

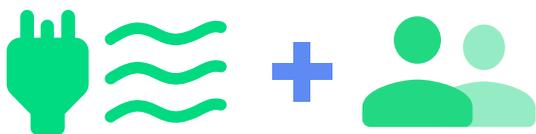
Why an inefficient cloud provider will put your business at risk

There's a reason why organisations try to avoid migrating to a new IT service provider. Whether this involves moving servers between data centres, or migrating to a cloud infrastructure or a SaaS platform, the process can be risky, costly in terms of fees, and distracting – slowing the business down. It's little wonder that in our experience businesses don't tend to undertake a second migration for at least three years.

Avoiding a nasty surprise

No business wants to discover a year down the line that the service it bought into isn't credible or sustainable, resulting in the need to move again. Nor do they want to be faced with costs that escalate way above the headline price the vendor fed them in order to get through the door.

This is why it's essential that every business evaluates the efficiency of a cloud provider or data centre operator before they take the leap.



The two largest costs involved in a data centre or cloud deployment are **power and people**, so these in turn represent a large percentage of the price of your service. For the service price to remain competitive and sustainable, you need to be certain that the provider's power and people costs are low and controlled.

Power: the impact of PUE on your TCO

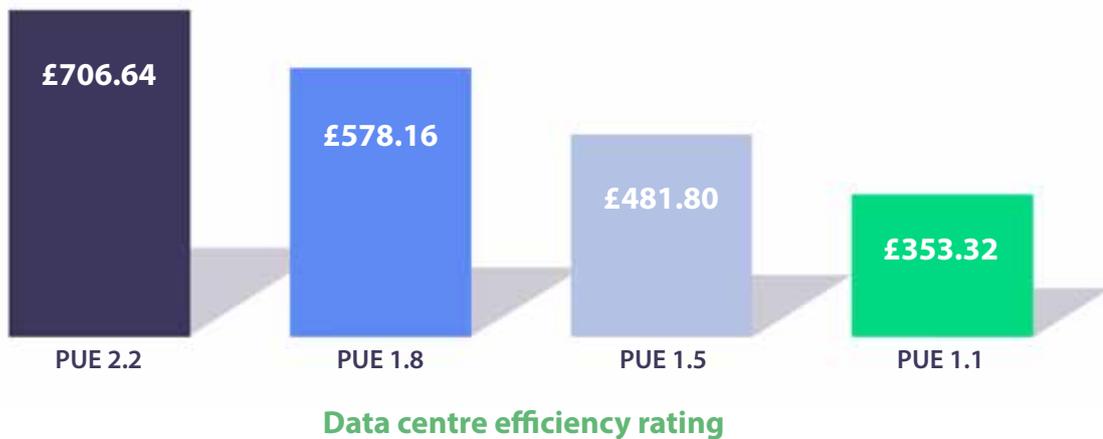
Colocation customers are usually more focused on the power costs of a data centre because this forms an obvious part of their monthly bill. But power efficiency also has a direct impact on the cost of an IaaS, SaaS or PaaS environment. It's great to get a good deal or initial offer, but power prices will fluctuate, and a lack of efficiency will see prices rising in the future as providers pass rising costs on to their customers.

You can get a handle on this by establishing the PUE (Power Usage Effectiveness) ratio of your provider. PUE describes how efficiently a data centre uses energy. Those operating at a higher PUE will be less efficient with power, and more affected by fluctuations in power costs within the marketplace.

PUE is calculated by determining how much power is required to cool and provide resilience for a server, per 1 unit of power used in operation.

For a data centre with a PUE of 2.2, this means that for every 1 unit of power used, 1.2 units are expended in cooling and resilience. This isn't particularly efficient, so it's likely the data centre is more than 10 years old, has a very compromised design, or is not very full and therefore unable to run efficiently. In this situation, the cost of power for a 4kW rack would be just over £700 per month.

Illustration: Monthly power cost of 4kw rack



A data centre with a PUE of 1.8 – which means 0.8 units of power are expended per 1 unit used – is quite typical in the marketplace for a five to 10 year old data centre. The cost to the business in this scenario would be £578pm for a 4kW rack. If efficiency is boosted to 1.5, the cost lowers again to £480pm, until you get to a fully optimised data centre – with features such as indirect cooling, a double conversion UPS system and full generator pre-heating – which will be in the 1.1 to 1.2 PUE region. The cost of power for such a highly efficient site would be £353pm; half that of an inefficient data centre. In addition, it would be much better able to cushion the blow of a fluctuating power market.

People: the benefits of automation – and a digital octopus

In an efficient data centre, repetitive manual tasks are carried out using automation wherever possible. This not only reduces people costs, it also improves service delivery and the availability of expert advice.

Too many data centres still rely on human beings walking around a 100,000 square foot facility and writing down what the power usage is for each rack. Or they operate their access control systems by printing large spreadsheets, and manually programming security passes when clients come to site. These are all activities you ultimately end up paying for as part of the service.

An efficient operator will task their people with service and support roles, leaving the more routine jobs to automation. This releases teams to focus on helping customers – doing the things humans do best, such as providing expert advice rooted in an understanding of the client's business, responding more quickly to client requests, and proactively addressing issues before they can affect the service.

Centralised control portals can dynamically take care of functions such as:

- ✓ measuring and tracking performance across all clouds and the data centre
- ✓ control of the cooling for the data centre
- ✓ managing access for clients
- ✓ automated billing
- ✓ making quick, remote adjustments to virtual machines
- ✓ real-time reporting on power usage.

One particularly effective approach to monitoring and managing infrastructure is the introduction of a centralised control portal that can be used by both the customer and provider. This centralises and simplifies all tasks, requests and communication – a bit like an octopus with tentacles in every part of the operation.

At Amito, we estimate our customer user portal saves us around 30 percent in human resources costs. It also streamlines the client experience, with both an admin and

customer interface into the system that allows access around the clock from anywhere and via any device.

If your cloud provider or data centre operator is inefficient, the cost of your service will be less stable. Support response times will be longer, and there will be less oversight of your service. In short, it will be harder and more expensive to get things done.

You should ask the provider three key questions:

- 1 What's the operating PUE (Power Usage Effectiveness) of your data centre?
- 2 How do you use automation within your business – and what tasks will we be responsible for?
- 3 Where and how can I view and manage my infrastructure?
Do you have one centralised portal, or will we need to switch between several?

Three critical questions

Evaluating the efficiency of a potential new provider is important, but it can also be a worthwhile undertaking with existing vendors. If you can challenge them on the way they run their operation, this may help to drive improvements – and help you to avoid having to make another migration.

Your performance and growth are inextricably linked to the efficiency of your provider. If efficiency is poor, you're exposed not only to unstable pricing models for your existing infrastructure come renewal time, but also to uncompetitive upgrade pricing during the contract period.

Alongside a reduction in running costs and the total cost of ownership of your infrastructure, partnering with an efficient service provider makes it more likely you'll receive an attentive and proactive service, because teams won't be bogged down in admin and mundane tasks. Changes and urgent support requests will be handled more quickly, and experts will have the time to consider new options or research new technologies that might benefit you – resulting in a service that's tailored to your needs rather than generic.

Ultimately, this means you're less likely to have to do another costly, risky and distracting migration in the near future to get what you need.

Author: Paul Tacey Green, Director of Cloud & Connectivity

Further reading: [10 questions you need to be asking your cloud provider >](#)