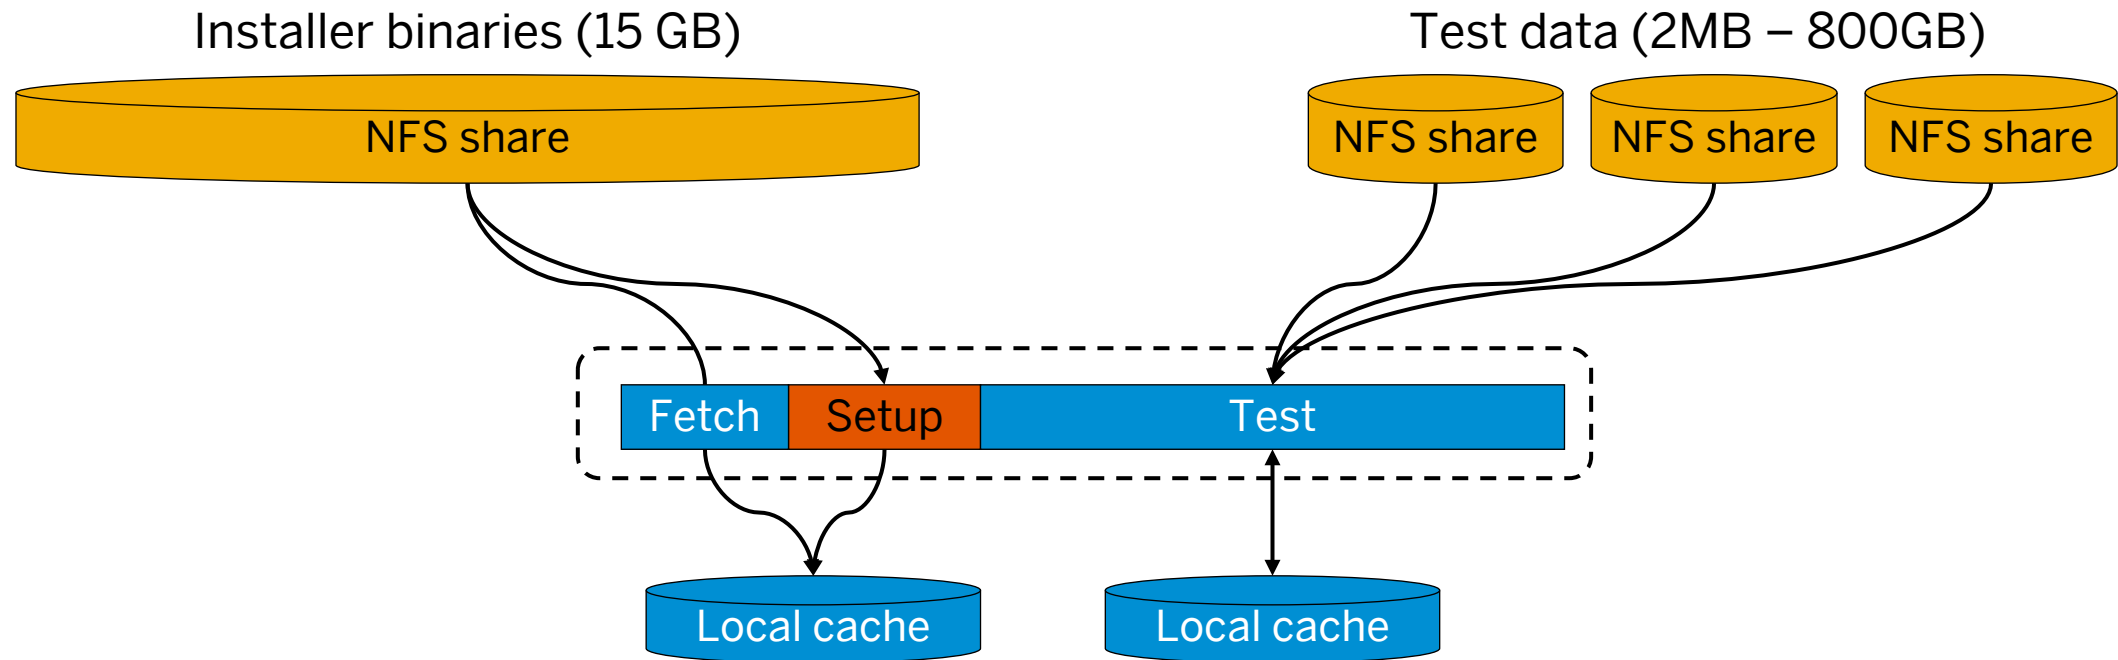
- Nightly test runs should be complete next morning
- Test bug fixes before new features
- Finish testing of 95% tested commit



Test Scheduling: Queuing



Artifacts




9PB data transfer per week
hitting bandwidth limits




66% less traffic **to 3PB**
per week

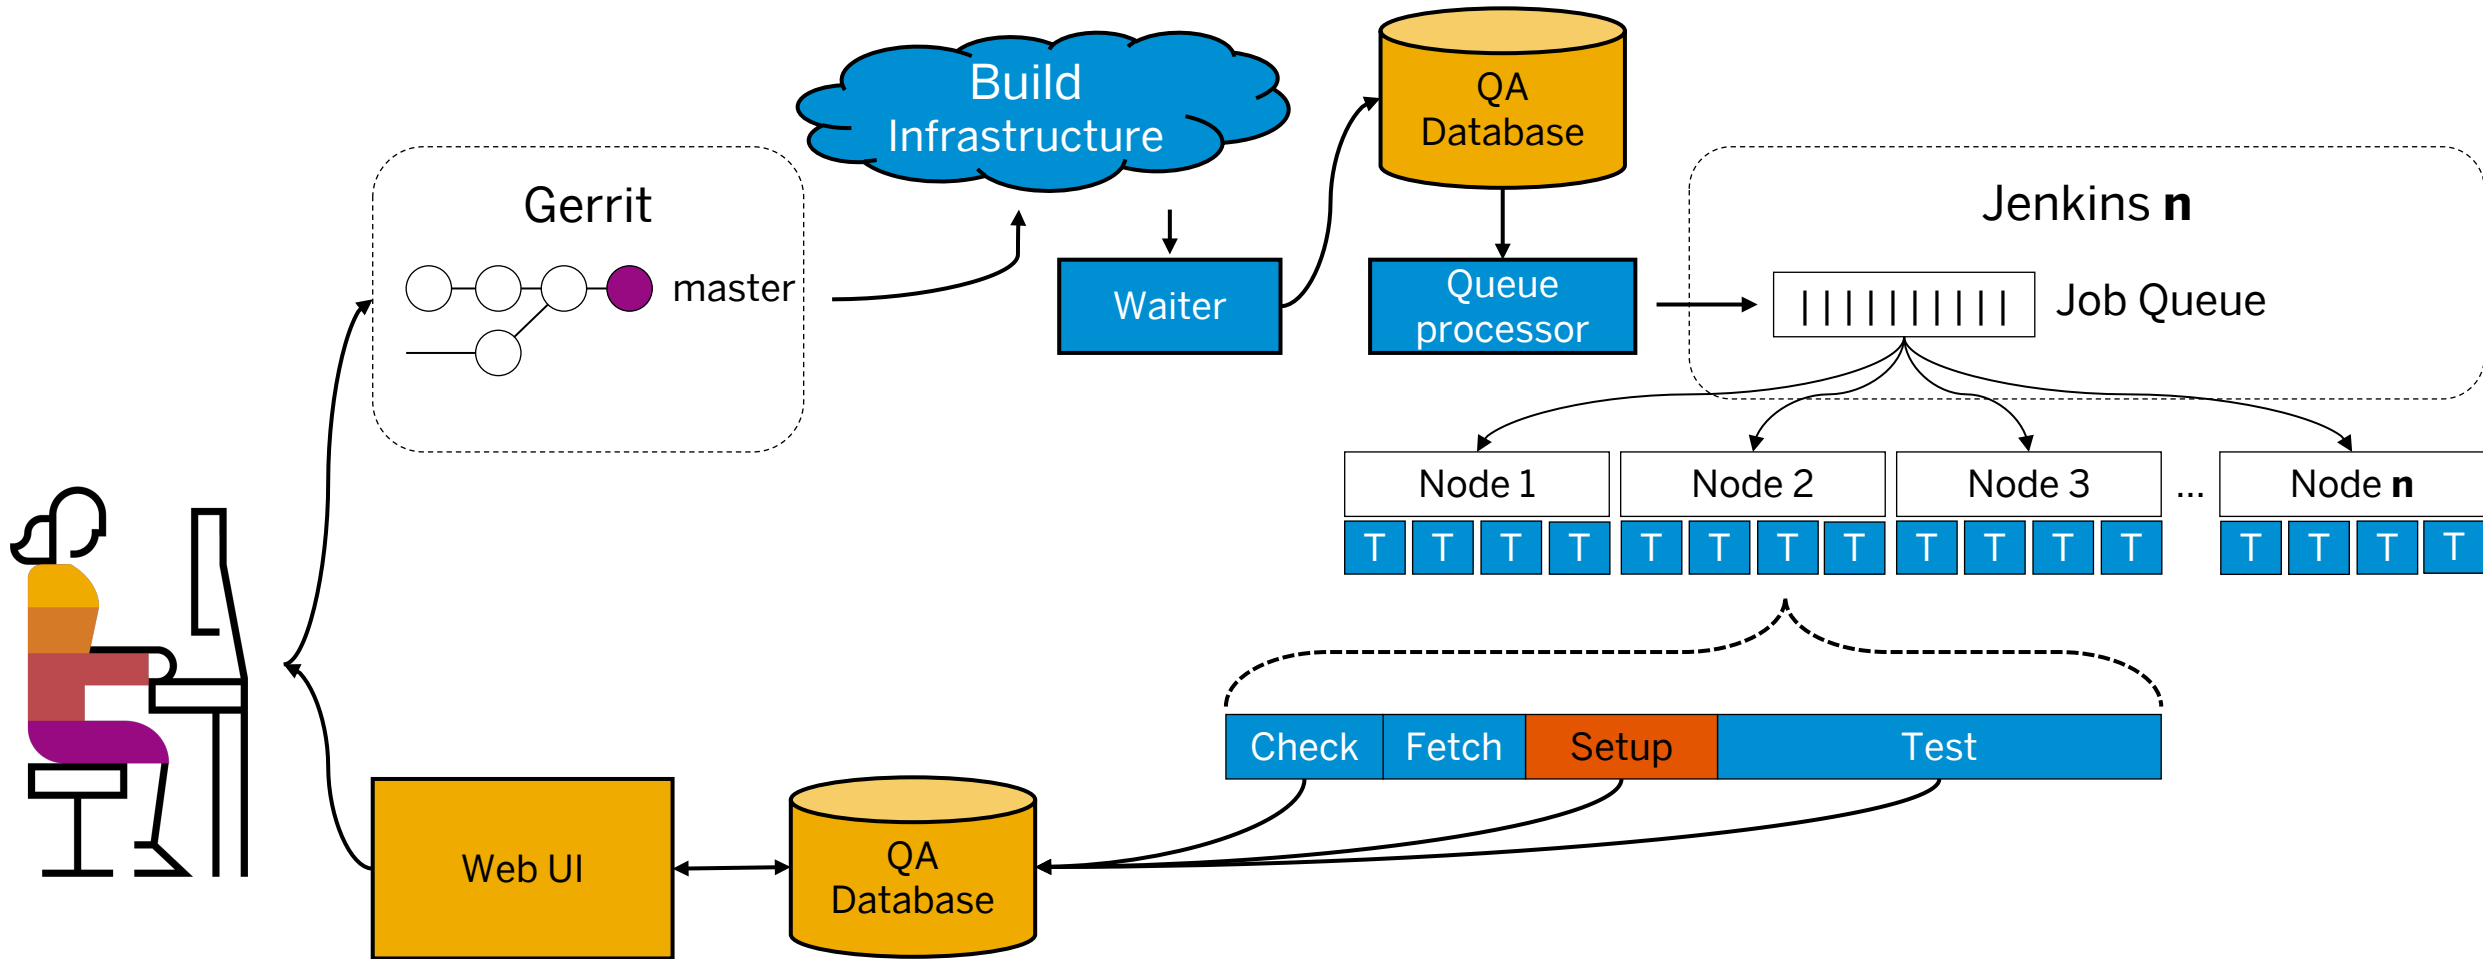
Healthy Test Environment




- External dependencies
 - will fail
- Parallel testing on host
 - Noisy neighbors
- Health check before and while testing
 - Failure marks test as invalid
 - Reschedule
 - Invalid tests are invisible for developers



Implementation – 2017





Learning curve


Non-developers can write tests



Community

virtualenv, pip, Fabric, Sentry etc.

Python



Development velocity

idea to production <1 day




Platform independent

CPU architecture + OS

Outlook

- Scale for ~3000 nodes
 - Resource based scheduling with Apache Mesos
- Isolation with Linux Containers
 - Guaranteed resources: Memory, CPU, IO
- Python 3
 - asyncio, type annotations
 - even better standard library



Thank you.

Christoph Heer

christoph.heer@sap.com
@ChristophHeer



We are hiring!

Open Python positions

(Junior) Python Developer:
<https://jobs.sap.com/s/kEui5u>

(Junior) Quality Engineer:
<https://jobs.sap.com/s/IKieYY>

(Senior) Quality Engineer:
<https://jobs.sap.com/s/2Cz9B8>

Linux System Administrator:
<https://jobs.sap.com/s/5uvE3u>

© 2017 SAP SE or an SAP affiliate company.

This presentation is licensed to the public under Attribution-NonCommercial-ShareAlike 4.0 International (CC BY-NC-SA 4.0).

The information contained herein may be changed without prior notice. Some software products marketed by SAP SE and its distributors contain proprietary software components of other software vendors. National product specifications may vary.

These materials are provided by SAP SE or an SAP affiliate company for informational purposes only, without representation or warranty of any kind, and SAP or its affiliated companies shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP or SAP affiliate company products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.

In particular, SAP SE or its affiliated companies have no obligation to pursue any course of business outlined in this document or any related presentation, or to develop or release any functionality mentioned therein. This document, or any related presentation, and SAP SE's or its affiliated companies' strategy and possible future developments, products, and/or platform directions and functionality are all subject to change and may be changed by SAP SE or its affiliated companies at any time for any reason without notice. The information in this document is not a commitment, promise, or legal obligation to deliver any material, code, or functionality. All forward-looking statements are subject to various risks and uncertainties that could cause actual results to differ materially from expectations. Readers are cautioned not to place undue reliance on these forward-looking statements, and they should not be relied upon in making purchasing decisions.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP SE (or an SAP affiliate company) in Germany and other countries. All other product and service names mentioned are the trademarks of their respective companies. See <http://global.sap.com/corporate-en/legal/copyright/index.epx> for additional trademark information and notices.