


Enterprise monitoring

Continuous Performance

...as a self-service

...with fully-automated feedback loops





© 2002 How Stuff Works

Continuous Integration

“...is the practice, in software engineering, of merging all developer working copies with a shared mainline several times a day.”

Pragmatically:

- autonomous, pre-scheduled build and deploy for DEV and QA
- inclusive of an automated “build verification test”
- is rarely autonomous or unattended

Implementation basics:

- powered-by Visual Studio, Jenkins, Bamboo
- integrated with code version control and config mgmt
- originated with eXP and test driven development (TDD)

Continuous Delivery

“the practice of continuous delivery further extends CI by making sure the software checked in on the mainline is always in a state that can be deployed to users and makes the actual deployment process very rapid.”

Pragmatically:

- automates release management between QA and PROD
- highly automated and orchestrated deploy/roll-back
- automated tests validate the code is “release ready”
- requires more rigorous, elaborate checking


Implementation basics:

- powered-by Puppet, Chef, VMWare, Home-grown
- defined as a part of eXP – refined by Jez Humble/Dave Farley


Continuous Performance Lifecycle



Continuous Performance Lifecycle



Promotional Flow



Promotional Flows


From DEV to PERF (via Continuous Integration)

- Performance regression testing – repeated testing and trending
- New code and app config changes
- New applications and integrations to be tested
- Unit-level performance information data

From PERF to OPS (via Continuous Delivery)

- Validation that a release will meet performance demand
- Performance guidance – top ten “worst offenders”
- Performance threshold update – re-configure APM monitors
- Performance test report generated and published to RM

Feedback Flow



Feedback Flow

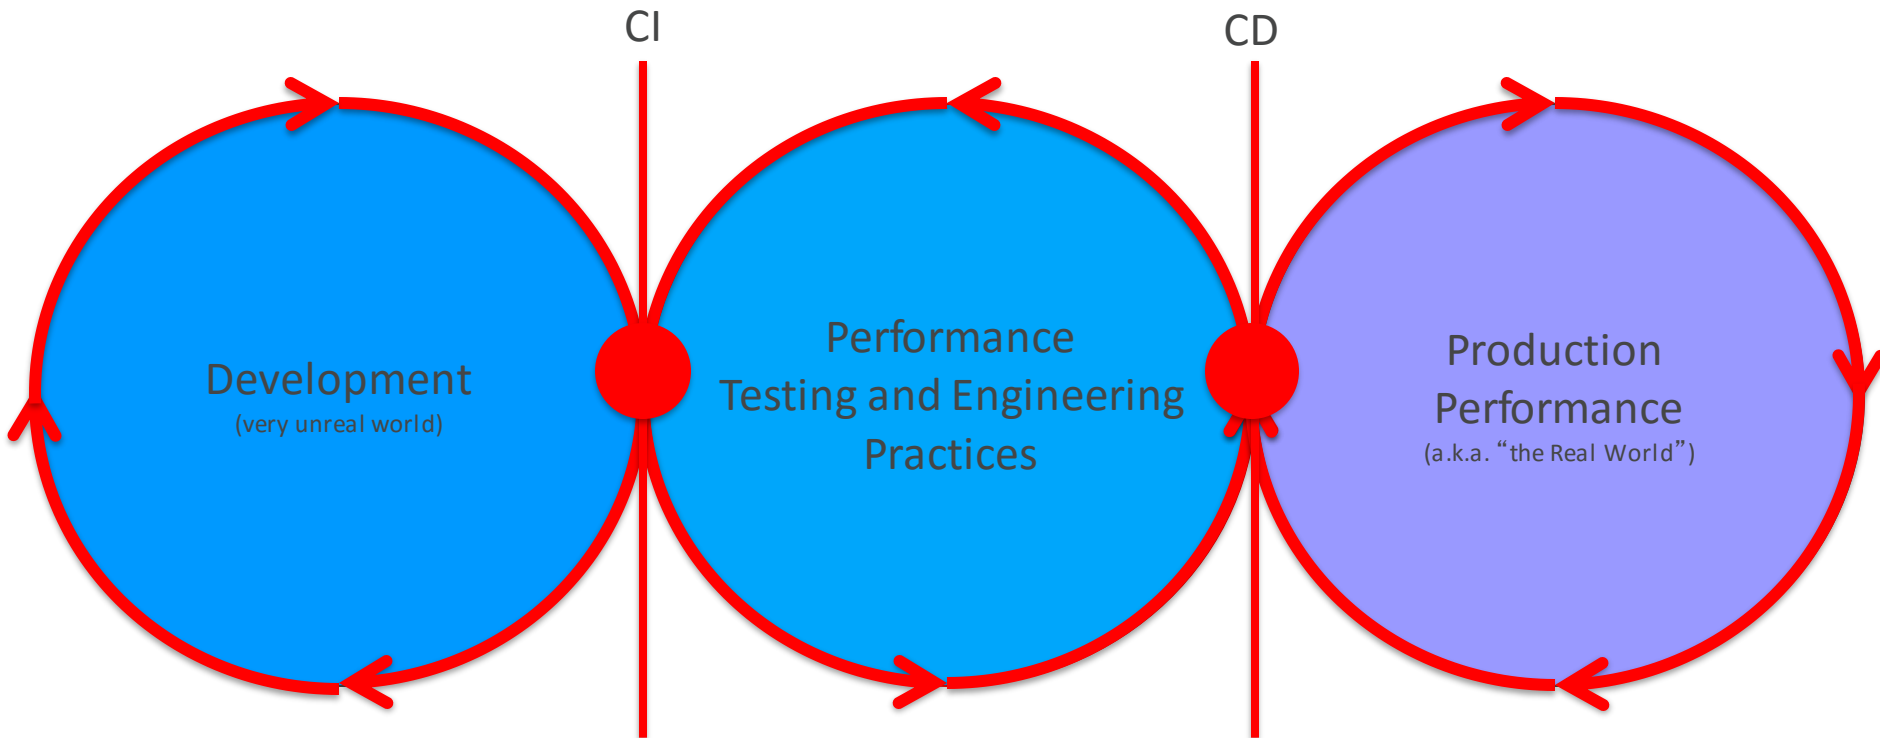
From OPS -> PERF

- Monitoring performance in PROD, setting new trends/thresholds
- Synchronizing test simulation with PROD (volume, mix, throughput)
- Synchronizing thresholds with PROD (resources, app metrics)
- Production issue repro/remediation – performance escalation

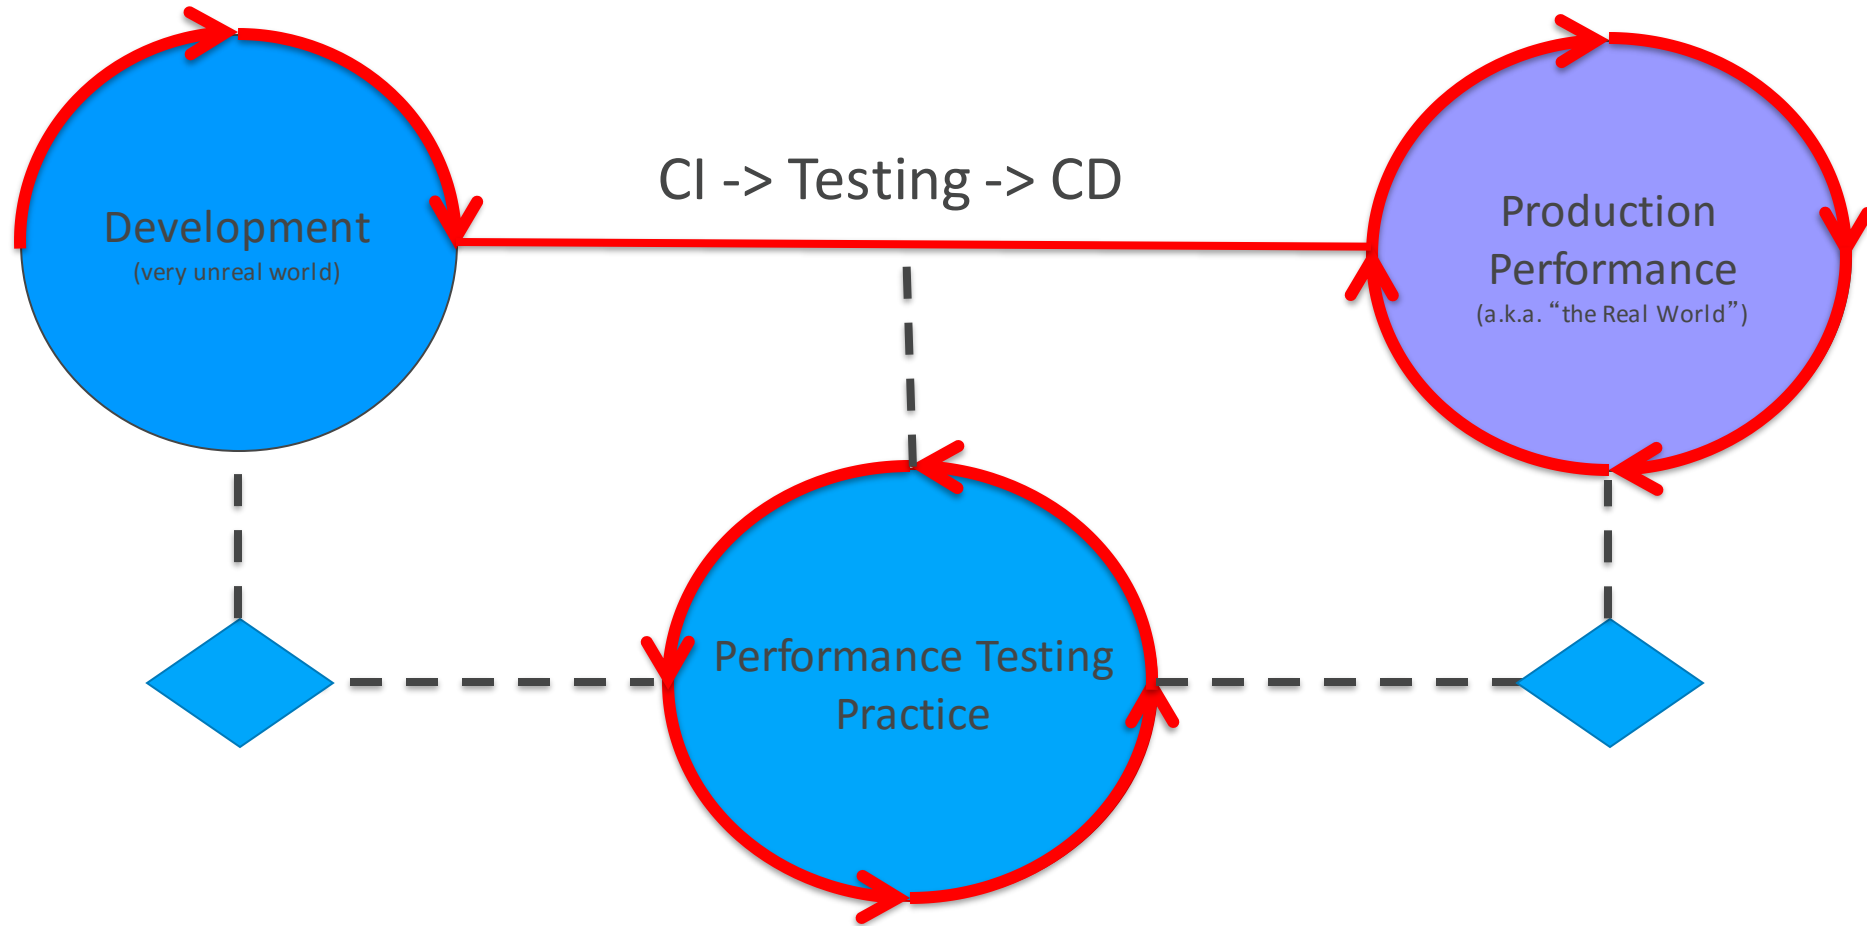
PERF->DEV

- Performance defects – bottlenecks and proposed remediation
- Early Performance Testing – unit-level performance results
- Strategic performance measurements – for architecture
- System future estimations – for PM's and Biz

Continuous Performance Flows



Out-of-Flow Continuous Performance



Performance Decisions

How do we make decisions about performance?

- We make decisions based on the information we have – not with the information we wish we had
- We make decision based on our understanding of the information we have – with our default perspective
- If the information is limited – the decision will be limited
- If the information isn't timely – our decision will be inaccurate
- If we don't know what came before – we can't estimate what will come next

With every sacrifice and limitation, our value decreases.

Performance Decisions

We can analyze performance at each point of the flow:

- Every time we move code along the promotional flow
- Every time we accept feedback about performance
- Every time we choose to reverse our decision to promote
- Every time we decide to repeat a test
- Every time we decide to log a bug
- Every time we change our plans according to new information

If the performance information is not timely:

- A wrong decision may delay competitive features for the biz
- A right decision may come too late to avoid disaster
- Any decision can and should be revisited in the face of change

Continuous Performance Tooling

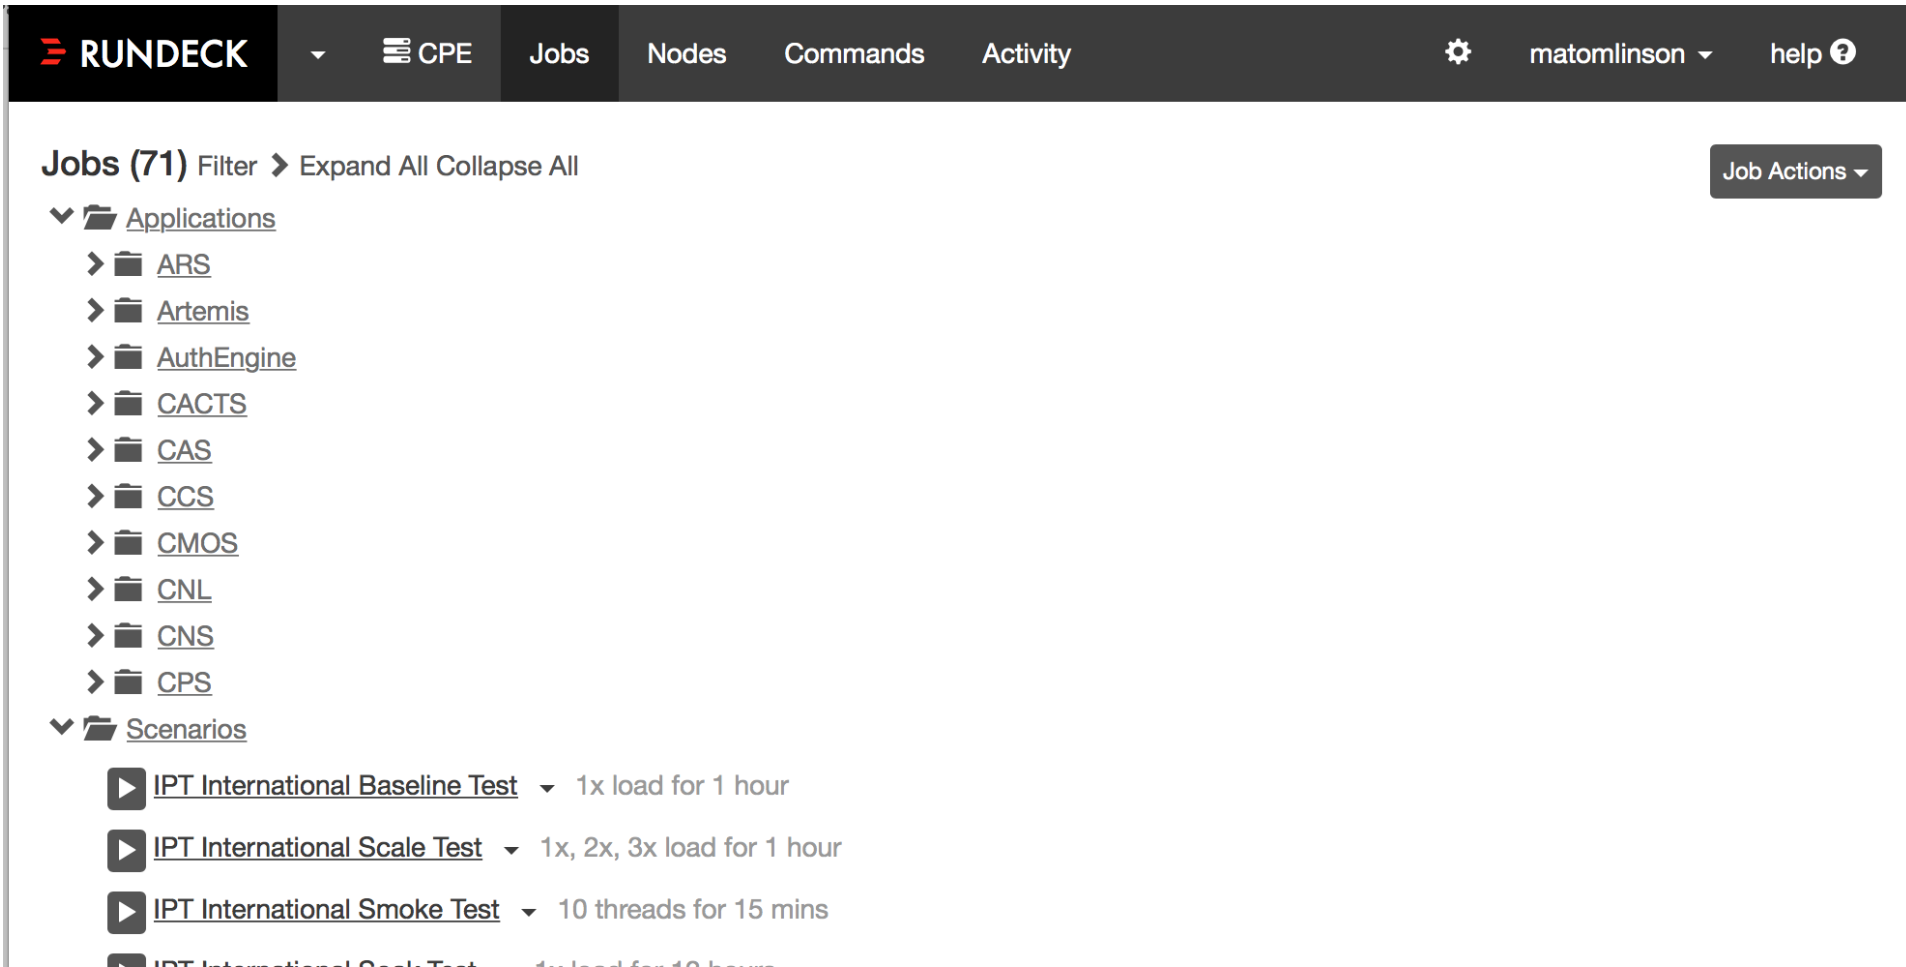
Continuous Performance Tooling

There are 3 primary “new tools” to add to your typical performance testing effort/team:

- tools for unattended automation
- tools for data trending/visualization
- tools for test notifications



DEMONSTRATION: UNATTENDED AUTOMATION WITH RUNDECK



The screenshot displays the Rundeck web interface. The top navigation bar is dark grey with the 'RUNDECK' logo on the left and user information 'matomlinson' and a 'help' link on the right. The main menu includes 'CPE', 'Jobs', 'Nodes', 'Commands', and 'Activity'. The 'Jobs' page is active, showing a list of jobs categorized under 'Applications' and 'Scenarios'. The 'Applications' section lists various services like ARS, Artemis, AuthEngine, CACTS, CAS, CCS, CMOS, CNL, CNS, and CPS. The 'Scenarios' section lists three tests: 'IPT International Baseline Test', 'IPT International Scale Test', and 'IPT International Smoke Test', each with a dropdown menu for configuration options.

RUNDECK CPE Jobs Nodes Commands Activity matomlinson help

Jobs (71) Filter Expand All Collapse All Job Actions

▼ Applications

- ARS
- Artemis
- AuthEngine
- CACTS
- CAS
- CCS
- CMOS
- CNL
- CNS
- CPS

▼ Scenarios

- IPT International Baseline Test ▼ 1x load for 1 hour
- IPT International Scale Test ▼ 1x, 2x, 3x load for 1 hour
- IPT International Smoke Test ▼ 10 threads for 15 mins
- IPT International Peak Test ▼ 1x load for 10 hours

REVIEW: UNATTENDED AUTOMATION WITH RUNDECK

- Using Rundeck we achieve the following:
 - Separation of jobs for scripts, scenarios, utilities
 - Ability to schedule automatic start
 - Ability to “halt the automation” (e.g. a hold button)
 - Ability to run a script or scenario manually
 - Job status notification
 - Job activity history

DEMONSTRATION: JMETER CUSTOMIZATIONS

The screenshot displays the Apache JMeter 3.2 interface with a custom test plan named 'ARSRequestorUK'. The left sidebar shows the test plan structure, including modules like 'Randomize Data Values', 'Transaction Mixer', and three thread groups. The right pane shows the 'Test Plan' configuration, including 'User Defined Variables' and test options.

Test Plan

Name: ARSRequestorUK

Comments:

User Defined Variables

Name:	Value
samplers	\${_P(samplers,ARSGetAccountData\,ARSGetStatementTr...
sampler_mix	\${_P(sampler_mix,012345)}
g1_threads	\${_P(g1_threads,1)}
g1_startup_delay	\${_P(g1_startup_delay,0)}
g1_rampup	\${_P(g1_rampup,30)}
g1_loop_count	\${_P(g1_loop_count,999999999)}
g1_duration	\${_P(g1_duration,300)}
g1_tpm	\${_P(g1_tpm,18)}
g1_data_offset	\${_P(g1_data_offset,0)}
g2_threads	\${_P(g2_threads,0)}
g2_startup_delay	\${_P(g2_startup_delay,0)}
g2_rampup	\${_P(g2_rampup,0)}
g2_loop_count	\${_P(g2_loop_count,0)}
g2_duration	\${_P(g2_duration,0)}
g2_tpm	\${_P(g2_tpm,0)}
g2_data_offset	\${_P(g2_data_offset,0)}
g3_threads	\${_P(g3_threads,0)}
g3_startup_delay	\${_P(g3_startup_delay,0)}

☐ Run Thread Groups consecutively (i.e. run groups one at a time)

☐ Run tearDown Thread Groups after shutdown of main threads

☐ Functional Test Mode (i.e. save Response Data and Sampler Data)

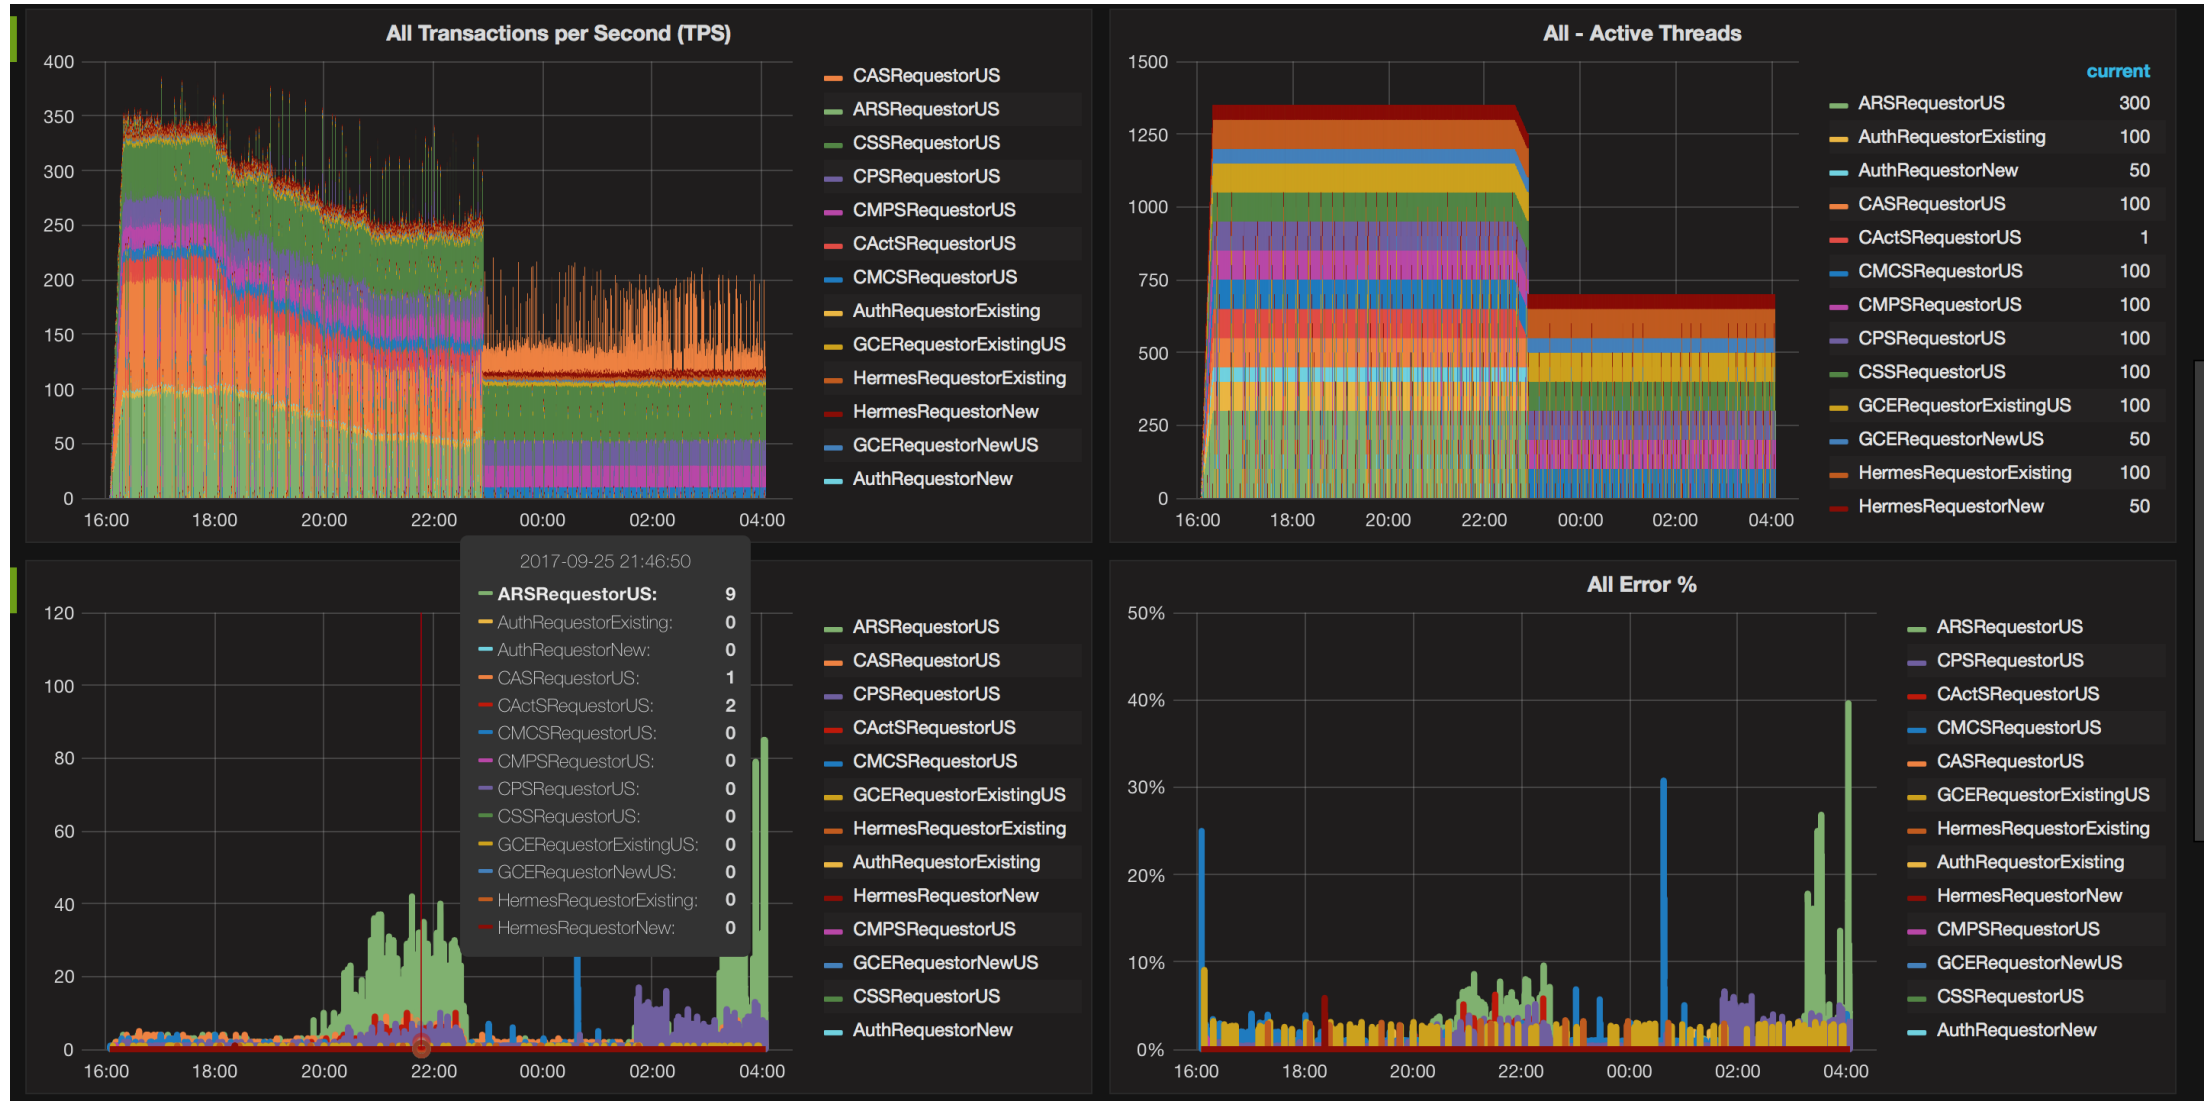
Selecting Functional Test Mode may adversely affect performance.

Add directory or jar to classpath

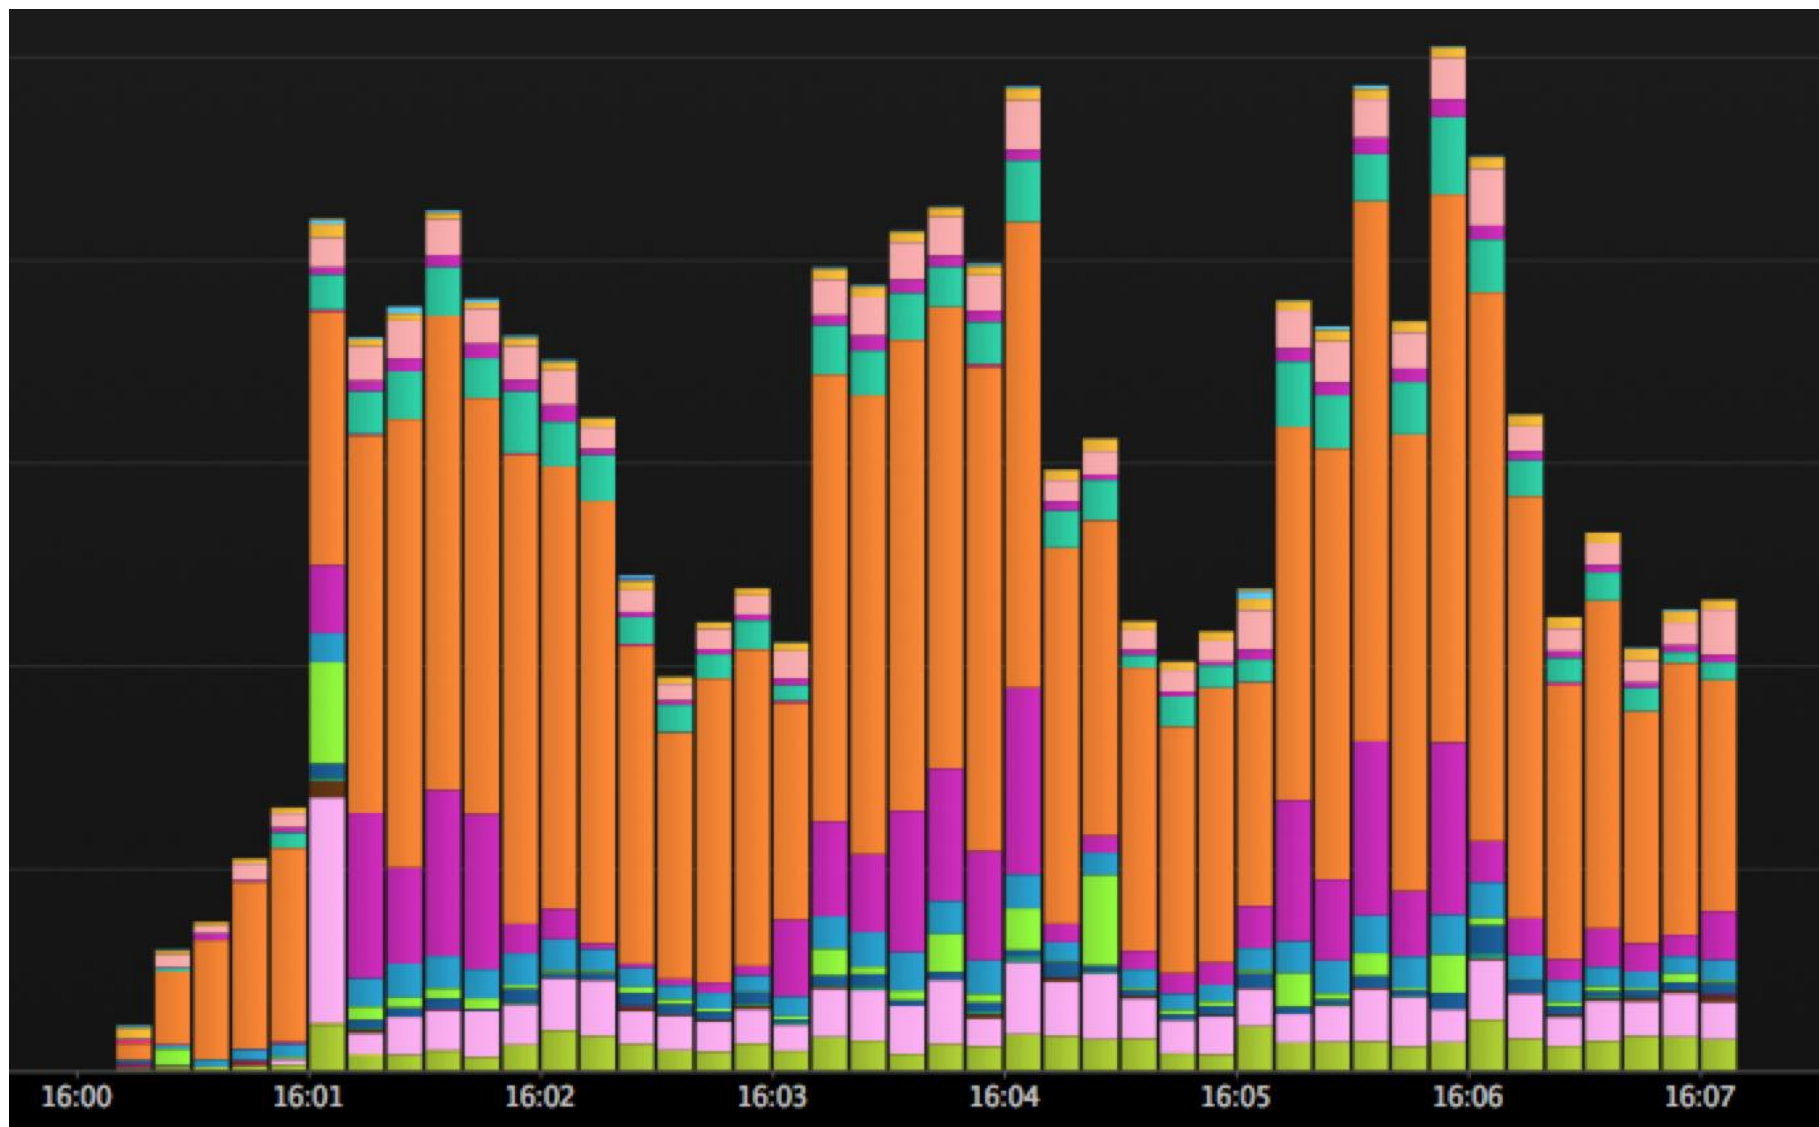
REVIEW: JMETER CUSTOMIZATIONS

- By Extending the testing tool we achieve the following:
 - Portability to different environments
 - Adaptability to serve multiple, different scenarios
 - Improved code-reuse and collaboration
 - Results extensibility to external data stores


DEMONSTRATION: PERFORMANCE VISUALIZATION



DEMONSTRATION: PERFORMANCE VISUALIZATION



DEMONSTRATION: PERFORMANCE VISUALIZATION




REVIEW: PERFORMANCE VISUALIZATIONS

We reviewed the following in Grafana and Dynatrace:

- Connection to a time-series database (e.g. InfluxDB)
- Separation of dashboards by app
- Combined dashboard for all traffic
- Drill-down capability: digging deeper into metrics
- Consistent coloring, helps to remember metrics
- Advanced features for viewing timeframe and refresh
- Advanced metrics querying, granularity and aggregation

DEMONSTRATION: NOTIFICATION AND COMMUNICATION



PayPal

Mark Tomlinson

⋮

All Unreads

🔍

All Threads

★

Starred

help-credit-lnp

lt-1054

lt-1056

lt-1069

lt-db-sync

1

lt-ipt-domestic

lt-ipt-international

lt-pre-push

🔒 timonium-lt-te...

3

John Pfeifer IV

saiprasad setty

William Safee

1

+

Channels

+

#lt-ipt-domestic

★ | 27 | 0 | Add a topic

📞

ℹ️

⚙️


🔍 Search

@


★

⋮


Yesterday




saiprasad setty 9:27 AM
— CNR + flashback
in **LT01**




saiprasad setty 9:33 AM ☆
uploaded this image: **Not a perfect Soak Run**






saiprasad setty 9:33 AM
ARS requestor failed and error % was touching 40% for quite good amount of time
investigating the last night run



Terrance APP 9:38 AM

DEMONSTRATION: NOTIFICATION AND COMMUNICATION

 Load Test / LOADTEST-1054


Test Artemis utils - 2.2.0-develop-201709141749130353

EditCommentAssignMoreOn HoldResolve



Export

Activity

AllCommentsWork LogHistoryActivityTransitions

>  _CPT_USER added a comment - 5 days ago LOAD TEST SUMMARY: LOADTEST-1054 BaselineTest R...

5 older comments

  _CPT_USER added a comment - Yesterday

LOAD TEST SUMMARY: [LOADTEST-1054](#) BaselineTest RUN5

Timeframe: Sep 27, 2017 02:41AM to Sep 27, 2017 02:58AM GMT

Project Link: <https://jira.paypal.com/browse/LOADTEST-1054>

Comment:

GRAFANA DASHBOARDS

- <https://lvs1-cpesg-u01.lvs.its.paypalcorp.com:3000/dashboard/db/1-ipt-overview?from=1506480118690&to=1506481104290>
- <https://lvs1-cpesg-u01.lvs.its.paypalcorp.com:3000/dashboard/db/2-cpu-overview?from=1506480118690&to=1506481104290>
- <https://lvs1-cpesg-u01.lvs.its.paypalcorp.com:3000/dashboard/db/3-authengine-dashboard?>

Time Tracking

Estimated: 1d

Remaining: 1d

Logged: Not Specified

Agile

[View on Board](#)

REVIEW: NOTIFICATION AND COMMUNICATION

We reviewed the following in Slack:

- Individual channels for apps/teams
- Combined channels for all traffic
- Real-time notification of test events
- Bot-identity ..making the automation more personable


We reviewed the following in JIRA Comments:

- Automatic summary comment from the test results
- Links to drill-down into the results tools/dashboards

OTHER TOOLS

Consider a few other things you might need:

- Managing test data state continuously (flashbacks/restore)
- Service Virtualization - simulating external systems
- Support for different load generation tools
- Application log analysis – grepping for exceptions/errors
- Storage for all the results data...can get HUGE over-time




WHERE TO START...

WHERE TO START...

DevOps Handbook Part 2 “Where to Start”

- Greenfield vs. “Brownfield”
- Systems of Record vs. Engagement
- Sympathetic and Innovative Teams
- Expanding DevOps/Continuous-ness



Consider the following questions at the start:


- are your stakeholders/customers ready for continuous feedback? How will they react? Will they value faster feedback?
- do you use information radiators/visual mediums to deliver the performance trends visually?
- do you have notification systems (email, chatops, Slack) for communication to results and trends?
- are your test environments configured to support unattended, non-stop test runs?
- do you have the tools and licenses ready for more executions?

GRADUAL ADOPTION:

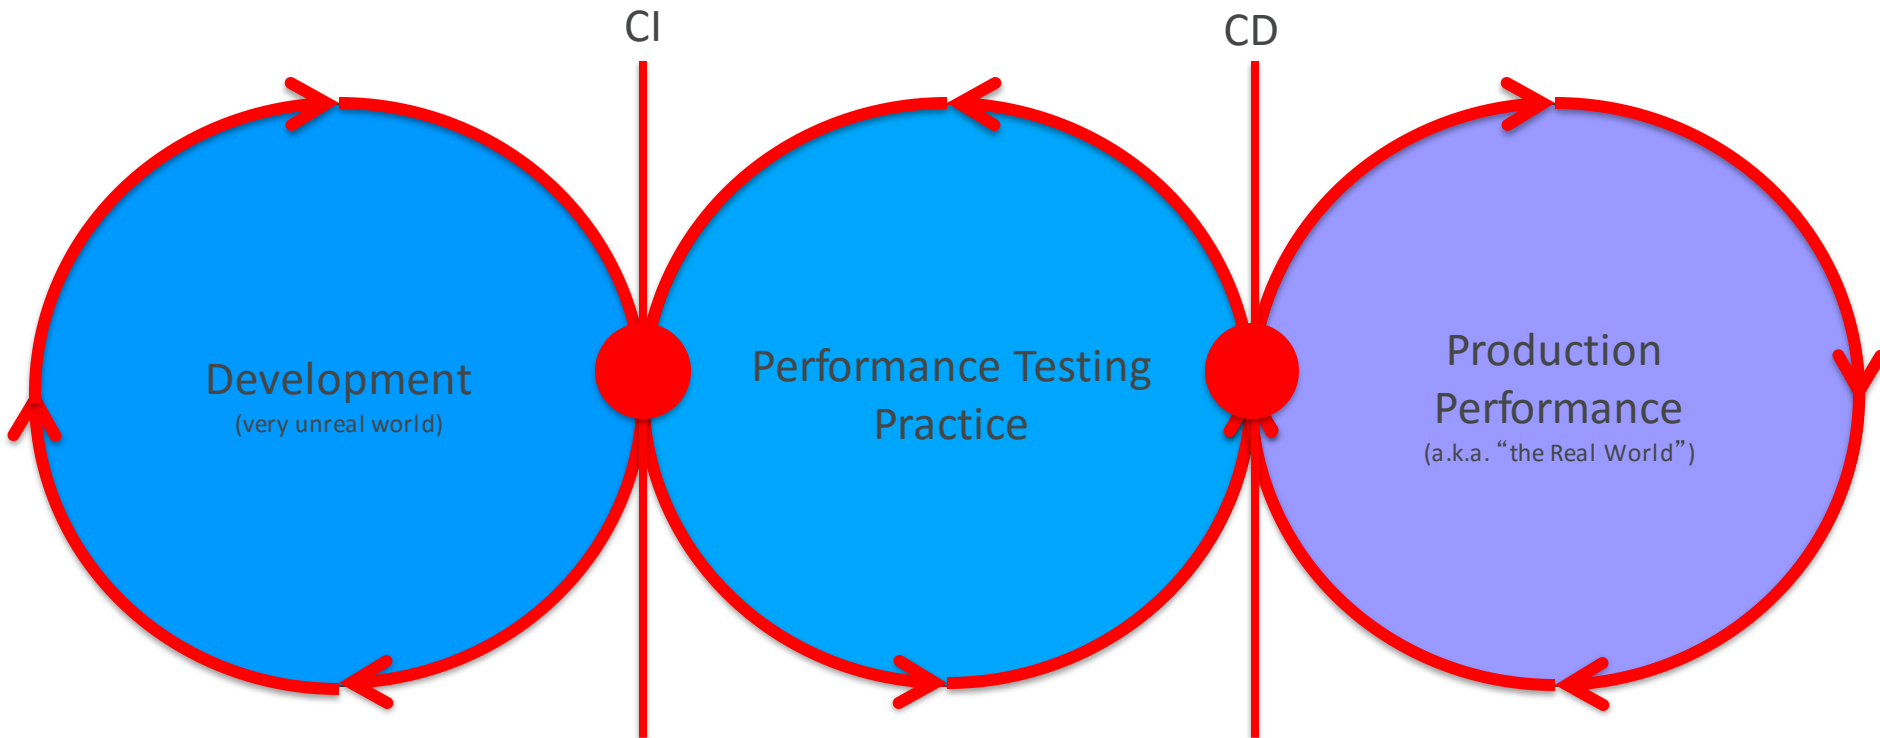
Slowly introduce continuous performance:

- Consider just 1 day a week to run non-stop testing
- Consider a team that already dominates your schedule
- It takes time for people to adjust to frequent feedback, how to react and behave
- You will also receive feedback on your tooling/visuals – take time to improve

Start with Dev...



...then add Production



SUMMARY

- Performance can be measured continuously
- Know the difference: promotion and feedback
- One-time performance tests vs. repeated testing
- Absolute vs. relative measurements
- Performance decisions are made continuously
- Decisions and analysis are based on data
- Find ways to capture metrics across the lifecycle

Mark Tomlinson, Performacologist
mtomlins@perfbytes.com
mark-on-task.blogspot.com @mark_on_task
PerfBytes Podcast: @perfbytes

Thank you

