Attention IT leaders...

Stop overspending in AWS

Cost Optimisation

aws PARTNER Well-Architected Partner Program

zen

Cost Optimisation

In today's fast-paced cloud environment, managing AWS costs effectively is essential for those looking to maximize their budgets without sacrificing performance.

The benefits of cloud's flexibility can help organisations realise long-term cost-savings. It can also make way for cloud-cost-creep which, if left unchecked, can quickly grow out of control.

We know it's easy to get caught up in the day to day, without having time to look up from your desk, never mind to look at the bigger picture.

In fact, taking time to understand some common cost challenges and how they can be overcome can help to save your business a significant amount.

This practical guide to Cost Optimisation in AWS will help point you and your team in the right direction so that you can start getting a better handle on your costs.

If you're looking for an experienced AWS Partner to help you identify and implement cost saving measures, Zen's Cost Optimisation Service is something you can benefit from.

Are you being overly generous?

Having delivered cost optimisation reviews for organisations big and small, we're sharing our key learnings for free, so you're saving already!

We know our stuff

When it comes to operating your resources in a cost-effective manner, our mantra is

Turn it off

Reserve it

Monitor it

The best way to save money in AWS is to make sure you're only running the resources you need. Simple right? Yet this is one of the most serious cost challenges that businesses face.

If you're running instances when you don't need to or you're overspeccing unnecessarily you could be spending significantly more money for no gain.

But how do you ensure that you're operating in the most cost-effective way? That's where our mantra comes in.

Turn it off

If you're running resources that can be turned off, this is perhaps the number one way to save costs in AWS, but it's also the one that most teams forget about.

An example might be if one team only needs access to an instance between 9-5, Monday to Friday, that leaves a lot of time when your resources are going unused – and a lot of time where you could be saving easy money. AWS Instance Scheduler lets you control this process automatically – saving you up to 70% when compared to running it 24 hours a day.

One to watch carefully, we've seen businesses rack up thousands of pounds in costs unknowingly, due to simply forgetting to turn off test environments. So remember those too, and if you're likely to forget, monitor everything!

Reserve it

Effectively, this is all about scheduling and reserving the resources you need. And that starts with appropriate provisioning. It's important to consider the right size to configure your resources – instance size, disk size, performance - essentially having the lowest cost resources that still meet your requirements.

If you know what level of demand you need to meet (and when), scheduling resources or using reserved instances can help to save a good amount of cash. **Reserved instances save up to 72% off on-demand pricing – which could be quite a staggering saving.**

Monitor it

Knowing when to turn resources on or off, when to schedule and knowing how much resource to reserve are all dependent on having a robust monitoring environment. It can be quick and easy to make changes in AWS but how do you know whether your instances are overspecced or you're operating in the most cost-effective way?

The tools we've recommended on the next page can help you keep track.

Are you making use of all the tools at your disposal?

These can help you keep a close eye on usage for little or no cost:

AWS Compute Optimiser – an opt-in tool that tracks resource utilisation across your compute (EC2, Lambda, EBS and Fargate) resources and makes recommendations for downsizing to reduce costs.

AWS Trusted Advisor – a tool that by default tracks various areas of your AWS account(s) periodically to identify areas of improvement. Note that access to Cost Optimisation checks requires Business Support or above.

AWS Cost Explorer – available in all AWS accounts, Cost Explorer is a powerful tool that can be used to visualise, query and report on costs and usage in your AWS Organisation and accounts.

AWS Budgets – a simple and straightforward way to avoid big surprises; receive alerts when costs and utilisation of reservations and savings plans cross a defined threshold.

CloudWatch Alarms – CloudWatch is utilised by AWS Budgets and other services to track usage information, so if you want to explore other metrics that can be monitored and alerted on or want to be a bit more creative with what triggers an alert, start here.

Service Control Policies – to prevent users from mistakenly provisioning unapproved, expensive resources such as large instance types you can configure Service Control Policies that can restrict this. In addition, you can ensure that testing and development accounts are using low-cost resources.

Cost Anomaly Detection – Combining CloudWatch Metrics, Alarms and Anomaly Detection this service that can be found under the Cost Explorer dashboard applies anomaly detection to cost alarms, allowing you to be alerted when anything looks out of the ordinary.

Are you right sized?

We've briefly touched on right sizing, but this important topic really does deserve a section of its own.

As we've mentioned, it helps to think of right sizing as applying the lowest cost resources that still meet the requirements of your workload...Unfortunately, there's not a button for it, it is very much an iterative process. So, you're unlikely to find the right size at day one and never have any need to change (in fact, never changing is inadvisable). But through continual review and monitoring you'll be properly reserving the right sized instances – and making cost savings in the process.

By carrying out a cost benefit analysis – even a basic one – you can ensure that your time and effort is spent on making changes and optimisations where you're most likely to get a good return. There's little point spending hours or days optimising something that might return a few dollars savings a month – you'd probably be better off putting that time to better use elsewhere. Our advice is to look at the top three or so AWS services where you're spending the most money and focus on those.

Perhaps the best time to carry out a right sizing review is when there's a notable change in usage patterns (which you'll be able to identify through CloudWatch), when a new AWS service is released or when there's a price change.

For Example:

We've found that new families of EC2 instances are generally cheaper than the older generations – an obvious trigger to conduct a review and determine whether you could be getting the same or better performance for less money.

And remember to keep an eye on all the components of your resource. That means not just the CPU and memory but also network and disk usage and beyond.

In our experience, this knowledge of every facet is vitally important. We've previously helped customers who have spent vast amounts on increasing instance sizes with little benefit when the fix lay in a simple change elsewhere.

Are you under committed?

You're trying to operate cost effectively and you know the importance of right sizing, but the success or otherwise of your cost saving efforts will ultimately be determined by the commitments you make.

On-demand instances – For short term, unpredictable workloads where interruption and unavailability aren't an option, on-demand instances are a great choice for a few weeks or maybe months. When you're able to commit for much longer, on-demand instances are not recommended. There's no long-term commitment, but there's also no discount.

Reserved instances – Giving you all the hardware you'd get with an on-demand instance, a reserved instance requires a one- or three-year commitment from you and you'll save up to 72% when compared to the on-demand equivalent, by simply committing. With a scheduled reserved instance – if time of day or day of the week is not important – you can make greater savings still by reserving the instance for particular times (overnight or weekends, say).

Spot instances – In our experience, spot instances are a great way to save money – up to 90% against on-demand costs. And because AWS have made it so easy and seamless to bring spot instances into your workload, they're a really viable option. Spot instances are ideal for testing or batch processing – really, anything bursty that you don't mind being slightly susceptible to interruption. Just remember that AWS can terminate your spot instance at any time; you'll get a 2 minute termination warning!

Managed services – This is another option that many users don't think about. Managed services through AWS partners like Zen or even AWS themselves move the responsibility for patching, backups and general maintenance into the hands others, freeing up your team to spend more time on business-critical issues. Wherever it makes sense for your business, we recommend making use of managed services.

Geography – Another easily overlooked decision is where your workloads are hosted. Some workloads need to be hosted in a particular region, but others don't. If it doesn't matter where in the world your workload is, take a look at the prices of other regions. It could be an easy way to make some real savings.

Are you matching supply with demand?

You can ease your decision making by using the right tools to match demand with supply.

The approach you take will be determined by the type of workload and the consistency or seasonality of demand.

Demand-based

Typically, demand-based scaling fits best in meeting the changeable demands of web traffic. As traffic enters your load balancer, auto scaling can ensure that you're always running at the right capacity. That auto scaling by the way is determined by your own predetermined parameters. We recommend using CloudWatch to help ensure your requirements are accurately mapped.

Buffer-based

When there's no urgency over a given workload, we would recommend you use buffer-based scaling instead. If you can afford to wait before you process that workload, a service like AWS Batch can scale out your resources and processes at a time that works best for you.

Time-based

With time-based scaling, you can ensure that your instances scale up to meet a known trend (seasonal or weekly demand for example), then ramp down for lower costs at all other times.

We're here to help

There's so much more to say about AWS cost optimisation. Although some of AWS' services automatically find cost-saving opportunities, it still takes an experienced human to identify the less obvious areas where some of the biggest cost savings can be found.

Zen's Cost Optimisation Service was created for this very purpose, allowing our customers to benefit from our extensive experience reducing AWS costs for customers of all shapes and sizes. We'll thoroughly review your AWS environment, providing advice on what improvements you can make and how you can make them.

For more details on AWS cost optimisation from Zen or to book your free cost optimisation analysis, get in touch:

Visit: zen.co.uk/aws
Email:BusinessSales@zen.co.uk
Call our team: 01706 902579

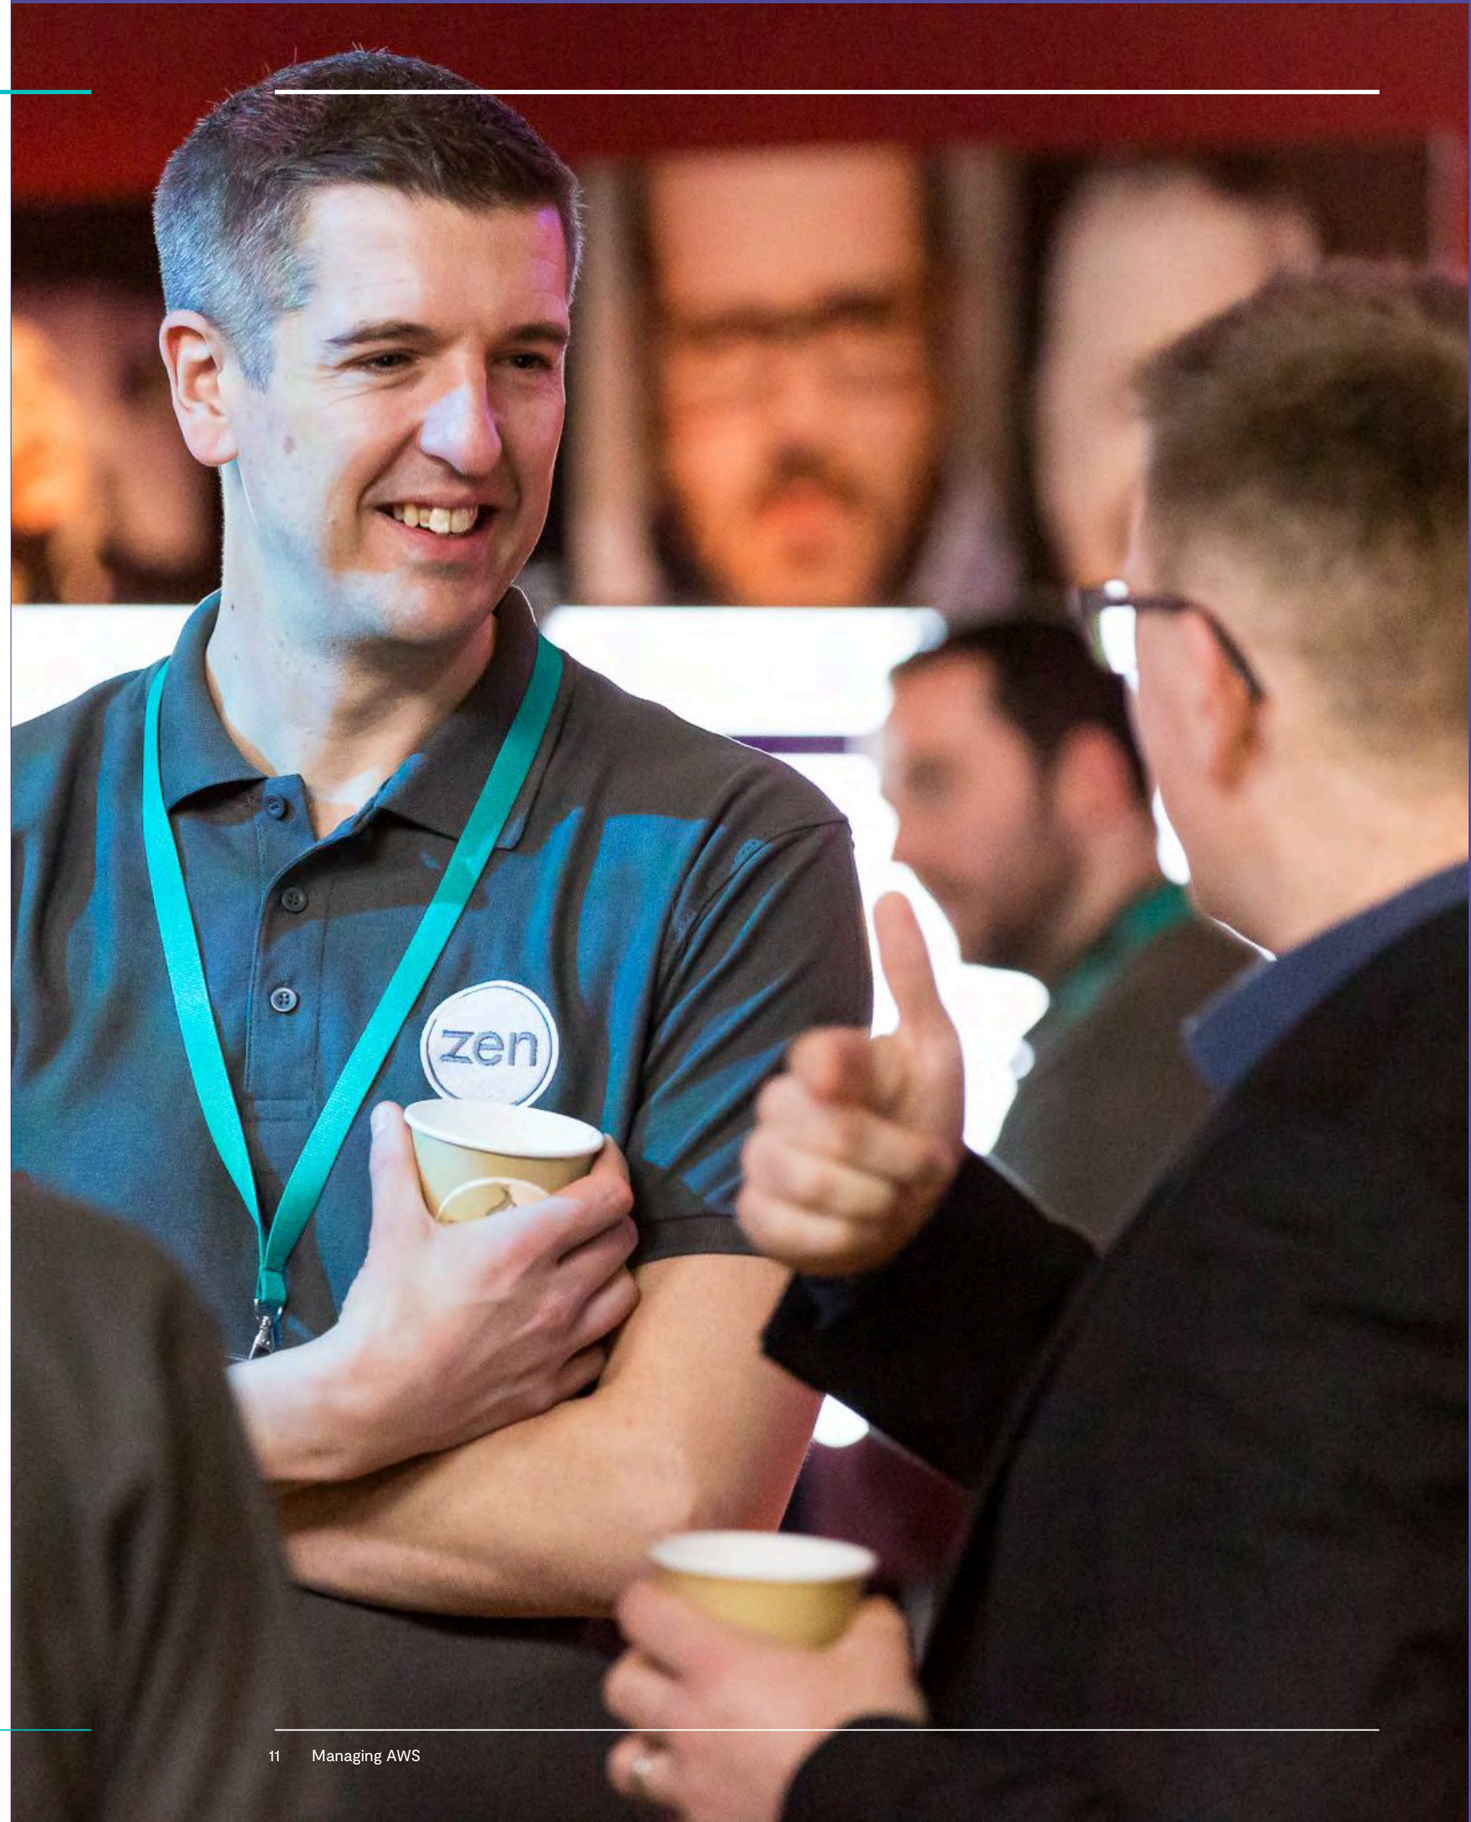

Keep on top of the cloud

Not ready to give us a call just yet?
Find out more about AWS at:

www.zen.co.uk/aws

aws PARTNER Well-Architected Partner Program

zen