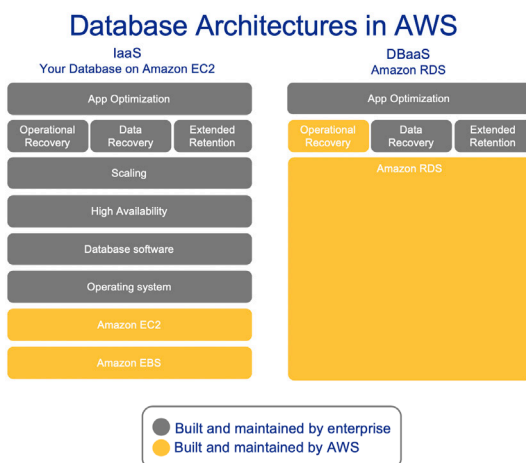


Removing Data Protection Roadblocks for Amazon RDS

Traditionally, enterprises are known for building their own data centers and hosting operational database systems to store, process, and retrieve data for business applications. These databases were indispensable and could never go down, and had to be nurtured accordingly. These were meticulously built, managed, and upgraded with utmost care. Two megatrends of the last decade--cloud computing and open-source software--are disrupting the way databases are now set up and managed. The popularity of database as a service (DBaaS) offerings from cloud service providers (CSPs) has enabled organizations to treat database systems like a service utility--you spin up as many database instances as needed and spin them down when you are done.

The DBaaS delivery model liberates organizations from setting up and managing what is arguably the most complex and difficult-to-manage layer of enterprise applications--the database. Instead of worrying about provisioning, scaling, monitoring, and managing databases by employing dedicated database administrators (DBAs), now engineering teams can focus on building line-of-business applications with the agility and efficiency from these database engines managed by cloud service providers. Amazon leads the market with Amazon Relational Database Services (RDS).

What's Been Holding RDS Adoption Back?



RDS has gone mainstream, but some enterprises continue to build their own databases in IaaS environments such as Amazon EC2. At the first look, it is easy to hypothesize that the trend here is to minimize modifications to existing applications. We interviewed the line-of-business leaders and cloud architects from several fortune 100 companies, and a pattern quickly emerged on this last line of resistance in embracing DBaaS offerings such as RDS for enterprise applications. The bottom line is that customers have two main concerns regarding moving to DBaaS--the first is wanting to control their own data, and the second is concerns about exposure to security breaches and compromised accounts.

Enterprises Want to Control Their Data

Data from business applications outlive the infrastructure where they currently reside. Enterprises want to ensure that the data is always in their control and may be ported across platforms. Two of the top three concerns raised by enterprises tie into the fear of vendor lock-in. While enterprises welcome the agility of cloud computing, and yearn to stay away from building and maintaining data centers, they worry that DBaaS may lock their data onto cloud services providers thereby making it difficult to repatriate data off cloud or migrate to another cloud services provider.

This fear is further exacerbated by limitations in extended retention capabilities of data in DBaaS offerings to meet industry regulations and company policies. The way to perform backups for databases served by Amazon RDS is snapshots which, in turn, is tied to AWS platform. Thus, even if they managed to move off production data with the help of professional or specialized service offerings from third party vendors, they still need to worry that retrieval of data stored long term is still tied to AWS services. Hence they need to maintain those accounts to meet compliance and respond to ad-hoc requests such as legal hold and eDiscovery.

Exposure to Data Breaches

The DBaaS offerings are tied to accounts the customer owns in CSPs. In the case of RDS, the database instance sits in a virtual private cloud (VPC) in a customer's AWS account. The key backup method of RDS is snapshots which co-exist with RDS instances in the same account. If an AWS account is compromised by a malicious actor, be it an insider or outsider, backed up snapshots can be deleted or modified along with production instances. This is a terrifying reality that many companies now face and need better tools and processes to mitigate.

There are ways to share and copy full blown snapshots onto alternate accounts, however those solutions have challenges:

Unpredictable costs:

Snapshots for a given RDS instance are maintained as a chain. The first snapshot is a full copy of RDS and subsequent ones are incremental. This helps in optimizing the storage costs associated with snapshots. However, once you start copying snapshots out of the account, the snapshots chain is broken and everything out of the account is fully blown snapshots. This results in unpredictable costs.

Operational Complexity:

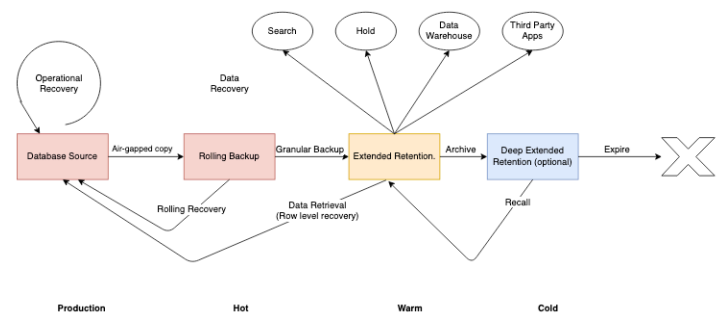
If the primary account is breached, there need to be guardrails in place to make sure that the alternate account is inaccessible from the information/credentials/keys that are already breached. This process of 'air-gapping' two accounts from one another requires significant investments in building scripts, key management, and procedures to be carefully implemented and regularly tested.

Reduces agility:

One of the key benefits of moving to the cloud is agility. With cross-account snapshot management and air-gapping needs to protect from breaches, the change control processes for production become procedural, rather than staying agile.

The Best of Both Worlds - RDS for Enterprises Without the Challenges

In order to design a data lifecycle management system that is suitable for enterprise, we need to look at the requirements of enterprises. Luckily, the requirements can be easily gleaned from how they used to manage data lifecycle in on-premises. Data from production database systems goes through the following lifecycle.



Backups for Operational Recovery

If the availability of a database serving mission critical applications is affected due to software glitches, infrastructure failures or human errors, backups stored close to the database act as the first line of defense. The recovery point objective here is minutes to days, while recovery time objective is minutes to hours--depending on the nature of the applications it is serving.

Recover Quickly From Data Breaches

Rolling backup is essentially a time lagged standby of the production database. Transactions on production are replicated to a standby database in an alternate site, but deliberately time lagged behind production by a predefined number of hours. This enables quick recovery without a full restore from backup, in case the production copy is affected by accidental or deliberate human intrusion.

Extended Retention and Records Retrieval

Enterprises store data beyond the operational recovery window to satisfy regulatory requirements, company policies, and historical data analytics. Organizations in regulated industries such as finance and healthcare create self-contained full backups on tapes on a periodic basis (e.g. every 14 days or every month) and store them offsite. They may also be making use of offsite storage services such as those from Iron Mountain™. During the last five years, cloud based object storage solutions such as Amazon S3/ Glacier are becoming a viable alternative to tape storage.

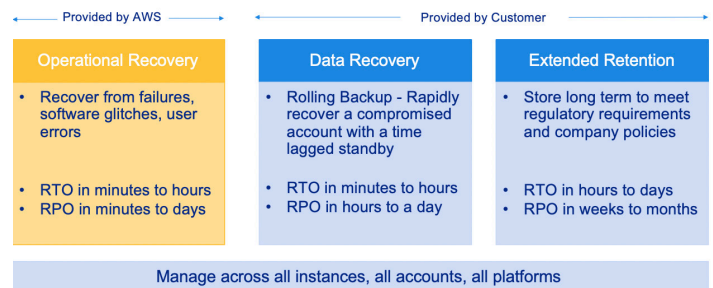
For enterprises, data is the new oil. Many industries prefer to retain analytical data for extended periods. One of the characteristics of extended retention copies is that those are used for historical data retrieval, rather than actual database recovery. However, organizations end up recovering the entire database onto a sandbox environment where the required database software stack is present so that they could use standard database access tools to retrieve the desired historical data. Thus, recovery here is a means to the end goal of data retrieval. Amazon RDS has built-in snapshots capabilities (known as RDS automated backups) that serve the first line of defense needed by enterprises. The rest of the capabilities need to be provided by the customer to meet the [shared responsibility model](#) of AWS. This can be made possible by a solution set that helps to manage the built-in first line defense capabilities of RDS while also helping to fulfil the shared responsibility for data lifecycle.

Deep Extended Retention

The last, and optional, phase of data is cold storage. Data is stored simply for the purpose of meeting regulatory requirements. The recall of data is a rare event. If and when data needs to be recalled, it is moved thawed into a warm tier for data retrieval workflows.

Amazon's shared responsibility model dictates that operational backups is AWS's responsibility. This is because operational recovery tools such as snapshots are part of the RDS platform itself. Once the data leaves the operational stage of its lifecycle, it is the responsibility of the customer to take care of data. Furthermore, as AWS accounts are the security perimeters of organizations or projects, rolling backups to alternate accounts is the customer's responsibility as well. This is depicted in the graphic below.

Making Amazon RDS ready for Enterprise Applications Meeting Enterprise Data Management Requirements



Conclusion

Amazon RDS brings agility and operational simplicity to operational database management systems. With the acceleration of workloads moving to the cloud, enterprises can make use of Amazon RDS to operational needs without compromising on regulatory and security requirements if they also adopt a database data lifecycle model that works in conjunction with the shared responsibility model of AWS.

Clumio unleashes the full power of the cloud to deliver secure backup and recovery for your data--wherever it needs to be. Clumio helps to fulfill the customer's responsibility of data protection in the cloud by closing the gaps in data protection and lifecycle management tasks. Clumio also helps to recover quickly in case of account breaches with the use of backups that are air-gapped away from production accounts. And, finally all the phases of data cycle management from maintaining operational recovery copies to extended retention can be easily achieved by unified protection policies that can be applied across all supported workloads, accounts and platforms.

About the Author

Abdul Rasheed is Sr. Director of Product at Clumio, where he is focused on building data management solutions for DBMS, NoSQL and Big Data. With his heart in technology and mind in business, he brings over a decade of experience in solving business critical problems through technology solutions. Prior to Clumio, Rasheed held positions in various independent software vendors focussed on solutions targeted for enterprise applications. #UMNproud from Carlson School of Business, University of Minnesota.