

Adaptive Memory Scaling for Robust Airflow Pipelines

Cyrus Dukart, David Sacerdote &
Jason Bridgemohansingh





What we are talking about today

1. Who we are
 2. Our Data Challenge
 3. Our solution
- 



Who we are?

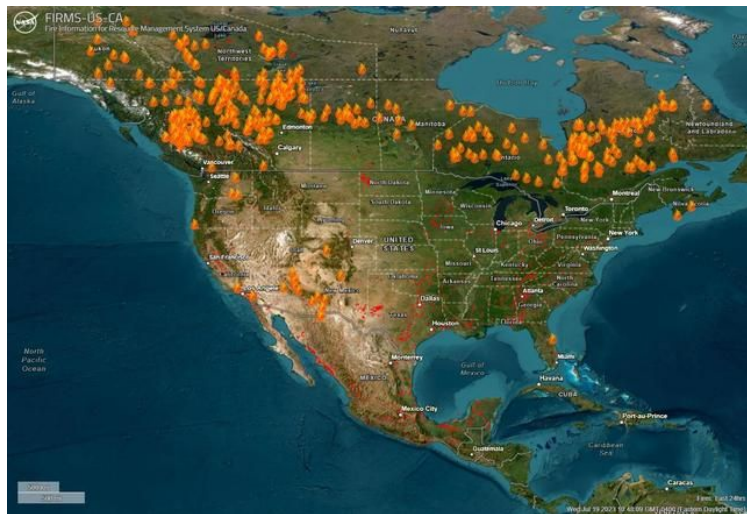
Vibrant Planet

The first common operating picture for
wildfire & ecosystem resilience



Climate Scientists Warn of a 'Global Wildfire Crisis'

Worsening heat and dryness could lead to a 50 percent rise in off-the-charts fires, according to a United Nations report.



Massive wildfires are a new threat to Chile. Here's why they're so deadly.



By Scott Dance

Updated February 5, 2024 at 7:56 p.m. EST | Published February 5, 2024 at 4:35 p.m. EST



Increasingly hotter fires

Catastrophic effects to soils and streams

Caldor Fire:
Lake Tahoe,
2021





Individual tree health to forest resilience


Using Synthetic Canopy Height Models to Segment individual Trees



Individual tree health to forest resilience

Using Synthetic Canopy Height Models to Segment individual Trees






No Treatment

Thinning + Prescribed Fire


Thinning Only

**The answer:
Speeding and scaling
what we know works**

What we use Airflow for




Created by Tobi Gull from Noon Project



Created by Hassanul Munir from Noon Project




Created by kareem1000 from Noon Project



Created by RinaBiflexia



Created by Shakesh Ch. from Noon Project




ML Models

Created by Heed Mujumbe from Noon Project




Transformations

Created by George Gull from Noon Project



Fire Treatment Recommendations



Data Layers

Disturbance Hazards

- Fire
- None

Strategic Areas, Resources, and Assets

Assets

- Structures
- Utilities

Safety

- Community Transmission Zone
- Protection
- Services

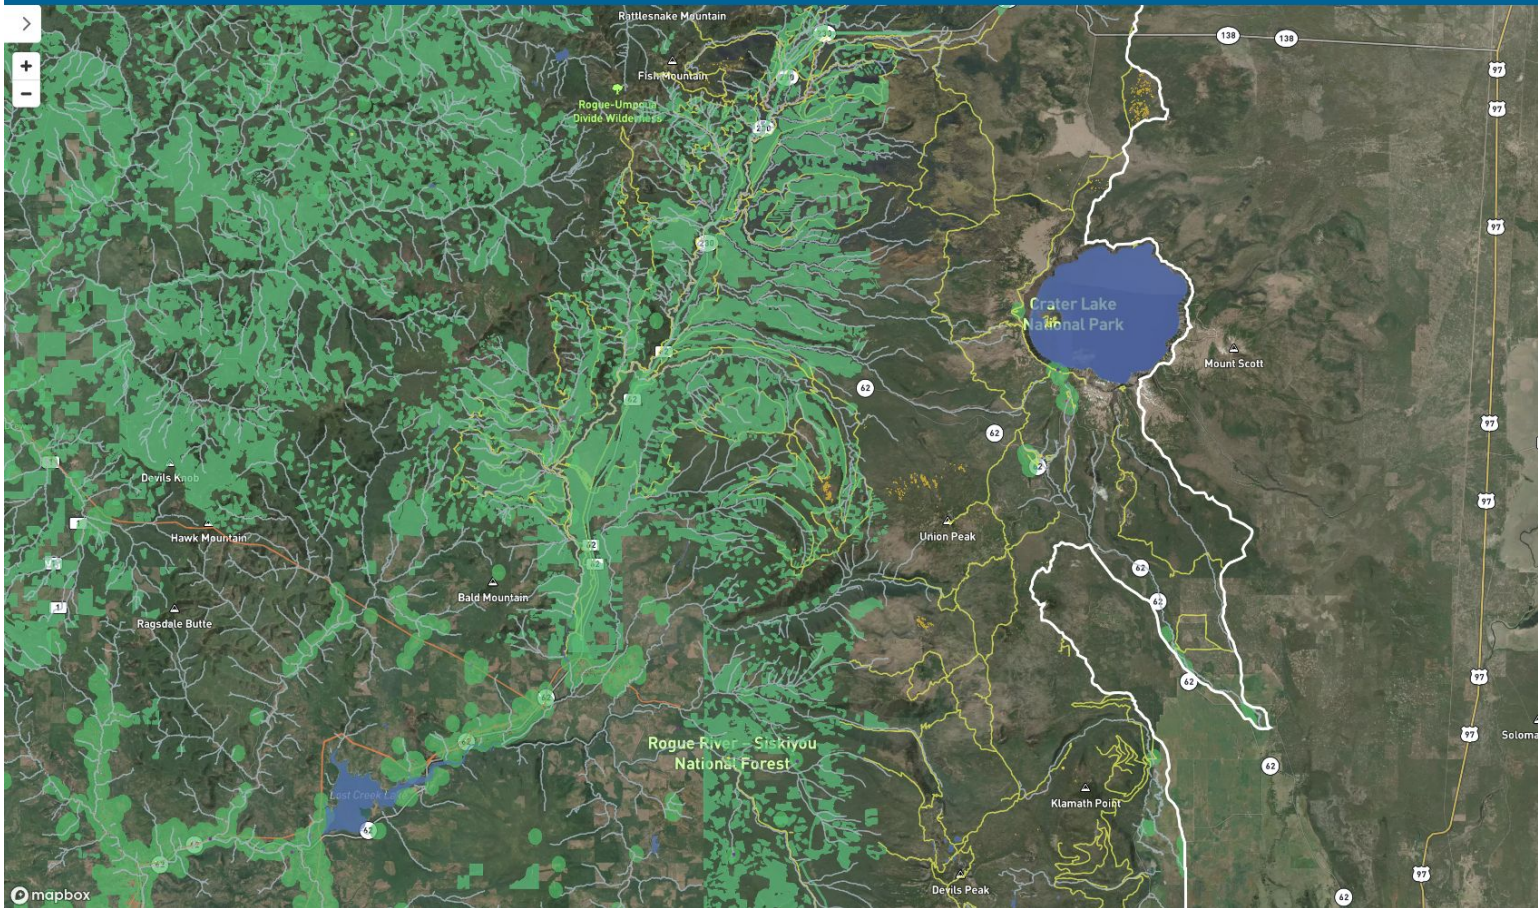
Recreation

- Recreation Areas
- Recreation Infrastructure

Biodiversity

- Animal Species/Communities
- Plant Species/Communities
- Whitebark Pine

Ecological Commodity



Data Layers

Strategic Areas, Resources, and Assets

Assets

- Structures
- Structures
- Utilities
- Energy Facilities
- Transmission or Distribution Lines
- Water Facilities

Safety

- Community Transmission Zone
- Protection
 - Wildland Urban Interface (Defense Zone)
- Services
 - Communication Infrastructure
 - Emergency Service Facilities

Recreation

- Recreation Areas
- Recreation Areas
- Recreation Infrastructure
- Trails

Biodiversity

- Animal Species/Communities


Source: Public Domain

© Mapbox © OpenStreetMap. Improve this map © Maxar

New Planning Area


Draw on map

Use the polygon drawing tool by clicking on the map to create consecutive vertices. Close and complete the polygon by clicking again on the first vertex.



Upload

SAVE PLANNING AREA CLEAR



Data Layers

Planning Area
My Planning Area

I'M FEELING RESILIENT

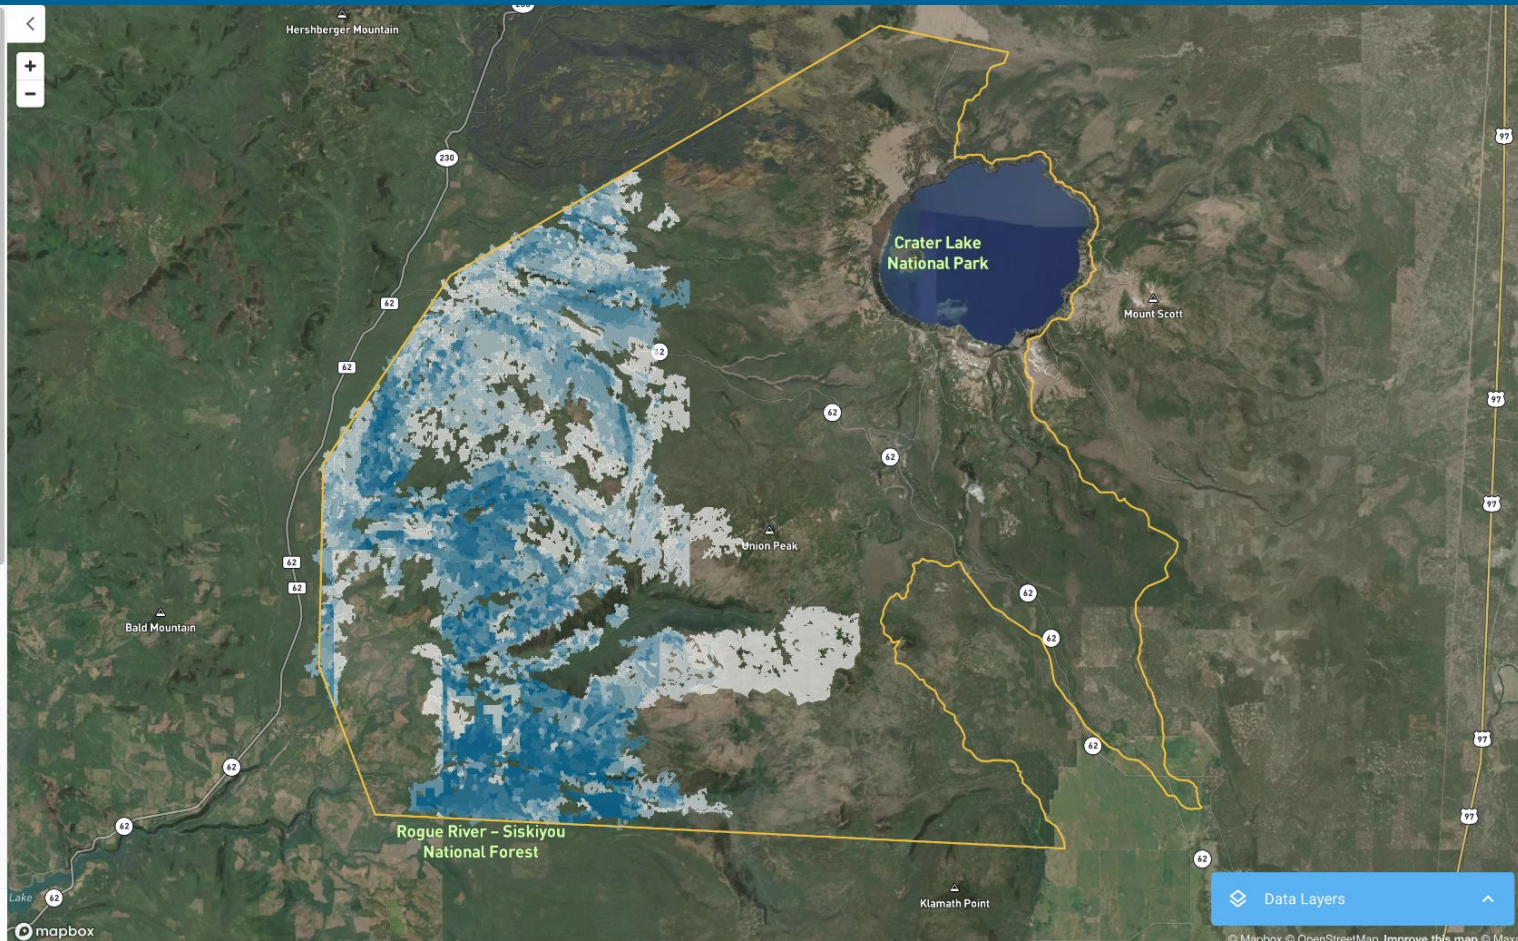
1. Set Priorities

Weigh Objectives **RESET**



- Assets
- Safety
- Recreation
- Biodiversity
- Ecological Commodity
- Carbon
- Water
- Science & Culture

CREATE SCENARIO



Data Layers

Planning Area
My Planning Area

Select Scenario
cyrus - 2024-05-10 00:20

- 1. Set Priorities
- 2. Set Constraints
- 3. Explore scenario

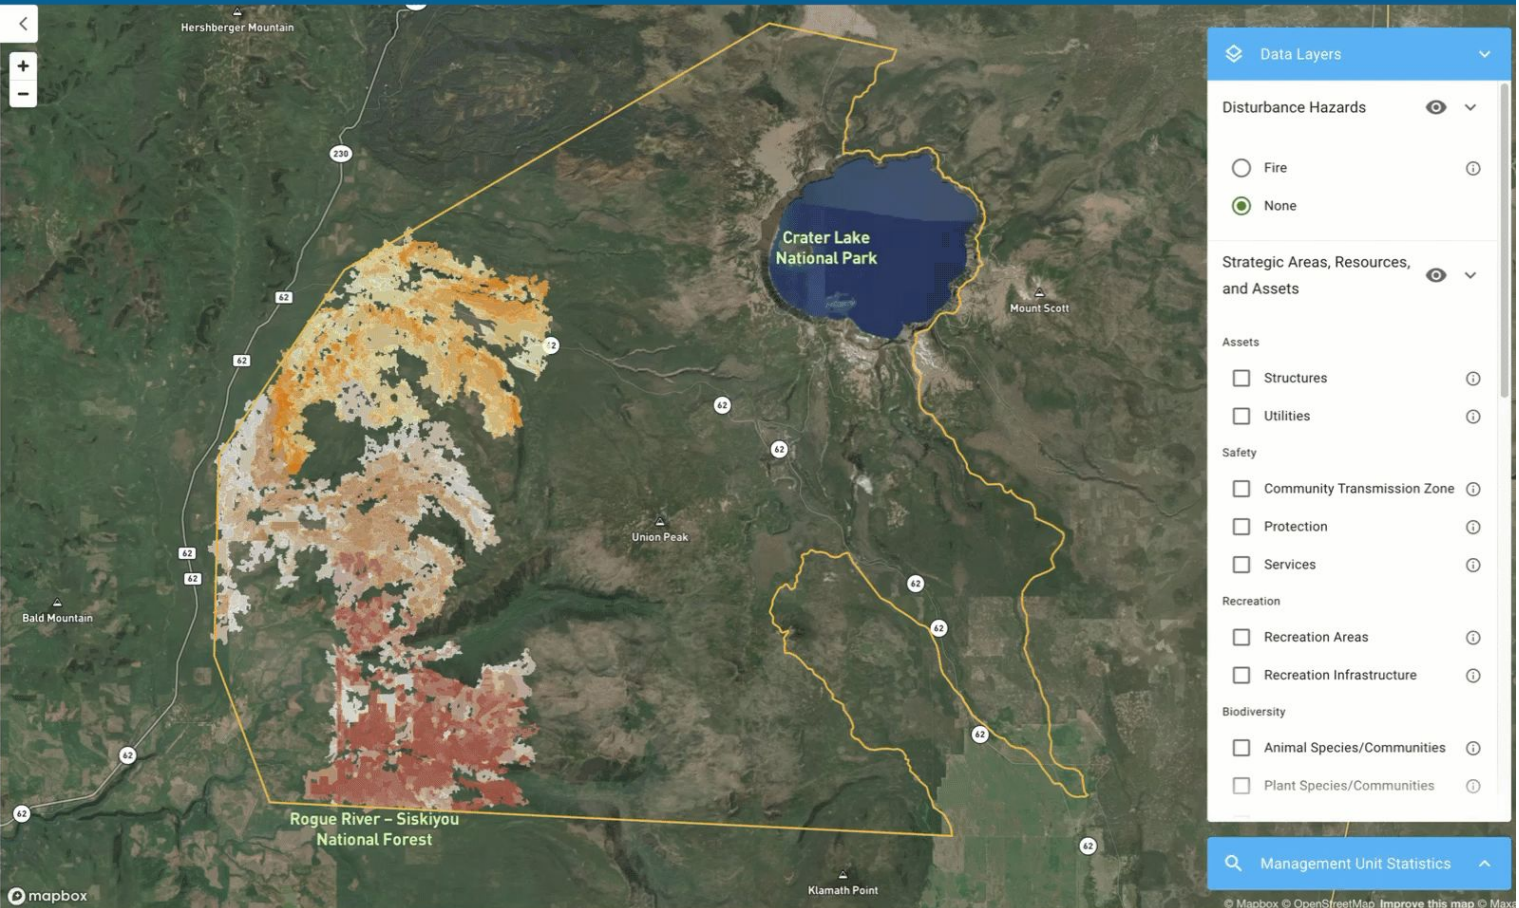
- Projects
- Project 1
 - Project 2
 - Project 3

Project statistics (3/3)

Total project area	45,003 Acres
Total coverage	21%
Estimated cost	\$72,507,000
Prioritized RROI	91%

ITERATE

CREATE COMPARISON



Data Layers

Disturbance Hazards

- Fire
- None

Strategic Areas, Resources, and Assets

Assets

- Structures
- Utilities

Safety

- Community Transmission Zone
- Protection
- Services

Recreation

- Recreation Areas
- Recreation Infrastructure

Biodiversity

- Animal Species/Communities
- Plant Species/Communities

Management Unit Statistics

The image features a green-tinted landscape. In the foreground, there is a dense forest of evergreen trees. In the background, a mountain range is visible, with a large, billowing plume of white smoke or ash rising from one of the peaks, dominating the upper half of the frame. The overall scene is dramatic and atmospheric.

Our Data Challenge?

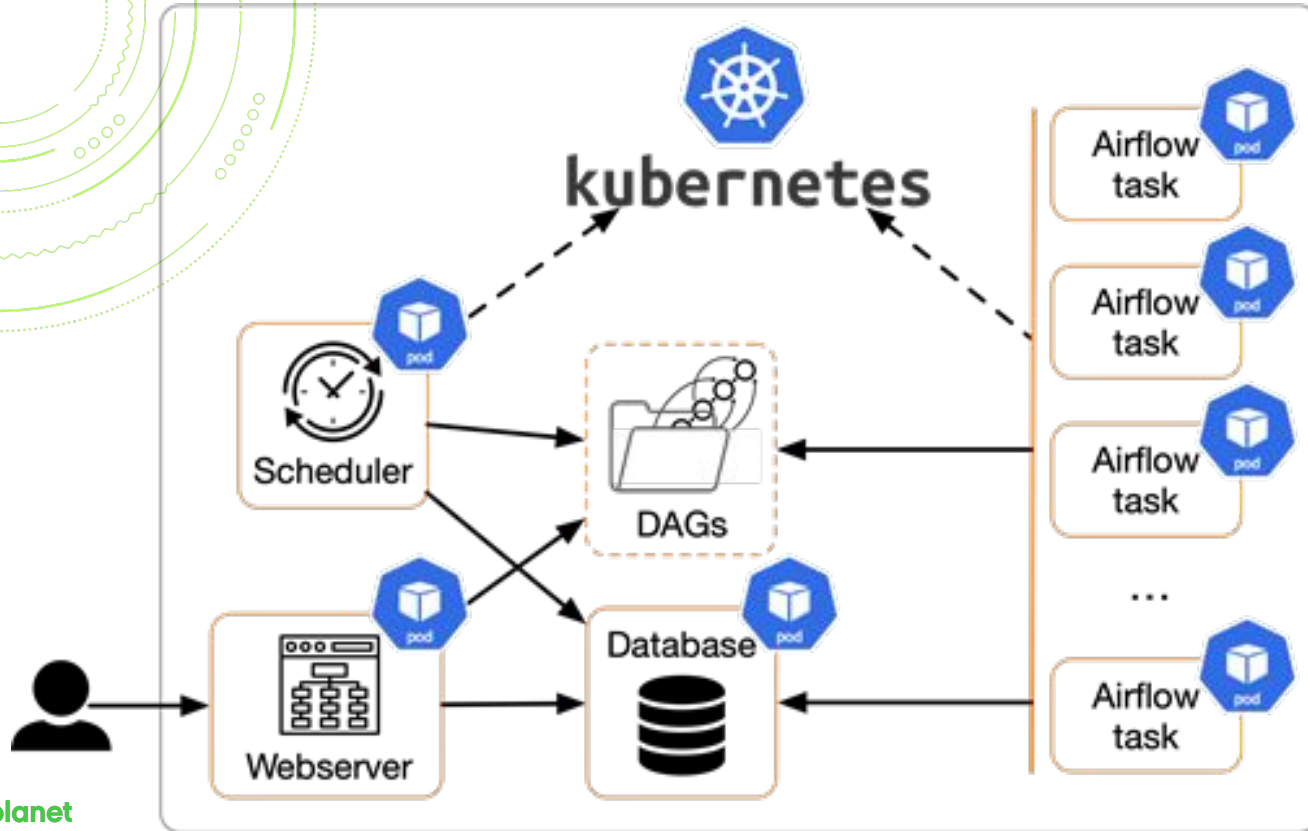
Airflow Executor Options



VS

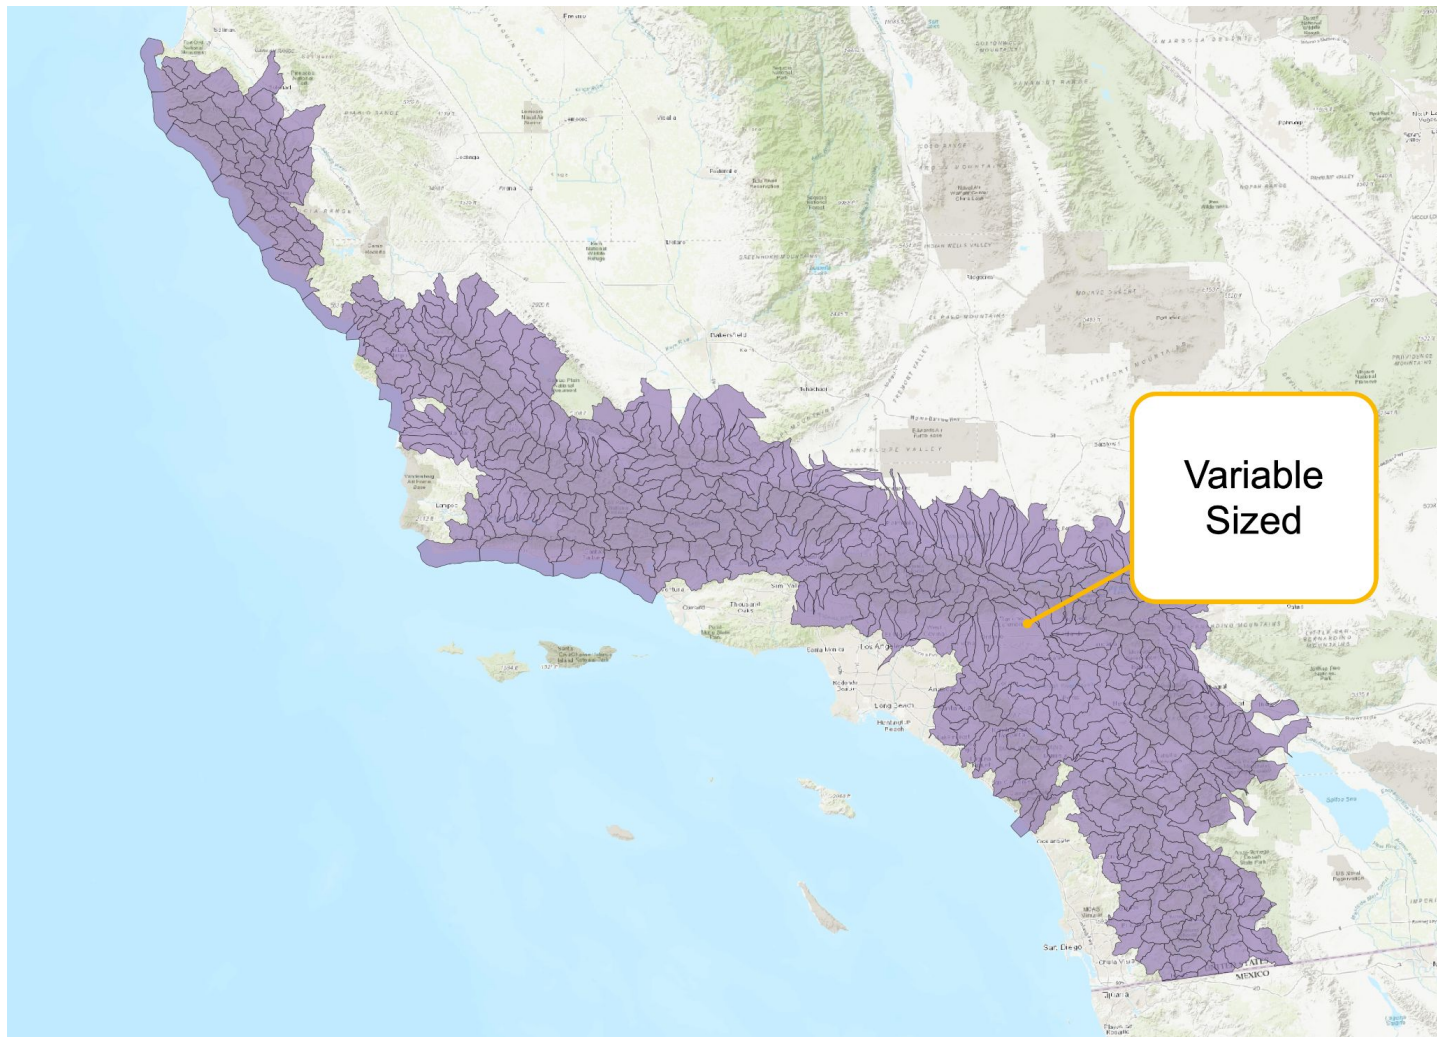


How Airflow allocates resources



Our Partitioning Scheme

Variable
Memory
Reqs

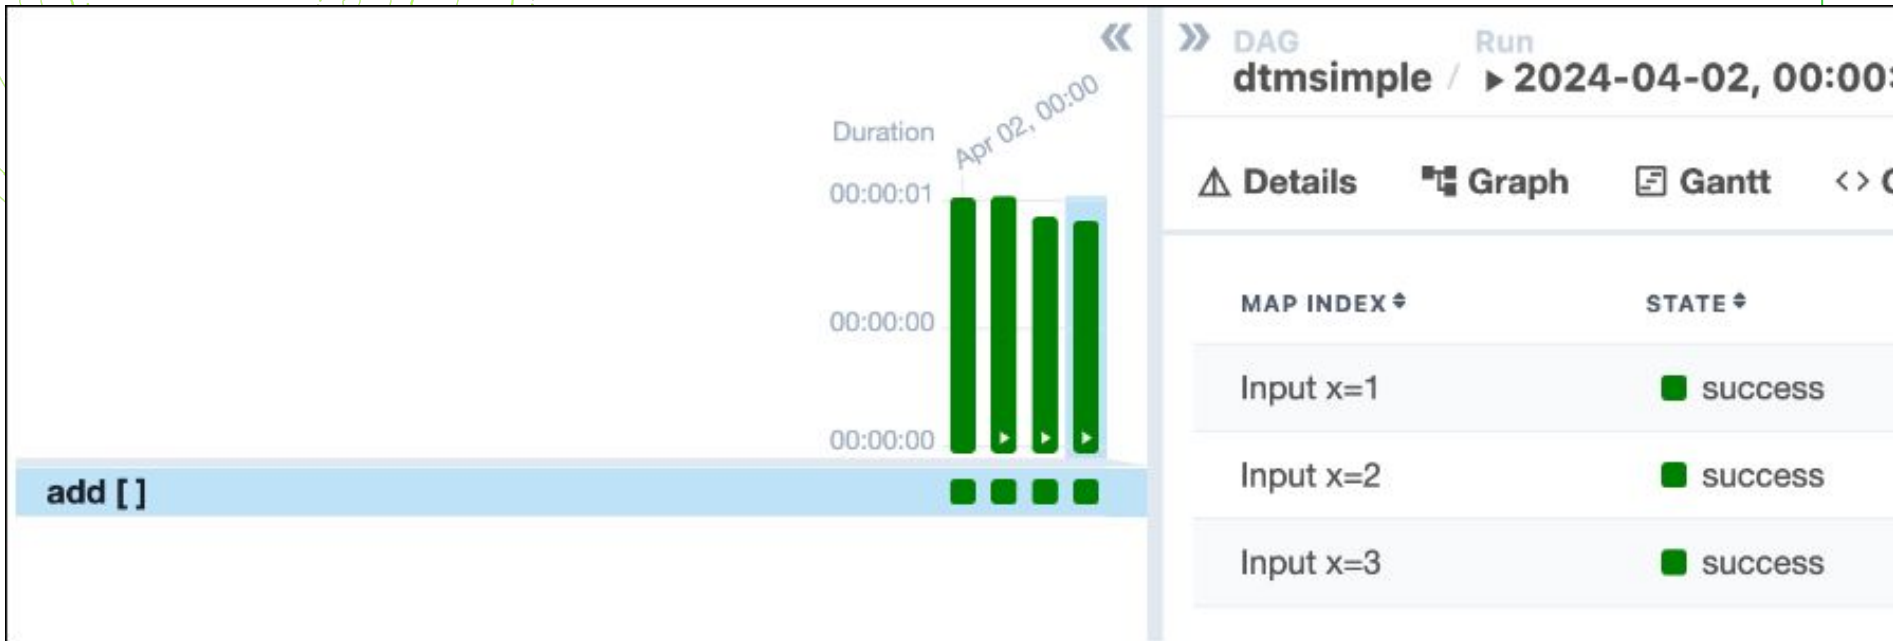


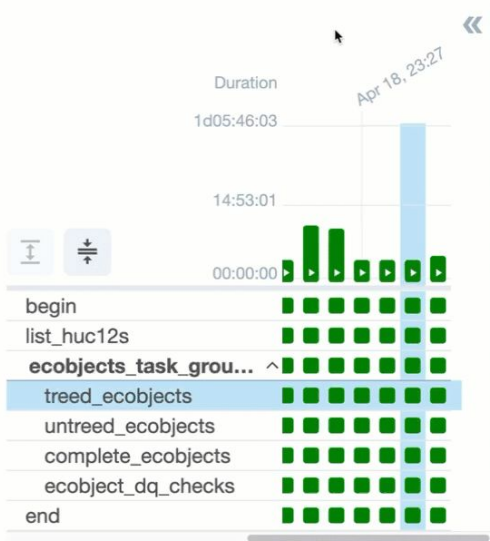
Dynamic Task Mapping

```
@task()
def add(x: int):
    logging.info(f"Running for {x}")

expansion_example = add.expand(x=[1, 2, 3])
```

Dynamic Task Mapping





» DAG Run Task
ecobjects / ▶ 2024-04-29_trinity_naip2022_w2w_taos_down / treed_ecobjects []

△ Details Graph Gantt <> Code [] Mapped Tasks



The Challenge

Default is
Same Memory
per
TaskInstance

The screenshot shows a DAG run interface for 'ecobjects' on '2024-04-29_trinity_naip2022_w2w_taos_down'. The 'Mapped Tasks' tab is active, displaying a table with columns for INDEX, STATE, and DURATION. The 'begin' task is highlighted in blue. A task list on the left shows 'ecobjects_task_g' with a grid of green dots representing task instances. Three yellow arrows point from the callout box to the 'success' state of task instances 3, 4, and 5 in the table.

INDEX	STATE	DURATION
1	success	00:22:16
2	success	00:42:16
3	success	00:24:16
4	success	00:17:31
5	success	00:25:31
6	success	00:47:31
7	success	00:15:46
8	success	00:19:31
9	success	00:21:46
10	success	00:15:16

What's in a k8s pod spec

- Contains a full description of the environment and resource limits

```
resources:
```

```
limits:
```

```
  cpu: '12'
```

```
  memory: 24Gi
```

```
requests:
```

```
  cpu: '12'
```

```
  memory: 24Gi
```

Example Task w Executor Config Override

```
from kubernetes.client import models as k8s
```

```
@task(
```

```
    executor_config={
```

```
        "pod_override": k8s.V1Pod(...)
```

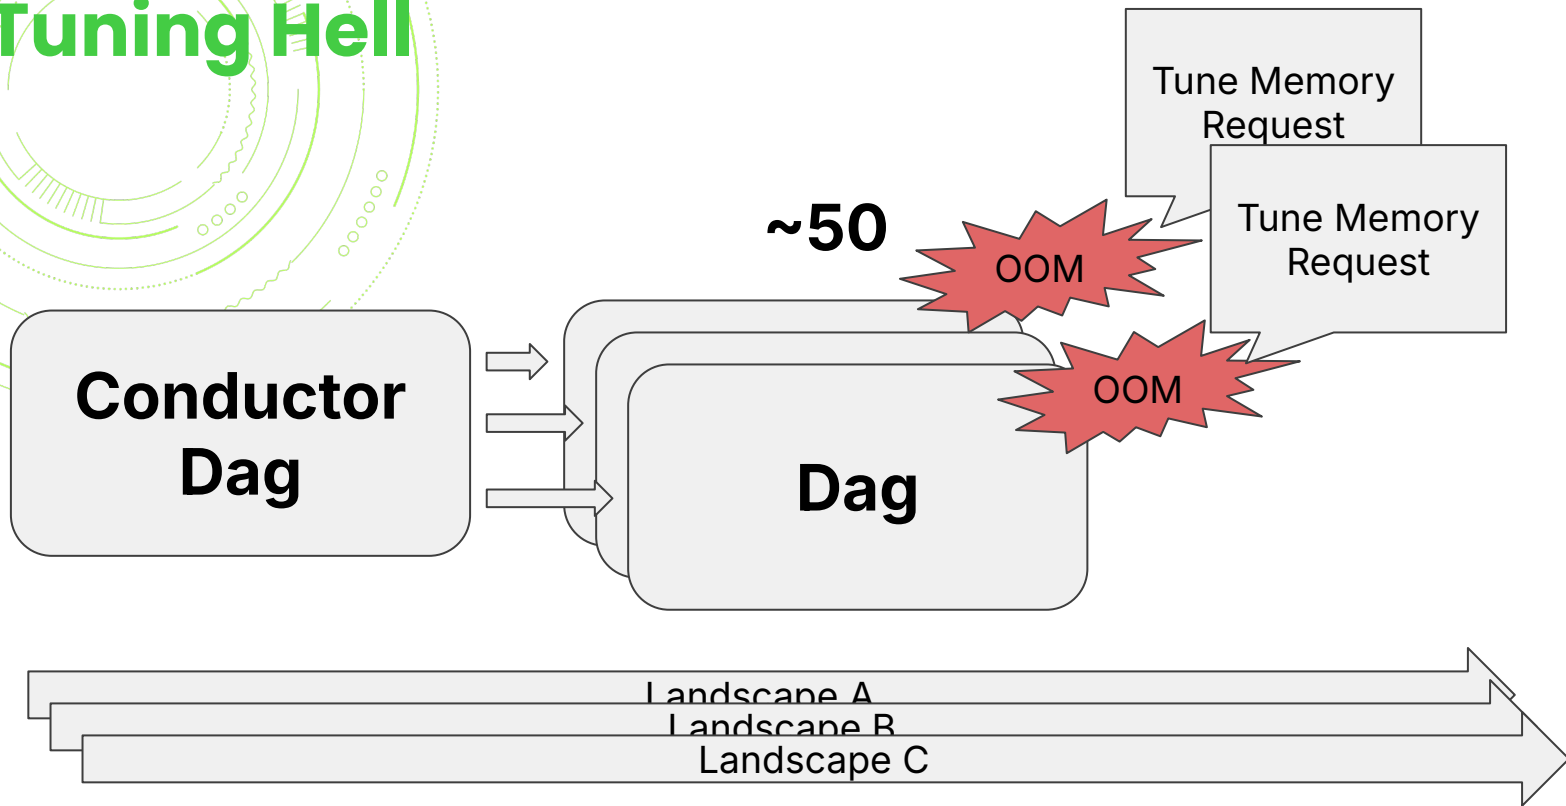
```
    }
```

```
)
```

```
def do_something():
```

```
    print("my resources were customized!")
```

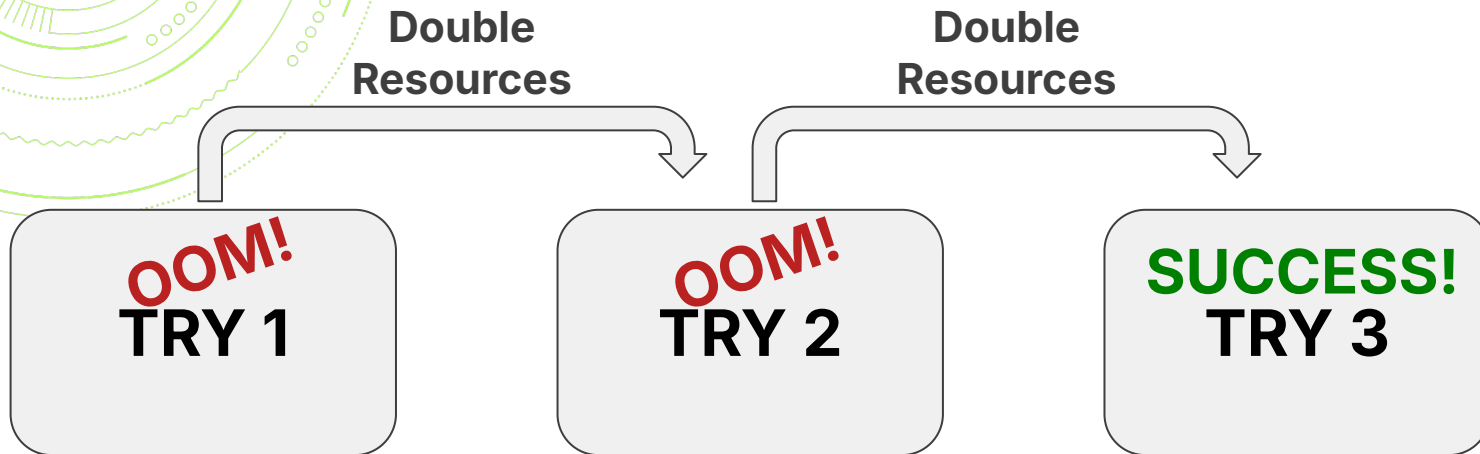
Tuning Hell



The image features a monochromatic green color scheme. In the foreground, there is a dense forest of evergreen trees. Behind the forest, a range of low mountains or hills is visible. The sky is filled with large, billowing clouds of white smoke or ash, which are rendered in a lighter shade of green. The overall scene suggests a natural event, such as a volcanic eruption, with a focus on environmental impact or a solution related to such events.

Our Solution

Make Self Healing Pipelines



Solution: Mutate Pod Executor Config

```
from kubernetes.client import models as k8s

def update_task_instance_pod_override(
    ti: TaskInstance,
    executor_config: Dict[str, V1Pod]
) -> None:
    with create_session() as session:
        ti.executor_config = executor_config
        session.add(ti)
    ti.refresh_from_db()
```

Solution: Example Executor Config

```
executor_config = {
    "pod_override": k8s.V1Pod(
        spec=k8s.V1PodSpec(
            containers=[
                k8s.V1Container(
                    name="base",
                    resources=k8s.V1ResourceRequirements(
                        requests=K8sResources(cpu="4", memory="8Gi"),
                        limits=K8sResources(cpu="4", memory="8Gi"),
                    ),
                ),
            ],
        ),
    ),
}
```


Solution: Example Dag

```
@dag()
def example_dag():

    @task(
        retries=3,
        on_retry_callback=double_memory()
    )
    def failing_task():
        raise ValueError("Failing to Memory Ramp")
```



vibrant



A landscape photograph featuring a dense forest of evergreen trees in the foreground. In the middle ground, a mountain range is visible, with a prominent peak on the left. A massive, billowing plume of white smoke or steam rises from the mountain range, filling much of the sky. The sky is a clear, bright blue. The overall scene suggests a volcanic eruption or a large-scale fire.

But wait there's more

What about initial Memory Allocation




The Problem

Default is
Same Memory
per
TaskInstance

The screenshot shows a DAG run interface for 'ecobjects' on '2024-04-29_trinity_naip2022_w2w_taos_down'. The 'Mapped Tasks' tab is active, displaying a table of task instances. A callout box highlights that the default is 'Same Memory per TaskInstance', with arrows pointing to the 'success' status of several task instances in the table.

INDEX	STATE	DURATION
1	success	00:22:16
2	success	00:42:16
3	success	00:24:16
4	success	00:17:31
5	success	00:25:31
6	success	00:47:31
7	success	00:41:01
8	success	00:15:46
9	success	00:19:31
10	success	00:21:46
11	success	00:15:16

Set Initial Resources per TI



Implementation: first try fail

Because we can't set memory allocation before we run, the first attempt to run any task is automatically failed as a means of setting the correct amount of memory

```
[2024-02-23, 23:41:01 UTC] {vp_python.py:202} INFO - Running with the following memory settings:
[2024-02-23, 23:41:01 UTC] {vp_python.py:203} INFO -     Total memory from Executor Config: 9.0 GiB
[2024-02-23, 23:41:01 UTC] {vp_python.py:206} INFO -     Total memory available: 9.0 GiB
[2024-02-23, 23:41:01 UTC] {vp_python.py:209} INFO -     Total memory limit on process: 8.1 GiB
[2024-02-23, 23:41:01 UTC] {airflow_helpers.py:200} INFO - Recommending default memory for hucl2_id =
[2024-02-23, 23:41:01 UTC] {airflow_helpers.py:435} INFO - Setting Memory for Executor Config to: 4.0
GiB
[2024-02-23, 23:41:02 UTC] {airflow_helpers.py:498} INFO - #####
[2024-02-23, 23:41:02 UTC] {airflow_helpers.py:499} INFO - # EXPECTED FIRST TRY FAIL #
[2024-02-23, 23:41:02 UTC] {airflow_helpers.py:500} INFO - #####
```

How we make an initial estimate

Task Run before?	Data (HUCs) Run Before?	Action
✓	✓	Allocate prior memory consumption + safety margin
✓	✗	Run a linear regression to estimate memory needs; add safety margin
✗	✗	Make a guess that works for a lot of tasks

Custom Decorator

We implement this using a custom decorator, `@task.vp`

```
@task.vp
def save_the_forest():
    from utils.forest_save import forest_saver_2000
    forest_saver_2000()
```




Custom Decorator Goodies

1. Stats Logging
2. Initial Memory Recommendation
3. Memory Ramp on Retry
4. OOM Detection



Q&A



Appendix



vibrant planet

Self Healing Pipelines (OOM Resilience)





OOM detection and retry

OOM frequently results in a process death with no direct notification

So we wrap our tasks in a parent process which detects OOM situations and kicks off a retry with double the amount of memory

This parent process also collects memory and CPU utilization statistics during process execution and logs them to a database upon completion.

Small job optimization

Jobs with low memory needs (eg: smaller than the memory-estimation stub) get launched inside the same pod and memory footprint of the memory-estimation stub, rather than needing a failure and retry. This reduces both both the time and expense of running them.

Short-running jobs are often batched together, avoiding the overhead of launching new containers.