Energy Consumption Study

REV 1

PEAKPROCESSING

Document Status

Job Reference:	Ogden Placement – Summer 2026
Issue Date/Version :	15 Sept 2026
Client :	Internal
Project :	Energy Consumption Study
Document Author :	Bobby Jinks & Jack Osborn
Approved By :	Steven Lloyd
Reason for Issue :	First issue.

Contents

DATE	VERSION	CHANGE
15 Sept 2026	1	First issue

1. Introduction

Peak Processing are offshore geophysical data processing and geohazard specialists, with one of their offices based in Wroxham. As part of an Ogden Trust placement, the authors (Bobby Jinks and Jack Osborn) have been tasked with researching how energy is being used across the company's small data centre, office space and other equipment, and finding ways in which energy efficiency could be improved.

At the start of the project, we had a briefing on the task that we would be undertaking and the systems that are used within the office, such as the server room, solar panels, and the desk setups. We drafted a plan for the next 4 weeks and created a list of questions to answer over the course of the project.

- What area uses the most energy? What is the total power consumption?
- What is the average energy produced by the solar panels?
- What is the most efficient server setup: 1 High Usage or multiple medium usage? What is the annual cost of energy?
- How can we collect and interpret the Octopus Energy data?

This report provides an overview of the project, with a copy of the initial brief and the final presentation contained within appendices.

2. Project Timeline

2.1. Week 1

During our first week, we created a skeletal plan of our project, including what we needed to monitor, the devices needed to monitor and collect data, and how we could present the data. We discussed using a program called LibreNMS, an open-source network monitoring system which discovers devices on the network automatically using SNMP and polls the devices to track their status and performance.

We created a closed off network using a switch, laptop, and a redundant backup server to download LIbreNMS and test if it functioned correctly. Following the installation guide, we downloaded and set up an Ubuntu 24.04 virtual machine, where we could install dependencies for different parts of LibreNMS, as well as create the database functionality using MariaDB and MySQL. We then configured PHP-FPM and nginx for the web server.

Once all the configuration was completed, we added the test server onto LIbreNMS to see how the polling system worked. Once we had LibreNMS installed, we moved it onto the network. Once that was complete, we went around every server pulling the id tag out and writing down the name and the idrac once we got all of that down, we then added all the servers onto to LibreNMS. Once the servers were added they started to ping back data which we monitored over the next two weeks, where we would write down things such as power usage spikes, drops, and average power consumption.

2.2. Week 2

During the second and third week we had also been tracking the energy production rate of the solar panels to see how much they were helping. We wrote down how much energy was produced every 30 minutes, 1 hour and each day over the course of the 2 weeks then discovered The Samil PowerSoftware that could also help to track the data without manual collection.

2.3. Week 3

By the end of the third week, we had added most electronics onto LibreNMS and started analysing the data collected from polling, however we were still missing things such as the aircon. Which we started to track by using the PDU we had ordered during the second week which had just arrived.

Once acquiring the PDU, we connected it to two AC units to track their power usage. We then discovered a script to connect the PDU to the network, which meant we could add it to LIbreNMS . After we ensured it was polling data, we left it over the weekend, so we had a good amount of information to analyse. However, only analysing data from the weekend would not give us the full picture, as the office is empty and the servers are not working, resulting in lower heat production and less AC usage.

2.4. Week 4

At the start of week four, we had collected enough data from the AC units and decided to move the PDU onto the garage mirror setup to track its power usage.

With both of these areas monitored, were able to make calculations based on assumptions and measured values for the equipment throughout the data centre.

3. Power Monitor (PDU) Selection

During week two, we started looking for a PDU to order for the project. A PDU (Power Distribution Unit) is a heavy-duty power strip built for server racks. The type of PDU we required was a metered PDU, which measures the electrical load passing through it and reports the results back to the user, on a built-in LED screen or over SNMP. We had a few requirements that the PDU needed to meet before we could decide. One of the main things we were looking for was SNMP support so that we could monitor and manage the PDU over the network.

We also needed a PDU with enough outlets for the equipment we planned to connect. Along with this, we looked at the price of each option to make sure it was within our budget. Delivery time was another important factor because we did not want the project to be delayed while waiting for equipment to arrive.

Our main goal was to find a multi-outlet metered PDU that met our technical requirements without costing too much. By comparing different products, their features, and their availability, we were able to identify a shortlist of suitable options for the project.

Once we got a few suitable options, we produced pros and cons of each to decide on which one had the best value. In the end we decided on one PDU (APC Metered Rack AP7850B PDU) which was affordable, SNMP enabled and had enough outlets for our task. We then sent our research over to our supervisor (Steve), who has been mentoring us for this project, for him to place the order.

The PDU was configured and added into LibreNMS to provide a portable power monitor which we could attach equipment to for a period of time and monitor alongside the other SNMP equipment.

4. LibreNMS & SNMP Monitoring

LibreNMS is an open-source network monitoring system which can be installed onto a server. Once it is installed, you can connect it to devices on the same network such as switches, routers, printers and servers. It can also automatically discover devices within the network and start to poll them to collect data about their health and activity.

After this data has started collating, you can access graphs from each device detailing minimums and maximums, as well as averages for aspects like current, voltage, temperature, and CPU usage. There are other features such as an alert system for when devices go down or other issues.

SNMP or Simple Network Management Protocol is a type of network protocol that is the standard way for network devices to report information about themselves to a monitoring system. Network protocols are established sets of rules and formats that determines how data processing works between different devices on a network.

SNMP is vital for this project as it allows us to connect the servers to LibreNMS to poll data from the servers and collect data from other sources that require manual collection at the same time.

The images below show the LibreNMS graphical interface, with some typical energy consumption graphs.

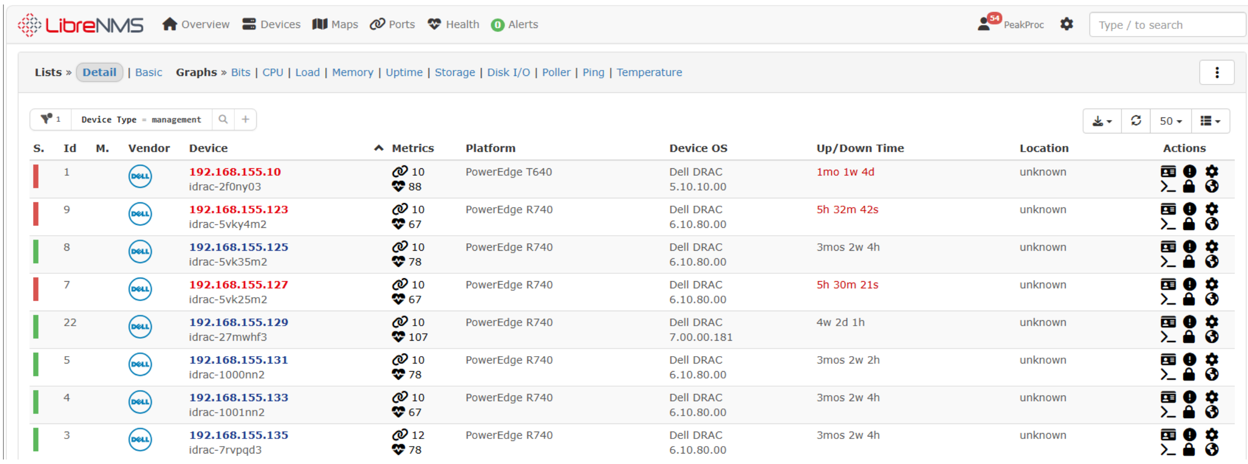

S.	Id	M.	Vendor	Device	Metrics	Platform	Device OS	Up/Down Time	Location	Actions
	1		DELL	192.168.155.10 idrac-2f0ny03	10 88	PowerEdge T640	Dell DRAC 5.10.10.00	1mo 1w 4d	unknown	
	9		DELL	192.168.155.123 idrac-5vky4m2	10 67	PowerEdge R740	Dell DRAC 6.10.80.00	5h 32m 42s	unknown	
	8		DELL	192.168.155.125 idrac-5vk35m2	10 78	PowerEdge R740	Dell DRAC 6.10.80.00	3mos 2w 4h	unknown	
	7		DELL	192.168.155.127 idrac-5vk25m2	10 67	PowerEdge R740	Dell DRAC 6.10.80.00	5h 30m 21s	unknown	
	22		DELL	192.168.155.129 idrac-27mwhf3	10 107	PowerEdge R740	Dell DRAC 7.00.00.181	4w 2d 1h	unknown	
	5		DELL	192.168.155.131 idrac-1000nn2	10 78	PowerEdge R740	Dell DRAC 6.10.80.00	3mos 2w 2h	unknown	
	4		DELL	192.168.155.133 idrac-1001nn2	10 67	PowerEdge R740	Dell DRAC 6.10.80.00	3mos 2w 4h	unknown	
	3		DELL	192.168.155.135 idrac-7rvpqd3	12 78	PowerEdge R740	Dell DRAC 6.10.80.00	3mos 2w 4h	unknown	

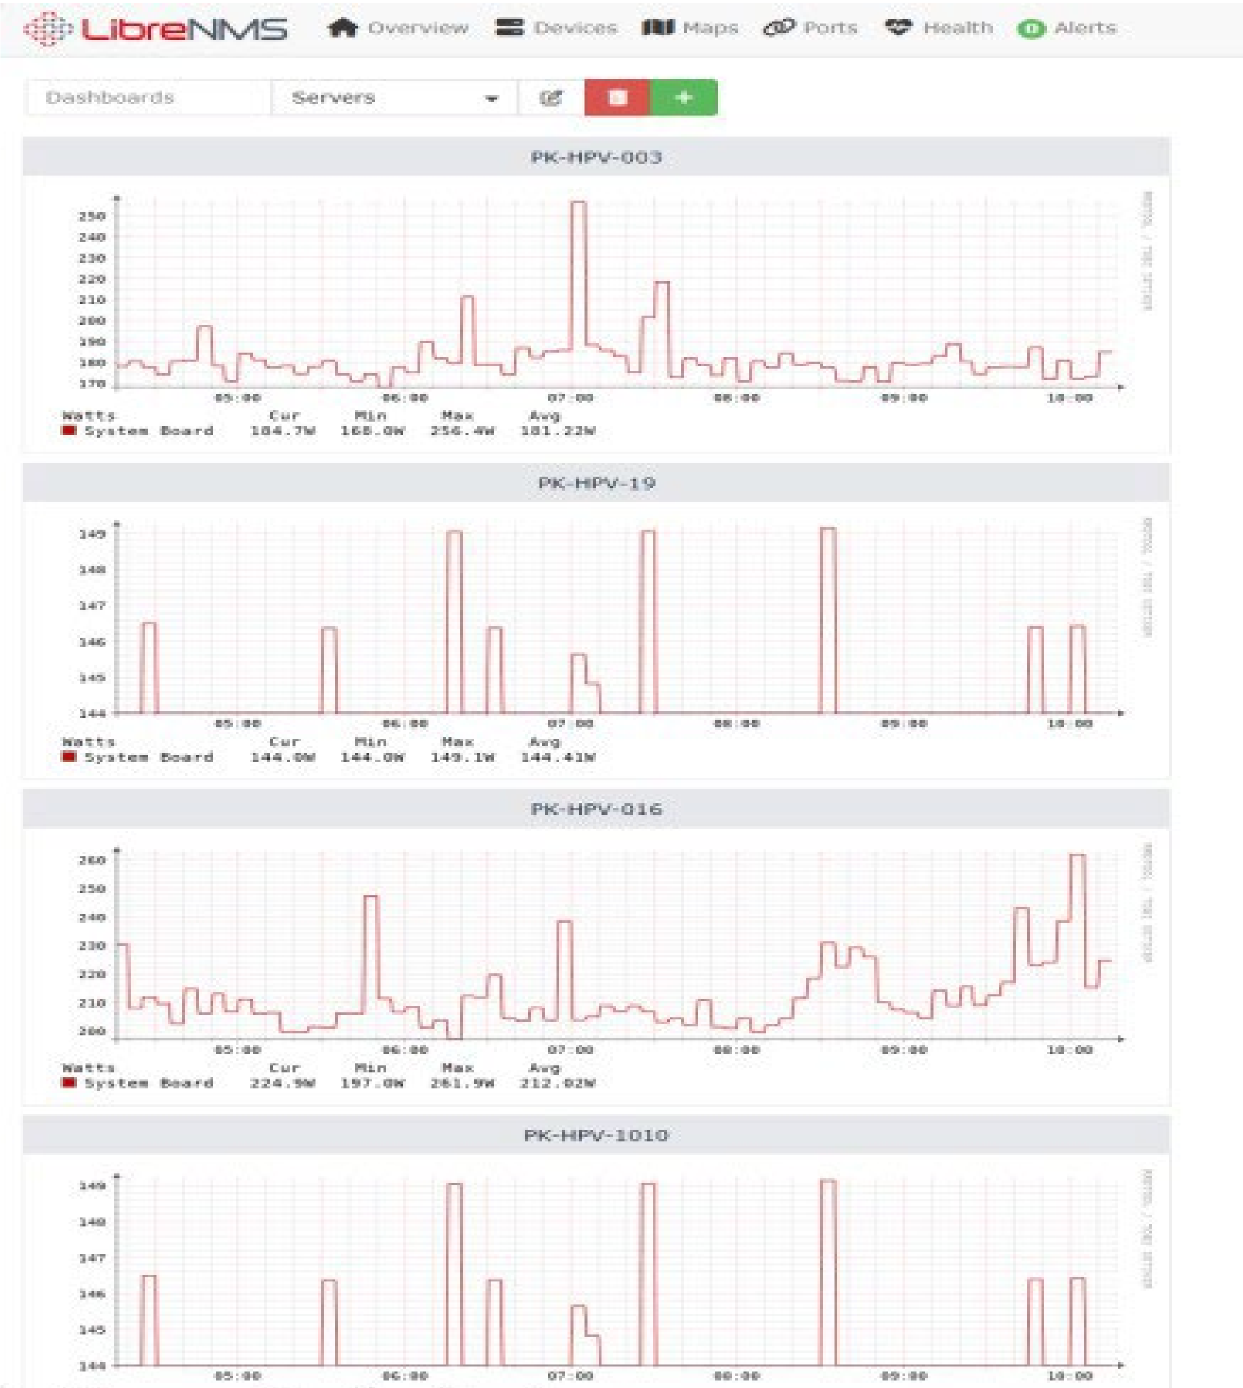
LibreNMS
Overview
Devices
Maps
Ports
Health
Alerts
Dashboards
Servers
PK-HPV-003
Watts Cur Min Max Avg
System Board 184.7W 168.0W 256.4W 181.22W
PK-HPV-19
Watts Cur Min Max Avg
System Board 144.0W 144.0W 149.1W 144.41W
PK-HPV-016
Watts Cur Min Max Avg
System Board 224.9W 197.0W 261.9W 212.02W
PK-HPV-1010

5. Solar Energy

The site's photovoltaic array is controlled by a pair of inverters which have the ability to log data at five-minute intervals to the Samil Power software, as seen in the image below :

In addition, manual reading were taken from a dedicated solar generation meter.

The two inverter exports provide eight complete overlapping days, 13 to 20 August, with combined measured generation of 217.4 kWh, or 27.2 kWh per complete day. Missing days were estimated from the measured daily solar output and the outside-temperature range. Temperature can distinguish some clear warm days from duller days, but cloud cover, orientation, shading and inverter availability are not captured. The 30-day solar result is therefore suitable for this study, but fuller monitoring would be beneficial.

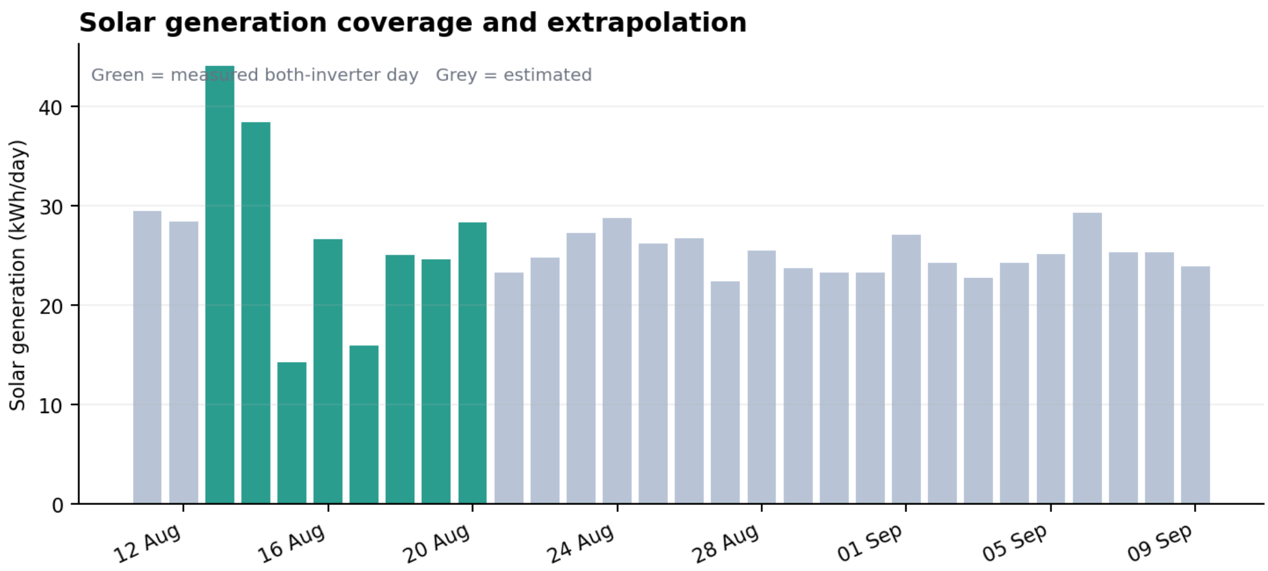

The principal finding concerns self-consumption, whereby the energy produced by the solar panels is never being exported to the grid as it is always less than the base load of the office. The economic consequence of this is actually favourable. Generated units displace imported electricity at the site's blended rate of approximately 23.6 p/kWh rather than being exported at a lower tariff, so annual generation represents an avoided cost of around £1,900 per year.

Additional capacity would likewise be absorbed by existing demand rather than exported.

6. Grid Energy

A spreadsheet was made available to us which imported the incoming Grid Energy reading from the Octopus Smart Meter using an API.

	A	B	C
1	consumption	interval_start	interval_end
2	5.635	2026-08-11T00:30:00+01:00	2026-08-11T01:00:00+01:00
3	5.846	2026-08-11T00:00:00+01:00	2026-08-11T00:30:00+01:00
4	5.916	2026-08-10T23:30:00+01:00	2026-08-11T00:00:00+01:00
5	5.819	2026-08-10T23:00:00+01:00	2026-08-10T23:30:00+01:00
6	5.854	2026-08-10T22:30:00+01:00	2026-08-10T23:00:00+01:00
7	5.865	2026-08-10T22:00:00+01:00	2026-08-10T22:30:00+01:00
8	6.262	2026-08-10T21:30:00+01:00	2026-08-10T22:00:00+01:00
9	6.812	2026-08-10T21:00:00+01:00	2026-08-10T21:30:00+01:00
10	6.905	2026-08-10T20:30:00+01:00	2026-08-10T21:00:00+01:00
11	7.064	2026-08-10T20:00:00+01:00	2026-08-10T20:30:00+01:00
12	3.664	2026-08-10T19:30:00+01:00	2026-08-10T20:00:00+01:00
13	3.986	2026-08-10T19:00:00+01:00	2026-08-10T19:30:00+01:00
14	3.726	2026-08-10T18:30:00+01:00	2026-08-10T19:00:00+01:00
15	3.641	2026-08-10T18:00:00+01:00	2026-08-10T18:30:00+01:00
16	3.806	2026-08-10T17:30:00+01:00	2026-08-10T18:00:00+01:00
17	3.64	2026-08-10T17:00:00+01:00	2026-08-10T17:30:00+01:00
18	6.969	2026-08-10T16:30:00+01:00	2026-08-10T17:00:00+01:00
19	7.151	2026-08-10T16:00:00+01:00	2026-08-10T16:30:00+01:00
20	7.148	2026-08-10T15:30:00+01:00	2026-08-10T16:00:00+01:00
21	6.889	2026-08-10T15:00:00+01:00	2026-08-10T15:30:00+01:00
22	6.918	2026-08-10T14:30:00+01:00	2026-08-10T15:00:00+01:00
23	6.14	2026-08-10T14:00:00+01:00	2026-08-10T14:30:00+01:00
24	6.157	2026-08-10T13:30:00+01:00	2026-08-10T14:00:00+01:00
25	5.756	2026-08-10T13:00:00+01:00	2026-08-10T13:30:00+01:00
26	5.825	2026-08-10T12:30:00+01:00	2026-08-10T13:00:00+01:00
27	5.63	2026-08-10T12:00:00+01:00	2026-08-10T12:30:00+01:00
28	5.562	2026-08-10T11:30:00+01:00	2026-08-10T12:00:00+01:00
29	6.669	2026-08-10T11:00:00+01:00	2026-08-10T11:30:00+01:00
30	6.885	2026-08-10T10:30:00+01:00	2026-08-10T11:00:00+01:00
31	6.969	2026-08-10T10:00:00+01:00	2026-08-10T10:30:00+01:00
32	6.406	2026-08-10T09:30:00+01:00	2026-08-10T10:00:00+01:00
33	6.083	2026-08-10T09:00:00+01:00	2026-08-10T09:30:00+01:00
34	5.509	2026-08-10T08:30:00+01:00	2026-08-10T09:00:00+01:00
35	4.967	2026-08-10T08:00:00+01:00	2026-08-10T08:30:00+01:00
36	5.321	2026-08-10T07:30:00+01:00	2026-08-10T08:00:00+01:00
37	5.682	2026-08-10T07:00:00+01:00	2026-08-10T07:30:00+01:00
38	5.639	2026-08-10T06:30:00+01:00	2026-08-10T07:00:00+01:00
39	5.847	2026-08-10T06:00:00+01:00	2026-08-10T06:30:00+01:00
40	6.922	2026-08-10T05:30:00+01:00	2026-08-10T06:00:00+01:00
41	6.902	2026-08-10T05:00:00+01:00	2026-08-10T05:30:00+01:00
42	7.311	2026-08-10T04:30:00+01:00	2026-08-10T05:00:00+01:00
43	9.788	2026-08-10T04:00:00+01:00	2026-08-10T04:30:00+01:00
44	10.018	2026-08-10T03:30:00+01:00	2026-08-10T04:00:00+01:00
45	9.978	2026-08-10T03:00:00+01:00	2026-08-10T03:30:00+01:00
46	10.211	2026-08-10T02:30:00+01:00	2026-08-10T03:00:00+01:00
47	10.109	2026-08-10T02:00:00+01:00	2026-08-10T02:30:00+01:00
48	10.391	2026-08-10T01:30:00+01:00	2026-08-10T02:00:00+01:00
49	10.533	2026-08-10T01:00:00+01:00	2026-08-10T01:30:00+01:00
50	7.518	2026-08-10T00:30:00+01:00	2026-08-10T01:00:00+01:00
51	7.571	2026-08-10T00:00:00+01:00	2026-08-10T00:30:00+01:00
52	7.646	2026-08-09T23:30:00+01:00	2026-08-10T00:00:00+01:00
53	7.534	2026-08-09T23:00:00+01:00	2026-08-09T23:30:00+01:00
54	7.535	2026-08-09T22:30:00+01:00	2026-08-09T23:00:00+01:00
55	7.485	2026-08-09T22:00:00+01:00	2026-08-09T22:30:00+01:00

This allowed us to analyse the import electricity for the full period of the study.

7. Battery operation and load shape

A set of Vicrton Inverters with large Lithium battery banks are installed. These primarily provide resilience to the power supply for the onsite equipment, but are also used to move some of the electrical load to battery from peak grid CO2 footprint periods (afternoon) to lower CO2 footprint periods (early morning).

The graph below shows the energy import from the Octopus meter, which clearly shows an overnight charging hump and a corresponding evening trough. The method identifies approximately 14.1 kWh per day on average.

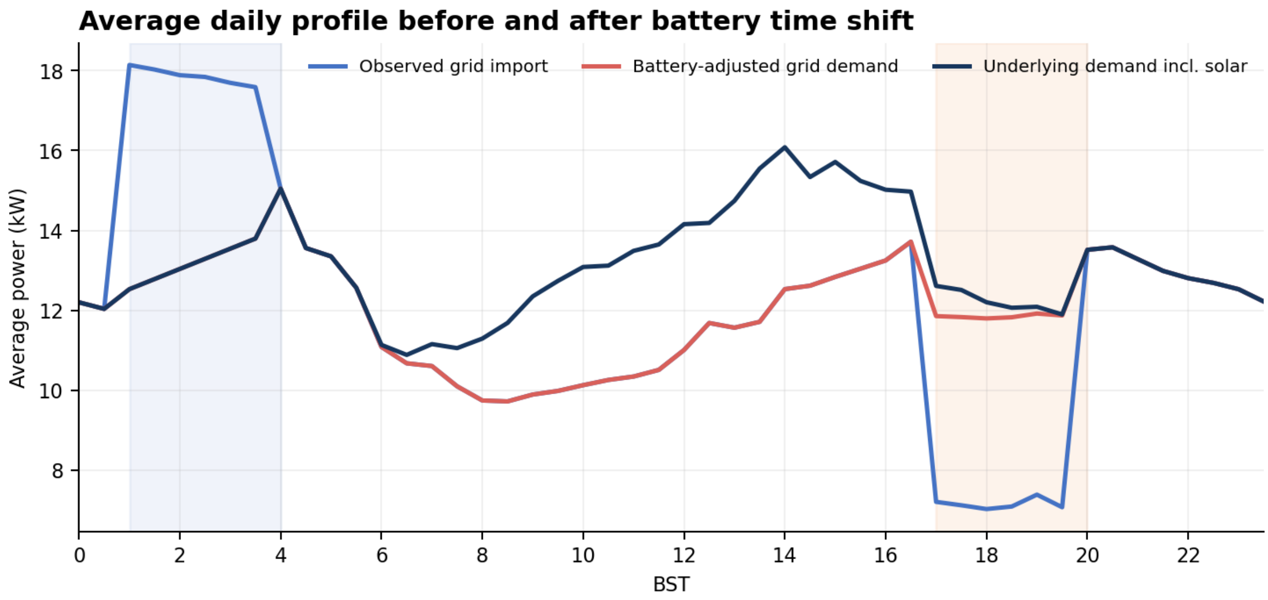

The strongest observed depression is approximately 17:00 to 20:00 BST, indicating the batteries are supplying the load between these periods.

8. Cooling and Ambient Temperature

Between 6:00 PM on Friday 14 August and 6:00 PM on Monday 17 August, the air conditioning unit was monitored continuously using the PDU, and the resulting power consumption data is shown in the graph below.

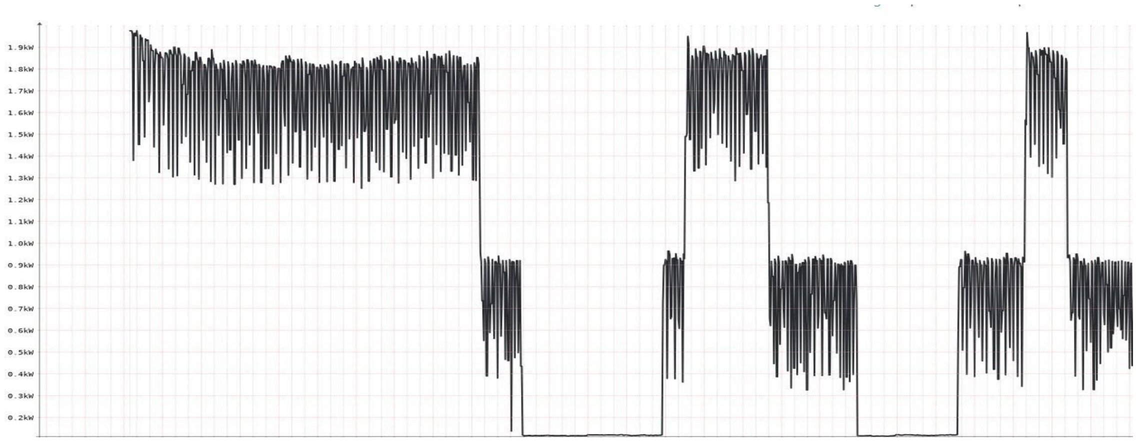

The graph shows that the air conditioning is typically drawing between 1.8 and 1.9 kW when cooling demand is high and around 0.8 to 0.9 kW during lighter operation. This variation is largely influenced by heat levels. As indoor and outdoor temperatures rise, the AC works harder and runs for longer periods, increasing its power consumption.

During cooler periods, the unit reduces output or cycles off, causing energy use to fall significantly.

The mean outside temperature was available at 15min intervals within a spreadsheet. These temperatures were used to estimate the air conditioning loads. The Server Room cooling model estimates 910 kWh per month for all four units. Office air conditioning is estimated at around 1,557 kWh over 30 days.

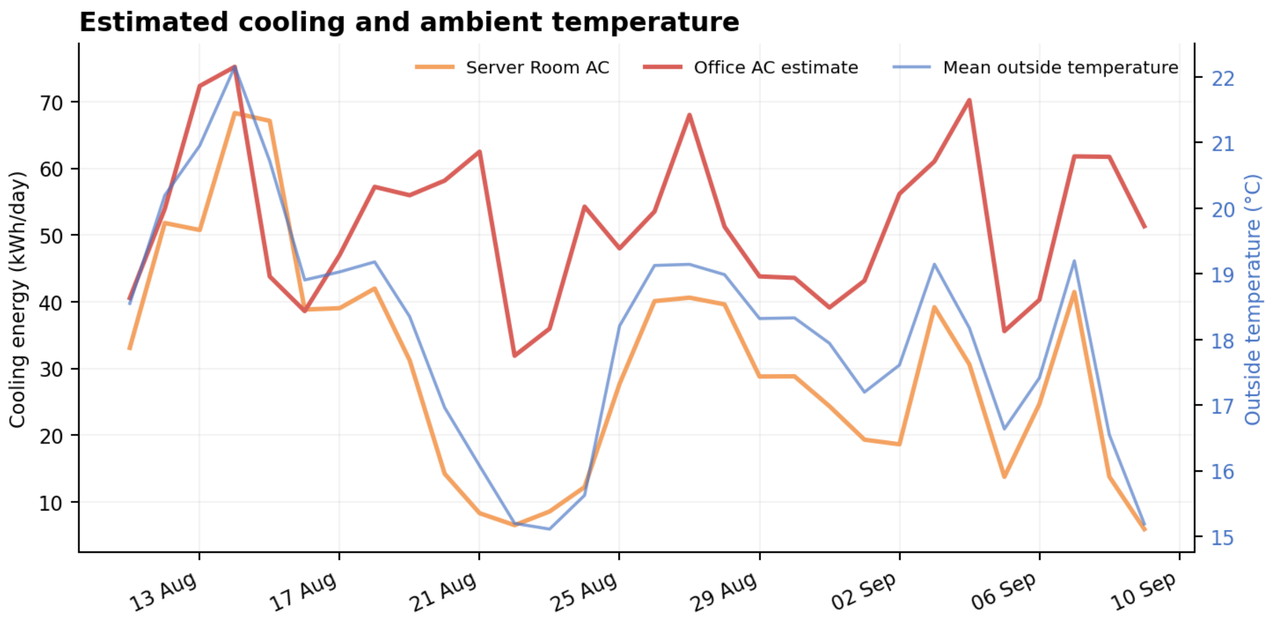

9. Conclusion

In conclusion, we met the project brief by identifying how energy is used across the building and where efficiency could be improved. Although setting up LibreNMS and collecting data from different areas was challenging, it gave us a clear view of the main sources of power consumption.

The results show that the servers are the largest energy user, followed by air conditioning and garage equipment. Smaller office devices use much less power but still contribute to the total daily consumption.

The solar panels are already being used effectively, with most generated energy consumed on site. Further improvements could include reviewing server utilisation, reducing unnecessary overnight loads, optimising air conditioning settings, and continuing long-term monitoring.

A presentation of our findings, as well as the original brief, are contained within the appendices of this report.

Appendix 1. - Initial Project Brief

PEAKPROCESSING

Work Placement Project Scope

Energy Consumption Study for Small Data Centre and Office Building

1. Project Aim

The aim of this project is to study how energy is used across the company's small data centre, office space, and supporting equipment. The students will identify the main areas of energy consumption, understand what drives that usage, and suggest practical ways to improve energy efficiency without reducing reliability, safety, or business performance.

The project should be suitable for an A-Level aged work placement students, with appropriate supervision. The students will not carry out electrical work or make changes to live systems without approval.

---

2. Project Objectives

The project should aim to answer the following questions:

1. **Where is energy being used?**

   Identify the main consumers of electricity, such as servers, storage, networking equipment, UPS systems, cooling, lighting, office equipment, and general building services.

2. **How much energy is each area using?**

   Estimate or measure the energy consumption of major equipment groups.

3. **How does IT workload affect energy use?**

   Compare the power consumption of servers at low load, normal load, and higher load.

4. **Is the current equipment usage efficient?**

   Consider whether fewer servers running at higher utilisation could be more efficient than more servers running at low utilisation.

5. **How significant is cooling?**

   Understand how air conditioning, ventilation, and office temperature settings affect energy consumption.

6. **What efficiency savings may be possible?**

   Propose realistic changes, such as server consolidation, shutdown schedules, cooling adjustments, equipment replacement, airflow improvements, or monitoring improvements.

7. **What are the risks or trade-offs?**

   Consider resilience, redundancy, staff comfort, hardware reliability, business continuity, and data security.

---

3. Scope of Study

The project should include:

IT Equipment

- Physical servers
- Storage arrays
- Network switches, firewalls, and routers
- Backup equipment (Garage)
- UPS system (Garage)

Cooling and Environment

- Air conditioning units
- Server room temperature
- Office temperature
- Ventilation and airflow
- Impact of raising cooling set-points slightly

Building and Office Loads

- Lighting
- Desktop PCs and monitors
- Printers and office equipment
- Kitchen or welfare equipment, if relevant
- Out-of-hours energy use

Exclusions

The project should not include:

- Opening electrical distribution boards
- Modifying mains wiring
- Changing server room cooling settings without approval
- Reconfiguring live IT systems without supervision
- Any work that affects systems, backups, security, or business continuity

---

4. General Methodology

Stage 1: Introduction and Background Research

The students should begin by understanding the purpose of the building and data centre.

They should research basic energy concepts including:

- Power, measured in watts or kilowatts
- Energy, measured in kilowatt-hours
- The difference between power draw and energy usage over time
- IT equipment load versus power consumption
- Why data centres require cooling
- The concept of Power Usage Effectiveness, or PUE
- Basic carbon and cost impact of electricity consumption

The students should also be introduced to the company's IT layout at a high level, including what equipment is business-critical.

Stage 2: Equipment Inventory

The students should create a simple inventory of major energy-consuming equipment.

For each item or equipment group, they should record:

- Equipment name or type
- Location
- Quantity
- Rated power, if shown on the label or specification sheet
- Estimated normal power use
- Whether it runs 24/7 or only during working hours
- Business purpose
- Whether it is critical, important, or non-critical

Example categories:

Area	Equipment	Quantity	Runs 24/7?	Estimated Power	Notes
Server room	Hyper-V servers	4	Yes	To measure	Business-critical
Server room	Network switches	3	Yes	To estimate	Always on
Office	Desktop PCs	10	No	To estimate	User equipment
Cooling	Air conditioning	1–2	Variable	To measure/estimate	Depends on temperature

Stage 3: Energy Measurement

Where possible, the students should collect real measurements rather than relying only on rated power.

Possible data sources include:

- Smart meter or half-hourly electricity readings
- UPS load readings
- Rack PDU readings, if available
- Server management tools such as iDRAC, iLO, or IPMI
- Virtualisation platform performance data
- Plug-in power meters for non-critical office equipment
- Air conditioning controller data
- Electricity bills
- Manual readings taken at regular intervals

The students should record:

- Typical daytime power use
- Overnight and weekend baseload
- Server room load
- Office load
- Cooling load, if measurable

- Any spikes or unusual patterns

A simple daily log could be used:

Date/Time	Total Load	Building	Server/UPS Load	Office Activity	Cooling Status	Notes
Monday 09:00				Busy	On	Normal working day
Monday 18:00				Low	On	After office hours
Saturday 12:00				None	On/Off	Weekend baseline

Stage 4: Workload and Server Utilisation Study

The students should review whether servers are being used efficiently.

They should compare:

- CPU utilisation
- Memory utilisation
- Disk activity
- Network activity
- Power consumption
- Time of day
- Business workload

The study should consider whether equipment is lightly loaded for long periods and whether workloads could be consolidated.

Example questions:

- Are some servers mostly idle?
- Do servers use almost the same power when idle as when moderately loaded?
- Would fewer, newer servers running at higher utilisation use less power?
- What redundancy would be lost if servers were consolidated?
- Are there systems that could be powered down outside working hours?
- Are there old servers that consume disproportionate power?

A useful comparison would be:

Scenario	Number of Servers	Average Load	Estimated Power	Resilience	Notes
Current setup	4	Low/medium		High	Existing arrangement
Consolidated option	3	Medium		Medium/high	Possible saving
Aggressive consolidation	2	High		Lower	Higher risk

Stage 5: Cooling and Environmental Study

The students should investigate the relationship between cooling and energy use.

They should record:

- Current server room temperature
- Office temperature
- Air conditioning set-points
- Temperature variation during the day
- Whether cooling runs continuously
- Whether server room and office cooling are linked
- Whether airflow is efficient
- Whether hot air is being recirculated into equipment intakes

Possible areas to investigate:

- Raising the cooling set-point by 1–2°C
- Improving airflow around racks
- Separating hot and cold air paths
- Closing gaps in racks
- Removing unnecessary heat sources
- Cleaning filters or checking maintenance schedules
- Avoiding over-cooling the office or server area

Any temperature change should be approved and monitored carefully.

Stage 6: Cost and Carbon Estimate

The students should convert energy use into estimated cost.

For each equipment group:

```
Energy per year = Power in kW × Hours used per year
```

```
Annual cost = Energy per year × Electricity cost per kWh
```

Example:

A device using 500 W continuously:

```
0.5 kW × 24 hours × 365 days = 4,380 kWh per year
```

At 25p per kWh:

```
4,380 × £0.25 = £1,095 per year
```

This helps compare small loads that run continuously with larger loads that only run occasionally.

The students may also estimate carbon impact using a published UK electricity carbon factor, if required.

Stage 7: Identify Efficiency Opportunities

The students should group potential savings into categories.

Quick Wins

- Turn off unused monitors, PCs, or office equipment
- Review out-of-hours power use
- Check power management settings
- Remove unused hardware
- Improve labelling and monitoring
- Adjust cooling timers or set-points, subject to approval

Medium-Term Improvements

- Consolidate lightly used servers
- Replace inefficient older hardware
- Improve rack airflow
- Improve environmental monitoring
- Use more efficient UPS or power distribution

- Schedule non-critical processing jobs efficiently

Long-Term Improvements

- Data centre redesign
- More efficient cooling system
- Cloud or hybrid hosting comparison
- Solar or battery storage review
- Formal energy monitoring system
- Hardware refresh strategy based partly on energy efficiency

---

5. Deliverables

The students should produce:

1. **Equipment energy inventory**

   A list of major energy-consuming equipment and estimated usage.

2. **Energy use summary**

   A simple breakdown of where electricity is being consumed.

3. **Graphs or charts**

   For example:

   - Day versus night energy use
   - Server load versus power consumption
   - Estimated annual cost by equipment group
   - Cooling impact over time

4. **Efficiency opportunities list**

   A ranked list of possible savings.

5. **Final short report**

   A clear summary of findings, assumptions, measurements, and recommendations.

6. **Presentation to the company**

   A short presentation explaining what was found and what could be improved.

---

6. Suggested Final Report Structure

1. Introduction
2. Description of the building and IT environment
   Method used
3. Equipment inventory
4. Energy measurements
5. Server utilisation findings
6. Cooling and environmental findings
7. Estimated costs
8. Efficiency opportunities
9. Risks and trade-offs
10. Recommendations
11. Further work

---

7. Safety, Supervision and Practical Constraints

The students should be supervised when working in the server room or around operational equipment.

They should not:

- Unplug equipment
- Change server settings
- Change air conditioning settings
- Open electrical cabinets
- Modify network or power connections
- Access customer data
- Install software without permission

The students may:

- Read labels
- Record observations
- Take supervised meter readings
- Use approved monitoring dashboards
- Use spreadsheet calculations
- Produce graphs and summaries
- Suggest changes for review

8. Expected Learning Outcomes

By the end of the placement, the students should understand:

- How energy is consumed in a small data centre and office
- Why IT equipment often consumes power even when lightly loaded
- How cooling contributes to total energy use
- How to measure and estimate electricity consumption
- How to convert technical measurements into business cost
- How to balance efficiency, resilience, comfort, and operational risk
- How to present technical findings clearly to a business audience

9. Success Criteria

The project will be successful if it produces:

- A clear picture of the main energy consumers
- A reasonable estimate of annual energy cost by category
- Evidence-based observations rather than guesses
- Practical recommendations for efficiency improvements
- A distinction between low-risk quick wins and higher-risk changes
- A final report that can be understood by both technical and non-technical staff

Appendix 2. - Presentation of Findings

Energy Usage

Where our energy goes across servers, cooling and electronics – and how we measure and reduce it.

INTRODUCTION

Every watt, accounted for

A look at how much energy our server room uses, where that energy goes, and the practical changes that could bring it down.

24/7
Always-on load

3 areas
Servers · cooling · electronics

Live data
Measured, not estimated

Why this review matters

- Servers, cooling and electronics run around the clock, so energy is one of our largest ongoing costs
- Rising energy prices mean even small inefficiencies add up to real money over a year
- Understanding where the power goes is the first step to cutting waste without risking reliability

What this deck covers

- Where energy is used across servers, cooling and other electronics, measured from real data
- The trade-offs behind our server and cooling choices, and where redundancy pays off
- How we monitor usage, what the solar panels contribute, and where we can save

AGENDA

Questions

- What takes up the most energy?
- What is the total power being used?
- How much energy do the solar panels produce on average?
- What is the energy usage of each area as a percentage?
- What are the advantages and disadvantages of using one high-usage server rather than two average-usage servers?
- What is the annual cost of energy by category?
- What equipment is needed to collect this data?
- What is the most optimal setup for AC units: One large unit or multiple smaller units?
- How can we collect and interpret the Octopus Energy data?

SECTION 01

Servers

The trade-offs of running one high-usage server, the risk of overheating, and how we build in resilience.

01

SERVERS · TRADE-OFFS

High-usage server

Advantages (easy version)

- **Better use of powerful server:** the high-usage server handles the main workload efficiently.
- **Cost savings:** the smaller server is cheaper, since it only handles light tasks.
- **Focused management:** easier to monitor and manage the server doing most of the work.
- **Some backup support:** the smaller server can help if the main server has minor issues.
- **Flexible usage:** the low-usage server can run experiments or secondary services without affecting main operations.

Disadvantages (easy version)

- **Main server can become a bottleneck:** if it goes down, most work stops.
- **Small server might be wasted:** often underused, which is less efficient financially.
- **Harder to handle spikes:** the main server may struggle if traffic suddenly increases.
- **Uneven load:** one server is very busy, the other is mostly idle, which can lead to performance imbalance.
- **Limited redundancy:** not a full backup setup; failover may need extra changes.
- **Maintenance problems:** fixing the busy server can hurt performance; the smaller server offers little protection.

SERVERS

Advantages of a high-usage server

- **Resource optimization:** the high-usage server is fully utilized for intensive tasks, giving maximum performance for critical operations.
- **Simplified management for the major workload:** concentrating the major load on one central server allows focused monitoring, tuning, and resource allocation.
-

- **Simplified Management and Maintenance**

Single OS Environment: One operating system, one backup policy, and unified fire wall settings reduce administrative overhead. Troubleshooting logs and performance metrics are centralized.

- **Lower Operational Complexity**: No need for load balancers, replication, or complex failover mechanisms, simplifying deployments and updates.

SERVERS

Disadvantages of a high-usage

- **Single point of failure:** the high-usage server becomes a bottleneck; if it fails, most operations are impacted, reducing system reliability.
- **Scalability limitations:** high-usage pressure on one server can limit the ability to handle sudden spikes in traffic or compute load without significant hardware upgrades.
- **Limited redundancy:** system redundancy is not fully optimized; true failover would require additional mechanisms to redistribute load during outages.
- **Maintenance risk:** maintenance on the heavily used server could degrade performance.

Dedicated High-Performance Resources

- A high end server typically dedicates all CPU, memory, storage and network bandwidth to a single tenant or workload.
- This **ensures optimal utilization without contention**, whereas multiple n ormal servers may have uneven load distribution and idle resources.

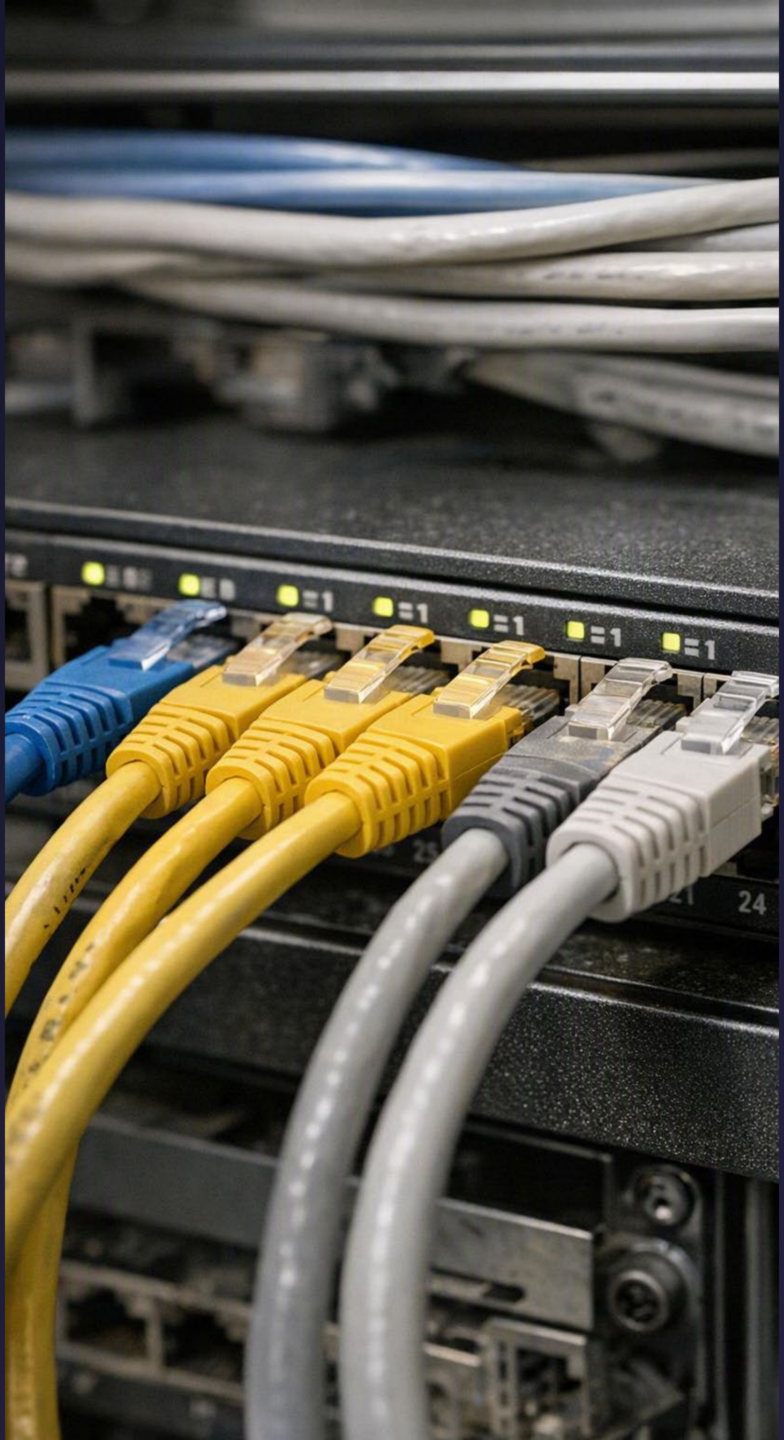

SERVERS · RISK

What happens when servers overheat?

If a server room runs too hot, the damage shows up in two serious ways — both costly and hard to undo.

Lost data

- Excess heat can corrupt storage drives and wipe out large amounts of saved data
- Much of that data may be impossible to recover, losing lots of data
- Backups and records can be damaged at the very moment they are needed most

Downtime

- Servers shut themselves down when they overheat to protect the hardware
- They can be left unable to run for a long stretch while they cool and recover
- Every hour offline means work stalls and normal operations stop

Conclusion for servers

Using one high-usage server can be appropriate for cost-sensitive setups with predictable workloads, but it introduces risks in performance bottlenecks, redundancy, and scalability. For higher reliability and better load balancing, one might consider clustering, distributed load balancing, or multiple medium-usage servers instead of a single high-usage server. This setup is a trade-off between cost, utilization efficiency, and resilience. Strategic monitoring, task allocation, and backup planning are essential to mitigate the disadvantages.

Building resilience beyond one server

The alternatives to a single high-usage server — how they spread load and keep us running when something fails.

Clustering & load balancing

- Clustering links several servers so they share one workload and cover for each other automatically
- A load balancer spreads incoming requests across nodes, so no single machine becomes the bottleneck
- If one node fails the others carry on, so a fault no longer stops the whole operation

Redundancy & failover

- N+1 redundancy keeps a spare server ready to take over the moment one drops out
- Regular backups and replication mean data lives in more than one place, not on a single drive
- Failover can be tested on a schedule, so we know recovery works before we ever need it

SECTION 02

Cooling

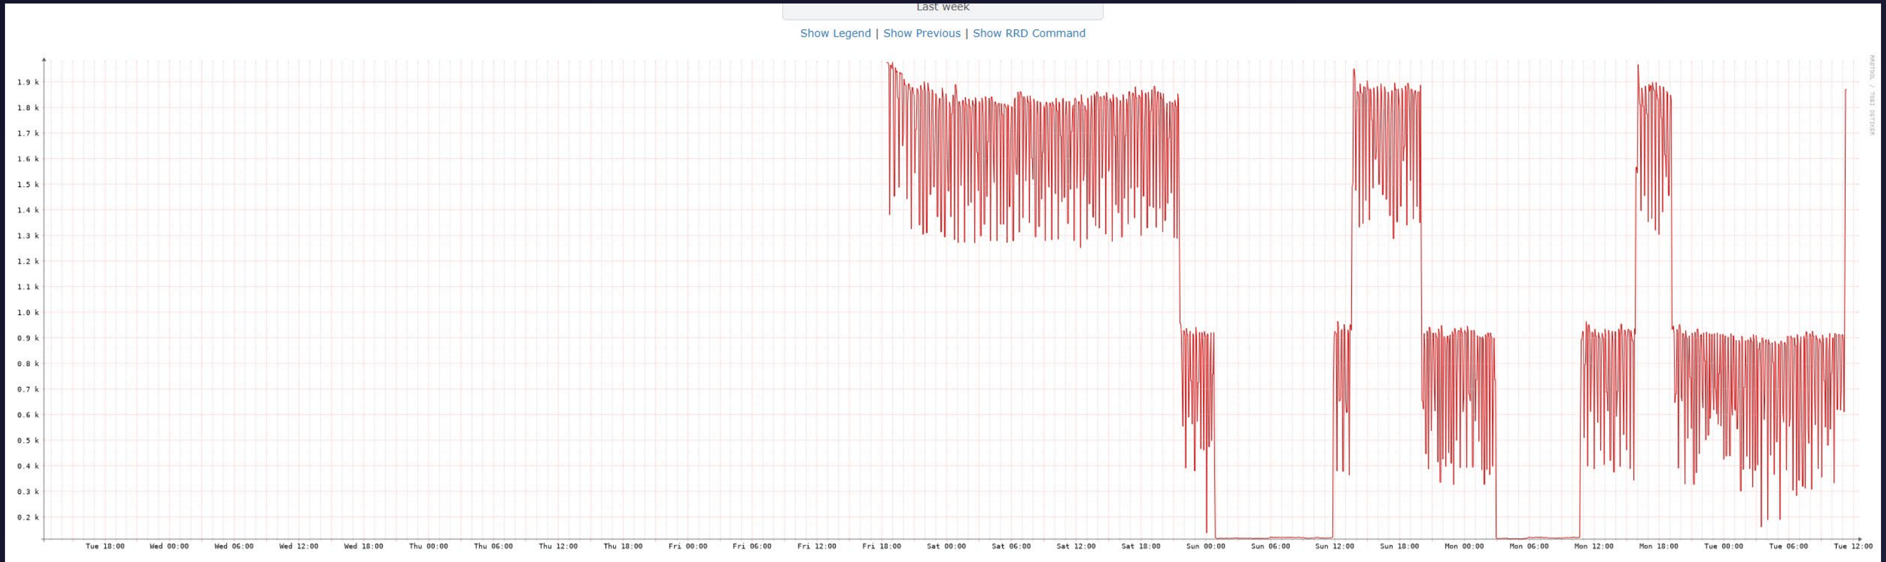

02

Keeping the room at a safe temperature — comparing aircon options and managing airflow and humidity.

COOLING · TEMPERATURE

About the cooling

How warm the server room gets — and how hard the AC has to work — comes down to two things we can measure and manage.

What drives the temperature

- Server workload: the harder the servers run, the more heat they pump out, forcing the AC to work harder to keep up
- Outside conditions: hot weather adds to the load, so the cooling system has to push much more on warm summer days
- Constant balancing: the AC ramps its output up or down automatically to hold a steady, safe room temperature

Keeping servers in the safe zone

- Summer target: the room is held around 21°C to keep the servers cool and running reliably
- Danger threshold: letting a server climb toward 30°C risks overheating, throttling and lasting hardware damage
- Why it matters: staying near the target protects performance and avoids costly emergency shutdowns

COOLING

Aircon

Key questions

- How much energy does it use?
- How many do we need?
- How powerful should it be?
- Does the energy usage compare to all the other energy users?

1. Thermal load and room size

Servers generate continuous and concentrated heat, requiring stable cooling 24/7. Large rooms or high-density racks benefit from systems that can distribute air evenly, such as CRAC units, ducted systems or multiple distributed units.

2. Advantages of a single large unit

- **Energy efficiency at high load:** operating near full capacity continuously for constant heat loads.
- **Simpler installation:** fewer components reduce setup complexity and cost.
- **Noise control:** one large compressor typically generates less noise than multiple smaller compressors.
- **Centralized maintenance:** easier monitoring and servicing of a single system.

COOLING · COMPARISON

One large unit vs multiple small units

Single large unit – disadvantages

- **Single point of failure:** if the unit fails, the entire cooling system is compromised, which is critical for server operations.
- **Air distribution:** large units may struggle to cool irregularly shaped spaces evenly, potentially creating hot spots.
- **Space and cost:** high-capacity units require more physical space and higher initial investment.

Multiple small units – advantages

- **Redundancy and reliability:** partial cooling continues if one unit fails, minimizing the risk of server overheating.
- **Flexibility and zoning:** units can target heat-prone areas, better managing thermal gradients within the room.
- **Efficiency at variable loads:** smaller units can modulate output or shut down individually when not needed, reducing energy waste.

Multiple small units – disadvantages

- **Higher installation and maintenance complexity:** multiple units require more piping, wiring, and service checks.
- **Potential noise increase:** running multiple units could create additive ambient noise.
- **Monitoring complexity:** requires robust controls to maintain consistent temperature and humidity.

Practical considerations for server rooms

- **Redundant cooling is standard:** data centers often implement N+1 or N+2 redundancy, so multiple smaller units are often preferred for reliability.
- **Temperature stability:** maintaining 20–24 °C and humidity between 40–60 % is crucial; multiple units improve even airflow.
- **Future expansion:** modular systems allow adding more units as server count increases, reducing the need for full replacement.
- **Hybrid approach:** for medium-sized server rooms, one high-capacity unit with one or two smaller units is often optimal – the large unit handles the base load while smaller units add redundancy and hot-spot management.w1

COOLING · RECOMMENDATION

Aircon energy usage

Summary recommendation

- **Small server rooms (1–2 racks):** 1–2 wall-mounted or cassette units are sufficient.
- **Medium to large rooms:** multiple smaller units (wall-mounted, ceiling cassettes, or ducted CRACs) are preferable to ensure redundancy and distributed cooling.
- **Mission-critical or high-density data centres:** adopt multiple units with N+1 redundancy for high reliability.

Conclusion

For server rooms, multiple smaller AC units are generally superior to a single large unit because they provide redundancy to minimise downtime, ensure even temperature distribution to prevent hot spots, offer flexible scaling for future expansion, and enable energy efficiency at variable loads through modular operation.

A hybrid approach that combines a primary large unit with auxiliary smaller units can balance energy efficiency and reliability, aligning with best practices in IT cooling management.

COOLING · AIRFLOW

Managing airflow and humidity

Getting cool air to the right place — and holding the right humidity — matters as much as raw cooling power.

Airflow and hot spots

- Hot and cold aisle layout keeps chilled air flowing to the intakes and pushes warm exhaust away
- Blanking panels in empty rack space stop hot and cold air mixing and wasting cooling
- Even distribution prevents hot spots, so no single rack overheats while the room reads normal

Humidity control

- Holding humidity between 40 and 60% keeps the room in the safe operating band
- Air that is too dry invites static discharge that can damage sensitive components
- Air that is too damp risks condensation and corrosion on boards and connectors

SECTION 03

Monitoring

03

Using SNMP to collect real power data from our devices and turn those readings into alerts we act on.

ENERGY · MONITORING

How to Figure Out Energy Usage Using SNMP

- SNMP (Simple Network Management Protocol) is used to monitor and manage network-connected devices
- It connects to smart power distribution units (PDUs), UPS systems, servers, and switches
- It collects real-time power and energy data from each device
- Organisations can track consumption, identify high-energy devices, and improve efficiency

Network

Server

Dashboard

ENERGY · MONITORING

The Process of Using SNMP

- We use LibreNMS, an open-source network monitoring tool, as the software that collects data from our devices over SNMP
- SNMP is switched on for each server, switch and PDU, and secured with a community string so only our monitor can read it
- LibreNMS automatically discovers these devices and polls them at set intervals for power and performance data
- The readings build up over time into dashboards, graphs and alerts the team can act on

LibreNMS monitoring dashboard

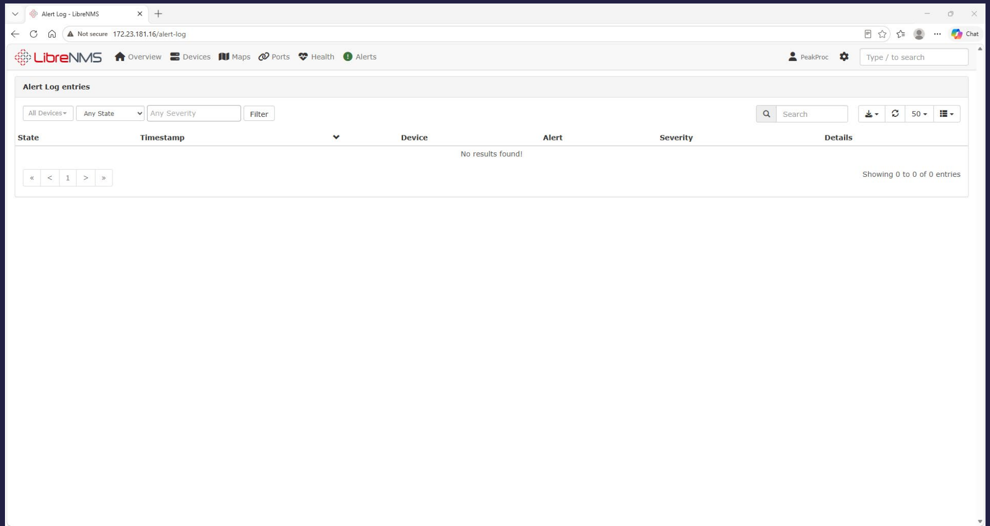

ENERGY · MONITORING

The process of using SNMP

LibreNMS installation on Ubuntu

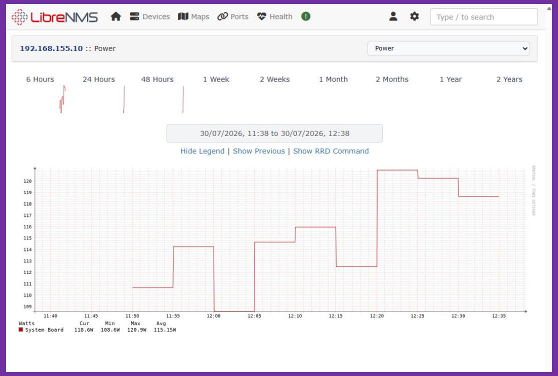

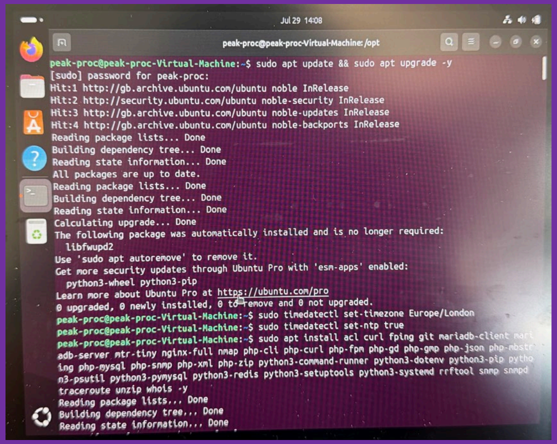

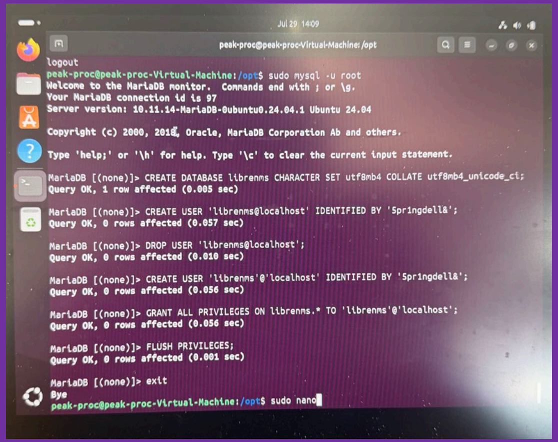

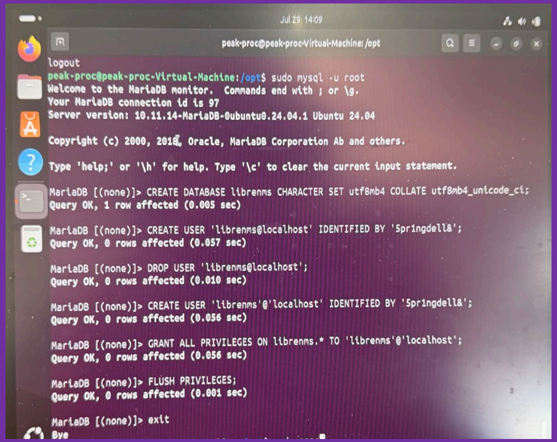

Installing LibreNMS

1. Download Ubuntu 24.04
2. Create a new VM in Hyper-V and install Ubuntu
3. Complete the setup wizard
4. Update the system
5. Install LibreNMS dependencies
6. Create the LibreNMS user and download LibreNMS
7. Install PHP dependencies using Composer
8. Set up the database using MariaDB and MySQL
9. Set the timezone
10. Configure PHP-FPM
11. Configure NGINX and the web server

ENERGY · MONITORING

Inside the SNMP polling cycle

How a single reading travels between the monitoring server and a device

1

The manager asks

The monitoring server sends a GET request to a device's SNMP agent over UDP port 161.

2

The agent looks up

It finds the value in its MIB, a catalogue where every metric has a unique numeric address called an OID.

3

The agent answers

It returns the current reading – such as interface status, CPU load or power draw – in a response packet.

4

Traps raise alarms

For urgent events the agent doesn't wait to be polled – it pushes an unsolicited trap to the manager on port 162.

ENERGY · MONITORING

Turning readings into alerts we act on

Collecting data is only half the job — thresholds and alerts are what let us catch problems before they cost us.

Thresholds that matter

- Set a temperature alert well below 30°C so we react long before hardware is at risk
- Flag any device whose power draw climbs sharply against its normal daily baseline
- Watch PDU and UPS load so we know before a circuit runs close to its limit

Acting on the data

- Route critical alerts to the team by email or message so nothing waits for the next check
- Trend the readings over weeks to separate a real problem from normal daily variation
- Trace each spike back to a specific device, then fix or retire the worst offenders

SECTION 04

Solar

04

What the panels generate, why solar is smart, clean power, and how payback and battery storage stack up.

15/09/2026

ENERGY · SOLAR

How much energy are the solar panels producing on average during the winter?

From the data we have been given from the 9th of November 2025 we can find out that

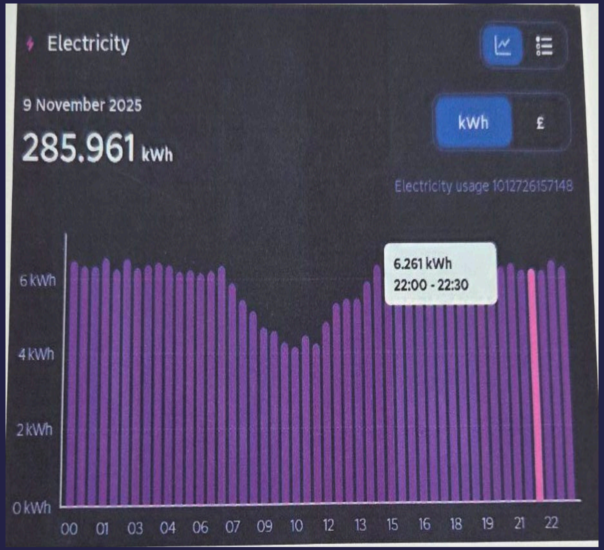

What this graph tells us

- There is a total energy used of 285.961kWh
- Each time interval is 30 minutes / half an hour
- Between roughly 7am and 3pm the solar panels are generating, offsetting grid draw and bringing power usage down to almost 4kWh.
- Overnight, with no solar input, there is a steady 6-6.5kWh used per hour - the servers and cooling running around the clock.

Solar panel output over the summer

- There are 36 solar panels installed across the site, feeding clean electricity back into our supply.
- Their output varies day to day with the weather - bright, clear days generate the most, while cloud, rain and shorter daylight cut production.
- We use the summer average because the longer, sunnier days give the panels their strongest and most representative output.
- Generation peaks around the middle of the day, lining up with when the servers and air conditioning draw the most power.
- To measure production accurately, fit a dedicated generation meter or standalone monitoring device rather than estimating it from how much extra grid power we buy.

ENERGY · SOLAR

Why solar is smart, clean power

Solar panels turn sunlight straight into electricity – giving us useful energy while barely releasing any CO_2.

Strong energy production

- Panels generate free electricity whenever the sun is up, cutting how much power we buy from the grid
- Output peaks in the middle of the day, lining up with when our servers and aircon work hardest
- Once installed they run for 25+ years with very little maintenance, so the energy keeps coming

Very low CO_2

- Generating electricity from sunlight produces no direct carbon emissions at all
- Every kWh from solar is a kWh we don't draw from fossil-fuel power, shrinking our carbon footprint
- It lowers energy bills and emissions at the same time, so it's good for the planet and the budget

ENERGY · SOLAR

Payback and battery storage

How the panels earn their cost back over time — and how batteries could stretch that clean energy further.

Paying for themselves

- Every kWh generated is a kWh we don't buy, so the panels chip away at the bill each day
- Output is strongest in summer daylight, exactly when servers and aircon draw the most
- With 25+ year panel life and low upkeep, savings keep building long after the payback point

Adding battery storage for solar panels

- Batteries capture surplus daytime generation instead of exporting it back to the grid
- Stored energy can run the overnight baseline load when the panels produce nothing
- Shifting solar into the evening cuts grid draw at its steady 6–6.5 kWh overnight rate

SECTION 05

Energy & cost

05

Total usage and cost, the breakdown by area and floor, what the Octopus data shows, and where we can save.

ENERGY · TOTALS

Figuring out total energy used and the price of it

1 GRID

We use the Octopus smart meter to determine the grid energy supplied.

8,692 kWh / month
~ 12 kW average

2 SOLAR

The grid only covers the power we buy, so we add the solar generation on top to capture the total energy we actually use on site.

778 kWh / month
~1 kW average

3 TOTAL

The underlying site demand is the sum of the two.

9,470 kWh / month
~ 13 kW average load

ENERGY · BREAKDOWN

Energy used in different areas (per month)

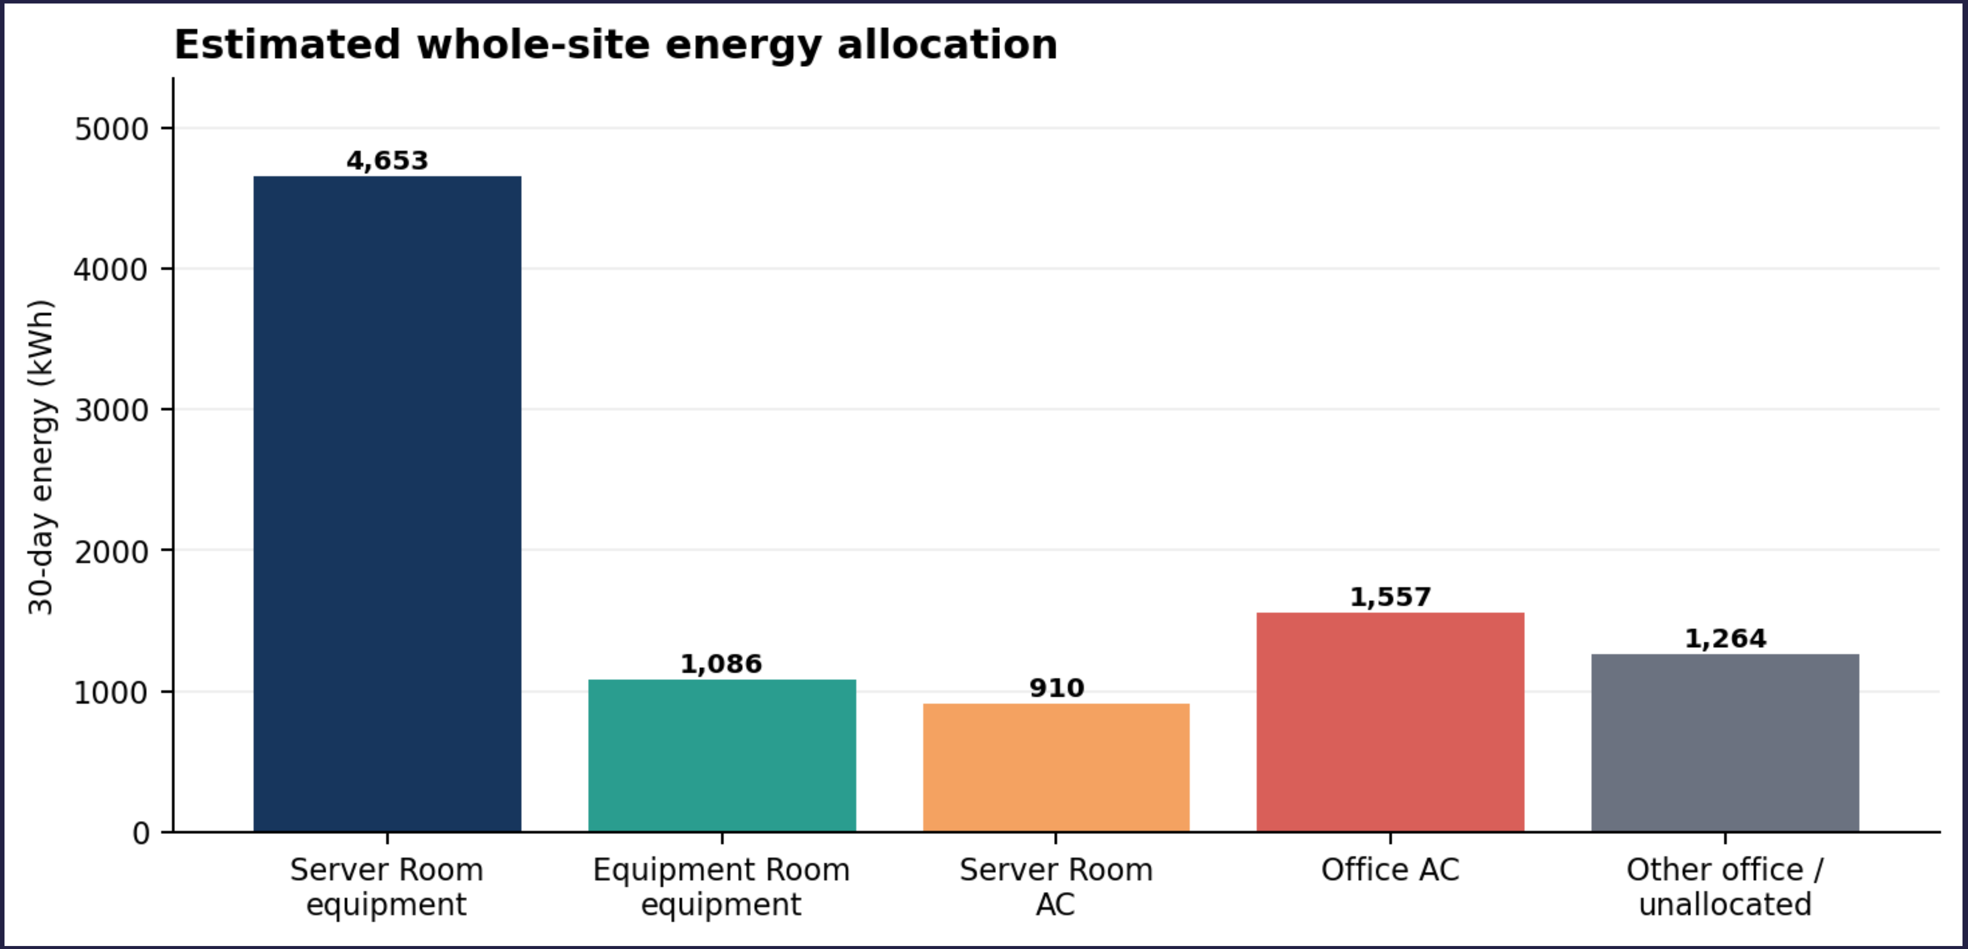

ENERGY · BREAKDOWN

Energy used in different areas

Servers

49.1%

of total energy

Server Room AC

9.6%

of total energy

Backup Servers

11.5%

of total energy

Office AC

16.4%

of total energy

Office Misc

13.3%

of total energy

ENERGY · OCTOPUS

What the Octopus data is telling us

Half-hourly smart-meter readings from Octopus give us an honest picture of exactly when and how much grid power we draw.

What the readings show

- A steady overnight baseline of 6–6.5 kWh per hour as the servers and cooling run around the clock
- A daytime dip between roughly 7am and 3pm as solar generation offsets the grid draw
- Clear summer peaks when hot weather pushes the air conditioning to work its hardest

How we use it

- Total the half-hourly readings across the period to get accurate weekly and annual usage
- Compare the profile against solar output to see how much daytime draw is already offset
- Spot unusual spikes early so we can trace them to a device and act before costs climb

ENERGY · COST

Annual cost of energy by category

Applying the Octopus unit rate to each area's share of grid draw turns kilowatt-hours into a yearly cost we can act on.

Where the money goes

- Servers dominate: running 24/7 at high load, they take the largest slice of the annual electricity bill
- Cooling follows: the AC works hardest in summer and tracks the server load, making it the second biggest cost
- Other electronics: switches, PDUs, lighting and monitoring add a smaller but steady year-round cost

Turning kWh into pounds

- Multiply each area's grid draw by the Octopus unit rate to get its share of the annual cost
- Solar offsets daytime draw, so every sunny kWh directly trims the servers' and aircon's running cost
- Tracking cost by category shows where a small efficiency gain saves the most money

ENERGY · EFFICIENCY

Practical ways to cut our energy use

Small, targeted changes across servers, cooling and monitoring add up to real savings without risking reliability.

Quick wins

- Consolidate light workloads and power down idle servers so we run only the capacity we need
- Hold the room near 21°C and use free cooling on cold days instead of running the AC harder
- Retire or replace the least efficient hardware flagged by the monitoring data

Longer-term moves

- Add cooling redundancy with smaller modular AC units that scale output to the actual heat load
- Lean on solar through daylight hours and explore battery storage to shift it into the evening
- Set weekly usage baselines from the Octopus data and track clear reduction targets over time

SUMMARY · KEY FINDINGS

Key findings at a glance

The four areas we looked at – and what the data and analysis told us

1

Servers

One high-usage server saves money but becomes a single point of failure. Clustering or several medium servers give better resilience.

2

Cooling

Several smaller AC units beat one large unit for redundancy and even cooling. A hybrid setup balances efficiency and reliability.

3

Monitoring

SNMP with LibreNMS gives real-time power data per device, so we can find the biggest energy users and track them over time.

4

Solar

36 panels cut grid draw during daylight with almost no CO_2, peaking when the servers and cooling work hardest.

Recommendations & next steps

What the findings point to, and the practical steps to act on them.

Recommendations

- Favour redundancy: several medium servers and multiple smaller AC units over one large unit.
- Hold the room near 21°C and react well before anything climbs toward 30°C.
- Lean on solar through daylight hours and use monitoring to target the heaviest energy users.

Next steps

- Roll out SNMP monitoring across every server, switch and PDU.
- Baseline weekly usage from the Octopus and solar data, then set clear reduction targets.
- Review cooling redundancy and plan any AC changes before next summer.