Administering BIG-IP and Configuring BIG-IP LTM : Local Traffic Manager Bundle

Duration: 5 Days Course Code: F5N_BIGIPAD_LTM_BDLE Version: 17.5 Delivery Method: Company Event

Overview:

Administering BIG-IP
Begin the BIG-IP learning journey with Administering BIG-IP. Learn how to set up and operate the BIG-IP system as it is commonly deployed in an application delivery network.
Through a combination of instructor-led lecture and hands-on labs, complete initial configuration including licensing, resource provisioning, networking, high-availability (HA), and more to establish a secure system.
Gain practical experience implementing traffic processing objects: pools, virtual servers, health monitors, network address translation (NATs and SNATs), and more.
Explore modifying traffic behavior with profiles that include SSL offload and re-encryption, and with persistence, including source address affinity and cookie persistence using the configuration utility graphical user interface (GUI) and TMSH (TMOS Shell) command line interface (CLI).
Build troubleshooting skills learning various logging methods including local, high-speed, and legacy remote logging combined with practice using troubleshooting tools like TCPDUMP.
Learn typical administrative and operational activities including user roles and administrative partition creation and management.

Configuring BIG-IP LTM: Local Traffic Manager
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 handson 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

Company Events

These events can be delivered exclusively for your company at our locations or yours, specifically for your delegates and your needs. The Company Events can be tailored or standard course deliveries.

Target Audience:

This course is intended for network administrators, operators, and engineers responsible for managing the normal day-to-day operation and administration of a BIG-IP application delivery network. This course presents the prerequisite knowledge for many other of F5's BIG-IP instructor-led training courses.

Objectives:

- After this course participants should be able to:
- Configure and administer F5 BIG-IP and Local Traffic Manager (LTM) in enterprise application delivery environments.
- Implement and manage core traffic services including virtual servers, pools, load balancing, health monitors, NATs, and SNATs.
- Optimize application delivery using profiles, SSL offload/re-encryption, persistence methods, caching, compression, and iRules.
- Deploy high availability (HA) and Device Service Clustering (DSC) solutions to ensure business continuity and application resilience.
- Monitor, troubleshoot, and secure BIG-IP environments using logging, TCPDUMP, performance monitoring, and LTM security features to maintain reliable and secure application services.

Prerequisites:

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

Administering BIG-IP

Chapter 1: Setting Up the BIG-IP System

- Introducing the BIG-IP System
- Initially Setting Up the BIG-IP System
- Configuring the Management Interface
- Activating the Software License
- Provisioning Modules and Resources
- Importing a Device Certificate
- Specifying BIG-IP Platform Properties
- Configuring the Network
- Configuring Network Time Protocol (NTP) Servers
- Configuring Domain Name System (DNS) Settings
- Configuring High Availability Options
- Archiving the BIG-IP Configuration
- Leveraging F5 Support Resources and Tools

Chapter 2: Traffic Processing Building Blocks

- Identifying BIG-IP Traffic Processing Objects
- Configuring Virtual Servers and Pools
- Load Balancing Traffic
- Viewing Module Statistics and Logs
- Using the Traffic Management Shell (TMSH)
- Understanding the TMSH Hierarchical Structure
- Navigating the TMSH Hierarchy
- Managing BIG-IP Configuration State and Files
- BIG-IP System Configuration State
- Loading and Saving the System Configuration
- Shutting Down and Restarting the BIG-IP System
- Saving and Replicating Configuration Data (UCS and SCF)

Chapter 3:

Using NATs and SNATs

- Address Translation on the BIG-IP System
- Mapping IP Addresses with NATs
- Solving Routing Issues with SNATs
- Configuring SNAT Auto Map on a Virtual Server
- Monitoring for and Mitigating Port Exhaustion

Chapter 4:

Monitoring Application Health

Chapter 7: Administering the BIG-IP System

- Configuring Logging
- Legacy Remote Logging
- Introducing High Speed Logging (HSL)
- High-Speed Logging Filters
- HSL Configuration Objects
- Configuring High Speed Logging
- Using TCPDUMP on the BIG-IP System
- Leveraging the BIG-IP iHealth System
- Viewing BIG-IP System Statistics
- Defining User Roles and Administrative Partitions
- Leveraging Multitenancy and vCMP (Legacy) vCMP

Chapter 8:Configuring High Availability

- Introducing Device Service Clustering (DSC)
- Preparing to Deploy a DSC Configuration
- Configuring DSC Communication Settings
- Establishing Device Trust
- Establishing a Sync-Failover Device Group
- Synchronizing Configuration Data
- Exploring Traffic Group Behavior
- Understanding Failover Managers and Triggers
- Achieving Stateful Failover with Mirroring

Configuring BIG-IP LTM: Local Traffic Manager

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

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

- Introducing Monitors
- Types of Monitors
- Monitor Interval and Timeout Settings
- Configuring Monitors
- Assigning Monitors to Resources
- Managing Pool, Pool Member, and Node Status Using the Network Map

Chapter 5: Modifying Traffic Behavior with Profiles

- Introducing Profiles
- Understanding Profile Types and Dependencies
- Configuring and Assigning Profiles
- Introducing SSL Offload and SSL Re-Encryption

Chapter 6:

Modifying Traffic Behavior with Persistence

- Understanding the Need for Persistence
- Introducing Source Address Affinity Persistence
- Managing Object State
- 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

Further Information:

For More information, or to book your course, please call us on Head Office 01189 123456 / Northern Office 0113 242 5931

info@globalknowledge.co.uk

www.globalknowledge.com/en-gb/

Global Knowledge, Mulberry Business Park, Fishponds Road, Wokingham Berkshire RG41 2GY UK