

**WITH GREAT  
SCALABILITY  
COMES GREAT  
RESPONSIBILITY**

Dana Engebretson







# Queen of the Amazons

@BigDana






**Amazon Warriors\***



**\*Our services in AWS**







**@BigDana**



**Dana Engebretson  
Performance Engineer  
SPS Commerce**



# Building a Data Pipeline

**Attempt 0: Python Multiprocessing on an EC2 Instance**

**Attempt 0.5: Python Multiprocessing on a Spot Instance**

**Attempt 1: Lambdas Orchestrated by Step Functions**

**Attempt 2: Using Event Triggers with S3**

**Attempt 3: Rate limiting with Queues and Cloudwatch triggered Lambdas**





# Building a Data Pipeline


**Attempt 0: Python Multiprocessing on an EC2 Instance**

**Attempt 0.5: Python Multiprocessing on a Spot Instance**

**Attempt 1: Lambdas Orchestrated by Step Functions**

**Attempt 2: Using Event Triggers with S3**

**Attempt 3: Rate limiting with Queues and Cloudwatch triggered Lambdas**





# Building a Data Pipeline


**Attempt 0: Python Multiprocessing on an EC2 Instance**

**Attempt 0.5: Python Multiprocessing on a Spot Instance**

**Attempt 1: Lambdas Orchestrated by Step Functions**

**Attempt 2: Using Event Triggers with S3**

**Attempt 3: Rate limiting with Queues and Cloudwatch triggered Lambdas**





# Building a Data Pipeline


**Attempt 0: Python Multiprocessing on an EC2 Instance**

**Attempt 0.5: Python Multiprocessing on a Spot Instance**

**Attempt 1: Lambdas Orchestrated by Step Functions**

**Attempt 2: Using Event Triggers with S3**

**Attempt 3: Rate limiting with Queues and Cloudwatch triggered Lambdas**





# Building a Data Pipeline


**Attempt 0: Python Multiprocessing on an EC2 Instance**

**Attempt 0.5: Python Multiprocessing on a Spot Instance**

**Attempt 1: Lambdas Orchestrated by Step Functions**

**Attempt 2: Using Event Triggers with S3**

**Attempt 3: Rate limiting with Queues and Cloudwatch triggered Lambdas**





# Building a Data Pipeline

**Attempt 0: Python Multiprocessing on an EC2 Instance**

**Attempt 0.5: Python Multiprocessing on a Spot Instance**

**Attempt 1: Lambdas Orchestrated by Step Functions**

**Attempt 2: Using Event Triggers with S3**

**Attempt 3: Rate limiting with Queues and Cloudwatch triggered Lambdas**



**The Oracle\***



**\*one of our monitoring vendors**




**The Oracle\***




**\*one of our monitoring vendors**

**Hermes: the messenger God\***




**\*Our Vendor's Api**











**The Oracle\***



**\*one of our monitoring vendors**

**Hermes: the messenger God\***



**\*Our Vendor's Api**

**FETCH ME  
ALL MY DATA  
FOR ALL MY  
AMAZON  
WARRIORS  
PLEASE!**



**The Oracle\***



**Hermes: the messenger God\***



**\*Our Vendor's Api**

**\*one of our monitoring vendors**



**The Oracle\***



**\*one of our monitoring vendors**

**Hermes: the messenger God\***



**\*Our Vendor's Api**



**The Oracle\***



**\*one of our monitoring vendors**

**Hermes: the messenger God\***



**\*Our Vendor's Api**







**The Oracle\***



**Hermes: the messenger God\***



**FETCH ME ALL MY  
AMAZON WARRIOR TROOPS\*  
PLEASE!**

**\*one of our monitoring vendors**

**\*Our Vendor's Api**

**\*grouped by micro-service**



**The Oracle\***



**\*one of our monitoring vendors**

**Hermes: the messenger God\***



**\*Our Vendor's Api**



**The Oracle\***



**\*one of our monitoring vendors**

**Hermes: the messenger God\***



**\*Our Vendor's Api**



**The Oracle\***



**\*one of our monitoring vendors**

**Hermes: the messenger God\***



**\*Our Vendor's Api**

**HERE ARE YOUR  
AMAZON WARRIOR TROOPS!**

**troop 3**




**The Oracle\***



**\*one of our monitoring vendors**

**Hermes: the messenger God\***



**\*Our Vendor's Api**



**The Oracle\***



**\*one of our monitoring vendors**

**Hermes: the messenger God\***



**\*Our Vendor's Api**

troop 1

troop 2

troop 3

**FOR TROOP 1, FETCH ME MY  
AMAZON WARRIORS PLEASE!**



**The Oracle\***



**\*one of our monitoring vendors**

**Hermes: the messenger God\***



**\*Our Vendor's Api**

**troop 1**

**troop 2**

**troop 3**



**The Oracle\***



**\*one of our monitoring vendors**

**Hermes: the messenger God\***



**\*Our Vendor's Api**

**troop 1**

**troop 2**

**troop 3**

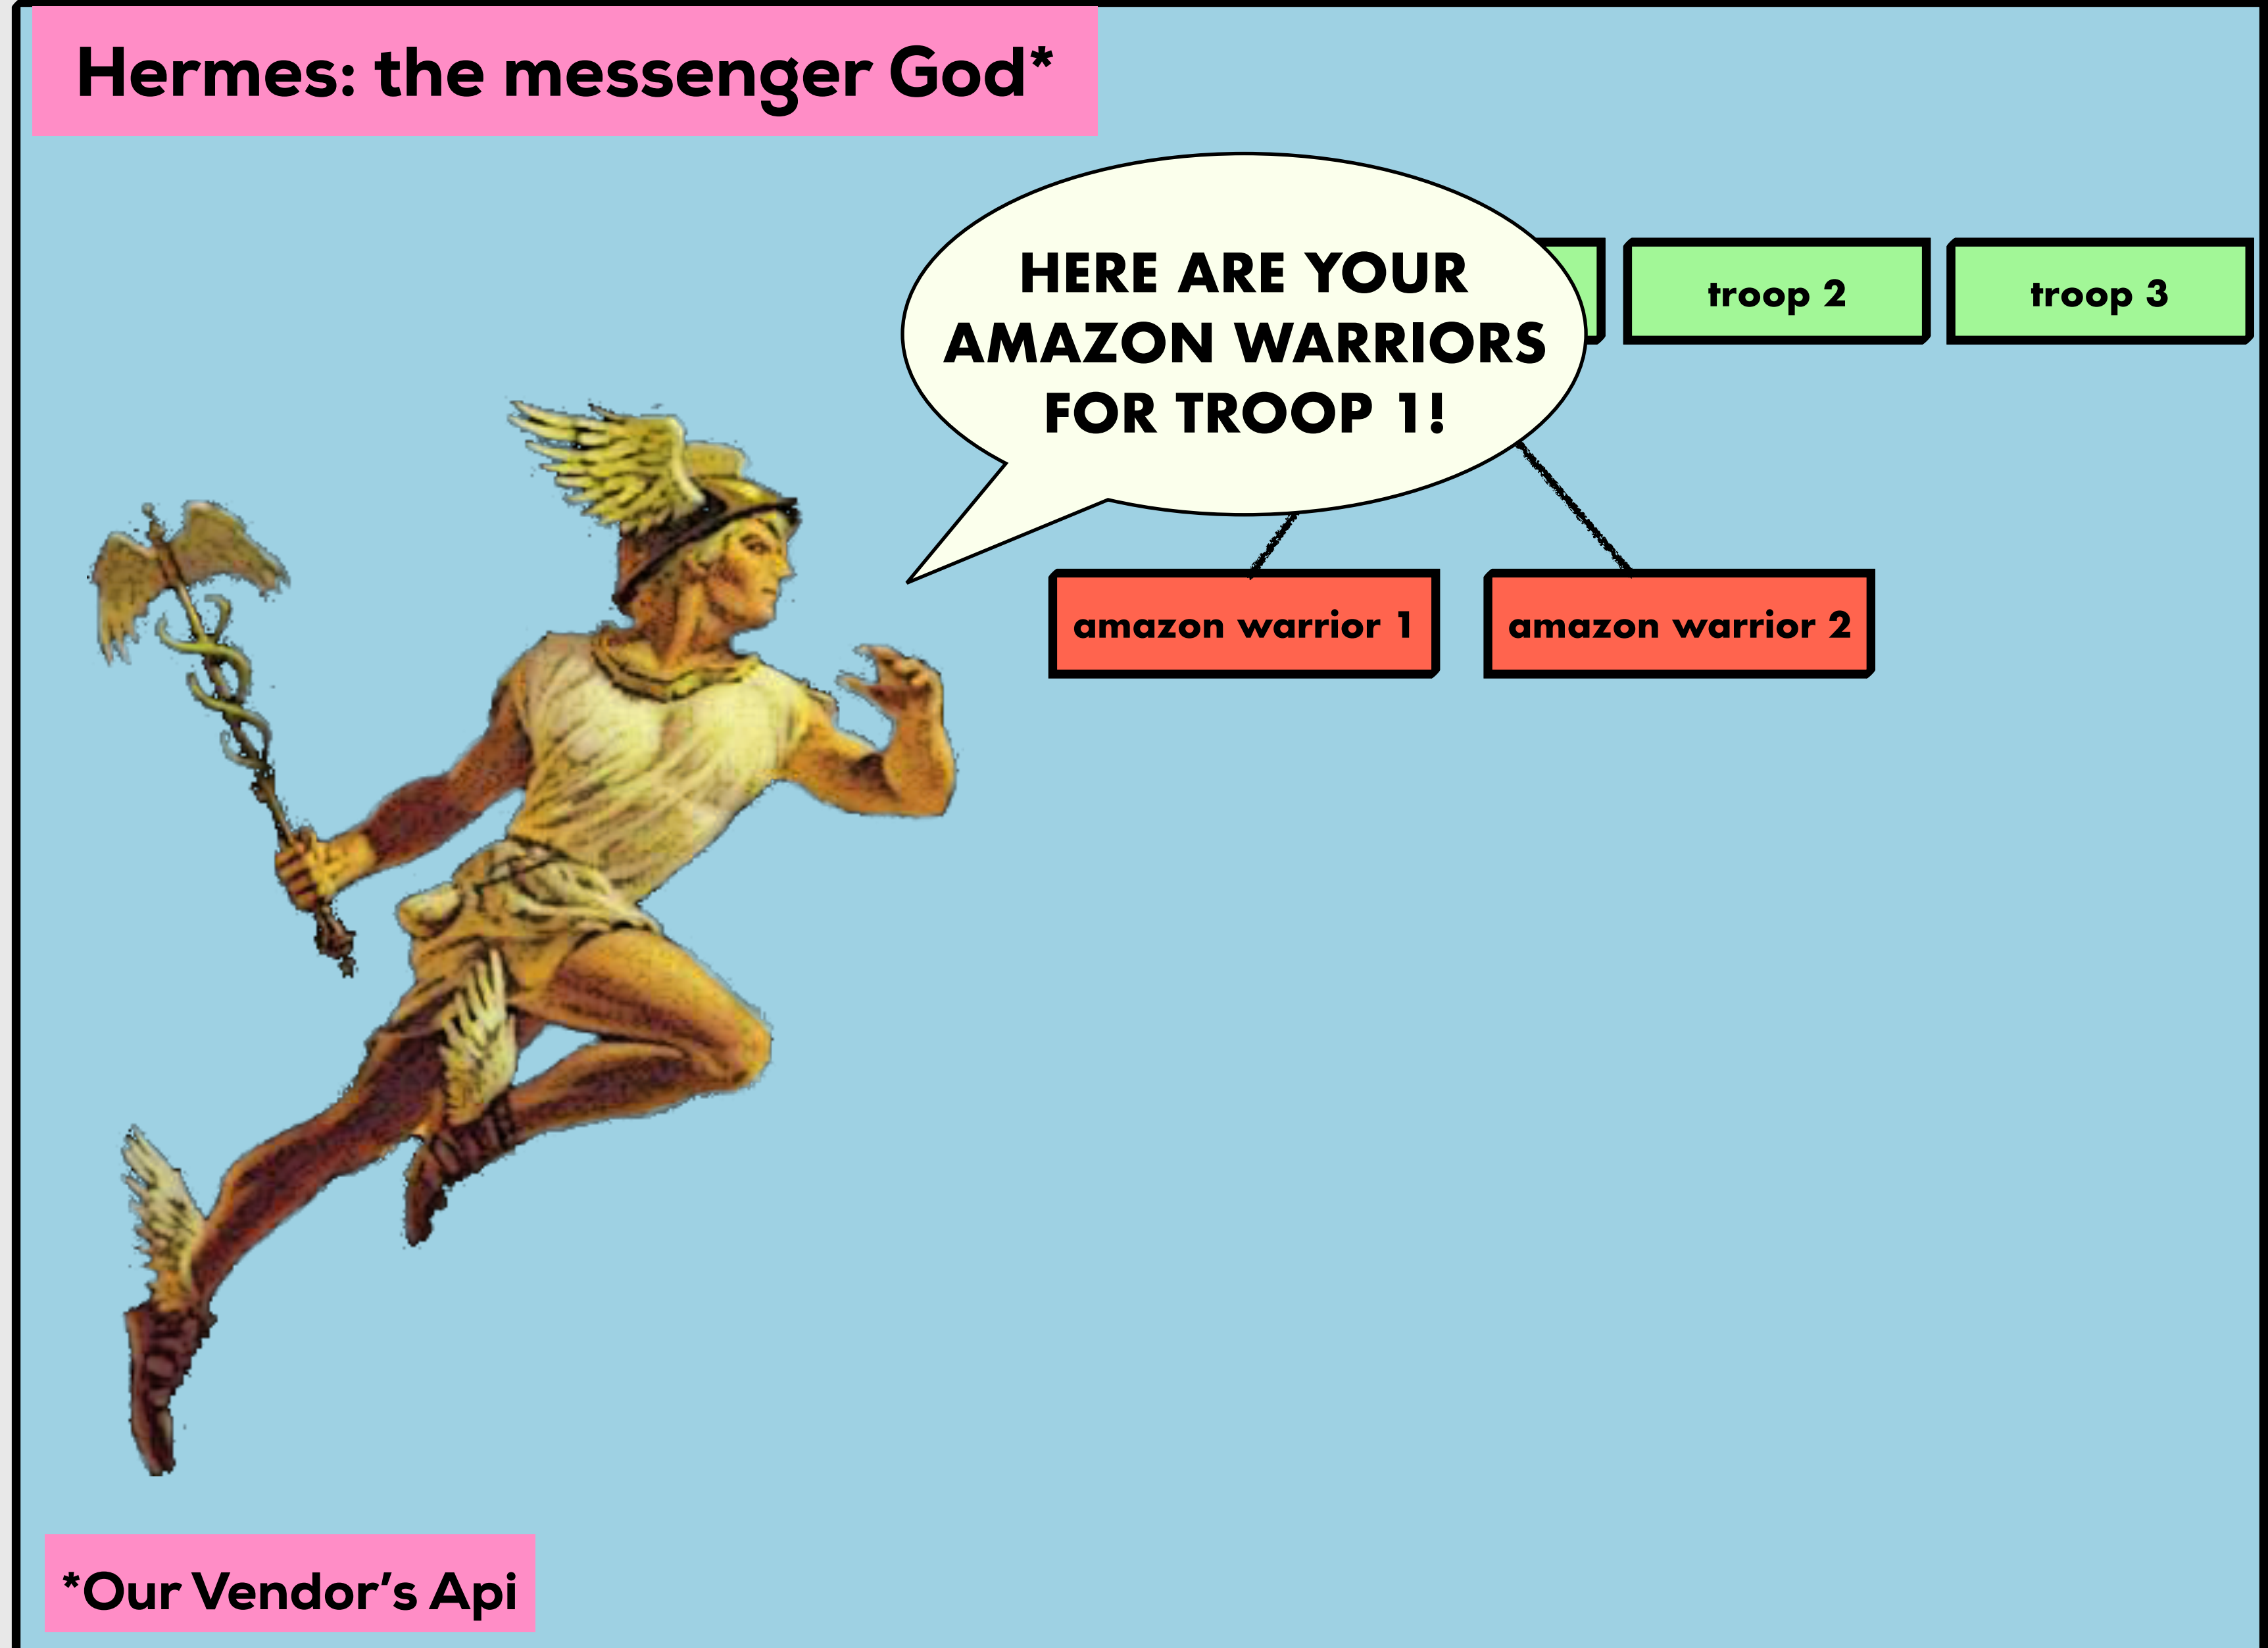


**The Oracle\***



**\*one of our monitoring vendors**

**Hermes: the messenger God\***



**\*Our Vendor's Api**



**The Oracle\***

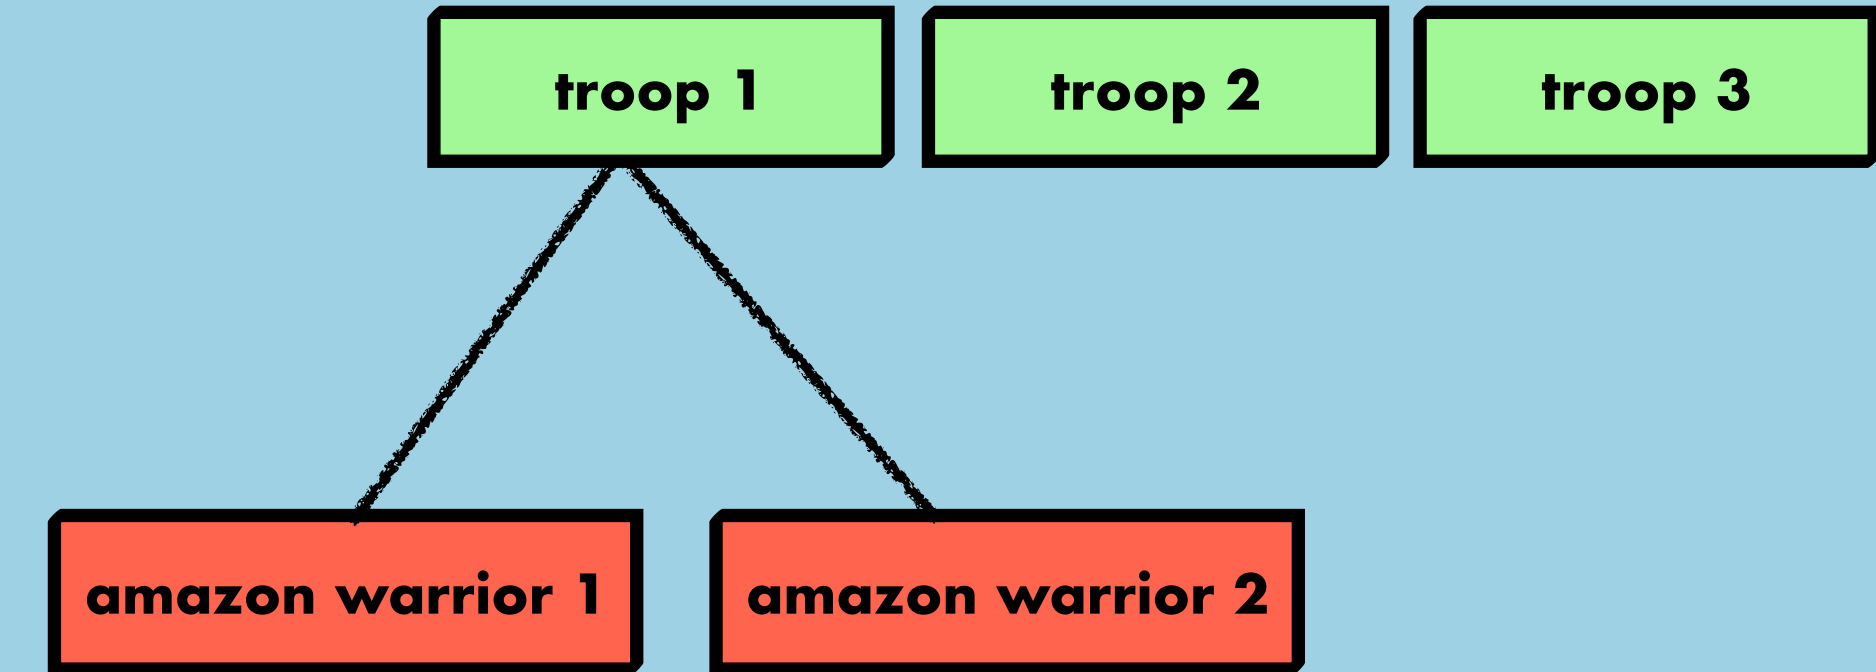


**\*one of our monitoring vendors**

**Hermes: the messenger God\***



**\*Our Vendor's Api**





**The Oracle\***



**\*one of our monitoring vendors**

**Hermes: the messenger God\***

An illustration of Hermes, the messenger god, running. He is wearing a winged helmet and a white tunic. He is holding a caduceus (a staff with a snake and wings) in his right hand. The background is a light blue gradient.

troop 1    troop 2    troop 3

amazon warrior 1    amazon warrior 2

**FOR AMAZON WARRIOR 1,  
FETCH ME MY DATA GROUPS  
PLEASE!**

**\*Our Vendor's Api**



**The Oracle\***

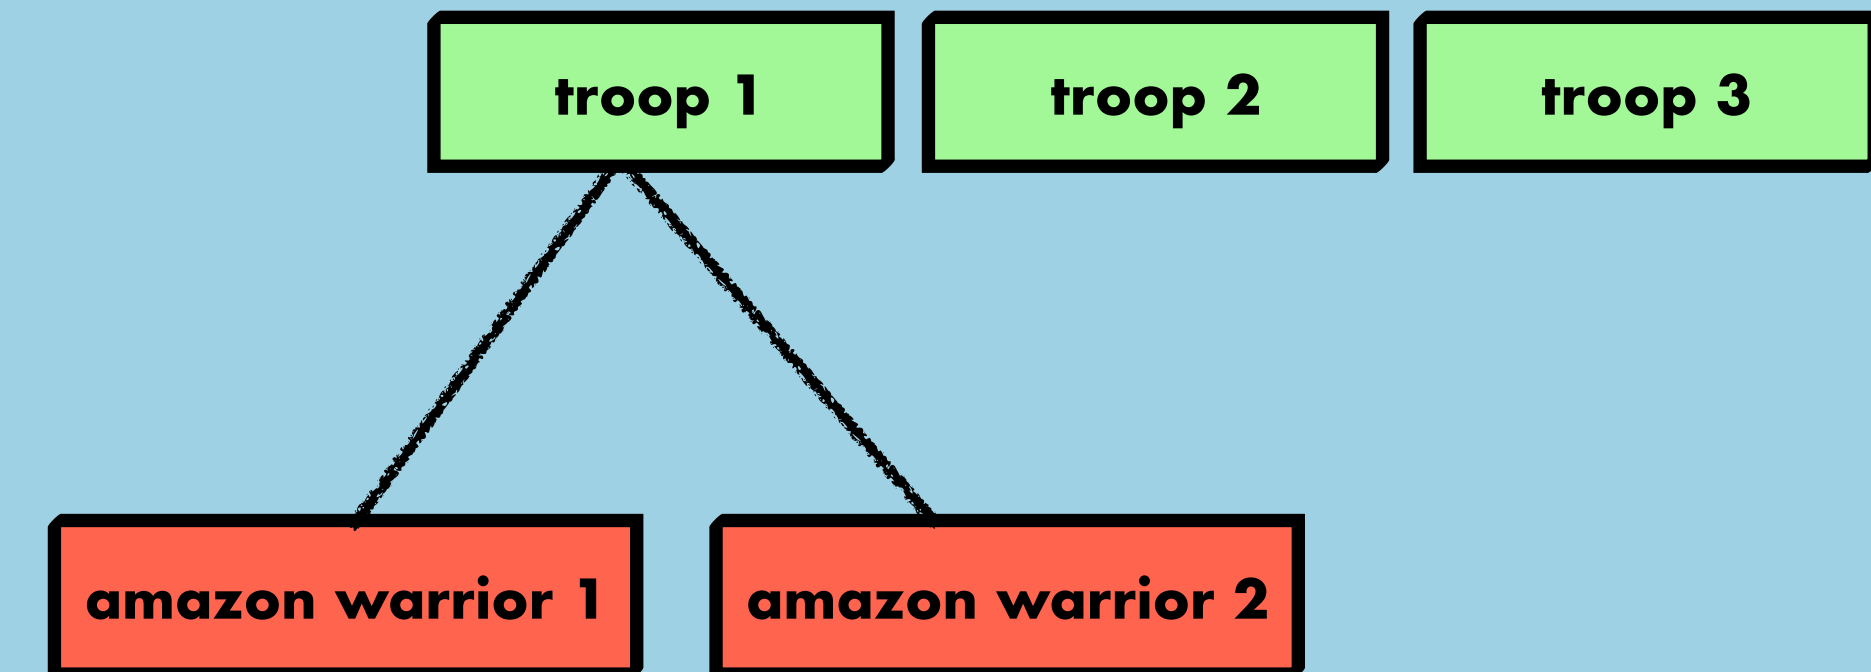


**\*one of our monitoring vendors**

**Hermes: the messenger God\***



**\*Our Vendor's Api**



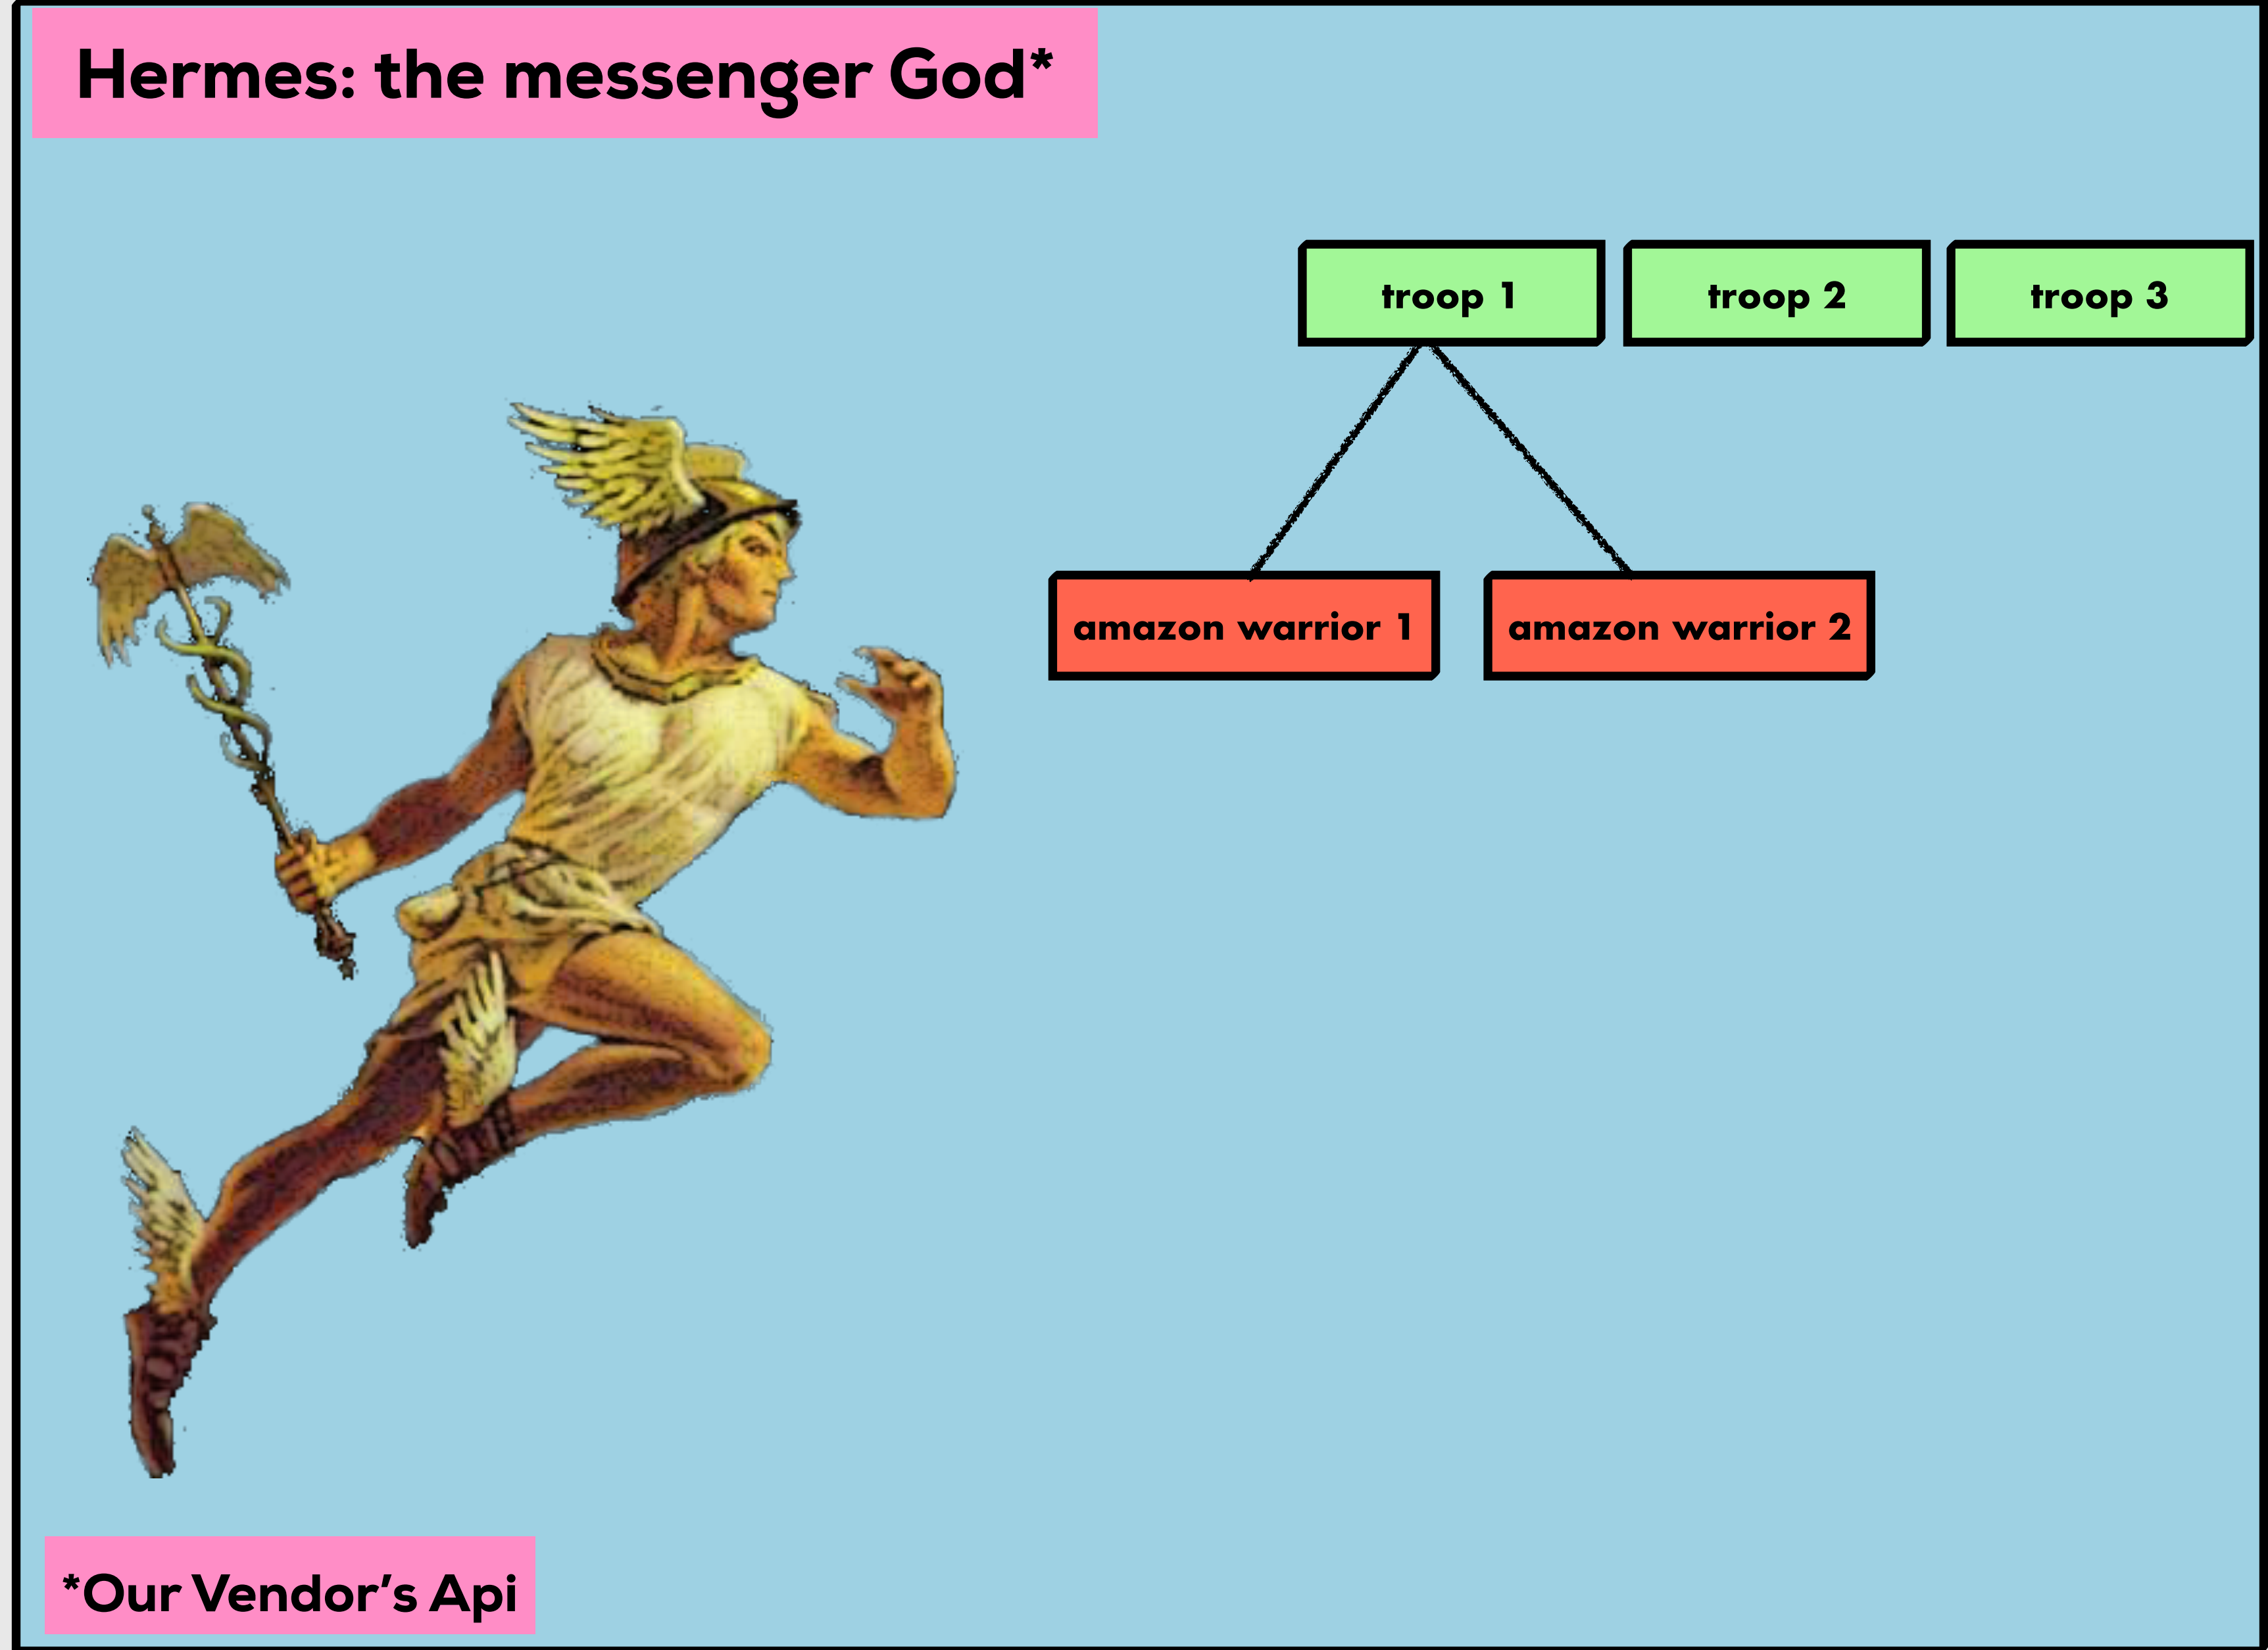


**The Oracle\***



**\*one of our monitoring vendors**

**Hermes: the messenger God\***



**\*Our Vendor's Api**

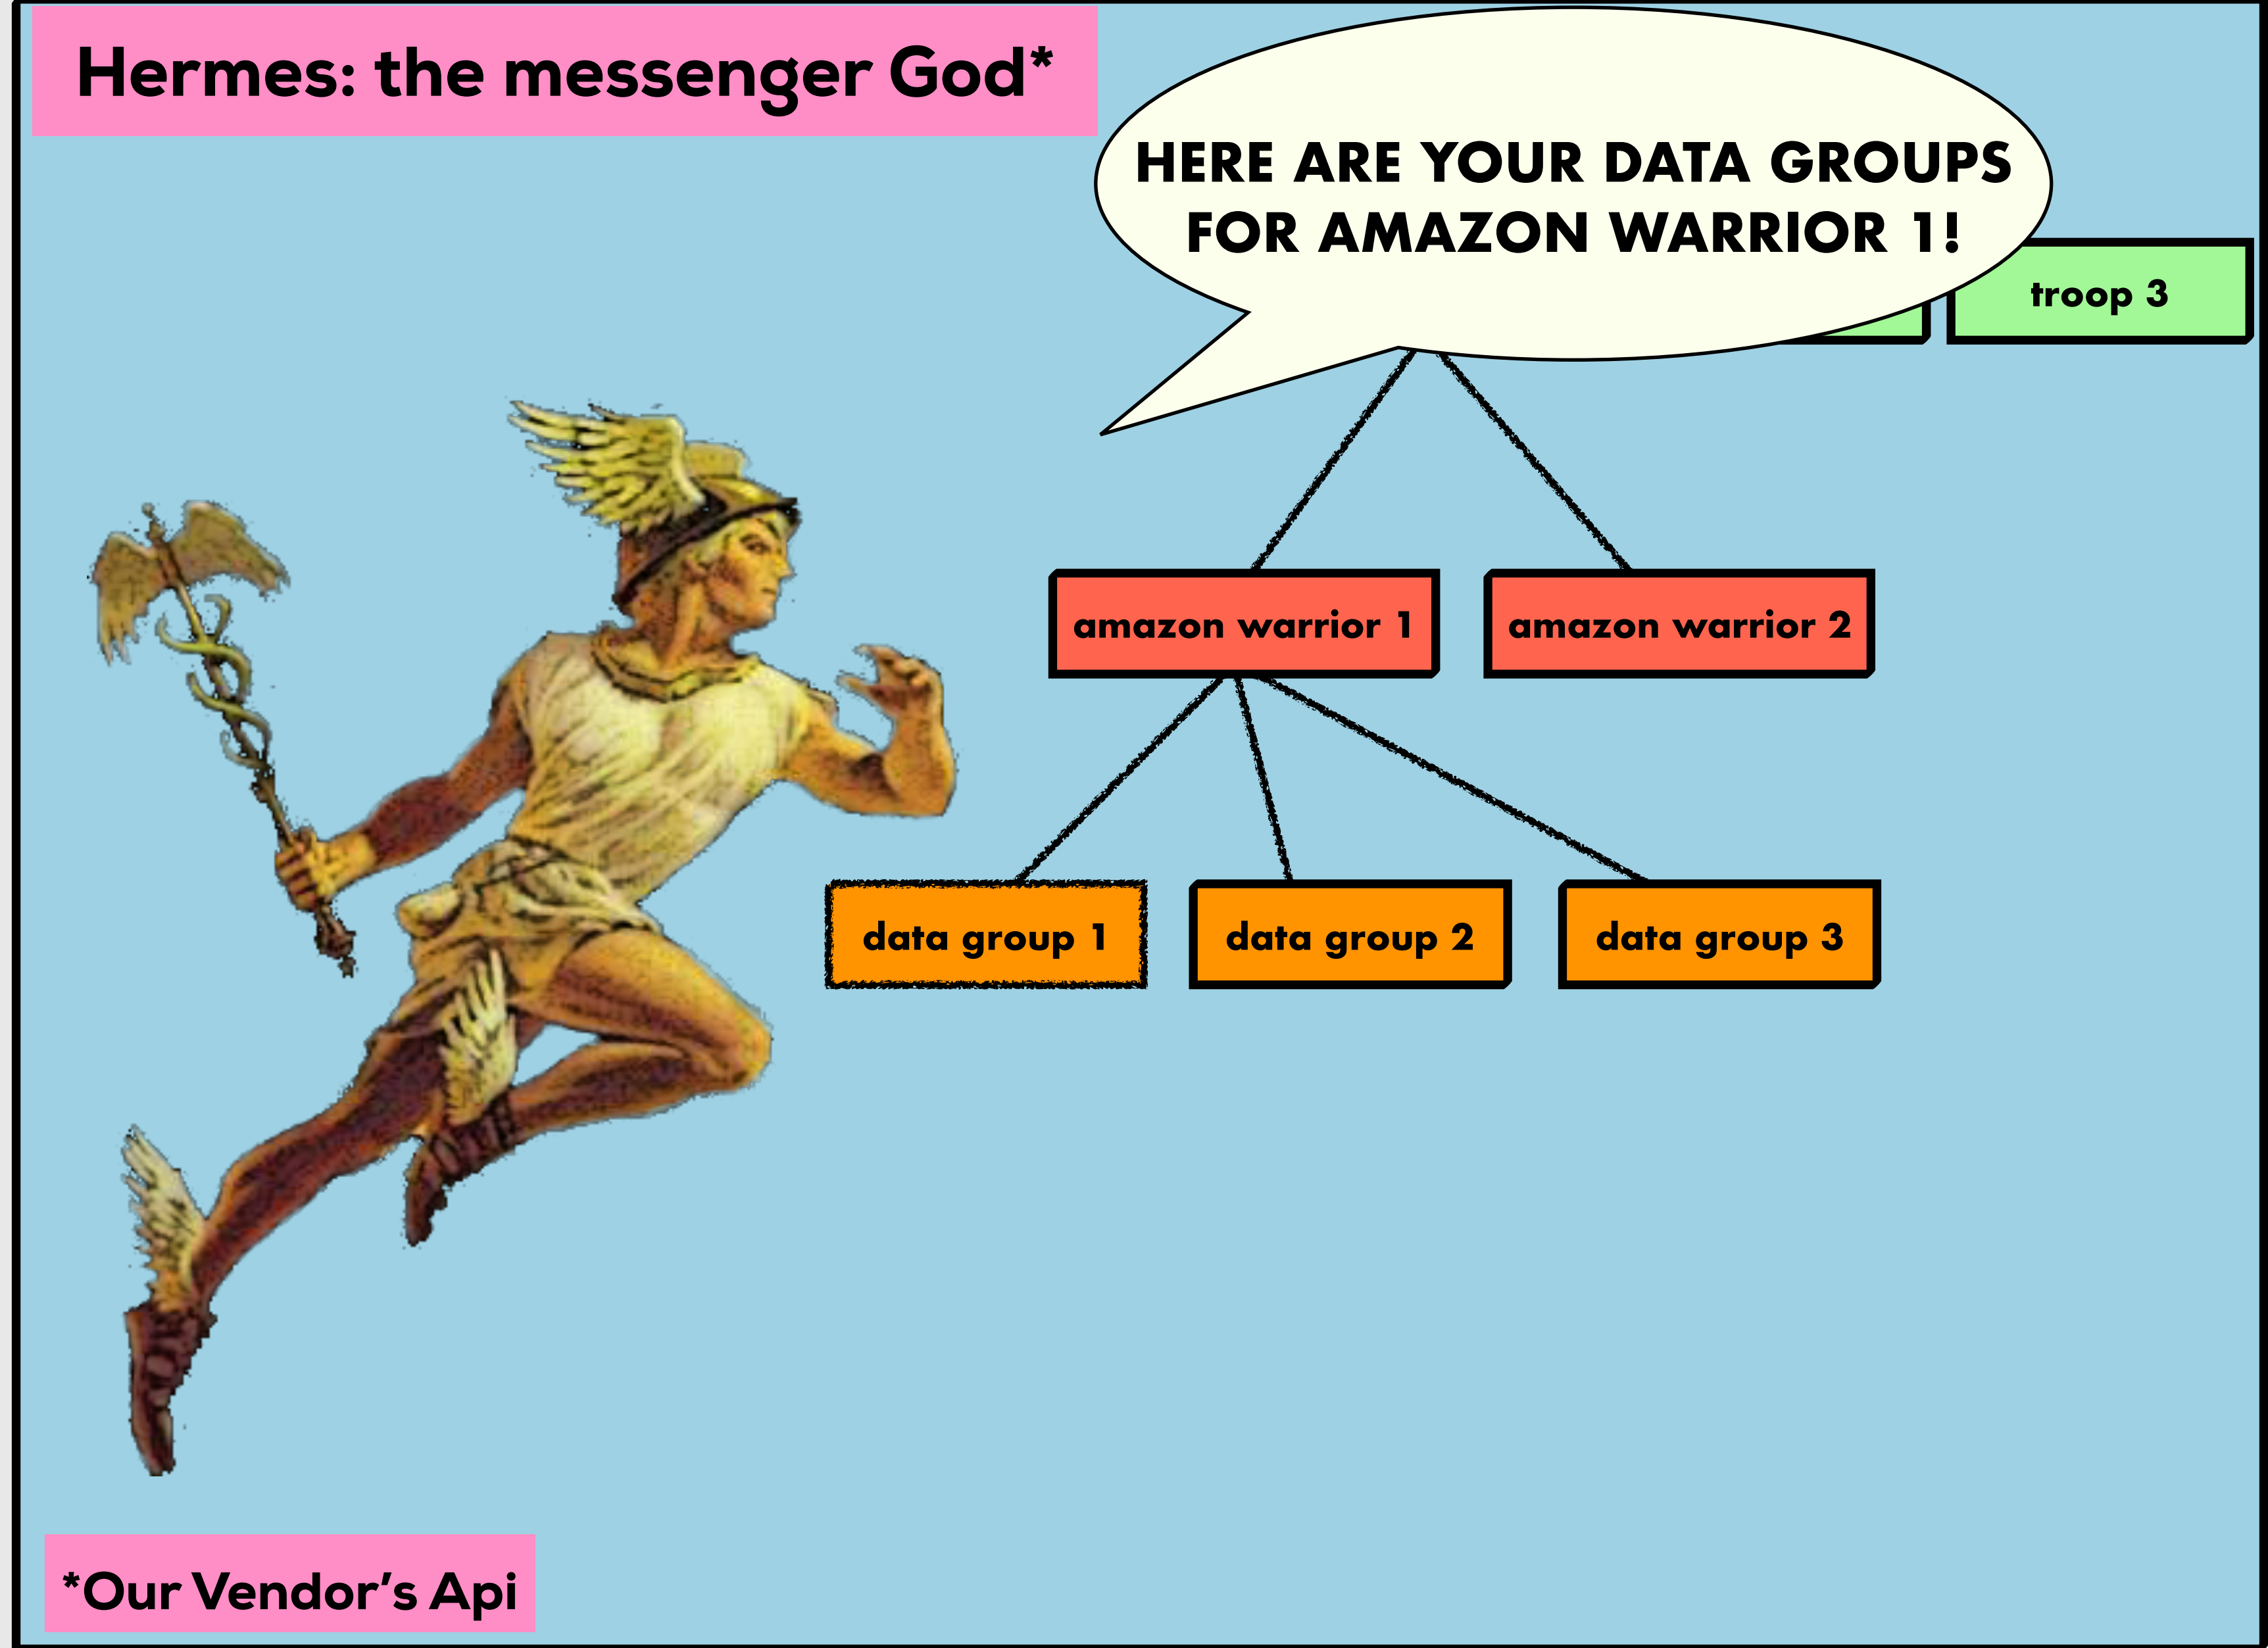


**The Oracle\***



**\*one of our monitoring vendors**

**Hermes: the messenger God\***



**\*Our Vendor's Api**




### The Oracle\*



\*one of our monitoring vendors

### Hermes: the messenger God\*



\*Our Vendor's Api

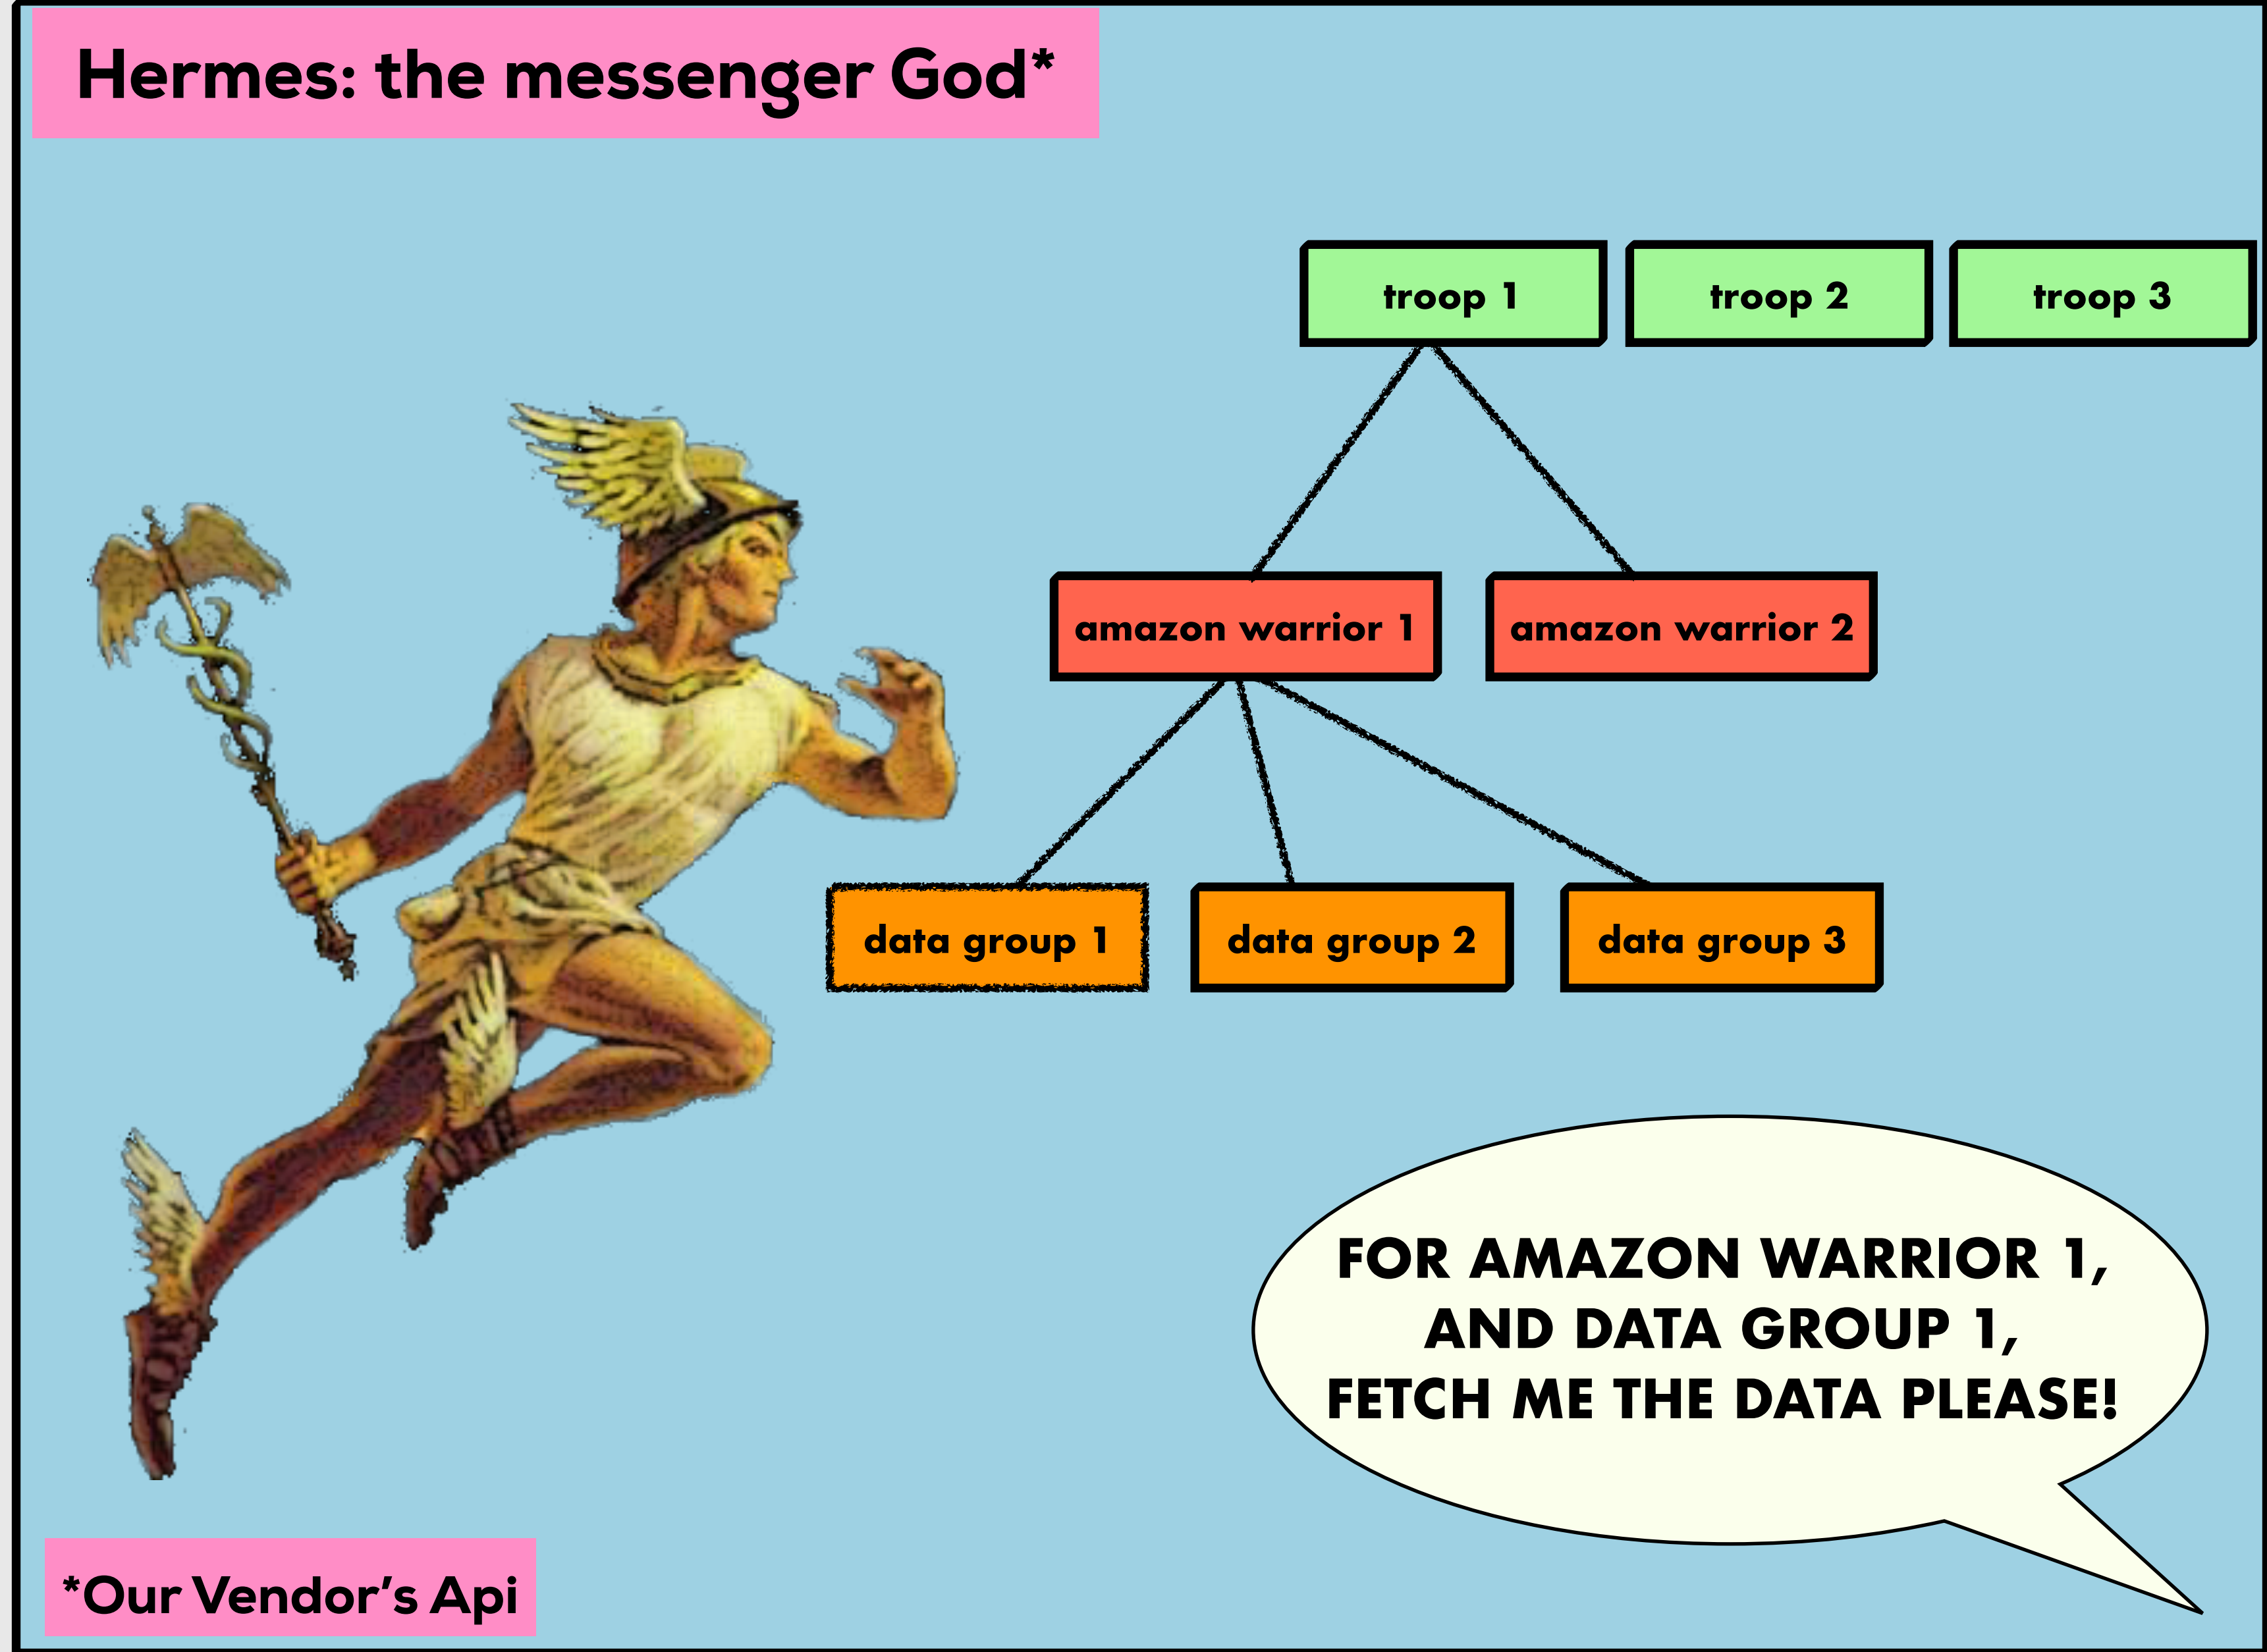


### The Oracle\*



\*one of our monitoring vendors

### Hermes: the messenger God\*



\*Our Vendor's Api

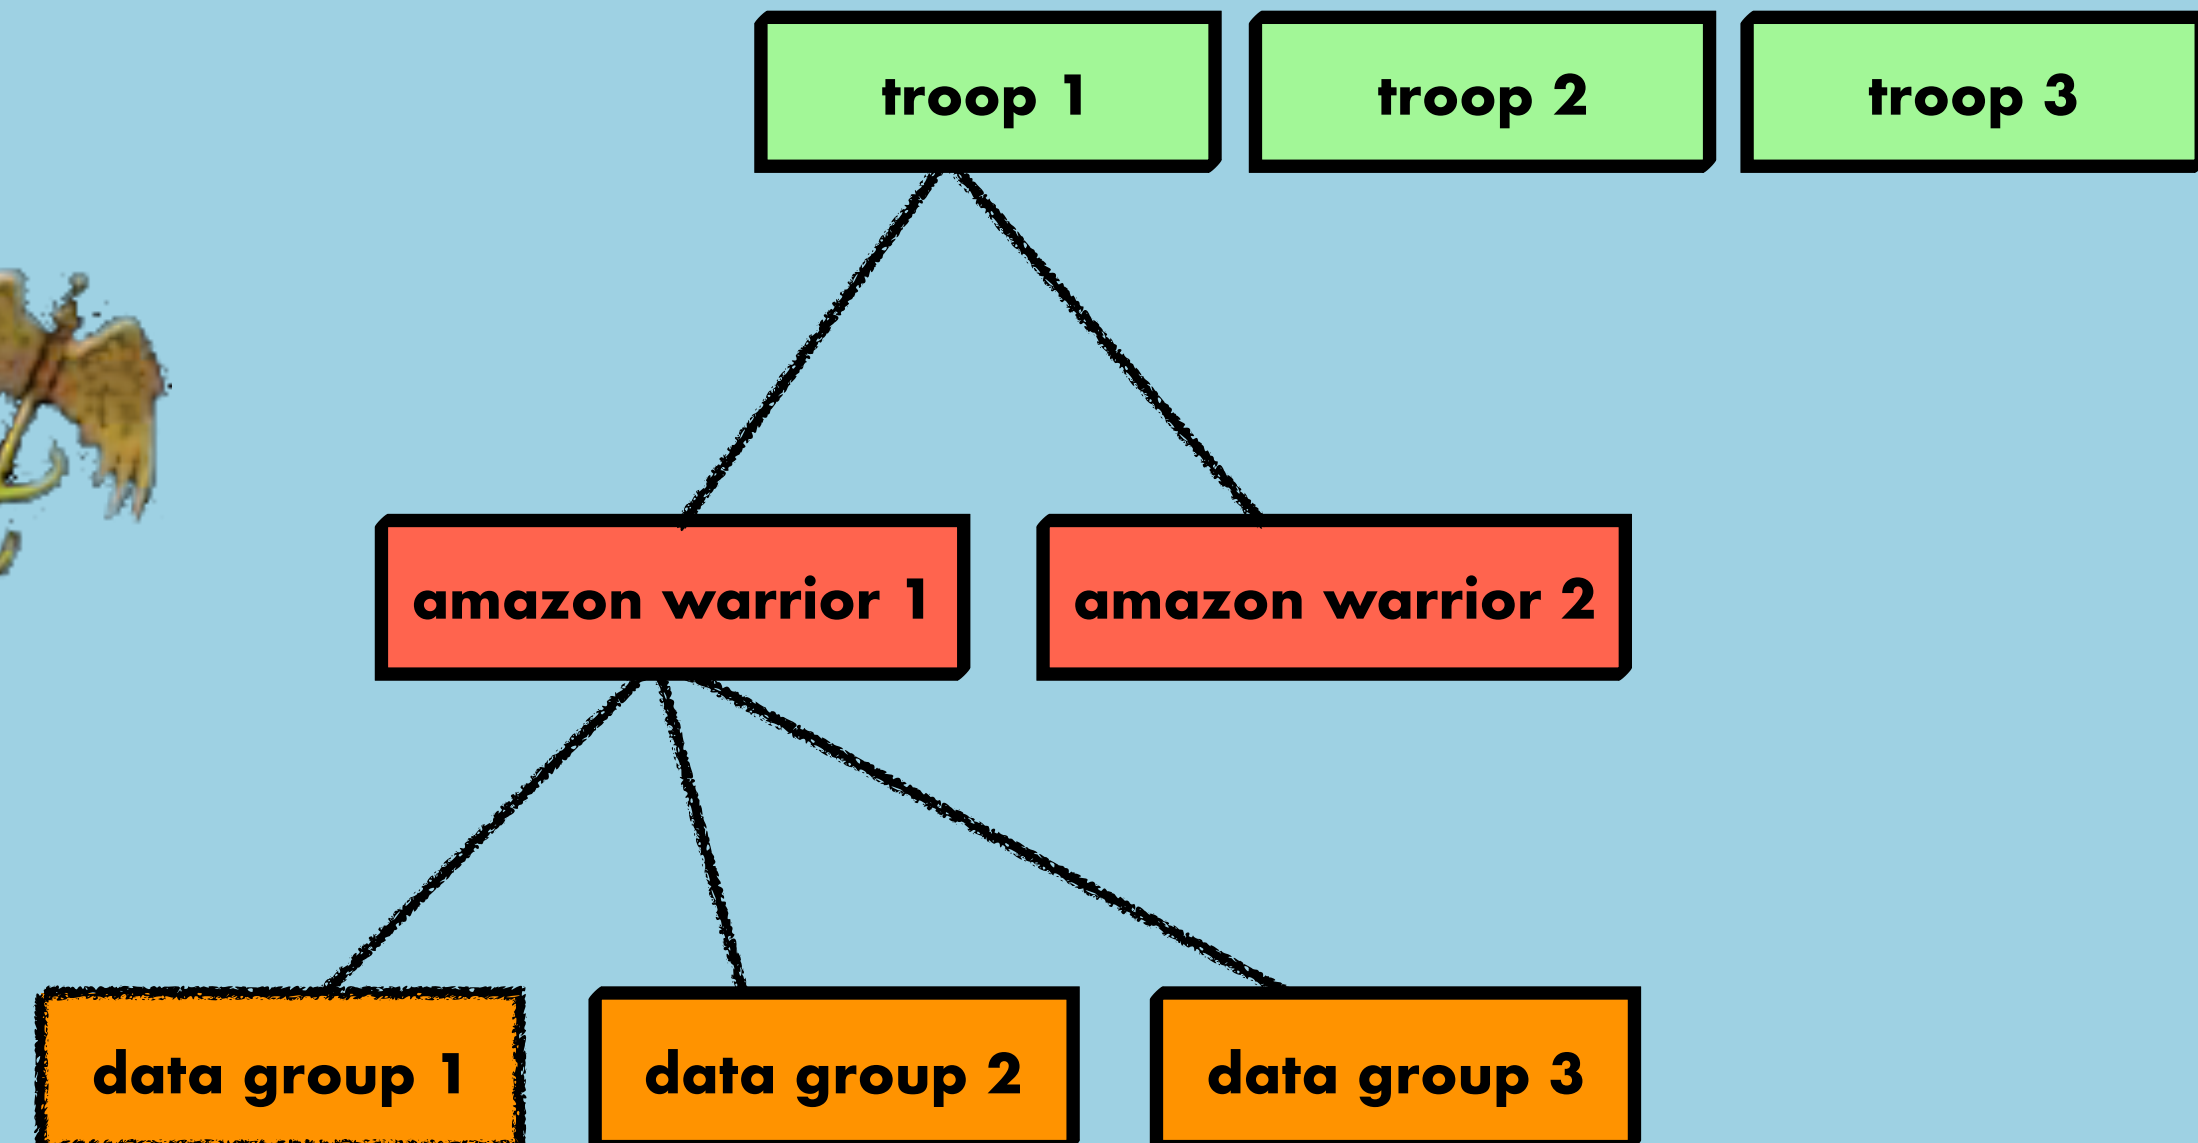


### The Oracle\*



\*one of our monitoring vendors

### Hermes: the messenger God\*



\*Our Vendor's Api

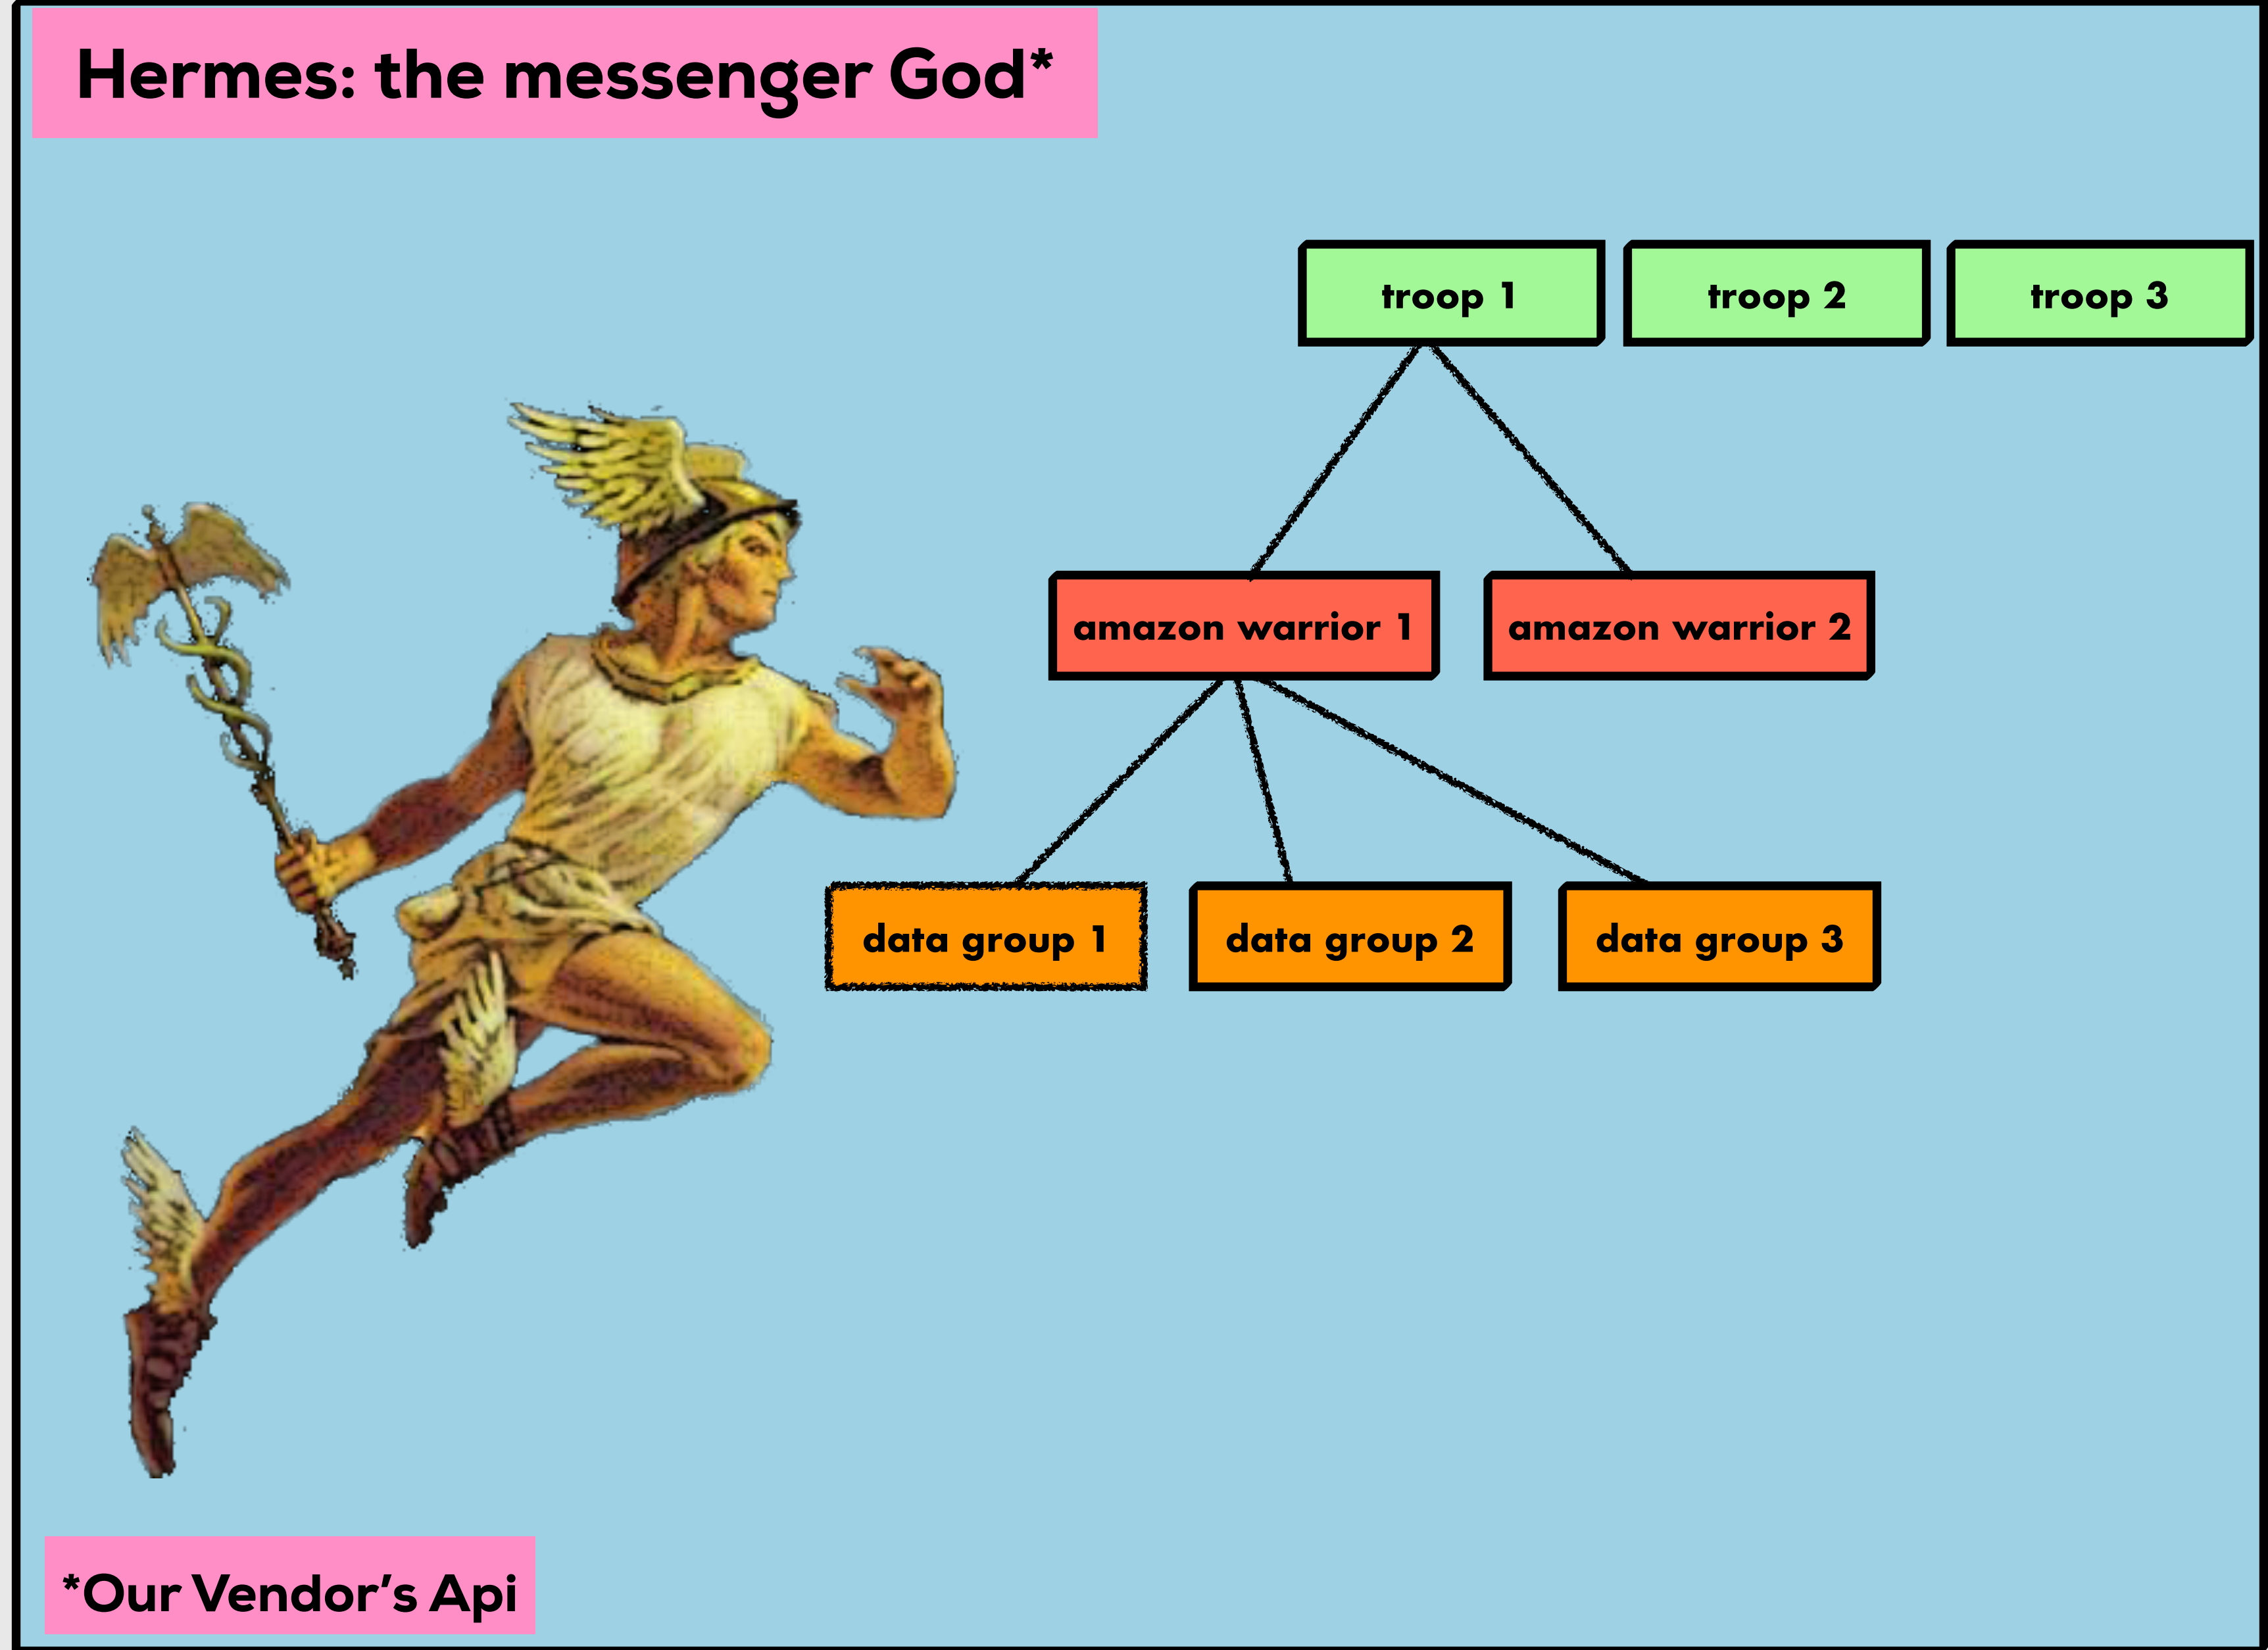


**The Oracle\***



**\*one of our monitoring vendors**

**Hermes: the messenger God\***



**\*Our Vendor's Api**

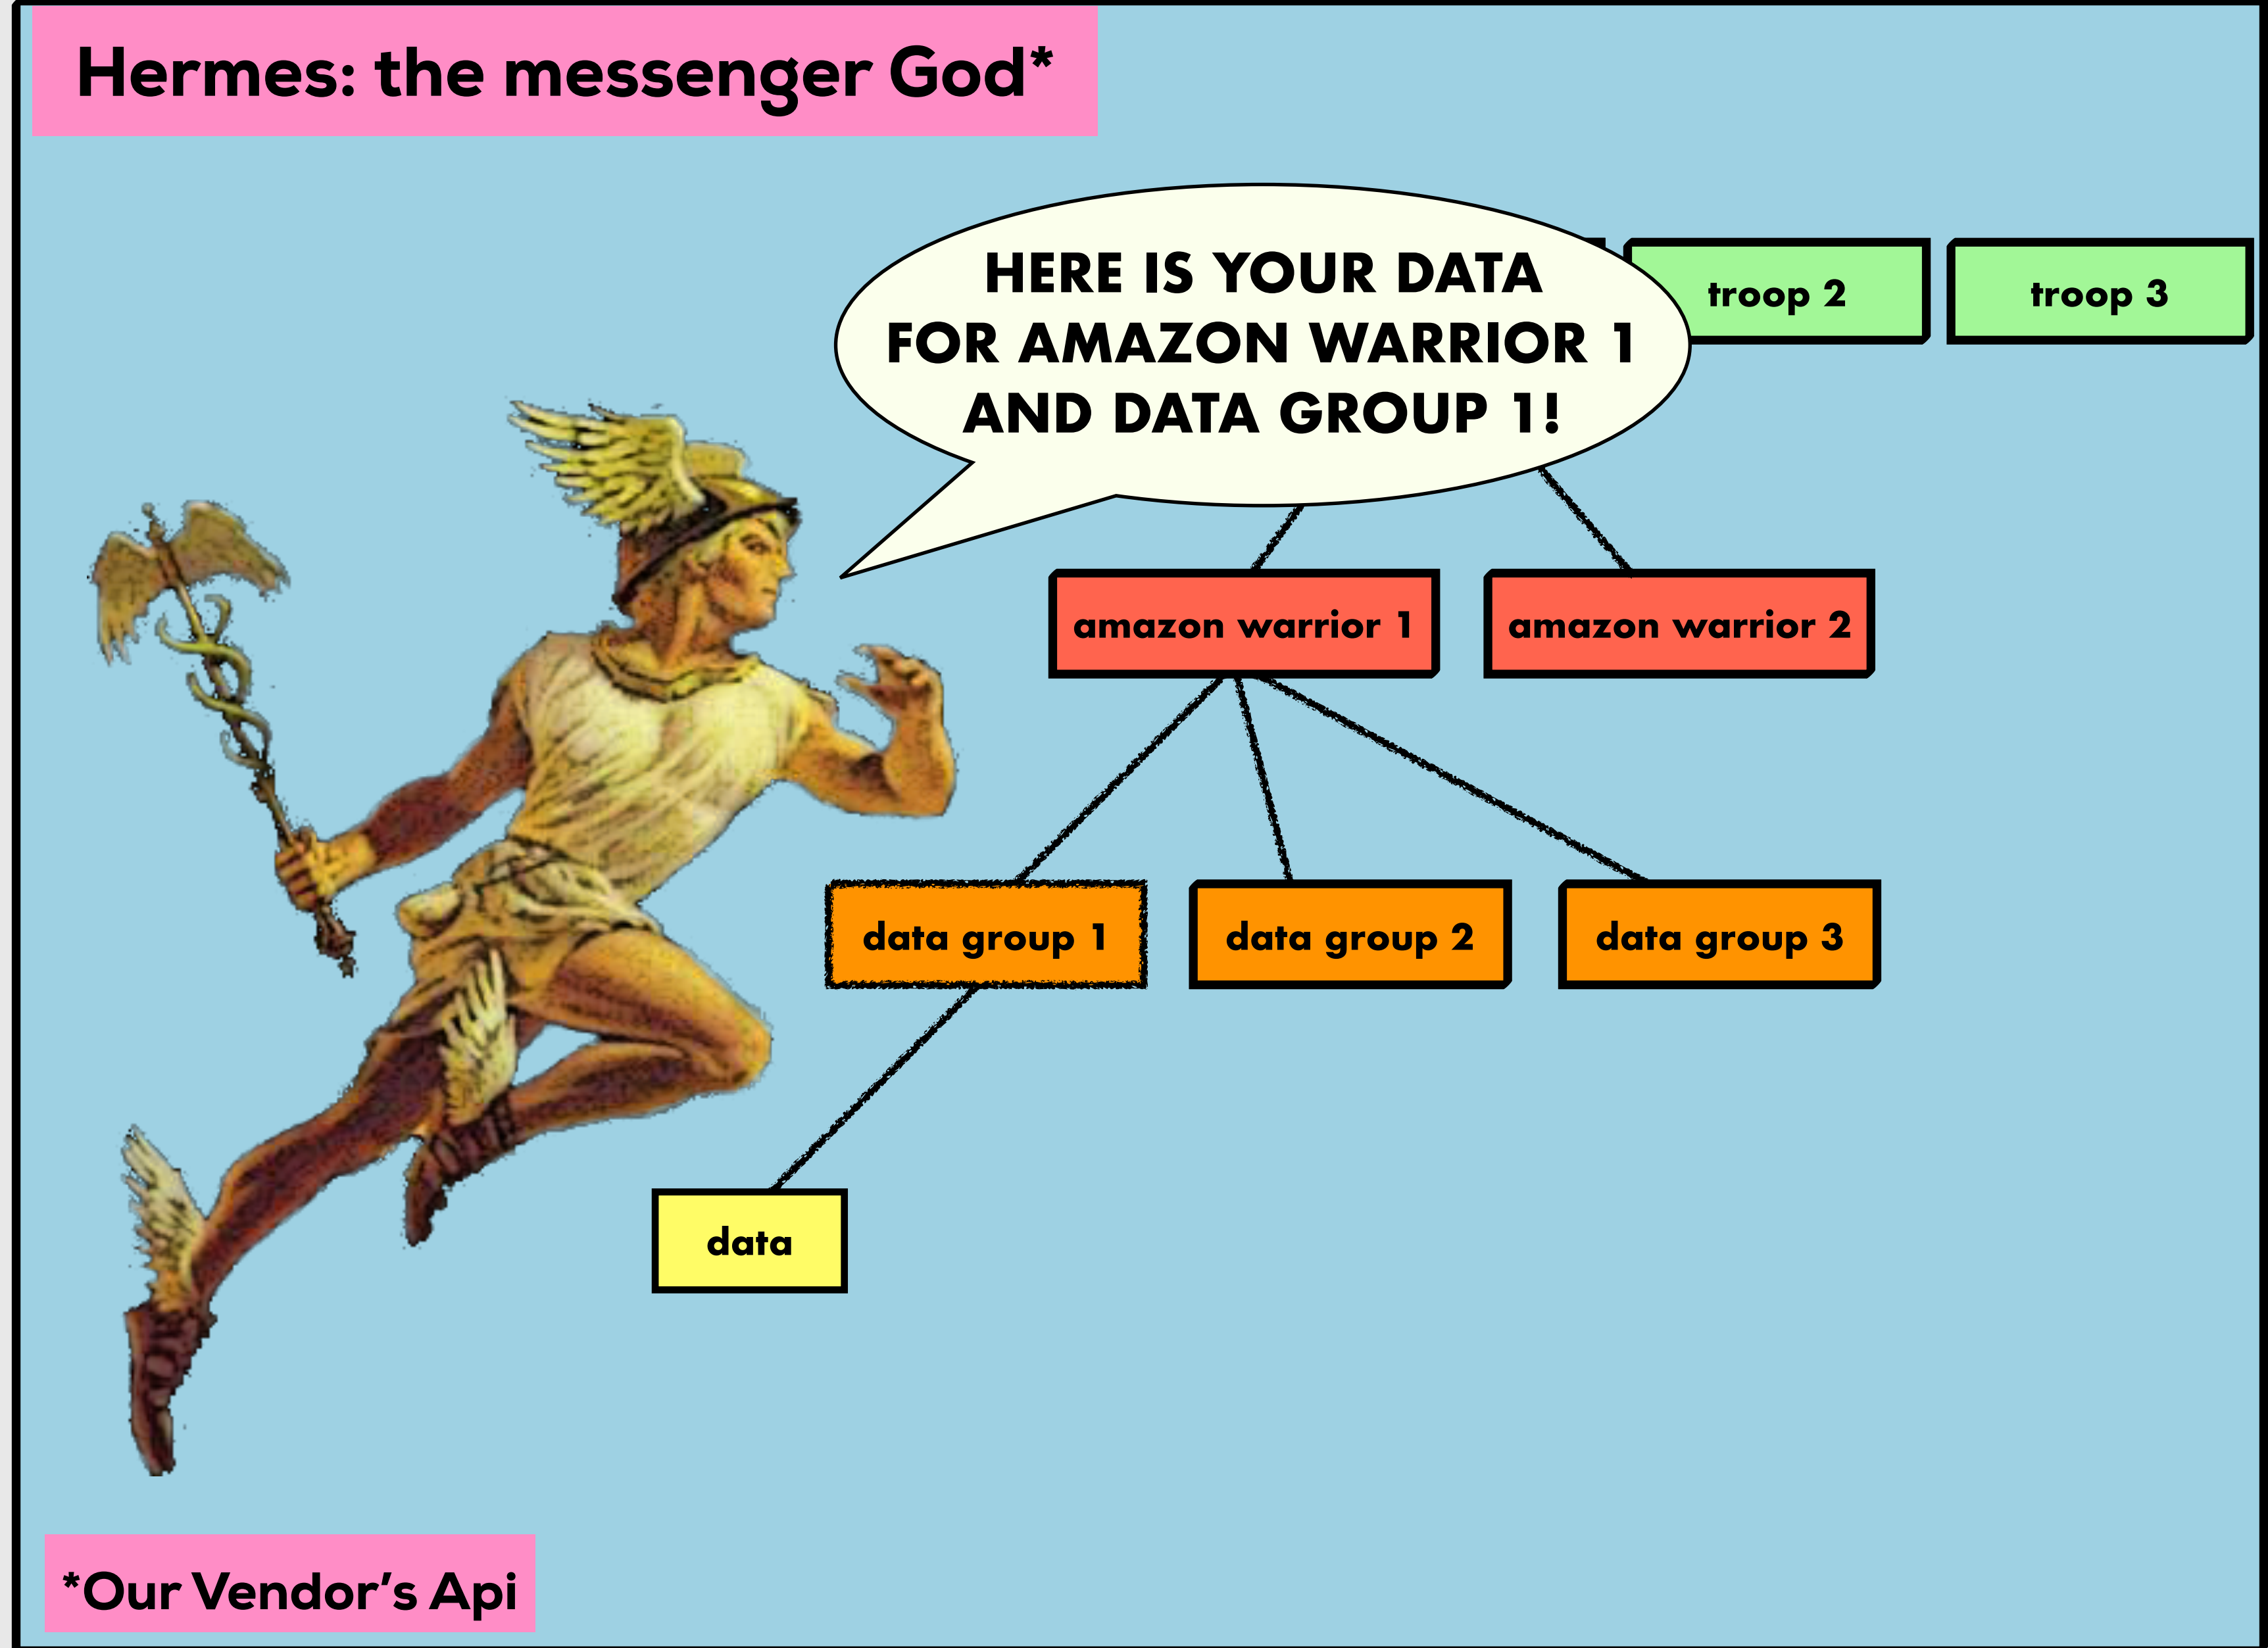


**The Oracle\***



**\*one of our monitoring vendors**

**Hermes: the messenger God\***



**\*Our Vendor's Api**

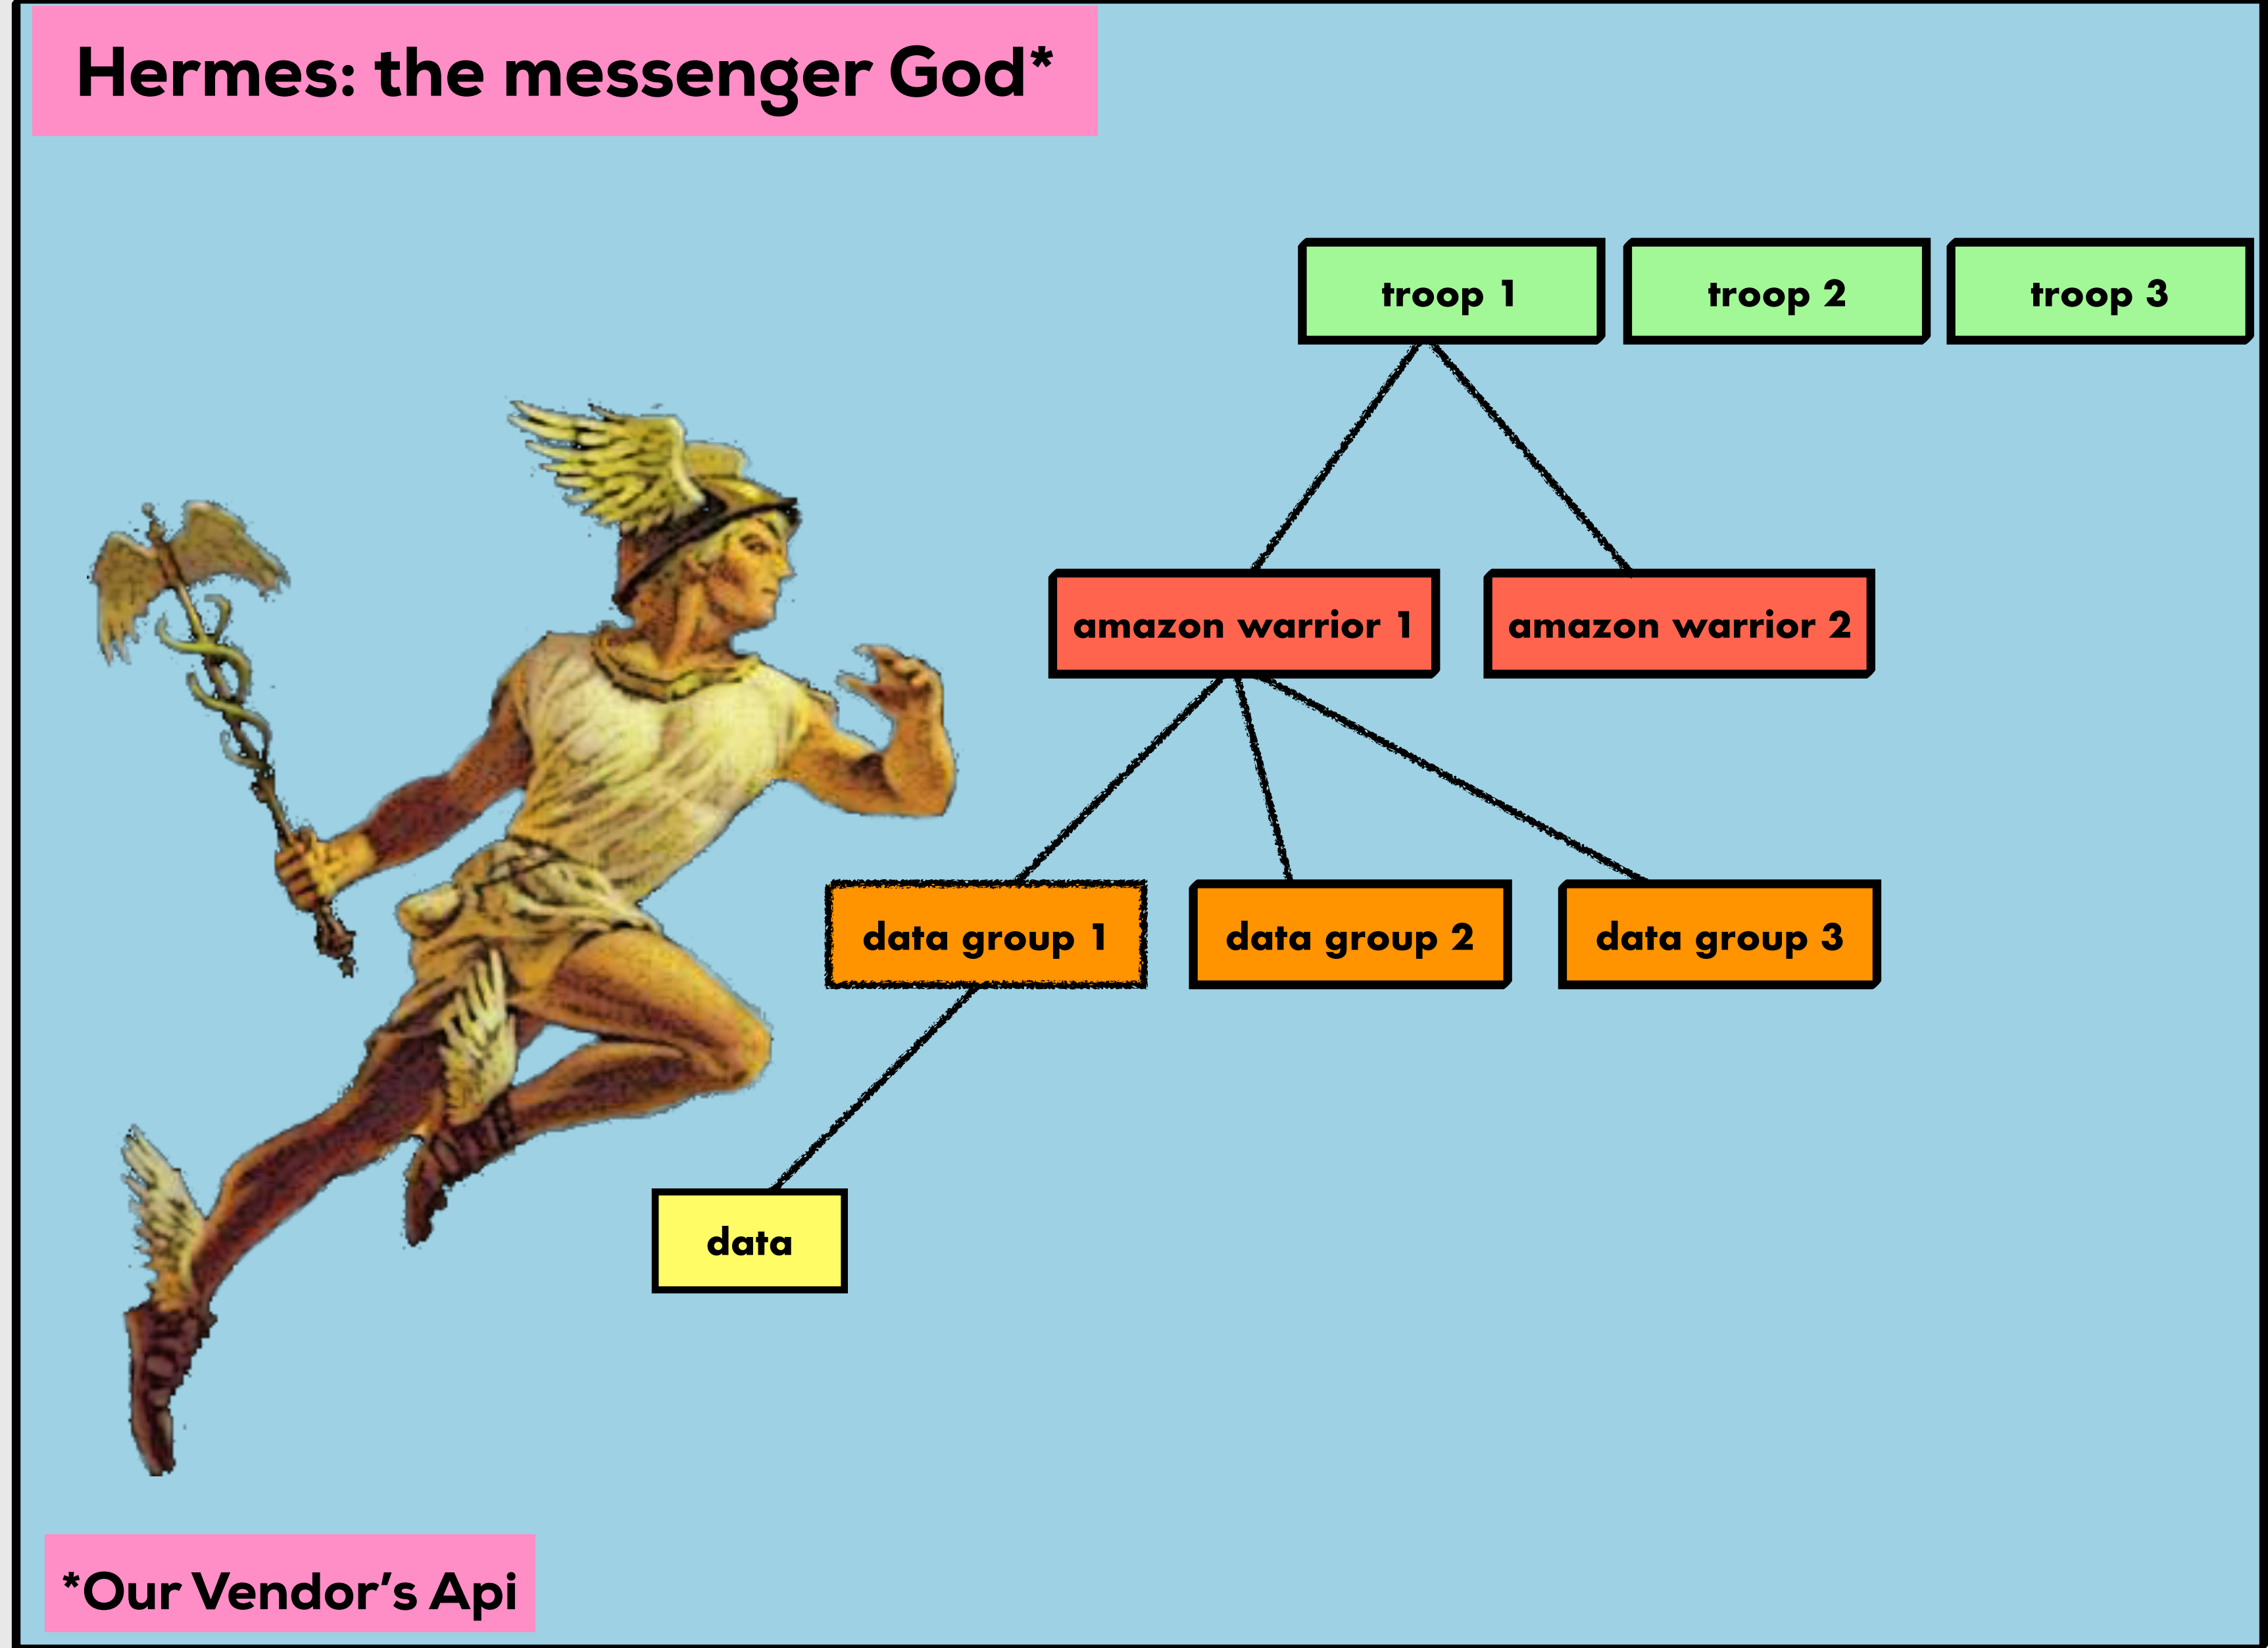


**The Oracle\***



**\*one of our monitoring vendors**

**Hermes: the messenger God\***




**\*Our Vendor's Api**



**35,000-65,000  
API CALLS!**




**HOW SHOULD I WRITE THE  
CODE TO COMMUNICATE WITH  
HERMES?**





# Building a Data Pipeline



**Attempt 0: Python Multiprocessing on an EC2 Instance**

**Attempt 0.5: Python Multiprocessing on a Spot Instance**

**Attempt 1: Lambdas Orchestrated by Step Functions**

**Attempt 2: Using Event Triggers with S3**

**Attempt 3: Rate limiting with Queues and Cloudwatch triggered Lambdas**





## Initial Solution:

### Python Multiprocessing on an EC2 instance - continuously running



**Easy for me to implement**



**not cheap**

**horizontal scaling seemed overkill**

**on-going maintenance  
and management**



## Initial Solution:

### Python Multiprocessing on an EC2 instance - intermittently running




**Easy for me to implement**  
**cheaper**



**still not cheap enough**  
**horizontal scaling seemed overkill**  
**on-going maintenance**  
**and management**



# Building a Data Pipeline



**Attempt 0: Python Multiprocessing on an EC2 Instance**

**Attempt 0.5: Python Multiprocessing on a Spot Instance**

**Attempt 1: Lambdas Orchestrated by Step Functions**


**Attempt 2: Using Event Triggers with S3**

**Attempt 3: Rate limiting with Queues and Cloudwatch triggered Lambdas**





# Building a Data Pipeline



**Attempt 0: Python Multiprocessing on an EC2 Instance**

**Attempt 0.5: Python Multiprocessing on a Spot Instance**

**Attempt 1: Lambdas Orchestrated by Step Functions**

**Attempt 2: Using Event Triggers with S3**

**Attempt 3: Rate limiting with Queues and Cloudwatch triggered Lambdas**





**Revision:**

**a Spot instance**

**Python Multiprocessing on an ~~EC2~~ instance**



**significantly cheaper**



**You need to configure  
how to restart where  
the previous process left off**



@BigDana

Should I try Serverless?





Owning your own server : Owning your own car

@BigDana





Owning your own car: Owning your own server

@BigDana



The Cloud: Leasing a car



Owning your own server : Owning your own car



The Cloud : Leasing a car



Spot Instances: Renting a car



Owning your own server: Owning your own car



@BigDana

Lambdas: car2go




Spot Instances: Renting a car



The Cloud: Leasing a car



# Building a Data Pipeline



**Attempt 0: Python Multiprocessing on an EC2 Instance**

**Attempt 0.5: Python Multiprocessing on a Spot Instance**

**Attempt 1: Lambdas Orchestrated by Step Functions**

**Attempt 2: Using Event Triggers with S3**

**Attempt 3: Rate limiting with Queues and Cloudwatch triggered Lambdas**





# Building a Data Pipeline


**Attempt 0: Python Multiprocessing on an EC2 Instance**

**Attempt 0.5: Python Multiprocessing on a Spot Instance**

**Attempt 1: Lambdas Orchestrated by Step Functions**

**Attempt 2: Using Event Triggers with S3**

**Attempt 3: Rate limiting with Queues and Cloudwatch triggered Lambdas**



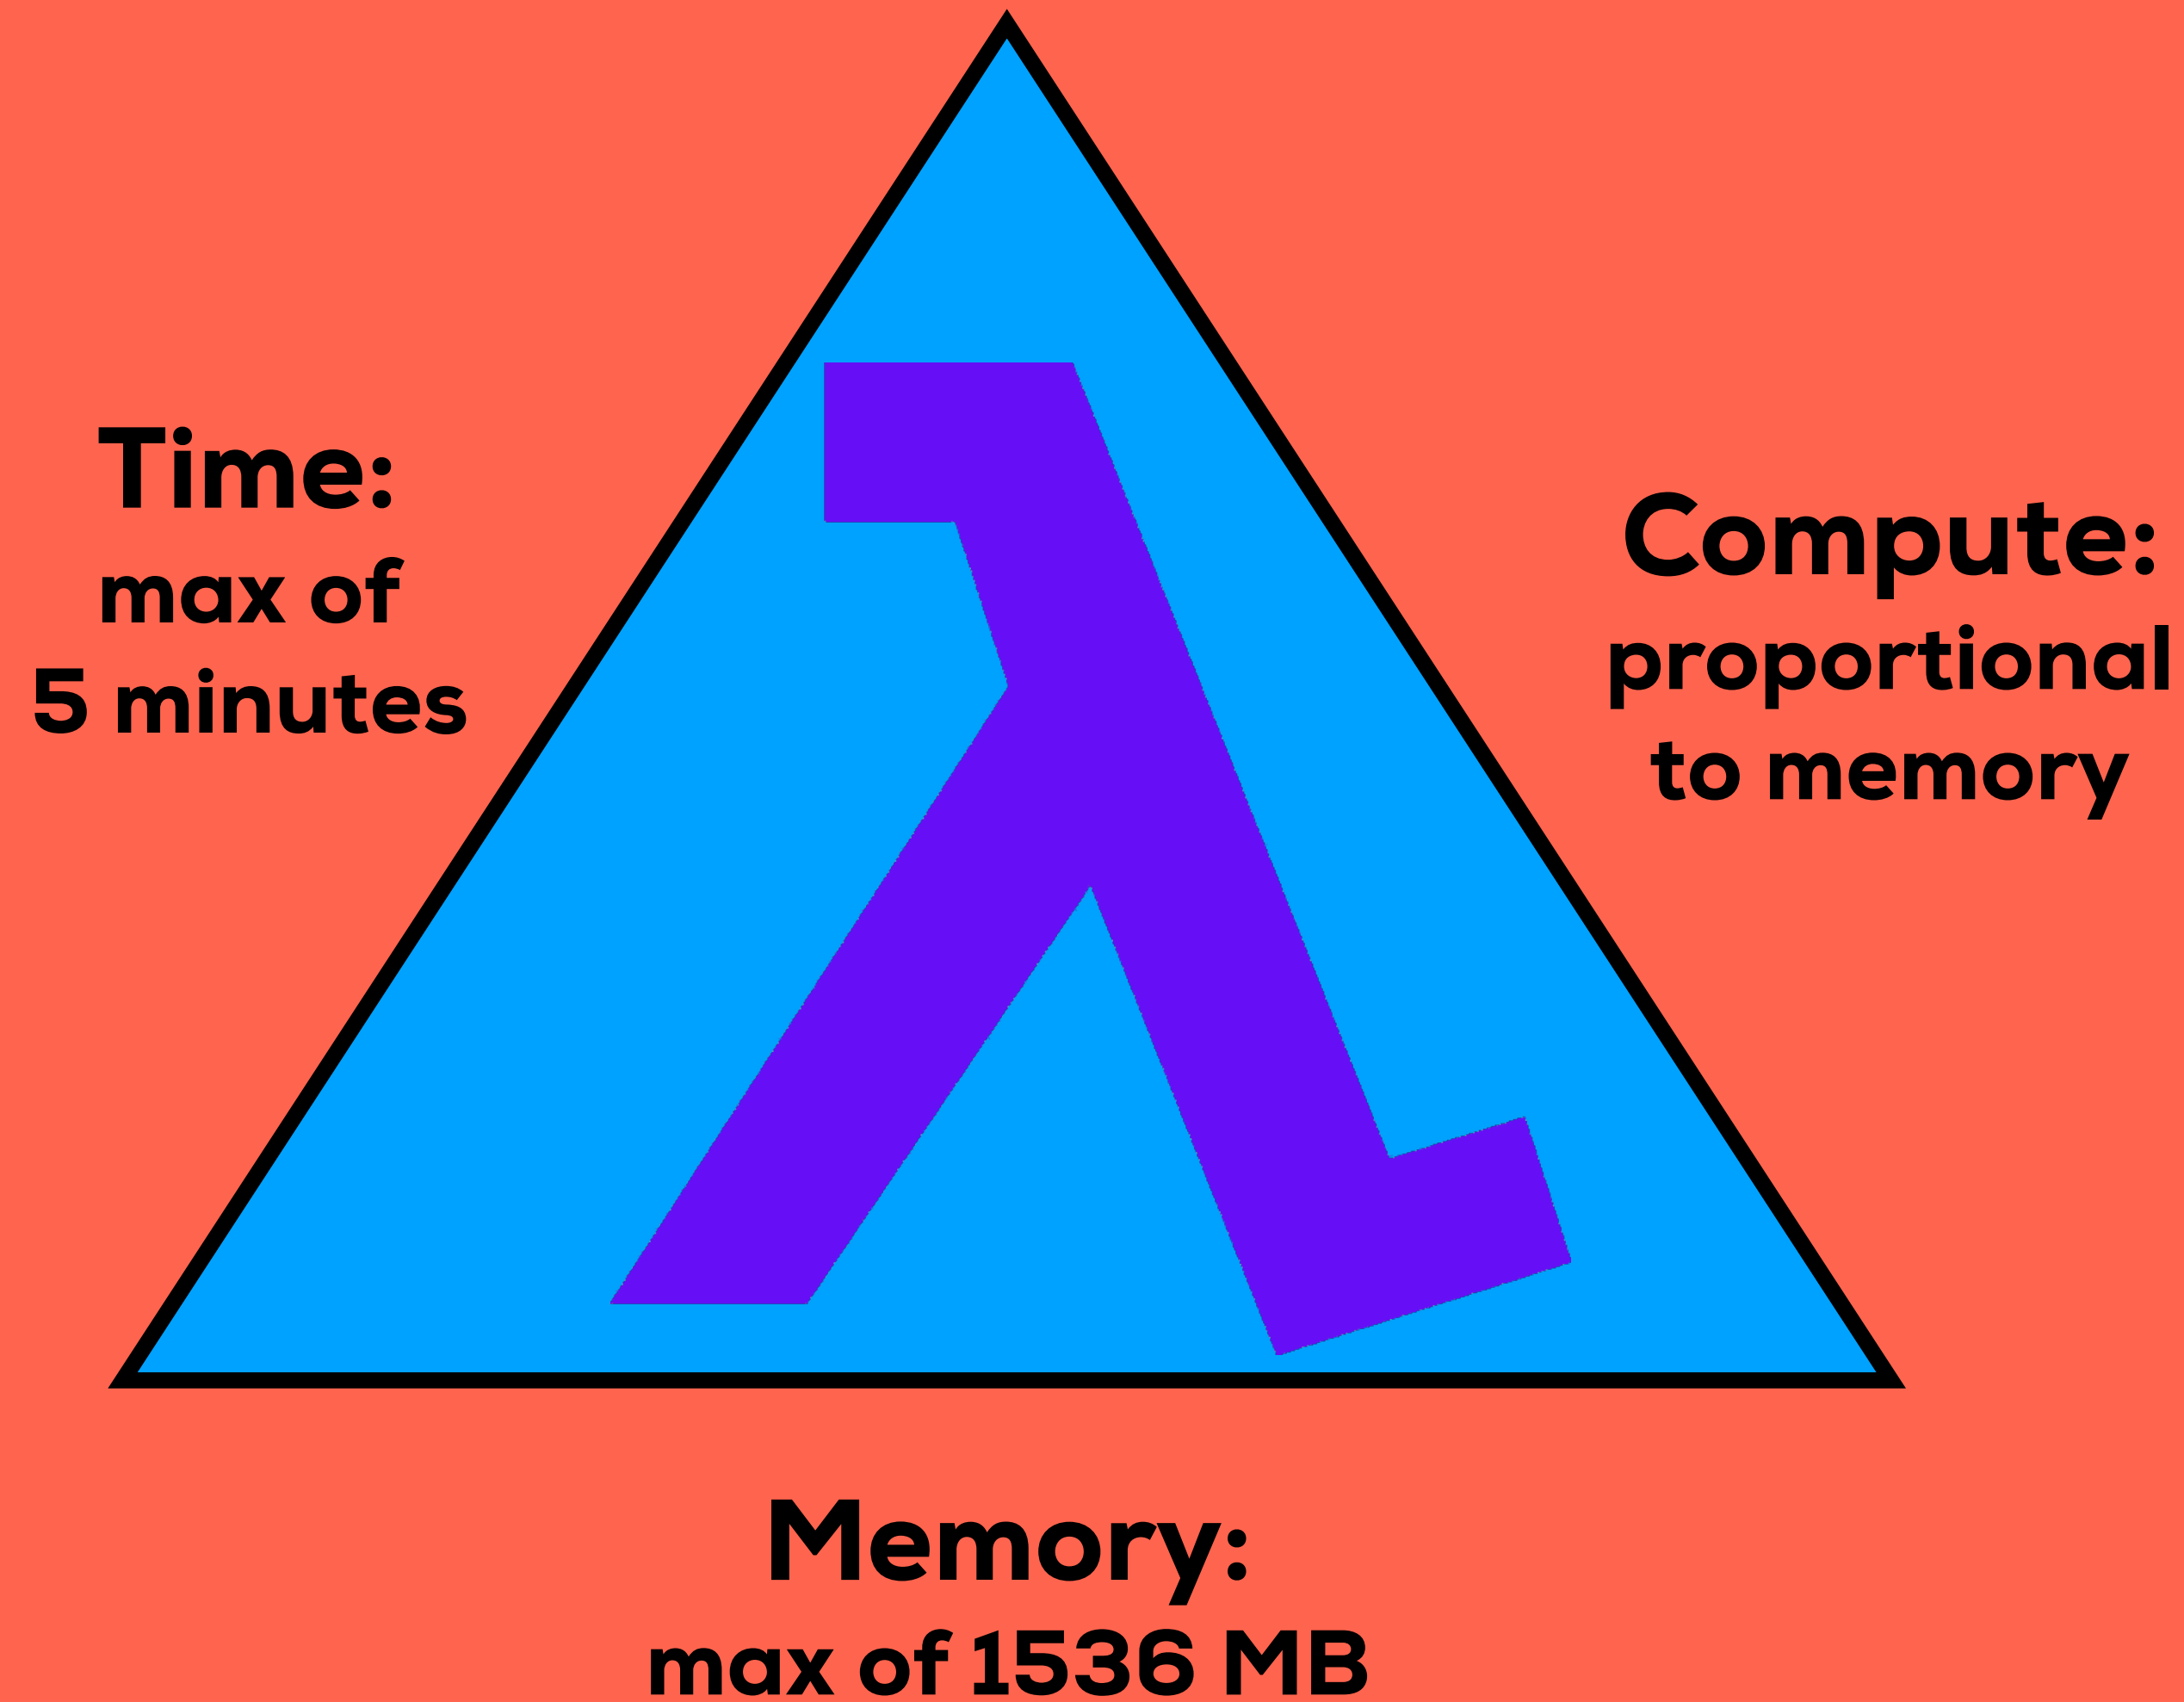


# Lambda: a Function as a Service






# Lambda: a Function as a Service





# Lambda: Function as a Service

Compute



**Dependency  
Zip File size**

250 MB uncompressed  
code/dependencies

Memory

Time



```
21 build:
22     docker build --no-cache -t flowers .
23
24 dev:
25     docker run -it -v `pwd`: /code flowers bash
26
27 libs:
28     conda install --yes --file /code/serverless/get_data_lambda/requirements.txt -v
29     mkdir -p serverless/get_data_lambda/libs
30     cp -a /opt/conda/pkgshare /code/serverless/get_data_lambda/libs
31
32 zip:
33     cp /code/serverless/get_data_lambda/get_data.py /code/serverless/get_data_lambda/libs/
34     rm -rf /code/tmpdir
35     mkdir -p /code/tmpdir
36     cp -r /code/serverless/get_data_lambda/libs/* /code/tmpdir
37     cd /code/tmpdir
38     zip deployment.zip *
```



```
21 build:
22     docker build --no-cache -t flowers .
23
24 dev:
25     docker run -it -v `pwd`:/code flowers bash
26
27 libs:
28     conda install --yes --file /code/serverless/get_data_lambda/requirements.txt -v
29     mkdir -p serverless/get_data_lambda/libs
30     cp -a /opt/conda/pkgs/. /code/serverless/get_data_lambda/libs
31
32 zip:
33     cp /code/serverless/get_data_lambda/get_data.py /code/serverless/get_data_lambda/libs/
34     rm -rf /code/tmpdir
35     mkdir -p /code/tmpdir
36     cp -r /code/serverless/get_data_lambda/libs/* /code/tmpdir
37     cd /code/tmpdir
38     zip deployment.zip *
```



```
21 build:
22     docker build --no-cache -t flowers .
23
24 dev:
25     docker run -it -v `pwd`:~/code flowers bash
26
27 libs:
28     conda install --yes --file /code/serverless/get_data_lambda/requirements.txt -v
29     mkdir -p serverless/get_data_lambda/libs
30     cp -a /opt/conda/pkgs/. /code/serverless/get_data_lambda/libs
31
32 zip:
33     cp /code/serverless/get_data_lambda/get_data.py /code/serverless/get_data_lambda/libs/
34     rm -rf /code/tmpdir
35     mkdir -p /code/tmpdir
36     cp -r /code/serverless/get_data_lambda/libs/* /code/tmpdir
37     cd /code/tmpdir
38     zip deployment.zip *
```



```
21 build:
22     docker build --no-cache -t flowers .
23
24 dev:
25     docker run -it -v `pwd`:/code flowers bash
26
27 libs:
28     conda install --yes --file /code/serverless/get_data_lambda/requirements.txt -v
29     mkdir -p serverless/get_data_lambda/libs
30     cp -a /opt/conda/pkgs/. /code/serverless/get_data_lambda/libs
31
32 zip:
33     cp /code/serverless/get_data_lambda/get_data.py /code/serverless/get_data_lambda/libs/
34     rm -rf /code/tmpdir
35     mkdir -p /code/tmpdir
36     cp -r /code/serverless/get_data_lambda/libs/* /code/tmpdir
37     cd /code/tmpdir
38     zip deployment.zip *
```



```
21 build:
22     docker build --no-cache -t flowers .
23
24 dev:
25     docker run -it -v `pwd`:/code flowers bash
26
27 libs:
28     conda install --yes --file /code/serverless/get_data_lambda/requirements.txt -v
29     mkdir -p serverless/get_data_lambda/libs
30     cp -a /opt/conda/pkgshare /code/serverless/get_data_lambda/libs
31
32 zip:
33     cp /code/serverless/get_data_lambda/get_data.py /code/serverless/get_data_lambda/libs/
34     rm -rf /code/tmpdir
35     mkdir -p /code/tmpdir
36     cp -r /code/serverless/get_data_lambda/libs/* /code/tmpdir
37     cd /code/tmpdir
38     zip deployment.zip *
```



▶	cache	Today, 4:53 PM	--	Folder
▶	cfffi-1.9.1-py27_0	Today, 4:53 PM	--	Folder
▶	conda-4.3.14-py27_0	Today, 4:53 PM	--	Folder
▶	conda-4.3.30-py27h6ae6dc7_0	Today, 4:53 PM	--	Folder
	conda-4.3.30-py27h6ae6dc7_0.tar.bz2	Today, 4:53 PM	519 KB	bzip2...archive
▶	conda-env-2.6.0-0	Today, 4:53 PM	--	Folder
▶	cryptography-1.7.1-py27_0	Today, 4:53 PM	--	Folder
▶	enum34-1.1.6-py27_0	Today, 4:53 PM	--	Folder
	get_data.py	Today, 4:53 PM	424 bytes	Python Script
▶	idna-2.2-py27_0	Today, 4:53 PM	--	Folder
▶	ipaddress-1.0.18-py27_0	Today, 4:53 PM	--	Folder
▶	libffi-3.2.1-1	Today, 4:53 PM	--	Folder
▶	openssl-1.0.2k-1	Today, 4:53 PM	--	Folder
▶	pip-9.0.1-py27_1	Today, 4:53 PM	--	Folder
▶	pyasn1-0.1.9-py27_0	Today, 4:54 PM	--	Folder
▶	pycosat-0.6.1-py27_1	Today, 4:54 PM	--	Folder
▶	pycparser-2.17-py27_0	Today, 4:54 PM	--	Folder
▶	pyopenssl-16.2.0-py27_0	Today, 4:54 PM	--	Folder
▶	python-2.7.13-0	Today, 4:54 PM	--	Folder
▶	readline-6.2-2	Today, 4:54 PM	--	Folder
▶	requests-2.12.4-py27_0	Today, 4:54 PM	--	Folder
▶	requests-2.14.2-py27_0	Today, 4:54 PM	--	Folder
	requests-2.14.2-py27_0.tar.bz2	Today, 4:54 PM	717 KB	bzip2...archive
▶	ruamel_yaml-0.11.14-py27_1	Today, 4:54 PM	--	Folder
▶	setuptools-27.2.0-py27_0	Today, 4:54 PM	--	Folder
▶	six-1.10.0-py27_0	Today, 4:54 PM	--	Folder
▶	sqlite-3.13.0-0	Today, 4:54 PM	--	Folder
▶	tk-8.5.18-0	Today, 4:54 PM	--	Folder
	urls	Today, 4:54 PM	3 KB	TextEd...ument
	urls.txt	Today, 4:54 PM	2 KB	Plain Text
▶	wheel-0.29.0-py27_0	Today, 4:54 PM	--	Folder
▶	yaml-0.1.6-0	Today, 4:54 PM	--	Folder
▶	zlib-1.2.8-3	Today, 4:54 PM	--	Folder




▶	cache	Today, 4:53 PM	--	Folder
▶	cfffi-1.9.1-py27_0	Today, 4:53 PM	--	Folder
▶	conda-4.3.14-py27_0	Today, 4:53 PM	--	Folder
▶	conda-4.3.30-py27h6ae6dc7_0	Today, 4:53 PM	--	Folder
	conda-4.3.30-py27h6ae6dc7_0.tar.bz2	Today, 4:53 PM	519 KB	bzip2...archive
▶	conda-env-2.6.0-0	Today, 4:53 PM	--	Folder
▶	cryptography-1.7.1-py27_0	Today, 4:53 PM	--	Folder
▶	enum34-1.1.6-py27_0	Today, 4:53 PM	--	Folder
	get_data.py	Today, 4:53 PM	424 bytes	Python Script
▶	idna-2.2-py27_0	Today, 4:53 PM	--	Folder
▶	ipaddress-1.0.18-py27_0	Today, 4:53 PM	--	Folder
▶	libffi-3.2.1-1	Today, 4:53 PM	--	Folder
▶	openssl-1.0.2k-1	Today, 4:53 PM	--	Folder
▶	pip-9.0.1-py27_1	Today, 4:53 PM	--	Folder
▶	pyasn1-0.1.9-py27_0	Today, 4:54 PM	--	Folder
▶	pycosat-0.6.1-py27_1	Today, 4:54 PM	--	Folder
▶	pycparser-2.17-py27_0	Today, 4:54 PM	--	Folder
▶	pyopenssl-16.2.0-py27_0	Today, 4:54 PM	--	Folder
▶	python-2.7.13-0	Today, 4:54 PM	--	Folder
▶	readline-6.2-2	Today, 4:54 PM	--	Folder
▶	requests-2.12.4-py27_0	Today, 4:54 PM	--	Folder
▶	requests-2.14.2-py27_0	Today, 4:54 PM	--	Folder
	requests-2.14.2-py27_0.tar.bz2	Today, 4:54 PM	717 KB	bzip2...archive
▶	ruamel_yaml-0.11.14-py27_1	Today, 4:54 PM	--	Folder
▶	setuptools-27.2.0-py27_0	Today, 4:54 PM	--	Folder
▶	six-1.10.0-py27_0	Today, 4:54 PM	--	Folder
▶	sqlite-3.13.0-0	Today, 4:54 PM	--	Folder
▶	tk-8.5.18-0	Today, 4:54 PM	--	Folder
	urls	Today, 4:54 PM	3 KB	TextEd...ument
	urls.txt	Today, 4:54 PM	2 KB	Plain Text
▶	wheel-0.29.0-py27_0	Today, 4:54 PM	--	Folder
▶	yaml-0.1.6-0	Today, 4:54 PM	--	Folder
▶	zlib-1.2.8-3	Today, 4:54 PM	--	Folder



# Lambda: Function as a Service

Compute



Memory

Time

**Dependency  
Zip File size**

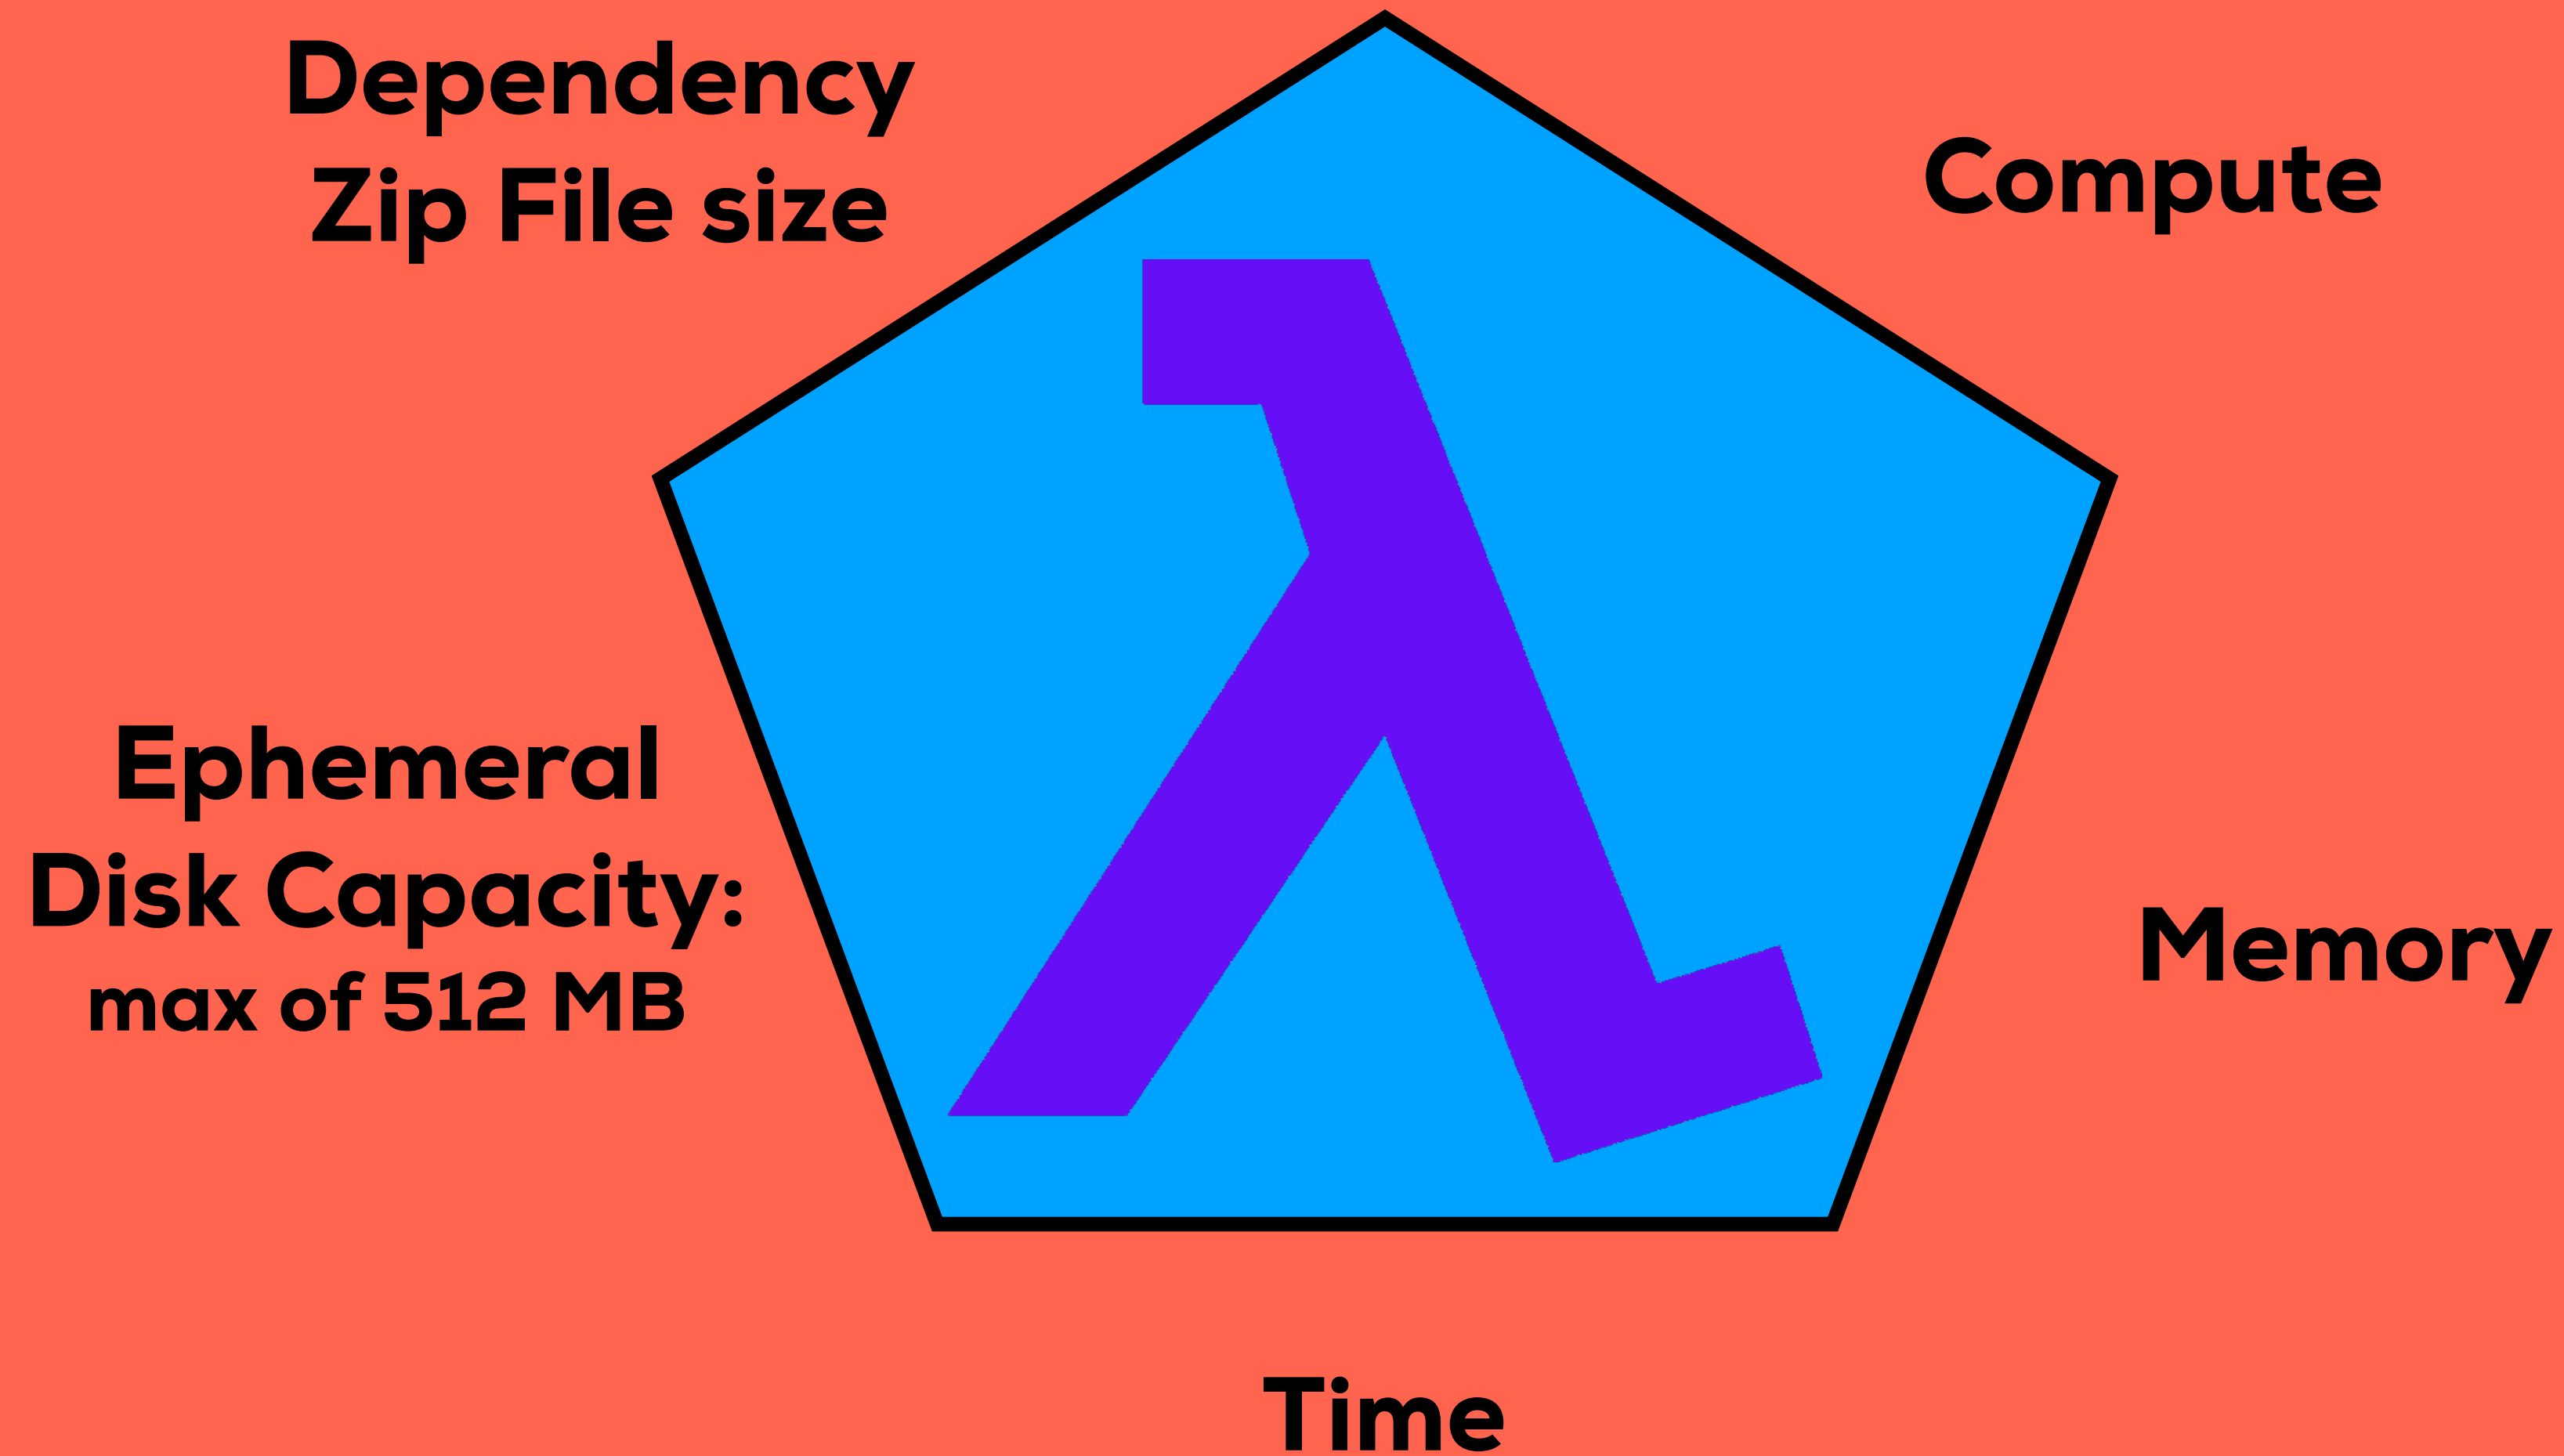
250 MB uncompressed  
code/dependencies



fastparquet  
didn't fit



# Lambda: Function as a Service



**Dependency**  
Zip File size

**Compute**


**Ephemeral  
Disk Capacity:**  
max of 512 MB

**Memory**

**Time**



# Lambda: Function as a Service

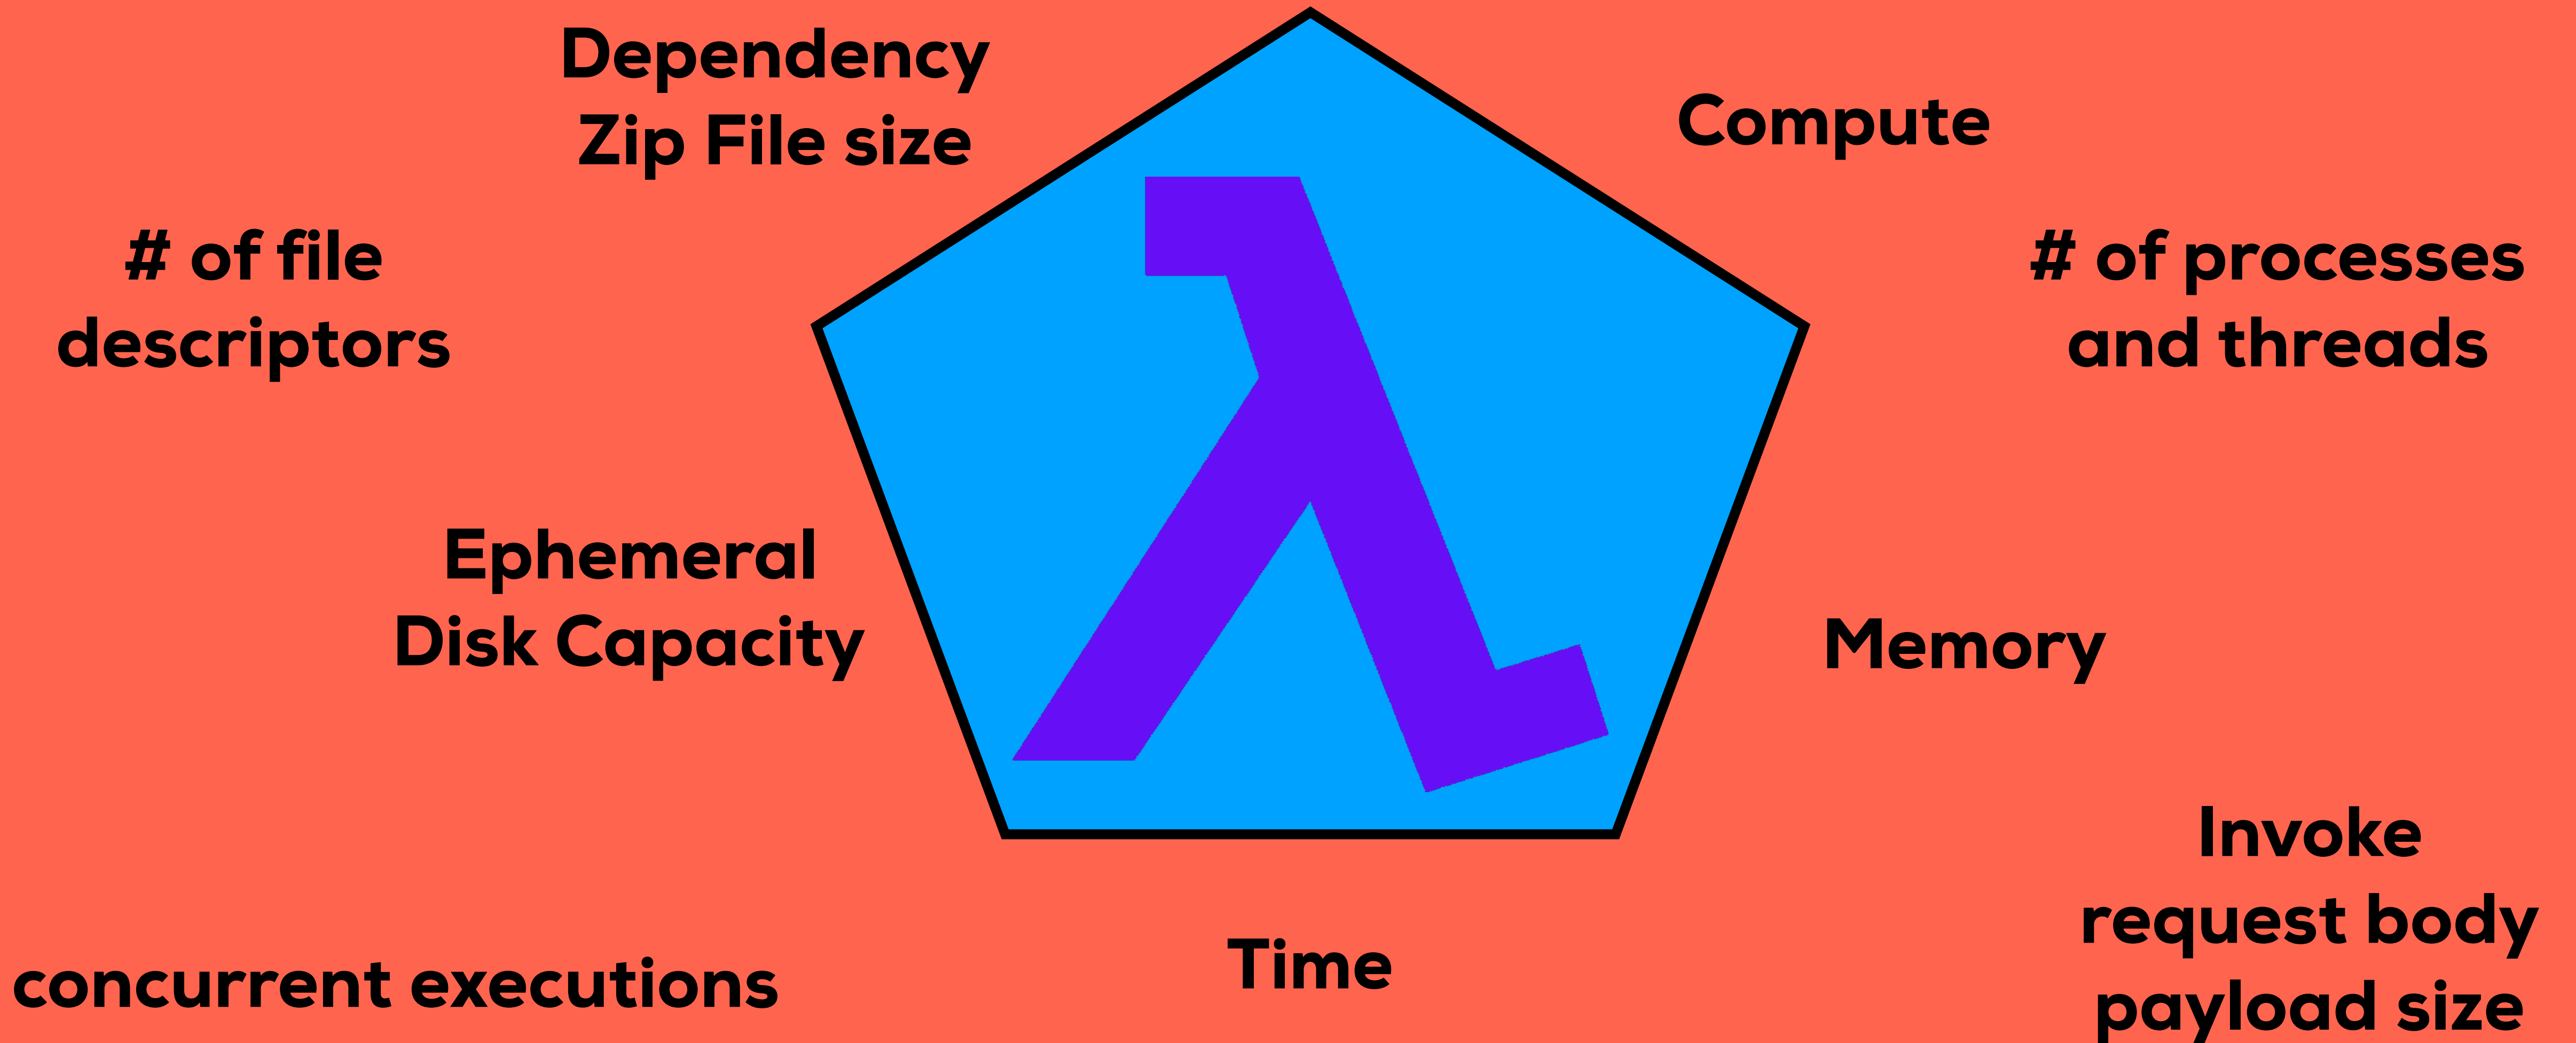


**concurrent executions:**  
default max of 1000 per account

**Time**

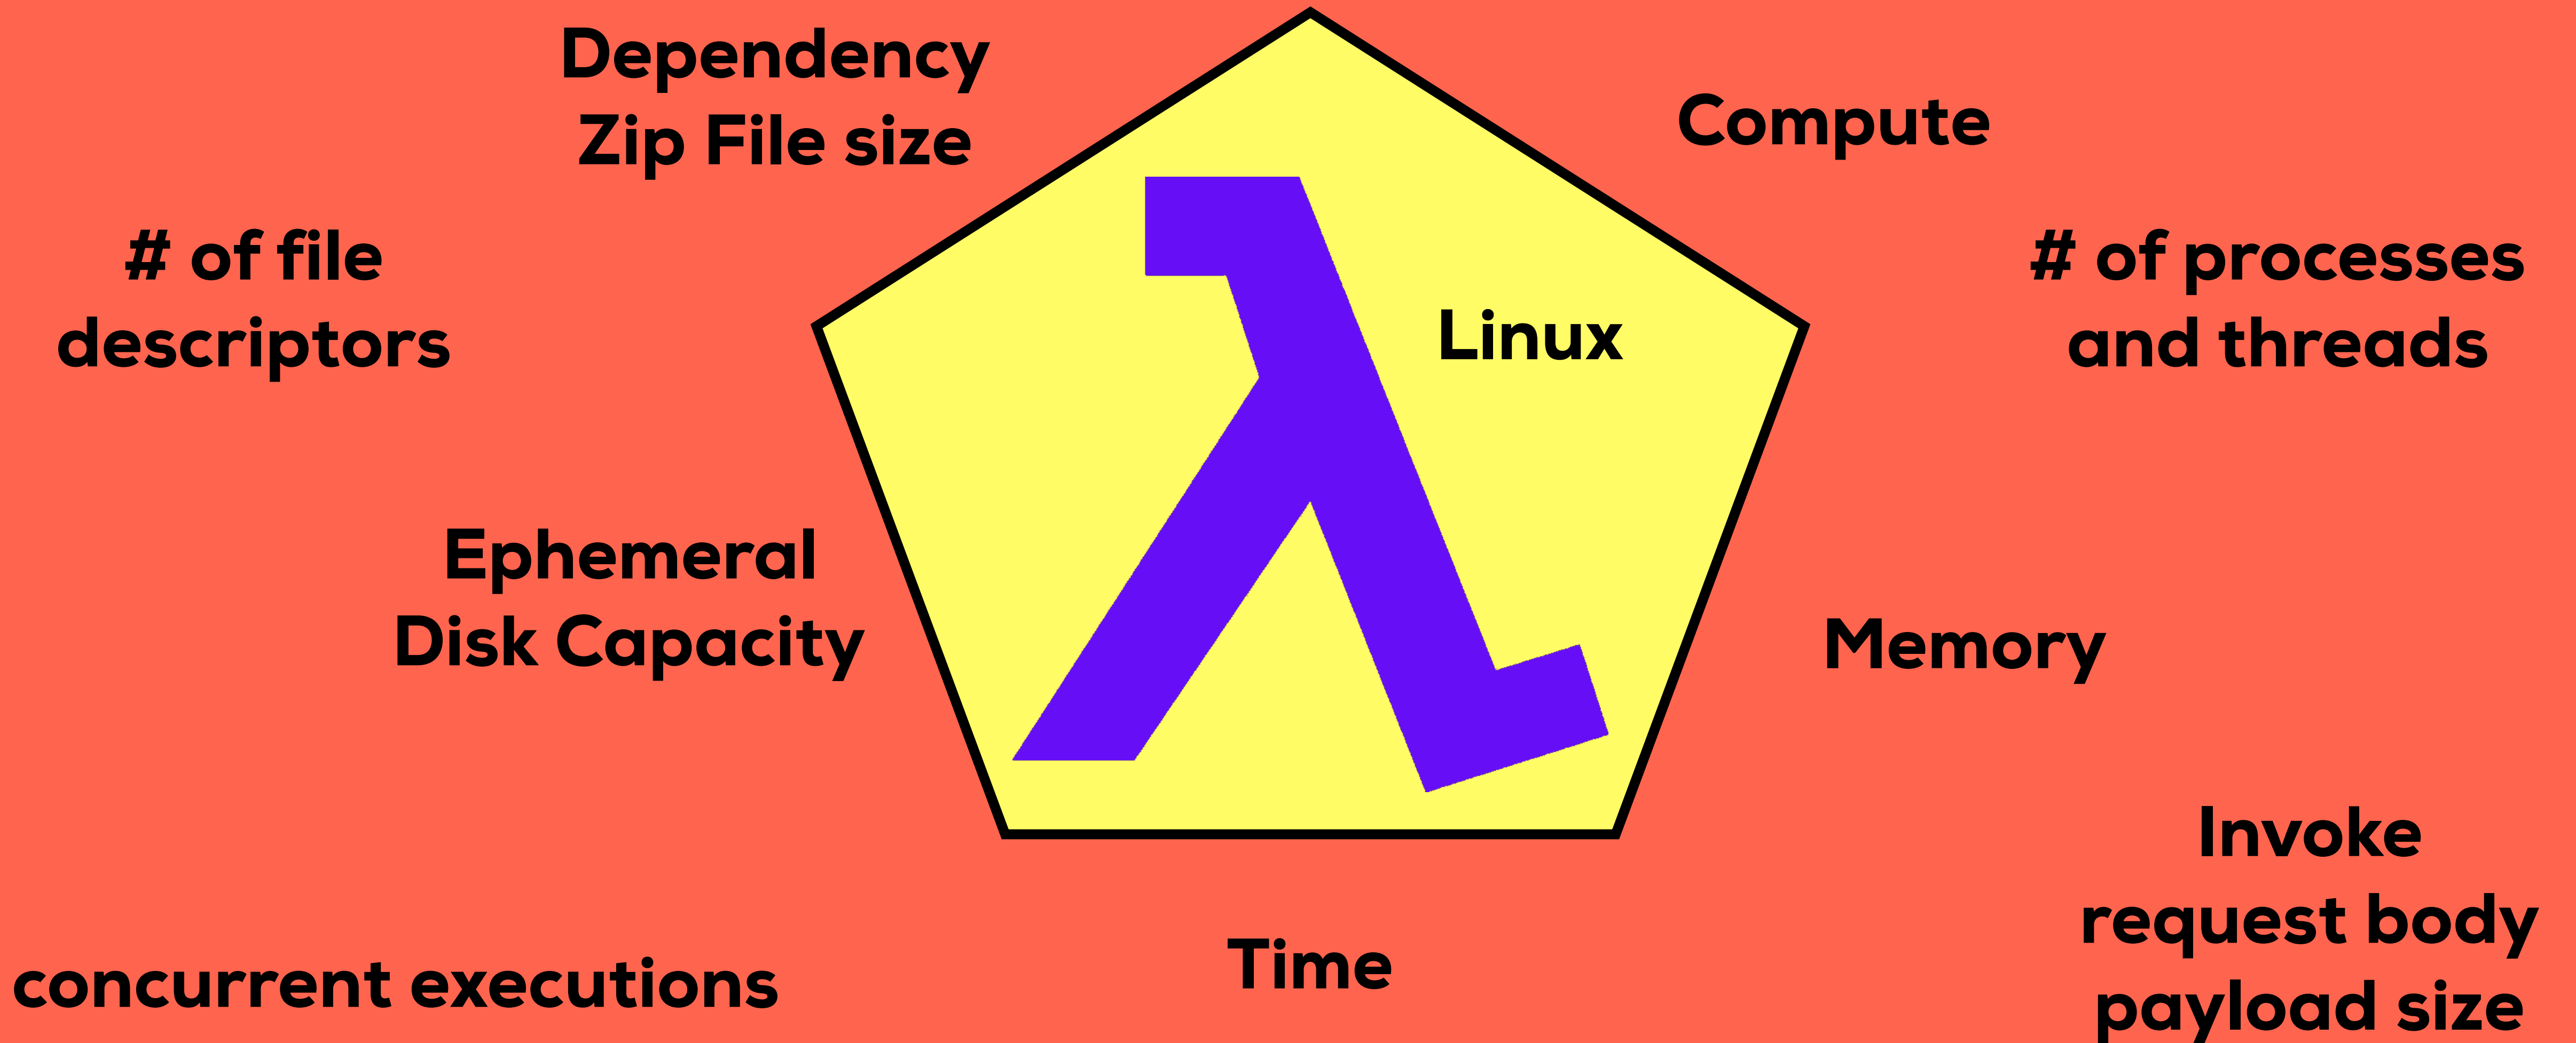


# Lambda: Function as a Service





# Lambda: Function as a Service





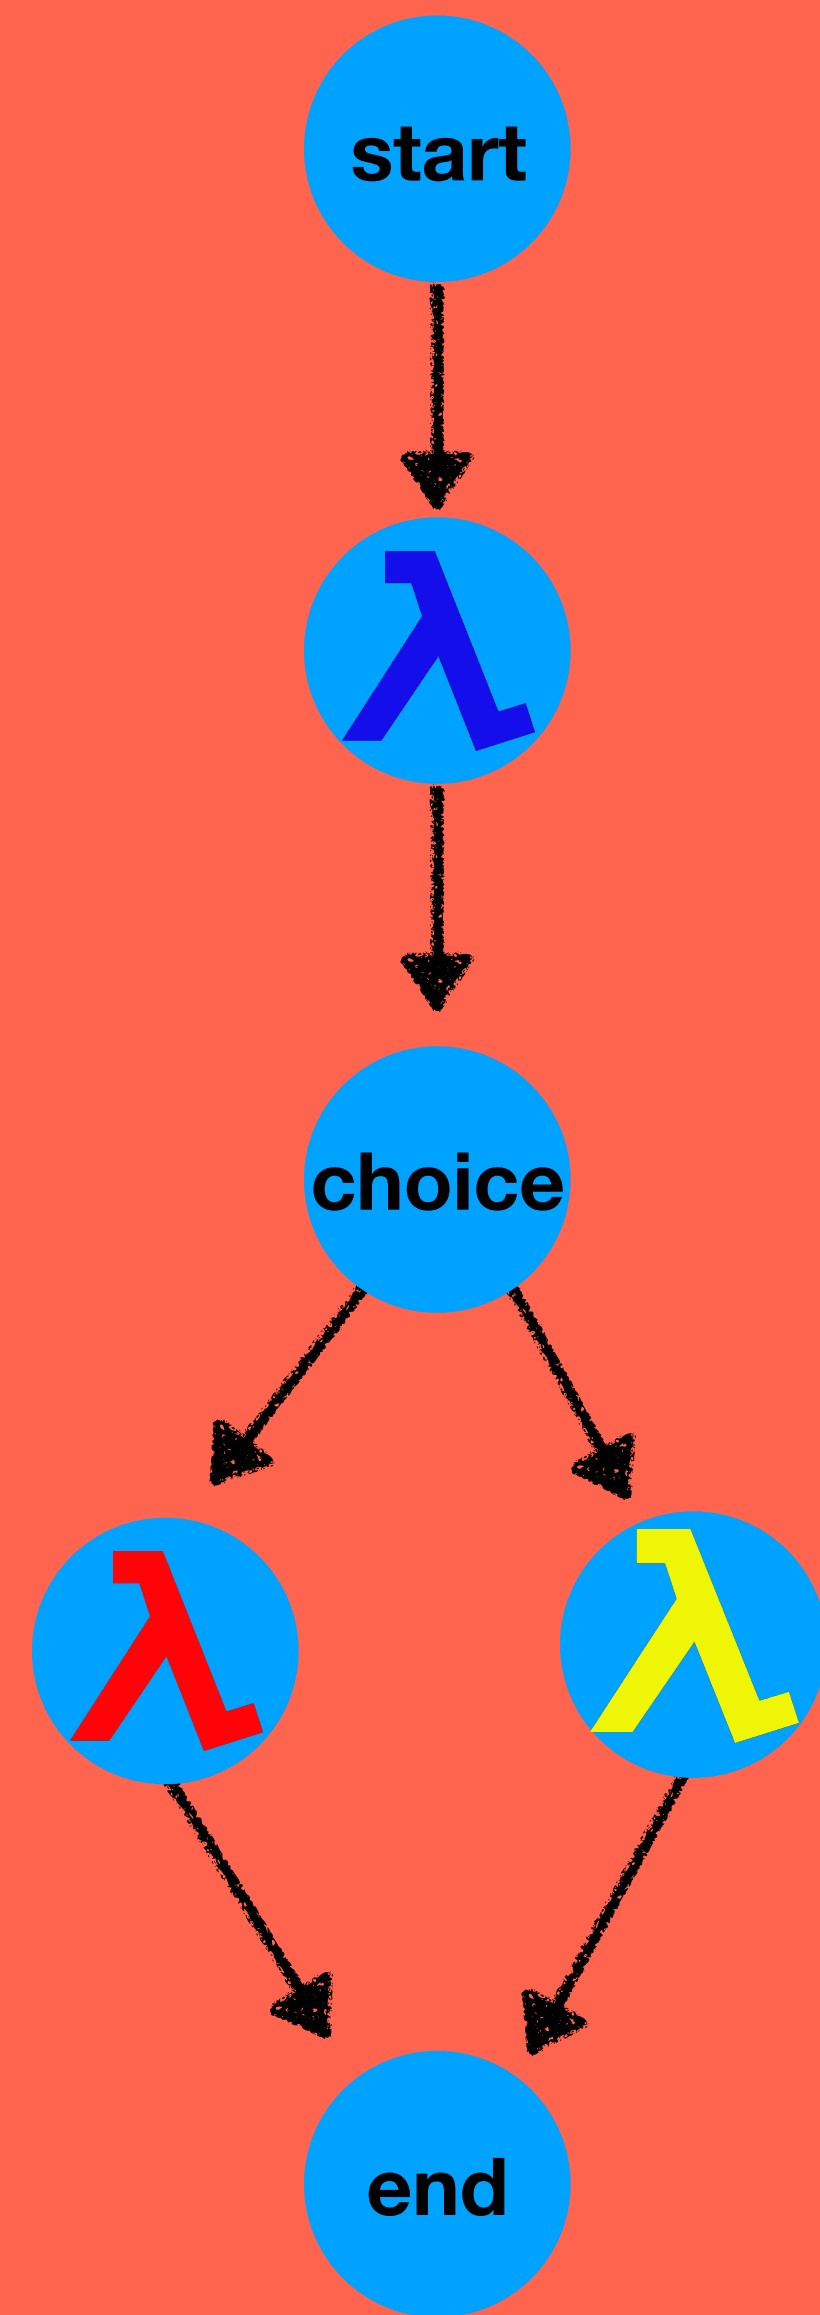






# Step Functions

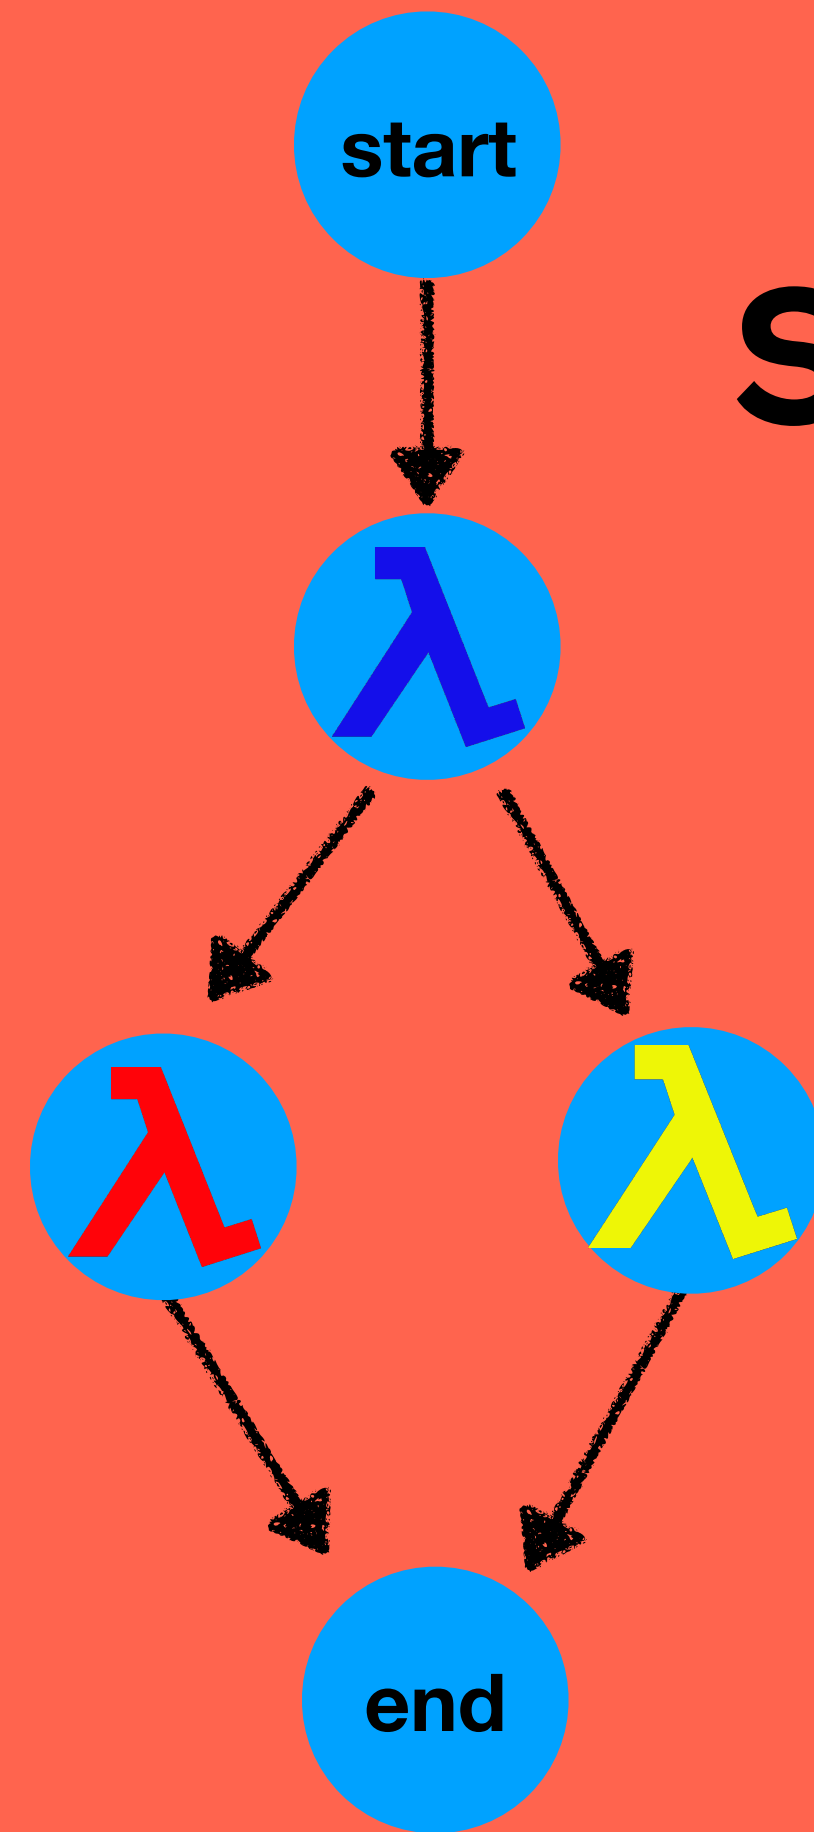




# Choice State

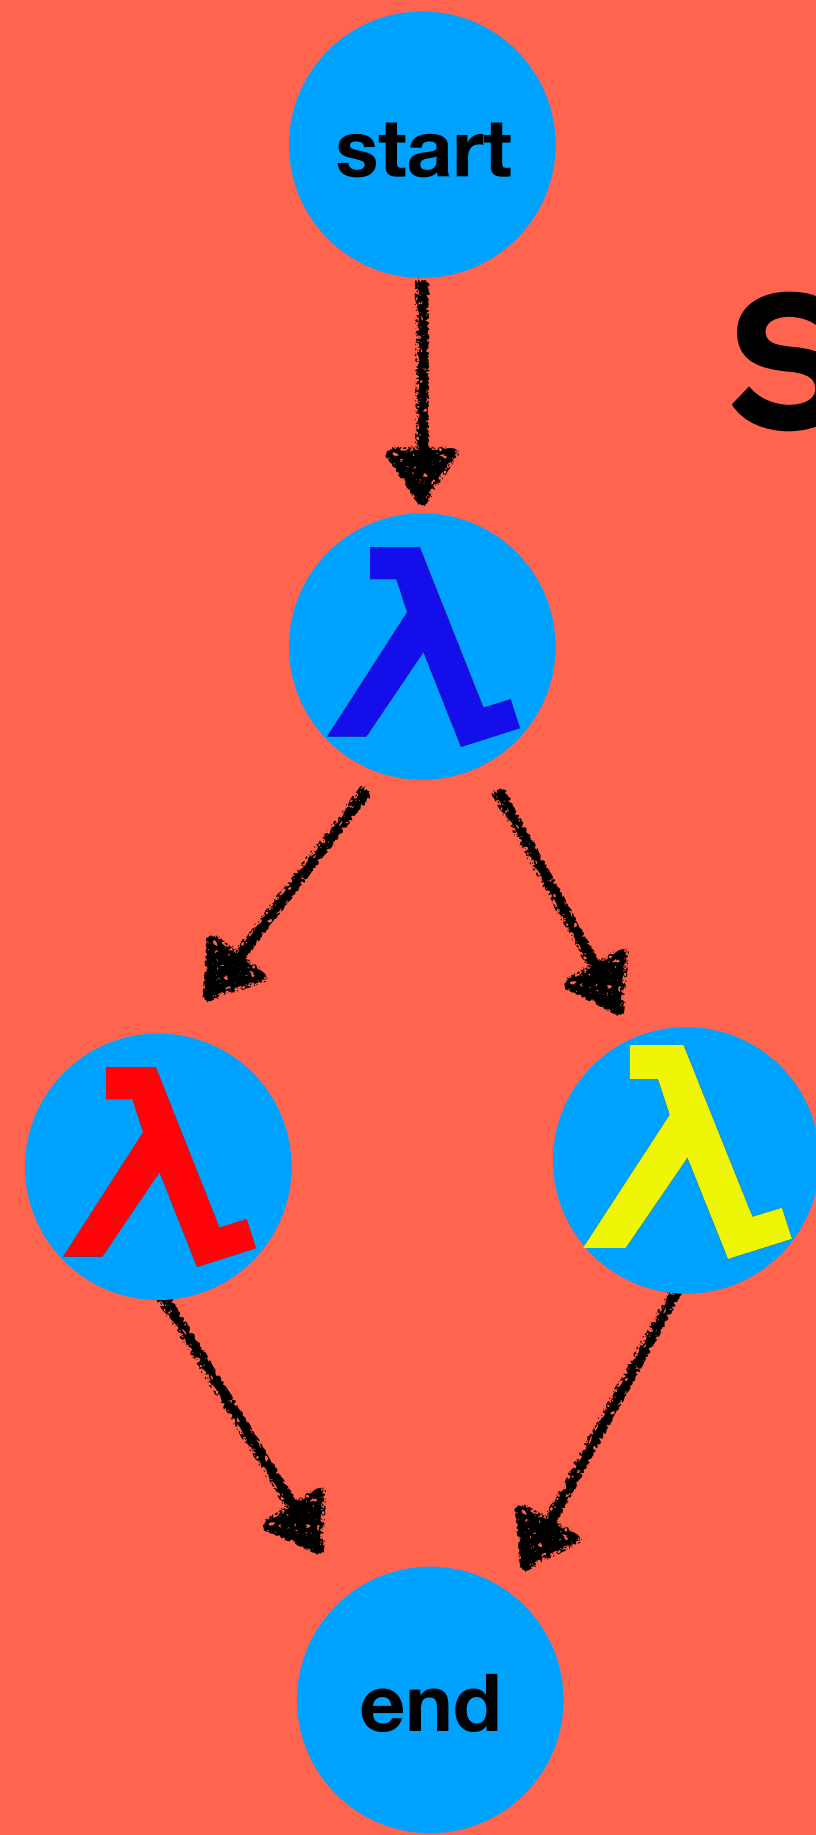


# Parallel Processing Step Functions can do

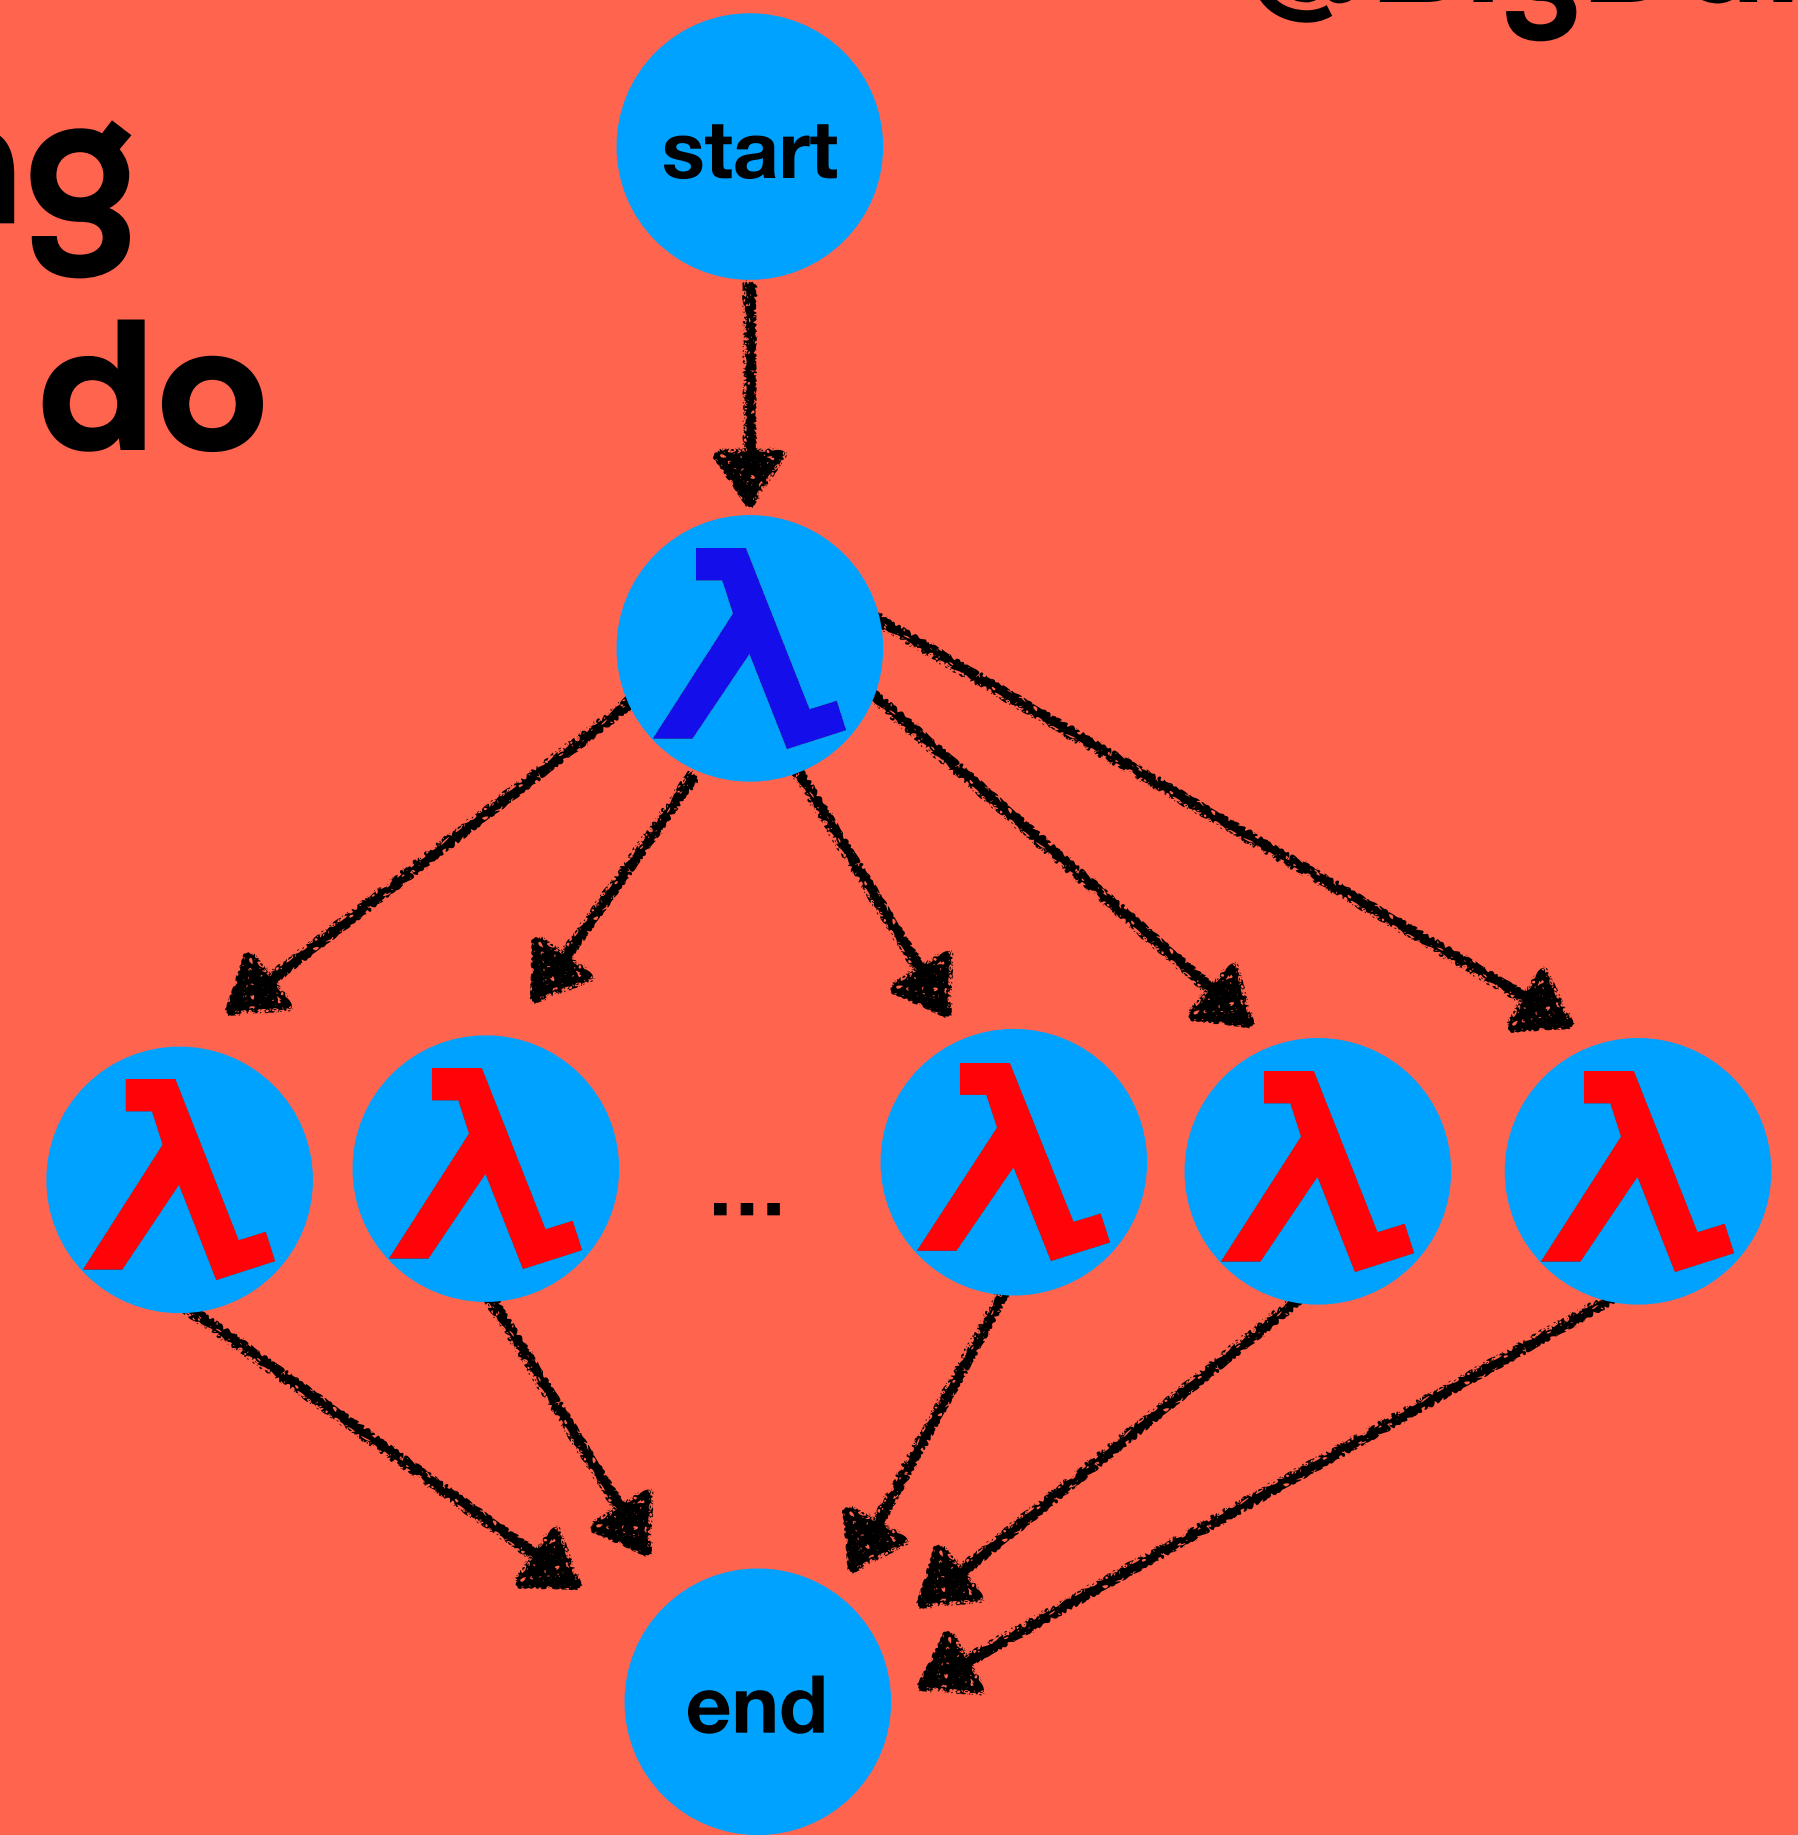




# Parallel Processing Step Functions can do



**vs**




## Dynamic Fan Out Parallel Processing











**VS**




**Defined Fan Out  
Parallel Processing**

**Dynamic Fan Out  
Parallel Processing**





**or**



**Defined Fan Out  
Parallel Processing**

**send to an EC2 instance**




# Step Functions



**easy to implement  
workflow management for  
lambdas and ec2 instances**

**built-in back off policy  
for retrying lambdas**



**doesn't yet natively  
support a dynamic  
one-to-many  
fan out architecture  
of lambdas**



WONDER WOMAN'S EYELID  
MUSCLES LOOSEN THE TAPE-BUT  
ALSO HER EYELASHES!

UN-UNH! MY FEMININE VANITY  
WON'T LET ME PULL OUT MY  
EYELASHES! I'LL HAVE TO  
ESCAPE BLINDFOLDED!





# Building a Data Pipeline


**Attempt 0: Python Multiprocessing on an EC2 Instance**

**Attempt 0.5: Python Multiprocessing on a Spot Instance**

**Attempt 1: Lambdas Orchestrated by Step Functions**

**Attempt 2: Using Event Triggers with S3**

**Attempt 3: Rate limiting with Queues and Cloudwatch triggered Lambdas**





# Building a Data Pipeline


**Attempt 0: Python Multiprocessing on an EC2 Instance**

**Attempt 0.5: Python Multiprocessing on a Spot Instance**

**Attempt 1: Lambdas Orchestrated by Step Functions**

**Attempt 2: Using Event Triggers with S3**

**Attempt 3: Rate limiting with Queues and Cloudwatch triggered Lambdas**








**S3 Bucket  
(Simple Storage Service)**



# Using Event Triggers

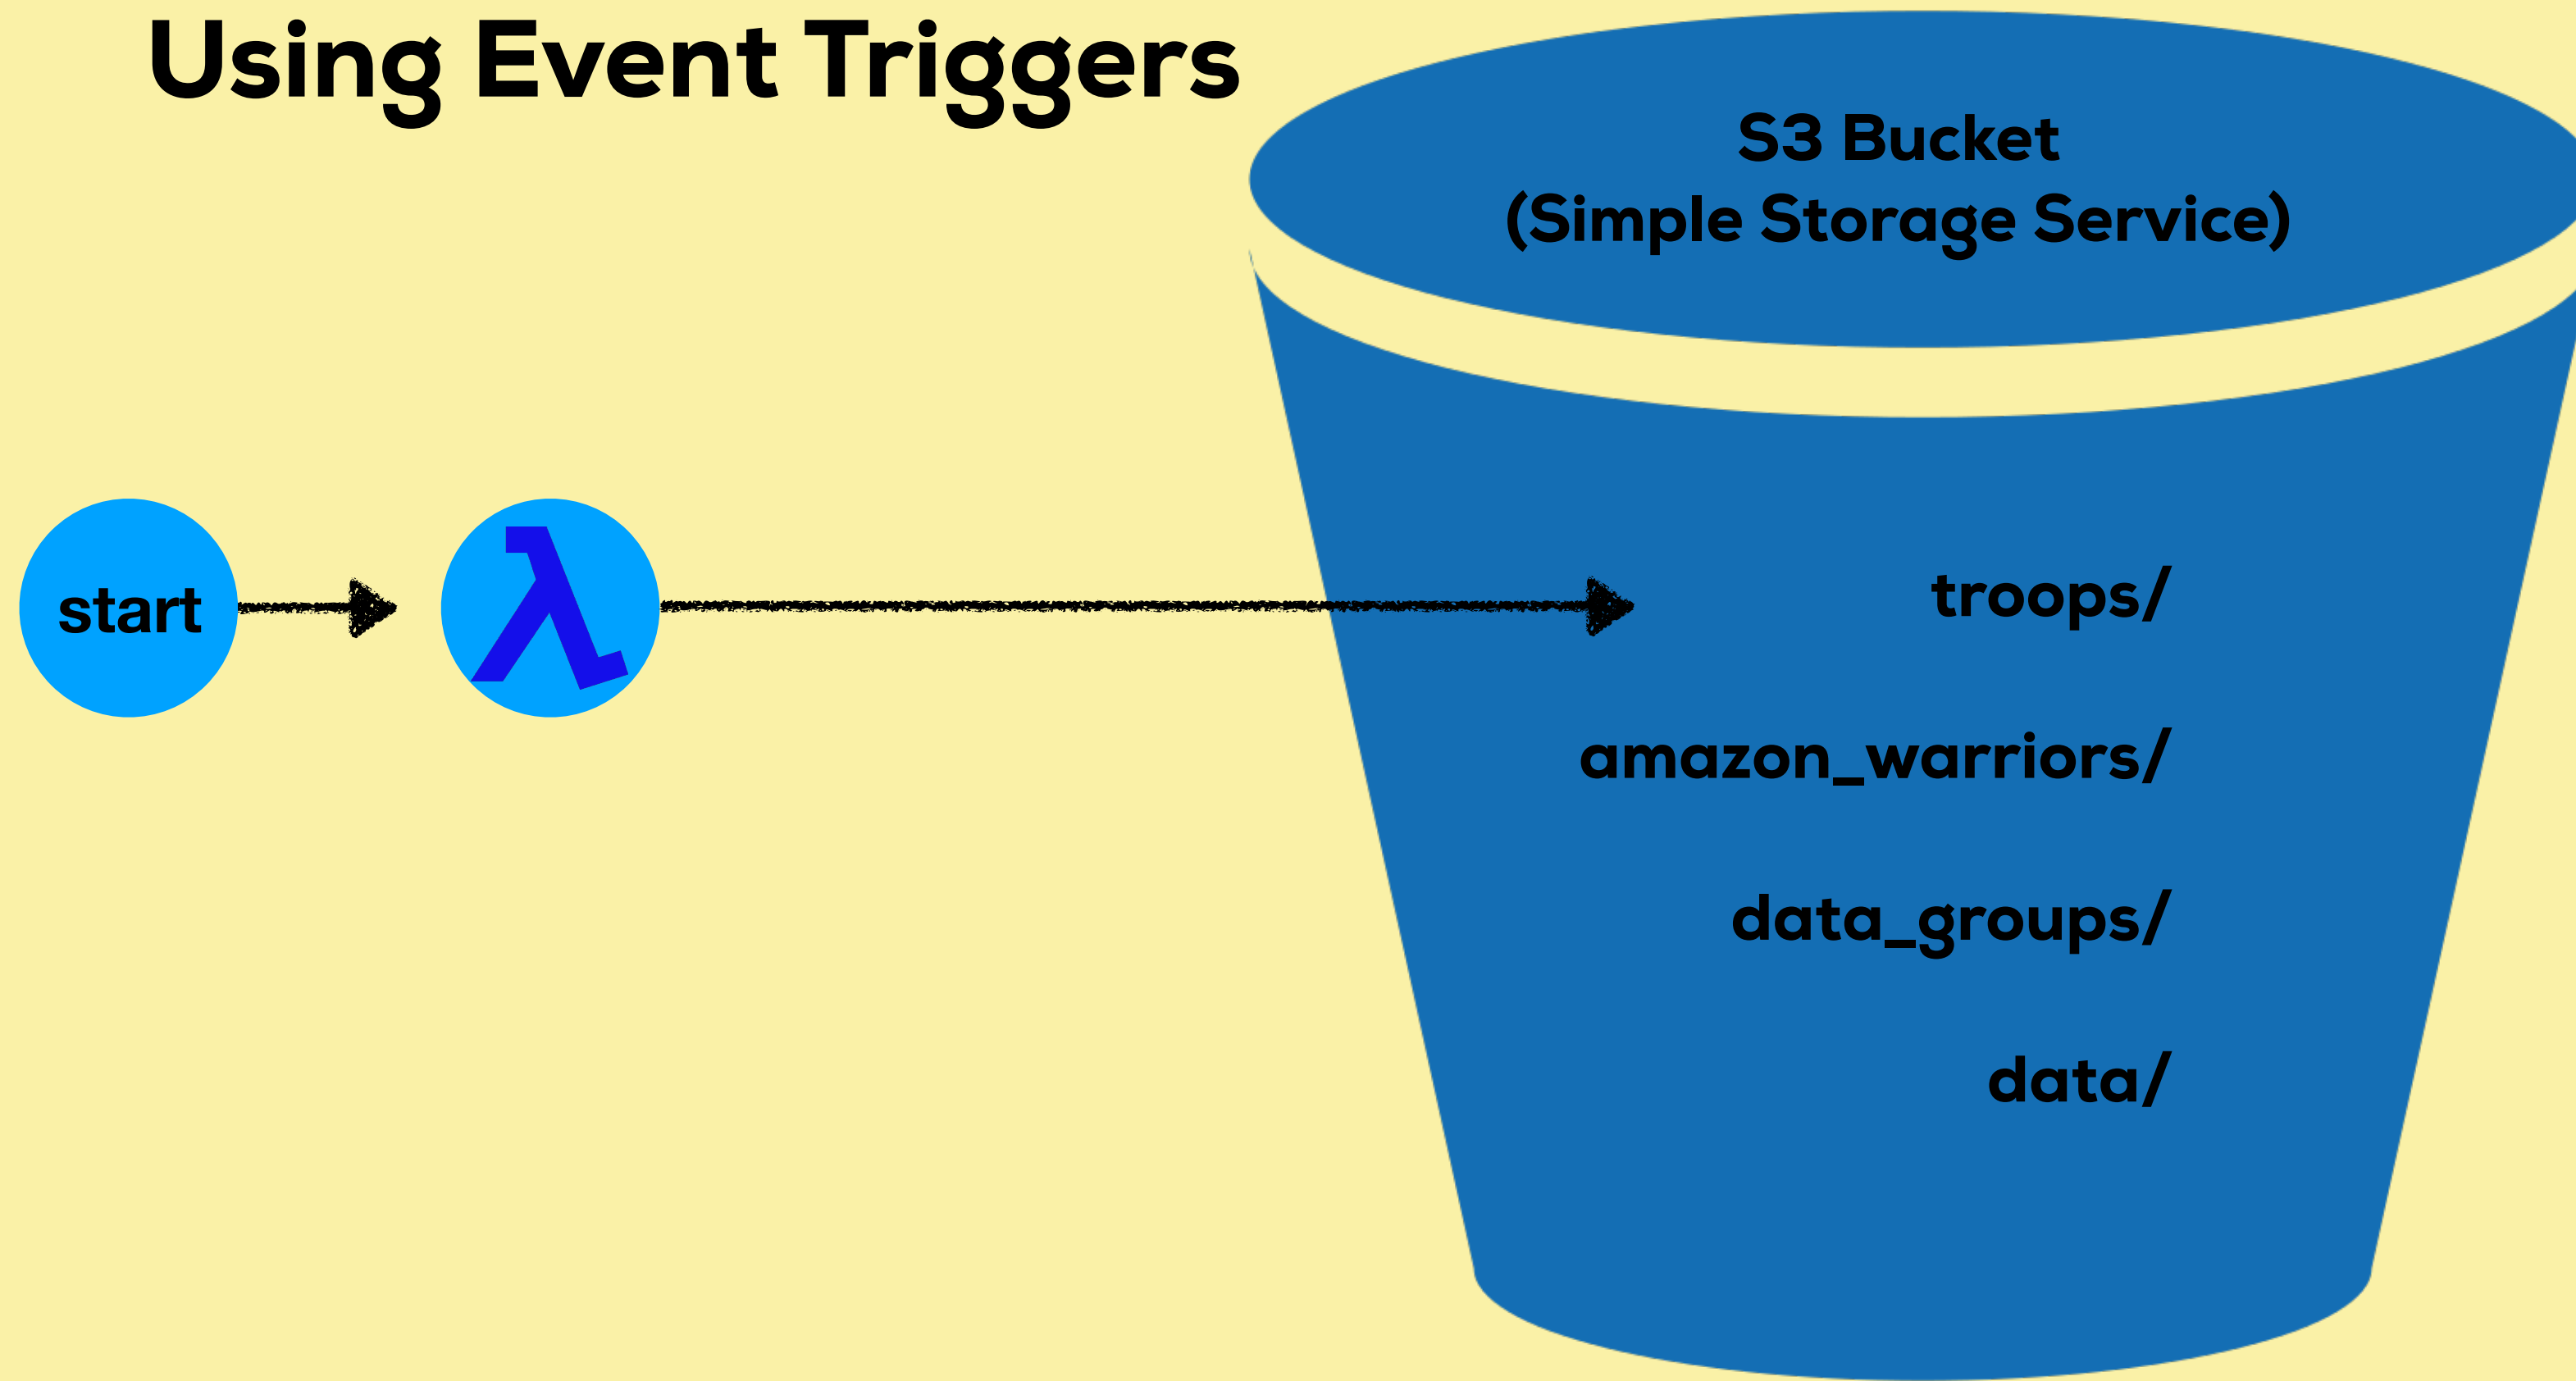
@BigDana





# Using Event Triggers

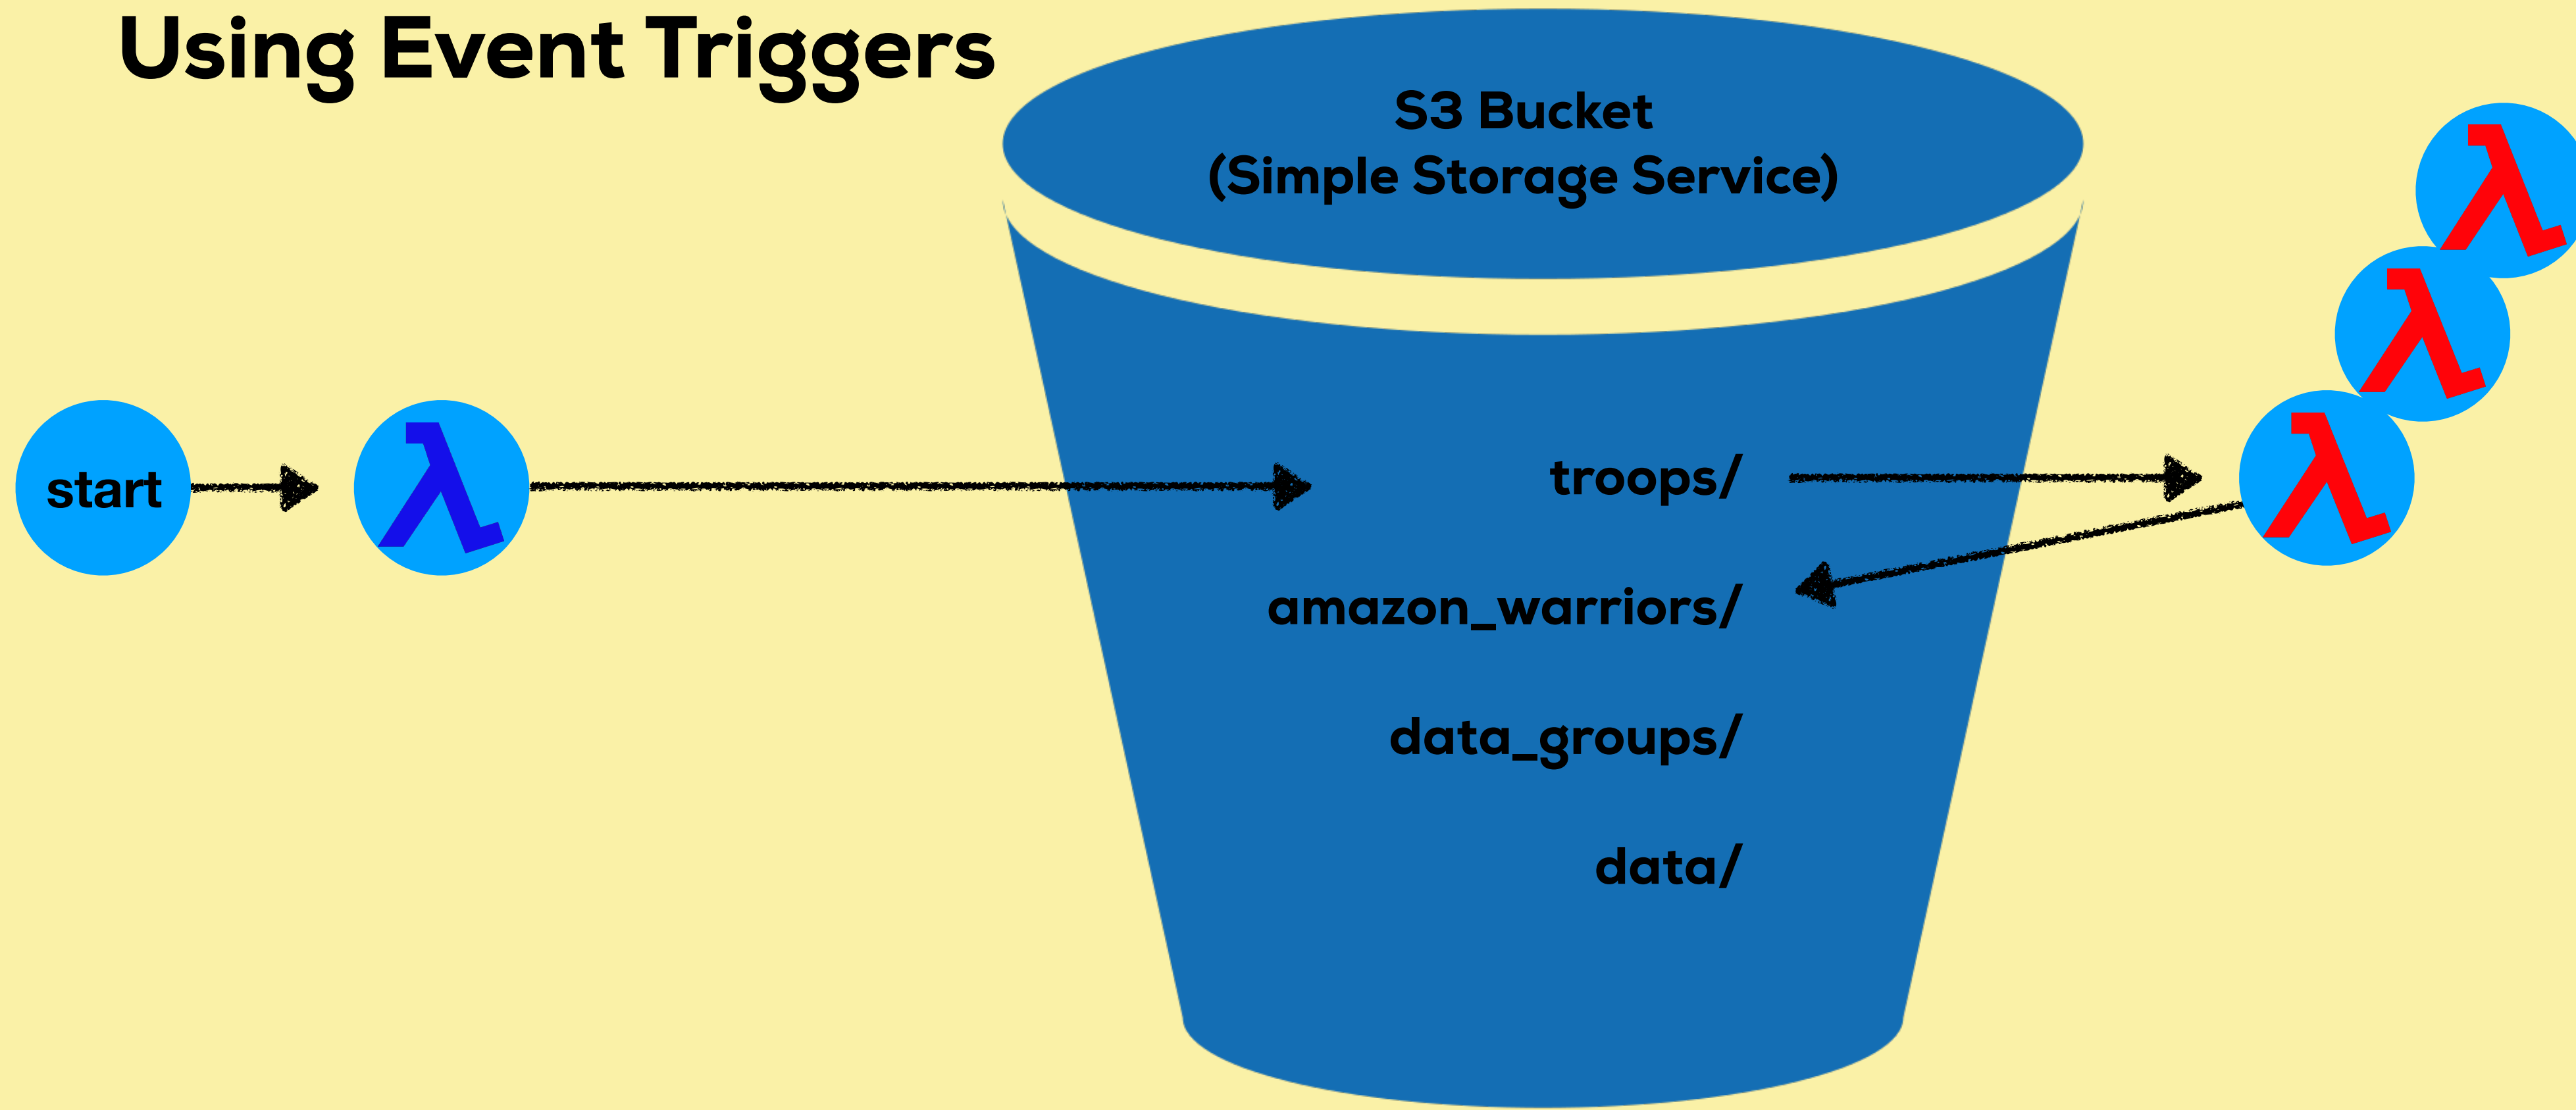
@BigDana





# Using Event Triggers

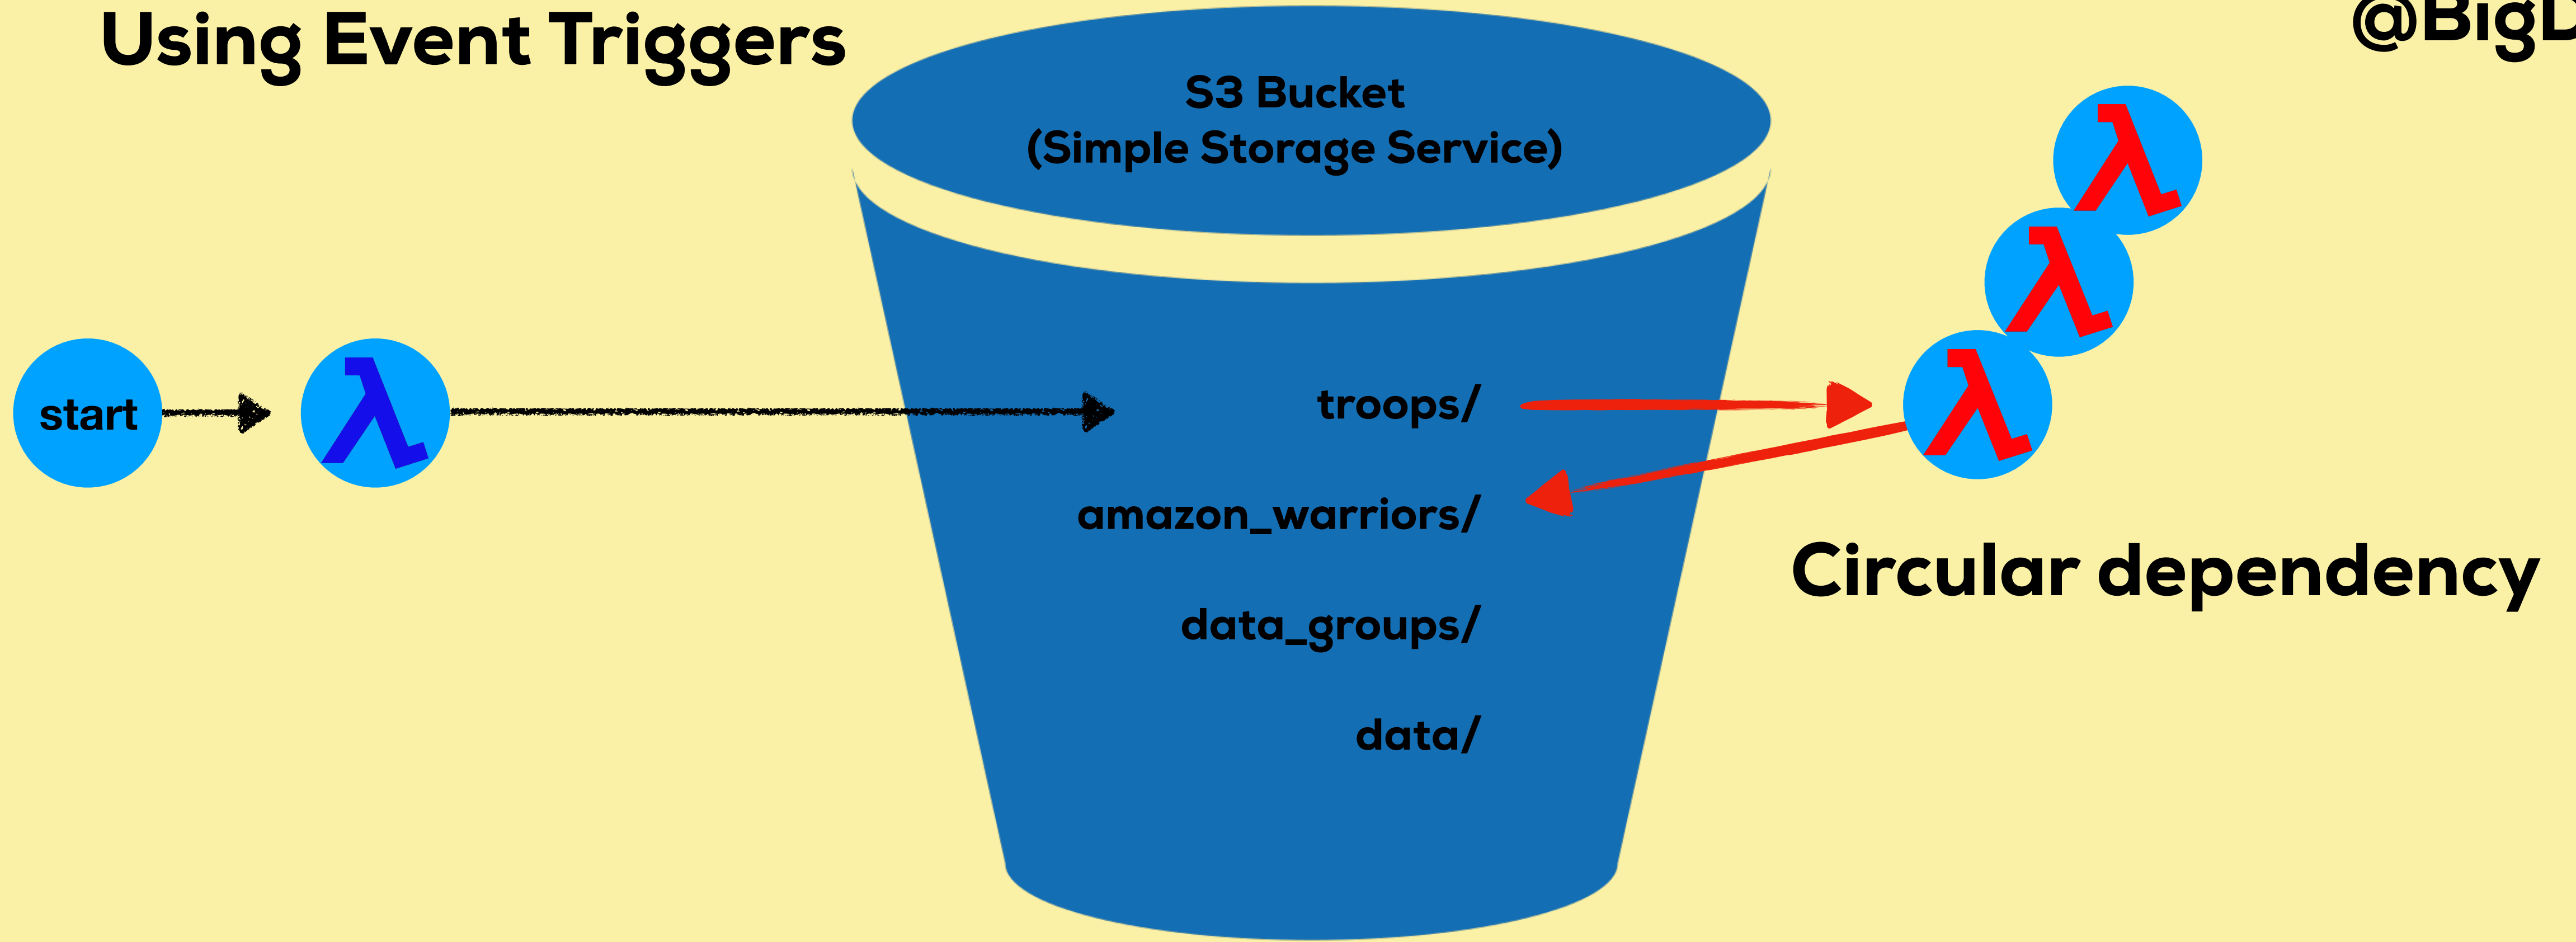
@BigDana





# Using Event Triggers

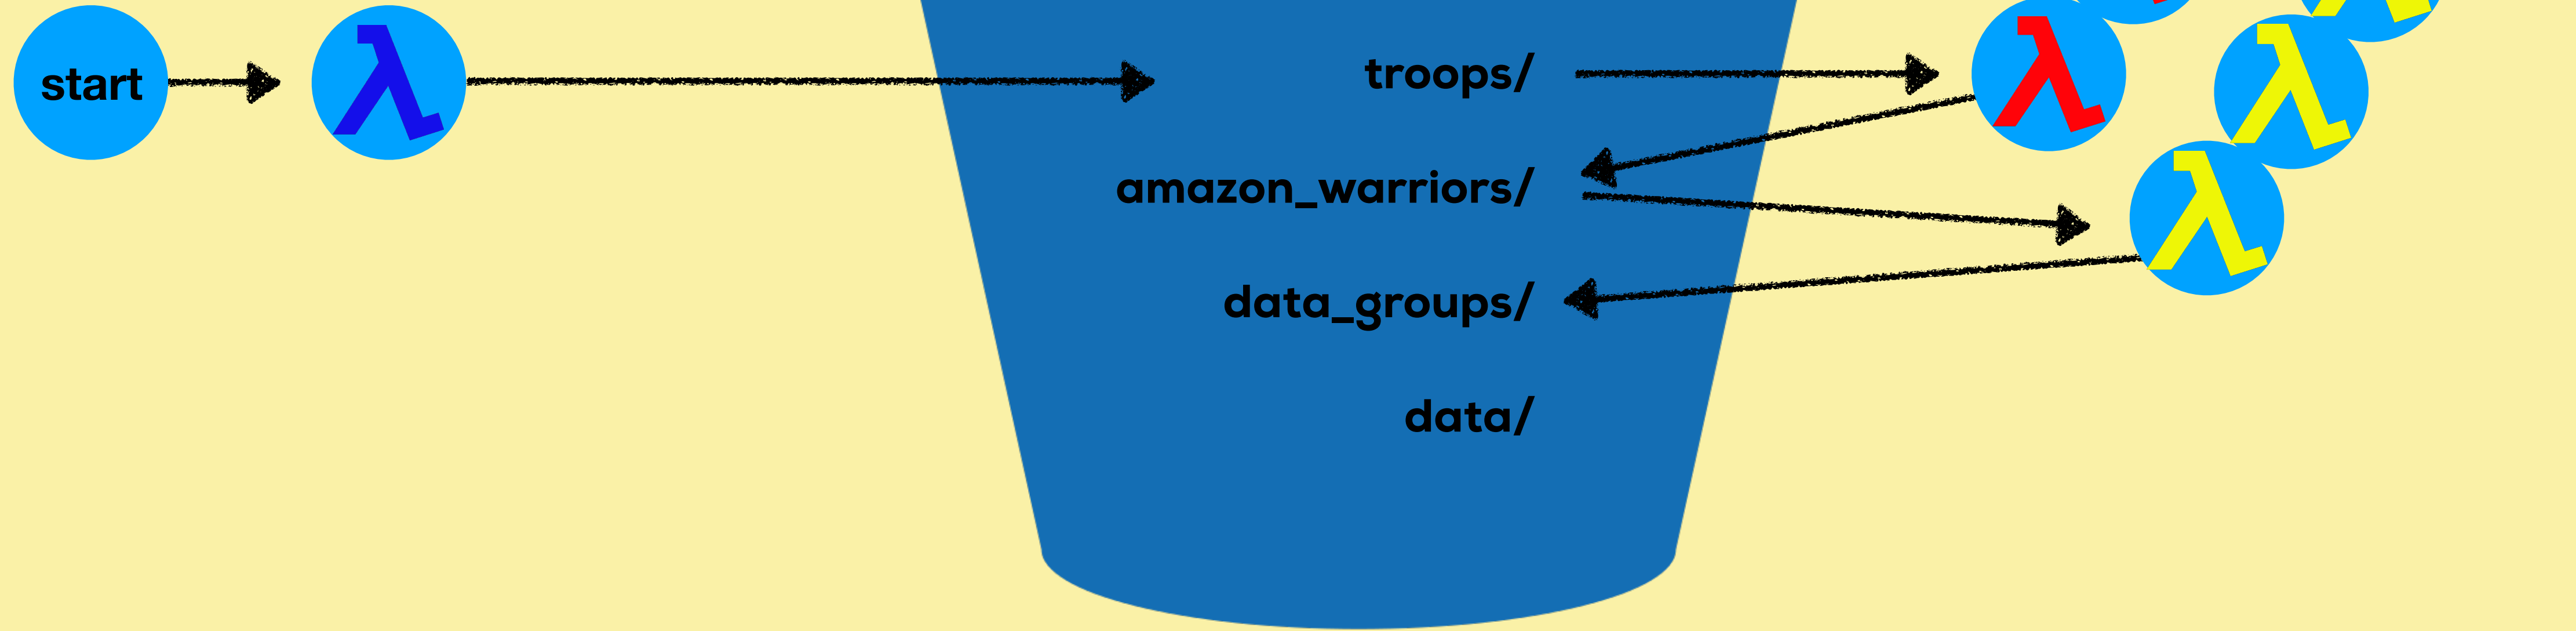
@BigDana





# Using Event Triggers

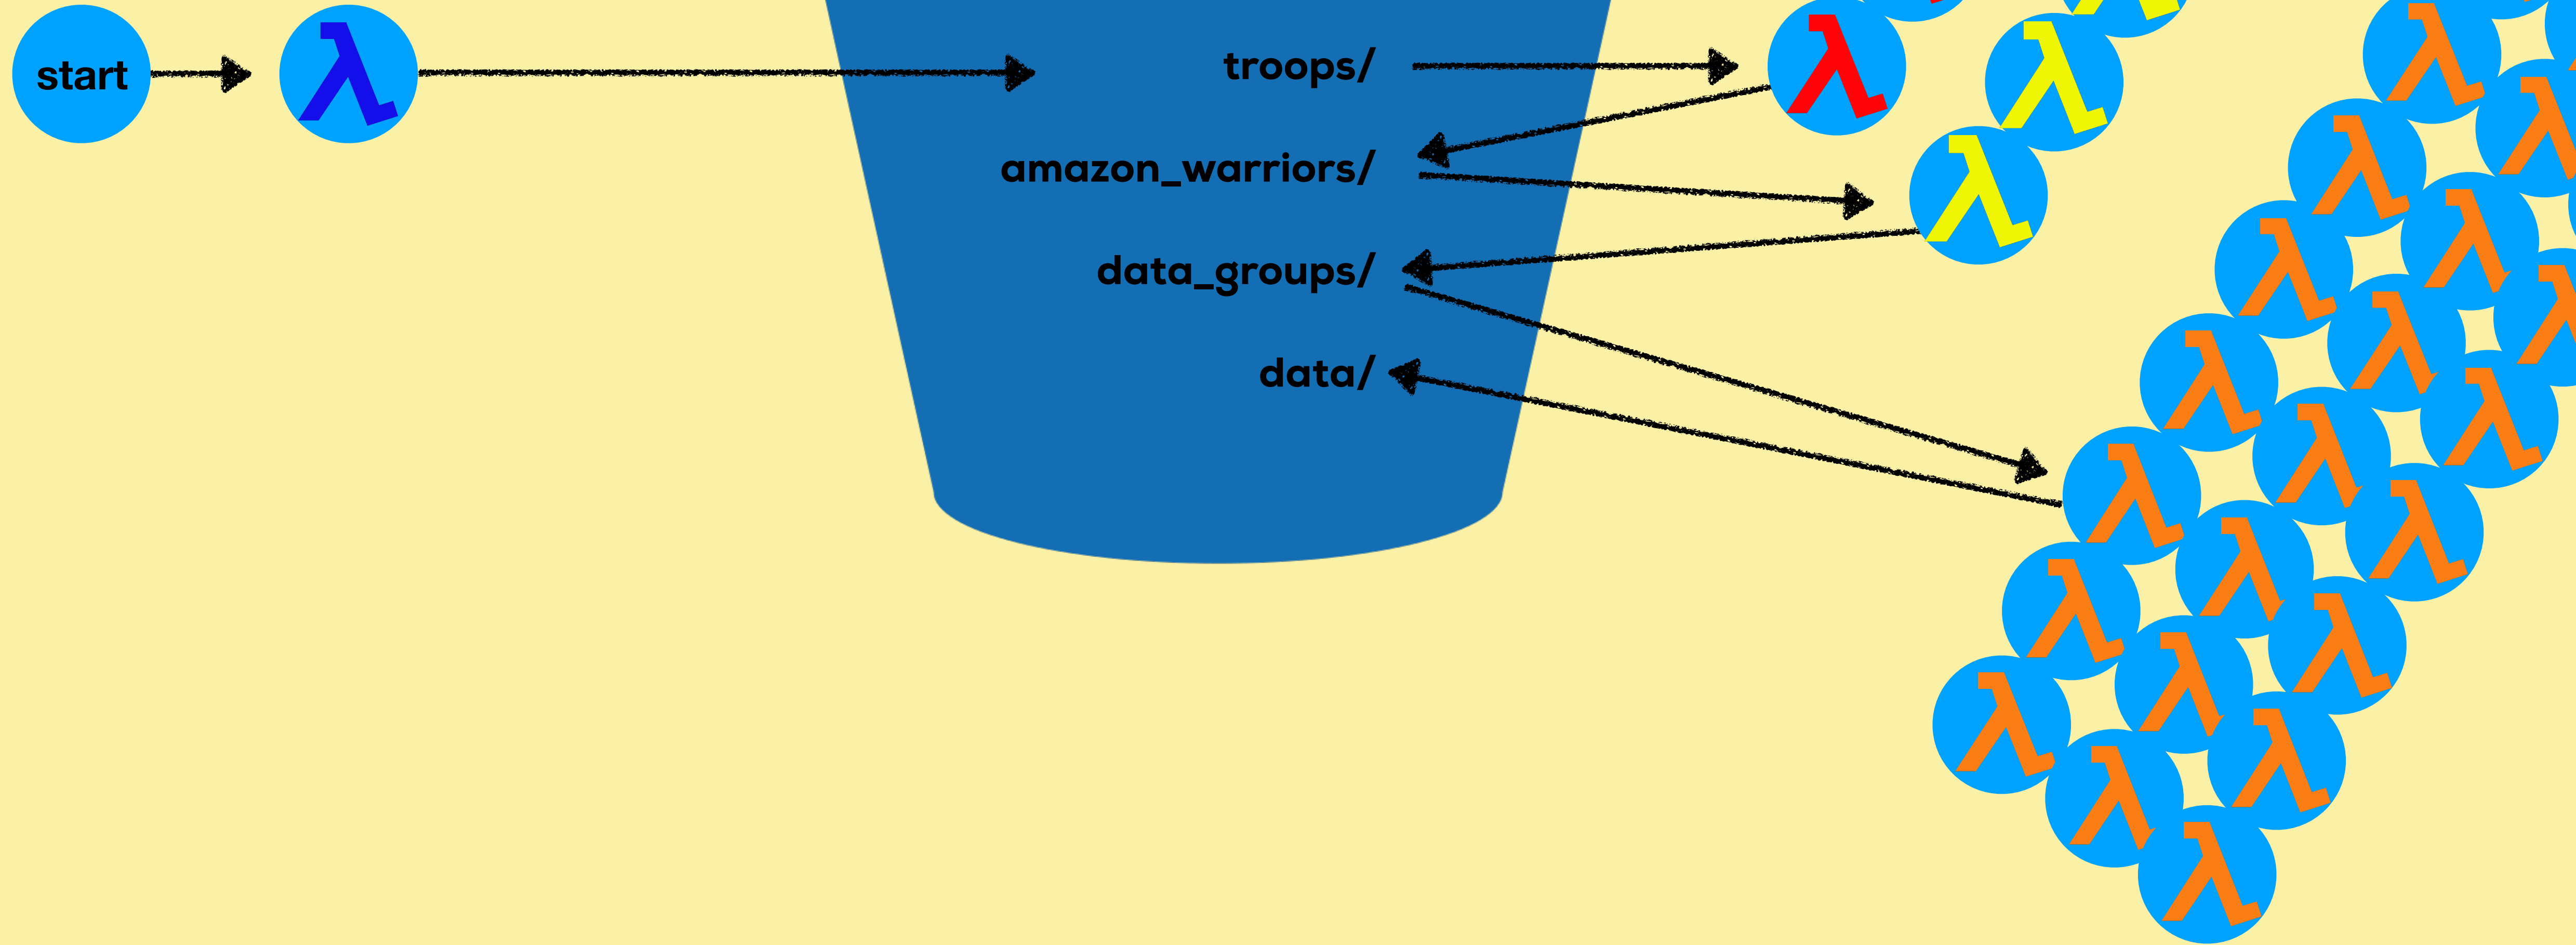
@BigDana





# Using Event Triggers

@BigDana



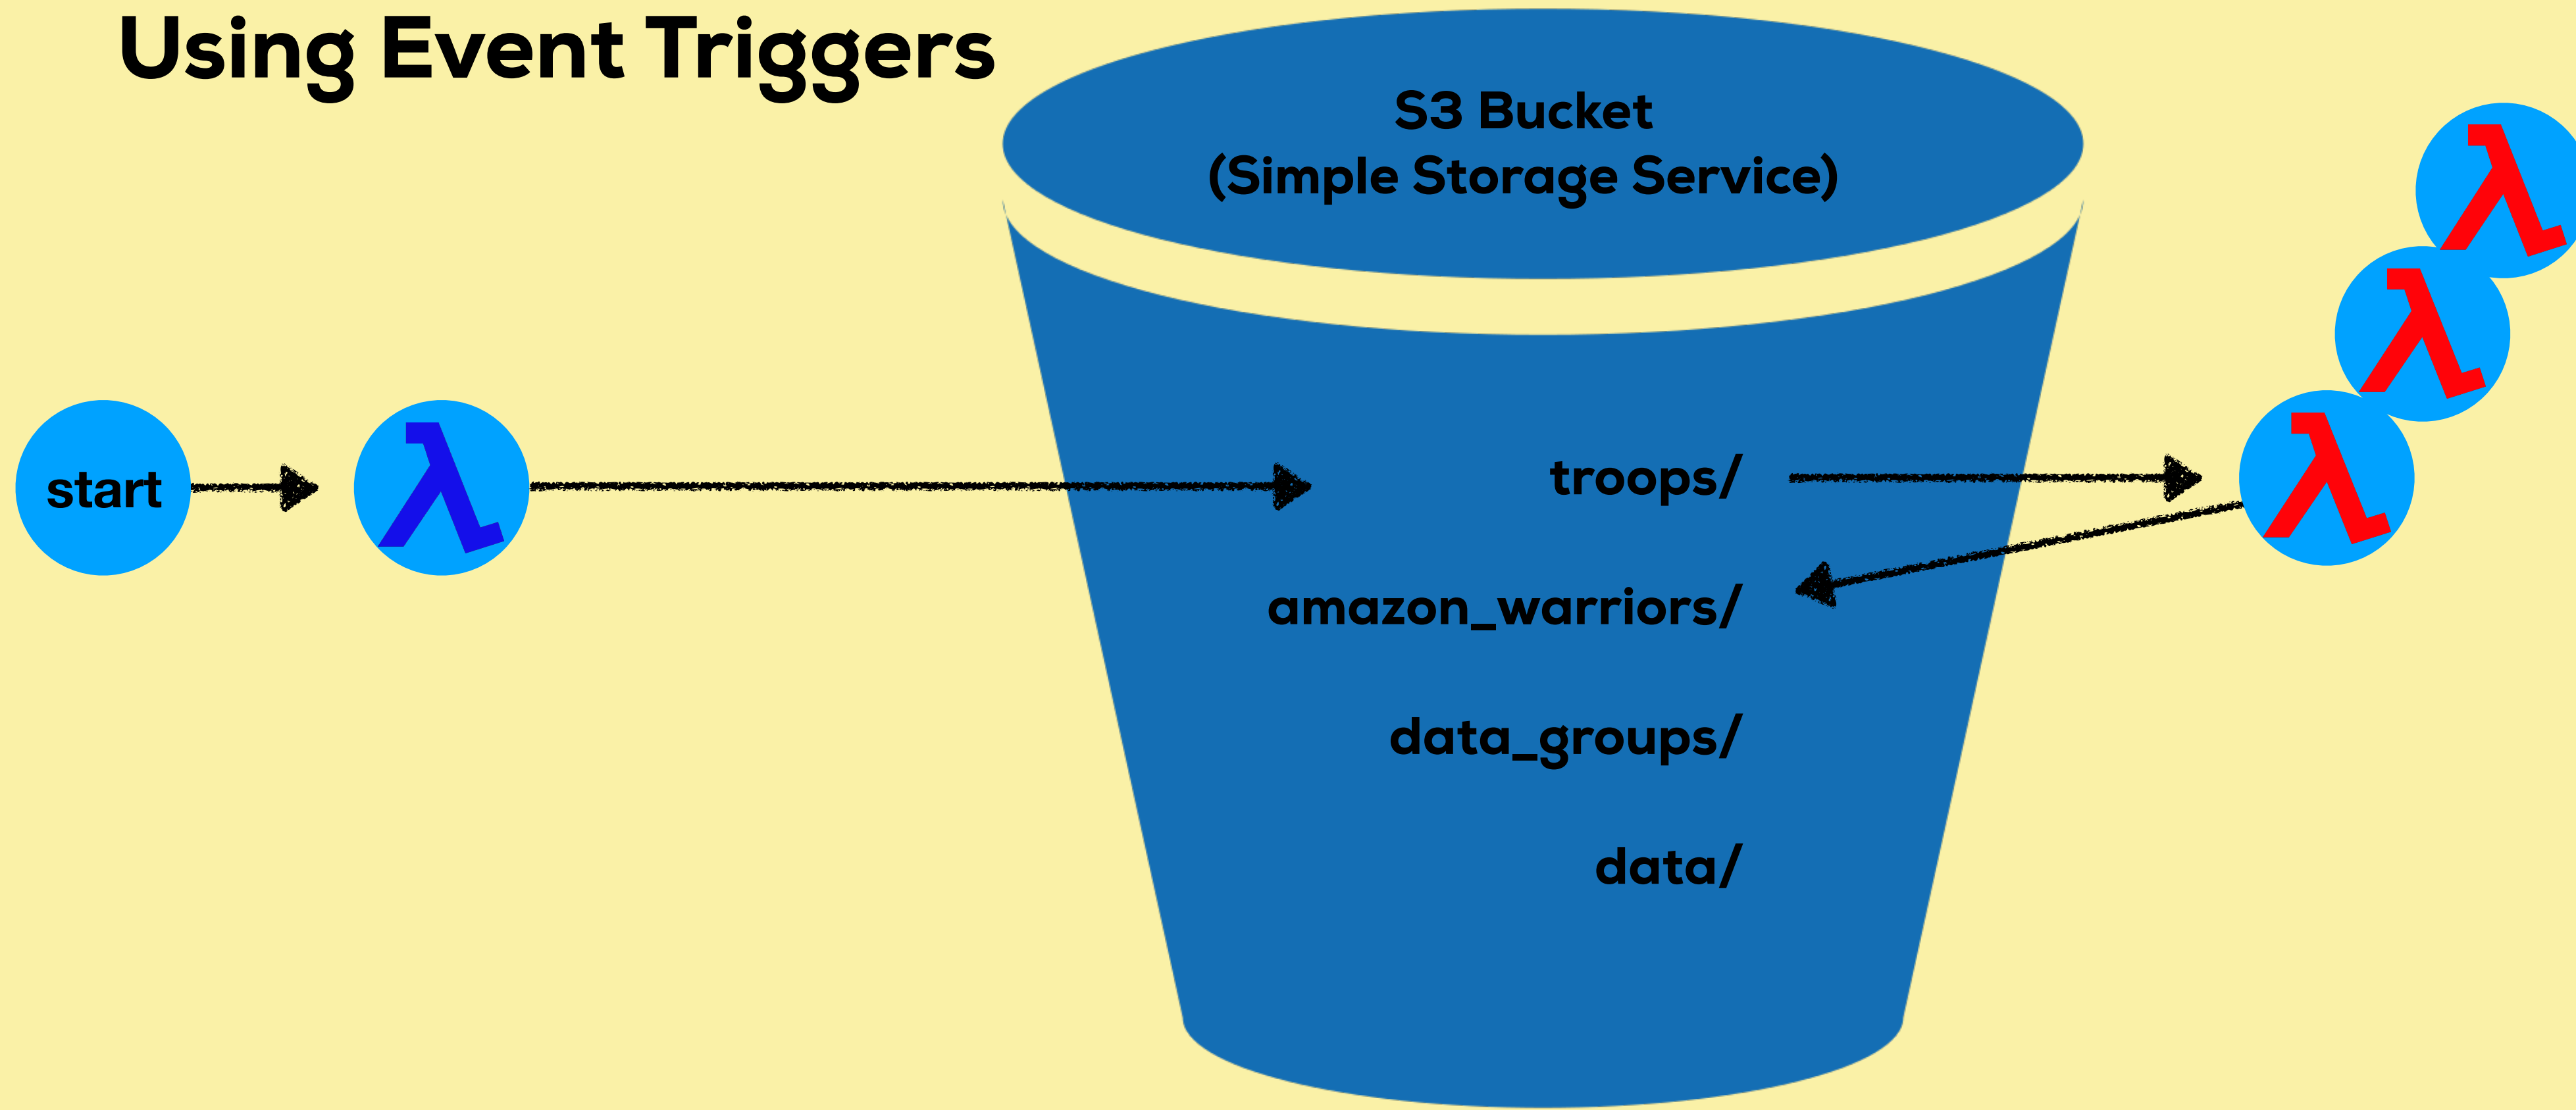






# Using Event Triggers

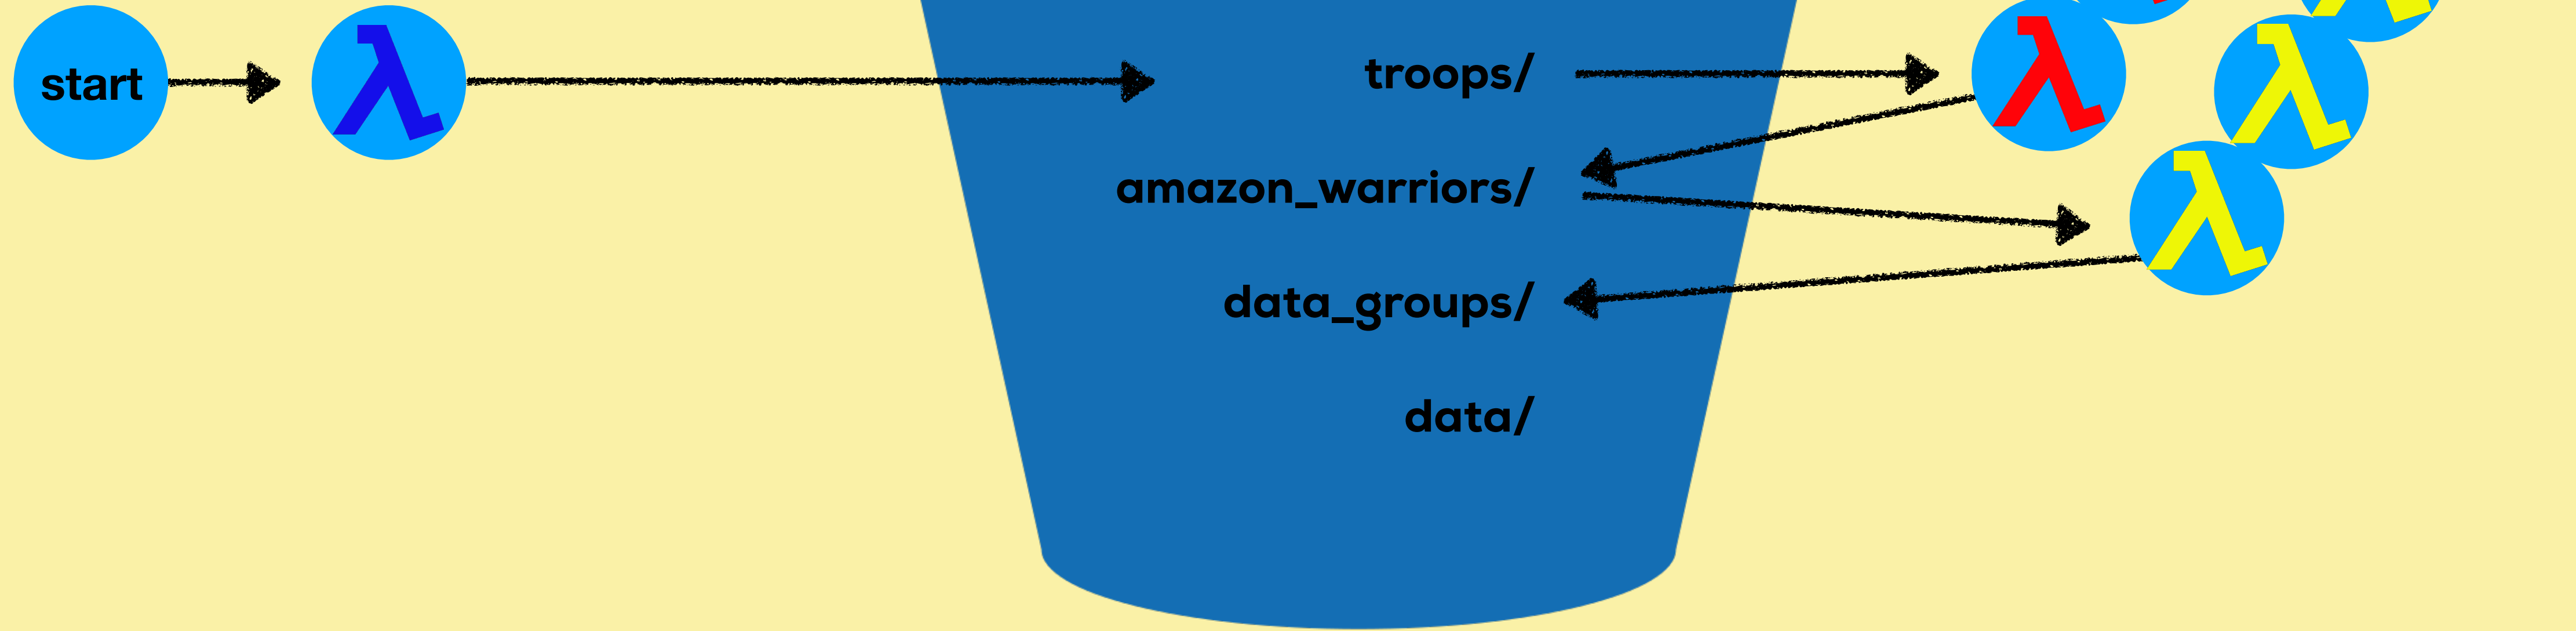
@BigDana





# Using Event Triggers

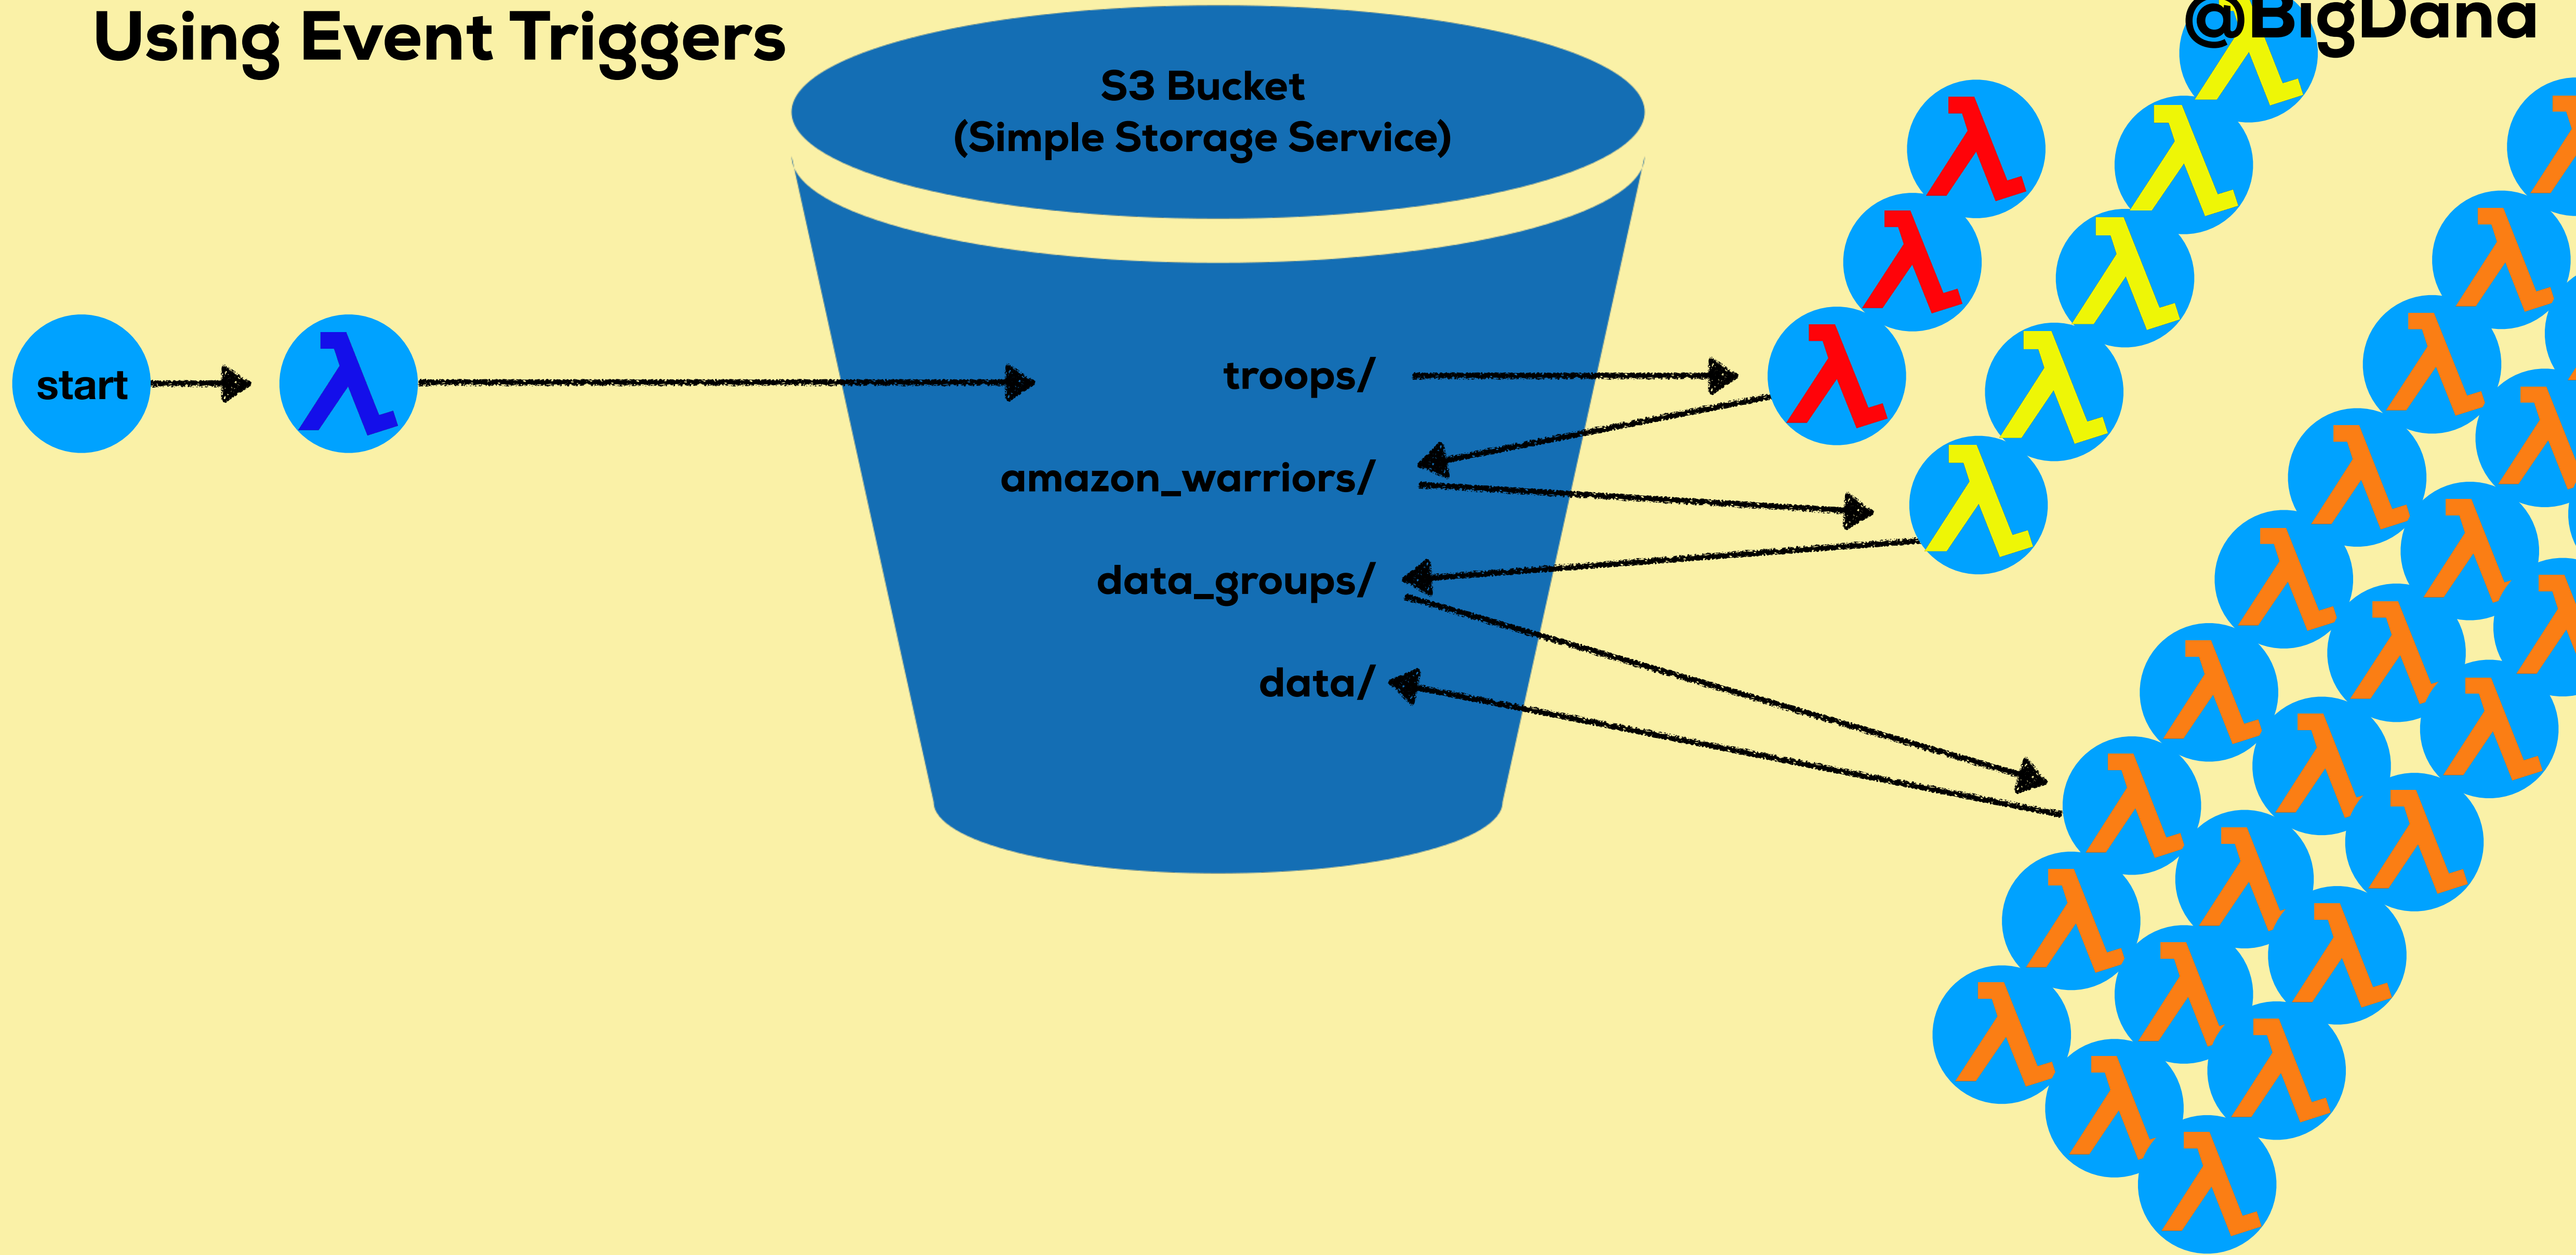
@BigDana





# Using Event Triggers

@BigDana







**Too Fast**



## Using Event Triggers with S3



**s3 is highly scalable and durable**

**Nested event triggers creates  
super fast multiprocessing**




**You can't control the rate  
the lambdas get spawned**

**It will only retry a lambda 2 times  
(can extend this manually)**

**You can lose in-flight tasks  
if there is a system outage**



WHY DID THEY BIND ME  
WITH SUCH SMALL CHAINS?  
IT'S AN INSULT!





**You can't run into a rate limit  
if you never scale.**





# Building a Data Pipeline


**Attempt 0: Python Multiprocessing on an EC2 Instance**

**Attempt 0.5: Python Multiprocessing on a Spot Instance**

**Attempt 1: Lambdas Orchestrated by Step Functions**

**Attempt 2: Using Event Triggers with S3**

**Attempt 3: Rate limiting with Queues and Cloudwatch triggered Lambdas**





# Building a Data Pipeline


**Attempt 0: Python Multiprocessing on an EC2 Instance**

**Attempt 0.5: Python Multiprocessing on a Spot Instance**

**Attempt 1: Lambdas Orchestrated by Step Functions**

**Attempt 2: Using Event Triggers with S3**


**Attempt 3: Rate limiting with Queues and Cloudwatch triggered Lambdas**






# SQS Queues



(Simple Queue Service)































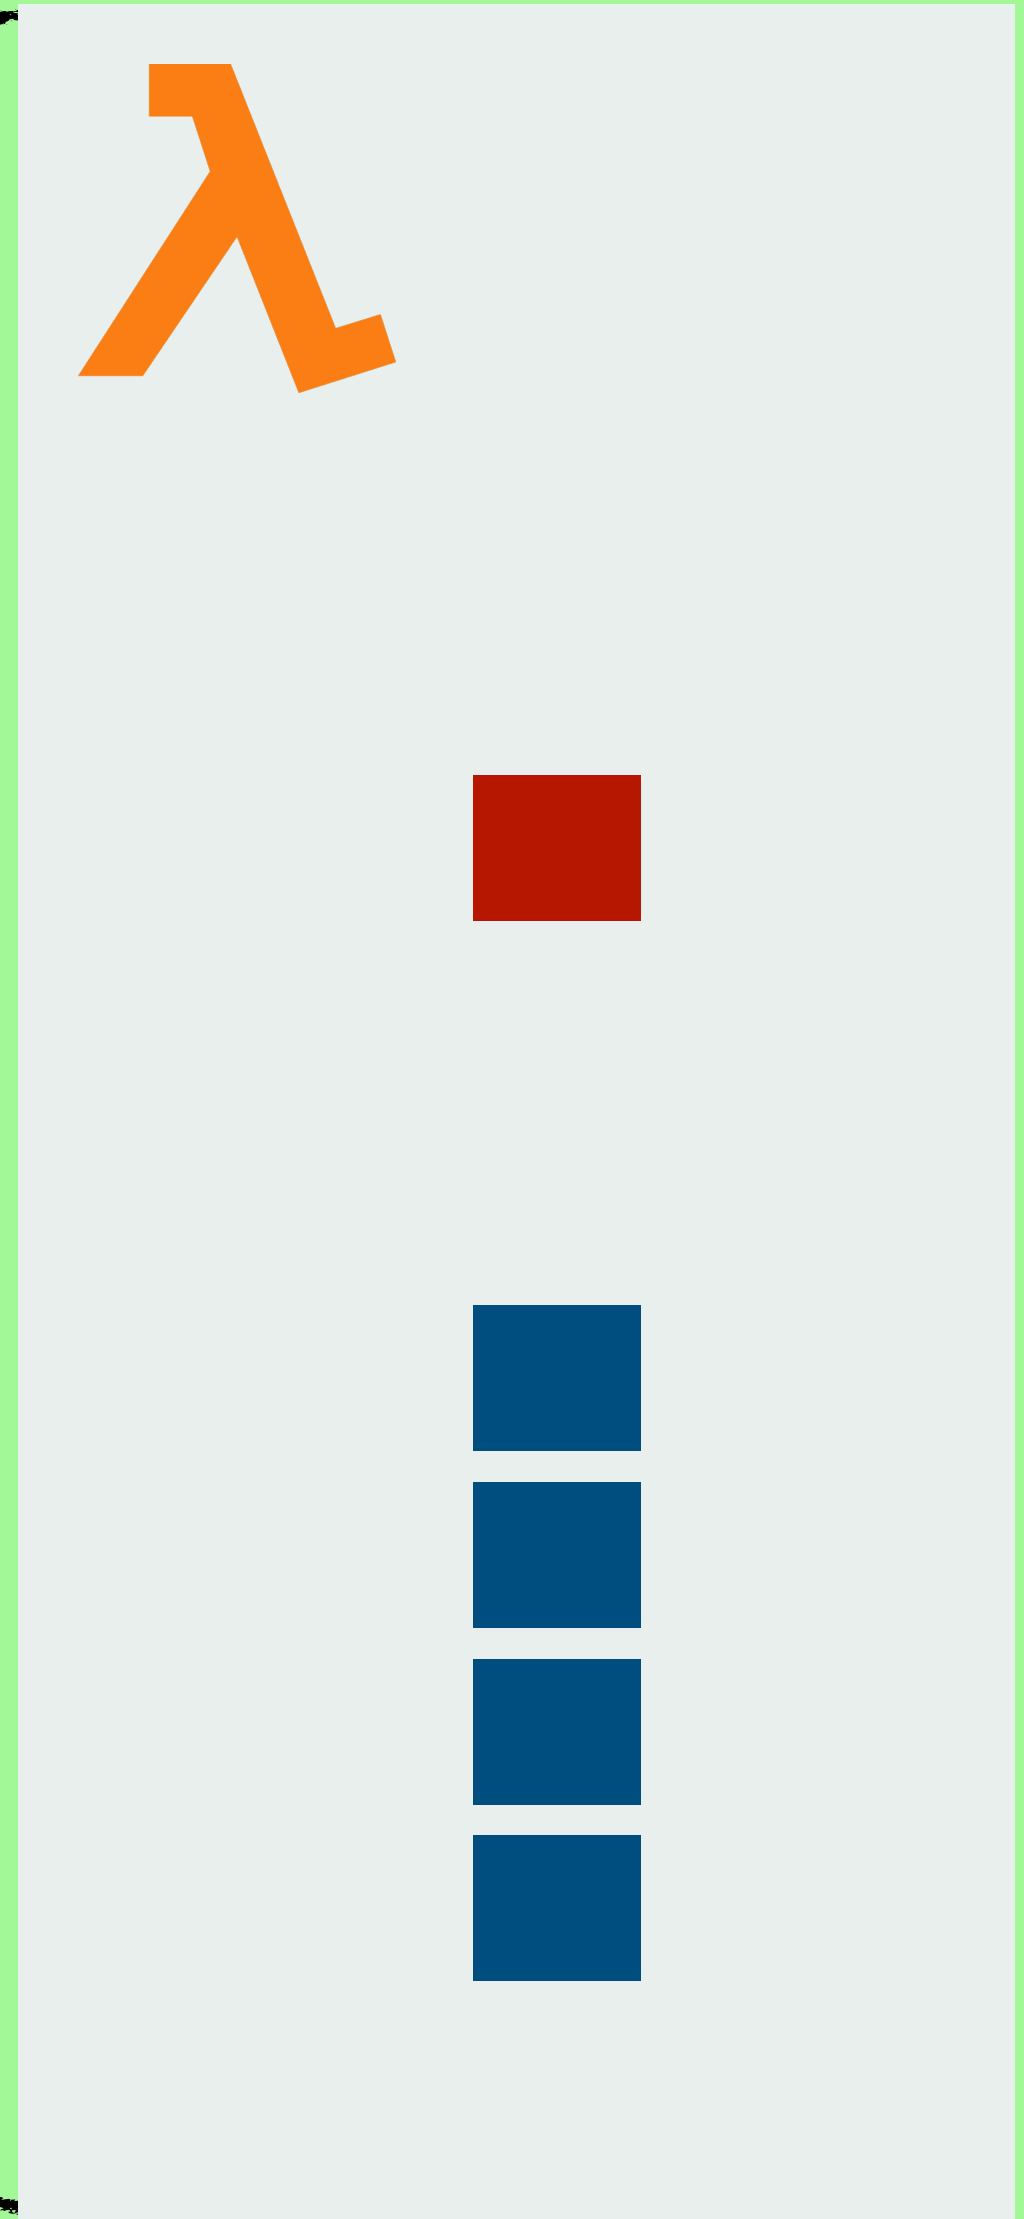





































@BigDana






@BigDana






@BigDana







A white rectangular box containing a vertical column of five dark blue squares on the left, an orange Lambda logo in the center, and a vertical column of six light blue squares on the right.

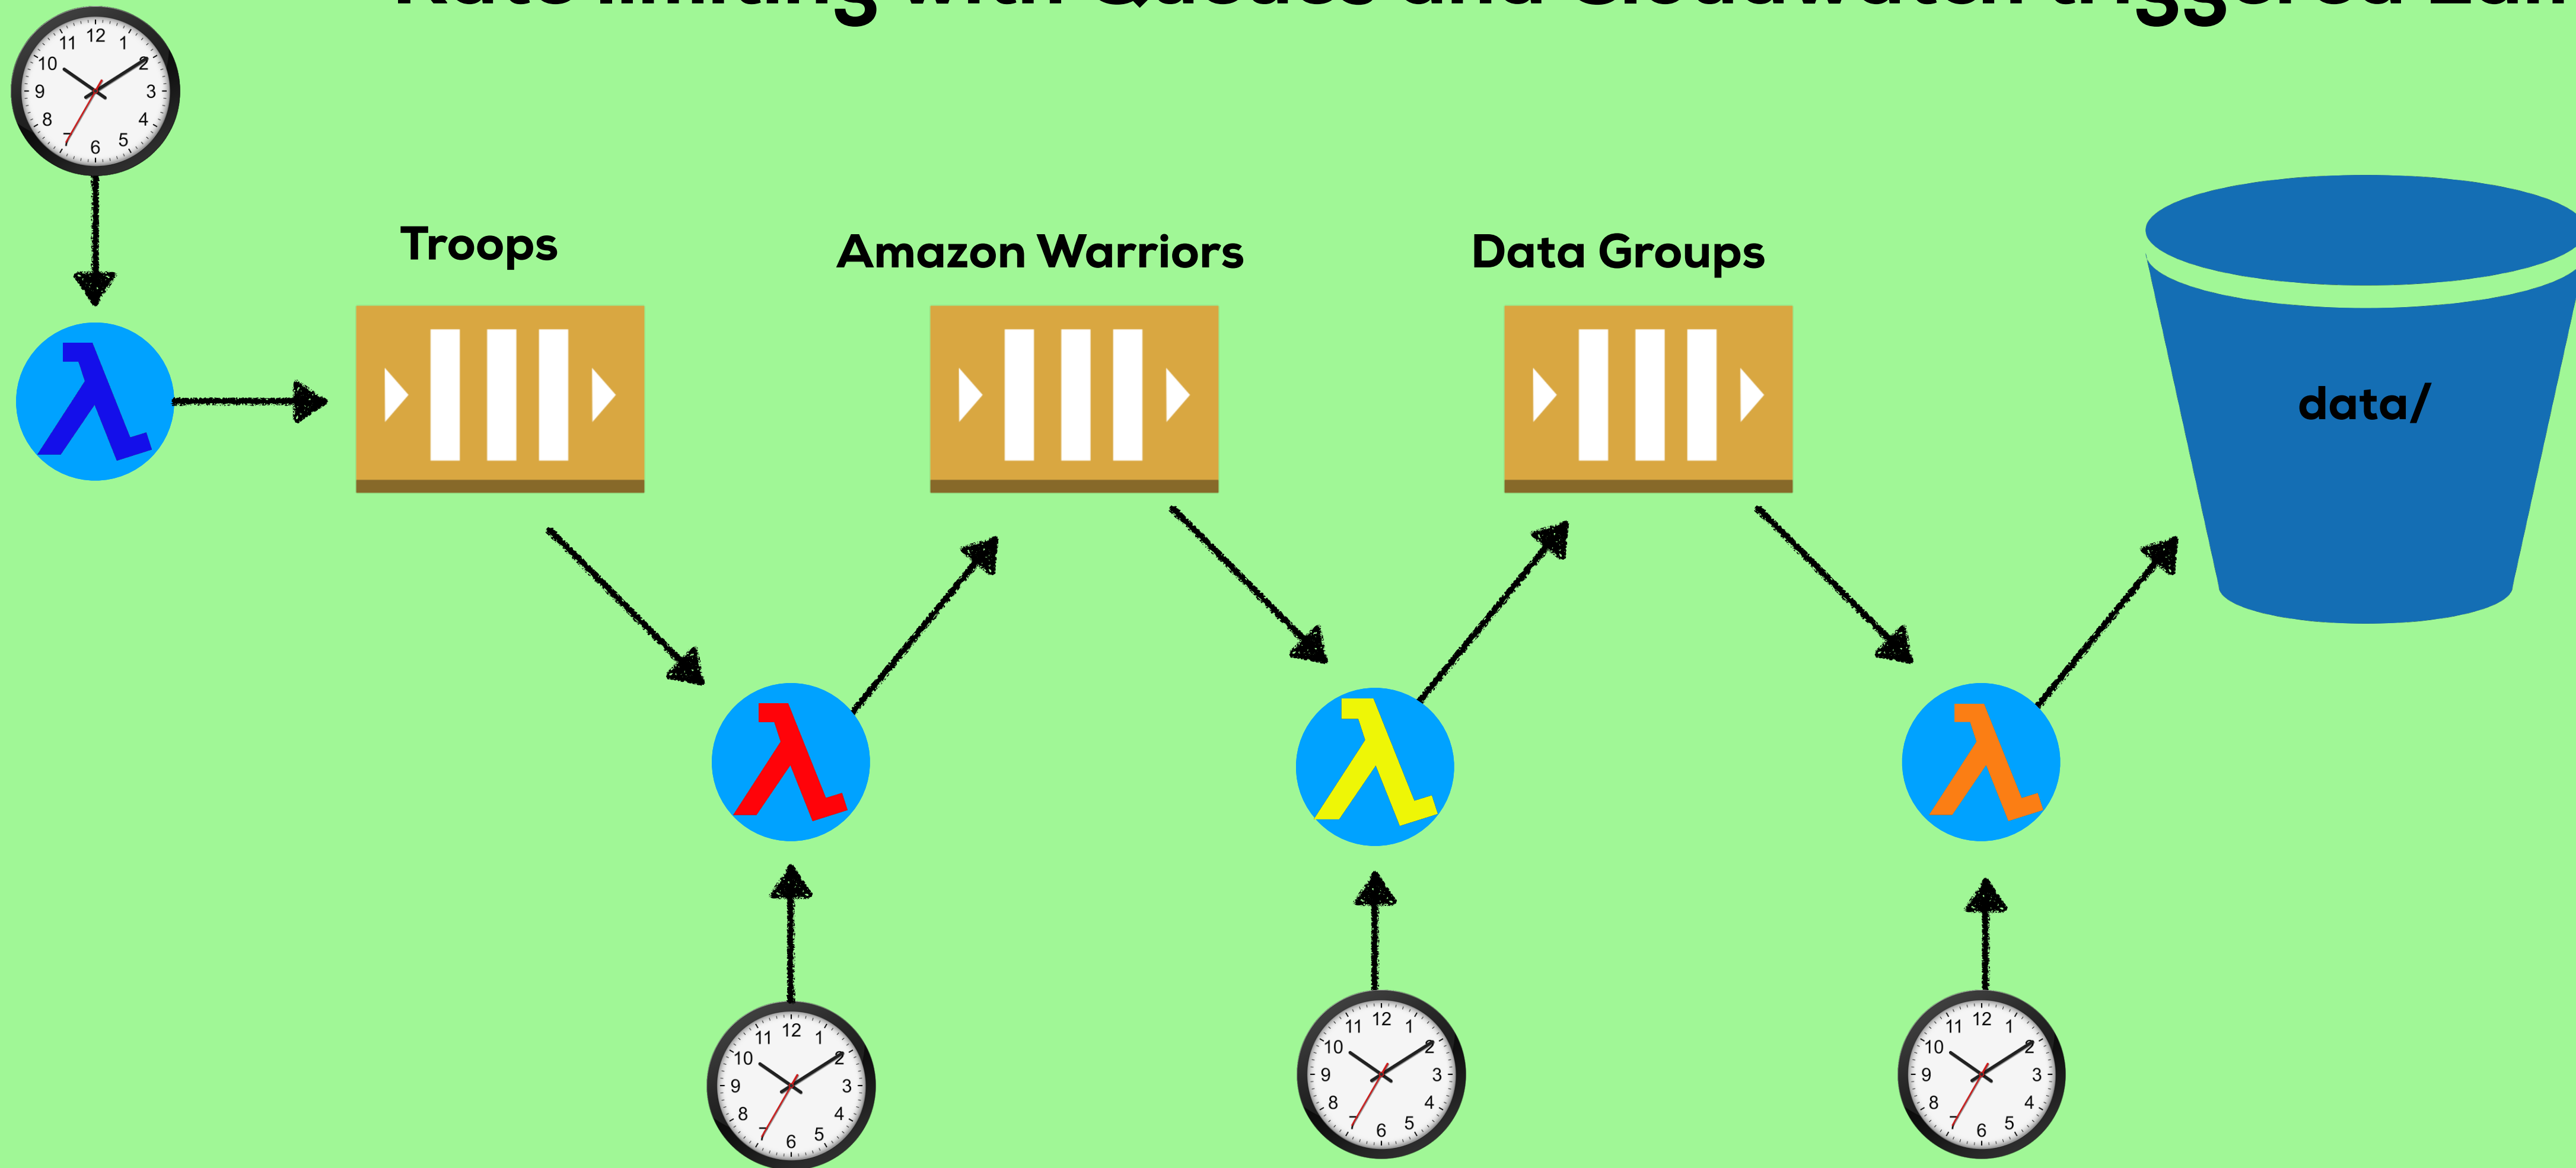


# Cloudwatch Rules



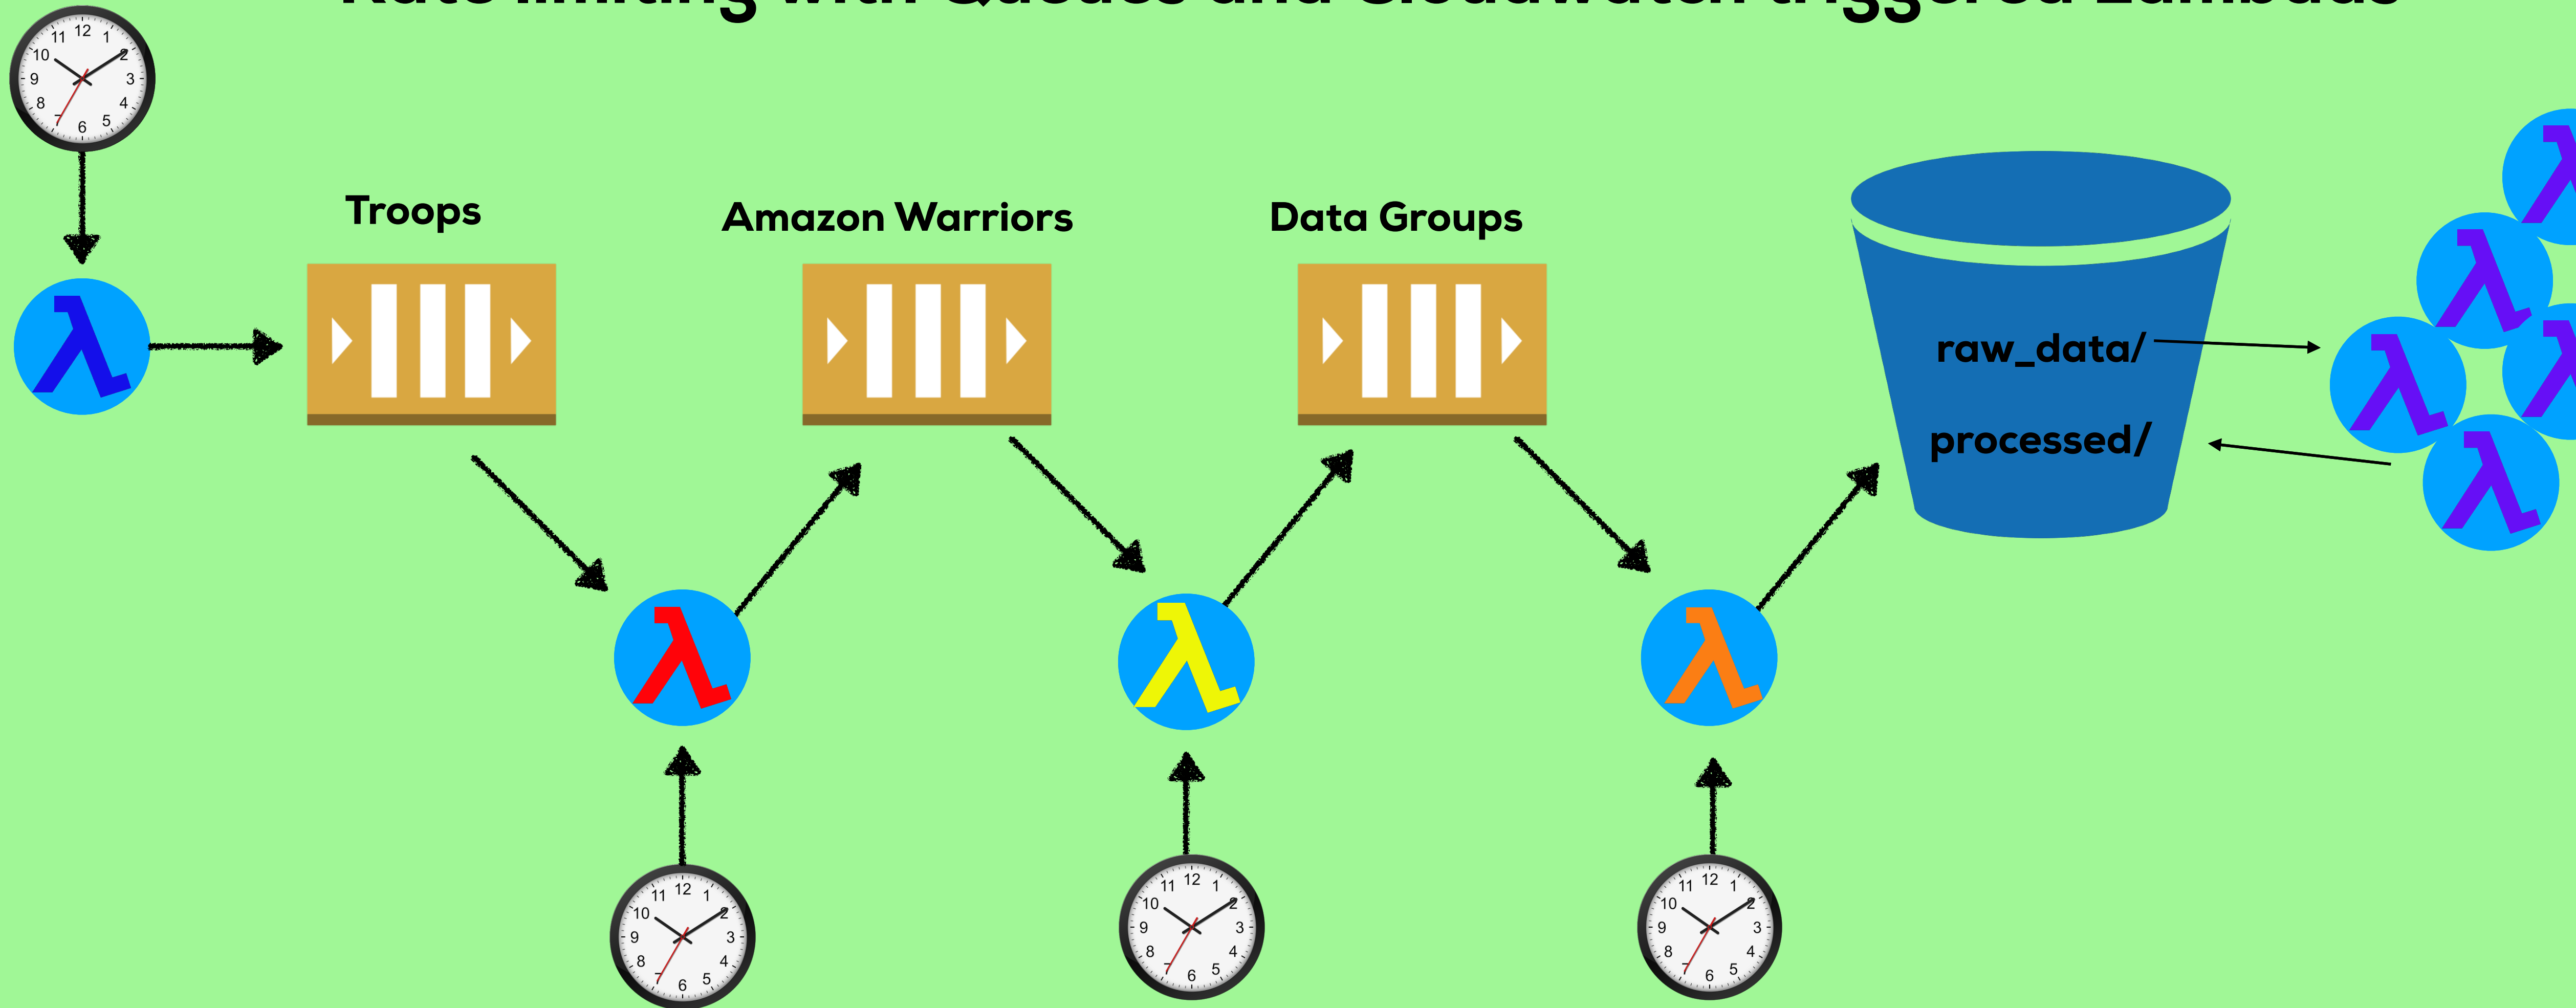


# Rate limiting with Queues and Cloudwatch triggered Lambdas






# Rate limiting with Queues and Cloudwatch triggered Lambdas



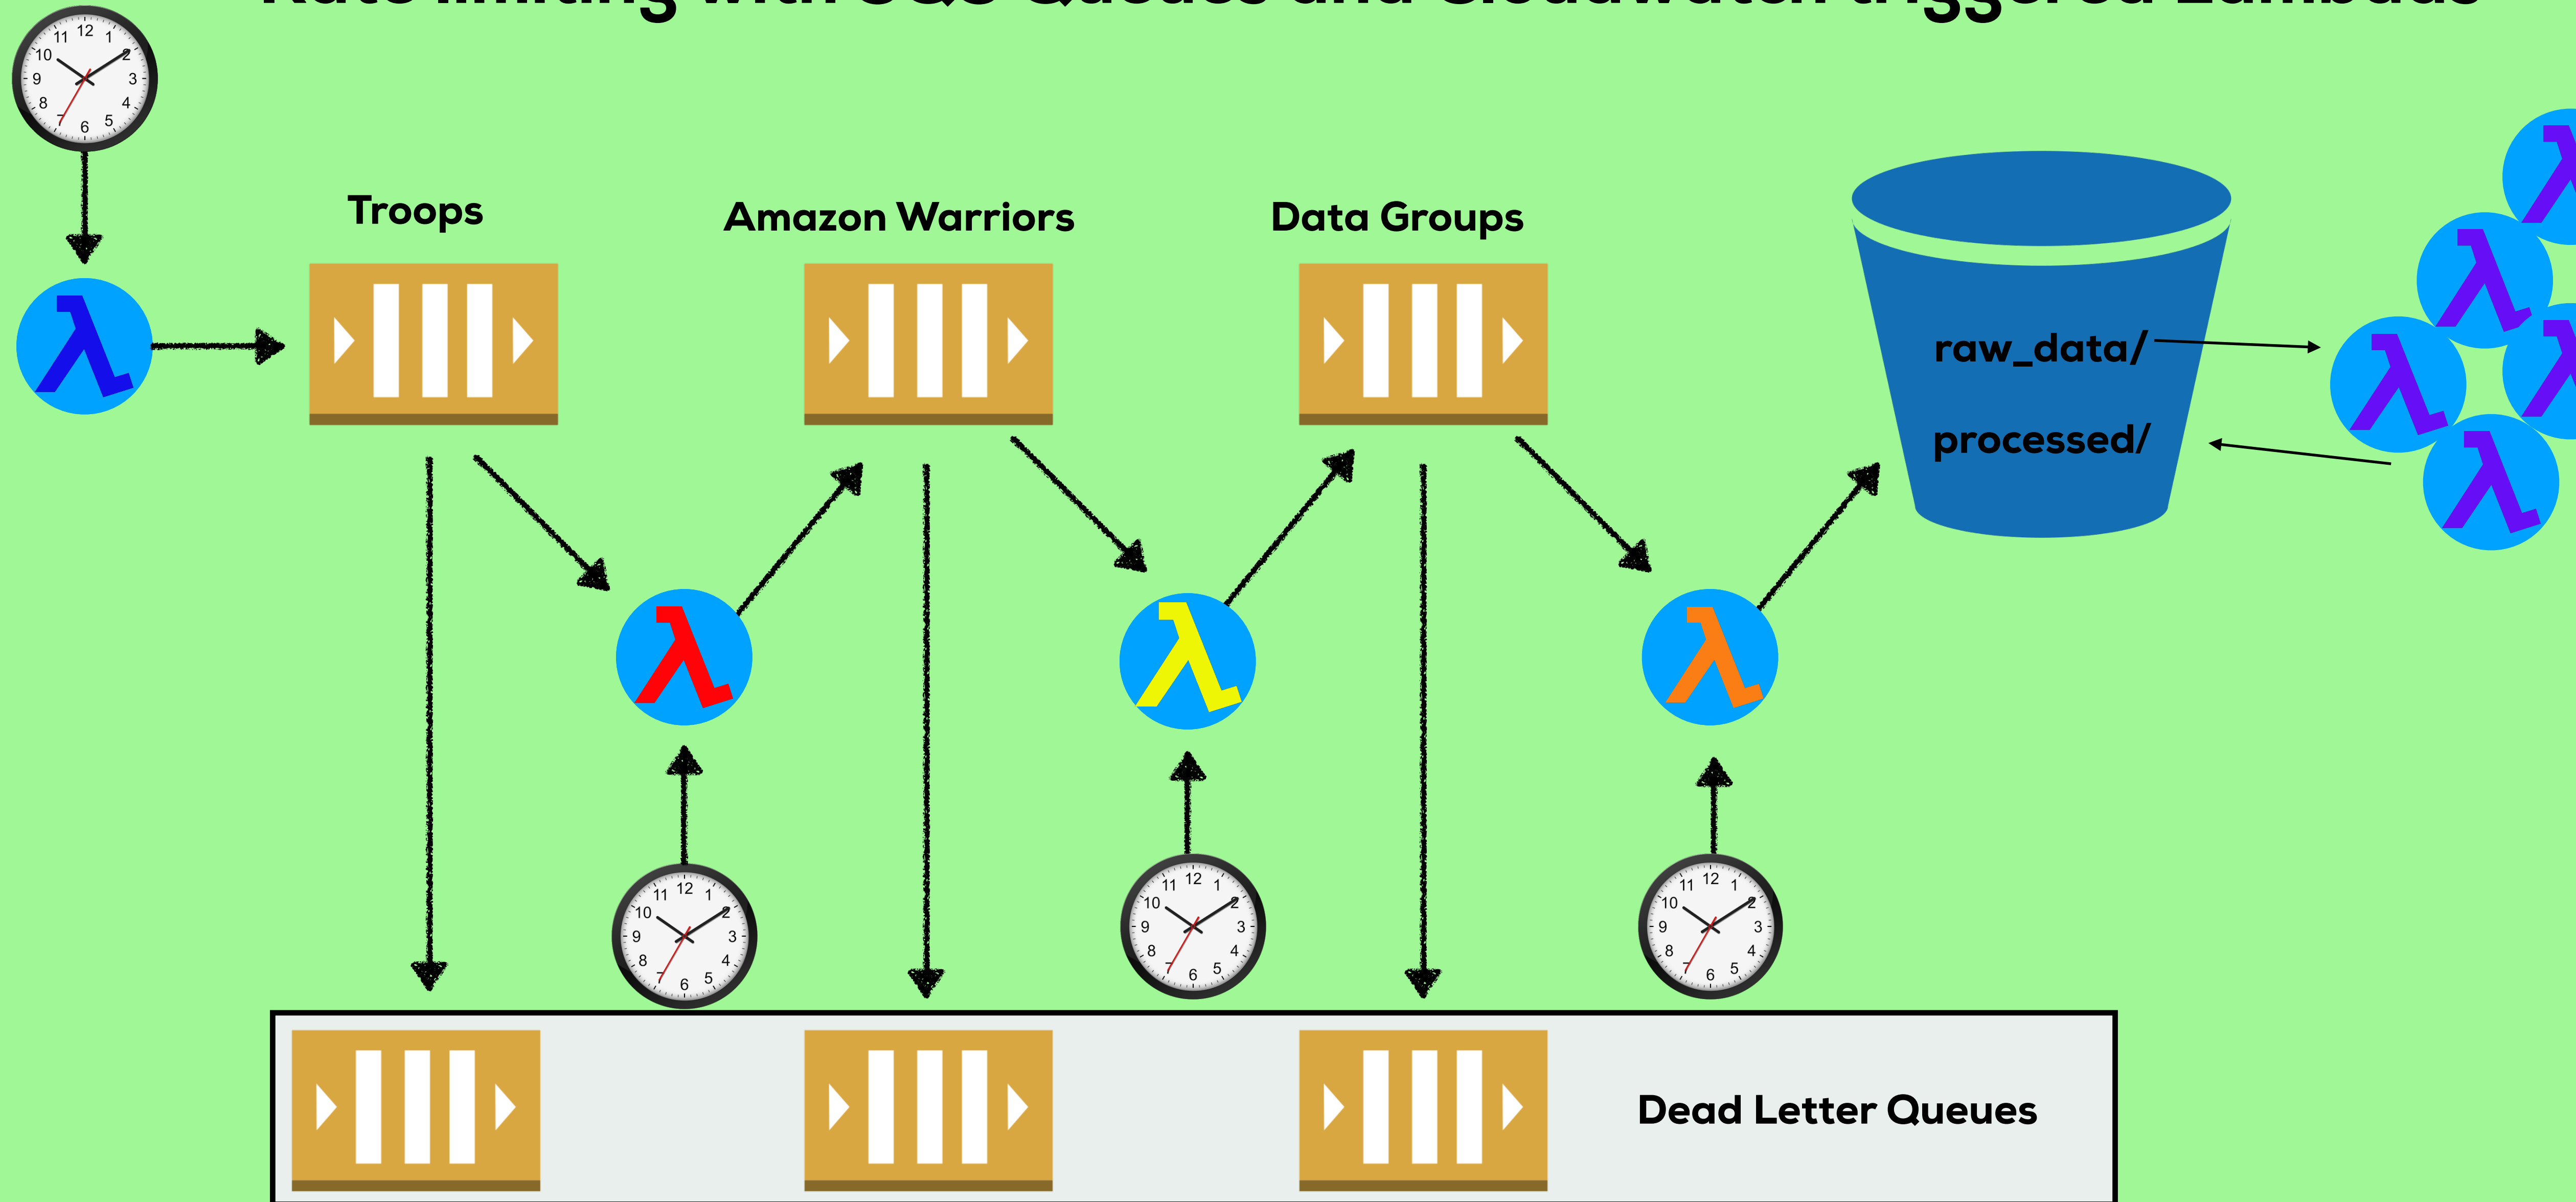


# Queue Depth Over Time






# Rate limiting with SQS Queues and Cloudwatch triggered Lambdas






# Queue Depth Over Time





# Queue Depth Over Time






## Rate limiting with Queues and Cloudwatch triggered Lambdas

```
18 while True:
19     try:
20         print('Reading up to 10 messages from device_queue...')
21         messages = device_queue.receive_messages(MaxNumberOfMessages=10)
22         got_messages = True
23     except botocore.exceptions.ClientError as e:
24         print('Error: failed to read from device_queue: {0}'.format(e))
25         got_messages = False
26
```



# Rate limiting with Queues and Cloudwatch triggered Lambdas






## Rate limiting with Queues and Cloudwatch triggered Lambdas

```
18 while True:
19     try:
20         print('Reading up to 10 messages from device_queue...')
21         messages = device_queue.receive_messages(MaxNumberOfMessages=10)
22         got_messages = True
23     except botocore.exceptions.ClientError as e:
24         print('Error: failed to read from device_queue: {0}'.format(e))
25         got_messages = False
26
27     if len(messages) == 0:
28         print("No messages on the queue, exiting lambda...")
29         break
```



# Rate limiting with Queues and Cloudwatch triggered Lambdas





## SQS Queues:



**highly scalable and resilient**

**can store failed tasks indefinitely**

**gives you control to rate limit the processing of the tasks in the queue**



**a small task can wait in line on a queue that already has many tasks**



## Lambdas triggered by Cloudwatch Rules



**allows for quite specific schedules**



**fastest schedule is  
every 1 minute**

**doesn't scale responsively to  
your queue depth**



IT TAKES REAL CHARACTER TO ADMIT ONE'S FAILURES -- AND NOT A LITTLE WISDOM TO TAKE YOUR PROFITS FROM DEFEAT. BUT REMEMBER, THIS MAN'S WORLD OF YOURS WILL NEVER BE WITHOUT RAIN AND SHEEPING UNTIL IT LEARNS L RIGHTS. KEEP ALL IN FRIEND GUN OF PERS

**IT TAKES REAL CHARACTER TO ADMIT ONES FAILURES – AND NOT A LITTLE WISDOM TO TAKE YOUR PROFITS FROM DEFEAT**

WONDERFUL, ANGEL!





# Takeaways

**Lambdas are not yet suitable for quick tasks if they require large dependencies**

**Step Functions does not yet provide dynamic fan out parallel processing**

**Event triggers don't yet allow you to rate limit**

**We can't yet trigger a Lambda based off an SQS Queue depth**



# Lessons Learned

**Expect Limits**

**Expect to Exceed those limits**

**Expect you will have to handle failures**



# Reflections for Managers

**Invest in training new hires**

**Decouple the learning process of your build deploy pipeline and building serverless architecture**



# Victoria Pierce

@thesecondshade

**June 5th 2017:**

**Joined SPS Commerce as a Site Reliability Engineer Intern**

**July 12th 2017:**

**Deployed her first lambda**

**As of November 15th 2017:**

**Maintains 6 serverless applications**





**@BigDana**

**Thank you!**

**Dana Engebretson  
Performance Engineer  
SPS Commerce**

**Questions?**

