F5 101 Notes by Suhesh Kasti

Section 1: OSI

1.01 – Explain, compare, and contrast the OSI layers

Layer 7 – Application

- Interacts with the user (FTP/HTTP/SMB/SSH/etc.)

Layer 6 – Presentation

- Converts information into data structures that are understandable by/useful to the system (XML/TLV/JSON)
- SSL/WEP/WPA

Layer 5 – Session

- Allows two endpoints to exchange data for a period of time.
- NetBIOS, TCP/IP Sockets, RPCs– Not necessarily the length of a TCP connection

Layer 4 – Transport

- Facilitates communication between multiple applications on diferent computers.
- Multiplexes and de-multiplexes multiple applications onto one network.
- Establishes, manages, and terminates connections.
- TCP/UDP ports.

Layer 3 – Network

- Routing and addressing.
- Internet Protocol (IPv4, IPv6)/OSPF/ICMP/ARP.
- Addresses packet for the final network destination.
- Uses IP addresses.

Layer 2 – Data-link

- Encapsulates upper layers into frames.
- Addresses frame for the destination device on the same network.

- Uses MAC addresses.
- 802.11/Ethernet/ATM/Frame Relay.

Layer 1 – Physical

- Converts bits to and from whatever transmission language is used by the physical hardware.
- Hubs/Repeaters/Fiber/UTP/Connectors.

1.02 – Explain Protocols and Technologies Specific to the Data Link Layer

Switch Forwarding Database

- Learns MAC addresses of devices on the local broadcast domain, and which port they were learned on.
- Helps the switch decide where to send each frame.ARP
- Helps a device discover the Layer 2 address of a corresponding Layer 3 address, in order to address its frame.
- "Glues" together the IP and Ethernet protocols.

MAC/Ethernet addresses

- All networking devices have a physical burned in address (BIA) which they use for communications on the local network.
- 6 bytes (48-bits) long
- First 3 bytes identify the vendor that created the device (or NIC), also known as the OUI (Organisationally Unique Identifier)
- Last 3 bytes uniquely identify the device
  - e.g. 78:31:c1:c8:22:8b

Broadcast domain

- A segment of the network where a device can transmit directly to another device without having to go through a router.
- 1 VLAN = 1 Broadcast domain.
- Broadcast frames cannot leave the broadcast domain.

VLANs

- A virtual broadcast domain created inside of a switch.
- Before VLANs, everything connected to a switch would be in the same broadcast domain.

- Now, a switch can create many virtual broadcast domains (VLANs) and allocate different ports to different broadcast domains.
- Devices on different VLANs can communicate through a router or a Layer 3 switch.
- A trunk port can carry multiple VLANs.

Link Aggregation

- F5 BIG-IP calls port-channels, or link aggregation "trunks".
  - A trunk is an etherchannel, or link aggregation.
  - A VLAN trunk is a regular trunk (multiple VLANs traversing one link).– Link aggregation allows you to "bond" or "bundle" multiple physical interfaces into one logical interface.
- Traffic is load balanced over the links using a hash of the source and destination L2 addresses.
- Link Aggregation Control Protocol (LACP, 802.3ad) detects errors on a link agg, and allows the two devices to negotiate parameters for the connection.

1.03 – Explain protocols and apply technologies specific to the network layer

Subnetting

- Basically, subnetting and how IPv4 addresses are structured and work.
- How do routing decisions get made?
  - Longest match wins.
  - ip.dst=10.1.3.24.
  - Routes: 10.0.0.0/8, 10.1.0.0/16, 10.1.3.0/24, 10.1.3.0/25.
  - 0.1.3.0/24 would be the route that gets matched, because it has the longest prefix/mask.

Routing Protocols

- Distance Vector routing protocols (RIP/BGP).
- Link State routing protocols (OSPF/ISIS).

Fragmentation

- If a packet needs to traverse a network with an MTU which is smaller than its size, it must be fragmented so that it fits.

- Typical MTU on an Ethernet network is 1500 bytes.
- A new IP header is put onto each fragment, indicating its size and fragment number.

TTL– Counts down from a value with each L3 hop a packet takes.

- Stops traffic from getting into an endless routing loop.
- When TTL reaches 0, the traffic is dropped.
- Maximum TTL is 255, because it is an 8-bit field in the IP header.
- Show source/dest IP/MAC at each hop
- Make sure you know how the L2 and L3 src/dst addresses change at each point in the network.
- When leaving a router, the L2 src address will be the router's interface address, and the L2 dst will be the address of the next router, or the destination host.
- The L3 src/dst addresses never change unless NAT is involved.

1.04 – Explain the features and functionality of protocols and technologies specific to the transport layer

MTU vs. MSS

- MTU is maximum size frame that can be sent at L2/L1.
- MSS (maximum segment size) is a parameter determined by two L4 hosts when they establish a connection.
- Neither host will send a TCP segment larger than the MSS.
- MSS is supposed to be used to limit the amount of fragmenting that is required, as fragmenting uses additional resources.
- Decrease the MSS to decrease the likelihood of fragmentation.

TCP

- TCP does error checking with checksums added to the TCP header.
- Each packet goes in order, and is accounted for by using acknowledgements.
- If a packet isn't acknowledged, it must be resent.

- Three-way handshake: SYN,SYN-ACK,ACK.

UDP

- Unreliable.
- Lower overhead than TCP.
- No acknowledgements.

Ports

- Allow multiple applications to run using one L3 address
- FTP and HTTP can run on the same IP address, using diferent ports (21 and 80)

TCP Reset (RST)

- Allows a participant in a TCP session to abort the connection.
- Typically used by a client when no acknowledgements are being received, and the connection appears unusable.

Delayed Binding

- The F5 BIG-IP will delay binding, which means it waits until the TCP session with the client is complete (3way handshake done) before it will connect to the server, and bind the client's session to the server.
- Protects the server from SYN flood attacks.

1.05 – Explain the features and functionality of protocols and technologies specific to the application layer

HTTP

- TCP 80
- HTTP 1.0 defined GET/POST/HEAD.
- HTTP 1.1 defined OPTIONS/PUT/DELETE/TRACE/CONNECT.

HTTP STATUS CODES

- 1xx – Informational
  - 100 Continue: Server received initial part of request; client should continue.
  - 101 Switching Protocols: Client requested server to switch protocols.

- 2xx – Success
  - 200 OK: Request succeeded.
  - 201 Created: Resource created as a result of the request.
  - 202 Accepted: Request accepted for processing, but not yet completed.
  - 203 Non-Authoritative Information: Response from a proxy, not original server.
- 3xx – Redirect
  - 300 Multiple Choices: Multiple options for resource (client chooses).
  - 301 Moved Permanently: Resource has been moved to a new URI.
  - 302 Found: Resource temporarily at a different URI.
  - 307 Temporary Redirect: Request method should not change on redirect.
  - 308 Permanent Redirect: Resource moved permanently (method should not change).
- 4xx – Client error
  - 400 Bad Request: Server could not understand the request due to invalid syntax.
  - 401 Unauthorized: Authentication is required and has failed or not provided.
  - 402 Payment Required: Reserved for future use (often experimental).
  - 403 Forbidden: Server understood the request but refuses to authorize it.
  - 404 Not Found: Resource could not be found.
  - 405 Method Not Allowed: Request method is known but not allowed for this resource.
- 5xx – Server error
  - 500 Internal Server Error: Generic server error.
  - 501 Not Implemented: Server does not support requested functionality.
  - 502 Bad Gateway: Server received invalid response from upstream server.
  - 503 Service Unavailable: Server is currently unavailable (overloaded

or down).

- 504 Gateway Timeout: Server did not receive a timely response from upstream server.

HTTP Methods:

- HEAD – Just returns response headers, no body.
- GET – Request data to the server.
- POST – Data is sent to the server with the request (submit forms, etc.).
- HTTP keep-alive used to re-use an existing HTTP connection instead of creating a new one.

DNS

- UDP 53.
- Resolves names into IP addresses.
- Hierarchical distributed naming system.

SIP (Session Initiation Protocol)

- UDP/TCP 5060 + 5061.
- Voice connection over the network.
- Allows video conferencing, presence, IM and voice.
- Enables unified communications.

FTP

- Used to transfer files between hosts.
- Uses separate control and data connections.
- Can use authentication, but is in clear-text.
- Also allows you to connect anonymously.
- Control port is generally TCP 21.
- In *Active mode*, data port is TCP 20.
- *Active mode*, the client specifies a port it is listening on for the server to connect the data channel on (server initiates data channel to client).
- *Passive mode*, the server tells the client a random high (>1023) port to connect to for the data channel.
- *Passive mode* is easier on the client's firewall, as no inbound connections need to be allowed.

SMTP (Simple Mail Transfer Protocol)

- Mail delivery protocol.
- TCP 25.
- HELO (say hi).

- EHLO (say hi, and use extended mode).
- MAIL FROM: (sender).
- RCPT TO: (recipient).
- DATA (body).

Cookies

- State information stored by the web server on the user's disk, to be later retrieved by the server.
- Name-Value pairs.
- Allows servers to remember you, or information about you/your session, regardless of your IP address.

The Name Resolution Process

http://www.tcpipguide.com/free/t_DNSNameResolutionProcess-2.htm

- User tries connecting to suhesh.com.np
- User's system looks in host file.
  - *Linux*: /etc/hosts
  - *Windows*: C:\Windows\System32\drivers\etc\hosts
- User's system looks in local DNS cache.
- User's system queries its local DNS server (LDNS) for suhesh.com.np. *For a home user it can be your router.*
- If still not found, user's system queries the recursive DNS server for suhesh.com.np.
  - *This can your ISP's DNS server or if you have set a private DNS server like 1.1.1.1 or 8.8.8.8*
- Recursive looks in its cache.
  - *It is called recursive DNS server because from now on this server will reach various DNS servers to find the requested doman name*
- Recursive DNS queries root servers for server that is authoritative for .com.np
- Root server will return the nameserver of Mercantile (*A company responsible for managing .com.np domain*)
- Recursive DNS now queries nameserver of mercantile to provide the authoritative nameserver (*the nameserver that has zone file for the domain*) for suhesh.com.np

- Recursive DNS now queries the authoritative nameserver (in my case cloudflare) for the A record for the domain suhesh.com.np
- Recursive DNS caches the response, sends it to the user, who also caches the response, and connects to the IP address.

  URLs
- URL – Uniform Resource Locator– A type of a Uniform Resource Identifier (URI)
- URL includes the protocol used to access the resource, URIs do not necessarily
- protocol://[user:pass@]host:port/path/to/resource?query#fragment

Section 2: F5 Solutions and Technology

2.01 – Articulate the role of F5 products

Application Acceleration Manager (AAM)

- Application optimisation.

  Advanced Firewall Manager (AFM)
- Firewall/anti-DDoS/traffic management/app security/DNS security

  Access Policy Manager (APM)
- Manages access to applications/resources.
- VPN/authentication server.
- Supports Citrix, VMware view, RDP and more.

  Application Security Manager (ASM)
- Protects web apps, enables regulatory compliance.

  Global Traffic Manager (GTM)
- DNS load balancing.
- Distributes DNS responses between DCs or locations based on policy or load.
- DNSSEC.

  Local Traffic Manager (LTM)

- Full proxy that sits between users and servers.
- Allows you to secure, optimise, and load balance application traffic.
- Health monitoring for the servers.– SSL offloading.

2.02 – Explain the purpose, use, and advantages of iRules

- Scripts that allow you to interact more directly with the live traffic.
- Allows you to play with the header and/or payload of the traffic as it flows through the F5.
- Tool command language (Tcl) is the language.
- Uses the Universal Inspection Engine (UIE) to search the packet.
- Lives in bigip.conf.
- Configured through CLI or GUI.
- Precompiled into byte code for better performance.
- Applied to virtual server.
- Use events to trigger an iRule at a particular stage of the flow.
- Custom logging/redirect/modifications/masking.

2.03 – Explain the purpose, use, and advantages of iApps

- A template and analytics system to ease deployment of new apps/configurations.
- Customisable.
- Comes with bundled configs.

2.04 – Explain the purpose of and use cases for full proxy and packet forwarding/packet based architectures

Proxies

- Acts on behalf of the user when talking to a server.– Can manipulate requests or response.
- Often logs connections.

Forward Proxies

- Generally web proxies.
- Act as a boundary between networks, sending requests on behalf of the user, where the user may not be allowed to.

Reverse Proxies

- Load balancers.
- Sit in front of applications and process the client requests in order to take load of the server.

Half Proxies

- One side of the conversation goes through the proxy (requests), but the other does not (responses).
- Can increase performance.
- Another use for the term "half proxy" is when the initial connection to the server is delayed by the proxy, in order to inspect the content, or make decisions, but once the initial request has gone through, subsequent traffic is not delayed, and passes normally.

Full Proxies

- Acts as a man in the middle on the connections.
- Connections from the client terminate at the proxy.
- Connections from the server terminate at the proxy.
- The proxy stitches the client and server connection together, and is able to inspect or take action on any part of the connection.
- Users can be completely air gapped from the servers by using a full proxy.
- F5 BIG-IP is a full proxy.

Packet forwarding architectures are generally faster than full proxies, because they do not understand the full protocol, and are able to forward the packets without completely interpreting the protocol, thus they have less power.2.05

Explain the advantages and configurations of high availability (HA)

- A redundant setup is a configuration that allows traffic to flow even if one of the F5's goes down.
- The redundant box automatically takes over from the other box when it goes down.
- This process is called failover.
- While operating, the config is synced from the main unit to ensure that the peer operates identically.

  ### Active/Active
- An active/active configuration is one where multiple units are actively processing connections at the same time.
- This can be used to increase throughput, however it is also recommended to have standby devices available in case one of the active units fails, otherwise the other active units may not be able to take up the slack of the failed device.
- Device Service Clustering (DSC) from BIG-IP 11.0.0 allows for more than 2 devices in a HA setup.

  ### Active/Standby
- Active/standby is when one device is active, and another is standing by, ready to take over if the active fails.
- Both devices sync their configurations so they are ready to take over at any point.

  "**Sync-Failover**" device group type on F5.
- Configuration objects are automatically added to the default traffic group, which floats to the standby when the active fails.
- "Failback" is when the device that was previously in a standby state, transitions back to a standby state.

Section 3: Load Balancing Essentials

3.01 – Discuss the purpose of, use cases for, and key considerations related to load balancing

- When an application grows in size, a single server may no longer be able to serve all of the users on its own.
- Load balancing distributes the load of the application between multiple servers.

  ### LTM Load Balancing Methods

  LTM load balances to a pool of member servers using one of the below load balancing methods.

  ### Round Robin (default)

- Evenly distributes load across all members.
  - One for A, one for B, one for A, one for B…

  ### Ratio (member/node)

- Like round robin, but each server has a weight.
  - One for A, one for B, one for A, one for A, one for B, one for A…
  - In the above, A has a weight of 2, and B has a weight of 1.
- This can be set up on the member level (port on server), or the node level (server IP).

  ### Dynamic Ratio (member/node)

- Like ratio, but works of of the current load on the server (CPU/mem/etc.).
  - Software must be installed on the servers so that the BIG-IP can detect the current load level.
- Typically uses WMI/SNMP.

  ### Fastest (node/application)

- Selects the server with the least number of L7 responses outstanding.
  - Virtual server must have a L7 profile attached.

  ### Least connections (member/node)

- Selects the server with the least number of active connections.

  ### Weighted least connections (member/node)

- Same as above, but each member/node has a connection limit set in order to calculate a weight for each member/node.

- The server at the lowest relative capacity gets the connection.

Observed (member/node)

- Ranks servers over time based on the number of connections, not often needed or used.
- Not recommended for large pools.

Predictive (member/node)

- Similar to observed, but analyses the rank trend and predicts where it is going. Not often needed or used.
  - Also not recommended for large pools.

Least sessions

- Selects the server with the least entries in the persistence table.
- Persistence must be enabled.

Ratio least sessions

- Same as least sessions, but each server has a ratio assigned (like in "Ratio")

Persistent vs Persistence

Persistent

- Persistent is used to describe the behaviour of HTTP, TCP and database connections.
- Persistence is related to TCP/HTTP connection handling, normally in relation to load balancing.
- Persistent connections are ones that are kept open, and able to be re-used.
- HTTP connections are persistent, as of HTTP 1.1
- Less load generated by TCP 3way handshakes and teardowns when a connection is reused.
- Less delay until a request can be served because the connection already exists.

Persistence (aka. stickiness/server affinity)

- Allows the load balancer to send requests from a particular user to the same server.
- Used when the user has state information (a shopping cart) stored on a particular server.

- If the user went to a diferent server, he would not be able to access his cart.
- Can be based of of many attributes, including cookies and IP addresses/subnets.

3.02 – Differentiate between a client and server

- Client requests resources/actions of the server.
- Server generally higher powered, and used by multiple clients.

Section 4: Security

4.01 – Compare and contrast positive and negative security models

Positive Security Model	Negative Security Model
Defines what is allowed and implicitly denies everything else	Defines what is not allowed and implicitly allows everything else
Network firewalls, Access Control Lists (ACLs)	Intrusion Prevention Systems (IPS), Web Application Firewalls (WAFs, e.g., F5 ASM)
Blocks unknown (0-day) attacks, assuming their behavior is not whitelisted	Allows everything unless there's a match with the blacklist or signatures
Can be difficult to ensure all legitimate behaviors are allowed, risking some legitimate behavior denial	Easier to implement as it allows all except explicitly blocked actions
Access lists can be turned into a negative security model by adding a permit all statement at the bottom.	Typically designed to block specific known threats, no whitelist conversion

4.02 – Explain the purpose of cryptographic services

Signing

- A mathematical way of demonstrating/verifying the authenticity of a message.
- Ensures the sender is who they say they are, and that the message is what they actually said (it wasn't modified).
- Use asymmetric crypto.

Encryption

- Can provide authentication, confidentiality, integrity, and non-repudiation (proves the sender really sent the message).
- Three types of crypto: Symmetric, asymmetric and hashing.
- Unencrypted data is referred to as plaintext, encrypted data is referred to as cipher text.

Certificates/certificate chains

Applications trust a root certificate authority.

- Any certificate that is signed by a trusted CA is also trusted by the client/application.
- Root CAs generally delegate their signing duties to an issuing CA in order to introduce a layer of separation.
- If an issuing CA is revoked, all of its issued certificates are now invalid.

Private/Public keys

- Private key encryption is when the key must remain a secret.
- The key is used to both encrypt and decrypt messages.
- Public key encryption involves the use of two keys, one to decrypt, and one to encrypt.
- The key used to decrypt is kept secret, and the key used to encrypt is made public so that anybody can use it to encrypt a message to the recipient.

Symmetric/Asymmetric encryption

- Symmetric encryption uses only private keys.
- Whoever has the key can both encrypt and decrypt any messages created with it.
- Asymmetric encryption uses a pair of private/public keys.

4.03 – Describe the purpose and advantages of authentication

- Authentication is determining if the user is who they say they are.
- Authentication methods can include something the user knows (password), that they have (RSA token), or something they are (fingerprint).

  Single Sign On (SSO)
- Reduce the number of accounts a user has to remember.
- Uses a single account (e.g. AD account) for multiple applications/services.

  Multifactor authentication (MFA)
- More than one form of authentication is required in order to verify identities.
- A password as well as a security token may be required.
- If a single factor is compromised (user tells somebody their password), the other person will not be able to get in unless they also have the security token.
- More secure than single factor authentication (SFA)

  AAA (Authentication, Authorization and Accounting)
- Authentication verifies the user is who they say they are.
- Authorization verifies whether the user is allowed to do what they are trying to do (issue commands, access a page).
- Accounting logs the resources that a particular user accesses.
- Servers interface with a AAA server using *Remote Authentication Dial-In User Service (RADIUS)* or *DIAMETER*.

4.04 – Describe the purpose, advantages, and use cases of IPsec and SSL VPN

IPsec

- Protocol suite that authenticates and encrypts each IP packet in a session.

- Operates at the Network layer.
- Often used for a permanent tunnel between two offices.
- Not generally used for secure remote user access, as SSL VPNs are much easier to use and deploy.
- Best solution for site to site tunnels.

SSL VPN

- Best solution for remote users.
- Can operate as a VPN in the browser.
- Much easier to deploy than IPsec VPNs.
- Often operate as ActiveX or Java clients, launched from the browser.

Section 5: Application Delivery Platforms

5.01 Describe the purpose, advantages, use cases, and challenges associated with hardware based application delivery platforms and virtual machines

Virtual

- Virtual editions are quicker, and more flexible to deploy.
- Can be run in the cloud easily on something like Amazon AWS.
- F5 offers BIG-IP Virtual Editions (VEs) which can run on most hypervisors.
- Much quicker to scale.
- Speed of SSL transactions per second (TPS) does not compare with 4000 series hardware and above.
- Can do up to 4Gbps of SSL/compression throughput, 10Gbps non-SSL/non-compressed.

Hardware

- Hardware solutions offer a single vendor solution, so there is no software vendor blaming the hardware vendor for issues.

- Hardware solutions also offer higher performance, as it is built for the purpose of application delivery.
- Hardware allows you to offload processing (SSL, compression) from the servers, meaning less servers are required.
- Takes longer to acquire than VMs.
- Biggest appliance can do 40Gbps SSL or compression throughput.
- VIPRION blades can do up to 20Gbps SSL/compress each.

5.02 – Describe the purpose of the various types of advanced acceleration techniques

Optimization

- Implement TCP features missing from the client or server in order to speed up application delivery.
- Selective TCP acknowledgements can be used to overcome packet loss.
- Pool server-side TCP connections to reduce server load associated with creating sessions.
- BIG-IP can spoof different hosts so that the client opens more connections to speed up the page load process.
- Client is trying to get 30 objects for suhesh.com.np, but will only open 2 sessions per domain.
- F5 can rewrite some of the objects so that they appear to exist on suhesh.com.np, and the client will open another 2 connections to that domain.

HTTP Caching

- Saves content close to the user once it has been loaded at least once.
- The client browser contains a cache that it can use to load objects that it has seen before, and that are still fresh.
- Specific cache devices can also be deployed in the data centre or at the remote branch to ease load on the WAN link[s]

Compression

- GZIP is supported by almost every browser and server.
- Acceleration devices can offload compression from servers.

Pipelining

- Sends multiple requests over a connection at once, instead of sending one and waiting for the reply before sending the next.
- Saves the round trip time (less waiting).
- Servers still need to send back the responses in parallel, and wait for acknowledgements, so the effect of pipelining is minimal.
- Can be more of a security risk than the minimal optimisation is worth.

Glossary:

- **BIG-IP:** F5's software and hardware offerings. ie. "BIG-IP Virtual Edition"or "BIG-IP 4000 Series".
- **Member:** Server that has been added to a pool, includes destination port.
- **Node:** Server that does not exist in a pool, but should have traffic sent to it.
- **Pool:** A group of members that are attached to a virtual server, and will be used to serve clients.
- **TMOS:** The operating system running on a BIG-IP.
- **VIPRION:** F5's blade based hardware chassis.
- **Virtual Server (VIP):** The server that clients connect to on the BIG-IP to have their traffic handled by a server.

This Note is made by Suhesh Kasti and there are many other useful stuff in my site. Click on my name to reach my site.

If you have any queries or you have any ideas improvements to make you can contact me from here