

MARIADB MAXSCALE

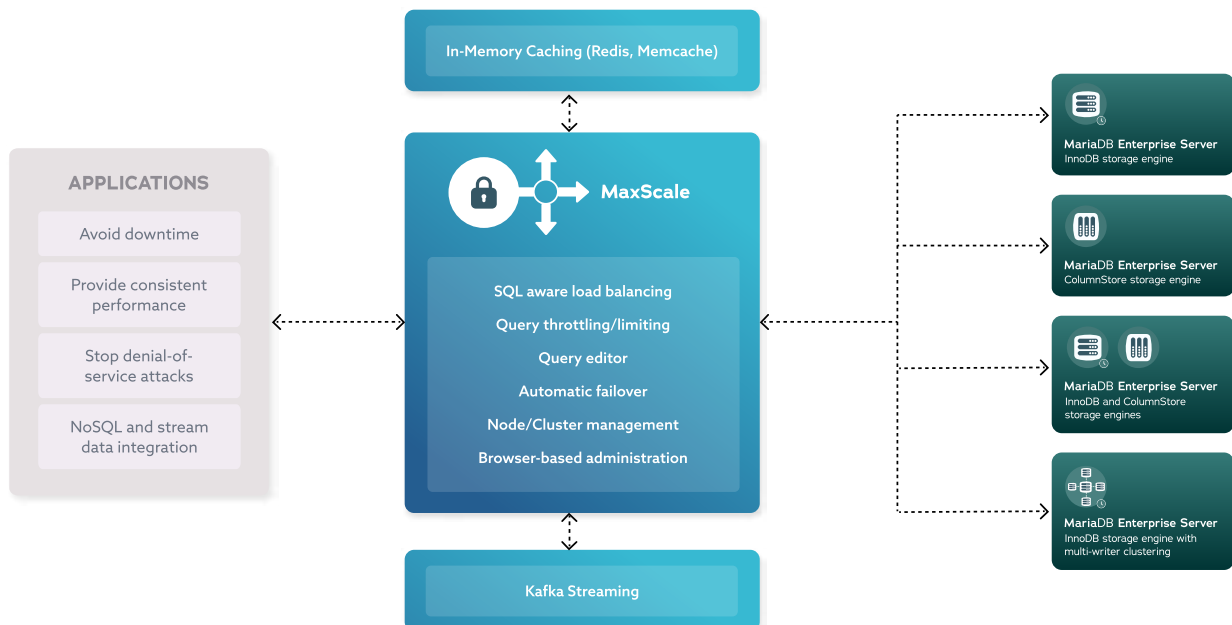
Technical Overview

MariaDB MaxScale is an advanced database proxy, a core feature of MariaDB Enterprise Server – powering its enterprise high availability, scalability, security and integration services. MaxScale is an intelligent, low-latency load balancer, expressly designed to be SQL-aware yet impose little overhead. MaxScale understands the complexities of backend database architectures and insulates client applications from them. It routes application requests based on a combination of defined algorithms, component state, request content, and session state.

INFRASTRUCTURE ABSTRACTION

MaxScale simplifies development and administration by abstracting away database infrastructure from the applications interacting with it – enabling developers to build applications without having to know what the underlying database topology is, and for operations teams and database administrators (DBAs) to make changes to it without impacting application uptime or requiring application configuration changes. As a result, applications will never know a failover has occurred, a switchover has been performed or a replica has been added.

In addition, MaxScale can route different queries to different database instances. For example, when deployed for hybrid transactional/analytical workloads, it can route transactional queries to instances with replicated, row-based storage and analytical queries to instances with distributed, columnar storage.




ENTERPRISE HIGH AVAILABILITY AND AUTOMATIC FAILOVER

MaxScale minimizes database service downtime and guarantees database uptime by extending the replication and clustering capabilities of MariaDB Enterprise Server with enterprise high availability services. The automatic failover service, when used with advanced features such as last transaction replay, nearly instant session restore and connection migration provides zero-interruption failover – an automatic failover that is completely transparent and nondisruptive to live applications with production traffic.



If the primary database fails, MaxScale will automatically promote an available replica, either in the same zone or a different one (even across regions), and begin routing writes to it – instantaneously, all while continuing to load-balance parallel reads across the remaining replicas.

MaxScale avoids being a single point of failure itself as it can be mirrored and coordinates between nodes of itself without application code.



ENTERPRISE SCALABILITY

Maximize performance without sacrificing consistency and simplify application development by taking advantage of enterprise scalability services such as transparent read/write splitting. It enables replicas to be used for read-scaling without modification to application code or configuration, and when causal reads are enabled, uses global transaction IDs (GTIDs) to enforce read-your-writes consistency. In addition, it prevents write conflicts and collisions when clustering is used, routing all writes to a single node. And finally, if adaptive load balancing is chosen, it continuously routes reads to the fastest database instance based on current response times.

DENIAL-OF-SERVICE PROTECTION

In addition to data breaches, denial-of-service (DoS) attacks are a prevalent threat as complexity, volume and duration of these attacks are increasing. These attacks not only cause site outages and damage reputation, they can cost businesses millions of dollars. The DoS protection features in MaxScale protect the database from being rendered unavailable due to malicious attacks or accidental queries that, if allowed to proceed, would overwhelm the database infrastructure with a result equal to a failure.

Query throttling

The query throttling filter prevents erroneous query loops or DoS attacks from rendering the database unavailable by placing a threshold on the maximum number of queries per second allowed. There can be short bursts within a defined window, but in order to protect the database, an application's session will be closed if it continues to exceed the threshold defined.

Result limiting

The result limiting filter prevents accidental or malicious queries from rendering the database unavailable or exposing large amounts of data by limiting the number of results a query can return – for example, to prevent a query from returning all ten million rows in a table. It not only eliminates the overhead (e.g., network) of returning ten million rows, but prevents significant amounts of data from being exposed to an attacker who finds a way to query the database.

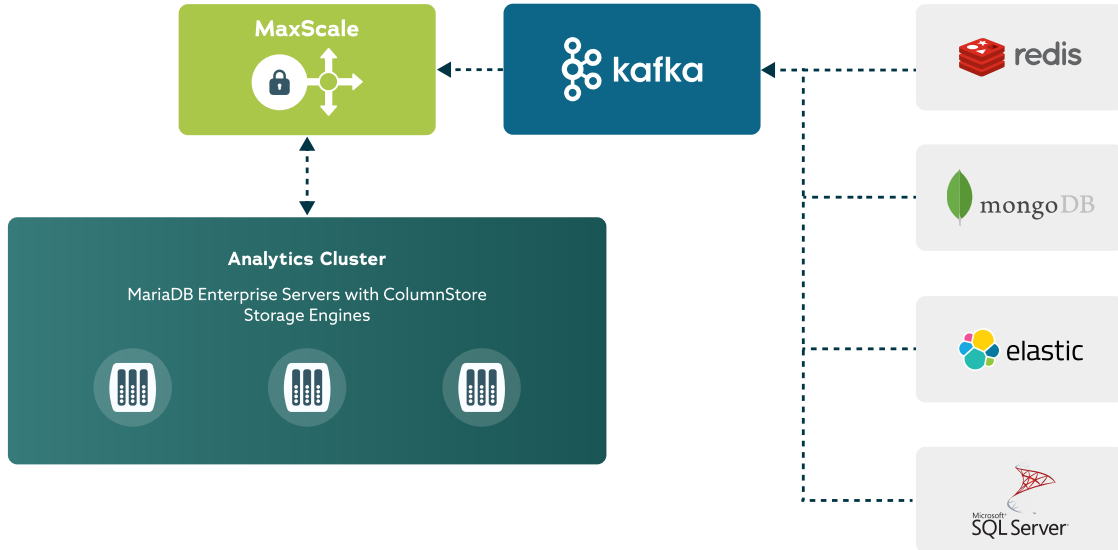
SECURE CONNECTIONS AND AUTHENTICATION

MaxScale supports SSL/TLS for secure connections both between applications and itself as well as between itself and other databases. In addition, it supports PAM and GSSAPI (i.e., Kerberos) for authentication, and it can be configured to use the Proxy Protocol to pass client information to MariaDB Server to further simplify user management.



ENTERPRISE INTEGRATION

Replicate data from MariaDB Enterprise Server to external systems such as Apache Kafka and Redis. The Kafka change data capture (CDC) service streams database changes to Kafka by reading binlog events from the primary MariaDB Enterprise Server instance, transforming them into JSON documents and publishing them to a Kafka topic. The caching service can store query results in Redis, using external memory to improve the performance of repeat queries, allowing multiple MaxScale instances to share the same cache and reducing the load on MariaDB Enterprise Server read replicas or cluster nodes.



ENTERPRISE ADMINISTRATION

MaxScale can be configured and managed either through MaxCtrl, a command-line utility, or through MaxGUI, a browser-based administrative tool. Both use a MaxScale-specific REST API, accessible through local host or remotely. Key MaxGUI features include visualization of MaxScale configuration of nodes, replication and Galera clusters, simple, rapid loading of SQL files, a built-in query editor with multiple query tabs, drag and drop replication, and log collection and archiving.

The screenshot shows the MaxGUI interface with a query editor, a results table, and a visualization. The query is as follows:

```

1 SELECT 'q'.airline,
2       'q'.volume AS 'flight_count',
3       ROUND(100 * 'q'.volume / SUM('q'.volume) OVER (
4         PARTITION BY 'q'.airline
5         ORDER BY 'q'.volume DESC) AS 'cancelled_shr',
6       ROUND(100 * (q.cancelled / q.volume), 2) AS 'cancelled_pct'
7 FROM
8   (SELECT 'a'.airline,
9         COUNT(+) AS 'volume',
10        SUM('diverted') AS 'diverted',
11        SUM('cancelled') AS 'cancelled'
12   FROM 'flights' 'f'
13   JOIN 'airlines' 'a' ON 'f'.carrier = 'a'.iata_code
14   WHERE 'f'.year = 2020
15   GROUP BY 'a'.airline) AS 'q'
16 ORDER BY 'flight_count' DESC;

```

The results table shows the following data:

#	airline	group	flight_count	group	market_share_pct	group	cancelled_pct	group	diverted_pct	group
1	Southwest Airlines		961276		19.14		8.10		0.11	
2	Skywest Airlines		597060		11.89		4.18		0.22	
3	Delta Airlines		581101		11.57		4.95		0.14	
4	American Airlines		569806		11.34		6.01		0.16	

The visualization is a horizontal bar chart showing the market share percentage for each airline. The Y-axis is labeled 'airline' and the X-axis is labeled 'market_share_pct'. The bars are ordered by flight count, with Southwest Airlines having the highest market share at approximately 19.14%.