



PagerDuty

Leading through AI-powered Complexity

Balancing Business Priorities, Innovation, and Resources

November 2023

Intro: João (spelling / ʒw.ˈẽw /)

General Manager & Engineering Lead



BACKGROUND

- CTO, GM @ Defined.ai
 - ML, Data Eng., Platform (DevOps, SRE), ProdDev, IT, Sales Eng.
- SDE, Manager @ Microsoft
 - Language Expansion for Xbox, Office, Windows
 - Search, Big Data
- PhD in Speech Tech and HCI
 - 5 patents, +40 publications, 1 book



The growth mandate for the modern enterprise



Win
and protect
revenue

The growth mandate for the modern enterprise



Win
and protect
revenue



Dramatically
compress costs

The growth mandate for the modern enterprise



Win
and protect
revenue



Dramatically
compress costs



Mitigate
Operational
Risk

The growth mandate for the modern enterprise



Win
and protect
revenue



Dramatically
compress costs




Mitigate
Operational
Risk



Harness
automation
and AI

Why does AI/ML and Generative AI matter?

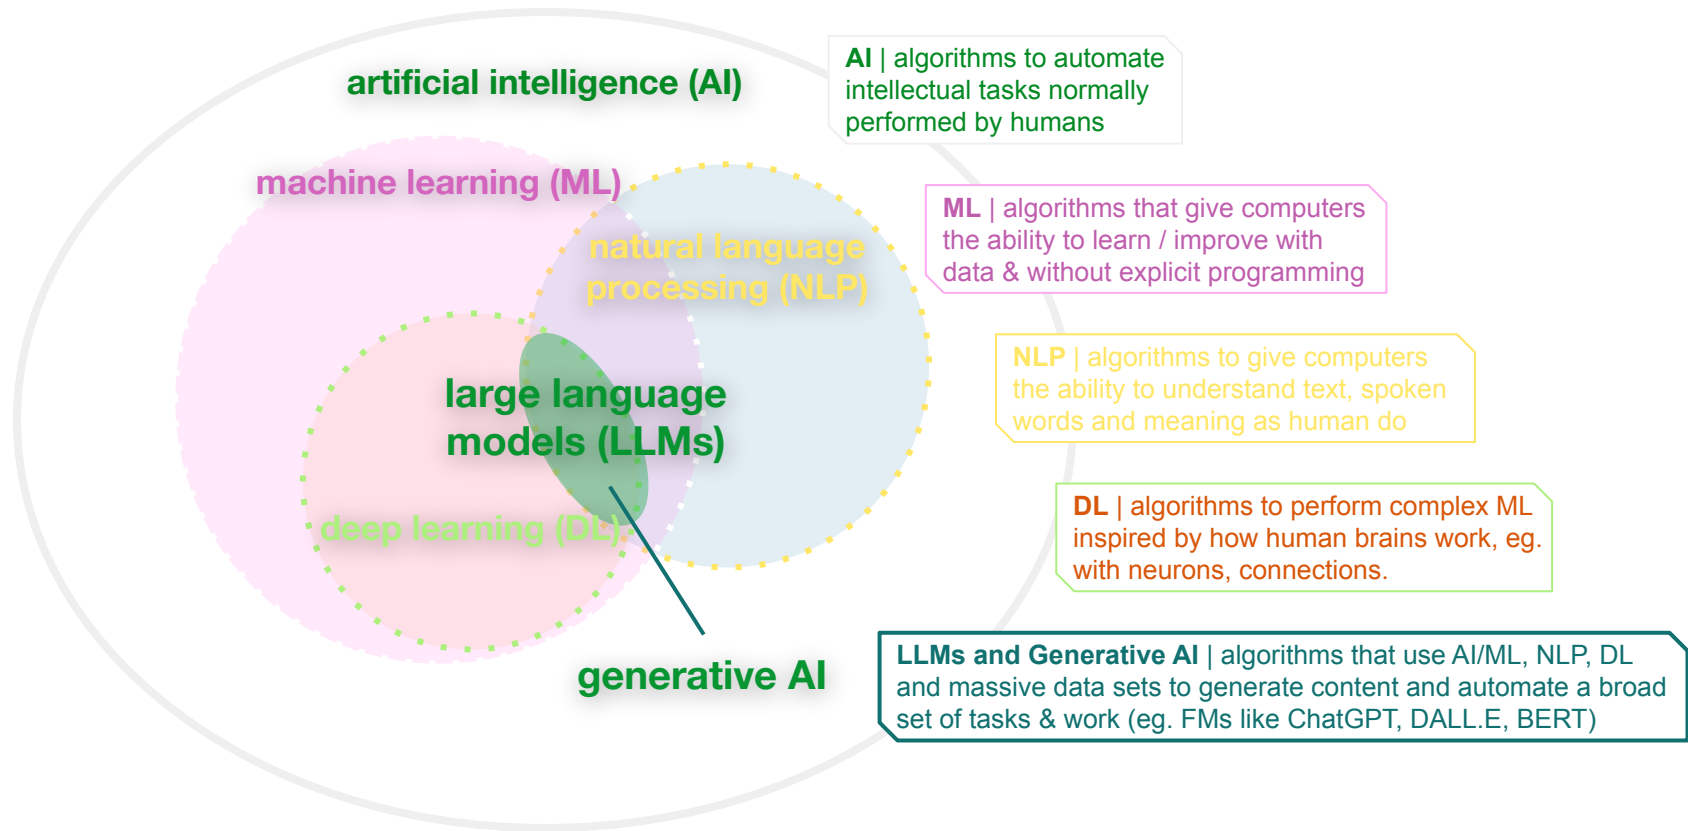
US Nonfarm Business Sector Output Per Hour (Y/Y, 5 Yr. Rolling Avg)




Source: AlphaWise CIO Survey, Morgan Stanley Research

Thesis: ~every company will aspire + invest to become an AI company and every end-user and consumer will expect AI built-in to their digital products


AI is a story of human progress



Hardware acceleration in 2000s & 2010s



LLM models leap forward post-2017



Transformers



Attention Is All You Need

Ashish Vaswani*
Google Brain
avaswani@google.com

Noam Shazeer*
Google Brain
noam@google.com

Niki Parmar*
Google Research
nikip@google.com

Jakob Uszkoreit*
Google Research
usz@google.com

Llion Jones*
Google Research
llion@google.com

Aidan N. Gomez* †
University of Toronto
aidan@cs.toronto.edu

Lukasz Kaiser*
Google Brain
lukaszkaizer@google.com

Illia Polosukhin* ‡
illia.polosukhin@gmail.com

The Uniform Bar Exam



The SAT



Sommelier examinations

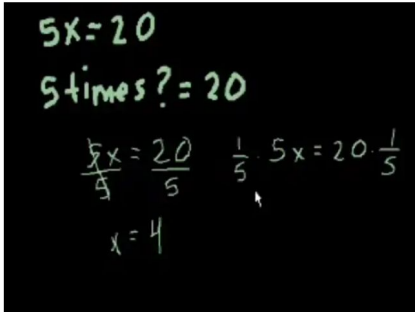


Wharton MBA exam



ChatGPT 4 at ~80-90 IQ


AMC Exams



AP Exams



GRE



USA Biology Olympiad Semifinal Exam



We are watching the rapid pace of generative AI progress

Midjourney Text to Graphics Evolution ~1 Year & ~10 people



V1

Released
February 2022



V2

Released
April 12, 2022



V3

Released
July 25, 2022



V4

Released
November 5, 2022



V5

Released
March 15, 2023



V5.1

Released
May 3, 2023

“Machine Learning sucks!” Y. LeCunn

Supervised Learn. -> large labeled datasets


Reinforcement Learn. -> high number of trials

Self-Supervised Learning -> very large datasets

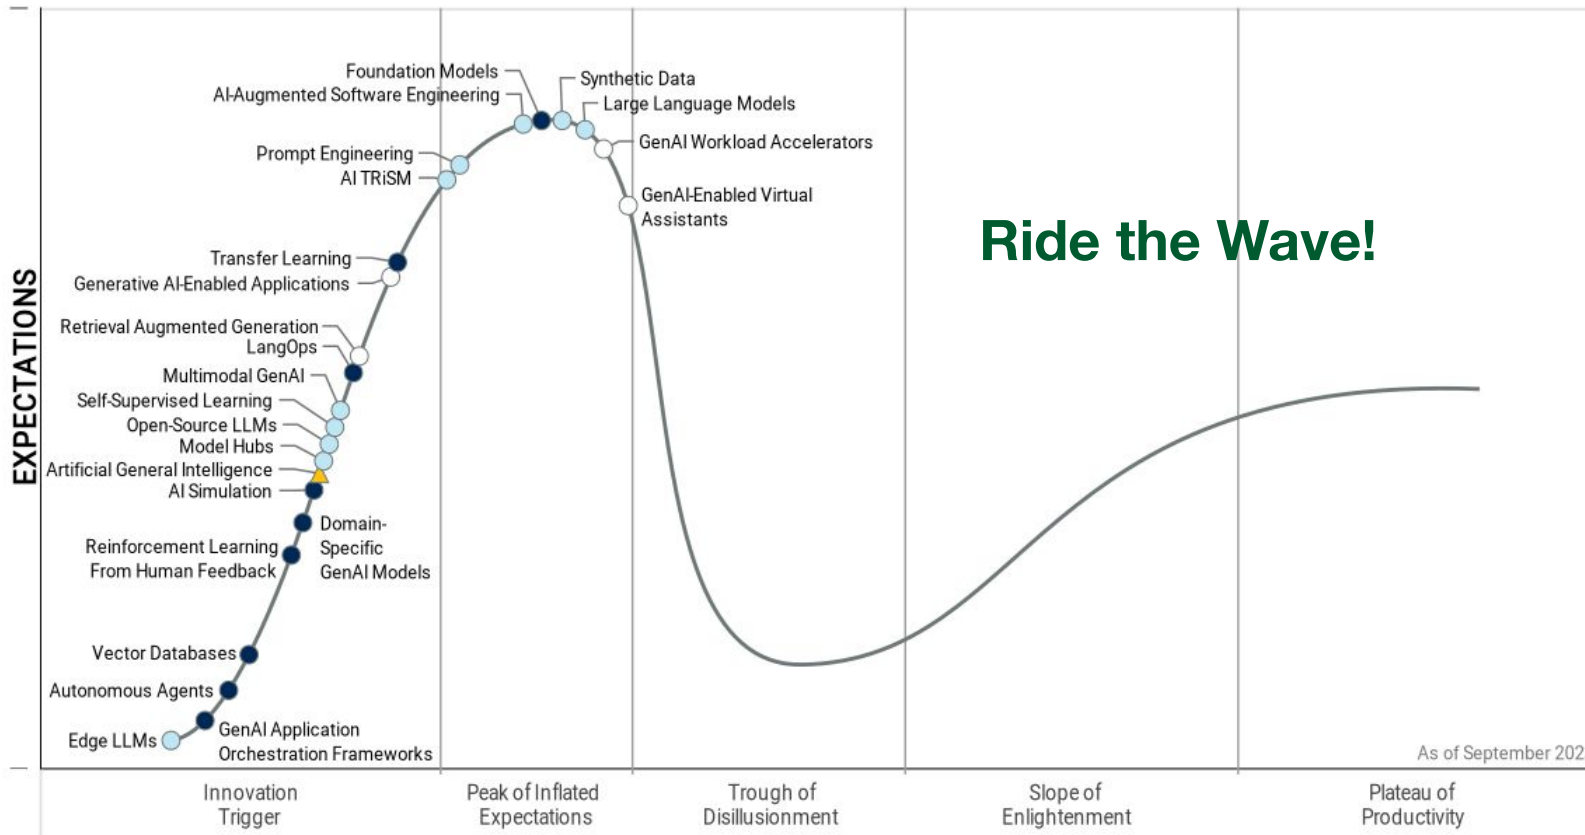
NO REASONING!
NO PLAN!

LLMs Performance is great but....

- Factual errors, logical errors, inconsistency, limited reasoning, toxicity...
- No knowledge of the underlying reality
- No common sense / can't plan their answer



Hype Cycle for Generative AI, 2023




Plateau will be reached: ○ <2 yrs. ● 2-5 yrs. ● 5-10 yrs. ▲ >10 yrs. ⊗ Obsolete before plateau

Invention x IMPACT = INNOVATION

IMPACT

*requires
Distribution
and Adoption*




INVENTION

INVENTION

INVENTION

INVENTION



PagerDuty

Strategy, Business and Resources

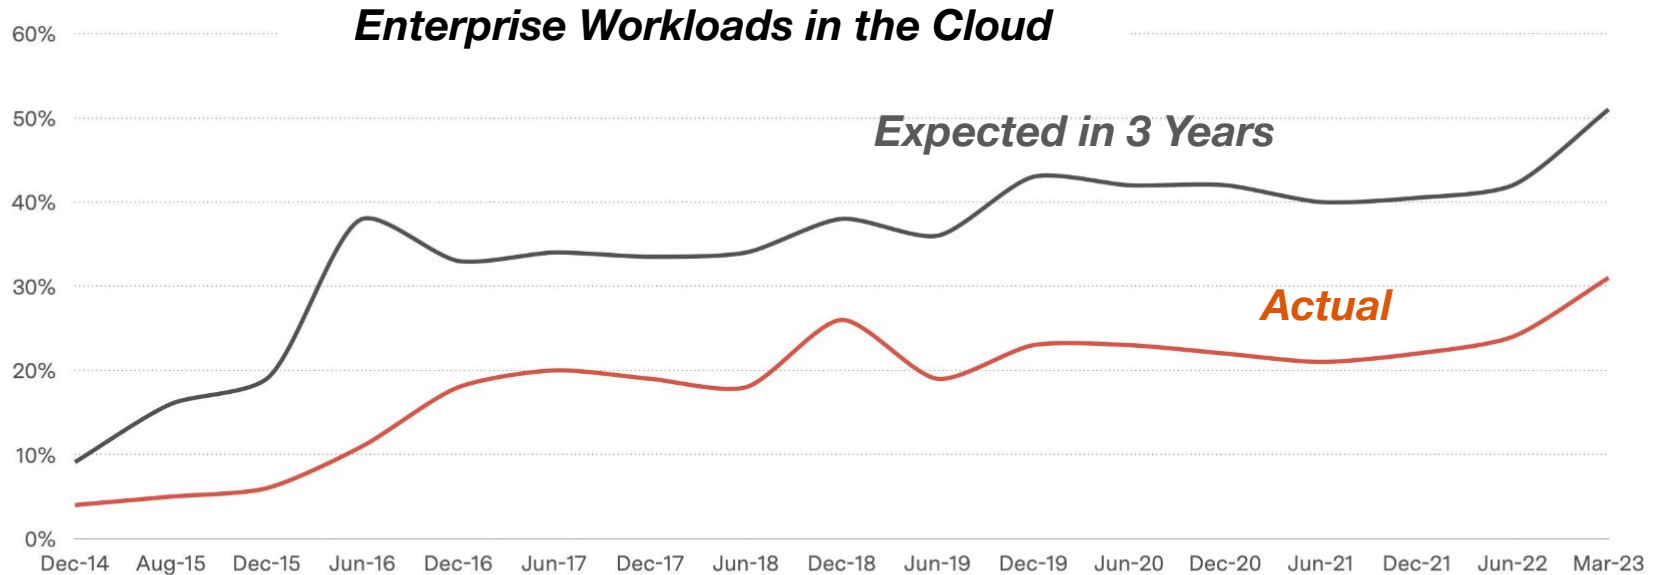
Proprietary & Confidential

Strategy: Long-term vs Short-term



Performance & Productivity Zones = Run the Business
Incubation & Transformation Zones = Growth the Business

Let's also remember.. generational technology shifts *take a while* in the enterprise



Demand for Talent

cfo.com/news/demand-for-gen-ai-skills-explodes-in-2023-artificial-intelligence-job-posting-hiring/699702/



Opinion Library Events Press Releases

Corporate Finance Human Capital People Regulation & Compliance Risk Strate

Demand for GenAI Skills Explodes in 2023

More than 10,000 jobs seeking wherewithal in generative AI have been posted so far this year.

Published Nov. 20, 2023

hrdive.com/news/demand-generative-ai-skills-explodes/697586/



Deep Dive Opinion Library Events Press Releases 2023 Report

Compliance Talent Learning Comp & Benefits Diversity & Inclusion HR Mgmt

DIVE BRIEF

Report: Demand for generative AI skills has exploded 1,848% since 2022

While generative AI demand is small compared to other AI skill sets, the increase reflects employers' rising interest in adopting AI tech, experts said.

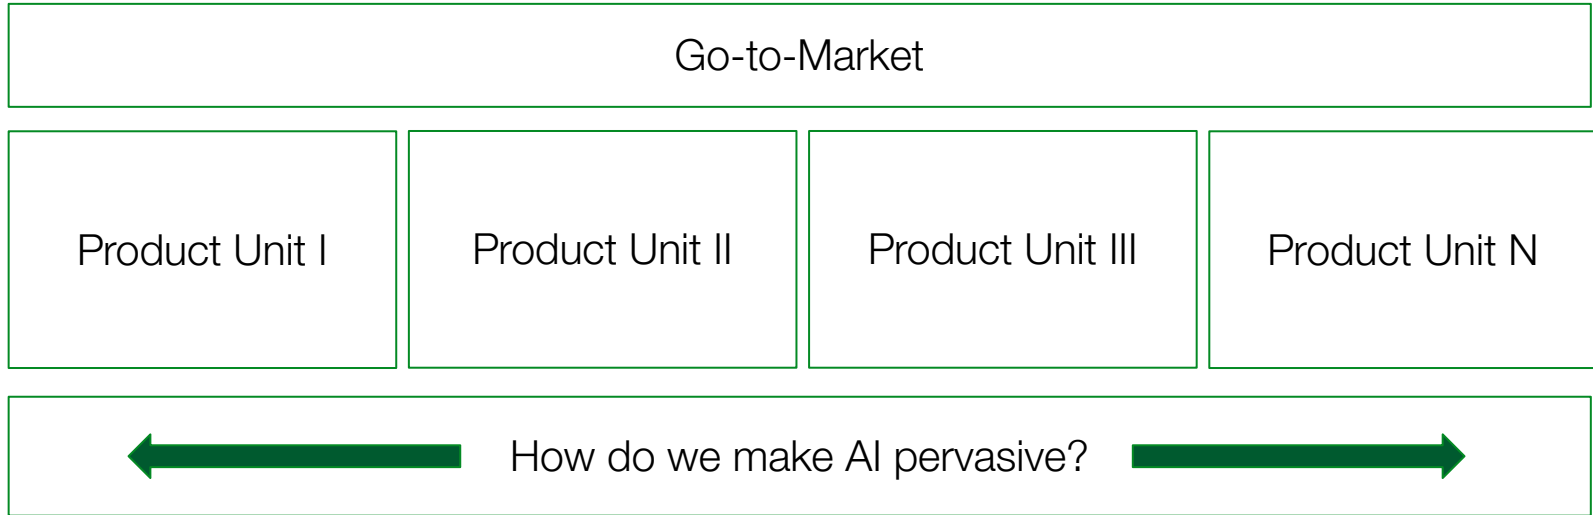
Published Oct. 26, 2023

By Laurel Kalsner
Contributor



**ML Talent is expensive
and hard to find!!**

AI in a Multi-Product Organization



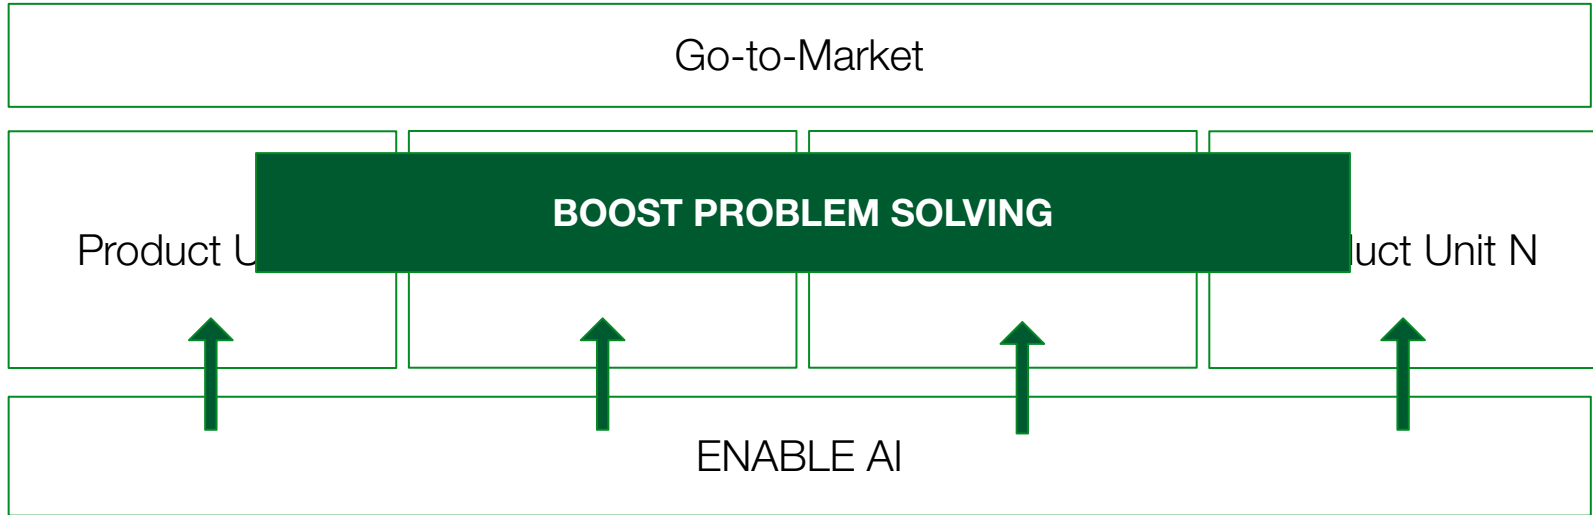
AI in a Multi-Product Organization



Work backwards from the customer problem


Developing scalable, sustainable use cases that deliver on key performance indicators

AI in a Multi-Product Organization



Team Impact


- Incentivize risk + innovation + disruption
- Where does Innovation fits in terms of your priorities?
- What are you NOT doing (to go disrupt)?
- Leadership support is essential



“Shoot for the moon. Even if you miss, you'll land among the stars” Norman V. Peale

What Problems should we go solve?

Focus, focus and focus

- Vector databases
 - Foundational LLMs
 - LLM Quality (bias, toxicity, etc.)
 - Scalability
 - Latency
 - Hosting
- 
- Buy vs Build
 - Volatile / Vendor Race
 - Leverage partners, 3rd parties
 - How are others solving GenAI blockers?

Manage Risk

- Privacy and data security
- Enhanced attack efficiency
- Misinformation
- Fraud and identity risks



Explainability/
Model Monitoring



AI Application
Security



ModelOps



Privacy

Extracting Training Data from ChatGPT

AUTHORS

Milad Nasr^{*1}, Nicholas Carlini^{*1}, Jon Hayase^{1,2}, Matthew Jagielski¹, A. Feder Cooper³, Daphne Ippolito^{1,4}, Christopher A. Choquette-Choo¹, Eric Wallace⁵, Florian Tramèr⁶, Katherine Lee^{+1,3}

¹Google DeepMind, ² University of Washington, ³Cornell, ⁴CMU, ⁵UC Berkeley, ⁶ETH Zurich. * Joint first author, +Senior author.

PUBLISHED

November
28, 2023

READ:

[[arxiv](#)]

We have just [released a paper](#) that allows us to extract several megabytes of ChatGPT's training data for about two hundred dollars. (Language models, like ChatGPT, are trained on data taken from the public internet. Our attack shows that, by querying the model, we can actually extract some of the exact data it was trained on.) We estimate that it would be possible to extract ~a gigabyte of ChatGPT's training dataset from the model by spending more money querying the model.



<https://not-just-memorization.github.io/extracting-training-data-from-chatgpt.html>

Encourage experimentation and adoption but...

- Set guidelines and guardrails for the use of GenAI
- Create awareness and training
- Internal vs customer use

For example:

- Can I use any provider?
 - What about extensions, plugins, etc.?
- What type of data can I send to an LLM?
- Can I use LLM in coding activities?
 - Can I include production-level code in the query?
- Can I launch product features that call LLM X?
- Opt-in vs Opt-out, which one to adopt?
- Customer Terms of usage and business model



Internal



External

Prototyping vs Product GA

1. Talk to customers -> ensure you are actually solving a problem
2. You don't need a fully mature system (reliable, scalable) to get feedback
3. Don't boil the ocean = move fast, adapt, course correct
4. Don't chase the buzz -> remember who your customers are

Important: Manage risk without killing innovation!

**... move fast
through AI
complexity**





PagerDuty Copilot

DATA

Generative AI for your critical
operations work.

✓ Incident resolved

Revenue Impact



Thank you

jfreitas@pagerduty.com

