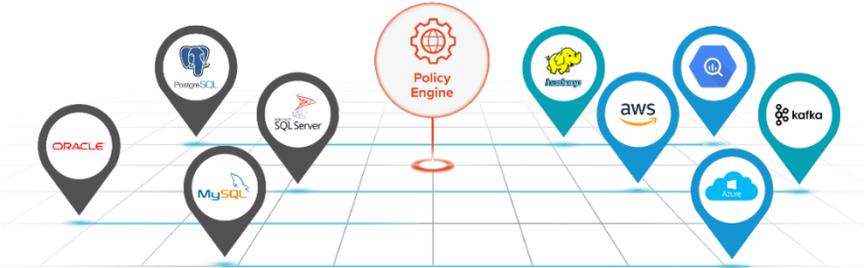# Hybrid Cloud Operations with Griddable.io

As enterprises embrace hybrid cloud as the new long-term reality, database architects need a simple approach to migrate enterprise data to the cloud, re-engineer for cloud-first architectures, and ensure portability of data across multiple public clouds to avoid lock-in. The Griddable.io smart grid for enterprise data continuously synchronizes data across heterogeneous clouds, databases, and schemas.

Key Benefits

- *Data portability across heterogeneous databases and clouds.*
  Rapidly respond to changing price/performance demands without service interruptions while avoiding cloud or database vendor lock-in

- *Selective, continuous synchronization in any topology*
  Synchronize and selectively filter operational databases to many simultaneous innovation or analytics projects to enable the distributed use of continuously up-to-date system-of-record data.

- *Selective data transformation customized to each destination*
  Transform data in transit across database versions, schema, and platforms. Synchronize legacy and new cloud-first systems and meet privacy requirements by protecting personal or sensitive data from unauthorized access or distribution

Hybrid cloud operational challenges

The Griddable.io smart grid for enterprise data guarantees transaction consistency across any number of clouds, databases, and schemas

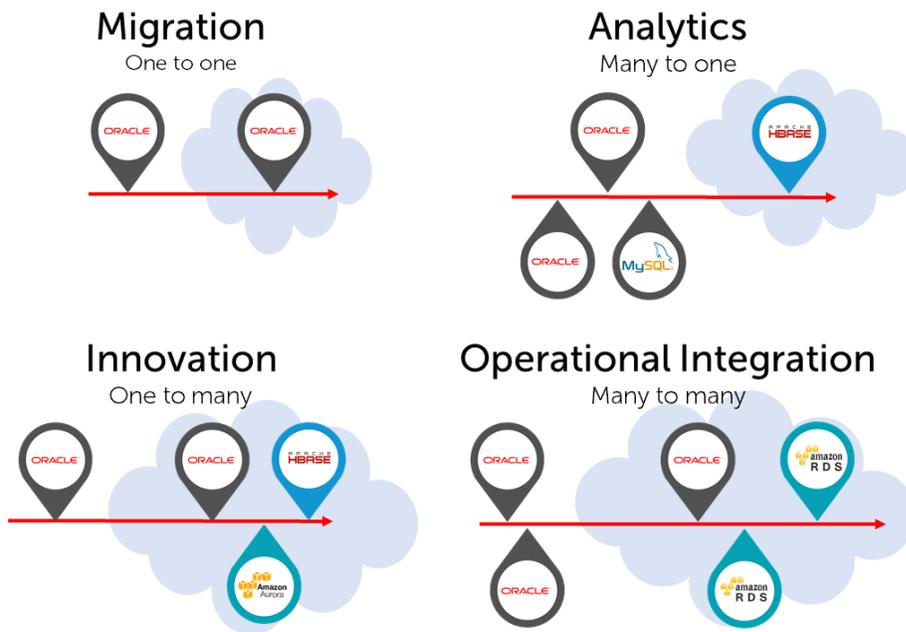The hybrid cloud operation cannot be effectively managed with legacy frameworks and tools.

- One-time database lift-and-shift processes based on database backup lead to long service disruptions, and database copies used in dev-ops environments quickly become stale and irrelevant.
- Current point-to-point data database replication tools perform a like-to-like full database replication unless heavily customized. These products were designed for on-premise installation, are difficult to change without service disruptions, and do not deliver advanced features like scaling and high availability.
- Cloud vendor data migration services are specific to one cloud and risk lock-in to that cloud platform. Over 80% of organizations have moderate to high levels of concern about public cloud lock-in[1].

Data portability across databases and clouds

Continuous non-disruptive data migration

Griddable.io automatically and continuously synchronizes data across heterogeneous combinations of clouds and databases. Griddable.io data synchronization is performed in real-time, resulting in a grid of databases which are all time synchronized, consistent, and providing a customized data set appropriate to each use case.

Migration
One to one

Analytics
Many to one

Innovation
One to many

Operational Integration
Many to many

Create sophisticated one to one, many to one, one to many, and many to many topologies in minutes using Griddable.io

Each Griddable grid can synchronize databases across clouds, enabling operation teams to rapidly respond to changing price/performance demands without service interruptions.

Using Griddable.io, topologies can be modified at runtime without service interruptions to add new sources or destinations dynamically. Individual topologies can be combined to create hyper-connected grids. High availability and scalability options are available to easily customize out-of-box functionality with advanced features.

Avoid vendor lock-in with data portability

Griddable.io is a SaaS synchronization service available on a variety of public clouds. Each subscription creates an isolated environment individually owned and controlled for maximum flexibility and security across clouds. This flexibility is key to reaching the ROI available in the cloud. Cloud customers can save up to 74% of

74

cloud costs with an AWS+1 strategy by utilizing AWS and at least one other cloud.[2]

Selective, continuous synchronization

The Griddable.io smart grid for enterprise data enables continuous, selective data filtering to the row and column level on both the source and destination. Cloud operations data architects can easily synchronize operational databases to several simultaneous innovation or analytics projects, avoiding copy sprawl and stale data copies throughout the organization. Each destination data set can be defined to be a precisely-crafted collection of rows, columns, and tables, or can dynamically change using tags which refer to live data sets.

Using Griddable selective filtering and synchronization, cloud operators can accelerate movement of key data to best-fit platforms without requiring special skills. Decisions regarding data placement can be made dynamically during production and adjusted to operational parameters.

Selective data transformations

Transform data across versions and schemas
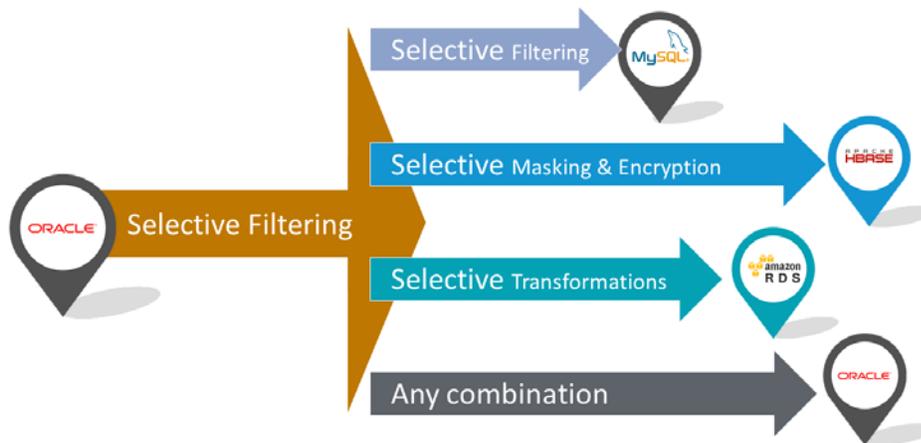
Griddable.io performs data filtering and transformations at the source or destination, or both, and the data on each destination is specific to that destination.

Griddable.io automatically transforms data across heterogeneous combinations of source and destination databases. All data transformations are performed in grid services without impacting the source data, eliminating the need to manually copy and transform data using backup copies or snapshots.

Using Griddable.io, operation teams can quickly complete simultaneous schema upgrades and database version changes across the enterprise that took weeks or months to complete in the past. Data transformations are specified in the Griddable.io policy, which is then used during data synchronization, automating repetitive and manual steps typically required on each database upgrade.

Griddable's cloud-first architecture

Griddable.io uses a loosely coupled distributed architecture of relays and consumers to synchronize and transform data. Relays interface with source databases, posting changes to an in-memory pub/sub buffer. Consumers read changes from relays and enter them in their destination datastores.

Between relays and consumers is the Griddable.io Change History Server, which archives changes and makes them available later. If a consumer falls behind due to load or network delay, the Change History Server provides required events without imposing delay on the relay, other consumers, or the source database.

The benefits of this loosely coupled architecture are reliability and scale. Consumers can pull changes from any relay, or the Change History Server, and relays are free to operate at optimal speed. There are no limitations in the system like the capacity of the file system to hold change data files. Both relays and consumers scale by simply adding additional relays or consumers and partitioning the load across them.

For further information on Griddable's cloud-first architecture, see the Griddable Architecture white paper, available on the Resources tab of company web portal.

Re-engineering projects across the hybrid cloud

Fully 81% of enterprises now have a multi-cloud strategy[3], with 48% of organizations using three or more cloud vendors[4]. Using Griddable.io transactions from the source database are continuously synchronized and transformed to any number of destinations in real-time across any number of clouds.

Data re-engineering scenarios can be created and tested on the fly by simply updating the Griddable.io policy which controls both data synchronization and data name and value transformations to each destination. And, since Griddable policies can transform data uniquely to each destination, several data re-engineering layouts can be materialized and tested simultaneously.

Meet privacy and regulatory requirements

Many emerging privacy regulations such as GDPR require personal or sensitive data to be anonymized or remain in the locale of origin. Isolating personal data throughout the enterprise and limiting its distribution is extremely difficult and when accomplished with traditional database ETL and backups is nearly impossible to demonstrate to auditors.

Griddable.io provides complete control on how personal data is synchronized throughout the enterprise without reorganizing or disrupting source databases.

griddable.io

1 Stratoscale Hybrid Cloud Survey, July 25, 2017
2 AWS+1 will become default cloud strategy in 2017, 451 Research, November 29, 2016
3 Rightscale State of the Cloud Report, January 2018
4 Hosting and Cloud Study, 451 Research, 2016