PARTNER USE CASE

Configuring Reverse Proxy Access to Microsoft Lync Using F5 BIG-IP Local Traffic Manager (LTM)

Best Practices

Ryan Korock, Michael Shimkus, and James Hendergart | F5 Networks
Stephane Taine, Yves Pitsch, and Rick Kingslan | Microsoft Corporation

Lync 2013

Contents

Introduction

Many Microsoft Lync Server deployments make use of what has been referred to as the "reverse proxy." It provides corporate users who are outside the office (not connected by VPN) access to all of the Lync functionality that a corporate user would have if operating inside the local area network. The reverse proxy also is required if you plan on supporting Lync Mobility Service.

Lync Edge Servers provide remote users with IM, voice, and other services. However, without a reverse proxy, the users will miss out on all of the functionality provided by Lync Web Services. That includes access to meeting content and the Lync Web Access client. For the complete list of features that are enabled by the reverse proxy, see the Microsoft TechNet articles "Setting Up Reverse Proxy Servers[1]" and "Components Required for External User Access.[2]" For these reasons, Lync reverse proxy should be considered critical to every enterprise or multi-tenant deployment of Lync.

Traditionally, organizations have used Microsoft Forefront Threat Management Gateway (TMG) to act as the Lync reverse proxy. Now that TMG end of sale has been announced, companies are looking for alternatives, and BIG-IP® Local Traffic Manager™ (LTM) from F5 can be that alternative. For most organizations, BIG-IP LTM currently load balances Lync Edge Servers, enabling them to deploy a reverse proxy for Lync without incremental capital expense.

Architecture

As an illustration, take a look at the "before" and "after" architecture. In Figure 1, you can see the external BIG-IP LTM device load balancing the Edge Servers in the network perimeter (also called the DMZ), and TMG acting as the reverse proxy.

1 "Setting Up Reverse Proxy Servers." Microsoft Corporation. **http://technet.microsoft.com/en-us/library/gg398069.aspx**

2 "Components Required for External User Access." **Microsoft Corporation. http://technet.microsoft.com/en-us/library/gg425779.aspx**

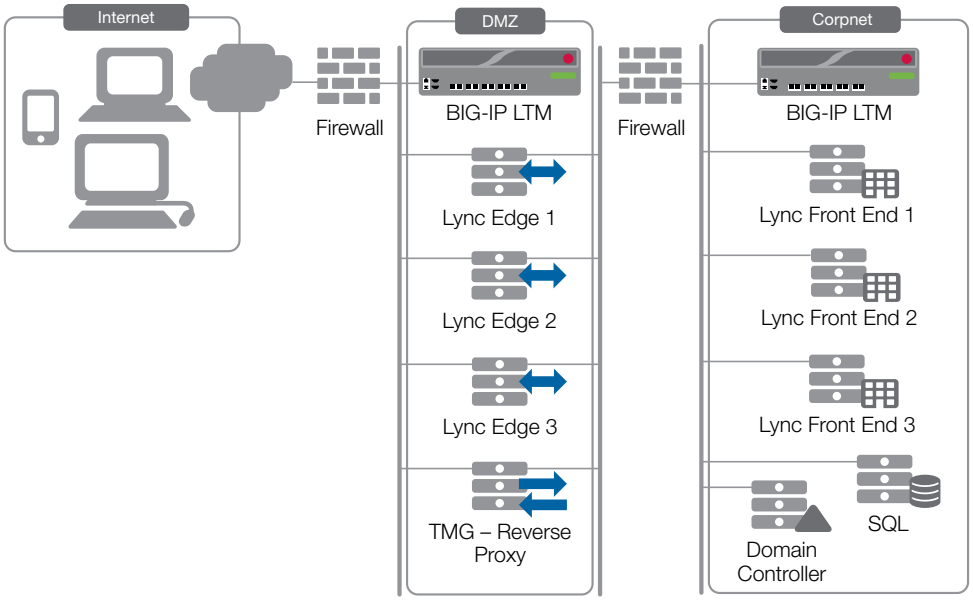

Figure 1: TMG reverse proxy.

Now in the next picture, TMG is removed, and the reverse proxy functionality is moved to the external BIG-IP LTM device.

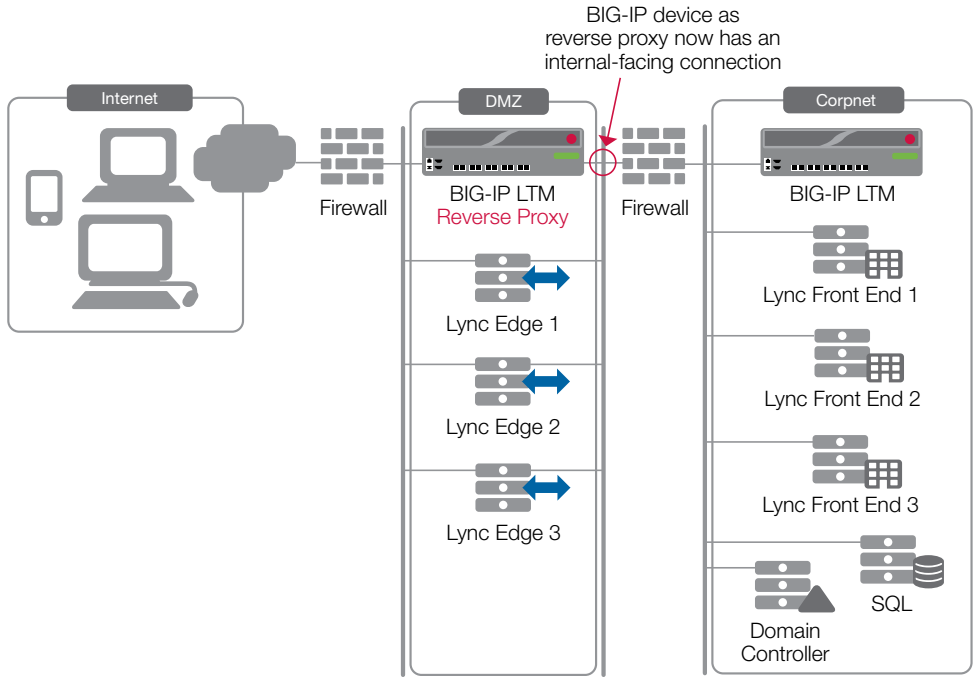

Figure 2: BIG-IP Local Traffic Manager reverse proxy.

Configuration

Use the **F5 Deployment Guide for Microsoft Lync** as the foundation for deployment and customize it based on how you answer these three questions for a given Lync reverse proxy configuration:

1. Traditional enterprise or multi-tenant deployment?
2. One-tier or two-tier approach?
3. Automated or manual configuration?

Traditional Enterprise vs. Multi-Tenant Deployment

Traditional enterprise deployments will have a single "instance" of Lync deployed, and the enterprise is the only tenant. Client traffic originates from the internal network or externally. For enterprises, the perimeter network protects the enterprise from connections originating externally.

Multi-tenant deployments, by design, host multiple tenants simultaneously. The semantics of what is external versus internal are unique in a multi-tenant deployment that is hosted by a third-party service provider because in essence, every client for every tenant is "external." Even though the basic Lync/Lync Edge topology fits both types, all client traffic originates from external networks and will enter through the reverse proxy or the Lync Edge Servers.

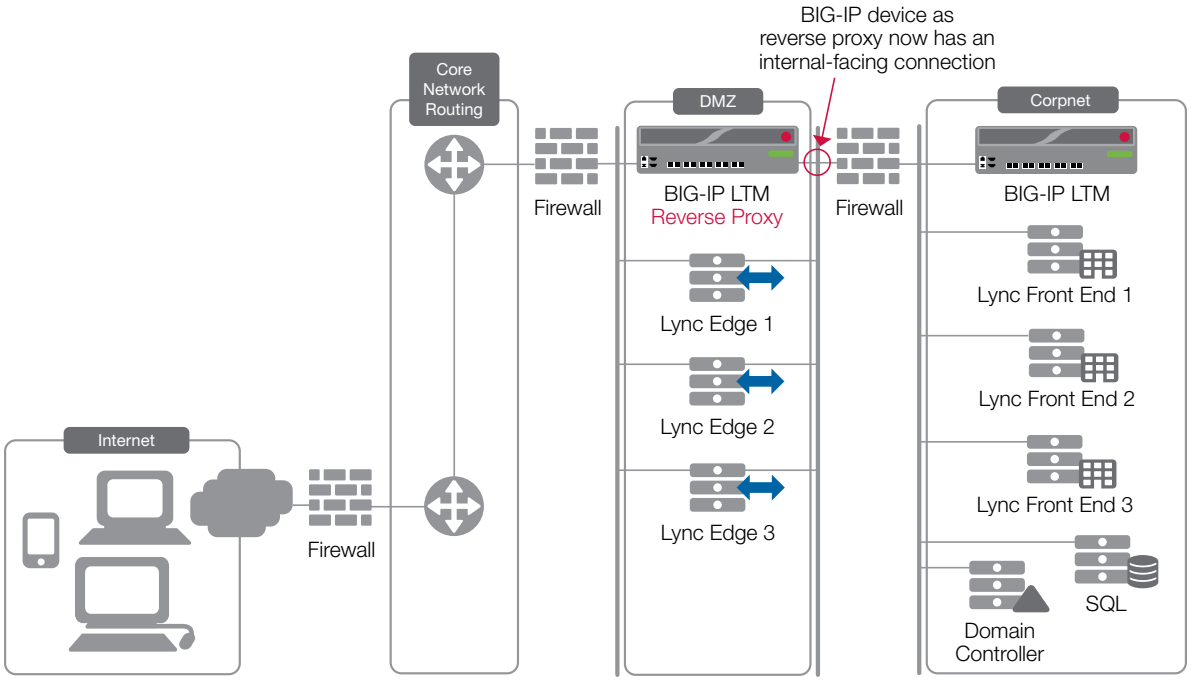

Figure 3: Placement of reverse proxy devices with respect to multi-tenant network zones.

One-Tier vs. Two-Tier Approach

Since the Lync deployment spans both the perimeter network and the internal network, organizations must decide whether or not to deploy BIG-IP® devices in one tier or two tiers.

A tier is defined as a high availability pair of BIG-IP devices. The two-tier design includes a pair of devices at the perimeter and another in the internal network. Physical device separation is often a security requirement and clearly demarks traffic by network zone.

A one-tier design consists of one pair of BIG-IP devices configured with VLANs to separate perimeter traffic from internal network traffic. In this case, the BIG-IP device spans multiple network zones. Traffic is secured through VLAN configuration.

If using a one-tier approach, please go to page 17 of the F5 Deployment Guide and follow the steps listed in that section.

Automated vs. Manual Configuration

Organizations using BIG-IP v11 or later can automate configuration and reduce or eliminate errors using F5 iApps™ templates. The iApp for Lync is pre-set to configure BIG-IP LTM, including reverse proxy settings, IP addresses, and SSL certificate/key names by asking the administrator a minimum of questions and then automatically building the configuration in seconds. See page 2 of the F5 Deployment Guide to acquire the latest F5 iApp for Microsoft Lync and to learn more about how this powerful technology can reduce your operational costs and deployment time.

The F5 iApp for Lync is a customizable extension of the BIG-IP platform for Lync that includes reverse proxy functionality. The operational integrity of iApps enables organizations to reduce common configuration errors by encapsulating all pertinent network settings for Lync into a discreet, manageable object. It can be saved, moved, and reused across all types of BIG-IP devices. In this way, the BIG-IP platform is automatically configured for Lync in the same way for your network, over and over—a feature particularly helpful for multi-tenant deployments.

iApps are customizable so that, as improvements are made to the environment, the current, best BIG-IP configuration is always carried forward. The F5 iApp for Lync is a flexible, effective feature unique to the BIG-IP platform, and there is no easier way to reliably automate device configuration for Lync. Figure 4 illustrates the simplicity of the Lync iApp user interface.

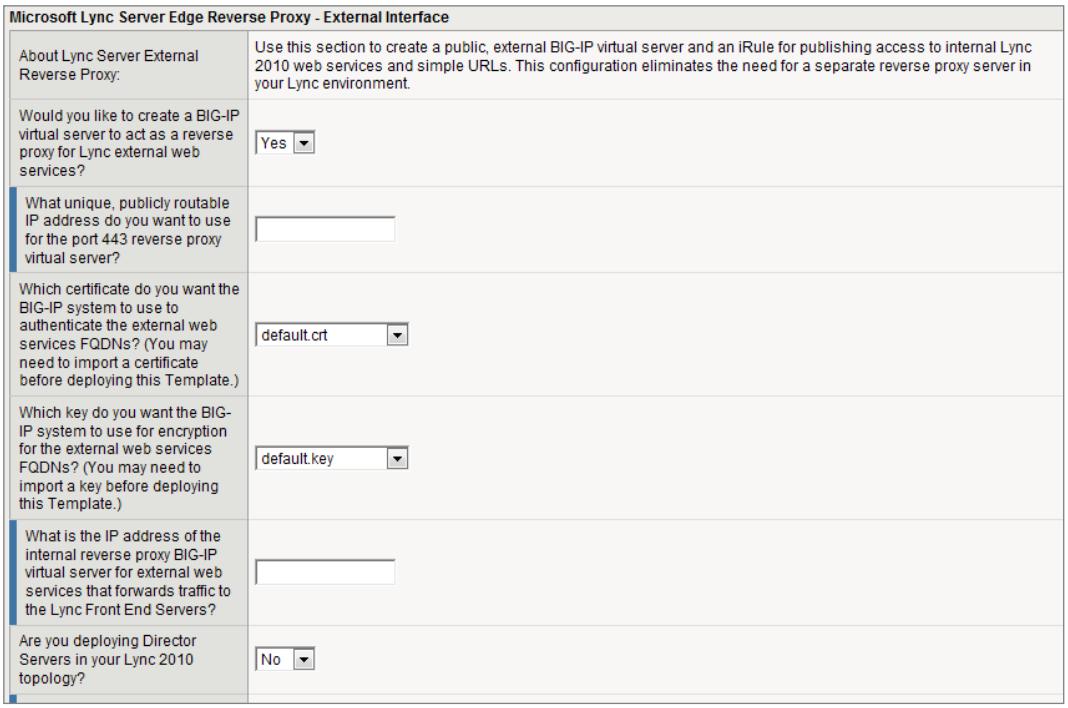

Figure 4: F5 iApp automates Lync reverse proxy configuration.

If using the iApp, refer to page 4, and then pages 9 and 14, of the F5 Deployment Guide, and follow the steps listed there. Be sure to note the additional steps required after applying the F5 iApp. If manually configuring the BIG-IP device, refer to the manual configuration tables on pages 22–25 of the guide.

Figure 5 provides a quick reference to the iApp and reverse proxy sections of the F5 Deployment Guide (based on Version 2.6 of the guide).

Using a single BIG-IP system for reverse proxy	Page 17
Getting the latest iApp for Lync	Page 4
iApp for reverse proxy on Lync Edge Servers—internal interface	Page 9
iApp for reverse proxy on Lync Edge Servers—external interface	Page 14
Manual configuration	Pages 21–25

Figure 5: Reference to sections in the F5 Deployment Guide for Lync.

Key Considerations

You have a few options for which technology to use for the reverse proxy, and here are some considerations when making your decision.

Reuse What You Already Have

BIG-IP LTM is already in your network, and the same BIG-IP LTM devices may be able to act as the reverse proxy without any additional hardware or software costs.

Make Sure Your Solution Is Scalable

The reverse proxy will be doing some IP/port translation, as well as URL filtering, and it needs to be able to scale appropriately. BIG-IP LTM ships with custom hardware designed for doing this type of traffic manipulation at speeds that are unmatched in other solutions that use commodity PC hardware architectures.

> A PC-based solution can be a typical application server running reverse proxy software (such as TMG) or it can be a hardware appliance offering network traffic features that is based on a client PC architecture, lacking the fundamental hardware design (bus architecture, chip sets, and integrated circuits) and firmware that provide computing speeds and capacity required for this type of manipulation for real-time communications workloads such as Lync.

Make Sure Your Solution Is Secure

It is very important to keep in mind that the reverse proxy lets external users inside your corporate (internal) network. TMG is a certified firewall, and its replacement should be, too.

This is *critical*. BIG-IP LTM is ICSA Labs certified as a firewall, and it includes the proper firewalling functionality to help secure the network. If a device is not certified, it should not be your reverse proxy.

Make Sure Your Solution Is Easy to Configure

Without an iApp engine, any Lync load balancing/reverse proxy solution is going to be complex, difficult to configure, and prone to misconfiguration. BIG-IP LTM has solved this problem.

Conclusion

Configuring BIG-IP Local Traffic Manager as a reverse proxy for Microsoft Lync can be accomplished by following the Deployment Guide F5 has published. This guidance includes load balancing and reverse proxy features for "internal" Lync servers as well as the Lync Edge Servers that reside in a network perimeter.

For multi-tenant deployments, every client is "external" to the network hosting Lync application services. Therefore, every client must pass its traffic through the reverse proxy or Lync Edge Servers when using any server-based feature.

Using this high-level configuration document, you can quickly determine which sections of the F5 Deployment Guide to use for your organization's deployment. Using the F5 iApp for Lync will speed and simplify the configuration.

Learn More

You can explore this topic further with these resources:

Deployment Guide

F5 Deployment Guide for Lync

Blogs

Best Practices for Lync Edge Server Network Design

TMG2F5 Series: BIG-IP LTM as the Lync Reverse Proxy

Ready to talk to F5? Visit the How to Buy page on f5.com to get in touch with an F5 representative.

F5 Networks, Inc. 401 Elliott Avenue West, Seattle, WA 98119 888-882-4447 www.f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com