

DevOps & Software Delivery in a Global Pandemic








John Allspaw

@allspaw

Normal 0%



Work-as-imagined versus work-as-done



3:00 AM · Apr 28, 2016 · Twitter for iPhone

performance described
VS
performance derived



Jeremy Meiss
Director, DevRel & Community



 @IAmJerdog

Dataset

257 mil+

workflows

44,000+

orgs

290,000+

projects

1,000x

Larger than surveys

Four classic metrics





Deployment
frequency

Lead time
to change

Change failure
rate

Recovery
from failure
time

CI/CD Benchmarks for high performance teams

		Suggested Benchmarks
	Throughput The average number of workflow runs per day	Merge on any pull request
	Duration The average length of time for a workflow to run	10 minutes
	Mean time to recovery The average time between failures & their next success	Under 1 hour
	Success rate The number of successful runs / the total number of runs over a period of time	90% or better on default branch



The Data




Photo by: Matthew Henry


Throughput

the average number of workflow runs per day


Throughput



Throughput



Throughput



***Most teams are not deploying
dozens of times per day***

Image by Pawan Kolhe from Pixabay




Duration

the length of time it takes for a workflow to run

Duration

Image by [Pawan Kolhe](#) from [Pixabay](#)


Duration



Duration

Image by [Pawan Kolhe](#) from [Pixabay](#)

Duration



Duration

Image by [Pawan Kolhe](#) from [Pixabay](#)

Duration






Photo by Brett Sayles from Pexels

Mean time to recovery

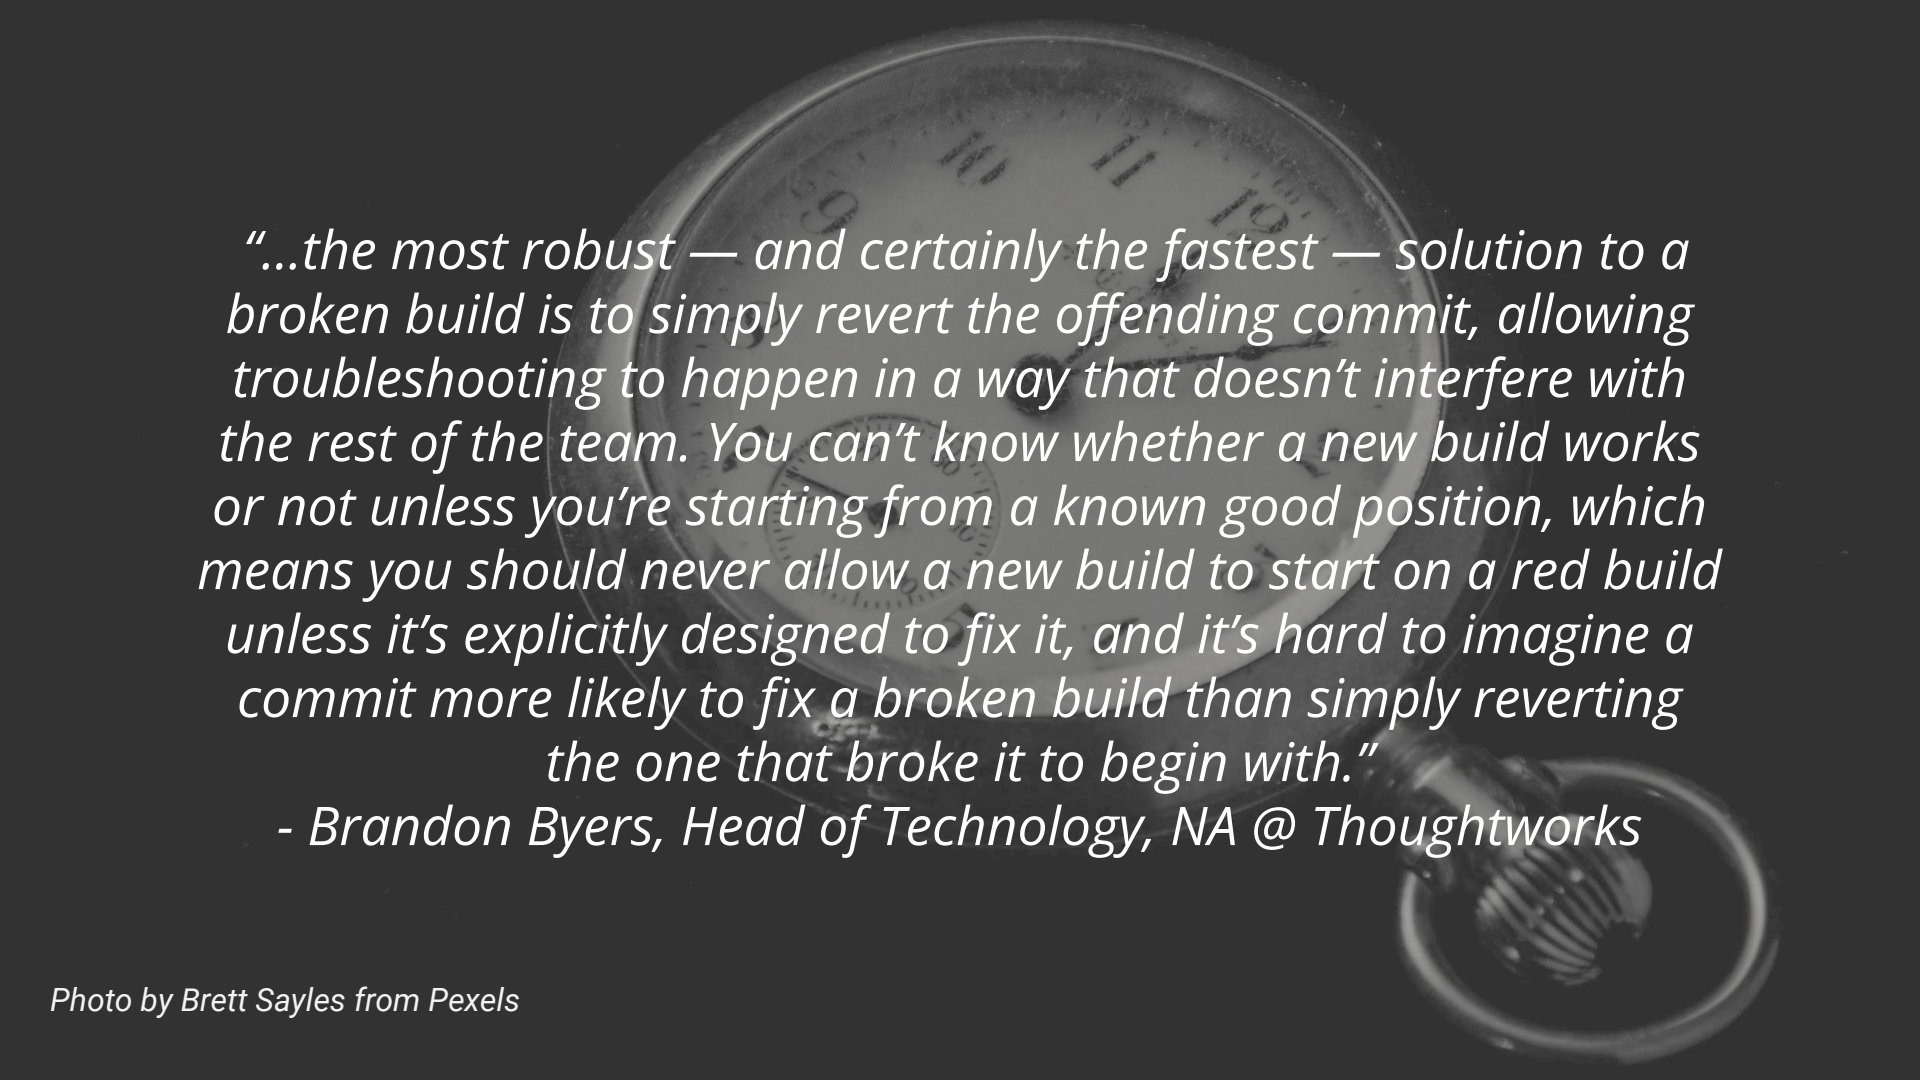


average time between a pipeline's failure and its next success

Mean time to recovery



shortest MTTR ∞ Duration


A vintage pocket watch with a lightbulb inside, symbolizing time and ideas. The watch face is visible, showing numbers and hands. The lightbulb is positioned at the bottom right, with its base visible. The background is dark, making the watch and lightbulb stand out.

"...the most robust — and certainly the fastest — solution to a broken build is to simply revert the offending commit, allowing troubleshooting to happen in a way that doesn't interfere with the rest of the team. You can't know whether a new build works or not unless you're starting from a known good position, which means you should never allow a new build to start on a red build unless it's explicitly designed to fix it, and it's hard to imagine a commit more likely to fix a broken build than simply reverting the one that broke it to begin with."

- Brandon Byers, Head of Technology, NA @ Thoughtworks


Recovery Time

TTR




Recovery Time

TTR



Recovery Time

TTR



Recovery Time

TTR




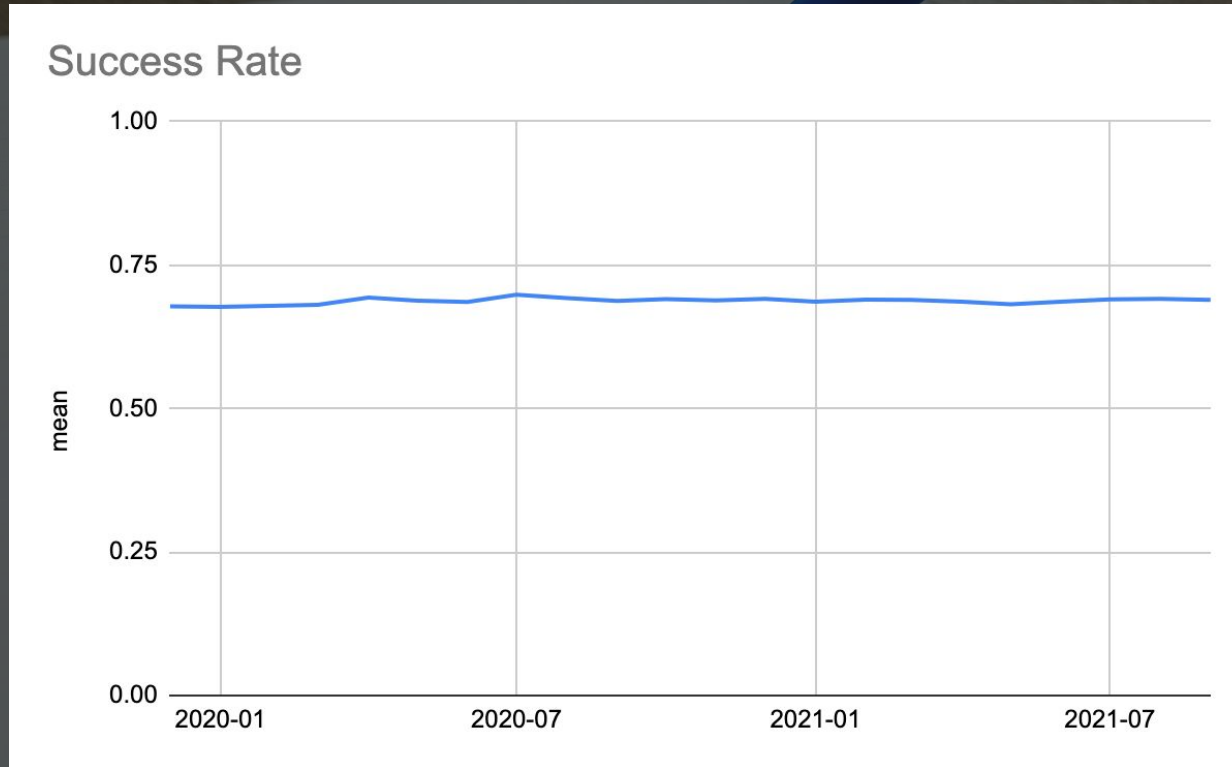


Photo by Lukas from Pexels

Success rate


The number of passing runs \div total number of runs
over a period of time

Success rate




Success rate

Success Rate



Success rate



	2019 (median)	2020 (median)	This Year (median)	Benchmark
Duration The average length of time for a workflow to run	3.38 min	3.96 min	3.7 min	5-10 minutes
TTR The average time between failures & their next success	52.5	55.11	73.6 min	< 60 minutes
Success rate The number of successful runs / the total number of runs over a period of time	60%	61%	77%	Average should be +90% on default branch
Throughput The average number of workflow runs per day	0.80/day	0.70/day	1.43/day	As often as your business requires - not a function of your tooling

Extra Insights



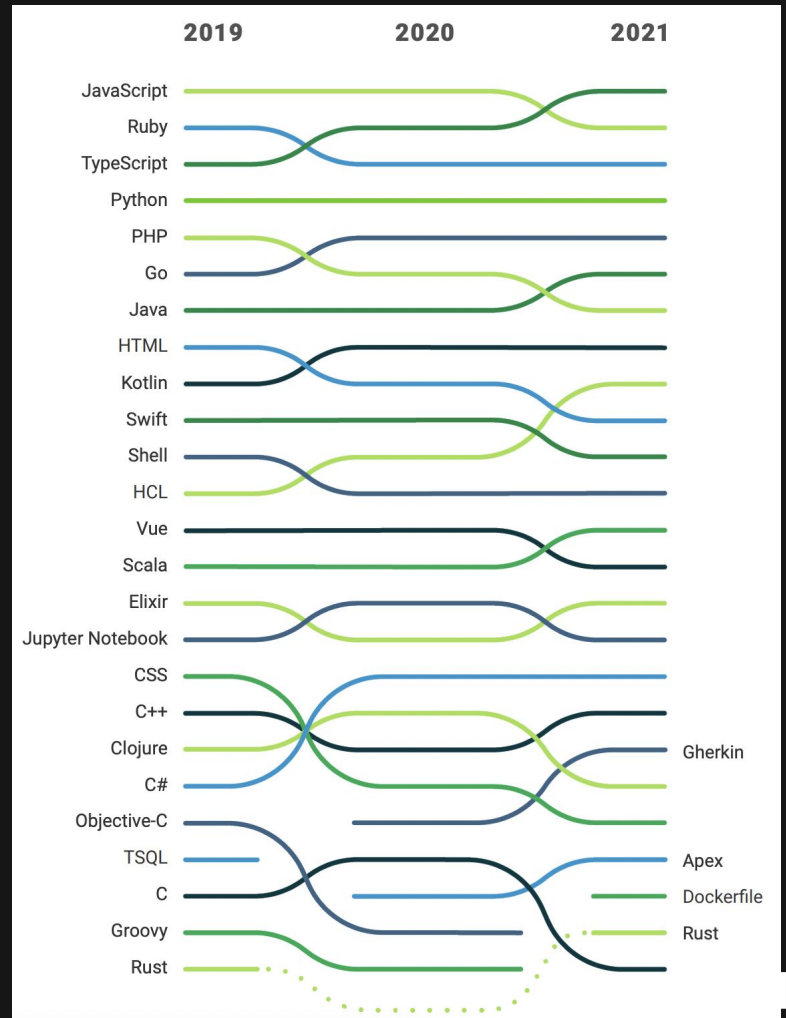
202x has been
a year.

**“Don’t deploy on Friday” is
not a thing.**

“Don’t Deploy on Friday” is not a thing

- 70% less **Throughput** on weekends
- 11% less **Throughput** on Friday (UTC)
- 9% less **Throughput** on Monday (UTC)

Language shifts over the last few years



Duration

- | | | |
|---------------|-----------------------|-------------------|
| 1. Batchfile | 10. Dockerfile | 19. Go |
| 2. SaltStack | 11. PLSQL | 20. Starlark |
| 3. Makefile | 12. Jinja | 21. API Blueprint |
| 4. Smarty | 13. Elm | 22. Roff |
| 5. Jsonnet | 14. Lua | 23. HTML |
| 6. Shell | 15. Liquid | 24. R |
| 7. Mustache | 16. VCL | 25. Python |
| 8. HCL | 17. Open Policy Agent | |
| 9. FreeMarker | 18. Groovy | |

MTTR

- | | | |
|---------------|----------------------|-----------|
| 1. Gherkin | 10. Python | 19. Java |
| 2. HCL | 11. PHP | 20. Swift |
| 3. JavaScript | 12. Perl | 21. Apex |
| 4. Go | 13. Shell | 22. CSS |
| 5. Clojure | 14. Kotlin | 23. C++ |
| 6. C# | 15. Elixir | 24. Rust |
| 7. Vue | 16. HTML | 25. C |
| 8. TypeScript | 17. Scala | |
| 9. Ruby | 18. Jupyter Notebook | |

Success Rate

- | | | |
|---------------|----------------------|--------------|
| 1. Dockerfile | 10. C# | 19. PLpgSQL |
| 2. Vue | 11. HCL | 20. Kotlin |
| 3. Shell | 12. JavaScript | 21. Ruby |
| 4. Go | 13. Elixir | 22. Makefile |
| 5. SCSS | 14. Clojure | 23. Groovy |
| 6. HTML | 15. Jupyter Notebook | 24. TSQL |
| 7. TypeScript | 16. Java | 25. Gherkin |
| 8. PHP | 17. Scala | |
| 9. Python | 18. CSS | |





Throughput

- | | | |
|---------------------|----------------|------------------|
| 1. Hack | 10. Apex | 19. Blade |
| 2. Slim | 11. TypeScript | 20. Scala |
| 3. Elm | 12. Swift | 21. Python |
| 4. Mustache | 13. Ruby | 22. LookML |
| 5. Haskell | 14. Dart | 23. Lua |
| 6. Jinja | 15. Elixir | 24. CoffeeScript |
| 7. Gherkin | 16. Go | 25. Clojure |
| 8. Jsonnet | 17. C# | |
| 9. Jupyter Notebook | 18. Kotlin | |

Vertical splits

Duration in minutes		MTTR in minutes		Throughput		Success Rate	
Industry	50th percentile	Industry	50th percentile	Industry	50th percentile	Industry	Average
Health Care Providers & Services	4.46	Hotels, Restaurants & Leisure	60.4	Health Care Providers & Services	1.68	Media	71.9%
Hotels, Restaurants & Leisure	3.96	Internet Software & Services	72.1	Hotels, Restaurants & Leisure	1.68	Consumer Discretionary	70.1%
Professional Services	3.88	Diversified Financial Services	73.2	Real Estate	1.68	Real Estate	69.9%
Internet Software & Services	3.80	Media	84.0	Consumer Discretionary	1.61	Internet Software & Services	69.0%
Real Estate	3.72	Professional Services	85.0	Diversified Financial Services	1.61	Professional Services	68.5%
Consumer Discretionary	3.66	Health Care Providers & Services	99.5	Internet Software & Services	1.46	Diversified Financial Services	68.0%
Diversified Financial Services	3.55	Consumer Discretionary	273	Professional Services	1.04	Health Care Providers & Services	66.5%
Media	3.21	Real Estate	318	Media	1.00	Hotels, Restaurants & Leisure	61.1%

Elite Performer validation

Software delivery performance metric	Elite	High	Medium	Low
 Deployment frequency For the primary application or service you work on, how often does your organization deploy code to production or release it to end users?	On-demand (multiple deploys per day)	Between once per week and once per month	Between once per month and once every 6 months	Fewer than once per six months
 Lead time for changes For the primary application or service you work on, what is your lead time for changes (i.e., how long does it take to go from code committed to code successfully running in production)?	Less than one hour	Between one day and one week	Between one month and six months	More than six months
 Time to restore service For the primary application or service you work on, how long does it generally take to restore service when a service incident or a defect that impacts users occurs (e.g., unplanned outage or service impairment)?	Less than one hour	Less than one day	Between one day and one week	More than six months
 Change failure rate For the primary application or service you work on, what percentage of changes to production or released to users result in degraded service (e.g., lead to service impairment or service outage) and subsequently require remediation (e.g., require a hotfix, rollback, fix forward, patch)?	0%-15%	16%-30%	16%-30%	16%-30%

50th percentile on CircleCI fit into the “Elite performer” category on the 2021 State of DevOps report

2020 Report



<https://circle.ci/ssd2020>

Full 2022 Report



<https://circle.ci/ssd2022>

Thank you.

For feedback and swag: circle.ci/jeremy



Timeline.jerdog.me



IAmJerdog



jerdog



/in/jeremymeiss



 @IAmJerdog