

Improving Application Access Control

with DNS Client Query Filtering (CQF)

Contents

- Overview
- Network Segmentation Down to Individual User
- The Filtering Process in Client Query Filtering (CQF)
- Key Benefits
- Rich Client Identification in CQF
- Easy and Scalable CQF List Management
- Conclusion

Access control to applications can be performed at multiple levels in accordance with the security policies in place within the organization. For most, the main level in place nowadays is Authentication and Authorization at the application level through credentials - meaning that normally no application is accessible without user screening.

But is that really enough? Can a user with no access to an application get access to the login page? If self registration is not an option for this application, which is mainly the case in organizations, then why expose access to its infrastructure from the network?

There are some very important applications that require specific access and run on a dedicated infrastructure with no sharing of main components. Filtering at the network level is an option to consider, whereby routing access lists and firewall rules are an implicit solution. However, by adding filtering at the DNS level, you raise the security level even higher. This leaves no possibility to resolve the application technical IP addresses, no network level and no credentials, so is a far better approach to security in a Zero Trust environment.

EfficientIP brings this granular network segmentation functionality with its DNS Client Query Filtering (CQF) feature. By having the ability to dynamically update the CQF lists with either application or client entries, security is automatically raised to the appropriate level, limiting the application's exposure and data visibility to unknown or non authorized users.

Network Segmentation Down to Individual User

Filtering at the DNS level with a DNS Firewall solution is mainly utilized for offering a first line of defense to any kind of user, device and application. Based on either a manually or dynamically managed list of destination domains to either allow or deny access to, it provides a very quick and easy way to protect client software when connecting to applications or services, based on their name or IP address.


However, DNS Filtering needs to be more specific. Applying the same filter to any device or user on the network tends to lower the security level to the most common rules to be applied. But some user groups require a higher level of security, while others need a lower level in order to perform their activities effectively. In addition, IoT devices would benefit from having access to only authorized applications and resources, based on an allow list (whitelist). Most of the other devices generally have some restrictions applied, based on a deny list (blacklist). But how can you perform such DNS filtering without the ability to differentiate the important part of the transaction: the client itself?

Options in the DNS engines do already exist, such as views that may help distinguish the client based on its IP address, but is that enough? Is it convenient? Are you sure you can still use this approach for devices roaming from site to site, with new devices on the network installed every day, perhaps directly by business entities?

Ideally, what is required is more granularity and a better way to classify clients, in order to apply the appropriate level of filtering and therefore the adapted level of security. This is the path to application filtering, to segmentation required by zero trust approach, to parental control for telcos, and also to governmental and regulatory filtering.

The Filtering Process in Client Query Filtering (CQF)

CQF brings a new facet to DNS filtering, with security based on the source client information mapped to the requested domain, rather than filtering based only on the domain. A specific filtering policy can therefore be applied just to specific clients requesting access to specific applications. This brings DNS security to a higher level, by combining client and destination information with allow and deny lists, therefore enabling application security enhancement.



The main components required by CQF in order to perform rich DNS filtering are: 1) a list of client identifiers 2) a list of domains to analyze and 3) the operation to perform, either allow, deny or apply countermeasure.

Each DNS request is compared with the content of the list of applications and domains for applying the relevant policy. The lists are either local to each Guardian DNS server and managed manually - which is useful for testing purposes - or centralized and managed globally, which is ideal for security policy enforcement. The domain list is a standard RPZ zone that can be maintained in the SOLIDserver through GUI actions and API calls, but can also be subscribed to from a threat intelligence provider. Distribution of each list to all the Guardian DNS servers is performed through standard replication mechanisms, scalable and in real time allowing automation scenarios with the security ecosystem and with OSS/BSS solutions.

The filtering process is the heart of the CQF feature and enables rich security usages. Having the ability to use and manipulate large amounts of information in the lists provides a real advantage when it comes to applying security to multiple groups of clients which are complex to identify. This management is made possible by the high performances of the Guardian DNS engine and its integration in the whole DDI ecosystem of the SOLIDserver.

Key Benefits

1

More Granular Filtering

Improve network segmentation down to the individual client

2

Better Application Access Control

Enable DNS-based client access control to vital apps and infrastructure

3

Early Security Barrier

Detect anomalies at the earliest point in the flow to reduce exposure risk

4

New Business Opportunities

Enable new B2B2C offers (e.g. parental control for telcos)

5

Stronger Security Ecosystem

Allowing immediate modifications through API and standard DNS zone manipulations

Rich Client Identification in CQF

DNS Clients are commonly identified by their IP address, but in some more complex scenarios, another field or a combination of fields extracted from each query can be used for this identification. For example, we can use the extended DNS Client Subnet field to identify either client groups located on the same subnet or each individual when used with a full subnet mask. We can also use a combination of the CPE (Customer Premise Equipment) identification on the telco network and the mac address of the device on the consumer network as a unique identification key. This variety of identification methods enables CQF to be used in conjunction with cascaded DNS servers, with DNS over HTTPS external engines or with ISP DNS relay embedded in the CPE. It can therefore be used by an organization with a local network but also by a more complex telecommunications network or service provider.

Easy and Scalable CQF List Management

Filtering at the DNS level is only viable if lists of domains and client identifiers are easy to manage and can scale with almost no limit. The DNS firewall solution uses the standard RPZ zone transfer for its list management, this enables loading large volumes of filtered domains with the help of the DNS zone transfer which is really efficient. This principle has been extended for the use of client query filtering: both domain and client identifier lists can be managed either locally or by using the DNS zone transfer mechanism.

Content of both lists can therefore be managed in a DNS server for centralized and automatic update of the content on the Guardian engines, which provides an efficient way to adapt security measures to the global policies. By using the SOLIDserver DNS service as the backend for managing the list content, you can leverage all the available and well known solutions proposed including the web GUI and the rich API set. The whole ecosystem of security and automation can then add and remove records from the security DNS zones which are pushed to the Guardian engines and immediately applied to the incoming traffic. Security is therefore enhanced for all the valuable solutions within the security and IT ecosystem.

Conclusion

SOLIDserver DNS's granular visibility over internet traffic offers precious contextual information for behavioral threat detection. Combining this with enhanced DNS security functions and smart DNS filtering creates a security barrier early in the traffic flow to protect the critical infrastructure, apps and services of organizations, and enable them to put their zero trust strategies in place.



REV: C-210504

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 © 2021 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.