F5 LTM Essentials

Anthony Graber
DISA Solutions Engineer
Chris Sheehan
SI Solutions Engineer

WE MAKE APPS GO FASTER. SMARTER. SAFER.

Agenda

- Set Up and Licensing
- TMOS and Networking
- Basic LTM Components
- Load Balancing and Monitoring
- Profiles and Persistence
- Security and SSL Offload
- Acceleration Technologies
- Device Service Clusters (High Availability)

Install

- **Initial System Setup**
- **Using the Setup Utility**
- **Using tmsh**

BIG-IP System Initial Setup

1. Set up the management port

2. Run the Setup Utility
   - License the BIG-IP system
   - Provision modules
   - Configure the platform
   - Optionally, setup a failover pair

Management Port Defaults

IP Address	192.168.1.245/24
Username/Password	Web: admin/admin CLI: root/default

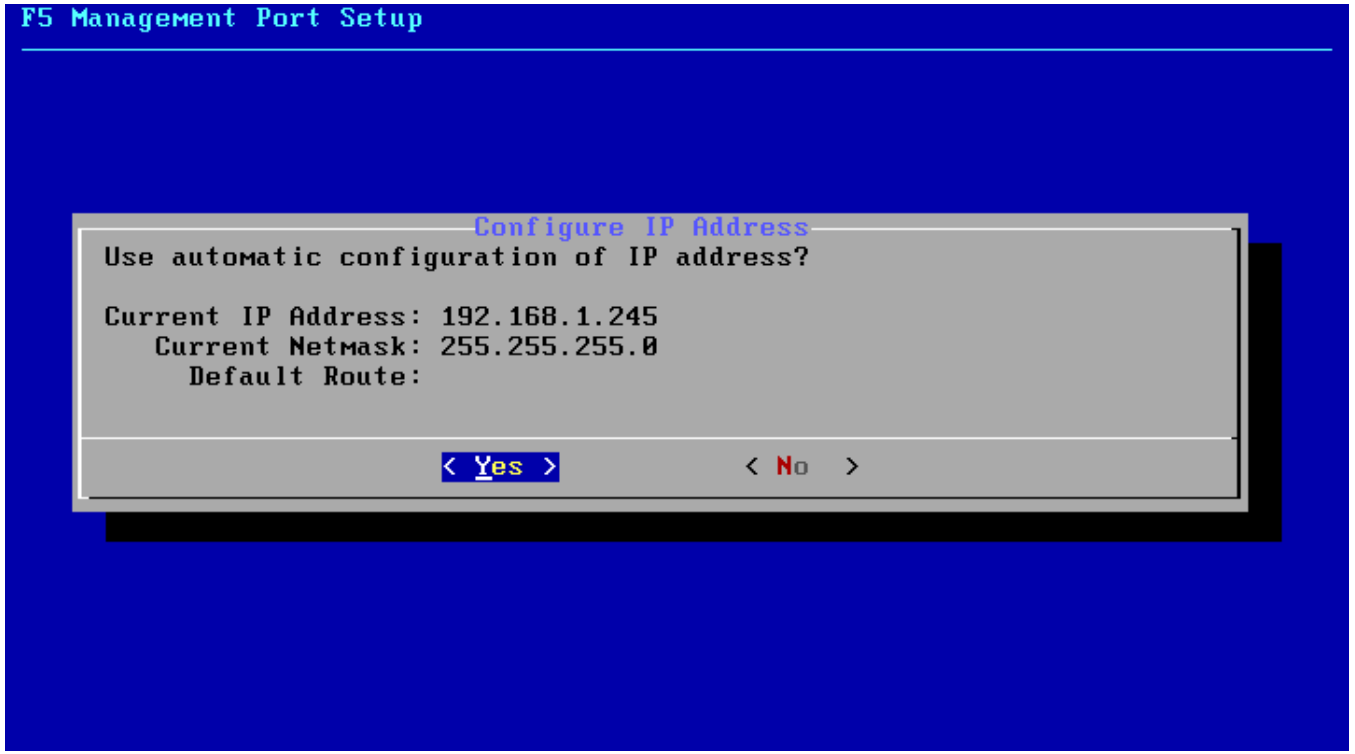

Configuring Management Port with CLI

1. Log into the CLI using root/default
2. Type “**config**” at the CLI prompt

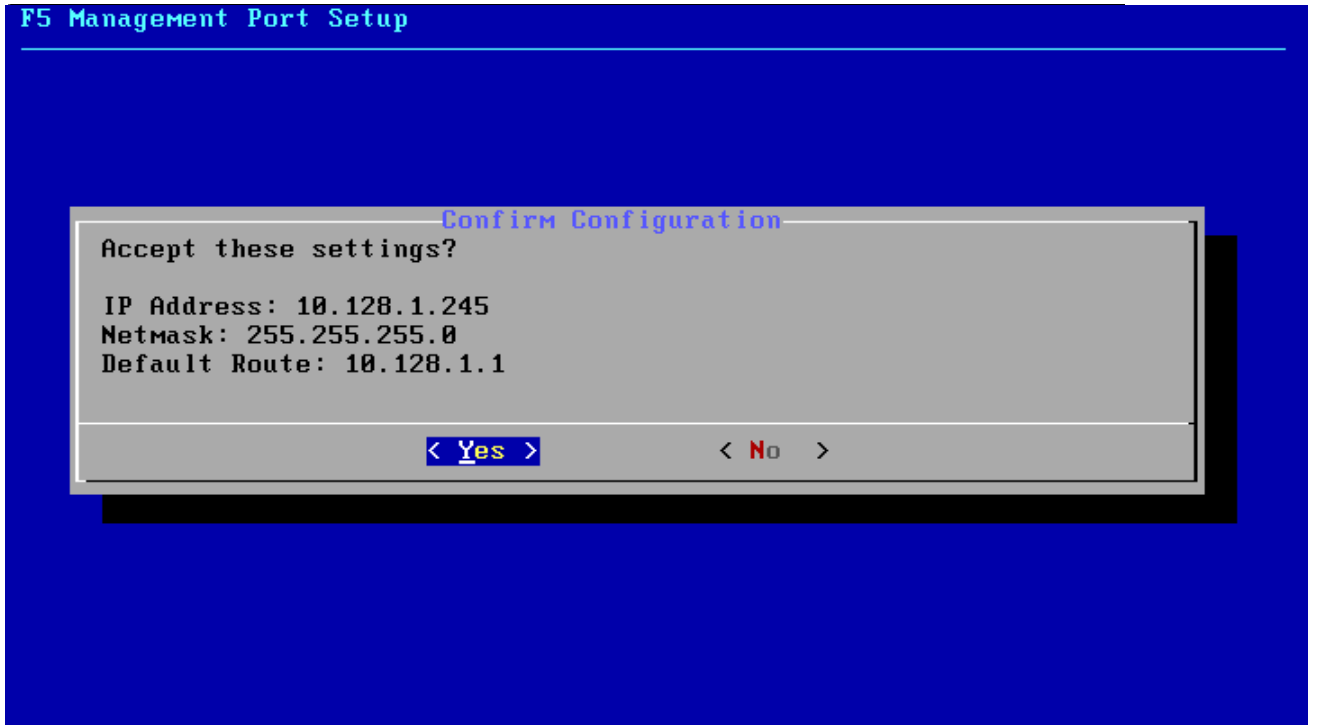

Accessing the BIG-IP Setup Utility

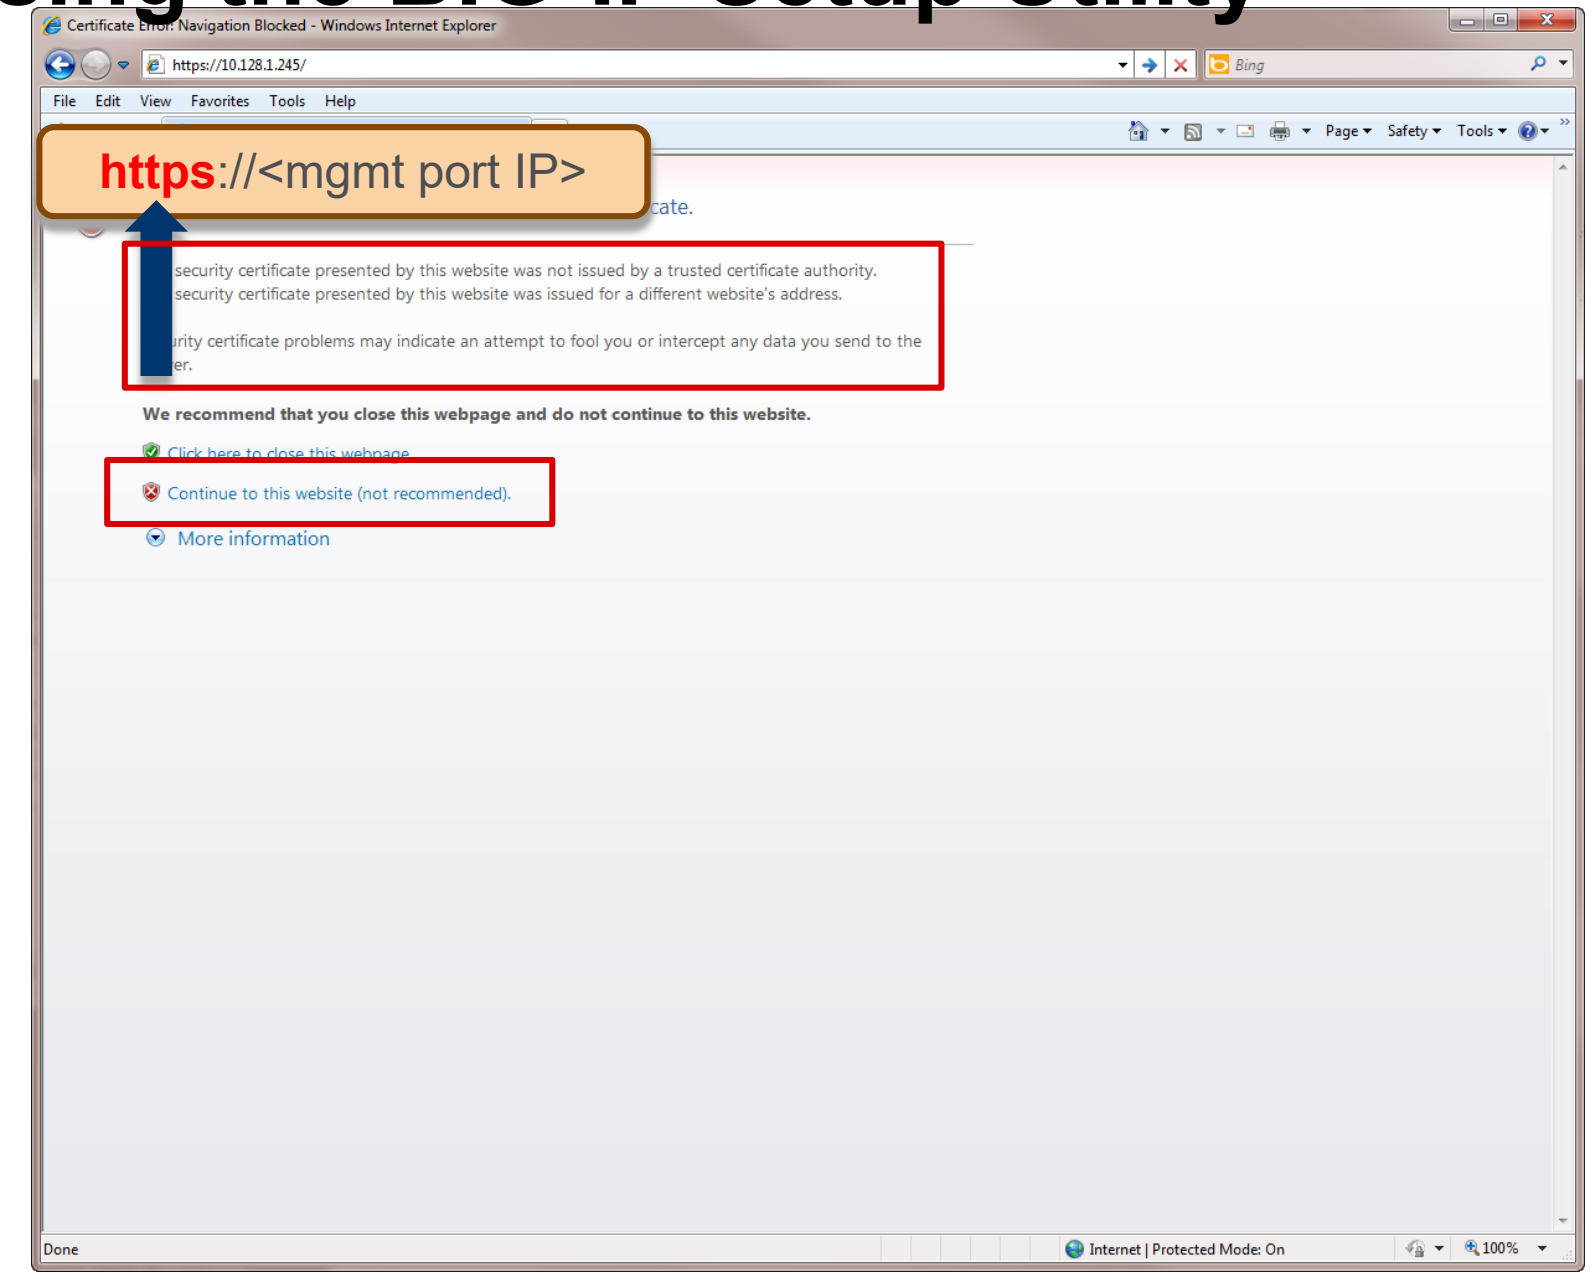

Install

- Initial System Setup
- Using the Setup Utility
- Using tmsh

Setup Utility

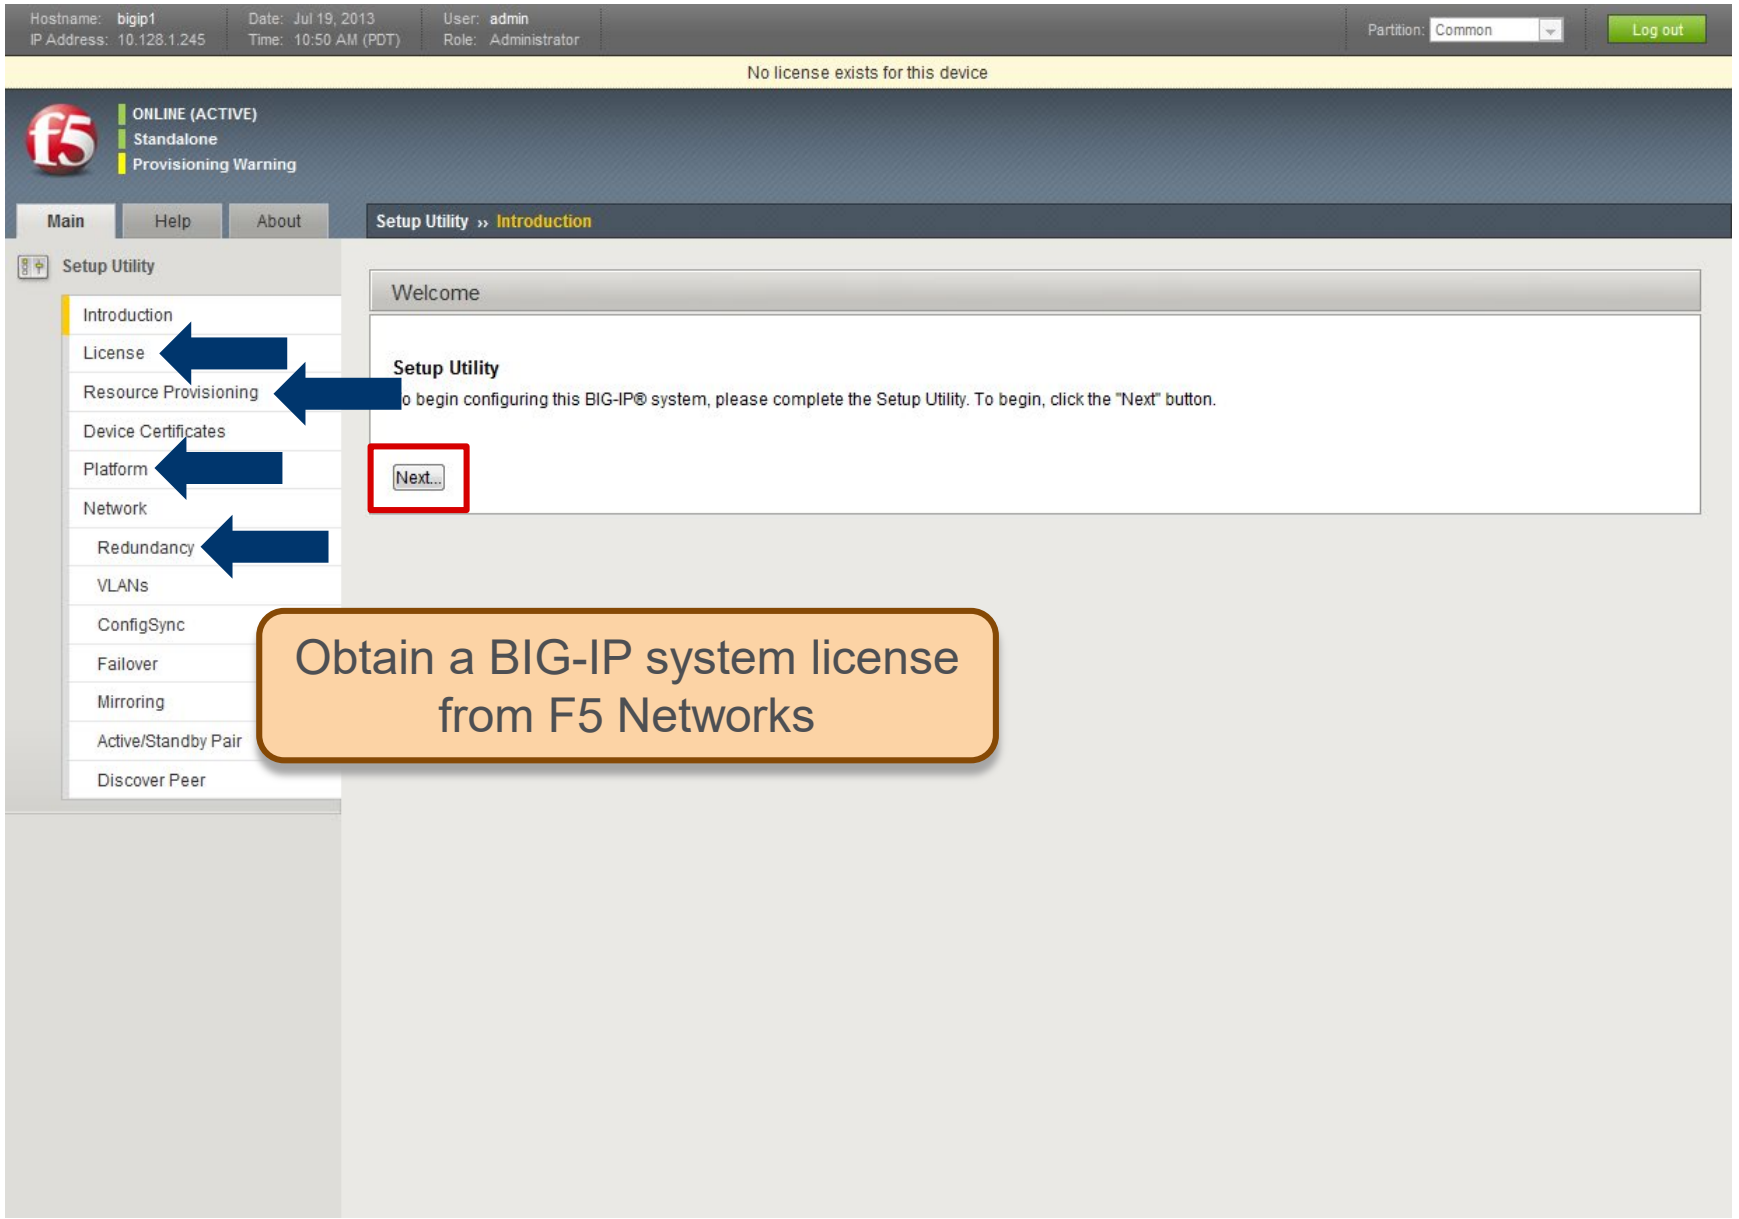

Licensing Methods

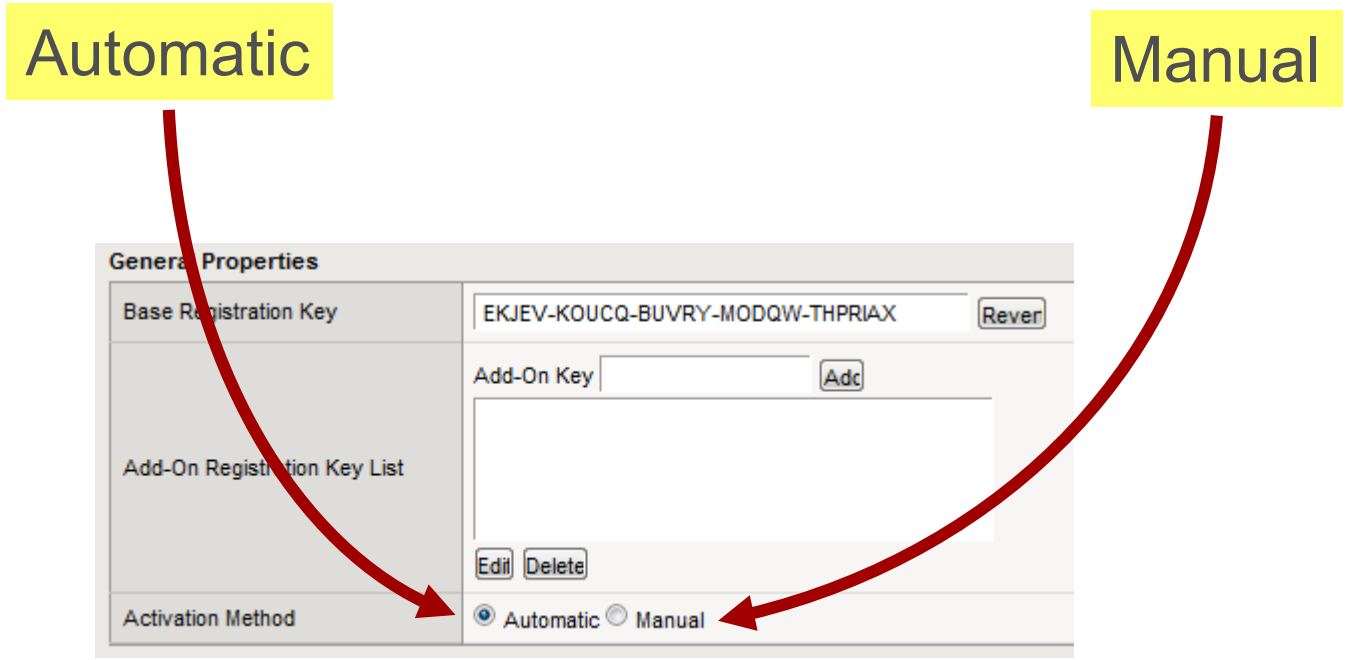

Using Automatic Licensing

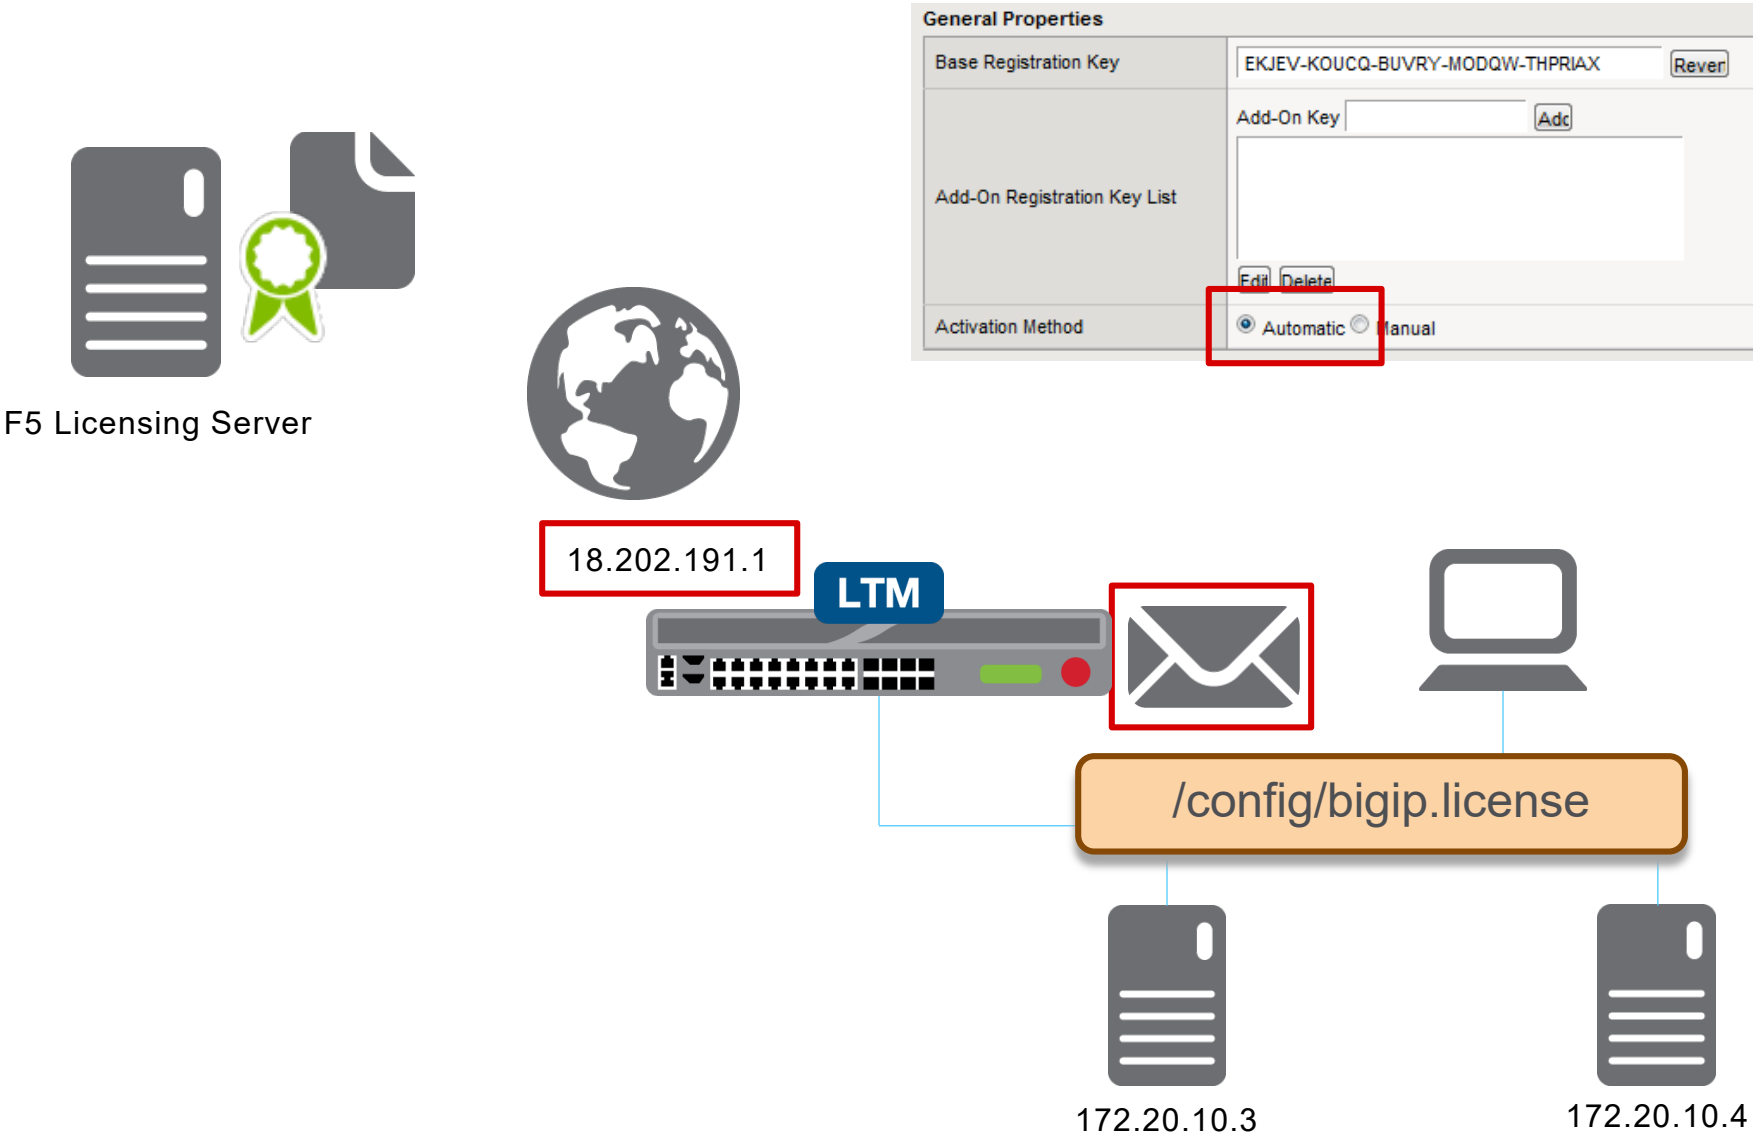

Using Manual Licensing

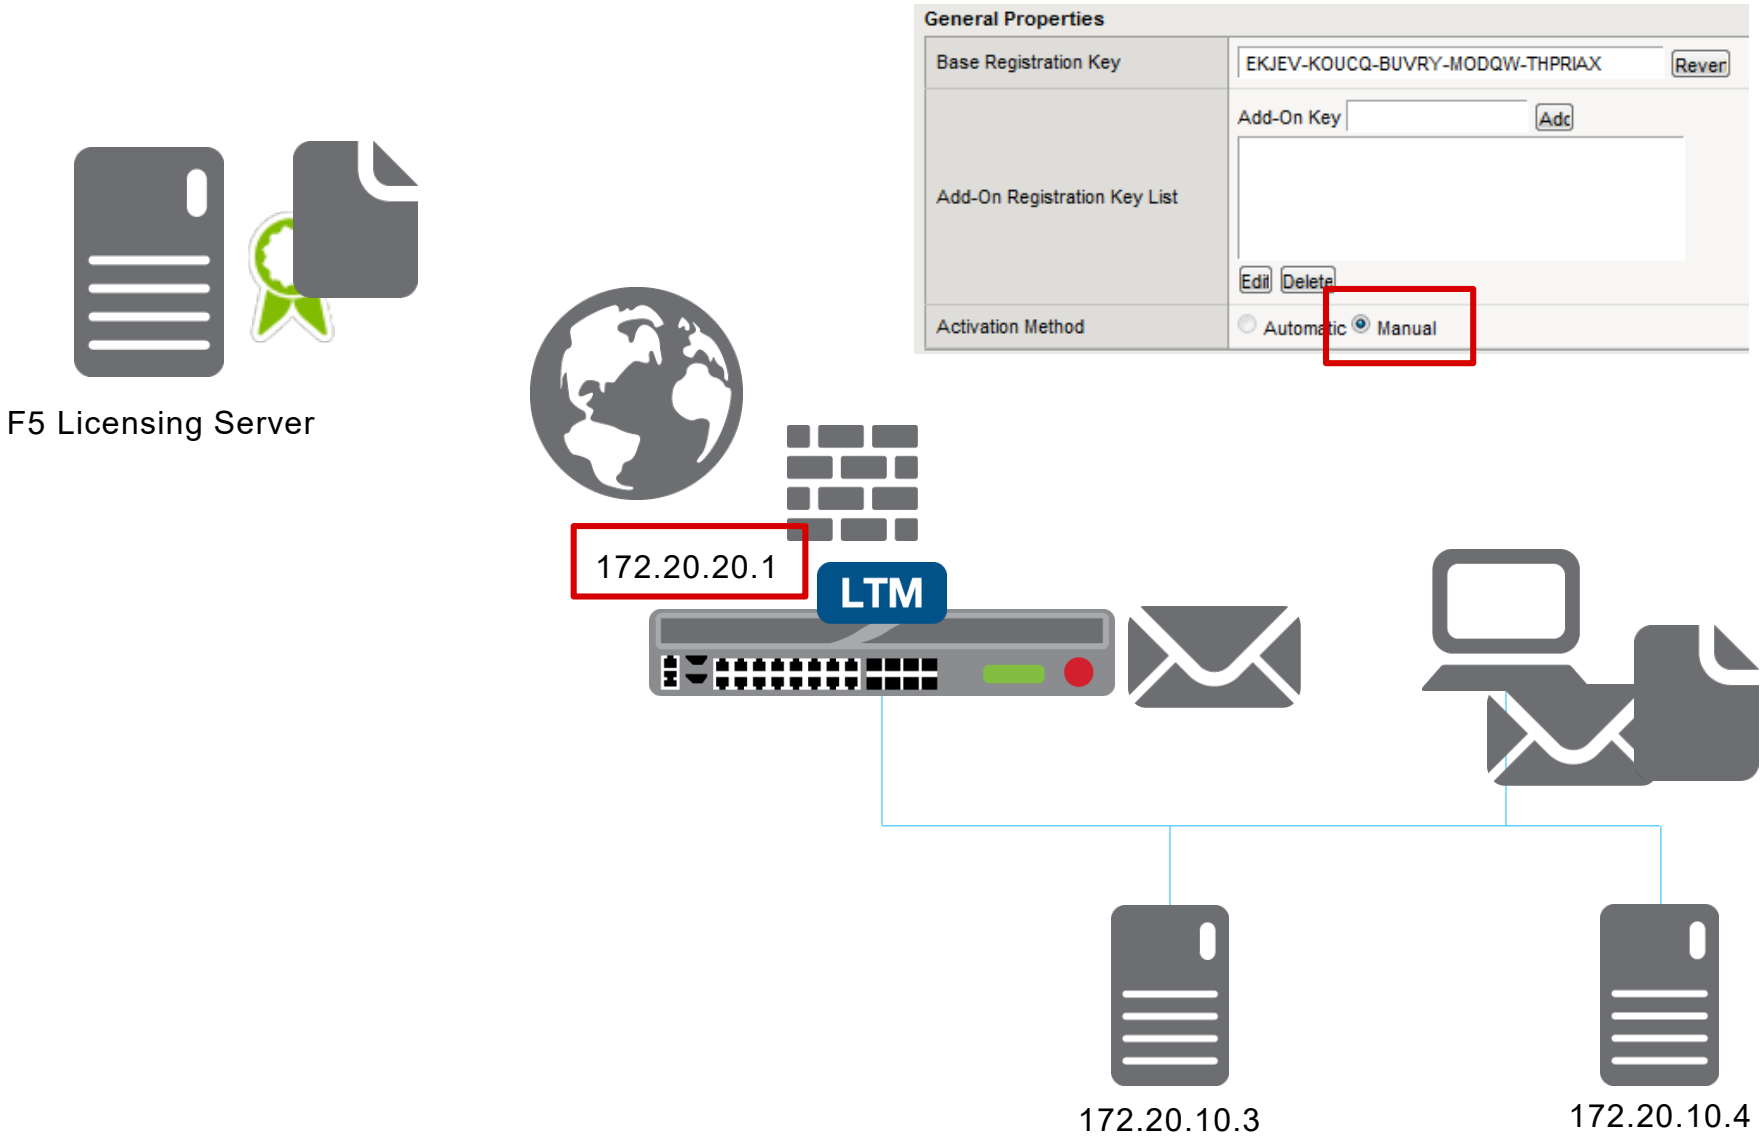

Two Methods for Manual Licensing

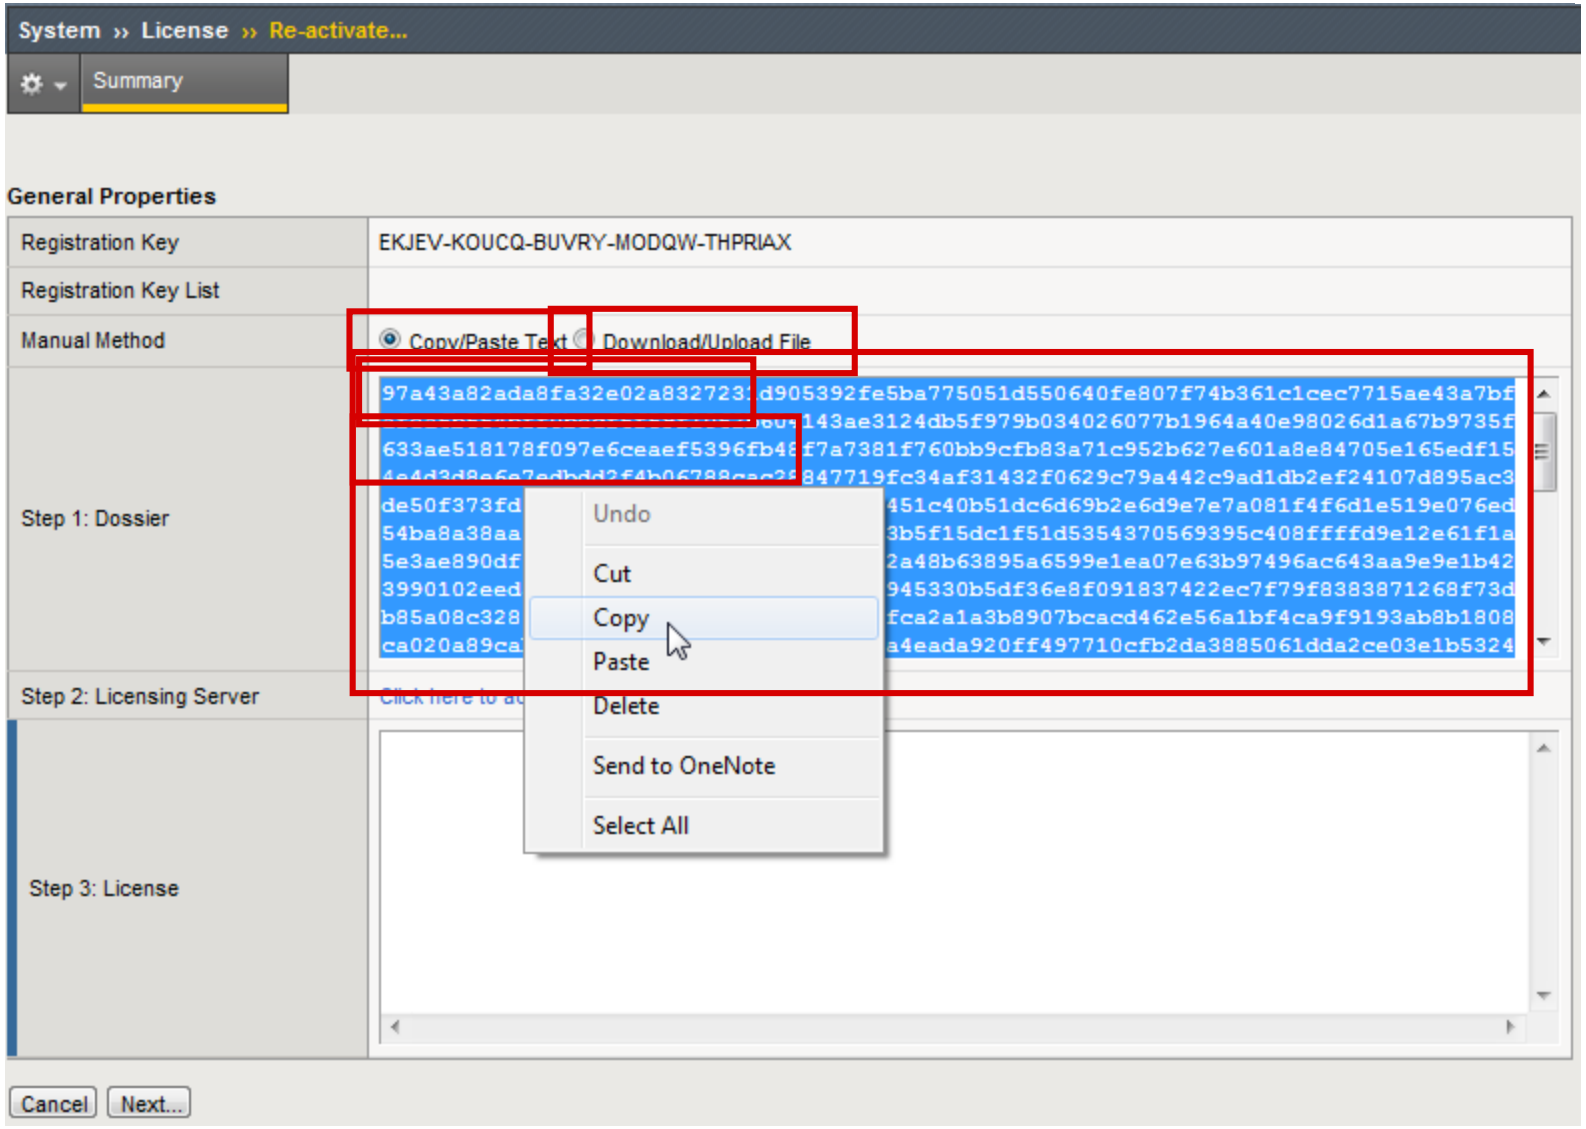

Using the F5 Licensing Server Web Site

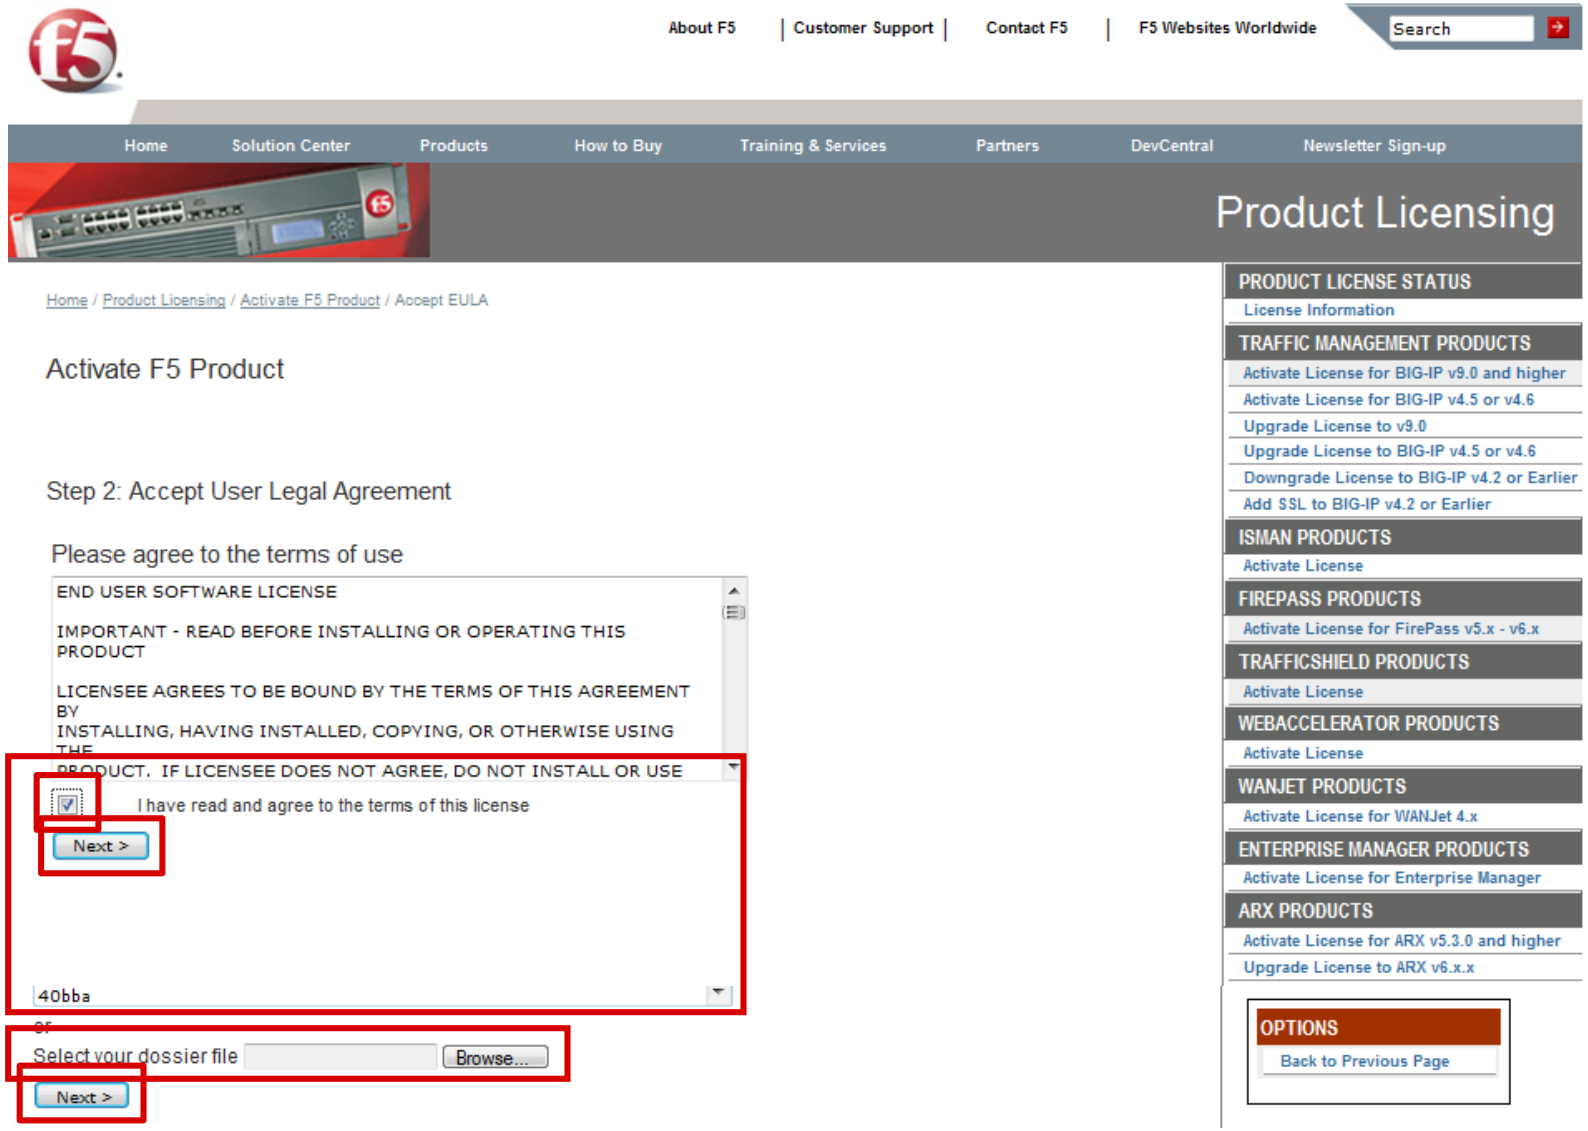

Download or Copy the F5 License

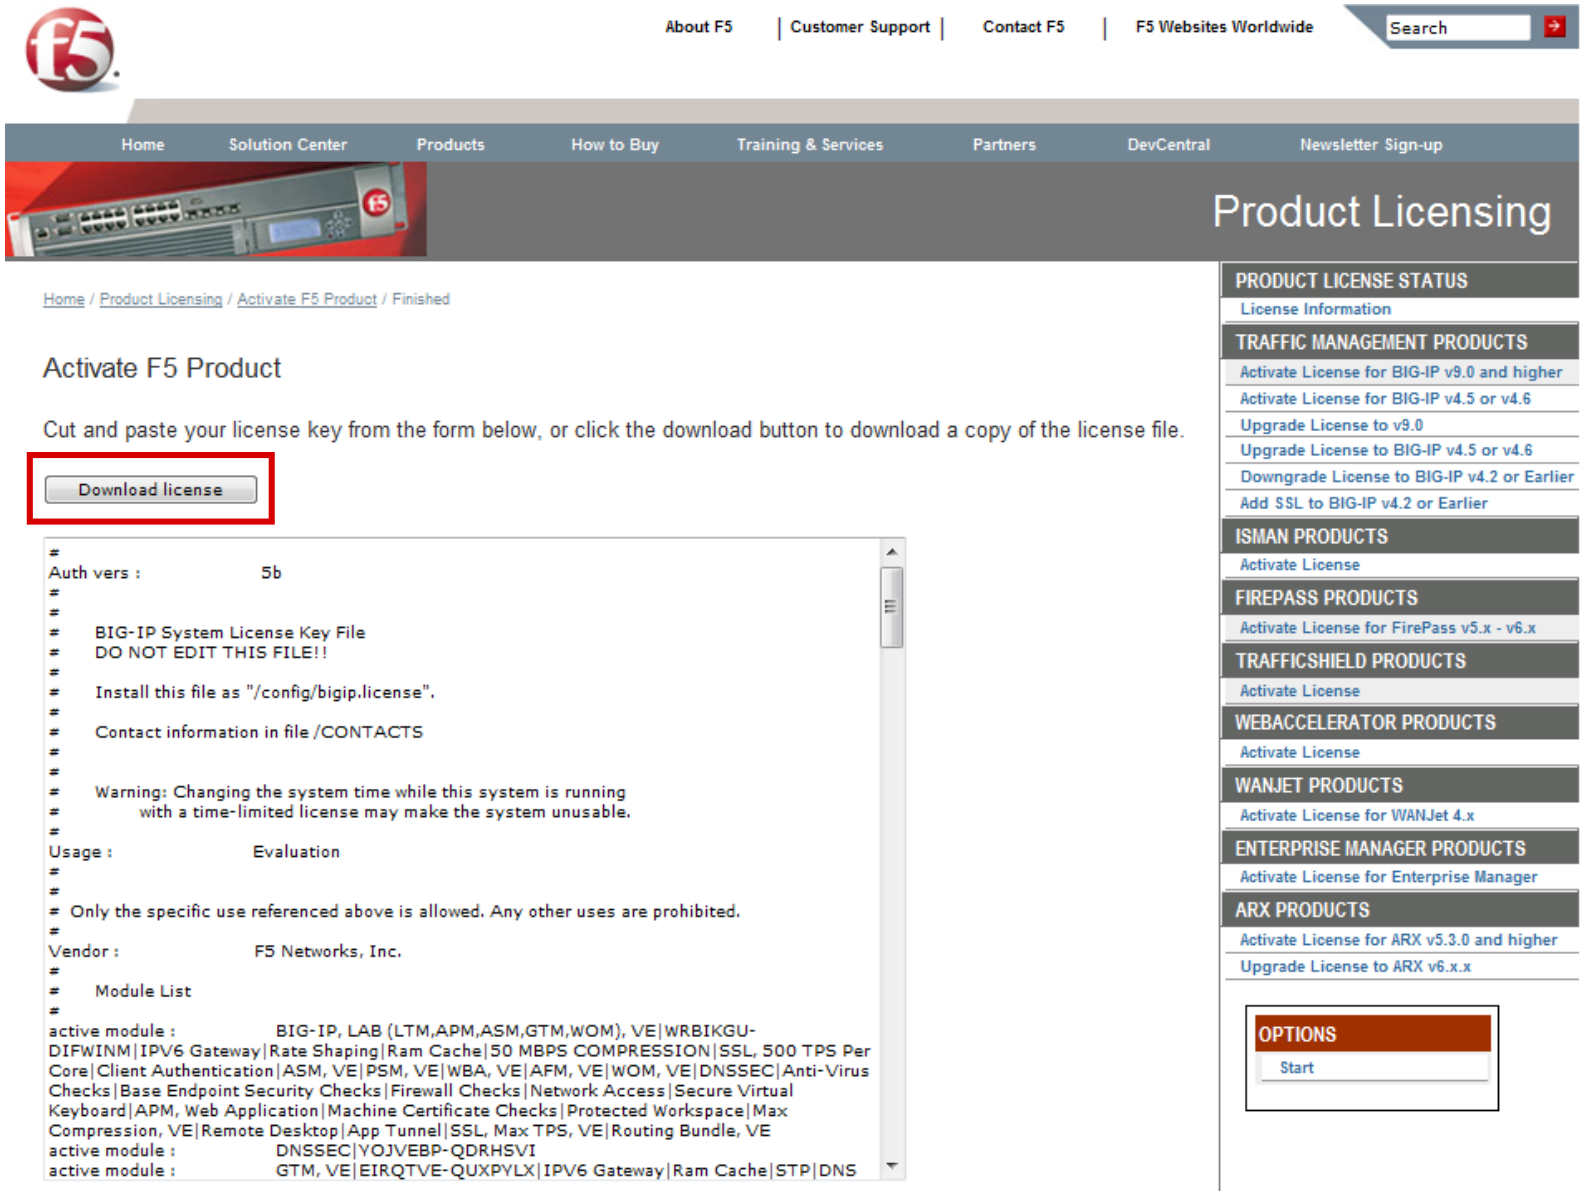

Paste the License on the BIG-IP

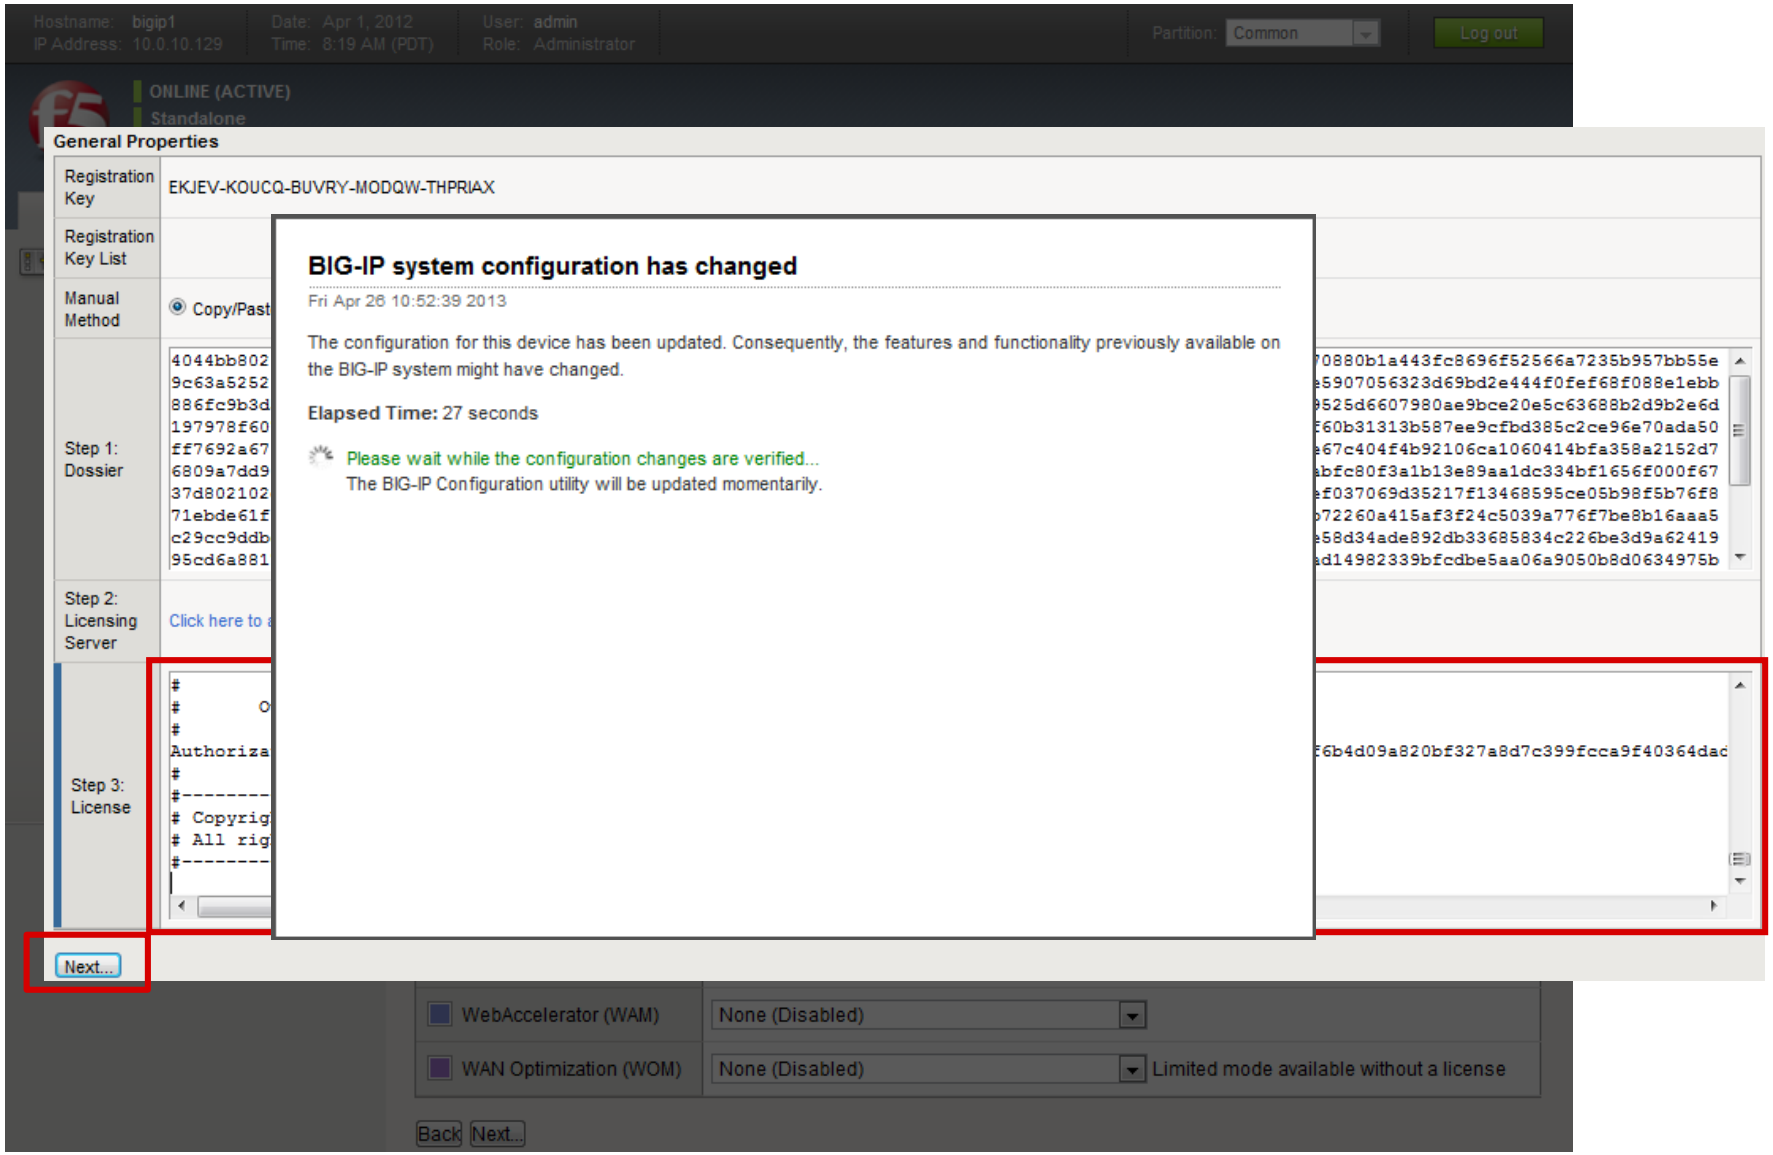

Resource Provisioning

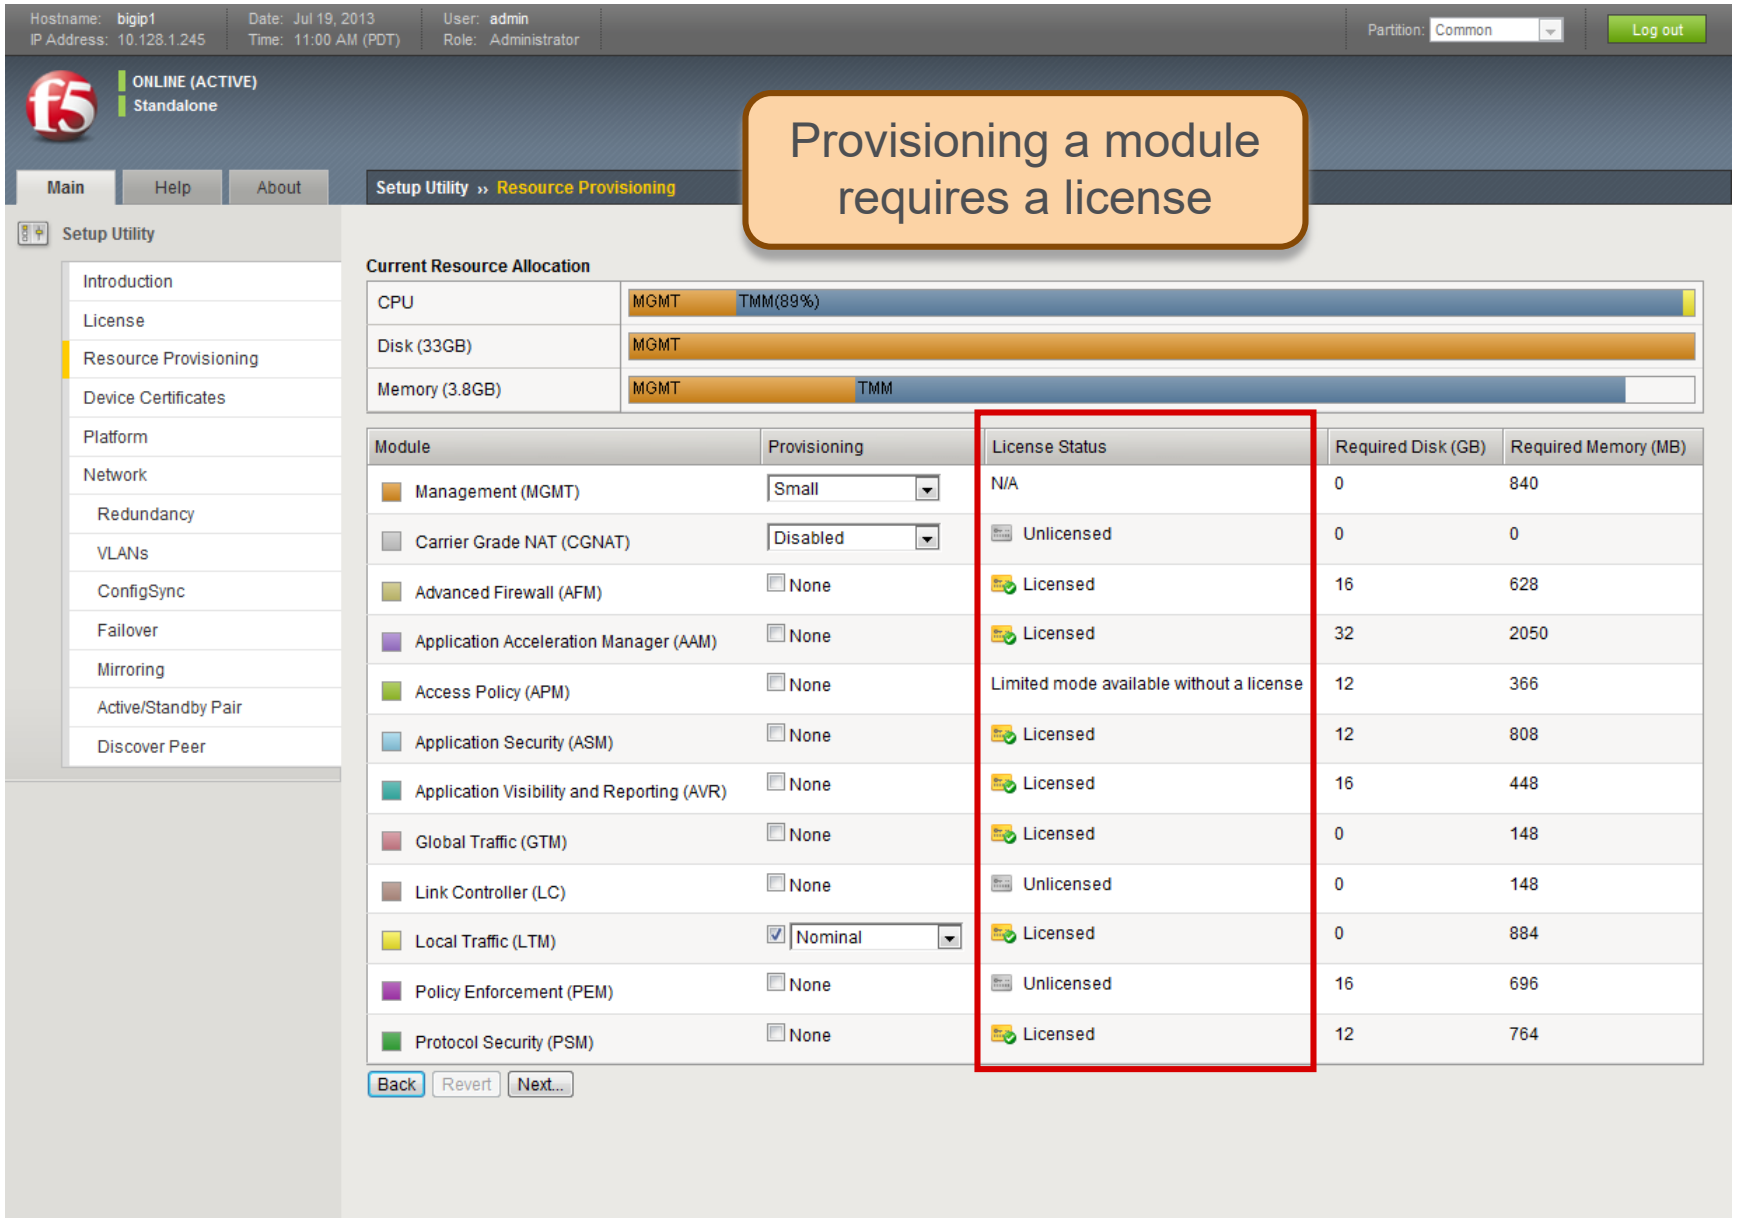

Setup Utility – Platform Page

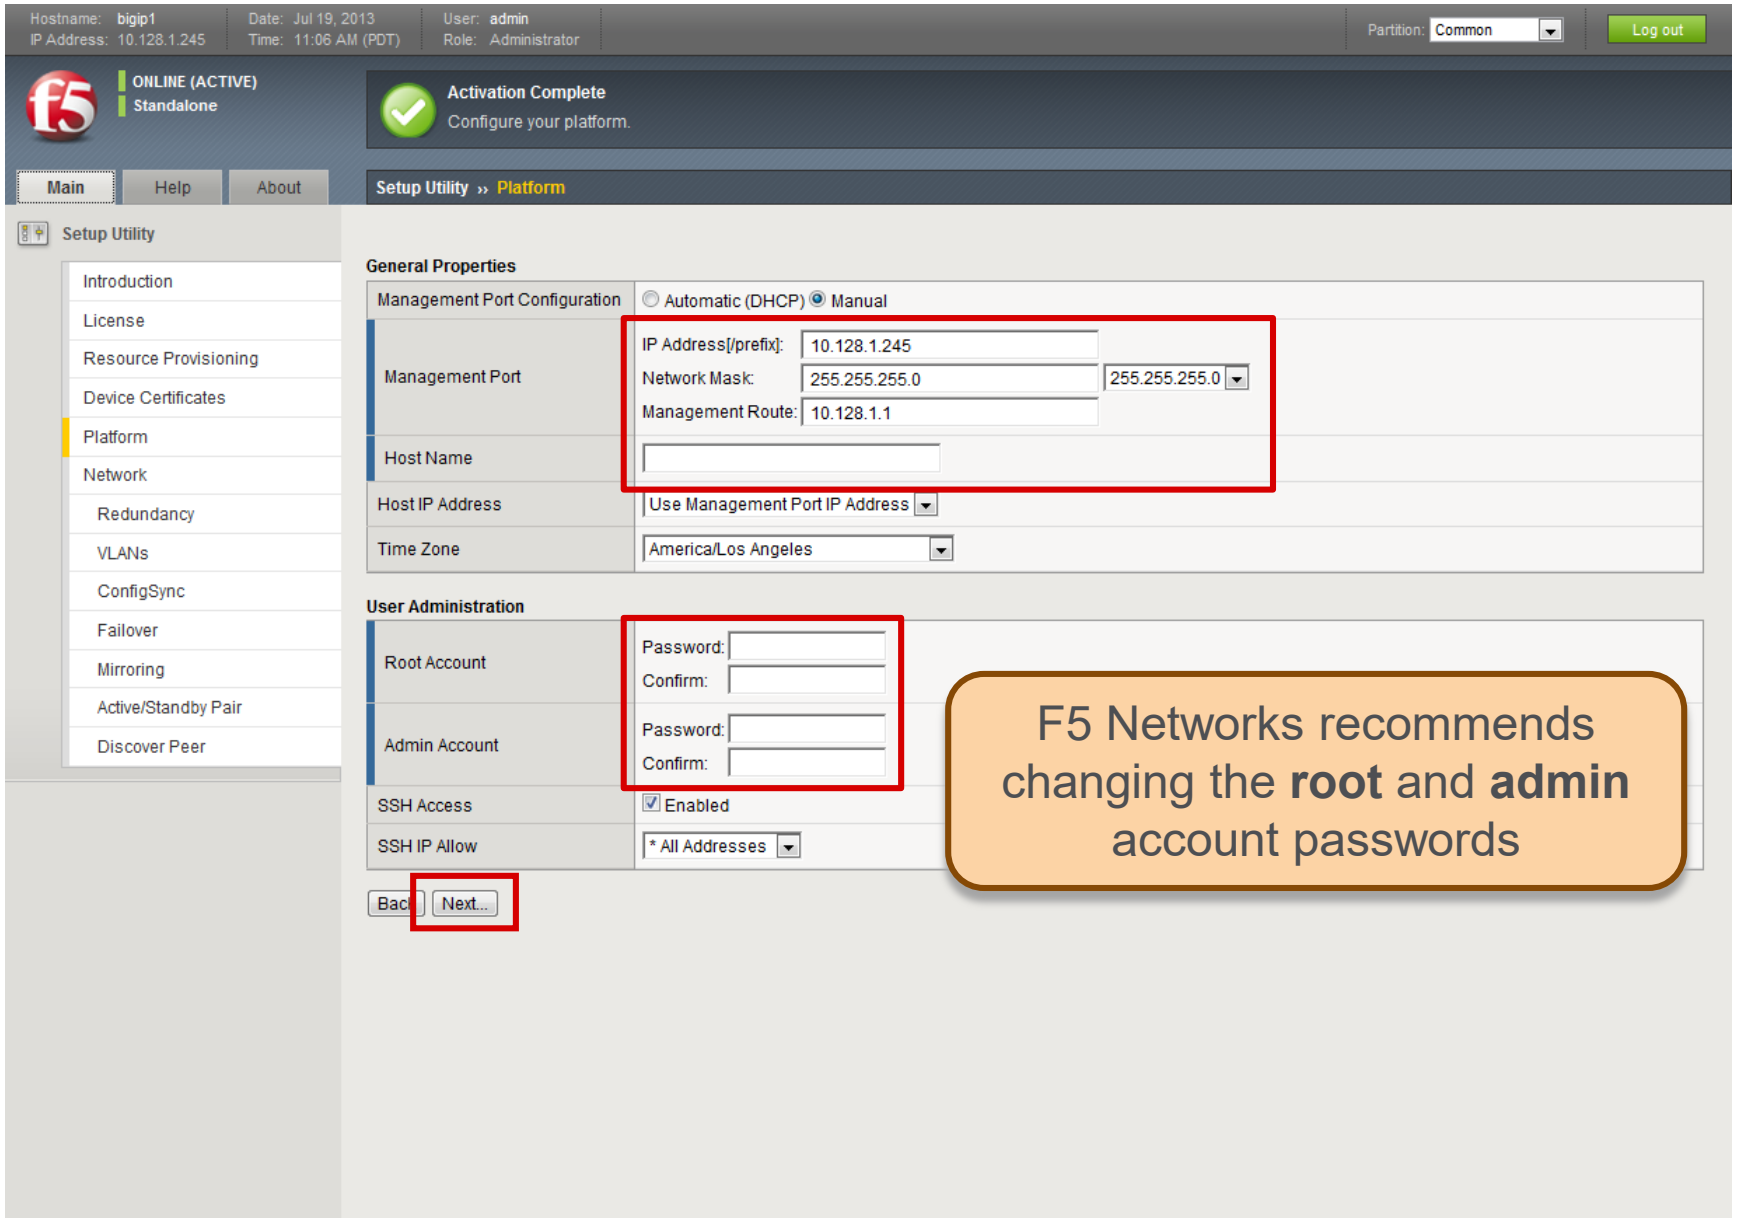

Setup Utility – Standard Network

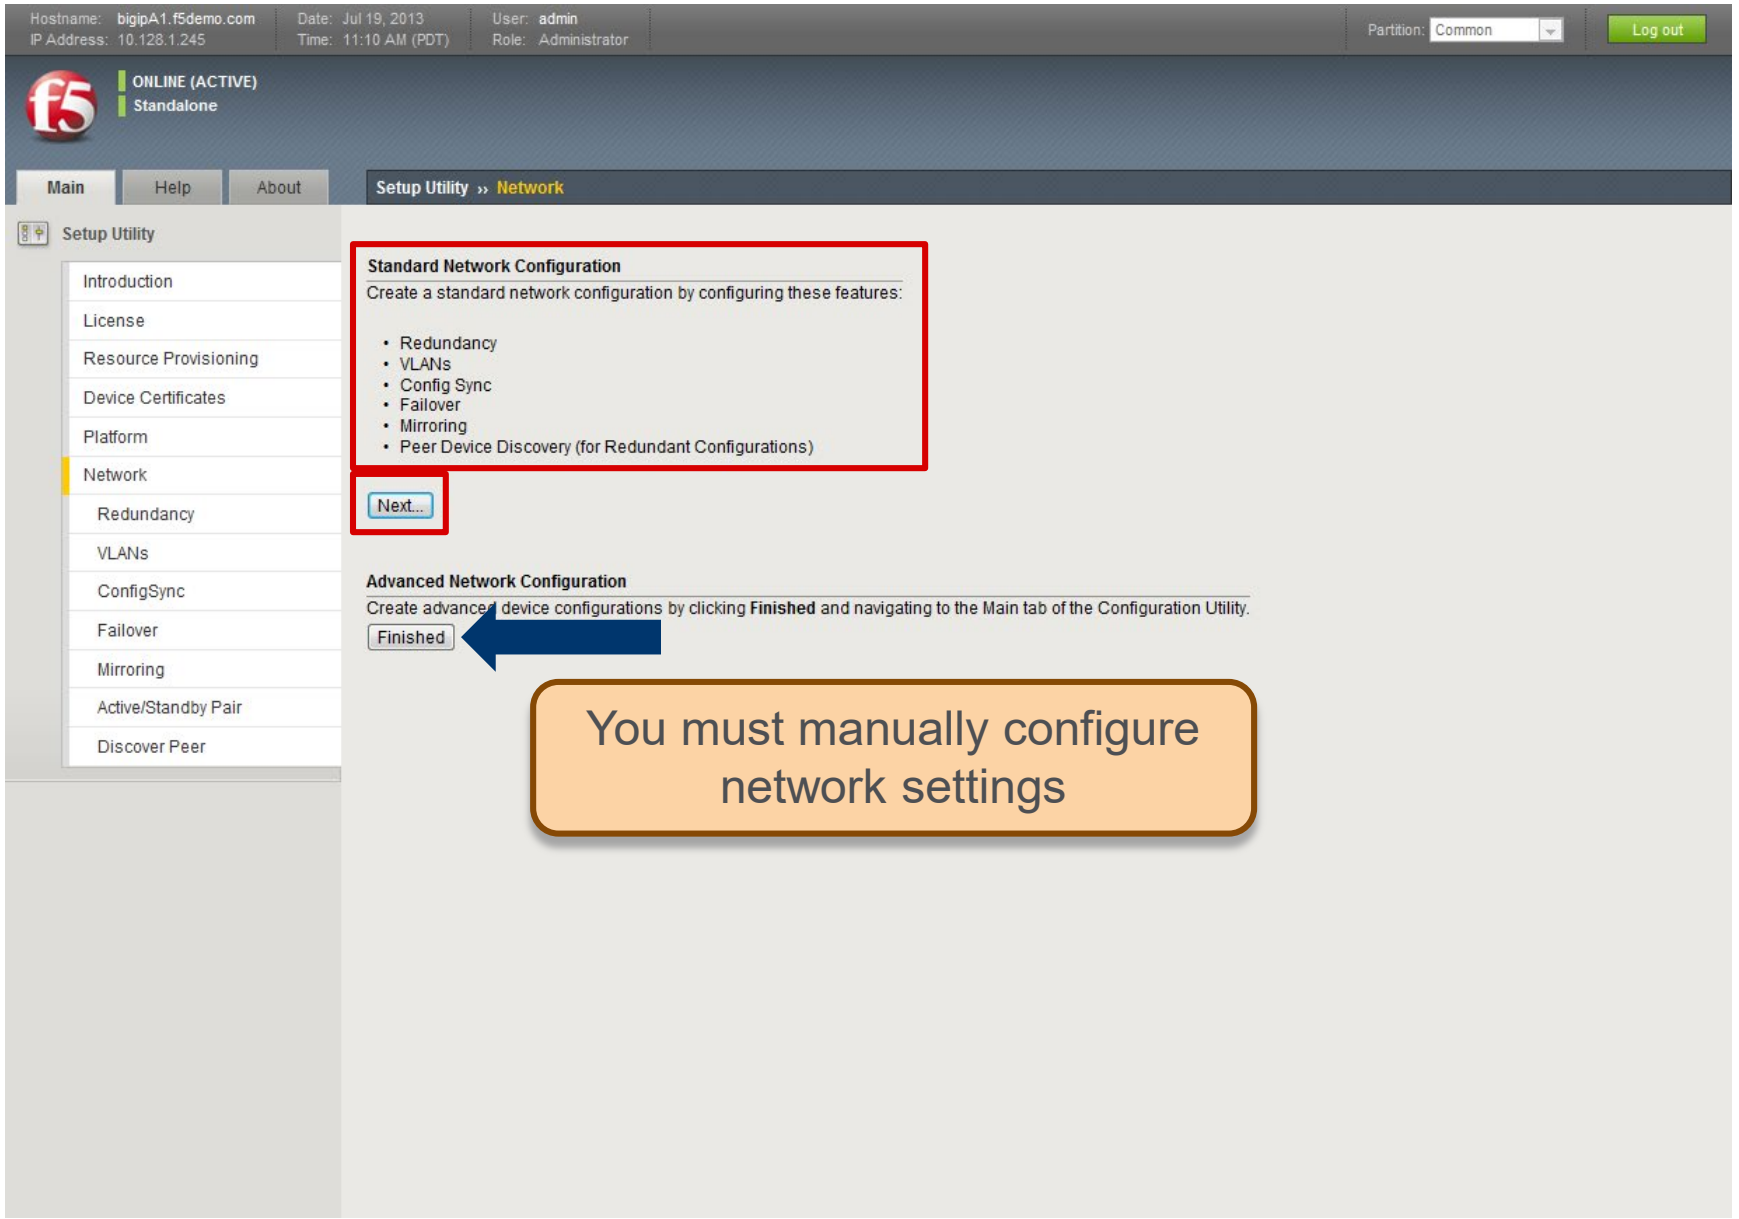

Setup Utility – Internal Network Configuration

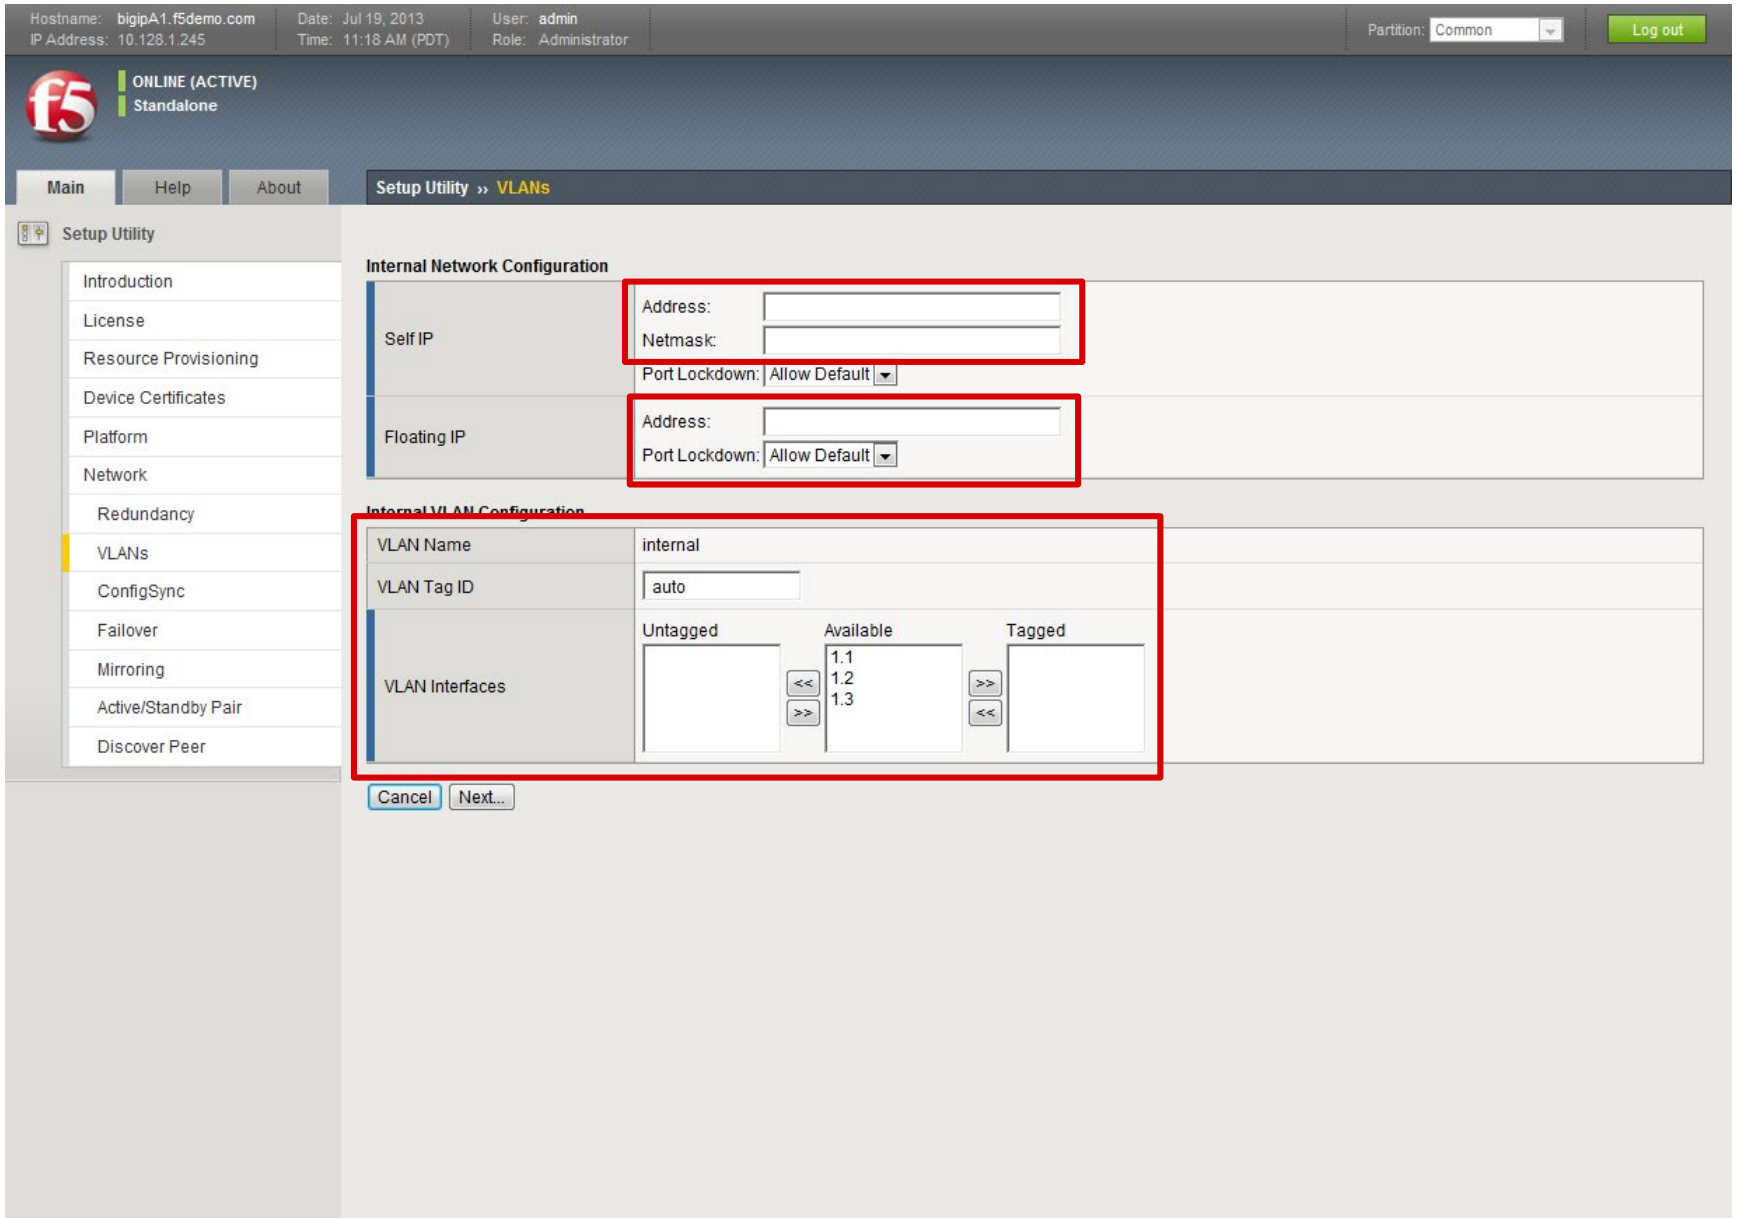

Setup Utility – External Network Configuration

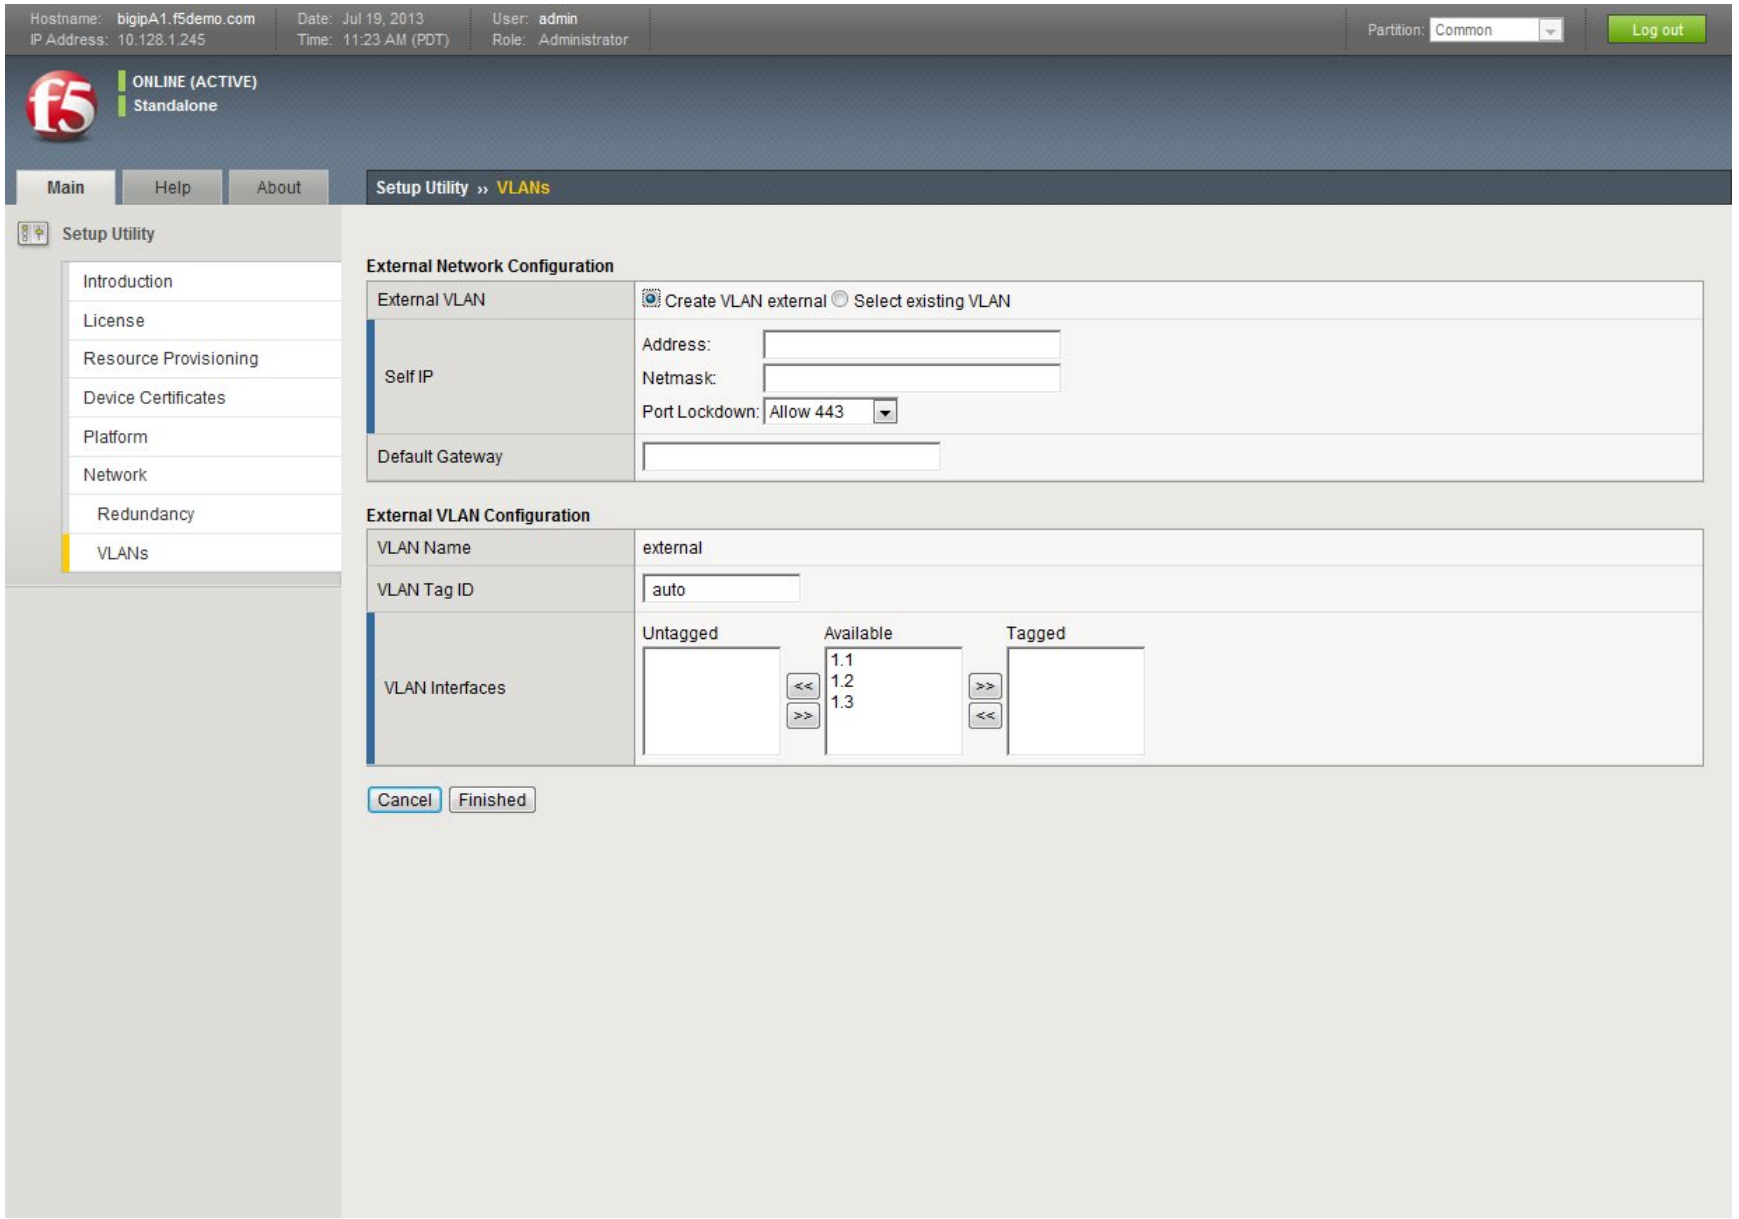

Using the Configuration Utility

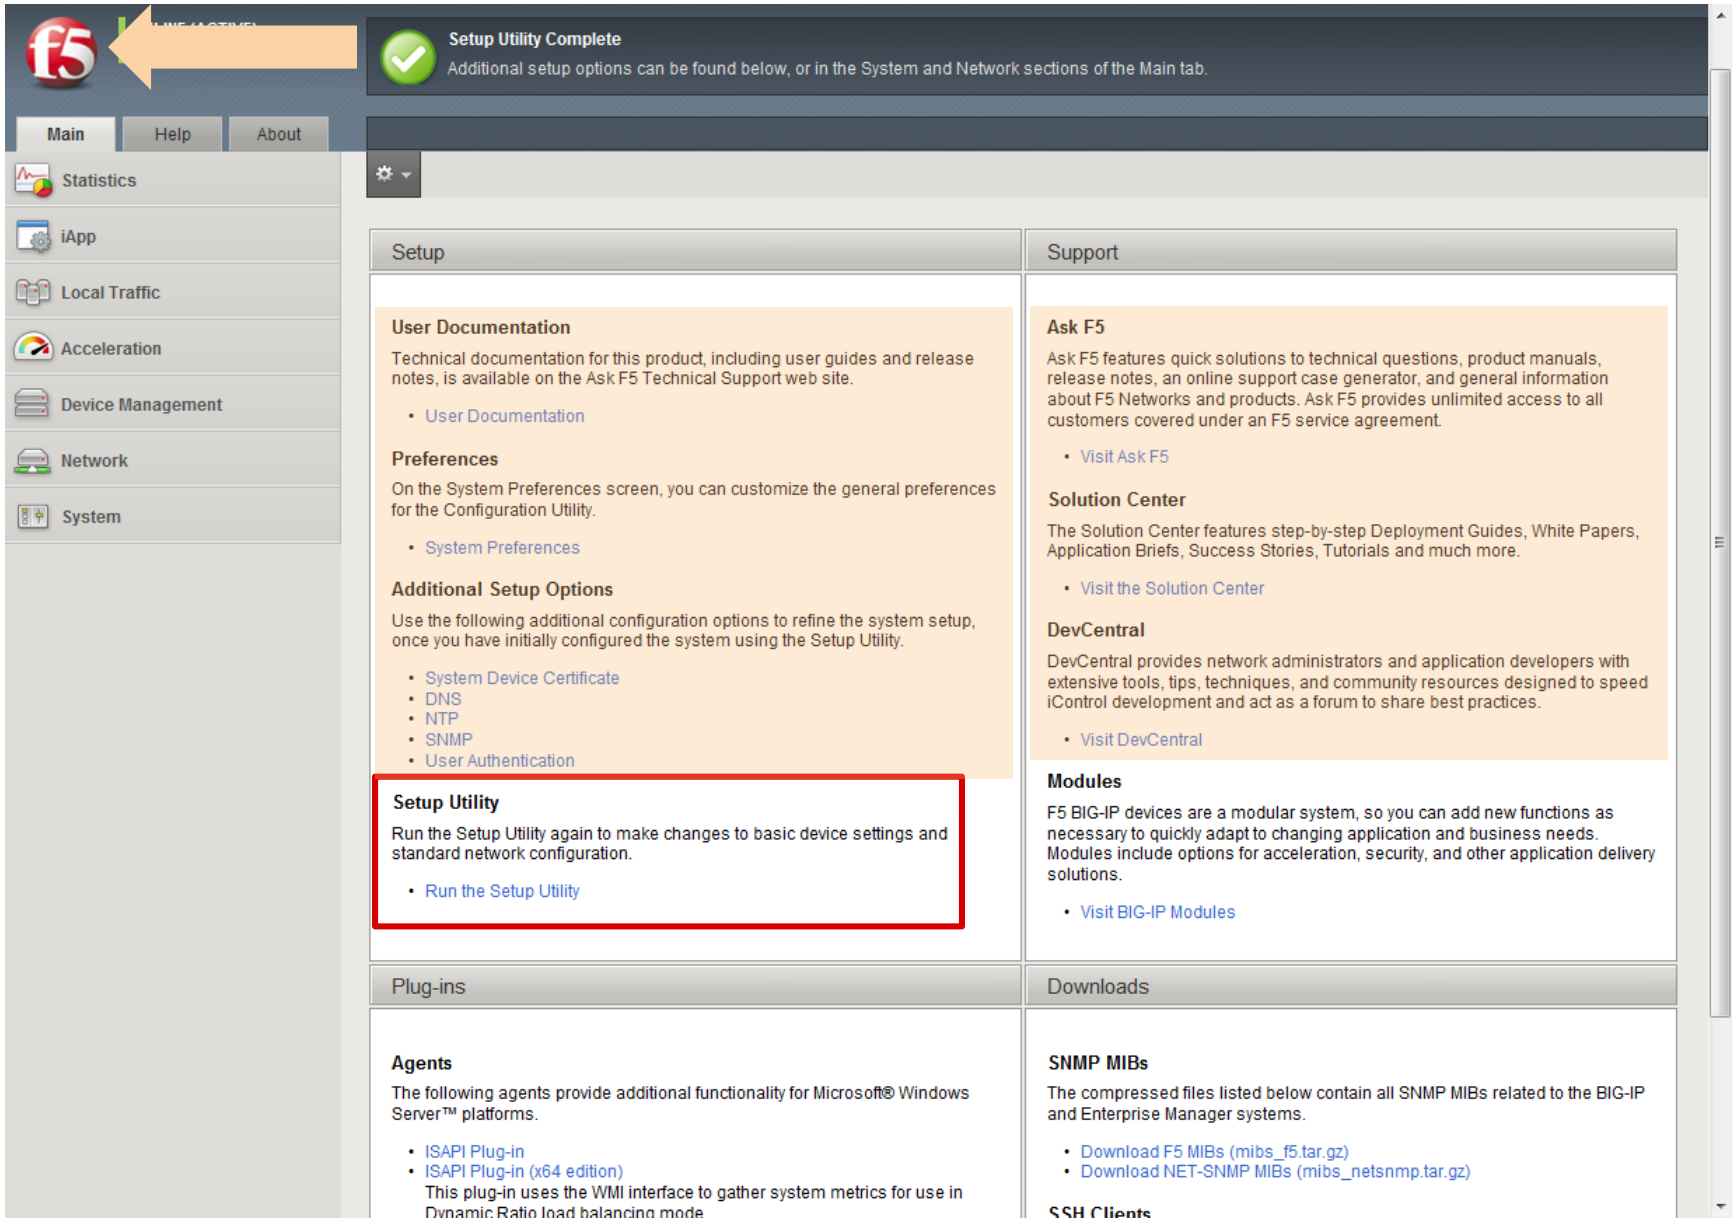

Configuration Utility User Interface

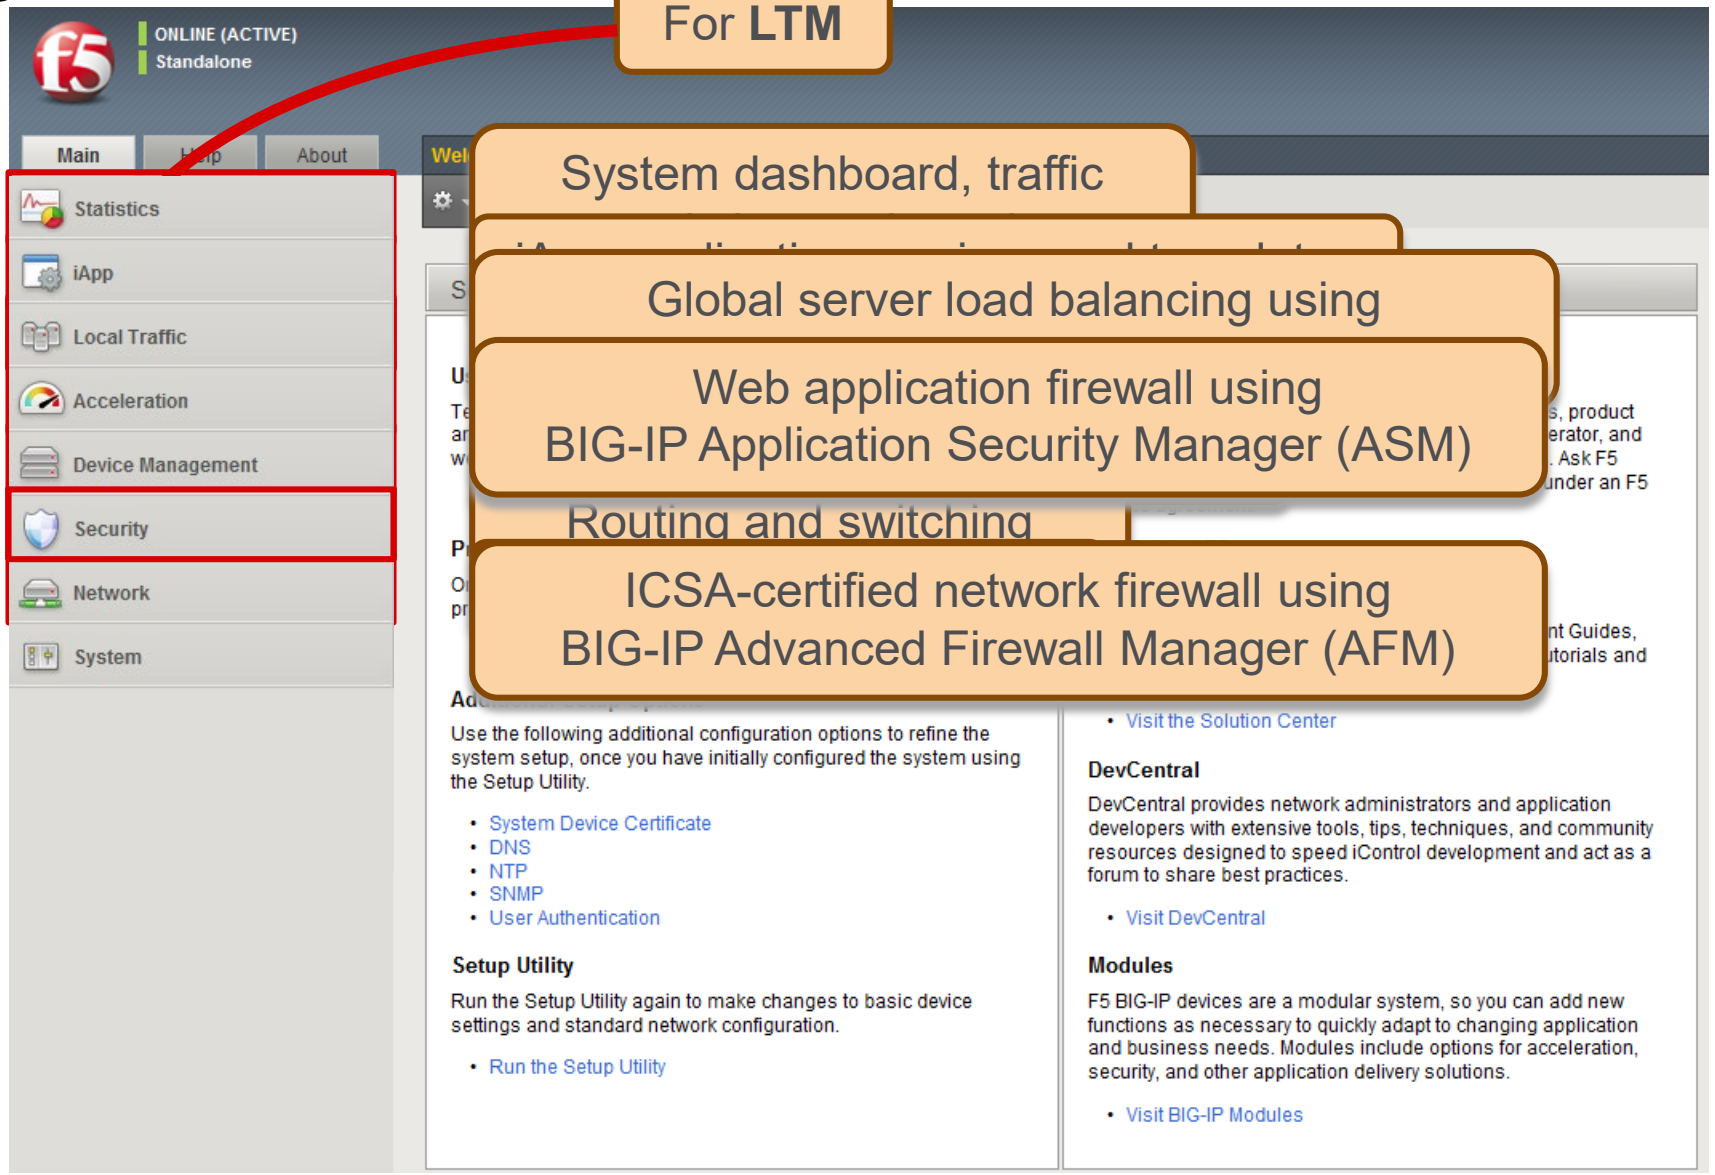

Accessing Archives Page

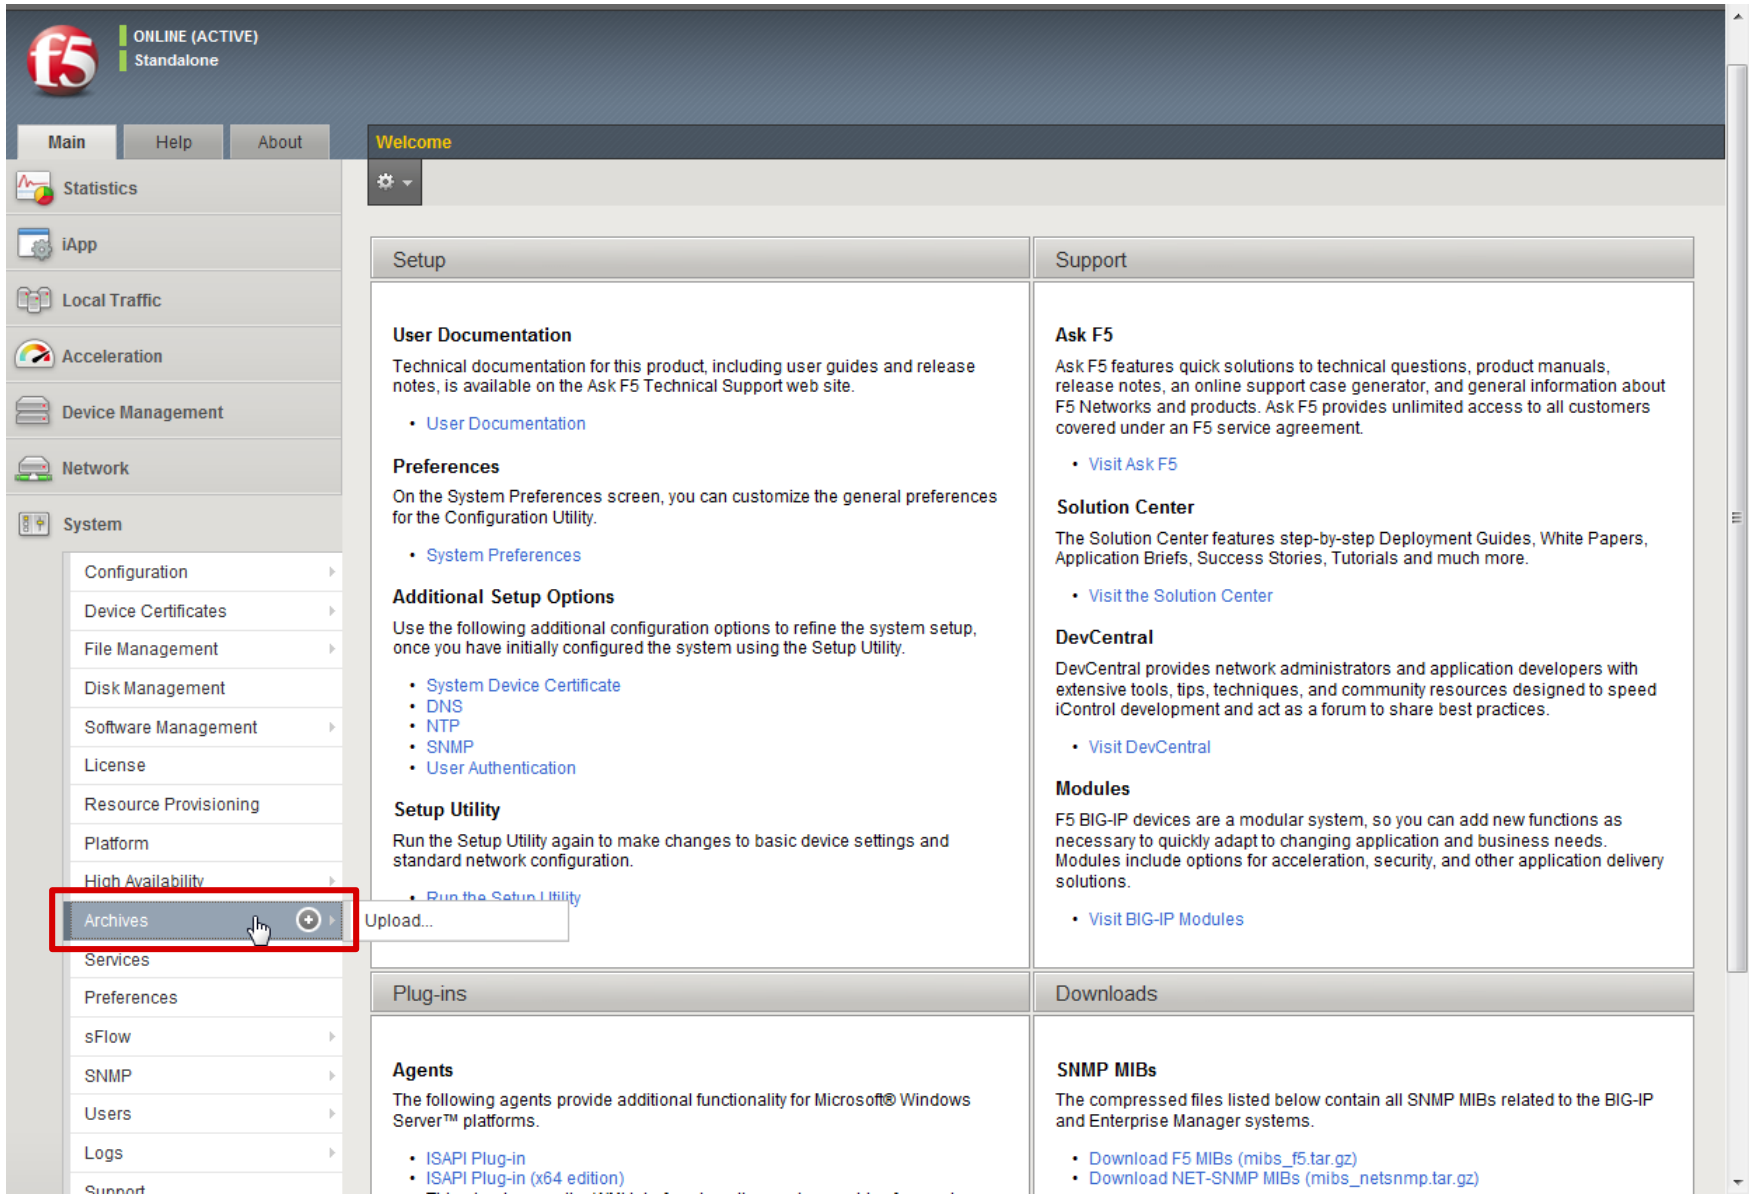

BIG-IP NETWORKING

Overview of Networking

- **TMOS is a full proxy architecture**
  - Traffic must pass through BIG-IP to gain the benefits of TMOS

- **Routed mode (recommended)**
  - Real servers are on an internal network behind the BIG-IP
  - The BIG-IP is default gateway for the servers
  - The virtual servers are on a external network
    - Accessible by the clients

- **Source Network Address Translation (SNAT) Mode**
  - Also know as, One-Armed mode
  - Allows a BIG-IP to be inserted into existing networks without changing the existing IP address structure

How Routed Mode Works

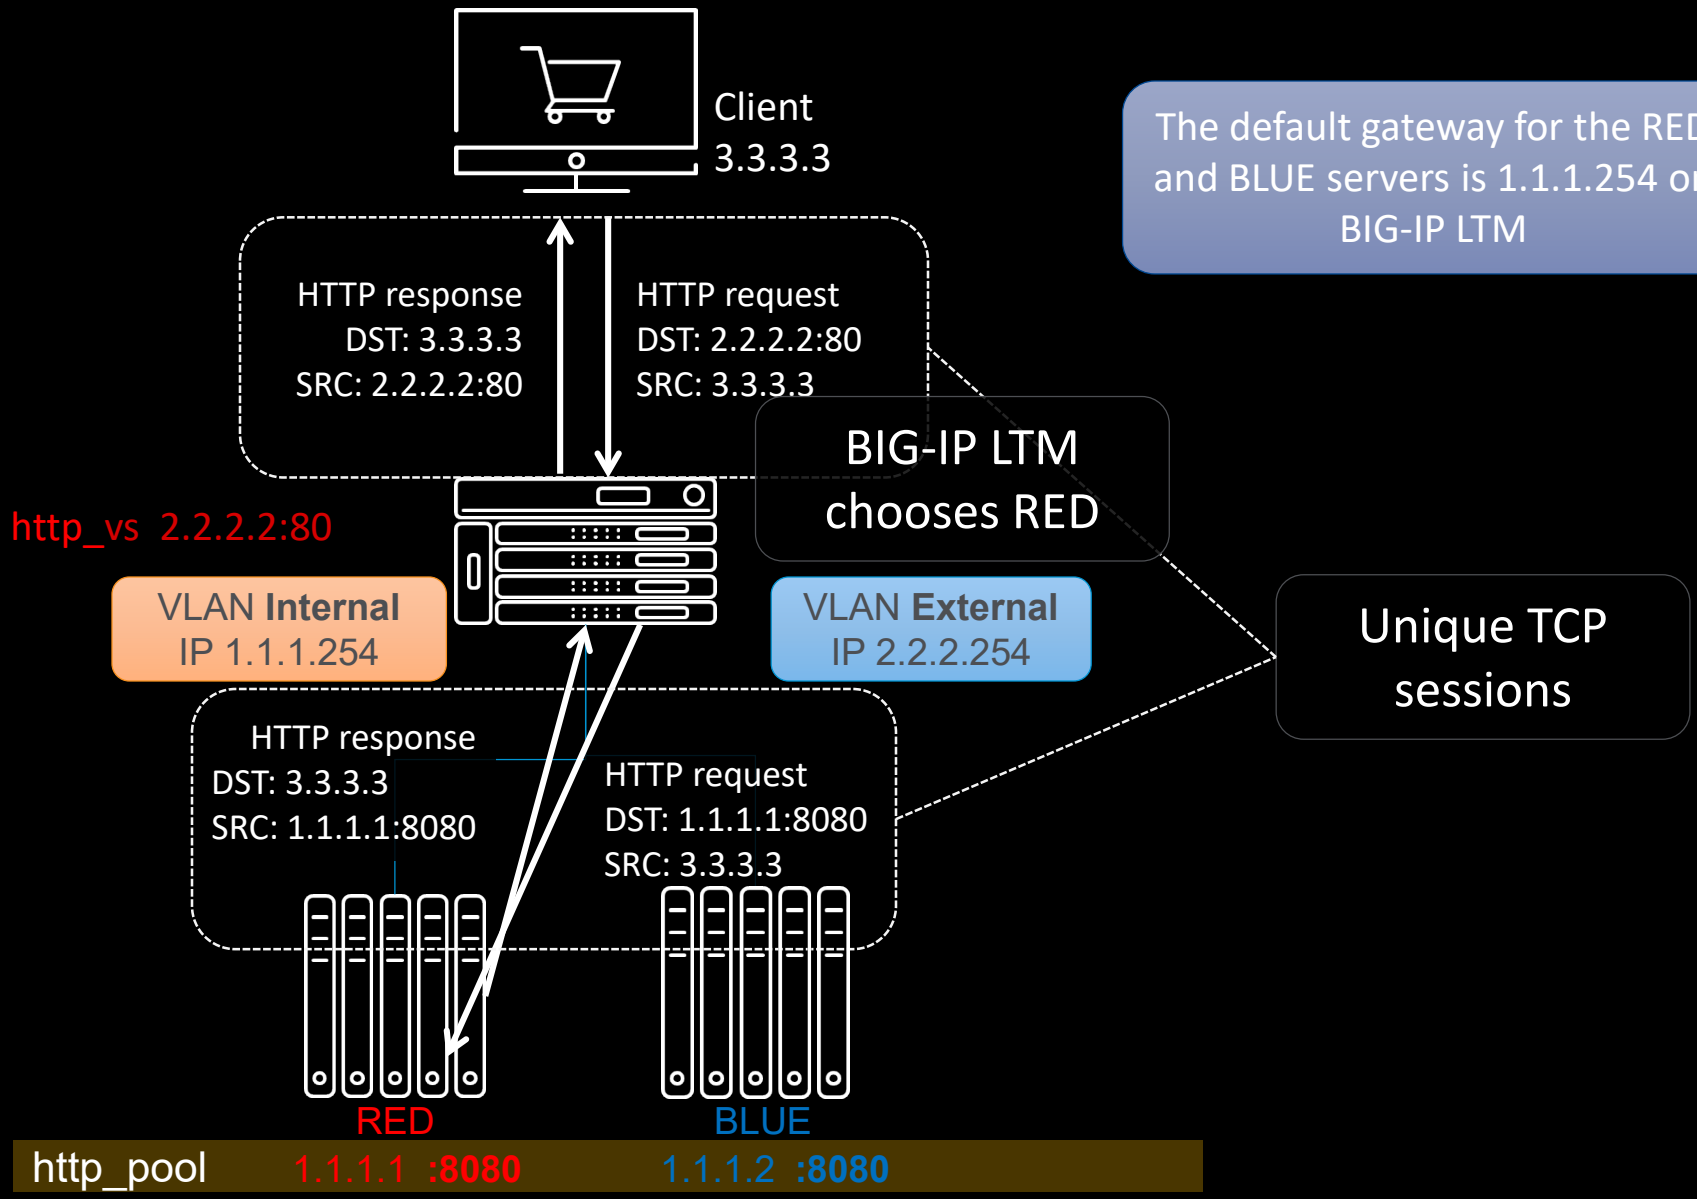

SNAT or One Armed Setup

- **SNATs translate the source IP to an IP owned by BIG-IP**

- **Advantages**
  - No changes to your servers or network
  - Easy option for quick proof of concept testing
  - OneConnect can operate in its most aggressive mode
  - Requires no BIG-IP configuration for direct access to real servers

- **Disadvantages**
  - Servers see the BIG-IP as the source IP, you lose the original source IP
    - May impact logging mechanisms and harder to debug
  - Allows direct access to servers, which reduces security

- **The BIG-IP can insert the original IP address into the headers using:**
  - An x-forward for the header through an HTTP profile (RFC compliant)
  - A custom header through iRules

Why SNATs May be Required

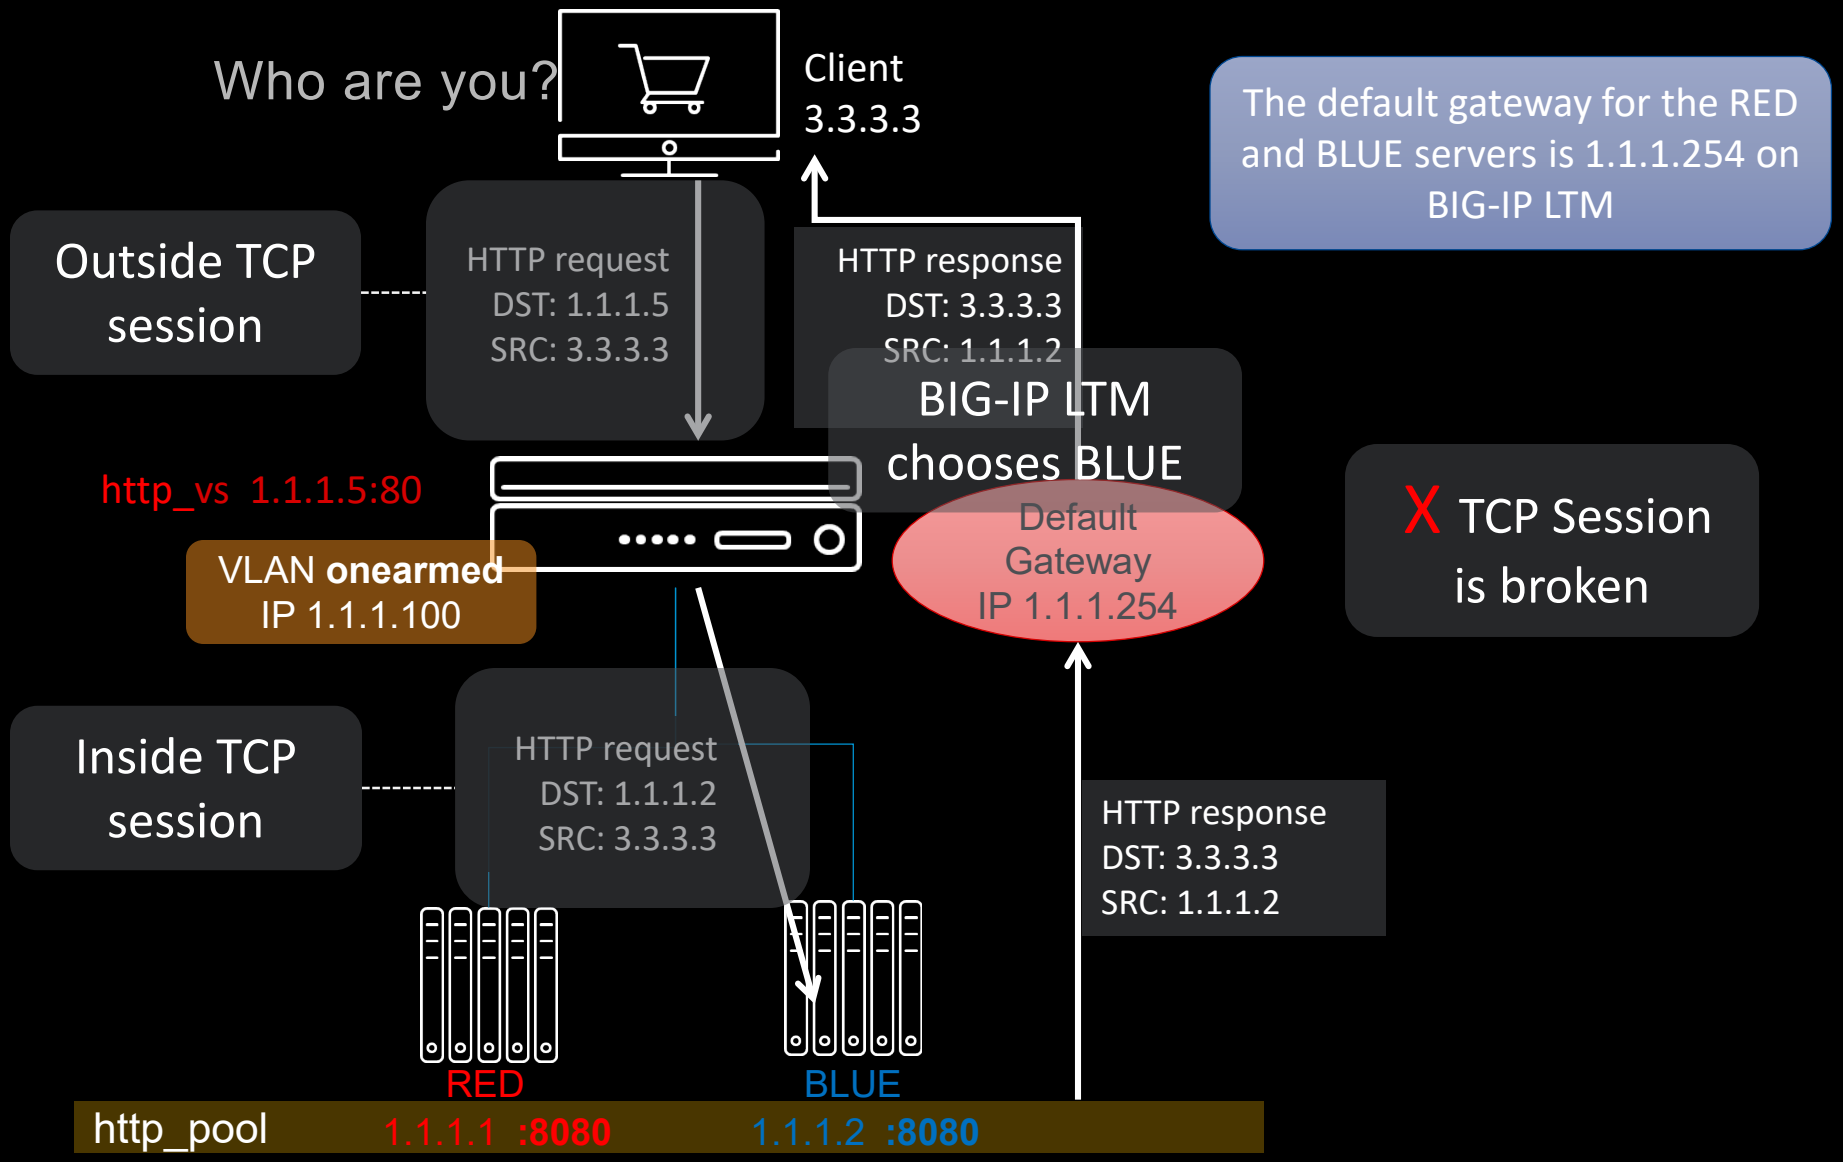

How SNAT Mode Works

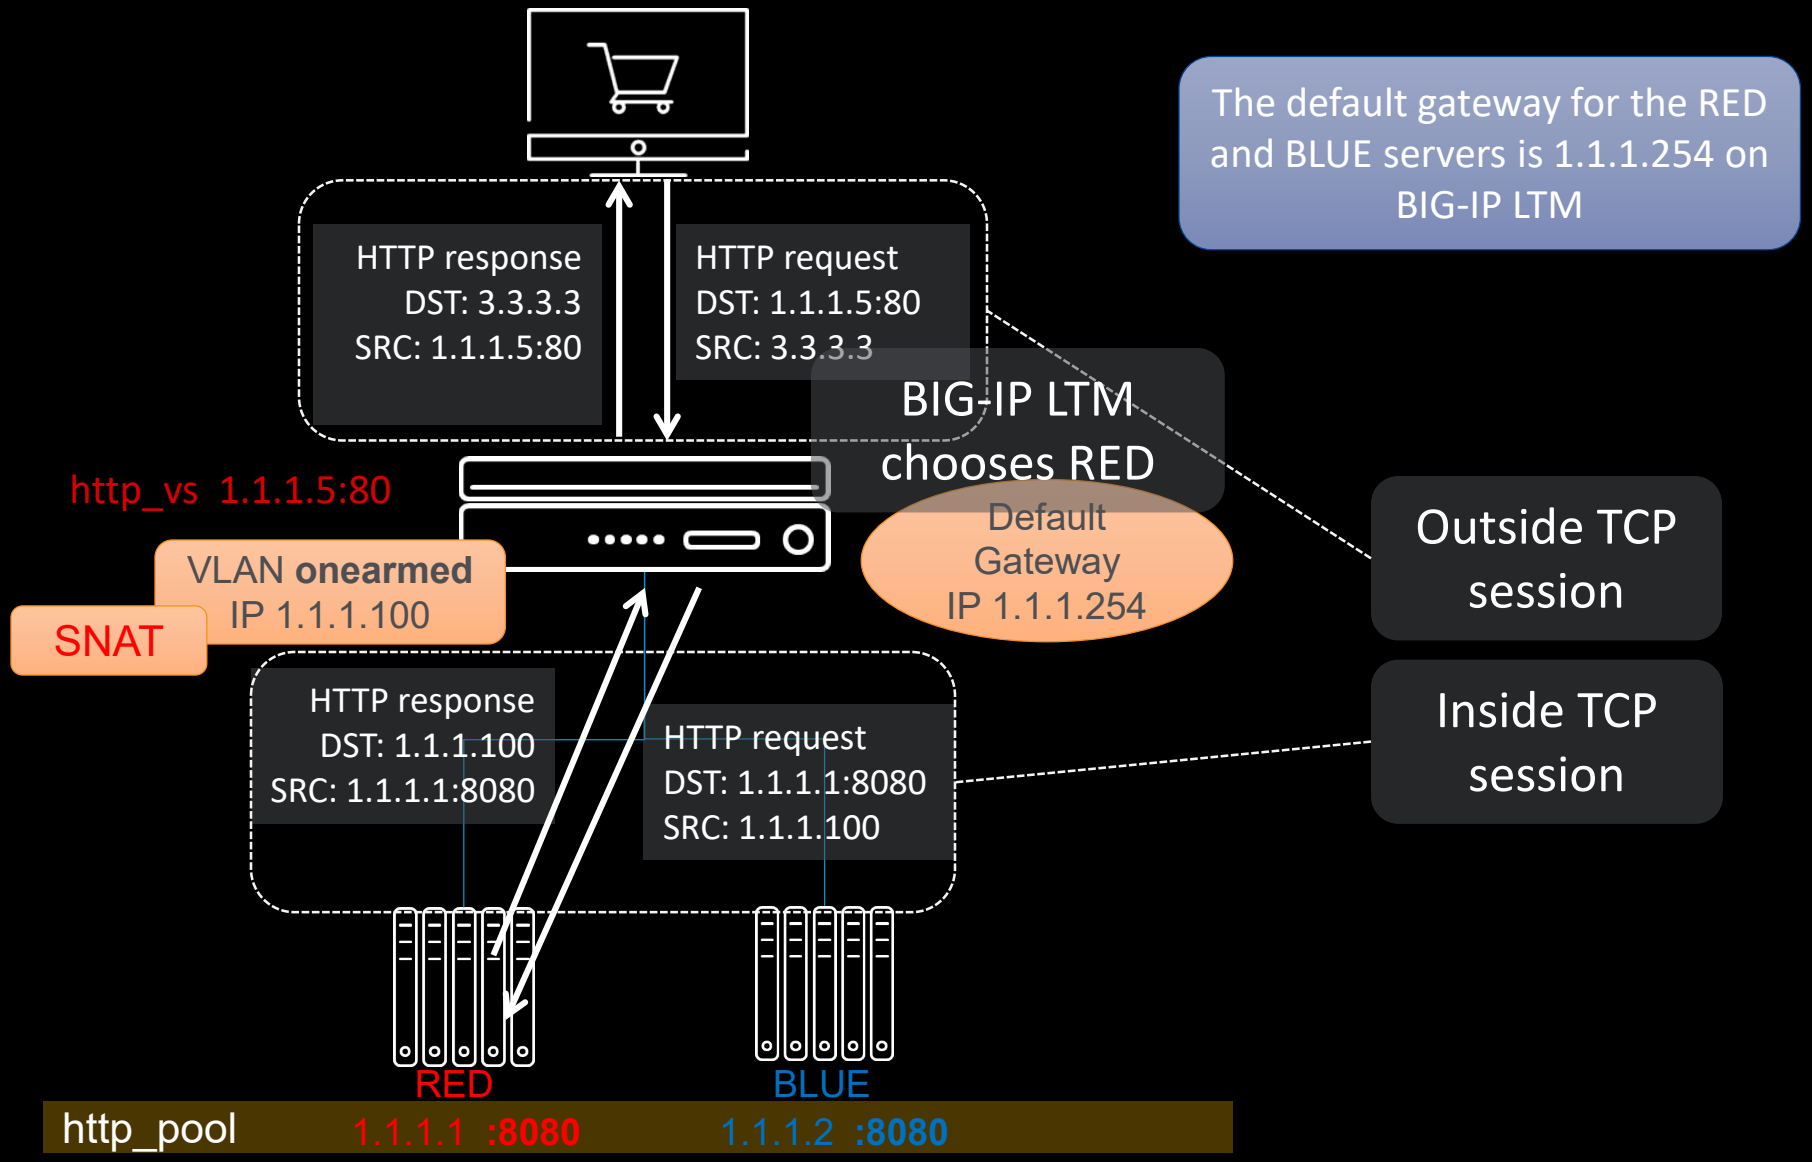

LTM COMPONENTS

BIG-IP LTM Components: Nodes

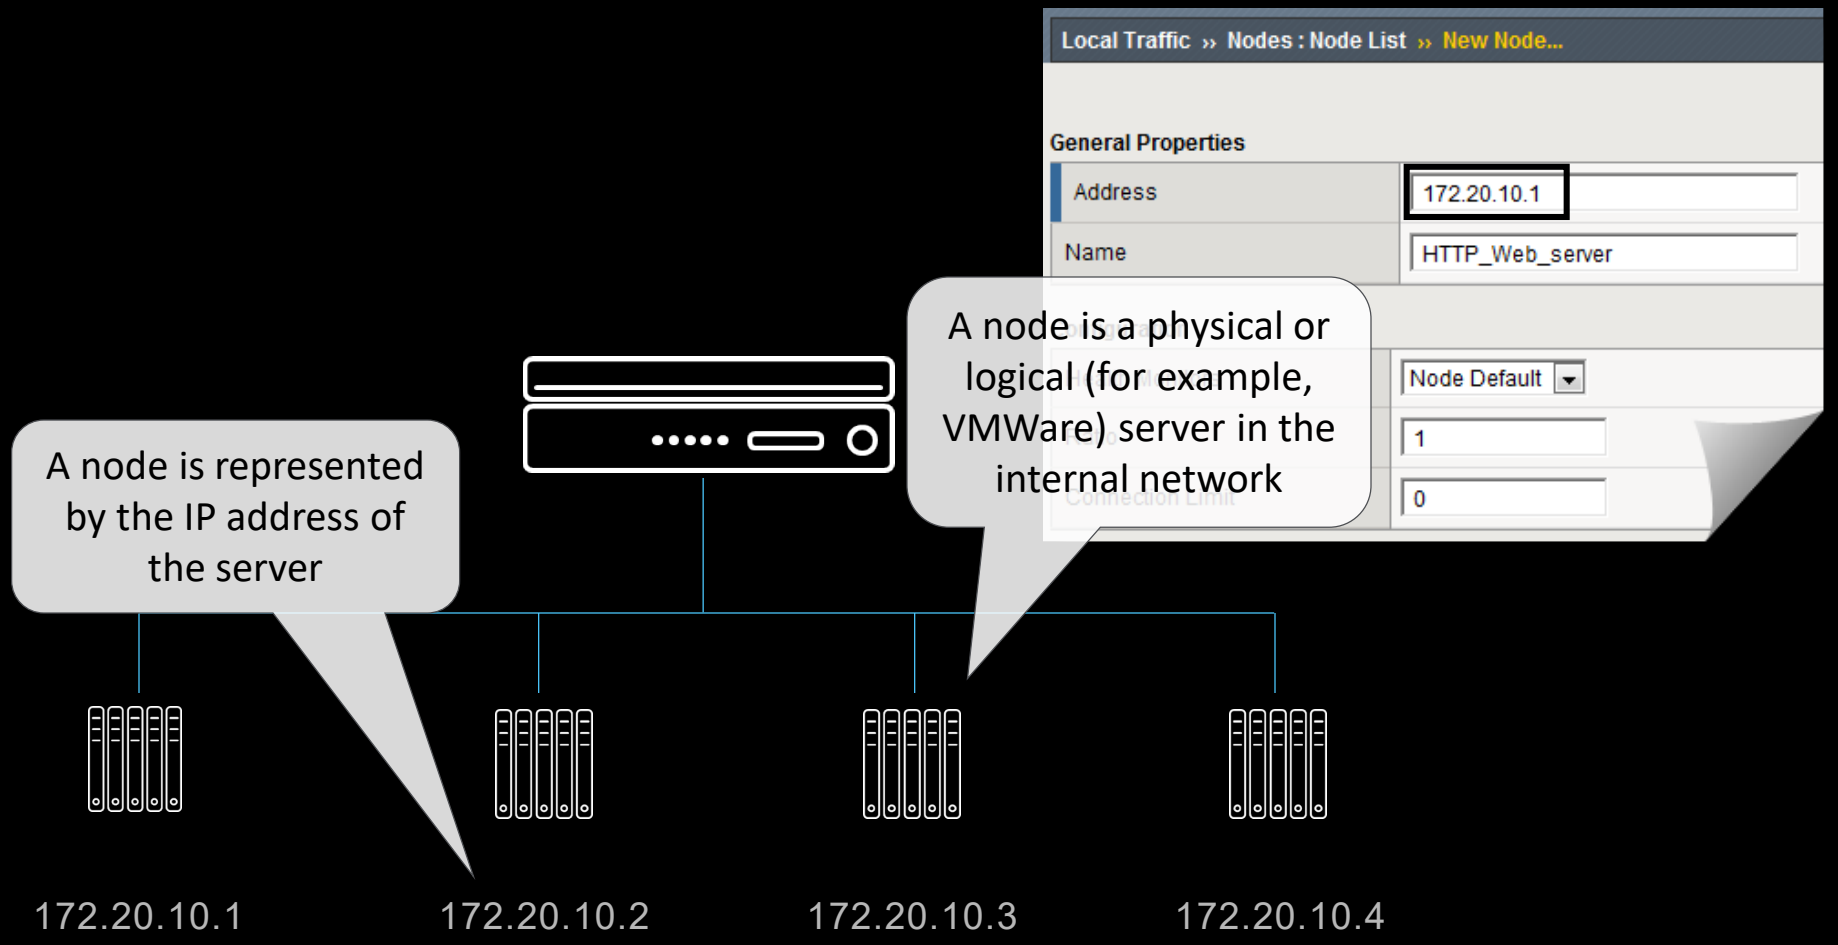

BIG-IP LTM Components: Pool Members

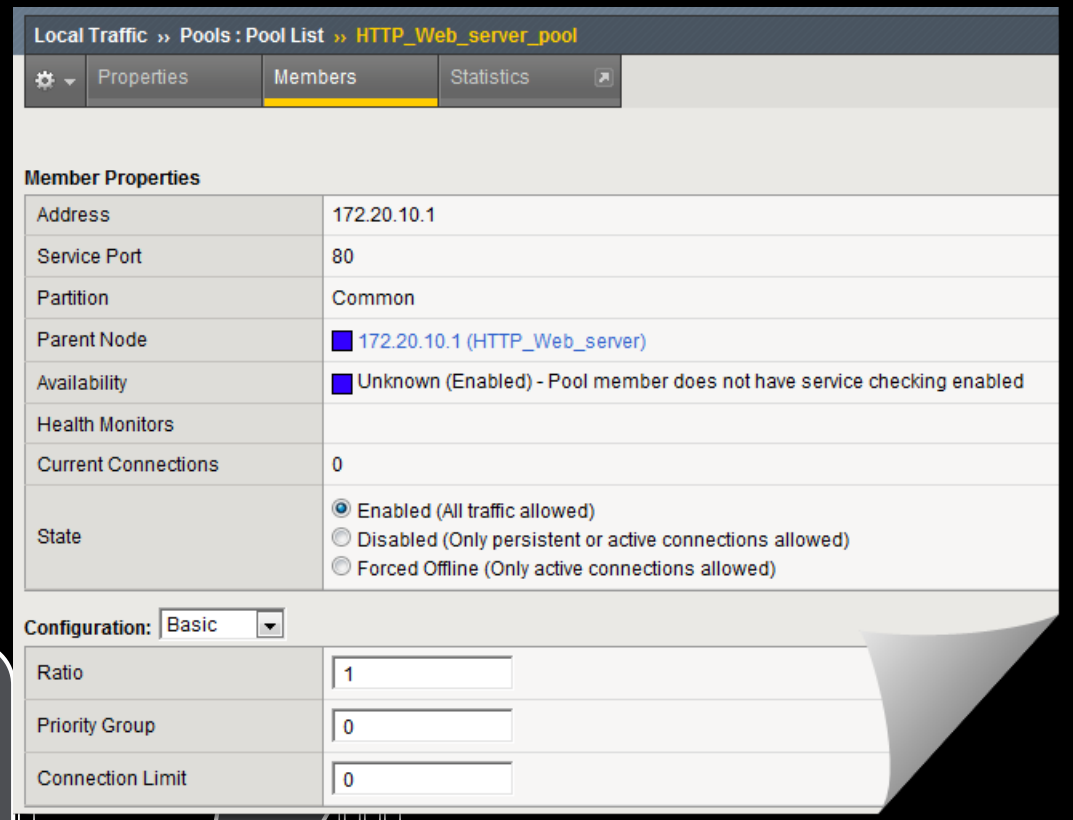

A pool member is a service running on a node, represented by the IP address of the node and service (port) number

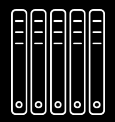

172.20.10.1
172.20.10.1:80

172.20.10.2
172.20.10.2:80
172.20.10.2:443

172.20.10.3
172.20.10.3:80
172.20.10.3:443

172.20.10.4
172.20.10.4:443

BIG-IP LTM Components: Pools

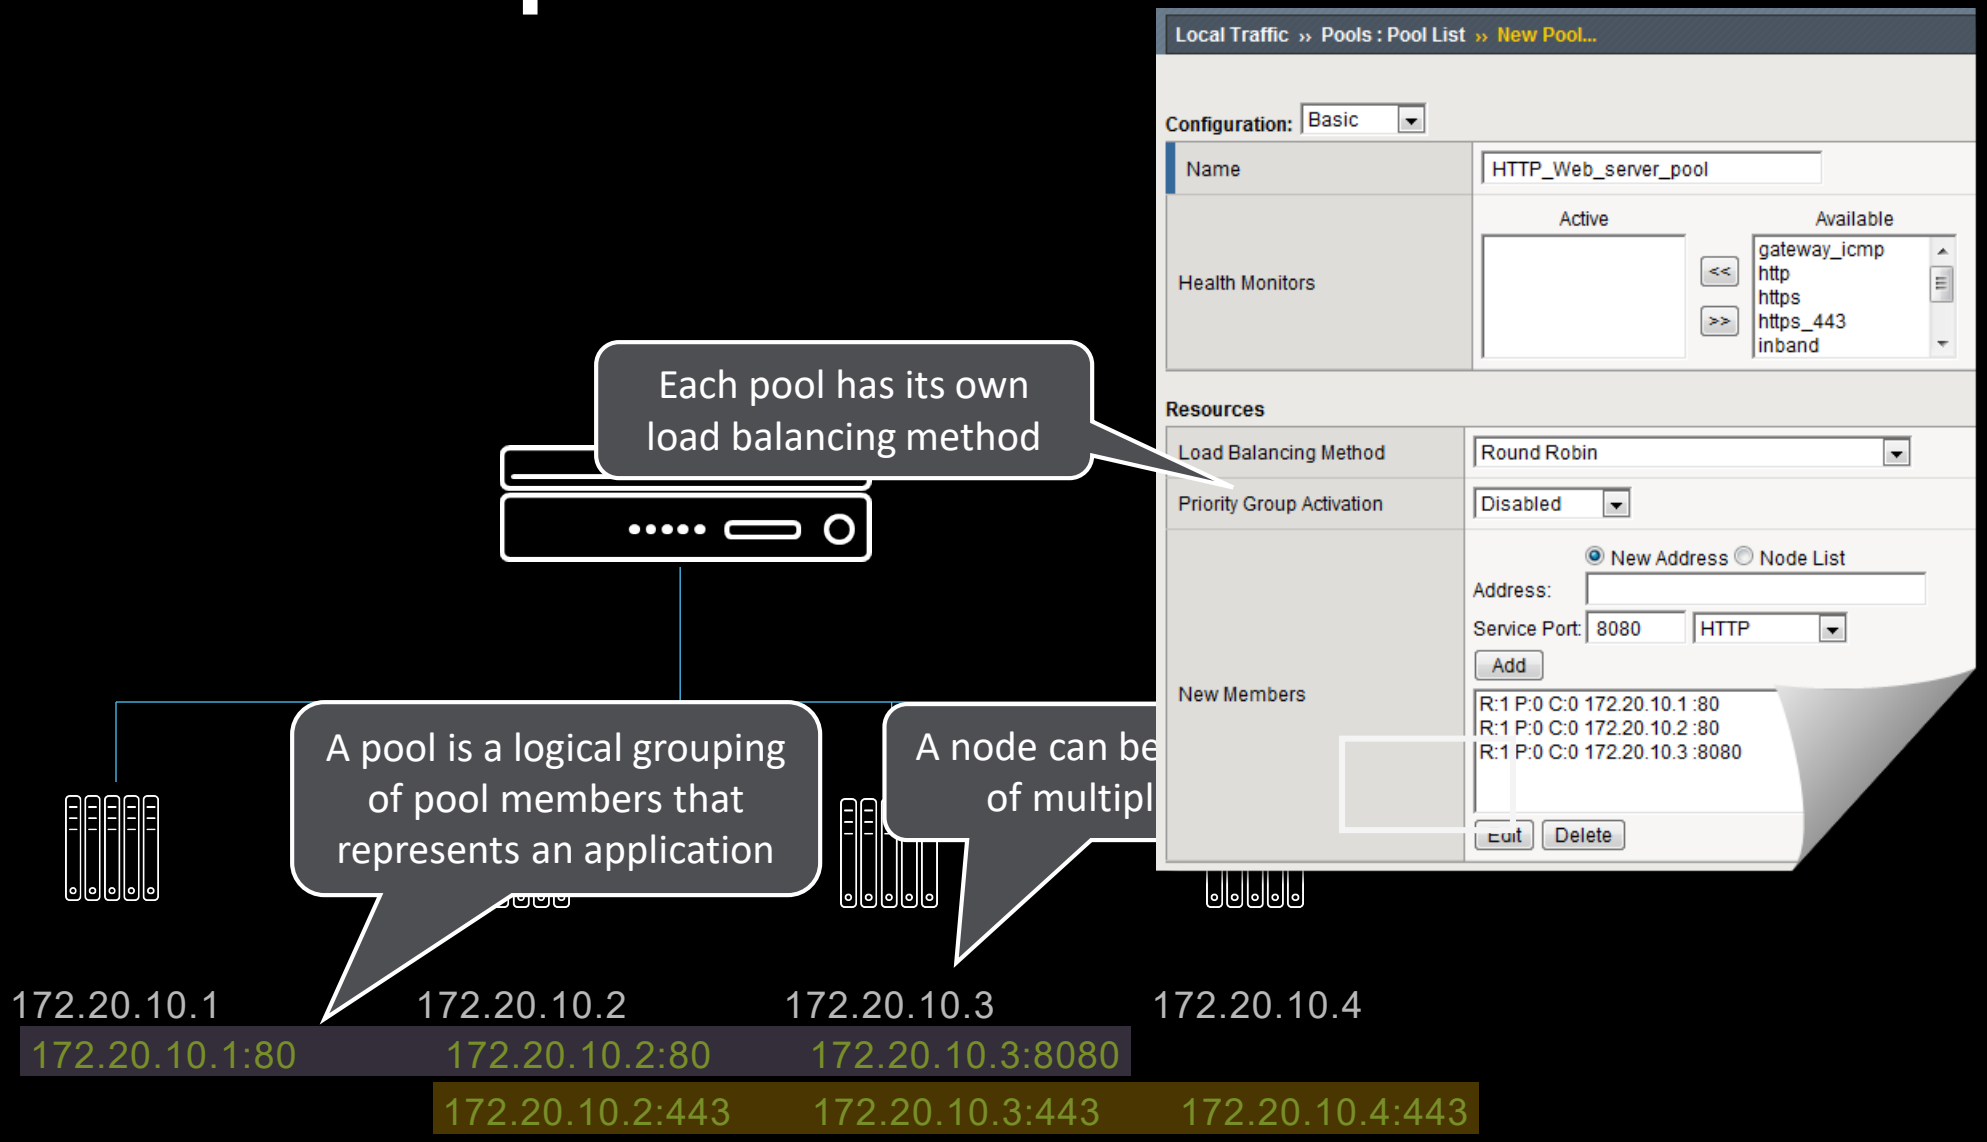

BIG-IP LTM Components: Virtual Servers

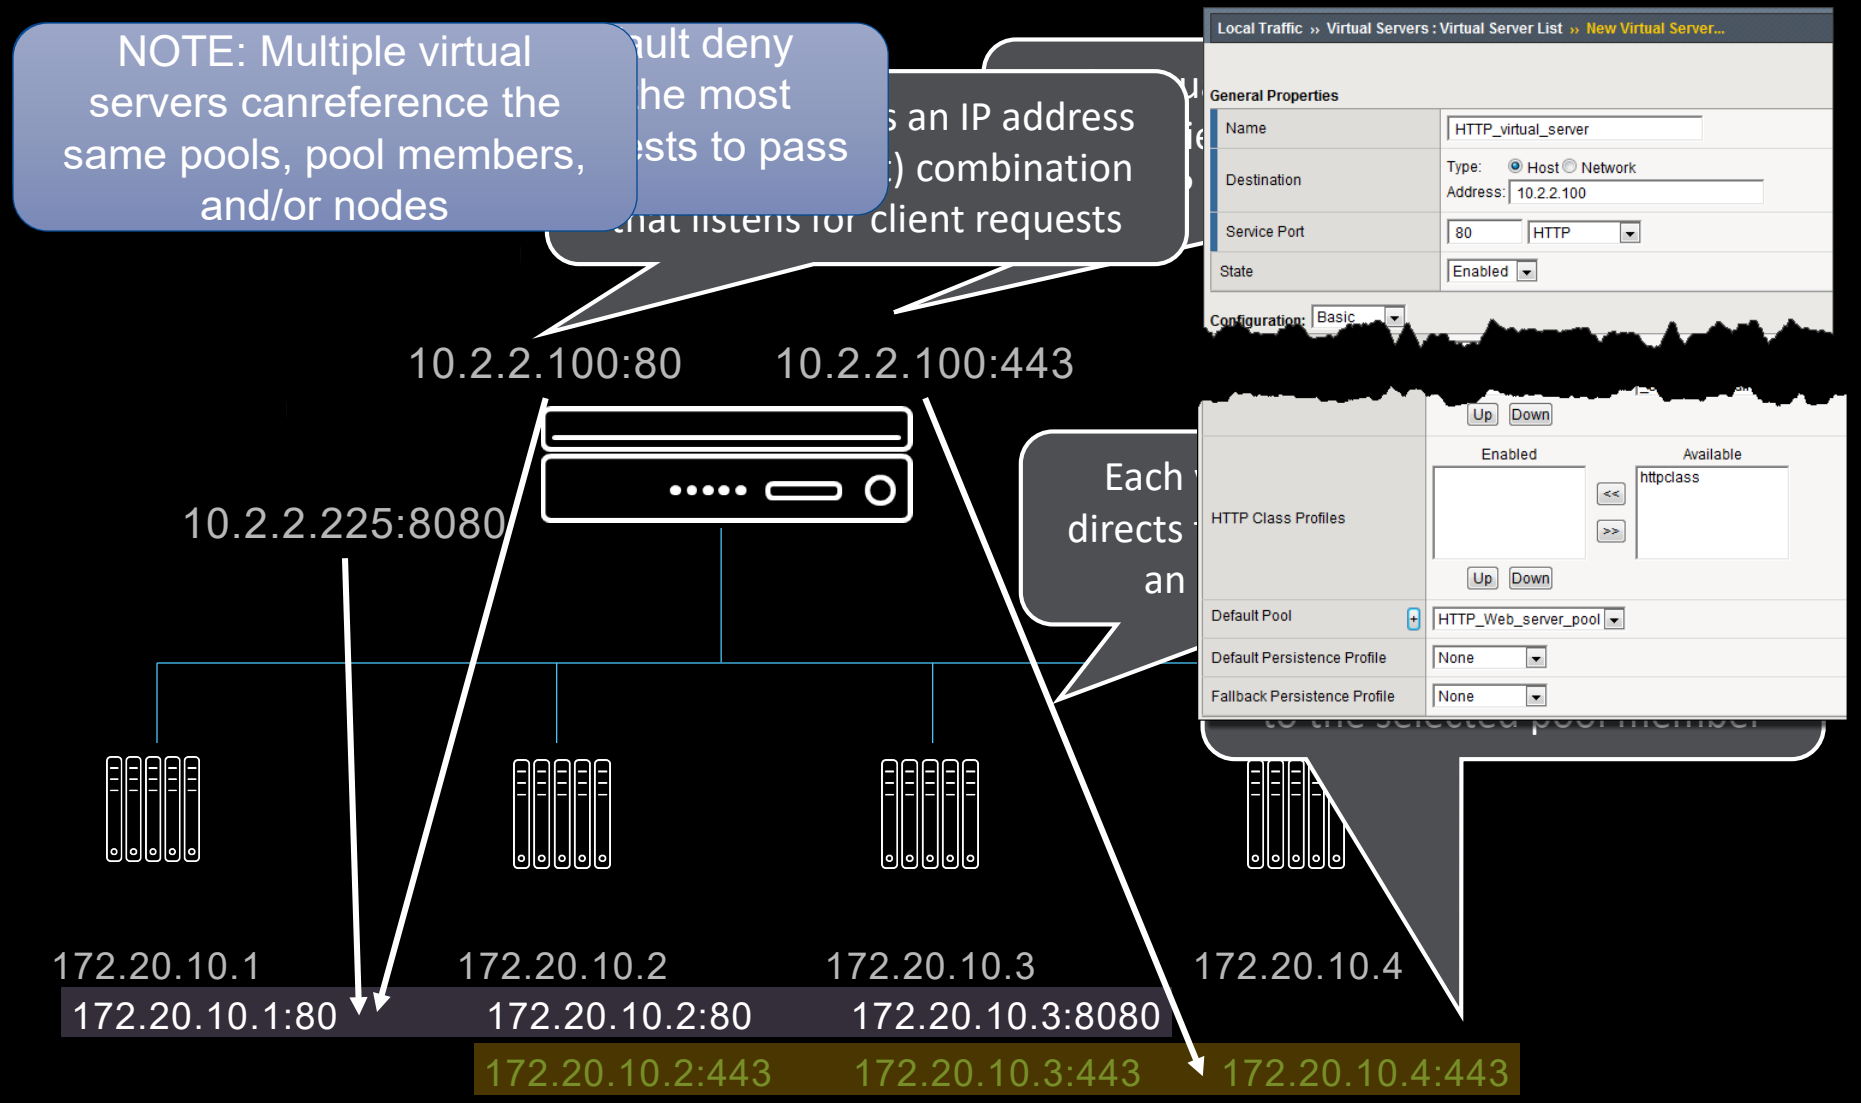

Virtual Servers

- **One of the most important configuration components**
  - Determines what traffic is to pass
  - Where the traffic goes
  - How it is viewed/manipulated/validated (mostly via profiles)
  - So in the last slides we saw virtual server basics (in and out) …..

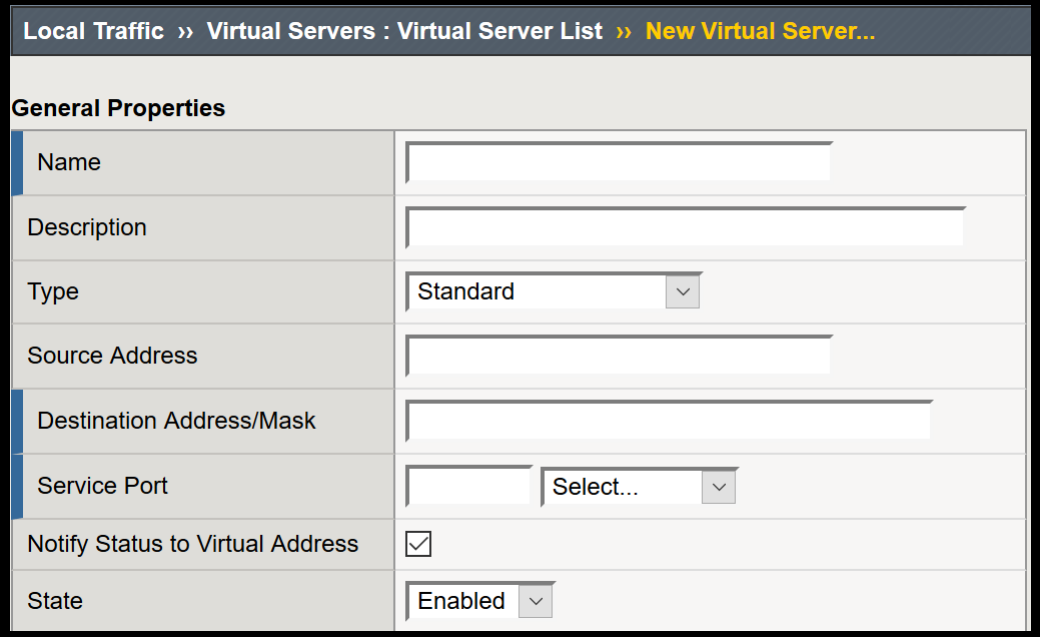

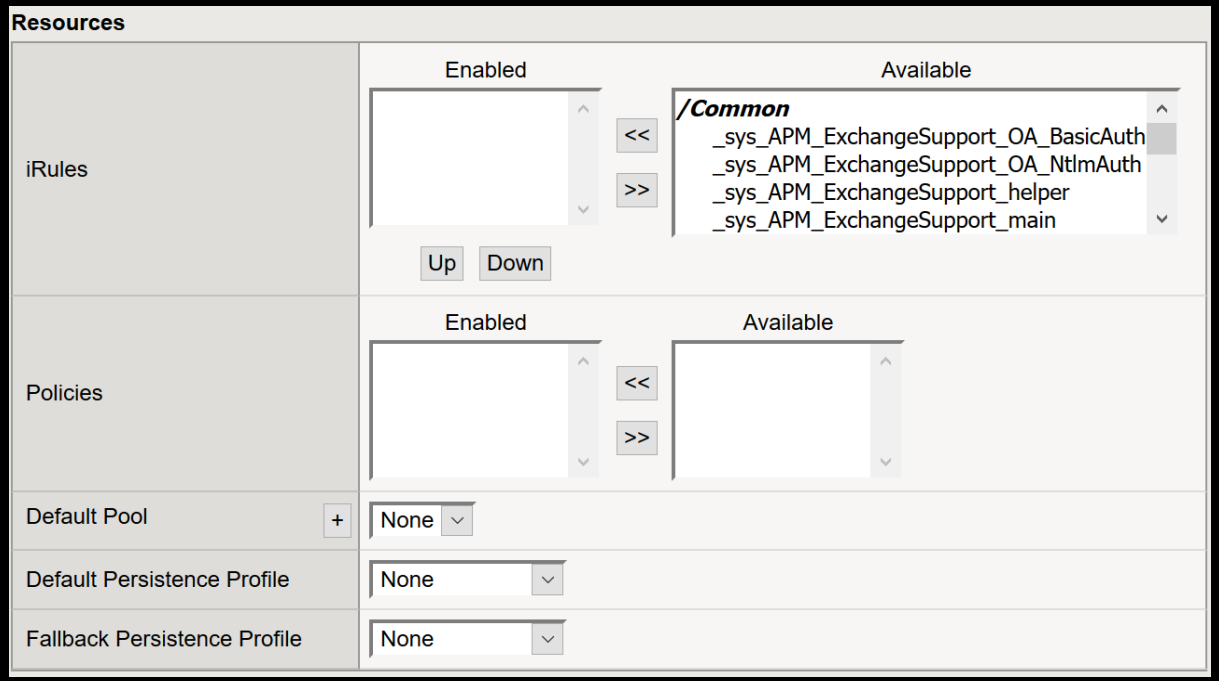

But there is so much more….

- **And this is just the basic menu…**
  - Layer 4-7 profiles
  - Restrictions on traffic
  - Source Address Translation
    - Destination IP and Port Translation in Advanced

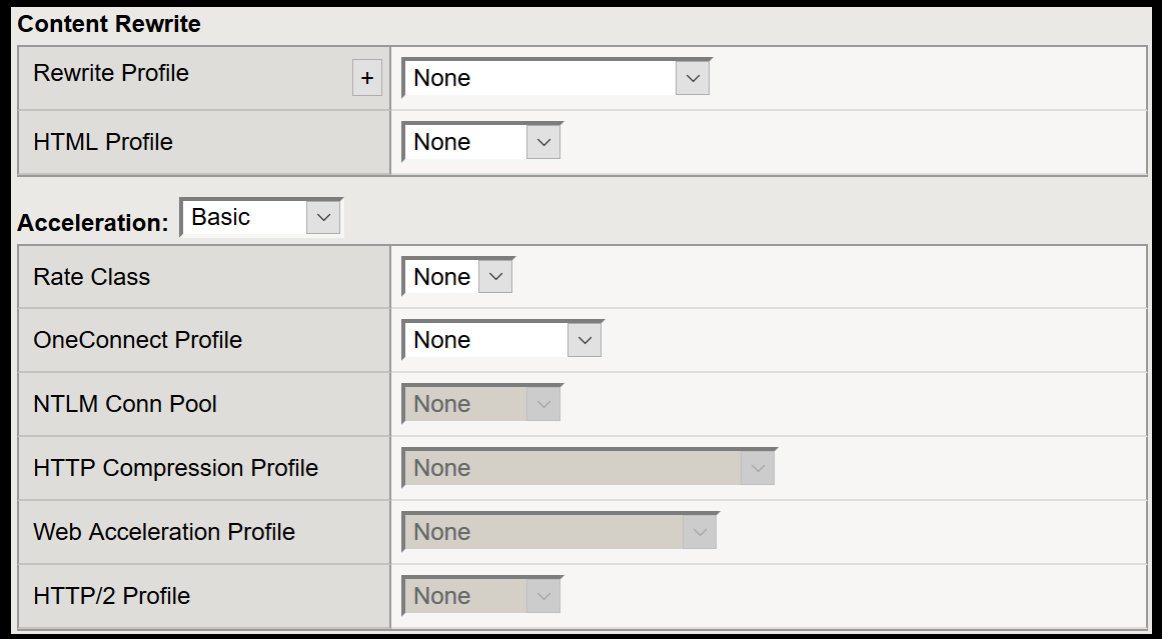

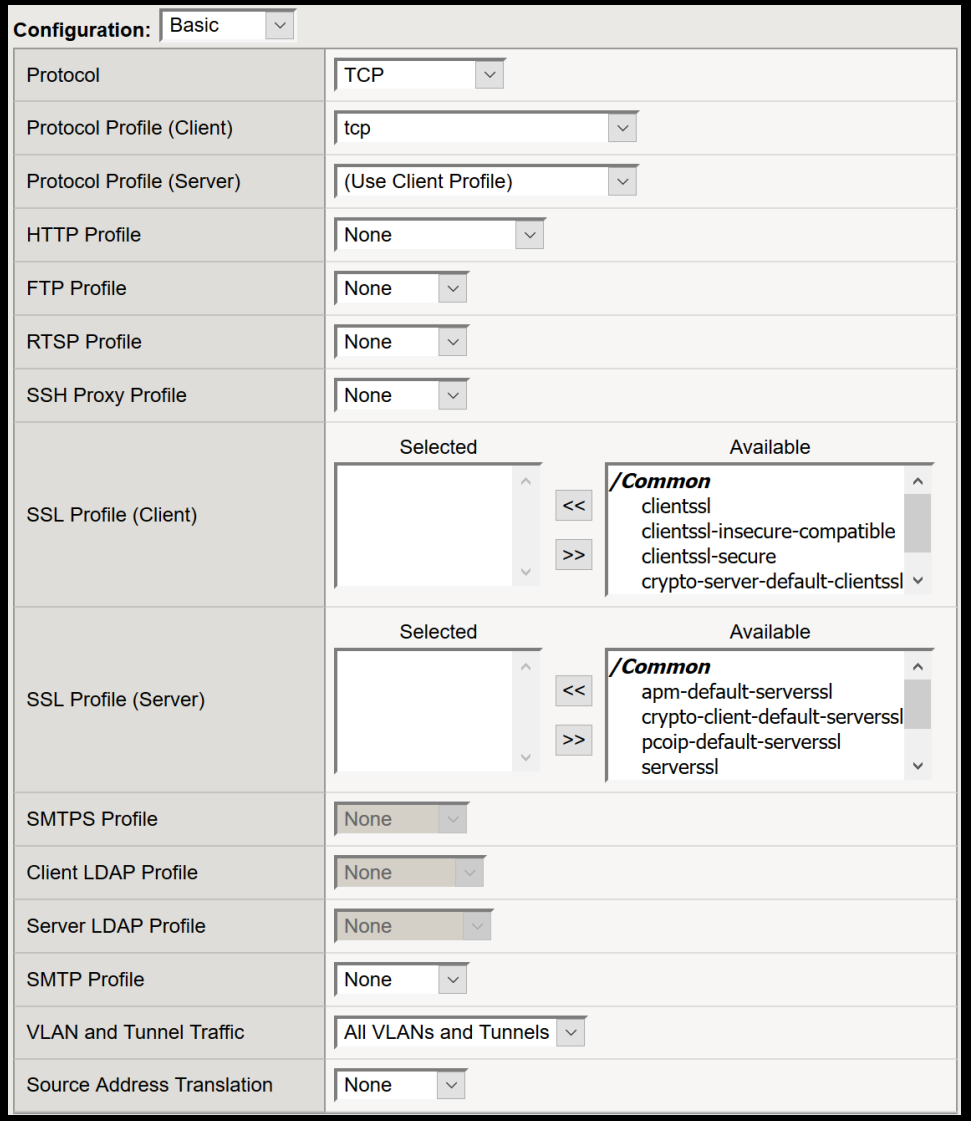

© 2016 F5 Networks

How Does a BIG-IP Handle Inbound Traffic

- **But virtual server isn't the only listener**

- **Listeners are**
  - Self IPs (Port Lockdown: None)
  - SNATs (Source initiated)
  - NATs (Two-way)
  - Virtual Servers

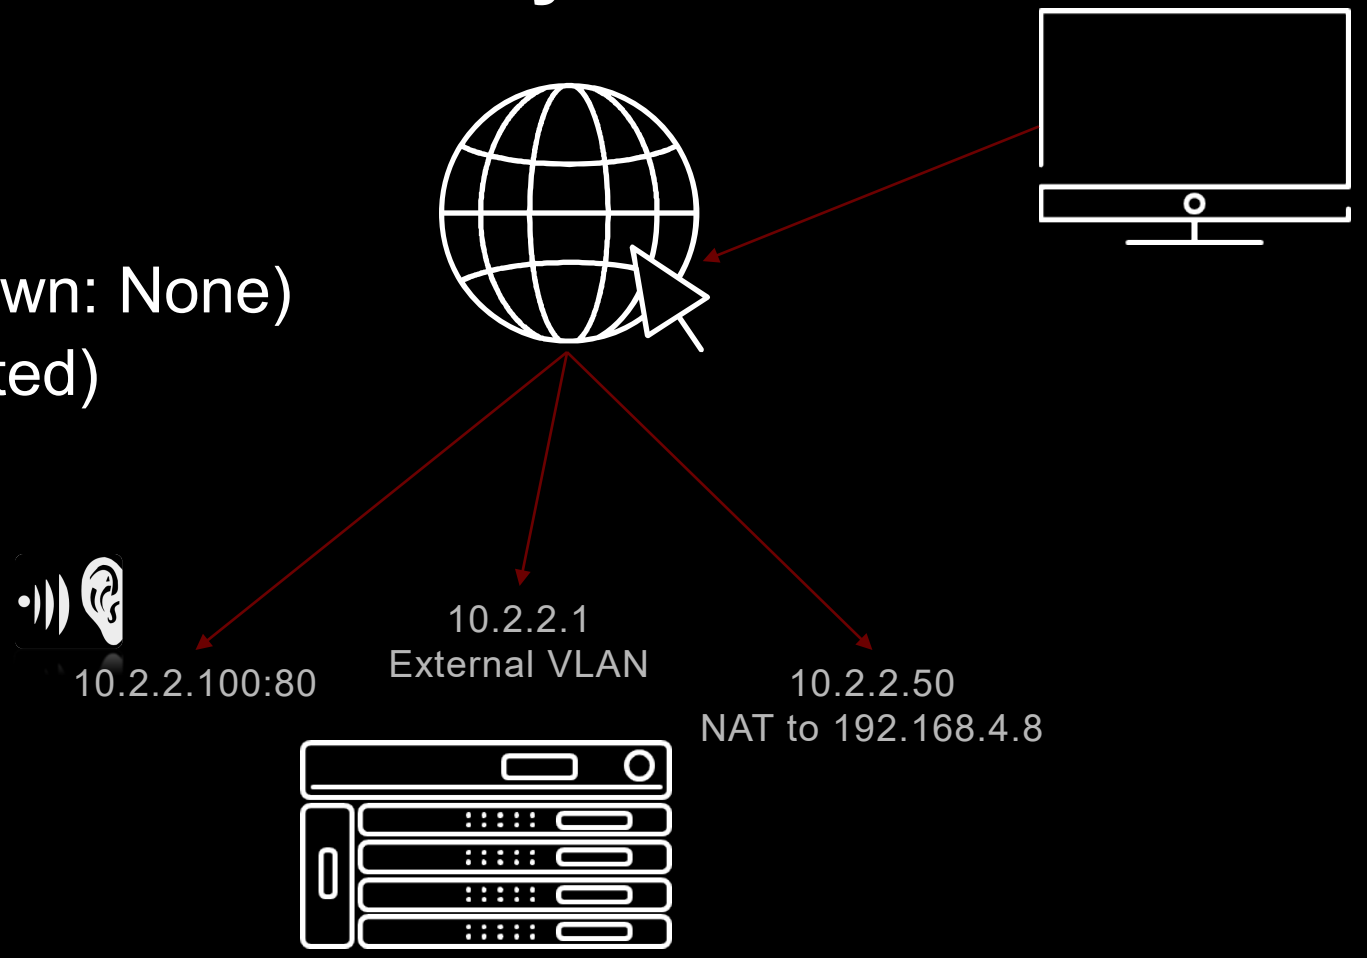

LOAD BALANCING

Introduction to Load Balancing

- **A load balancing method is an algorithm or formula used to determine which pool member to send traffic to**
  - Load balancing is connection based

- **Static load balancing methods distribute connections in a fixed manner**
  - Round Robin (RR)
  - Ratio (Weighted Round Robin)
    - Distributes in a RR fashion for members/nodes whose ratio has not been met

- **Dynamic load balancing methods take into account one or more factors, such as the current connection count**

- **It is important to experiment with different load balancing methods and select the one that offers the best performance in your particular environment**

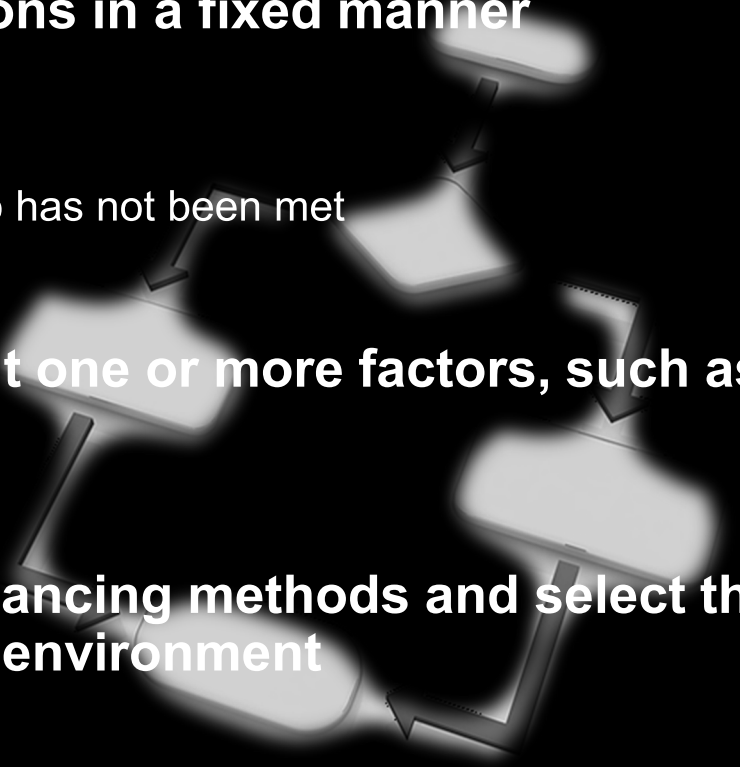

Dynamic Load Balancing Methods

- **Least Connections**
  - Fewest L4 connections when load balancing decision is being made
  - Recommended when servers have similar capabilities
  - Very commonly used

- **Fastest**
  - Balances based upon the number of outstanding L7 requests and then L4 connections
  - Requires a L7 profile on the virtual server, else its just Least Connection
  - Recommended when servers have similar capabilities

- **Observed**
  - Calculates a ratio each second based on the number of L4 connections
  - Not recommended for large pools

*SOL6406 - Change in Behavior: Fastest, Observed, and Predictive load balancing modes

Dynamic Load Balancing Methods

- **Predictive**
  - Calculates ratio base on the change between the previous connection counts and the current connection counts
  - Not recommended for large pools

- **Weighted Least Connections**
  - Based on how close the number of connections are to meeting the connection limit for a pool member or node
  - Requires connection limits be set on pool member or node
  - Recommended when servers have different capabilities

- **Dynamic Ratio**
  - Dynamically weights servers based on the results of SNMP/WMI queries
  - Requires SNMP_DCA , SNMP_Base, or WMI pool monitoring
  - Recommended when custom calculations are needed

Load Balancing a Service (Member)

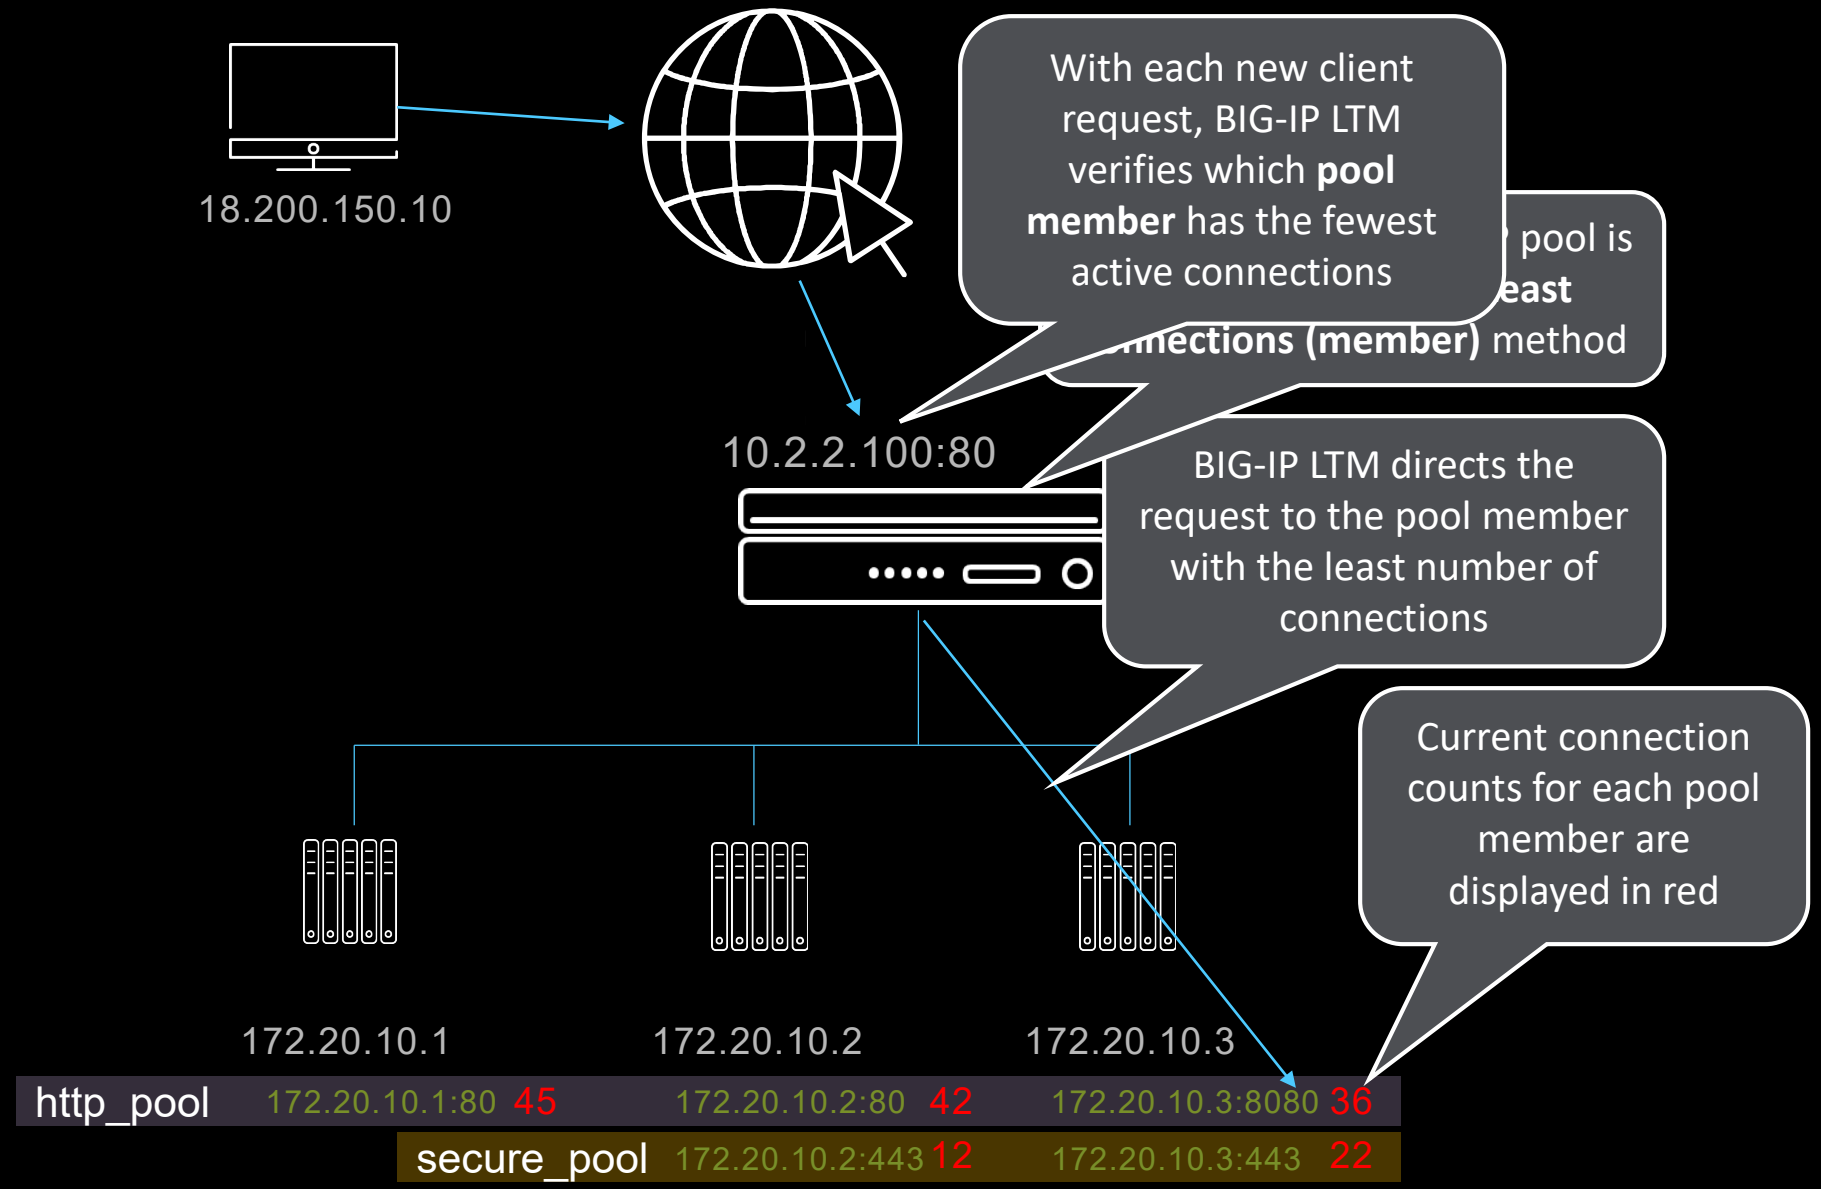

Load Balancing an IP Address (Node)

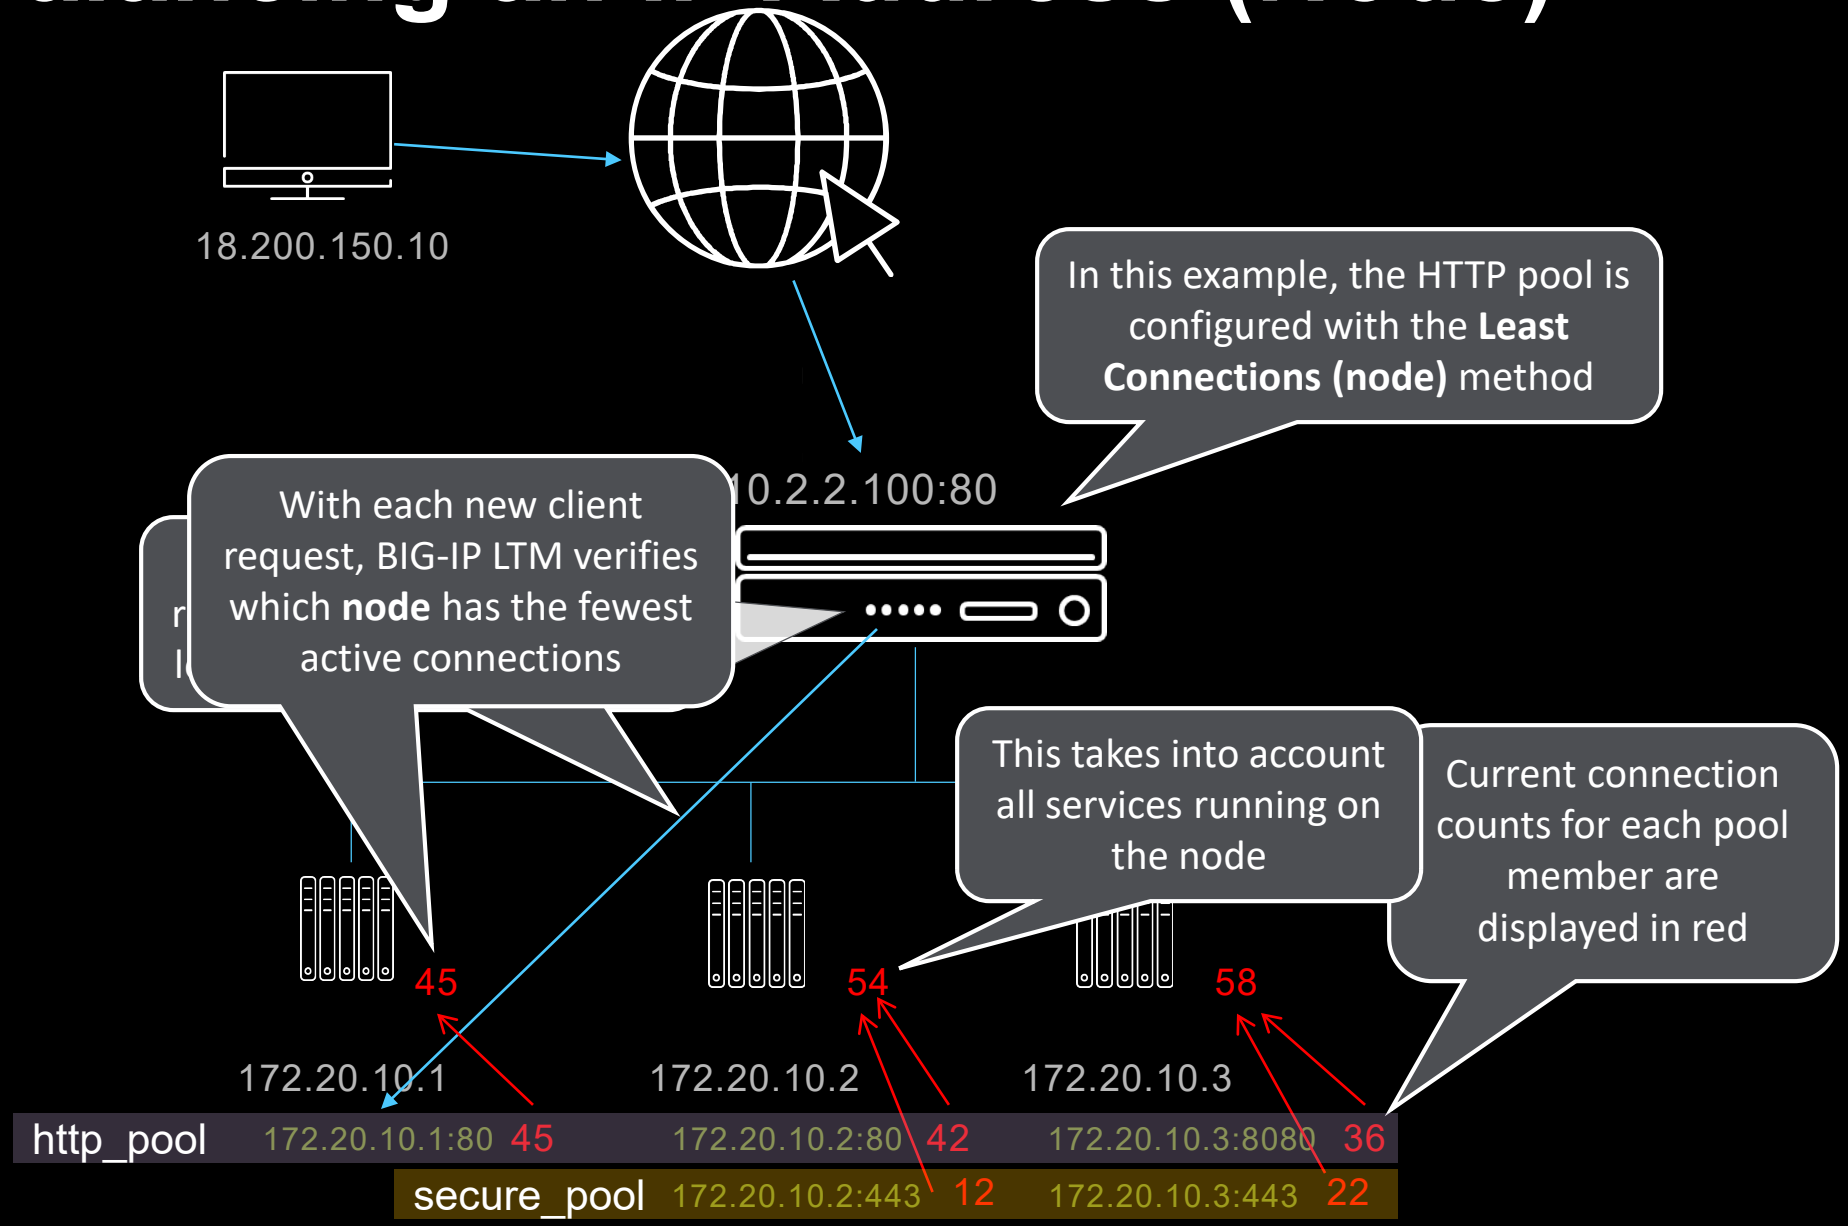

Pool Failure Mechanisms

- **Fallback Host (for HTTP and HTTPS applications)**
  - Is the server of last resort if all pool members are unavailable
  - Returns HTTP redirect (http 302) to client
  - Configured in the HTTP profile, the fallback host is not monitored

- **Priority Group Activation**
  - Can dynamically pull in new members into the pool
  - Pulls lower priority groups into higher priority groups
  - Pulls in all members of a priority group together

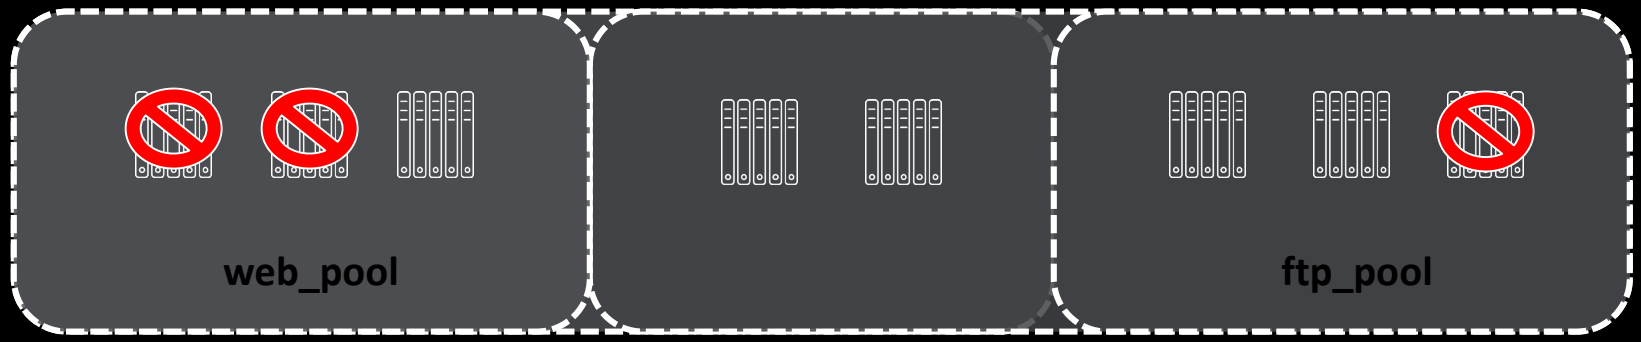

Priority Group Activation

Preferred and backup sets of pool members

To meet client traffic demands

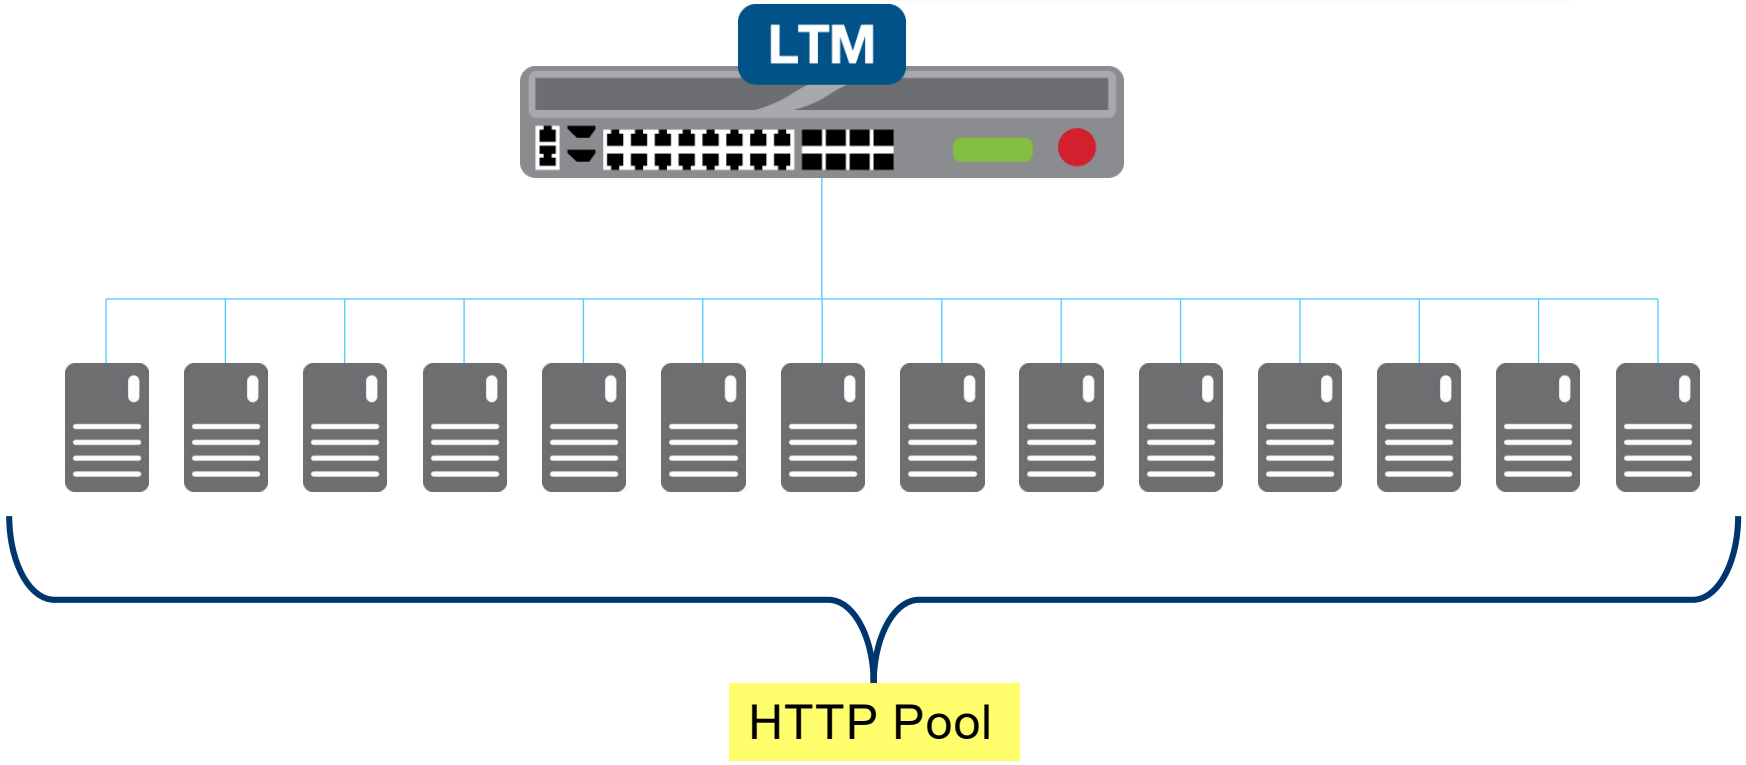

Using Priority Group Activation

BIG-IP LTM uses members with the highest priority number first

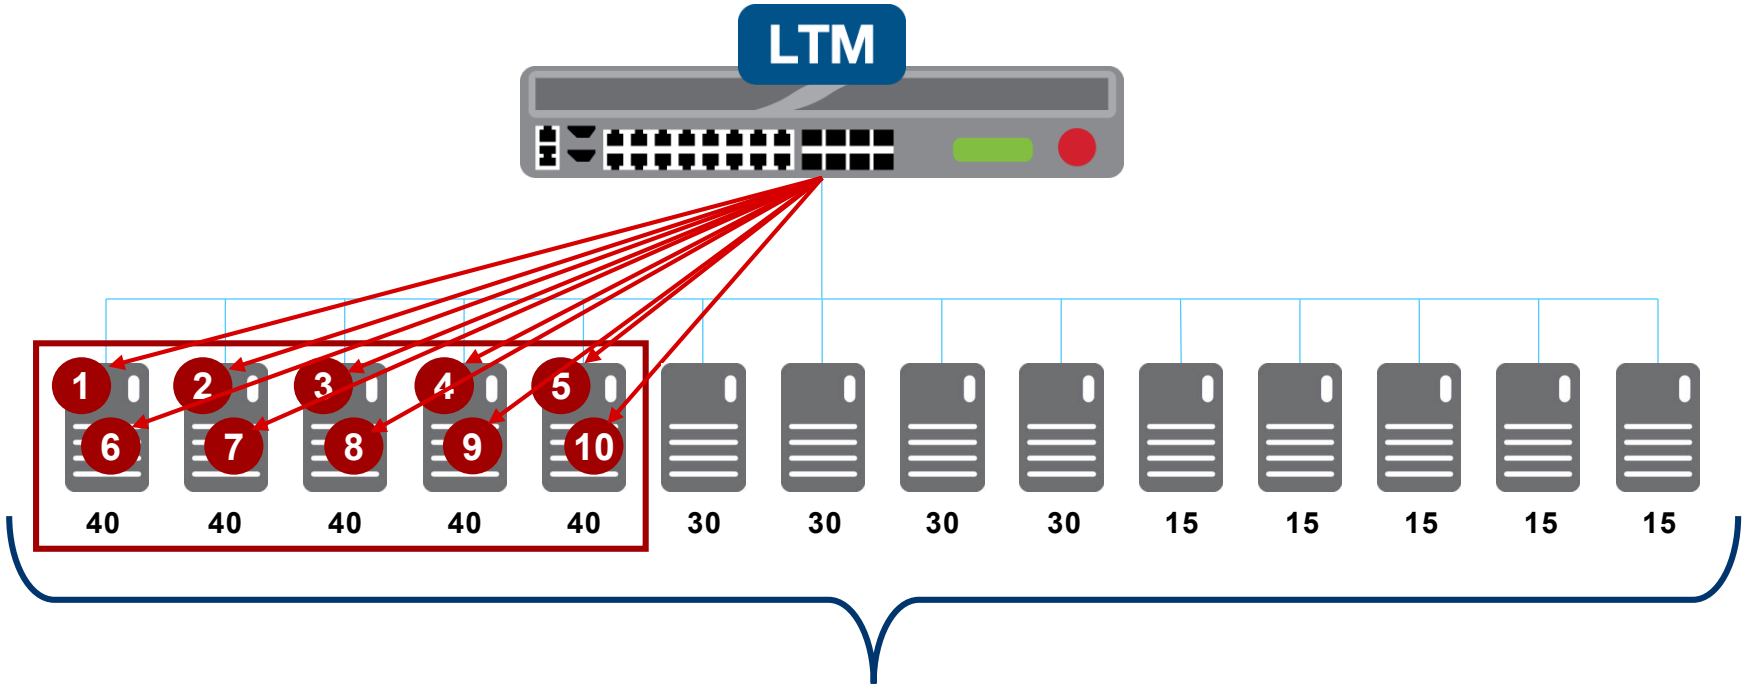

HTTP Pool

Priority Group Activation: Enabled

Less than 5 members

How Priority Group Activation Works

Priority Group Activation ensures that a pool doesn't go **below** a threshold

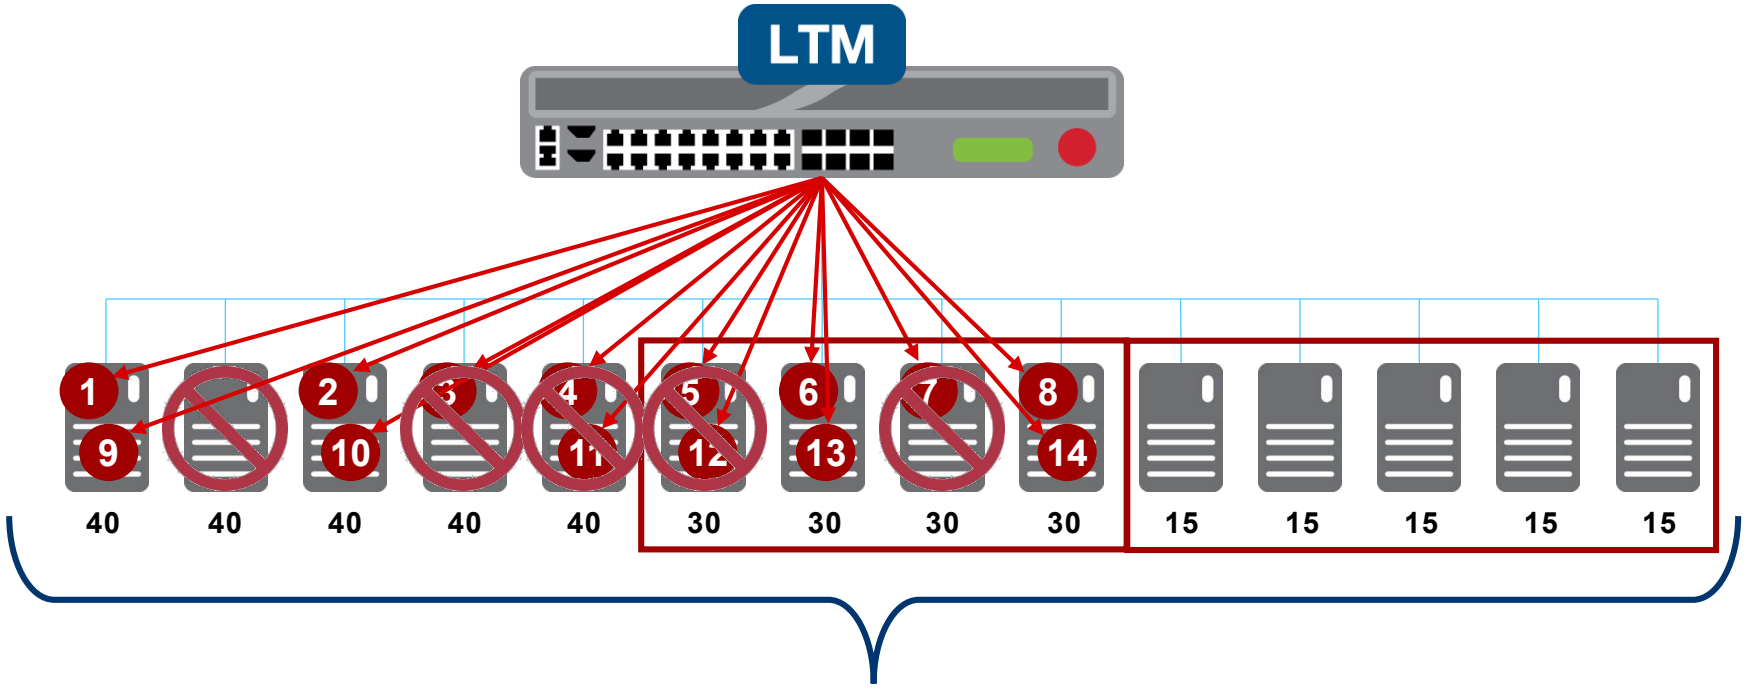

HTTP Pool

Priority Group Activation: Enabled

Less than 5 members

MONITORS

Types of Monitor Checks

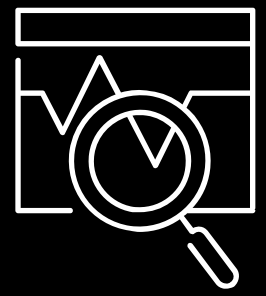

- **Node checks**
  - Determines availability of all services for a particular node
    - For example, ICMP check determines if the node is pingable
  - When a check fails, the node is pull from all pools it has membership in
- **Service checks**
  - Checks connectivity to services/ports
    - For example, HTTP check determines if port 80 can be opened
- **Content checks**
  - Queries the service and checks the contents of the query
    - For example, HTTP GET / determines if the page returns with correct content
  - Content checks can involve username and passwords
- **Path checks**
  - Are transparent monitors that check devices outside the pool
- **Interactive checks**
  - Custom scripts that interact with application

Monitor Types in the GUI

- **Simple monitors**
  - ICMP
  - Gateway ICMP
  - TCP Echo
- **Extended Content Verification (ECV) monitors**
  - TCP
  - HTTP
  - HTTPS
- **Extended Application Verification (EAV) monitors**
  - External
  - FTP
  - IMAP
  - LDAP
  - MSSQL
  - NNTP
  - Oracle POP3
  - RADIUS
  - Real Server
  - SIP
  - SMTP
  - SNMP DCA
  - SNMP
  - SOAP
  - UDP
  - WMI

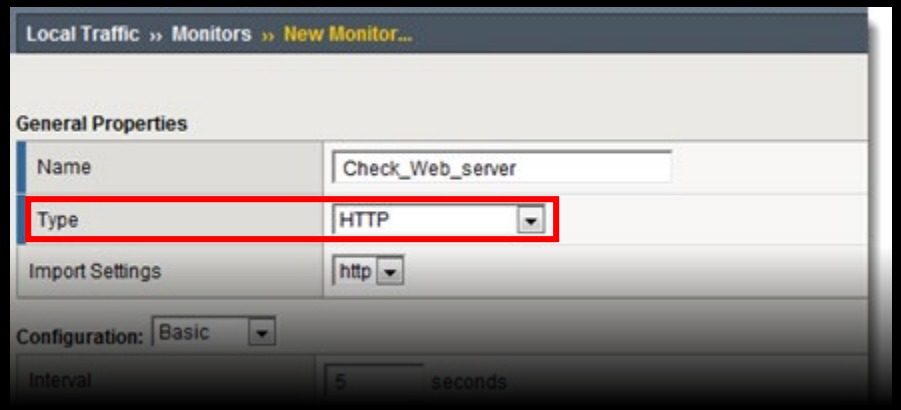

How Active Monitors Work

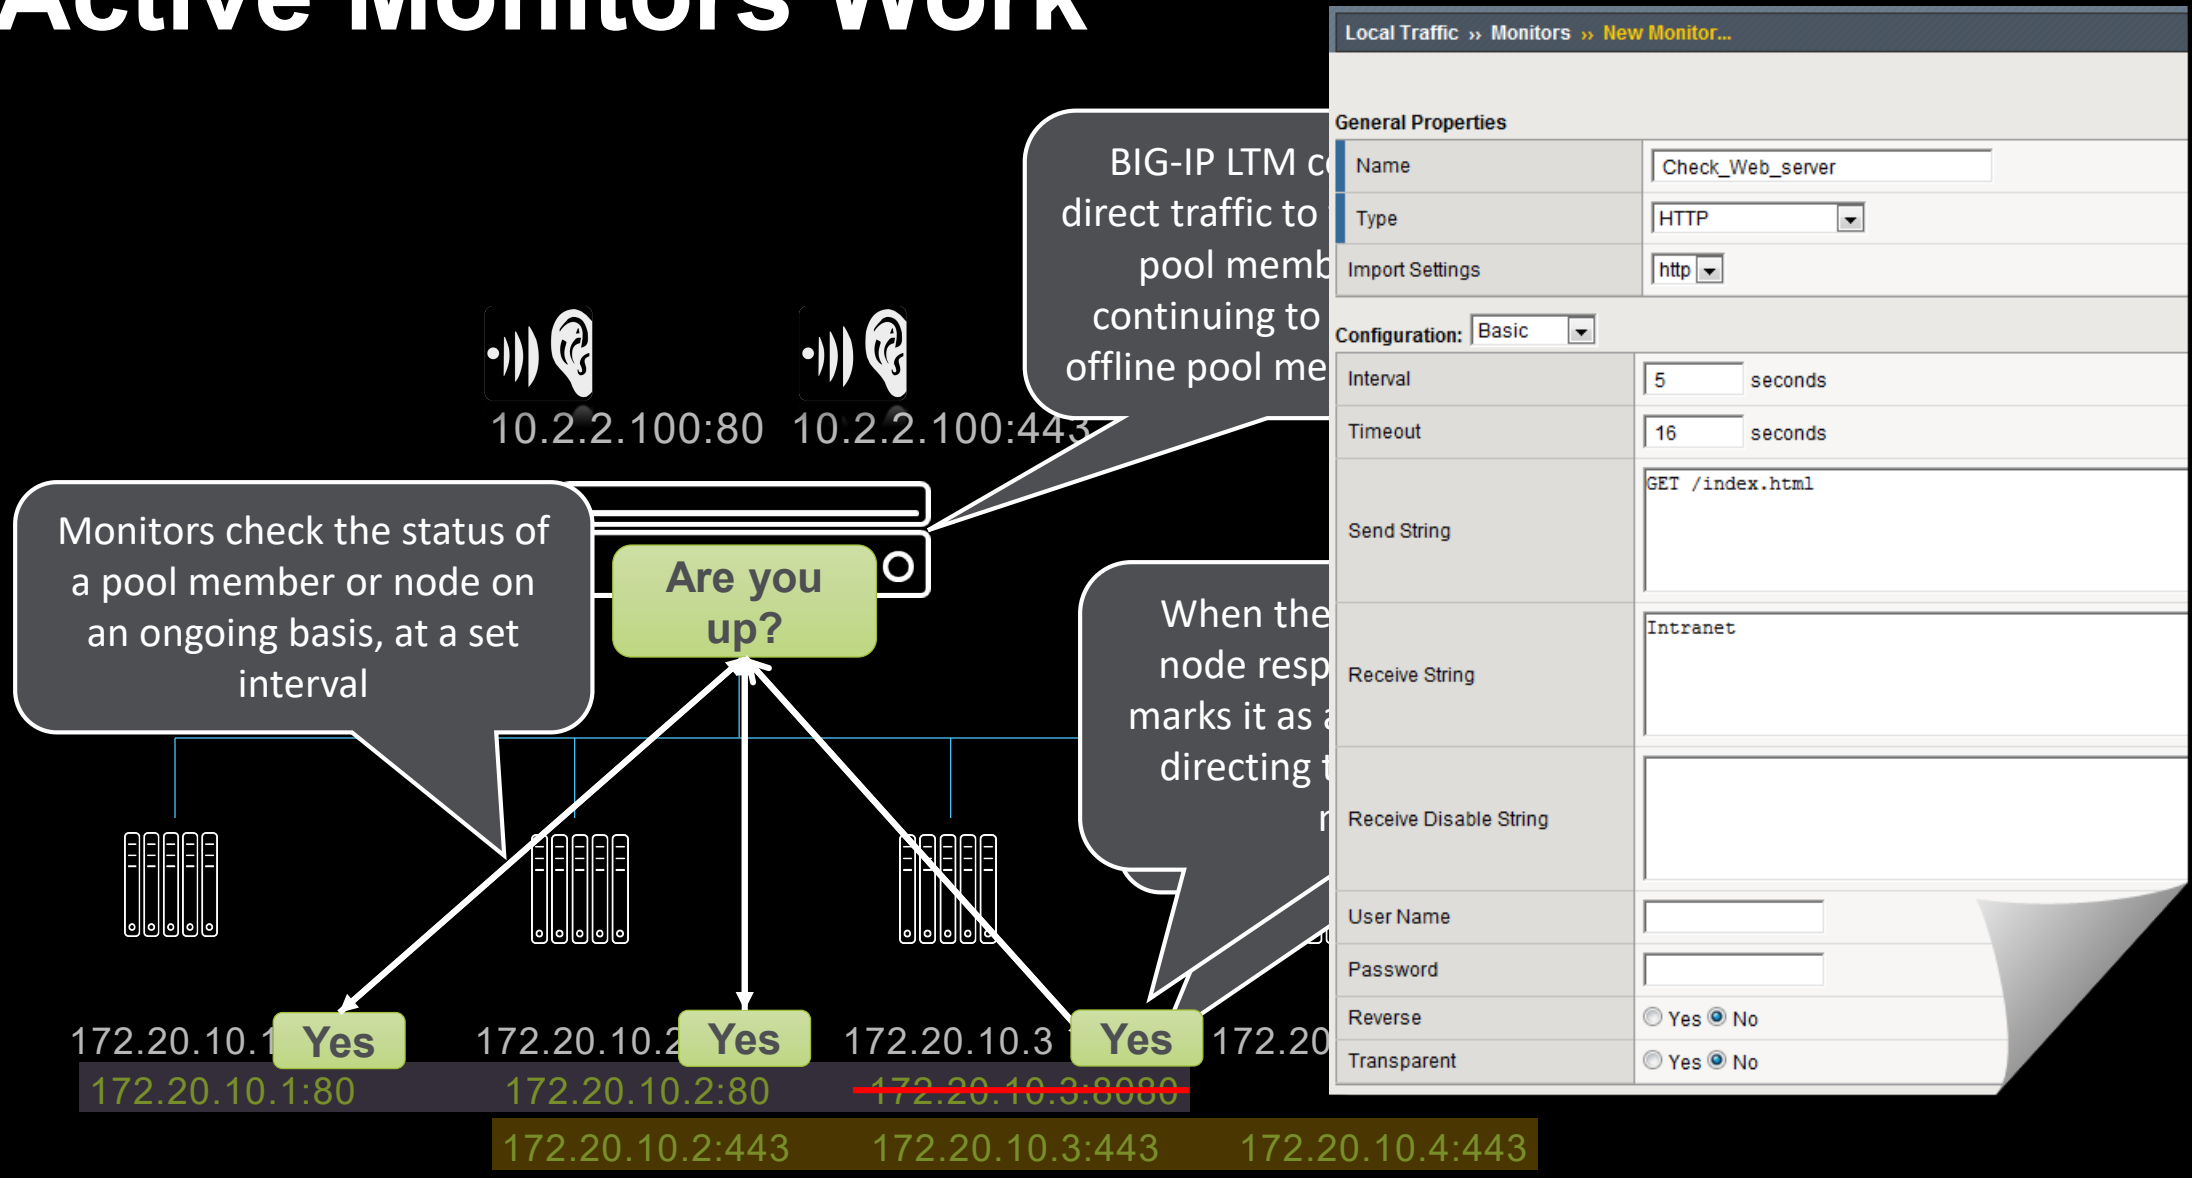

Monitor Status Reporting

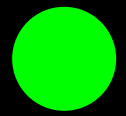

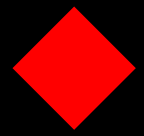

Status	Status Definition	
Available	General: Child	• Monitor successful
	General: Parent	• At least one child is Green
	Child –Node	• Most recent monitor successful
	Pool Member	• Most recent monitor successful
	Pool	• **At least one pool member is available**
	Virtual Server	• At least one pool is available
Unknown	General: Child	• No associated monitor (or timeout of first check not reached)
	General: Parent	• All child objects are unknown (blue)
	Node	• No associated monitor (or timeout of first check not reached and not successful)
	Pool Member	• No associated monitor (or timeout of first check not reached and not successful)
	Pool	• All pool members are unknown (blue)
	Virtual Server	• All pools are unknown (blue)
Offline	General: Child	• Monitor failed
	General: Parent	• At least one child red AND no green or yellow children available
	Node	• Most recent monitor failed (no successful checks within timeout period)
	Pool Member	• Most recent monitor failed (no successful checks within timeout period)
	Pool	• One or more members are offline and no members are available
	Virtual Server	• One or more pools offline and no members available

Inband (Passive) Monitors

- **Inband monitors use client request to see if the pool member:**
  - Connects (SYN-SYN/ACK-ACK)
  - If there is a L7 profile on the virtual server (ie. HTTP), checks for a response to the L7 request
- **Inband monitors require pools and virtual servers that:**
  - Are Standard or Performance (Layer 4)
  - Use the TCP or SCTP protocol profile
- **Inband monitors are effective in marking a pool member down**
  - Not as efficient in marking a member up
  - Member remains uncheck for a period of time
- **Active monitors are more effective in marking members up**
- **Active and inband monitors can be combined**
  - For effective monitoring with a lower overhead monitoring
  - Monitors at a different interval if the pool member is available

Receive Disabled String

- **This setting works like Receive String**
- **If the response matches Receive Disable String AND NOT the Receive String**
  - The node or member **Disabled**
- **Can be used to allow web admins to disable members without needing BIG-IP access**
- **To use this setting, you must specify both**
  - **Receive Disable String** and **Receive String**.

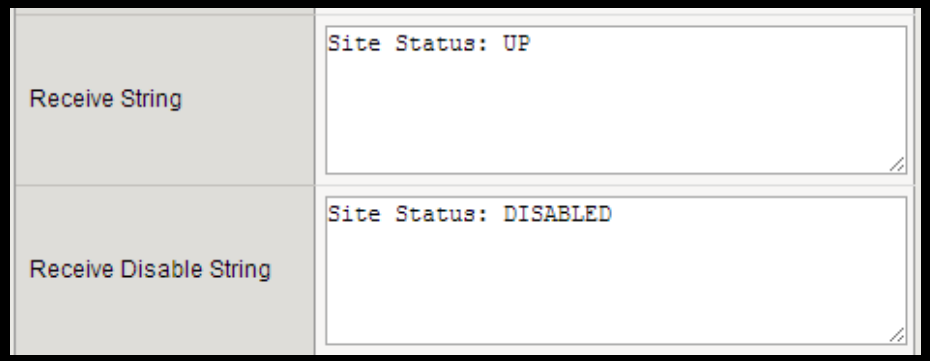

PROFILES & PERSISTENCE

Profiles

- **Are a configuration tool that aids you in managing application traffic**

- **A profile defines how a virtual server processes packets it receives**
  - Based on which profiles are assigned to the virtual server
  - Based upon the profile's configured parameters
  - The same profile can be associated with one or more virtual servers

- **Different profile types, different traffic processing capabilities**
  - Protocol profiles, such as, TCP and UDP
  - SSL profiles, for client-side and server-side certificates and keys
  - Service (L7) profiles, such as, HTTP, FTP, DNS
  - And many more…….

- **Profiles have a parent/child relationship**
  - Changes to a parent profile are passed down to the child profile(s)

Persistence

- **Persistence**
  - Directs a client back to the same server after the initial load balancing decision has been made
  - Is required for stateful applications
    - such as e-commerce shopping carts
  - May skew load balancing statistics
- **Basic persistence methods**
  - Simple Persistence
  - SSL Session ID
  - Cookie Persistence (Recommended for HTTP)
- **Other persistence methods**
  - Session Initiated Protocol (SIP)
  - MSRDP
- **Universal Persistence**
  - iRules can create persistence records based on anything in the clients request
    - Such as, jsessionid, username, etc.

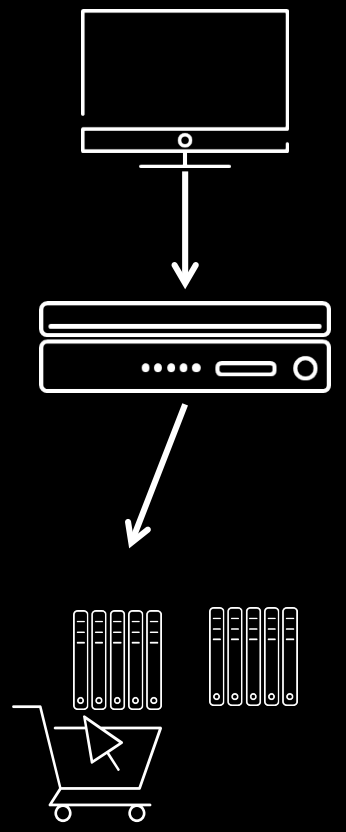

Persistence

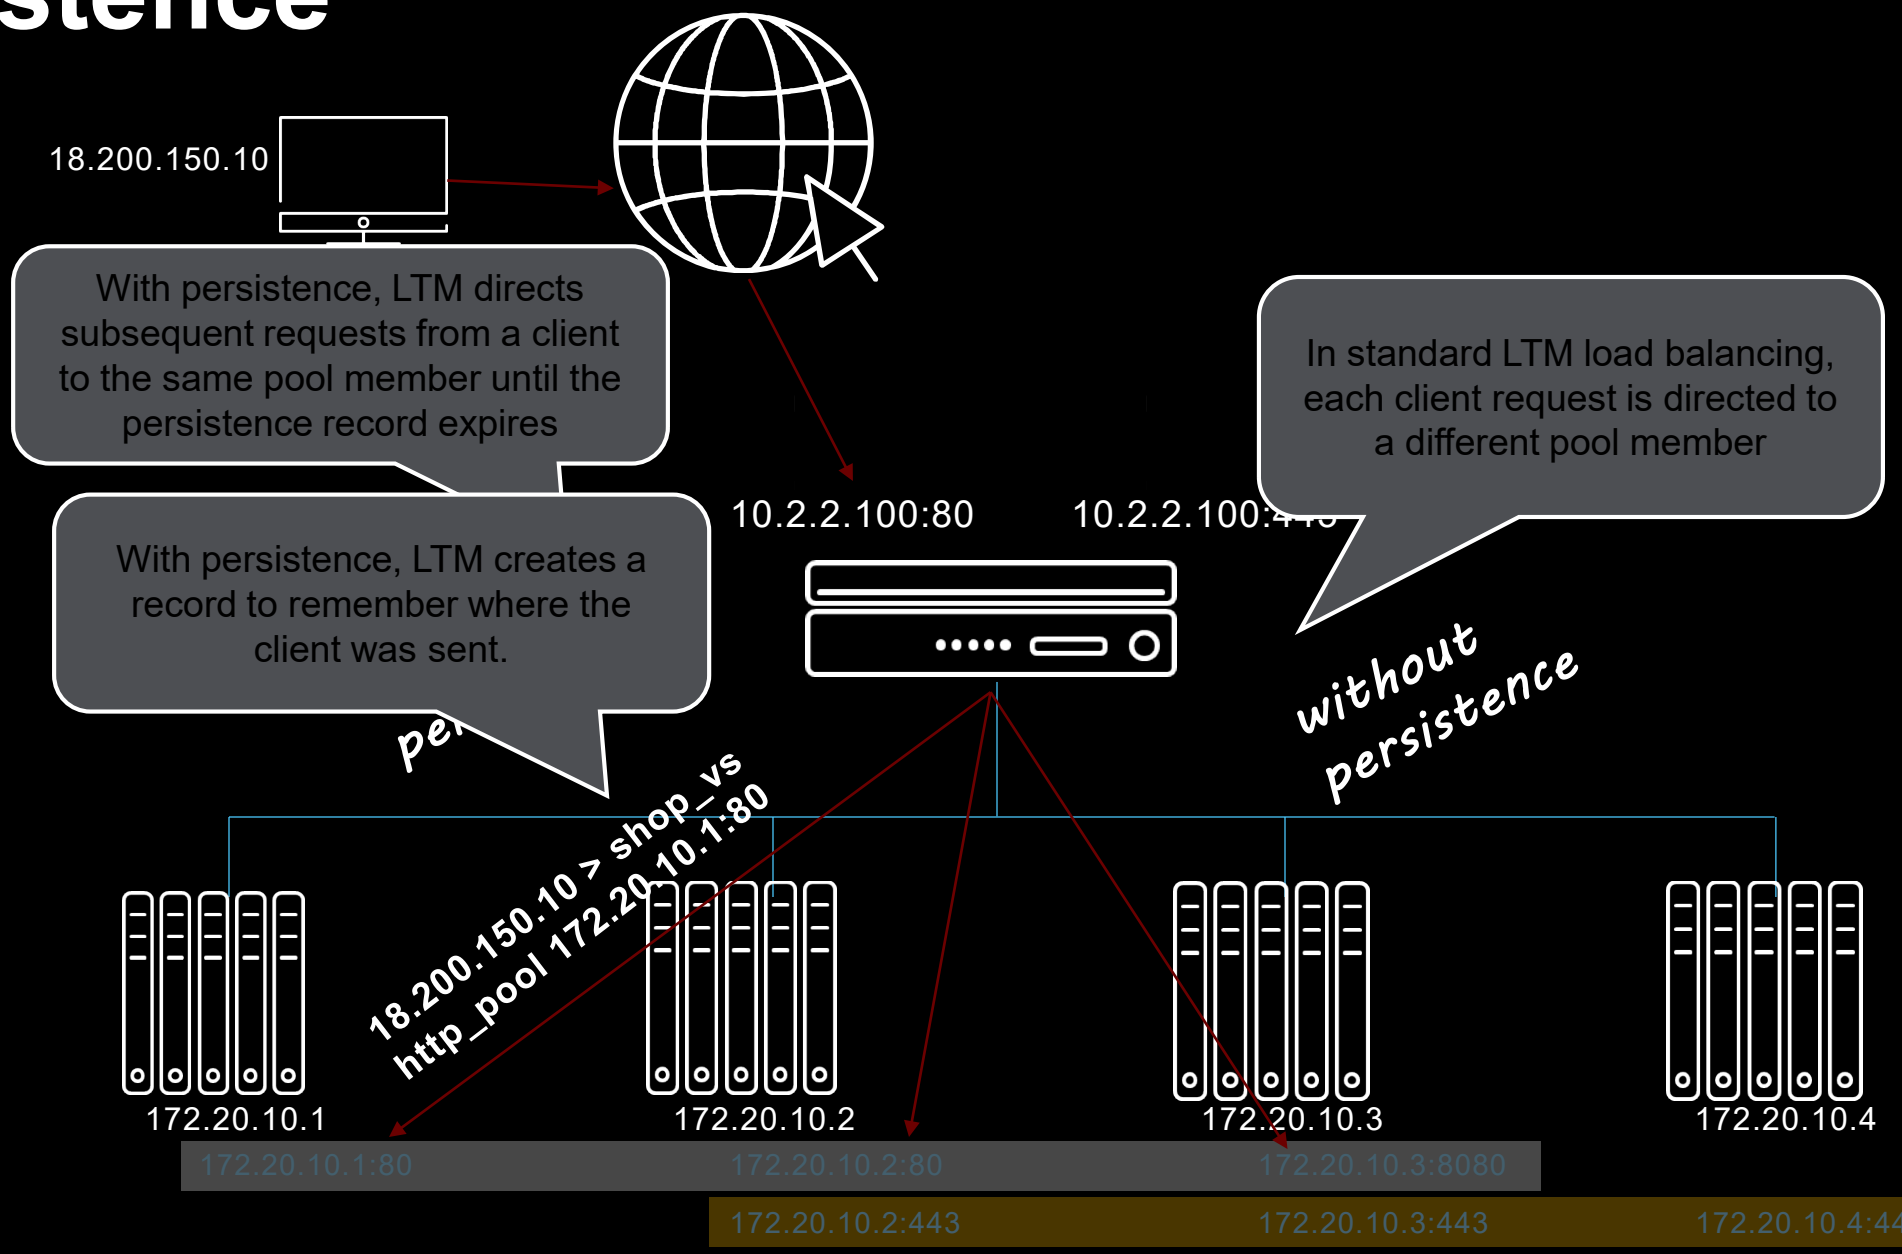

Persistence Settings

Match Across Services

- When enabled, specifies that all persistent connections from a client IP address that go to the same virtual IP address also go to the same pool member

Timeout

- Specifies the duration of the persistence entries
- Resets on a new connection

Override Connection Limit

- Allows new connections to be established if the connection limit is reached, if there is a persistence record

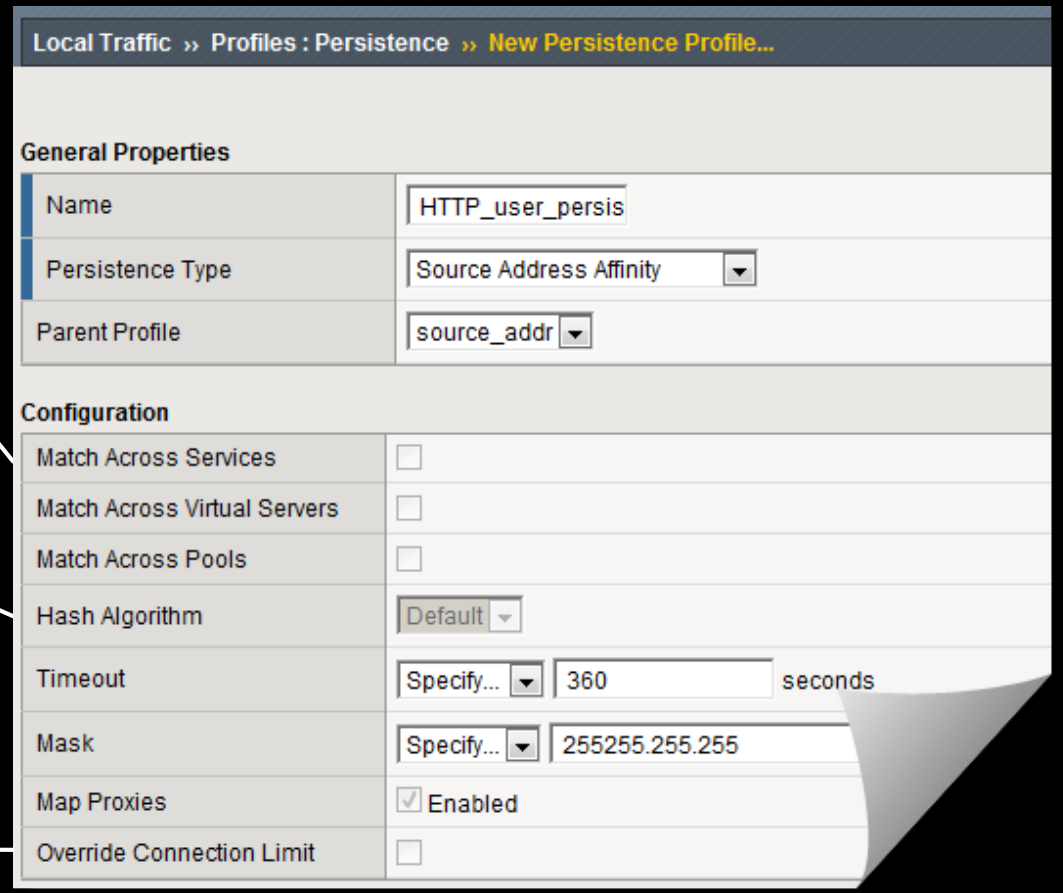

Cookie Persistence

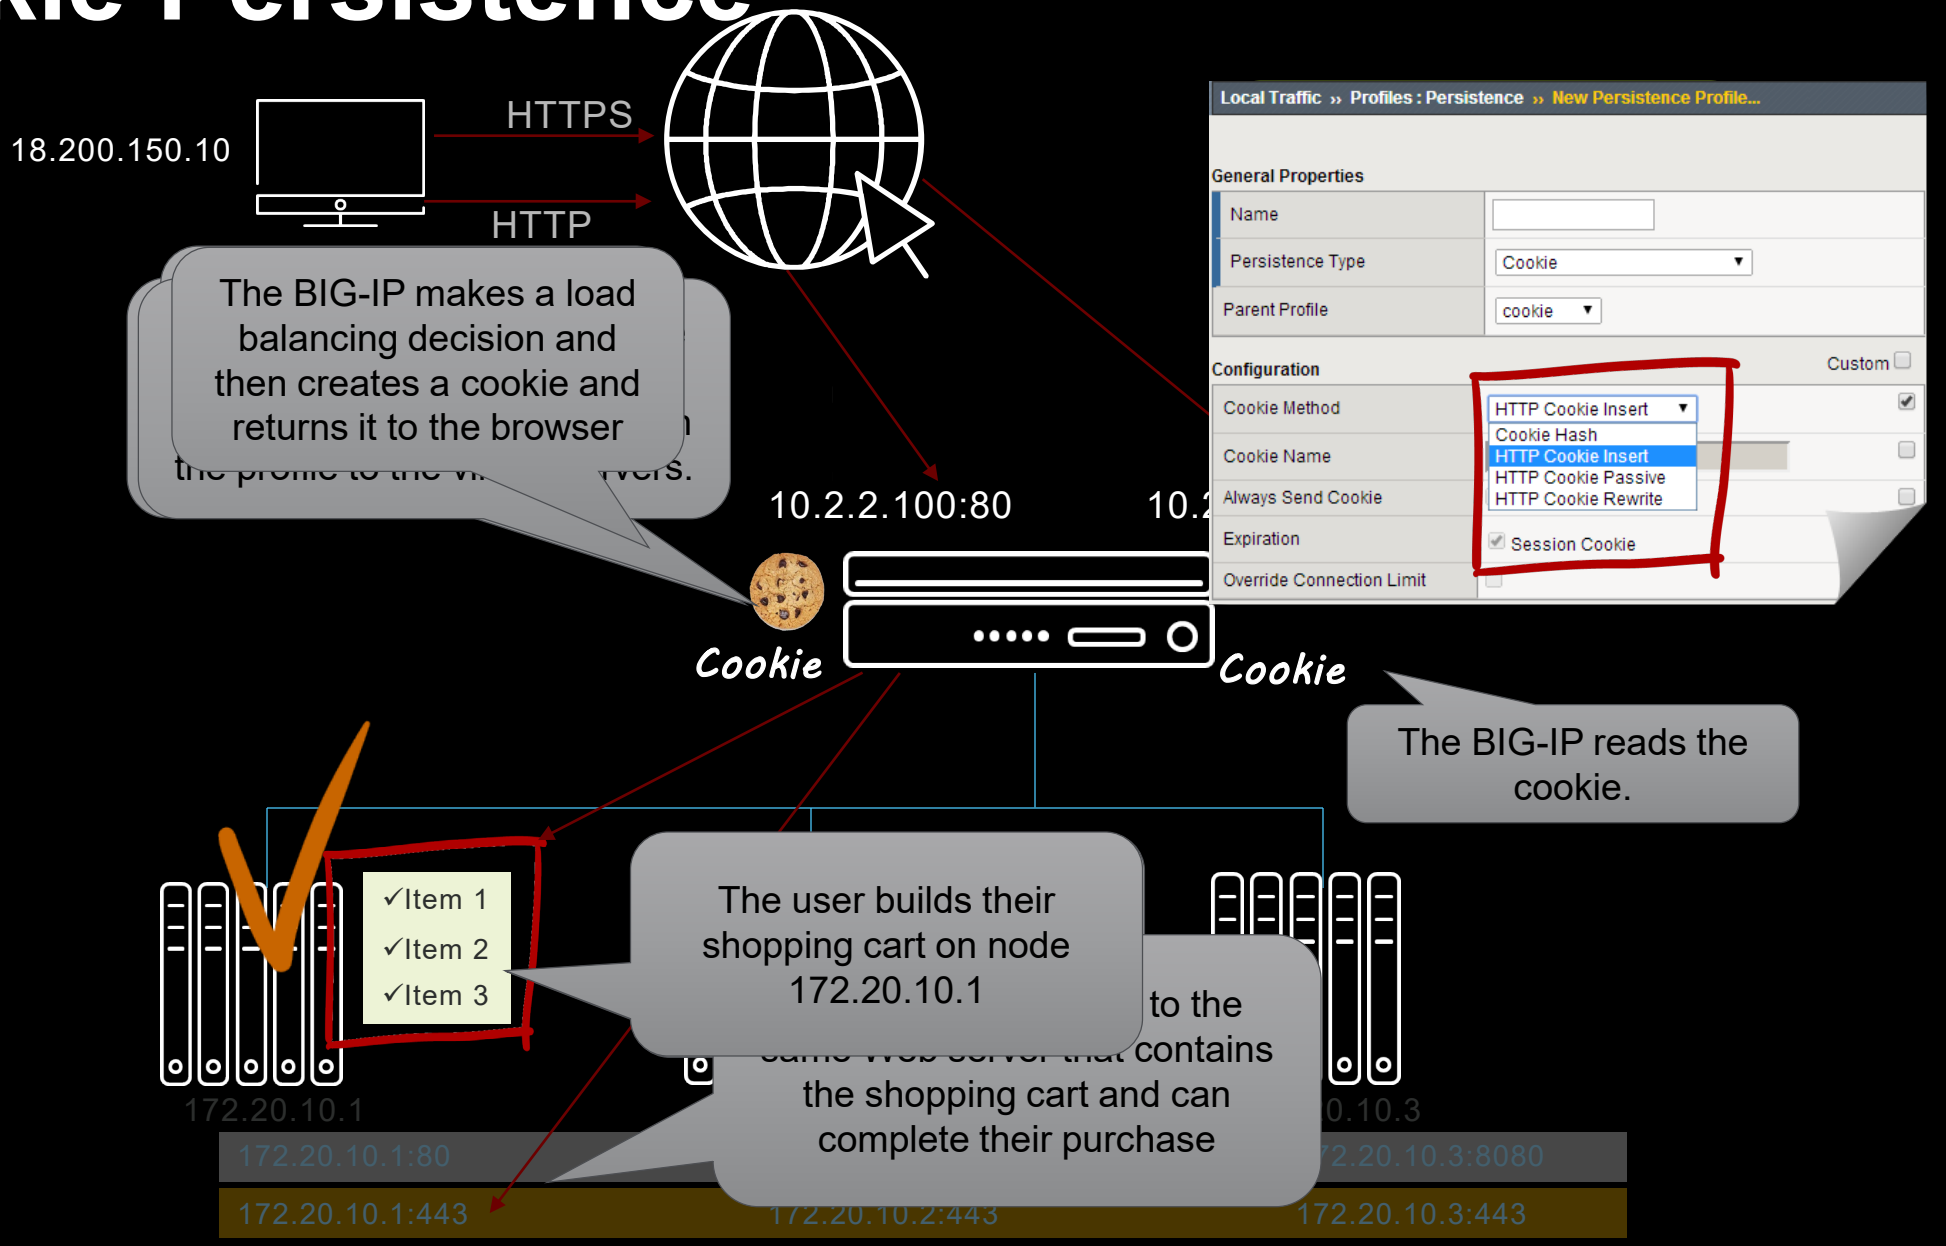

SSL OFFLOAD

SSL Offload (aka SSL Visibility)

- **Terminates the SSL connection at BIG-IP**
  - BIG-IP has full visibility into the application
  - Enables the use of iRules, Profiles, et al
  - Can decrypt/encrypt for 3rd party security devices (ie. IPS/IDS)
  - Can free up valuable server resources
- **Includes:**
  - Consolidated certificate and key management
  - Support for FIPS hardware-based key security
- **Can selectively insert/retrieve SSL client certificate information to be used in traffic management decisions**
- **BIG-IP has hardware ASIC to perform:**
  - SSL key exchange
  - **SSL bulk encryption**

7	APPLICATION
6	PRESENTATION
5	SESSION
4	TRANSPORT
3	NETWORK
2	DATA LINK
1	PHYSICAL

Certificates and Keys

- **BIG-IP allows you to create, import and view certificates**

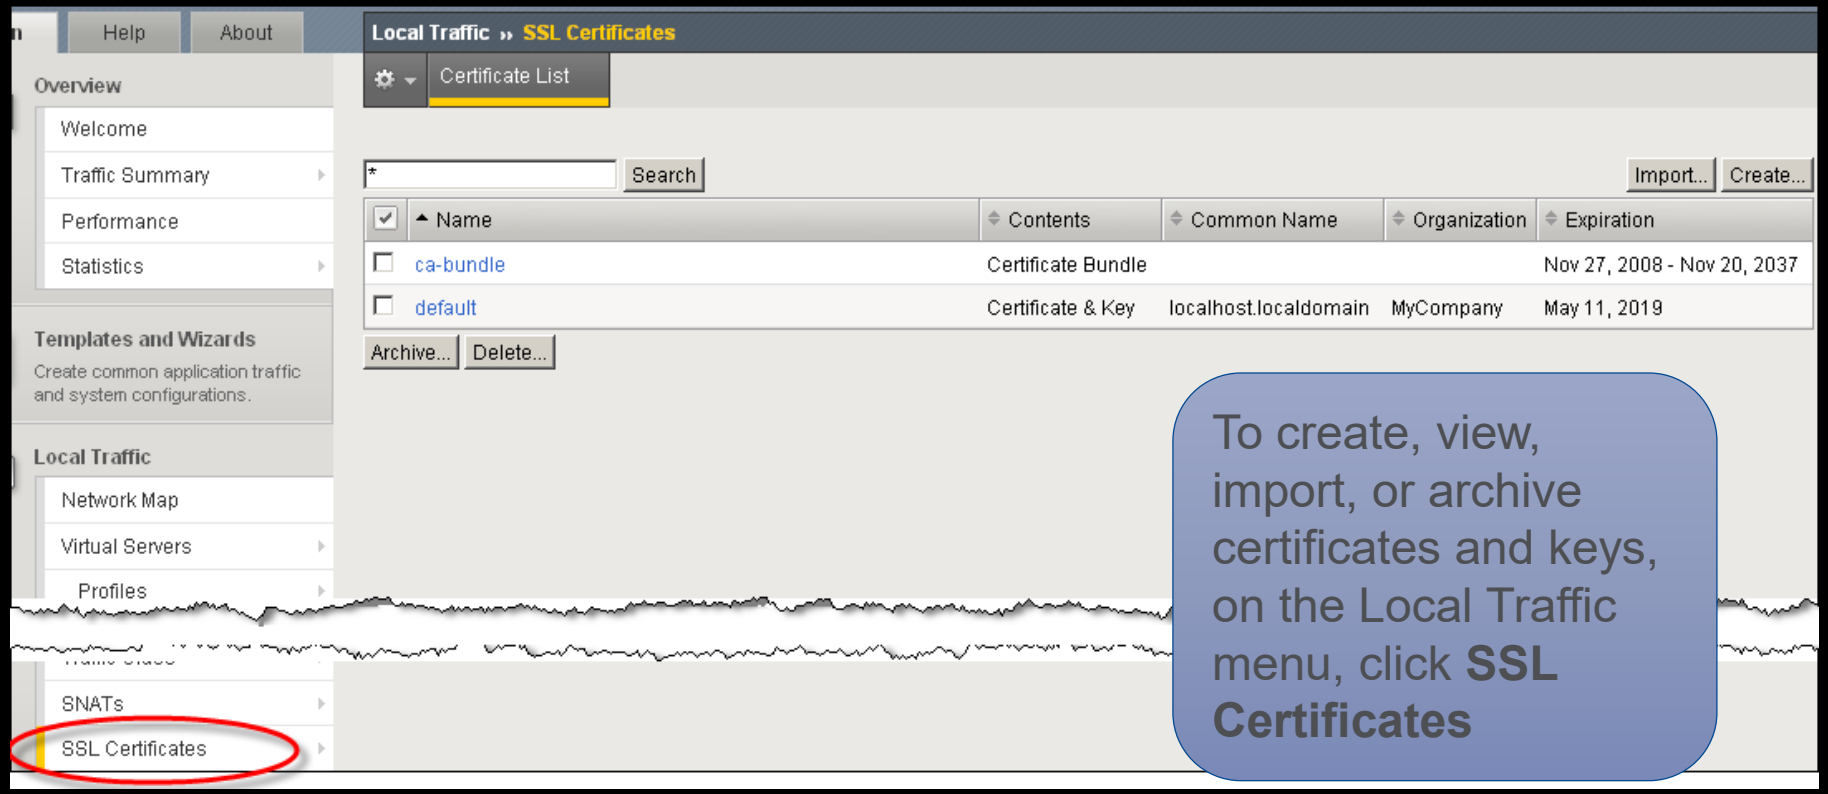

Certificates and Keys

You can create self-signed certificates or work with a Certificate Authority

You can import certificates into BIG-IP LTM

- Both PEM and PKCS certificates
- Or you can import previously archived certificates

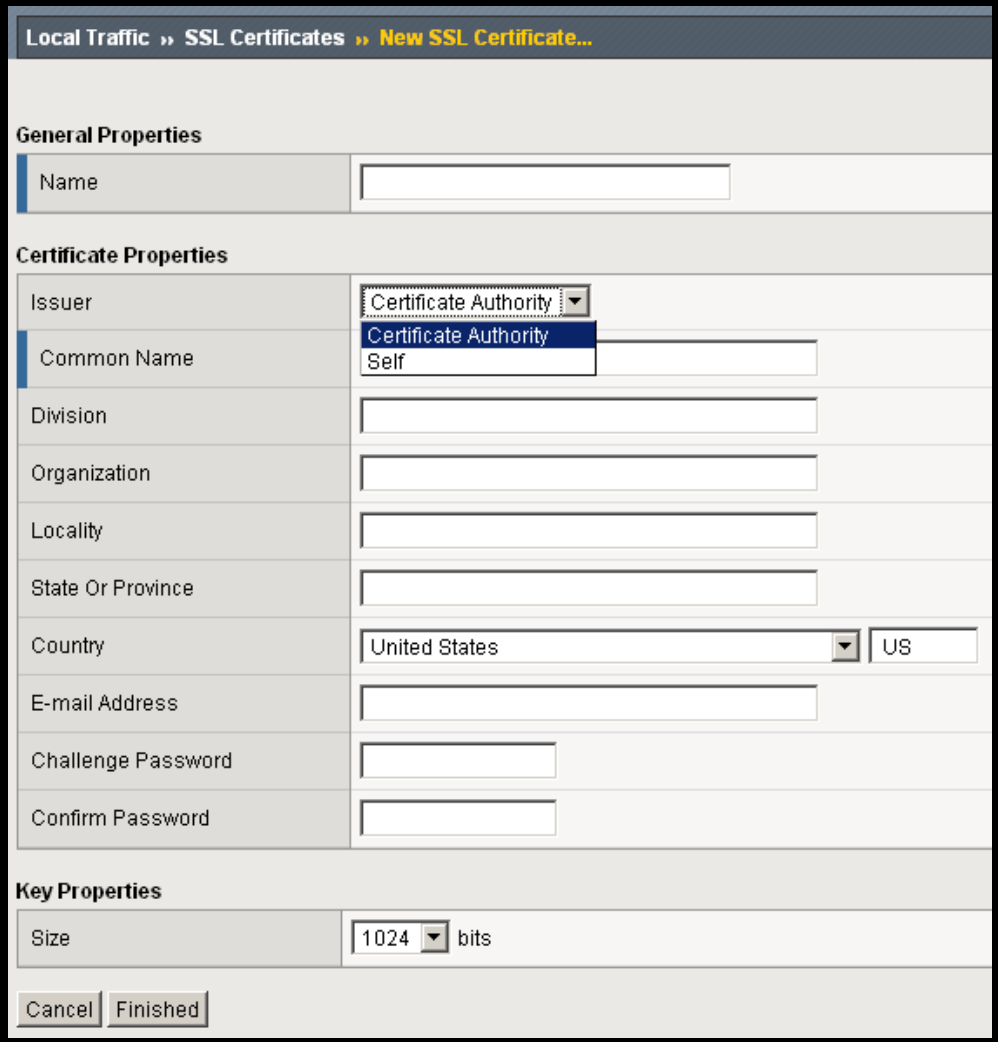

Processing SSL Traffic on the Client-SIde

- **To configure a virtual server to process HTTPS:**
  - Import/Create certificate and key
  - Create a client SSL profile,
    - Attach the certificate and key
  - Create a virtual server
    - In the **SSL Profile (Client)** box choose the SSL profile

- **How client-side processing works**
  - Client connects to a virtual server that is configured with the client SSL profile
  - The client and BIG-IP perform a key exchange and establish an encrypted session
  - The virtual server receives the client traffic
    - Decrypts traffic
    - Performs traffic management functions
      - For example, iRules or cookie persistence
  - The BIG-IP then sends the unencrypted traffic to the chosen pool member

Processing SSL Traffic on the Server

- **Use SSL server profiles for highly secure environments**
  - Configure a server-side SSL profile
    - ***Certificate could be self signed of lower encryption***
  - Attach the SSL Server profile to the virtual server

- **How server-side processing works**
  - Client connects to the virtual server using the cert and key in the client SSL profile
  - They establish an encrypted session
  - The virtual server receives and decrypts the traffic
  - Performs traffic management functions
  - An encrypted session is established between BIG-IP LTM and the selected pool member.
    - Using the certificate and key in the SSL Server profile
  - The data is re-encrypted and sent to the pool member

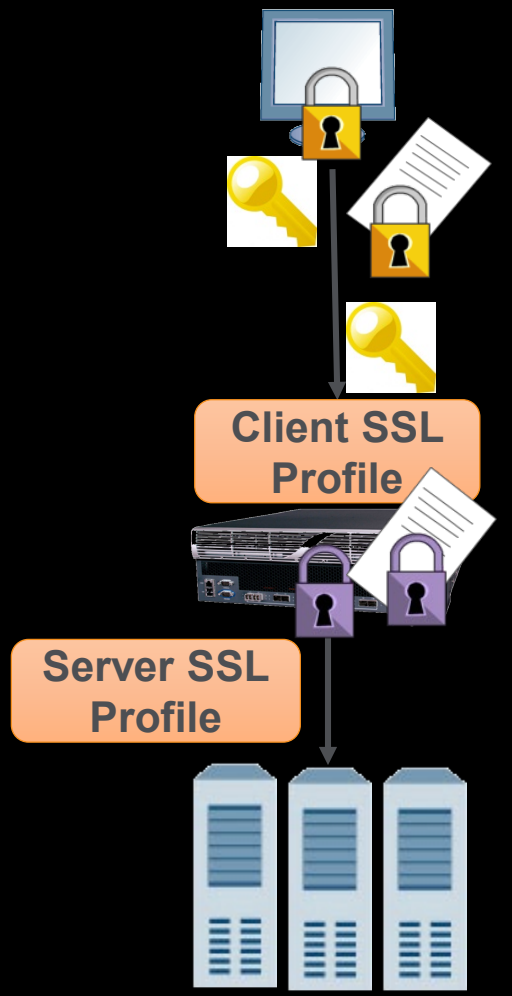

SSL visibility and inspection

Benefits of SSL visibility

SSL visibility provides:

- Malware protection
- Corporate compliance
- Productivity monitoring
- Intellectual property protection
- Customer experience enhancement
- Decreased cost and complexity of content security functions

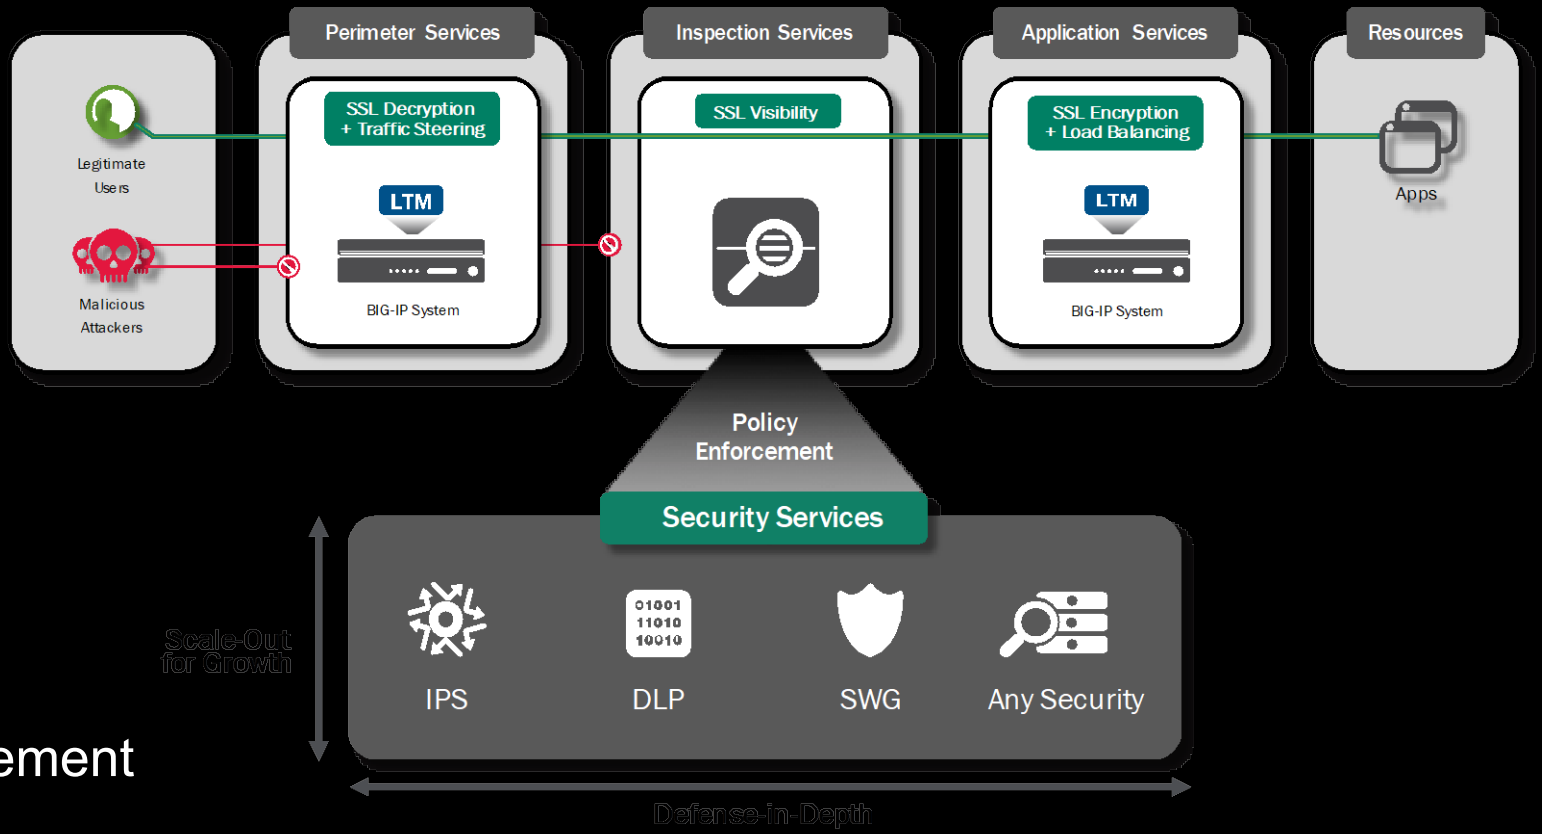

ACCELERATION

LTM (TMOS) Core Acceleration Features

- **Core TMOS acceleration technologies**
  - TCP Express
  - OneConnect
  - HTTP compression
  - Fast Cache
  - iSessions
  - Rate Shaping
  - Bandwidth Controller
  - SPDY Gateway
  - HTTP/2 Gateway

Acceleration Interface

- Prior to V11.4 acceleration profiles appeared only under Local Traffic >> Profiles
- After V11.4 acceleration profiles are also under the Acceleration tab, along with other features

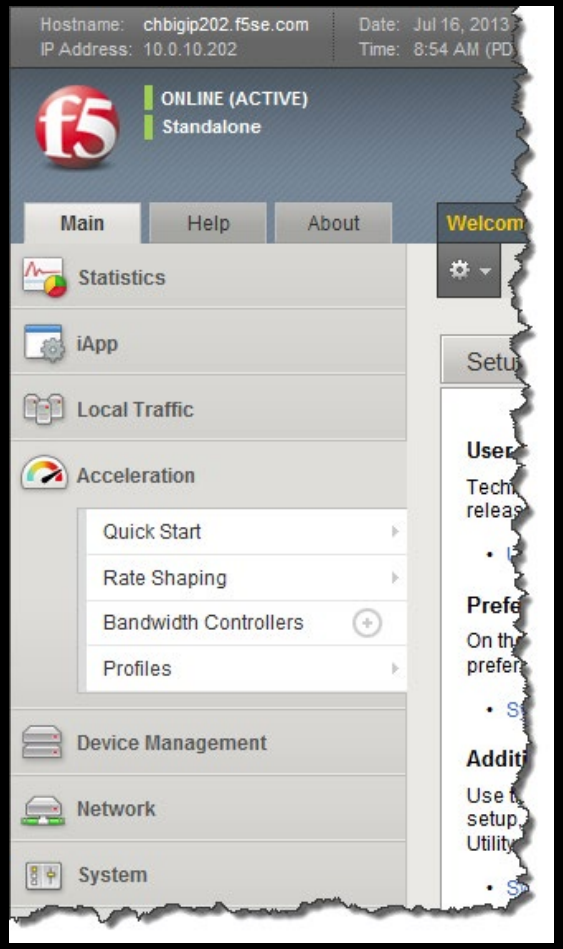

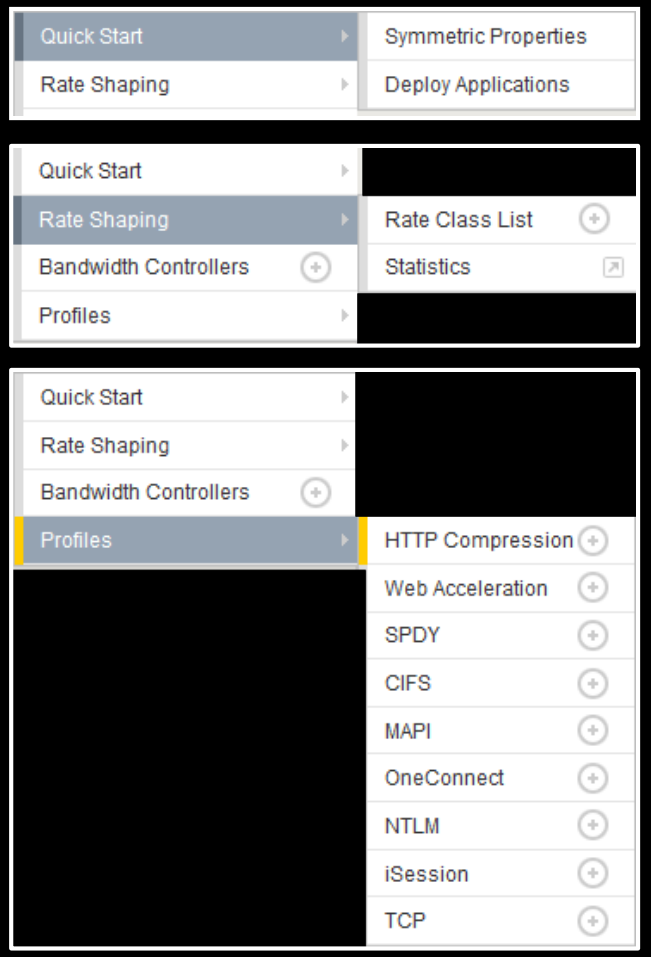

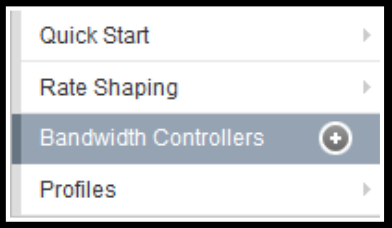

Problems with Traditional TCP

- **TCP/IP inefficiencies, coupled with WAN latency and packet loss, all contribute to slow application performance**
  - Inflated client response times
  - Reduced bandwidth utilization (ability to “fill the pipe”)

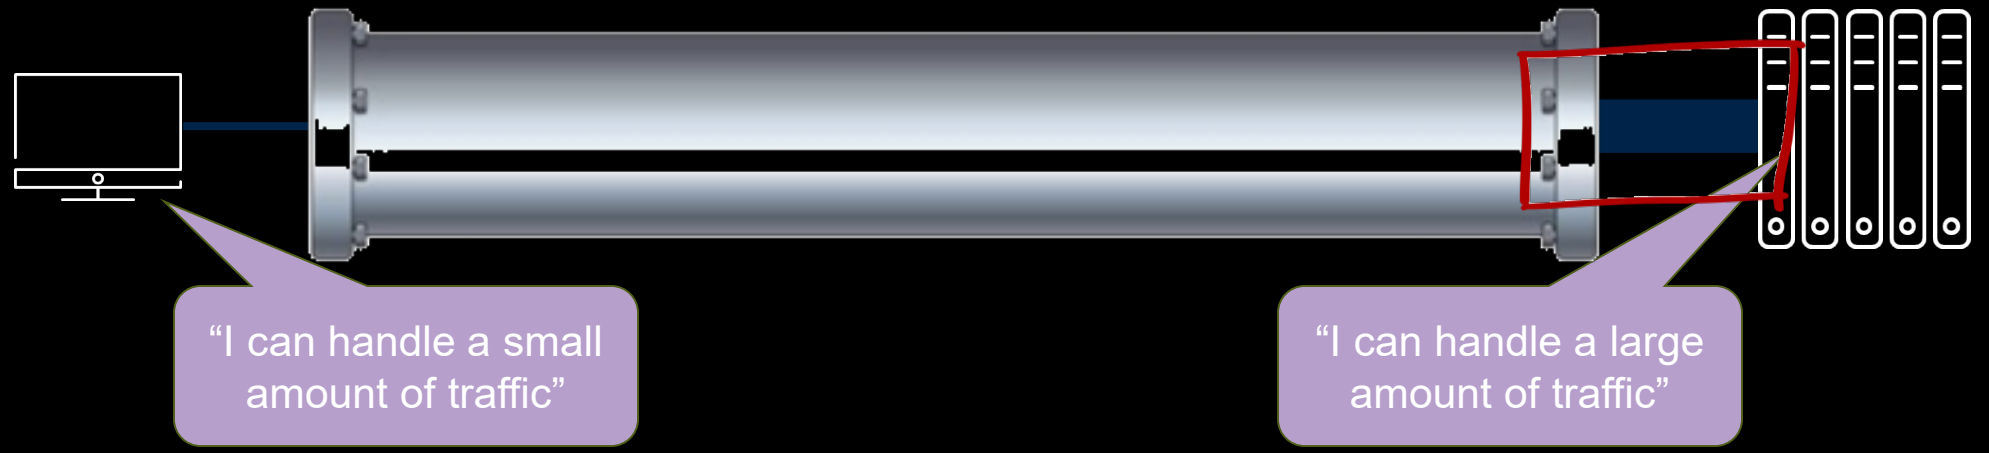

TCP Acceleration Features

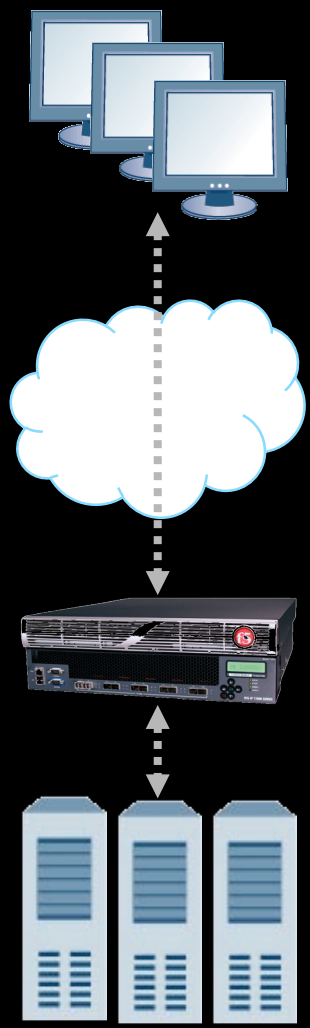

Goal: To improve the client experience

TCP Express (or TCP optimization)

- Adaptive congestion windows
- Fast retransmits
- Selective acknowledgements
- Congestion notification

TCP Client-Side Profiles

Goal: Server offload

Content Buffering

- Content spooling

Connection Management

- OneConnect

TCP Server-Side Profiles

What Is the HTTP Profile?

- **Layer 7 profile, notifies the virtual server to parse the HTTP protocol, headers and request/response sequences**
- The all HTTP profiles are under **Local Traffic > Profile > Services**
- The HTTP acceleration profiles are also under **Acceleration > Profiles**

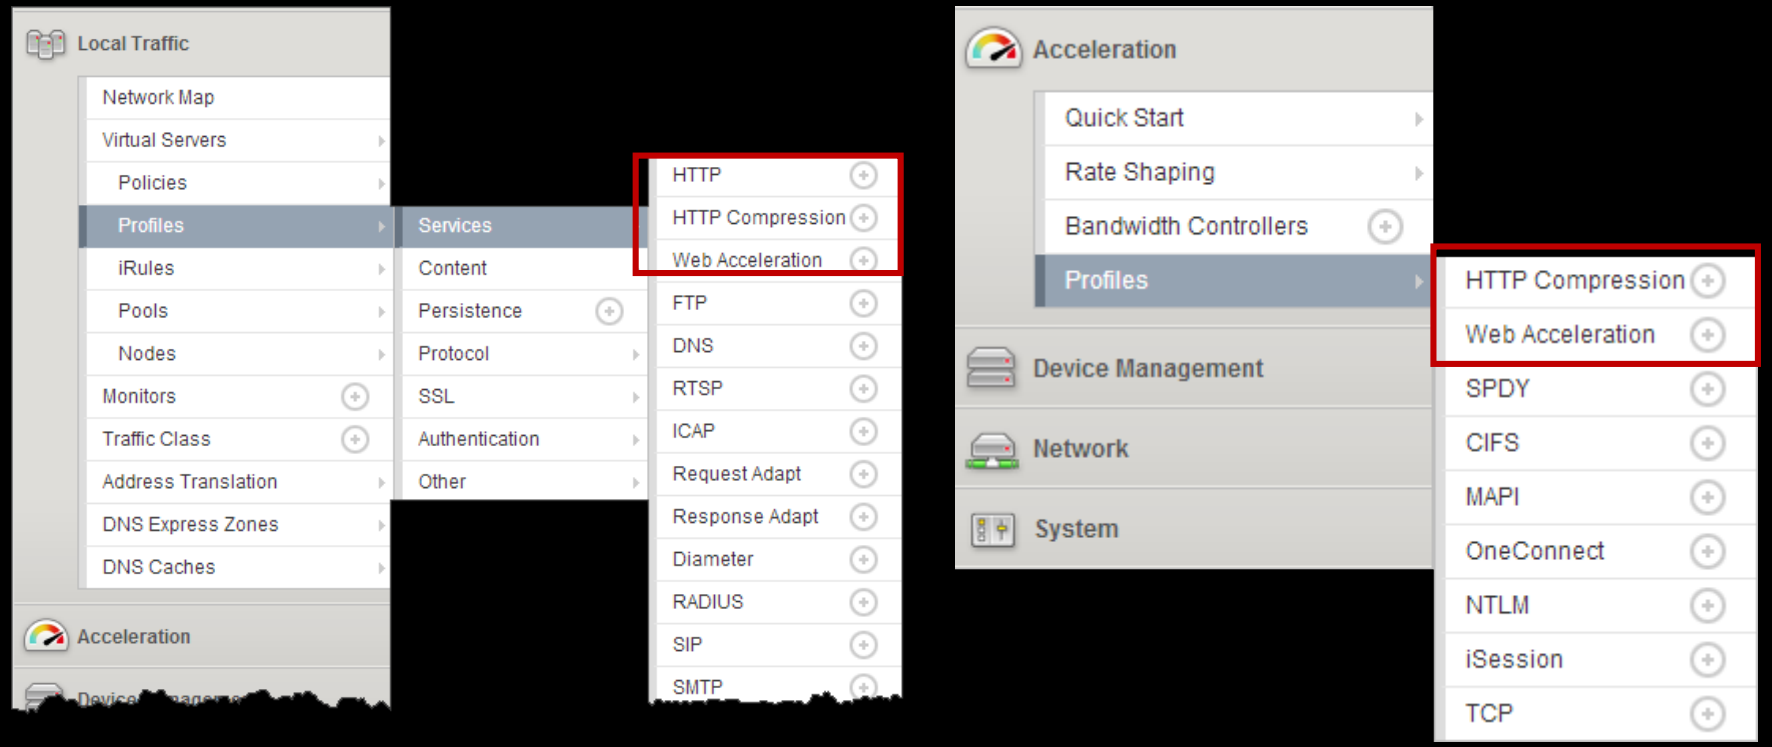

HTTP Profil

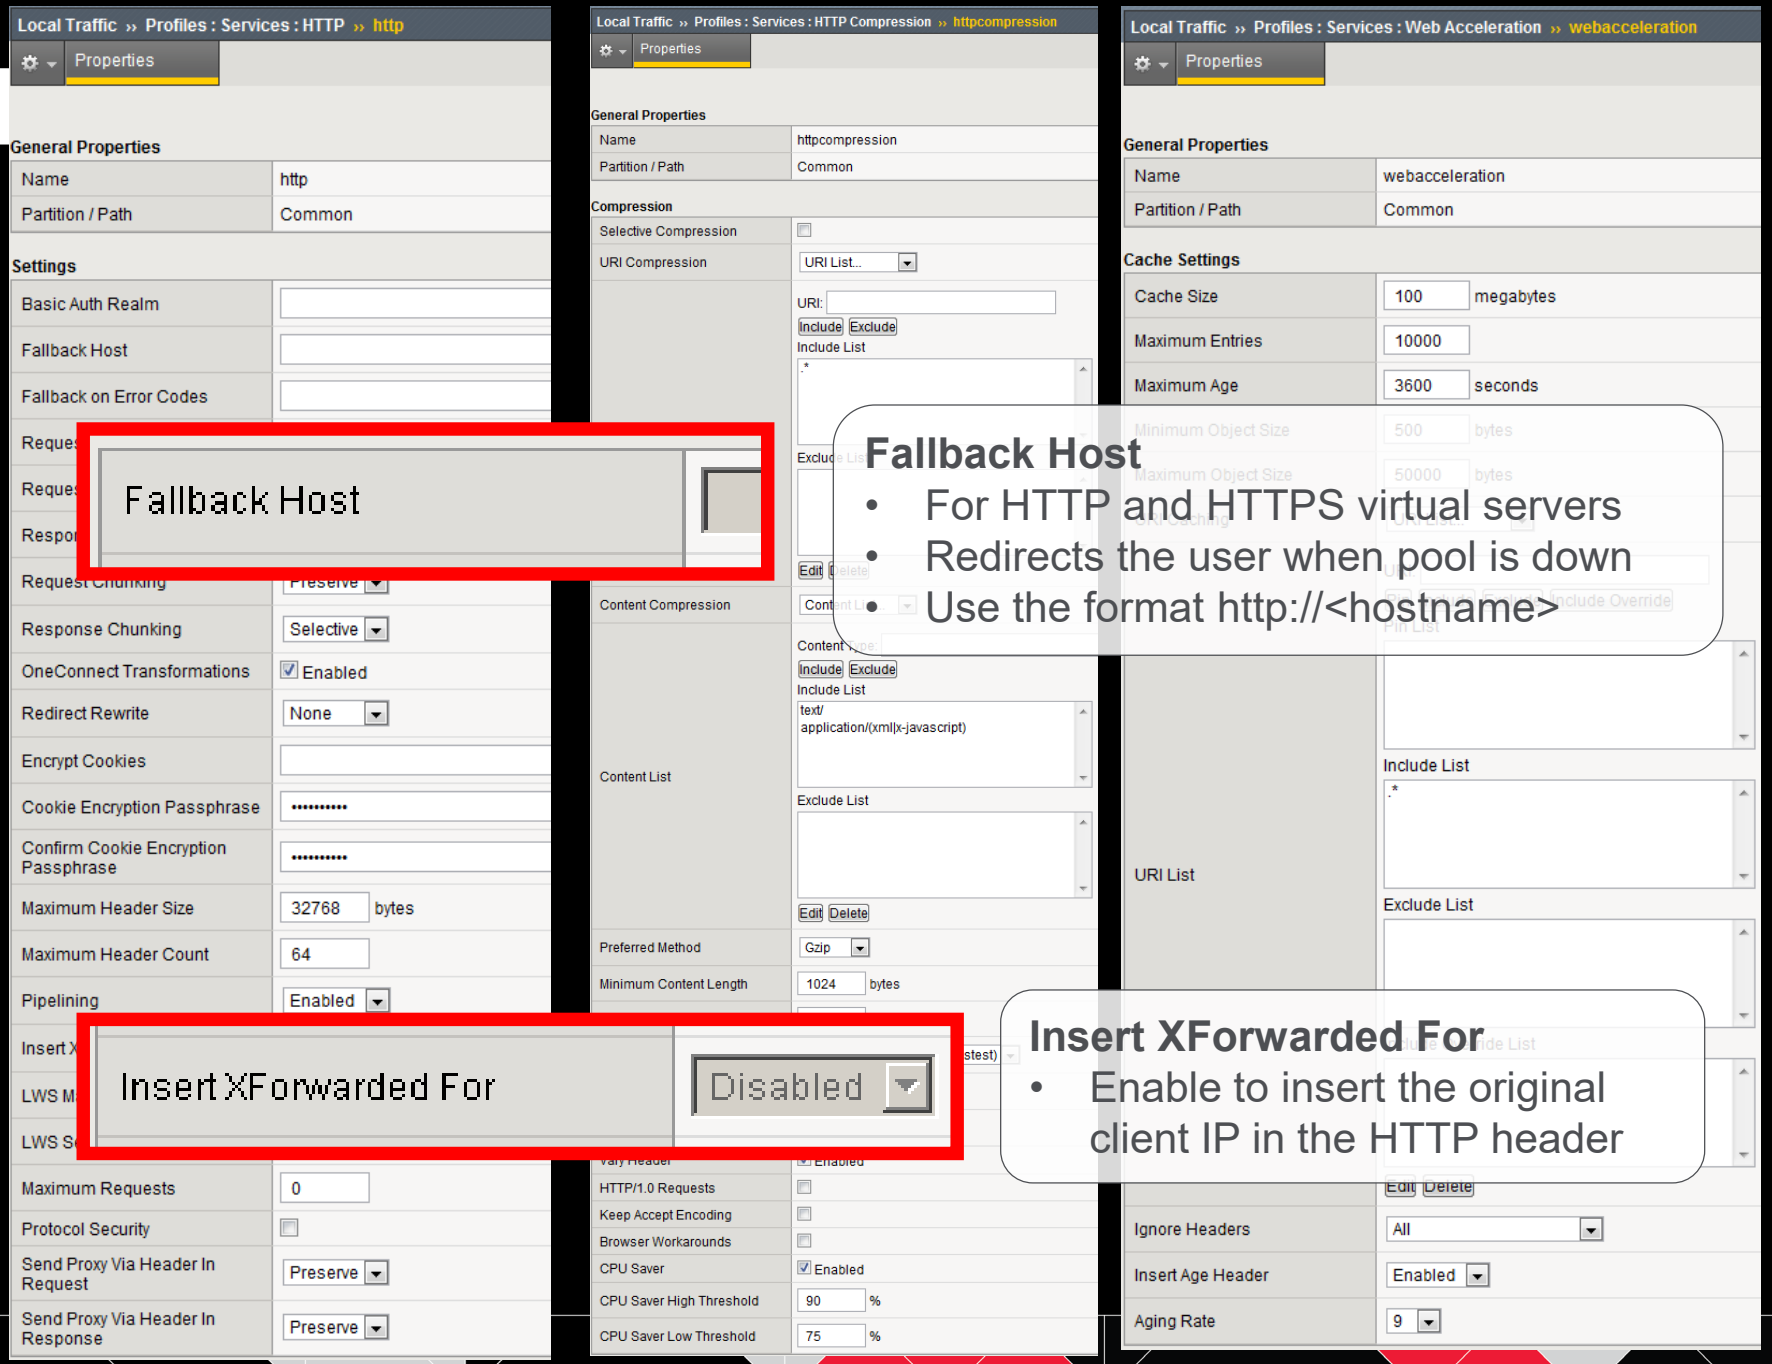

Fallback Host

- For HTTP and HTTPS virtual servers
- Redirects the user when pool is down
- Use the format http://<hostname>

Insert XForwarded For

- Enable to insert the original client IP in the HTTP header

COMPRESSION

Using Standard HTTP Compression

- **Benefits:**
  - Clients get data more quickly
  - Reduces bandwidth usage
  - Less data to encrypt

- **Drawbacks:**
  - Adds an extra load to both the client and server
  - Often requires compression software/hardware on every web server

BIG-IP Intelligent HTTP Compression

- **All the benefits previously discussed**
- **Offload compression from the servers, reducing CPU**
  - This can be accomplished even if the servers are currently compression
  - We will review this on the next slide
- **If iRules or BIG-IP policies are required to manipulated the data**
  - HTTP headers are not compressed
- **Compress on a per virtual server**
- **Compress based on URI or file type**
- **Compression can be scaled based on CPU load**
  - Important if you are doing compression through the software
- **Base Licensing**
  - On current hardware, max compression part of base license
  - Compression in hardware on X200 series.

DEVICE SERVICE CLUSTER (DSC)

Device Service Cluster (DSC)

- **DSC is a series BIG-IPs supporting each other**
  - May also be referred to as Centralized Management Infrastructure (CMI)

- **Each BIG-IP has a Device Object for itself containing;**
  - Unique device information
  - A Certificate for building trusts
  - Device HA and failover settings for the local device

- **BIG-IPs are then placed in Device Trust Groups**
  - Exchange certificates for secure communications
  - Exchange HA settings

- **BIG-IPs in a Trust Group are combined into Device Groups**
  - A device group may support config sync and failover
  - Or synchronization of selected configuration items only

Sync Only Device Groups

- **Allows flexible membership**
  - Different hardware platforms
  - Different license/modules

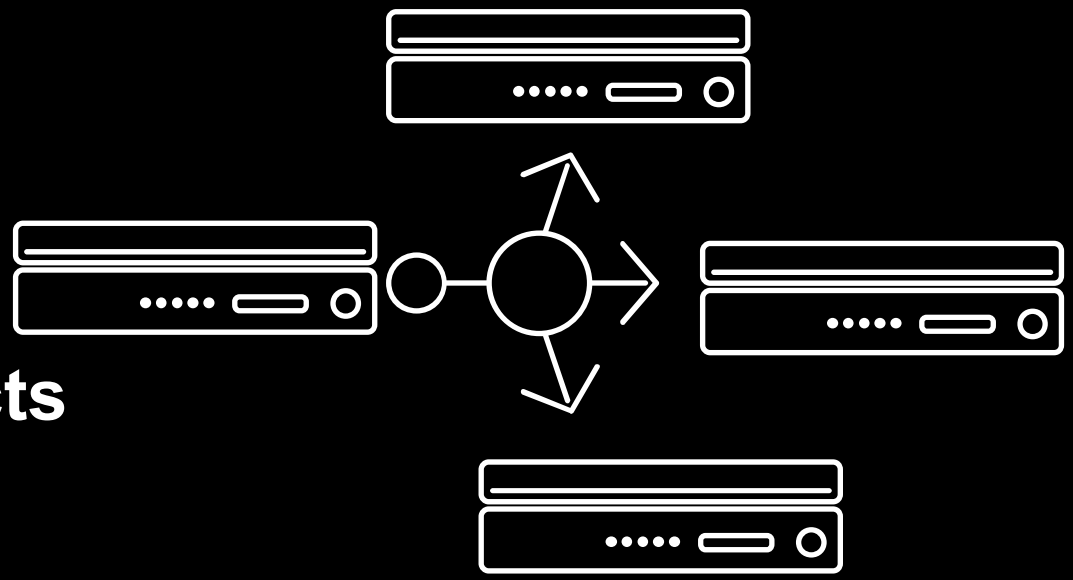

- **Can be configured to auto-sync objects**
  - Certificates
  - CRL
  - Data groups
  - External monitors
  - iApps
  - iRules
  - Policies
  - Profiles
- **Max of 32 Sync-Only groups are supported**
- **Device trust uses built-in sync-only group “device_trust_group”**
  - Auto-sync enabled
  - Adding devices to trust-domain auto-adds to device_trust_group

Traffic Groups

- **A group of listeners (IP addresses)**
  - Virtual address
  - Self IP address
  - SNAT
  - NAT

- **Two different types**
  - Non-floating
  - Floating

Going Beyond Two Devices

If a customer wishes to add devices to the sync-failover group

- Set up the configsync and HA configuration and add the new device to the trust
- Assign the new to the sync-failover device group and sync to the new device
- Adjust traffic group

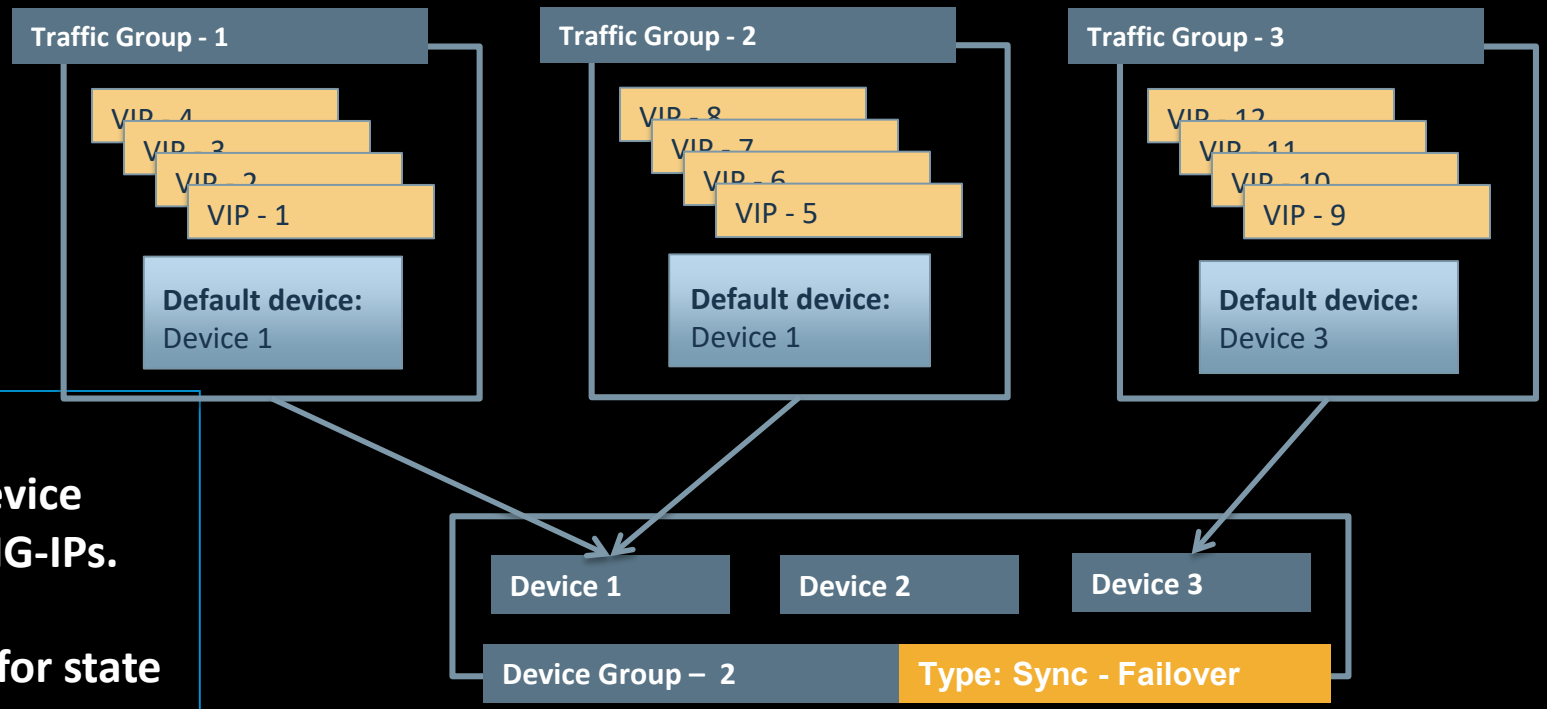

Caveats:

- **GTM can only monitor device groups of two or fewer BIG-IPs.**
- **V11.4 or higher required for state mirroring.**

Oh! In case you missed it!

- **Automatic synchronization is now available for sync-failover**
  - Disabled by default

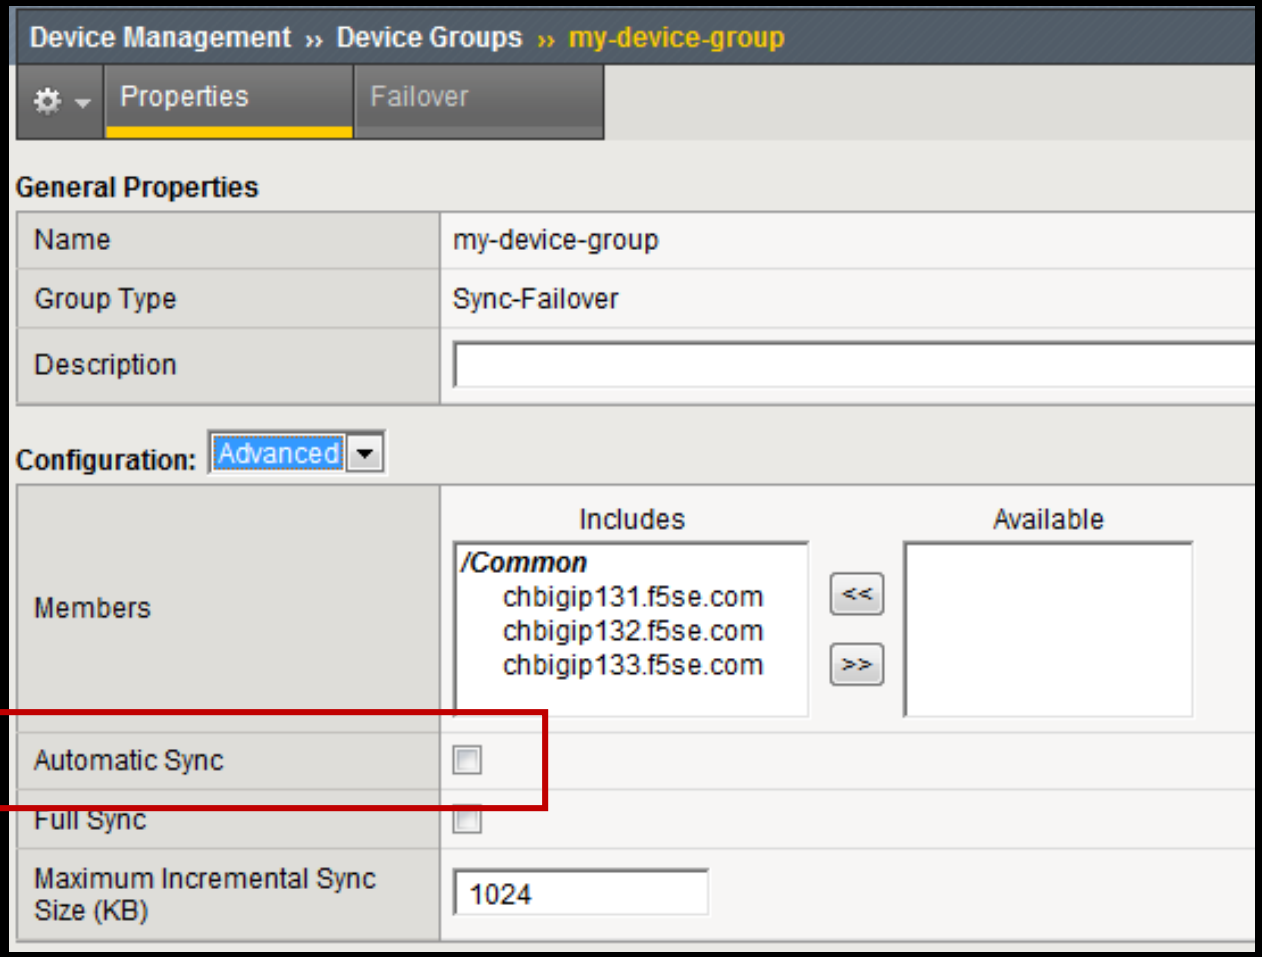

F5 DoD Account Team

Air Force		AE / East	Eddie Augustine	e.augustine@f5.com	301-717-4131
		AE / West	Dustin Purkey	D.Purkey@F5.com	714-501-4815
		SE / East	Arnulfo Hernandez	A.Hernandez@f5.com	202-360-1984
		SE / West	Paul Deakin	p.deakin@f5.com	949-395-0051
DISA		AE	David Thomas	d.thomas@f5.com	703-930-9623
		AE	Thomas Ries	T.Ries@f5.om	703-850-4654
		SE	Anthony Graber	anthony.graber@f5.com	443-987-6487
Navy Marine Corps		AE / East	John Manning	j.manning@f5.com	703-898-4135
		AE / West	Archie Newell	a.newell@f5.com	858-922-2654
		SE /East	Paul Simmons	p.simmons@f5.com	843-300-7392
		SE / West	Jimmy Jennings	j.jennings@f5.com	951-334-8558
Pentagon Defense Agencies		AE	Mark Oldknow	m.oldknow@f5.com	512-410-9462
		SE	August Weinerstein	a.winterstein@f5.com	301-660-9644
Army		MAM / West	Brig Lambert	B.Lambert@f5.com	801-319-1221
		MAM / East	Todd Favakeh	t.favakeh@f5.com	847-334-5610
		SE /East	Shaun Simmons	s.simmons@f5.com	412-329-8366
		SE / West	Michael Slavinsky	M.Slavinsky@f5.com	206-637-2056

F5 DoD Virtual User Group (DoDVUG) Schedule

Date	Title	F5 DoDVUG Topic
Apr 9th Thursday@ 1500	F5 DoD Virtual User Group #1	F5 Access Policy Manager with remote access, network tunneling, and CAC/PIV Authentication.
April 23rd Thursday@ 1500	F5 DoD Virtual User Group #2	Get Your SaaS in Gear Enterprise Application Strategy
May 7th Thursday@ 1500	F5 DoD Virtual User Group #3	Ghastly Wealth Compliance using F5 ASM
May 21st Thursday@ 1500	F5 DoD Virtual User Group #4	Automation/Orchestration - F5 A/O Toolchain
June 4th Thursday@ 1500	F5 DoD Virtual User Group #5	SCCA / SACA
June 18th Thursday@ 1500	F5 DoD Virtual User Group #6	SSLO Orchestrator

Thank You