

What We Learned from Four Years of Sciencing the Crap Out of DevOps

Nicole Forsgren, PhD
CEO and Chief Scientist
DevOps Research and Assessment (DORA)

@nicolefv



Outline

How to make your data suck less

- Writing good survey questions
- Making sure the survey questions are good
- with SCIENCE
- (These methods apply to your system and log data)

What we found... that we did (AND didn't) expect
Things about Continuous Delivery
Things about Management

Not all data is created equal

Who here thinks surveys are sh*t?




@nicolefv

Not all data is created equal

Who here thinks surveys are sh*t?

Who here LOVES the data from their log files?

What is a Latent Construct?



@nicolefv

We use

PSYCHOMETRICS

to make our survey data good*

**or give us a reasonable assurance that it's telling us what we think it's telling us (& some of this can also apply to your log data)*

@nicolefv

Psychometrics includes:

Construct creation (manual)

- When possible: use previously validated constructs
- Based on definitions and theory, carefully and precisely worded, card sorting task, pilot tested

Construct evaluation (statistics)

- Establishing Validity: discriminant and convergent
- Establishing Reliability

Psychometrics Writing

Example: Culture

- Does it matter to our study?
 - More than just intuition?
- What KIND of culture?
 - National identity and norms
 - Adaptive culture
 - Value learning (2014 study)
 - Value information flow and trust (2014 and 2015 studies -- Westrum culture)

Psychometrics Writing

Example: Culture

- Does it matter to our study?
 - More than just intuition?
- What KIND of culture?
 - National identity and norms
 - Adaptive culture
 - Value learning (2014 study)
 - Value information flow and trust (2014 and 2015 studies -- Westrum culture)


Try writing items yourself!

Westrum typology


Use strong statements with clear language.

| Pathological <i>Power-oriented</i> | Bureaucratic <i>Rule-oriented</i> | Generative <i>Performance-oriented</i> |
|---------------------------------------|--------------------------------------|---|
| Low cooperation | Modest cooperation | High cooperation |
| Messengers shot | Messengers neglected | Messengers trained |
| Responsibilities shirked | Narrow responsibilities | Risks are shared |
| Bridging discouraged | Bridging tolerated | Bridging encouraged |
| Failure leads to scapegoating | Failure leads to justice | Failure leads to inquiry |
| Novelty crushed | Novelty leads to problems | Novelty implemented |

Westrum Culture Items



Found
to be
valid &
reliable



Predictive of
IT Performance &
Organizational
Performance

- On my team, information is actively sought.
- On my team, failures are learning opportunities, and messengers of them are not punished.
- On my team, responsibilities are shared.
- On my team, cross-functional collaboration is encouraged and rewarded.
- On my team, failure causes inquiry.
- On my team, new ideas are welcomed.

Psychometrics Analysis

Example: Notification of Failure

Original
in 2014,
but there
was a
surprise.
Can you
spot it?

At my organization...

- We are primarily notified of failures by reports from customers.
- We are primarily notified of failures by the NOC.
- We get failure alerts from logging and monitoring systems.
- We monitor system health based on threshold warnings (ex. CPU exceeds 100%).
- We monitor system health based on rate-of-change warnings (ex. CPU usage has increased by 25% over the last 10 minutes).

Psychometrics Analysis

Example: Notification of Failure

At my organization...



- We are primarily notified of failures by reports from customers.
- We are primarily notified of failures by the NOC.
- We get failure alerts from logging and monitoring systems.
- We monitor system health based on threshold warnings (ex. CPU exceeds 100%).
- We monitor system health based on rate-of-change warnings (ex. CPU usage has increased by 25% over the last 10 minutes).

More data tests!

Plus, we test to make sure the survey doesn't have other problems.

- Common method variance (CMV) (aka CMB for Bias)
- Early vs. late responders
- Survey drop-off rates and bias

Okay **NOW** we can look at the data and how it relates to each other

@nicolefv

A note about analysis methods

One of three conditions must be met:

1. Longitudinal (no, this is cross-sectional)
2. Randomized, experimental design (no, this is a non-experimental)
3. **Theory-based design**

When this condition was not met, only correlations were tested and reported

KEY FINDING: IT Performance and its behavior

A combination of *throughput* and *stability*

- lead time for changes
- release frequency
- time to restore service
- change fail rate

Forsgren, N., J. Humble (2016). "DevOps: Profiles in ITSM Performance and Contributing Factors." In the Proceedings of the Western Decision Sciences Institute (WDSI) 2016, Las Vegas, NV.





<http://bit.ly/2014-devops-report/>

KEY FINDING: IT performance matters!

“Firms with high-performing IT organizations were twice as likely to exceed their profitability, market share and productivity goals.”

IT Performance is predictive of organizational performance.

Forsgren, N., J. Humble (2016). “The Role of Continuous Delivery in IT and Organizational Performance.” In the Proceedings of the Western Decision Sciences Institute (WDSI) 2016, Las Vegas, NV.



<http://bit.ly/2015-devops-report/>

@nicolefv

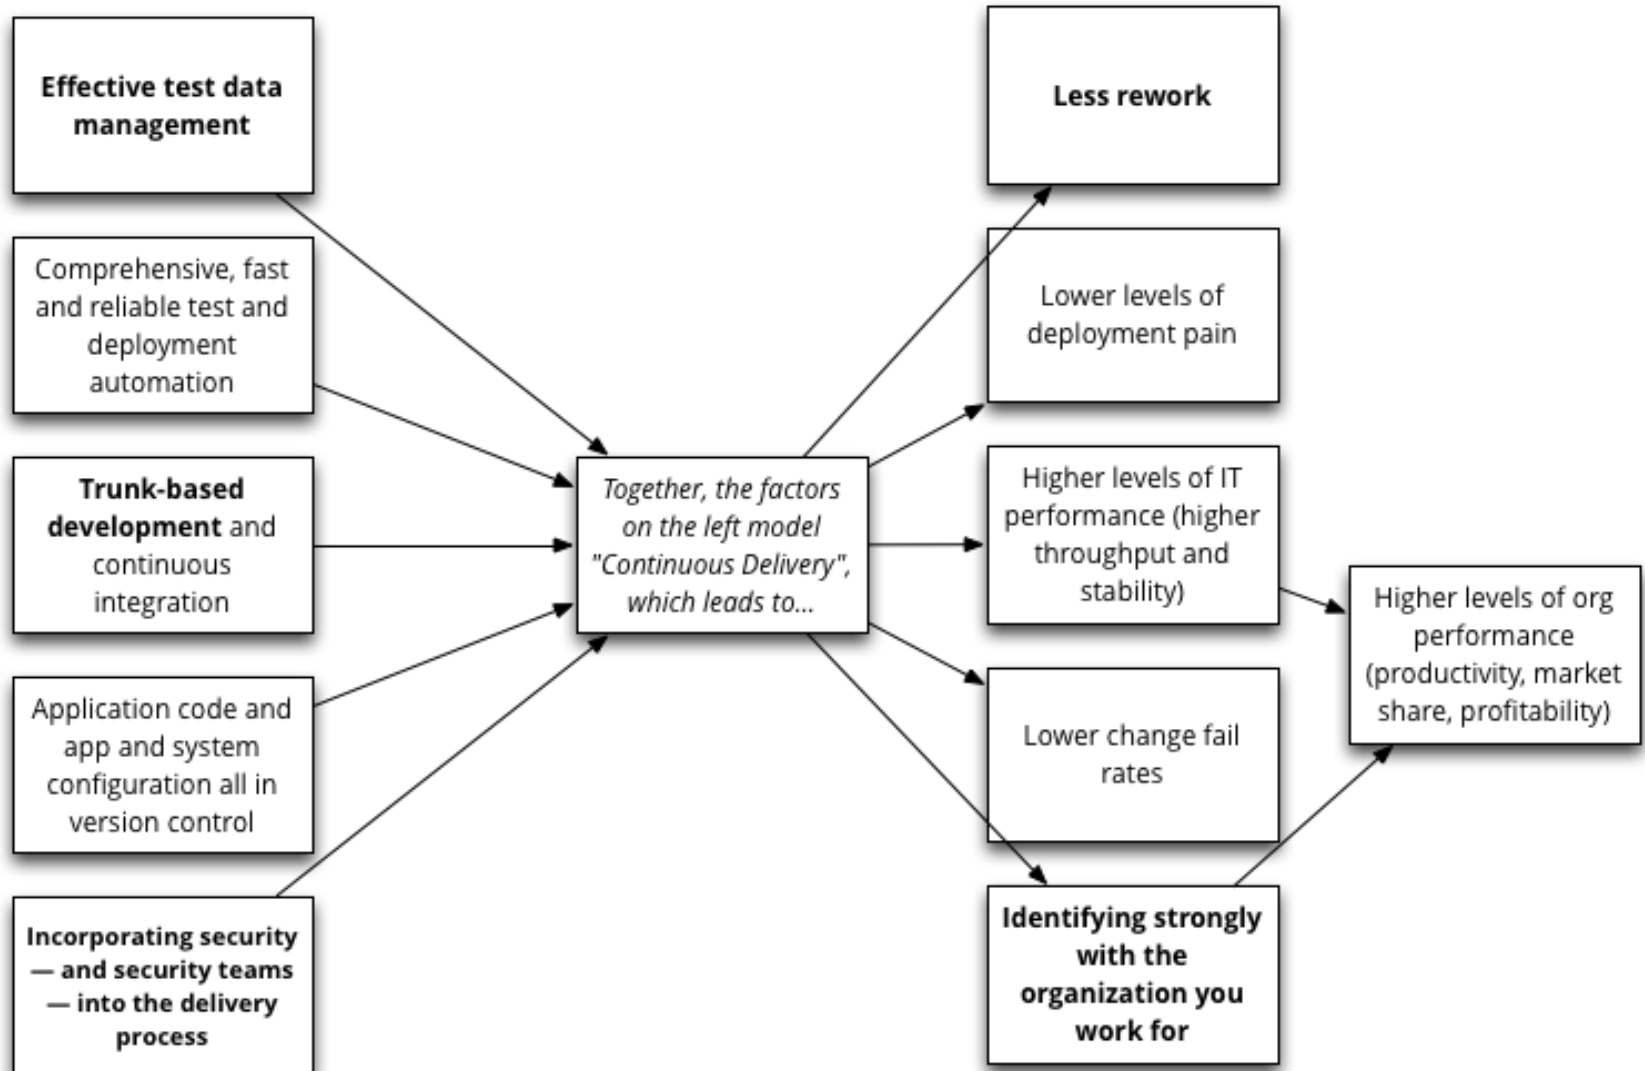
 **DORA**
DEVOPS RESEARCH & ASSESSMENT

2016 IT Performance by Cluster

| | High IT Performers | Medium IT Performers | Low IT Performers |
|---|---|---|---|
| Deployment frequency <i>For the primary application or service you work on, how often does your organization deploy code?</i> | On demand (multiple deploys per day) | Between once per week and once per month | Between once per month and once every 6 months |
| Lead time for changes <i>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 commit to code successfully running in production)?</i> | Less than one hour | Between one week and one month | Between one month and 6 months |
| Mean time to recover (MTTR) <i>For the primary application or service you work on, how long does it generally take to restore service when a service incident occurs (e.g., unplanned outage, service impairment)?</i> | Less than one hour | Less than one day | Less than one day* |
| Change failure rate <i>For the primary application or service you work on, what percentage of the changes either result in degraded service or subsequently require remediation (e.g., lead to service impairment, service outage, require a hotfix, rollback, fix forward, patch)?</i> | 0-15% | 31-45% | 16-30% |

@nicolefv

* Low performers were lower on average (at a statistically significant level), but had the same median as the medium performers.



@nicolefv

some surprises



@nicolefv

Which of these measure effective test practices?

Developers primarily create & maintain acceptance tests

QA primarily create & maintain acceptance tests

Primarily created & maintained by outsourced party

When automated tests pass, I'm confident the software is releasable

Test failures are likely to indicate a real defect

It's easy for developers to fix acceptance tests

Developers share a common pool of test servers to reproduce failures

Developers create on demand test environments

Developers use their own dev environments to reproduce failures

Which of these measure effective test practices?

Developers primarily create & maintain acceptance tests

QA primarily create & maintain acceptance tests

Primarily created & maintained by outsourced party

When automated tests pass, I'm confident the software is releasable

Test failures are likely to indicate a real defect

It's easy for developers to fix acceptance tests

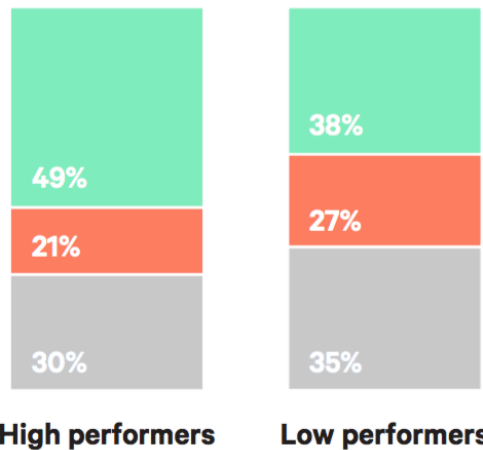
Developers share a common pool of test servers to reproduce failures

Developers create on demand test environments


Developers use their own dev environments to reproduce failures

Key Finding: Unplanned work

New work vs. unplanned work



High performers spend 20% more time on new work than low performers, and 22% less time on unplanned work and rework.



@nicolefv

Surprises with Culture

- We wanted to add additional measures of culture
 - Google study
 - Identity
 - Retain Westrum culture

Identity & Google items

- I am glad I chose to work for this organization rather than another company.
- I talk of this organization to my friends as a great company to work for.
- I am willing to put in a great deal of effort beyond what is normally expected to help my organization to be successful.
- I find that my values and my organization's values are very similar.
- In general, the people employed by my organization are working toward the same goal.
- I feel that my organization cares about me.

Adapted from adapted from Atreyi Kankanhalli, Bernard C.Y. Tan, and Kwok-Kei Wei (2005), "Contributing Knowledge to Electronic Knowledge Repositories: An Empirical Investigation," MIS Quarterly, 29, 113-143.



Identity & Google items

- I am glad I chose to work for this organization rather than another company.
- I talk of this organization to my friends as a great company to work for.
- I am willing to put in a great deal of effort beyond what is normally expected to help my organization to be successful.
- I find that my values and my organization's values are very similar.
- In general, the people employed by my organization are working toward the same goal.
- I feel that my organization cares about me.

Adapted from adapted from Atreyi Kankanhalli, Bernard C.Y. Tan, and Kwok-Kei Wej (2005), "Contributing Knowledge to Electronic Knowledge Repositories: An Empirical Investigation," MIS Quarterly, 29, 113-143.

@nicolefv

Westrum items

1 **Psychological Safety**
Team members feel safe to take risks and be vulnerable in front of each other.

2 **Dependability**

3 **Structure & Clarity**
Team members have clear roles, plans, and goals.

4 **Meaning**
Work is personally important to team members.

5 **Impact**
Team members think their work matters and creates change.

re:Work


Now for management stuff

We all know managing WIP is important, right?



Now for management stuff

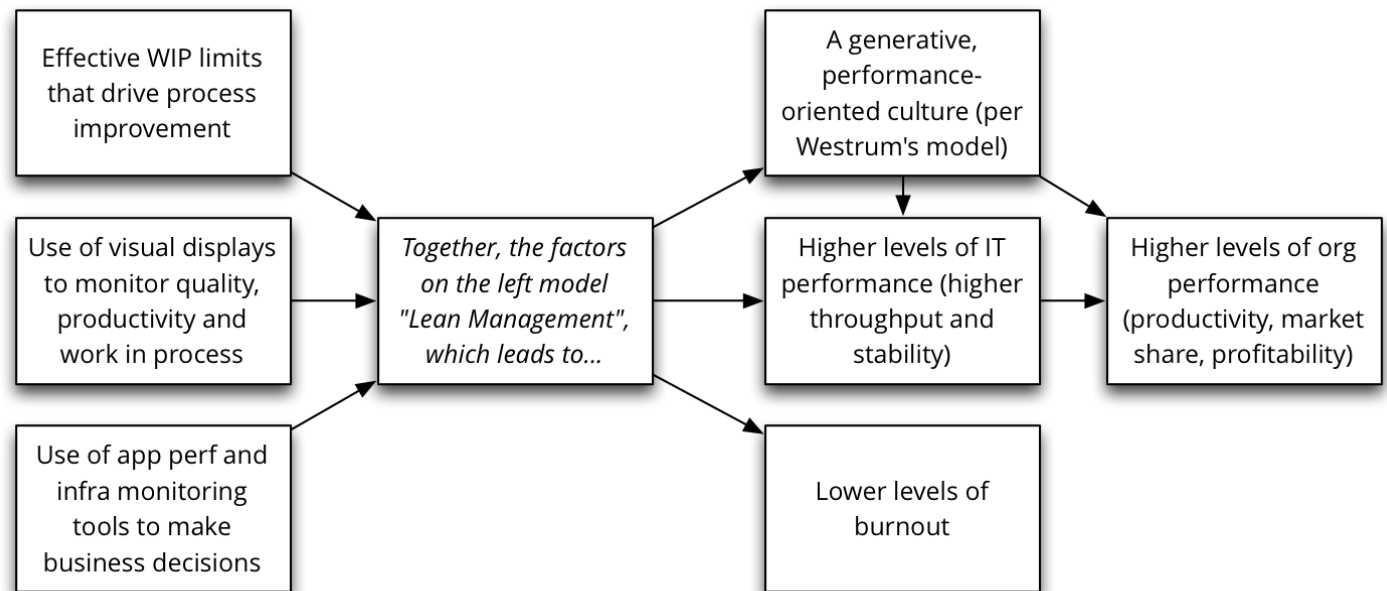
We all know managing WIP is important, right?
Correlation between WIP and ITPerf is negligible



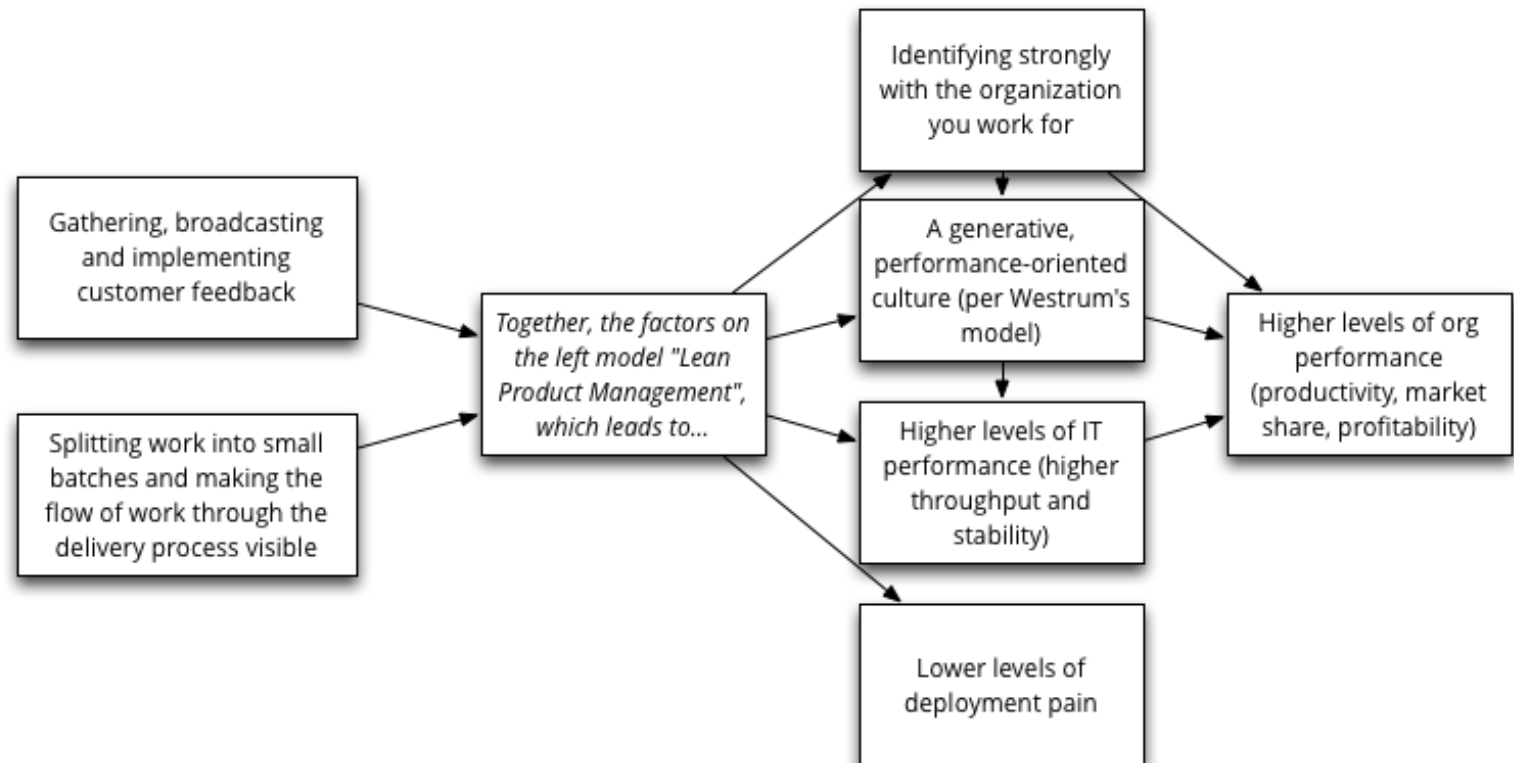
Now for management stuff

We all know managing WIP is important, right?
Correlation between WIP and ITPerf is negligible
What's going on?

Lean management SEM



Also lean product management



Conclusions

- Even if you think it's obvious, TEST WITH DATA.
 - (if the results don't surprise you, you're doing it wrong)
 - (if you don't also confirm some things you expected, you're doing it wrong)
- We CAN have it all, or at least throughput AND stability.
- IT matters (but you have to do it right)
- DevOps culture & practices have a measurable impact on IT & org perf

For more science-ing...



Sign up for our ROI whitepaper &
get peer-reviewed research
devops-research.com

@nicolefv



Thank you

devops-research.com

@nicolefv

