Carbon Accounting for UK Research Computing: proof of concept & beyond

NetDRIVE

Measuring Data Centre Carbon Costs

Allocating Carbon Costs to Payloads

IPMI Vs PDU

What can we learn?

Think like an Accountant!

Speaker: Dr Alex Owen

r.a.owen@qmul.ac.uk

Netzero DRI Vision & Expertise

https://eng.ox.ac.uk/netdrive

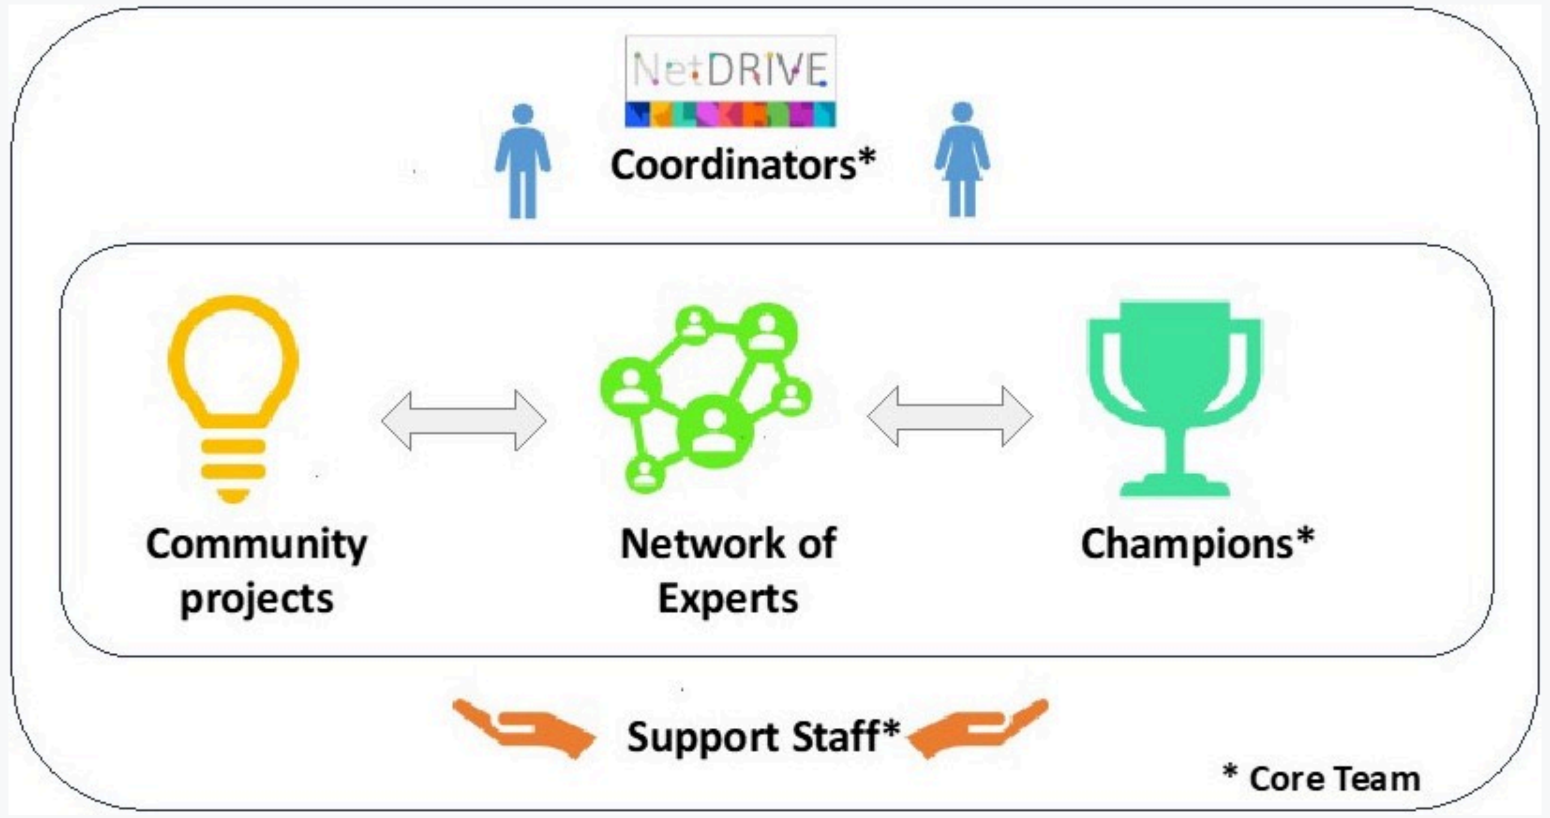

Actionable, authoritative advice for UKRI & Community

Coordinators	Martin Juckes, Sarah Sparrow
Champions	Loïc Lannelongue, Michael Rudyard, **Alex Owen**, Kirsty Pringle, Andy Turner,Lorna Smith, Jessica Huntley, Erinma Ochu, Liz Ing-Simmons
Core Team	Amani Maalouf, Graeme Smith, Jole Cutting

13 Community Projects

- Delivering training and generating an evidence base and actionable resources

3 Workingroups

4 Workshops

3 Community Meetings

1 Early Career Sandpit

Procurement WG - reported

Metrics WG **Training Materials WG**

Projects - March 2027

NetDRIVE 2025 - March 2028 ~£3M

Measuring Data Centre Carbon Costs

2023

IRISCAST
Audit of Carbon Costs

Jon Hays (QMUL)
Nic Walton (Cambridge)
Adrian Jackson (Edinburgh)
Alison Packer (STFC)

Alex Owen (QMUL)
Alex Ogden (Cambridge)
Anish Mudaraddi (STFC)

Dan Traynor (QMUL)
Derek Ross (STFC)
Alexander Dibbo (STFC)
Jon Roddom (STFC)
Martin Summers (STFC)
Jacob Ward (STFC)
Dan Whitehouse (Imperial)
Alastair Basden (Durham)

https://doi.org/10.5281/zenodo.7692451
https://doi.org/10.1051/epjconf/202429507015

Queen Mary University of London
Science and Technology Facilities Council | Scientific Computing
epcc
UNIVERSITY OF CAMBRIDGE

Allocating Carbon Costs to Payloads

2024

iris CMP

Payload = Batch Job or Cloud VM
GPU’s were out of scope

Alex Owen (QMUL)
Jon Hays (QMUL)
Daohai Li (QMUL)
Alex DIbbo (STFC)

https://doi.org/10.5281/zenodo.10966001
https://doi.org/10.1051/epjconf/202533701202

Queen Mary University of London
Science and Technology Facilities Council | Scientific Computing
GridPP

iris

https://www.iris.ac.uk/

eInfrastructure for Research and Innovation for STFC

IRIS is a cooperative community brining together (mainly) STFC computing interests

Works closely with STFC but run by the community

Formed bottom up by science communities and compute providers

IRIS Science Director is Prof J Hays @QMUL

What is IRISCAST?

IRIS Carbon Audit SnapshoT
24 Hour snapshot

The Challenge

Estimate carbon costs for scientific computing across a broad heterogeneous landscape

Identifying the key drivers for carbon costs

Identifying the hurdles and barriers

Communicating the carbon costs to drive change

Working coherently across different communities

The Project

Work together coherently across different facilities with different remits, tooling, and capabilities.

Learn by doing!

Document the gaps, the barriers and the issues, drive requirements for future work and decision making

Communicate across our communities and build a foundation for future action

Carbon Model

IRISCAST

Audit of Carbon Costs

$$C_t^p = \boxed{C_a^p} + \boxed{C_e^p}$$

Carbon Cost (C) for a period (p) is sum of active carbon (a) and embedded carbon (e)

$$\boxed{C_a^p} = CM_e^p \left(E_{nodes}^p + E_{network}^p + \boxed{E_{cooling}^p} + E_{building}^p \right)$$

Carbon Intensity of Power (Grid)

Measure Energy Usage (E)
Obtain Carbon Grid Intensity (CM)

UPDATE

We can gather Carbon Intensity at 30 minute intervals in the UK, so should probably measure energy usage at 30 min intervals too.

$$\boxed{C_e^p} = p\left(\boxed{\frac{C_{e\_node1}}{L_{node1}} + \frac{C_{e\_node2}}{L_{node2}} + \ldots + \frac{C_{e\_network}}{L_{network}} + \ldots + \frac{C_{e\_building}}{L_{building}} + \ldots} \right)$$

Q_{ef} **Embedded Carbon Amortisation Rate**

Inventory of Equipment
Obtain Embedded Carbon
Estimate Lifetime (L)

UPDATE

We have started to think about a relational database model for the carbon inventory and to track cluster configuration. But nothing to share yet!

24 Hour Snapshot at Six Facilities

IRISCAST
Audit of Carbon Costs

Learning By Doing

Facilities

Cambridge IRIS HPC/Cloud
STFC SCD Cloud
STFC SCARF
QMUL GridPP Tier 2
Imperial GridPP Tier 2
DiRAC (Durham)

Good embedded data are hard to find!

Summary Inventories

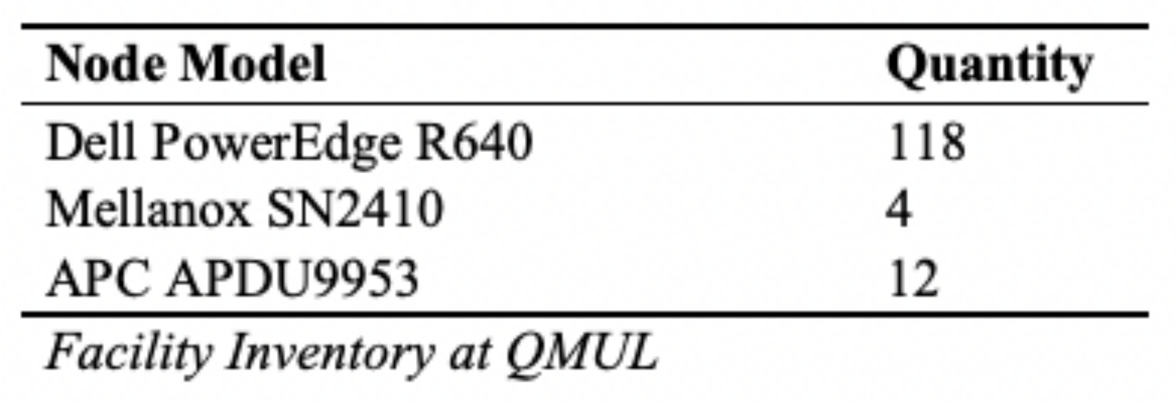

Node Model	Quantity
Dell PowerEdge R640	118
Mellanox SN2410	4
APC APDU9953	12

Facility Inventory at QMUL

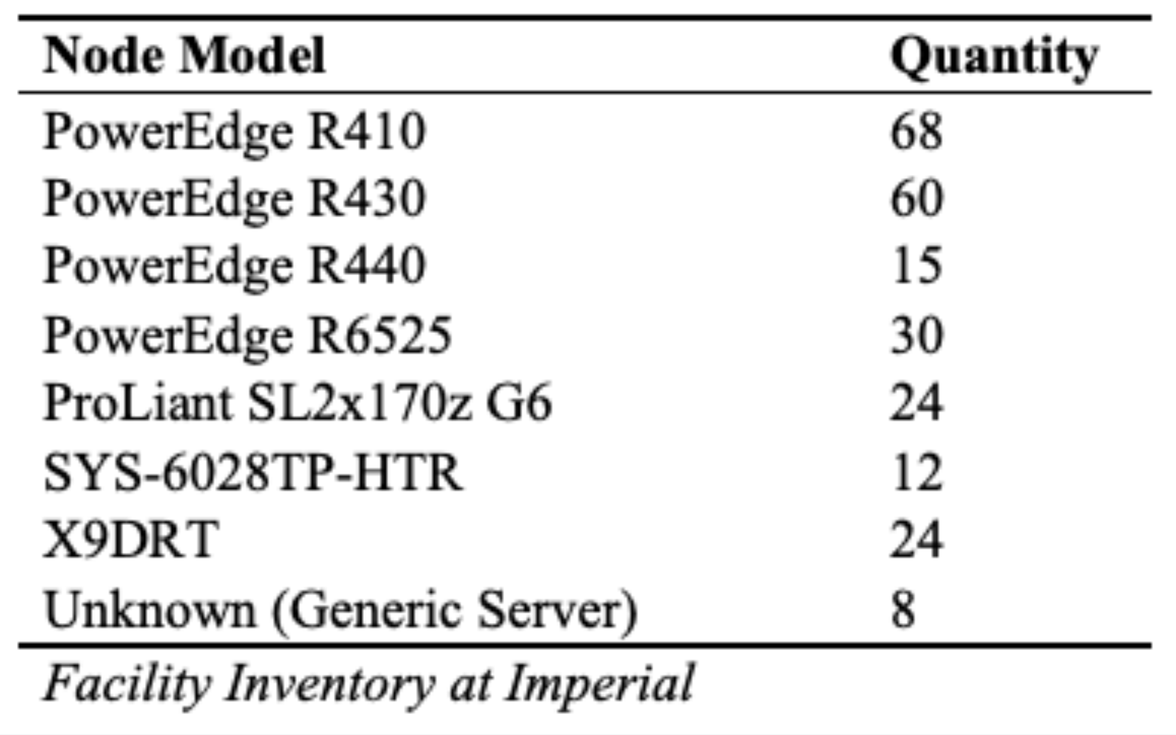

Node Model	Quantity
PowerEdge R410	68
PowerEdge R430	60
PowerEdge R440	15
PowerEdge R6525	30
ProLiant SL2x170z G6	24
SYS-6028TP-HTR	12
X9DRT	24
Unknown (Generic Server)	8

Facility Inventory at Imperial

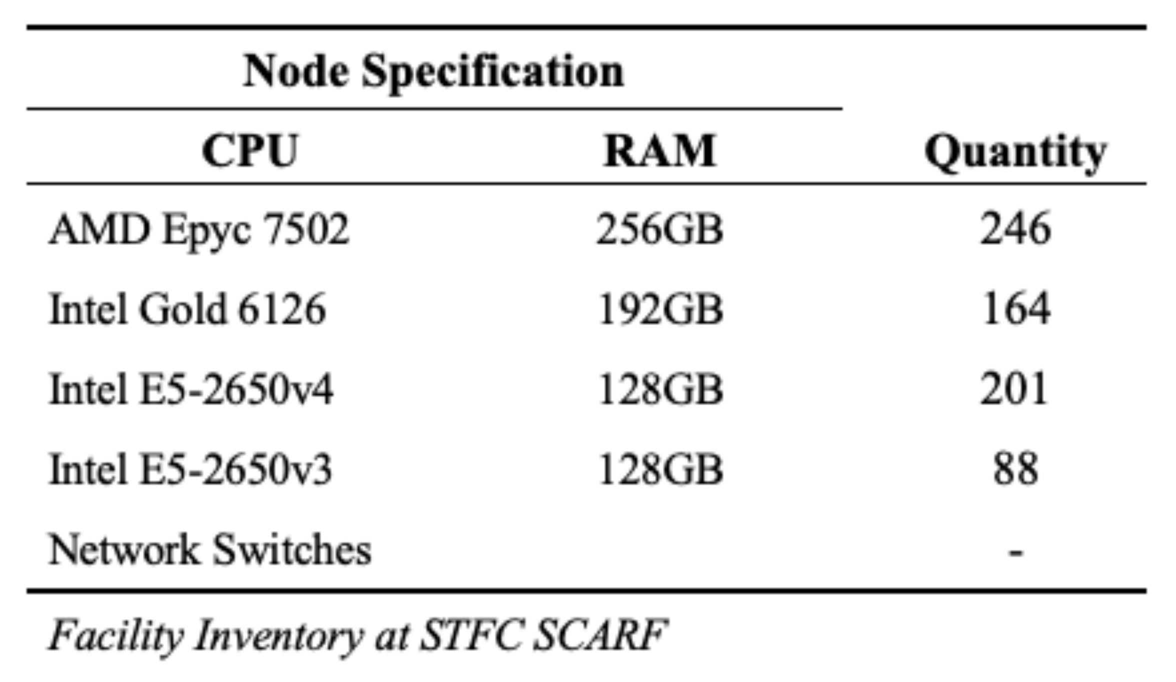

Node Specification		
CPU	RAM	Quantity
AMD Epyc 7502	256GB	246
Intel Gold 6126	192GB	164
Intel E5-2650v4	128GB	201
Intel E5-2650v3	128GB	88
Network Switches		-

Facility Inventory at STFC SCARF

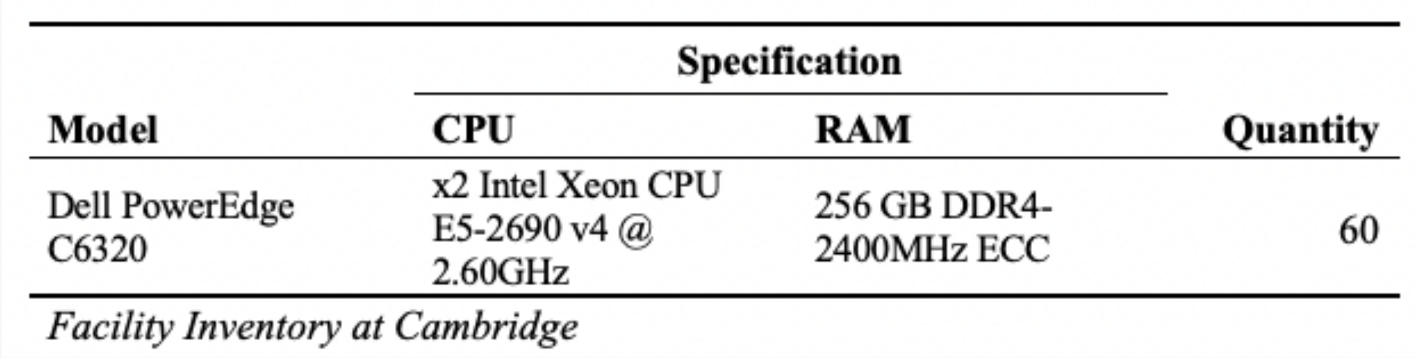

	Specification		
Model	CPU	RAM	Quantity
Dell PowerEdge C6320	x2 Intel Xeon CPU E5-2690 v4 @ 2.60GHz	256 GB DDR4-2400MHz ECC	60

Facility Inventory at Cambridge

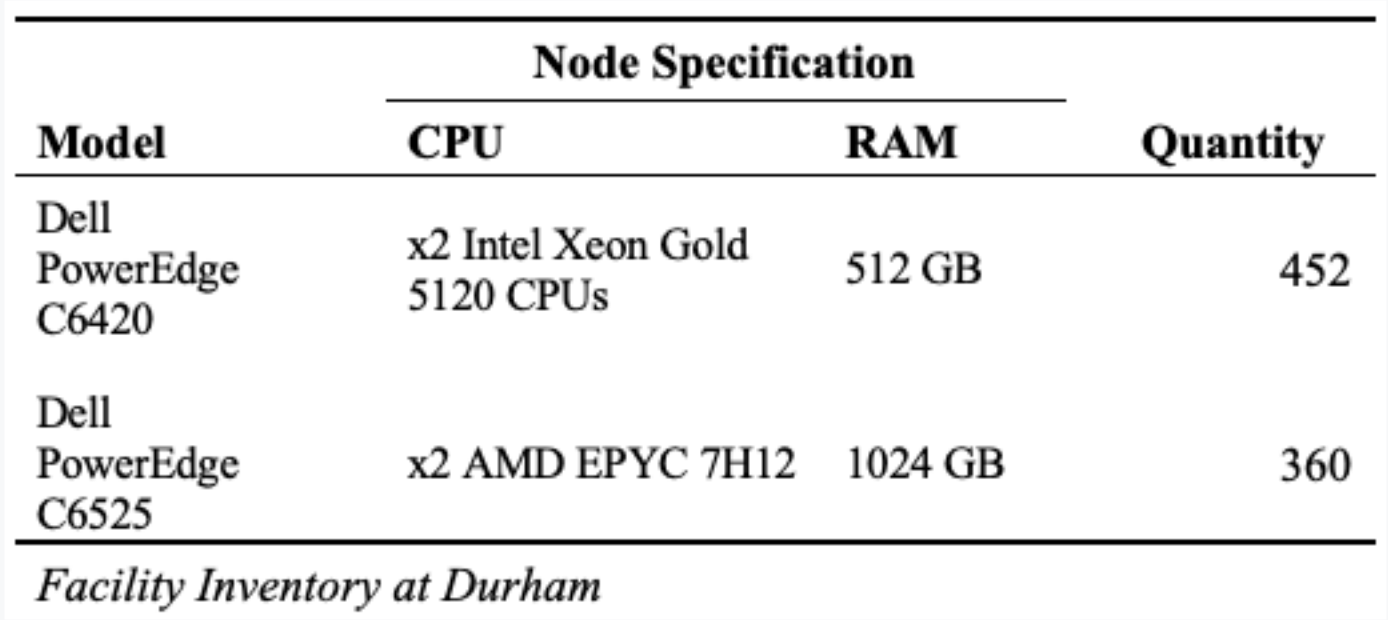

	Node Specification		
Model	CPU	RAM	Quantity
Dell PowerEdge C6420	x2 Intel Xeon Gold 5120 CPUs	512 GB	452
Dell PowerEdge C6525	x2 AMD EPYC 7H12	1024 GB	360

Facility Inventory at Durham

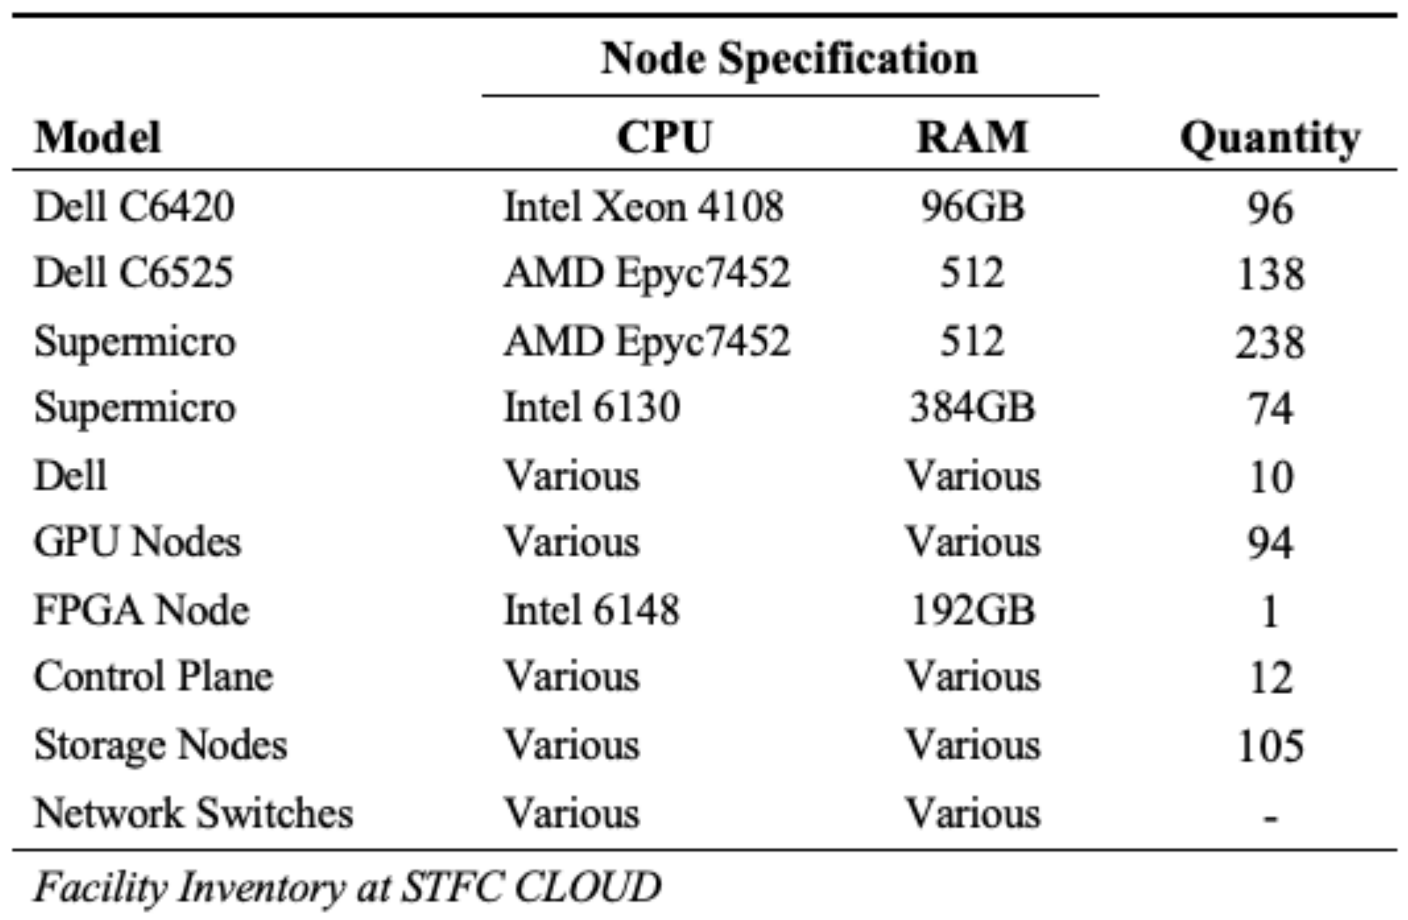

	Node Specification		
Model	CPU	RAM	Quantity
Dell C6420	Intel Xeon 4108	96GB	96
Dell C6525	AMD Epyc7452	512	138
Supermicro	AMD Epyc7452	512	238
Supermicro	Intel 6130	384GB	74
Dell	Various	Various	10
GPU Nodes	Various	Various	94
FPGA Node	Intel 6148	192GB	1
Control Plane	Various	Various	12
Storage Nodes	Various	Various	105
Network Switches	Various	Various	-

Facility Inventory at STFC CLOUD

2022

Energy
Vs
Power

For active carbon we need to know about ENERGY usage

Time Stamped Energy usage is more robust than instantaneous power

APC PDU's SNMP query	✓	`PowerNet-MIB::rPDU2DeviceStatusEnergy.1`	**Hoctowatthour (100 wh) units**
APC port n SNMP query	✓	`PowerNet-MIB::rPDU2OutletMeteredStatusEnergy.n`	
IPMI query (freeipmi)	✓	`ipmi-oem dell get-power-consumption-data`	
IPMI query (ipmitool)	?	`ipmitool sensor list`	**Probably Power Measurements Energy is better!**

PDU
Vs
IPMI

IRISCAST
Audit of Carbon Costs

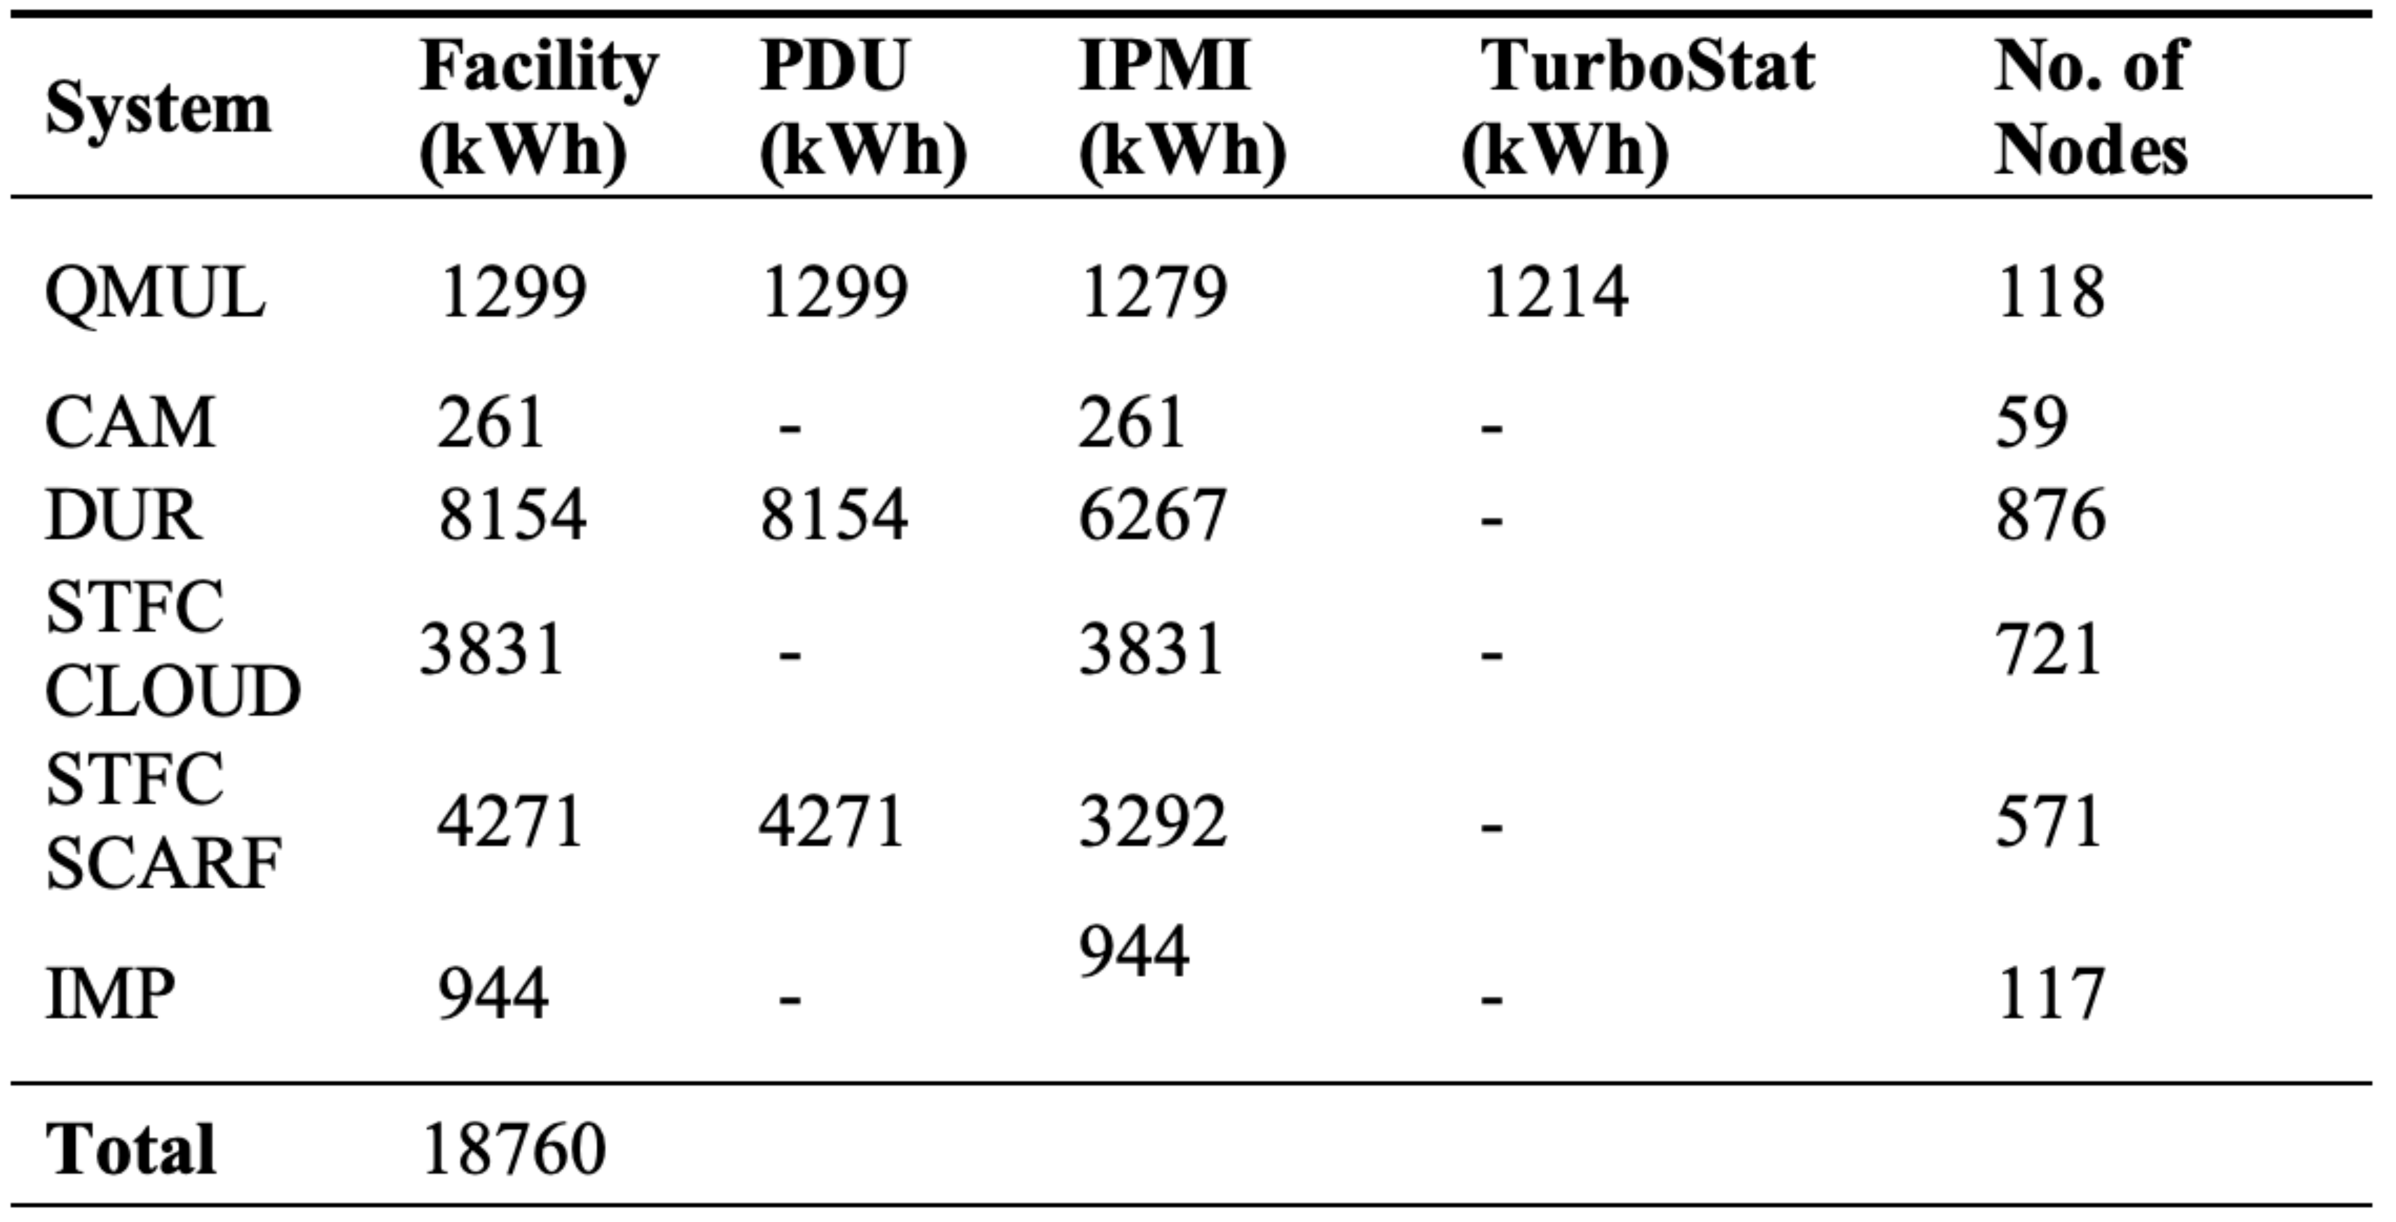

System	Facility (kWh)	PDU (kWh)	IPMI (kWh)	TurboStat (kWh)	No. of Nodes
QMUL	1299	1299	1279	1214	118
CAM	261	-	261	-	59
DUR	8154	8154	6267	-	876
STFC CLOUD	3831	-	3831	-	721
STFC SCARF	4271	4271	3292	-	571
IMP	944	-	944	-	117
Total	18760				

IPMI approx 20% low cf PDU Except at QMUL where 1.5% (APDU9953)

More questions than answers: Check your calibrations!

UPDATE

IPMI protocol is standardised. The sensors are not!
Measurement details are model specific.

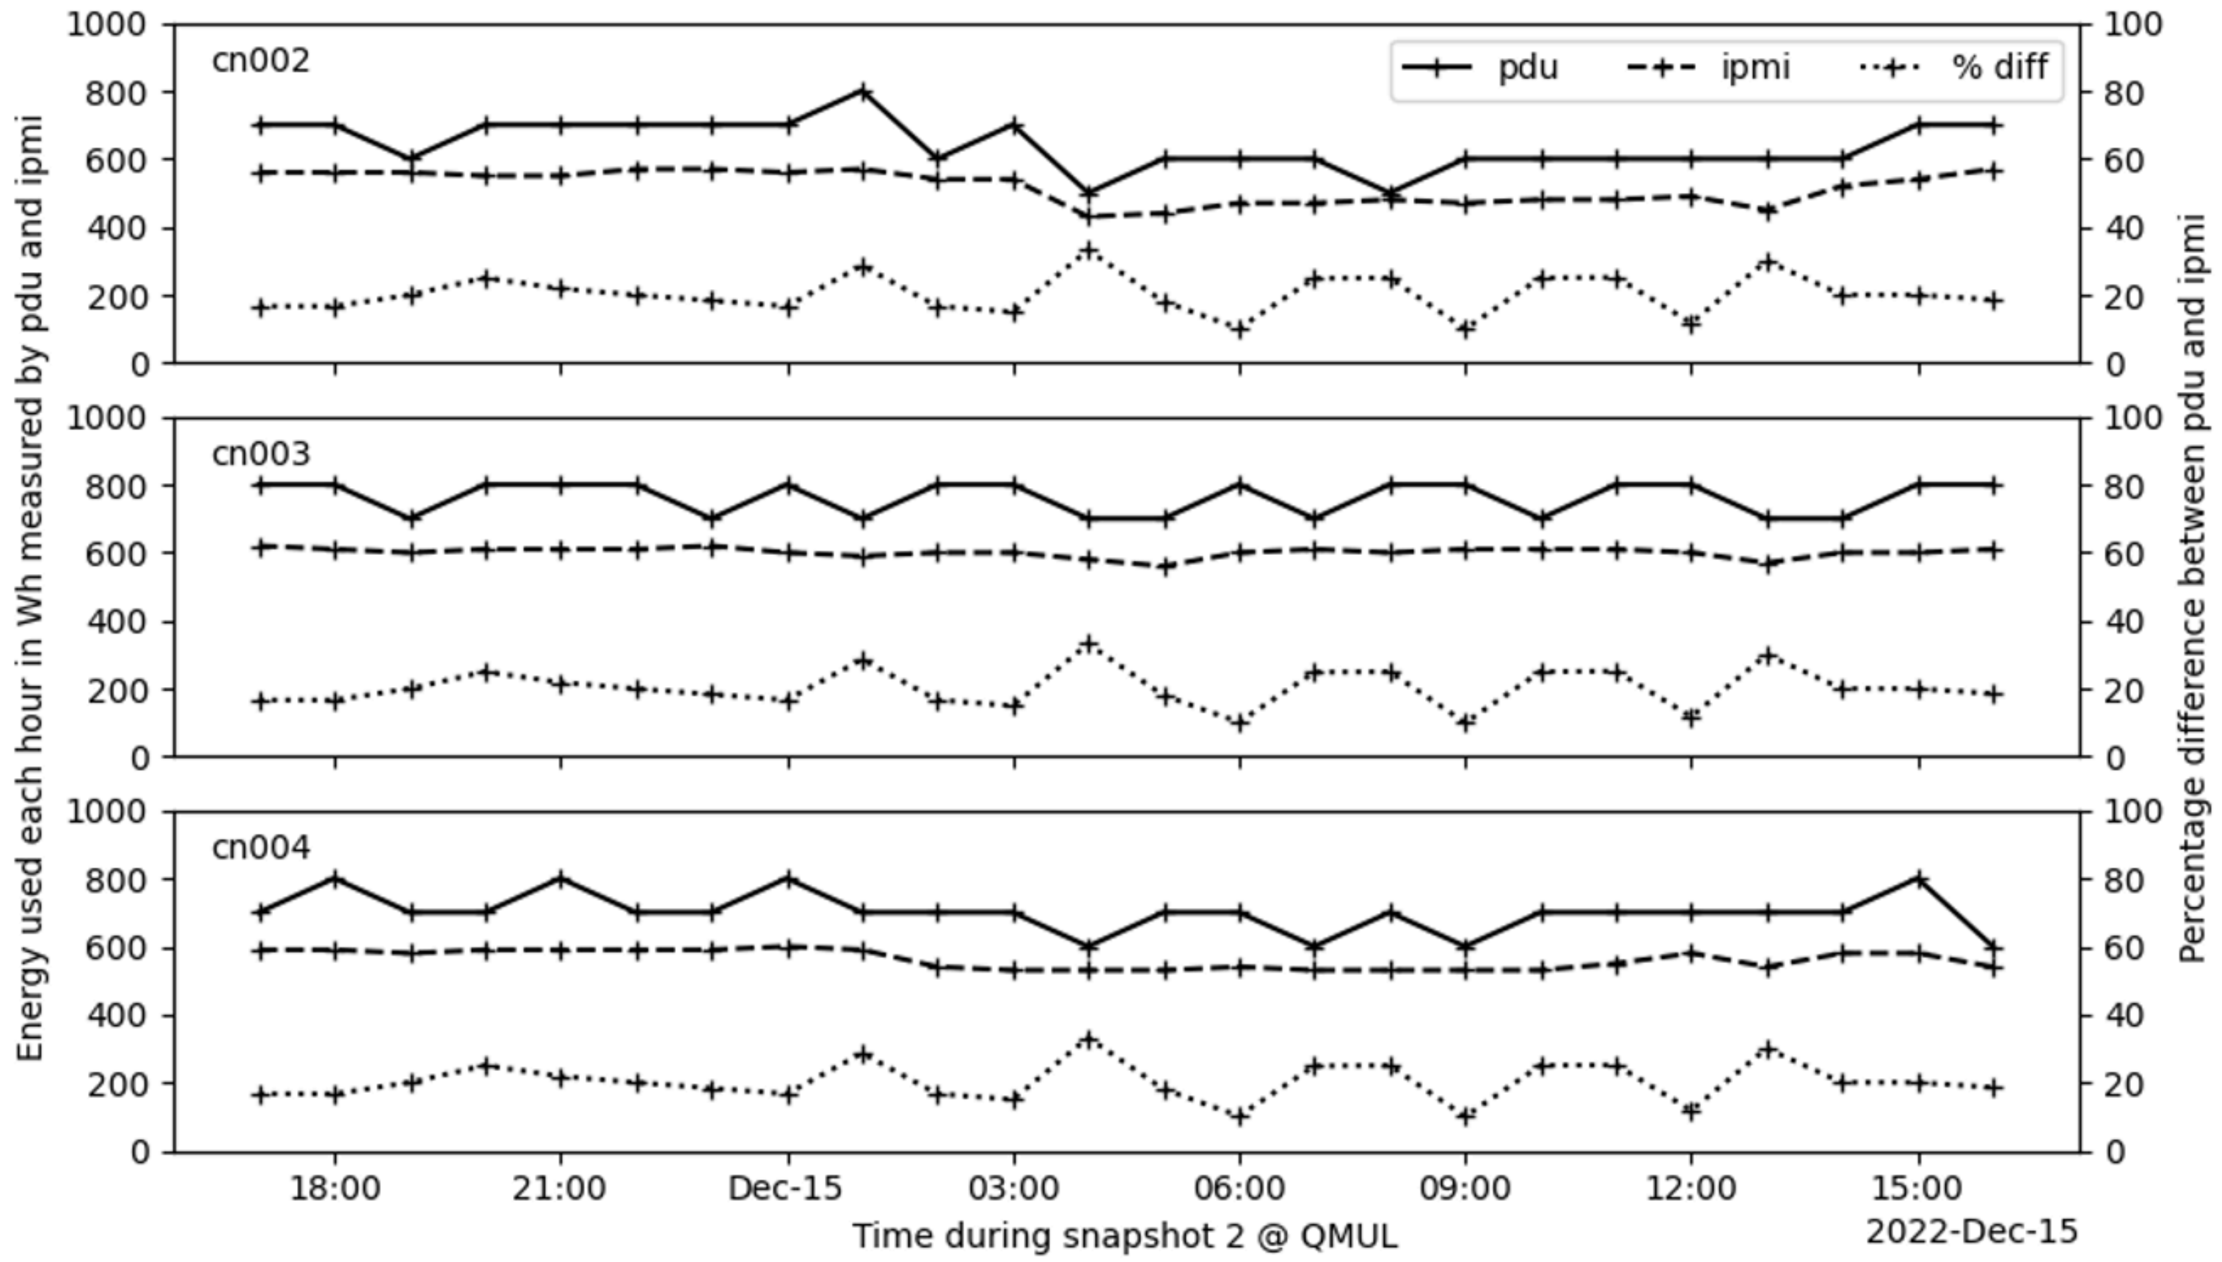

20% difference between APC AP8459WW port measurements and IPMI measurements

Re test at QMUL with AP8459WW per port PDU 20% difference!

Carbon Model

Model a range of scenarios

IRISCAST

Audit of Carbon Costs

Measuring computer energy usage is the easy bit.

Cooling energy usage/PUE less well known.

Computer embedded carbon figures hard to find.

Other equipment embedded carbon figures even harder to find.

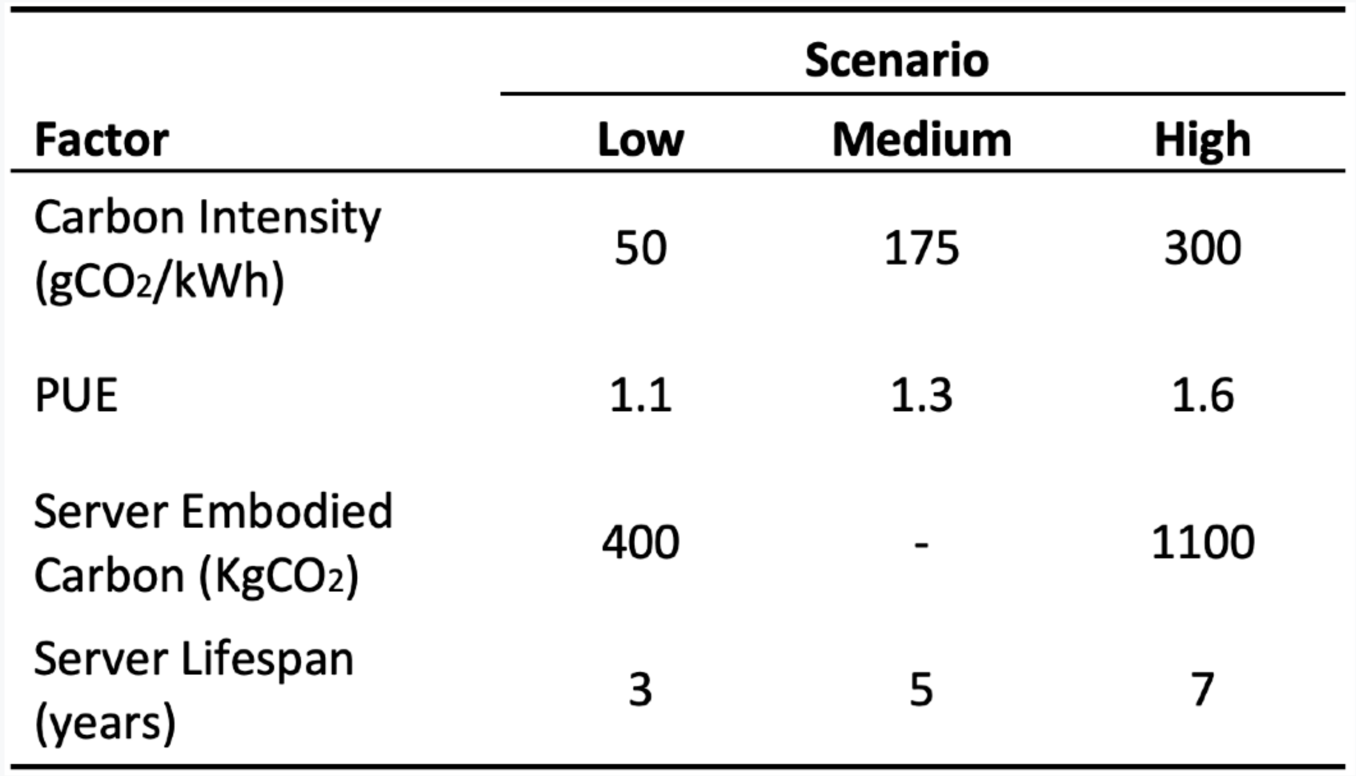

	Scenario		
Factor	**Low**	**Medium**	**High**
Carbon Intensity (gCO_2/kWh)	50	175	300
PUE	1.1	1.3	1.6
Server Embodied Carbon ($KgCO_2$)	400	-	1100
Server Lifespan (years)	3	5	7

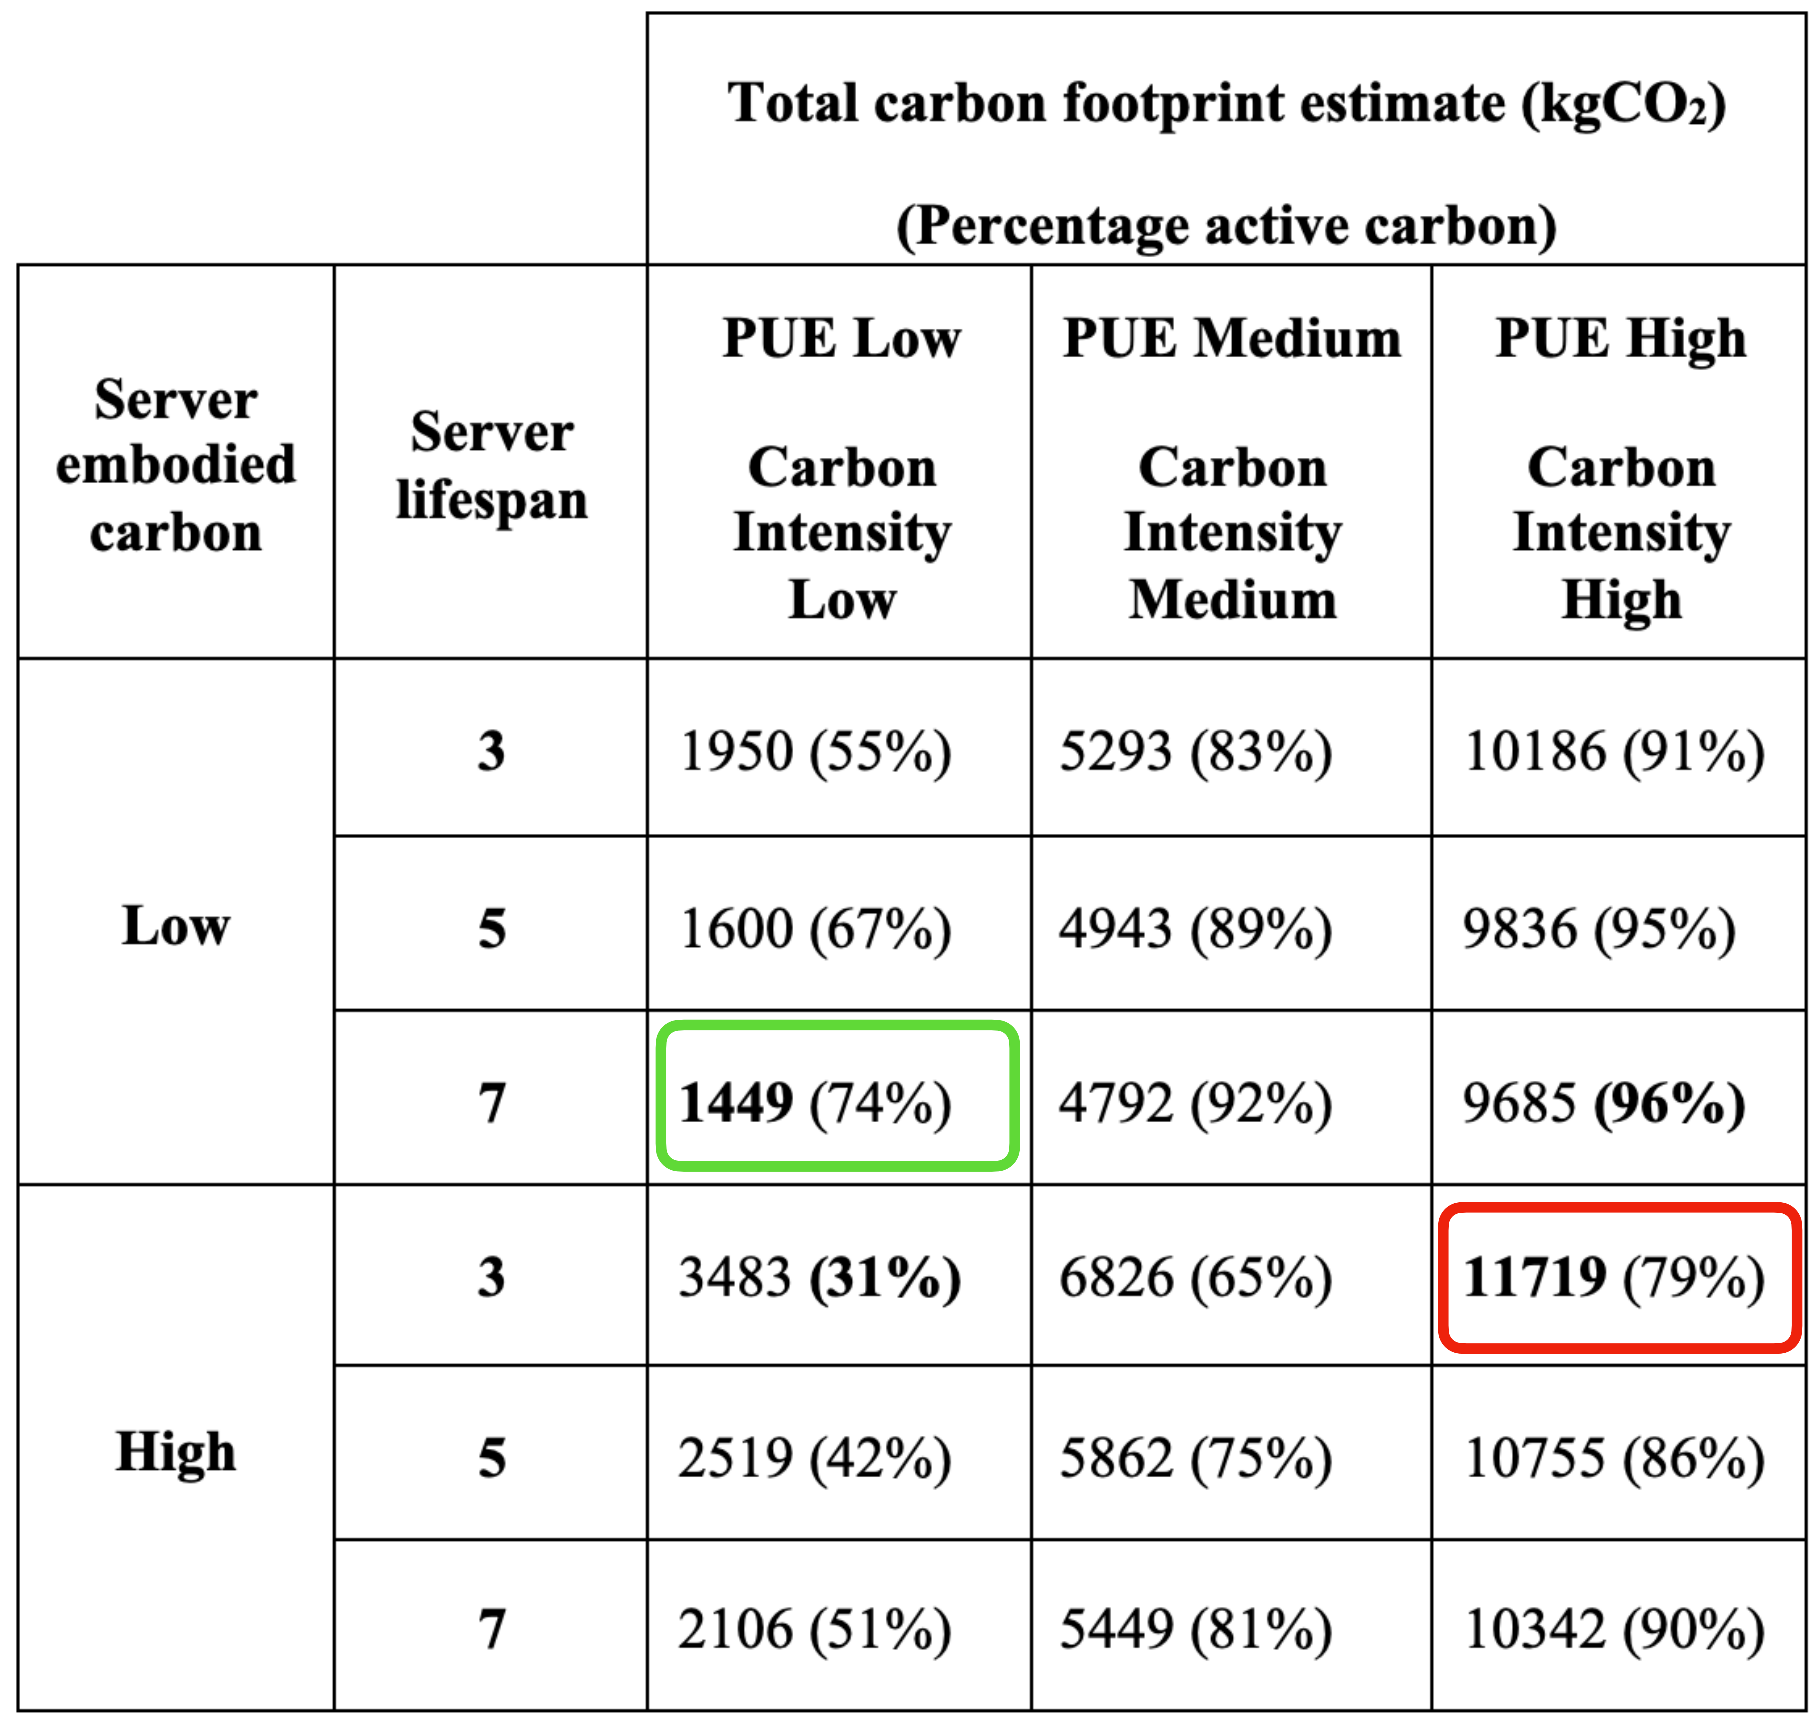

		Total carbon footprint estimate ($kgCO_2$) (Percentage active carbon)		
Server embodied carbon	**Server lifespan**	**PUE Low Carbon Intensity Low**	**PUE Medium Carbon Intensity Medium**	**PUE High Carbon Intensity High**
Low	**3**	1950 (55%)	5293 (83%)	10186 (91%)
	5	1600 (67%)	4943 (89%)	9836 (95%)
	7	**1449** (74%)	4792 (92%)	9685 **(96%)**
High	**3**	3483 **(31%)**	6826 (65%)	**11719** (79%)
	5	2519 (42%)	5862 (75%)	10755 (86%)
	7	2106 (51%)	5449 (81%)	10342 (90%)

IRISCAST 24 hour snapshot roughly 1-4 people on 12 hour return Jet

Potential to reduce carbon emissions by an order of magnitude!

We can measure Carbon Costs of a computer service ✓

IRISCAST Proposes

High Level Feedback

Carbon Equivalent per month

Low Level Feedback

Figure of merit per Job → Carbon Equivalent per Job

1. Future DRI procurement to include a score based on embedded carbon costs and equipment energy usage.
2. New computer hardware to include energy measurement capability such as IPMI (or per port PDUs) and require the supplier to provide best estimates of embedded carbon costs.
3. Measure energy used by cooling infrastructure and the computing infrastructure.
4. Facilities to keep an inventory of equipment including embedded carbon cost and idle power draw.
5. Monthly (or other periodic) reporting of carbon usage by facilities based on 3 and 4 above. Roll into standard grant reporting regime.
6. Collect per job (or VM) energy usage by using tools like Slurm (correctly configured). Combined this with embedded carbon from inventory and electricity carbon intensity to feedback job carbon cost to the end user to drive improvements in user code and workflow.
7. Identify user communities and the authors of community codebases so that useful feedback can be given to them to drive the development of more efficient code and workflows.

We can allocate Carbon Costs to Payloads

?

iris Carbon Mapping Project

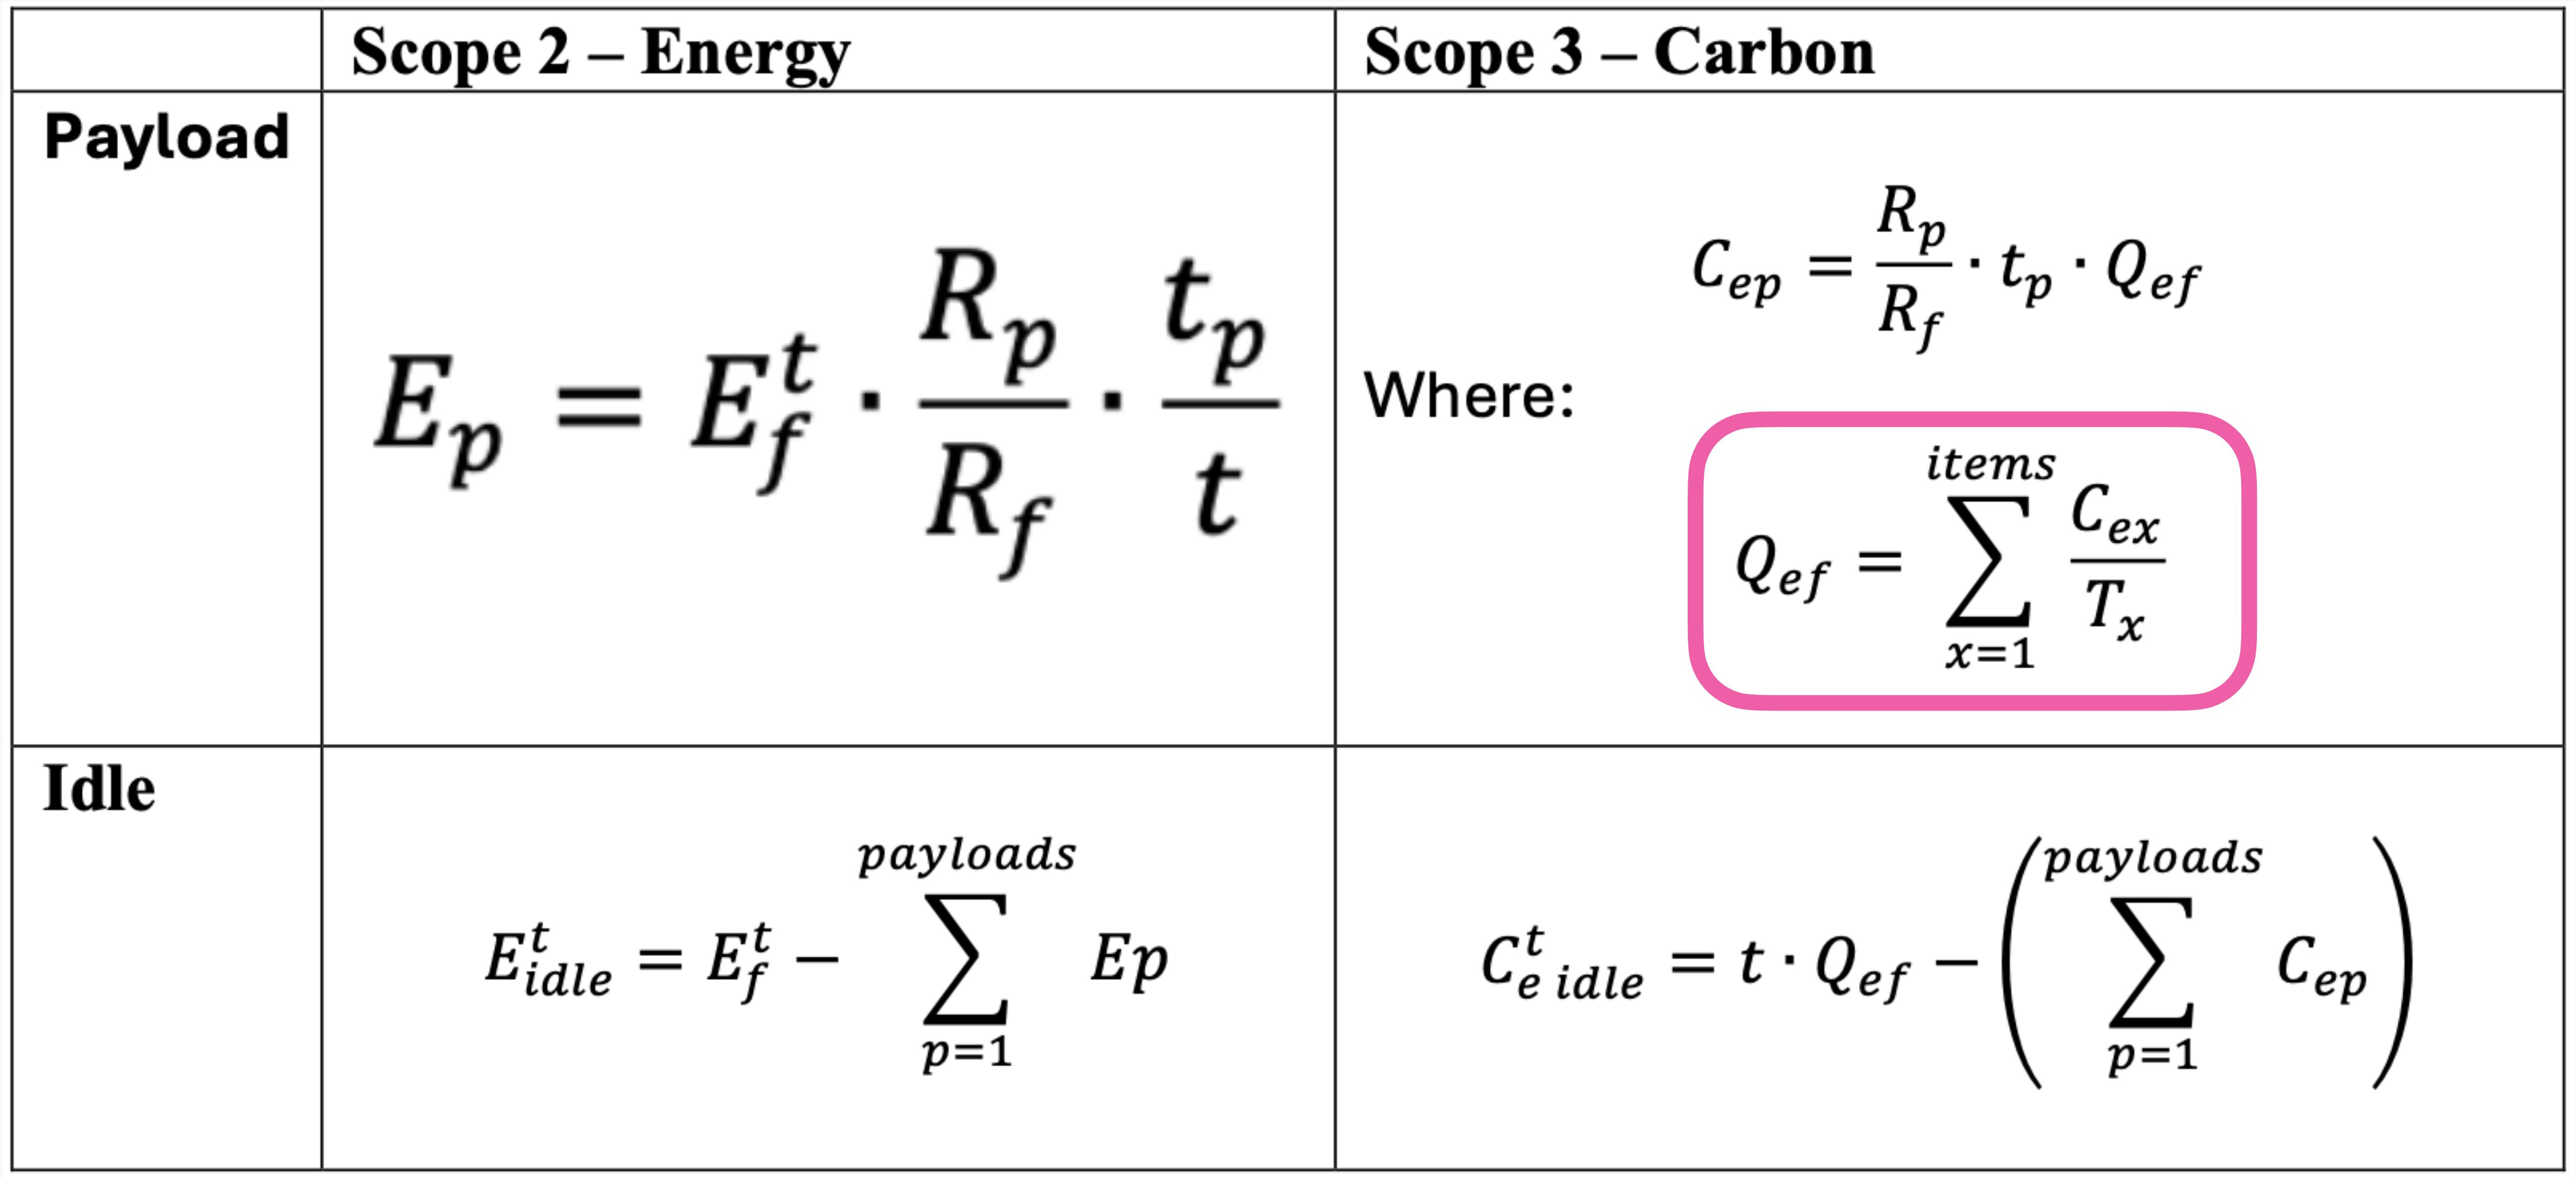

	Scope 2 – Energy	Scope 3 – Carbon
Payload	$E_p = E_f^t \cdot \frac{R_p}{R_f} \cdot \frac{t_p}{t}$	$C_{ep} = \frac{R_p}{R_f} \cdot t_p \cdot Q_{ef}$ Where: $Q_{ef} = \sum_{x=1}^{items} \frac{C_{ex}}{T_x}$
Idle	$E_{idle}^t = E_f^t - \sum_{p=1}^{payloads} Ep$	$C_{e\,idle}^t = t \cdot Q_{ef} - \left(\sum_{p=1}^{payloads} C_{ep} \right)$

Table 1: Summary of the Simple Payload Model showing allocations of Scope 2 energy and Scope 3 carbon to user payloads and the remaining idle allocation to the provider.

Allocating Carbon Costs to Payloads

CPU Payloads on Batch & Cloud

Simple Payload Model

Apportion by Real Time

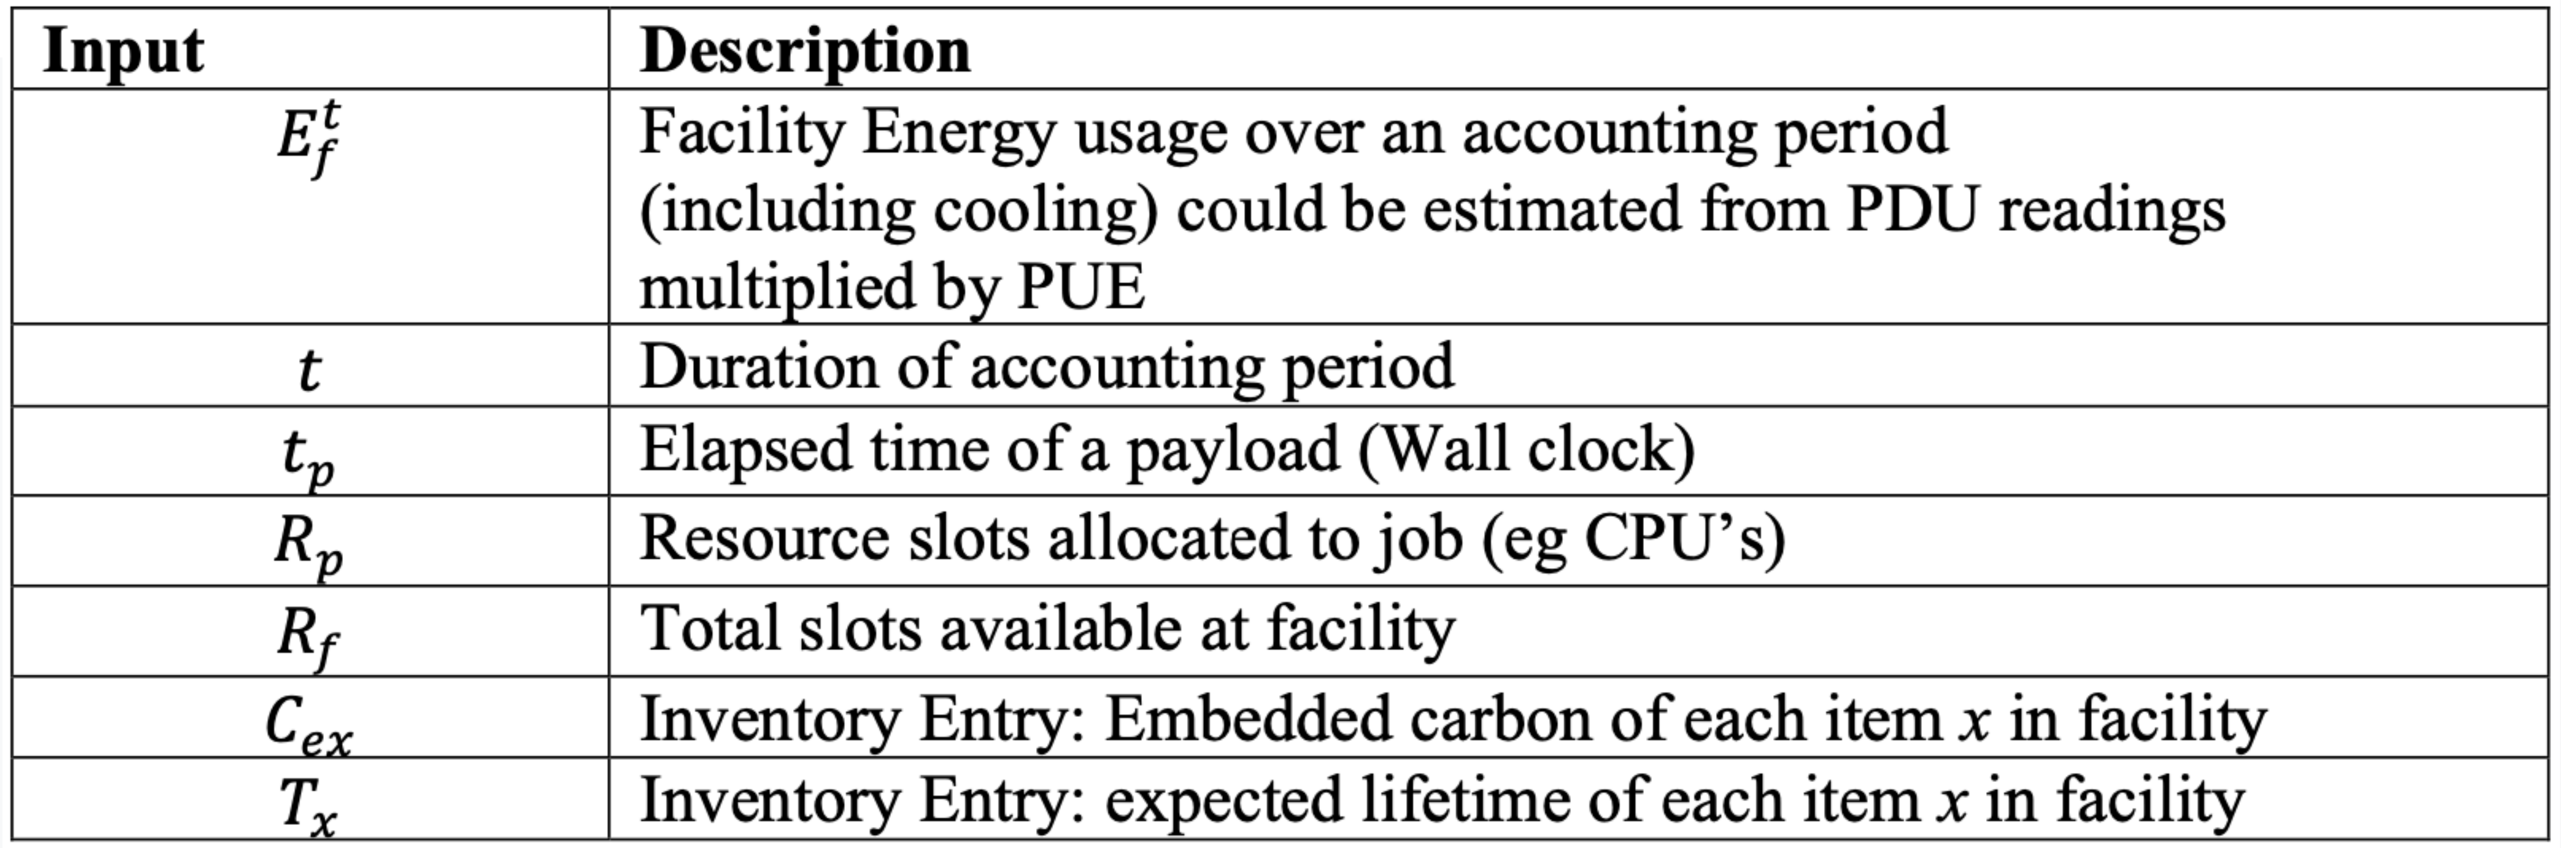

Input	Description
E_f^t	Facility Energy usage over an accounting period (including cooling) could be estimated from PDU readings multiplied by PUE
t	Duration of accounting period
t_p	Elapsed time of a payload (Wall clock)
R_p	Resource slots allocated to job (eg CPU's)
R_f	Total slots available at facility
C_{ex}	Inventory Entry: Embedded carbon of each item x in facility
T_x	Inventory Entry: expected lifetime of each item x in facility

Table 2: Summary of the inputs needed to evaluate the Simple Payload Model.

iris Carbon Mapping Project

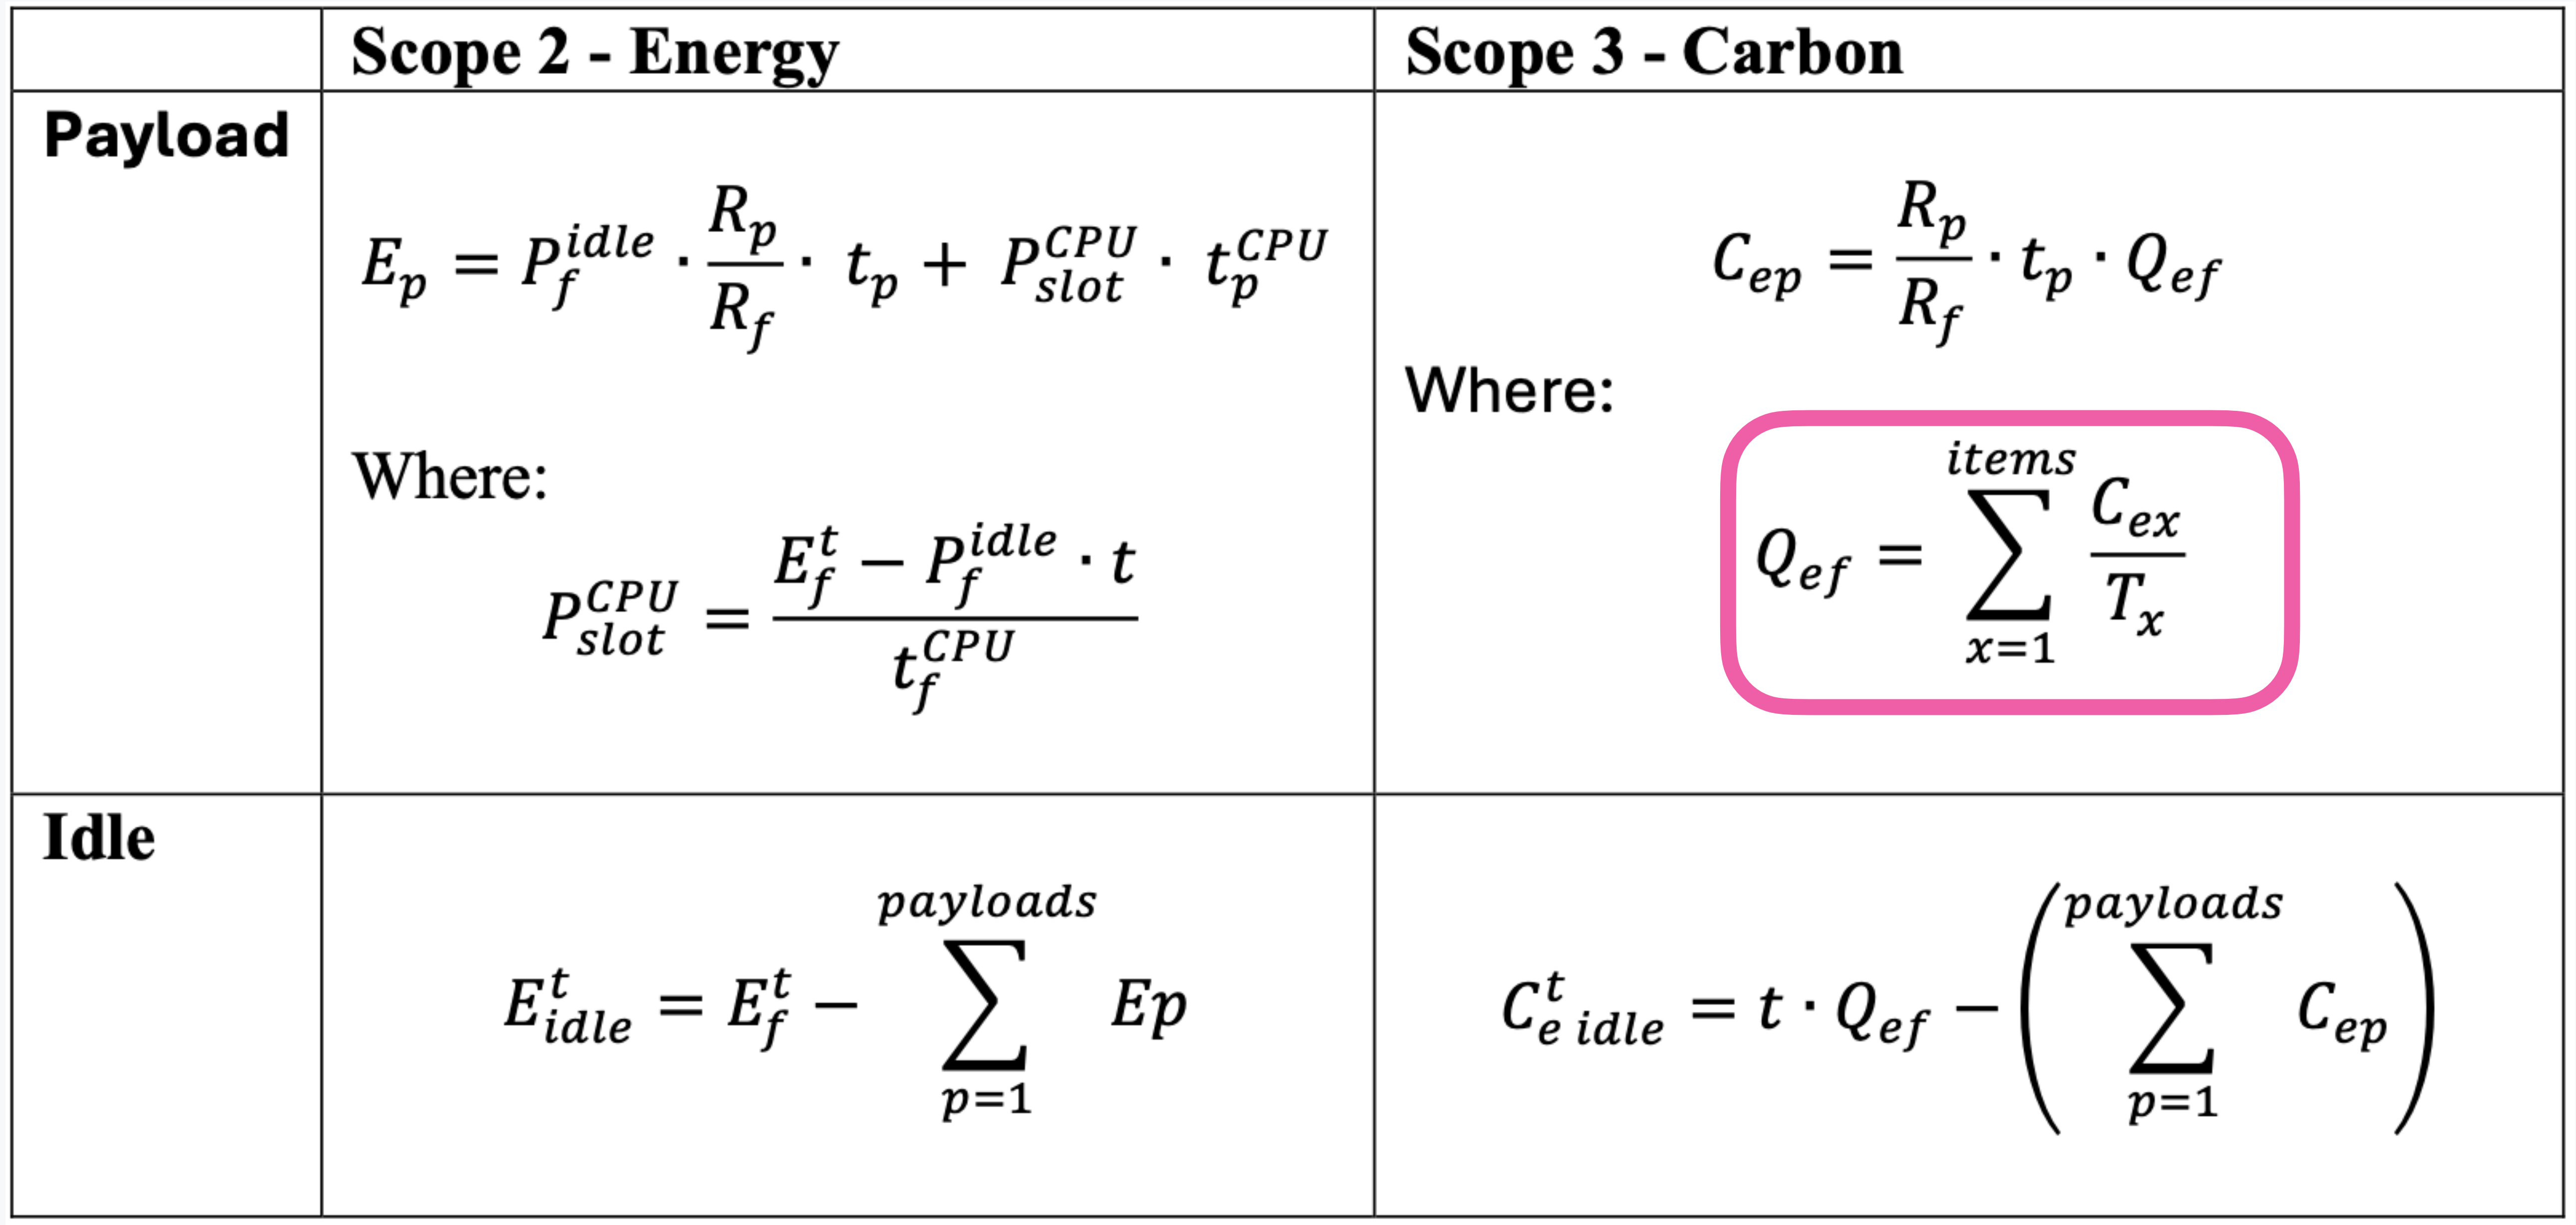

	Scope 2 - Energy	Scope 3 - Carbon
Payload	$E_p = P_f^{idle}\cdot\frac{R_p}{R_f}\cdot t_p + P_{slot}^{CPU}\cdot t_p^{CPU}$ Where: $P_{slot}^{CPU} = \frac{E_f^t - P_f^{idle}\cdot t}{t_f^{CPU}}$	$C_{ep} = \frac{R_p}{R_f}\cdot t_p\cdot Q_{ef}$ Where: $\boxed{Q_{ef} = \sum_{x=1}^{items}\frac{C_{ex}}{T_x}}$
Idle	$E_{idle}^t = E_f^t - \sum_{p=1}^{payloads} Ep$	$C_{e\ idle}^t = t\cdot Q_{ef} - \left(\sum_{p=1}^{payloads} C_{ep}\right)$

Table 3: Summary of the Enhanced Payload Model showing allocations of Scope 2 energy and Scope 3 carbon to user payloads and the remaining idle allocation to the provider

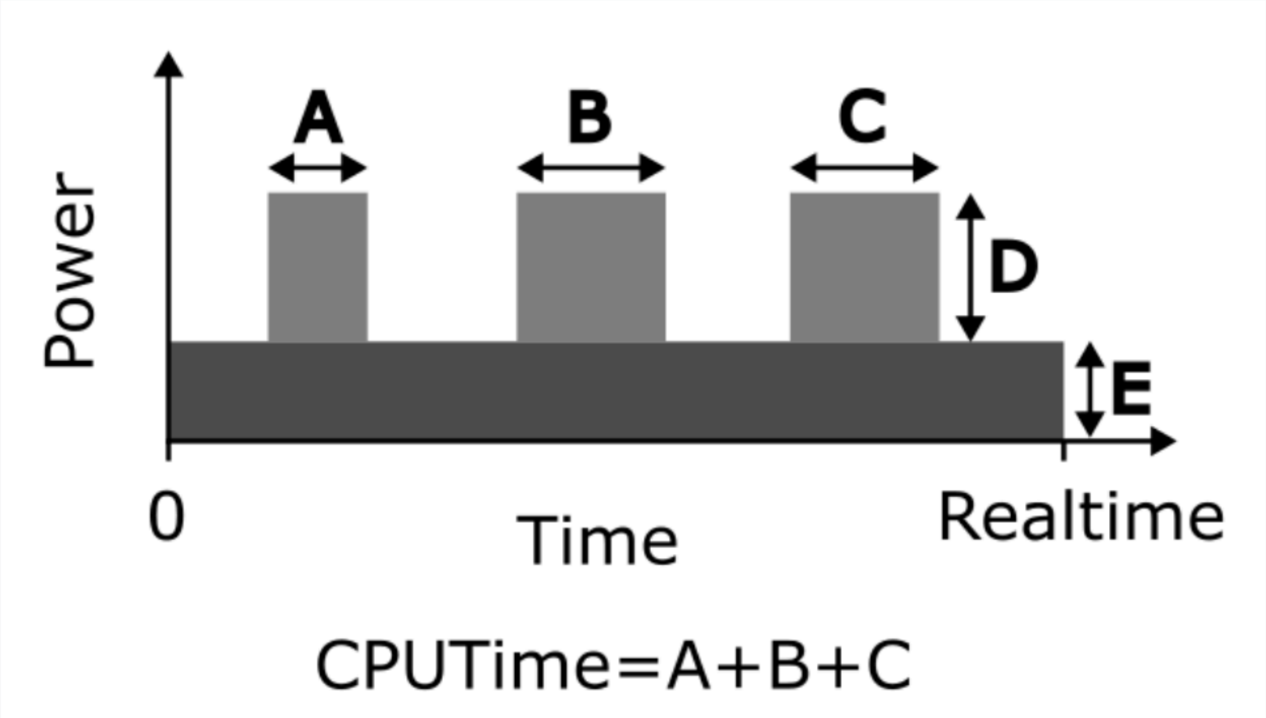

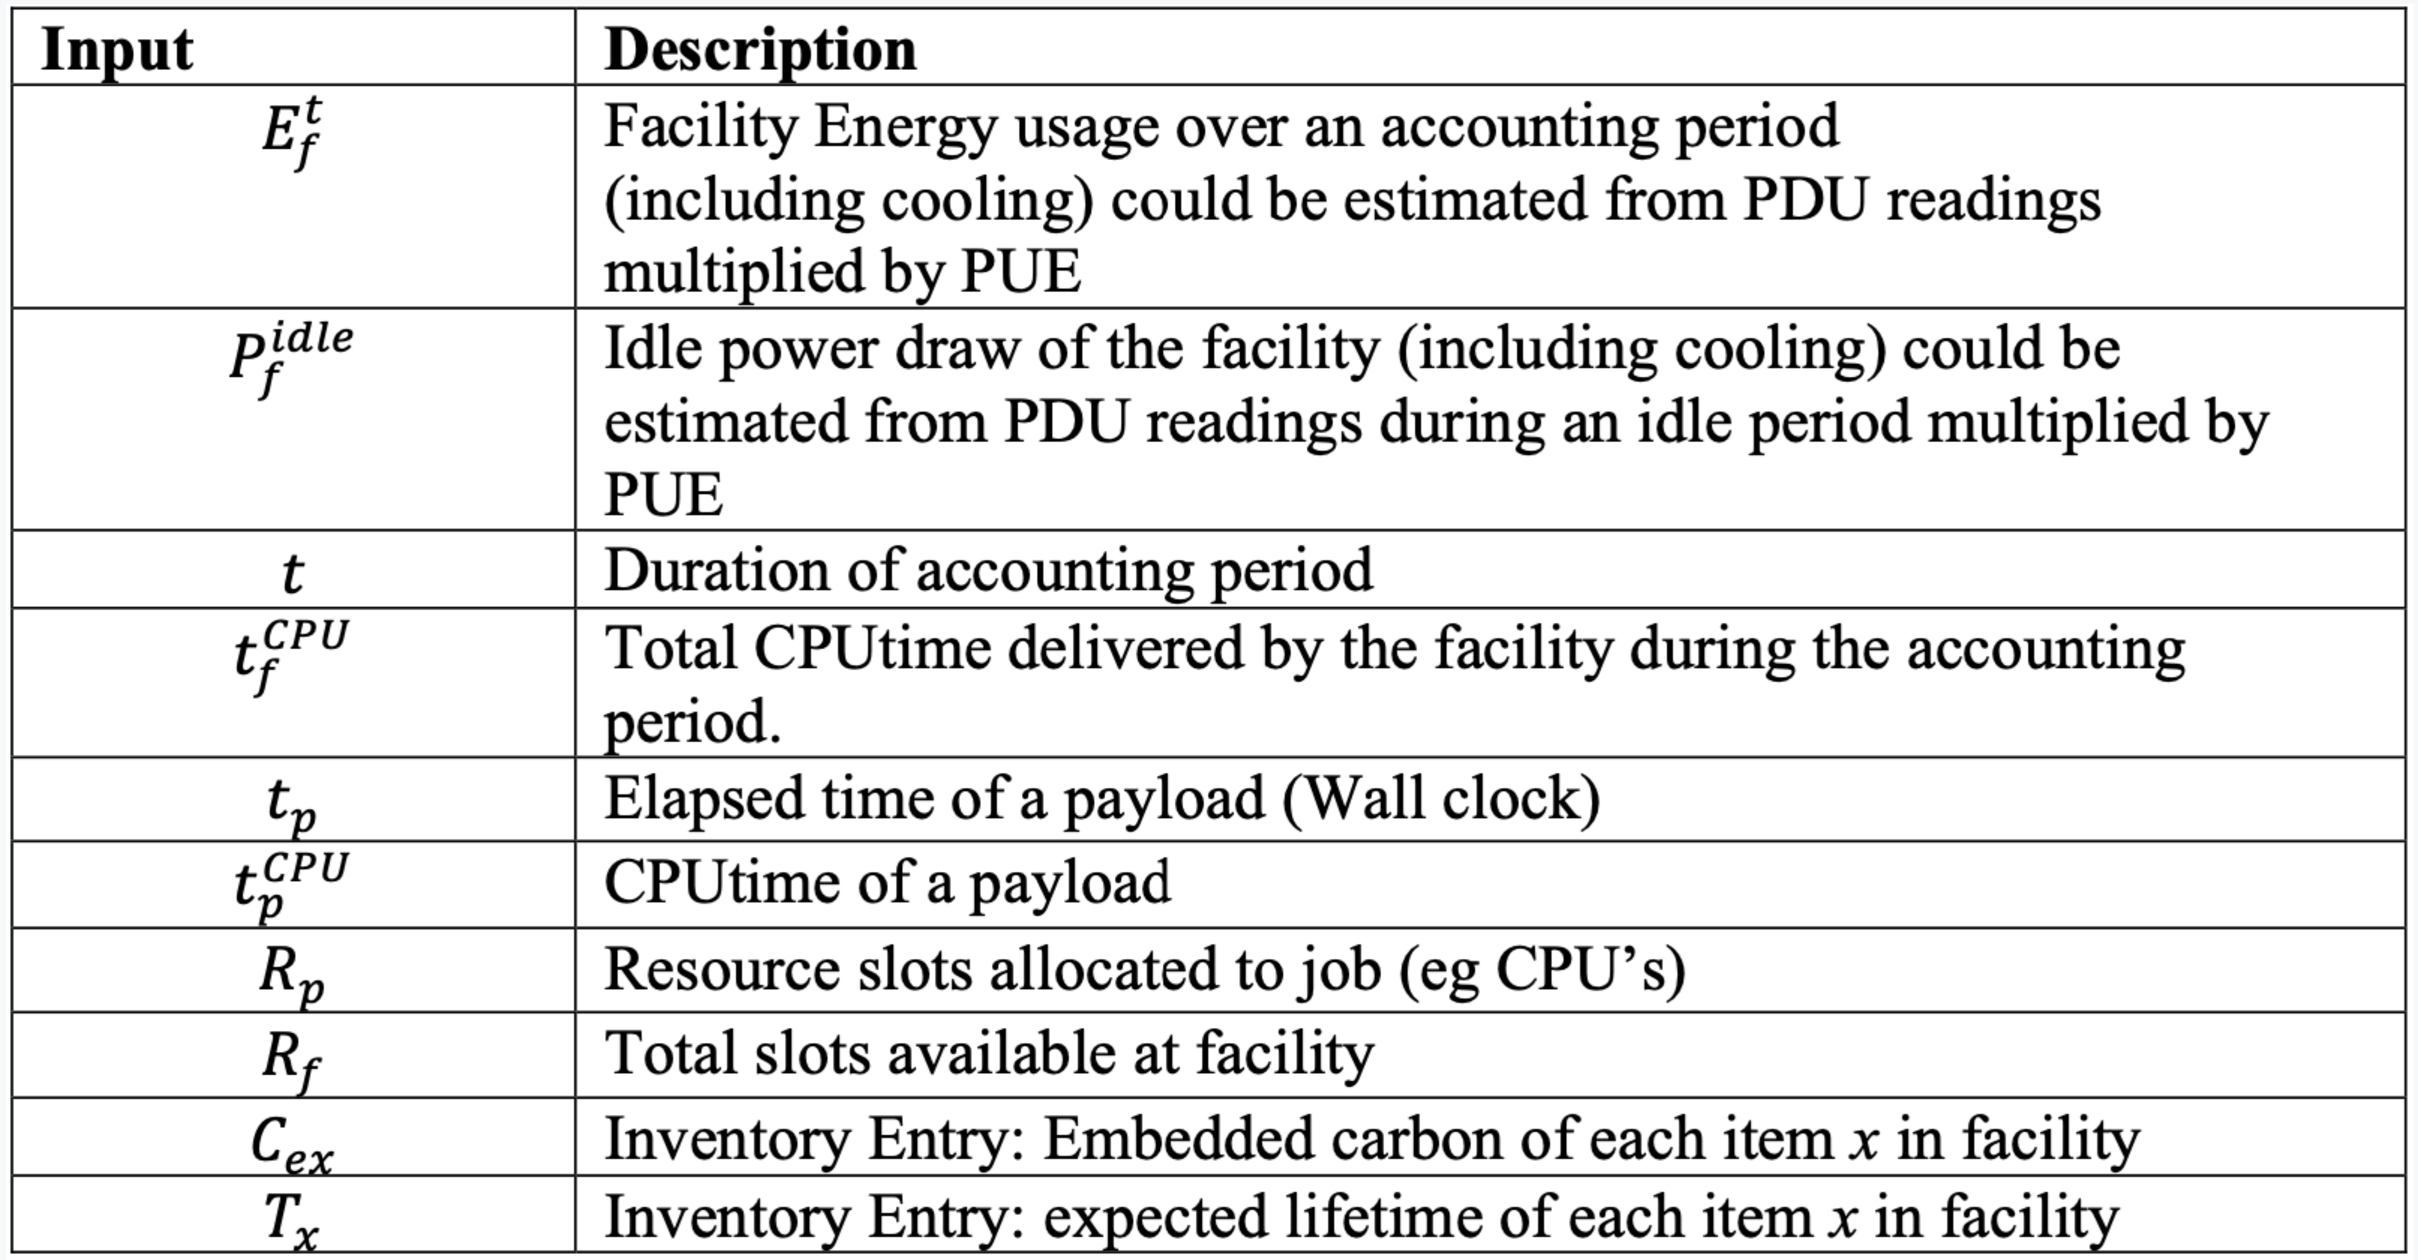

Input	Description
E_f^t	Facility Energy usage over an accounting period (including cooling) could be estimated from PDU readings multiplied by PUE
P_f^{idle}	Idle power draw of the facility (including cooling) could be estimated from PDU readings during an idle period multiplied by PUE
t	Duration of accounting period
t_f^{CPU}	Total CPUtime delivered by the facility during the accounting period.
t_p	Elapsed time of a payload (Wall clock)
t_p^{CPU}	CPUtime of a payload
R_p	Resource slots allocated to job (eg CPU's)
R_f	Total slots available at facility
C_{ex}	Inventory Entry: Embedded carbon of each item x in facility
T_x	Inventory Entry: expected lifetime of each item x in facility

Table 4: Summary of the inputs needed to evaluate the Enhanced Payload Model.

Enhanced Payload Model

Know your idle power?

Know your CPUtime?

iris Carbon Mapping Project

Apportion by Quota

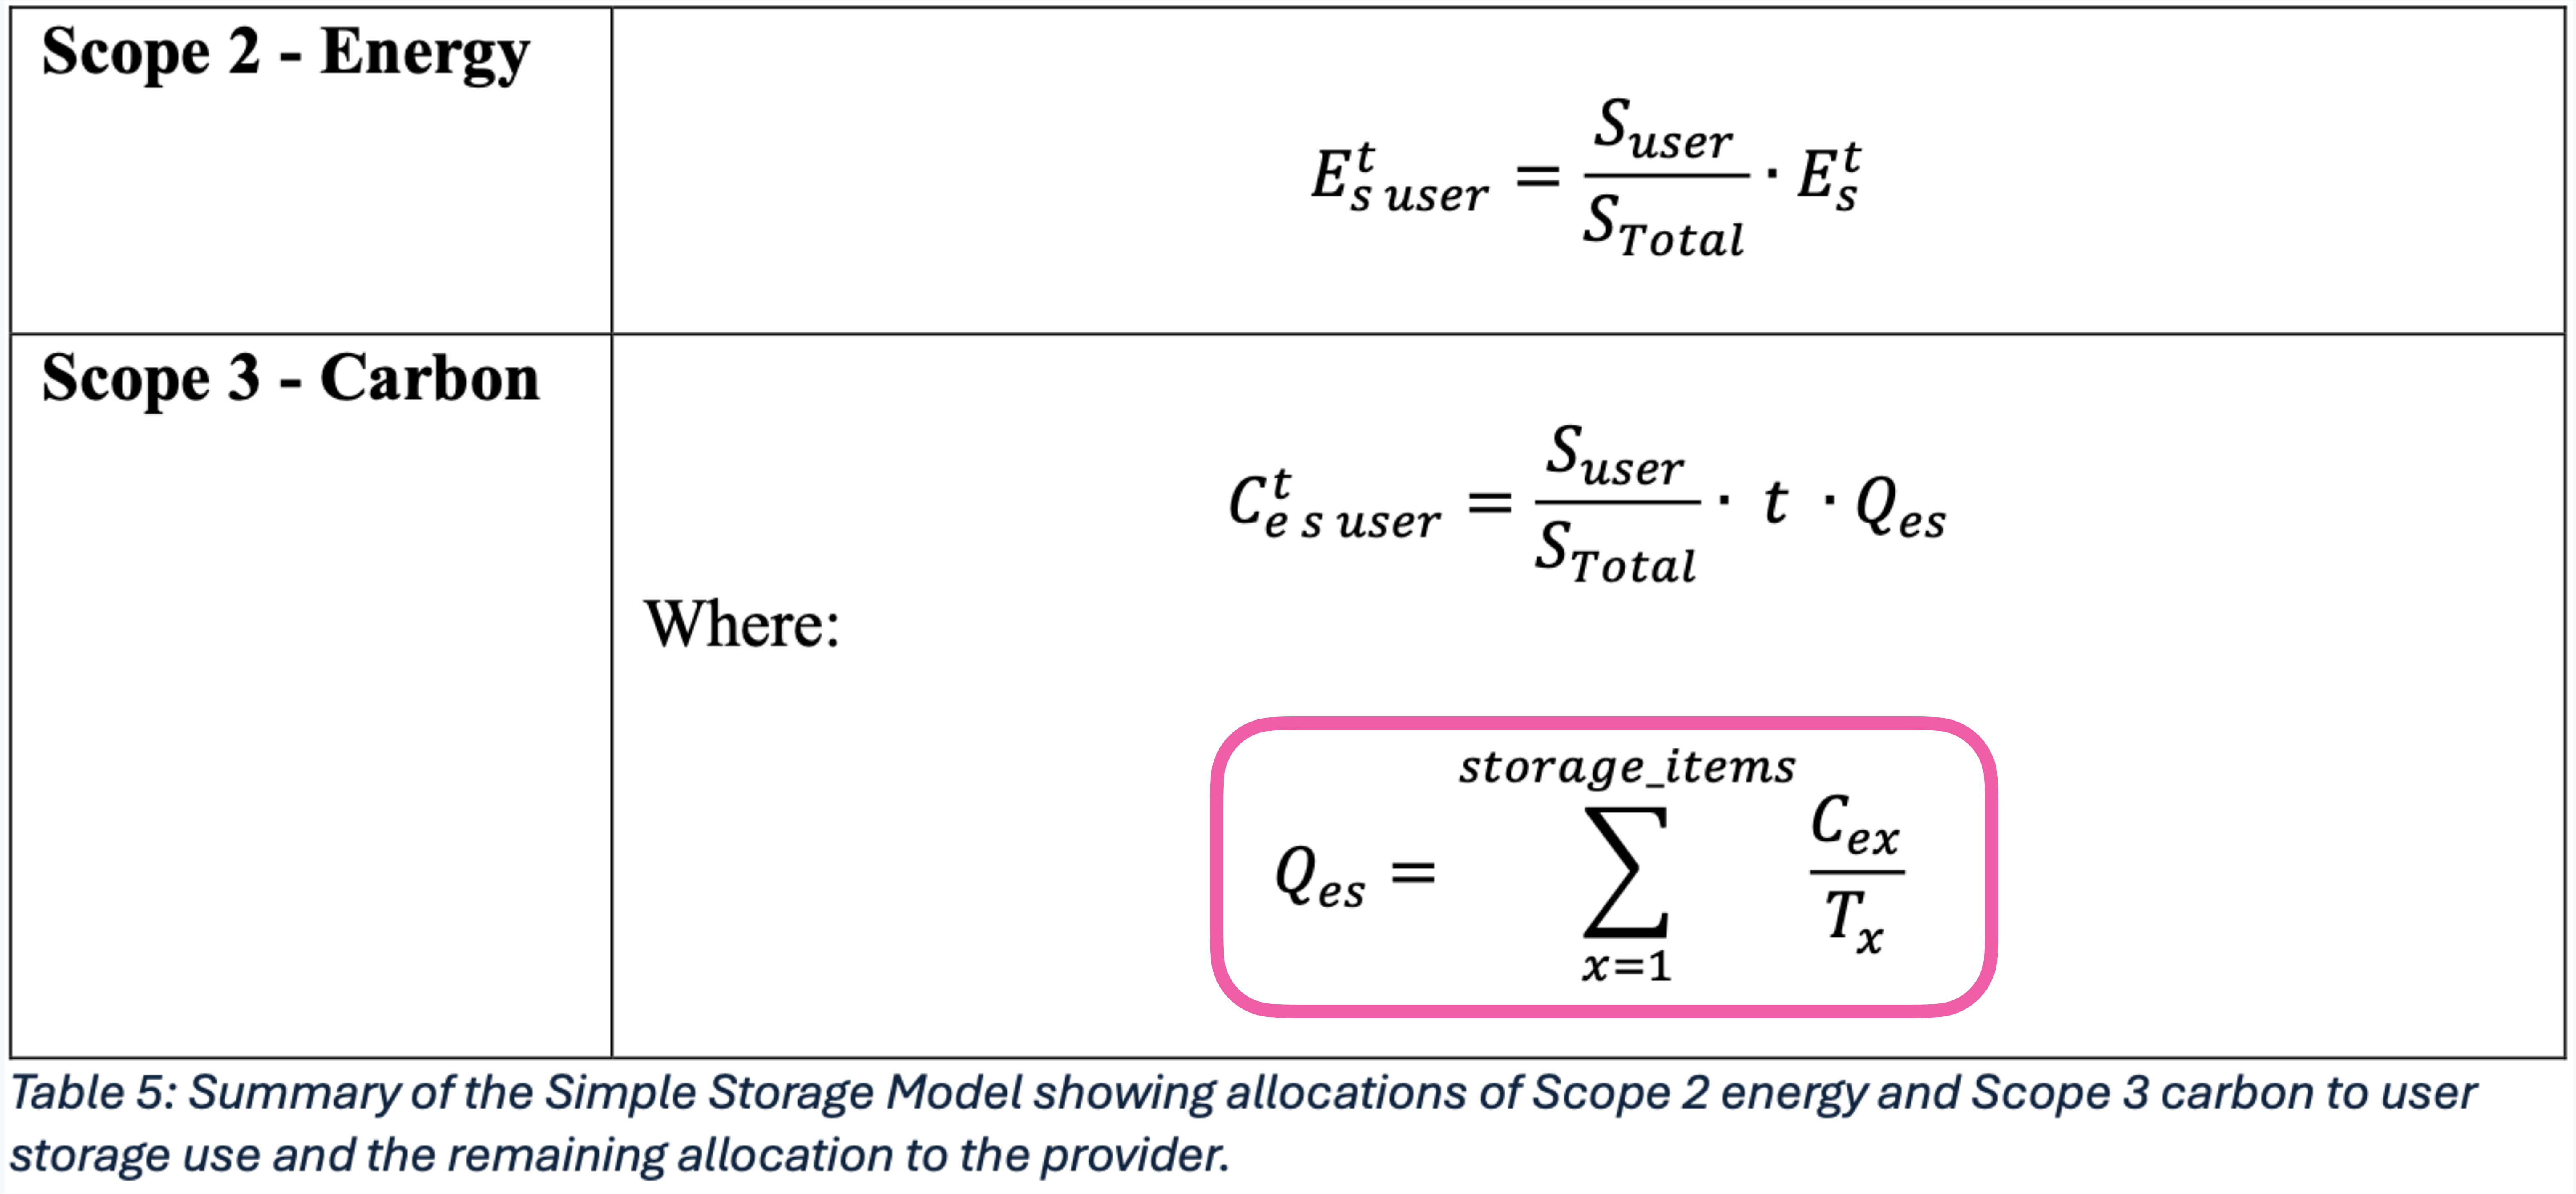

Scope 2 - Energy	$E^t_{s\,user} = \frac{S_{user}}{S_{Total}} \cdot E^t_s$
Scope 3 - Carbon	$C^t_{e\,s\,user} = \frac{S_{user}}{S_{Total}} \cdot t \cdot Q_{es}$ Where: $\boxed{Q_{es} = \sum_{x=1}^{storage\_items} \frac{C_{ex}}{T_x}}$

Table 5: Summary of the Simple Storage Model showing allocations of Scope 2 energy and Scope 3 carbon to user storage use and the remaining allocation to the provider.

Simple Storage Model

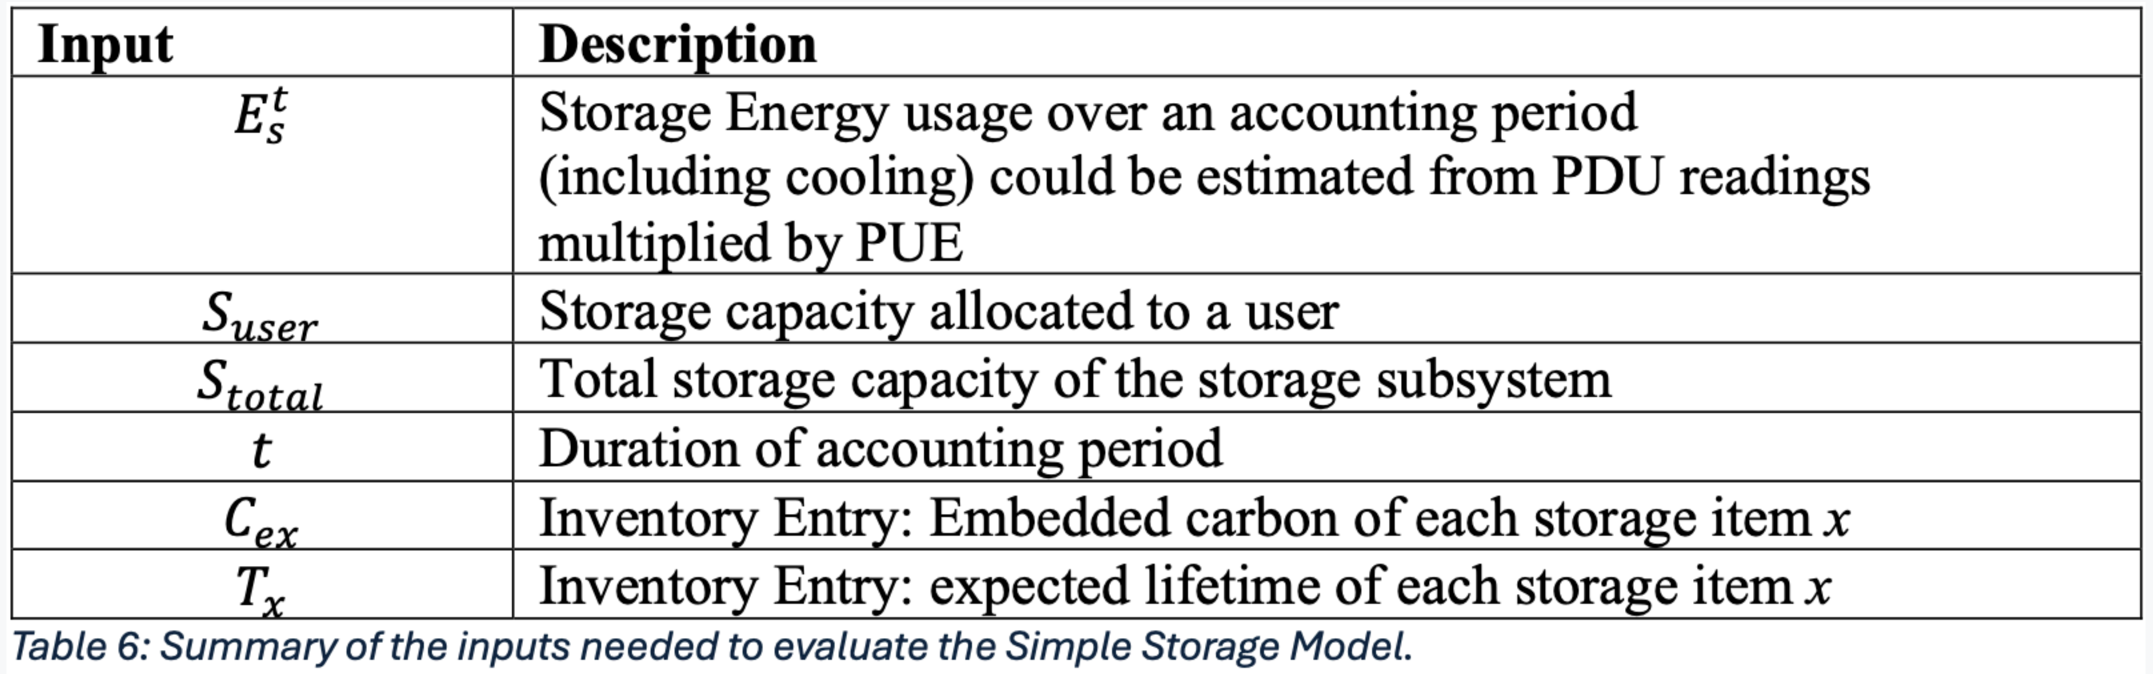

Input	**Description**
E^t_s	Storage Energy usage over an accounting period (including cooling) could be estimated from PDU readings multiplied by PUE
S_{user}	Storage capacity allocated to a user
S_{total}	Total storage capacity of the storage subsystem
t	Duration of accounting period
C_{ex}	Inventory Entry: Embedded carbon of each storage item x
T_x	Inventory Entry: expected lifetime of each storage item x

Table 6: Summary of the inputs needed to evaluate the Simple Storage Model.

iris Carbon Mapping Project

Enhanced Storage Model

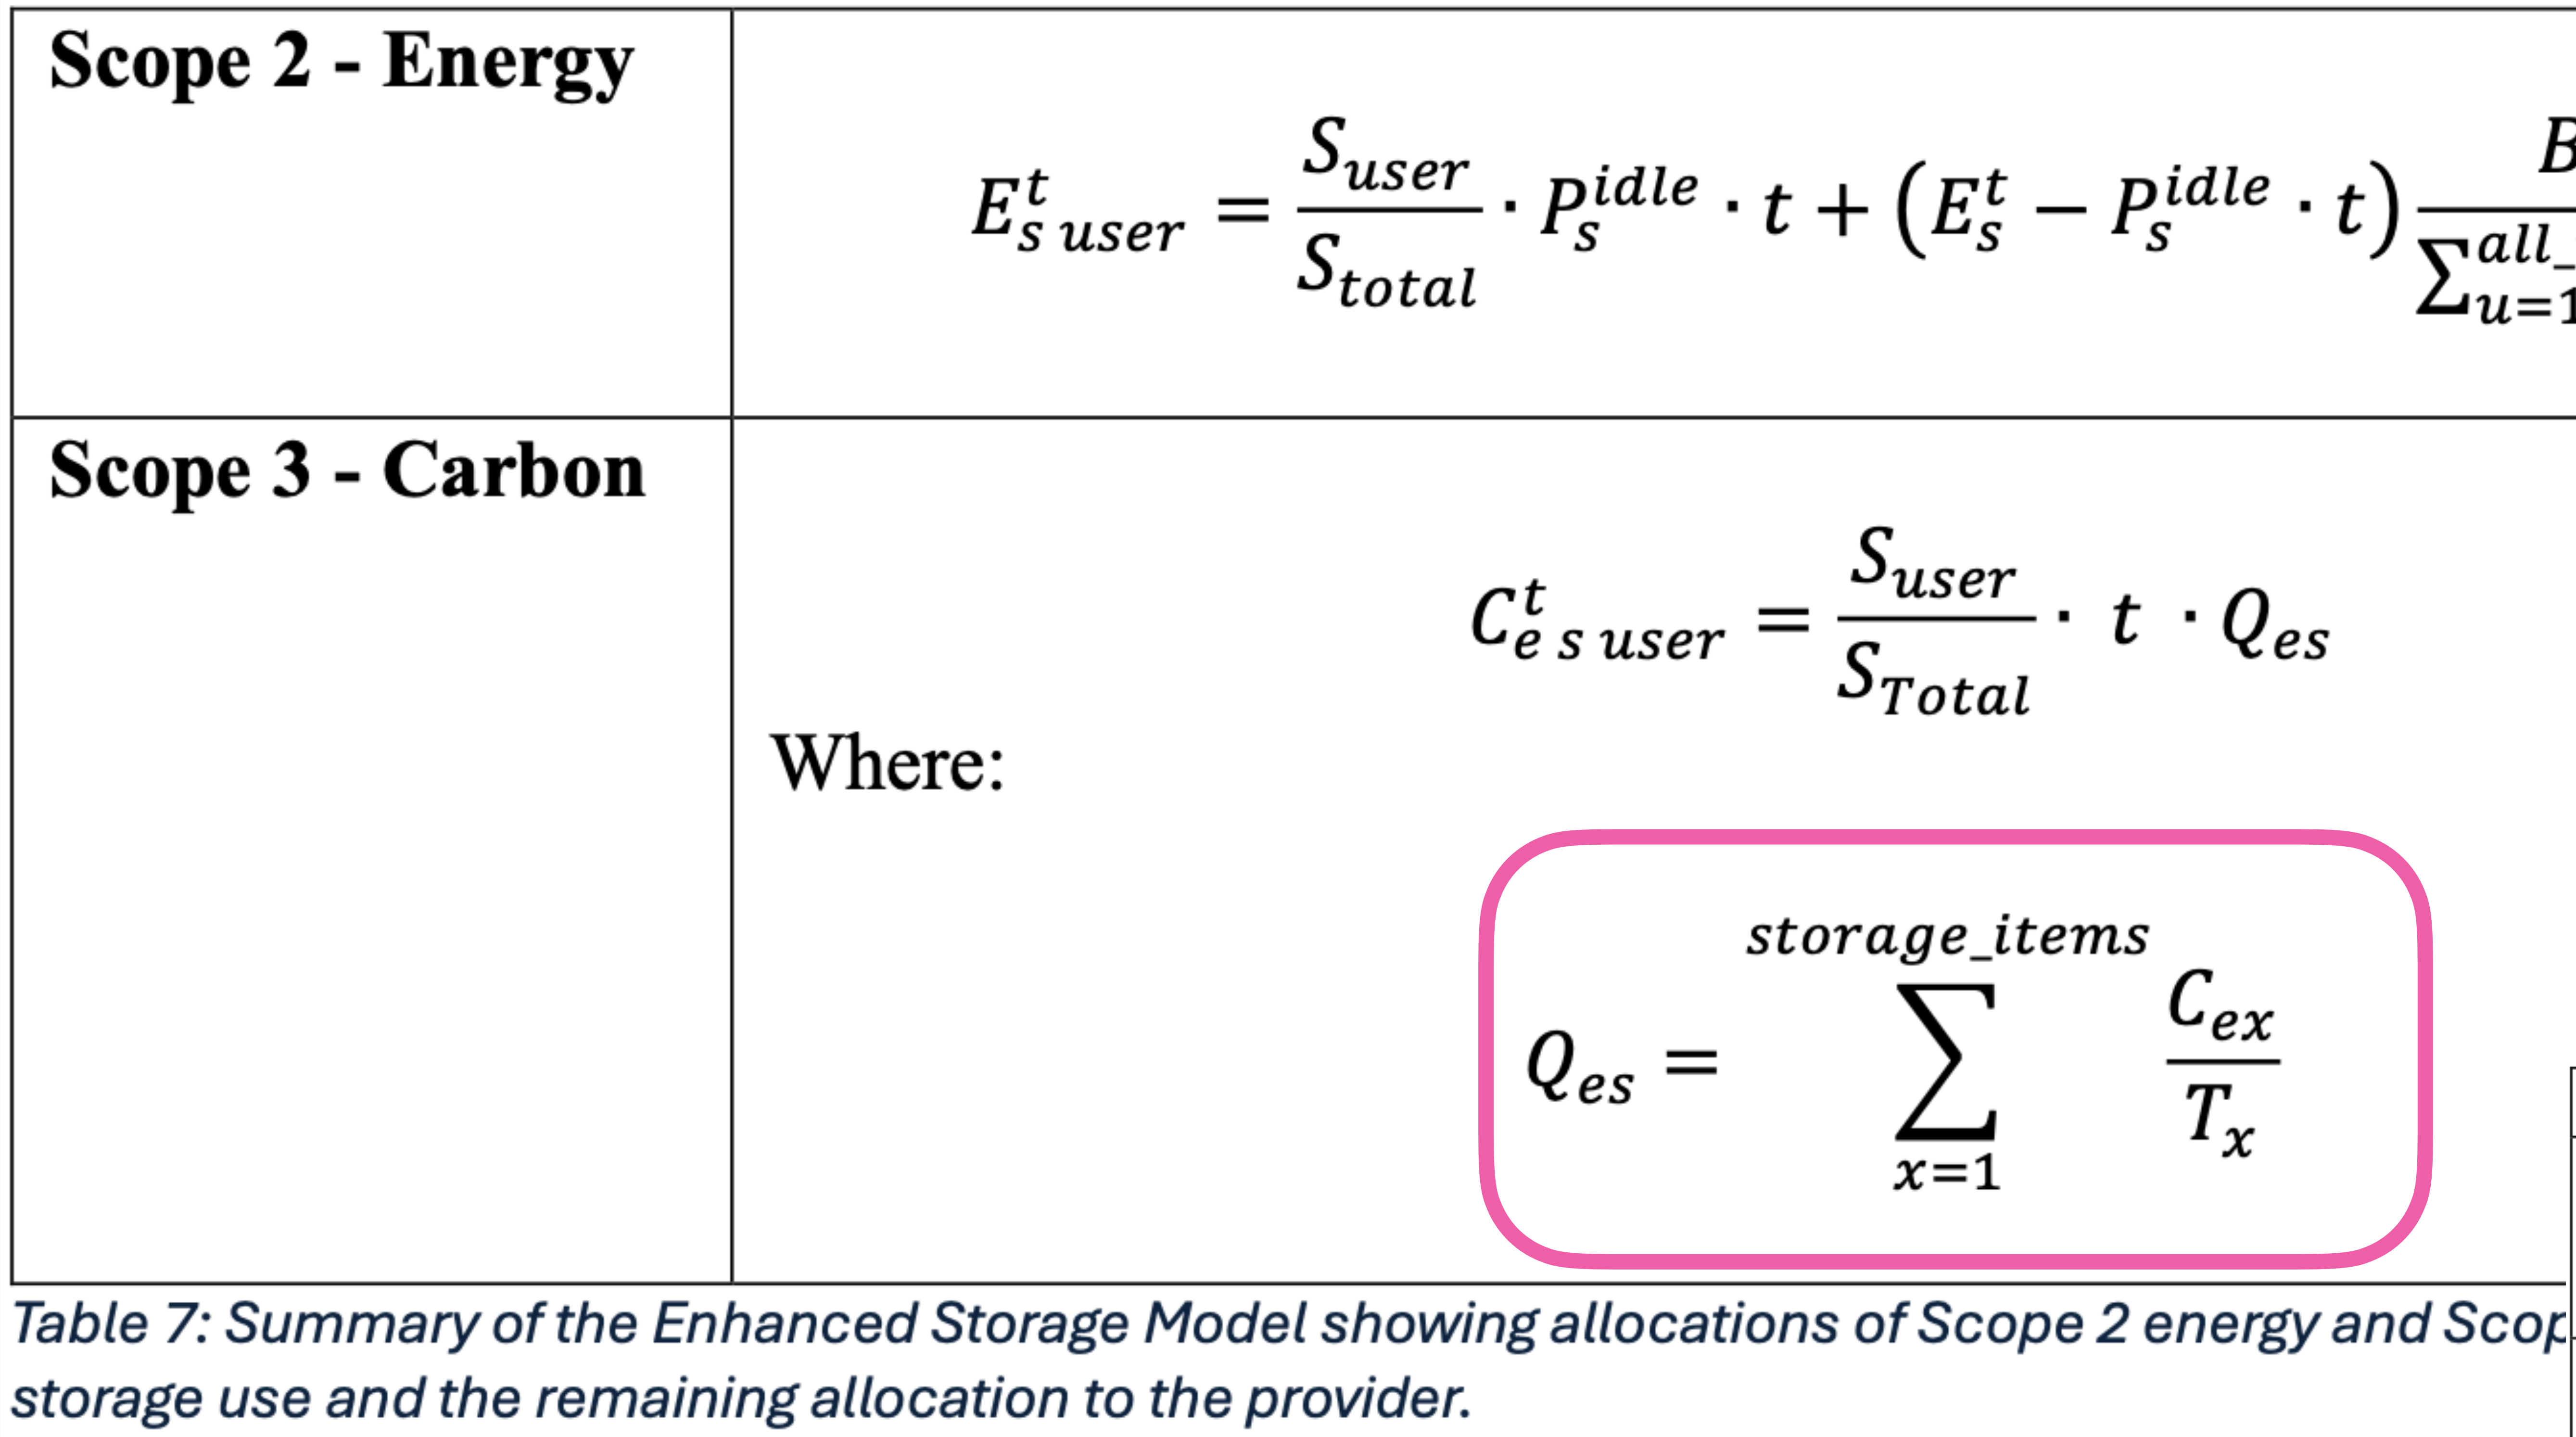

Scope 2 - Energy	$E^t_{s\,user} = \frac{S_{user}}{S_{total}} \cdot P^{idle}_s \cdot t + \left(E^t_s - P^{idle}_s \cdot t\right)\frac{B_{user}}{\sum_{u=1}^{all\_users} B_u}$
Scope 3 - Carbon	$C^t_{e\,s\,user} = \frac{S_{user}}{S_{Total}} \cdot t \cdot Q_{es}$ Where: $\boxed{Q_{es} = \sum_{x=1}^{storage\_items} \frac{C_{ex}}{T_x}}$

Table 7: Summary of the Enhanced Storage Model showing allocations of Scope 2 energy and Scop storage use and the remaining allocation to the provider.

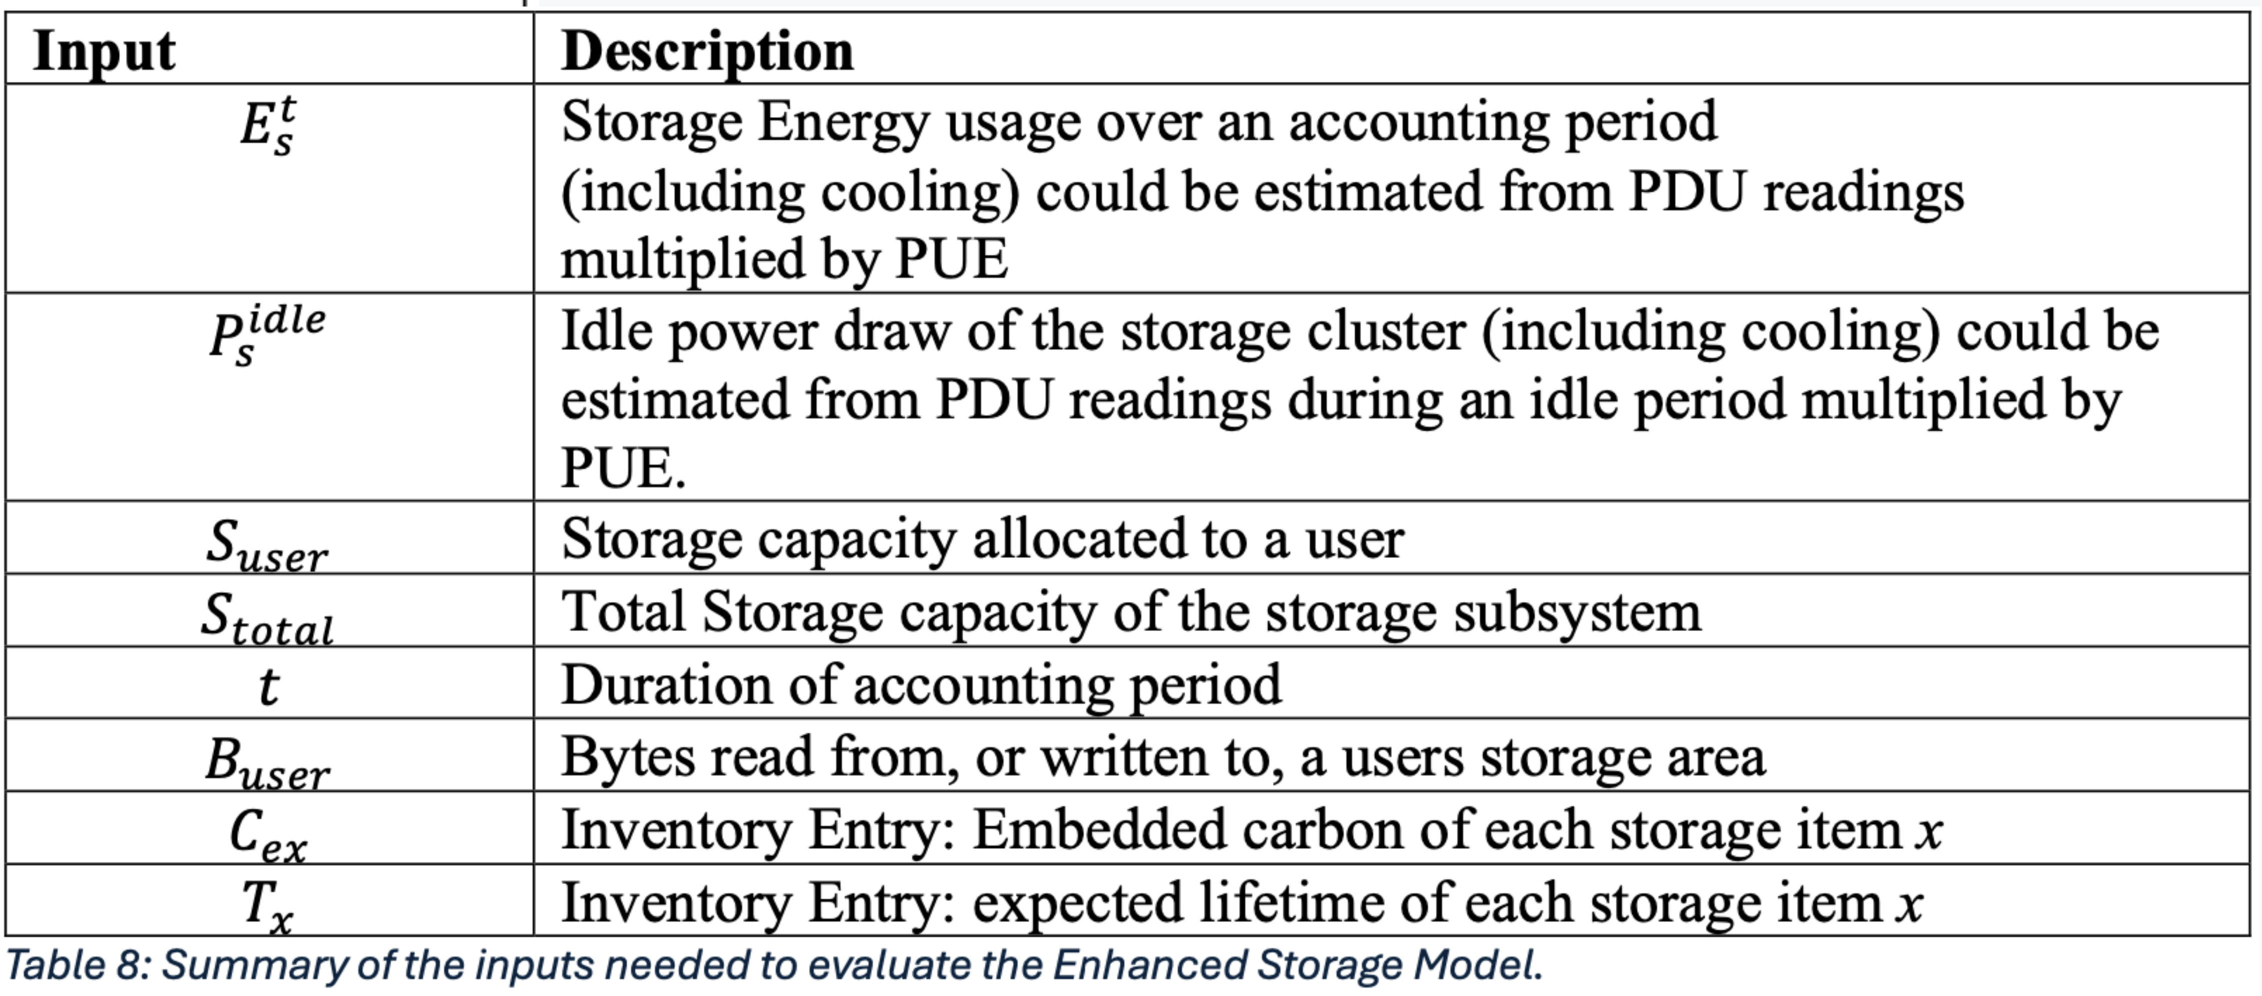

Input	Description
E^t_s	Storage Energy usage over an accounting period (including cooling) could be estimated from PDU readings multiplied by PUE
P^{idle}_s	Idle power draw of the storage cluster (including cooling) could be estimated from PDU readings during an idle period multiplied by PUE.
S_{user}	Storage capacity allocated to a user
S_{total}	Total Storage capacity of the storage subsystem
t	Duration of accounting period
B_{user}	Bytes read from, or written to, a users storage area
C_{ex}	Inventory Entry: Embedded carbon of each storage item x
T_x	Inventory Entry: expected lifetime of each storage item x

Table 8: Summary of the inputs needed to evaluate the Enhanced Storage Model.

Know your idle power?

Know your bytes?

iris Carbon Mapping Project

Testing the Payload Models for Batch

Idle Power →

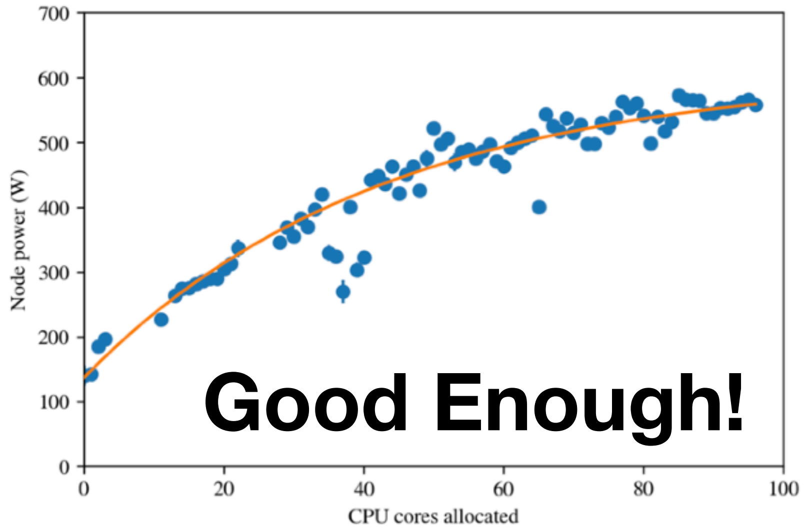

Good Enough!

Figure 2: Node power P_N as a function of allocated CPU cores in a node N_{CA}. The Blue dots and error bars show the mean power and std-deviation for each discrete N_{CA}. Regression line y-intercept gives the idle power of a node $P_{idle} \approx 137.1W$.

User	**Simple Payload Model** **kWh**	**Enhanced Payload Model** **kWh**
prdatl	1204.79	1191.95
pillhcb	159.08	242.24
pilcms	76.83	71.28
pilatl	48.86	51.58
Pilmoe	10.75	16.86
Pildune	2.46	0.61
Others	0.08	0.04
Sub total	1502.86	1574.57
Idle(provider)	94.14	22.43
Total	1597	1597

Table 10: Results of evaluating the Simple and Enhanced Payload models on QMUL batch payloads the 24 hour period of 2024-03-07.

Input	**Value**	**Slurm name**	**Description**
E_f^t	1597 kWh	-	Facility Energy usage. In this four rack example the PDU cumulative energy readings were used to calculate this.
P_f^{idle}	16.45 kW	-	Idle power draw of the facility. In this example the 137.1W per node was multiplied by 120 nodes.
t	86400 s	-	Duration of accounting period. In this case 24 hours.
t_f^{CPU}	-	$\sum TotalCPU$	Total CPUtime delivered by the facility during the accounting period. Sum of the TotalCPU figures for all payloads
t_p	-	Elapsed	Elapsed time of a payload (Wall clock)
t_p^{CPU}	-	TotalCPU	CPUtime of a payload
$Slots_p$	-	AllocCPUS	Resource slots allocated to job (eg CPU's)
$Slots_f$	11520	-	Total slots available at facility. In this case 120 nodes with 96 cores each.

Table 9: Measured and derived constants and Slurm accounting data names used to evaluate the payload models for QMUL batch payloads.

Works for Batch!

iris Carbon Mapping Project

Testing the Payload Models for Cloud

Works for Cloud too!

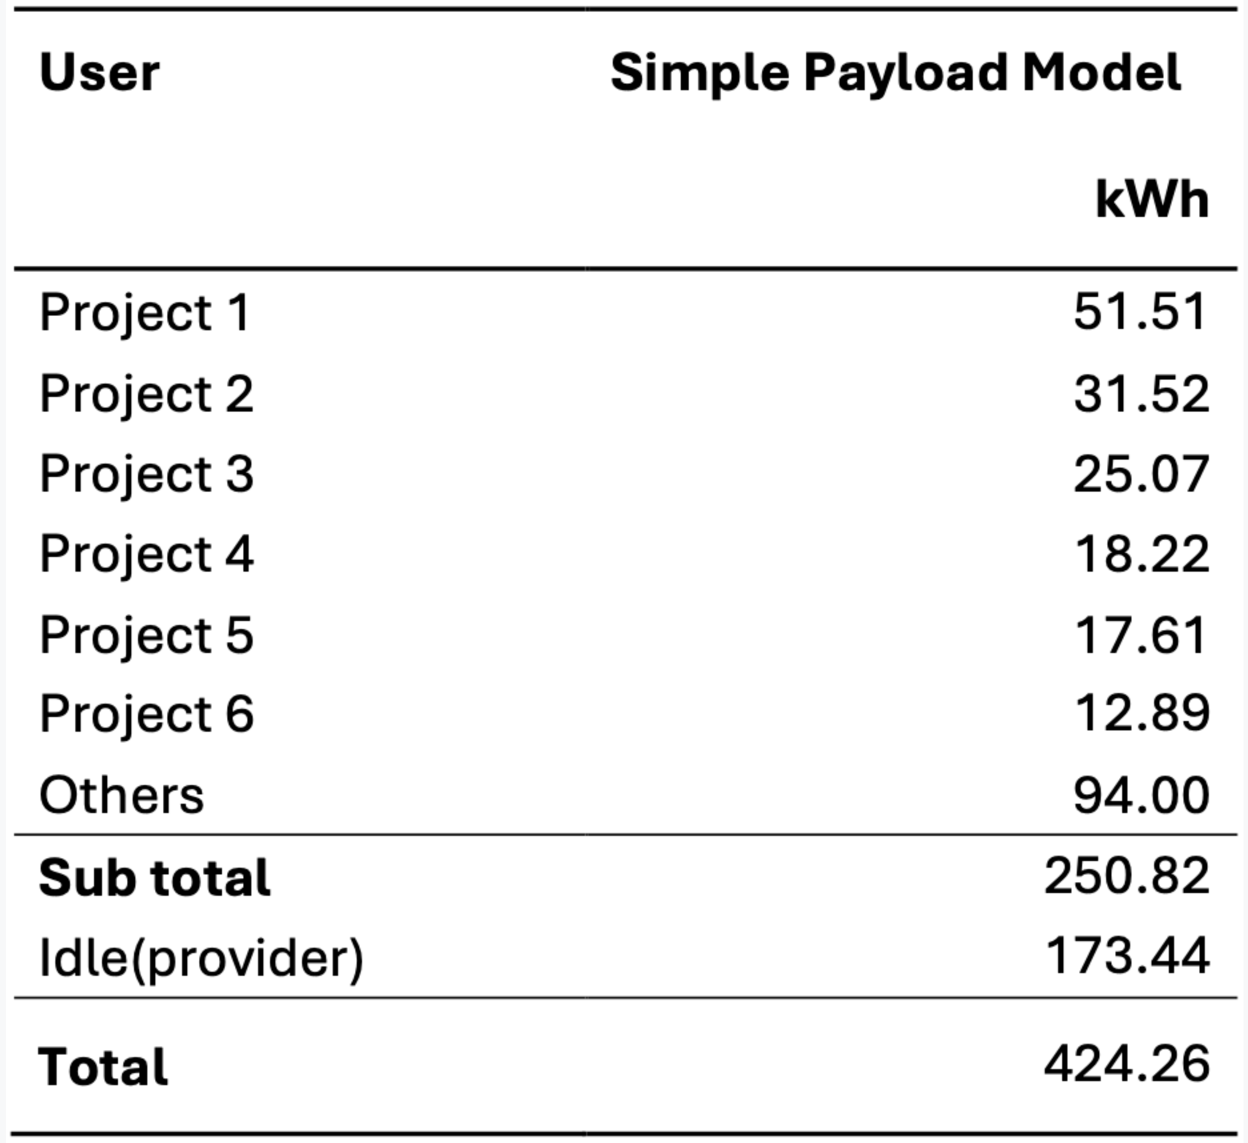

User	Simple Payload Model kWh
Project 1	51.51
Project 2	31.52
Project 3	25.07
Project 4	18.22
Project 5	17.61
Project 6	12.89
Others	94.00
Sub total	250.82
Idle(provider)	173.44
Total	424.26

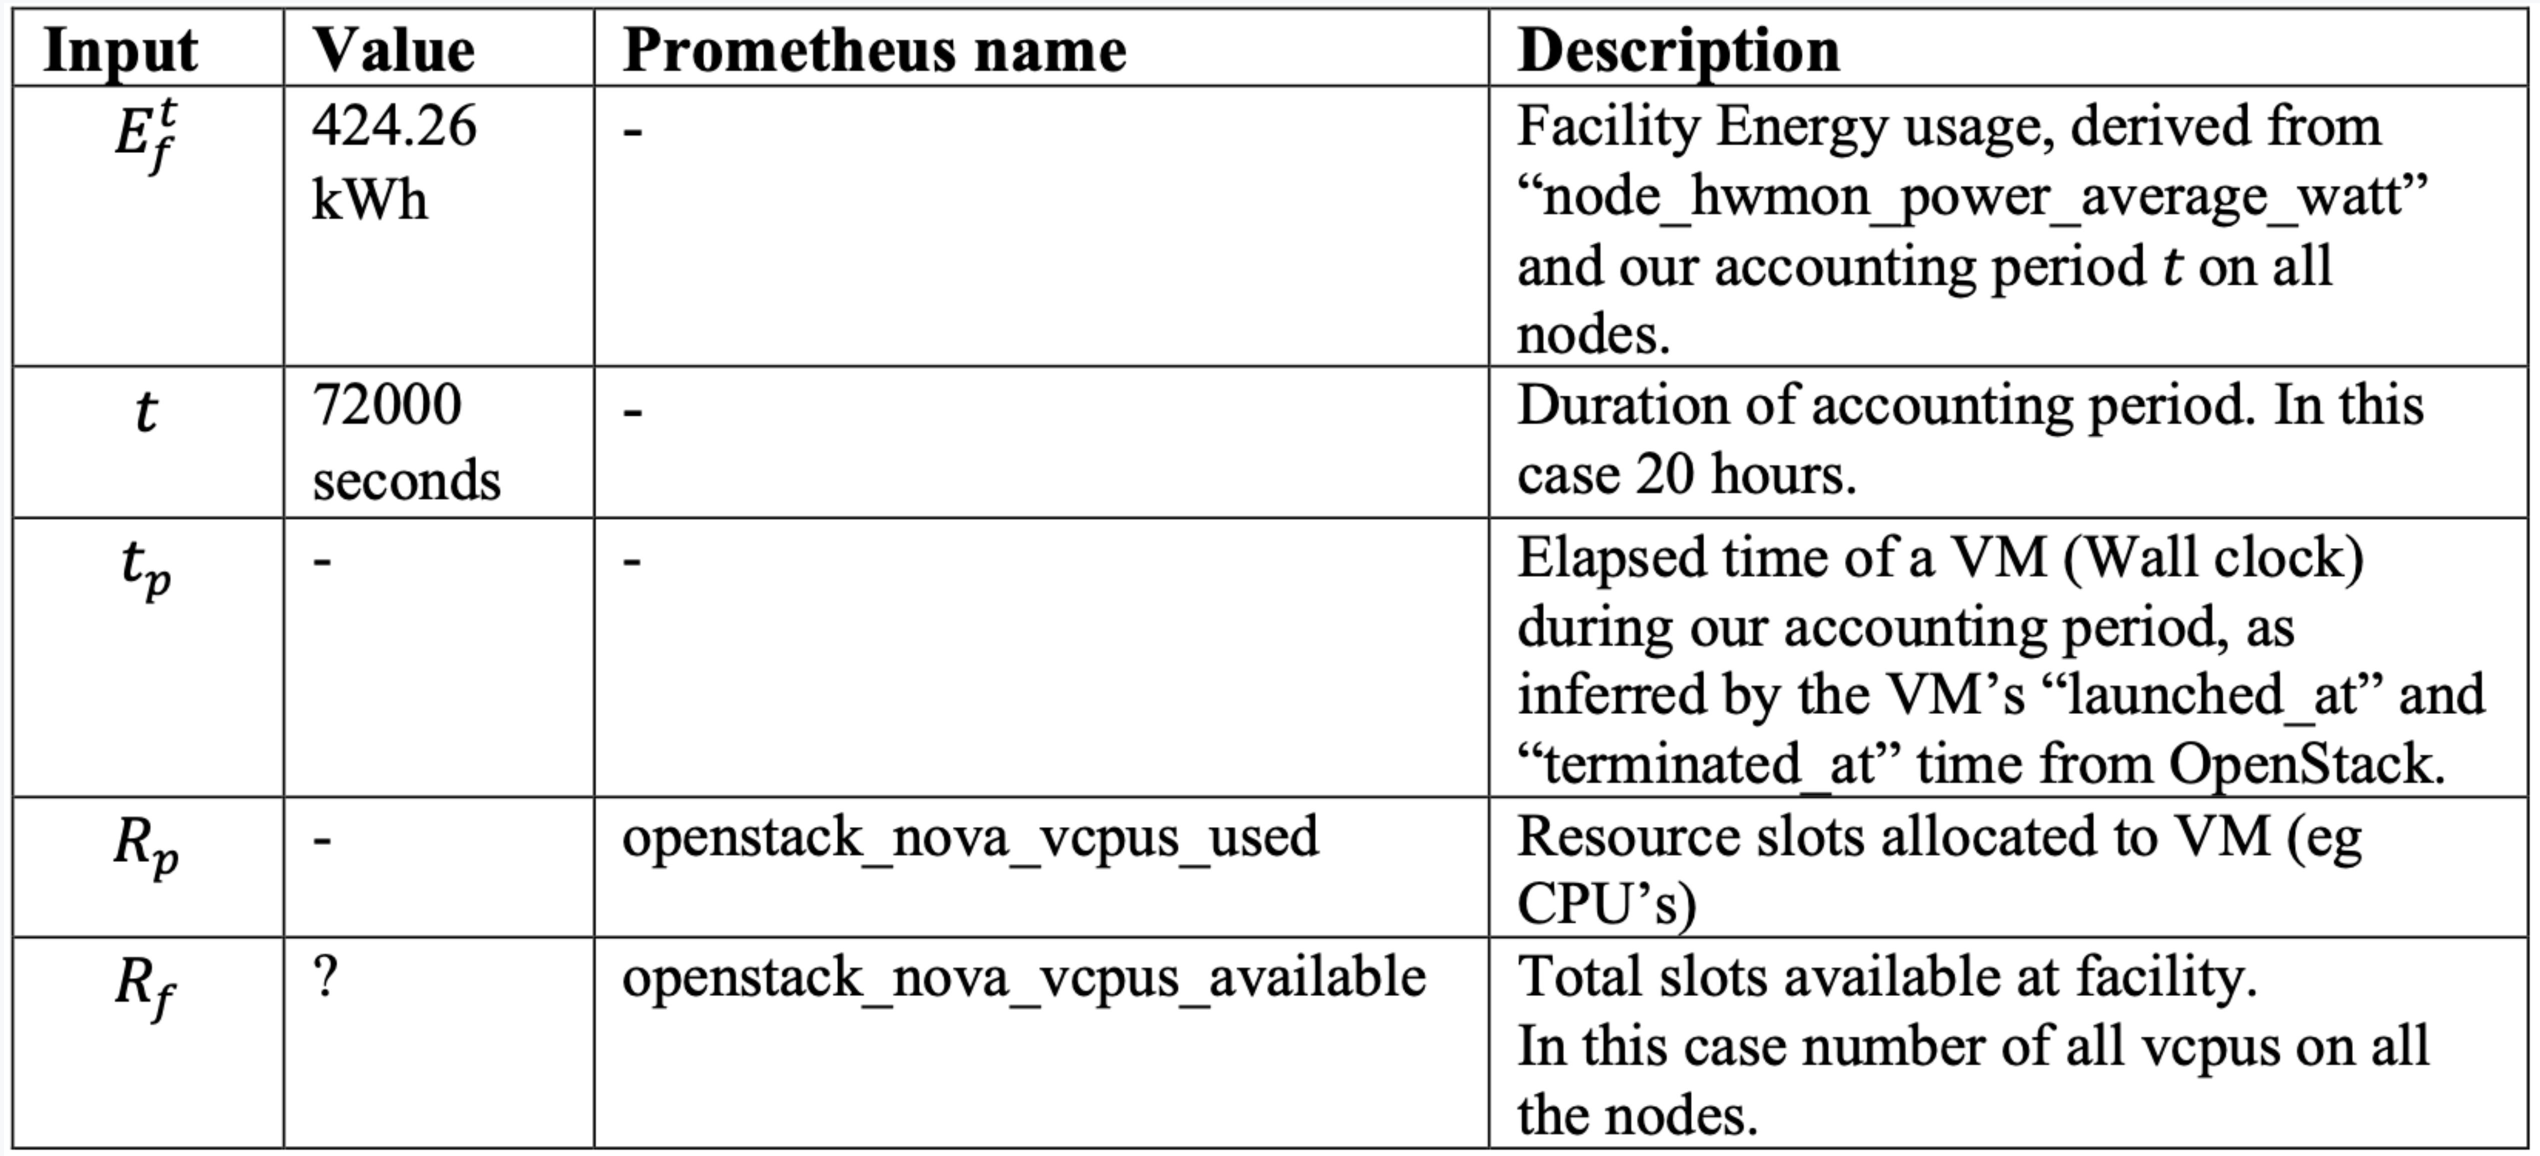

Input	Value	Prometheus name	Description
E_f^t	424.26 kWh	-	Facility Energy usage, derived from "node_hwmon_power_average_watt" and our accounting period t on all nodes.
t	72000 seconds	-	Duration of accounting period. In this case 20 hours.
t_p	-	-	Elapsed time of a VM (Wall clock) during our accounting period, as inferred by the VM's "launched_at" and "terminated_at" time from OpenStack.
R_p	-	openstack_nova_vcpus_used	Resource slots allocated to VM (eg CPU's)
R_f	?	openstack_nova_vcpus_available	Total slots available at facility. In this case number of all vcpus on all the nodes.

Table 13: Measured and derived constants and Prometheus accounting data names used to evaluate the simple payload model for STFC Cloud payloads.

Simple does…

Enhanced should too…

iris Carbon Mapping Project

Testing the Storage Models

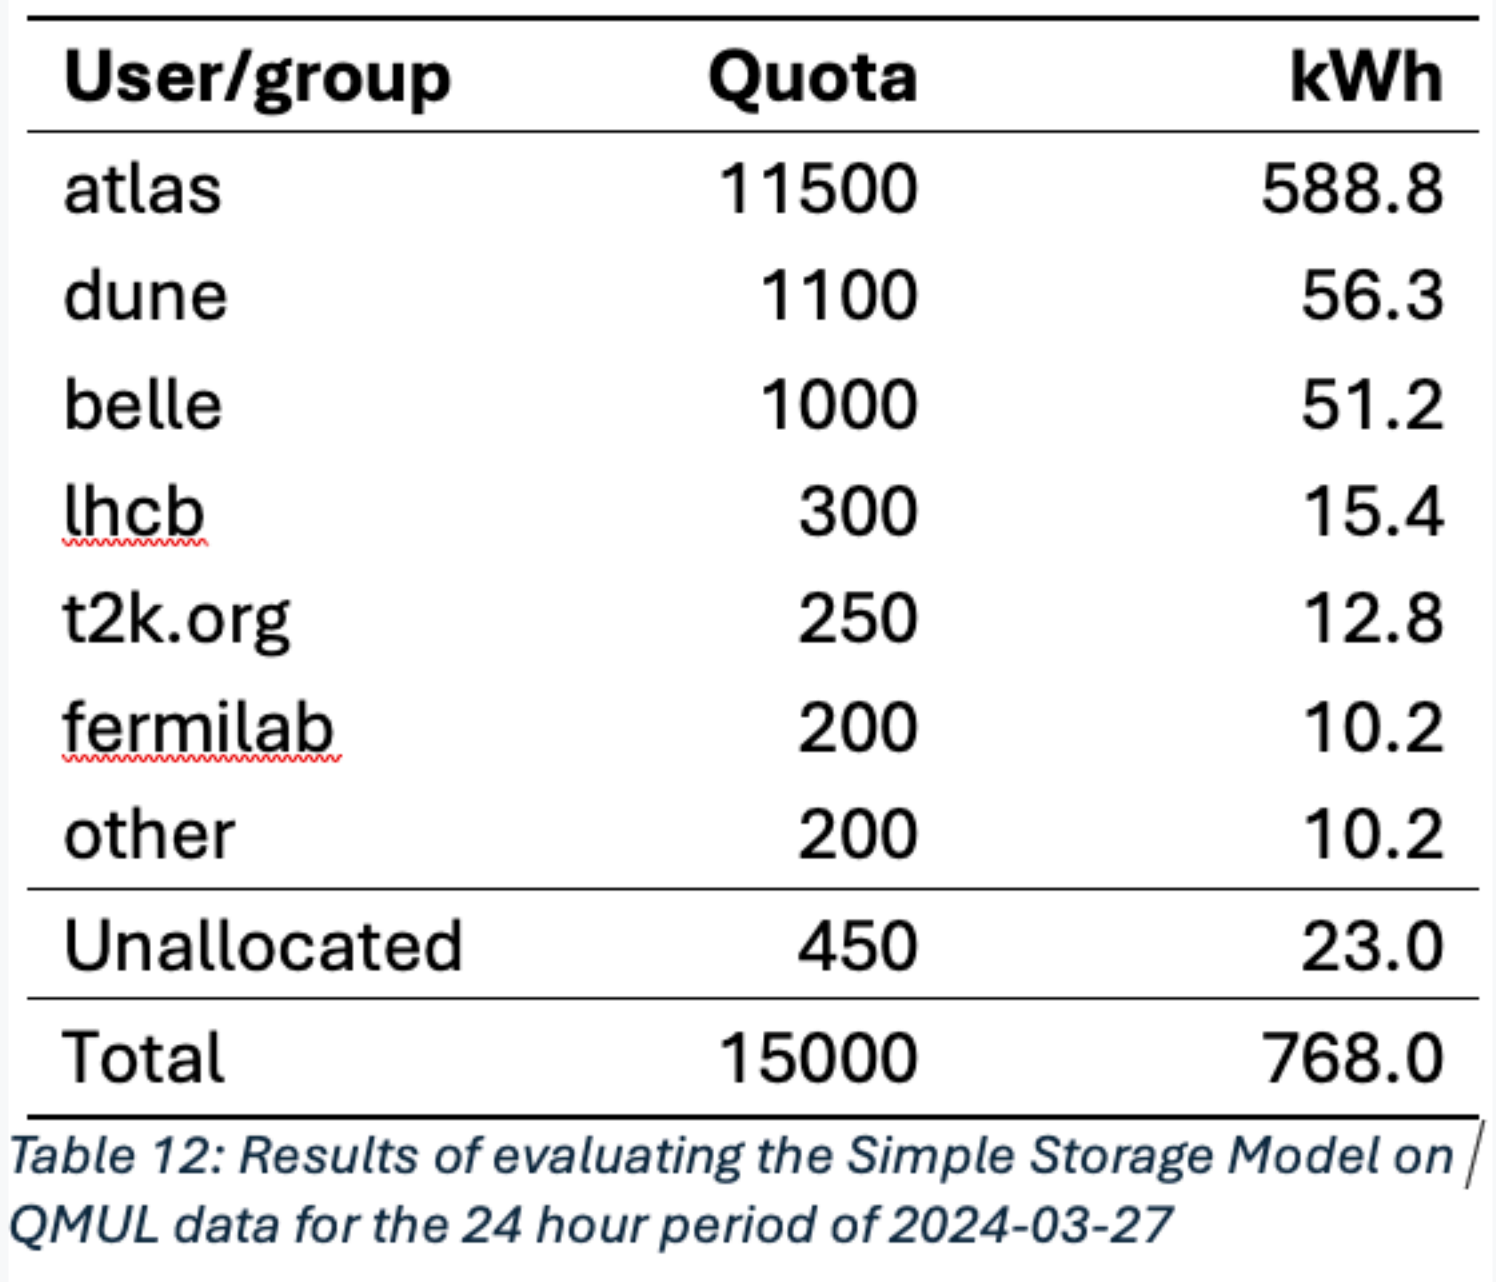

User/group	Quota	kWh
atlas	11500	588.8
dune	1100	56.3
belle	1000	51.2
lhcb	300	15.4
t2k.org	250	12.8
fermilab	200	10.2
other	200	10.2
Unallocated	450	23.0
Total	15000	768.0

Table 12: Results of evaluating the Simple Storage Model on QMUL data for the 24 hour period of 2024-03-27

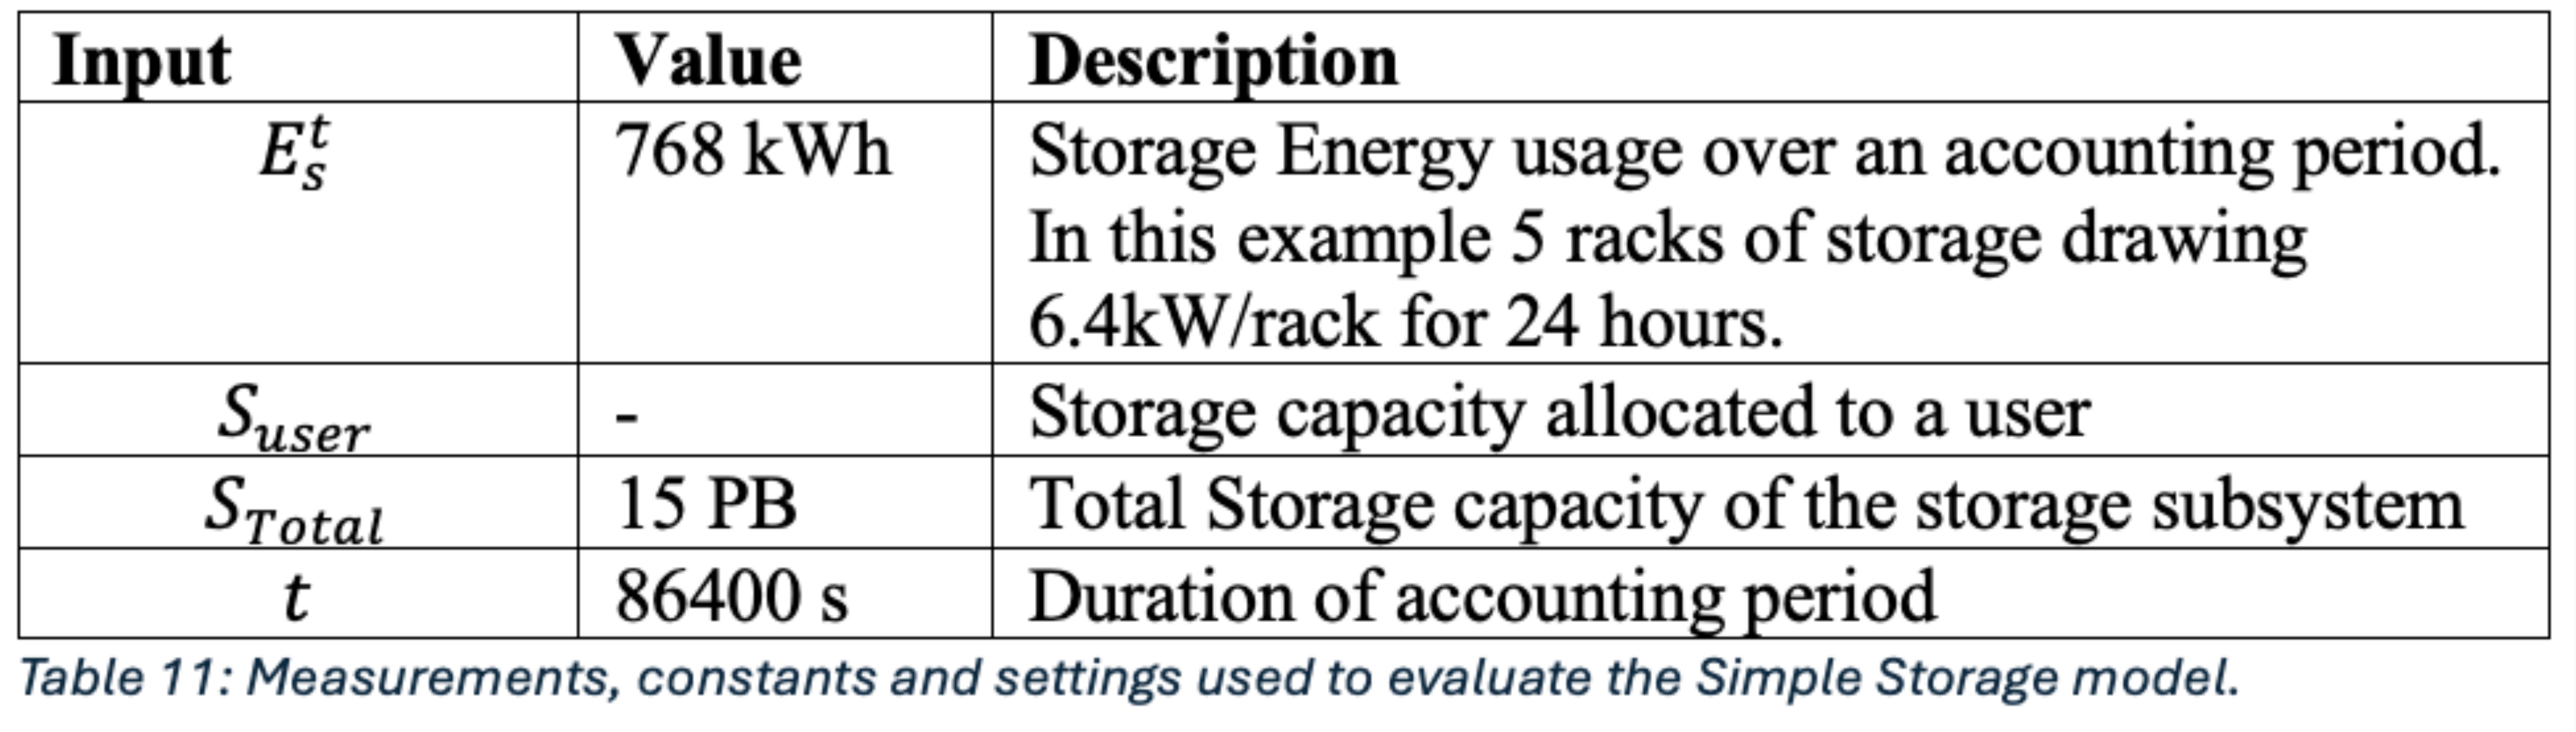

Input	Value	Description
E_s^t	768 kWh	Storage Energy usage over an accounting period. In this example 5 racks of storage drawing 6.4kW/rack for 24 hours.
S_{user}	-	Storage capacity allocated to a user
S_{Total}	15 PB	Total Storage capacity of the storage subsystem
t	86400 s	Duration of accounting period

Table 11: Measurements, constants and settings used to evaluate the Simple Storage model.

Ran the numbers of simple model on QMUL Batch Farm

Should also work for Cloud

Need to extract per user usage figures for Enhanced model

iris Carbon Mapping Project

Which Payload Model is Best?

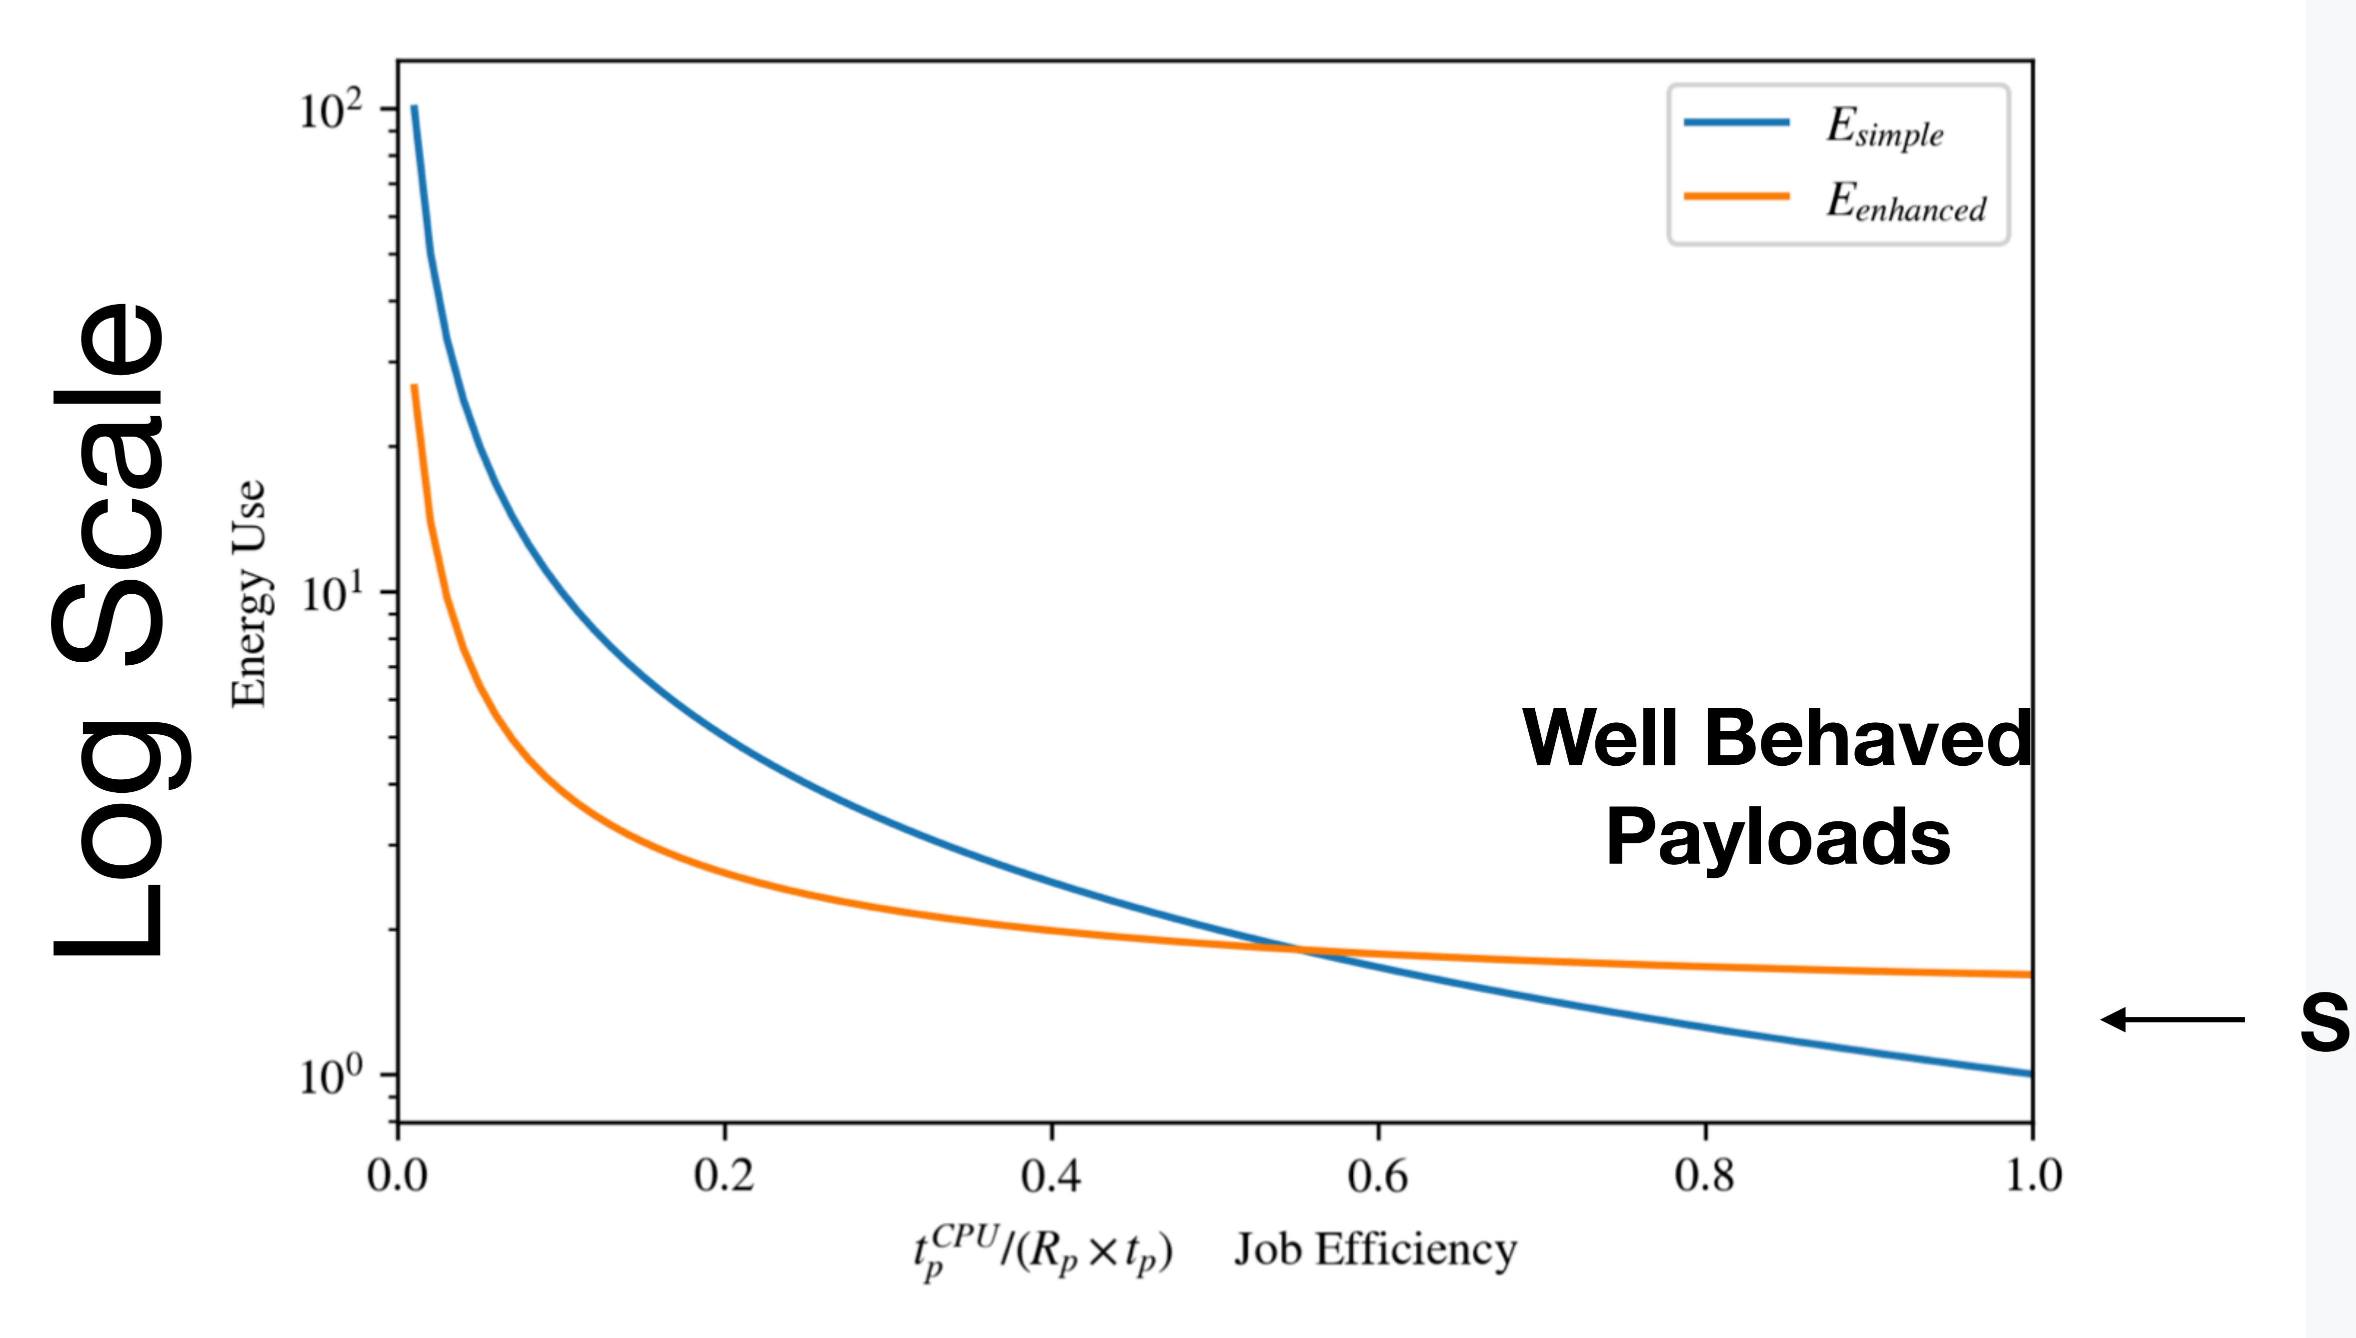

Log Scale

Well Behaved Payloads

Small Delta

Figure 3: Behaviour of the Simple and Enhanced payload models for a fixed amount of work (constant CPUtime) varying with Job Efficiency. Plotted on a log scale.

Keep it simple!

Not much to choose between them.

Both encourage more efficient code

Enhanced reduces Allocation to Providers

Which Storage Mode is Best? -> Can we get bytes read/written?

Meter
Vs
PDU
Vs
IPMI

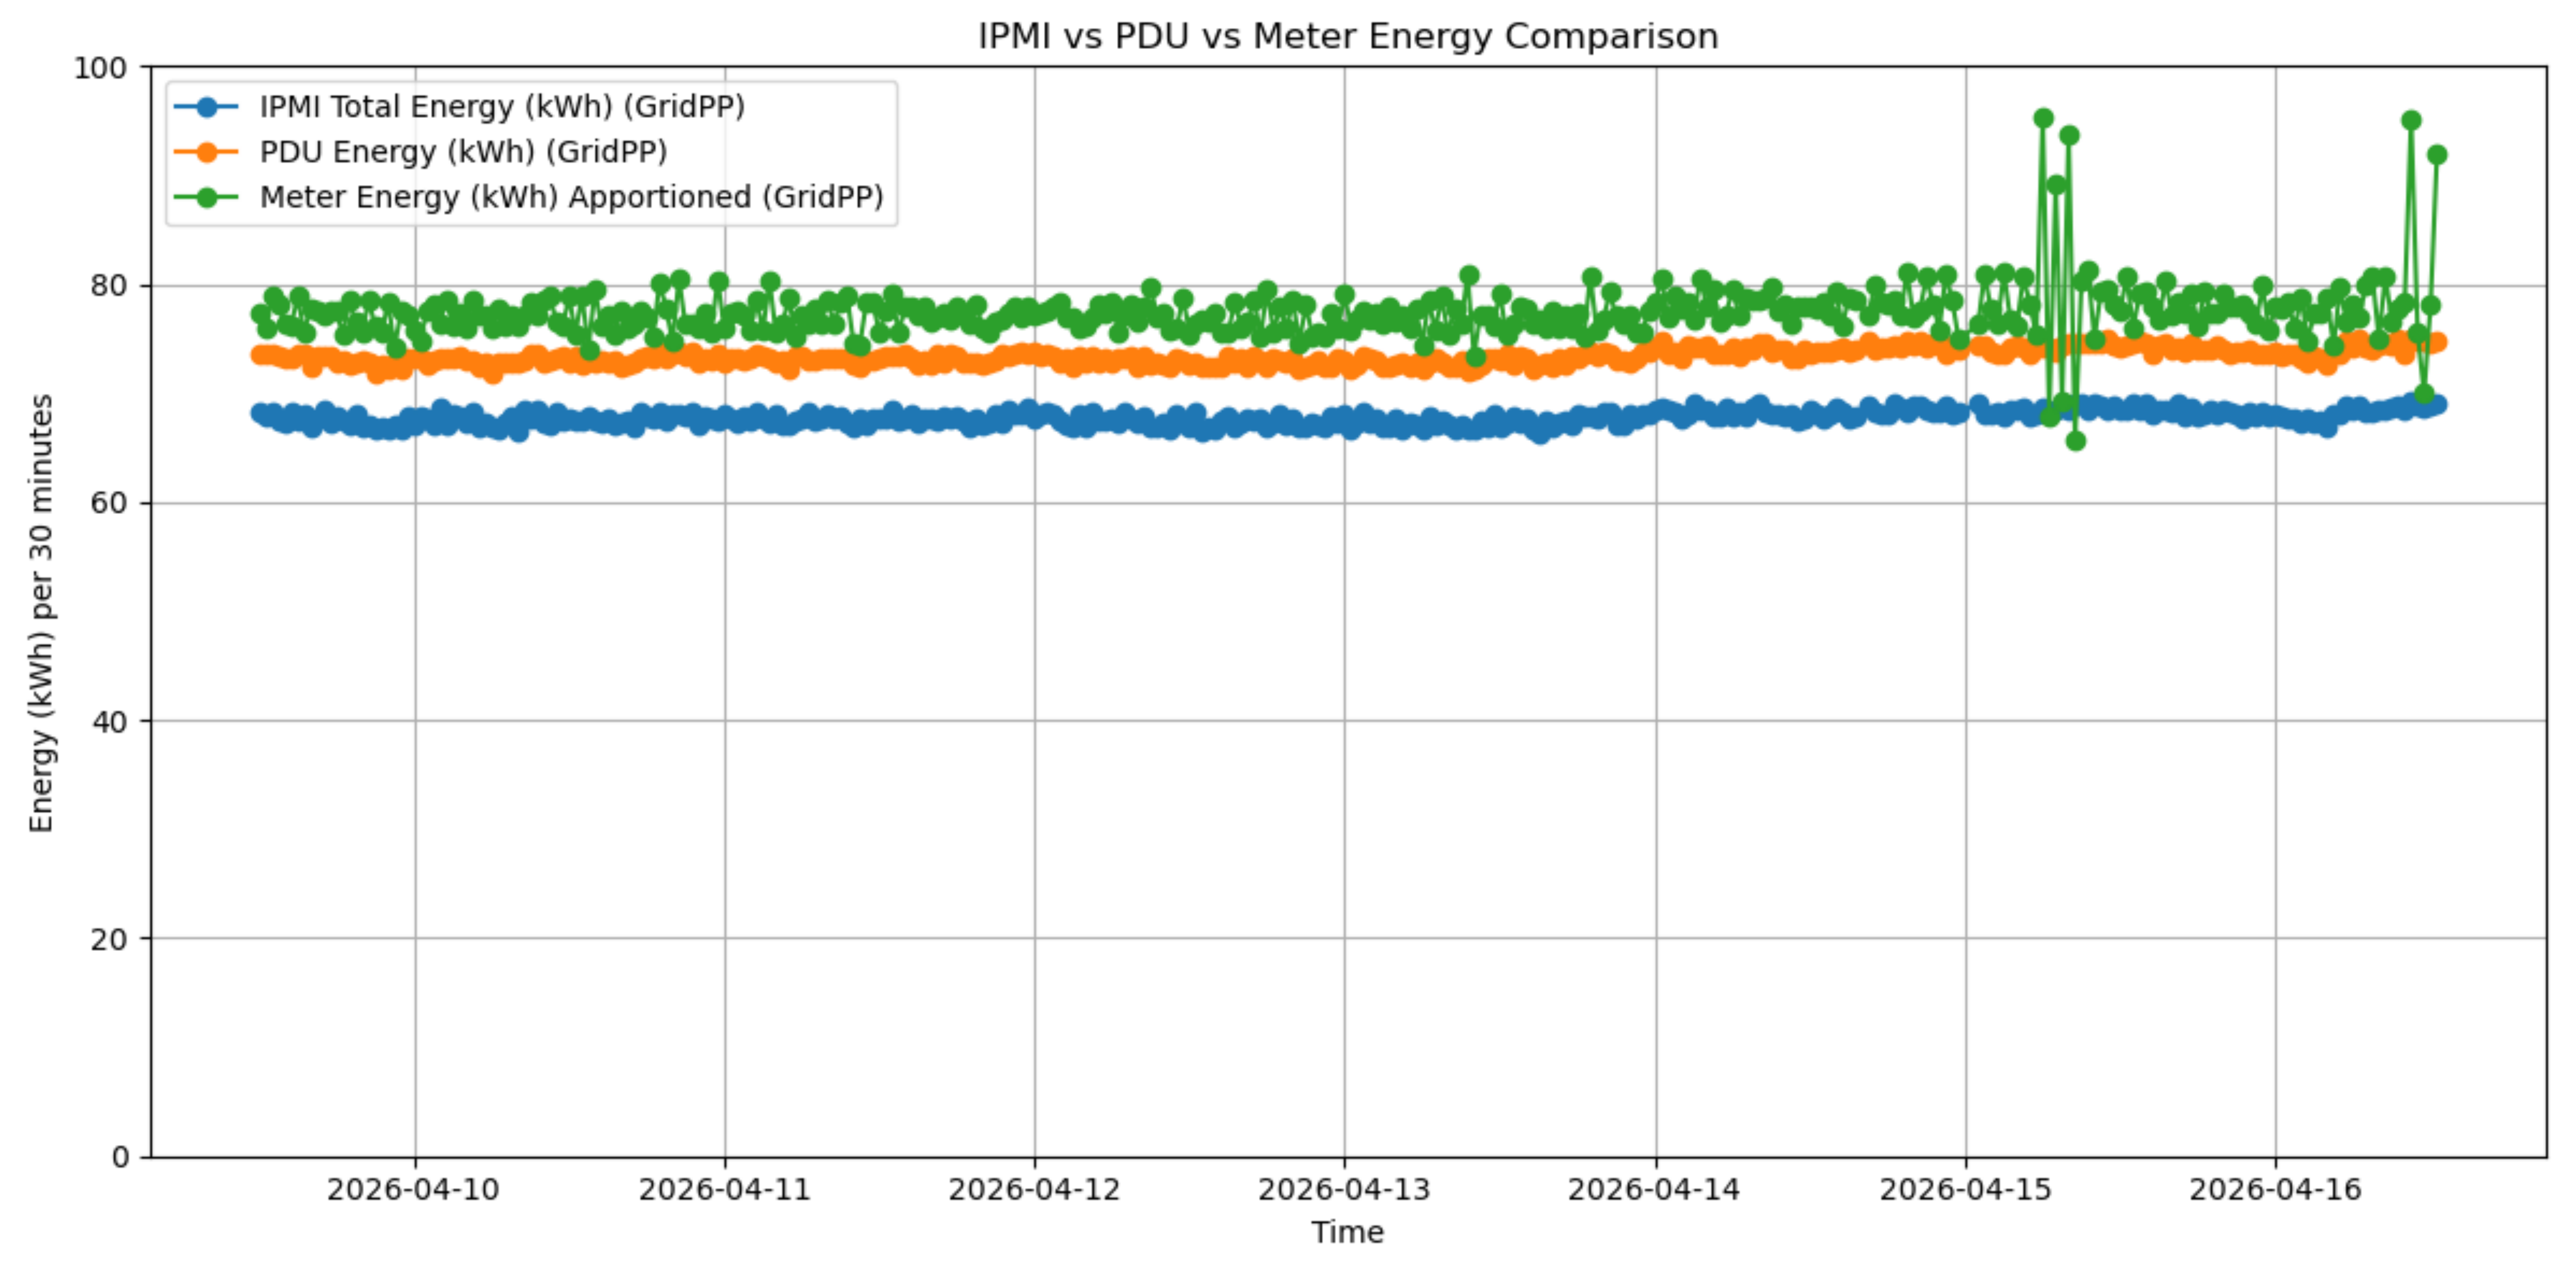

UPDATE

PDU’s seem better.
IPMI only on traditional server hardware:
What about switches?
What about ARM/RISC etc

UPDATE

Think like an accountant and apportion meter readings.

Apportion Carbon Costs

Carbon Accounting for UK Research Computing: proof of concept & beyond

Think like an accountant

Allocate Carbon Costs to Payloads

Simple models are probably sufficient

Embedded data is hard to find!

Keep an eye on NetDRIVE

Start simple & learn by doing

Thanks for listening

NetDRIVE

Speaker: Dr Alex Owen

r.a.owen@qmul.ac.uk

Queen Mary
University of London

Carbon Accounting for UK Research Computing: proof of concept & beyond

Backup Slides

Speaker: Dr Alex Owen

r.a.owen@qmul.ac.uk

NetZero The Big Picture

IRISCAST

Audit of Carbon Costs

iris CMP

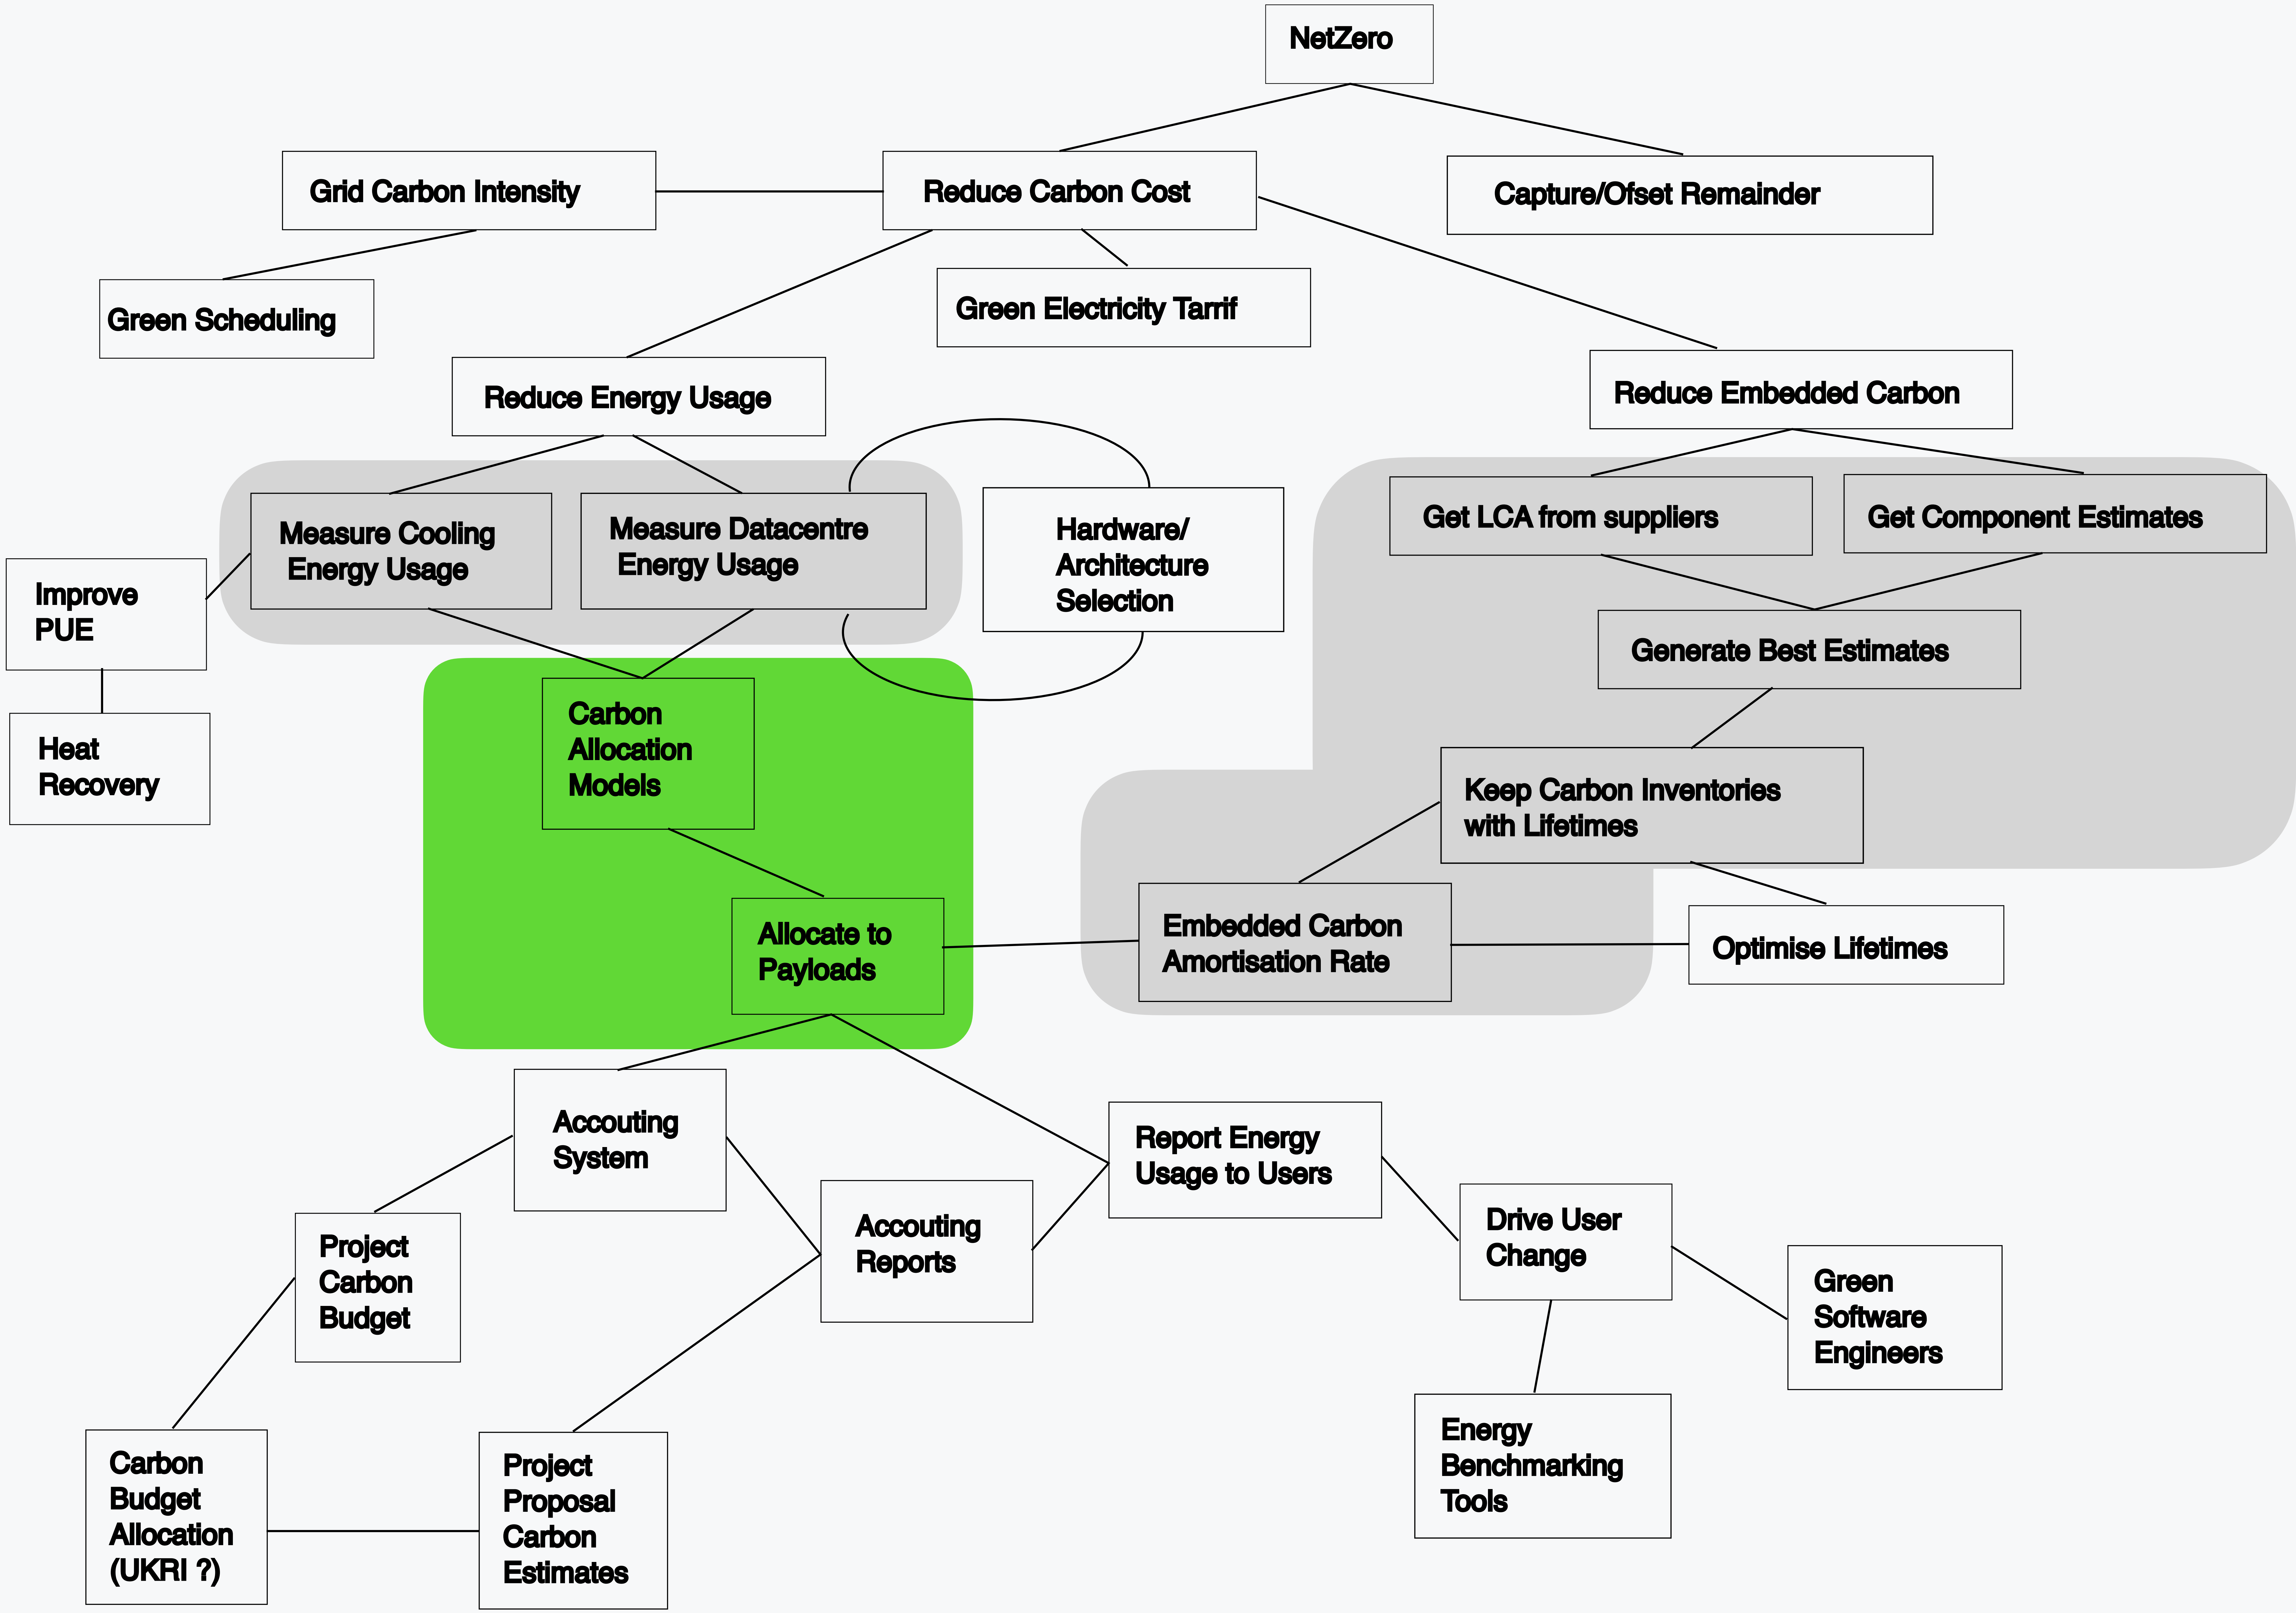

NetZero

Grid Carbon Intensity

Reduce Carbon Cost

Capture/Ofset Remainder

Green Scheduling

Green Electricity Tarrif

Reduce Energy Usage

Reduce Embedded Carbon

Measure Cooling Energy Usage

Measure Datacentre Energy Usage

Hardware/ Architecture Selection

Get LCA from suppliers

Get Component Estimates

Improve PUE

Heat Recovery

Generate Best Estimates

Carbon Allocation Models

Keep Carbon Inventories with Lifetimes

Allocate to Payloads

Embedded Carbon Amortisation Rate

Optimise Lifetimes

Accouting System

Report Energy Usage to Users

Project Carbon Budget

Accouting Reports

Drive User Change

Green Software Engineers

Energy Benchmarking Tools

Carbon Budget Allocation (UKRI ?)

Project Proposal Carbon Estimates

Based on figure in IRIS-CMP Report

GridPP

ECO-Grid Model

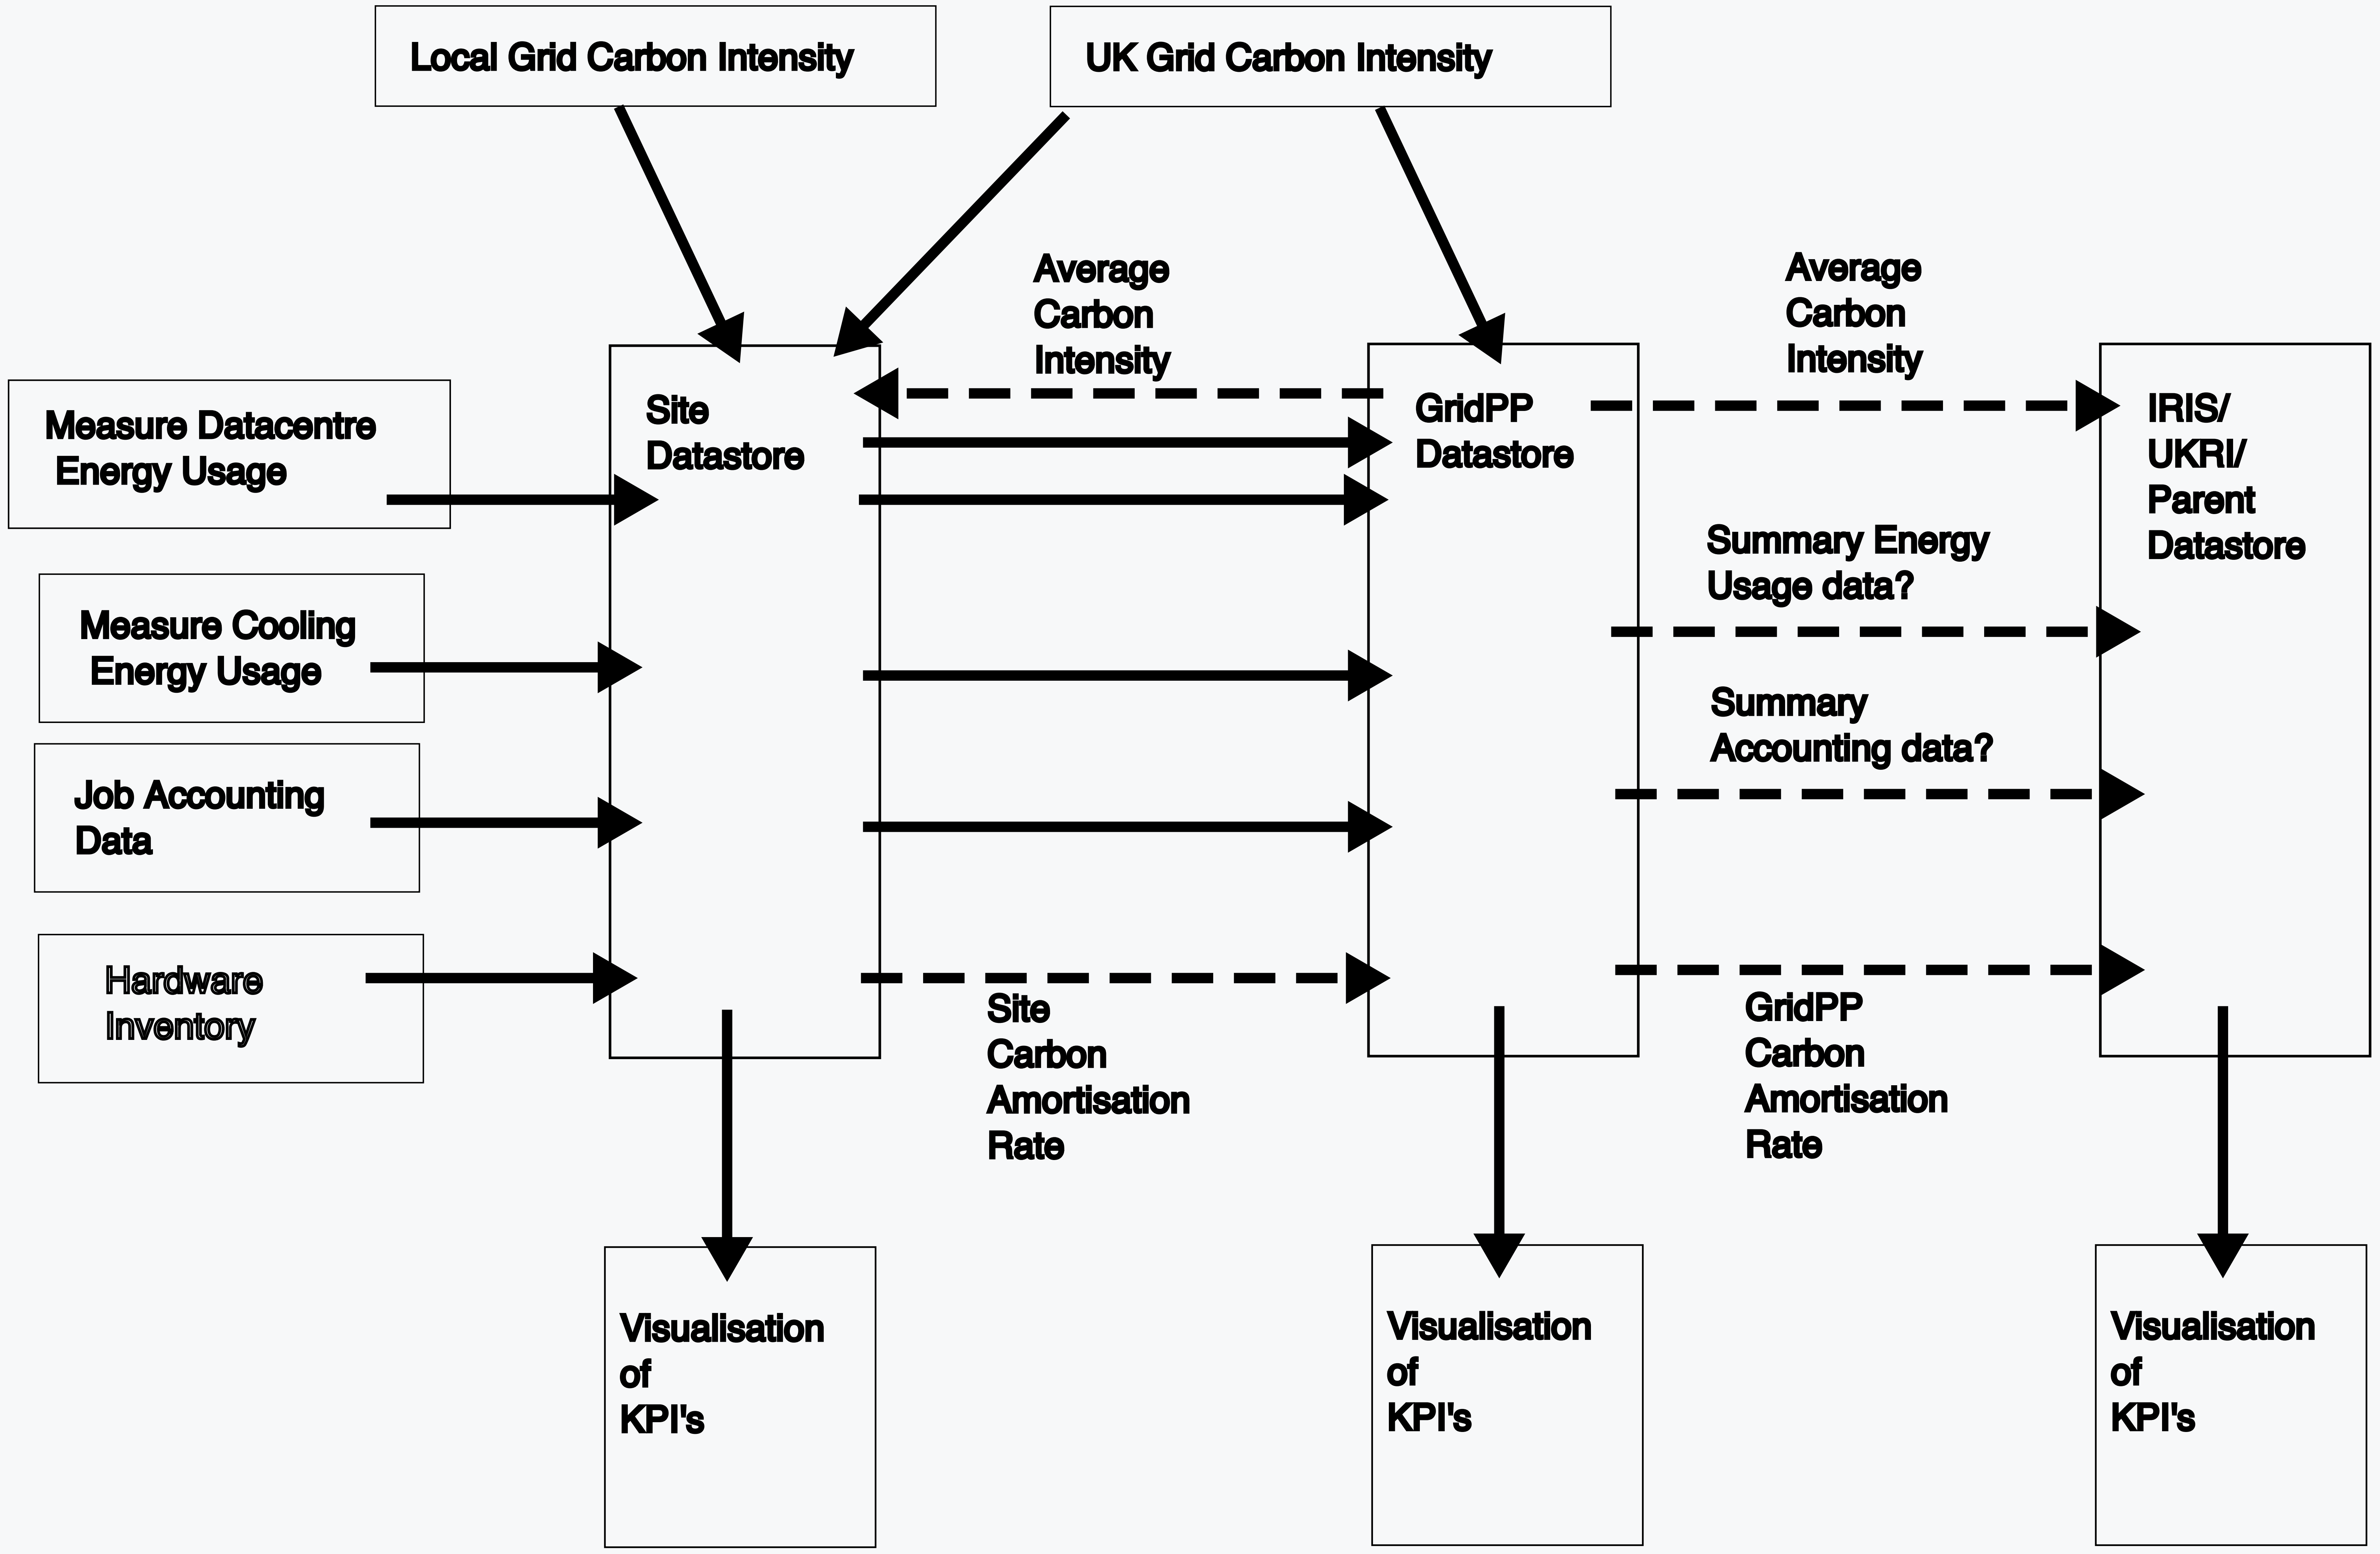

Local Grid Carbon Intensity

UK Grid Carbon Intensity

Measure Datacentre Energy Usage

Measure Cooling Energy Usage

Job Accounting Data

Hardware Inventory

Site Datastore

Average Carbon Intensity

Site Carbon Amortisation Rate

GridPP Datastore

Average Carbon Intensity

Summary Energy Usage data?

Summary Accounting data?

GridPP Carbon Amortisation Rate

IRIS/ UKRI/ Parent Datastore

Visualisation of KPI's

Visualisation of KPI's

Visualisation of KPI's

V0.0

iris

Carbon Mapping Project

Reporting Requirements / Concerns

Federation

Carbon costs of IRIS activity/providers broken down into scope 2 and scope 3.

Carbon costs of IRIS supported projects broken down into scope 2 and scope 3.

Carbon saved by being a federation

Reporting upwards:
 Benefit realisation, infrastructure efficiency
 Demonstrate right mix of platforms/tech
 Value of heterogeneity in the federation

Present success while continuing research
 Power used per hepspec
 Fossil power used per hepspc

Try to lead the narrative

Providers

Carbon costs of a provider’s service broken down by scope.

Allocate service carbon cost to users and idle/provider

Ease of implementation

Allocate maximum to users (minimum to idle/provider)

Users

Energy per job

Average IRIS Carbon Intensity

Average Embedded carbon factor

Try to avoid motivating behaviour that increase federation carbon costs.