

DNS Global Server Load Balancing

Optimize User to Application Access From The Edge

Highlights:

- Reduce capex and opex with combined DNS/GSLB functionality on same server
- Ensure application availability thanks to constant health checking
- Guarantee disaster recovery and business continuity due to automatic failover, even in multi-cloud
- Simplify deployment of multi-region datacenters with native geolocalization
- Improve user experience by accelerating application response times
- Enforce policy control with unified management of DDI and GSLB functionality
- Gain cross-platform visibility of applications via centralized repository
- Strengthen network security by simplifying architecture

DNS service is a master piece for enabling communication between IP devices and applications. It enables technical translation of host and domain names to IP addresses, thus is at the core of any IP conversation between nodes and applications. Digital transformation is modifying the landscape with distributed applications, users are becoming more mobile, and datacenters are evolving towards hybrid clouds topology. So to ensure business continuity and meet user experience expectations in this challenging context, it's important to prioritize application availability, performance and latency.

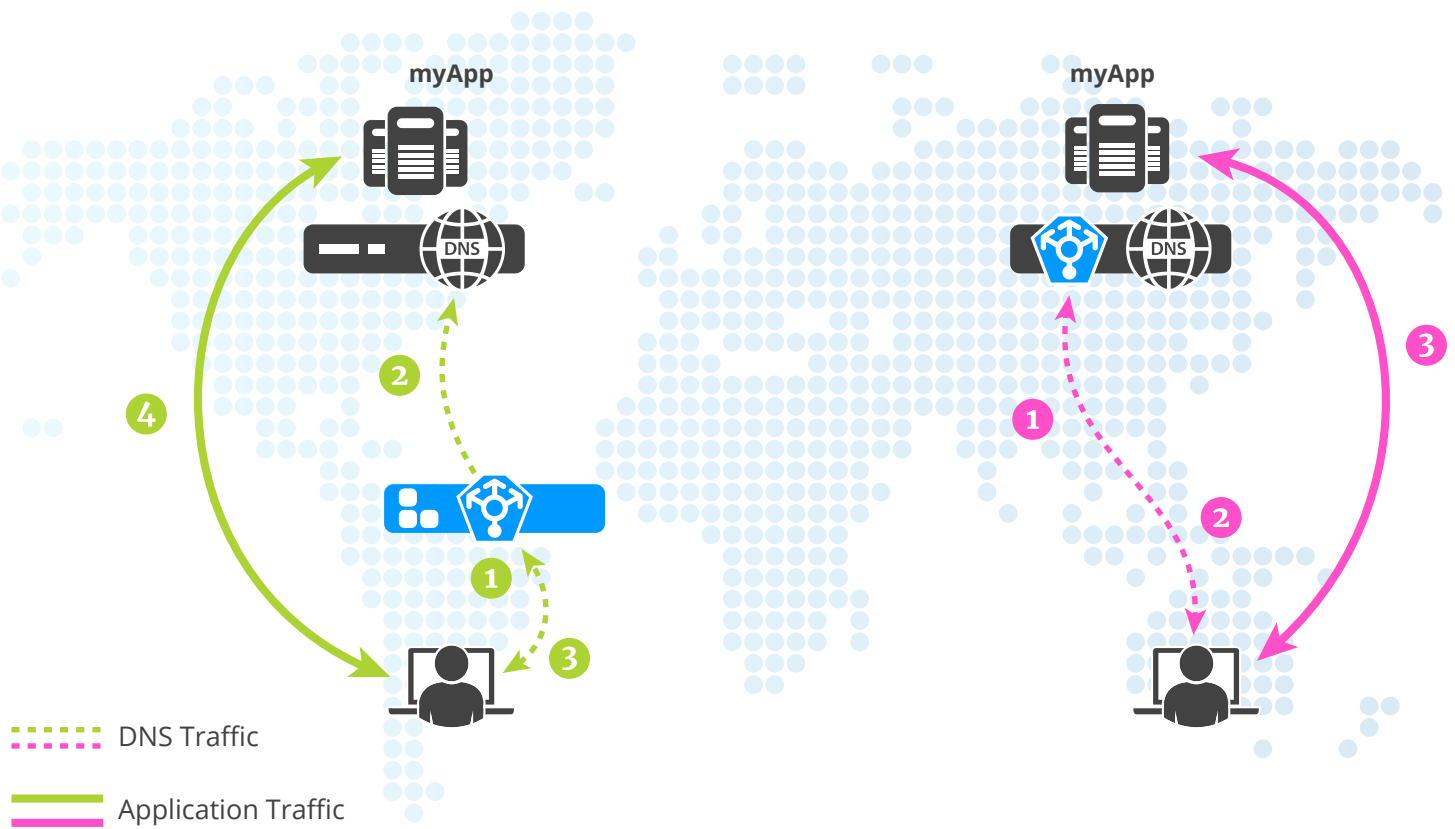
EfficientIP DNS GSLB (Global Server Load Balancing) brings a simple and efficient way to load balance traffic from the network edge, taking into account geographical dispersion and availability of resources. The solution is built into DNS servers to simplify management and reduce costs. It can be used as an alternative or a complement to traditional load balancers/application delivery controllers, being deployable on authoritative as well as recursive servers for bringing unique edge GSLB capability.

In addition, EfficientIP IPAM (IP Address Management) can be leveraged to provide unified management of IP resources and applications from the same console, bringing significant time savings for network administrators. Having a centralized "app repository" offers cross-platform visibility of your applications, helping to ensure error-free configurations which comply to corporate policies.

Offer Edge Load Balancing with GSLB Built-in to Recursive and Authoritative DNS appliances

EfficientIP's DNS GSLB is the world's first edge load balancer, providing smart, robust GSLB functionality incorporated into recursive DNS servers, as well as authoritative. This allows unique capability for application traffic routing decisions to be taken even at the network edge, hence much closer to each user. Many issues related to geolocation, breakout and peering are eliminated as a direct consequence.

Enabling GSLB with EfficientIP DNS removes the requirement to move domains towards specific load-balancing solutions such as Application Delivery Controller (ADC). All DNS domains and records remain in a single management system easing transition to load-balancing, while at the same time simplifying deployment, lowering administration cost and avoiding administration dispersion in multiple teams.



- | | | |
|----------|---|----------|
| 1 | DNS request for application IP resolution to recursive GSLB server | |
| 2 | If DNS entry not in cache, request to authoritative DNS server | |
| | DNS request for application IP resolution to authoritative GSLB server | 1 |
| 3 | GSLB answers the request with the most suitable application server IP address | 2 |
| 4 | Application data flow to selected server IP address | 3 |

Recursive DNS Appliances with Built-in GSLB bring load balancing capability to the network edge

Ensure App Availability and Accelerate Response Time

The heart of DNS GSLB is about load balancing traffic based on DNS queries from applications and users. Each application exposing a specific DNS entry (FQDN) automatically turns on the GSLB feature for this entry. Any DNS server with GSLB feature enabled will compute the answer to a resolution request for each client with the objective of directing the traffic according to the load balancing strategy.

Pools and Nodes

For distributing user traffic amongst multiple servers, the servers (known as nodes) are grouped into pools. Every pool supports the load balancing strategy. Any node can be defined as being «unmanaged» - it will be contained in the pool but not included in traffic distribution - this is useful when preparing a future operation or performing planned maintenance on the node. Changing the managed status of a node could be performed through GUI by the administrator and directly by the application through APIs.

Flexible Load Balancing Strategies

A load balancing strategy is defined for each application's pool of back-end servers. Choosing the appropriate strategy depends on the way the application is providing service to its users. Each pool implements one of the 3 available strategies for load-balancing. Every DNS request on the application FQDN will follow the strategy rules.

- **Round-Robin:** distribution uses a circular mode on all the available back-end nodes in the pool.
- **Weighted Round-Robin:** this mode allows a weighted distribution on all the available nodes. This can be seen as an evolution of the round-robin mode where all nodes share the same weight.
- **Best Response Time:** the response time of each node in the pool is used to select a subset of the best ones to which the traffic is distributed. Best response time strategy can be based on advanced metrics like connection number, CPU load, or users connected.

Any node included in a pool can be tagged as a backup node. It will only participate in distribution if all active nodes are unavailable or unmanaged. This allows easy handling of primary/backup architecture for applications unable to handle session sharing or data center resiliency.

Smart Application Health Checks

By enabling health checking, the load balancing strategies are enhanced with the current health of each node supporting the application. Any node failing its health check will be removed from the selection process.

When GSLB is deployed at the edge of the network on recursive DNS servers close to the users, health checking ensures accurate evaluation of application response time to select the best suitable server for each user.

The available health check types are:

- **None:** no health check is set for this node, it will always be selected in the distribution process if managed and in active mode. This type has precedence over any other health check mode, even the OK mode.
- **OK:** the node is inserted in the health-check process, announced as always available and with a latency of 0ms. In latency load-balancing strategy, it will be preferred over any other node with a higher latency.
- **Ping:** a connectivity check is performed from the GSLB server with ICMP ECHO requests.
- **TCP:** a connectivity check is performed using TCP session establish.
- **HTTP(S):** a page is gathered from the node using HTTP and optionally using TLS (HTTPS). In addition to connectivity and response time, this health check could request a specific page, check the HTTP status code and verify the presence of specific content in the response body.
- **Custom:** can be used for specific check on proprietary application and advanced application health checking.

Session Persistence

Even if serverless is used in modern development architecture, most applications require the user to use the same server during a session. In order to enable this behavior within the dynamic traffic distribution authorized with GSLB, a session persistence option can be activated on each pool, regardless of the load balancing strategy. This allows each client to be redirected to the same server during its entire session. It also enables all the clients to be spread across all nodes based on the strategy in place.

This powerful feature does not require any modification at application level and is totally transparent for the user, making it easy to install and use on all environments including development, test and qualification.

Decision Matrix

	Round-Robin	Weighted Round-Robin	Best Response Time
Managed status	Y/N	Y/N	Y/N
Active mode	Active or Backup		
Node eligible subset	Within all nodes in managed status 1- all nodes in active mode 2- if no active mode, one backup node		Within all nodes in managed status, uses max preferred node limit nodes with the lowest latency returned from health check
Health check	All		
Weight	All nodes have a weight of 1	Specified by node	

Unify Management Using Application Repository

EfficientIP’s SOLIDserver IPAM solution already uses the repository of IP information such as addresses, subnets, devices, DNS records and DHCP pools. It is considered as the IP network source of truth. The GSLB solution brings a new dimension with the application repository. Each application entry is linked to a fully qualified name that will be used in DNS requests to load-balance the traffic.

Thanks to the advance Class feature of SOLIDserver IPAM, it is possible to enhance the Application object in the database with specific meta-data in order to fit the way each enterprise manages its applications. For example, adding IT organizational information like zone or district, or adding a reference to another repository such as the CMDB brings significant value for the IT ecosystem. The application repository is fully manageable through SOLIDserver API set for automatic updates.

Strengthen Network Security By Simplifying Architecture

Implementing DNS GSLB strengthens network security as it simplifies architecture by removing the need for separate DNS servers and ADCs (Application Delivery Controllers). Additionally, when deployed at the edge of the network, GSLB improves resiliency and user experience, specifically in a multi-region/multi-cloud environment. Incorporating EfficientIP DNS Security features then adds best protection against DDoS and other DNS attack types to ensure service continuity, protect against malware and safeguard data.

Improve application deployment velocity in distributed environment

In the context of multi-cloud strategy used today by most enterprises, and development approaches becoming more agile using DevOps and automation, EfficientIP GSLB enables users to access enterprise applications regardless of their location.

With application and GSLB module fully configurable through REST API, infrastructure as code opens the door to application auto-configuration. In a dynamic server infrastructure where elastic resources are used, GSLB configuration is automatically updated to add and suppress nodes from a pool in a seamless fashion. At scale up, new application nodes are automatically available to users. At scale down, applications can disable the node from the application pool allowing the session to terminate and therefore shut down.



REV: C-190515

As one of the world’s fastest growing DDI vendors, EfficientIP helps organizations drive business efficiency through agile, secure and reliable network infrastructures. Our unified management framework for DNS-DHCP-IPAM (DDI) and network configurations ensures end-to-end visibility, consistency control and advanced automation. Additionally, our unique 360° DNS security solution protects data confidentiality and application access from anywhere at any time. Companies rely on us to help control the risks and reduce the complexity of challenges they face with modern key IT initiatives such as cloud applications, virtualization, and mobility. Institutions across a variety of industries and government sectors worldwide rely on our offerings to assure business continuity, reduce operating costs and increase the management efficiency of their network and security teams. Copyright © 2022 EfficientIP, SAS. All rights reserved. EfficientIP and SOLIDserver logo are trademarks or registered trademarks of EfficientIP SAS. All registered trademarks are property of their respective owners. EfficientIP assumes no responsibility for any inaccuracies in this document or for any obligation to update information in this document.