F5 ADMINISTERING & CONFIGURING BIG-IP LOCAL TRAFFIC MANAGER (LTM)

DURATION - 5 DAYS

COURSE OVERVIEW

This 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.

WHO SHOULD ATTEND?

This course is intended for network administrators, network operators and network engineers with no or minimal experience with F5 who would be responsible for installation and set up of BIG-IP LTM systems.

PREREQUISITES

Students should understand TCP/IP addressing and routing . However, no prior F5 experience required

COURSE OBJECTIVES

- Back up the BIG-IP system configuration for safekeeping
- Configure virtual servers, pools, monitors, profiles, and persistence objects
- Test and verify application delivery through the BIG-IP system using local traffic statistics
- Configure priority group activation on a load balancing pool to allow servers to be activated only as needed to process traffic
- Compare and contrast member-based and node-based dynamic load balancing methods
- Configure connection limits to place a threshold on traffic volume to particular pool members and nodes
- Differentiate between cookie, SSL, SIP, universal, and destination address affinity persistence, and describe use cases for each
- Describe the three Match Across Services persistence options and use cases for each
- Configure health monitors to appropriately monitor application delivery through a BIG-IP system
- Configure different types of virtual services to support different types of traffic processing through a BIG-IP system
- Configure different types of SNATs to support routing of traffic through a BIG-IP system
- Configure VLAN tagging and trunking
- Restrict administrative and application traffic through the BIG-IP system using packet filters, port lockdown, and virtual server settings
- Configure SNMP alerts and traps in support of remote monitoring of the BIG-IP system
- Use iRules and local traffic policies appropriately to customize application delivery through the BIG-IP system
- Configure the BIG-IP to detect and mitigate some common attacks at the network and application layers using LTM features such as SYN check, eviction policies, iRules and Local Traffic Policies

WHAT YOU'LL LEARN

- Getting started with the BIG-IP system & initial setup (licensing, provisioning, and network configuration)
- Traffic processing with LTM and SNATs
- Using the Traffic Management Shell (tmsh) command line interface
- Using NATs and SNATs
- Monitoring application health and managing object status
- Modifying traffic behaviour with profiles
- Modifying traffic behaviour with persistence, including source address affinity and cookie persistence
- Troubleshooting the BIG-IP system, including logging (local, high-speed, and legacy remote logging), and using tcpdump
- Always-On Management (AOM)
- User roles and administrative partitions
- vCMP concepts   Customising application delivery with iRules
- A review of BIG-IP local traffic configuration objects
- Using dynamic load balancing methods
- Monitoring application health with Layer 3, Layer 4, and Layer 7 monitors
- Processing traffic with virtual servers
- Configuring high availability (including active/standby and N+1 sync failover device groups, connection and persistence mirroring, and sync-only device groups)
- Deploying application services with iApps
- Customising application delivery with iRules and local traffic policies