

Configuring BIG-IP LTM: Local Traffic Manager

Length 3 Days

Description

This 3 day course gives network professionals a functional understanding of BIG-IP Local Traffic Manager, introducing students to both commonly used and advanced BIG-IP LTM features and functionality. Incorporating lecture, extensive hands-on labs, and classroom discussion, the course helps students build the well-rounded skill set needed to manage BIG-IP LTM systems as part of a flexible and high performance application delivery network.

Participants

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

Prerequisites

IMPORTANT INFORMATION:

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

- Administering BIG-IP instructor-led course
- F5 Certified BIG-IP Administrator

F5 have advised that with the introduction of v13 there is no longer any overlap between the BIG-IP Admin course and the Configuring LTM course.

This has started to present many issues for delegates attending the LTM course without having sufficient BIG-IP Admin skills.

To help solve this F5 US has brought out a BIG-IP Admin exam which can be taken before attending the LTM course. Indeed, this test should be taken by anyone wanting to attend any of the F5 specialist classes (DNS, APM, ASM, AFM etc.) if they have not attended the BIG-IP Admin or LTM class.

Although this is not mandatory, F5 would like to stress the importance of anyone attending F5 specialist classes if they have not attended an Admin class or passed the test below with 70% or greater. Obviously, a 201 BIG-IP Administrator exam certification is an instant prerequisite.

Step 1: get an account on F5 University (<https://university.f5.com>)

Step 2: goto: <https://wbt.training.f5.com/wbt/go?id=2110>

Take the test. Pass mark is 70%.

Step 3: take a screen shot as proof of your results.

The following free web-based courses, although optional, will be very helpful for any student with limited BIG-IP administration and configuration experience:

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

The following general network technology knowledge and experience are recommended before attending any F5 Global Training Services instructor-led course:

- OSI model encapsulation
- Routing and switching
- Ethernet and ARP
- TCP/IP concepts
- IP addressing and subnetting
- NAT and private IP addressing
- Default gateway
- Network firewalls
- 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

Objectives

- BIG-IP initial setup (licensing, provisioning, and network configuration)
- A review of BIG-IP local traffic configuration objects
- Using dynamic load balancing methods
- Modifying traffic behavior with persistence (including SSL, SIP, universal, and destination address affinity persistence)
- Monitoring application health with Layer 3, Layer 4, and Layer 7 monitors (including transparent, scripted, and external monitors)
- Processing traffic with virtual servers (including network, forwarding, and reject virtual servers)
- Processing traffic with SNATs (including SNAT pools and SNATs as listeners)
- Configuring high availability (including active/standby and N+1 sync failover device groups, connection and persistence mirroring, and sync-only device groups)
- Modifying traffic behavior with profiles (including advanced HTTP profile options, caching, compression, and OneConnect profiles)
- Advanced BIG-IP LTM configuration options (including VLAN tagging and trunking, SNMP features, and packet filters)
- Deploying application services with iApps
- Customizing application delivery with iRules and local traffic policies

By the end of this course, the student should be able to use both the Configuration utility, TMSH, and Linux commands to configure and manage BIG-IP LTM systems in an application delivery network. In addition, students should be able to monitor the BIG-IP system to achieve operational efficiency, and establish and maintain high availability infrastructure for critical business applications.

Programme

Chapter 1: Setting Up 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 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

- Virtual Servers Concepts
- Path Load Balancing
- Introducing Auto Last Hop

Chapter 7: Processing Traffic with SNATs

- Overview of SNATs
- SNAT Auto Map

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

Chapter 8: Configuring High Availability

- Sync-Failover Group Concepts
- Synchronization, State and Failover
- Traffic Group Concepts
- N+1 Concepts

Chapter 9: Configuring High Availability Part 2

- Failover Triggers and Detection
- Stateful Failover
- Device Group Communication
- Sync-Only Device Groups

Chapter 10: Modifying Traffic Behavior with Profiles

- Profiles Overview
- Common Protocol Profile Types and Settings
- TCP Express Optimization
- Performance Improvements
- Configuring and Using Profiles
- HTTP Profile Options
- OneConnect
- Offloading HTTP Compression to BIG-IP
- HTTP Caching
- Stream Profiles
- F5 Acceleration Technologies
- Analytics

Chapter 11: Selected Topics

- VLAN, VLAN Tagging, and Trunking
- Restricting Network Access
- SNMP Features
- Internet Protocol Version 6 (IPv6)

Chapter 12: Deploying Application Services with iApps

- Simplifying Application Deployment with iApps
- Using iApps Templates
- Deploying an Application Service
- Reconfiguring an Application Service
- Leveraging the iApps Ecosystem on DevCentral

Chapter 13: Customizing Application Delivery with iRules and Local Traffic Policies

- Getting Started with iRules
- Triggering an iRule
- Introducing iRule Constructs
- Leveraging the DevCentral Ecosystem
- Deploying and Testing iRules
- Getting Started with Local Traffic Policies
 - o What Can You Do with a Local Traffic Policy?
 - o How Does a Local Traffic Policy Work?
 - o Understanding Local Traffic Policy Workflow
 - o Introducing the Elements of a Local Traffic Policy
 - o Specifying the Matching Strategy
 - o What Are Rules?
 - o Understanding Requires and Controls
- Configuring and Managing Policy Rules
 - o Configuring a New Rule
 - o Including Tcl in Certain Rule Settings

Chapter 14: Final Lab Project

Chapter 15: 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 Support