

Enabling Database-as-a-Service (DBaaS) within Enterprises or Cloud Offerings

Introduction

Accelerating time to market, increasing IT agility to enable business strategies, and improving cost efficiencies are fueling the adoption of cloud computing today. While many of the scalability, availability, and agility infrastructure challenges have been solved – for compute, networking, and storage for example – cloud database offerings remain complex, rigid, and slow to respond to the needs of the market.

With the DBaaS market projected to reach \$1.8 billion by 2016¹, the need for an agile data tier in the cloud has become paramount. The problem with today's legacy database architectures lies in the forced 1:1 connection between applications and databases. Availability, scalability, and high-performance features of the database architecture must be coded into applications. This significant time investment – both initial and on-going – slows app development and QA.

In addition, databases continue to be a substantial source of downtime, both planned (maintenance windows for patching, schema updates, hardware enhancement) and unplanned (outages). Such downtime affects all users of the application or service.

Enabling database-as-a-service, with its inherent flexibility and scalability, can significantly improve uptime and lower the cost of operations. Enterprises have the opportunity to reduce database sprawl via consolidation, allow for rapid provisioning and automation, simplify patch management with zero-downtime maintenance, increase security, and enable quality of service (QoS) with multi-tenancy.

Historically, building DBaaS has been challenging, but new technologies, such as ScaleArc's load balancing management software, vastly simplify the task of enabling DBaaS in the enterprise, for private, public, or hybrid deployments.

.....
Enabling database-as-a-service (DBaaS), with its inherent flexibility and scalability, can significantly improve uptime and lower the cost of operations.
.....

¹ DBaaS poised to drive next-generation database growth, Matt Aslett, 451 Research, Aug 2013

Challenges in Building a DBaaS Offering

Software as a Service (SaaS) is a well understood and broadly adopted architecture in which an application supports multi-tenancy and can scale with the addition of more application servers behind a web load balancer or Application Delivery Controller. Databases, in contrast, remain deployed in support of a single application.

Building a DBaaS offering is extremely complex. Databases inherently do not perform multi-tenancy well, and scaling takes far more than simply adding more database servers. Applications have to be re-engineered to use more than one database server, and few apps today are designed to do so. As a result, most cloud providers have chosen to offer managed database software hosted on a virtual machine (VM). Such offerings – really nothing more than “Database-on-a-VM” – have been labeled DBaaS, but these solutions do not actually support “as a service” capabilities. Instead, they continue to put the onus of scalability, resiliency, and performance on the applications, requiring extensive engineering in the application to enable such functionality.

Other challenges exist as well with the Database-on-a-VM service. Databases can face severe performance degradation on virtualized platforms, and it is difficult to provision more instances and grow capacity on the fly. Also, Database-on-a-VM requires a separate copy of the OS and database for each VM, resulting in higher resource requirements and in some cases yielding the opposite result of the “as a service” model, with less consolidation and performance instead of improvements in those areas. Each instance of the database now has to be managed, patched, and troubleshooted independently, increasing application downtime.

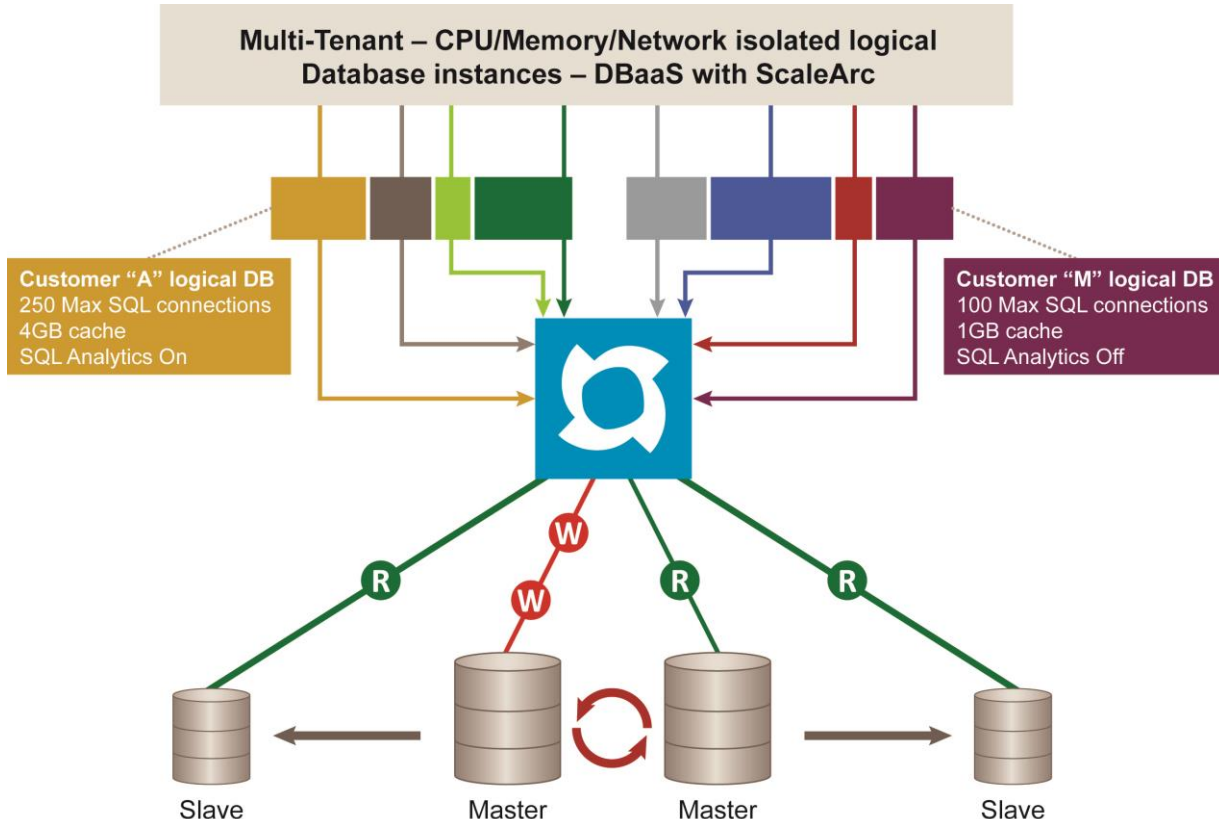
Amazon Web Services’ (AWS) Relational Database Service (RDS) offering is a perfect example of this Database-on-a-VM managed service. It provides only a single database instance and nothing more – no automatic load balancing or clustering, no transparent caching, no SQL instrumentation, and no resource-based billing. Such features would constitute DBaaS in a cloud computing context. RDS enables customers to provision multiple slave databases alongside a master for scale out, but the client applications still must be re-engineered to support this architecture. The app has to be updated to perform functions such as sending writes to the master and reads to the slaves or monitoring replication lag.

With the current approach, cloud-based database offerings force customers to accept severe compromises in scale and performance, significantly reduced value, and additional complexity. It also takes more hardware to provide the same performance since virtualization creates inefficiencies – the consequential increase in VM counts, in turn, compounds management costs.

.....
Databases inherently do not perform multi-tenancy well, and scaling takes far more than simply adding more database servers.
.....

A New Approach: Database Load Balancing Software

Database load balancing software is a new class of software that drops in transparently between application or web servers and the database. It looks to the app like a database and to the database like an app. It performs a range of SQL management capabilities, including read/write split, replication monitoring, and auto failover that dramatically improve the availability and performance of apps.



These sophisticated load balancing features make it incredibly easy to take large physical database clusters and provision logical customer databases within them – instant DBaaS for an enterprise.

ScaleArc offers database load balancing platforms for SQL Server, MySQL, and Oracle. By centralizing and automating database management tasks, a DBaaS enabled by ScaleArc software delivers zero downtime for apps, increased app performance, and lower operational costs. Perhaps most compellingly, the software delivers these capabilities with no changes needed in the application or database.

The ScaleArc software features built-in multi-tenancy, supporting complete data isolation between logical customer databases. Each logical database can also have its own logical cluster within ScaleArc – with complete network, CPU, and capacity control. Enabling this provisioning of logical databases lets cloud providers benefit from economies of scale and lowers the needed hardware and energy costs, while

at the same time dramatically improving application uptime and providing 2x to 3x the performance compared to hosting each database individually on VMs.

In addition to being able to quickly and economically enable DBaaS for their organization, ScaleArc customers can also enable the following value-added services – potentially generating additional revenue streams:

- SQL-aware dynamic database load balancing and high-availability services
- SQL surge queue services and zero-downtime maintenance environments
- Query response caching services to dramatically improve performance and availability
- Advanced database architecture services such as sharding and multi-master deployments
- SQL query filtering services
- Advanced real-time SQL analytics, instrumentation and reporting

ScaleArc is available as on-premise software that can be hosted on bare-metal x86 servers to create appliances or deployed as a VM on a standard hypervisor. ScaleArc was built from the ground up for cloud service providers and enterprises and has full multi-tenancy capabilities built in. ScaleArc offers license-agreement pricing for cloud service providers so that expenses can be matched to revenues for a highly profitable DBaaS business. ScaleArc can also enable DBaaS for enterprises to build a highly efficient internal private DBaaS offering.

.....
The ScaleArc database load balancing software provides for zero-downtime maintenance by allowing patch management of the database servers without incurring a maintenance window and by preventing application errors.
.....

The Top 5 Reasons to Use ScaleArc for DBaaS Offering(s)

1. Consolidation and multi-tenancy

ScaleArc reduces database sprawl by enabling consolidation of many database servers into a smaller set. Organizations gain significant operational efficiency by maintaining a small number of highly available database servers. ScaleArc supports DBaaS by offering the ability to provision multiple tenants over the consolidated servers, with different QoS levels for each tenant. The ScaleArc database load balancing software provides for zero-downtime maintenance by allowing patch management of the database servers without incurring a maintenance window and by preventing application errors. ScaleArc has IP/CPU/memory-level isolation between tenants (clusters) on the same instance. Each cluster has its own memory space, and customers can create clusters without sharing CPUs. ScaleArc also has a built-in wire-speed SQL filtering engine, allowing for firewall-like regular expression rules to be configured via the RESTful API or GUI. With the ability to set connection limits on a per-cluster basis, ScaleArc can enforce different QoS levels for the different tenants.

2. Rapid provisioning of new and tiered services

ScaleArc's transparent approach provides for an easier integration of new applications and services, accelerating revenue and time to market for new applications and services. ScaleArc can be completely managed by a RESTful API. Since new service creation depends on simply a handful of API calls into ScaleArc, providing a customer with a new database infrastructure setup and configuration takes just minutes. ScaleArc's

transparent caching and real-time analytics, applied on a per-tenant basis, enable new value-added services to support tiered service offerings. With ScaleArc, operators can build an easy-to-use platform that app developers and users can securely leverage.

3. Resiliency and auto scaling

ScaleArc's auto failover increases availability of the consolidated database servers – essential for achieving 99.999% uptime for SLAs. Staff can add database capacity automatically, making it available to applications instantly without requiring any changes to the apps. ScaleArc has built-in support for read/write split, which enables use of the scale-out architectures of MySQL master/slave or SQL Server AlwaysOn topologies. For example, an online gaming customer – with traffic that fluctuates wildly depending on time of day, viral factors, and promotional events – uses ScaleArc to provide for an elastic on-demand database infrastructure. With ScaleArc's SQL surge queue to buffer unexpected query bursts, application performance remains acceptable even during peak loads. ScaleArc's query routing enables incoming SQL traffic to be switched based on user name, database name, table name, source IP address, and other parameters – allowing for a single endpoint for all database access for reporting as well as online queries.

4. Faster time-to-market for new applications

Development time and expenses are reduced significantly, as is time to market for new apps, since ScaleArc does not require any application changes to scale, optimize, or add resiliency. Enterprises or cloud service providers who use ScaleArc's database load balancing software to offer hosted database and DBaaS will be able to sell this value to their customer base, allowing their customers to achieve business goals faster with accelerated application development. With ScaleArc's multi-tenant, QoS-enabled DBaaS, organizations can easily host non-production databases such as test, development, or training along with production and mission-critical databases.

.....
Development time and expenses are reduced significantly, as is time to market for new apps, since ScaleArc does not require any application changes to scale, optimize, or add resiliency.
.....

5. Superior performance with transparent caching

Successfully scaling databases requires real-time analysis at the query level to apply granular caching rules with continuous query-level performance reporting. These capabilities are either not natively available on databases or affect performance when reporting or logging is turned on.

ScaleArc provides a transparent query response caching solution that requires no changes to existing applications or database servers, with no development effort needed to realize the caching benefits. Even organizations that have already written their own home-grown software solutions based on Memcached benefit from ScaleArc's one-click caching as a complement. ScaleArc caching can reduce server disk I/O significantly and can improve response times up to 60x while reducing server loads up to 95% – depending on the application.

Conclusion: Faster, Simpler, Better DBaaS with ScaleArc

ScaleArc's database load balancing software creates an agile data tier that instantly eliminates the challenges inherent in today's application and database architectures. ScaleArc abstracts the database environment away from the applications to provide scalability, uptime, and performance while enabling a DBaaS environment. ScaleArc's multi-tenancy capabilities allow enterprise IT organizations and cloud service providers to quickly and easily offer DBaaS while offering additional value-added services such as caching and analytics to quickly grow a profitable DBaaS business and create additional value for their customers.



2901 Tasman Drive, Suite 205
Santa Clara, CA 95054
Phone: 1-408-780-2040
Fax: 1-408-427-3748
www.scalearc.com



ScaleArc is the leading provider of database load balancing software. The ScaleArc software inserts transparently between applications and databases, creating an agile data tier that provides continuous availability and increased performance for all apps. With ScaleArc, enterprises also gain instant database scalability and a new level of real-time visibility for their application environments, both on prem and in the cloud. Learn more about ScaleArc, our customers, and our partners at www.ScaleArc.com.

© 2015 ScaleArc. All Rights Reserved. ScaleArc and the ScaleArc logo are trademarks or registered trademarks of ScaleArc in the United States and other countries. All brand names, product names, or trademarks belong to their respective holders.