




F5 CSP 2018 Hands-on Lab Guide

F5 CSP Big Data Training

https://github.com/jarrodlucia/bigip_elasticsearch/



Contents:

1	Welcome	5
2	Getting Started	7
3	Prerequisites	9
3.1	UDF Blueprint	9
4	Lab Topology	11
5	Class 1: BIG-IP AVR (BIG-IP Goodness)	13
5.1	Module 1: REST API Basics	13
5.2	Module 2: F5 Application Visibility and Reporting	21
6	Class 2: Introduction to ELK Stack (ELK Coolness)	55
6.1	Module 1: ELK Stack Build Ubuntu Server	55
6.2	Module 2: Kibana and Visualisation	77
7	HOWTOs: Index	87
7.1	HOWTO - how to do stuff	87

Welcome

Welcome to F5's Service Provider AVR and Big Data hands on training series. The intended audience for these labs are Service Provider engineers that would like to leverage the power of F5 data visibility and integrate this immense data capability into open source tools such as Elasticsearch, Hadoop and others.

Getting Started

Please follow the instructions provided by this documentation to start your lab and access your lab.

Note: All work for this lab will be performed exclusively from the Linux Jumphost and Linux Client Machines. All required access and services needed to perform classes and labs are provided by the UDF. No installation or interaction with your local system is required.

Prerequisites

In order to complete this series of training classes you will need to utilize the provided blueprint for the course session. To access the UDF sessions you will need to have the following prerequisites met.

- Current Access to UDF
- SSH key of your access machine in UDF
- Windows or MAC ssh client working with UDF

All pre-built environments implement the lab-topology shown below.

3.1 UDF Blueprint


Please follow the instructions provided by your lab instructor to access your lab environment. The lab environment will be delivered via UDF blueprints to each student.

Note: Please deploy and start your lab as soon as you have access to the class as the lab takes some time to boot all the components.

Lab Topology

The network topology implemented for this lab based on the Service Provider Gi lan path. The focus of the lab is Control Plane programmability and Data Plane elements, so this lab will focus at both parts at different time. The following components have been included in your lab environment:

- 1 x F5 BIG-IP VE (v13.0 HF2)
- 1 x Linux Jumphost (ubuntu 16.04 - mate)
- 2 x Linux Clients (ubuntu 16.04 - mate)
- 1 x Linux Server (ubuntu 16.04)



The following table lists VLANs, IP Addresses and Credentials for all components:

Table 4.1: Lab Network Information

Component	VLAN	IP Address	Credentials
Linux Jumphost	Mgmt	10.1.1.20	
BIG-IP	Mgmt	10.1.1.4	admin/admin
	Internal	10.1.10.5	
	External	10.1.20.5	
	Control	10.1.30.5	
Client 00	Mgmt	10.1.1.9	udfclient/S3rv1ceP0weR
	Internal	10.1.10.25	
Client 01	Mgmt	10.1.1.7	udfclient/S3rv1ceP0weR
	Internal	10.1.10.30	
ELK Stack	Mgmt	10.1.1.5	ubuntu/default
	Control	10.1.30.15	

Class 1: BIG-IP AVR (BIG-IP Goodness)

This class covers the following topics:

- Module 1
 - REST API Basics
- Module 2
 - F5 BIG-IP AVR
 - Configuring AVR
 - Navigating AVR
 - Modify AVR Reports

Expected time to complete: **30 mins**

5.1 Module 1: REST API Basics

In this module you will learn the basic concepts required to interact with the BIG-IP iControl REST API. Additionally, you will walk through a typical Device navigation.

This is a cut down version of the F5 Programmability Super Net Ops training.

Note: The Lab Deployment for this lab includes a single BIG-IP devices. For most of the labs we will configuring the BIG-IP device (management IP and licensing have already been completed).

Note: It's beneficial to have GUI/SSH sessions open to BIG-IP devices while going through this lab. Feel free to verify the actions taken in the lab against the GUI or SSH. You can also watch the following logs:

- BIG-IP:
 - /var/log/ltn
 - /var/log/restjavad.0.log
-


5.1.1 Lab 1.1: Exploring the iControl REST API

Task 1 – Explore the API using the TMOS Web Interface

In this lab we will explore the API using an interface that is built-in to TMOS. This utility is useful for understanding how TMOS objects map to the REST API. The interfaces implement full Create, Read, Update and Delete (CRUD) functionality, however, in most practical use cases it's far easier to use this interface as a 'Read' tool rather than trying to Create objects directly from it. It's usually far easier to use TMUI or TMSH to create the object as needed and then use this tool to view the created object with all the correct attributes already populated.

1. Open Google Chrome and navigate to the following bookmarks: BIG-IP

GUI. Bypass any SSL errors that appear and ensure you see the login screen for each bookmark.






1. Navigate to the URL <https://10.1.1.4/mgmt/toc> (or click the BIG-IP REST TOC bookmark). The '/mgmt/toc' path in the URL is available on all TMOS versions 11.6 or newer.
2. Authenticate to the interface using the default admin/admin credentials.
3. You will now be presented with a top-level list of various REST resources. At the top of the page there is a search box that can be used to find items on the page. Type 'net' in the search box

Table of Contents

iControl REST Resources

[net](#)

and then click on the 'net' link under iControl REST Resources: **Traffic Management**

4. Find the `/mgmt/tm/net/route-domain` **Collection** and click it.
5. You will now see a listing of the **Resources** that are part of the route-domain(s) collection. As you can see the default route domain of 0 is listed. You can also create new objects by clicking the  button. Additionally resources can be deleted using the  button or edited using the  button.
6. Click the 0 resource to view the attributes of route-domain 0 on the device:

/mgmt/tm/net/route-domain/~Common~0

name	0
partition	Common
fullPath	/Common/0
connectionLimit	0
id	0
strict	enabled

Take note of the full path to the resource. Here is how the path is broken down:

```
/ mgmt / tm / net / route-domain / ~Common~0
| Root | OC | OC | Collection | Resource
*OC=Organizing Collection
```


5.1.2 Lab 1.2: REST API Authentication & 'example' Templates


One of the many basic concepts related to interaction with REST API's is how a particular consumer is authenticated to the system. BIG-IP supports two types of authentication: HTTP BASIC and Token based. It's important to understand both of these authentication mechanisms, as consumers of the API will often make use of both types depending on the use case. This lab will demonstrate how to interact with both types of authentication.

Task 1 - Import the Postman Collection & Environment

In this task you will Import a Postman Collection & Environment for this lab. Perform the following steps to complete this task:




1. Open the Postman tool by clicking the  icon of the desktop of your Linux Junphost
2. Click the 'Import' button in the top left of the Postman window




3. Click the 'Import from Link' tab. Paste the following URL into the text box and click 'Import'

```
https://raw.githubusercontent.com/jarrodlucia/bigip_elk_server/develop/postman_collections/F5_Automation_Orchestration_Intro.postman_collection.json
```




4. You should now see a collection named 'F5 Automation & Orchestration Intro' in your Postman Collections sidebar:



5. Import the Environment file by clicking 'Import' -> 'Import from Link' and pasting the following URL and clicking 'Import':

```
https://raw.githubusercontent.com/jarrodlucia/bigip_elk_server/develop/postman_collections/INTRO_Automation_Orchestration_Lab.postman_environment.json
```

6. To assist in multi-step procedures we make heavy use of the 'Environments' capability in Postman. This capability allows us to set various global variables that are then substituted into a request before it's sent. Set your environment to 'INTRO - Automation & Orchestration Lab' by using the menu at the top right of your Postman window:



Task 2 – HTTP BASIC Authentication

In this task we will use the Postman tool to send API requests using HTTP BASIC authentication. As its name implies this method of authentication encodes the user credentials via the existing BASIC authentication.


tion method provided by the HTTP protocol. The mechanism this method uses is to insert an HTTP header named 'Authorization' with a value that is built by Base 64 encoding the string <username> : <password>. The resulting header takes this form:

```
Authorization: Basic YWRtaW46YWRtaW4=
```


It should be noted that cracking the method of authentication is TRIVIAL; as a result API calls should always be performed using HTTPS (F5 default) rather than HTTP.

Perform the following steps to complete this task:

1. Click the 'Collections' tab on the left side of the screen, expand the 'F5 Automation & Orchestration Intro' collection on the left side of the screen, expand the 'Lab 1.2 – API Authentication' folder:



2. Click the 'Step 1: HTTP BASIC Authentication' item. Click the 'Authorization' tab and select 'Basic Auth' as the Type. Fill in the username and password (admin/admin) and click the 'Update Request' button. Notice that the number of Headers in the Headers tab changed from 1 to 2. This is because Postman automatically created the HTTP header and updated your request to include it. Click the 'Headers' tab and examine the HTTP header:



3. Click the 'Send' button to send the request. If the request succeeds you should be presented with a listing of the /mgmt/tm/ltn Organizing Collection.
4. Update the credentials and specify an INCORRECT password. Send the request again and examine the response:




Task 3 – Token Based Authentication

One of the disadvantages of BASIC Authentication is that credentials are sent with each and every request. This can result in a much greater attack surface being exposed unnecessarily. As a result Token Based Authentication (TBA) is preferred in many cases. This method only sends the credentials once, on the first request. The system then responds with a unique token for that session and the consumer then uses that token for all subsequent requests. Both BIG-IP and iWorkflow support token-based authentication that drops down to the underlying authentication subsystems available in TMOS. As a result the system can be configured to support external authentication providers (RADIUS, TACACS, AD, etc) and those authentication methods can flow through to the REST API. In this task we will demonstrate TBA using the local authentication database, however, authentication to external providers is fully supported.

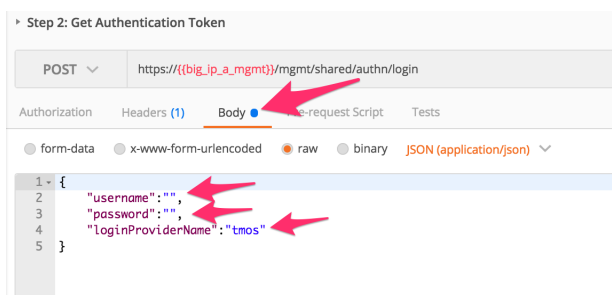
For more information about external authentication providers see the section titled “**About external authentication providers with iControl REST**” in the iControl REST API User Guide available at <https://devcentral.f5.com>

Perform the following steps to complete this task:

1. Click the 'Step 2: Get Authentication Token' item in the Lab 1.2 Postman Collection
2. Notice that we send a POST request to the `/mgmt/shared/authn/login` endpoint.



3. Click the 'Body' tab and examine the JSON that we will send to BIG-IP to provide credentials and the authentication provider:



4. Modify the JSON body and add the required credentials (admin/admin). Then click the 'Send' button.
5. Examine the response status code. If authentication succeeded and a token was generated the response will have a 200 OK status code. If the status code is 401 then check your credentials:

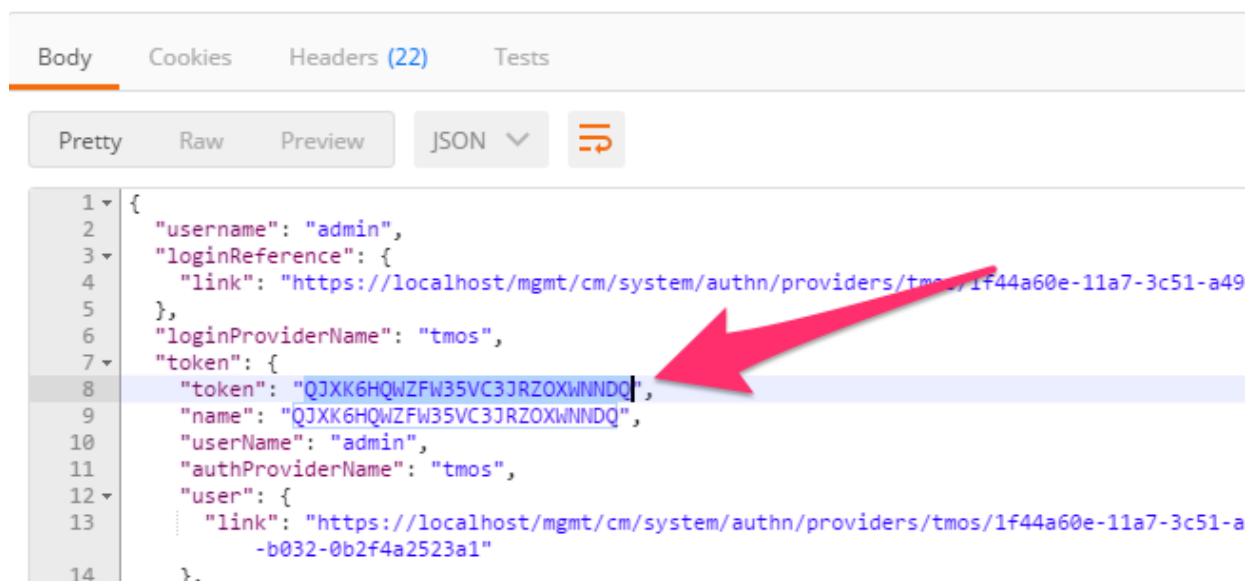
Successful:




Unsuccessful:



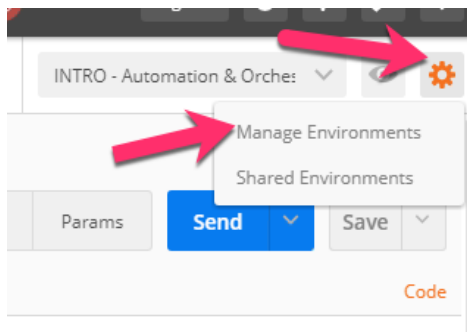
- Once you receive a 200 OK status code examine the response body. The various attributes show the parameters assigned to the particular token. Find the 'token' attribute and copy it into your clipboard (Ctrl+c) for use in the next step:




- Click the 'Step 3: Verify Authentication Works' item in the Lab 1.2 Postman collection. Click the 'Headers' tab and paste the token value copied above as the VALUE for the X-F5-Auth-Token header. This header is required to be sent on all requests when using token based authentication.



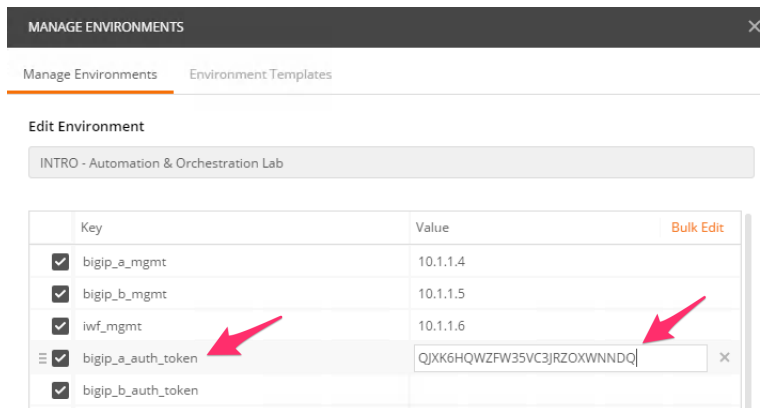
- Click the 'Send' button. If your request is successful you should see a '200 OK' status and a listing of the `ltm` Organizing Collection.
- We will now update your Postman environment to use this auth token for the remainder of the lab. Click the Environment menu in the top right of the Postman window and click 'Manage Environments':



10. Click the 'INTRO – Automation & Orchestration Lab' item:

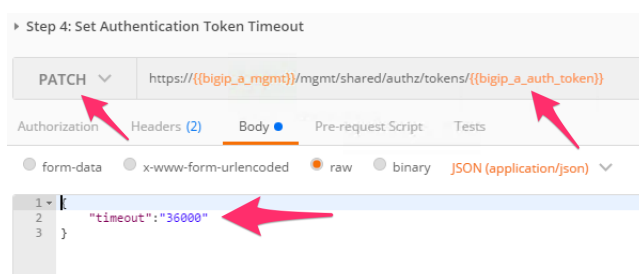


11. Update the value for `bigip_a_auth_token` by Pasting (Ctrl-v) in your auth token:



12. Click the 'Update' button and then close the 'Manage Environments' window. Your subsequent requests will now automatically include the token.

13. Click the 'Step 4: Set Authentication Token Timeout' item in the Lab 1.2 Postman collection. This request will PATCH your token Resource (check the URI) and update the timeout attribute so we can complete the lab easily. Examine the request type and JSON Body and then click the 'Send' button. Verify that the timeout has been changed to '36000' in the response:




Task 4 – Get a pool ‘example’ Template

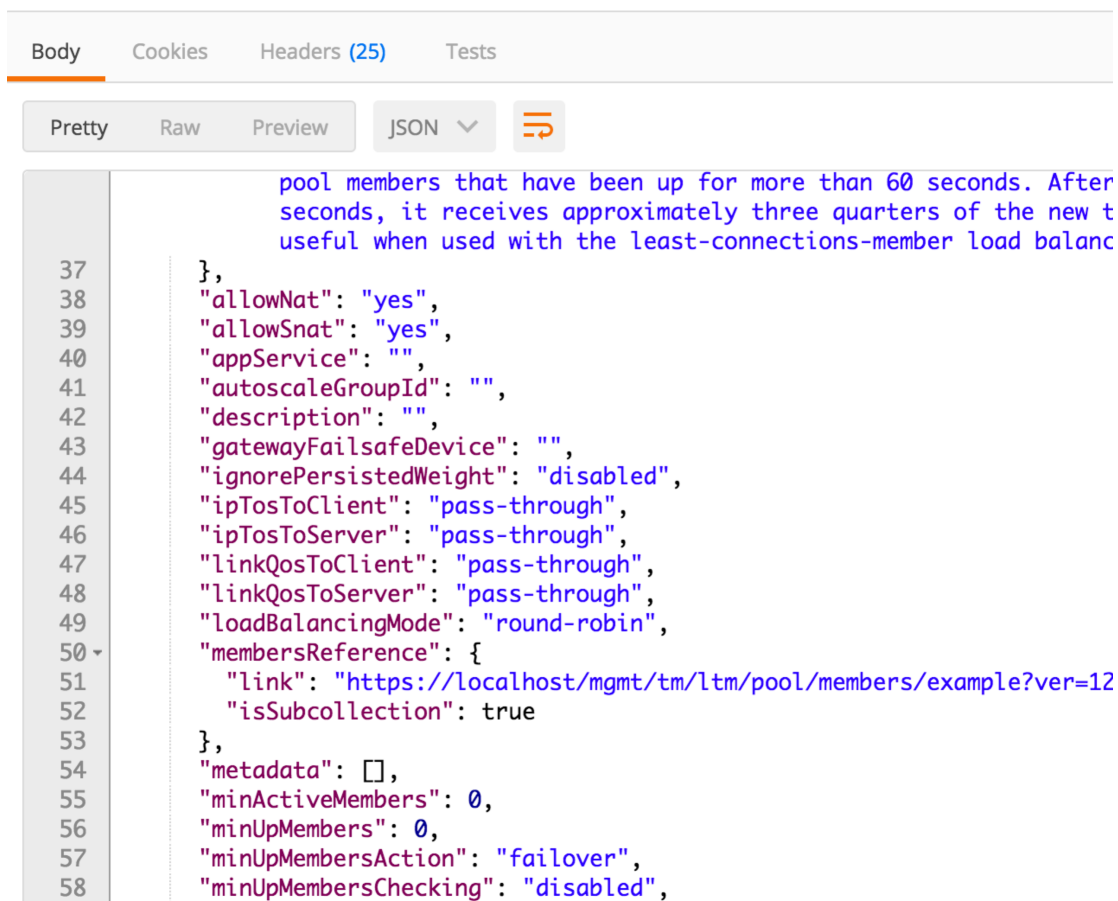
In order to assist with REST API interactions you can request a template of the various attributes of a Resource type in a Collection. This template can then be used as the body of a POST, PUT or PATCH request as needed.

Perform the following steps:

1. Click the ‘Step 5: Get ‘example’ of a Pool Resource’ item in the Lab 1.2 Postman collection
2. Examine the URI. Notice the addition of example at the end of the collection name:



3. Click ‘Send’ and examine the FULL response. You will see descriptions and then all the attributes for the Pool resource type. The response also shows the default values for the attributes if applicable:



5.2 Module 2: F5 Application Visibility and Reporting

In this module we will explore how to configure and use F5's Application Visibility and Reporting to provide application reporting. Analytics (also called Application Visibility and Reporting) is a module on the BIG-IP® system that you can use to analyze the performance of services and applications. It provides detailed

metrics such as transactions per second, server and client latency, request and response throughput, and sessions. You can view metrics for applications, virtual servers, pool members, URLs, specific countries, and additional detailed statistics about application traffic running through the BIG-IP system.

The labs in the module will focus on the high level features of AVR. These will include Analytics profiles, configuration and navigation of the AVR reports and information generated.

The BIG-IP in the lab is preconfigured with DNS / PEM / and AFM provisioned and configured. Please explore the current f5 config to familiarise yourself with this lab.

Confirm the following main config items to verify your BIG-IP lab is on working order:

1. Checked Provisioned Modules reflects the below image (DNS / PEM / AFM and AVR).

Setup Utility >> Resource Provisioning

Modified Resource Allocation (prior to redistribution)

Module	Provisioning	License Status
Management (MGMT)	Small	N/A
Carrier Grade NAT (CGNAT)	Enabled	Licensed
Local Traffic (LTM)	None	Licensed
Application Security (ASM)	None	Unlicensed
Fraud Protection Service (FPS)	None	N/A
Global Traffic (DNS)	Nominal	Licensed
Link Controller (LC)	None	Unlicensed
Access Policy (APM)	None	Limited mode available without a license
Application Visibility and Reporting (AVR)	Nominal	Licensed
Policy Enforcement (PEM)	Nominal	Licensed
Advanced Firewall (AFM)	Nominal	Licensed
Application Acceleration Manager (AAM)	None	Unlicensed
Secure Web Gateway (SWG)	None	Unlicensed
iRules Language Extensions (iRulesLX)	None	Licensed
URLDB Minimal (URLDB)	None	Unlicensed
DDOS Protection (DOS)	None	Unlicensed

Back Revert Next...

1. Check VLAN setup. Make sure Interval VLAN is set to source (SP DAG).

Network » VLANs : VLAN List

⚙️ **VLAN List** VLAN Groups

✱

☒ ▲ Name

☐ Control

☐ External

☐ Internal

Configuration: **Advanced** ⬆

Source Check	<input type="checkbox"/>
MTU	<input type="text" value="1500"/>
MAC Address	<input type="text" value="02:e2:db:5f:18:98"/>
Fail-safe	<input type="checkbox"/>
Auto Last Hop	<input type="text" value="Default"/>
CMP Hash	<input type="text" value="Source Address"/>
DAG Tunnel	<input type="text" value="Outer"/>
DAG Round Robin	<input type="checkbox"/>
Hardware SYN Cookie	<input type="checkbox"/>

1. Verify SELF-IP's and routes are present.

Network » Self IPs

⚙️ **Self IP List**

✱

☑ Name	Application	IP Address	Netmask	VLAN / Tunnel
<input type="checkbox"/> Control		10.1.30.5	255.255.255.0	Control
<input type="checkbox"/> External		10.1.20.5	255.255.255.0	External
<input type="checkbox"/> Internal		10.1.10.5	255.255.255.0	Internal

Network » Routes

⚙️ **Route List**

☑ Name

Application	Destination	Netmask	Route Domain	Resource Type	Resource
	Default IPv4		Partition Default Route Domain	Gateway	10.1.20.1

1. Check that PeM Data plane is setup, you should see four PEM data plane VS as below.

Policy Enforcement >> Data Plane Listeners										
Data Plane										
Data Plane Virtual Server Groups										
Add Group										
<input checked="" type="checkbox"/>	Name	Type	Protocol	Source Address	Destination	Mask	Port	VLAN	Address Translation	Profiles
<input type="checkbox"/>	/CommonUDF	PEM Profile: /CommonUDF_pem_profile								Delete Group Add Virtual
<input type="checkbox"/>	/CommonUDF_ANY_IP	Standard	All	0.0.0.0/0	0.0.0.0	0.0.0.0	0	Internal	None	
<input type="checkbox"/>	/CommonUDF_HTTP	Standard	TCP	0.0.0.0/0	0.0.0.0	0.0.0.0	80	Internal	None	/Commonhttp-mobile-optimized
<input type="checkbox"/>	/CommonUDF_L4_1	Standard	TCP	0.0.0.0/0	0.0.0.0	0.0.0.0	0	Internal	None	/Commonhttp-mobile-optimized
<input type="checkbox"/>	/CommonUDF_L4_2	Standard	UDP	0.0.0.0/0	0.0.0.0	0.0.0.0	0	Internal	None	/Commonudp_decrement_ttl
Delete										

1. Check that DNS Listener is configured.

DNS >> Delivery : Listeners : Listener List										
Listener List										
Statistics										
Search										
<input checked="" type="checkbox"/>	State	Name								
<input type="checkbox"/>	Enabled	SP_DNS								
Enable Disable Delete...										
			Destination	Protocol	Partition / Path					
			10.1.10.6	UDP	Common					

Note: Explore the rest of the configuration. Please look at the DNS setup (cache / monitor) and AFM CGNAT (NAPT) configurations.

5.2.1 Lab 2.1: Configure AVR Profiles

An Analytics profile is a set of definitions that determines the circumstances under which the system gathers, logs, notifies, and graphically displays information regarding traffic to an application or service. The Analytics module requires that you select an Analytics profile for each application you want to monitor. You associate the Analytics profile with one or more virtual servers used by the application / service.


In the Analytics profile, you customize:


- What statistics to collect
- Where to collect data (locally, remotely, or both)
- Whether to capture the traffic itself
- Whether to send notifications.

Task 1 - Import the Postman Collection & Environment

In this task you will Import a Postman Collection & Environment for this lab. Perform the following steps to complete this task:

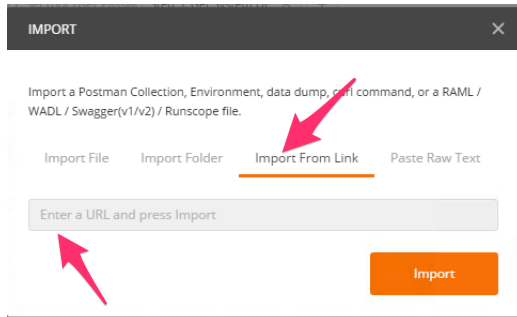


1. Open the Postman tool by clicking the  icon of the desktop of your Linux Jumphost (Postman should be open from previous Lab)
2. Click the 'Import' button in the top left of the Postman window

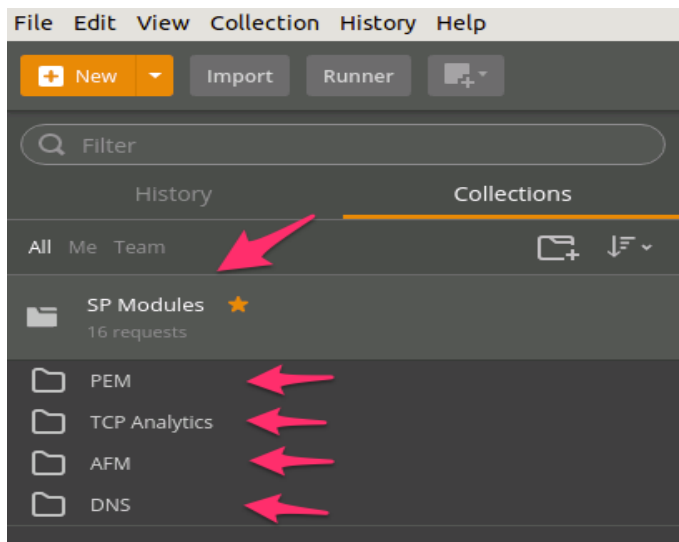


1. Click the 'Import from Link' tab. Paste the following URL into the text box and click 'Import'

`https://raw.githubusercontent.com/jarrodlucia/bigip_elk_server/develop/postman_collections/SPModules.postman_collection.json`




1. You should now see a collection named 'SP Modules' in your Postman Collections sidebar:



2. Import the Environment file by clicking 'Import' -> 'Import from Link' and pasting the following URL and clicking 'Import':

`https://raw.githubusercontent.com/jarrodlucia/bigip_elk_server/develop/postman_collections/F5SPDevOps.postman_environment.json`

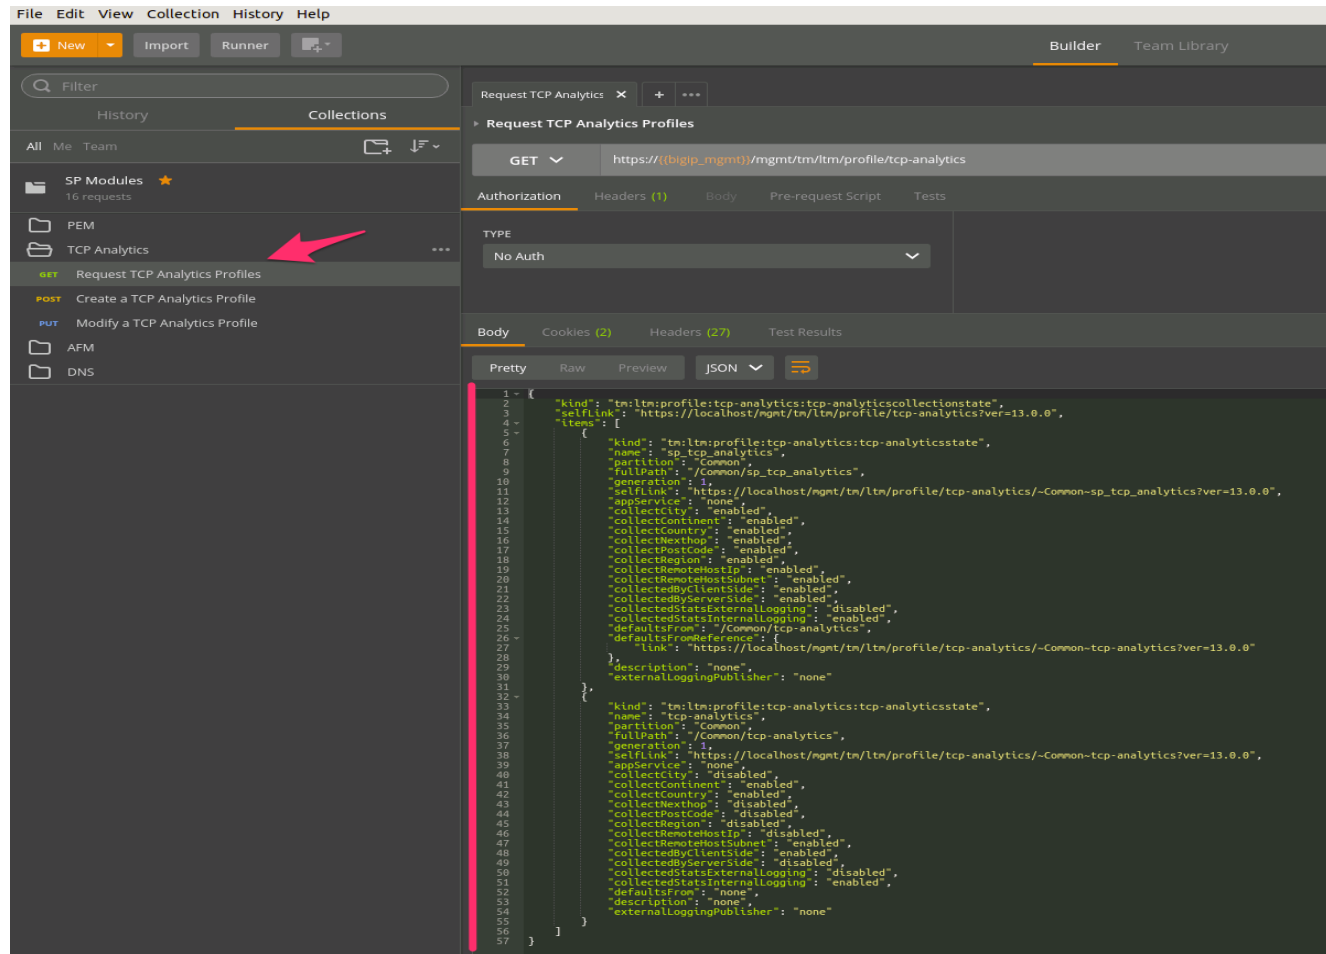


Task 2 – Configure TCP Analytics

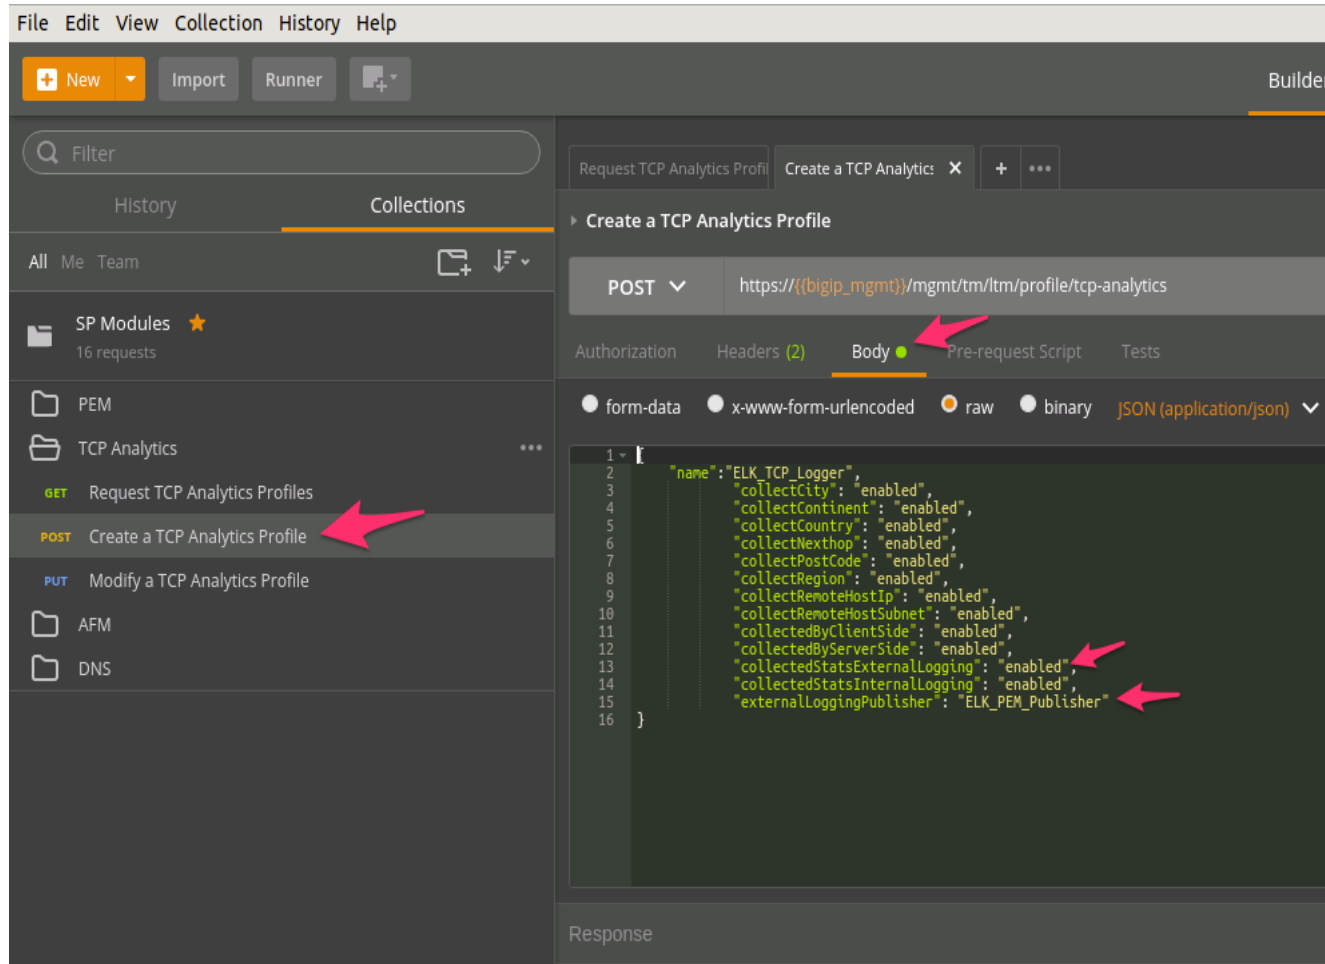
In this task we will query and configure TCP AVR profile. This will be done using REST API (explored in previous Lab)

Perform the following steps to complete this task:

1. Click the 'TCP Analytics' item in the SP Module Postman Collection
2. Notice that we are sending a GET request to the `/mgmt/tm/ltn/profile/tcp-analytics` endpoint. Check the body returned and observe the default values.



3. Click on the 'Create TCP Analytics Profile', check the body message for ELK_PEM_Publisher (We will use the PEM index in ELK for logging TCP Optimisation)



4. Verify in BIG-IP TMUI that the new profile was created.

Local Traffic » Profiles : Analytics : TCP Analytics » **ELK_TCP_Logger**


⚙️ Properties

General Configuration

Profile Name	ELK_TCP_Logger
Partition / Path	Common
Parent Profile	tcp-analytics
Description	
Statistics Logging Type	<input checked="" type="checkbox"/> Internal <input checked="" type="checkbox"/> External
Remote Publisher	ELK_PEM_Publisher
Statistics Collection	<input checked="" type="checkbox"/> Client side <input checked="" type="checkbox"/> Server side Note: If unchecked, collection is done through iRule selection only.

Associated Virtual Servers

<input checked="" type="checkbox"/>	Name	Destination

Virtual Servers  Add... Delete

Statistics Gathering Configuration

Collected Entities	<input checked="" type="checkbox"/> Virtual Server <input checked="" type="checkbox"/> Remote Host IP Address <input checked="" type="checkbox"/> Remote Host Subnet <input checked="" type="checkbox"/> Next Hop Ethernet Address <input checked="" type="checkbox"/> Continent <input checked="" type="checkbox"/> Country <input checked="" type="checkbox"/> Region <input checked="" type="checkbox"/> City <input checked="" type="checkbox"/> Postcode
--------------------	---

Cancel Update **Note:** Changes you make might take up to 10 minutes to be reflected in the charts.

5. Add in the VS manually (This is not available in REST API currently)

Select Virtual Server

<input type="checkbox"/>	Name	Destination	Service Port	Partition / Path
<input type="checkbox"/>	UDF_HTTP	any	80	Common
<input type="checkbox"/>	UDF_L4_1	any	0	Common

Total Entries: 2



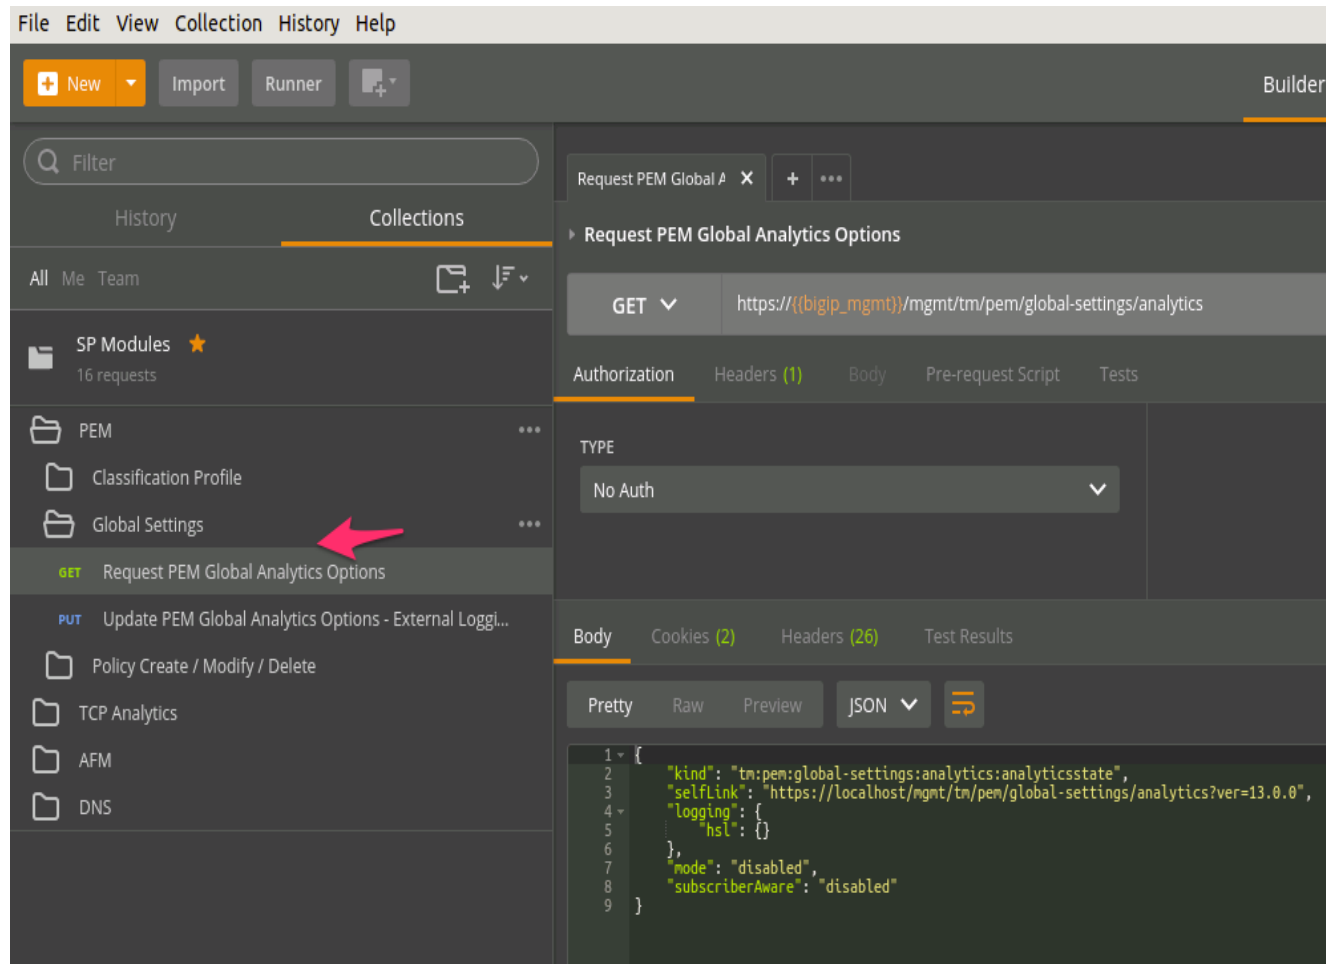
Add both VS to profile

Task 3 – Configure PEM Analytics

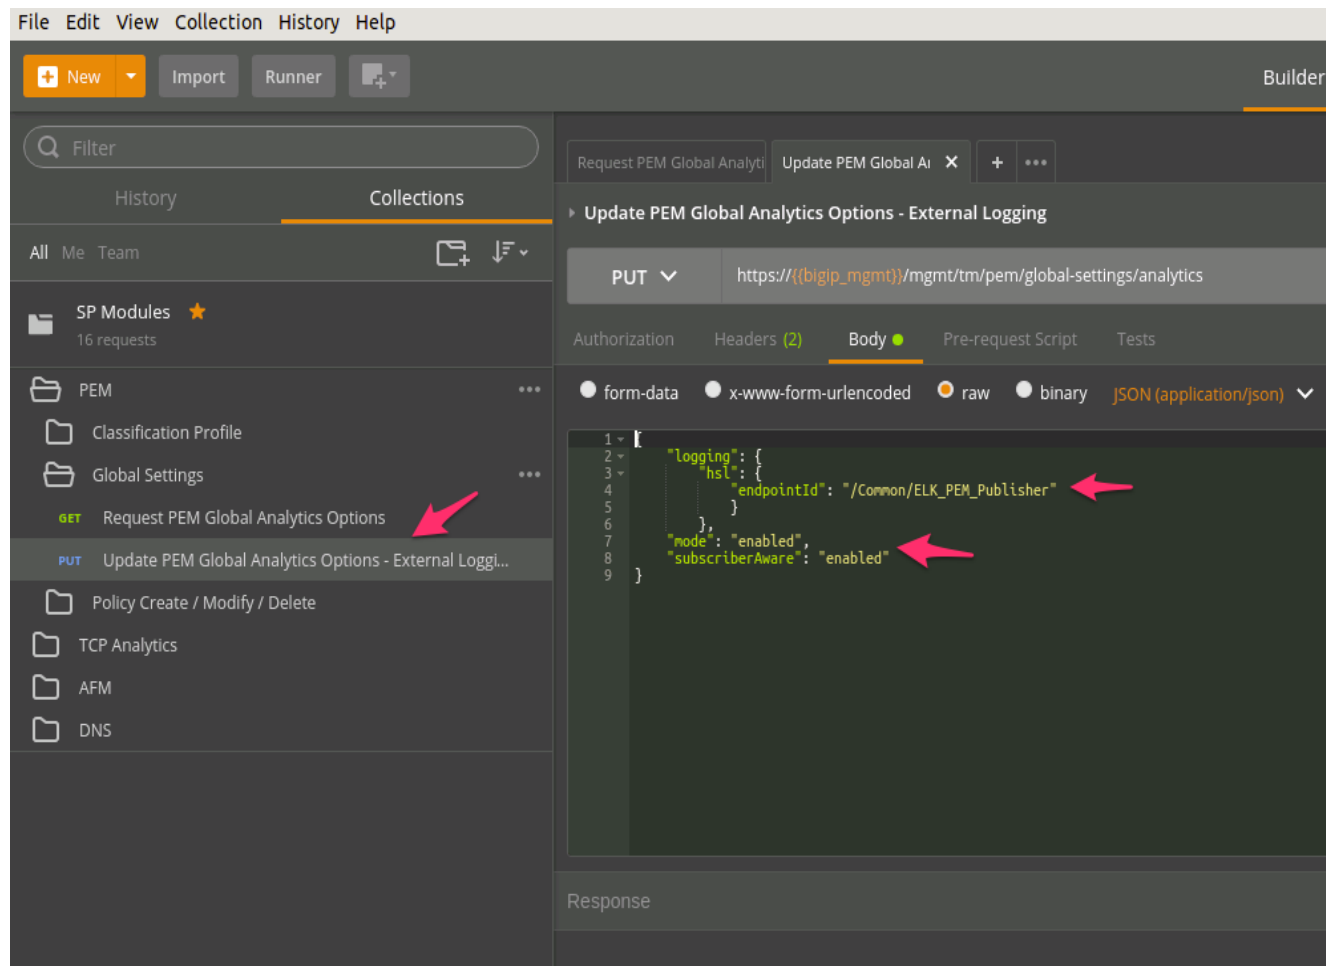
In this task we will query and configure PEM AVR profile. This will be done using REST API (explored in previous Lab)

Perform the following steps to complete this task:

1. Click the 'PEM' item in the SP Module Postman Collection
2. Notice there are two sections we must update Global and Classification. We will do Global first, click on 'Request PEM Global Analytics Options' we are sending a GET request to the `/mgmt/tm/pem/global-settings/analytics` endpoint. Check the body returned and observe the default values.



3. Click on the 'Update PEM Global Analytics Options - External Logging' , check the body message for ELK_PEM_Publisher.



4. Verify in BIG-IP TMUI that the new updates were changed in PEM global options.
5. Click on 'Request PEM Classification Profile' we are sending a GET request to the `/mgmt/tm/ltm/profile/classification/classification_pem` endpoint. Check the body returned and observe the default values.

The screenshot shows the Postman application interface. On the left sidebar, under the 'Collections' tab, the 'Classification Profile' folder is highlighted with a red arrow. Below it, the 'Request PEM Classification Profile' endpoint is selected. The main panel shows the details for this request, including the URL, headers, and the JSON body.

Request PEM Classification Profile

GET [https://\(bigip_mgmt\)/mgmt/tm/ltn/profile/classification/classification_pem](https://(bigip_mgmt)/mgmt/tm/ltn/profile/classification/classification_pem)

Authorization Headers (1) Body Pre-request Script Tests

TYPE

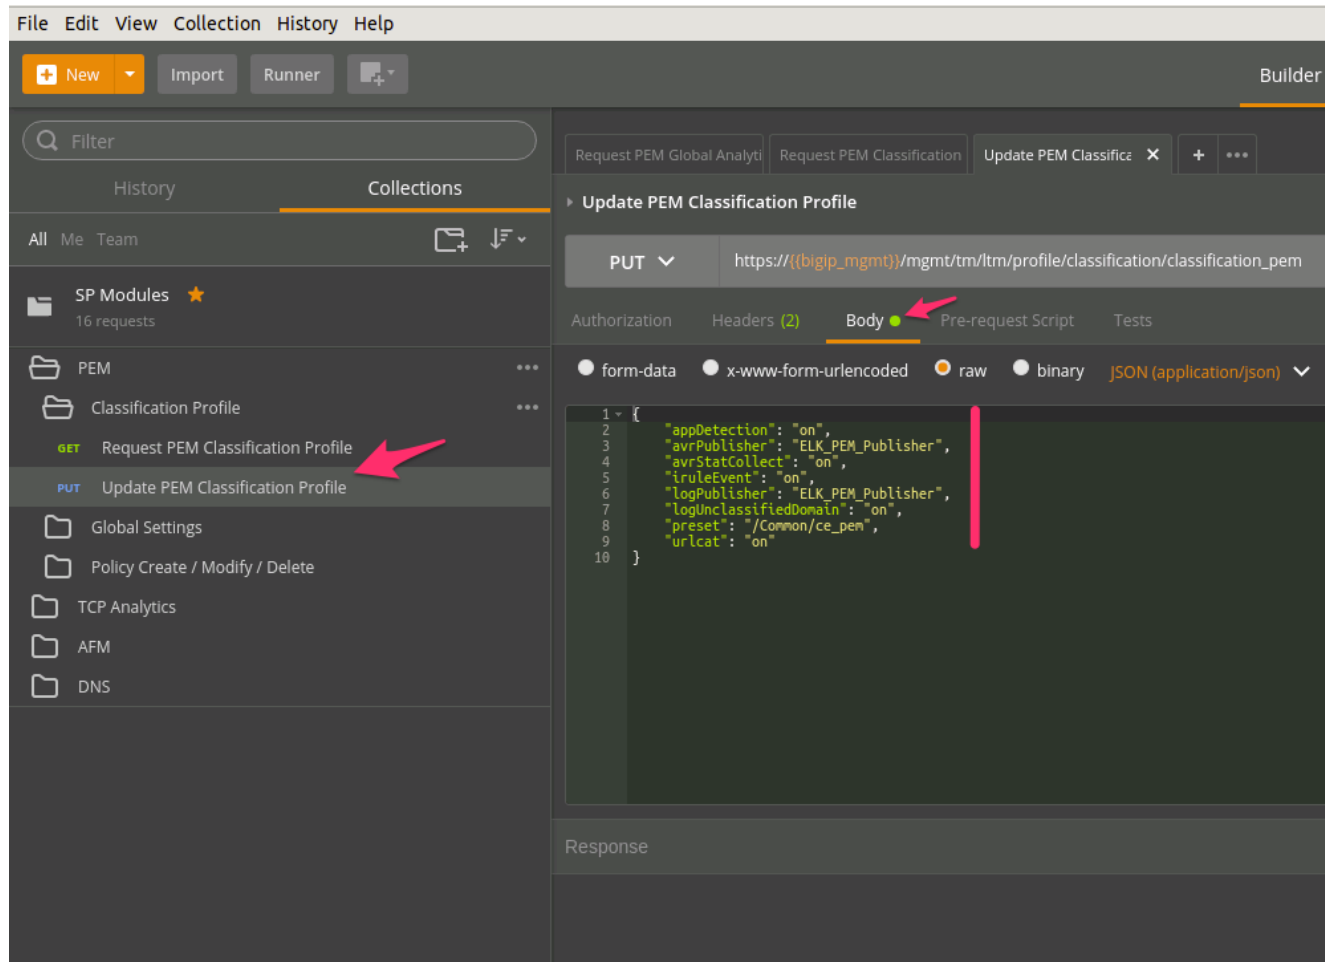
No Auth

Body Cookies (2) Headers (27) Test Results

Pretty Raw Preview JSON

```
1 {
2   "kind": "tm:ltn:profile:classification:classificationstate",
3   "name": "classification_pem",
4   "fullPath": "classification_pem",
5   "generation": 1,
6   "selfLink": "https://localhost/mgmt/tm/ltn/profile/classification/classification_pem?ver=13.0.0",
7   "appDetection": "on",
8   "avrPublisher": "none",
9   "avrStatCollect": "off",
10  "defaultsFrom": "/Common/classification",
11  "defaultsFromReference": {
12    "link": "https://localhost/mgmt/tm/ltn/profile/classification/~Common-classification?ver=13.0.0"
13  },
14  "description": "none",
15  "iruleEvent": "on",
16  "logPublisher": "none",
17  "logUnclassifiedDomain": "off",
18  "preset": "/Common/ce_pem",
19  "presetReference": {
20    "link": "https://localhost/mgmt/tm/ltn/classification/ce/~Common-ce_pem?ver=13.0.0"
21  },
22  "urlcat": "on"
23 }
```

6. Click on the 'Update PEM Classification Profile', check the body message for ELK_PEM_Publisher.



7. Verify in BIG-IP TMUI that the new updates were changed in PEM Classification.

Task 4 – Configure AFM Analytics

In this task we will query and configure AFM AVR profile and Logging. This will be done using REST API (explored in previous Lab)

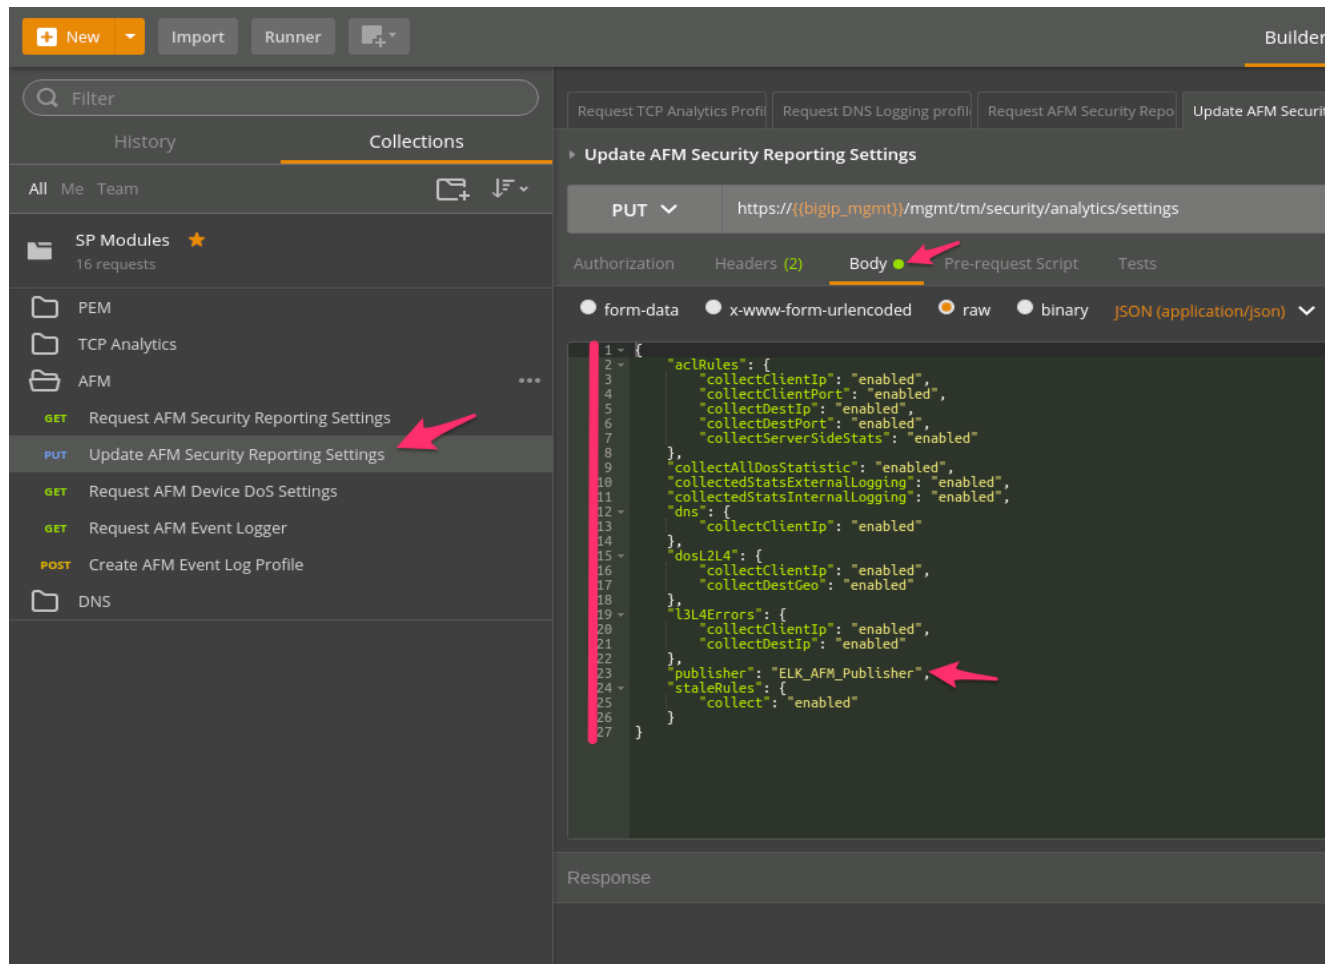
Perform the following steps to complete this task:

1. Click the 'AFM' item in the SP Module Postman Collection
2. Notice there are two sections we must update Security Reporting and Event Logging. We will do Security Reporting first, click on 'Request AFM Security Reporting Settings' we are sending a GET request to the `/mgmt/tm/security/analytics/settings` endpoint. Check the body returned and observe the default values.

The screenshot displays the Postman application interface. On the left sidebar, under the 'Collections' tab, the 'AFM' collection is expanded, and a red arrow points to the 'Request AFM Security Reporting Settings' item. The main panel shows the details of this request, including the URL 'https://{{bigip_mgmt}}/mgmt/tm/security/analytics/settings', the 'GET' method, and the 'No Auth' authorization type. The 'Body' tab is selected, showing a JSON response with the following structure:

```
1 {
2   "kind": "tm:security:analytics:settings:settingsstate",
3   "selfLink": "https://localhost/mgmt/tm/security/analytics/settings?ver=13.0.0",
4   "aclRules": {
5     "collectClientIp": "enabled",
6     "collectClientPort": "disabled",
7     "collectDestIp": "enabled",
8     "collectDestPort": "enabled",
9     "collectServerSideStats": "disabled"
10  },
11  "collectAllDosStatistic": "disabled",
12  "collectedStatsExternalLogging": "disabled",
13  "collectedStatsInternalLogging": "enabled",
14  "dns": {
15    "collectClientIp": "enabled"
16  },
17  "dosL2L4": {
18    "collectClientIp": "enabled",
19    "collectDestGeo": "enabled"
20  },
21  "l3L4Errors": {
22    "collectClientIp": "enabled",
23    "collectDestIp": "enabled"
24  },
25  "staleRules": {
26    "collect": "disabled"
27  }
28 }
```

3. Click on the 'Update AFM Security Reporting Settings' , check the body message for ELK_AFM_Publisher.



4. Verify in BIG-IP TMUI that the new updates were changed in AFM Report Settings.

Note: Request AFM Device DoS Settings - Can be used to report on settings currently set, however REST API cannot be used to update these settings at this time.

1. Click on 'Request AFM Event Logger' we are sending a GET request to the `/mgmt/tm/security/log/profile/` endpoint. Check the body returned and observe the default values.

The screenshot shows the Postman interface with the 'Request AFM Event Logger' collection selected. The left sidebar lists various collections, and a red arrow points to the 'Request AFM Event Logger' item. The main panel displays the details of this request, including the URL 'https://(t(bigip_mgmt))/mgmt/tm/security/log/profile/', the 'No Auth' authorization type, and the JSON body of the request. The JSON body is a complex object representing the AFM Event Log Profile configuration.

```

1 {
2   "kind": "tm:security:log:profile:profilecollectionstate",
3   "selflink": "https://localhost/mgmt/tm/security/log/profile?ver=13.0.0",
4   "items": [
5     {
6       "kind": "tm:security:log:profile:profilestate",
7       "name": "Log all requests",
8       "partition": "Common",
9       "fullPath": "/Common/Log all requests",
10      "generation": 1,
11      "selflink": "https://localhost/mgmt/tm/security/log/profile/~Common-Log%20all%20requests?ver=13.0.0",
12      "description": "Default logging profile for all requests",
13      "intelligence": {
14        "aggregateRate": 4294967295,
15        "logRtBb": "disabled",
16        "logScrubber": "disabled",
17        "logShun": "disabled",
18        "logTranslationFields": "disabled"
19      },
20      "net": {
21        "endInboundSession": "disabled",
22        "endOutboundSession": {
23          "action": "disabled"
24        },
25        "errors": "disabled",
26        "isLegacyMode": "disabled",
27        "quotaExceeded": "disabled",
28        "rateLimit": {
29          "aggregateRate": 4294967295,
30          "endInboundSession": 4294967295,
31          "endOutboundSession": 4294967295,
32          "errors": 4294967295,
33          "quotaExceeded": 4294967295,
34          "startInboundSession": 4294967295,
35          "startOutboundSession": 4294967295
36        },
37        "startInboundSession": "disabled",
38        "startOutboundSession": {
39          "action": "disabled"
40        }
41      },
42      "portMisuse": {
43        "aggregateRate": 4294967295
44      },
45      "trafficStatistics": {
46        "activeFlows": "disabled",
47        "missedFlows": "disabled",
48        "reapedFlows": "disabled",
49        "syncCookies": "disabled",
50        "syncCookiesWhitelist": "disabled"
51      },
52      "networkReference": {
53        "link": "https://localhost/mgmt/tm/security/log/profile/~Common-Log%20all%20requests/network?ver=13.0.0",
54        "isSubcollection": true
55      },
56      "protocolDnsReference": {
57        "link": "https://localhost/mgmt/tm/security/log/profile/~Common-Log%20all%20requests/protocol-dns?ver=13.0.0",
58        "isSubcollection": true
59      }
60    ]
61  }

```

2. Click on the 'Create AFM Event Log Profile' , check the body message for ELK_AFM_Publisher.

The screenshot shows the Postman application with the 'Create AFM Event Log Profile' request selected in the 'Collections' list. The request is a POST to `https://(bigip_mgmt)/mgmt/tm/security/log/profile/`. The 'Body' tab is active, showing a JSON payload. Red arrows point to the following fields in the JSON:

- `"name": "ELK_AFM_Logger"` (line 2)
- `"dosNetworkPublisher": "/Common/ELK_AFM_Publisher"` (line 3)
- `"nat": {` (line 12)
- `"portMisuse": {` (line 35)
- `"protocolTransfer": {` (line 49)

3. Additional Steps are required for AFM as not all REST commands can configure all sections at this time. Go to TMUI on BIG-IP and navigate to Security / Event Logs / Logging Profiles. Change Publishers and tick events to log.

f5 Firewall: Consistent
ONLINE (ACTIVE)
Standalone

Main Help About

Security » Event Logs : Logging Profiles » Edit Logging Profile

Statistics
iApps
DNS
Local Traffic
Traffic Intelligence
Acceleration
Policy Enforcement
Subscriber Management
Device Management
Security
Network
System

Overview
Protocol Security
Network Firewall
Network Address Translation
DoS Protection
Event Logs
Reporting
Options

Logging Profile Properties

Profile Name	ELK_AFM_Logger
Partition / Path	Common
Description	
Protocol Security	<input checked="" type="checkbox"/> Enabled
Network Firewall	<input checked="" type="checkbox"/> Enabled
Network Address Translation	<input checked="" type="checkbox"/> Enabled
DoS Protection	<input checked="" type="checkbox"/> Enabled

Protocol Security Network Firewall Network Address Translation DoS Protection

HTTP, FTP, and SMTP Security

Publisher	ELK_AFM_Publisher
-----------	-------------------

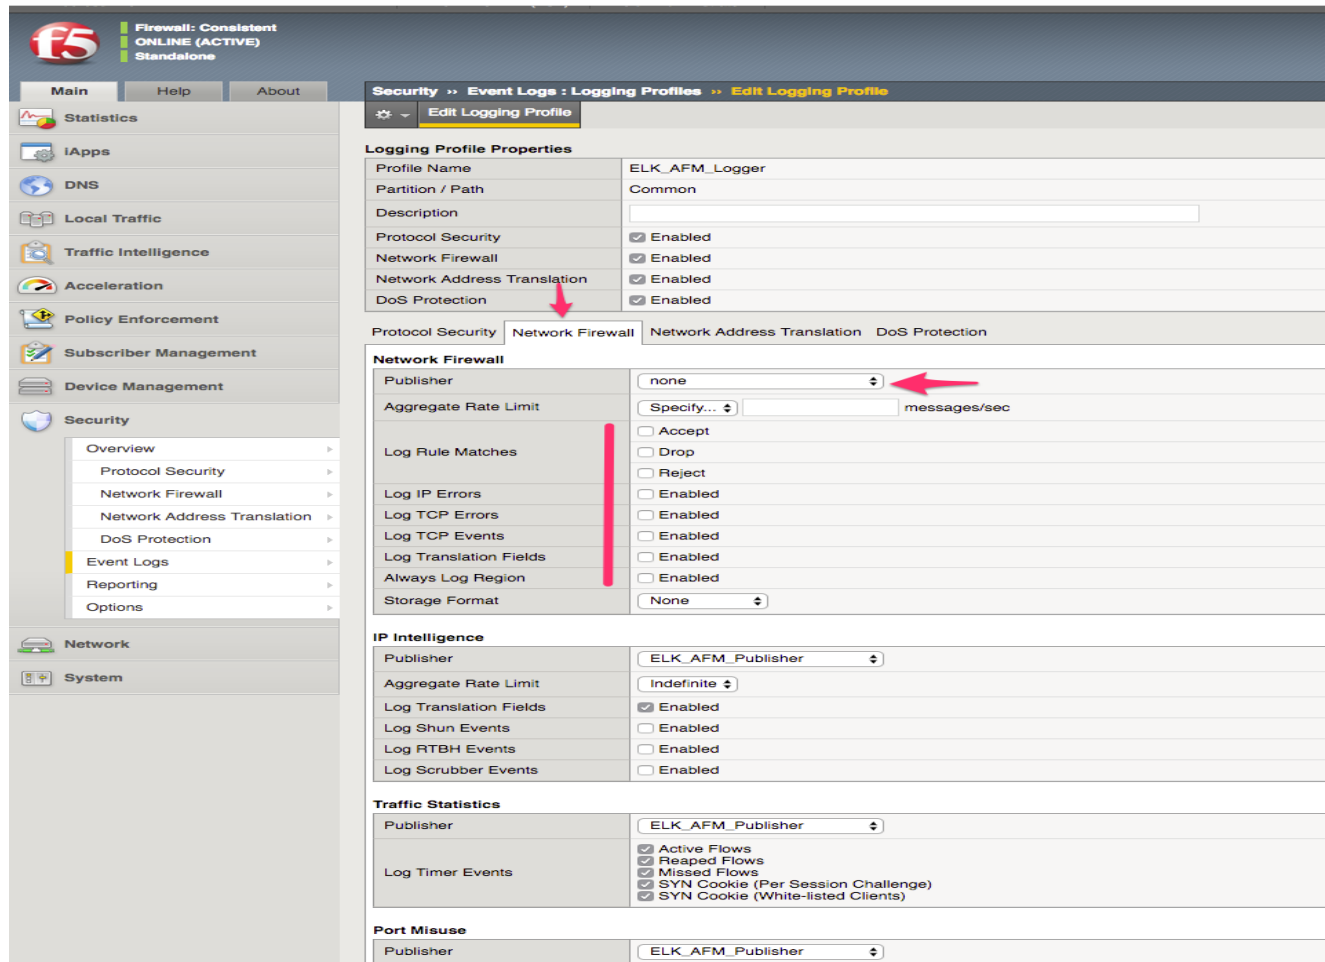
DNS Security

Publisher	ELK_AFM_Publisher
Log Dropped Requests	<input type="checkbox"/> Enabled
Log Filtered Dropped Requests	<input type="checkbox"/> Enabled
Log Malformed Requests	<input type="checkbox"/> Enabled
Log Rejected Requests	<input type="checkbox"/> Enabled
Log Malicious Requests	<input type="checkbox"/> Enabled
Storage Format	None

SIP Security

Publisher	none
Log Dropped Requests	<input type="checkbox"/> Enabled
Log Global Failures	<input type="checkbox"/> Enabled
Log Malformed Requests	<input type="checkbox"/> Enabled
Log Redirection Responses	<input type="checkbox"/> Enabled
Log Request Failures	<input type="checkbox"/> Enabled
Log Server Errors	<input type="checkbox"/> Enabled
Storage Format	None

Update Network Firewall tab and click update.

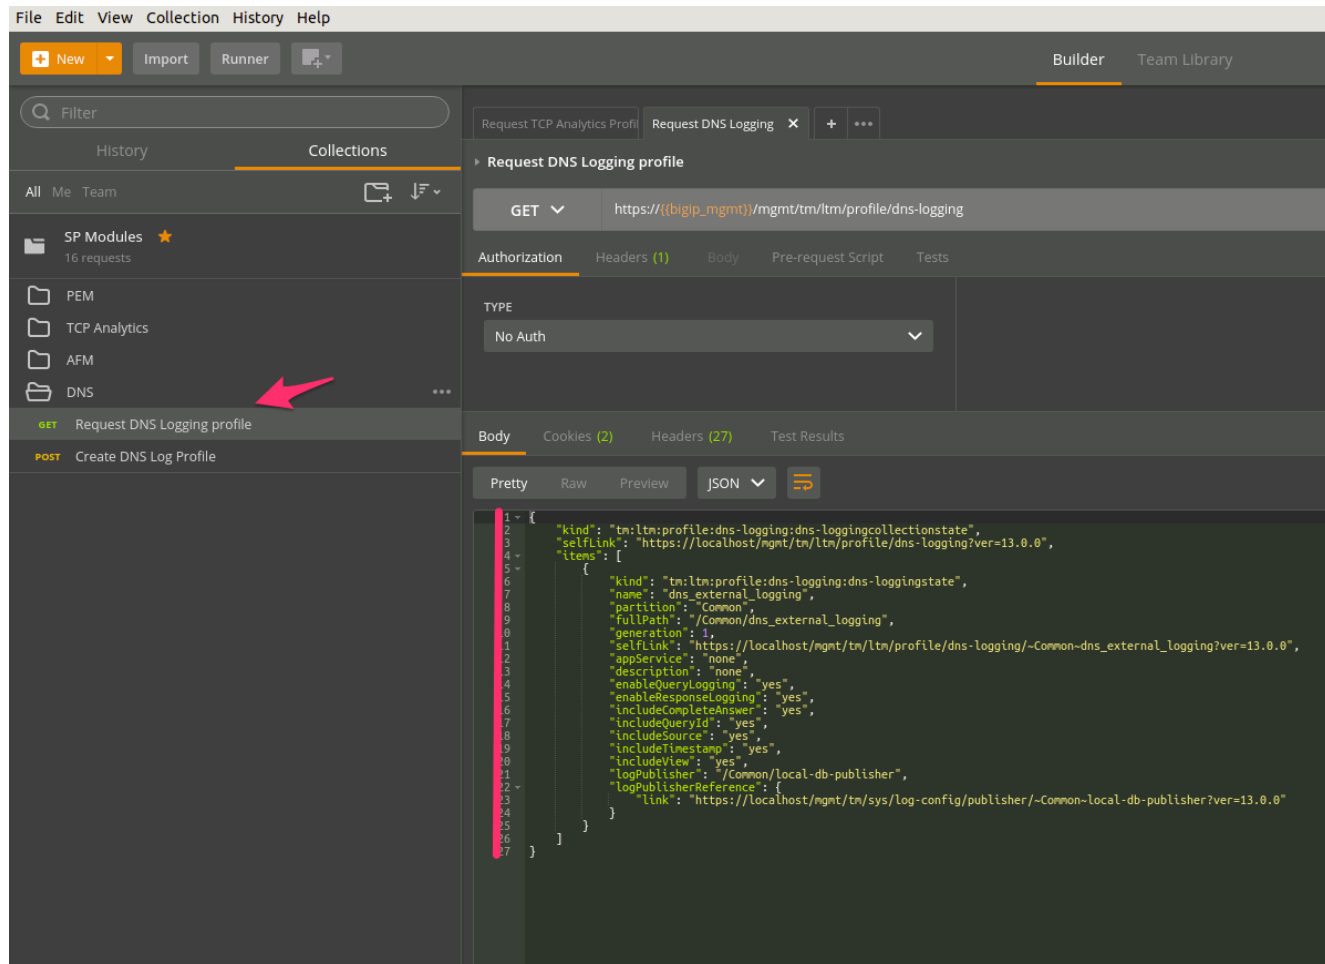


Task 5 – Configure DNS Analytics

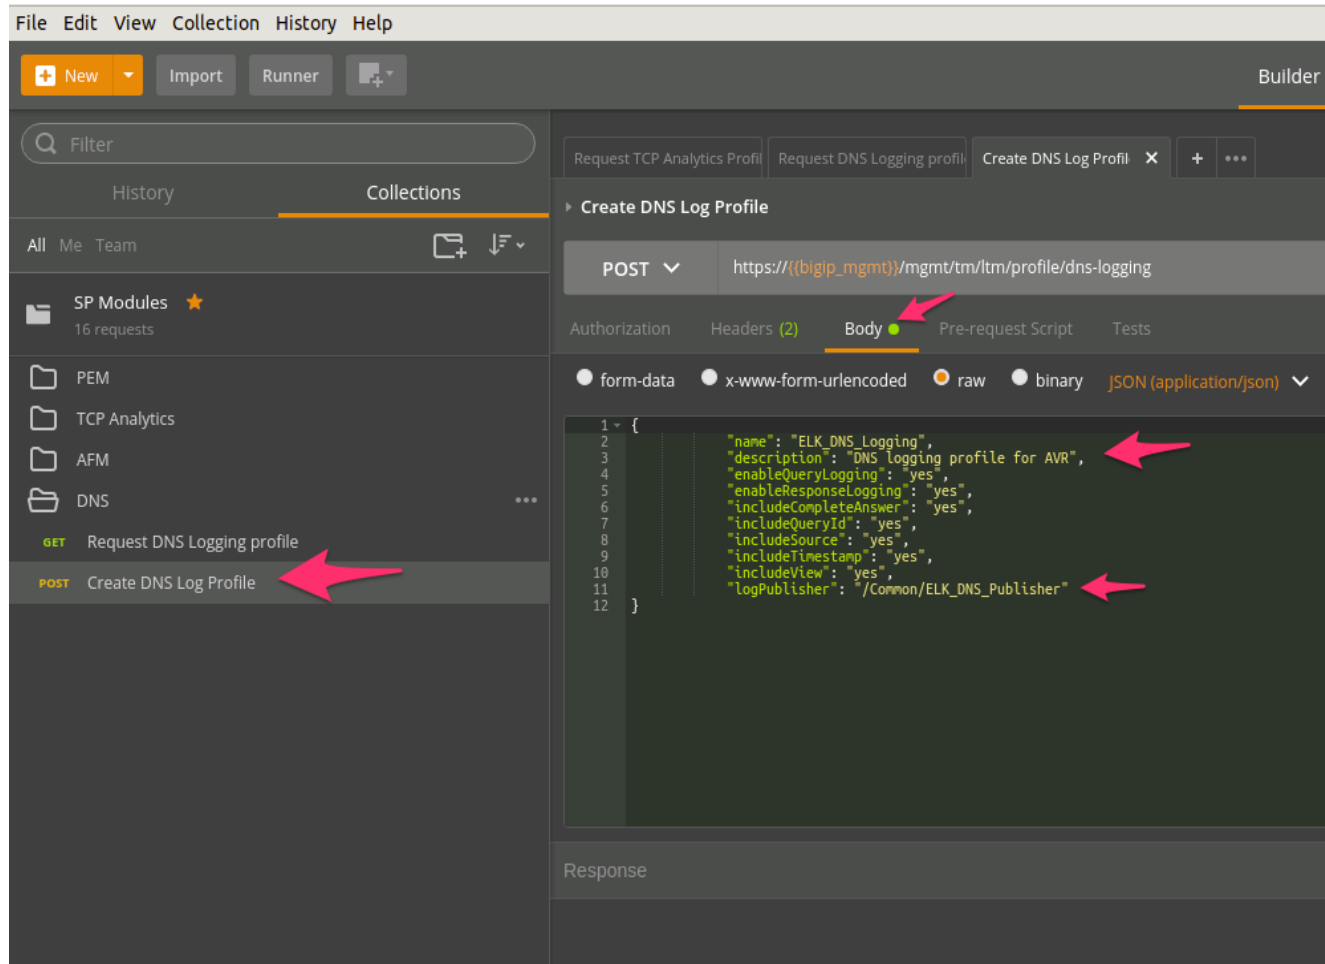
In this task we will query and configure DNS AVR profile. This will be done using REST API (explored in previous Lab)

Perform the following steps to complete this task:

1. Click the 'DNS' item in the SP Module Postman Collection
2. Notice that we are sending a GET request to the `/mgmt/tm/ltn/profile/dns-logging` endpoint. Check the body returned and observe the default values.



3. Click on the 'Create DNS Log Profile', check the body message for ELK_DNS_Publisher.



4. Verify in BIG-IP TMUI that the new profile was created.

5.2.2 Lab 2.2: Access Clients and Generate Traffic

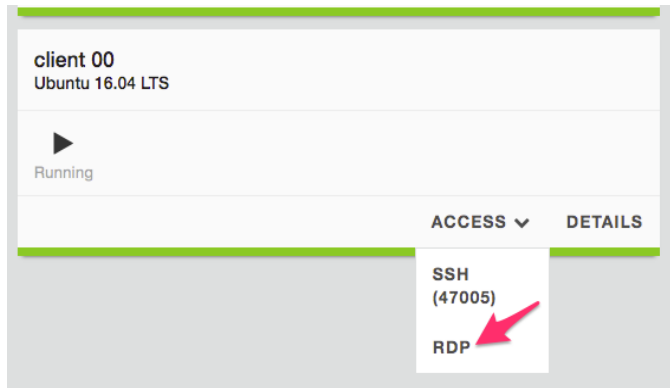
In this lab you will walk through re-configuring the Clients to USE the F5 for traffic. This will generate traffic for PEM / DNS / and AFM for AVR and logging to ELK Stack.

Task 1 - Configure Client Networking & Traffic Generation


In this task we will configure and use the Client UDF machines. These Clients are required to be reconfigured to utilise the network and DNS from the F5..

Perform the following steps to complete this task:


1. Click the on the RDP access for UDF for each client.



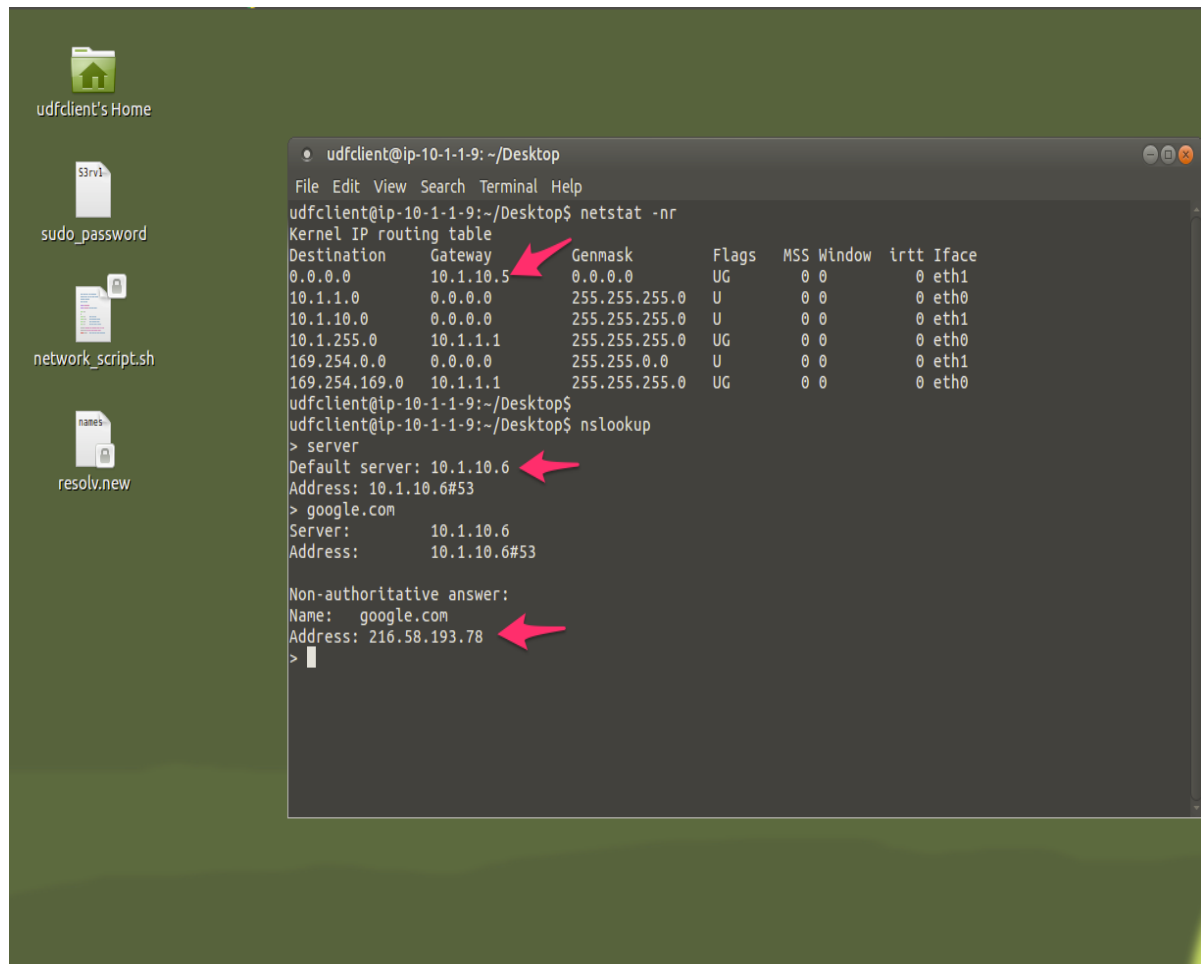
Accept warning always



2. Click on the networking script (this will prompt for Sudo password)



3. Once the script has completed, check `netstat -nr` and `nslookup` to verify you have traffic passing the F5.








4. Verify in BIG-IP TMUI that you see traffic on the F5 VS's

Statistics » Module Statistics : Local Traffic » **Virtual Servers**

Display Options

Statistics Type	Virtual Servers
Data Format	Normalized
Auto Refresh	Disabled <input type="button" value="Refresh"/>

<input checked="" type="checkbox"/>	Status	Virtual Server
<input type="checkbox"/>		SP_DNS
<input type="checkbox"/>		UDF_ANY_IP
<input type="checkbox"/>		UDF_HTTP
<input type="checkbox"/>		UDF_L4_1
<input type="checkbox"/>		UDF_L4_2

5. Apply the same fix for the other client.
6. Once both clients are fixed, generate traffic by opening applications and webpages (Leave the applications open so traffic generation continues)



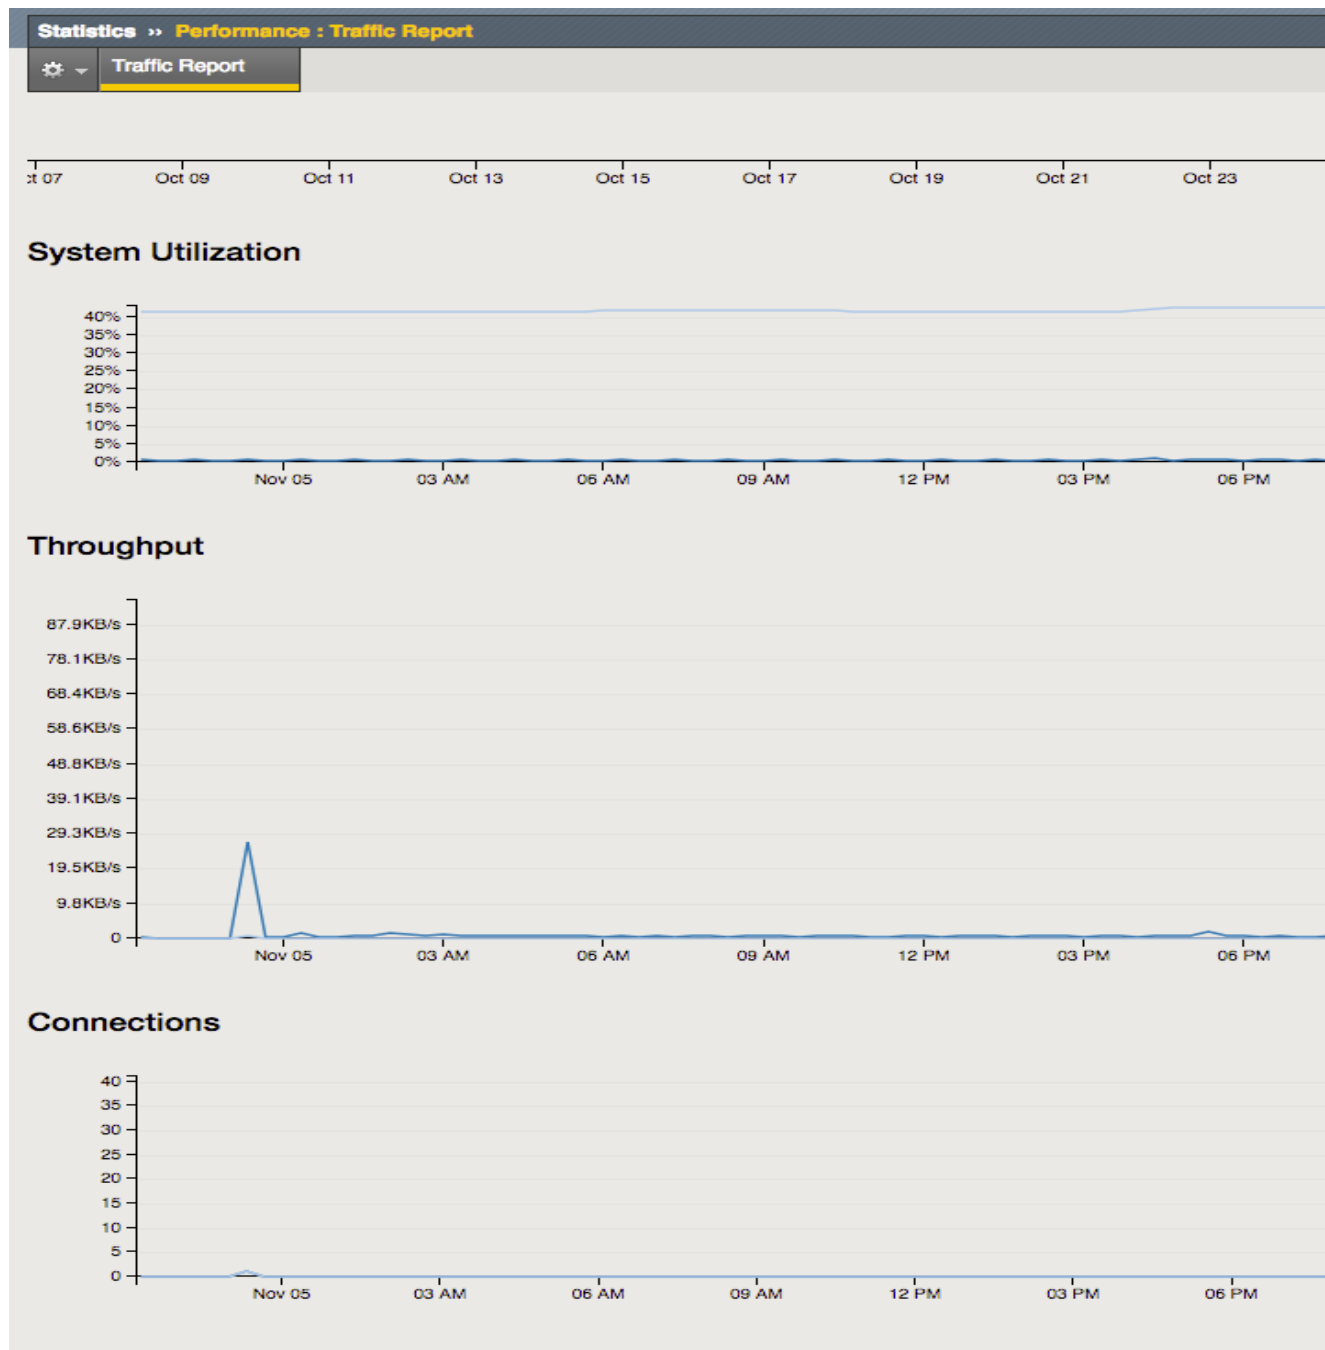
5.2.3 Lab 2.3: Navigating AVR

Navigating and viewing AVR reports.

Task 1 – BIG-IP Performance Report

Perform the following steps to complete this task:

1. Navigate to Performance Report under Statistics.

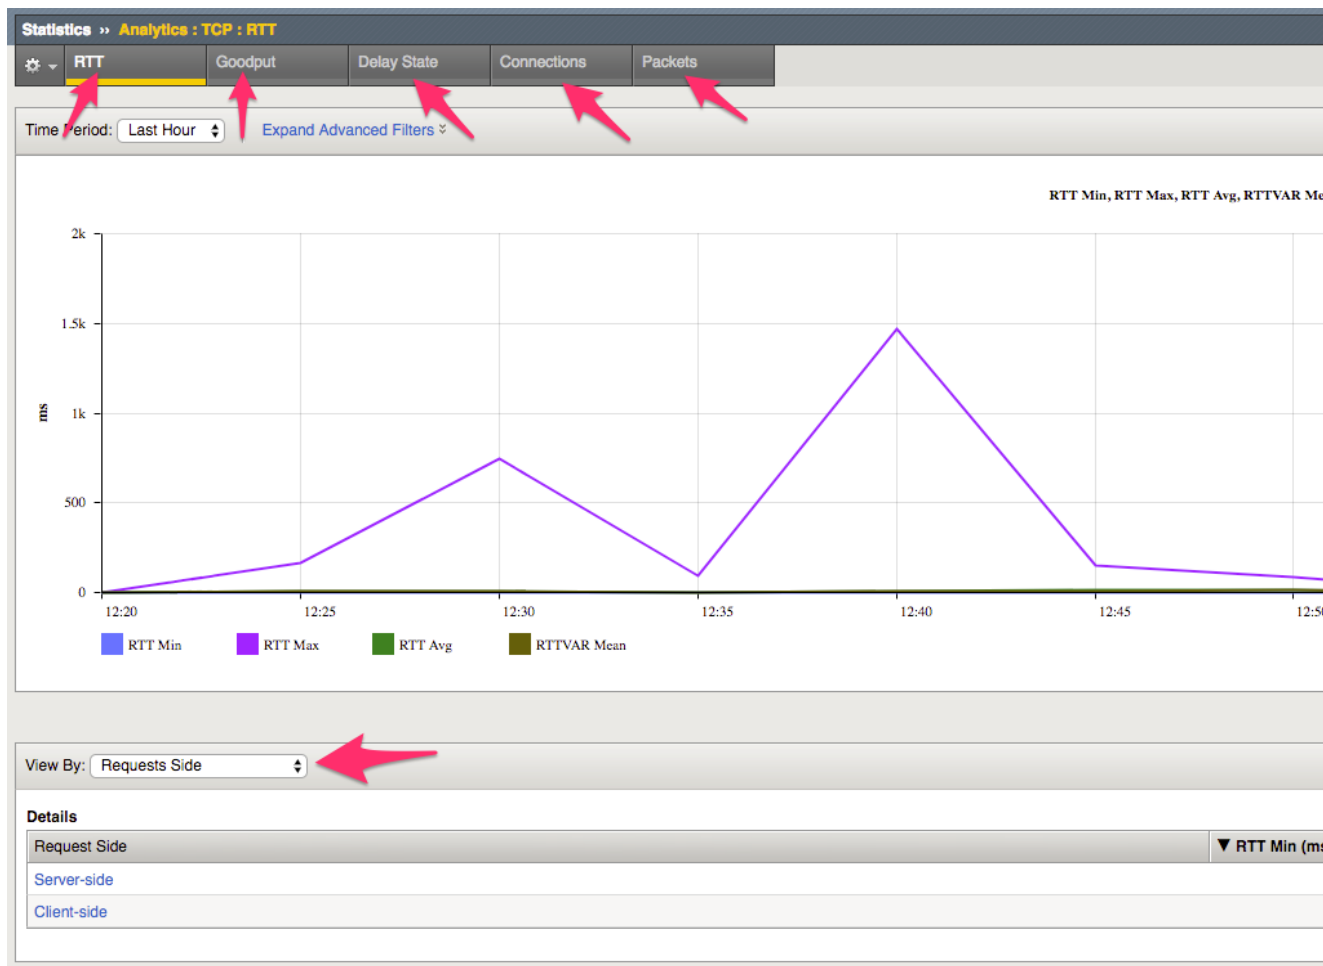


Explore the interface with the sliding bar, and tick and untick options.

Task 2 – AVR TCP Optimisation

Perform the following steps to complete this task:

1. Navigate to Analytics TCP Statistics.



2. Explore the different display options by clicking around the dashboard.


See the following link for further TCP AVR information:

https://support.f5.com/kb/en-us/products/big-ip_analytics/manuals/product/analytics-implementations-12-1-0/9.html

Task 3 – AVR Traffic Classification

Perform the following steps to complete this task:

1. Navigate to Traffic Classification Analytics.

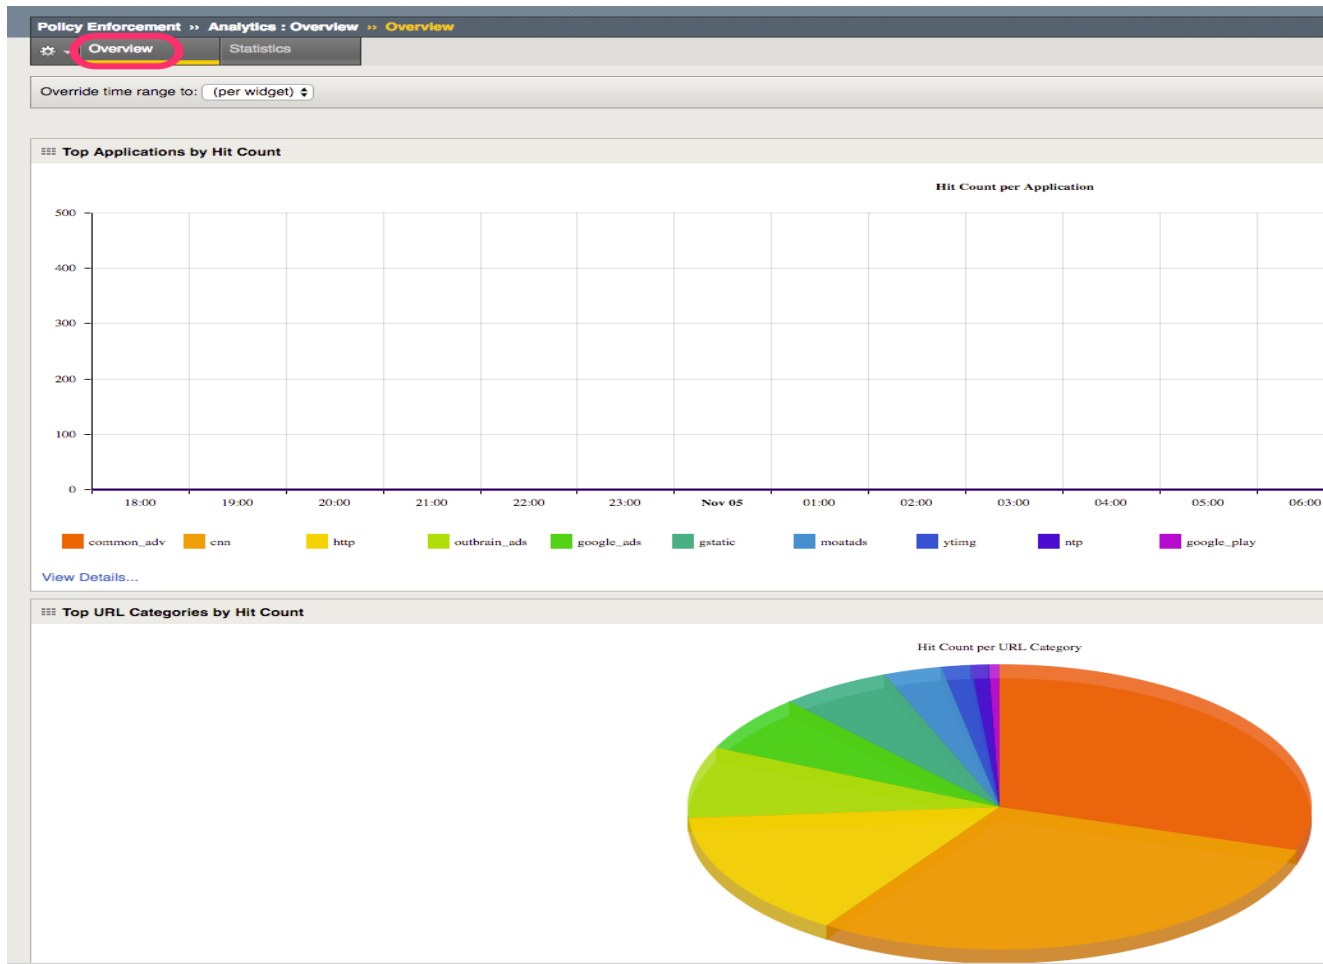


2. Explore the different display options by clicking around the dashboard.


Task 4 – PEM Analytics Report

Perform the following steps to complete this task:

1. Navigate to Policy Enforcement Analytics Overview.



2. Navigate to Policy Enforcement Analytics Statistics.



Explore the different screens and options available for display. See the following link for further AVR information:

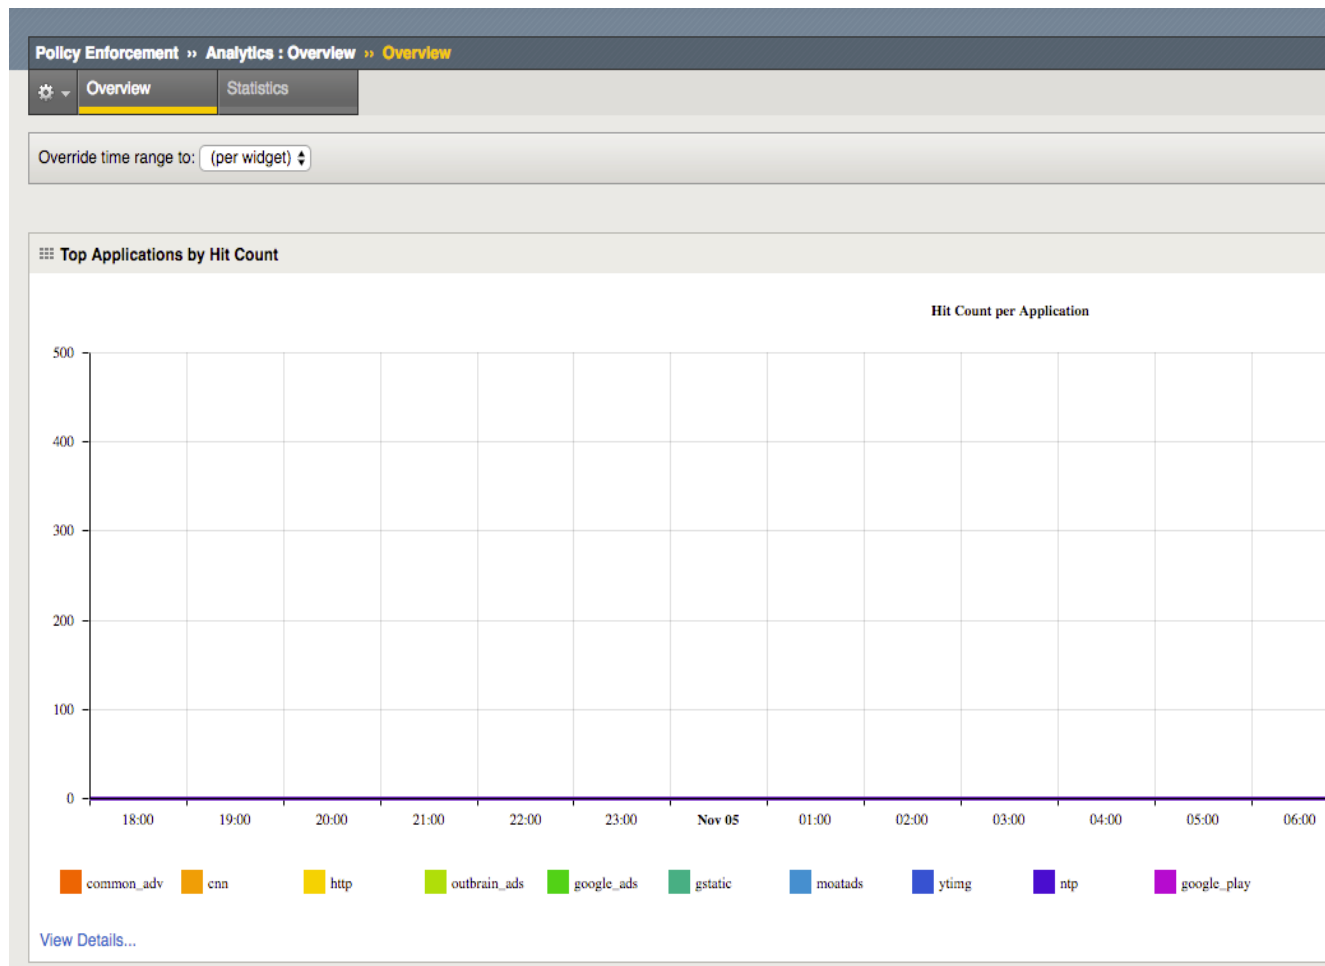
<https://support.f5.com/kb/en-us/products/big-ip-pem/manuals/product/pem-implementations-13-0-0.html>

Task 5 – Modify PEM AVR Dashboard / Export AVR Report

In this task we will modify and add widgets to default dashboard, and export an Analytics dashboard to a PDF report.

Perform the following steps to complete this task:

1. Navigate to Policy Enforcement Analytics.



2. Click on Add Widget
3. Create a New Widget of your choice.
4. Explore the options within the Dashboard widgets for display
5. Click on Export, select PDF to generate report.

Task 6 - PEM Scheduled Reports

In this task we will configure a Scheduled PEM report.

Perform the following steps to complete this task:

1. Navigate to Policy Enforcement Analytics Scheduled Reports.

Policy Enforcement » Analytics : Scheduled Reports » New Reporting Schedule...

Chart Schedule Properties

Name	<input type="text"/>						
Send To (E-Mails)	<input type="text"/> <input type="button" value="Add"/> <div style="border: 1px solid #ccc; height: 60px; margin-top: 5px;"></div> <input type="button" value="Delete"/>						
SMTP Configuration	<input type="text"/> <input type="button" value="Create..."/>						
Reporting Module	Policy Enforcement Manager <input type="button" value="v"/>						
Chart	<div> Filter <table border="1"> <tr> <td>Time Period</td> <td>Last Day <input type="button" value="v"/></td> </tr> <tr> <td>Show Results</td> <td>Top 10 <input type="button" value="v"/></td> </tr> </table> </div> <div> Chart Path <table border="1"> <tr> <td>Please select top report criteria</td> <td>Action <input type="button" value="v"/> <input style="float: right;" type="button" value="+"/></td> </tr> </table> <div style="display: flex; justify-content: space-between; margin-top: 10px;"> <div> Selected measures (up to 6): <div style="border: 1px solid #ccc; height: 60px; margin-top: 5px;"></div> <input type="button" value="↑"/> <input type="button" value="↓"/> </div> <div> Available measures: <div style="border: 1px solid #ccc; padding: 5px; margin-top: 5px;"> Hit Count Total Bytes In Total Bytes Out </div> <input type="button" value="←"/> <input type="button" value="→"/> </div> </div> </div> <div> <input checked="" type="checkbox"/> Include Overall </div>	Time Period	Last Day <input type="button" value="v"/>	Show Results	Top 10 <input type="button" value="v"/>	Please select top report criteria	Action <input type="button" value="v"/> <input style="float: right;" type="button" value="+"/>
Time Period	Last Day <input type="button" value="v"/>						
Show Results	Top 10 <input type="button" value="v"/>						
Please select top report criteria	Action <input type="button" value="v"/> <input style="float: right;" type="button" value="+"/>						
Mail Frequency	Send once every <input type="text" value="Day"/> starting at <input type="text" value="2017-11-05 19:00"/> <input type="button" value="Calendar"/>						

2. Explore the options for scheduled reporting.

Class 2: Introduction to ELK Stack (ELK Coolness)

This class covers the following topics:

- ELK Stack Overview
- ELK Stack build on Ubuntu
- F5 logging to ELK Stack
- ELK Stack:
 - Indexes
 - Navigation
 - Searches
 - Visualisations
 - Dashboards

Expected time to complete: **1.5 hours**

6.1 Module 1: ELK Stack Build Ubuntu Server

ELK stack from the previous module is made up of three key components:

- Logstash,
- Elasticsearch,
- Kibana.

There are many ready made ELK stack services that can be used:

- Docker <https://elk-docker.readthedocs.io/>
- AWS <https://aws.amazon.com/elasticsearch-service/>
- Azure <https://azuremarketplace.microsoft.com/en-us/marketplace/apps/bitnami.elk?tab=Overview>
- Elastic Cloud <https://www.elastic.co/cloud>

However, it is important to understand how the ELK stack is build, the configuration files and their purposes. This module will guide you through the installation of ELK stack onto a ubuntu server.

External Reference Documentation:

<https://www.elastic.co/guide/index.html>

6.1.1 Lab 1.1: Install the Ubuntu Base

In this lab you will walk through installing the ubuntu base ready for ELK stack

Task 1 - GIT Clone Repo onto the Server

git clone https://github.com/jarrodlucia/bigip_elk_server <directory of choice>

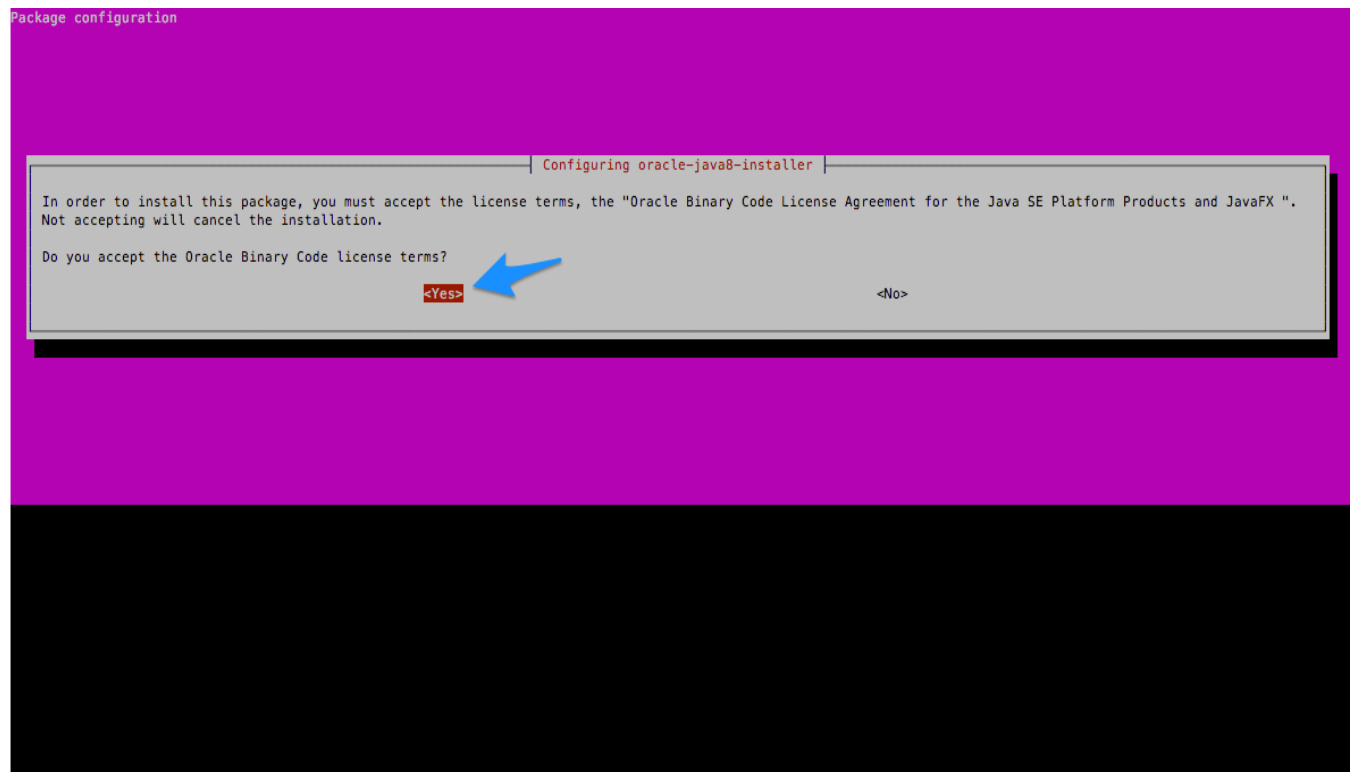
Task 2 - Install additional software required for ELK Stack

```
sudo apt-get install software-properties-common
sudo apt install curl
```

Task 3 - Add and install Java

```
sudo add-apt-repository -y ppa:webupd8team/java
sudo apt-get update
sudo apt-get -y install oracle-java8-installer
```

Accept the Oracle License Agreement



Fix Below for Java8 Error (If Required)

```
sudo apt-get -y install oracle-java8-installer
sudo sed -i 's|JAVA_VERSION=8u144|JAVA_VERSION=8u152|' oracle-java8-installer.*
sudo sed -i 's|PARTNER_URL=http://download.oracle.com/otn-pub/java/jdk/8u144-b01/
↪090f390dda5b47b9b721c7dfaa008135/|PARTNER_URL=http://download.oracle.com/otn-pub/
↪java/jdk/8u152-b16/aa0333dd3019491ca4f6ddbe78cdb6d0/|' oracle-java8-installer.*
sudo sed -i 's|SHA256SUM_TGZ=
↪"e8a341ce566f32c3d06f6d0f0eeea9a0f434f538d22af949ae58bc86f2eeaae4"|SHA256SUM_TGZ=
↪"218b3b340c3f6d05d940b817d0270dfe0cfd657a636bad074dcabe0c111961bf"|' oracle-java8-
↪installer.*
sudo sed -i 's|J_DIR=jdk1.8.0_144|J_DIR=jdk1.8.0_152|' oracle-java8-installer.*
sudo apt-get -y install oracle-java8-installer
```

6.1.2 Lab 1.2: Install Elasticsearch

This lab will install the Elasticsearch component, It is recommended to install Elasticsearch as the first module.

Task 1 Install Repo and Keys

1. Download and install the public signing key:

```
wget -qO - https://artifacts.elastic.co/GPG-KEY-elasticsearch | sudo apt-key add -
```

1. Save the repository definition to /etc/apt/sources.list.d/elastic-5.x.list:

```
echo "deb https://artifacts.elastic.co/packages/5.x/apt stable main" | sudo tee -a /  
↳etc/apt/sources.list.d/elastic-5.x.list  
sudo apt-get update
```

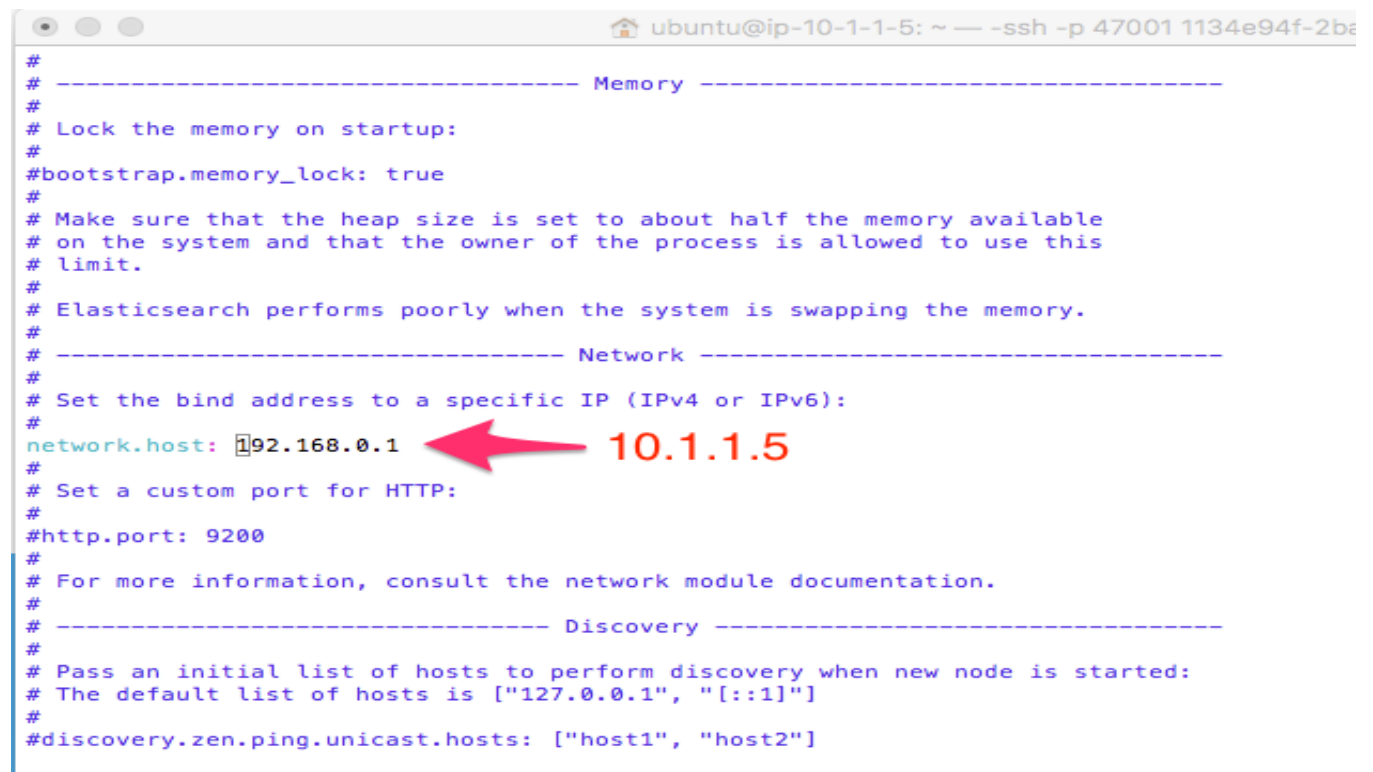
Task 2 Install elasticseach and setup system

1. Install Elasticsearch

```
sudo apt-get install elasticsearch
```

1. Edit config file to change bind address to Host address 10.1.1.5

```
sudo vi /etc/elasticsearch/elasticsearch.yml
```



```
ubuntu@ip-10-1-1-5: ~ — -ssh -p 47001 1134e94f-2ba  
#  
# ----- Memory -----  
#  
# Lock the memory on startup:  
#  
#bootstrap.memory_lock: true  
#  
# Make sure that the heap size is set to about half the memory available  
# on the system and that the owner of the process is allowed to use this  
# limit.  
#  
# Elasticsearch performs poorly when the system is swapping the memory.  
#  
# ----- Network -----  
#  
# Set the bind address to a specific IP (IPv4 or IPv6):  
#  
network.host: 192.168.0.1  
#  
# Set a custom port for HTTP:  
#  
#http.port: 9200  
#  
# For more information, consult the network module documentation.  
#  
# ----- Discovery -----  
#  
# Pass an initial list of hosts to perform discovery when new node is started:  
# The default list of hosts is ["127.0.0.1", "[::1]"]  
#  
#discovery.zen.ping.unicast.hosts: ["host1", "host2"]
```

1. Install additional plugins

```
sudo /usr/share/elasticsearch/bin/elasticsearch-plugin install ingest-geoip
```

1. Restart Elastic Search

```
sudo systemctl restart elasticsearch
```

1. Configure the system to start at boot

```
sudo /bin/systemctl daemon-reload  
sudo /bin/systemctl enable elasticsearch.service
```

1. Checking Start / Stop / Status

```
sudo systemctl start elasticsearch.service
sudo systemctl stop elasticsearch.service
sudo systemctl status elasticsearch.service
```

```
ubuntu@ip-10-1-1-5:~$ sudo systemctl status elasticsearch.service
● elasticsearch.service - Elasticsearch
   Loaded: loaded (/usr/lib/systemd/system/elasticsearch.service; enabled; vendor preset: enabled)
   Active: active (running) since Mon 2017-11-06 11:24:11 UTC; 17s ago
     Docs: http://www.elastic.co
   Main PID: 4425 (java)
   CGroup: /system.slice/elasticsearch.service
           └─4425 /usr/bin/java -Xms2g -Xmx2g -XX:+UseConcMarkSweepGC -XX:CMSInitiatingOccupancyFraction=75 -XX:
Nov 06 11:24:11 ip-10-1-1-5 systemd[1]: Starting Elasticsearch...
Nov 06 11:24:11 ip-10-1-1-5 systemd[1]: Started Elasticsearch.
ubuntu@ip-10-1-1-5:~$
```

6.1.3 Lab 1.3: Install Kibana

In this lab we will install Kibana

Task 1 Install Kibana

1. Install Kibana

```
sudo apt-get install kibana
```

2. Change config file to set Outside IP address

```
sudo vi /etc/kibana/kibana.yml
```

Note: Kibana is served by a back end server. This setting specifies the port to use. Server port is set as default Kibana Port 5601. Server host should be set to the UDF Management IP address 10.1.1.5 as we will be accessing this via the Linux Jump host. The URL of the Elasticsearch instance to use for all your queries.

- server.port: 5601
- server.host: "10.1.1.5"
- elasticsearch.url: "http://10.1.1.5:9200"

```
ubuntu@ip-10-1-1-5: ~ -- ssh -p 47001 1134e94f-2ba0-4477-899a-f7b1549a1581.access.udf.f5.com -- 15:
Kibana is served by a back end server. This setting specifies the port to use.
server.port: 5601

# Specifies the address to which the Kibana server will bind. IP addresses and host names are both valid values.
# The default is 'localhost', which usually means remote machines will not be able to connect.
# To allow connections from remote users, set this parameter to a non-loopback address.
server.host: "10.1.1.5"

# Enables you to specify a path to mount Kibana at if you are running behind a proxy. This only affects
# the URLs generated by Kibana, your proxy is expected to remove the basePath value before forwarding requests
# to Kibana. This setting cannot end in a slash.
#server.basePath: ""

# The maximum payload size in bytes for incoming server requests.
#server.maxPayloadBytes: 1048576

# The Kibana server's name. This is used for display purposes.
#server.name: "your-hostname"

# The URL of the Elasticsearch instance to use for all your queries.
elasticsearch.url: "http://10.1.1.5:9200"

# When this setting's value is true Kibana uses the hostname specified in the server.host
# setting. When the value of this setting is false, Kibana uses the hostname of the host
# that connects to this Kibana instance.
#elasticsearch.preserveHost: true

# Kibana uses an index in Elasticsearch to store saved searches, visualizations and
# dashboards. Kibana creates a new index if the index doesn't already exist.
#kibana.index: ".kibana"

# The default application to load.
#kibana.defaultAppId: "discover"

# If your Elasticsearch is protected with basic authentication, these settings provide
# the username and password that the Kibana server uses to perform maintenance on the Kibana
# index at startup. Your Kibana users still need to authenticate with Elasticsearch, which
# is proxied through the Kibana server.
#elasticsearch.username: "user"
#elasticsearch.password: "pass"

"/etc/kibana/kibana.yml" 104L, 4644C
```

1. Kibana restart

```
sudo systemctl restart kibana.service
```

2. To configure Kibana to start automatically when the system boots up, run the following commands:

```
sudo /bin/systemctl daemon-reload
sudo /bin/systemctl enable kibana.service
```

3. Kibana Control

```
sudo systemctl start kibana.service
sudo systemctl stop kibana.service
```

4. Check Kibana is running via command-line:

6.1.4 Lab 1.4: Install Logstash

Install Logstash

Task 1 - Install Logstash

1. Install Logstash


```
sudo apt-get install logstash
```

1. Install Additional Plugins

```
sudo /usr/share/logstash/bin/logstash-plugin install logstash-filter-dns  
sudo /usr/share/logstash/bin/logstash-plugin install logstash-filter-geoip
```

Note: Be patient with plugin install it can take a few moments

```
[ubuntu@ip-10-1-1-5:~$ sudo apt-get install logstash  
Reading package lists... Done  
Building dependency tree  
Reading state information... Done  
The following packages were automatically installed and are no longer required:  
  linux-headers-4.4.0-36 linux-headers-4.4.0-36-generic linux-image-4.4.0-36-generic  
Use 'sudo apt autoremove' to remove them.  
The following NEW packages will be installed:  
  logstash  
0 upgraded, 1 newly installed, 0 to remove and 0 not upgraded.  
Need to get 105 MB of archives.  
After this operation, 192 MB of additional disk space will be used.  
Get:1 https://artifacts.elastic.co/packages/5.x/apt stable/main amd64 logstash amd64 5.6.3-1 [105 MB]  
Fetched 105 MB in 4s (23.3 MB/s)  
Selecting previously unselected package logstash.  
(Reading database ... 148182 files and directories currently installed.)  
Preparing to unpack .../logstash_1%3a5.6.3-1_all.deb ...  
Unpacking logstash (1:5.6.3-1) ...  
Setting up logstash (1:5.6.3-1) ...  
Using provided startup.options file: /etc/logstash/startup.options  
Successfully created system startup script for Logstash  
[ubuntu@ip-10-1-1-5:~$ sudo /usr/share/logstash/bin/logstash-plugin install logstash-filter-dns  
Validating logstash-filter-dns  
Installing logstash-filter-dns
```



1. Copy or Create new file to Directory /etc/logstash/conf.d/

```
sudo cp <git clone directory>/config_files/logstash.conf /etc/logstash/conf.d/  
↪ logstash.conf  
sudo vi /etc/logstash/conf.d/logstash.conf
```

1. Logstash restart

```
sudo systemctl restart logstash.service
```

1. Check logstash started correctly with no errors from logstash.conf file

```

ubuntu@ip-10-1-1-5:~$ tail -f /var/log/logstash/logstash-plain.log
[2017-11-06T12:50:24,971][INFO ][logstash.outputs.elasticsearch] Running health check to see if an
p://localhost:9200/, :path=>"/"}
[2017-11-06T12:50:24,976][WARN ][logstash.outputs.elasticsearch] Restored connection to ES instance
[2017-11-06T12:50:24,980][INFO ][logstash.outputs.elasticsearch] Using mapping template from {pat
[2017-11-06T12:50:24,981][INFO ][logstash.outputs.elasticsearch] Attempting to install template {:
1, "settings"=>{"index.refresh_interval"=>"5s"}, "mappings"=>{"_default_"=>{"_all"=>{"enabled"=>tr
=>{"path_match"=>"message", "match_mapping_type"=>"string", "mapping"=>{"type"=>"text", "norms"=>f
pe"=>"string", "mapping"=>{"type"=>"text", "norms"=>false, "fields"=>{"keyword"=>{"type"=>"keyword
{"type"=>"date", "include_in_all"=>false}, "@version"=>{"type"=>"keyword", "include_in_all"=>false
>"ip"}, "location"=>{"type"=>"geo_point"}, "latitude"=>{"type"=>"half_float"}, "longitude"=>{"type
[2017-11-06T12:50:24,985][INFO ][logstash.outputs.elasticsearch] Installing elasticsearch template
[2017-11-06T12:50:25,001][INFO ][logstash.outputs.elasticsearch] New Elasticsearch output {:class=
9200"}]
[2017-11-06T12:50:25,007][INFO ][logstash.filters.geoip      ] Using geoip database {:path=>"/usr/sha
oip-4.3.1-java/vendor/GeoLite2-City.mmdb"}
[2017-11-06T12:50:25,037][INFO ][logstash.filters.geoip      ] Using geoip database {:path=>"/usr/sha
oip-4.3.1-java/vendor/GeoLite2-City.mmdb"}
[2017-11-06T12:50:25,044][INFO ][logstash.pipeline          ] Starting pipeline {"id"=>"main", "pipe
ch.delay"=>5, "pipeline.max_inflight"=>500}
[2017-11-06T12:50:25,188][INFO ][logstash.pipeline          ] Pipeline main started
[2017-11-06T12:50:25,212][INFO ][logstash.agent            ] Successfully started Logstash API endp

```

1. To configure Logstash to start automatically when the system boots up, run the following commands:

```

sudo /bin/systemctl daemon-reload
sudo /bin/systemctl enable logstash.service

```

1. Logstash Control

```

sudo systemctl start logstash.service
sudo systemctl stop logstash.service
sudo systemctl status logstash.service

```

logstash.conf

```

1  input {
2    tcp {
3      port => 5516
4      type => afm
5    }
6    tcp {
7      port => 5515
8      type => dns
9    }
10   tcp {
11     port => 5514
12     type => pem
13   }
14 }
15
16 filter {
17   if [type] == 'pem' {
18     kv {
19       source => "message"
20       field_split => ", "
21     }
22   }
23   if [type] == 'afm' {
24     kv {
25       source => "message"

```



```

26         field_split => ",",
27     }
28     geoip {
29         source => "SourceIp"
30         target => "SourceIp_geo"
31         add_field => [ "[geoip][coordinates]", "%{[geoip][longitude]}"
↵]
32         add_field => [ "[geoip][coordinates]", "%{[geoip][latitude]}"
↵]
33     }
34     geoip {
35         source => "DestinationIp"
36         target => "DestinationIp_geo"
37         add_field => [ "[geoip][coordinates]", "%{[geoip][longitude]}"
↵]
38         add_field => [ "[geoip][coordinates]", "%{[geoip][latitude]}"
↵]
39     }
40     mutate {
41         convert => [ "[geoip][coordinates]", "float"]
42     }
43 }
44 if [type] == 'dns' {
45     kv {
46         source => "message"
47         field_split => ",",
48     }
49 }
50 }
51
52 output {
53     if [type] == 'pem' {
54         elasticsearch {
55             hosts => ["10.1.1.5:9200"]
56             index => "pem-%{+YYYY.MM.dd}"
57             template_name => "pem"
58         }
59     }
60     if [type] == 'afm' {
61         elasticsearch {
62             hosts => ["10.1.1.5:9200"]
63             index => "afm-%{+YYYY.MM.dd}"
64             template_name => "afm"
65         }
66     }
67     if [type] == 'dns' {
68         elasticsearch {
69             hosts => ["10.1.1.5:9200"]
70             index => "dns-%{+YYYY.MM.dd}"
71             template_name => "dns"
72         }
73     }
74     stdout {}
75 }

```


6.1.5 Lab 1.5: Configure elasticsearch templates

Templates are used to create mappings between logstash and elasticsearch. Without the mappings elasticsearch will create automatic mappings however these will be elasticsearch's best guess as to the field. In most cases this will default to `text`. This means many of the fields such as IP address's will be searchable but not able to be used in Visualisations.

Upload elasticsearch templates and mappings. There are multiple way this can be achieved. The most common ways are cURL and a REST based program such as POSTMAN. Feel free to use whichever method you are most comfortable with.

Note: RECOMMENDATION Use cURL for the uploading of the templates with json file. POSTMAN is useful for Elasticsearch management once the template are in place.

Task 1 Option1 - Install module templates in Elasticsearch via cURL

1. Install Index Templates into Elastic Search for the required modules

`cd <git clone directory>/json/` **git clone directory from Lab 1**

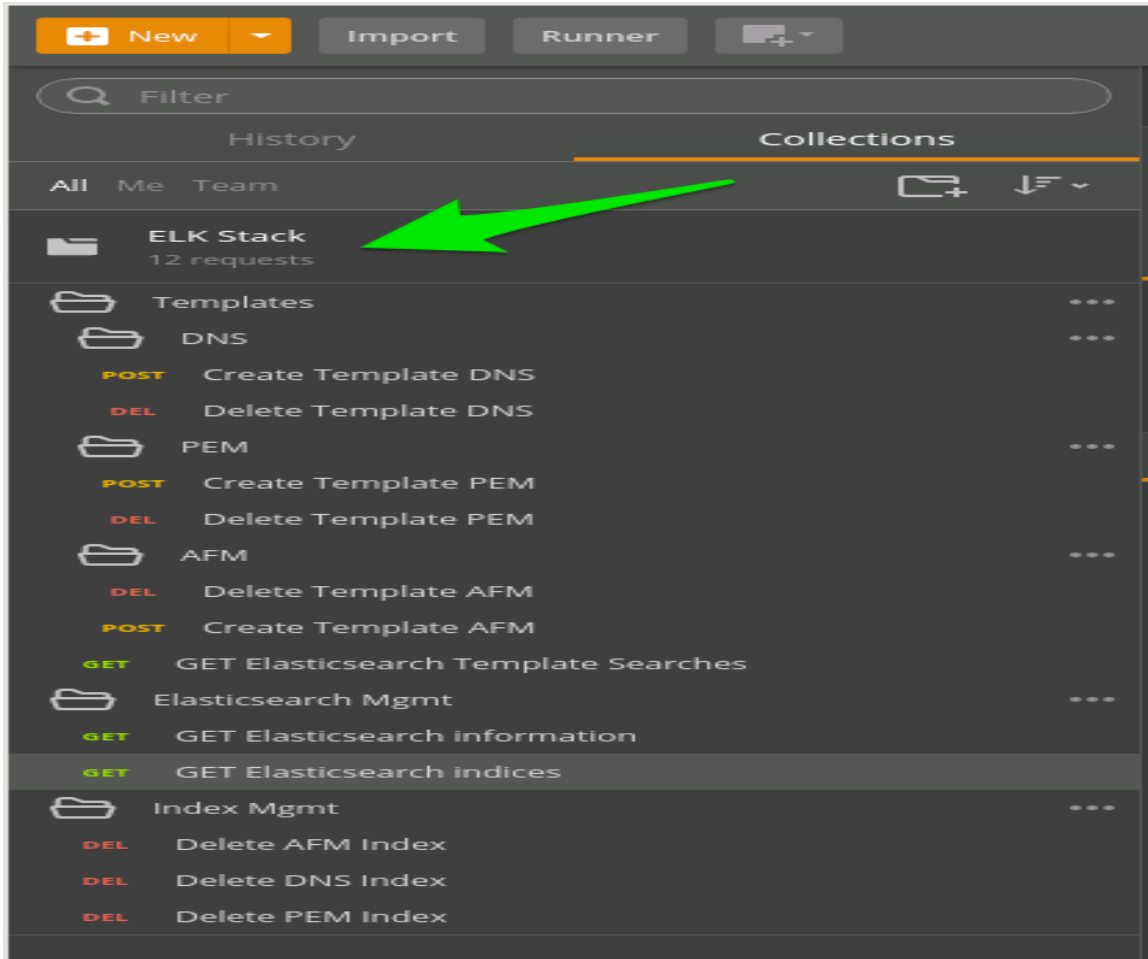
```
curl -XPUT http://localhost:9200/_template/pem?pretty -d @pem_mapping.json
curl -XPUT http://localhost:9200/_template/afm?pretty -d @afm_mapping.json
curl -XPUT http://localhost:9200/_template/dns?pretty -d @dns_mapping.json
```

Task 1 Option1 - Install module templates in Elasticsearch via POSTMAN

1. Import ELK Postman Collection and Environment
2. Click the 'Import from Link' tab. Paste the following URL into the text box and click 'Import'


```
https://raw.githubusercontent.com/jarrodlucia/bigip\_elk\_server/develop/postman\_collections/ELKStack.postman\_collection.json
```

3. You should now see a collection named 'F5 ELK' in your Postman Collections sidebar:



4. Import the Environment file by clicking 'Import' -> 'Import from Link' and pasting the following URL and clicking 'Import':

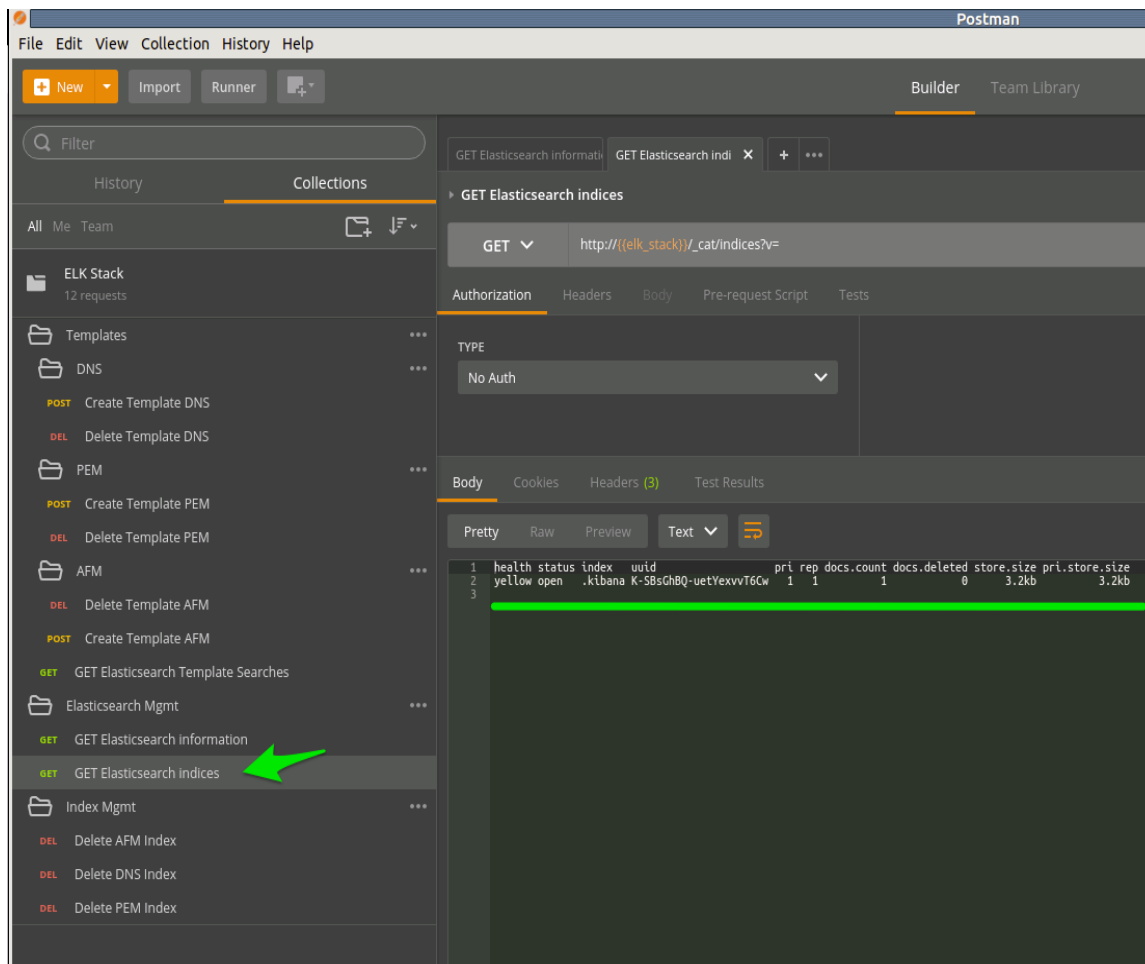
`https://raw.githubusercontent.com/jarrodlucia/bigip_elk_server/develop/postman_collections/F5ELKEnv.postman_environment.json`



1. Click on GET Elasticsearch information, **HIT SEND**.

You should see cluster information regarding elasticsearch

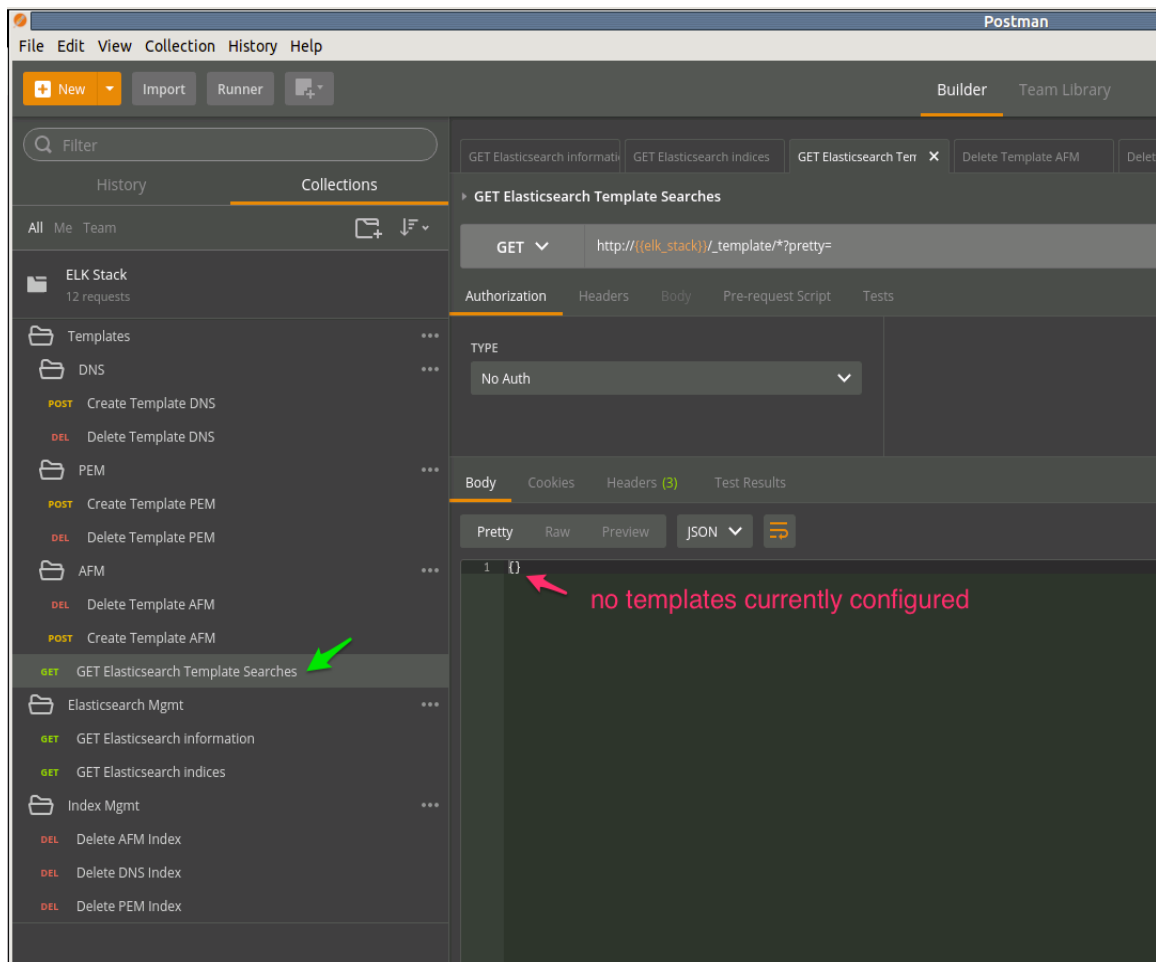
1. Click on GET Elasticsearch indices, **HIT SEND**.



You should see the current index's and information regarding each index.

We will use this command to observe the creation of new indexes

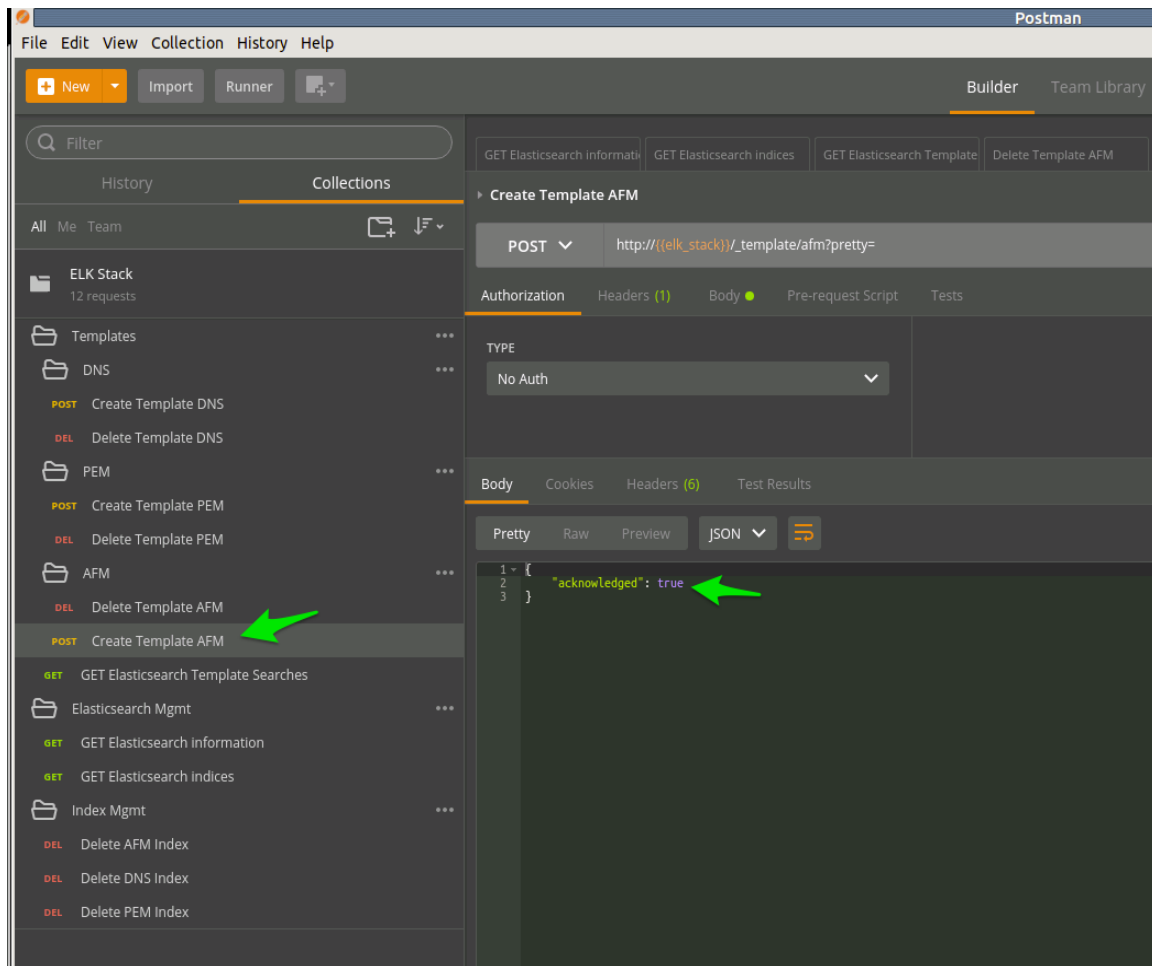
1. Click on GET Elasticsearch Template Searches, **HIT SEND**



You should see any current templates listed.

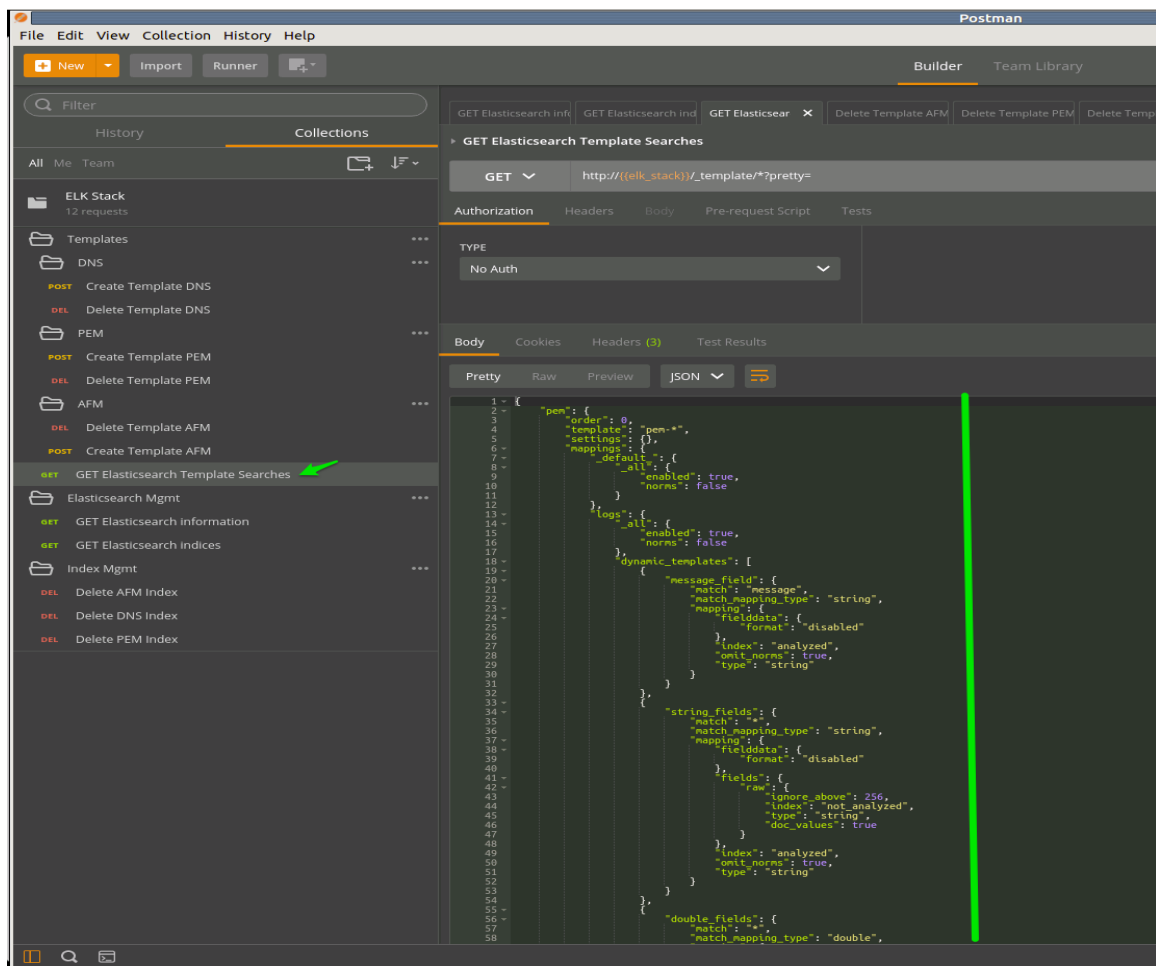
Note: New Install will **NOT** contain any templates showing {}

1. Click on Create Template AFM + PEM + DNS **Install all templates**



Note: Create all templates from the POSTMAN collection

1. Verify templates created and exist. Click on GET Elasticsearch Template Searches



Note: Look through the template JSON outputted by POSTMAN. Verify and check that the three templates created are present.

6.1.6 Lab 1.6: Send Logs to ELK Stack

Configure f5 for logging to new ELK stack

Check that data is arriving at ELK stack

Task 1 - Confirm BIG-IP is sending logs to ELK Stack

1. Confirm via TMUI that the setup from **Class 1 Lab 2.1**

Update AFM Reporting to include what was not included in previous lab.

Security » Network Address Translation : Policies » **Subscriber_CGNAT_Policy**

Active Rules Policies Source Translation Destination Translation

Unsaved changes to the policy!
One or more policy rules have been modified but not committed to the system. Changes must be committed to the system before taking effect.
[Commit Changes to System](#) [Cancel Changes](#)

Properties

Name Subscriber_CGNAT_Policy

Description

Filter Policy List

ID	Name	State	Protocol	Source	Destination
<input type="checkbox"/> 1	sub01	enabled	any	Addresses 10.1.10.30/32	Any
2	sub00	Enabled	Any	Addresses 10.1.10.25/32	add new destination Add

Description

add new source Add

Done Editing Cancel

Note:

Make sure the correct port is allocated as per previous Logstash configuration

- Pool = tcp server:5514 - PEM
 - Pool = tcp server:5515 - DNS
 - Pool = tcp server:5516 - AFM/CGNAT
-

1. Confirm Data is arriving on server

`sudo tcpdump -i eth1 port 5514`

1. Check that Data is arriving in the Index

`curl 'localhost:9200/_cat/indices?v'`

```
ubuntu@ip-172-31-9-140:~$ curl 'localhost:9200/_cat/indices?v'
```

health	status	index	uuid	pri	rep	docs.count	docs.deleted	store.size	pri.store.size
yellow	open	afm-2017.10.30	KsbHmgC5T_-Pb6zCiWq04A	5	1	99099	0	96mb	96mb
yellow	open	afm-2017.10.20	vNiFLuhoT3yZcQLlLNng	5	1	91377	0	93.4mb	93.4mb
yellow	open	pem-2017.11.06	WxNj_Bn6S0qgXka4pqboRw	5	1	107905	0	68.8mb	68.8mb
yellow	open	dns-2017.10.23	6V_oD6vgS5CI_FNZKQ81mw	5	1	23192	0	6.7mb	6.7mb
yellow	open	pem-2017.11.01	6e8HsrsaR9W9ZswMBi1jQg	5	1	241562	0	154.5mb	154.5mb
yellow	open	dns-2017.10.26	e8pb_CnwQg2Kb_wTuDkS9A	5	1	23354	0	6.2mb	6.2mb
yellow	open	afm-2017.11.01	gsYtMwETQcyPXWEWuIQkVg	5	1	124940	0	125.7mb	125.7mb
yellow	open	pem-2017.10.25	jcZYTnKQQ5myhMONIGqe3Q	5	1	197893	0	134.5mb	134.5mb
yellow	open	pem-2017.10.24	klMyNpnHQJakQhaipQcdEQ	5	1	164781	0	112.4mb	112.4mb
yellow	open	dns-2017.10.27	FIls_BWLRQuYnsDQQYJkgQ	5	1	42133	0	12.2mb	12.2mb
yellow	open	afm-2017.10.29	AeT2IKi3S_OuxS0QKylkkg	5	1	109143	0	107.4mb	107.4mb
yellow	open	pem-2017.10.20	n4u6ln-hRc05dK0ly_Uj5w	5	1	97583	0	76mb	76mb
yellow	open	dns-2017.10.19	2gkvW3ibRLyxsHeKBejqTw	5	1	33939	0	9.7mb	9.7mb
yellow	open	dns-2017.10.20	Xudl6KAFTEqIKlj7oATUGg	5	1	42111	0	12.7mb	12.7mb
yellow	open	pem-2017.10.31	GCMJjdPTSp-kFhUkLayG2A	5	1	184746	0	135.6mb	135.6mb
yellow	open	afm-2017.10.21	_SS3GGdBSA0EvXMRWZexAA	5	1	120292	0	120.8mb	120.8mb
yellow	open	dns-2017.10.25	Ujjj6kT2IR_uZ02ZlpT9GAA	5	1	30585	0	8.5mb	8.5mb
yellow	open	afm-2017.10.25	oSVqB6cvSuy2-0zvD_EHCQ	5	1	119819	0	111.9mb	111.9mb
yellow	open	pem-2017.11.02	rl0XsTJSR0KhLmUuzWLxLA	5	1	360479	0	222.5mb	222.5mb
yellow	open	dns-2017.10.31	y144oZxYRDqONMIN5W56dA	5	1	35611	0	10.8mb	10.8mb
yellow	open	dns-2017.11.02	BSRK9Yj6QNmYwWBAAnSEGwA	5	1	19036	0	6mb	6mb
yellow	open	afm-2017.11.02	2oZ_ckfvSYSZgM0AvwURJQ	5	1	105112	0	121.9mb	121.9mb
yellow	open	dns-2017.11.06	47rN2rM7QyScu7WvPWihZA	5	1	4558	0	2.4mb	2.4mb
yellow	open	dns-2017.10.24	p_GXqXSLRNaeACEkr3IEkQ	5	1	25447	0	7.1mb	7.1mb
yellow	open	.kibana	GiC8XwNBTGKXnsw-fWYevg	1	1	58	20	208.4kb	208.4kb

or via POSTMAN

The screenshot shows the Postman application interface. On the left sidebar, under the 'Collections' tab, there is a collection named 'ELK Stack' containing 12 requests. Below it, there are folders for 'Templates', 'DNS', 'PEM', and 'AFM', each with a 'Create Template' and 'Delete Template' request. At the bottom of the sidebar, there is a 'GET Elasticsearch Template Searches' request.

The main area of the interface shows a GET request to the endpoint `http://(elk_stack)/_cat/indices?v=`. A green arrow points to this endpoint. Below the endpoint, the 'Authorization' tab is selected, showing 'No Auth'. The 'Body' tab is also visible. At the bottom, the 'Text' tab is selected, showing the response body. The response body is a table with 10 columns: health, status, index, uuid, pri, rep, docs.count, docs.deleted, store.size, and pri.store.size. The response body is highlighted with a green box.

health	status	index	uuid	pri	rep	docs.count	docs.deleted	store.size	pri.store.size
1	yellow	open	afm-2017.11.06	-nL4PESRuCcB9juA0U6_Q	5	1	1816	0	2.3mb
2	yellow	open	afm-2017.11.06	9Nsb1EPWTTCC0bt6d6A8zcA	5	1	596	0	1.6mb
3	yellow	open	afm-2017.11.06	zshjeIuvRk6F63TFK1ciJA	5	1	2256	0	1.1mb
4	yellow	open	dns-2017.11.06	K-SBSGhBQ-uetYexvvT6Cw	1	1	4	0	50.9kb

6.1.7 Lab 1.7: Create Index and Import Pre-Configured

Index's are elasticsearch's way of storing documents in shards. When index's are created the mapping templates we uploaded before are used to map each of the fields to a type. This is only done once when the index is created

Note: If mappings are changed are updates required the “index” will have to be deleted, the template deleted and mapping changed and template added. At this point re-creating the index will remap to the new template

This Lab will focus on creating the index's for each module based on logstash in **Lab4**

We will import the prepared f5 module json kibana searches / virtuals / and dashboards.

Task 1 - Create Kibana Index's

1. Configure Indexes in Kibana

Configure the first and default index

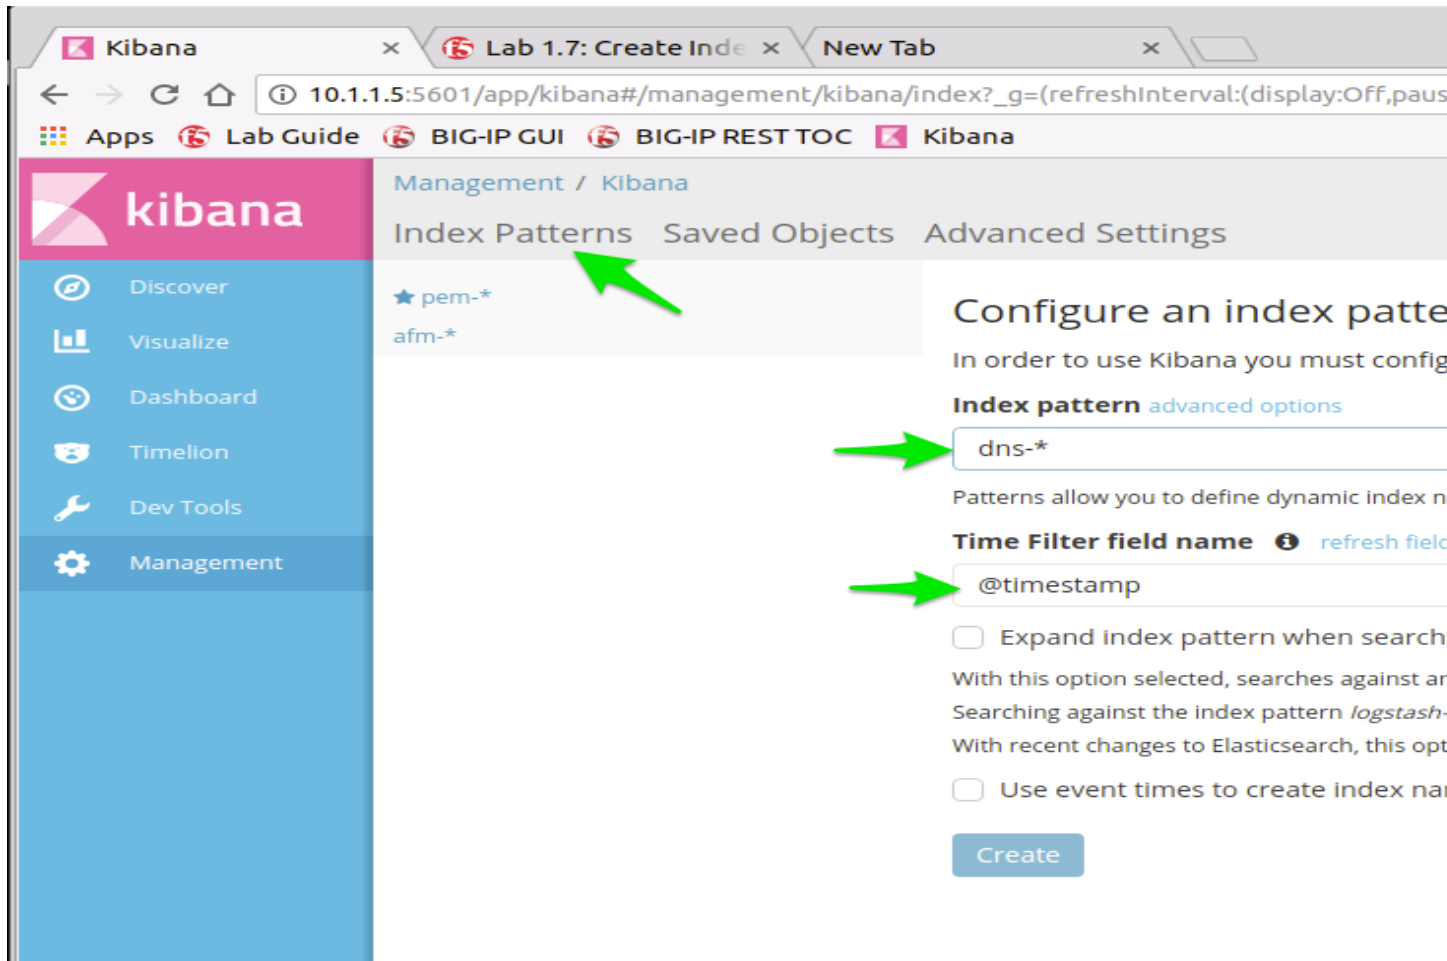
- index pattern = `logstash-*`
- select `@timestamp`

The screenshot shows the Kibana Management console interface. The left sidebar contains navigation links: Discover, Visualize, Dashboard, Timelion, Dev Tools, and Management. The main content area is titled 'Configure an index pattern'. It displays a warning message: 'Warning: No default index pattern. You must select or create one to continue.' Below the warning, the 'Index pattern' field is set to 'logstash-*'. The 'Time Filter field name' dropdown is set to '@timestamp'. There are two checkboxes: 'Expand index pattern when searching [DEPRECATED]' and 'Use event times to create index names [DEPRECATED]', both of which are unchecked. A blue button at the bottom says 'Time Filter field name is required'.

- index pattern = afm-*
- select @timestamps

Follow PEM example above for AFM

- index pattern = dns-*
- select @timestamps



Task 2 - Import preconfigured Kibana json's

Searches / Visualisation and Dashboards

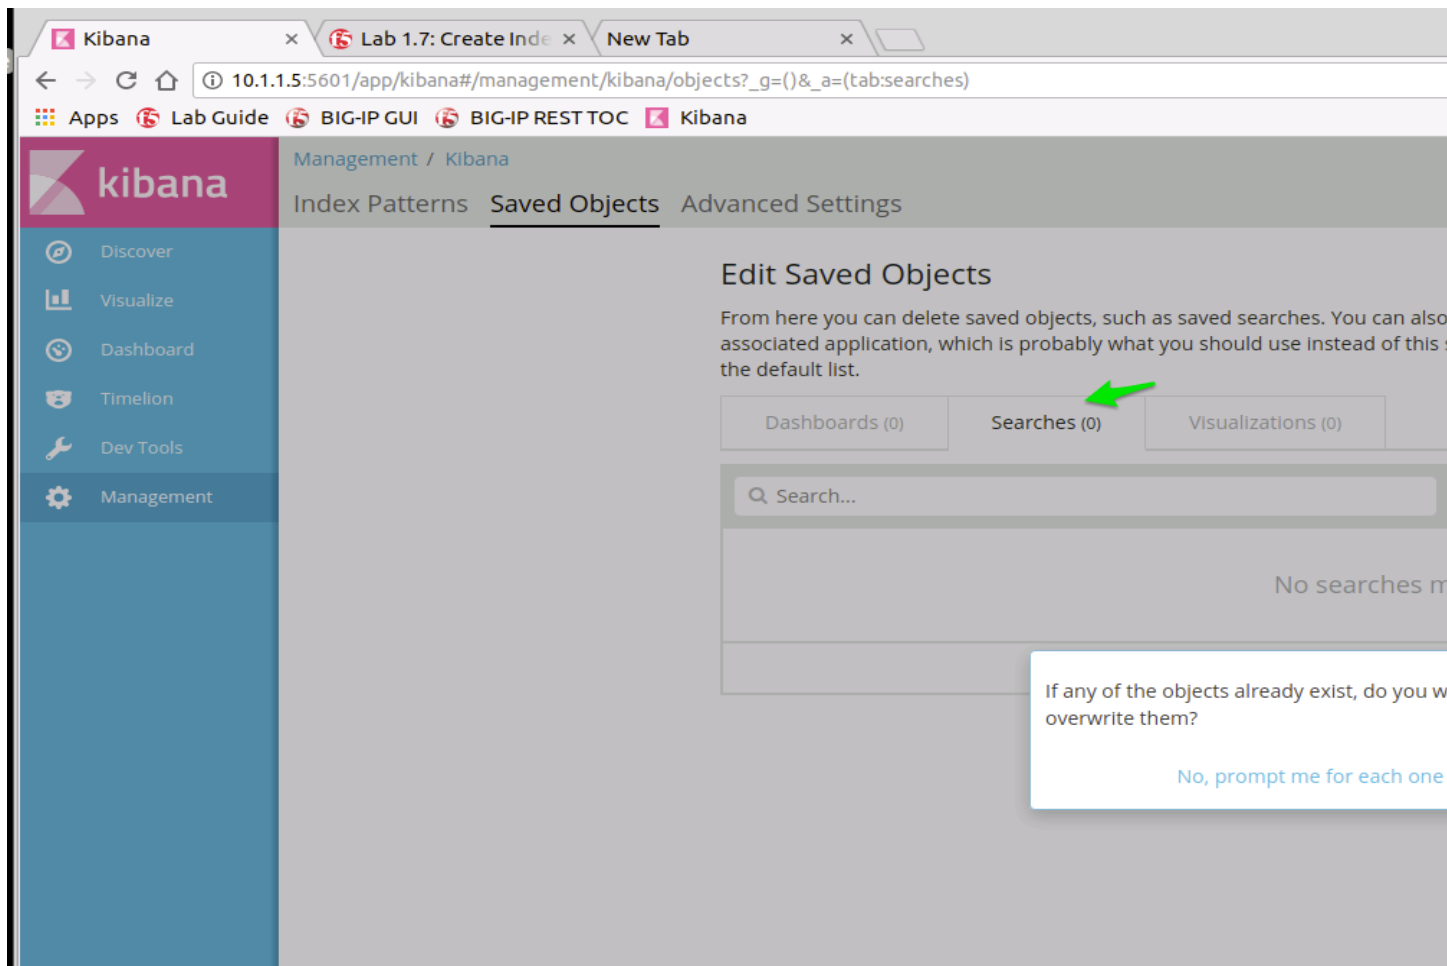
1. Import object data into Kibana

Import the JSON files in the following order:

- Searches
- Visualisations
- Dashboards

Searches

```
https://raw.githubusercontent.com/jarrodlucia/bigip_elk_server/develop/
json/elk_searches.json
```



Index Pattern Conflicts

The following saved objects use index patterns that do not exist. Please select the index patterns you'd like re-associated them with.

ID	Count	Sample of affected objects	New index pattern
AV84IG9n-zl4nyjF_fm6	15	URL Cat Name Subscriber ID Destination IP Source IP pem_subscriber	<div>pem → dns:*</div>
AV84JcD-zl4nyjF_fnV	3	All Logs AFM Translation Logs v2 AFM Translation Logs	<div>afm → dns:*</div>



[Cancel](#) [Confirm all changes](#)

Visuals

```
https://raw.githubusercontent.com/jarrodlucia/bigip_elk_server/develop/
json/elk_visualisations.json
```

Index Pattern Conflicts

The following saved objects use index patterns that do not exist. Please select the index patterns you'd like re-associated them with.

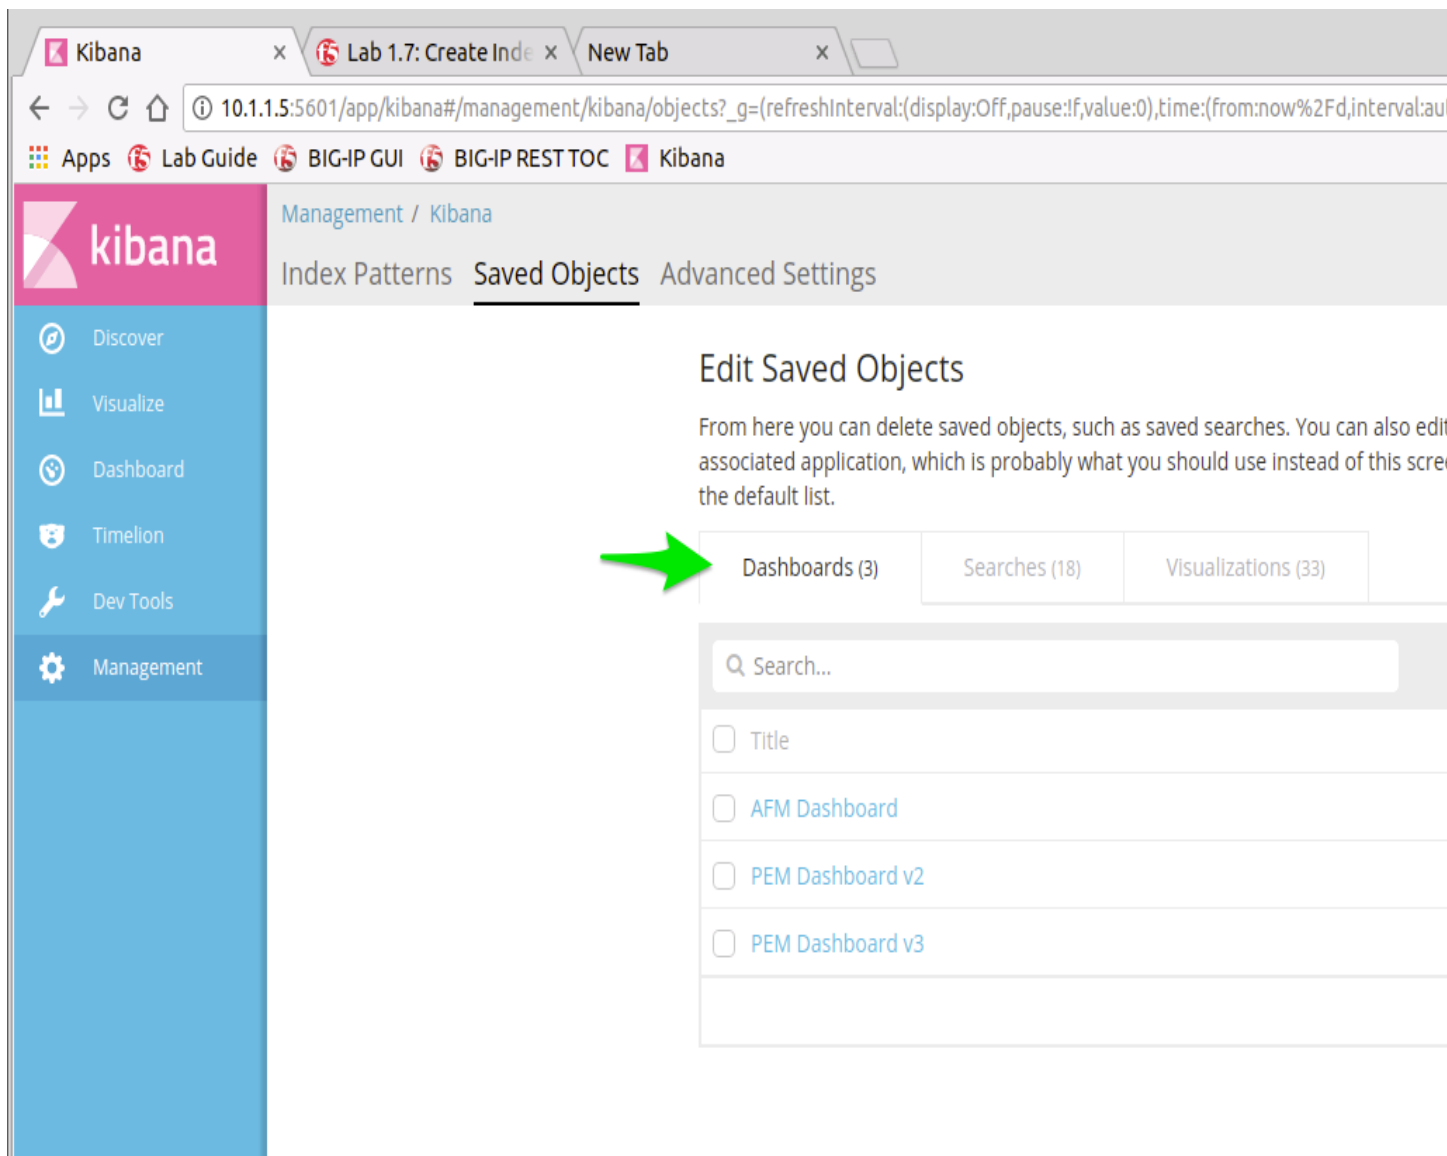
ID	Count	Sample of affected objects	New index pattern
AV84IjcD-zl4nyjF_fnV	17	AFM VS Context Program - Table Severity - Pie BIG-IP Devices Severity	afm-* 
AV84IG9n-zl4nyjF_fm6	2	Subscriber Names Subscriber IP List	pem-* 

Cancel

Confirm all changes

Dashboards

https://raw.githubusercontent.com/jarrodlucia/bigip_elk_server/develop/json/elk_dashboards.json



Note: The JSON files have been placed in the IN_CASE_OF_EMERGENCY folder on the desktop

6.2 Module 2: Kibana and Visualisation

Coolness of Kibana interface

- Navigation
- Searching
- Creating Searches
- Creating Visualisations
- Creating Dashboards

6.2.1 Lab 2.1 – Kibana Interface & Search

Kibana is the interface to elasticsearch and makes visualisation and dashboards available. It allows REST API calls for development of additional Customer interfaces.

This lab will look at the look and feel of the Kibana interface, and some key navigation hints and tips.

Task 1 - Kibana Interface Explantion

This task will focus on explation of the Kibana interface and navigating different aspects of the interface.

The screenshot shows the Kibana interface with several annotations:

- query bar**: Points to the search bar at the top with the text "Search... (e.g. status:200 AND extension:PHP)".
- index selection**: Points to the dropdown menu showing "afm-*".
- side navigation**: Points to the left sidebar menu.
- document table**: Points to the table of search results at the bottom.

The interface displays 52,460 hits. The left sidebar includes navigation options: Discover, Visualize, Dashboard, Timelion, Dev Tools, and Management. The main area shows a bar chart of counts over time (November 6th, 2017) and a table of documents with fields like device_version, device_product, hostname, and source_port.

Try changing the following:

- Time Range
- Index
- Dashboards

Note: Take your time to explore each of the interface elements.

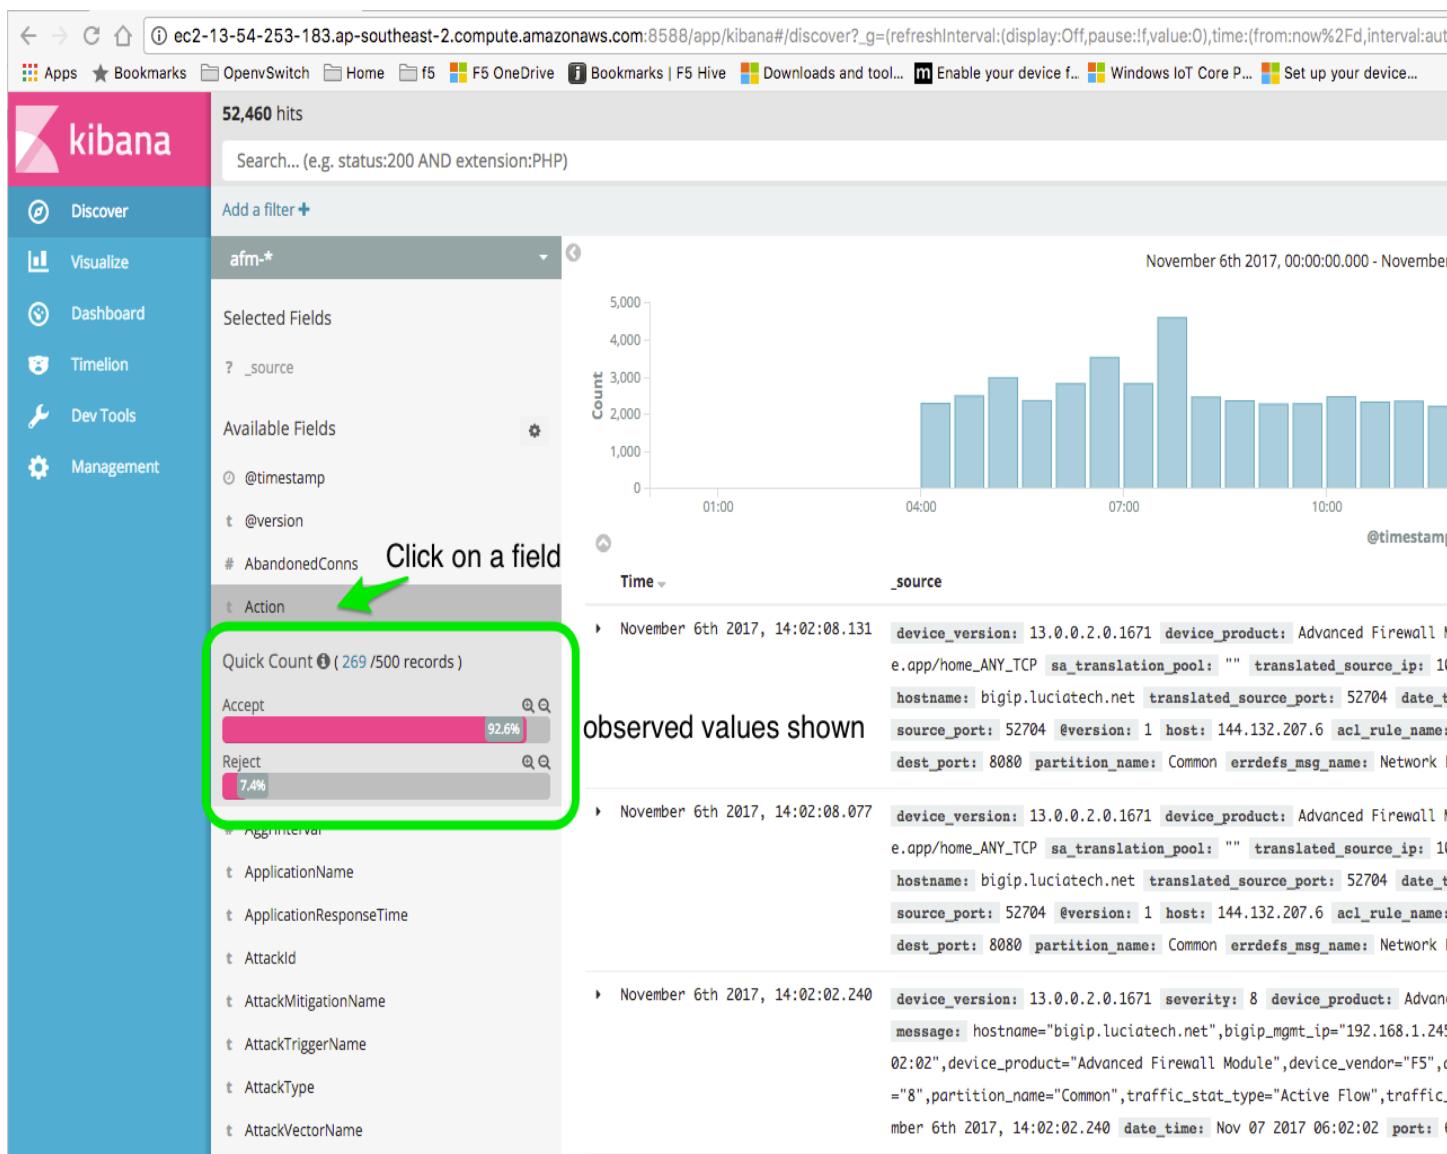
Task 2 - Searching Kibana

In this task we will use two example search types to see how Kibana uses elasticsearch. These example searches will be the following:

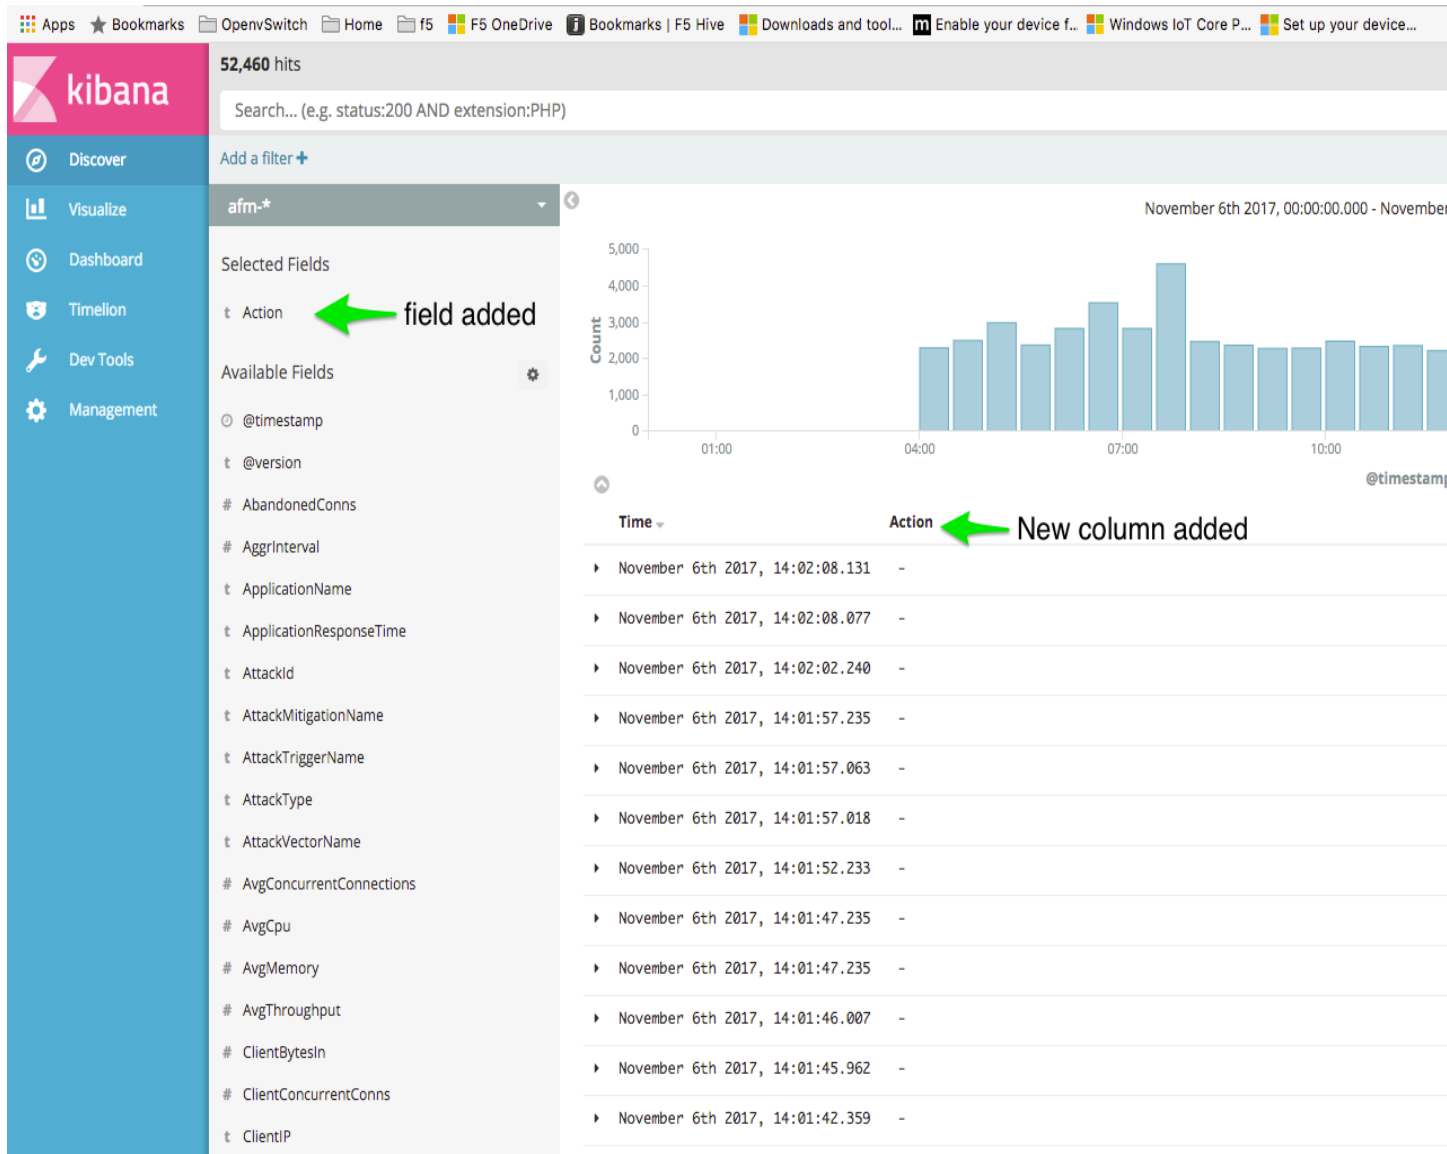
- Field Search
- Query Bar

Field Search Field searching is very useful in Kibana and can be used to see types of data and values that elasticsearch is indexing. To conduct field searching conduct the following:

1. Click on a field
2. Examine the expanded field, note the values that elasticsearch is indexing

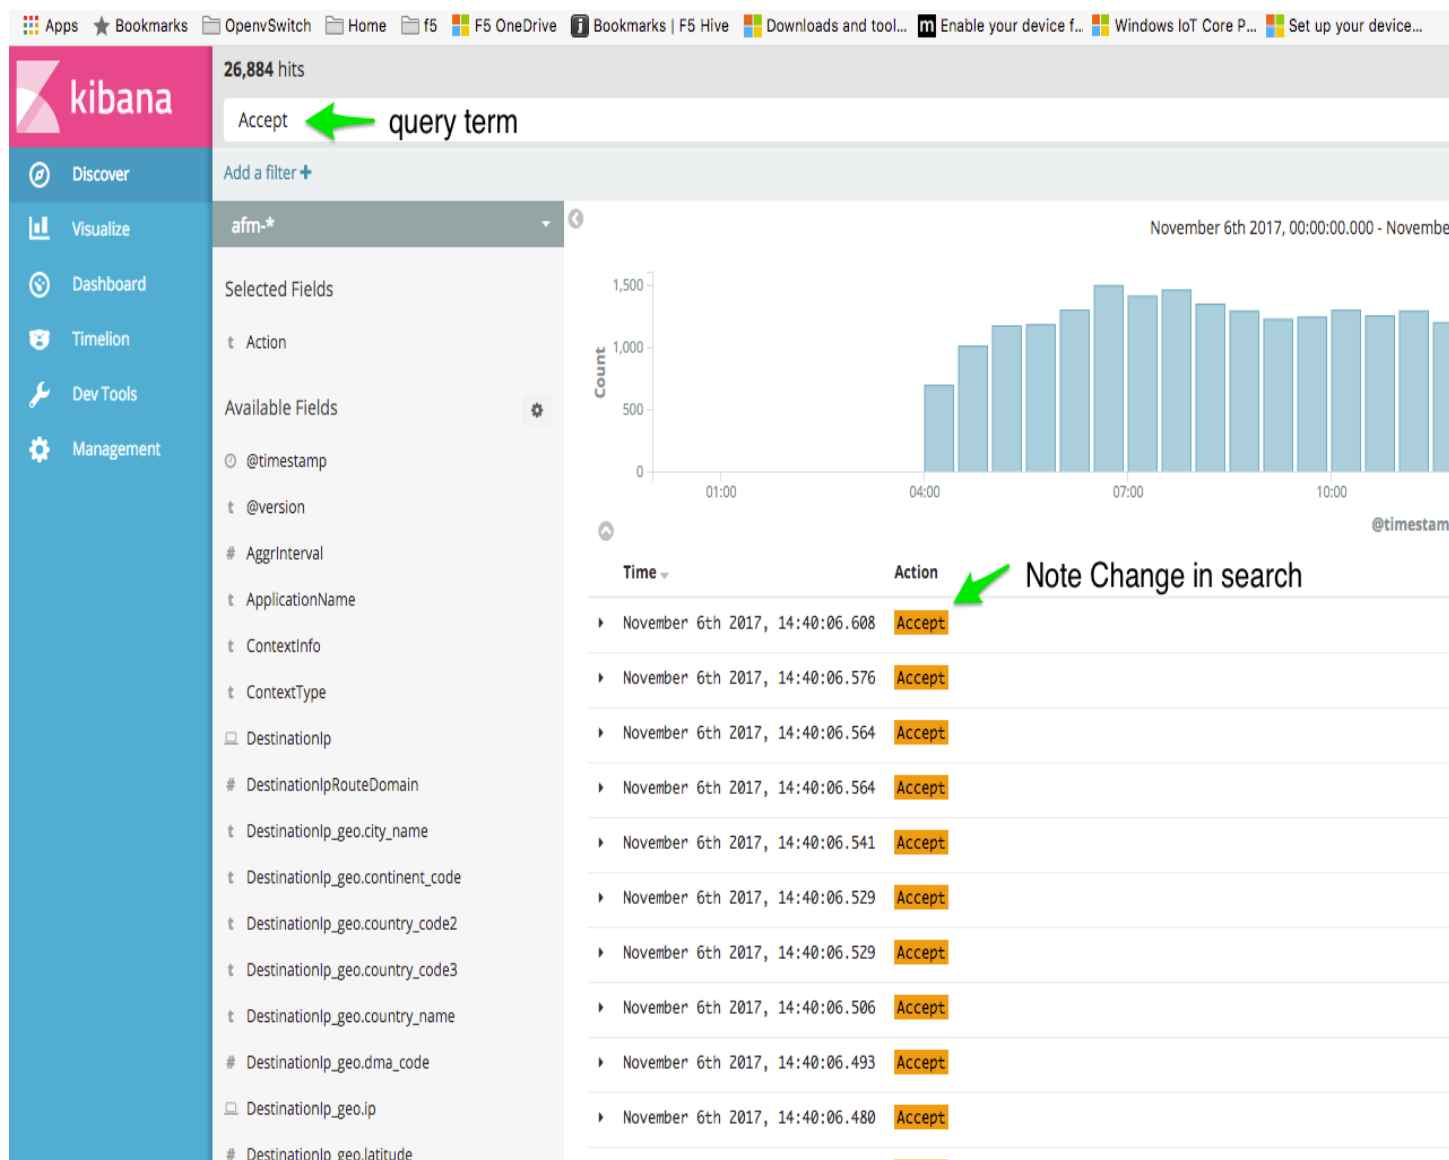


1. Click the add button.
2. Notice the field is in the Selected Field section.



Note: Take time to explore multiple field add to Selected field and build up a set of interesting columns.

Query Bar This type of searching is searching all data fields not only Selected fields as we did previously.



Note: Take time to explore multiple field add to Selected field and use Query terms to see the results.

6.2.2 Lab 2.2 – Creating Kibana Usefulness

This Lab will focus on creation of three key components of Kibana for useful display of information, namely:
- Searches - Visualisations - Dashboards

Task 1 - Creating Searches

Create and Save a 3 x search based on the previous lab.

Apps

★ Bookmarks

OpenvSwitch

Home

f5

F5 OneDrive

Bookmarks | F5 Hive

Downloads and tool...

m Enable your device f...

Windows IoT Core P...

Set up your device...

kibana

28,368 hits

Save Search

New Saved Search

Save

Accept

Add a filter +

afm-*

Selected Fields

t Action

Available Fields

@timestamp

@version

AggrInterval

t ApplicationName

t ContextInfo

t ContextType

DestinationIp

DestinationIpRouteDomain

Count

1,500

1,000

500

0

01:00

04:00

07:00

10:00

November 6th 2017, 00:00:00.000 - November

@timestamp

Time

Action

November 6th 2017, 15:15:07.497

Accept

November 6th 2017, 15:15:07.484

Accept

November 6th 2017, 15:15:07.456

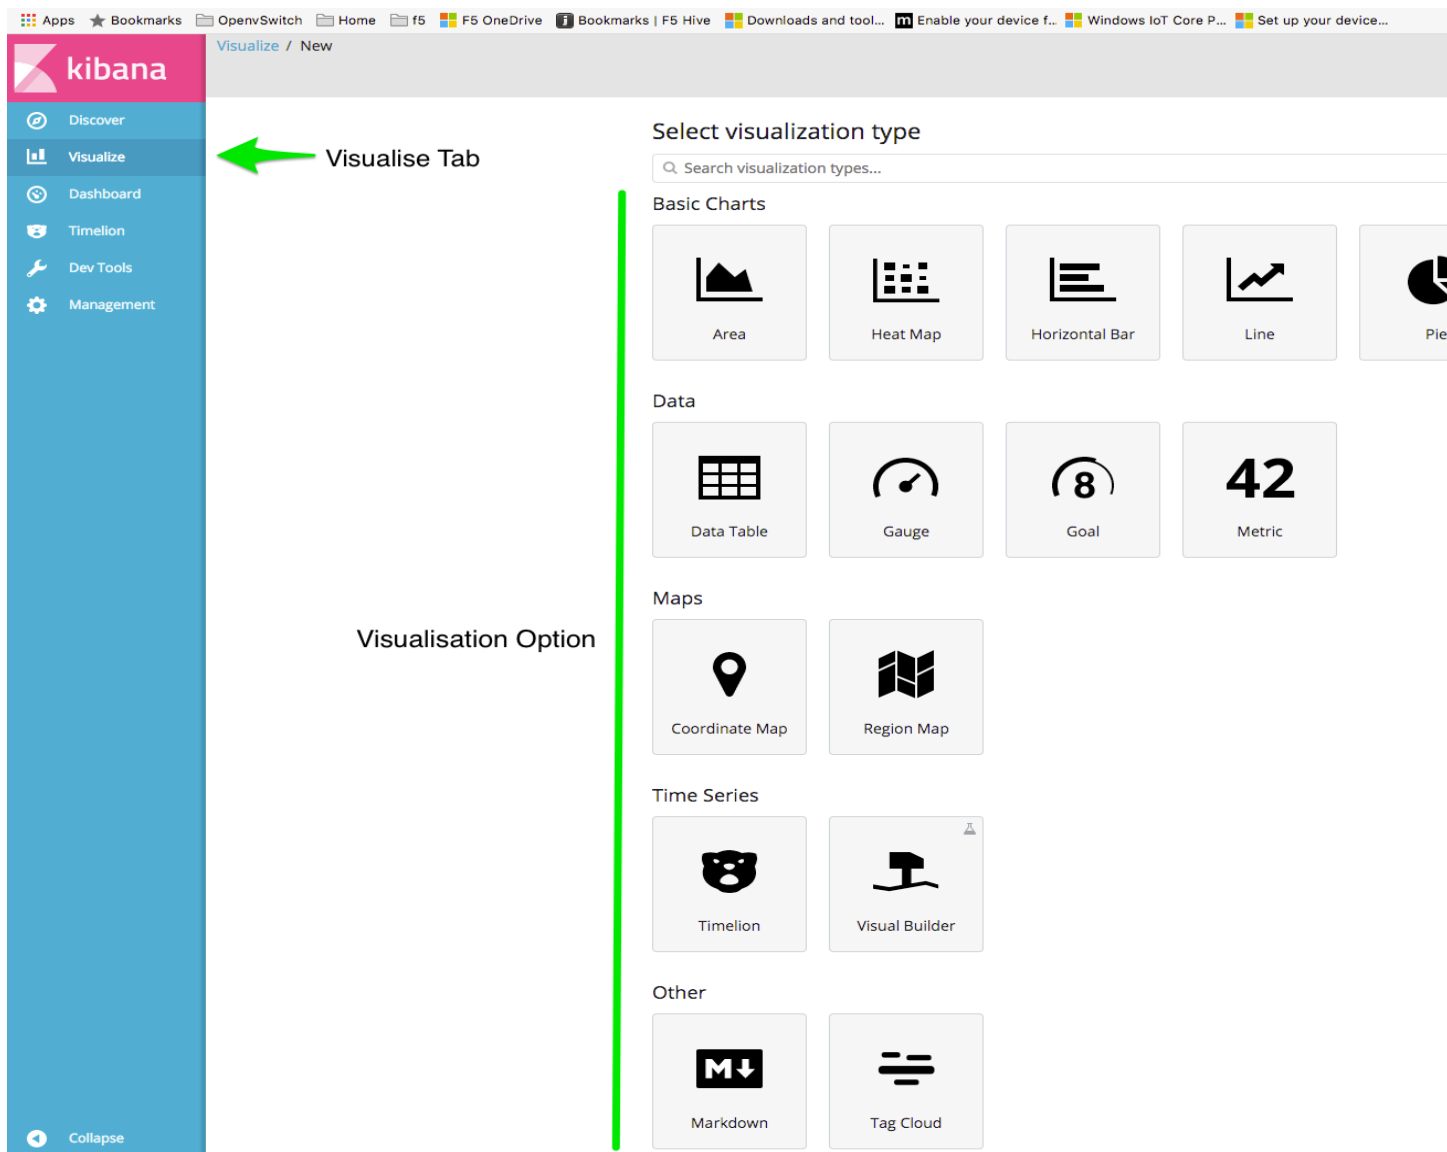
Accept

November 6th 2017, 15:15:07.456

Accept

Task 2 - Creating Visualisations

Create and Save a 3 x visualisation based on the above search or the previous lab.



Examine existing Visualisations to understand how some of the different visualisation are constructed.

Apps Bookmarks OpenvSwitch Home f5 F5 OneDrive Bookmarks | F5 Hive Downloads and tool... Enable your device f... Windows IoT C

kibana Visualize

Discover Visualize Dashboard Timelion Dev Tools Management


Examine existing visuals

Search...

- ☐ Name ▲
- ☐ AFM (Accept / Reject Rules)
- ☐ AFM Destination IP Map
- ☐ AFM VS Context
- ☐ Attacks per Destination table
- ☐ Attacks per Vserver table
- ☐ Average Data PEM
- ☐ Average Throughput
- ☐ BIG-IP Devices
- ☐ BIG-IP Processes
- ☐ Destination IP Top Ten
- ☐ Destination Port Top 10
- ☐ IP Protocol Traffic
- ☐ Logs over time
- ☐ Logs per Severity over time
- ☐ Program - Table
- ☐ Severity
- ☐ Severity - Pie
- ☐ Severity Bars
- ☐ Subscriber IP List
- ☐ Subscriber Names

Task 3 - Creating Dashboards

Create a dashboard from your 3 visualisations created above.



Search... (e.g. status:200 AND extension:PHP)

add saved visuals and searches



This dashboard is empty. Let's fill it

Click the **Add** button in the menu bar above to add a visualization

If you haven't set up any visualizations yet, [visit the Visualize app](#) to create

This section contains useful HOWTOs

7.1 HOWTO - how to do stuff

Twill put extra stuff into here

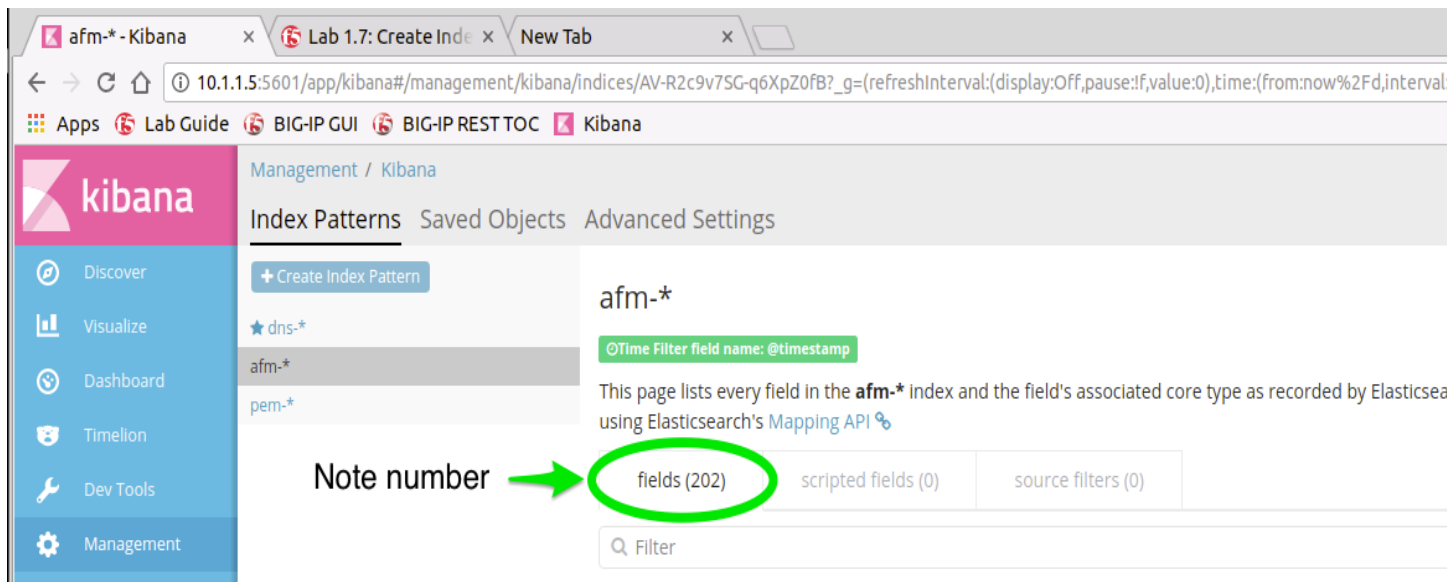
7.1.1 Task 1 – Update unknown index fields

At times new fields may appear in the index field based on software version or addiitonal logging from irules. It will be requiried to update the index to make these fields usable.

The screenshot shows the Kibana interface with the Discover page. The left sidebar has a 'Collapse' button at the bottom. The main panel displays a list of network events. The field list on the left includes various fields, with 'dest_port' highlighted by a green circle. The event details on the right show fields like Policy, Entity, Action, SaTranslationPool, DstUserName, SourceIpRouteDomain, SelfIp, ContextInfo, SrcCountry, Context, and Hostname.

Note: Notice the ? symbol next to the field.

Update by clicking on the refresh button



Note the increased change



7.1.2 Task 2 - Manual Index Changes

Index changes in json can be done manually if importing from another system.

1. Create a new search or visualisation
2. Export the new search json
3. Open the json and copy the index id
4. Open the json to be imported and paste the updated index id

WE MAKE APPS  FASTER.
SMARTER.
SAFER.

F5 Networks, Inc. | f5.com



US Headquarters: 401 Elliott Ave W, Seattle, WA 98119 | 888-882-4447 // Americas: info@f5.com // Asia-Pacific: apacinfo@f5.com // Europe/Middle East/Africa: emeainfo@f5.com // Japan: f5j-info@f5.com
©2017 F5 Networks, Inc. All rights reserved. F5, F5 Networks, and the F5 logo are trademarks of F5 Networks, Inc. in the U.S. and in certain other countries. Other F5 trademarks are identified at f5.com. Any other products, services, or company names referenced herein may be trademarks of their respective owners with no endorsement or affiliation, express or implied, claimed by F5. These training materials and documentation are F5 Confidential Information and are subject to the F5 Networks Reseller Agreement. You may not share these training materials and documentation with any third party without the express written permission of F5.