

Edge DNS GSLB Use Cases

Improving UX, DRP and Datacenter Agility

Contents:

- Use Case #01:
Improve Disaster Recovery
- Use Case #02:
Distribute Traffic Amongst
Multiple DCs
- Use Case #03:
Wan Failure Detection
- Use Case #04:
Test Application with Multiple
Servers

As explained in our solution paper *“Edge DNS GSLB: Complement Your Load Balancing and Multi-Cloud Strategy”*, the EfficientIP Edge DNS GSLB solution brings valuable enhancements to user experience, multi-site resiliency, disaster recovery planning, and datacenter scalability and agility.

As a reminder, Edge DNS GSLB brings a simple and efficient way to load balance traffic, taking into account geographical dispersion and resource availability. It enables application traffic routing decisions to be taken from the edge of the network, close to where the clients are located.

The use cases shown in this document describe challenges faced during typical enterprise scenarios and help explain the benefits brought by implementing Edge DNS GSLB.

Edge DNS GSLB Use Cases

Use Case #01 – Improve Disaster Recovery

Some applications are vital for the company, therefore specific servers hosting them are available at the DRP site in case of major problems on the primary site (destruction, fire, flood, power outage, human error).

These applications are used by many users in the company.


The association of the technical IP address of the servers hosting the applications and their names is performed at the authoritative DNS level. If a disaster occurs in the datacenter and the applications are not available anymore, the IP associations need to be changed in order to reflect the new location on the network. Changing all the DNS records requires time and proper documentation in order to avoid errors. Testing the DRP for an application requires switching of the service for all users to the DRP site, which impacts production and needs agreement from the business unit.

Edge DNS GSLB allows multiple architecture and switchover scenarios that can be combined at the application level. This brings many benefits and operations simplification for the I&O teams which may prove to be valuable in case of disaster recovery operation.

Solution 1 - Manual Failover:

Each application is created in the SOLIDserver Application Manager inventory and linked to its FQDN with a specific pool and 2 nodes, one pointing to the primary site and the other one pointing to the backup site. Using the main and backup settings on the nodes we can manage which will be selected in the application routing process.

Since disaster recovery is a situation that needs to be agreed at a high level in the organization, the failover will be decided and applied manually by disabling the main node of the pool for all applications that need to switch to the DR site. Automatically, all recursive DNS servers will answer with the backup site IP address for this FQDN. All users will be redirected automatically to the backup site at expiration of their cache for the DNS records.



Solution 2 - Orchestrated Failover:

In addition to the solution 1 that allows preparation of the disaster recovery plan with all the addresses already in place, the inventory completed and documented, we can easily add automation in order to simplify the switch at disaster time. Using the rich set of SOLIDserver APIs, a specific automation process can be set up on an external orchestration service already available. This will automatically change the status of the nodes in the GSLB configuration. Performing this task at the GSLB level and not at the authoritative DNS level allows the orchestration service to be simplified by avoiding IP addresses having to be contained on both live and backup servers. This approach enforces the dissociation of configuration and execution, which is a very common architecture and security pattern.

Solution 3 - Automatic Failover:

For some applications, it can be interesting to propose an automatic failover to the DR site. This requires introducing server health checks from the Edge DNS GSLB nodes in order to have a good understanding of the situation on both live and DR sites. The rules for automatically switching the target server when an event occurs will be tuned to avoid inconsistencies, mainly by adjusting the hysteresis timers. When all the active nodes for a specific application inside a pool are considered as down, the GSLB will ensure transition to the backup node. The switchover back to normal is automatically performed whenever the situation at the main site is recovered.

Use Case #02 - Distribute Traffic Amongst Multiple DCs

An application is hosted in three regional datacenters for redundancy and load distribution, the users need to be routed based on their respective location to the appropriate datacenter in order to maximize their user experience. The application is accessed through a standard browser and a single URL.

How to handle the routing of the application traffic from a client-specific location towards the appropriate data center? Static geographical distribution may not provide adequate user experience as the latency implied by the network is not the only one to take into consideration with modern networks and high bandwidth.

Through multiple views, the DNS service allows manipulation of authoritative records based on who is asking for resolution. This technique requires maintaining different copies of the zone, one for each view, and associating the requester with a view based on technical information such as source IP address. Each network evolution requires reconfiguration of zones or source IP address filters. This solution is not able to be dynamic and change upon network events or server failures.

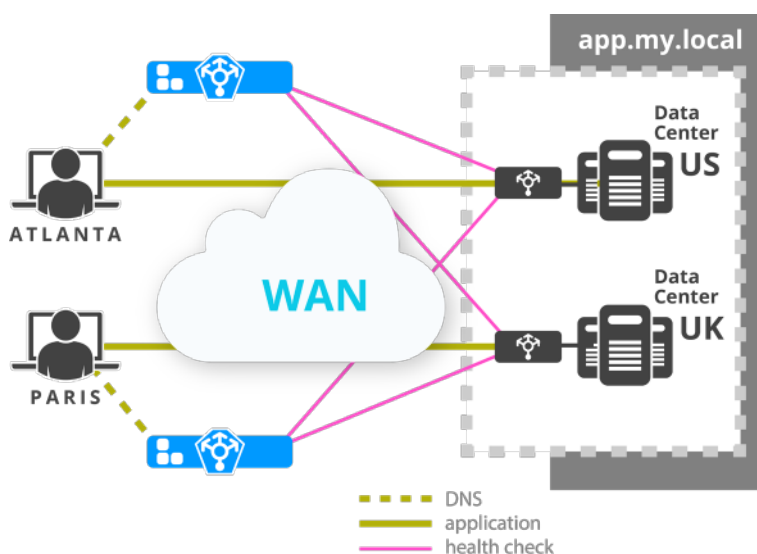
The purpose of standard GSLB is to distribute different resolution information based on the requester. It also allows automatic switching in case of server failure using server health checks and switchover scenarios. It simplifies the DNS configuration based on views but requires maintenance of a distribution policy based on the IP source address of all clients. Taking network events into account is generally not possible with GSLB based on authoritative servers.

Solution 1 - Static Distribution

With Edge DNS GSLB, the routing policy configuration can be applied to just a subset of the DNS recursive servers. Applying a different configuration for the same FQDN to different recursive servers (by their location), will allow different static traffic routing policies to be applied. If required, a backup or disaster recovery site can be added to each configuration model.

Solution 2 - Dynamic Distribution

Edge DNS GSLB can regularly check the status of the application server node in order to determine if the users can be directed to it. By comparing health checks results for all the application nodes inside the same application pool, each GSLB server can determine which one is the best for its users. Depending on the location of the Edge DNS GSLB server, results may be different, for example when using latency as a comparator the user will be routed towards the datacenter offering the best application round trip delay. The analysis can be based on network-only round trip delay, but can also be based on the application session protocol used by the clients, providing the best appreciation of the application response time. By using custom health checks it is possible to mix multiple metrics into a rich combined one that will be used to decide the application traffic routing.



Use Case #03 - Wan Failure Detection

An application is fully redundant, working in active/active mode and hosted in various regional datacenters. The users need to be directed to the application server that is the best and available for them at any time even in cases of network or WAN failure.


It is not possible to be absolutely certain from a central point of the network if a user located at a remote location can access an application hosted in a datacenter or not. The syndrome is generally known in the support team as «but it works for me».

Solution

From a central point or elsewhere on the network, you may have a different analysis and what is working at one end of the network may be broken at the other end. Issues can be related to IP routing (route disappeared), layer 2 rerouting increasing the latency (going on a backup link rather than the main / xDSL vs optical fiber), protocol optimisation system failing, firewall rule filtering out part of the traffic or SD-WAN incorrect appreciation of the application best route.

On Edge DNS GSLB, by enabling continuous health check at the edge of the network, near the user and with the same protocol as the one used by the application itself, it is possible to detect an anomaly in the network between the remote location and the application server. The GSLB does not need to know anything about the network topology, it takes the same path as the client to reach the application for its health checks and therefore can always take the correct routing decision.

In addition, events logged at the Edge DNS GSLB engine level can be used from a supervision standpoint to detect network issues early and without the need to deploy specific sensors at each location. With the increasing complexity of IP networks and optimization solutions to increase user experience, taking the exact same route as the client is a smart solution for always being at the highest level of knowledge.



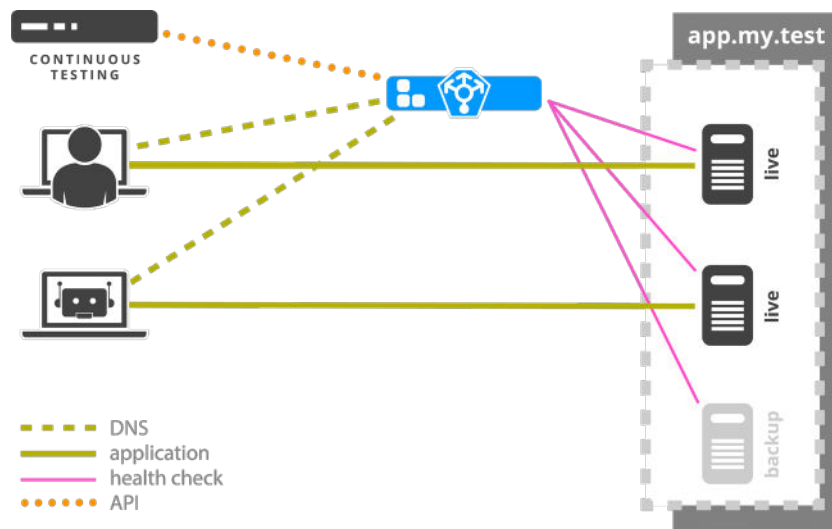
Use Case #04 - Test Application with Multiple Servers

Traditional load balancers are generally expensive components of the infrastructure, requiring lots of bandwidth since all the traffic crosses their interfaces. It may happen that for testing purposes of an application a load balancer strategy should be configured to validate how the application is behaving.

ADC are not available everywhere on the architecture, some testing environments are not covered by expensive hardware and virtual servers may not be easy to deploy, configure, or administer for development and test purposes.

Solution

Edge DNS GSLB and configuration through API can help in setting up simple strategies of traffic load-balancing without requiring specific routing through an ADC. Having the Edge DNS GSLB configuration only applied at a specific location can help constrain the test zone and thus not impact other dev or testing groups, even production. For example, the recursive server can be used by the dev team and a farm of Selenium testers that are performing automatic testing using real interaction through a browser



Conclusion

The above use cases are just a few examples of where Edge DNS GSLB can be used to complement load balancers, ADCs, SD-WAN and other network components related to distribution or routing of application traffic.

Considering the ever-increasing complexity of IP networks and the optimization solutions being introduced to enhance user experience, organizations would do well to leverage their DNS offering and consider making use of Edge DNS GSLB. The solution offers a smarter approach for controlling application traffic routing, is really simple to implement, and is most certainly efficient in use.