


Airflow @ Monzo

*Evolving our Data Platform
as the Bank Scales*

Jonathan Rainer
Ed Sparkes



Airflow Summit
Let's flow together

September 19-21, 2023,
Toronto, Canada

September 2022



Jonathan Rainer
Backend Engineer
(Data Platform)

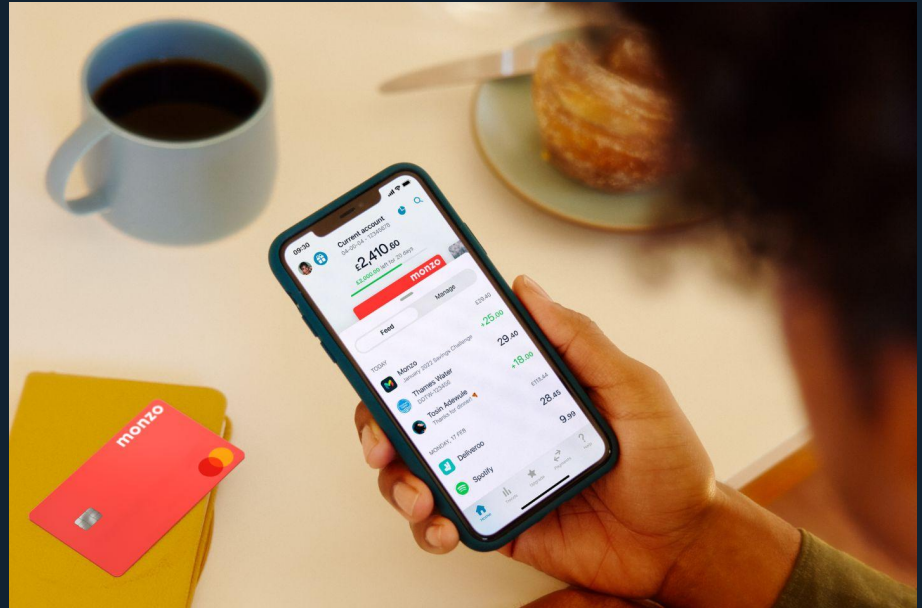


Ed Sparkes
Senior Backend Engineer
(Data Platform)

Who are we?

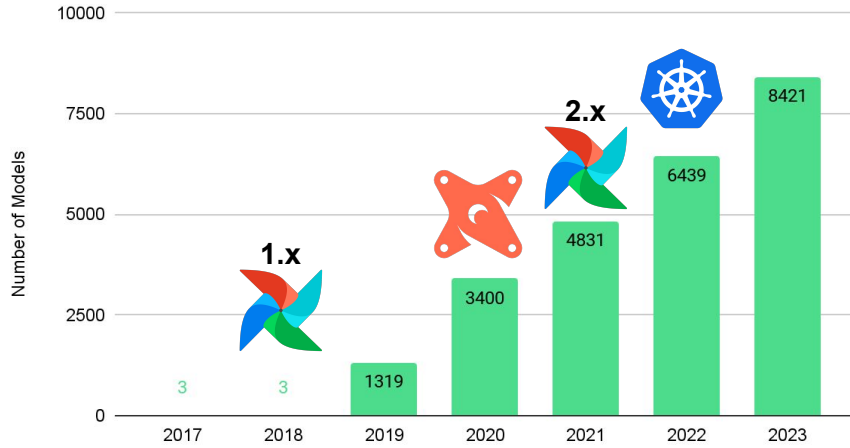
Making Money Work For Everyone

monzo




History of Data At Monzo

Number of Models



Number of Employees in Data Discipline





Events

Analytics
Events


Raw
Data
Store

Airflow

Data
Warehouse

BI
Platform

Airflow + dbt



28

8421

141

529

2

- Dags dynamically defined using dbt model selectors
- SQL pre-rendered in CI by dbt and run in Airflow with Python + BigQuery SDK
- Final DAG defined at runtime by building the graph from the model selector and interleaving python workloads
- BigQuery copy optimisation for shard merging
- Indirect selects allow configurable cross DAG dependencies

```

dag_builder = MonzoDAGBuilder(Config(
    start_date=pendulum.datetime(
        year=2019, month=5, day=20, hour=0, minute=0, tz="Europe/London"
    ),
    execution_timeout=timedelta(hours=6),
))

dag_builder.add_tasks(
    # We manually exclude expensive dependencies from running in near realtime, as it
    # takes longer than 15 minutes.
    # That means that if you model is not explicitly tagged as "run=near_realtime"
    # then this DAG will not run it.
    dbt_model_selectors="--models tag:run=near_realtime",
    cron_paths=[ "cron_a", "cron_b", "cron_c", "cron_d", "cron_e", "cron_f" ]
)

dag = dag_builder.build(
    name="near-realtime",
    schedule="*/30 5-23 * * *", # Every 30 minutes, between 05:00 and 23:00
)

```

Automatic Retry

Retrying Transient Errors

- Failure is inspected via callback and for certain classes of failure tasks are scheduled for retry
- Subsequent successes unblock pipelines automatically and enable us to meet SLAs
- Allow models on our core path alert to a common shared channel on failure
- Persistent failures trigger our internal paging system

monzo

Thread # data-pipeline-monitoring



Data Pipeline Monitoring APP 6 days ago

A task run failed within Airflow which is attributed to a model failure

Task Name:

my_first_model

DAG ID:

weekend_business_critical

Failure Reason:

Google Sheets Resources Exceeded

Auto retry:

(1 / 3)

Start Date: 2023-08-05T01:28:19

Airflow Logs

Github SQL

1

2 replies



Airflow Task Clearer APP 6 days ago

Automatically clearing task **my_first_model** (1 / 3)

Downstream tasks cleared: 12 of 12



Data Pipeline Monitoring APP 6 days ago

Task subsequently succeeded

Start Date: 2023-08-05T01:36:50Z

Airflow Logs

Github SQL

BigQuery

Notifying for quick incident identification

- Model owners can configure how and where they want to be notified when their models are run
- Contextual information to enable debugging is threaded
- Model degradation notifications

monzo

data-monitoring

- dbt Alerts APP 04:42
 my_model_1
 2 replies Last reply 16 days ago
- dbt Alerts APP 04:48
 my_model_2
 6 replies Last reply 16 days ago
- dbt Alerts APP 05:04
 my_model_3
 1 reply 16 days ago

```

{
  config(
    tags=[
      "pit=false",
      "model_type=entity",
      "run=nightly",
      "notify=#data-monitoring",
    ],
    execution_project=choose_execution_project()
  )
}

```

dbt Alerts APP 16 days ago

Log Table Docs DAG

Query Job ID: d868a059-7919-4e52-9f8a-3c28a560bf68

DAG

nightly_weekdays

Tests Passed

1/2

Source Compiled SQL

dbt Alerts APP 16 days ago

my_model_1

Test returned 2 failure(s). Contains warning(s)=True. Contains error(s)=True.

Source Compiled SQL

Upstream of Core Data Models Monitoring APP 13:43

Upstream of core models check for models on critical path increasing 7d average duration returned the following results:

model_to_check


Task: **my_model_4** Dag Id: **nightly_weekdays**

Duration 7D Average Mins: **82** Previous Duration 7D Average Mins: **4**


Duration 7D Average Difference Mins: **78**

Calculating Critical Paths From DAGs


- We want to ensure that our DAGs don't degrade in performance
- We built tooling to retrieve Airflow DAGs and then calculate their critical paths to focus optimisations



Landing Times From Critical Paths



- Our SLOs are based on model landing times
- We can re-use our critical path calculations, combined with estimates to give estimates on when a model will complete

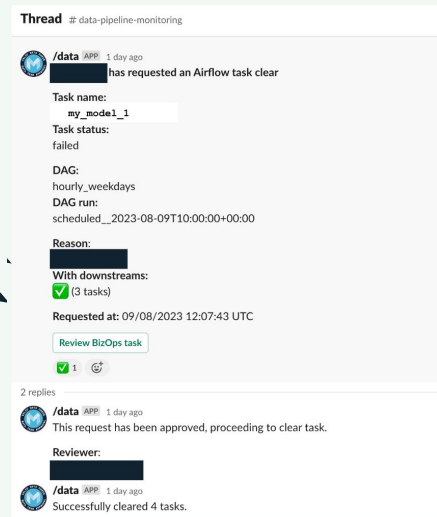
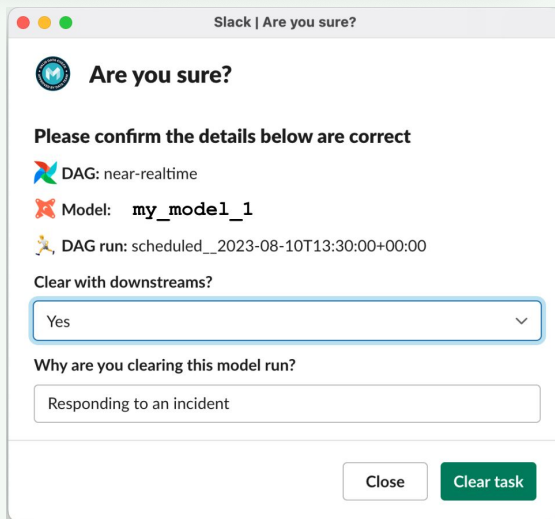
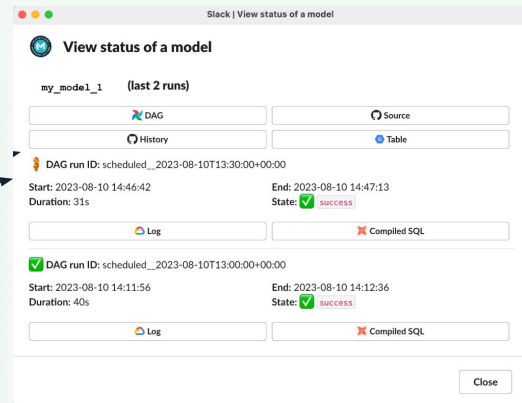
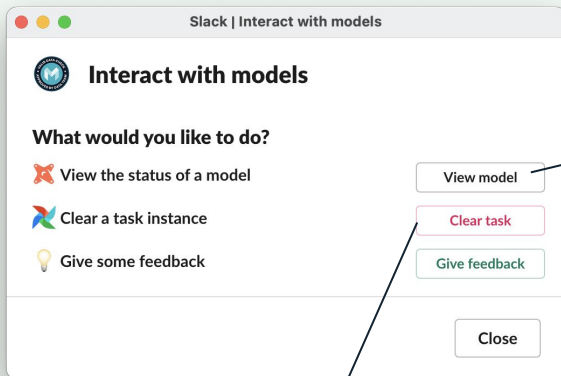
 /data APP 15:46
🕒 Landing Time Estimation (mean strategy) Task name: **model_a**

DAG: near-realtime DAG Run ID: scheduled__2023-08-10T13:30:00+00:00

Estimated to land at: 10 Aug 23 14:47 UTC

Exposing Our Tools

- We expose our tools via a Slack App to reduce the bar to entry for Data Discipline Members
- You can clear tasks, get statuses and estimate landing times
- Potentially destructive tasks (clearing tasks) protected by Multi Party Approval



Airflow Upgrades

We'd like to upgrade from 2.3.4 to take advantage of new features.

Further Cost Reductions

Autoscaling

Deferrable Operators & Triggers



dbt Upgrades & Better DAG Build

We would like to move to 1.5

Questions?

Optionally share some contact info like
email, blog or social media handles



How to navigate this deck:

Section 1 - Title slides

Section 2 - Written content

Section 3 - Tables & Timelines

Section 4 - Using Emojis

Section 5 - Using Graphs

Navigators

Slides in this colourway are navigation slides and not intended to be used as part of your presentations.

Sub copy section

~~Sub copy section
Over two lines~~

Longer Sub copy section

When using the pill, please try keep the copy as short as possible and on one line.

If necessary please extend the pill rather than drop copy onto second line.



1. Title slides

Below you will find a section of title slides to choose from.





September 2022

Hero title goes here

monzo

September 2022

Hero title goes here

monzo

Title slide

You can use these slides to break up your presentation into sections

September 2022

Longer headline copy written across two lines

Sub title copy

monzo

Headline message

Short intro to paragraph (H3)

This space can be used to speak in detail about your work. The paragraphs are broken down into two sections to allow for easy reading. The use of a shorter line length makes the copy more digestible.

By pressing the enter key you can create a paragraph break. The shorter we can keep copy the better.

Short intro to paragraph (H3)

This space can be used to speak in detail about your work. The paragraphs are broken down into two sections to allow for easy reading. The use of a shorter line length makes the copy more digestible.

By pressing the enter key you can create a paragraph break. The shorter we can keep copy the better.

Sub copy section

This space could be used as a short
introduction to your presentation.

monzo

Title slide

You can use these slides to break up your presentation into sections

Sub copy section

1,000,000+
customers

monzo

2. Written content

Below you will find a section slides which can be used to display copy.

Longer headlines can run across two lines

Short intro to paragraph (H3)

This space can be used to speak in detail about your work. The paragraphs are broken down into two sections to allow for easy reading. The use of a shorter line length makes the copy more digestible.

By pressing the enter key you can create a paragraph break. The shorter we can keep copy the better.

monzo

Short intro to paragraph (H3)

This space can be used to speak in detail about your work. The paragraphs are broken down into two sections to allow for easy reading. The use of a shorter line length makes the copy more digestible.

By pressing the enter key you can create a paragraph break. The shorter we can keep copy the better.

1

Space for copy step 1

Moving away from agencies and growing our capabilities to help us power growth

2

Space for copy step 2

Moving away from agencies and growing our capabilities to help us power growth

3

Space for copy step 3

Moving away from agencies and growing our capabilities to help us power growth

Sub copy section

Goals for next meeting.

monzo

- 1 Short and concise steps can help the reader digest information. Bring them on your journey.
- 2 Short and concise steps can help the reader digest information. Bring them on your journey.
- 3 Short and concise steps can help the reader digest information. Bring them on your journey.
- 4 Short and concise steps can help the reader digest information. Bring them on your journey.

Sub copy section

Headline message

Subheader paragraph (H3)

We're a species of skim readers. We've got more to read than any humans in history, on smaller and smaller screens, with less and less time to do it. So one of our toughest jobs is just to hold people's attention.

That's where subheadings come in.

monzo

Subheader paragraph (H3)

We're a species of skim readers. We've got more to read than any humans in history, on smaller and smaller screens, with less and less time to do it. So one of our toughest jobs is just to hold people's attention.

That's where subheadings come in.

Subheader paragraph (H3)

We're a species of skim readers. We've got more to read than any humans in history, on smaller and smaller screens, with less and less time to do it. So one of our toughest jobs is just to hold people's attention.

That's where subheadings come in.

Subheader paragraph (H3)

We're a species of skim readers. We've got more to read than any humans in history, on smaller and smaller screens, with less and less time to do it. So one of our toughest jobs is just to hold people's attention.

That's where subheadings come in.

Sub copy section

Longer headlines can run across two lines

Subheader paragraph (H3)

We're a species of skim readers. We've got more to read than any humans in history, on smaller and smaller screens, with less and less time to do it. So one of our toughest jobs is just to hold people's attention.

Subheader paragraph (H3)

We're a species of skim readers. We've got more to read than any humans in history, on smaller and smaller screens, with less and less time to do it. So one of our toughest jobs is just to hold people's attention.

Subheader paragraph (H3)

We're a species of skim readers. We've got more to read than any humans in history, on smaller and smaller screens, with less and less time to do it. So one of our toughest jobs is just to hold people's attention.

Subheader paragraph (H3)

We're a species of skim readers. We've got more to read than any humans in history, on smaller and smaller screens, with less and less time to do it. So one of our toughest jobs is just to hold people's attention.

With this slide you should be able to take these steps (H4)

1

Write in steps about your process and the journey. Use these blocks of text to bring the reader on a journey. Help them understand your process by walking them through the steps you took.

2

Write in steps about your process and the journey. Use these blocks of text to bring the reader on a journey. Help them understand your process by walking them through the steps you took.

3

Write in steps about your process and the journey. Use these blocks of text to bring the reader on a journey. Help them understand your process by walking them through the steps you took.

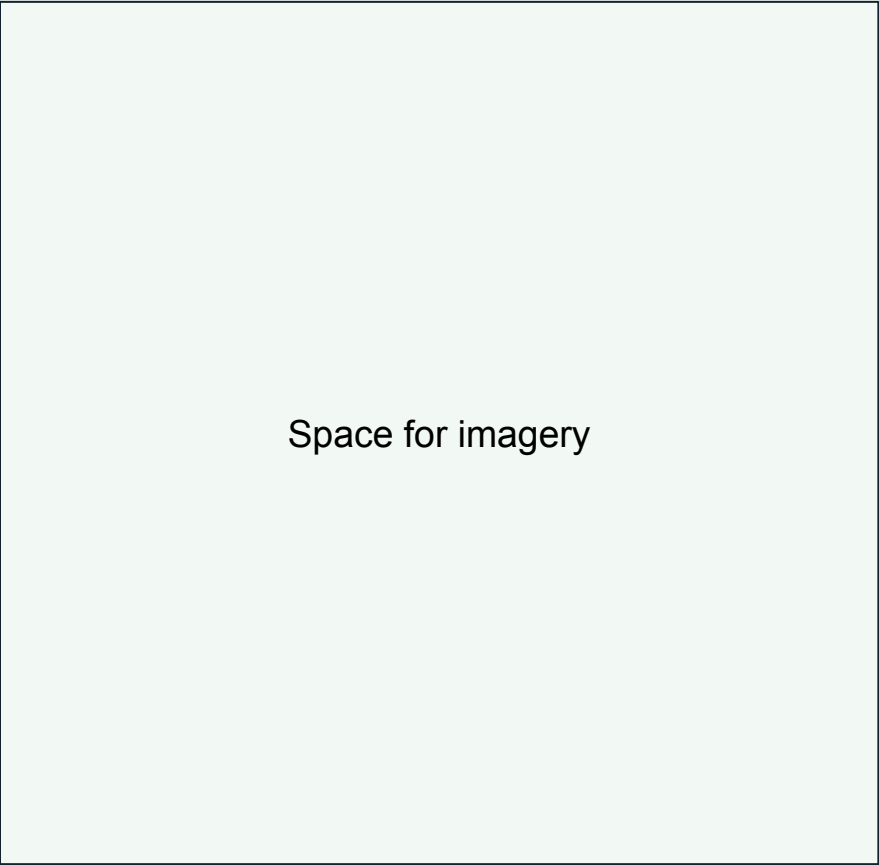
Sub copy section

Developing longer headline copy examples

Short intro to paragraph (H4)

On slides with large imagery you can insert any image you like. Simply click the Hot Coral square, then click the “replace image” button and choose your new image you would like to insert.

monzo



Space for imagery

Short lists of bullet points

Our advice on bullet points is pretty simple. Big lists of bullets are no easier to read than big blocks of text. So they should be:

- Text size drops to 11.5
- Spacing set to 1.5
- No more than a (short) sentence
- Under 6 per list
- Use two columns if needed
- Related to each other
- Consistent in style

Sub copy section

What you want to say, in steps

monzo

- 1 Short and concise steps can help the reader digest information. Bring them on your journey.
- 2 Short and concise steps can help the reader digest information. Bring them on your journey.
- 3 Short and concise steps can help the reader digest information. Bring them on your journey.
- 4 Short and concise steps can help the reader digest information. Bring them on your journey.

Short lists of bullet points

Our advice on bullet points is pretty simple. Big lists of bullets are no easier to read than big blocks of text. So they should be:

- Text size drops to 16
- Spacing set to 1.5
- No more than a (short) sentence
- Under 6 per list

Sub copy section

“Including a quote: Monzo changed my life, I am finally in control of my spending and it feels great!”

– Jake Peralta

monzo



To replace the above image just select the box and select the "replace image" tab above.

You can use this space for a little blurb about the new starter.

3. Tables & timelines content

Below you will find a section of slides which can be used to display information within tables and timelines to take the viewer on a journey.

	Category 1	Category 2	Category 3	Category 4
Level 1	7.895	5.874	9.670	2.565
Level 2	7.445	6.543	7.883	9.224
Level 3	4.663	2.668	8.437	9.548



Timeline No1

Secondary copy,
Description

Timeline No3

Secondary copy,
Description

Timeline No2

Secondary copy,
Description

Timeline No4

Secondary copy,
Description

Sub copy section

March

Sub copy section

Lorem ipsum dolor sit
 amet, consectetur
 adipiscing elit.

Lorem ipsum dolor sit
 amet, consectetur
 adipiscing elit.

Lorem ipsum dolor sit
 amet, consectetur
 adipiscing elit.

April

Sub copy section

Lorem ipsum dolor sit
 amet, consectetur
 adipiscing elit.

Lorem ipsum dolor sit
 amet, consectetur
 adipiscing elit.

Lorem ipsum dolor sit
 amet, consectetur
 adipiscing elit.

May

Sub copy section

Lorem ipsum dolor sit
 amet, consectetur
 adipiscing elit.

Lorem ipsum dolor sit
 amet, consectetur
 adipiscing elit.

Lorem ipsum dolor sit
 amet, consectetur
 adipiscing elit.

June

Sub copy section

Lorem ipsum dolor sit
 amet, consectetur
 adipiscing elit.

Lorem ipsum dolor sit
 amet, consectetur
 adipiscing elit.

Lorem ipsum dolor sit
 amet, consectetur
 adipiscing elit.

July

Sub copy section

Lorem ipsum dolor sit
 amet, consectetur
 adipiscing elit.

Lorem ipsum dolor sit
 amet, consectetur
 adipiscing elit.

Lorem ipsum dolor sit
 amet, consectetur
 adipiscing elit.

With this slide you should be able to take these steps (timeline)

Write in steps about your process and the journey. Use these blocks of text to bring the reader on a journey. Help them understand



Write in steps about your process and the journey. Use these blocks of text to bring the reader on a journey. Help them understand



Write in steps about your process and the journey. Use these blocks of text to bring the reader on a journey. Help them understand



Write in steps about your process and the journey. Use these blocks of text to bring the reader on a journey. Help them understand



Breaking content Into categories

Short intro to paragraph (H4)

On slides with large imagery you can insert any image you like. Simply click the Hot Coral square, then click the “replace image” button and choose your new image you would like to insert.

Category one

Lorem ipsum dolor sit amet. Quo tempora placeat aut enim blanditiis et voluptas eius aut architecto eaque eos quia corrupti

Category two

Lorem ipsum dolor sit amet. Quo tempora placeat aut enim blanditiis et voluptas eius aut architecto eaque eos quia corrupti

Category three

Lorem ipsum dolor sit amet. Quo tempora placeat aut enim blanditiis et voluptas eius aut architecto eaque eos quia corrupti

Category four

Lorem ipsum dolor sit amet. Quo tempora placeat aut enim blanditiis et voluptas eius aut architecto eaque eos quia corrupti

Category five

Lorem ipsum dolor sit amet. Quo tempora placeat aut enim blanditiis et voluptas eius aut architecto eaque eos quia corrupti

What you want to say (H4)

With this slide you can insert four images. You can insert any image you like. Simply click the Hot Coral square, then click the “replace image” button and choose your new image you would like to insert.



Write in steps about your process and the journey. Use these blocks of text to bring the reader on a journey. Help them understand



Write in steps about your process and the journey. Use these blocks of text to bring the reader on a journey. Help them understand



Write in steps about your process and the journey. Use these blocks of text to bring the reader on a journey. Help them understand



Write in steps about your process and the journey. Use these blocks of text to bring the reader on a journey. Help them understand

Sub copy section

**Category
One**

**Category
Two**

**Category
Three**

**Category
Four**

**Category
Five**

**Category
Six**

4. Using Emojis and illustrations

Below you will find a section of slides
which explains how we can best use
emojis and illustrations in our
presentations



Using emojis

Emojis set us apart and reflect what we're like as people: colourful, friendly and open to new ideas. So feel free to use them, especially if it helps the slide



Using emojis

Emojis set us apart and reflect what we're like as people: colourful, friendly and open to new ideas. So feel free to use them, especially if it helps the slide



Using emojis

Emojis set us apart and reflect what we're like as people: colourful, friendly and open to new ideas. So feel free to use them, especially if it helps the slide

Using emojis

To use more emojis you can use <https://emojipedia.org/>.
Simply type in your keyword and then copy the
Emoji and paste it into the doc.



Using emojis

Emojis can help to break up copy and
create a visual aid for your point.



Using emojis

Emojis can help to break up copy and
create a visual aid for your point.



Using emojis

Emojis can help to break up copy and
create a visual aid for your point.



Using emojis

Emojis can help to break up copy and
create a visual aid for your point.

Sub copy section



Sub copy section



monzo

Sub copy section



monzo


5. Using Graphs

Below you will find a section of slides which allows you to drop in your graphs and add a description. These slides use a whistle background, allowing you to drop in screenshots of your charts, graphs etc.



Pie chart example

Sub header

Most charts will tend to be created outside of slides and imported in on a white background. To insert your own chart you can click on the existing image, select “replace image” from the button above, and then choose your replacement image from your computer.



Space for a tag-line where you can give context for the graphs and help the viewer understand.



Sub copy section

Space for a tag-line where you can give context for the content and help the viewer understand.

5. Mockups

