

Machine Learning Development Lifecycle Reliability

Brian Hansen
Production Engineer
Manager, Meta

Yan Yan(PE)
Production Engineer
AdsML PE, Meta

**SRE
CON** ASIA
PACIFIC

**SYDNEY,
AUSTRALIA**

7–9 December, 2022

Overview

MLOps + SDLC \neq ML DLC

Why Now and Why PE/SREs?

ML DLC Deep Dive on Top 3 Challenges

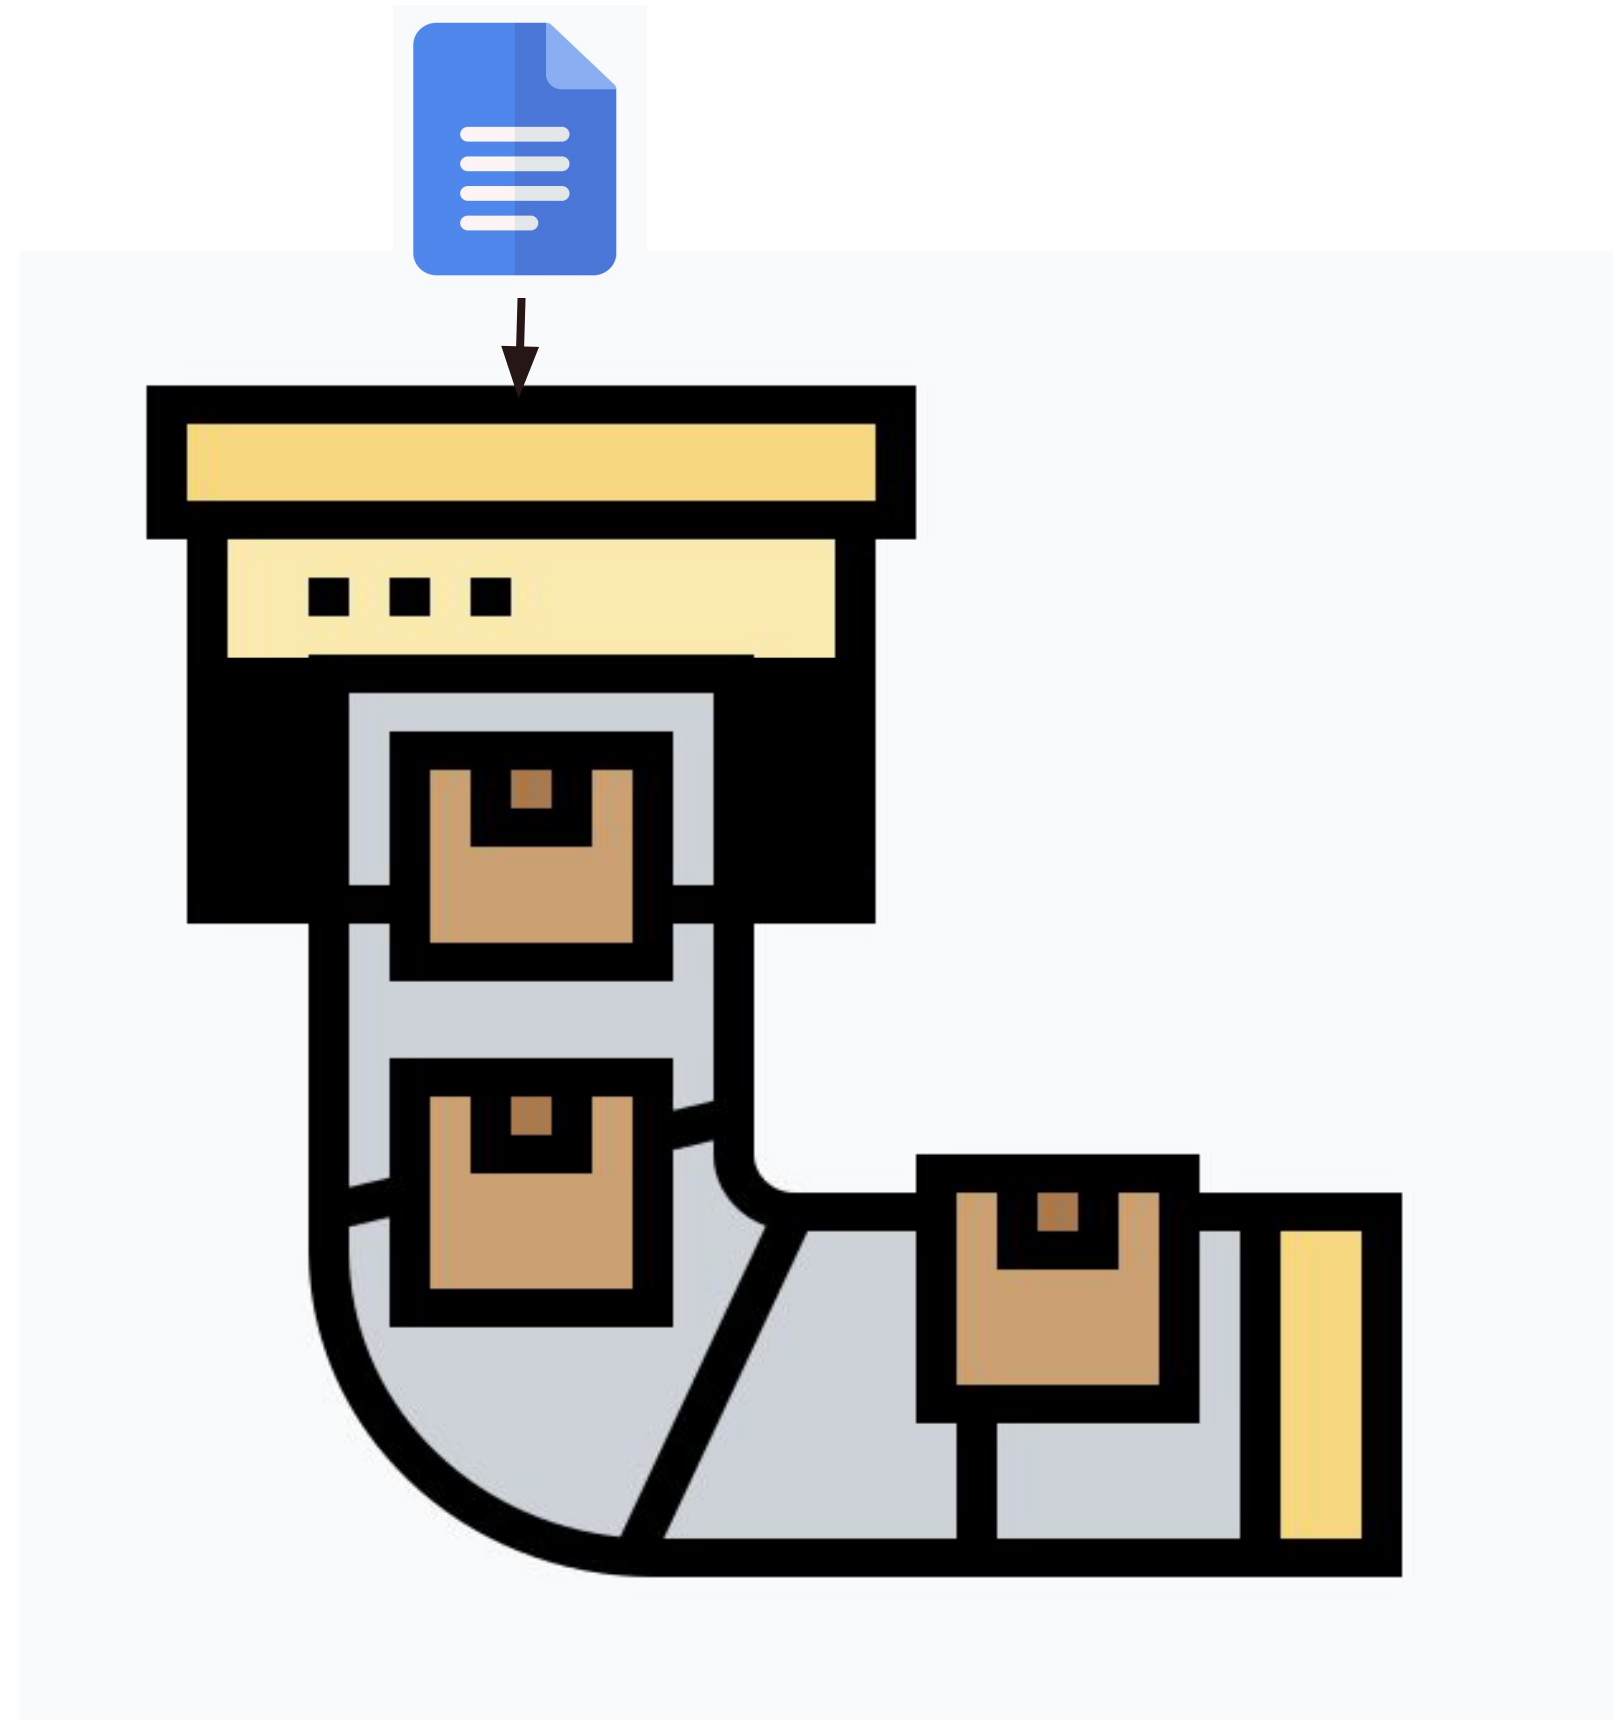
Recap

PE \approx SRE

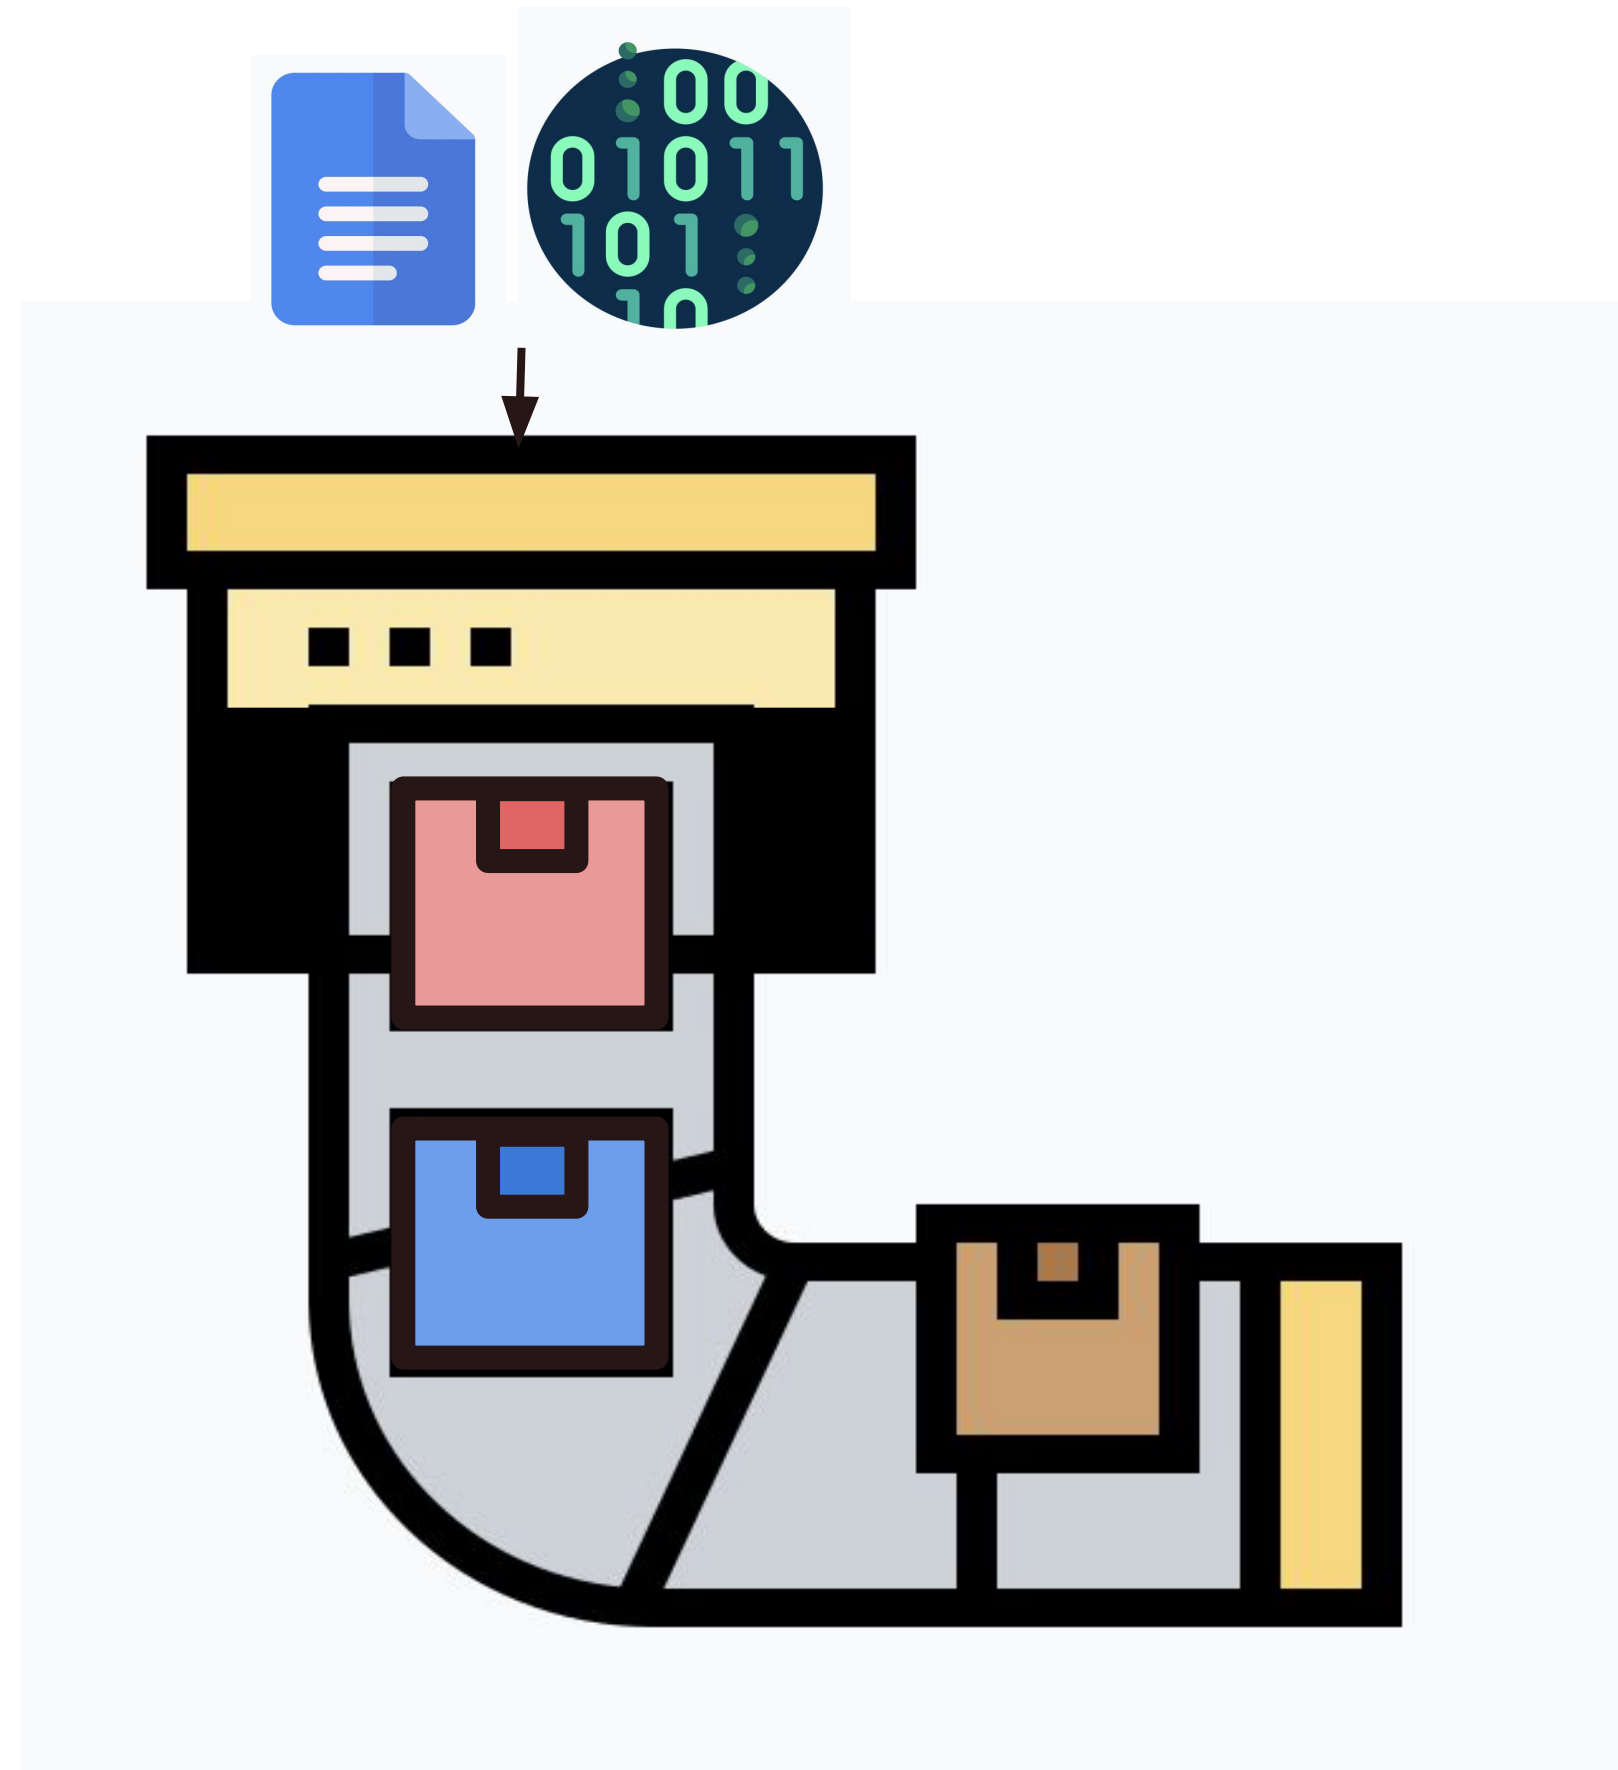
ML Ops

- MLOps seeks to increase **automation** and improve the **quality** of production models, while also focusing on **business** and **regulatory** requirements.
- While MLOps started as a set of best practices, it is slowly evolving into an independent approach to **ML lifecycle management**.

Software



Machine Learning



0 → 1

2011

Nature Reviews Drug Discovery

65% medical studies were
inconsistent when retested

2017

Journals in Economic Sciences

48% peer review
compliance rate

2020

Carnegie Mellon University survey

90% of 1500 ML
researchers believed in
crisis of reproducibility of
ML results

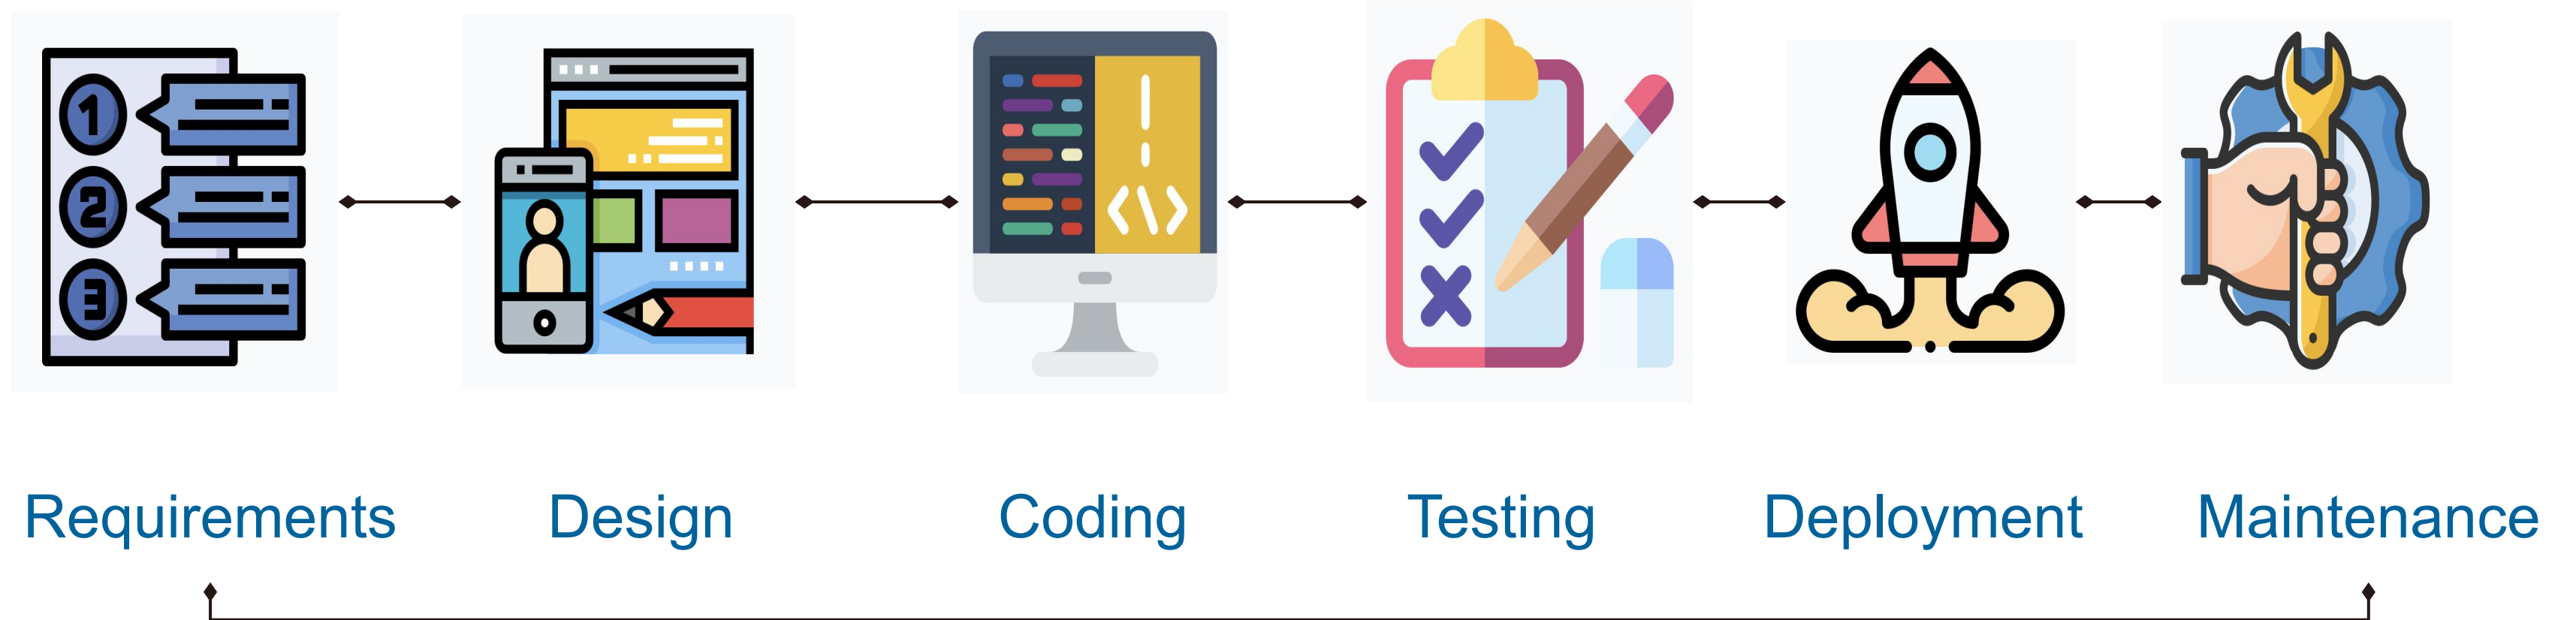
Prinz, F.; Schlange, T.; Asadullah, K. (2011). "Believe it or not: How much can we rely on published data on potential drug targets?". *Nature Reviews Drug Discovery*. 10 (9): 712. doi:10.1038/nrd3439-c1. PMID 21892149.

Vlaeminck, Sven; Podkrajac, Felix (2017-12-10). "Journals in Economic Sciences: Paying Lip Service to Reproducible Research?". *IASSIST Quarterly*. 41 (1-4): 16. doi:10.29173/iq6. hdl:11108/359.

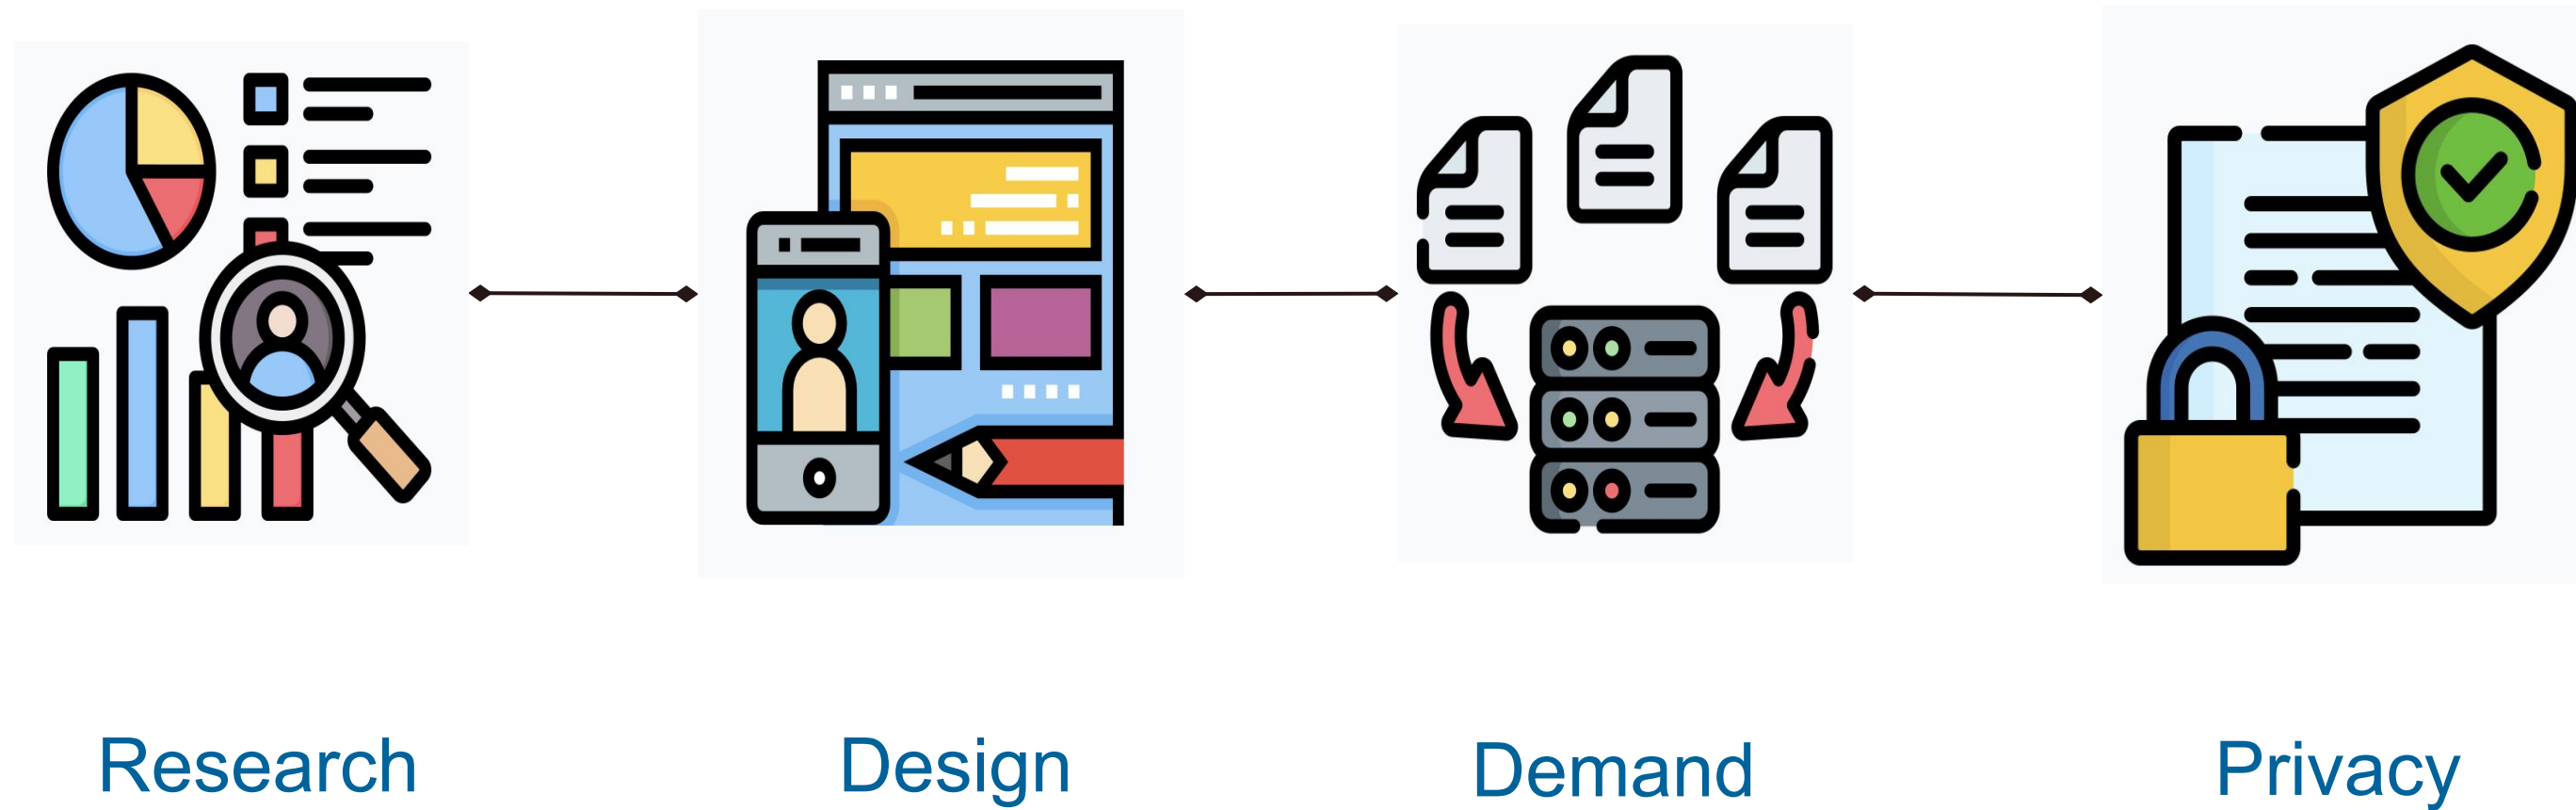
<https://blog.ml.cmu.edu/2020/08/31/5-reproducibility/>

1 → 10,000


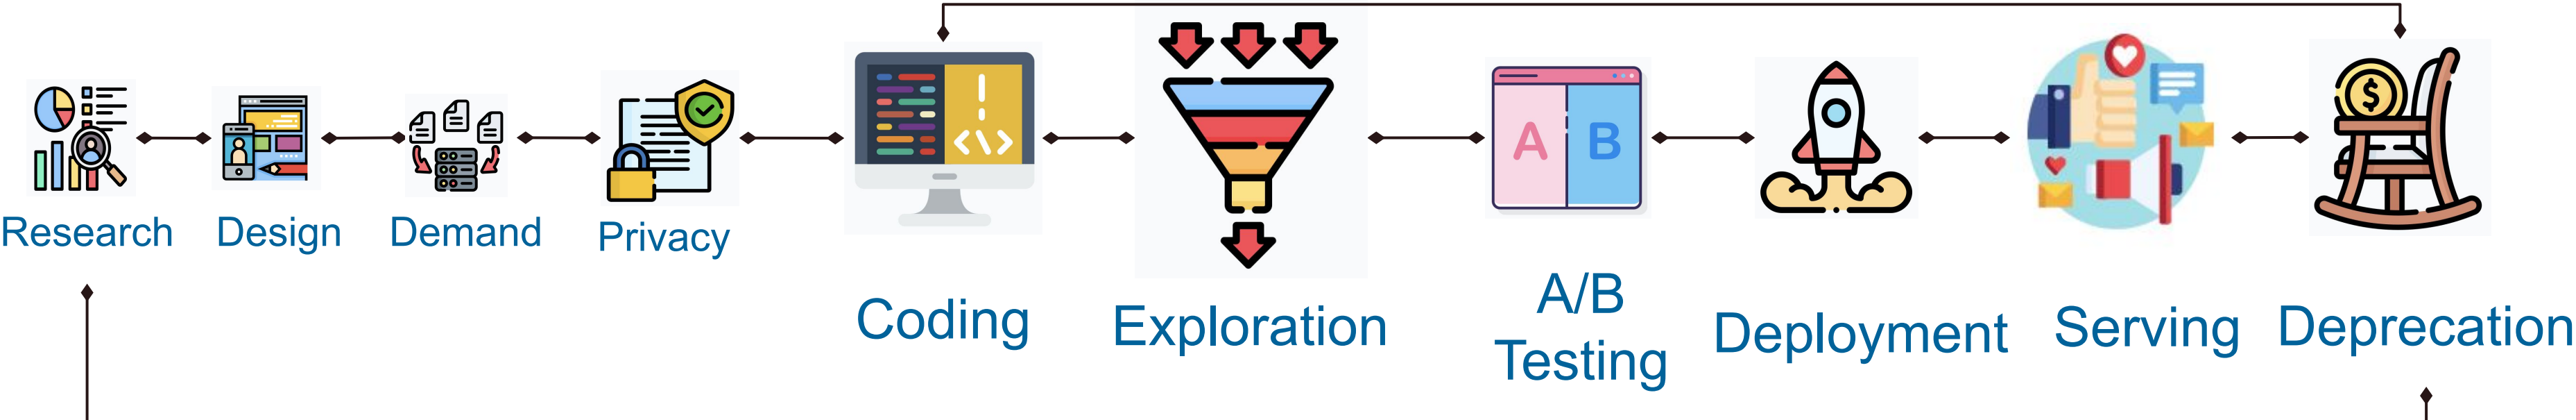
Software Development LifeCycle



ML Development LifeCycle



ML Development LifeCycle



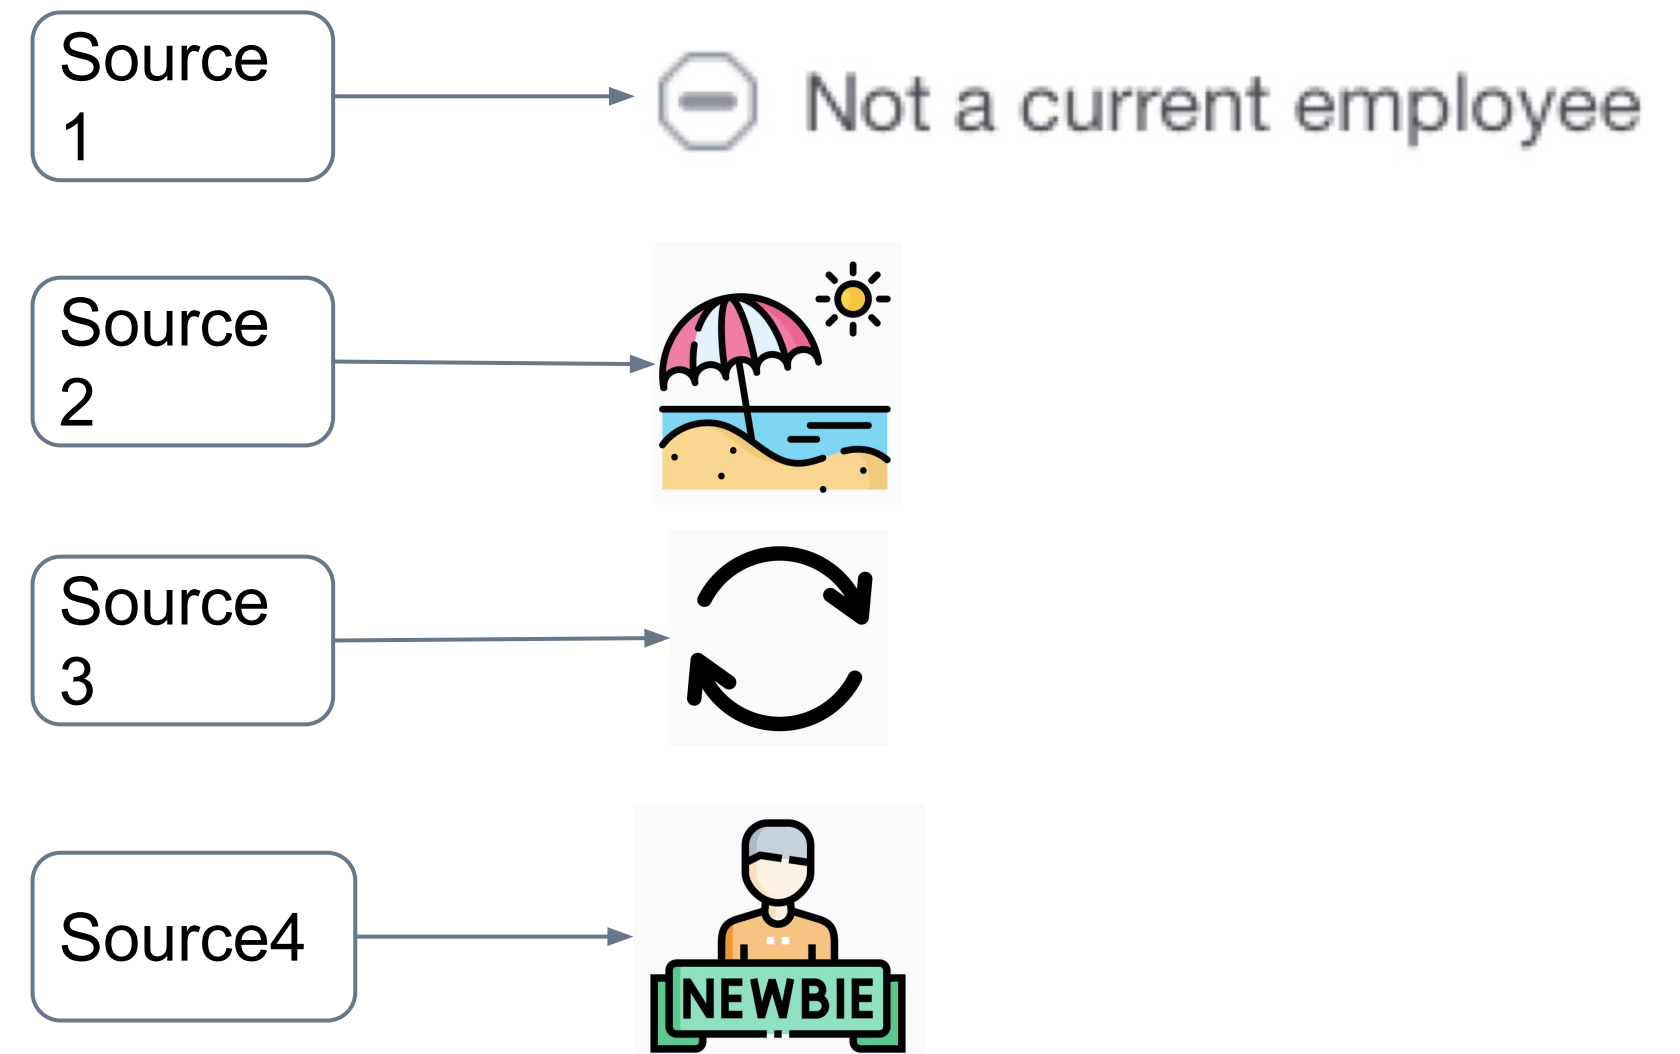

Icon made by [Freepik](#), [Uniconlabs](#), [prettycons](#), [Flat Icons](#), [monkik](#), [juicy_fish](#) from www.flaticon.com

Challenge 1: Observability

FANTASTIC
VO Models

AND WHERE
TO FIND THEM

Ownership Challenges



Ownership Solutions

- Centralize Ownership for consistency
- Replace Owners with Oncall Rotations
- Build Monitoring and Alerting
- Establish Auto Escalation

Production Models

Ads Instagram Reels Marketplace

Top Line Metrics

Table & Metric Data Last Updated: Dec 9, 9999

Total Models ⓘ 300	Pytorch ⓘ 100	Staled ⓘ 2	Not Training ⓘ 8	Average / Target Maintainability Index ⓘ 85 / 90	Launches This Month/Half ⓘ 0 / 100	Monthly Distinct Users ⓘ 2 k
------------------------------	-------------------------	----------------------	----------------------------	--	--	--

Vital Dashboards

Prediction Stability [↗](#) Revenue Stability [↗](#) Calibration Stability [↗](#) Model Staleness [↗](#) ML Diagnosability [↗](#)

Ongoing SEVs

SEV 3 Example: Training crashes

Search... Save

300 search results Download Post Create SEV Column Settings

Ranking Targeting Identity

CoreML

Model Type ⓘ ↑↓	Model ⓘ ↑↓	Launch Time ⓘ ↑↓	Snapshot Staleness ⓘ ↑↓	Maintainability Index ⓘ ↑↓	Revenue Importance ⓘ ↓	Training Type ⓘ ↑↓	Framework ↑↓	Architecture ↑↓	Package ↑↓	Fallback ⓘ ↑↓	Owner ⓘ ↑↓	Oncall ↑↓
example_model_type ↕	999 ▼ History Recent 10: 888 777 666 555 444 333 222 111	September 9, 9999	4h ago SLA: 14 hours	88	7%	Online ↗	Pytorch	SPARSENN	example_pkg:200	900	Example Owner	Current Oncall ads_coreml_m...

Screenshot with Example Data

North Star: One-stop shop of ALL Production Models

Benefits:

- Identify and mitigate incidents faster
- Build Centralized Monitoring and Alerting
- Integrate with other systems easily



ML artifacts




ML artifacts




Production ML Artifacts

Scalability



Tree Level View




Forest Level View

Challenge 2: Capacity




Capacity Crunch Debugging

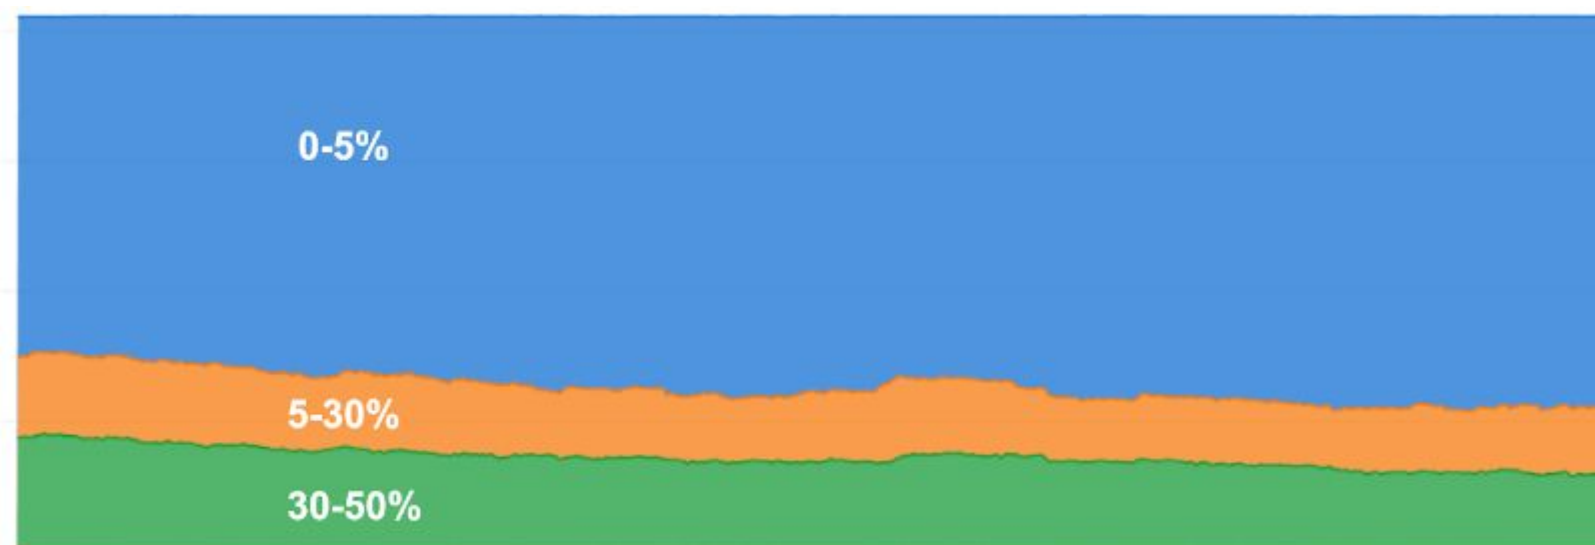
Model Backlog



Capacity Crunch Debugging



Training Fleet CPU/GPU Usage %



Training Fleet Memory Usage %




Capacity Crunch Debugging

Detailed Resource Usage: `$ cli_tool_name status --resource <job_id>`

Example Job:



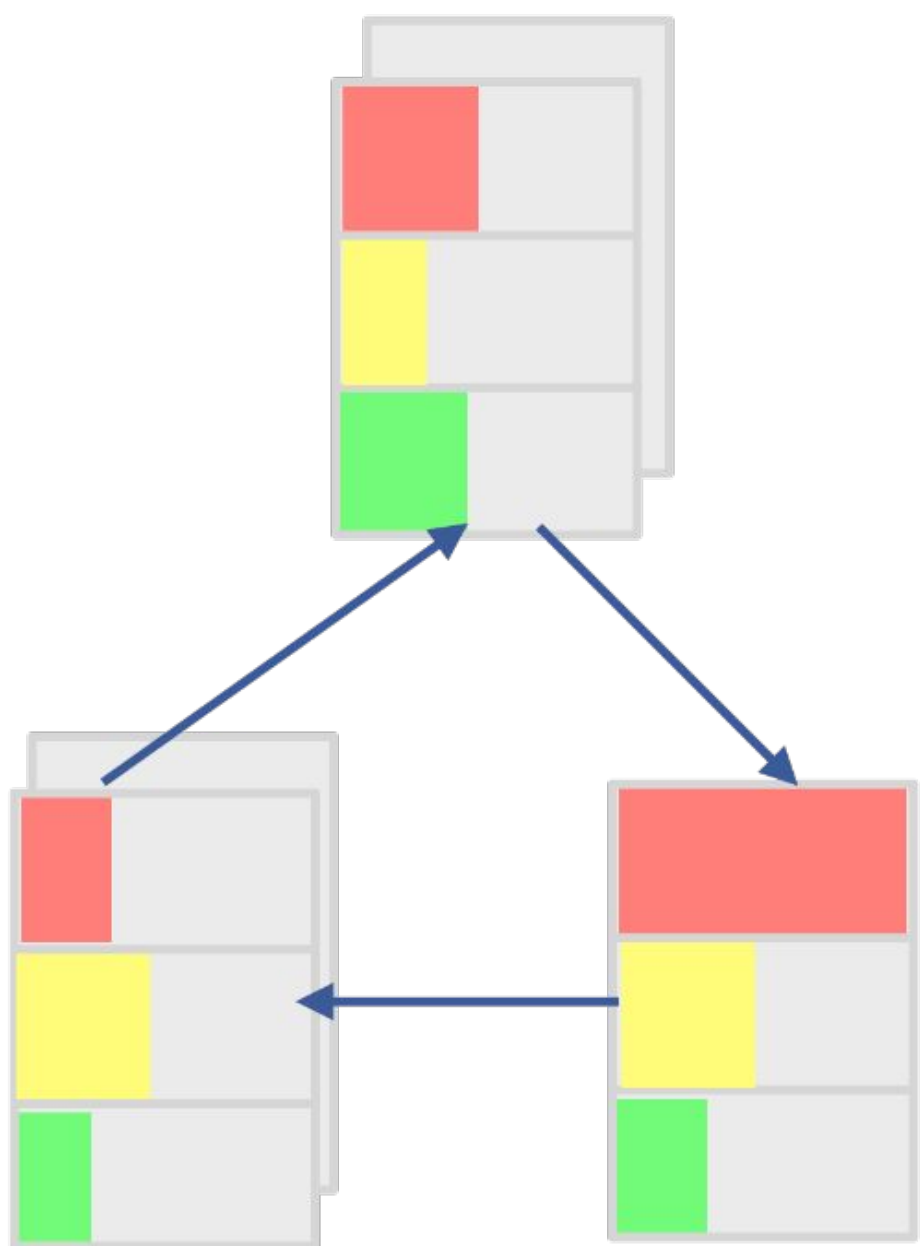
Example Machine:




Capacity Crunch Debugging

Find Best Practice: `$ cli_tool_name compare --job-ids "id1 id2 id3 ..."`

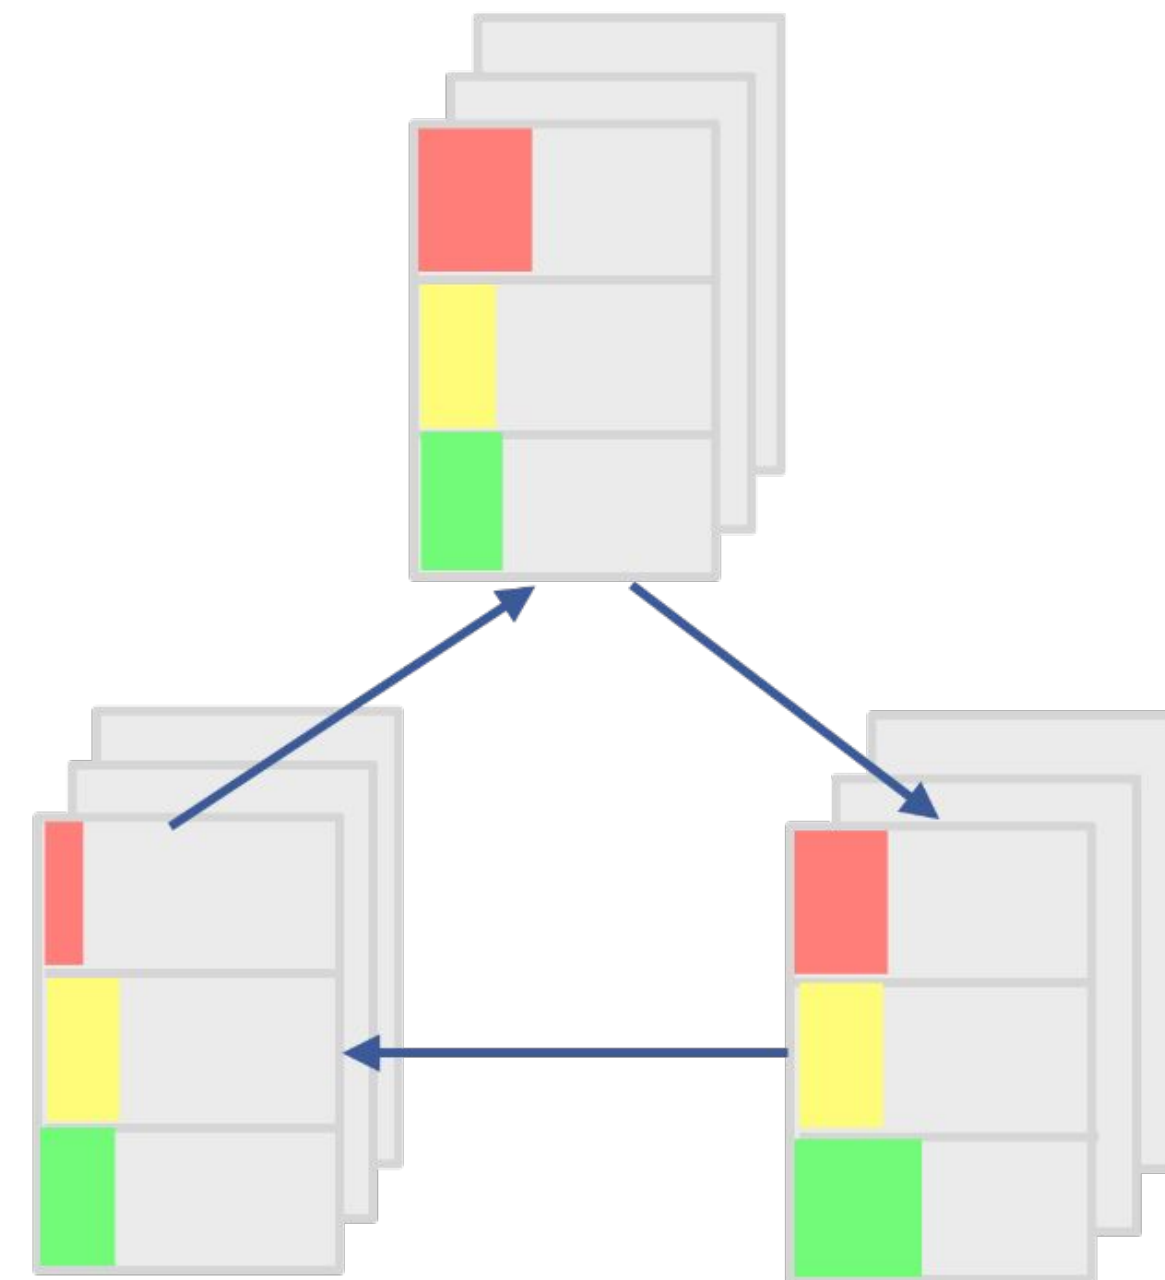
Job 1



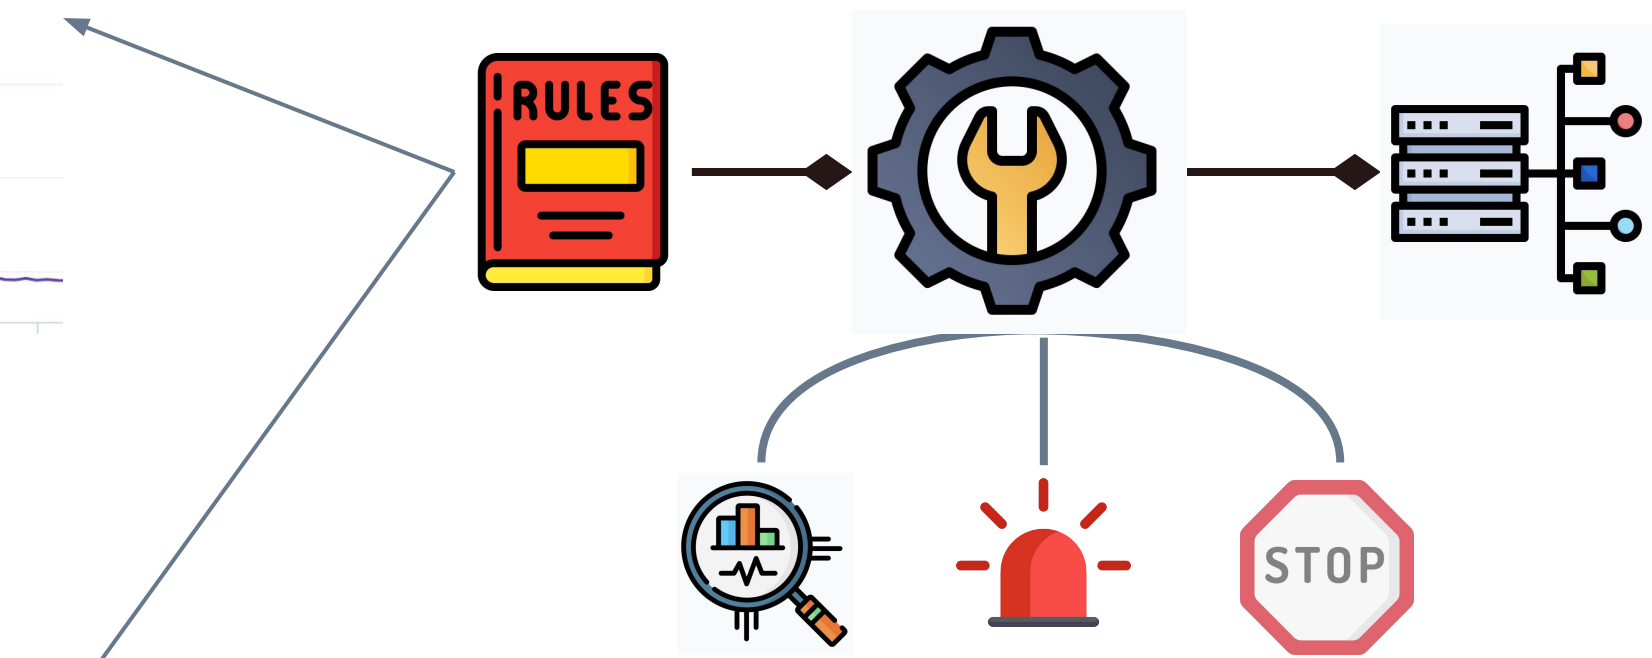

Job 2



Job 3




ML Guardian




Admission Control Evolution




Team- based




Project- based



ROI- based



- 10,000s jobs



- 1000s jobs



- 100s jobs



Reliability

90%

Reliability

$$.9 * .9 * .9 * .9 * .9 = 59\%$$

Reliability

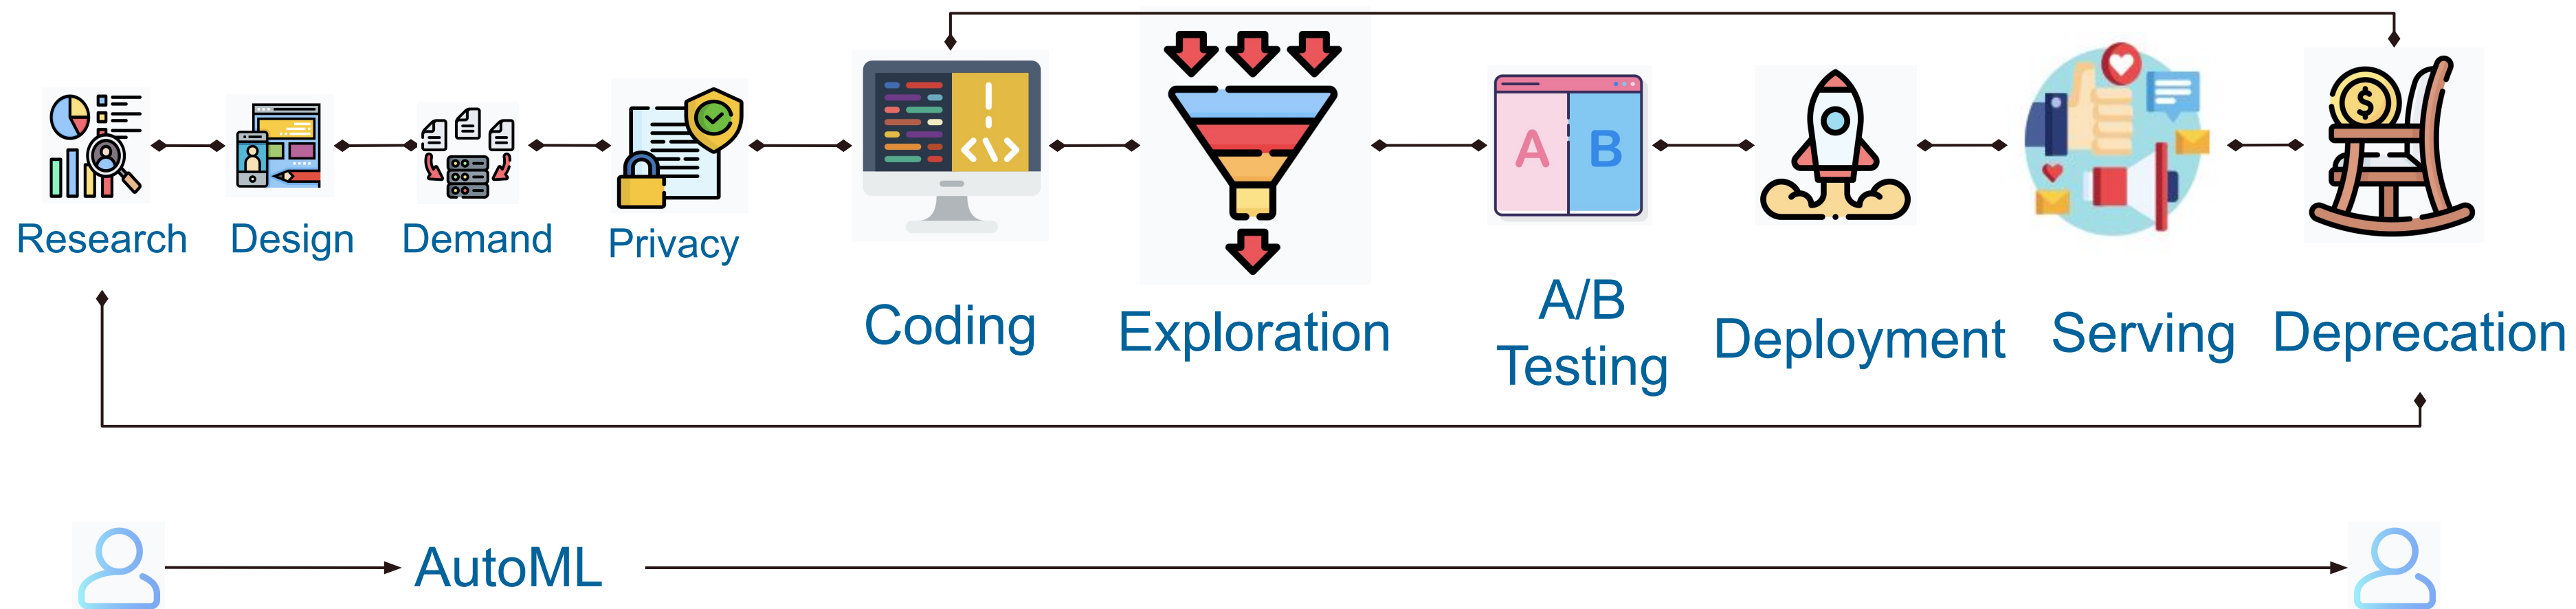
99%

Reliability

$$.99^* \cdot .99^* \cdot .99^* \cdot .99^* \cdot .99 = 95\%$$

Challenge 3: Reliability

ML Development LifeCycle



AutoML Package Release

7/13							7/12							7/11				7/10			7/9											
B3880	B3879	B3878	B3877	B3876	B3875	B3874	B3873	B3872	B3871	B3870	B3869	B3868	B3867	B3866	B3865	B3864	B3863	B3862	B3861	B3860	B3859	B3858	B3857	B3856	B3855	B3854	B3853	B3852	B3851	B3850	B3849	B3848
×	×	×	⊘	×	×	×	×	!	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	

↑
Build Failures

AutoML Package Release

8/5			8/4				8/3				8/2				8/1			7/31			7/30										
B3987	B3986	B3985	B3984	B3983	B3982	B3981	B3980	B3979	B3978	B3977	B3976	B3975	B3974	B3973	B3972	B3971	B3970	B3969	B3968	B3967	B3966	B3965	B3964	B3963	B3962	B3961	B3960	B3959	B3958	B3957	B3956
✓	✗	✗	✓	✗	⊘	✓	✓	✗	✗	✗	✓	✓	✓	✗	✗	✗	✗	✓	✓	✗	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✓
✓			✓			✓	✓				✓	✓	✓					✓	✓		✓	✓	✓	✓	✓	✓	✓	✓		✓	
✓			✓			✓	✓				✓	✗	✗					✓	✗		✗	✓	✓	✓	✓	✓	✓	✓		✗	
✓			✓			✓	✓				✓	✓	✓					✓	✓		✓	✓	✓	✓	✓	✓	✓	✓		✓	
✓			✓			✓	✗				✓	✗	✓						✗		✗	✓	✗		✓	!	✗	✓	✓	✓	
✓			✓			✓	✗				✗	✓	✓								✗	✗	✗		✓	✓	✓	✓	✓	✓	
✓			✓			✓	✓				✓	✓	✓					✓	✓		✗	✗	✓		✗	✓	✓	✓	✓	✓	
						✓						✓																✓			
						✓						✓																✓			

AutoML Package Release

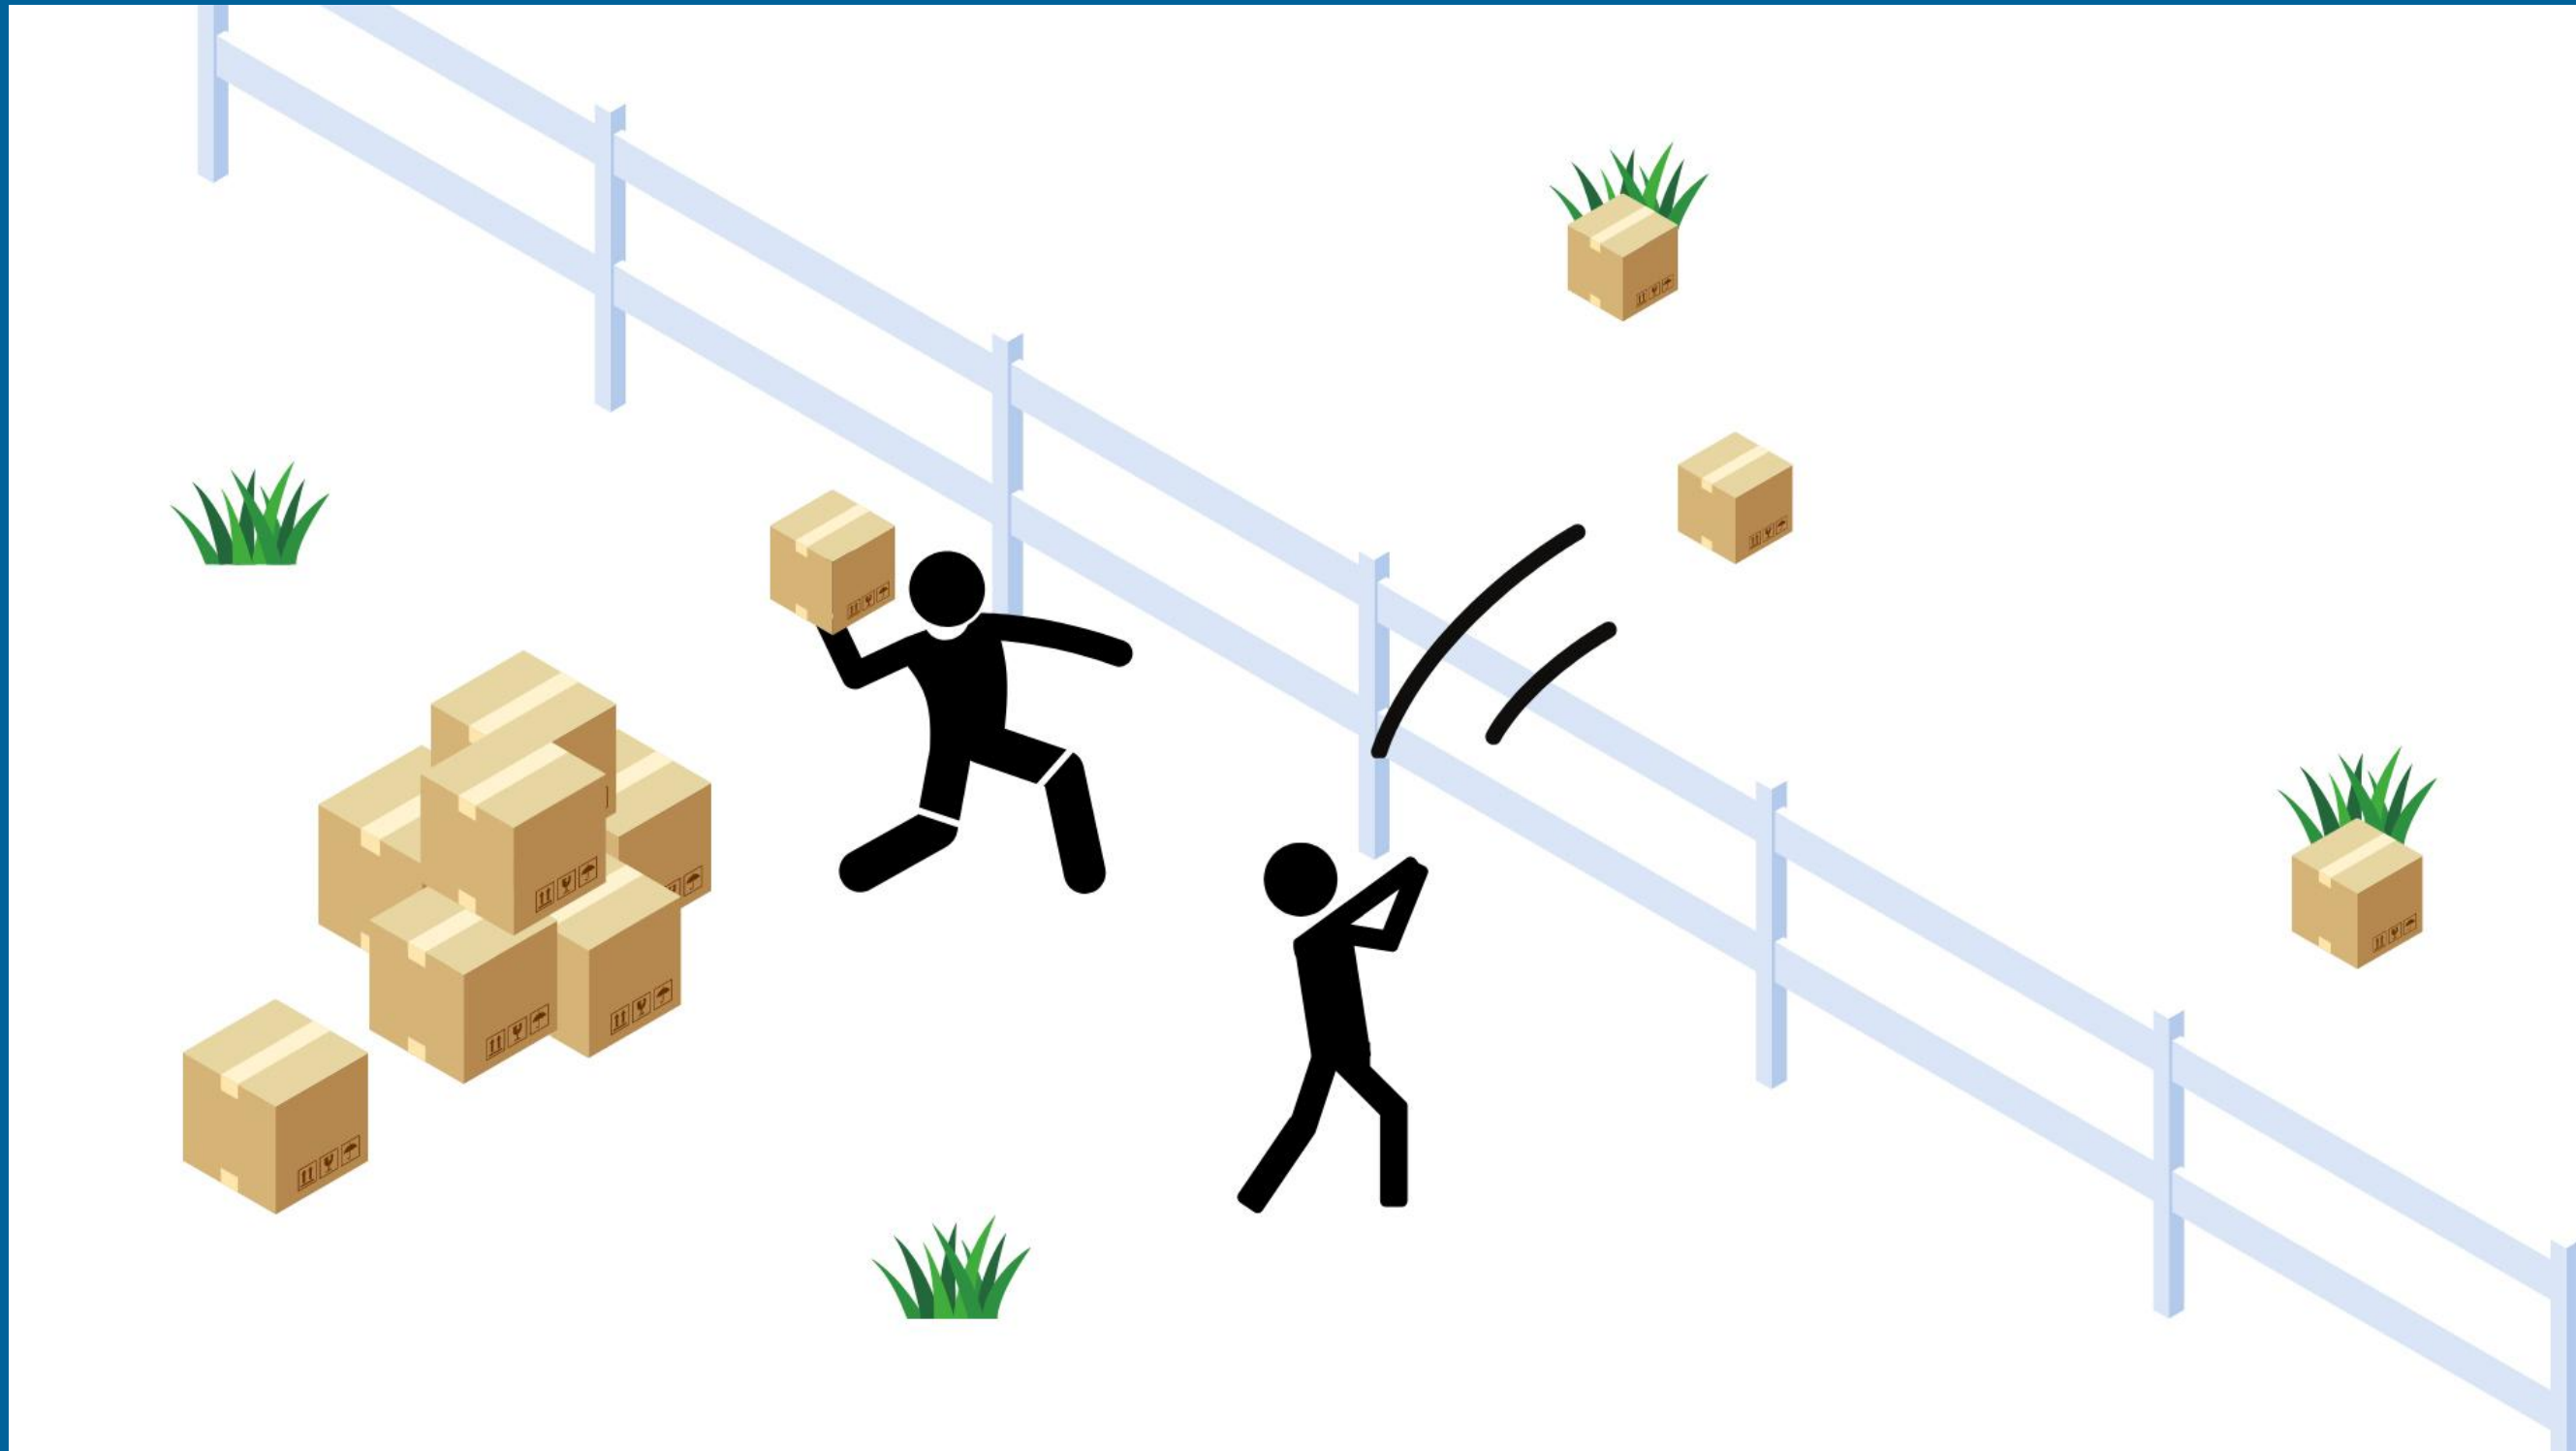
		11/30			11/29			11/28			11/27			11/26			11/25			11/24			11/23										
		B4487	B4486	B4485	B4484	B4483	B4482	B4481	B4480	B4479	B4478	B4477	B4476	B4475	B4474	B4473	B4472	B4471	B4470	B4469	B4468	B4467	B4466	B4465	B4464	B4463	B4462	B4461	B4460	B4459	B4458	B4457	B4456
		✓	✗	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	⊘	✓	⊘	✓	✗	✓	✗	
				✗	✗	✓		✓	✓	✓	✗	✓	✓	✓	✓	✗	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓		✓		✓	✓		
				✓	✓	✗		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗		✓		✓	✓		
				✓	✓	✓		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓							✓		✓	✓	✓		
				✓	✓	✓		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓							✓		✓	✓	✓		
				✓	✓	✓		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓							✓		✓	✓	✓		
				✓	✓	✓		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓							✓		✓	✓	✓		

↑
Tests

Takeaways

Nondeterministic

Takeaways



Takeaways

**LIVE
IN THE
FUTURE**

AdsML PE Team

Anamaya Sullery

Anup Mohan

Arun Paramadhathil

Brian Hansen

Daniel Thurman

David Egerton

Denny Zhang

Dmitry Khapilin

Haimo Chen

Jialing Zheng

Jiayi Zhou

Krish Kalai

Mark Klunder

Natasha Morando

Qingyu Wang

Sidd Choure

Tianchen Zhong

Yansong Chen

Yan Yan

Thank you!