datAvail
DATABASE SERVICES

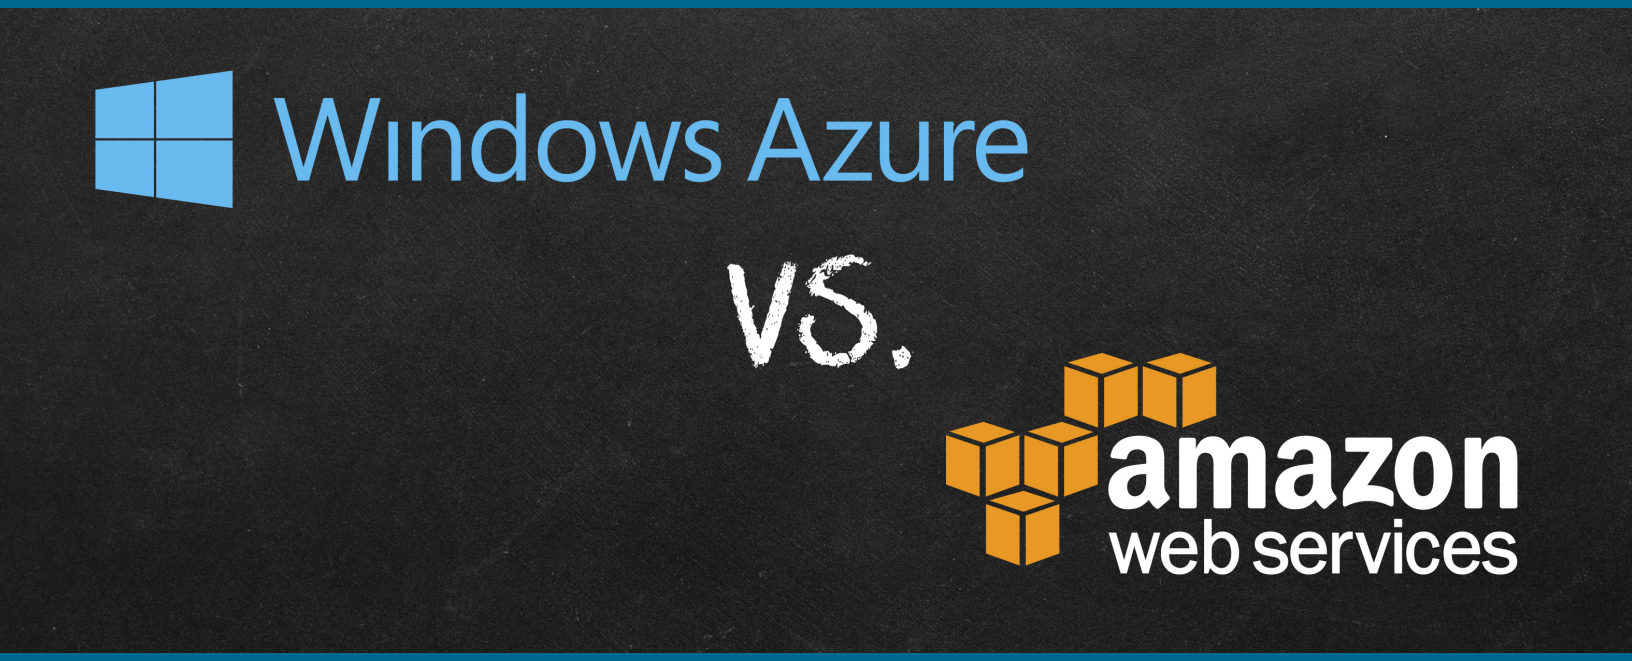

ABSTRACT:

By Dilip Kumar

Microsoft SQL Azure and Amazon RDS are marketed in remarkably similar ways. Both companies claim that their cloud database product makes it easy to migrate from on premise servers to their database-as-a-service cloud offerings; simply migrate your schema and data to the cloud, and then change a connection string in your application. They also emphasize that the same management tools used with an on premise database can be used in the cloud. The emphasis on these similarities seems to imply that the main competition facing the two companies is on-premise offerings, rather than each other. An examination of both products shows that it is very difficult to directly compare them because while the two products are marketed in a similar manner, the two vendors have taken an entirely different approach to architecture. The key difference arises from the fact that Amazon dedicates hardware resources to the user, while Microsoft shares resources among users. Perhaps the most significant result of this difference is the consequent disparity in pricing. However there are also other differences that may influence the decision of a consumer, and these are examined below.

There are many ways to run databases in the public cloud. Competition for your business is fierce, prices are steadily dropping, and the service offerings of the major providers change almost continuously. This whitepaper focuses on two popular offerings: Microsoft's Azure MySQL Databases and Amazon's Relational Database Service (RDS) running MySQL Server. We compare them on eight key features and suggests the database configurations that each system is better at handling.

datAvail

DATABASE SERVICES

Microsoft SQL Azure vs Amazon RDS

content

Cloud-based DBaaS

The advent of cloud-based computing promises several benefits for organizations and information technology professionals. The cloud frees organizations from needing costly fixed hardware assets, and offers/provides resources to run and maintain them. This allows them to reduce overhead and scale computing infrastructure responsively by taking advantage of on-demand resource pricing.

Cloud-based Database-as-a-Service (DBaaS aka PaaS) offerings extend this promise by providing virtual operating systems and databases that organizations can rent on an as-needed basis. As for pricing or price comparisons between the two services, there is no particular set "menu" available. Each situation is different, but to make it simpler for customers to estimate costs, both companies provide service calculators.

Baron Schwartz, an expert in MySQL and the author of High Performance MySQL, explained on ReadWriteWeb:

> Although company-run databases won't ever completely go away, the rise of DBaaS can free many organizations of all sizes from the drudgery of maintaining their own local database systems. The primary benefits are that it can streamline operations and reduce cost—freeing the business to focus on developing their core products or services.

Elsewhere, he adds:

> Overall, the service makes it easier to solve problems, correct mistakes, and transfer data from one system to the next. They can scale as large as necessary, fit the needs of the customers, and offer better availability and security than most in-house operations.

Interest in DBaaS cloud solutions continues growing for precisely these reasons, but which external provider has the digital data warehousing solution right for your organization?

There are many new services from which to choose, including two popular cloud-based offerings: Microsoft Azure SQL Databases and Amazon Relational Database Service (EC2.) With Amazon RDS, database administrators can run six different platforms: Amazon Aurora, Oracle, Microsoft SQL Server, PostgreSQL, MySQL and MariaDB. Because these are very similar services from established companies with the same goal—providing database administrators with the resources they need to establish and access one or more databases in the cloud—this white paper provides an overview, comparing and contrasting the high level attributes of these services.

Amazon Web Services or Microsoft Azure

The reason that AWS maintains its leadership in this market is because it provides 92 percent of the required features for enterprise use, while Azure only provides 75 percent. Microsoft has also gone to great lengths to support Linux on its Azure platform, including a command line interface, MySQL clustering, LAMP stack, Docker containers, gaming servers, and support for rpm-based and dpkg-based distributions. Azure also supports SUSE, openSUSE, CoreOS, and FreeBSD in its MSOpenTech VM Depot.

Developer tools and support are both important to businesses because many want to use public cloud infrastructure as test environments for production services. A full complement of tools is essential to this effort. Microsoft offers the online version of its Visual Studio suite and the Azure SDK, which is a powerful integrated development environment (IDE.) Using it, developers can write, test, and deploy code, plus make database connections and on-the-fly changes to code.

Amazon counters with its own set of tools, including CodeCommit for version control, CodeDeploy for automated application deployment and updating, CodePipeline for continuous code modelling, visualization, and automated releases, and Command Line Interface which provides a unified tool that allows developers to interact with all parts of AWS.

For the compute service, Microsoft Azure has two options: Cloud Services and Virtual Machines. Virtual machines are: dedicated amounts of compute power in the form of CPU cores (1 to 8), RAM (0.75 to 14GB), disk space (20 to 240GB), and an operating system/database server (Windows, Linux, Oracle, SQL Server.)

Alternatively, AWS offers full virtualization via its EC2 instances, containers via Amazon ECS, virtual private cloud, elastic load balancing, auto scaling, applications via AWS Elastic Beanstalk, and AWS Lambda, which is a highly available compute platform for backend web developers that is also a zero administration service.

Note: Features and specific offerings are always subject to change. Amazon has made numerous changes to Amazon RDS and Microsoft has updated the scaling tools available for its Azure SQL as this paper was in the process of being developed. Therefore, there may be newly-added tools or features not explicitly mentioned in this whitepaper.

What is Microsoft Azure SQL Databases?

Microsoft Azure is the Redmond-based tech giant's cloud computing platform. They tout it as "a growing collection of integrated services—analytics, computing, database, mobile, networking, storage, and web—for moving faster, achieving more, and saving money." In other words, it is designed to tick all the boxes an IT pro would need for development and productivity across several different types of environments. And it includes, the Azure MySQL Database. To that end, it offers hybrid database and storage solutions for those IT professionals who want some assets to remain fixed in their datacenters.

Microsoft continues to take great pains to differentiate its cloud-based product from Amazon Web Services, even devoting an entire page to the task. The database offering is but a fraction of the suite of cloud services both Microsoft and Amazon offer.

What is Amazon RDS?

Amazon's Relational Database Service (Amazon RDS) is a web service that makes it easier to set up, operate, and scale a relational database in the cloud. It provides cost-efficient, resizable capacity for an industry-standard relational database and manages common database administration tasks. The MySQL-based service was launched in 2009 and became generally available in 2013. According to the company, "It provides cost-efficient and resizable capacity while managing time-consuming database administration tasks, freeing you up to focus on your applications and business." It now offers six different database engines, which will be discussed as the features of both are explored later in this paper.

"When it comes to hosted database solutions, Amazon is at eye level or even ahead of the other giants in the industry, such as Oracle, Microsoft, IBM, SAP and Google," Matthias Gelbmann, co-founder of DB-Engines, told TechRepublic.

Microsoft Azure SQL Databases vs. Amazon RDS: 8 Key Features

Each of these cloud-based services has a slightly different approach to its database offering. Discussed below are eight features of keen interest to database administrators: the system architecture, its backup and restore capability, how scaling is handled, database size parameters, database management features and tools, container support, databases supported and pricing.

A few differences between these technologies has been showcased:

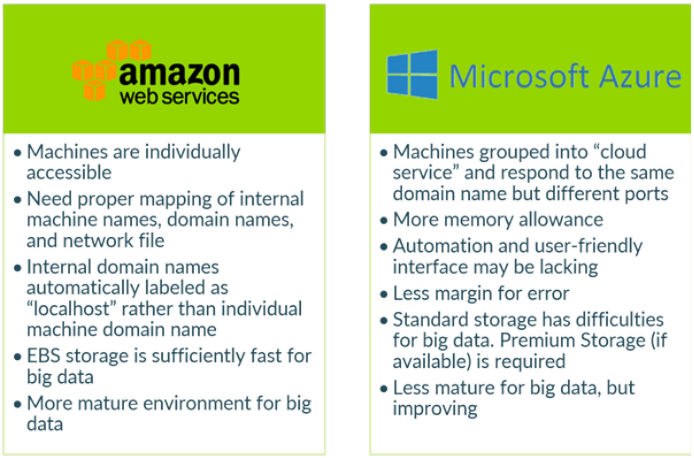

The prospective user will want to explore these key features in detail in order to determine whether a given service fits their organization's needs. As mentioned, the database offering is but a fraction of Microsoft's and Amazon's cloud services. There may be issues related to, for example, security or regulatory compliance that may merit consideration prior to making a selection.

Microsoft Azure vs Amazon AWS: Features and Services

Selecting one cloud over the other will come down to the wants and needs of each individual organization and the workloads being run/implemented. In fact, it is often the case that organizations will use both providers within different parts of their operations for different use cases. However there are a number of differentiating factors that separate the approaches of the two firms, which can help end-users consider which is right for them.

Start by looking at the services offered by the two firms. Both AWS and Azure offer similar capabilities to some degree along the lines of computing, storage, and networking. They share the common elements of a public cloud: self-service and instant provisioning, auto scaling, plus security, compliance, and identify management features. Both companies also continually invest in meeting demand for new cloud services. This has resulted in more mature analytics offerings. For example, support for Hadoop clusters are provided by both AWS (Elastic Map Reduce) and Azure (HDInsight). AWS and Microsoft have added machine learning tools

and a number of features targeted at the Internet of Things, while customers can tap either cloud to build a mobile app or even create high performance computing environment depending on the need. Meanwhile the recent buzz around containers is also addressed, with both providers supporting Docker services. Microsoft has also been creating its own Hyper-V container technology that it will run in on-premise datacenters.

1. Architecture

The primary difference between these offerings is their architecture. Amazon dedicates hardware resources to the user, while Microsoft shares resources among users. This difference contributes to the disparity in pricing, which is addressed as a separate issue below.

2. Backup and Restoration

Amazon RDS creates a storage volume snapshot of your DB instance, backing up the entire DB instance— not just individual databases. Amazon RDS provides two different methods for backing up your Amazon DB instances: automated backups and customer-initiated DB snapshots. Automated backups automatically back up your DB instance during a specific, user-definable backup window, and keeps the backups for a limited, user-specified period of time (called the backup retention period); you can later recover your database to any point in time during that retention period. DB snapshots are user-initiated backups that enable you to back up your DB instance to a known state, and restore to that specific state at any time. Amazon RDS keeps all DB snapshots until you delete them. Amazon RDS automated backups and DB snapshots are currently supported for all DB engines. For the MySQL DB engine, only the InnoDB storage engine is supported; use of these features with other MySQL storage engines, including MyISAM, may lead to unreliable behavior while restoring from backups. Specifically, since storage engines like MyISAM do not support reliable crash recovery, your tables can be corrupted in the event of a crash. For this reason, it may be wise to consider using the InnoDB storage engine. But in Azure, any SQL, or MySQL, the backup and restore feature requires the App Service plan to be in the Standard tier or higher. For more information about scaling your App Service plan to use a higher tier, see Scale a Web App in the Azure App Service. Note that Premium tier allows a greater number of daily backups than Standard tier. The Backup and Restore feature requires an Azure storage account and container that must belong to the same subscription as the web app that is being backed up. If a storage account has not yet been created, the Storage Account in the Backups blade of the Azure Portal, and then chooe the Storage Account and the Container from the Destination blade. For more information on Azure storage accounts, see the links at the end of this article. The Backup and Restore feature supports up to 10GB of website and database content. An error will be indicated if the backup feature cannot proceed because the payload exceeds this limit.

3. Scaling Up and/or Scaling Out

Scaling is a key issue for many database administrators. There is a challenge inherent with throttling resources to best meet an organization's current operational parameters. More CPUs may need to be added to accommodate growth, for example. Or performance can be increased by increasing the I/O in the environment, which is known as scaling out.

Both providers recently changed the products such that they now allow users to scale both up and out. Microsoft now provides its Elastic Database tool set available since April 2015.

4. Database Size

MySQL, PostgreSQL, and Oracle RDS database instances can now be created with up to 6TB of storage. Existing MySQL, PostgreSQL, and Oracle RDS database instances can be scaled to these new database storage limits without any downtime. Users of large transactional databases and data warehouses can now run even larger, higher performing workloads on a single database instance without needing to shard the data across multiple RDS instances. The new storage limit doubles the available storage for MySQL, PostgreSQL, and Oracle databases on RDS.

For Microsoft Azure SQL Databases, database storage is sold in three different tiers -- Basic, Standard, and Premium. The Basic tier has a maximum size of 117 GB, the Standard tier has a 1 TB maximum, and the Premium tier has a 750 GB maximum. The Premium tier can handle up to 1500 transactions per second, while the Standard and Basic tiers handle up to 1200 transactions per second.

5. Database Management Features and Tools

Amazon RDS supports DB instances running several versions of MySQL. First use the Amazon RDS management tools or interfaces to create an Amazon RDS MySQL DB instance. Then use the Amazon RDS tools to perform management actions for the DB instance, such as reconfiguring or resizing the DB instance, authorizing connections to the DB instance, creating and restoring from backups or snapshots, creating Multi-AZ secondarys, creating Read Replicas, and monitoring the performance of the DB instance.

Use standard MySQL utilities and applications to store and access the data in the DB instance.

In Azure using Azure portal, MySQL Instances can be created on the fly. As for being third party - yes, the service (MySQL-as-a-Service) is provided by a third party but it lives in Azure! It is only housed on the Azure servers, primarily as worker roles. The third party vendor is ClearDB— They provide MySQL as a service on top of Windows Azure. So, all your data is in the same Data Center as the client web site. Install the MySQL workbench to manage database connections, and resize the azure data disk via azure power shell script.

6. Container Support

Both Microsoft and Amazon offer container support with Docker integration. The Azure Docker VM extension can be used with the Azure Linux Agent to create virtual machines able to host containers for applications via Azure. Amazon EC2 Container Service, which also supports Docker, allows users to run applications on Amazon EC2 instances.

Containers are the hot new technology because they allow developers to spin up a small virtual system without the overhead of a virtual machine. Containers are much smaller and only use the resources they need. However, Containers have been primarily a Linux phenomenon that has slowly migrated to Windows, which gives Amazon an advantage. However, that does not mean Azure is a Docker-free zone. Quite the contrary. Microsoft worked with Docker and Mesosphere to create a container scheduling and orchestration service called the Azure Container Service. It will also come native in the Windows Server new version.

Azure Container service uses Mesosphere's Datacenter Operating System (DCOS), a massive machine that makes the servers appear as a single elastic pool of resources, allowing users to manage virtualized or non-virtualized datacenters. So slowly Microsoft is closing the gap in container/Docker offerings.

7. Databases Supported

When evaluating the database offerings of each service, there's no strict one-to-one comparison that can be made. It's a fruit salad of apples and oranges. Each service has a variety of databases and/or approaches for users. Does Amazon RDS have a MySQL offering? Yes. Does Azure have a MySQL offering? Well, yes, but it's not native.

Azure provides users with the ability to create a MySQL database using ClearDB. This is a fault-tolerant database-as-a-service (DBaaS) provider "that allows you to run and manage MySQL databases in Azure datacentres and connect to them from any application."

Amazon RDS was originally based on MySQL, but now offers a choice of database engines. Today, they recommend going with Innodb engines. These include Amazon's own Aurora, PostgreSQL, and Oracle, but DynamoDB, a NoSQL database responsible for spurring RDS user adoption, and Maria DB are also offered.

8. Pricing

Price is a critical differentiator for many organizations. Typically, costs need to be aligned with value such that there is a compelling business case for the technology that is justifiable to executives outside of IT. Costs also differ based on how the service will be used. The price for storage, for example, is not the same when using a cloud provider for compute time.

Comparing pricing for cloud services is notoriously difficult in an apples-to-apples manner. This is because there are differences in the configuration and measuring of computing units, as well as differences in the myriad services offered – compute, storage, database, traffic and so on.

For example, an AWS m3.large instance (21 CPU, 3.75 GB memory) costs $0.133/hour. The closest equivalent through Microsoft is the Medium VM (2 x 1.6GHz CPU, 3.5GB RAM) which costs $0.45/hour. Despite the tit-for-tat price cuts, Azure still seems to be more expensive than AWS for computing, but the former also offers good discounts for upfront, long-term payments. Read more at the respective pricing pages at azure.microsoft.com.

"RDS will actually cost less when dealing with a small CPU and memory allocation and having more than 10GB of data," state Ellis and Reed. "This pricing difference is significant enough that many organizations will make the choice solely on price."

Amazon made large-scale distributed computing cheaper and offers "very advanced data technology" with its cost-effective service. "Microsoft is playing catch-up, but it may be making some inroads" with Azure SQL.
And a neutral opinion was voiced by William Van Winkle, who contends both providers offer transparent pricing to users. Both also offer users the benefit of try-before-you-buy versions of their services.

Pricing Winner: SQL Azure

Due to architectural differences, the two companies have set up their own way of charging for their offerings. SQL Azure only charges by the amount of data stored (database size). There is no cost for dedicated CPU compute time or for memory use.

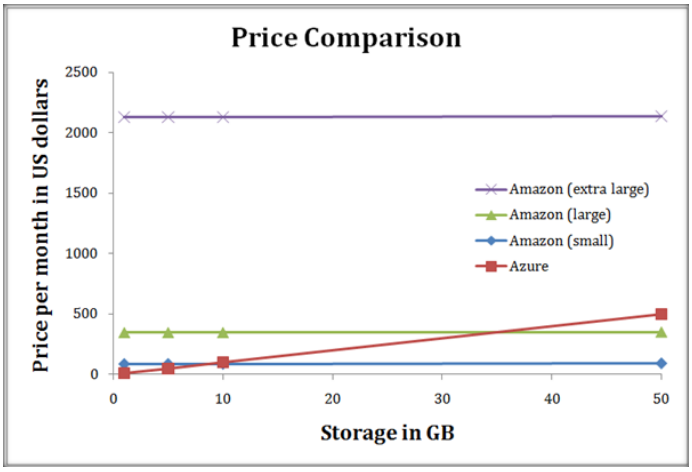

In contrast, the Amazon offering charges for CPU time, regardless of the size of the database. SQL Azure can be priced between $10 - $500/month, while RDS can cost between $84 - $2100/month. As the chart shows below, SQL Azure can cost significantly less than RDS, but it depends on the particular situation.

By zooming in on the graph, the data shows that RDS will cost less when the situation requires a small CPU, small memory allocation, and more than 10GB of data. This pricing difference is significant enough that many organizations will make the choice solely on price.

Backup/Restore Winner: RDS

Is this phone call familiar? "Uh, I accidentally deleted all the customers in the customer table… in production." That's when Point-in-Time restore will save the day. This category is an easy one to call because RDS offers backup/restore, while SQL Azure does not.

RDS's backup strategy allows for eight days of backups. It will also backup the logs and restore those logs to a specific point in time. The SQL Azure solution must be configured manually using existing on-premise tools which causes additional bandwidth charges. This is a significant drawback to SQL Azure which may prevent companies from migrating to this product.

Scaling Winner: RDS

When scaling a database, there are two choices: 'scaling up' or 'scaling out.' 'Scaling up' is adding more CPUs or memory to one box, allowing it to process requests faster. "Scaling out" is partitioning the database into smaller chunks to put it on more servers, increasing the I/O throughput. Because RDS runs on dedicated hardware, scaling up is possible by choosing how much processing power and memory the instance uses. Although that is expensive, it is available. SQL Azure doesn't offer the option of scaling up. Neither product offers an in-the-box scale out solution, but third-party software is available for both.

Database Size Winner: RDS

RDS allows you to have up to a 2TB database size with Innodb, while SQL Azure has a maximum capacity of 500GB. This will prohibit some companies from moving to SQL Azure.

Performance Winner: RDS

Datavail ran rudimentary performance tests from San Diego, including INSERTs and SELECTs against both products. Data showed that RDS was significantly faster on INSERT performance, while SQL Azure was slightly faster on SELECT performance. Since Amazon's datacenter is in Northern California, while Microsoft's datacenter is in San Antonio, the Datavail location may account for some of the performance difference.

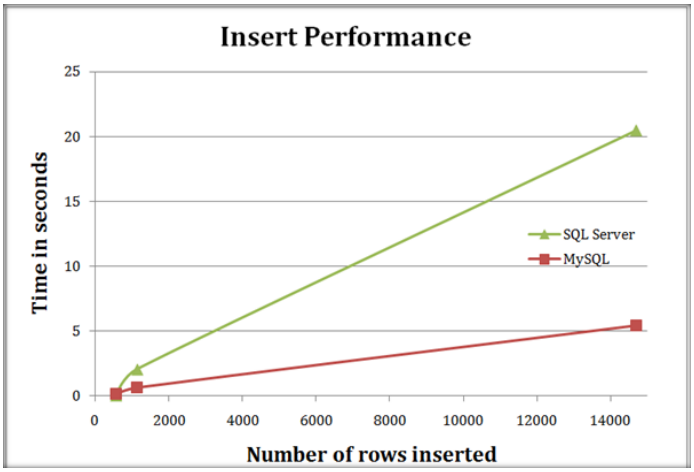

It might also be because RDS runs on dedicated hardware, while Microsoft is a shared service. During the study RDS was run on the smallest allotment of CPU and memory. With additional financial investment, RDS could be even faster. Again, that option isn't available on SQL Azure.

Feature Set Winner: RDS

RDS runs as an instance of a full-blown MySQL installation. As a result, RDS has every MySQL feature in existence. RDS users can also choose a specific MySQL version to deploy. SQL Azure has a subset of the features that ship with SQL Server 2008. While SQL Server 2008 is arguably more full-featured than MySQL; feature for feature SQL Azure is might come out ahead, users who are used to particular features from SQL Server 2008 may find that they are not available in SQL Azure. For example, SQL Azure does not support XML indexes, CLR objects, or SQL Server Profiler.

Tooling Winner: SQL Azure

As previously stated, both vendors enable the user to use tools that they are already comfortable with. SQL Azure can be managed using MySQL Windows server 2012 R2 or Linux VM with 5.6 MySQL release. Connection can also be achieved using MySQL workbench. Datavail used MySQL Workbench 5.2 to manage RDS, but all of the existing MySQL tools are not available. All the tools for both products worked seamlessly and without incident. Because the cloud allows users to experiment with different technologies and platforms without the burden of management, installation, and licensing, it was determined that a cloud database should provide some management software as a service (SaaS.) This would allow novices to get comfortable with a new environment. Microsoft addresses this with a Silverlight database tool called Database Manager (formally Project Houston.) Using this simple and small-featured tool, anyone can get started creating tables, adding data, and creating stored procedures without installing local software. Datavail ran it using Internet Explorer on Windows 7 and using Chrome on Mac OSX. While the fact that it even exists gives Microsoft a significant edge over RDS, Microsoft's tooling could be improved.

Disaster Recovery Winner: SQL Azure

RDS allows the creation of a standby replica of a database, but this is not done automatically. SQL Azure automatically creates standby servers when a new database is created. RDS's standby replicas can be in a different datacenter in the same geographic region (called a multiple availability zone deployment), unlike SQL Azure. In SQL Azure, there is no additional cost for standby replicas, while RDS replicas can double or even triple the cost.

Future Roadmap Winner: SQL Azure

The future of SQL Azure assures an even stronger cloud database offering. Well-published roadmaps promise new features including SQL Azure Reporting Services, SQL Azure Data Sync Services, and SQL Azure OData support. An online search reveals that Microsoft is addressing their woeful backup shortcomings and their limited maximum database size. In videos from PDC and TechEd Berlin, references to MySQL Server Analysis Services in the cloud were made as well as references to dedicated compute and memory SLAs. Meanwhile, Amazon has said that in Q2-2011, they will offer RDS using Oracle 11g as well as MySQL. While that is compelling, it is expected that they will raise prices even higher due to the licensing costs.

Conclusion

There are significant differences between Microsoft SQL Azure and running similar platforms on Amazon RDS. If price is not the determining factor, RDS offers greater performance and range of features. But SQL Azure has enough features to satisfy most users and it costs less. Microsoft users will be more comfortable and familiar with SQL Azure. Developers who work with MySQL will probably find RDS more comfortable.

There is no clear answer as to which product a given organization should adopt. The one that best fits an organization's history and needs will be optimal.

Biography

Dilip Kumar
Senior MySQL DBA for Datavail

Dilip has 12 years of experience in the IT industry, working as a MySQL database administrator. He specializes in MySQL database administration and technologies like MySQL, Cluster Services, Oracle Golden Gate, and scripting. He also specializes in database architecture, DBA, HA Solutions, performance tuning, and managing enterprise level database support.

About Datavail

Datavail Corporation is the largest provider of remote database administration (DBA) services in North America, offering database design and architecture, administration and 24x7 support. The company specializes in Oracle, Oracle E-Business Suite, Microsoft SQL Server, MySQL, MongoDB, DB2 and SharePoint, and provides flexible on-site/off-site, onshore/offshore service delivery options to meet each customer's unique business needs.

Contact Us

General Inquiries:
877-722-8247
Fax Number: 303-469-2399
Email: info@datavail.com

Corporate Headquarters:
Datavail Corporation
11800 Ridge Parkway
Suite 125
Broomfield, CO 80021

Database Operations Control Center:
Datavail Infotech Pvt. Ltd
3rd Floor, Unit No. B-3
Ashar IT Park, Road No. 16Z
Wagale Estate
Thane (West), Thane 400604
Direct Telephone Number: 022-61517000

Bangalore Office
Datavail Infotech Pvt. Ltd
Concept Business Park
#319/9, 1st floor, Block A
Hosur Main Road
Bommanahalli, Bangalore 560100

datAvail

DATABASE SERVICES

www.datavail.com | 877.634.9222