

AWS vs **AZURE**

The similarities and the differences between
the two most requested cloud providers.



Offis
MULTI-CLOUD SERVICES

Introduction

Offis is proudly vendor agnostic, working with all of the major cloud providers and technologies.

That means any recommendations we make are not influenced by preference deals or restricted options.

Amazon Web Services (AWS) and Microsoft Azure are still the two most requested cloud providers. AWS remains the gorilla of the industry, chalking up around US\$6 billion in revenue last year*.

AWS may be ten times larger than Azure, but adoption of Microsoft's cloud is accelerating with more than 10,000 new customers each week.

So, in the knockdown, drag-em out title fight for your business, who deserves to win? Does the sheer size of the market leader make this an automatic KO? Or could the fast-moving energy of the challenger put the heavyweight on the mat?

In this ebook, we outline the similarities and the differences between the two 'Infrastructure as a Service' (IaaS) platforms.

But first, a disclaimer. This ebook does not constitute direct advice. Every use case is different, and there are many other variables and considerations to take into account before deciding which cloud service

might be right for you.

This isn't a performance review, either. Instead, this is a feature comparison, with some extra commentary and observations originally presented at Offis' eClub event in July, 2015. [offis.com/community/eclub/events/eclub-23rd-july2015/] The goal is to explore the various choices available to you and any business keen to achieve more in the cloud.

Chy Chuawiwat - CEO



* - <http://www.forbes.com/sites/louiscolumbus/2015/04/15/sizing-microsoft-azure-and-amazon-aws-revenue>

AWS v Azure Global Regions



ROUND 1 *Compute* Amazon Elastic Compute Cloud (EC2)

Amazon compute calls their environment EC2, or Elastic Compute Cloud. There are many types of instances, each offering different levels of compute, memory, storage, network, and so on.

Choosing the right instance type comes down to your workload; some applications require more memory, some applications require more CPU.

The T2 and M3 instance types are intended to be general purpose, typically used for burstable compute requirements, and are ideal if your application experiences fluctuating demand over time. By contrast, the other instance types in the M range are commonly used for a more consistent running environment. If you're confident that the demands on the server will remain pretty steady, the M range may be enough for you. We recommend the newer M4 instance types which just been released as they tend to be faster and more cost effective.

With 36 CPUs and 60GBs, the C range of instances is designed for more compute intensive tasks. Maybe you're running high-end web servers, or you need to run complex application data analysis.

If C is for Compute, then G is for Graphics. The G range offers up to 4 GPUs (Graphics Processing Unit) with 32GB RAM. If you're running games, encoding complex graphics or rendering animations, you'll appreciate the increased RAM and GPU oomph.

No guesses what R stands for, as these instances pack 224GB of RAM.


The I range is designed to handle I/O intensive apps that need to read/write a large amount of data, with up to 32 vCPUs.

D2 brings density storage to the EC2 family and has only recently been made available. These are intended for Massively Parallel Processing (MPP) scenarios, such as data warehousing, MapReduce and Hadoop distributed computing, and log or data-processing applications.

AWS is always releasing new features, so something else may appear on the market by the time you read this.

“It comes down to your workload. Some applications require more memory, some applications require more CPU.”





At a Glance

- **T range:** General purpose, burstable performance instances
- **M range:** General purpose, fixed performance instances
- **C range:** Up to 36 vCPUs/60GB RAM for faster processing
- **G range:** Up to 4 GPUs/32 GB RAM for graphics intensive applications
- **R range:** 32 CPUs/224GB RAM designed for memory intensive apps
- **I range:** Up to 32 vCPU/224 GB/6.4TB SSD for IO intensive apps
- **D2 range:** Up to 36 CPUs/244 GB RAM/48 TB of local storage

NB: Instance types will vary when selecting the operating system & databases.

Latest updates and further information, please refer to:

<http://aws.amazon.com/ec2/instance-types>



ROUND 1 *Compute* Microsoft Azure

Like AWS, Microsoft is equally boring in its naming conventions. That's why some developers prefer to call the Azure G series 'Godzilla', because it's the biggest instance size they have. Unfortunately, the high performance Godzilla series, with 32 CPUs and 448GB of RAM, is only available in two regions—West US and East US—and is currently unavailable in Australia.

By contrast, the A series instances are much smaller, intended as quick and disposable, general-purpose machines. They're like little mayflies, popping up as you need them and disappearing when you don't, even if that's only a few hours. Ranked from A0 to A7, each instance has up to 8 CPUs, 16GB of RAM, and 605GB storage.

The A8 and A9 servers are optimised for faster networking between servers, with 40Gbit/s InfiniBand network and remote direct memory access (RDMA) technology. These are perfect for Message Passing Interface (MPI) applications,

high-performance clusters, modelling and simulations, video encoding, and other compute or network intensive scenarios.

The A10 and A11 servers leave out the Infiniband and other network optimisation features. Otherwise, they're virtually the same as the A8 and A9 flavours, with 16 CPUs and 112GB of RAM, and are intended as compute intensive machines. Just not networked.


The D Series is compute optimised. The D1-D14 servers have been re-architected with better CPUs, up to 16 cores, 112 GB RAM and 800 GB of faster, local SSD storage. Similar to the AWS M3 and M4 servers, this new architecture allows you to get up to 60% better performance for roughly the same price. That's the claim, anyway.

Storage optimised DS servers should be available from September 2015. If your needs expand beyond your server's primary C drive, this allows you to attach additional terabytes of storage to create a secondary drive.

Naturally, different regions have different price points. Australian prices are currently about 5-40% higher than the US, depending on instance type.

“It comes down to your workload. Some applications require more memory, some require more CPU.”





At a Glance

- **A0-A7:** Up to 8 CPUs/56GB RAM/605GB storage for general computing.
- **A8-A9:** 16 CPUs/112GB RAM plus 40GB/s for infiniband network.
- **A10-A11:** 16 CPUs/112GB RAM for compute intensive machines.
- **D1-D14:** Up to 16 CPUs/112GB RAM/800GB SSD storage for 60% faster performance.
- **G series:** Up to 32 CPUs/448GB RAM* Intended for big data projects.

*Currently unavailable in Australia.

For the latest updates and further information please refer to:

- <http://azure.microsoft.com/en-us/pricing/details/virtual-machines/>.



AWS has four main types of storage. The Simple Storage Service is nicknamed S3, and is a secure, durable, highly scalable, object storage service. Scattered across multiple availability zones, you can store and retrieve data from anywhere on the web with nine 9s (99.9999999%) durability and 99.99% availability. There are a number of use cases, including static website hosting, cloud applications, content distribution, backup and archiving, disaster recovery, and big data analytics.

Elastic Block Storage (EBS) is mountable storage for your EC2 machines. Each virtual machine (VM) already has a built in storage drive, but EBS allows you to mount additional drives for more capacity and speed. Of course, the more storage you provision in EC2, the higher your IOPS (rate of Input/Output Operations Per Second) will become. Provisioned IOPS (SSD) volumes support up to 20,000 IOPS and 320 MB/s of throughput. Host-based encrypted (AES 256) is also available for EBS.

Not to be confused with EBS, Elastic File Storage (EFS) is fairly new. It works as a scalable, shared file server for any or all of your instances in EC2, so you can share the same files across all of your instances without creating numerous local copies. This makes EFS a very cost effective file sharing service.

As the name implies, Glacier is veeeery slow storage. While this isn't what you want for real-time performance or back-ups, it's fine for archive data that needs to be kept, but rarely retrieved.

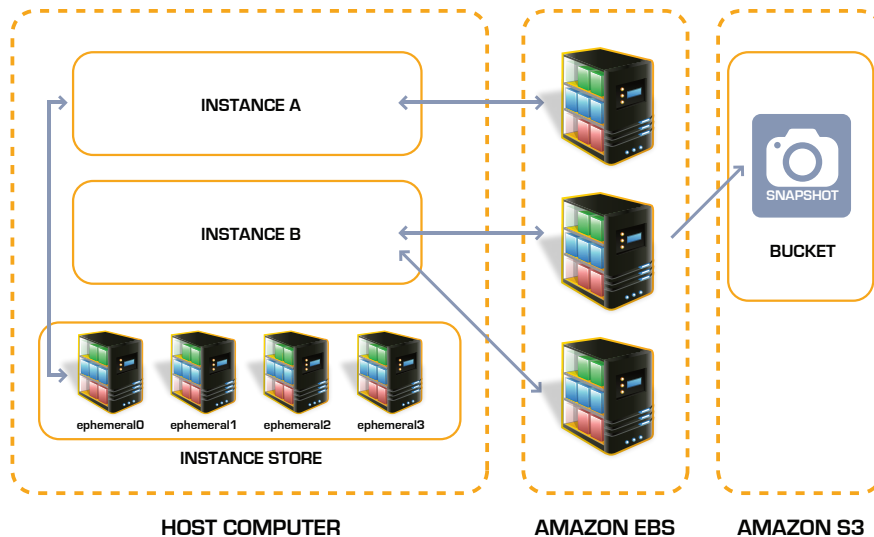
You can also select Ephemeral Storage, which is accessed via the AWS portal alongside your VM. However, as the name suggests, if you shut down the machine, the storage will disappear and cannot be retrieved. This is because ephemeral storage sits on the VM's host server, so if you terminate the machine for whatever reason and later decide to restart it, the VM will migrate to another host with its own storage.

Ephemeral storage is intended for storing temporary or frequently changing information, such as buffering, caching and paging files.

Finally, AWS Import/Export transfers your data directly from the AWS portable storage device to and from your storage devices, using Amazon's high-speed internal network and bypassing the Internet. However, this is currently NOT available in Australia.

“I wouldn't recommend using ephemeral storage for live data because you'll cry when it disappears!”





At a Glance

- **S3 (Simple Storage Service)**: Secure, durable, highly scalable, object storage service with 99.9999999% durability and 99.99% availability.
- **Elastic Block Storage (EBS)**: Block-level storage volumes you can attach to a running instance, for data that requires frequent and granular updates.
- **Elastic File System (EFS)**: Elastic shared file storage service for Amazon EC2 that behaves as a common data source.
- **Amazon Glacier**: Low-cost storage service that is very slow to access. Cost-effective for data archiving and replacing redundant magnetic tape systems.

- **For the latest updates and further information:**

AWS Storage: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Storage.html>

AWS S3: <http://aws.amazon.com/s3/details/>

EBS: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AmazonEBS.html>

EFS: <http://aws.amazon.com/efs/>



ROUND 2 *Storage* Microsoft Azure

Every Azure subscription can create up to 100 storage accounts, each of which is a unique namespace for your storage resources, so you don't have to run any VMs to access your data. You can use whatever name structure makes sense to keep everything organised—my stuff, your stuff, blue stuff, red stuff.

There are 2 types for Storage accounts Standard and Premium. A Standard storage account includes Blob, Table, Queue, and File storage. You can also choose how your data will be replicated when setting up the account.

Locally Redundant Storage (LRS) is the discount option that maintains three copies of your data to protect it from hardware failures. But as these copies are stored within the same facility, your data is still at risk if the data centre experiences a failure.

More durable than LRS, **Zone Redundant Storage (ZRS)** also maintains three copies of your data, but this time it is replicated across multiple facilities

and regions. ZRS is currently available only for storage Blobs.

However, for maximum durability and resilience, we recommend **Geo-Redundant Storage (GRS)**. GRS replicates your data three times in the primary region and another three times in a geographically distant secondary region. So, if a nuclear bomb suddenly hit Sydney, the system would automatically failover to Melbourne to keep everything running. Awesome.

With regular GRS, the data in the secondary region is only accessible in failover mode. So, if you want to access and read data from Melbourne as well as Sydney, then you'll need to select the Read-Access flavour of GRS (RA-GRS).

For maximum durability and resilience, we recommend (GRS) Geo-Redundant Storage which is enabled by default.

You can change how your data is replicated after your storage account is set up. However, if you choose GRS when creating your account, then you cannot switch to another replication type later. Instead, you would need to create the new storage type separately and migrate the data across. Transfer fees would apply.

Premium Storage is currently in preview mode, giving VMs up to 32TB of SSD-backed storage with 50k IOPS and 1ms latency. If you attach it to a VM, it's blindingly fast.


Data import from your on-premises storage to Azure is NOT available in Australia.

Fault Domains define the group of VMs that share a common power source and network.

So VMs in a given **Availability Set** would be placed in different Fault Domains.



At a Glance



- **Blob storage (default):** Used for any unstructured txt or binary data, such as a document, media file, or application installer.
- **Table storage:** Handles structured datasets within a NoSQL key-tribute data store. Good for rapid development and fast access to large quantities of data.
- **Queue storage:** Reliable messaging for workflow processing and for communication between components of cloud services.
- **File storage (Preview):** Shared storage for legacy applications using the standard SMB 2.1 protocol. Allows a developer to move legacy apps easily into Azure.
- **Locally Redundant Storage (LRS):** Maintains three copies of your data, replicated within a single facility in a single region.
- **Zone-redundant storage (ZRS):** Maintains 3 copies of your data, replicated across 2 to 3 facilities, across 1 or 2 regions.
- **Geo-Redundant Storage (GRS):** Maintains 6 copies of your data, across 2 geographically distant regions, with automatic failover.
- **Read Access Geo-Redundant Storage (RA-GRS):** The same as GRS, but with read access enabled for both primary and secondary regions, for greater flexibility.



Each account in AWS allows you to create a whole bunch of internal IP addresses, sub-nets and security groups, binding all of your machines together into one or more Virtual Private Clouds (VPC).

You can connect to your VPC with a site-to-site VPN to ensure your data is encrypted.

The Elastic Load Balancer, does what it says on the tin. It's a load balancer and it's elastic. You can run it as a public load balancer or as an internal load balancer. Either way, as traffic comes in, the load balancer directs it to one of your availability zones within AWS according to your settings. You might direct traffic to the closest zone or, if the zone is experiencing heavy demand, to another zone capable of providing better performance.

Route 53 is the DNS. If you're technically minded, you may recognise that 53 refers to the port used for DNS queries. You can run it as a public DNS, redirecting traffic from a URL to a different IP address.

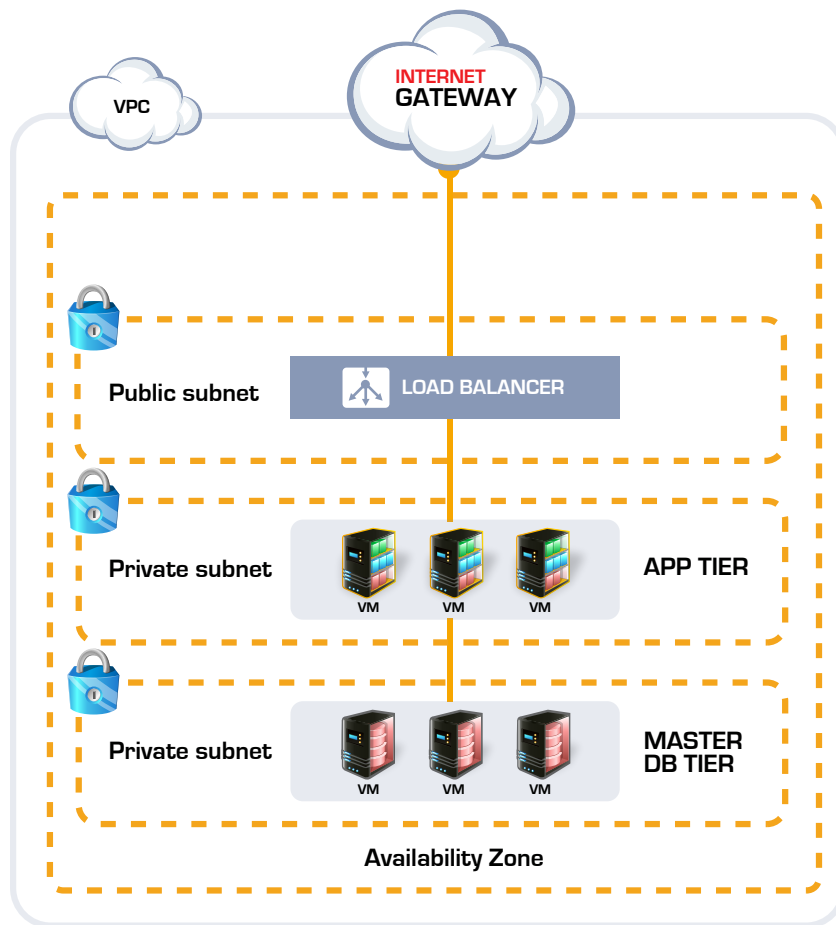
Route 53 can also be used as a private DNS, redirecting to other internal IP addresses within the same network, or to set up your own load balance or failover configuration across multiple regions.

Direct Connect is really just a layer two connection from your data centre or corporate office into AWS, allowing you to transfer gigabits of data faster and more easily. Direct Connect can be an efficient way to carry out a fast offsite backup of your corporate environment.

If you're looking for a data centre in Australia, both Global Switch and Equinix support Direct Connect. Meanwhile, Direct Connect network partners include AAPT, Optus, PacNet and Verizon.

“Every VPC has the ability to create a site-to-site VPN, giving you a virtual private link through the Internet to Amazon. Your data is encrypted so people can't sniff whatever you need to do.”





At a Glance

- **Virtual Private Cloud (VPC):** linking your machines together into a single network.
- **Virtual Private Network (VPN):** an encrypted link from one private network to another over a public network.
- **Direct Connect:** A two layer connection between your data centre and AWS, fast data transfer of large data, no bandwidth costs.
- **Elastic Load Balancer:** Directs web traffic to different machines or zones according to preset criteria.
- **Route 53:** DNS manager, suitable for both internal and public DNS tasks.

For the latest updates and further information, please refer to

ELB: <http://aws.amazon.com/elasticloadbalancing/>

Route 53: <http://aws.amazon.com/route53/>



Microsoft's Virtual Network came out in 2013 and is very similar to Amazon's Virtual Private Cloud concept.

Until recently, large enterprises struggled with a whole range of security issues in Azure. Thankfully, Microsoft has invested heavily in Azure's security and resolved most of these, including all of the most serious or most requested fixes. One of the most popular new developments is the new multi NIC feature, allowing up to sixteen network interface cards (NICs) per Virtual Machine.

Azure has two VPN services. Point-to-site VPN is a secure connection from your computer to the Virtual Network. You can connect to an extensive list of devices via the VPN, thanks to multiple protocols.

Meanwhile, a site-to-site VPN connects your virtual machines so that they communicate with each other. Up until very recently, the site-to-site VPN could only give you a 100MB per second connection, even with a gigabit router on your end. This has improved and is now capable of up to 200MB per second.

There's also ExpressRoute, offering 1GB per second connections. ExpressRoute is Azure's version of AWS's Direct Connect. However, more Aussie providers currently support ExpressRoute; including networking partners Telstra, Verizon and Optus, as well as data centres Equinix and NextDC. It's more expensive than the regular VPNs, but could be worth it if you have a lot of data to transfer. Security is taken care of with an encrypted, private link dedicated to your account, so there's no sharing.

Traffic Manager comprises of the DNS and load balancer within your private network. You can set up load balancing between multiple virtual machines or virtual networks, to create geographical redundancy. If you have services in multiple regions, Traffic Manager can also determine how each request will be directed, based on the location of the user or customer, in a similar way to AWS's Route 53.

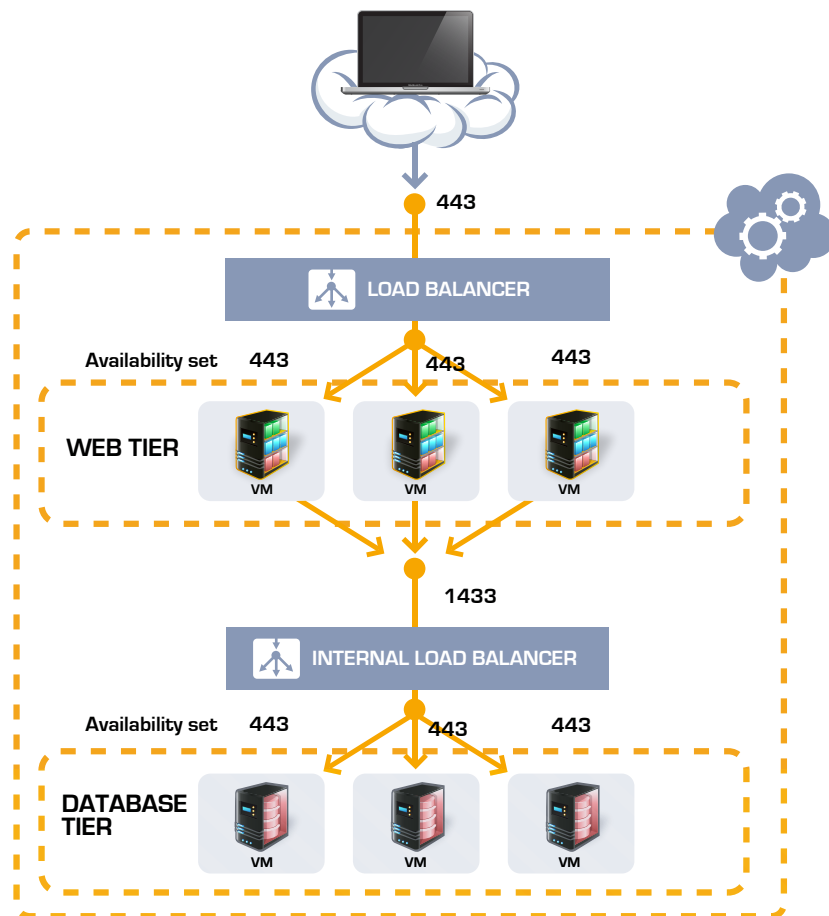
However, the DNS component of Traffic Manager is very basic. In fact, both Azure and AWS only provide the most fundamental DNS features. But, Azure does allow you to bring in your own DNS if you do want anything more comprehensive. There are a few other DNS providers offering a variety of high-level features, so we would recommend having a look at what is available.

There are always pricing differences between the US and Australia, and local prices for Azure services are expected to rise by 26% in late 2015. Plus there are other regional service variations that can impact your bill from a cloud provider. For example, a SoftLayer customer used to pay \$35 per 5TB of outbound bandwidth. The rate is the same at \$35, but the amount of traffic allowed has been reduced to only 250GB, blowing the charge for 5TB of outbound bandwidth out to \$6151.

*“Microsoft has invested heavily in Azure's security.
In fact, they're clearly intent on world domination.”*



At a Glance



- **Availability Set:** If two VMs are part of a Availability Set **(AS)**, it is guaranteed that one of the VMs will always be available even if there is a hardware, network or power failure or hosts are getting updated through planned maintenance.
- **Virtual Network:** Allows you to link your machines together into your own network.
- **Point-to-site, Site-to-site:** Secure client and network VPNs with up to 200GB per second data transfer speeds.
- **ExpressRoute:** 1GB per second data transfer via a layer two connection between your data centre and Azure.
- **Traffic Manager (DNS):** Very basic. However, you can use your preferred DNS if necessary.
- **Traffic Manager (Load Balancer):** Directs web traffic across multiple machines, data centres or geographic regions.

A full list of possible VPN connections is available at:

<https://azure.microsoft.com/en-us/documentation/articles/vpn-gateway-about-vpn-devices/>



Post Fight Analysis



The final bell rings. The fighters retreat to their corners. The blood is wiped from the canvas. So who won?

Of course, it depends. In the absence of a clear knockout, the judges can't agree on points. The final tally is up to you.

And that's exactly as it should be. There can be no clear winner, except the one that's right for your unique circumstances.

Are you ready to enter the ring?

Online business is one long tournament. In every industry and every niche, today's digital champion constantly defends its title from a hundred eager startups, full of new ideas and fast-moving innovations. And then there are the older, heavyweight brands on the comeback trail, ditching the old moves for new strategies, always looking for another shot at recapturing the market.

Maybe your business is the heavyweight. Or maybe you're the bantam weight startup. You might even be the reigning champion in your field. Whether you're competing to seize the advantage from your opponents, or defending your title from all challengers, you need to develop digital muscle in all the right places. The speed, stamina and power of your digital business depend on choosing the right infrastructure and the right maintenance regime to keep your websites and applications fighting fit.

If you're still not sure whether AWS or Azure should be your cloud champion, talk to us. But whichever you choose, Offis Multi-Cloud Services will be in your corner with a sponge and a towel.

"Ding, Ding"

