Configuring BIG-IP LTM: Local Traffic Manager

Duration: 3 Days Course Code: ARR_F5N_BIG-LTM-CFG-3 Version: 17.5 Delivery Method: Virtual Learning

Overview:

Learn how to configure and manage BIG-IP Local Traffic Manager (LTM) as it is commonly deployed in an application delivery network to achieve operational efficiency and maintain critical business applications. Through a combination of lecture and hands-on labs, explore features and functionality to process and modify traffic behavior using profiles, persistence, caching, compression, and source network address translation (SNAT).

Monitor application health at layers 3, 4, and 7, and implement dynamic load balancing methods. Use traffic management shell (TMSH), the Configuration utility, and Linux commands to create traffic processing and monitoring objects, observe the resulting traffic statistics, and effectively operate the BIG-IP LTM system. Customize application delivery with iRules, establish application security, and harden system security using BIG-IP LTM functionality.

Updated November 2026

Virtual Learning

This interactive training can be taken from any location, your office or home and is delivered by a trainer. This training does not have any delegates in the class with the instructor, since all delegates are virtually connected. Virtual delegates do not travel to this course, Global Knowledge will send you all the information needed before the start of the course and you can test the logins.

Target Audience:

This course is intended for system and network administrators responsible for installation, setup, configuration, and administration of the BIG-IP LTM system.

Objectives:

- **After this course participants should be able to:**
- Configure and manage BIG-IP Local Traffic Manager (LTM) in an application delivery network.
- Implement load balancing, persistence, health monitoring, and traffic processing services.
- Modify traffic behavior using profiles, SNATs, caching, compression, and virtual servers.
- Create and manage traffic and monitoring objects using TMSH, Configuration Utility, and Linux commands.
- Customize and secure application delivery using iRules, local traffic policies, and LTM security features.

Prerequisites:

Students must complete one of the following F5 prerequisites before attending this course:

- Administering BIG-IP (ILT)
- F5 Certified BIG-IP Administrator

Suggested Prework

The following free Self-Directed Training (SDT) courses, although optional, are helpful for any student with limited BIG-IP administration and configuration experience:

- Getting Started with BIG-IP
- Getting Started with Local Traffic Manager (LTM)

General network technology knowledge and experience are

recommended before attending any F5 Global Training Services instructor-led course, including OSI model encapsulation, routing and switching, Ethernet and ARP, TCP/IP concepts, IP addressing and subnetting, NAT and private IP addressing, NAT and private IP addressing, default gateway, network firewalls, and LAN vs. WAN.

The following course-specific knowledge and experience is suggested before attending this course:

- Web application delivery
- HTTP, HTTPS, FTP, and SSH protocols
- TLS/SSL

Follow-on-Courses:

- *None*

Content:

Programme

Chapter 1: Introducing the BIG-IP System

- Introducing the BIG-IP System
- Initially Setting Up the BIG-IP System
- Archiving the BIG-IP Configuration
- Leveraging F5 Support Resources and Tools

Chapter 2: Reviewing Local Traffic Configuration

- Reviewing Nodes, Pools, and Virtual Servers
- Reviewing Address Translation
- Reviewing Routing Assumptions
- Reviewing Application Health Monitoring
- Reviewing Traffic Behavior Modification with Profiles
- Reviewing the TMOS Shell (TMSH)
- Reviewing Managing BIG-IP Configuration Data

Chapter 3: Load Balancing Traffic with LTM

- Exploring Load Balancing Options
- Using Priority Group Activation and Fallback Host
- Comparing Member and Node Load Balancing

Chapter 4: Modifying Traffic Behavior with Persistence

- Reviewing Persistence
- Introducing Cookie Persistence
- Specifying Default and Fallback Persistence
- Introducing SSL Persistence
- Introducing SIP Persistence
- Introducing Universal Persistence
- Introducing Destination Address Affinity Persistence
- Using Match Across Options for Persistence

Chapter 5: Monitoring Application Health

- Differentiating Monitor Types
- Customizing the HTTP Monitor
- Monitoring an Alias Address and Port
- Monitoring a Path vs. Monitoring a Device
- Managing Multiple Monitors
- Using Application Check Monitors
- Using Manual Resume and Advanced Monitor Timer Settings

Chapter 6: Processing Traffic with Virtual Servers

- Understanding the Need for Other Virtual Server Types
- Forwarding Traffic with a Virtual Server
- Understanding Virtual Server Order of Precedence
- Path Load Balancing

Chapter 7: Processing Traffic with SNATs

- Overview of SNATs
- Using SNAT Pools
- SNATs as Listeners
- SNAT Specificity
- VIP Bounceback
- Additional SNAT Options
- Network Packet Processing Review

Chapter 8: Modifying Traffic Behavior with Profiles

- Profiles Overview
- TCP Express Optimization
- TCP Profiles Overview
- HTTP Profile Options
- HTTP/2 Profile Options
- OneConnect
- Offloading HTTP Compression to BIG-IP
- Web Acceleration Profile and HTTP Caching
- Stream Profiles
- F5 Acceleration Technologies

Chapter 9: Selected Topics

- VLAN, VLAN Tagging, and Trunking
- Restricting Network Access
- SNMP Features
- Segmenting Network Traffic with Route Domains

Chapter 10: Customizing Application Delivery with iRules

- Getting Started with iRules
- Understanding When iRules are Triggered
- Deploying iRules
- Constructing an iRule
- Testing and Debugging iRules
- Exploring iRules Documentation

Chapter 11: Customizing Application Delivery with Local Traffic Policies

- Getting Started with Local Traffic Policies
- Configuring and Managing Policy Rules

Chapter 12: Securing Application Delivery with LTM

- Understanding Today's Threat Landscape
- Integrating LTM Into Your Security Strategy
- Defending Your Environment Against SYN Flood Attacks
- Defending Your Environment Against Other Volumetric Attacks
- Addressing Application Vulnerabilities with iRules and Local Traffic Policies
- Detecting and Mitigating Other Common HTTP Threats

Chapter 13: Final Lab Project

- About the Final Lab Project
- Possible Solution to Lab 13.1

Chapter 14: Additional Training and Certification

- Getting Started Series Web-Based Training
- F5 Instructor Led Training Curriculum
- F5 Professional Certification Program

Appendix A: Troubleshooting

- Working with F5 Technical Support

Appendix B: Additional Resources

- Resources on MyF5, DevCentral and More

Further Information:

For More information, or to book your course, please call us on 0800/84.009

info@globalknowledge.be

www.globalknowledge.com/en-be/