

# Predicting Workflow Performance

Hugo Hiden

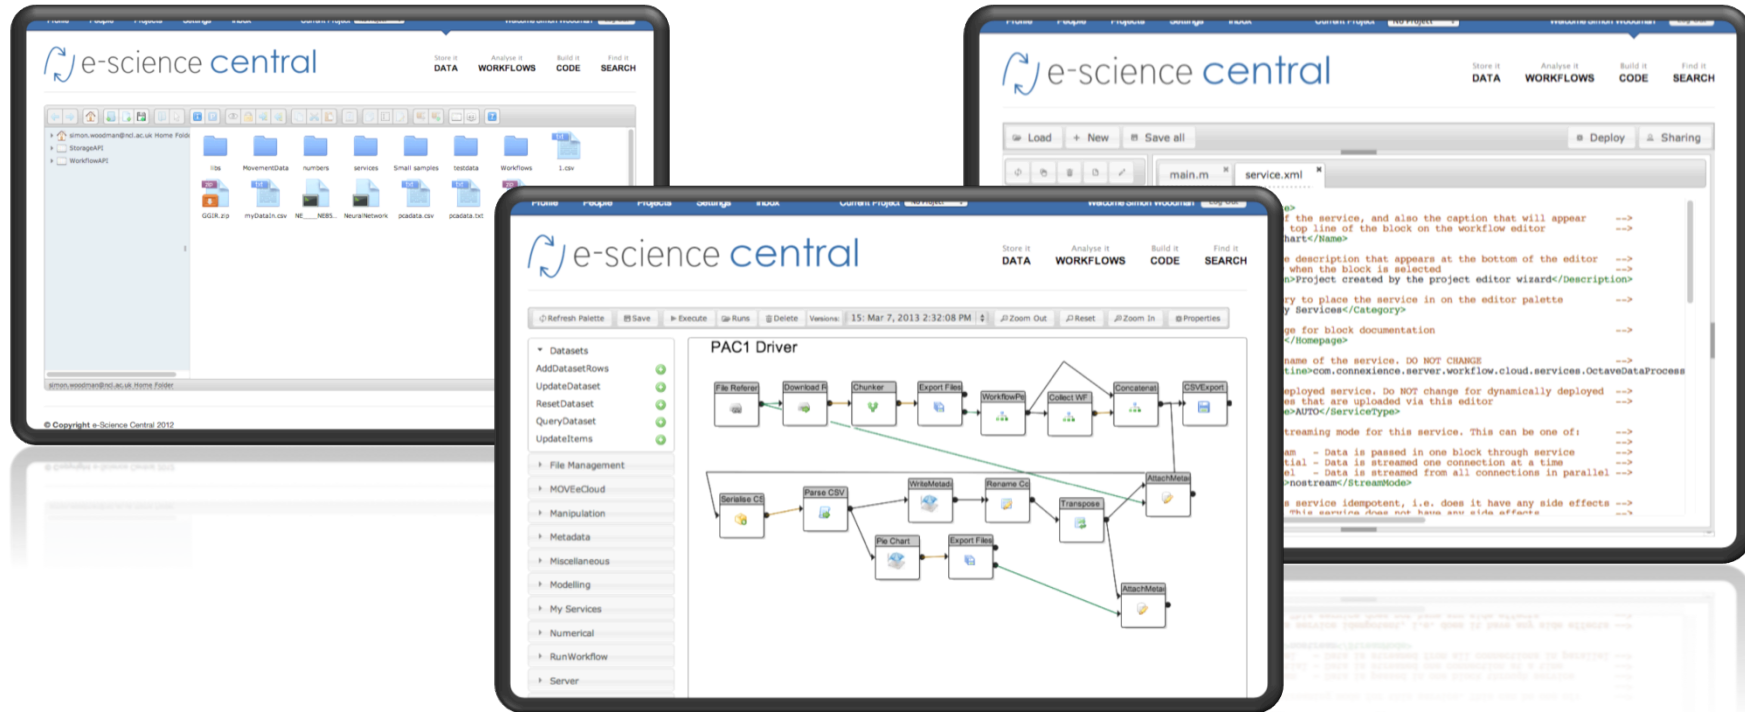
Simon Woodman

# Performance Prediction

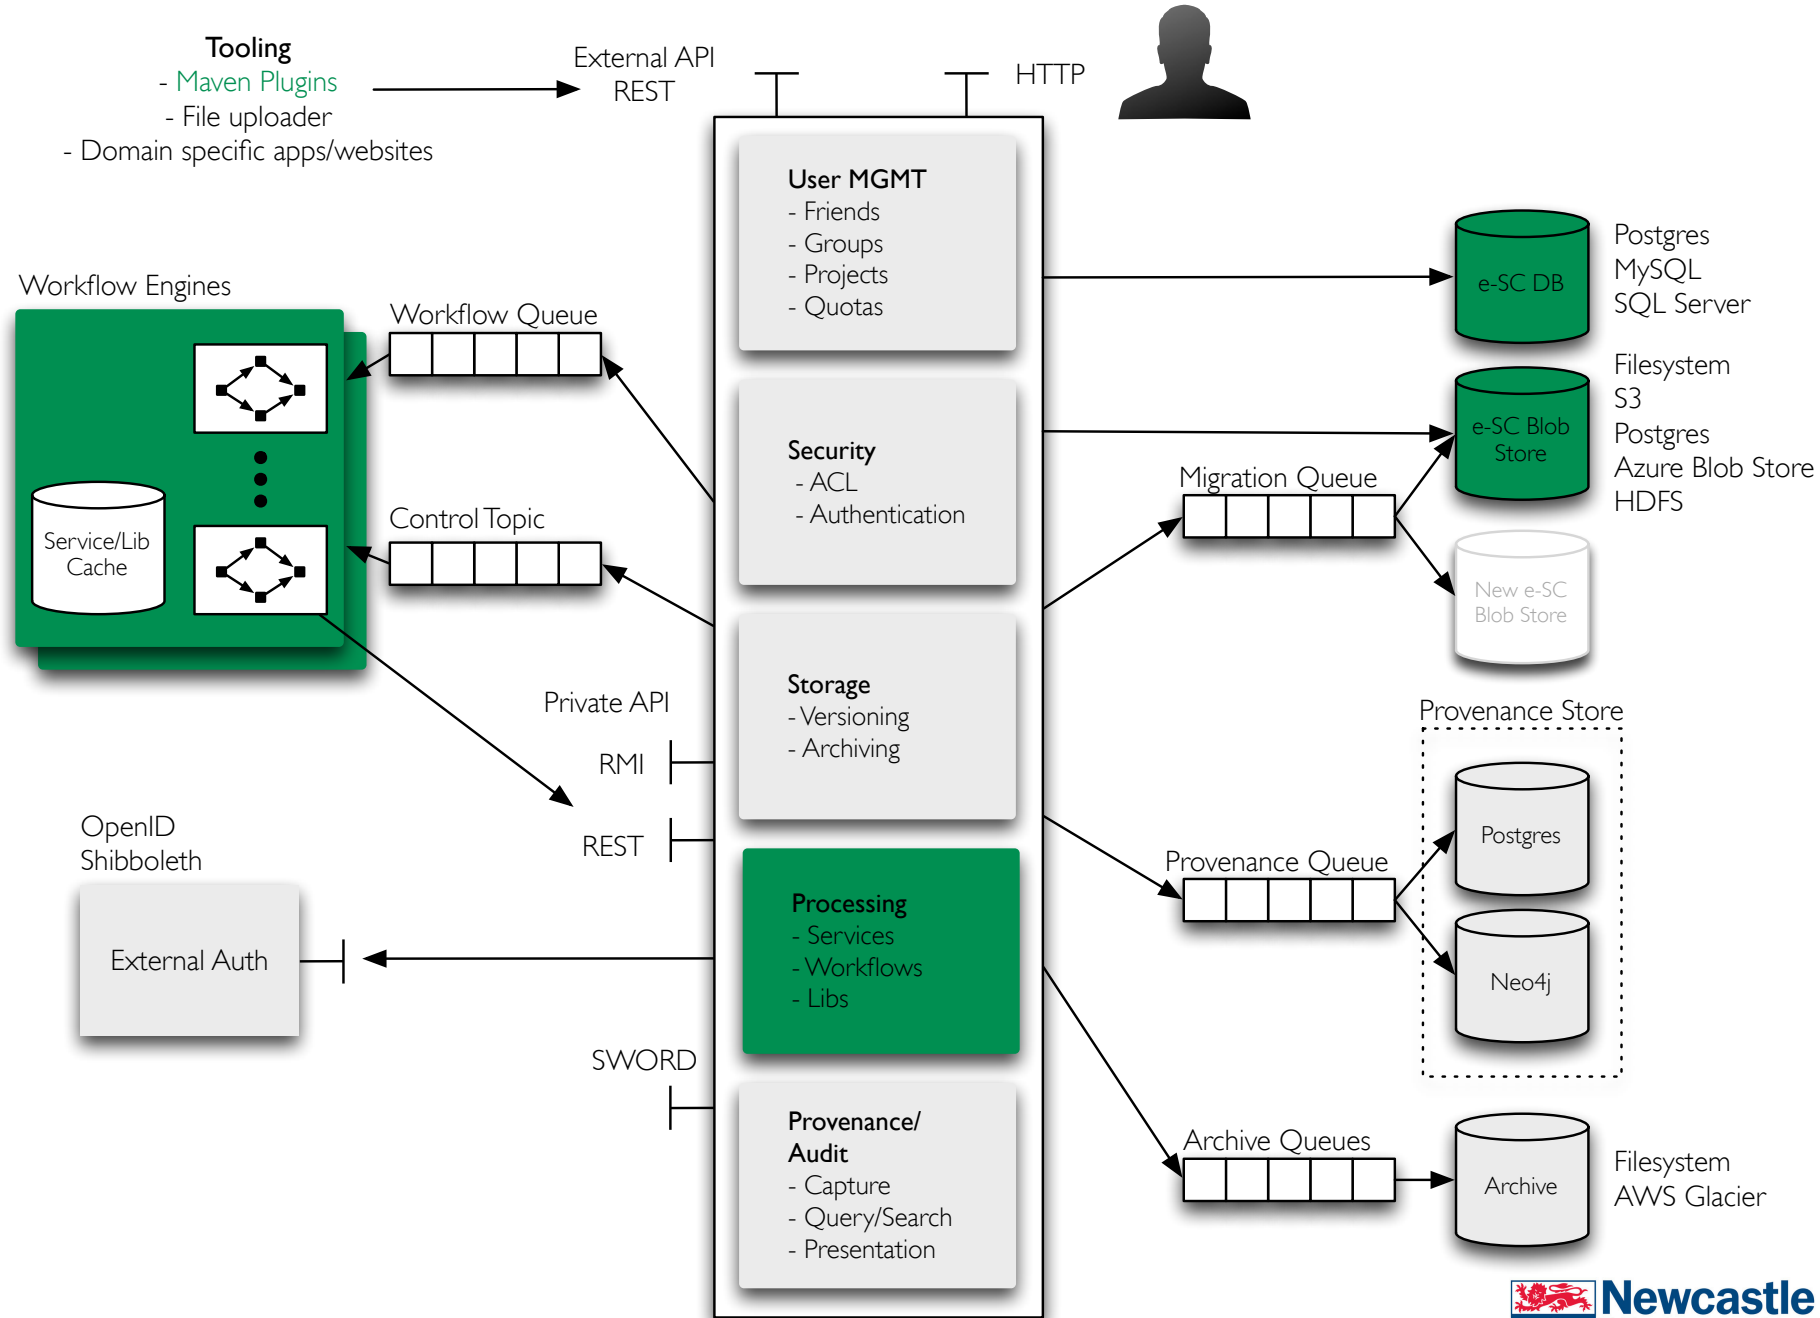
- Good because:
  - Cost in Cloud
  - Know when to expect results
  - Capacity planning
  - Other work needs it
    - Paul's security
    - Scheduling
  - Other possibilities

# e-Science Central

Platform for cloud based data analysis



Built as a distributed system...



# Workflows for data processing

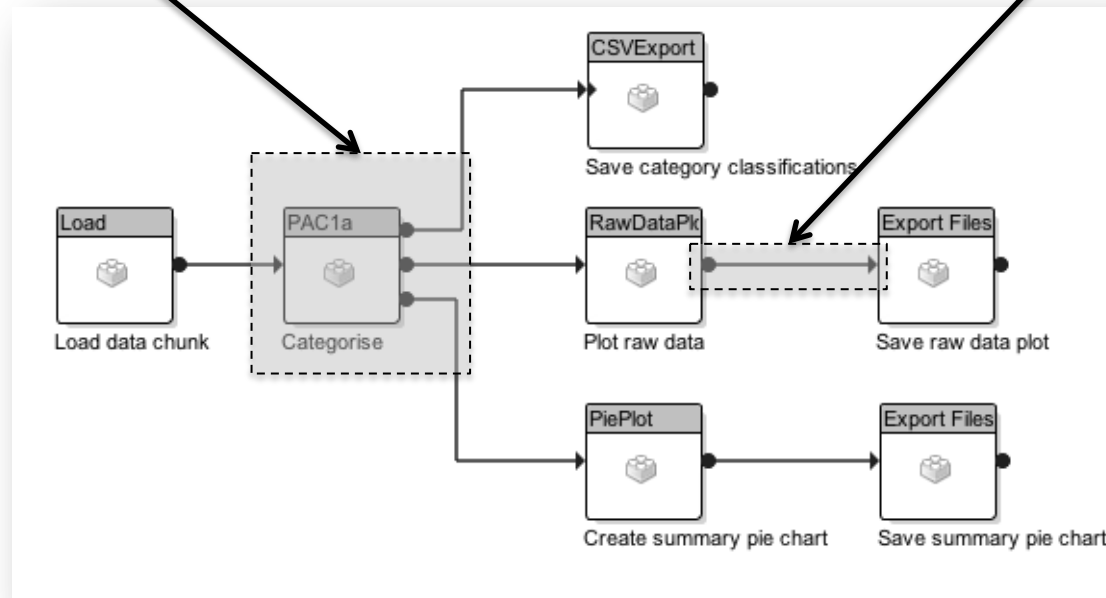
- Spreadsheets become unwieldy quickly
  - Emphasise **data** over **process**
  - Hard to see what has been done with data
    - Not obvious what calculations have been done
    - Hard to extract some of the the calculations and re-use them
  - Require everything to be done using the spreadsheets tools
    - May not include everything needed
- Workflows attempt to mitigate some of this
- Integrate different languages
  - Java, R, Octave, Javascript

# Anatomy of a workflow

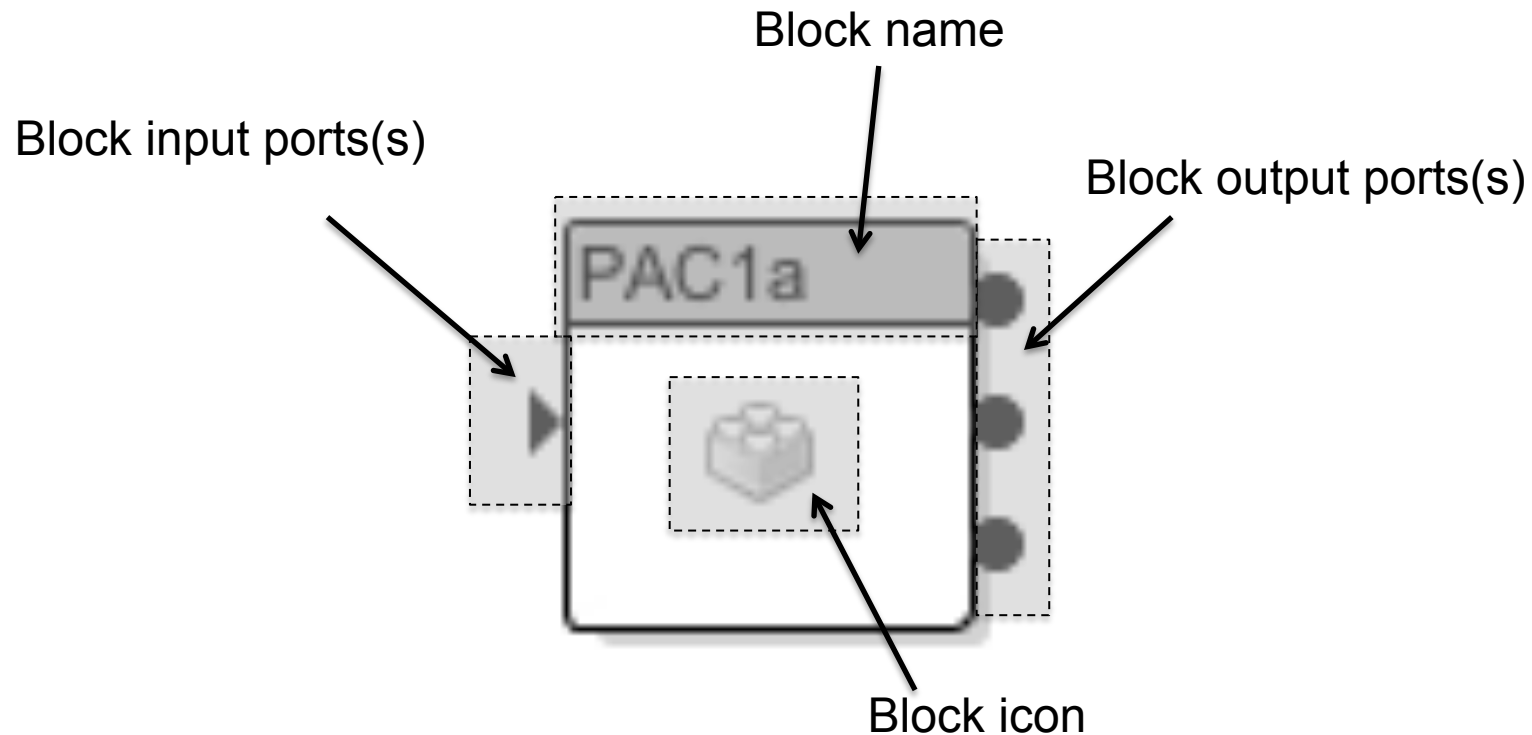
Workflows are made up of Blocks

Data flows between Blocks in the direction of the arrows

Workflow block



# Anatomy of a workflow Block



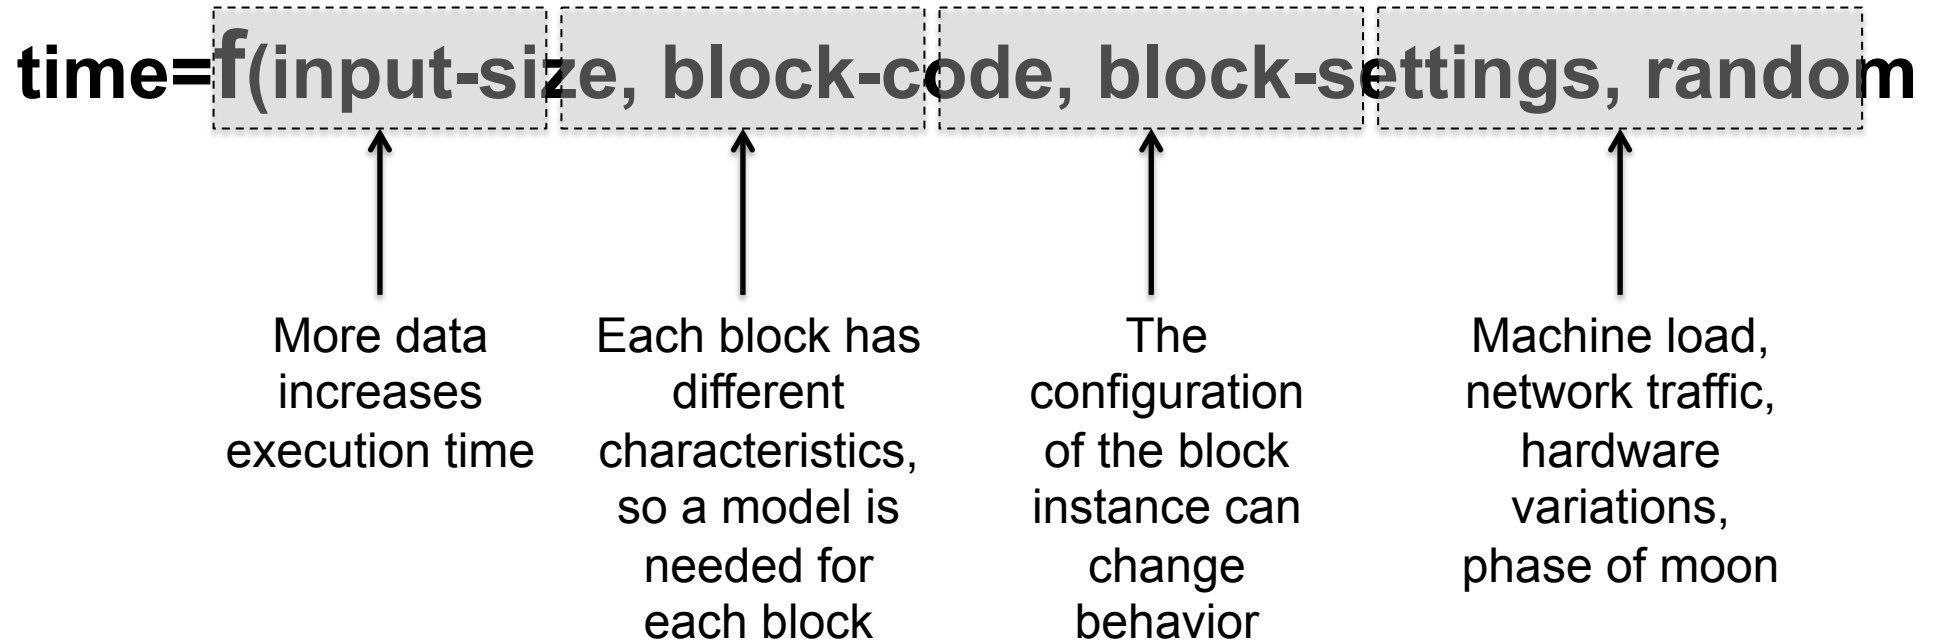
- Blocks read data from their input ports, process it and pass the results to their output ports
- Ports have specific meaning – e.g. on a block with multiple output ports, each port will typically contain a part of the result
- Each output port can be connected to multiple input ports on other blocks
- Only one connection is allowed per input port

# Factors influencing performance

- Variable execution time
  - dedicated machine < local server < cloud VM
- There are good predictors
  - The code itself
  - The configuration of the block
  - The input data sizes
  - The machine it is running on
- Predictable?



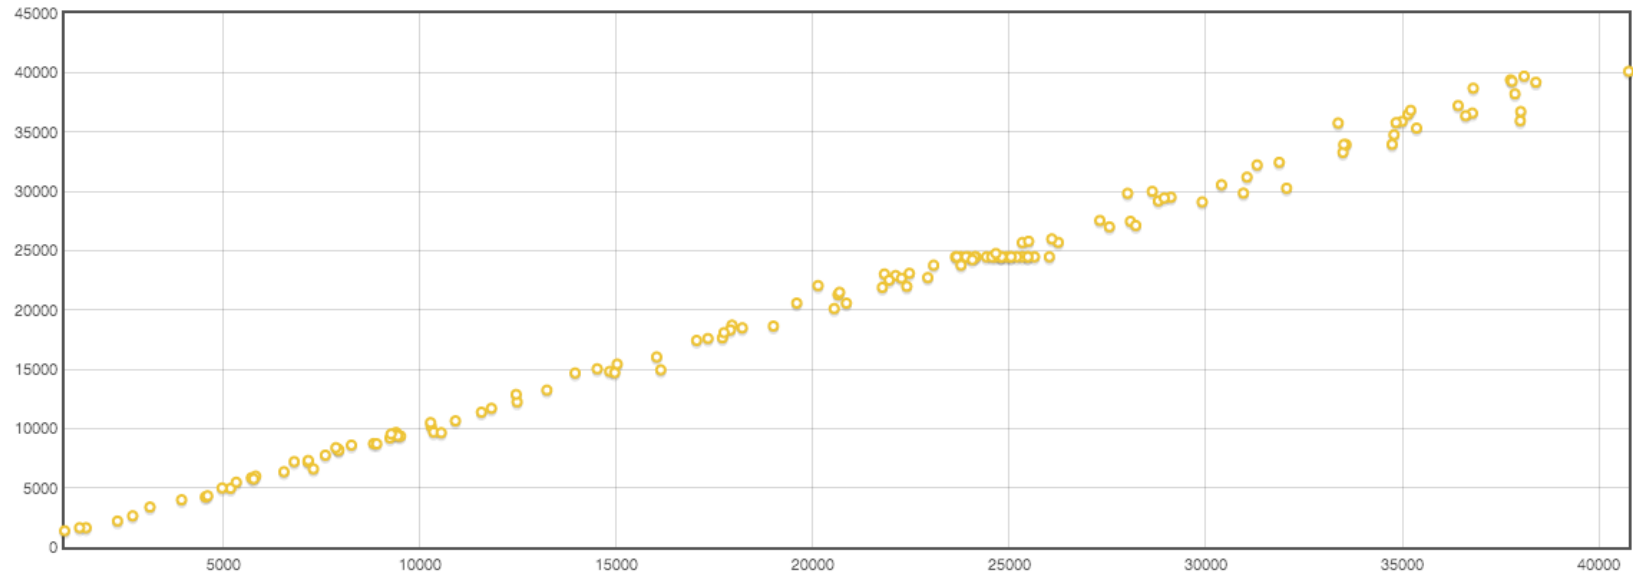
# Execution time of a block



A workflow is a connected pathway of blocks...

# A predictable block

## Actual -v- Predicted: Duration



Parameter Name

Parameter Type

Parameter Value

Intercept

ModelParameter

1356.713285822049


input-data

InputSize

0.0035241319375301345

# A less predictable block

**Actual -v- Predicted: Duration**



Parameter Name	Parameter Type	Parameter Value
Intercept	ModelParameter	1206.6135807933206
Source	ServiceProperty	2.4409403971602595E-5

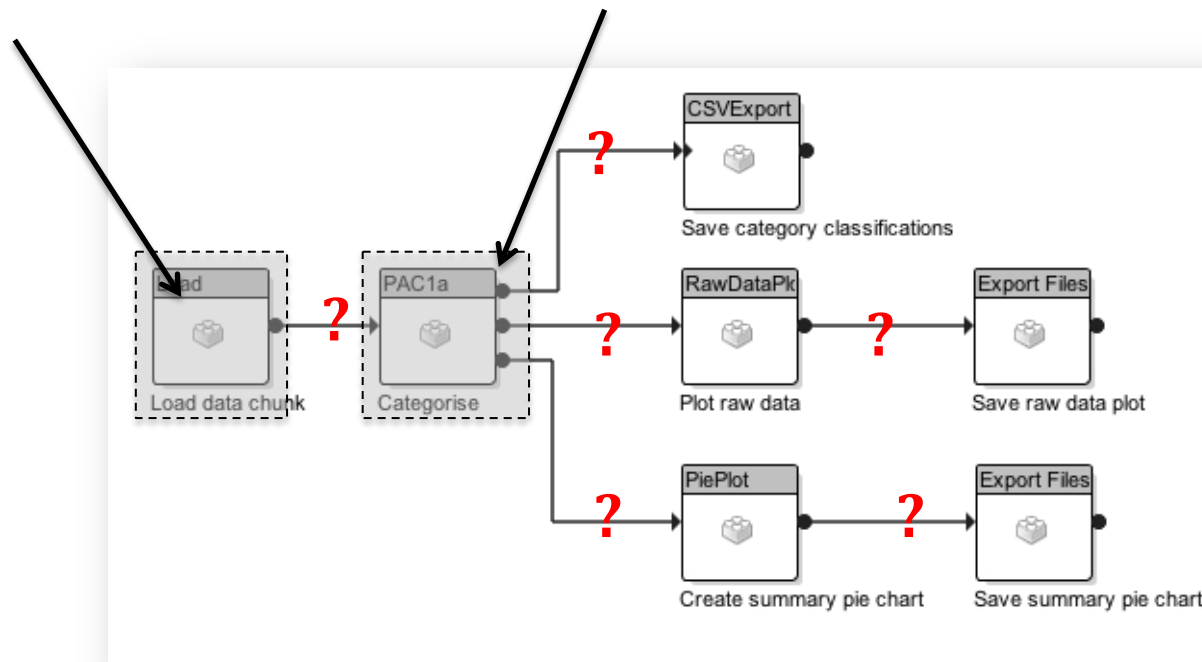
# Predicting Workflow duration

Modelling is complicated by connected nature of workflow

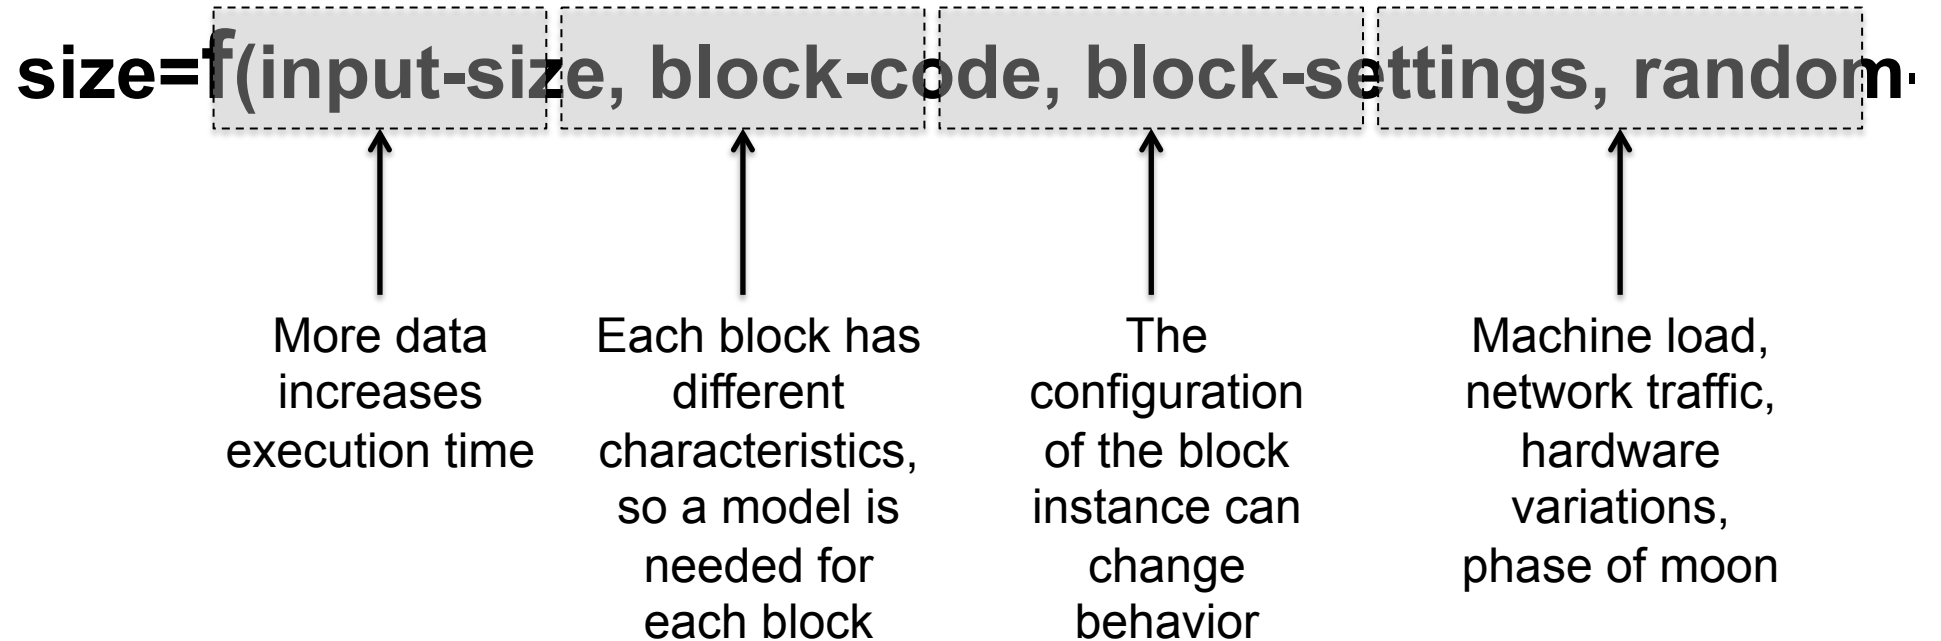
All data for model readily available...

... not the case here

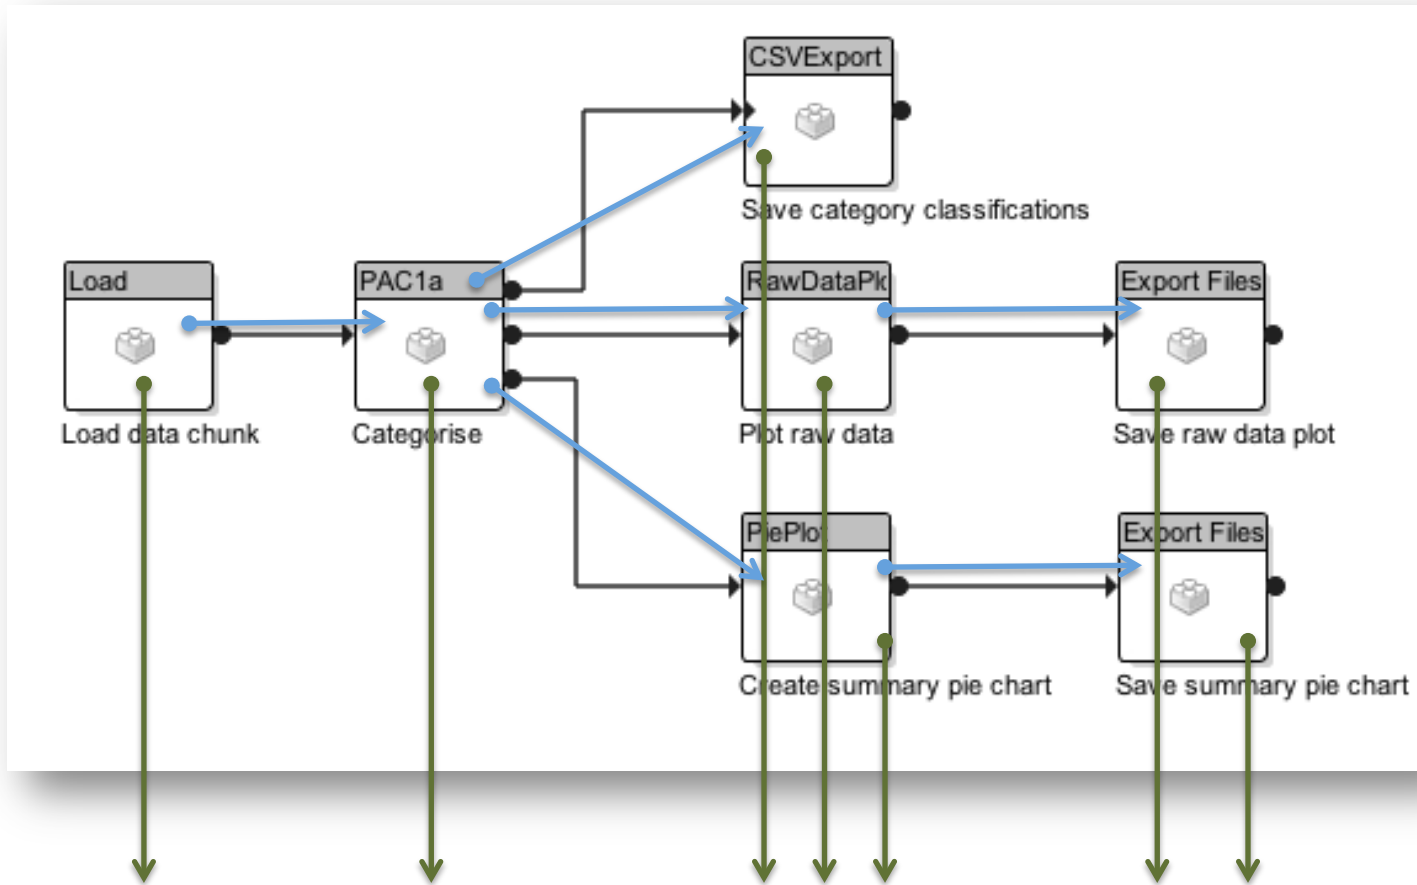
? how big are the intermediate data transfers



# Data volume produced by a block



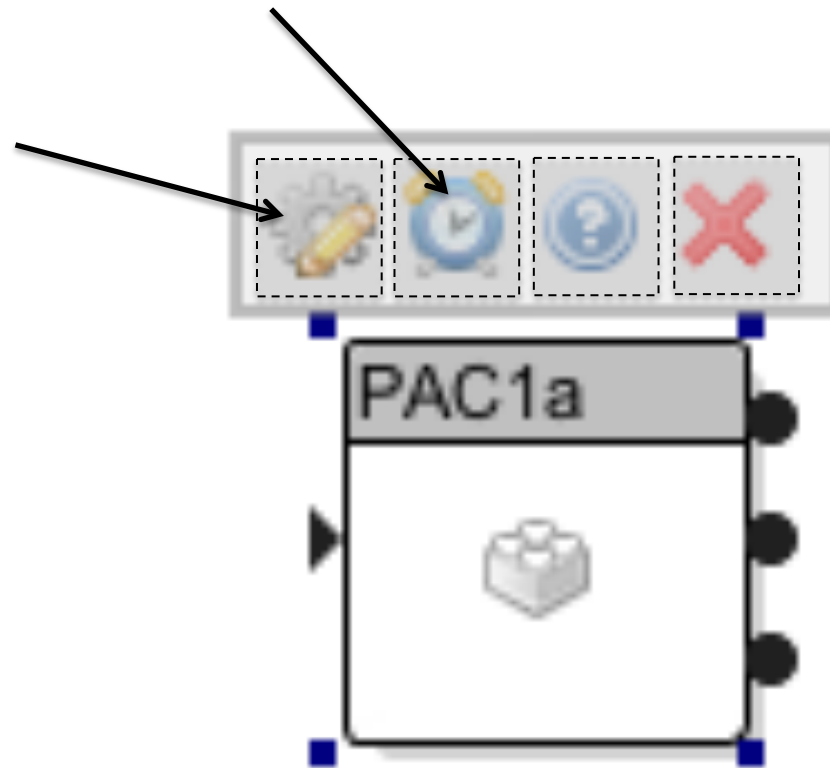
# Propagation of data sizes



# Configuring a workflow Block

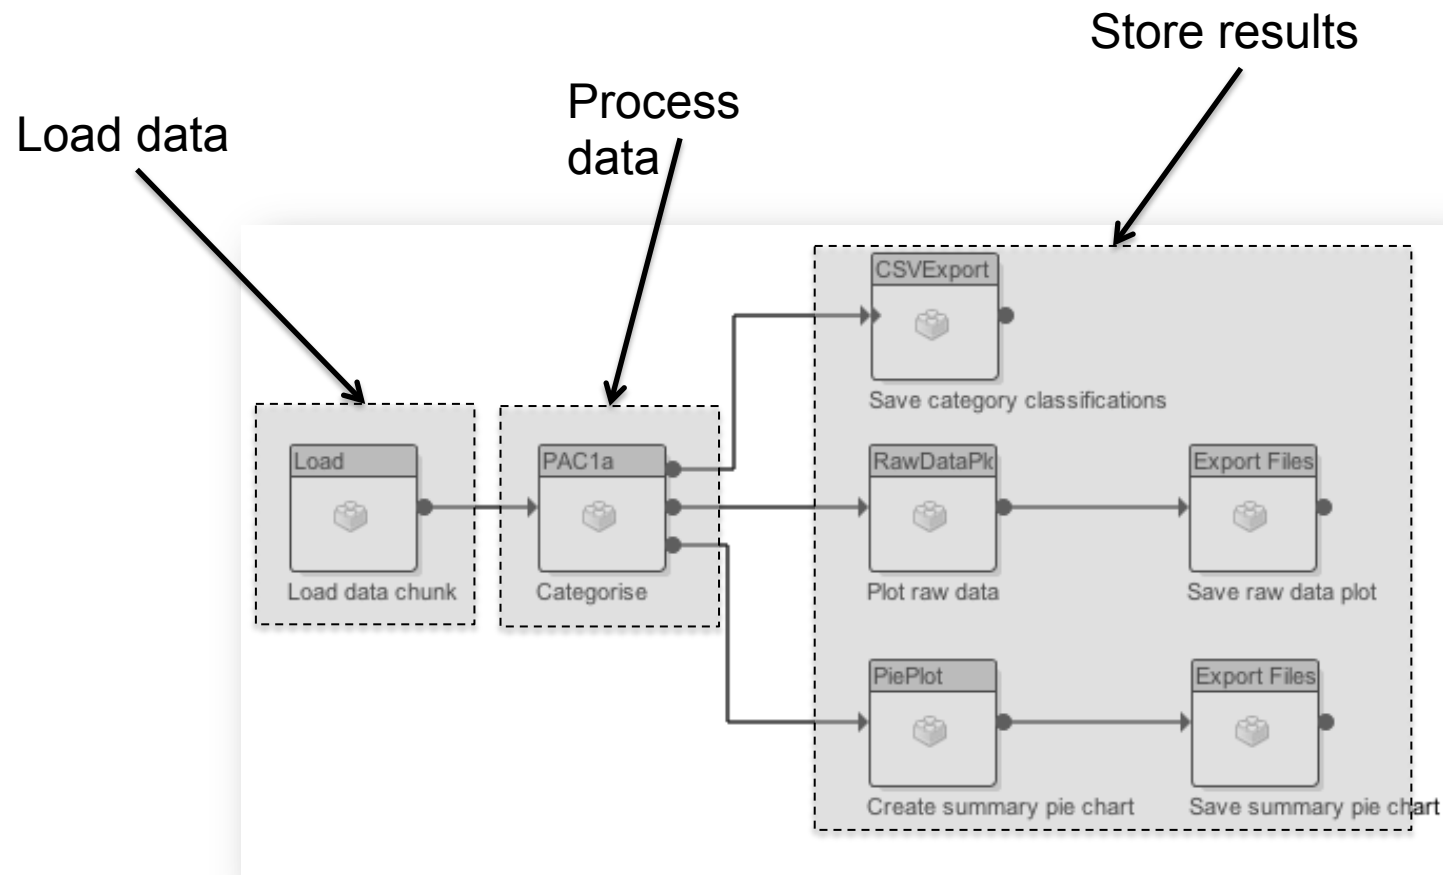
Select the version  
of the block

Configure  
block



# Anatomy of a workflow

Typical simple workflow – follows the standard pattern



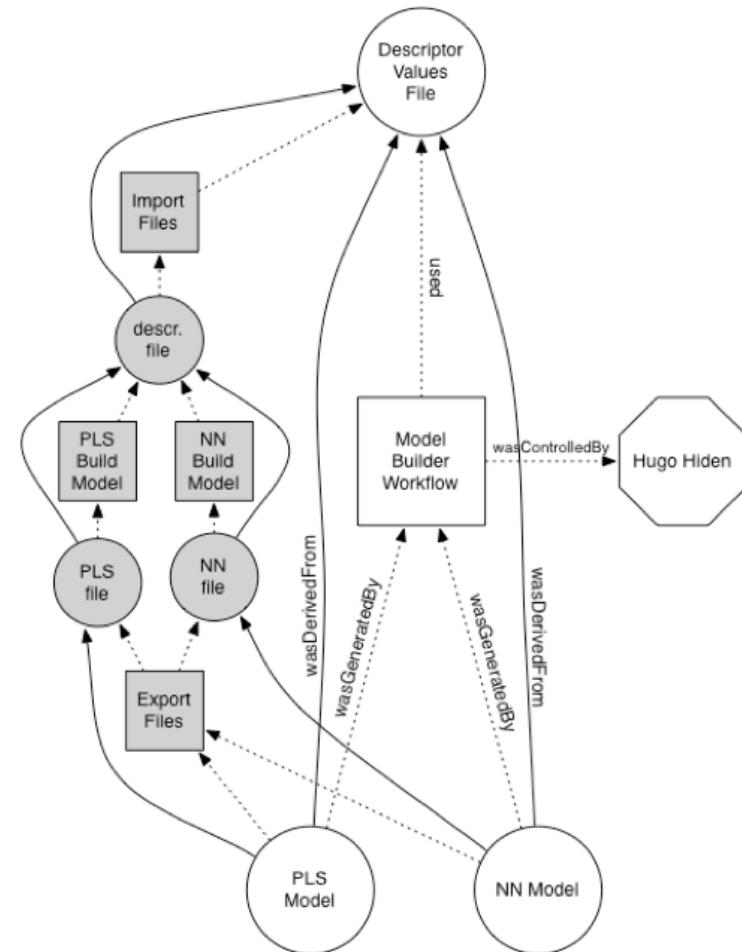


# Provenance/Audit Requirements

- How was data generated?
  - What algorithm?
  - What version?
- Are these results reproducible?
- How have bugs manifested?
  - Which data affected
  - How do we regenerate affected data?
- Performance Characteristics
- How do we deal with new data?

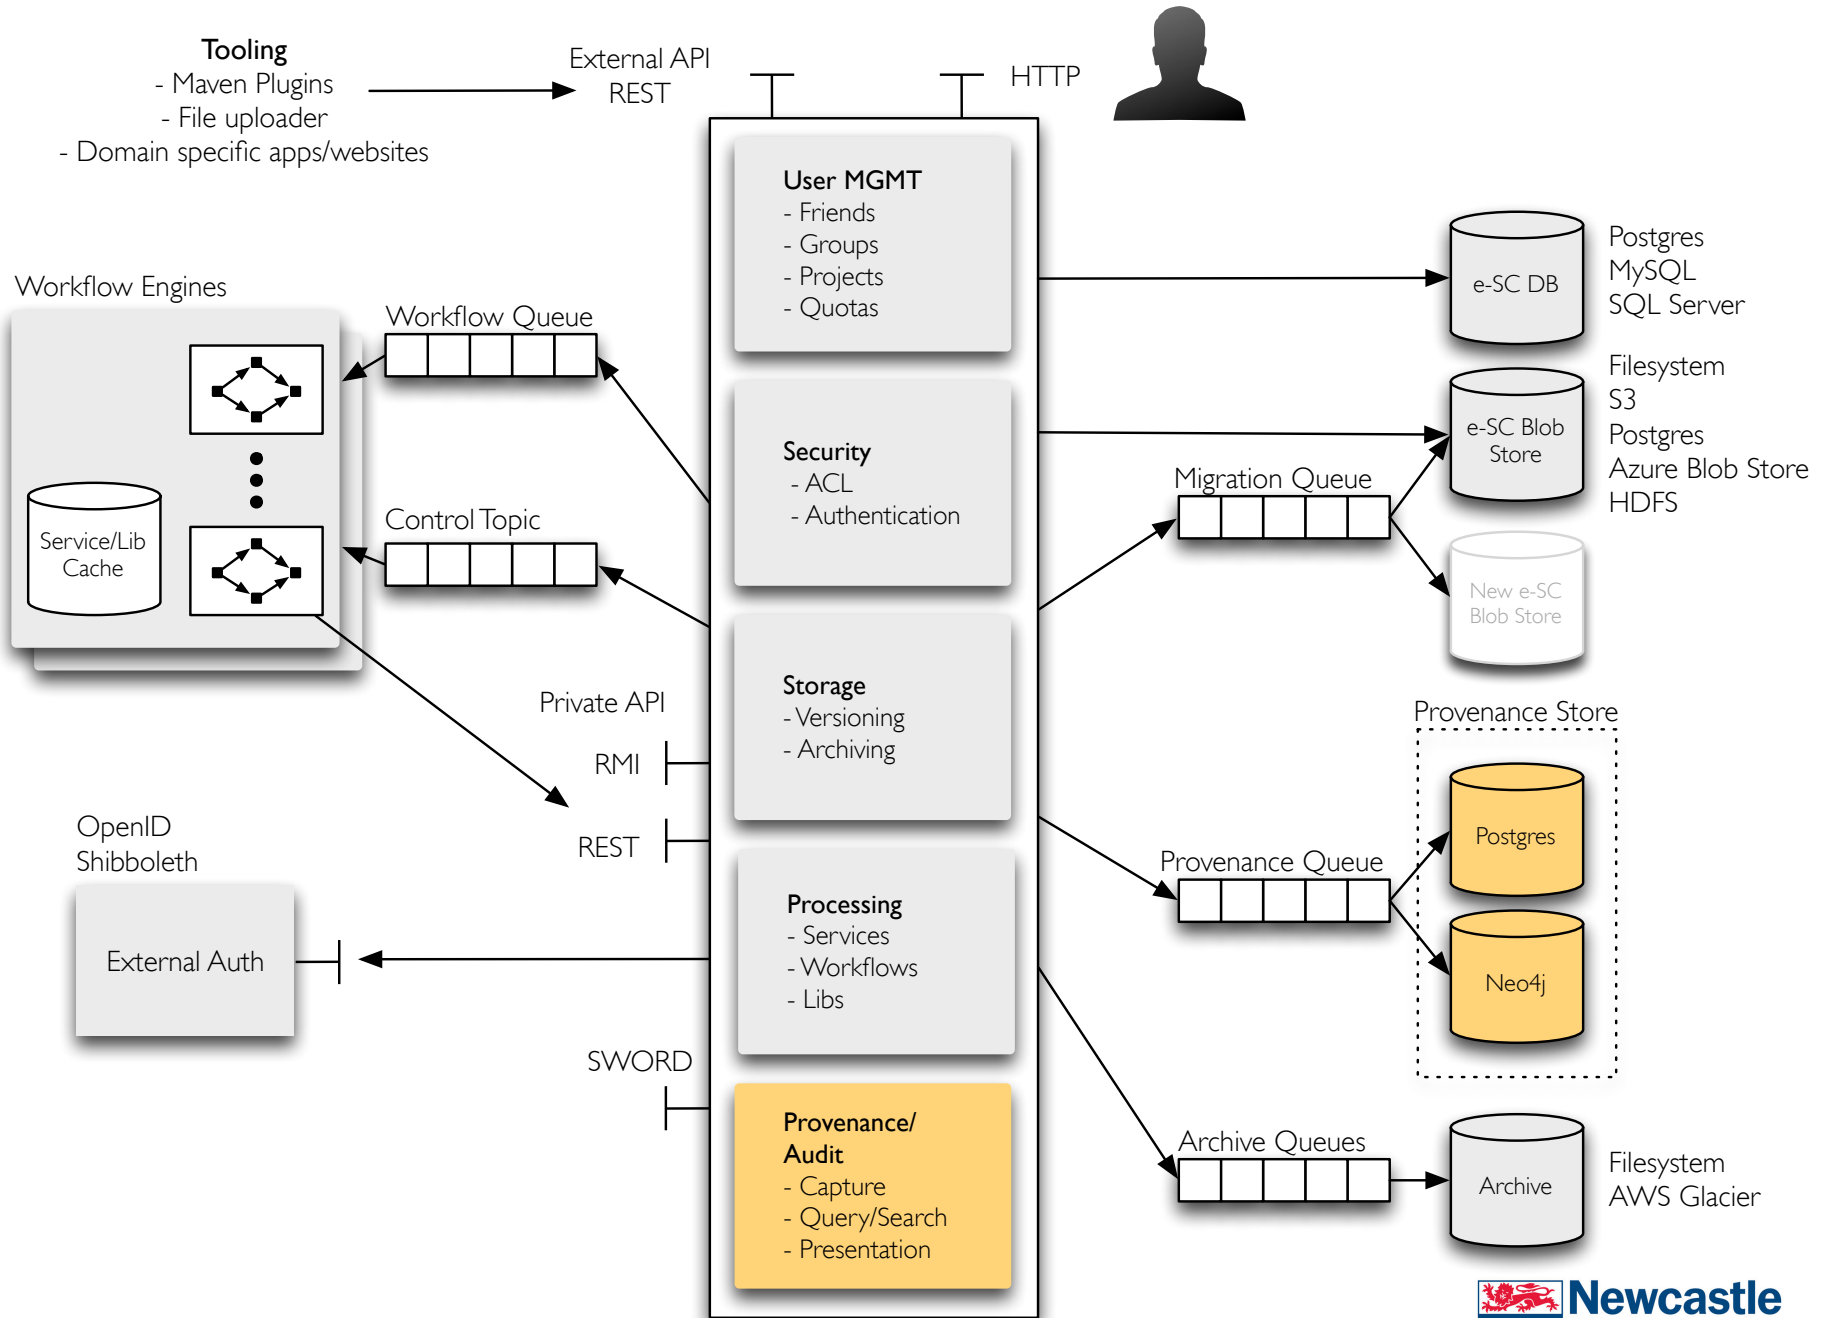
# Provenance Model

- Based on OPM
  - Processes, Artifacts, Agents
- Directed Graph
- Multiple views of provenance
  - Dependent on security privileges



# Storing Provenance

- Neo4j
  - Open Source Graph Database
  - Nodes/Relationships + properties
  - Querying/traversing
- Access
  - Java lib for OPM
  - e-SC library built on top of OPM lib
  - REST interface
- Options for HA and Sharding for performance



# Workflow Blocks

- Workflow blocks are units of code that execute as part of a workflow
  - They have a defined structure
  - Can be configured using properties
    - Strings, numbers, booleans, file references, lists etc
  - Can act on local files, data-sets, name-value pairs, serialized Java objects and links to files stored in e-SC
  - Need to be able to operate without user interaction

# Execution of Workflow Blocks

- Blocks execute as part of a workflow
  - The code is transferred to the machine executing the workflow
  - The block code is unpacked on the workflow machine
  - Dependencies are also downloaded and unpacked
  - The block is then executed in the workflow working directory
    - Properties are assigned
    - Initialisation code is executed
    - The block main code is then executed (potentially multiple times)
    - Termination code is then executed

# Accessing the e-SC Server


- Java and JavaScript blocks have access to an API that provides limited access to information held in the e-SC server
  - Upload / download files
  - Attach / read metadata
  - Update / query datasets
  - Execute additional workflows
- Actions performed using the API are carried out as the same user executing the workflow

# Case Studies



# QSAR - The Problem

What are the properties of this molecule?



Perform experiments



Time consuming

Expensive

Ethical constraints

# QSAR

## Quantitative Structure Activity Relationship

$$\text{Activity} \approx f(\text{Structure})$$

More accurately, Activity related to a *quantifiable* structural attribute

$$\text{Activity} \approx f(\log P, \text{number of atoms}, \text{shape}....)$$



Currently > 3,000 recognised attributes

<http://www.qsarworld.com/>


# The Alternative to Experiments

Predict likely properties based on **similar** molecules

CHEMBL Database: data on **622,824** compounds,  
collected from **33,956** publications

WOMBAT Database: data on **251,560** structures,  
for over **1,966** targets

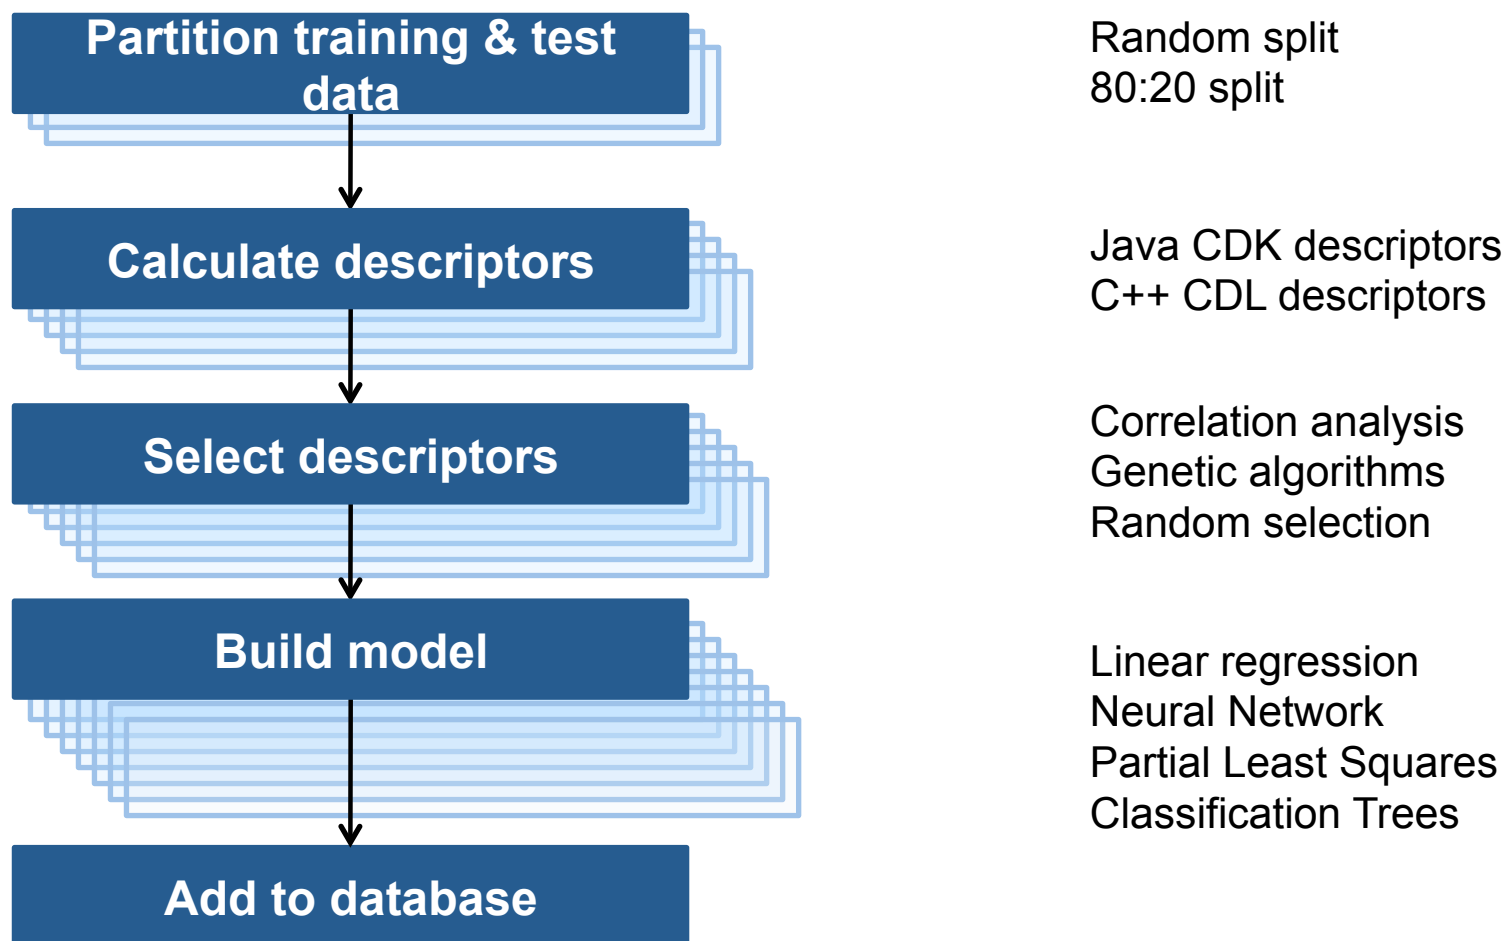
WOMBAT-PK Database: data on **1230** compounds,  
for over **13,000** clinical measurements

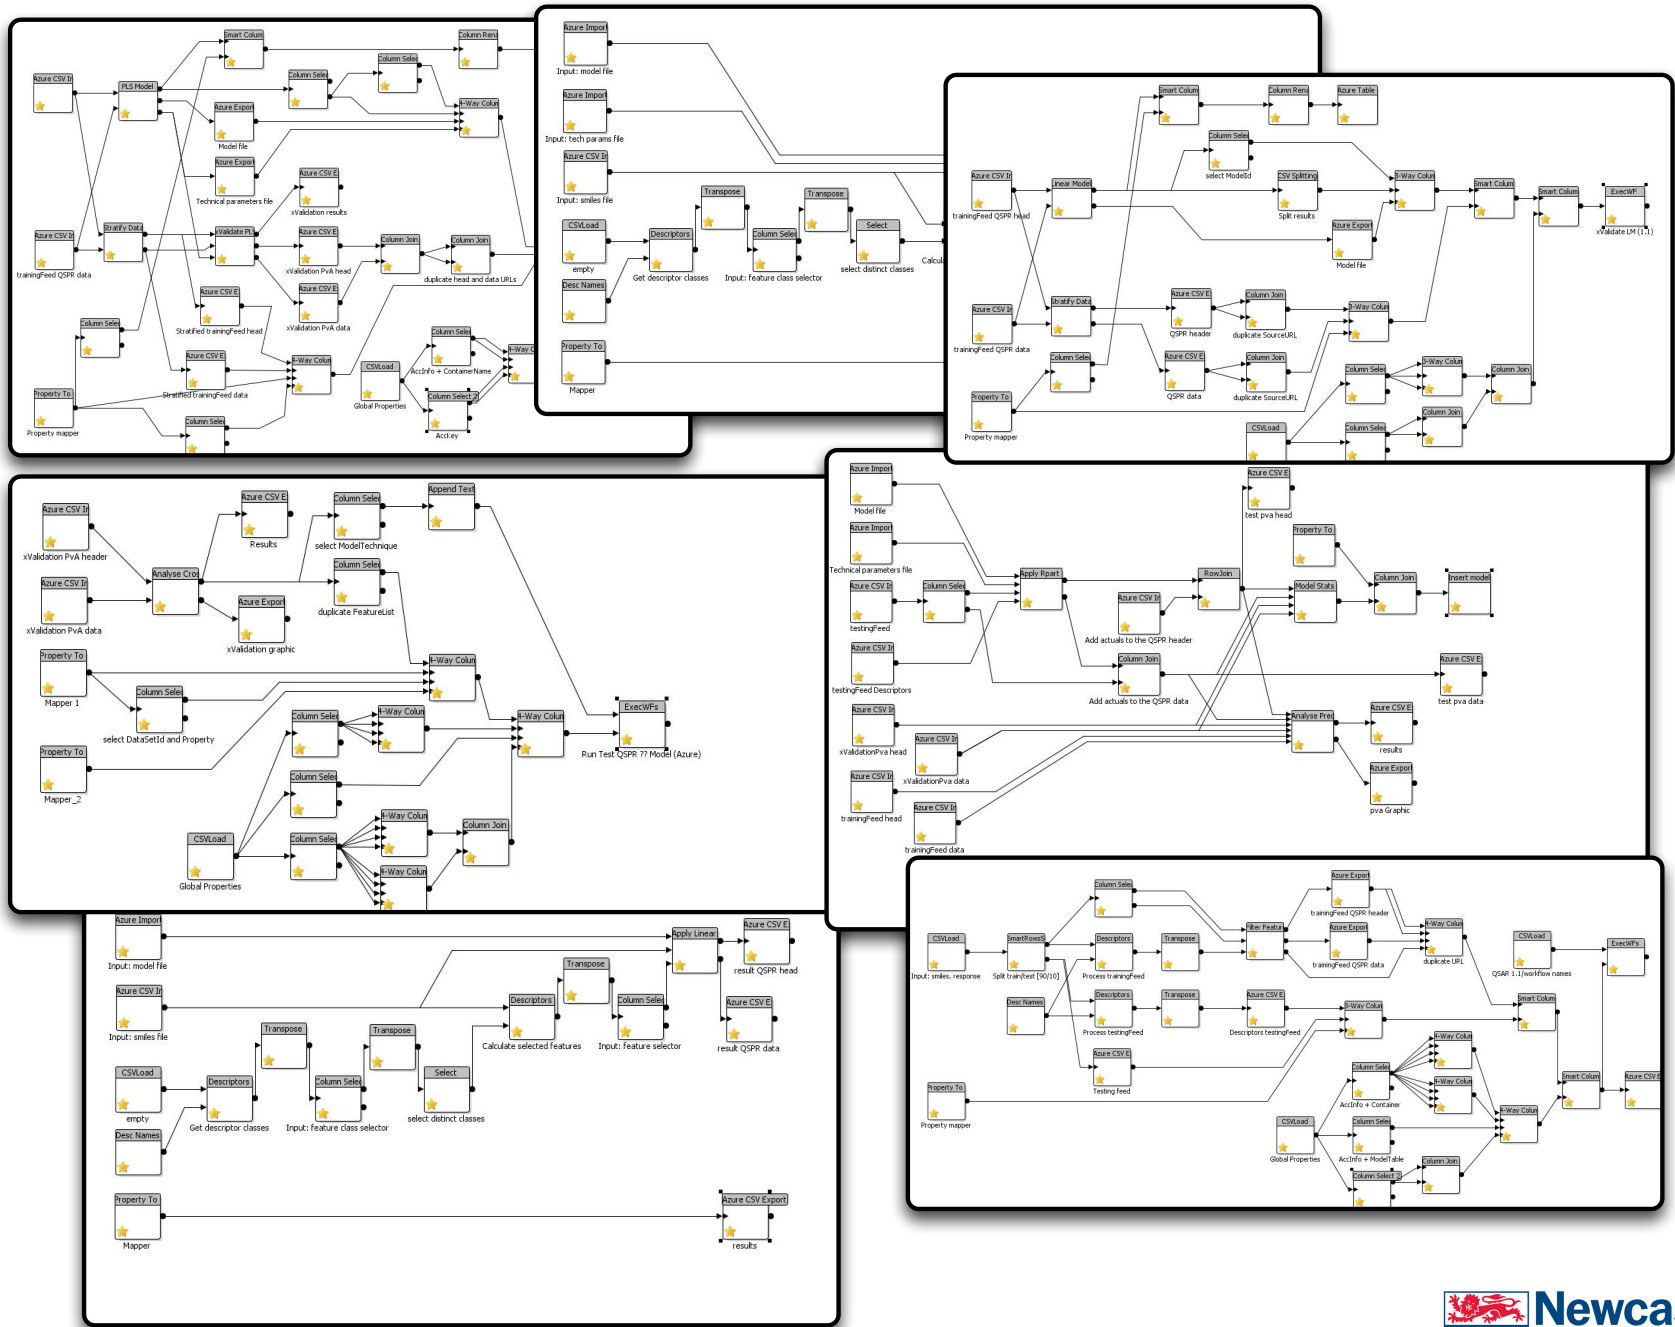


All these databases contain **structure** information and **numerical** activity data

What is the relationship between structure and activity?

# Branching Workflows






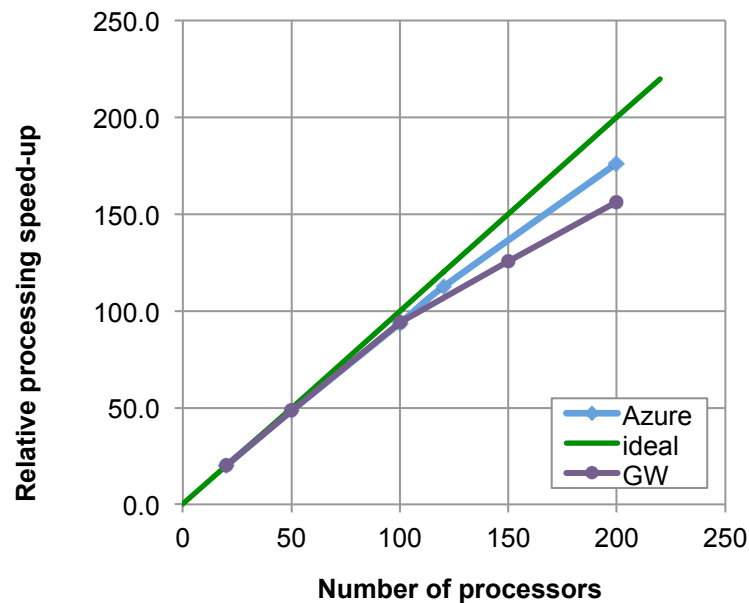
# Results

- 250k models
  - Linear Regression
  - PLS
  - RPartitioning
  - Neural Net
- 460K workflow executions
- 4.4M service calls
- QSAR Explorer
  - Browse
  - Search
  - Get Predictions

# Scalability: Large Scale QSAR

480 datasets sequential time: 11 days

	100 Nodes	200 Nodes
Response Time	3hr 19mins	1hr 50mins
Speedup	94x	156x
Efficiency	94%	78%
Cost	\$55.68	\$51.84



Performance is great but ...

Drug Development requires us to capture  
the **data** and the **process**

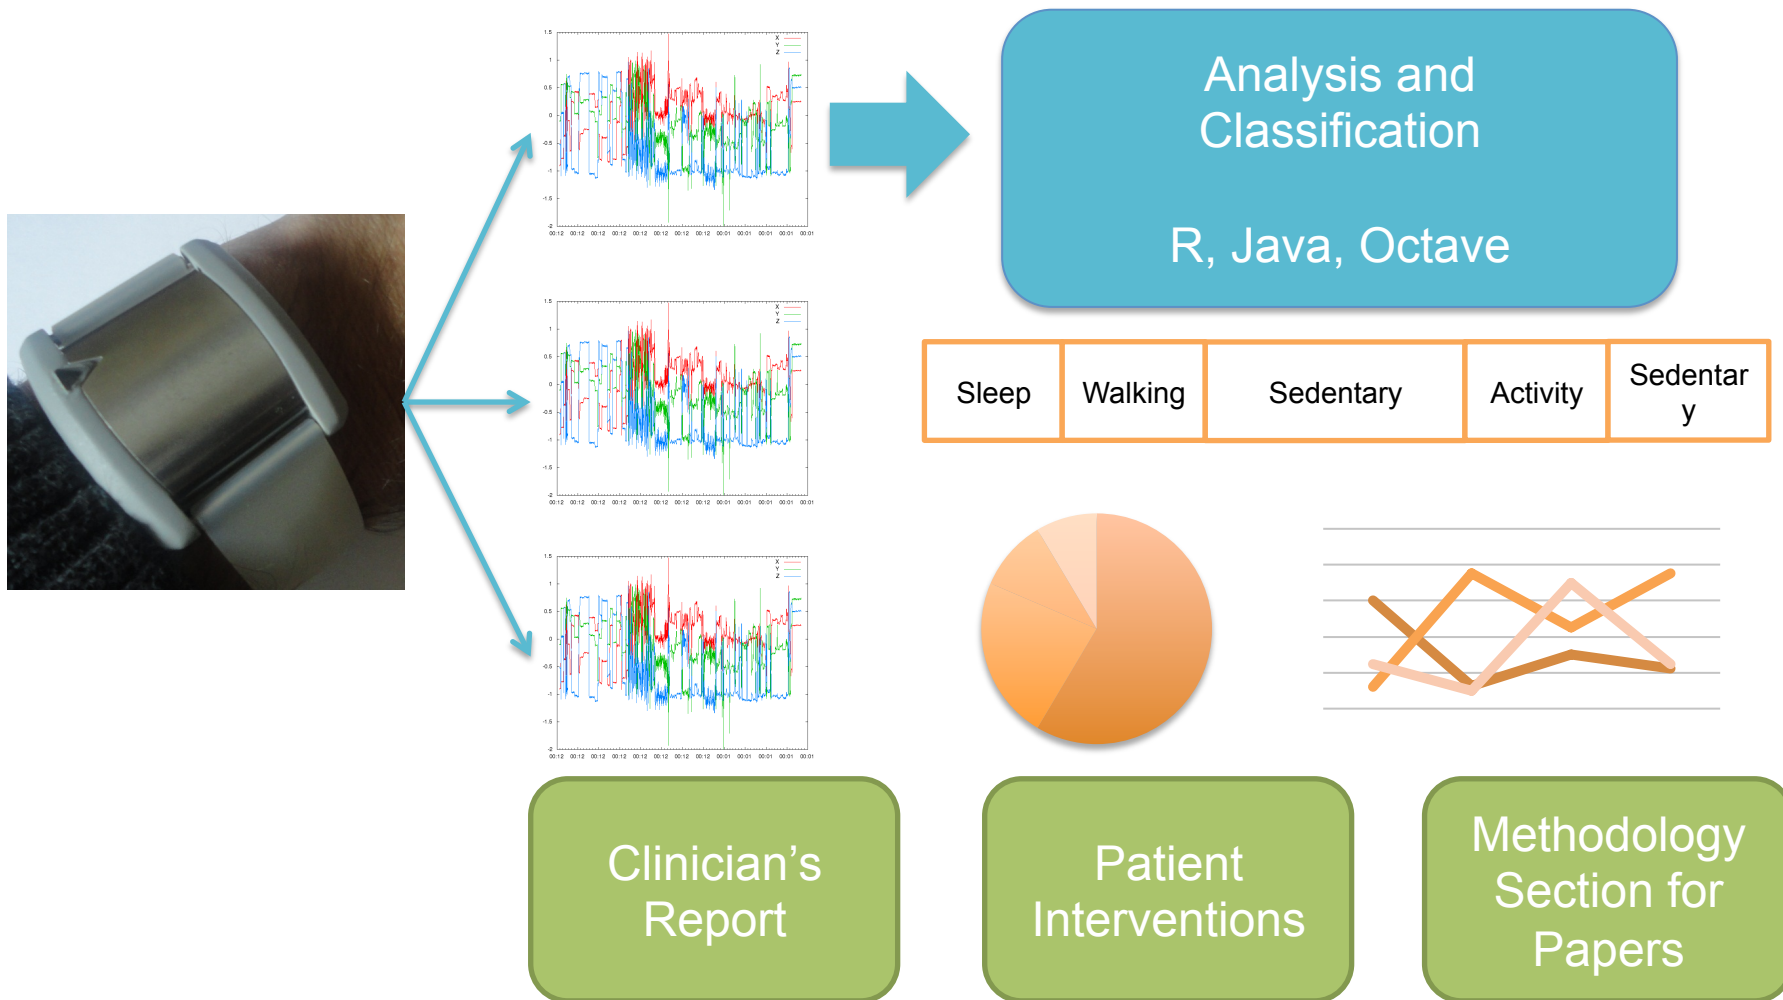


# MOVEeCloud Project

- Investigating the links between physical activity and common diseases – type 2 diabetes, cardiovascular diseases
- Wrist accelerometers worn over period
- Measures movement at 100Hz
- Processing ideal for Azure
  - Bursty data processing as new data
  - Embarrassingly parallel
  - Large datasets



# MOVEeCloud Process



# Data Sizes

100 samples / second

100 rows

3600 seconds / hour

360,000 rows

24 hours / day

8,640,000 rows

7 days / study

60,480,000 rows



/ patient / visit



Cohort size of 800 patients and multiple visits

# Working with larger data sets

- As we add more workflow engines server load increases
  - One server can cope 200 engines if files are small
- This is not the case with movement data
  - Only support 4 engines
- Increase the bandwidth to the engines
  - Clustering appserver /database?

# HDFS

- Implemented prior to Native HDFS on Azure
- Easy to integrate with e-sc
  - Java system just requires libraries included in e-sc
- Distributed store where bandwidth increases with number of machines
  - Bits of data spread around lots of machines
- Concept of data location
  - Potential to route workflows to execute as close as possible to storage
- Other applications also also built on top of HDFS
  - Open TSDB to store timeseries for movement data

# Initial Results

For a single data set processing went from 60 to 16 minutes using 4 workflow engines running HDFS

- 4 engines the limit for one e-sc server
  - Main server hit 100% CPU delivering data
  - No further improvements with more engines
- Using HDFS CPU was consistently below 5%
  - More like our earlier scalability results
- Once data had been chunked processing was the same for each chunk
- The improvement lay entirely in staging and uploading results

# Upcoming Challenges

- Process Newcastle 85+ and Whitehall Study Data
  - 6TB
- New TSB Project with RedHat and Arjuna
  - e-Science Central onto OpenShift
  - Integrate Arjuna Agility
  - Analyse traffic flow data from Newcastle

# Demo