

**STP 201**

# How Ginkgo Bioworks uses AWS to make organisms

**Dave Treff**

Head of IT & DevOps  
Ginkgo Bioworks

Dave Treff  
Head of IT & DevOps  
Ginkgo Bioworks, Boston  
MA



**GINKGO** BIOWORKS  
THE ORGANISM COMPANY











First ...





the second one ...





And then this one ...






And then at my first  
computer job ...






# I went on to ...

- EECS degree from UConn
- Compilers at multiple minicomputer startups
- Spent a decade (on and off) at DEC
  - Release 1 Ultrix (BSD 4.2) DecNet comms stack
  - Phase 5 DecNet “routers” and “bridges”
- Lead for OSF DCE project
- Trading systems at Boston Stock Exchange
- Vendor and investor due-diligence, expert witness
- Bank of America – Enterprise architecture
- AD DevOps @ Novartis
- VP of Tech Services @Qcentive





At the same time I was groggily running that all-night program, the guy on the left, Tom Knight, was getting his PhD.





**“The 21<sup>st</sup> Century will  
be the century of  
engineering biology.”**

Tom Knight.



PROPRIETARY AND  
CONFIDENTIAL




# GINKGO BIOWORKS

## THE ORGANISM COMPANY

**2008** Incorporated the company

**~290** Employees in Boston

**•(\$1e8)** Funding raised






# GINKGO BIOWORKS

## THE ORGANISM COMPANY

**~140** Robots

**#8** Y Combinator

**\$4.2B** Valuation




NOT FOR PUBLIC DISTRIBUTION

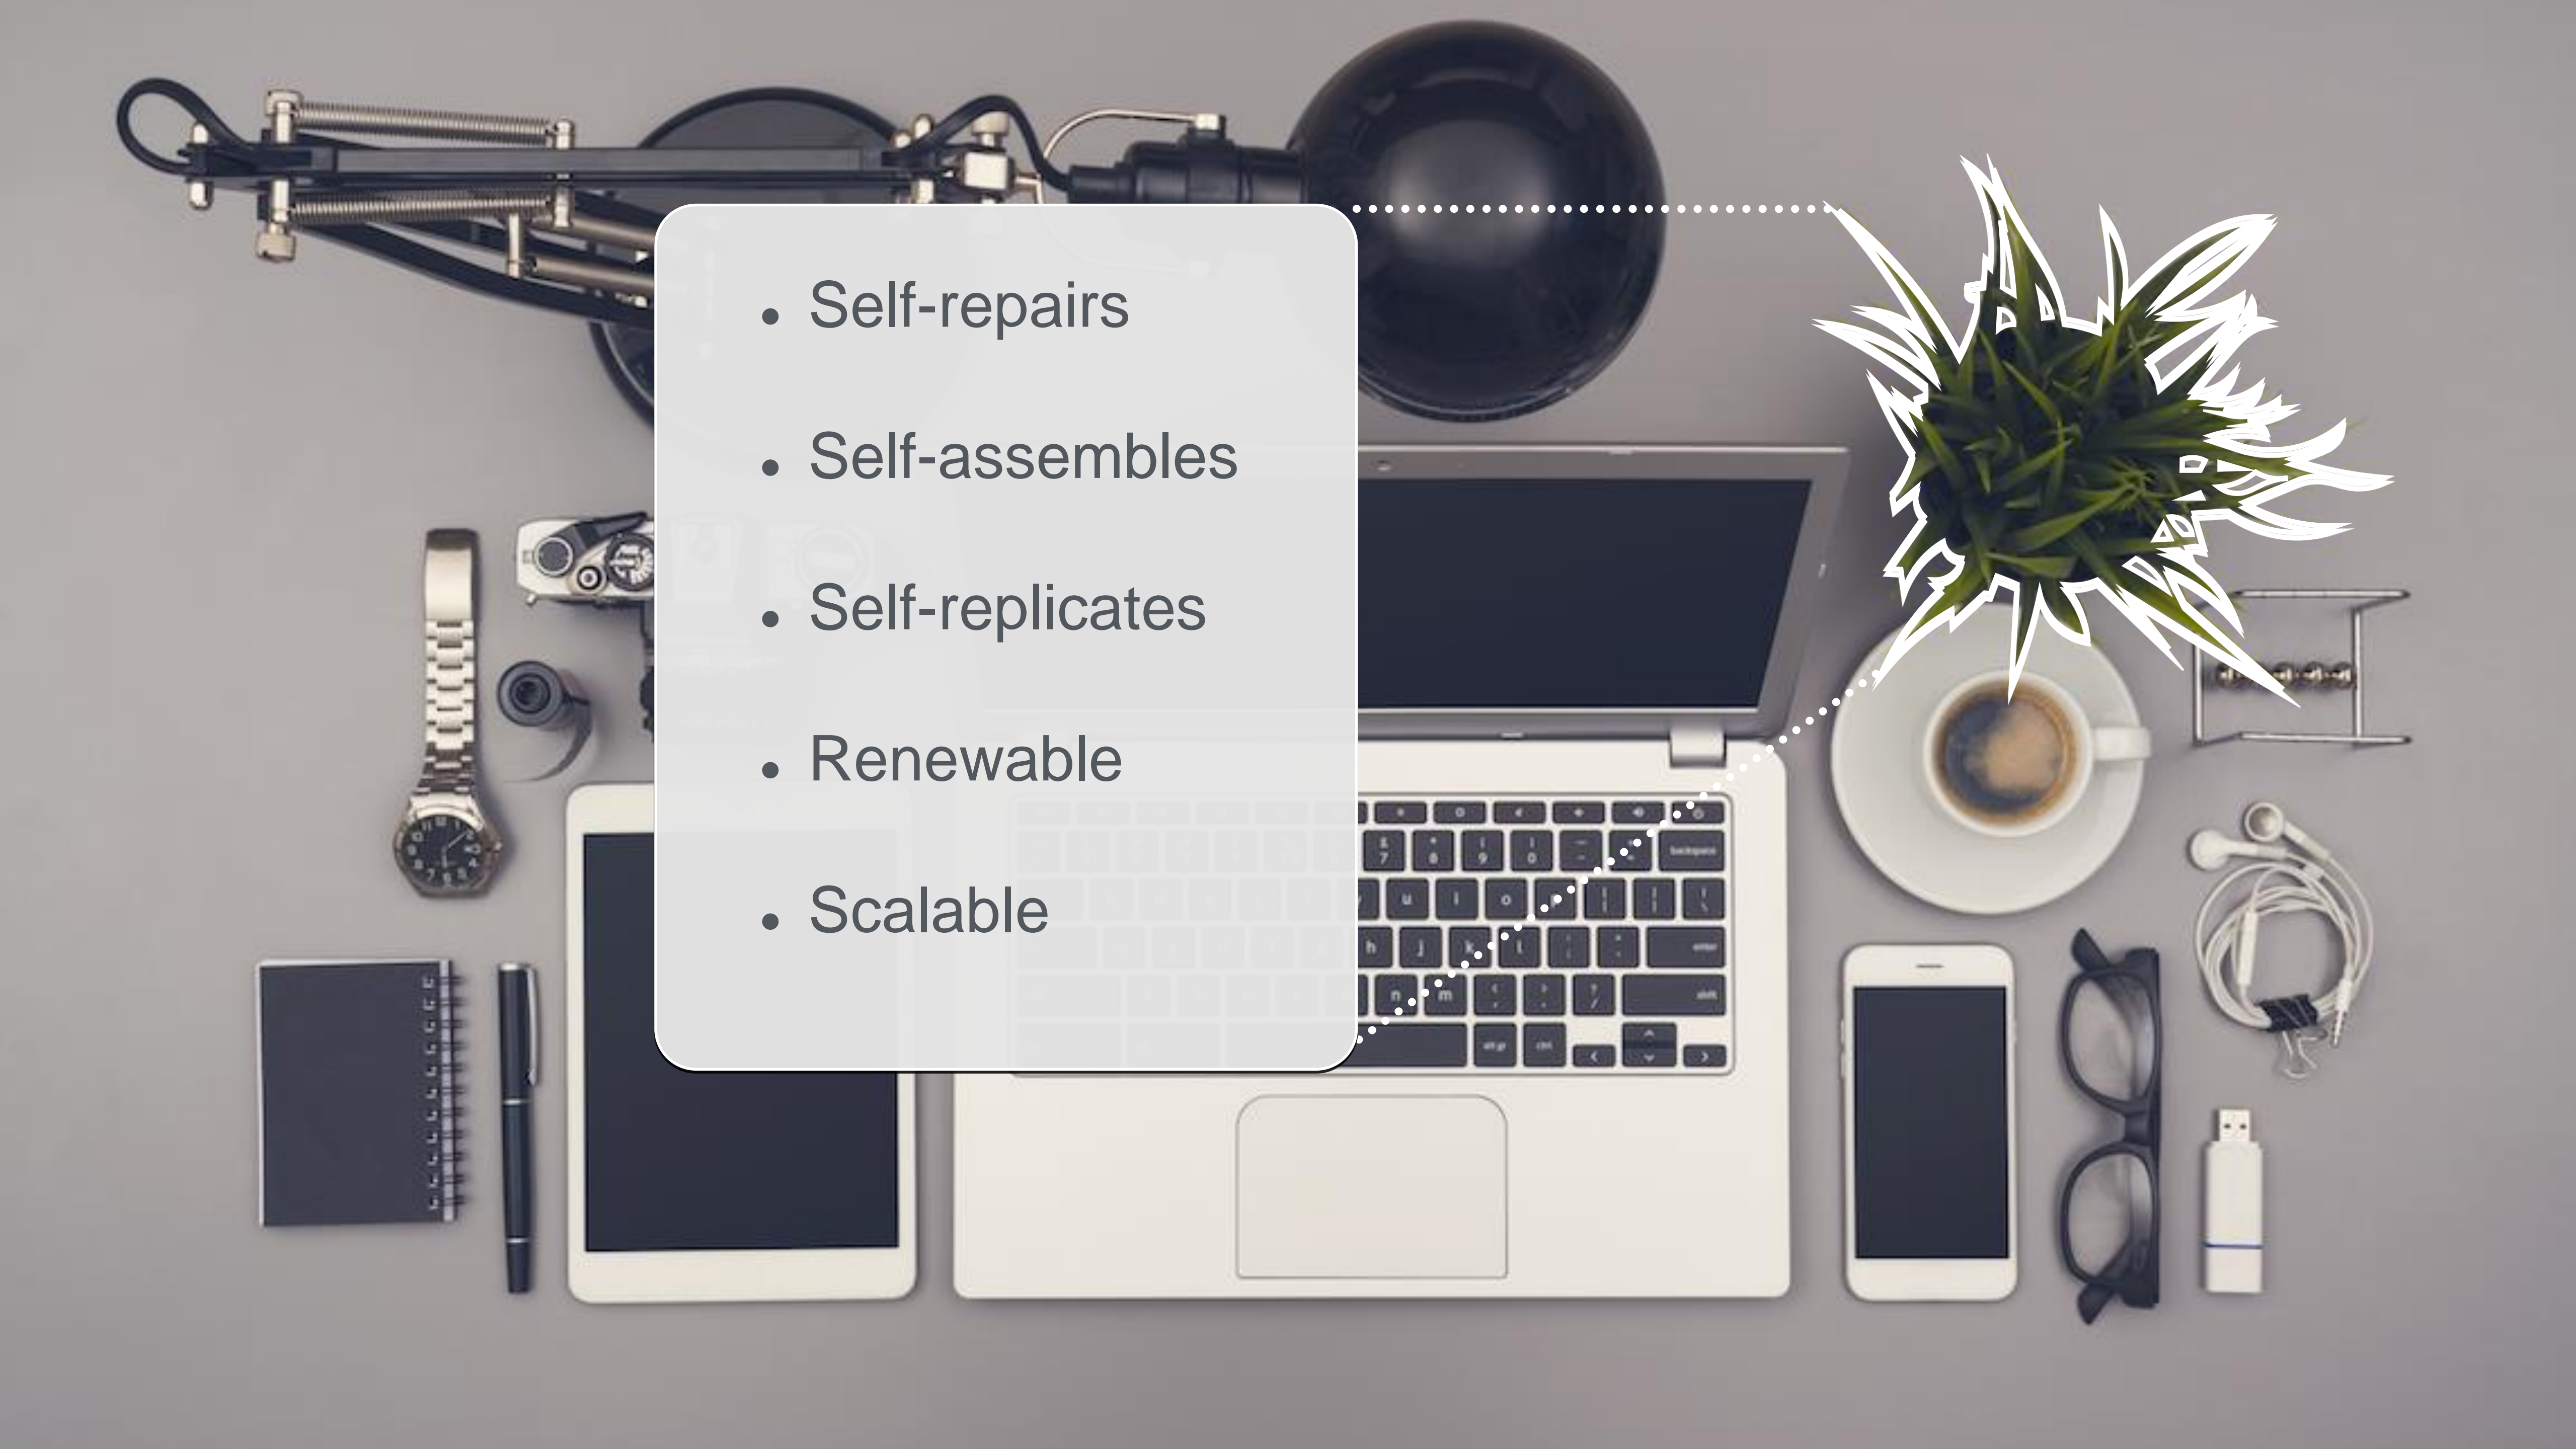


# Why Ginkgo Does Synthetic Biology







- 
- Self-repairs
  - Self-assembles
  - Self-replicates
  - Renewable
  - Scalable






“Make Biology easier to engineer”





# Make biology easier to engineer



- We engineer microorganisms for fermentation of small molecules and enzymes.
- Over 100 projects in 40+ different organisms
- Biosynthesis of specialty chemicals
  - Naturally occurring or novel
  - Known or unknown pathways
  - Flavors, drugs, vitamins, industrial chemicals, food ingredients
- Develop organisms that sense, respond to, and alter their environment



# Synthetic Biology @ Ginkgo



Ginkgo concentrates on solving vitally important problems with custom DNA programming to create organisms. These organisms create chemicals that

- Replace petroleum-sourced chemicals
- Replace animal-sourced food products
- Lessen reliance on endangered or at-risk plants
- Bioremediation of hazardous waste and non-degrading refuse
- Novel medicines and cell-based therapeutics
- Agriculture without harmful fertilizers and other chemicals



Ginkgo Bioworks is mission-driven:

We want to save the world with Biology



# Ginkgo Scale

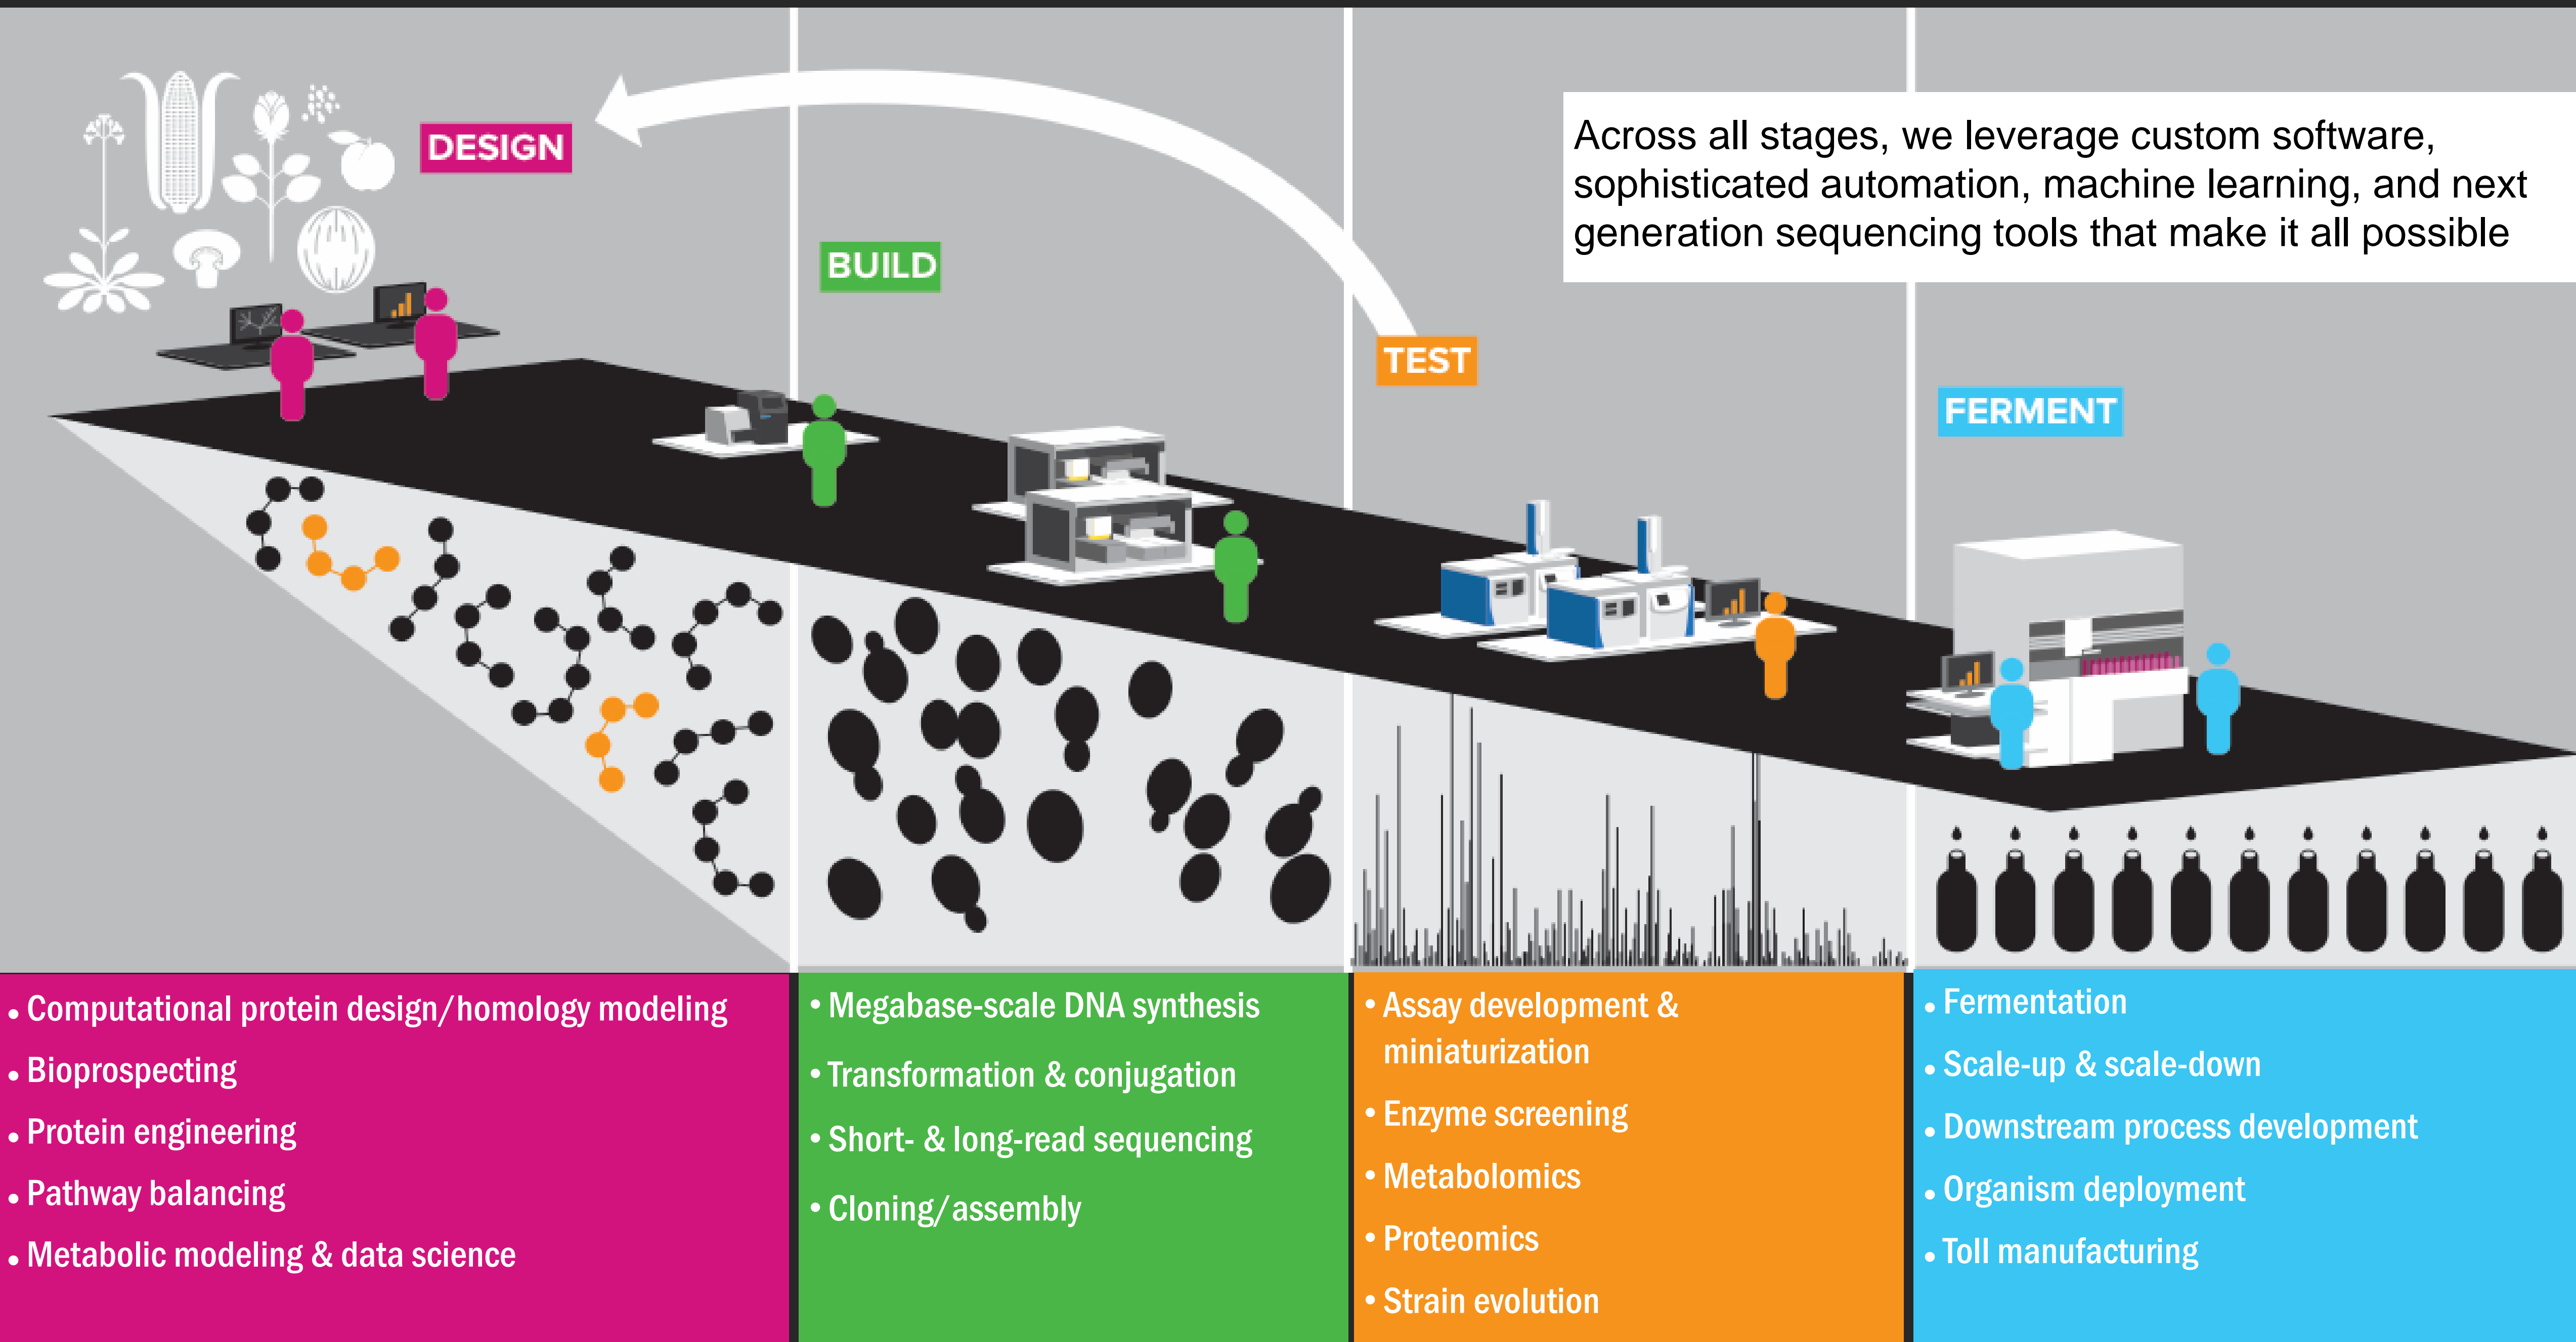


# Ginkgo's FOUNDRY compiles and debugs DNA code

Synthetic DNA is inserted into the genomes of cells and the cells are tested. This work is conducted on robotics - reducing costs compared to the status quo of scientists working by hand at the lab bench.




# Our foundry combines all the tools necessary for large-scale organism engineering under one roof











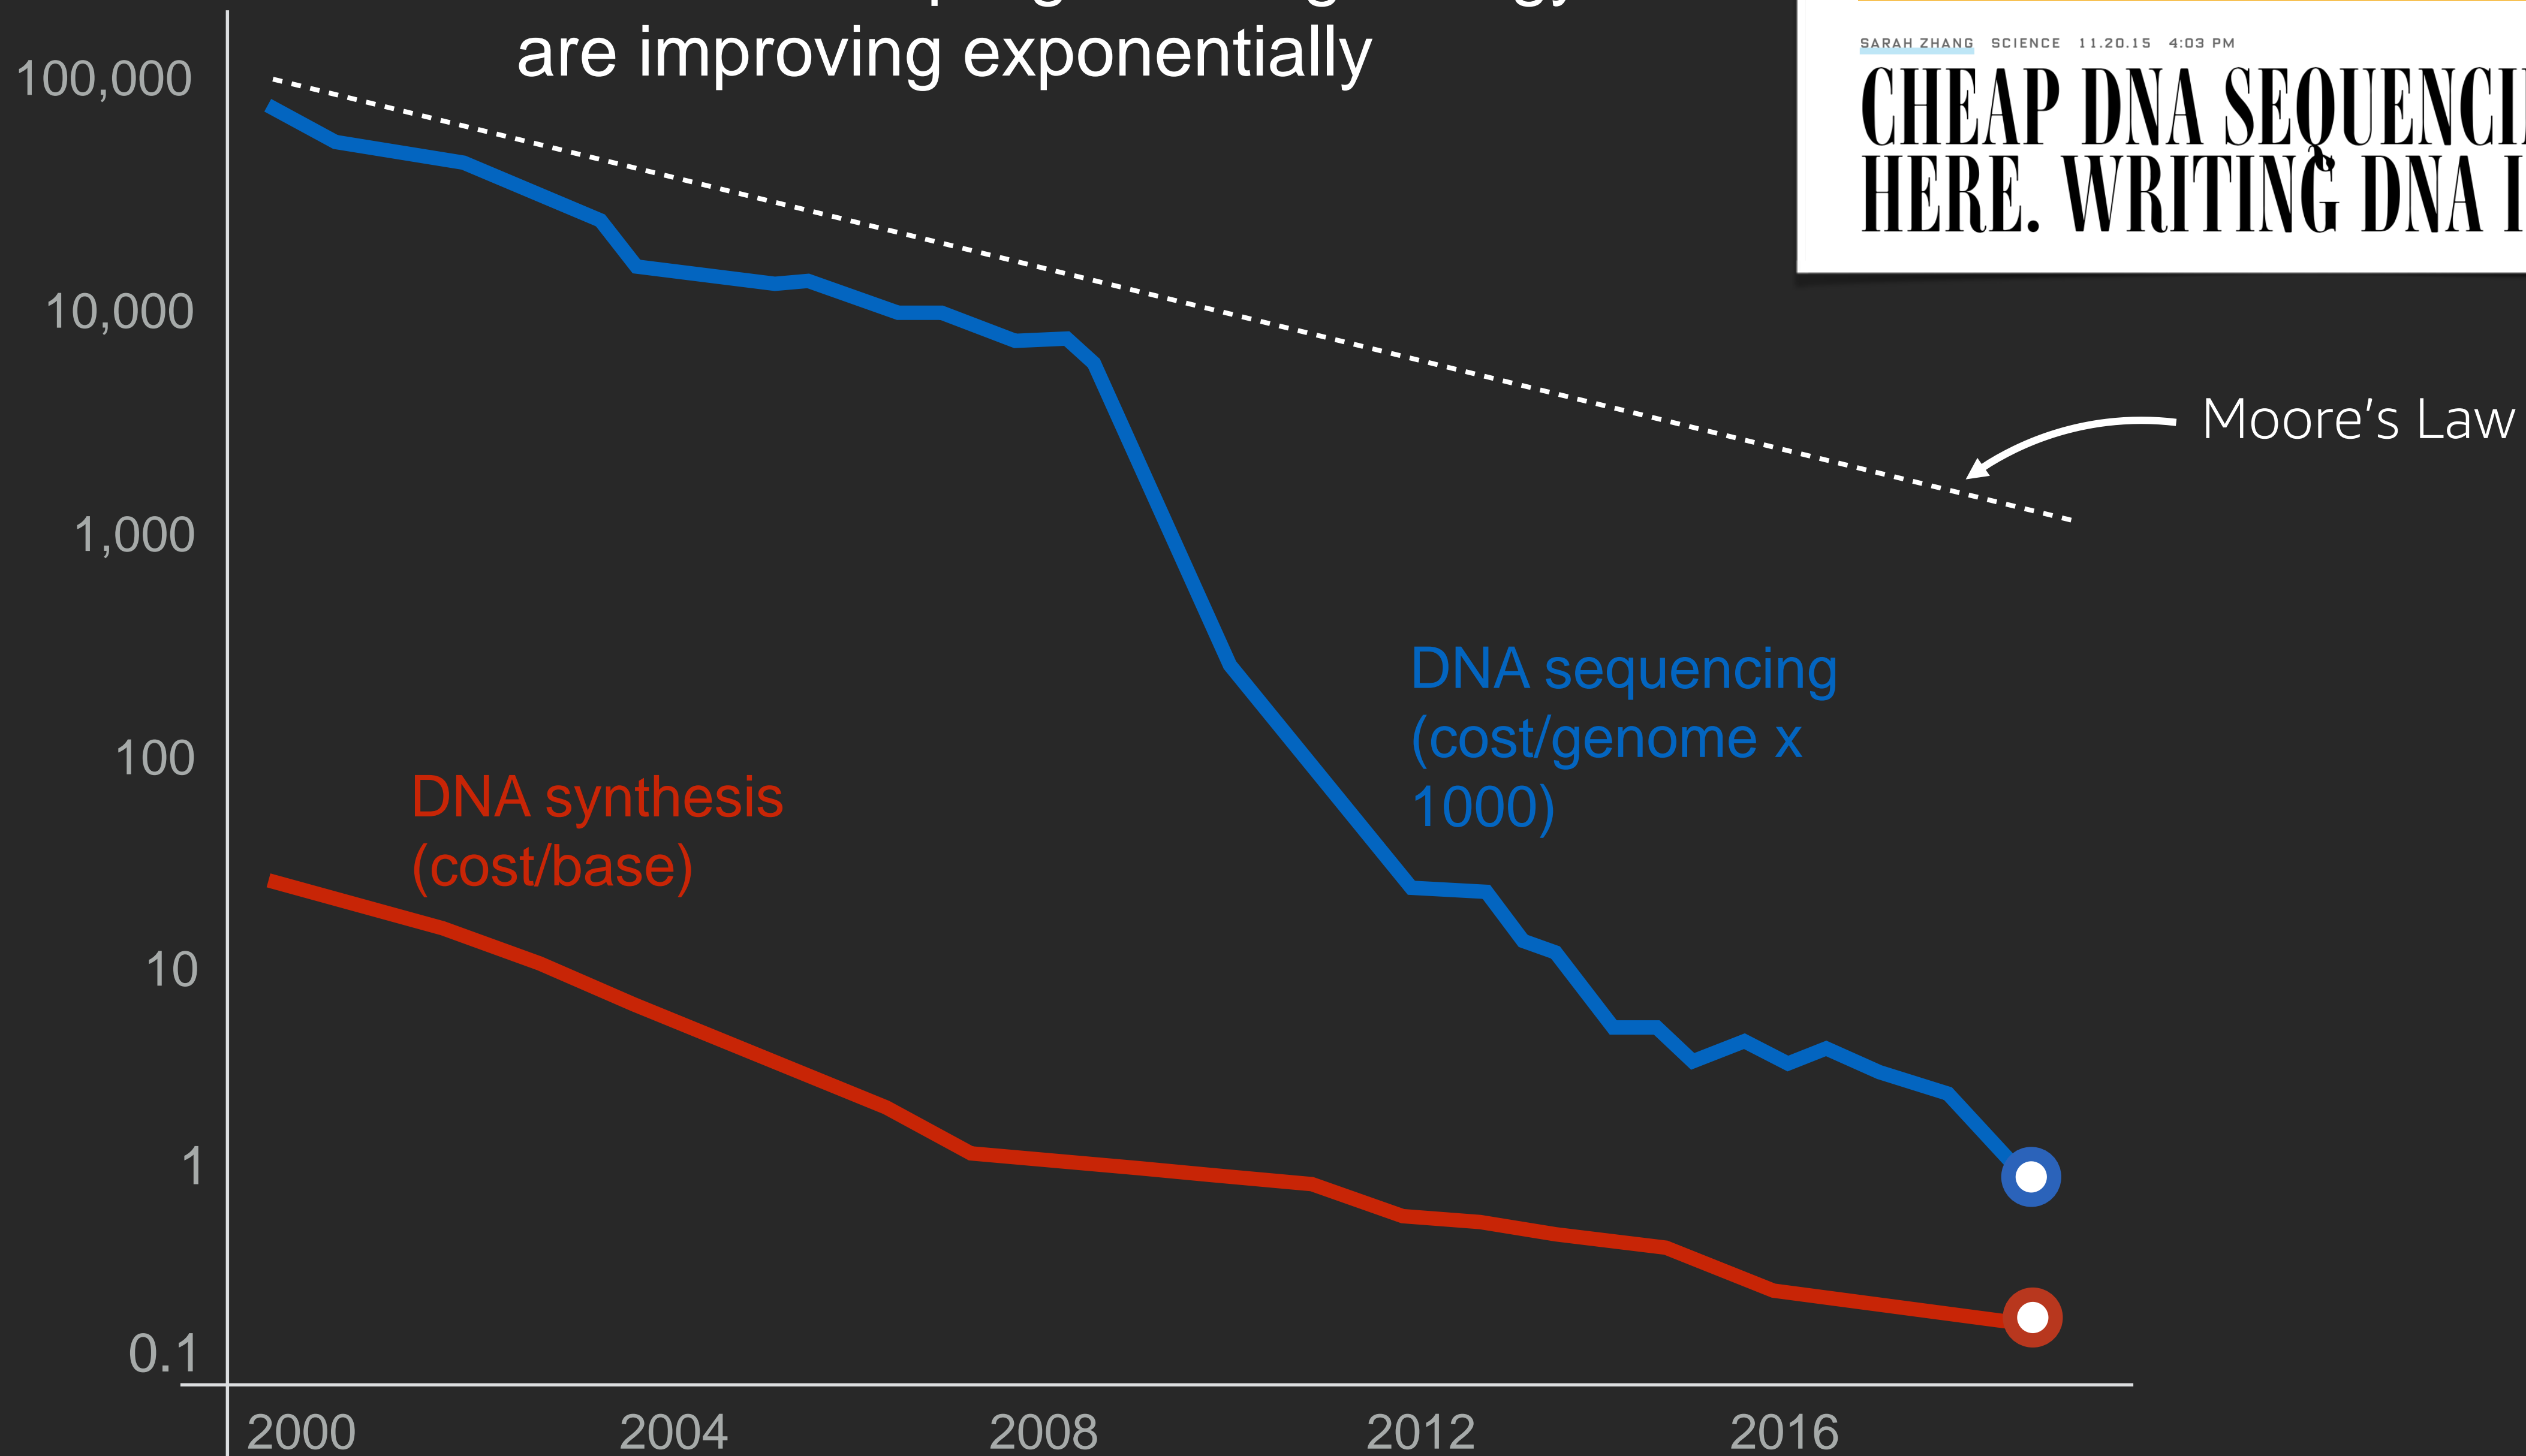
**FOUNDRY**



**CODEBASE**



The tools for programming biology  
are improving exponentially



WIRED

Why This Startup Is Betting on Writing—Not Editing—DNA

SARAH ZHANG SCIENCE 11.20.15 4:03 PM

**CHEAP DNA SEQUENCING IS  
HERE. WRITING DNA IS NEXT**






Knight's Law:

"The cost to genetically engineer a cell falls by 50% and the number of designs tested increases by 3X per year in Ginkgo's automated cell engineering foundries."

-GINKGO FOUNDER, TOM KNIGHT




# COST PER ENGINEERED CELL TESTED IS DROPPING 50% PER YEAR






# NUMBER OF STRAIN TESTS IS INCREASING BY 4X PER YEAR






DNA MOLECULES RECEIVED IS INCREASING BY 4X PER YEAR



**PRINT DNA** —————> **EXECUTE LAB OPERATIONS** —————> **MEASURE CELL PERFORMANCE**



NUMBER OF LIQUID TRANSFERS IS INCREASING BY 3X PER YEAR




PRINT DNA

EXECUTE LAB OPERATIONS

MEASURE CELL PERFORMANCE



NUMBER OF STRAIN TESTS IS INCREASING BY 4X PER YEAR



PRINT DNA —————> EXECUTE LAB OPERATIONS —————> MEASURE CELL PERFORMANCE





CRONOS  
G R O U P

 genomatica



# Ginkgo Startups



MarketWatch.com Sept 19, 2019 6:00 a.m. ET


“Most recently, Ginkgo established itself as a home for early-stage startups and entrepreneurs to build their biotech businesses. Through two new partnerships with Y Combinator and Petri, startups can access Ginkgo's platform and mentorship in exchange for equity.”



# Interlude: Ginkgo Stories



Story # 1: "So, how many did you guys do?"







Story #2: "Half a million"





# Story #3: "DNA"



# **IT And DevOps At Ginkgo Scale, or AWS The Ginkgo Way**



# **Imagine you're me, running IT-DevOps @Ginkgo**

- 1. Our output triples every year!!**
- 2. Sequencers can grab 1/2 bandwidth of the network core, and produce over a dozen Tb/day**
- 3. In the near future 1Pb/year. ... then 3Pb/year ... then 9 Pb/year...**



2014





2015





2016





2017






2019



2019





# Other Ginkgo-Scale Requirements

- **The foundry has a service architecture**
- **Scientists work on multiple projects simultaneously**
- **Pivot from one project to the other several times a day**
- **Can not handle any interruption in service**
- **Data from robots, lab equipment, and software need to be accessible immediately**
  - **No data latency, ever**
- **Multiple sessions of the same tools open at the same time (e.g. Jupyter notebooks)**
- **300 people, tens of thousands of sessions**
- **Hundreds of sensors**



Also, we're on a dock. And ...











MASSPORT  
POLICE

EMERGENCY SERVICE

ESU1

D STREET

SPEED  
LIMIT  
25



# @IT-DevOps

How do we support  
~300 Ginkgo scientists  
and everyone else with 8  
IT-DevOps people?

3 are all we have for  
desktop, lab computers,  
WiFi, printers, ...





# @DevOps

How do we support major AWS workloads from ...

- > 30 applications
- >20 software engineers
- >10 DNA designers
- A dozen DNA fabricators (2 shifts)
- >10 NGS scientists and operators
- A dozen robot engineers
- 140+ robots
- Lots of random scientists

... with 5 DevOps engineers?





We do it the Ginkgo  
Way

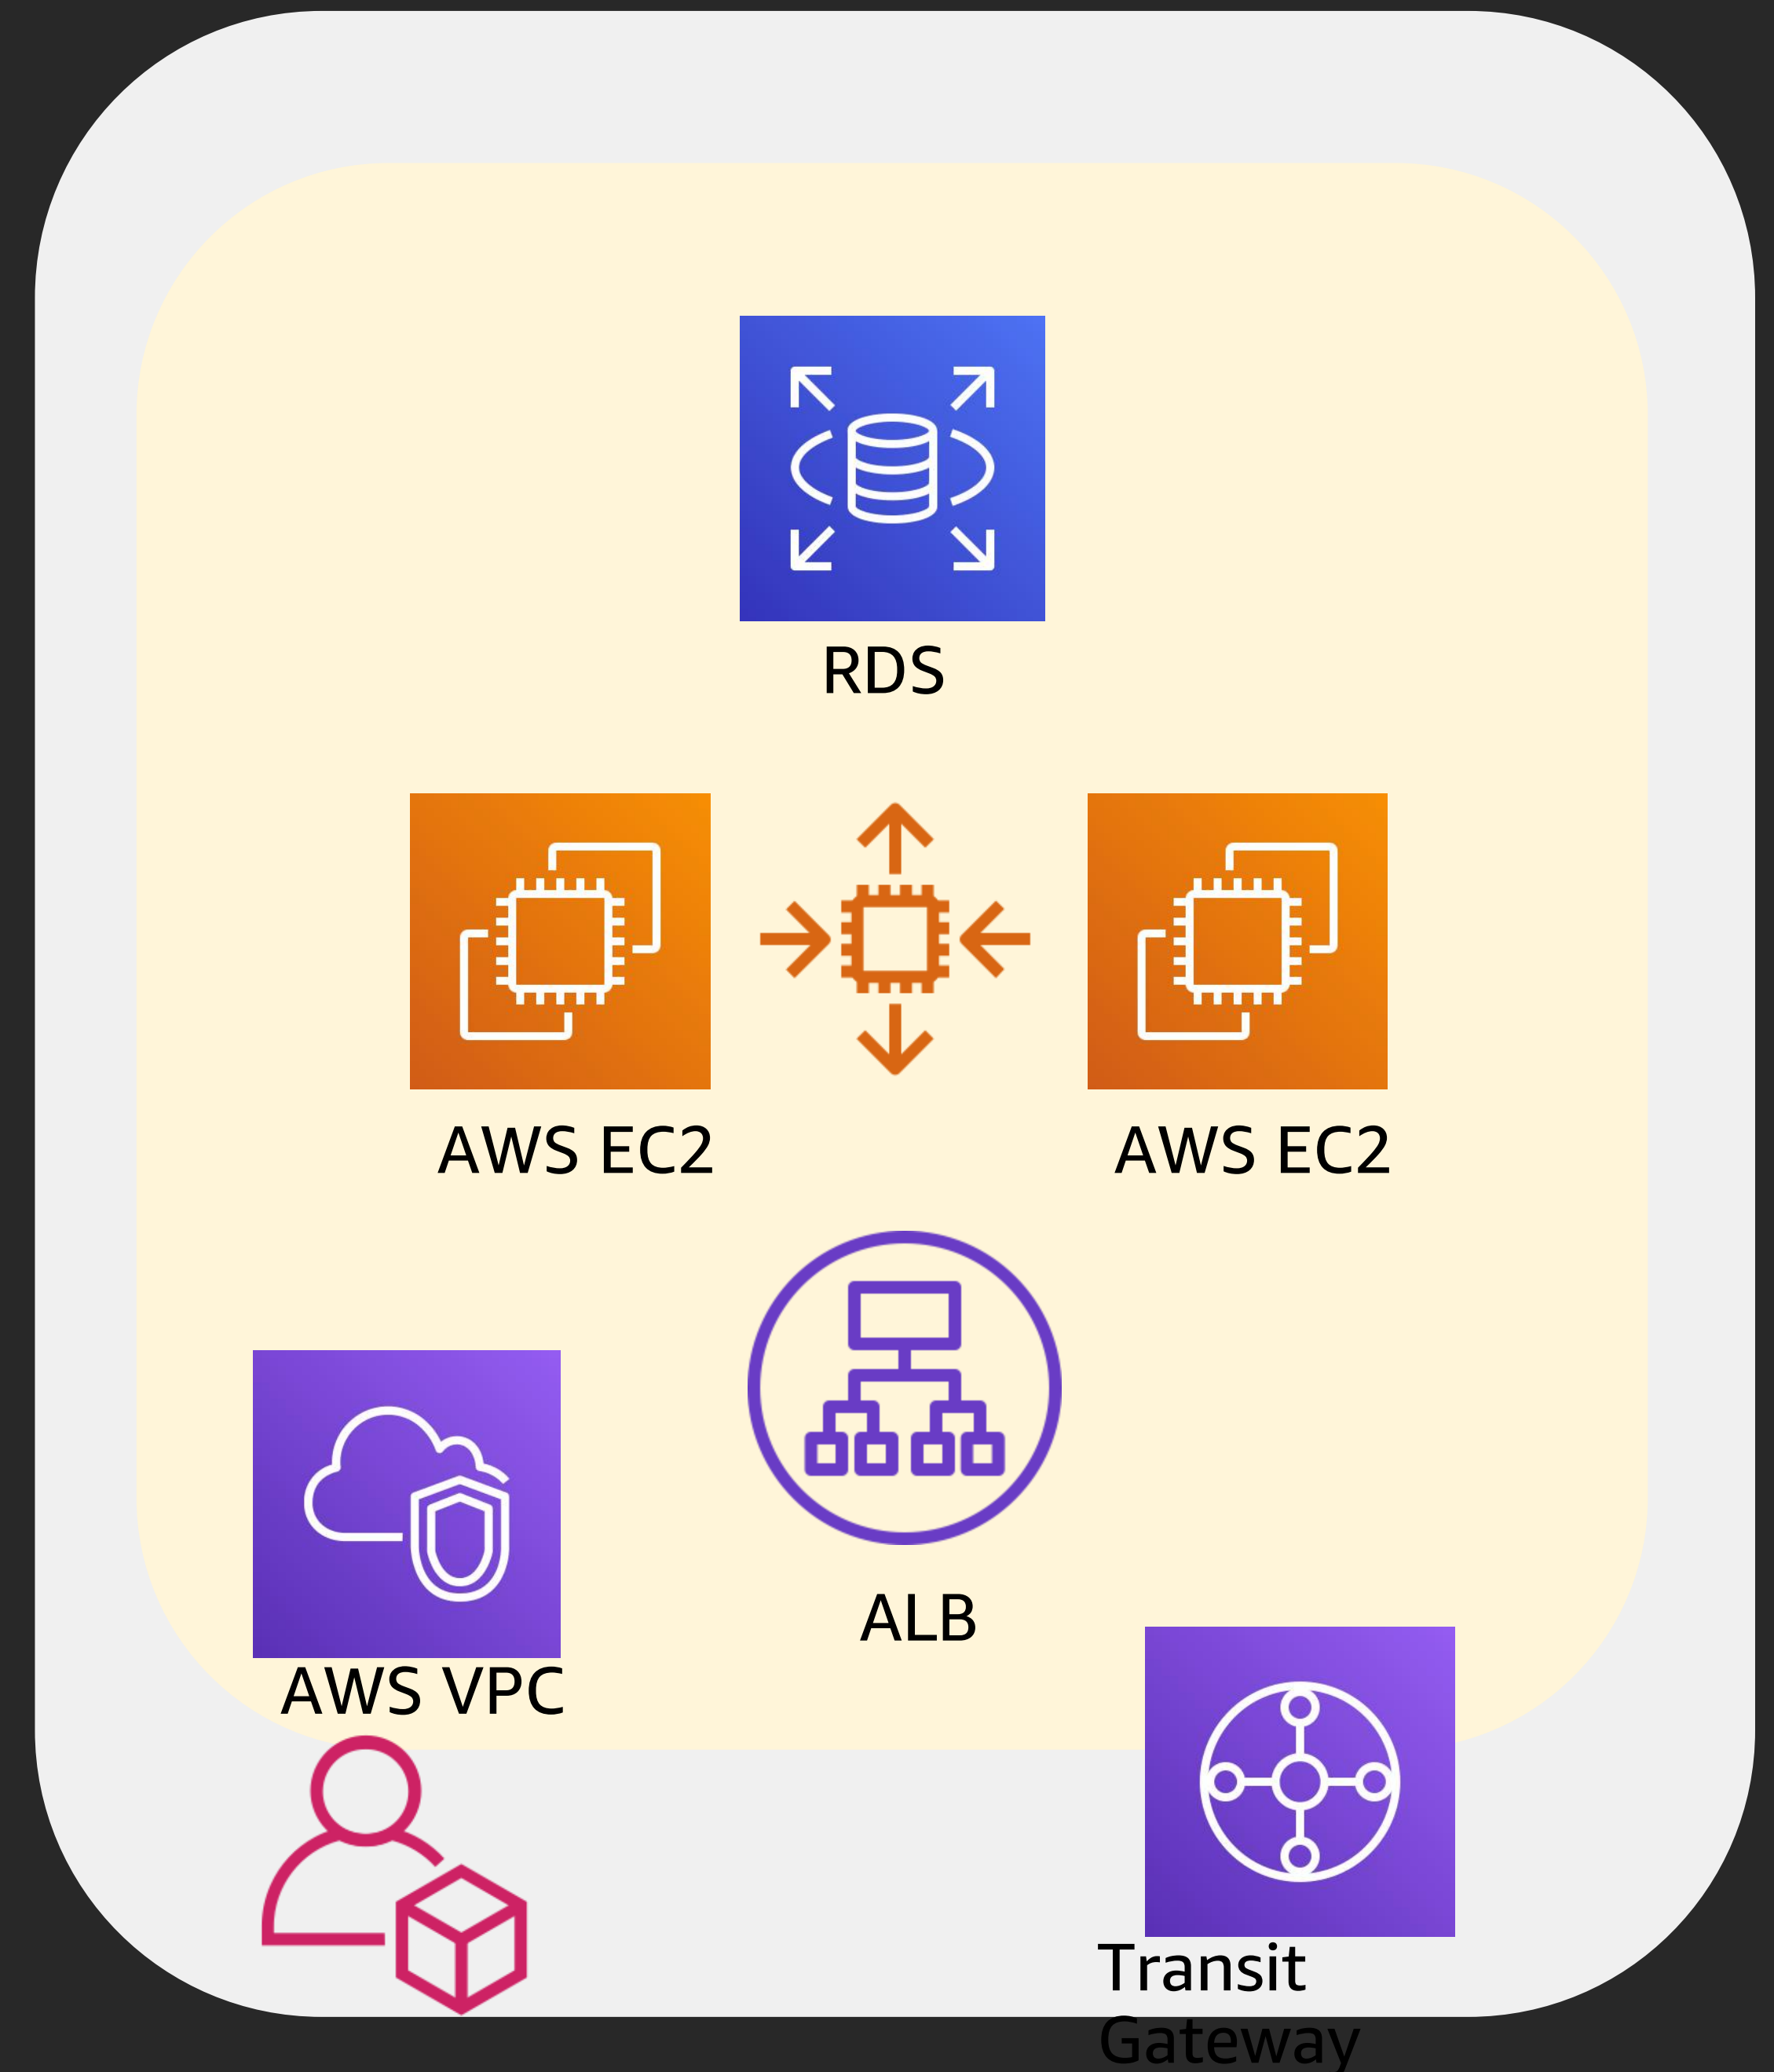
# Extreme Automation

```
TrailName:
  Fn::Sub:
    "${AWS::StackName}-global-trail"
EnableLogFileValidation: true
IncludeGlobalServiceEvents: true
IsLogging: true
IsMultiRegionTrail: true
S3BucketName:
  Ref: loggingBucket
S3KeyPrefix: cloudtrail
Tags:
  - Key: Team
    Value: IT
  - Key: Contact
    Value: "it@ginkgobioworks.com"
  - Key: Workloadtype
    Value: prod
  - Key: Project
    Value: IT
#####
# backup bucket
#
# Backup bucket keeps things for 90 days, then they
are
```



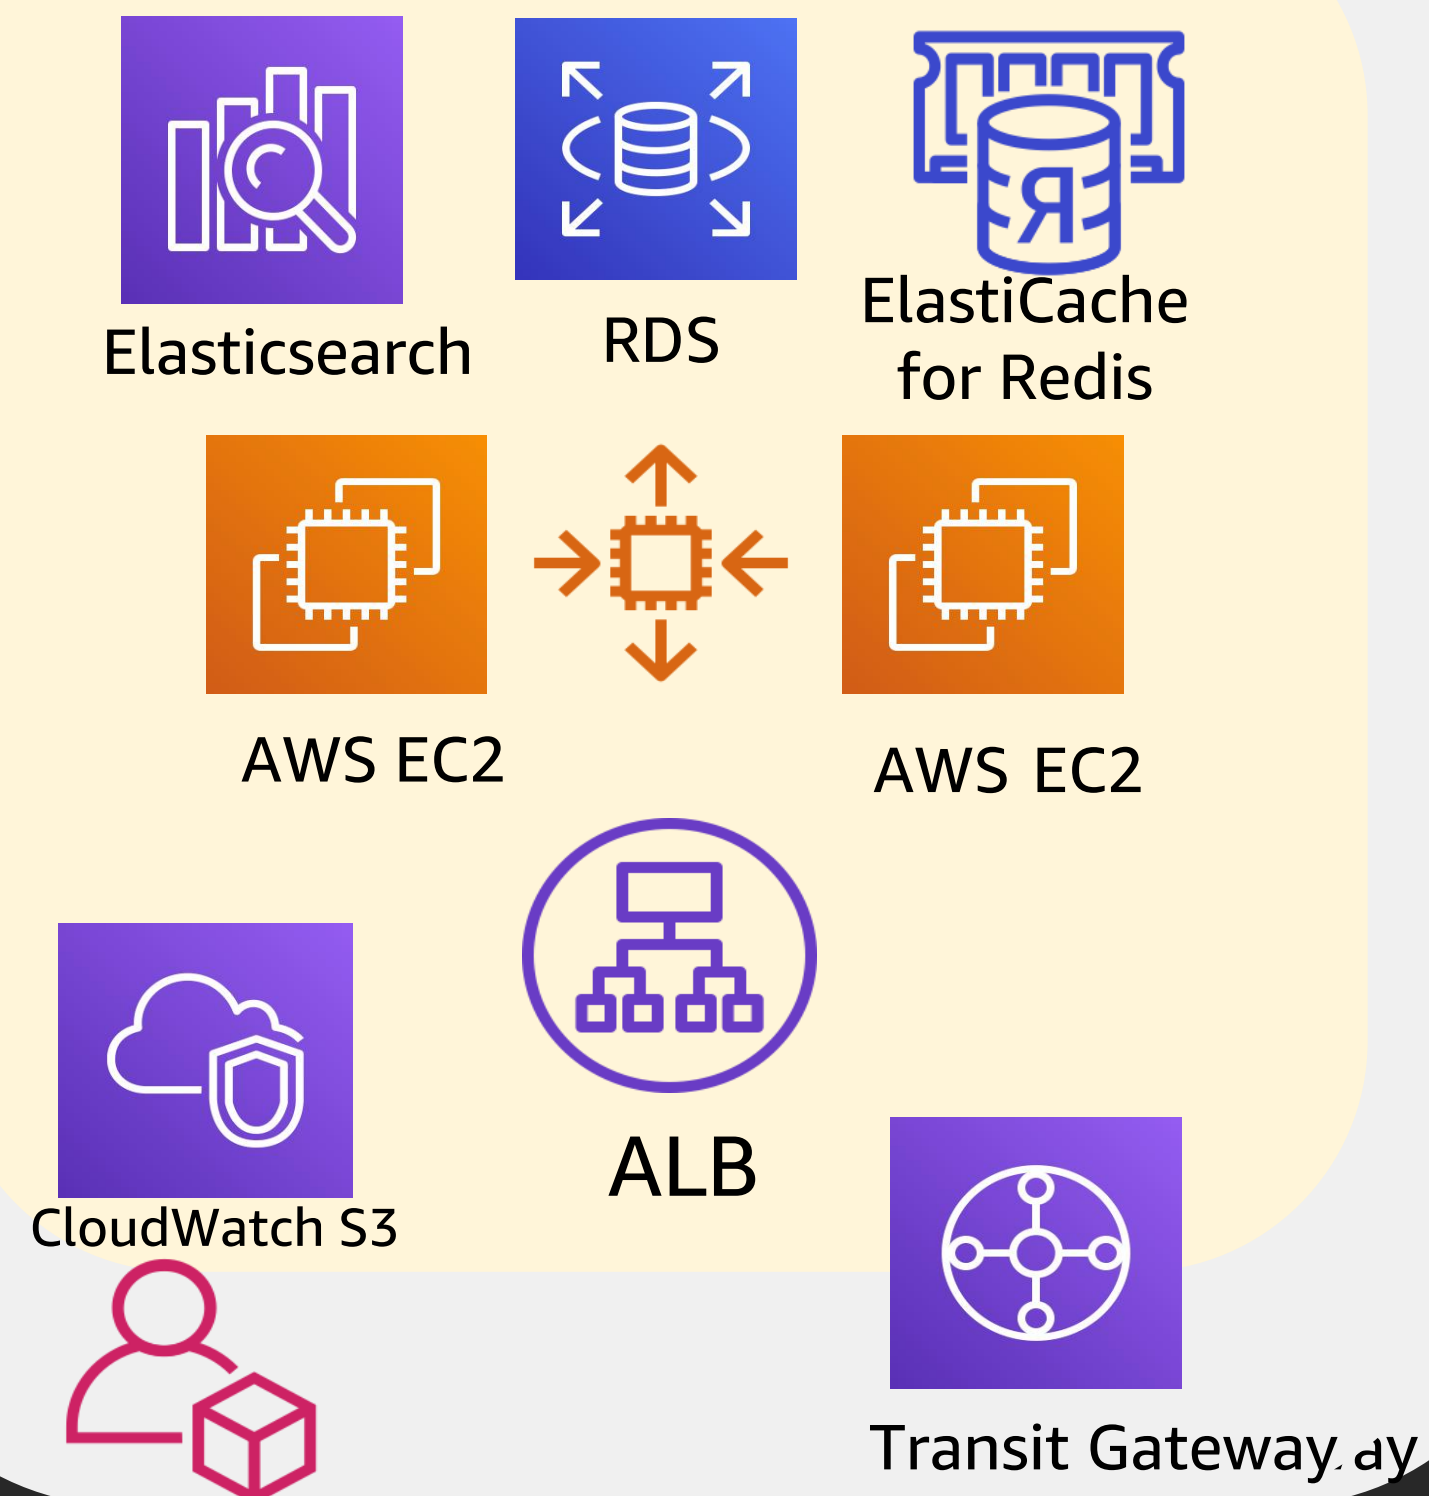
We build applications  
very simply.

We stamp them out  
like Grandma's  
holiday cookies





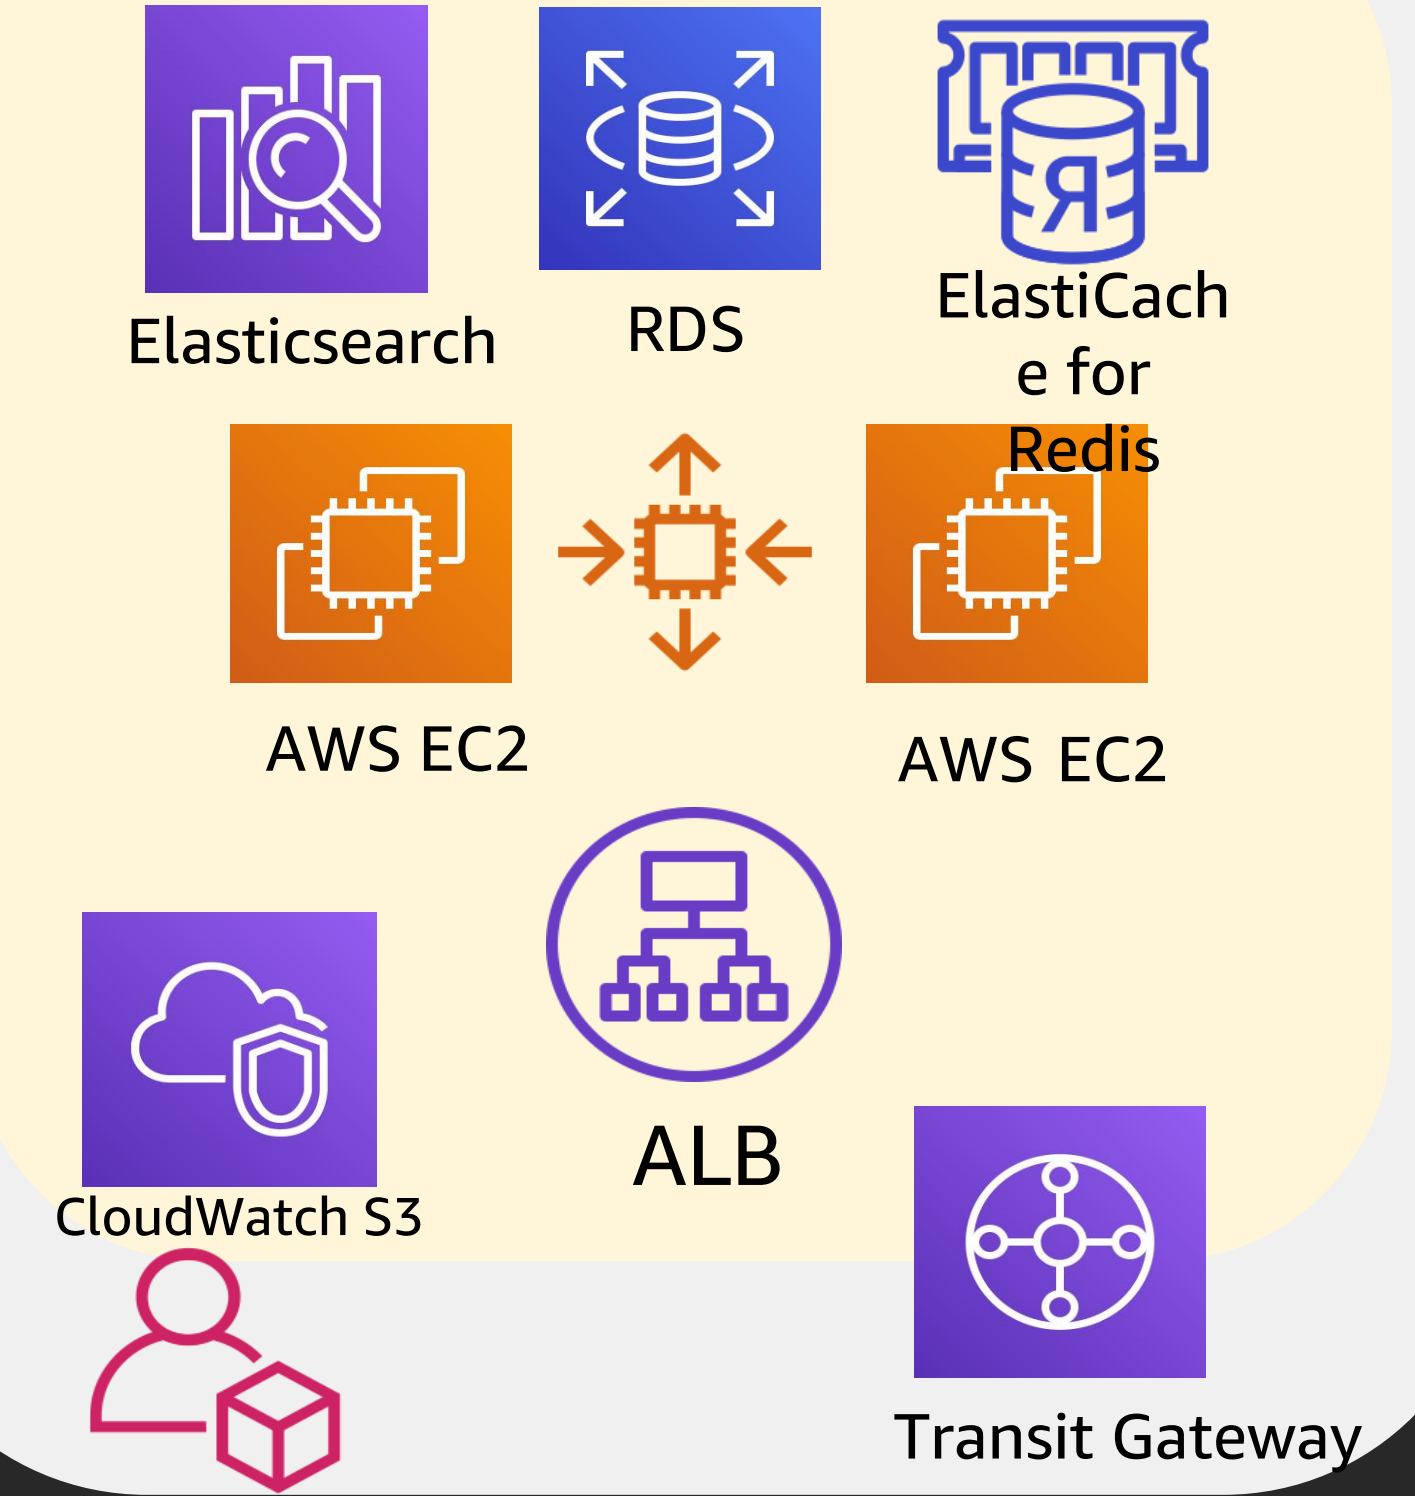
# Multi-account architecture



App Dev



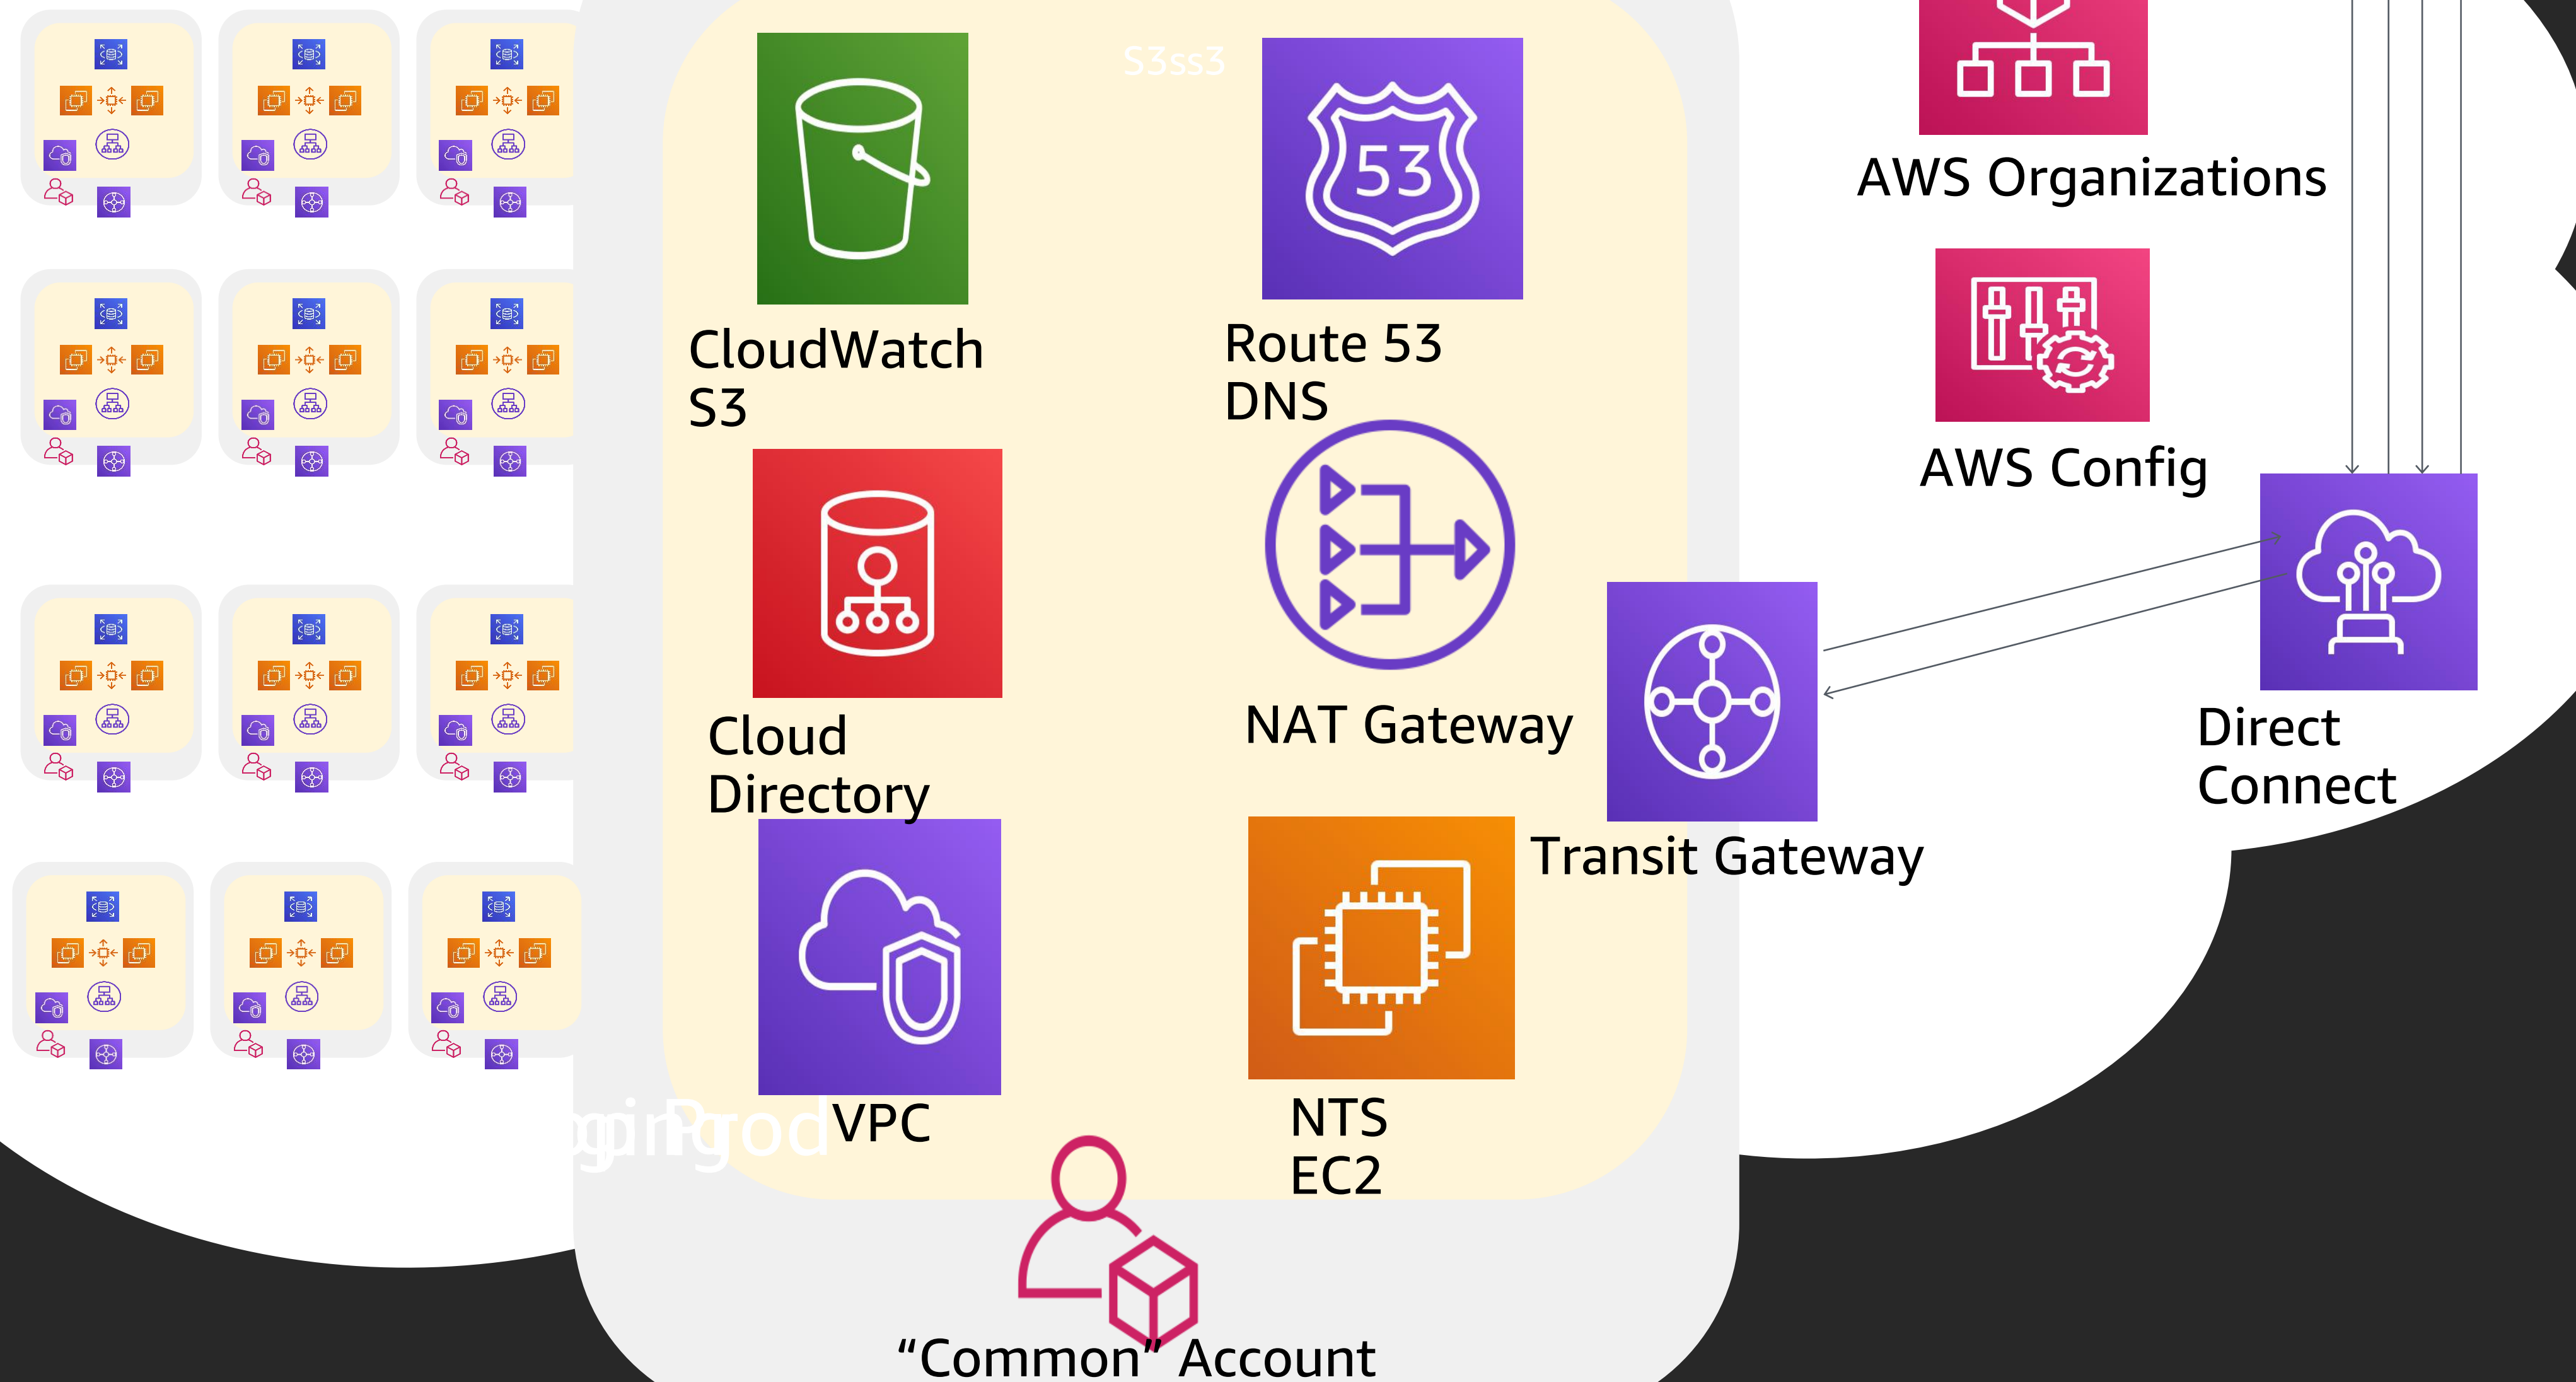
App Staging



App Prod



# Ginkgo AWS Architecture





NGS is one of IT-  
DevOps' favorite  
services





NGS is one of IT-  
DevOps' favorite  
services

1. We use it to size our  
network





NGS is one of IT-DevOps' favorite services

1. We use it to size our network
2. It demonstrates the awesome power of AWS services





"A very impressive attempt. However, the public demo is hardly usable and the development try-it-out deployment is broken due to dependencies that are incompatible. Unable to assess complex workflows or retry capability as the project is full of dated documentation and broken links. I think they simply couldn't keep up with the task at hand given the scope of their vision and the project is ultimately a failure now that the company was bought. Useful to know they spent millions developing it as a benchmark."

"Desktop based UI for creating workflows that allows user interaction and composition of tools. Apache is creating a web based system so it might be good to check on this project periodically. Not really clear if it will ever manage execution state well."

"UI for running jobs, but ultimately too simple"

"The first I found that is a cloud first workflow system. Handles failures well. Can create and update workflows in the UI. It looks very nice, but I know literally nothing about the people implementing it so I'm hesitant."

"No UI, command line driven. Would need to implement an introspection layer".

"A composer for CWL workflows that then submits the CWL to a task execution engine. Onus would be on a separate product to run the CWL and to manage its state."

**"Project is in maintenance mode it appears. <xxx> seems to be successor. On a separate note, I really didn't like how tasks were designed -- I prefer function based tasks instead of class based tasks because it's easier to eject from a framework and re-use code."**

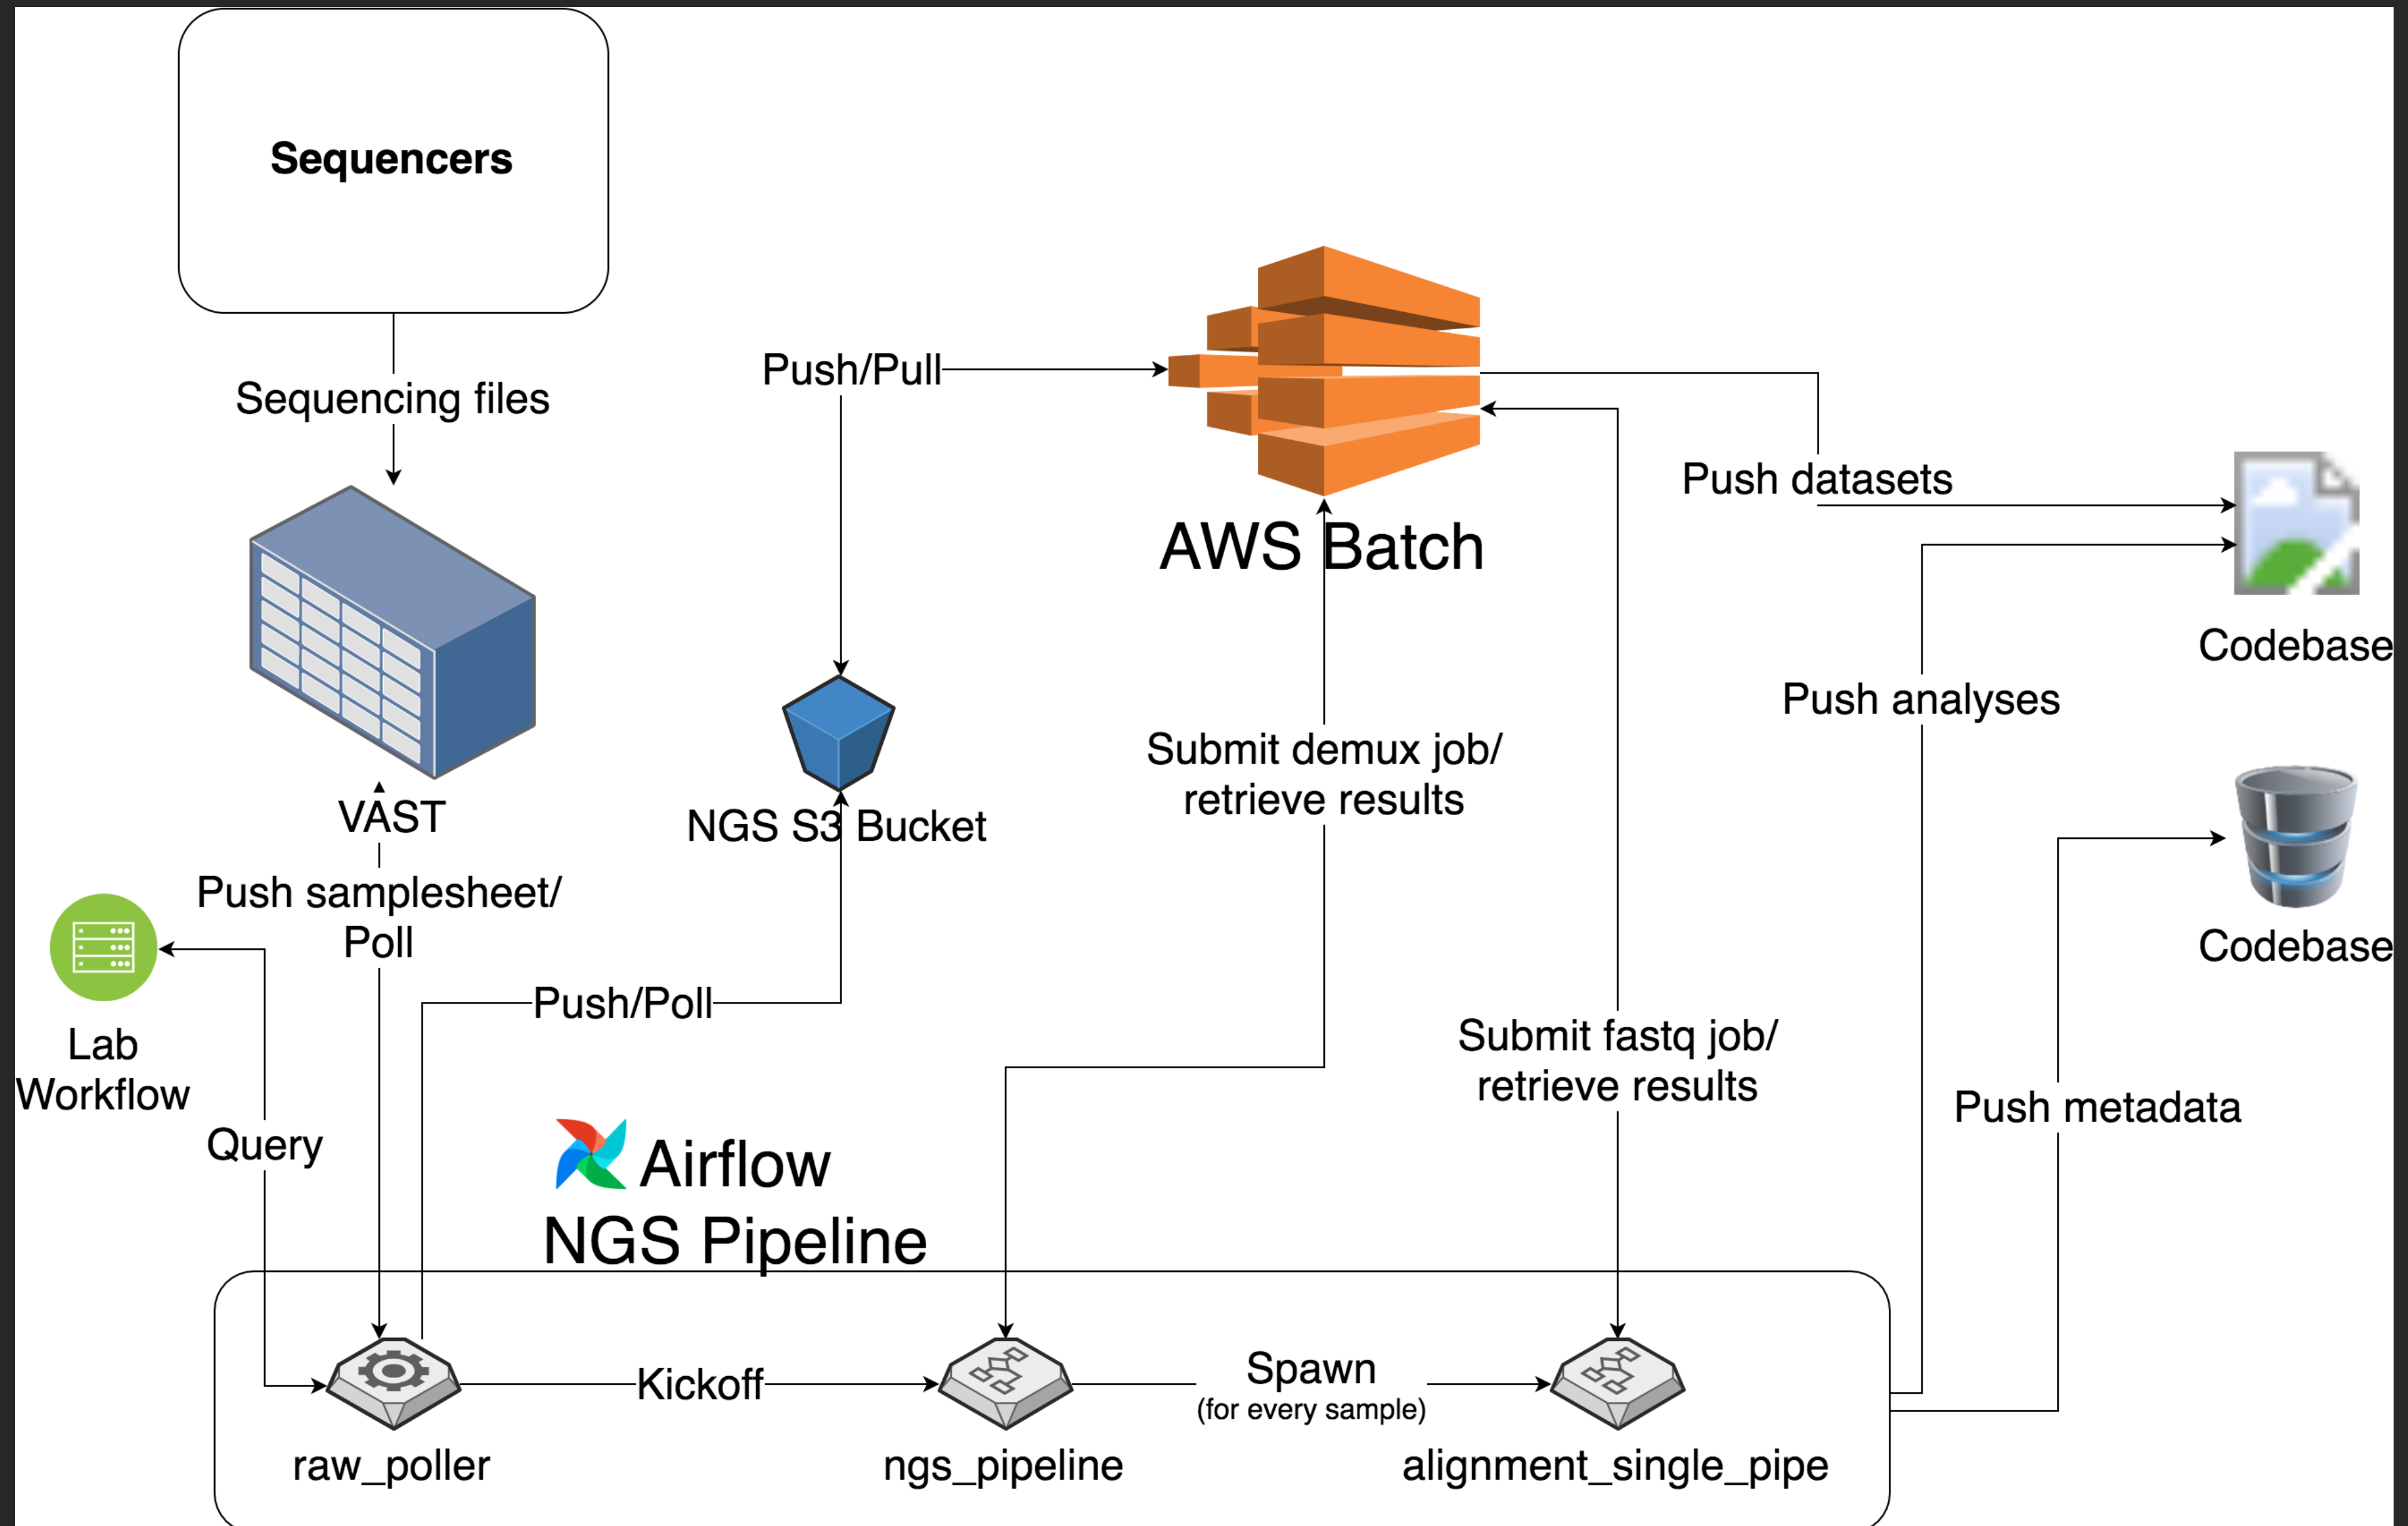
**"No UI. Academic project, their production build was broken for months. I have little faith in it." (edited)**

**"No UI, no introspection layer, and entirely command line driven ."**

**"I also don't really like things that use CWL/WDL, my outsider perspective is it's being built in an echo chamber by 3 groups."**



# So instead, we use... AWS Batch!





# Ginkgo's Favorite AWS Services




AWS CloudFormation



AWS Direct  
Connect




AWS Transit Gateway



AWS Systems  
Manager




Amazon Elastic Container  
Registry




AWS CloudTrail



Amazon CloudWatch



AWS Config



AWS Organizations



The clear winner is ...



AWS Batch




# Ginkgo could not scale without AWS

- No other cloud provider

has the type and multiplicity of services that Ginkgo needs

We don't just need Tesseract and Deep Mind  
We need EVERYTHING

AWS has the fanatical customer service orientation  
Ginkgo needs to continue to achieve Knight's Law





# Partners At Scale





# 3 Keys To AWS Success



# 3 Keys To AWS Success

## 1. Enterprise Service Agreement



# 3 Keys To AWS Success

1. 1. Enterprise Service Agreement

2. 2. Enterprise Service Agreement



# 3 Keys To AWS Success

1. Enterprise Service Agreement.
2. Enterprise Service Agreement!
3. Enterprise Service Agreement!!!!



# Why Enterprise Service Agreement?





A photograph of a city street during a heavy snowfall. The ground is covered in a thick layer of snow, with some patches of ice visible. In the background, there is a brick building with a sign that reads "75" and "A BIRD & GRILL". A traffic light is visible on the right side of the frame. A red text box is overlaid on the bottom right of the image.

**Why Enterprise Service Agreement?**

**Because AWS is *HARD!!!***



# Use AWS best-practices

- Multi-account strategy
- Compliance guidelines
- Cloud management services and guidelines
  - Organizations, Config, CloudWatch...
- AWS is just another datacenter. Security is still your problem.



# Next Level AWS

Try a service and don't like it?

- Complain to AWS (open a Case), tell them your requirements.
- Wait a year, try it again.





# GINKGO BIOWORKS

THE ORGANISM COMPANY

“Make Biology easier to engineer”





# GINKGO BIOWORKS

THE ORGANISM COMPANY

“Make Biology easier to engineer”

Come join us!!



# Thank you!

**Dave Treff**

dave@ginkgobioworks.com