

Design

- [F5 Container Ingress Service](#)
- [F5 APM Microsoft Exchange 2016](#)
- [F5 APM SSO Infos](#)
- [Tuning the OneConnect Profile](#)
- [Kerberos Delegation & Protocol Transition](#)
- [BigIP DNS \(Formerly GTM\)](#)

F5 Container Ingress Service

Installation

Installation über Helm Chart mit Beispiel values.yaml

<https://clouddocs.f5.com/containers/latest/userguide/cis-installation.html>

Manual Installation

<https://clouddocs.f5.com/containers/latest/userguide/kubernetes/>

Deployment Options

Via NodePort oder ClusterIP

<https://clouddocs.f5.com/containers/latest/userguide/config-options.html>

NodePort

Similar to Docker, BIG-IP communicates with an ephemeral port, but in this case the kube-proxy keeps track of the backend Pod (container). This works well, but the downside is that you have an additional layer of load balancing with the kube-proxy.

image-1641973699480.png

ClusterIP

The BIG-IP CIS also supports a cluster mode where Ingress traffic bypasses the Kube-proxy and routes traffic directly to the pod. **This requires that the BIG-IP have the ability to route to the pod.** This could be by using an overlay network that F5 supports (Flannel VXLAN, or OpenShift VXLAN). Leave the kube-proxy intact (no changes to underlying Kubernetes infrastructure).

image-1641973600651.png

F5 APM Microsoft Exchange 2016

Microsoft Exchange specific config options for integration with F5 APM

OWA 2010

Integrate Microsoft OWA and APM seems to be easy but the devil has to be found in the details.

Example: The customer want to give the options like the light version or public- or private-computer to the users like here in OWA 2010:

OWA Options.png

When you log in, OWA offers the following options:

- **This is a public or shared computer:** Allows you to be logged in for 1 hour before you are automatically logged out.
- **This is a private computer:** Allows you to be logged in for 24 hours you are automatically logged out.
- **Use the light version of Outlook Web App:** Removes some of the features of OWA.

This settings are described here: <https://docs.microsoft.com/en-us/powershell/module/exchange/client-access-servers/set-owavirtualdirectory>

OWA 2016

In OWA 2016 the logon screen look like the following:

OWA 2016 logon options.png

Config Options

Change the OWA logon options with the following settings:

```
Set-OwaVirtualDirectory -LogonPageLightSelectionEnabled <$true | $false>
```

```
Set-OwaVirtualDirectory -LogonPagePublicPrivateSelectionEnabled <$true | $false>
```

Direct URL for using the different OWA modes

Documentation

[Using Outlook Web App Web Parts](#)

[OWA Layout choice](#)

[How to compose a New Message or Event and Populate fields in OWA](#)

Layout Mode

Type one of the following parameters behind the /owa/ part to change the layout.

?layout=tnarrow	Single column layout optimized for small screens or when holding the tablet in portrait mode. This is the standard for recognized and supported smartphones.
?layout=twide	Multi column layout optimized for larger touch screens or when holding the tablet in landscape mode. This is the standard for recognized and supported tablets.
?layout=tmouse	Default layout optimized for mouse usage. This is the standard on normal desktops and devices with a supported browser. When the browser isn't supported, OWA Light will be loaded instead.
?layout=light	Light Layout for low bandwidth usage

Links

<https://blogs.technet.microsoft.com/ptsblog/2013/10/21/url-for-office-365-outlook-web-app-light/>

<https://hochwald.net/enable-adfs-authentication-on-exchange-2016/>

<https://hochwald.net/adfs-authentication-with-exchange-troubleshooting/>

<https://blogs.technet.microsoft.com/exchange/2017/12/06/announcing-hybrid-modern-authentication-for-exchange-on-premises/>

<https://asichel.de/2017/06/14/adfs-4-0-mit-exchange-2016-konfigurationsuebersicht/>

<https://proofid.com/blog/single-sign-on-to-outlook-web-access-using-pingfederate/>

<https://www.citrix.com/blogs/2014/03/31/owa-2010-login-options-on-aaa-login/>

<https://serverfault.com/questions/162433/single-signon-options-for-exchange-2010>

F5 APM SSO Infos

NTLM and APM

- [Microsoft Pass-Through Authentication](#)
- [Configuring APM client side NTLM Authentication](#)
- [F5 APM NTLM, Basic and SAML Seamless](#)
- [Use NTLM to bypass f5 APM login page](#)
- [APM Troubleshooting with ADTest](#)
- [\[APM VPE VarAssign UPN or Logonname Auth\]\(APM VarAssign UPN or Logonname\)](#)

Kerberos and APM

- [Basic Auth and Kerberos logon](#)
- [Kerberos servicePrincipalName \(SPN\)](#)
- [\[F5 Devcentral Kerberos is easy\]\[7\]](#)
- [\[Kerberos Survival Guide\]\[8\]](#)
- [\[Setting up Kerberos Authentication for a Website in IIS\]\[10\]](#)
- [\[Kerberos SSO across External trust \(KRB Constrained Troubleshooting\)\]\[11\]](#)
- [\[Kerberos Constraint Delegation White Paper\]\[12\]](#)
- [\[Microsoft: How the Kerberos Version 5 Authentication Protocol Works\]\[13\]](#)
- [\[Digicomp: Troubleshooting Kerberos Authentisierung\]\[14\]](#)
- [\[Microsoft: Configuring Kerberos authentication for load-balanced Client Access services\]\[15\]](#)

image-1614418749649.png

image-1614418755739.png

Links

[1]: <https://msdn.microsoft.com/en-us/library/cc237015.aspx>

[2]: <https://devcentral.f5.com/articles/configuring-apm-client-side-ntlm-authentication>

[3]: <https://devcentral.f5.com/articles/leveraging-big-ip-apm-for-seamless-client-ntlm-authentication>

[4]: <https://devcentral.f5.com/questions/how-to-use-ntlm-to-basically-bypass-the-f5-apm-login-page>

- [5]: https://support.f5.com/kb/en-us/products/big-ip_apm/manuals/product/apm-authentication-single-sign-on-11-5-0/9.html
- [6]: <https://www.msxfaq.de/windows/kerberos/kerberosspn.htm>
- [7]: <https://devcentral.f5.com/articles/kerberos-is-easy-part-1>
- [8]: <https://social.technet.microsoft.com/wiki/contents/articles/4209.kerberos-survival-guide.aspx>
- [9]: <https://f5guru.com/2015/07/02/apm-troubleshooting-with-adtest/>
- [10]: <https://blogs.msdn.microsoft.com/chiranth/2014/04/17/setting-up-kerberos-authentication-for-a-website-in-iis/>
- [11]: <https://devcentral.f5.com/questions/kerberos-sso-across-external-trust>
- [12]: <https://f5.com/Portals/1/Cache/Pdfs/2421/kerberos-constrained-delegation-and-protocol-transition-in-smart-card-pki-architecture-.pdf>
- [13]: [https://technet.microsoft.com/en-us/library/cc772815\(Ws.10\).aspx](https://technet.microsoft.com/en-us/library/cc772815(Ws.10).aspx)
- [14]: <https://news.digicomp.ch/de/2013/11/20/troubleshooting-kerberos-authentisierung/>
- [15]: <https://technet.microsoft.com/en-us/library/ff808312.aspx>

Tuning the OneConnect Profile

Description

OneConnect™ is a feature of the BIG-IP LTM system that improves web application performance and decreases server load by reducing the concurrent connections and connection rate on back-end servers.

[oneconnect-tuning-dg.pdf](#)

Kerberos Delegation & Protocol Transition

A very good video from Lightboard Lessons from F5:

<https://www.youtube.com/embed/UGWP4ewxcTA>

Link: <https://devcentral.f5.com/articles/lightboard-lessons-kerberos-delegation-protocol-transition-32686>

BigIP DNS (Formerly GTM)

Preface

The following is from the BigIP DNS Documentation: <https://support.f5.com/kb/en-us/products/big-ip-dns/manuals/product/bigip-dns-load-balancing-13-1-0/>

Introducing BIG-IP DNS

BIG-IP® DNS (formerly GTM™) is a system that monitors the availability and performance of global resources and uses that information to manage network traffic patterns. BIG-IP DNS uses load balancing algorithms, topology-based routing, and iRules® to control and distribute traffic according to specific policies.

About global server load balancing

BIG-IP® DNS provides tiered global server load balancing (GSLB). BIG-IP DNS distributes DNS name resolution requests, first to the best available pool in a wide IP, and then to the best available virtual server within that pool. BIG-IP DNS selects the best available resource using either a static or a dynamic load balancing method. Using a static load balancing method, BIG-IP DNS selects a resource based on a pre-defined pattern. Using a dynamic load balancing method, BIG-IP DNS selects a resource based on current performance metrics collected by the big3d agents running in each data center.

Static load balancing methods

This table describes the static load balancing methods available in BIG-IP® DNS.

Name	Description	Recommended Use	Wide IP Load Balancing	Preferred Method	Alternate Method	Fallback Method

Drop Packet	BIG-IP DNS drops the DNS request.	Use Drop Packet for the Alternate load balancing method when you want to ensure that BIG-IP DNS does not offer in a response a virtual server that is potentially unavailable.	No	Yes	Yes	Yes
Fallback IP	BIG-IP DNS distributes DNS name resolution requests to a virtual server that you specify. This virtual server is not monitored for availability.	Use Fallback IP for the fallback load balancing method when you want BIG-IP DNS to return a disaster recovery site when the preferred and alternate load balancing methods do not return an available virtual server.	No	No	No	Yes
Global Availability	BIG-IP DNS distributes DNS name resolution requests to the first available virtual server in a pool. BIG-IP DNS starts at the top of a manually configured list of virtual servers and sends requests to the first available virtual server in the list. Only when the virtual server becomes unavailable does BIG-IP DNS send requests to the next virtual server in the list. Over time, the first virtual server in the list receives the most requests and the last virtual server in the list receives the least requests.	Use Global Availability when you have specific virtual servers that you want to handle most of the requests.	Yes	Yes	Yes	Yes
None	BIG-IP DNS distributes DNS name resolution requests skipping either the next available pool in a multiple pool configuration or the current load balancing method. If all pools are unavailable, BIG-IP DNS returns an aggregate of the IP addresses of all the virtual servers in the pool using BIND.	Use None for the alternate and fallback methods when you want to limit each pool to a single load balancing method. If the preferred load balancing method fails, BIG-IP DNS offers the next pool in a load balancing response.	No	No	Yes	Yes
Ratio	BIG-IP DNS distributes DNS name resolution requests among the virtual servers in a pool or among pools in a multiple pool configuration using weighted round robin, a load balancing pattern in which requests are distributed among several resources based on a priority level or weight assigned to each resource.	Use Ratio when you want to send twice as many connections to a fast server and half as many connections to a slow server.	Yes	Yes	Yes	Yes

Return to DNS	BIG-IP DNS immediately distributes DNS name resolution requests to an LDNS for resolution.	Use Return to DNS when you want to temporarily remove a pool from service. You can also use Return to DNS when you want to limit a pool in a single pool configuration to only one or two load balancing attempts.	No	Yes	Yes	Yes
Round Robin	BIG-IP DNS distributes DNS name resolution requests in a circular and sequential pattern among the virtual servers in a pool. Over time each virtual server receives an equal number of requests.	Use Round Robin when you want to distribute requests equally among all virtual servers in a pool.	Yes	Yes	Yes	Yes
Static Persist	BIG-IP DNS distributes DNS name resolution requests to the first available virtual server in a pool using the persist mask with the source IP address of the LDNS and a hash algorithm to determine the order of the virtual servers in the list. This hash algorithm orders the virtual servers in the list differently for each LDNS that is passing traffic to the system taking into account the specified CIDR of the LDNS. Each LDNS (and thus each client) generally resolves to the same virtual server; however, when the selected virtual server becomes unavailable, BIG-IP DNS sends requests to another virtual server until the original virtual server becomes available. Then BIG-IP DNS again resolves requests to that virtual server.	Use Static Persist when you want requests from a specific LDNS to resolve to a specific virtual server.	No	Yes	Yes	Yes
Topology	BIG-IP DNS distributes DNS name resolution requests using proximity-based load balancing. BIG-IP DNS determines the proximity of the resource by comparing location information derived from the DNS message to the topology records in a topology statement you have configured.	Use Topology when you want to send requests from a client in a particular geographic region to a data center or server located in that region.	Yes	Yes	Yes	Yes

Dynamic load balancing methods

This table describes the dynamic load balancing methods available in BIG-IP® DNS.

Name	Description	Wide IP load balancing	Preferred method	Alternate method	Fallback method
Completion Rate	BIG-IP® DNS distributes DNS name resolution requests to the virtual server that currently maintains the least number of dropped or timed-out packets during a transaction between a data center and the client's LDNS.	No	Yes	No	Yes
CPU	BIG-IP DNS distributes DNS name resolution requests to the virtual server that currently has the most CPU processing time available.	No	Yes	No	Yes
Hops	BIG-IP DNS distributes DNS name resolution requests to a virtual server in the data center that has the fewest router hops from the client's LDNS. BIG-IP DNS uses the traceroute utility to track the number of router hops between a client's LDNS and each data center.	No	Yes	No	Yes
Kilobytes/Second	BIG-IP DNS distributes DNS name resolution requests to the virtual server that is currently processing the fewest number of kilobytes per second. Use Kilobytes/Second only with virtual servers for which BIG-IP DNS can collect the kilobytes per second metric.	No	Yes	No	Yes
Least Connections	BIG-IP DNS distributes DNS name resolution requests to virtual servers on BIG-IP® Local Traffic Manager™ (LTM®) that currently hosts the fewest connections. Use Least Connections only with LTM servers.	No	Yes	No	Yes
Packet Rate	BIG-IP DNS distributes DNS name resolution requests to the virtual server that is currently processing the fewest number of packets per second.	No	Yes	Yes	Yes
Quality of Service	BIG-IP DNS distributes DNS name resolution requests to virtual servers based on a score assigned to each virtual server that is calculated from current performance metrics. Use Quality of Service only when you have configured BIG-IP DNS to calculate an overall score for each virtual server based on performance metrics.	No	Yes	No	Yes
Round Trip Time	BIG-IP DNS distributes DNS name resolution requests to the virtual server with the fastest measured round trip time between a data center and a client's LDNS.	No	Yes	No	Yes
Virtual Server Score	BIG-IP DNS distributes DNS name resolution requests to virtual servers on LTM based on a user-defined ranking. Use Virtual Server Score only with LTM systems on which you have assigned scores to each virtual server.	No	Yes	Yes	Yes
Virtual Server Capacity	BIG-IP DNS distributes DNS name resolution requests to virtual servers in a list that are weighted by the number of available virtual servers in the pool. Use Virtual Server Capacity for load balancing virtual servers managed by LTM Systems. BIG-IP DNS selects a virtual server that has the most available (UP) members. When selecting a virtual server from a wide IP pool and two or more virtual servers result in equal scores, BIG-IP DNS will return one of the equal scored virtual servers randomly.	No	Yes	Yes	Yes