- [SUMIT PANDEY]

Module-1
Understand Splunk and its components

Splunk is log analysing and monitoring tool. It allows you to play with your data generated by your applications, network devices etc.

Index data → Search & Investigate → Add knowledge → Monitor, alert, report, visualize

Splunk is comprised of three main processing components –

1. Forwarder
2. Indexer
3. Search head

Indexer: Processes machine data and stores the results in indexes in the form of events. As the indexer indexes data, it creates several files organized in sets of directories by age.

Forwarder: It consumes and sends data to indexer. It resides on the machine where data originates.

Search Head: Allows users to search indexed data. Also, provide the facility of adding knowledge object on your search results and save them as an alert, report and dashboard.

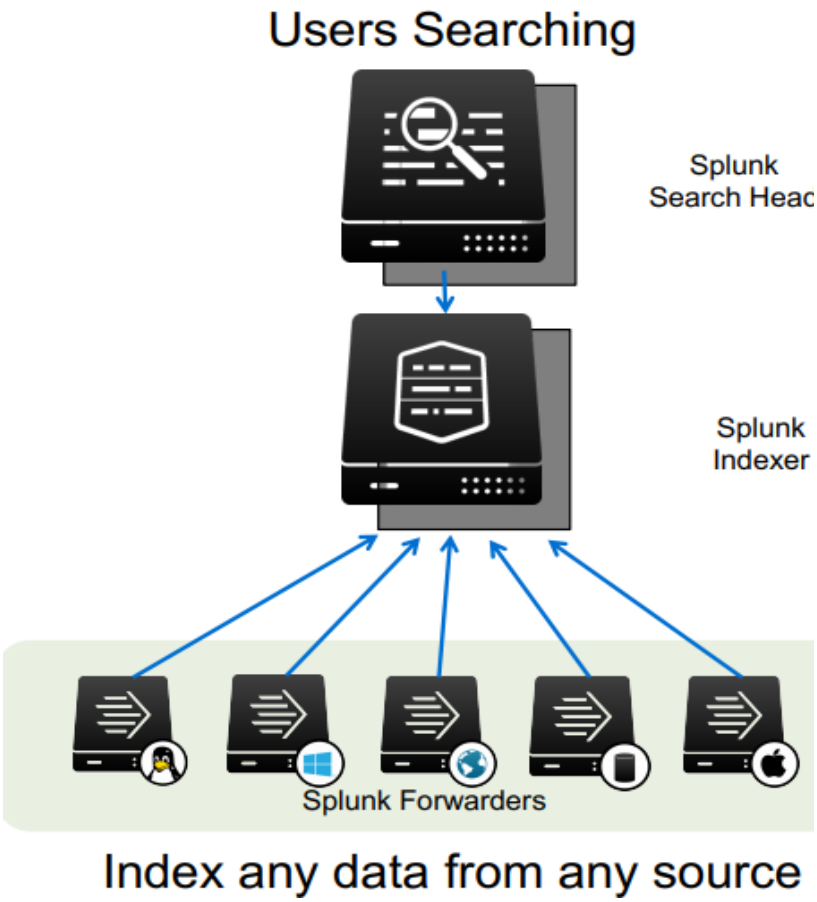

Additional Splunk components –

1. Deployment server
2. Cluster master
3. License master

Module-2
Installation of Splunk and Logging In

(a) On Windows:

Download the Splunk from Splunk's official website.
https://www.splunk.com/en_us/software/splunk-enterprise.html

Install it in the required directory as per your convenience, however, it's not advisable to install any software in C drive of the system as we have OS installed and running in C drive. And, failure or corruption of OS will lead to failures of installed softwares in the same directory.

Once Splunk is installed, make sure to check status of splunkd daemon service whether it's running or not. Splunkd service must be running to access Splunk from GUI.

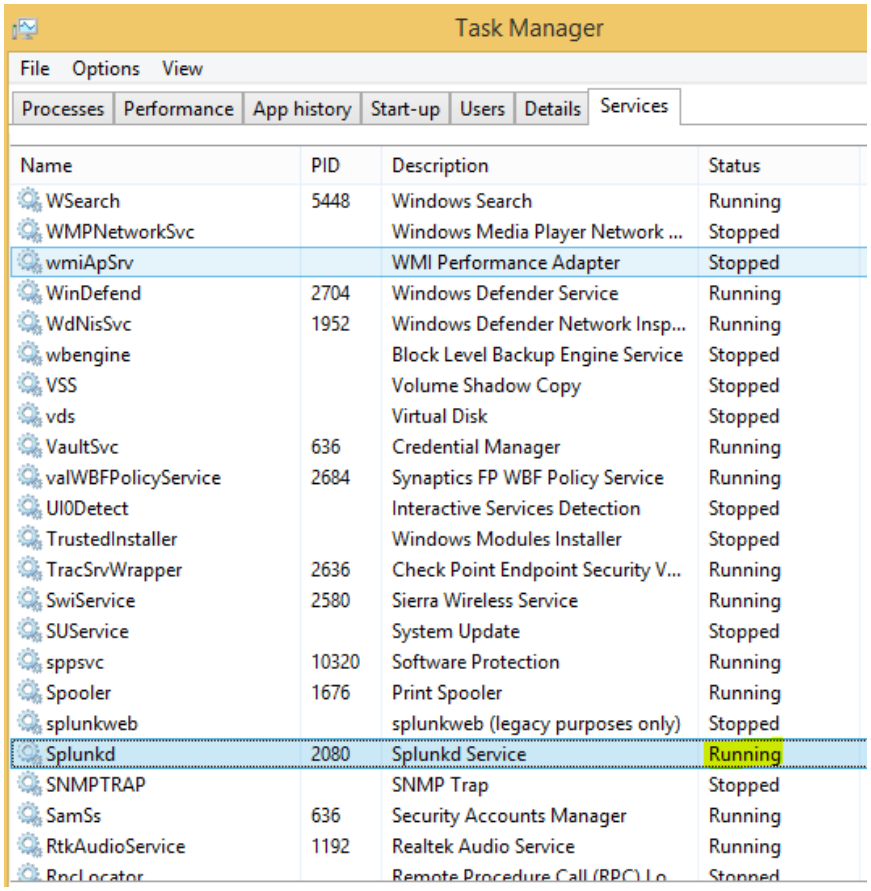

If Splunkd service is not up and running. Please follow the below steps to start the splunkd service.

1. Open command prompt
2. Go to bin folder of your Splunk installation directory.
   e.g. Cd C:\Program Files\Splunk\bin
3. Run the command
   splunk restart

NOTE: Splunk is the program in the bin directory to run the CLI.

COMMON SPLUNK COMMANDS

Command	Operation
splunk help	Display a usage summary
splunk [start][stop][restart]	Manages the Splunk processes
splunk start --accept license	Automatically accept the license without prompt
splunk status	Display the Splunk process status
splunk show splunkd-port	Show the port that the splunkd listens on
splunk show web-port	Show the port that the Splunk web listens on
splunk show servername	Sow the server name of this instance
splunk show default-hostname	Show the default host name used for all data inputs

splunk enable boot-start -user	Initialize script to run Splunk Enterprise as system start-up

Logging in –

Access the Splunk GUI using below URL:
<server-IP>:8000

e.g. 10.122.20.23:8000 OR localhost:8000
The web login page looks like below.

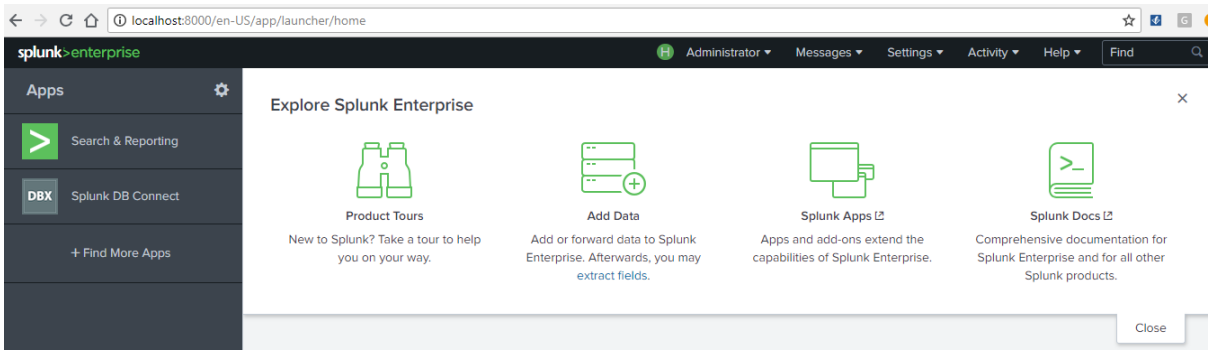

Congrats! Your Splunk set-up is done. Now, you may explore it's all available options.

To run the SPL (Splunk programming language) – Visit 'Search & reporting' button. And the view is like –

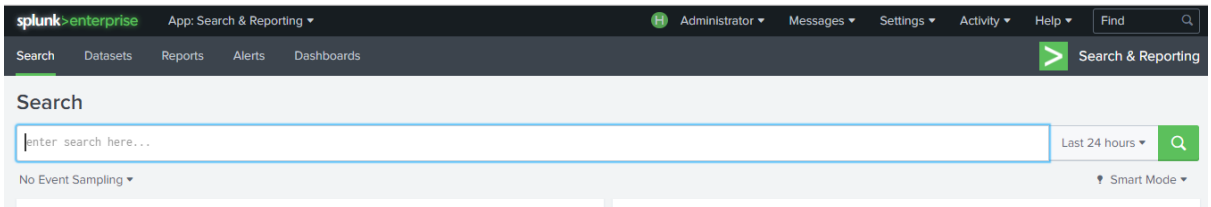

(b) <u>On Linux:</u>

(i) Download .tgz file and cancel when download begins.
(ii) Copy WGET command and execute at the Linux server.
Sudo <paste command here>
(iii) Once Splunk is copied on the server, copy it to opt folder by using below command.
Sudo cp splunk-7.11… /opt
(iv) Now Splunk file is successfully moved to /opt directory. Please go ahead, unzip and install.
sudo tar -xvzf splunk-7.11…
(v) Go to bin folder and execute –
Cd /splunk/bin
sudo ./splunk start –accept-license **OR** sudo ./splunk start --accept-license -yes

(To install from .rpm, please use rpm -i splunk…. And from .deb, use dpkg -i splunk….)

P.S: Providing sudo explicitly not required in each command if you're logged in as a root (sudo) user.

Also, make sure relevant ports are open for the instance. (TCP-8000, TCP-8089, TCP-9997 etc.)

Module-3
Installation of Splunk Forwarder

(a) On Windows:

Download Splunk universal forwarder from Splunk's official website and install it. We need to install the forwarder on the server wherein data originates.
Download from the link –
https://www.splunk.com/en_us/download/universal-forwarder.html

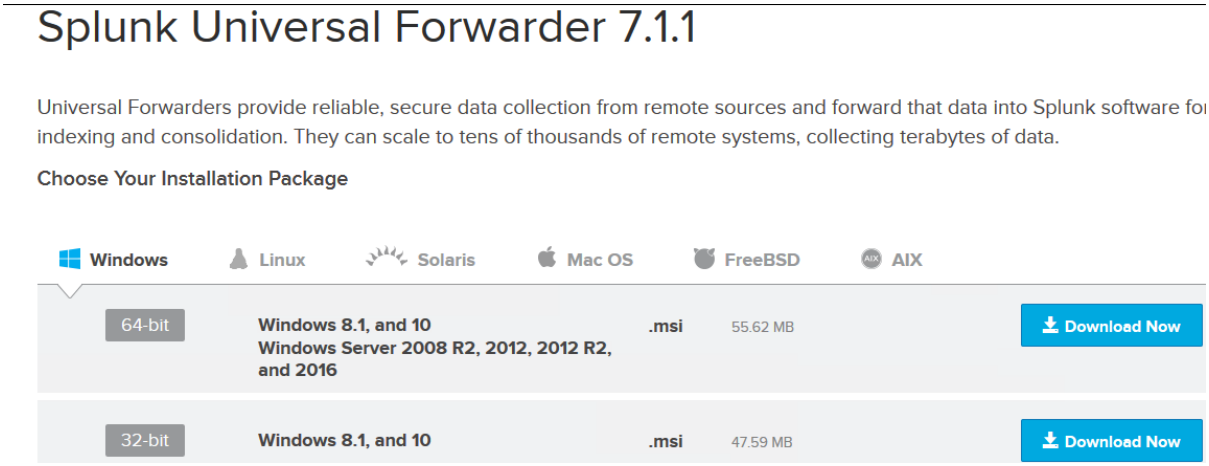

While installing universal forwarder, please mention your indexer server IP and port (the machine wherein you have installed your Splunk enterprise), however, this can be done later as well from CLI or editing outputs.conf at forwarder end as mentioned below -

(i) Splunk add forward-server 10.122.20.23:9997, list down if it's added –
Splunk list forward-server

OR

(ii) Go to /etc/system/local and mention in outputs.conf

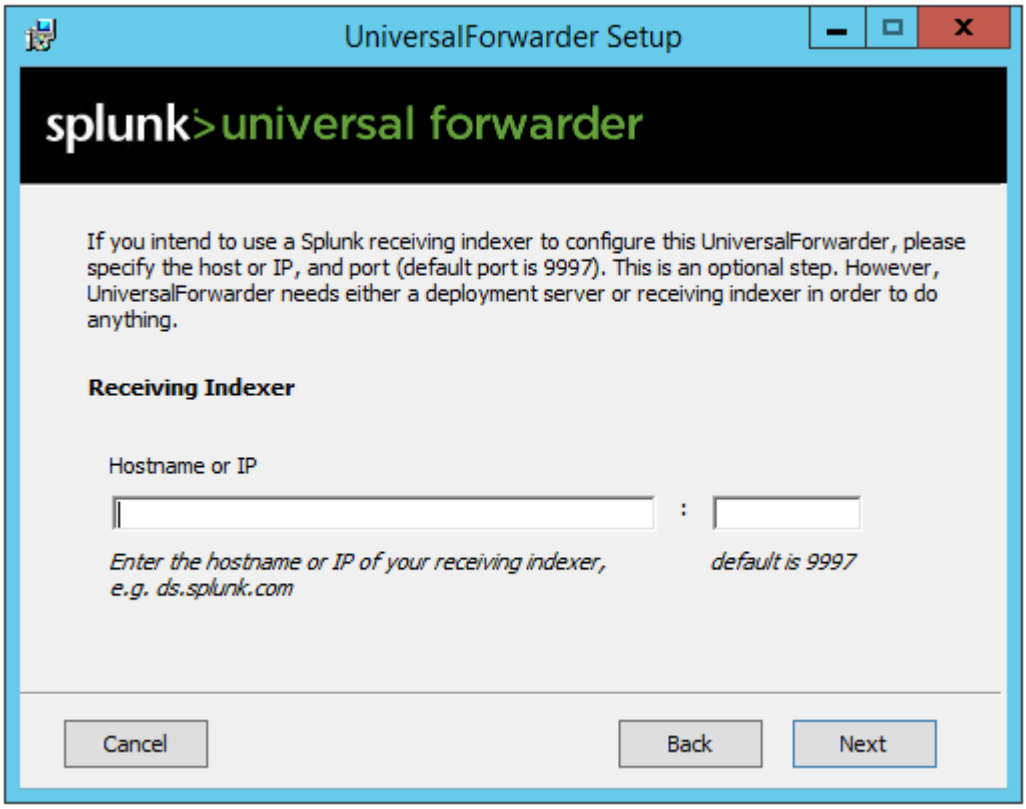

And click next. Your installation is done.
Before pushing the data from forwarder to indexer, you need to enable receiving at indexer side. That you can do in two ways –

1. Define receiving port from GUI. (Forwarder and receiving)

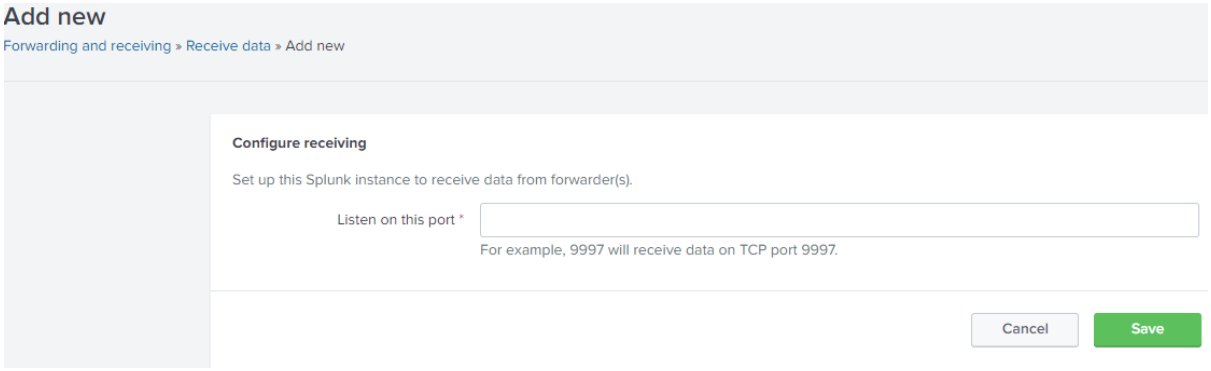

2. Make the change in inputs.conf at indexer to listen data over TCP port
   [splunktcp://9997]
   disabled = 0

   Now your forwarder set-up is ready. You may go ahead and configure inputs for monitoring shown as below:
   [default]
   host =
   index = test

   [monitor://path]

sourcetype =
disabled = false

OR

[default]
host =

[monitor://path]
sourcetype =
index=
disabled = false

NOTE: Please make a note that splunkd daemon service should be restarted whenever any change is made in any of the configurations file in order to take the implementation in place.

(b) On Linux: You may follow similar steps as Splunk installation on Linux explained above by downloading .tgz file or you can also try to install using .deb package as below –

(i) Download .deb file and cancel when download begins.
(ii) Copy WGET command and execute at the Linux server.
<paste command here>
(iii) Login as a root user - sudo su
(iv) dpkg -i splunkforwarder….
(v) It's installed. Check the cd /opt directory
(vi) Start it by going to bin folder
./splunk start
(vii) Accept license by selecting 'Y'

Module-4
License Management

There are few types of licenses in Splunk.

1. **Enterprise trial** – Download and install with product. Valid for 60 days and post that converts into Free License.
2. **Free License** – Disables a bunch of features like alerts, reports, clustering etc.
3. **Forwarder License** – allows Splunk instance to be installed as Heavy Forwarder.
4. **Splunk enterprise license** – Purchased from a Splunk sales representative. Daily volume depends on how much you pay.

What happens if you breach daily indexing quota: 5 Warnings in 30 days. Earlier Splunk used to disable searching but with new versions that has been stopped. So, you can still enjoy with your data while breaching daily indexing limit.
[Also, Splunk will never stop indexing of your data though you're violating the limit]
[Splunk Licensing - Buy a license from Splunk.com and upload it going to settings – licensing]

Module-5
Splunk Apps

Splunk apps are collection of configurations file. Among all the Splunk apps, Splunk DB connect is widely used app.

1. **DB connect:** To connect with the database.

   (i) Install it and create a user SplnunDBConect in the database.

   CREATE USER 'splunkDBConn'@'localhost' IDENTIFIED BY 'splunkdbpass';

   GRANT ALL ON *.* to 'splunkDBConn'@'localhost' IDENTIFIED BY 'splunkdbpass'
   GRANT ALL ON *.* to 'splunkDBConn'@'milsplunk01%' IDENTIFIED BY 'splunkdbpass'

   (ii) Modify settings:
   Add **driver** (Splunk to C:\Program Files\Splunk\etc\apps\splunk_app_db_connect\drivers [MySQL connector Java driver]

   (iii) JVM Options: -Ddw.server.applicationConnectors[0].port=9998
   Check if **dbx_task_server.yml** file is available at C:\Program Files\Splunk\etc\apps\splunk_app_db_connect\default

   (iv) Restart Splunk and try to connect.

2. **PagerDuty:** To connect with incident management tool.

3. **Splunk AWS app:** Install Splunk AWS app, Splunk add on AWS and Python for commuting 64 bit.
   Create one policy for all naming something like Splunk_AWS and role too using same policy.
   (Configure one policy containing permissions for all inputs)
   Now, assign this policy to EC2 instance and reboot the instance.
   Now you should be able to see data on AWS Splunk app.

Note – If any error is there on the app then increase the disk of the instance to minimum 100 GB.

NOTE: An app can also be installed from a file. Download an app from splunkbase.com and then install it from Splunk GUI.
e.g.: Splunk Dashboard Example

Module-6

Splunk Configuration Files

Splunk is collection of configuration files. Configuration files govern how Splunk works.

Configuration file structure

[Stanza header]
Attributes
(key=value)

Example – outputs.conf
[tcpout:splunk_indexer]
Server = 10.122.20.23 9997

Splunk ***configuration files*** in the default directories come with Splunk and are sample configuration files.
Note - Modifications should be made in the local directory. Don't change the configuration files in the default directory.

Configuration files Structure

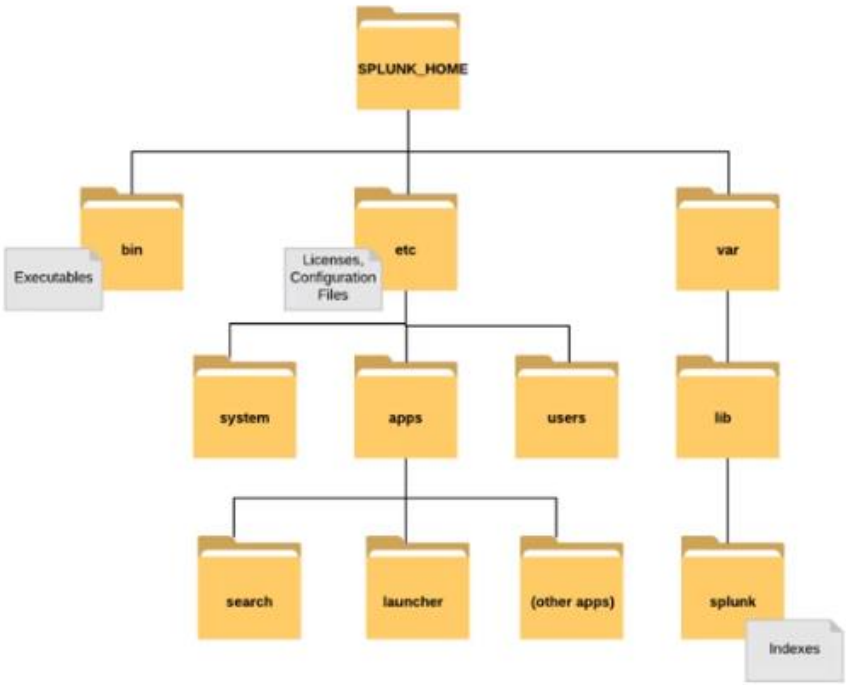

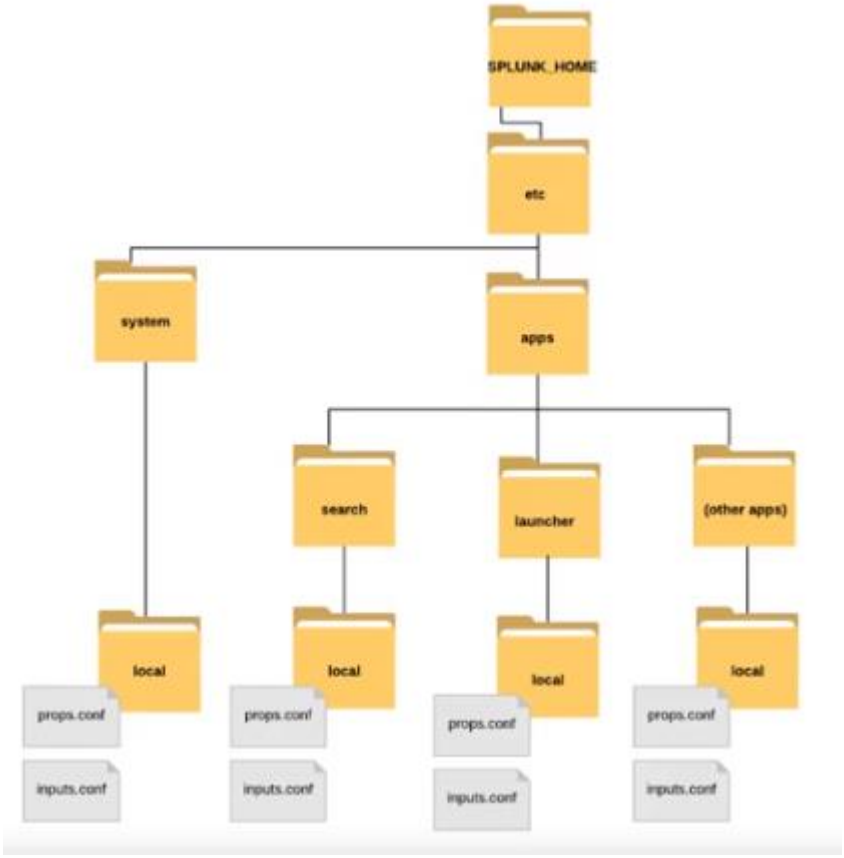

Configuration precedence in Splunk –

1. System /local directory
2. App /local directory
3. App /default directory
4. System /default directory

Here are important configuration files –

1. Inputs.conf – Defines data inputs
2. Outputs.conf – Governs forwarding behaviour
3. Props.conf – Indexer configurations, source type rules & more
4. Limits.conf – Defines limits for search commands

Below is how props.conf file looks like –

```
[index::bn]
TZ = Asia/Dhaka
[MIP_Test_Sourcetype]
TIME_PREFIX = \[
TIME_FORMAT = %Y-%d-%m %H:%M:%S
SHOULD_LINEMERGE=false
```

Module-7

Splunk Index Management

Splunk indexer process the incoming machine data and stores in indexes. Indexes are nothing, they're just directories in your Splunk installation path. Indexes helps in faster searching and access restrictions.

If you do not define any index in inputs.conf and create one manually, the data will be ingested in main index by default.

To create a new index –

Settings -> Indexes -> New Index

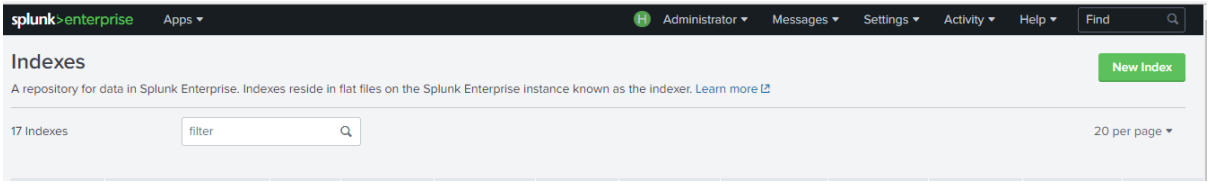

Also, you can add/delete any index directly by modifying indexes.conf file located at **C:\Program Files\Splunk\etc\apps\search\local**

These newly created indexes can be found at below location in your Splunk installation directory.

C:\Program Files\Splunk\var\lib\splunk

And raw data can be found under JOURNAL folder at below path:

C:\Program Files\Splunk\var\lib\splunk\<**index name**>\db\db_1530199890_1530199890_0\rawdata

How to delete data of any index:

To delete indexed data of any index, you must have "can_delete" roles permission. Write your SPL and then put PIPE symbol and give delete command to delete that indexed data.

Index=test earliest=-1d@d latest=@d|delete

Use case: Requires most of the time when your application /product goes out in live out of pilot or test.

Module-8
Splunk Users, Roles & Authentication

Before creating a new user in Splunk, you may set-up a different role based on the accesses new role requires because Splunk users are assigned roles. Roles determine capabilities and data access.

There are three main roles –

- Admin
- Power
- User

Splunk administrators can create additional roles.

1. Go to **Settings** – **Access control** – **Roles**. Create **New** –

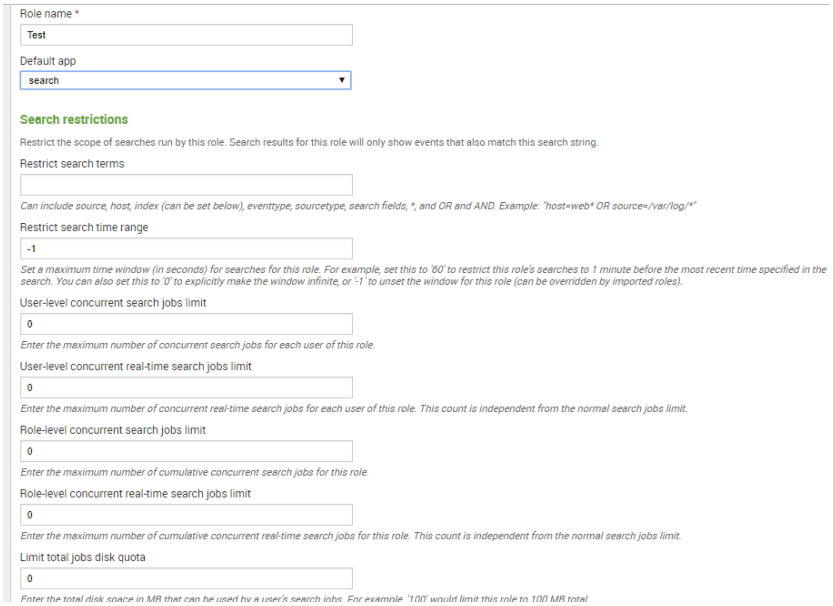

2. Assign roles, capabilities and indexes from the drop-down panels and click save.

3. Go back and create a new user and assign above created role to this user.

Module-9
Heavy Forwarder

Heavy forwarders mainly used in larger scale Splunk environments. It's capable enough of doing parsing and filtering of the data before forwarding it to Splunk indexer. In general, there is no separate application for heavy forwarder. Spunk enterprise software can be used as a heavy forwarder. In similar way as UF, you can add receiver indexer IP in outputs.conf or you may enable it from GUI itself under "forwarding & receiving". Once done, add an entry in inputs.conf and it will start forwarding the data to Indexer server.
Install the Splunk on the machine of which you want to forward the logs to Splunk indexer. Once installed and if you want to change the license group, please follow as -

1. Go to Settings – Licensing

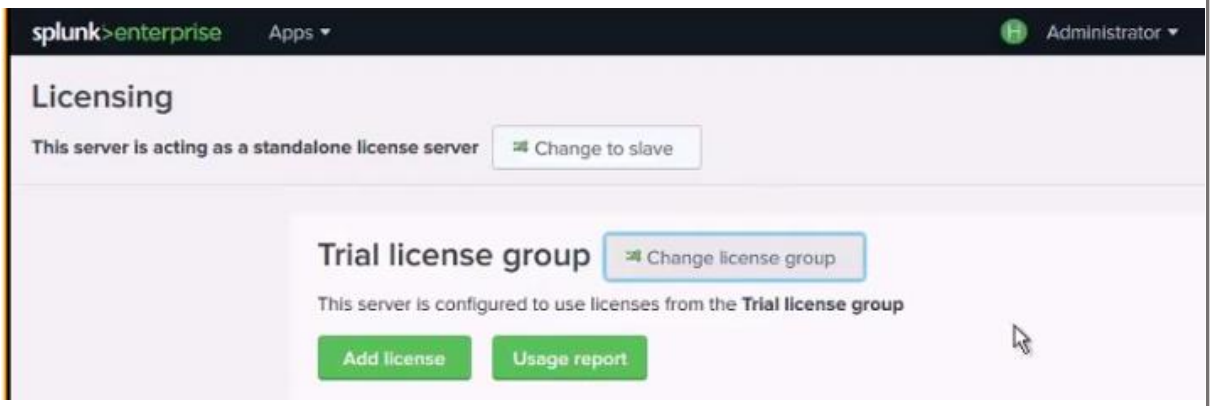

2. Select “Change license group”

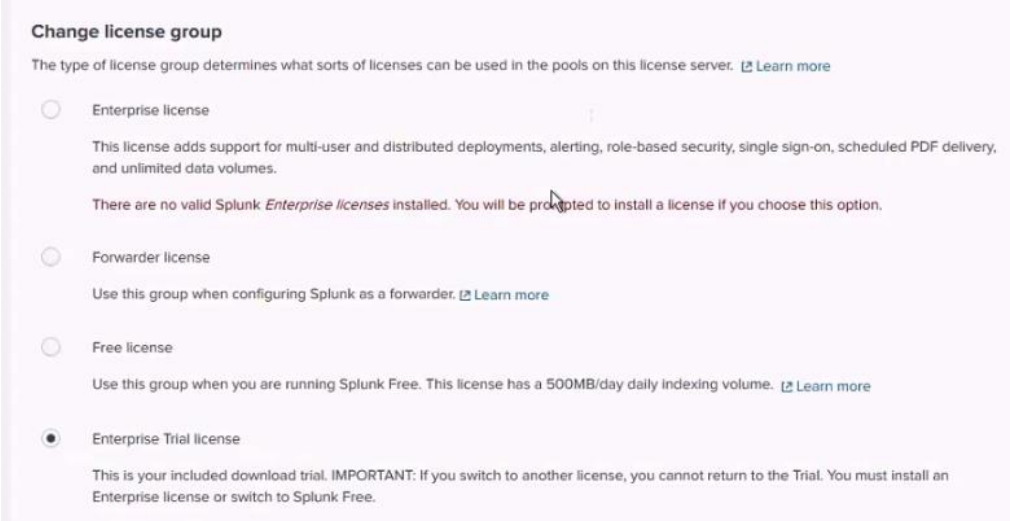

3. Restart this instance of Splunk to act as a forwarder.

Main difference between UF and HF –

Universal forwarder	Heavy forwarder
No GUI	GUI
No parsing feature	Parsing feature
Capable of sending unparsed data with little tagging – source, source-type, index	Capable enough to filter the data before forwarding
Event level viewing not possible	Event level viewing possible

Module-10
Forwarder Management

You can manage all your forwarder clients from Splunk GUI. Run below commands on Splunk UF server from CLI –

splunk set deploy-poll 10.122.20.71:8089

splunk restart

To list down forwarders –
Splunk list forward-server

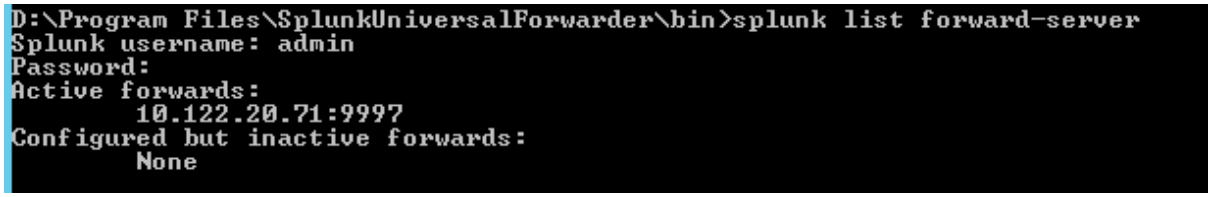

Module-11
Windows Agentless Inputs

Monitor data through "Windows Management Instrumentation". WMI provides access of remote machine's event logs and security data.
To access remote data through WMI, Splunk must be installed using **domain account**. And, save below file in local directory as wmi.conf

[WMI:TailApplicationLogs]
interval = 10
event_log_file = Application, Security, System
server = 10.122.20.28
disabled = 0
current_only = 1
batch_size = 10

Module-12
Upgrade Splunk – Windows

Upgrade Splunk Enterprise using the GUI installer – It's recommended to take backup of indexes and configuration files before running upgrade on Spunk.

a. Download the new MSI file from the Splunk download page.
b. Double-click the MSI file. The installer runs and attempts to detect the existing version of Splunk Enterprise installed on the machine. When it locates the older version, it displays a pane that asks you to accept the licensing agreement.
c. Accept the license agreement. The installer then installs the updated Splunk Enterprise. This method of upgrade retains all parameters from the existing installation. By default, the installer restarts Splunk Enterprise when the upgrade completes and places a log of the changes made to configuration files during the upgrade in %TEMP%

Module-13
Splunk Troubleshooting

CASE-1: Splunk not sending any email for scheduled report and alerts
Solution – Check in scheduler.log for an error being received while Splunk tried to trigger an email. Look for the reason.

savedsearch_name="BIMA SMS In Queue Piling Up - PD", priority=default, status=skipped, reason="The maximum number of concurrent running jobs for this historical scheduled search on this instance has been reached"

Try to check CPU usage / concurrent scheduled (running) searches. If no issues with those then go ahead and restart the splunkd.

Module-14
Splunk Queries (SPL)

Append – It adds new rows for sub-query along with fields if fieldnames are different from master query.

index=main sourcetype="mip_deduction_success" deduction_batch="0"|stats sum(amount) by offer_id|append[search index=main sourcetype="mip_deduction_success" deduction_batch="0"|stats sum(amount) by offer_cover_id]

You can rename field name like the field in master query to avoid adding new fields.

index=main sourcetype="mip_deduction_success" deduction_batch="0"|stats sum(amount) by offer_id|append[search index=main sourcetype="mip_deduction_success" deduction_batch="0"|stats sum(amount) by offer_cover_id|eval offer_id=offer_cover_id]

another example –

index=main sourcetype="mip_deduction_success" deduction_batch="0"|stats sum(amount) by offer_id|append[search index=main sourcetype="mip_deduction_success" deduction_batch="0"|stats count(msisdn) by offer_id]

Appendpipe – When you're grouping results e.g. count /sum by two fields and want summary of the master query result set. To append the search results of post process (subpipeline) of the current result set.

deduction_batch	offer_id	Total
0	2	9734.1
0	3	10890.2
0	4	7546.8
0	5	104.8
0	6	635.8
1	2	28502.4
1	3	27600.9
1	4	17163.9
1	5	405.7
1	6	1750.6
2	2	2788.2
2	3	2864.8
2	4	2624.9
2	5	46.2
2	6	211.6
3	2	2752.2
3	3	2760.1
3	4	2846.9
3	5	6.8
3	6	48.6

deduction_batch	offer_id	Total
0	2	9734.1
0	3	10890.2
0	4	7546.8
0	5	104.8
0	6	635.8
0	Total amount by this deduction_batch	28911.7
1	2	28502.4
1	3	27600.9
1	4	17163.9
1	5	405.7
1	6	1750.6
1	Total amount by this deduction_batch	75423.5
2	2	2788.2
2	3	2864.8
2	4	2624.9
2	5	46.2
2	6	211.6
2	Total amount by this deduction_batch	8535.7
3	2	2752.2
3	3	2760.1
3	4	2846.9
3	5	6.8
3	6	48.6
3	Total amount by this deduction_batch	8414.6

index=main sourcetype="mip_deduction_success" |stats sum(amount) as Total by deduction_batch,offer_id|appendpipe[stats sum(Total) as Total by deduction_batch|eval offer_id="Total amount by this deduction_batch"]|sort deduction_batch

Appendcols – It adds new columns with the output of sub-search. 1st row of first search will be merged with 1st row of second search. The number of outputs can be different for both the queries.

index=main sourcetype="mip_deduction_success" deduction_batch="0"|stats sum(amount) by offer_id|appendcols[search index=main sourcetype="mip_deduction_success" deduction_batch="0"|stats count(msisdn) by offer_cover_id]

Difference between appendcols and joins –

In appendcols it's not mandatory that there should be any matching column between outer search and inner search but join use to return matching result set.

Joins – Splunk supports two types of joins and they're inner and left /outer join. Inner join returns matching data from both the result sets while left /outer join returns all data from first data set and matching data from second data set. Important to note, in Splunk join, field name on which join is used should be same in both the data sets.

Inner Join – Will return matching records from both the data sets.

index=tz sourcetype=mip_deduct_charges offer_id|table msisdn,offer_id|join type=inner msisdn[search index=tz sourcetype=mip_deduction_success|table msisdn,amount]

Left /Outer Join – All data from 1st search while matching from second. Both left and outer joins are same.

index=tz sourcetype=mip_deduct_charges offer_id|table msisdn,offer_id|join type=left msisdn[search index=tz sourcetype=mip_deduction_success|table msisdn,amount]

Self-Join – Splunk self-join is completely different from SQL self-join. It works on same column and matches the previous row value with next row value. By default, return second result. You can use overwrite=FALSE to return first result. It's useful in cases wherein your process passes through different stage and you want to check only first or last stage. In this case process ID will remain same so you can use self-join on process id. Or, if you have duplicates in your result set and want first entry of the field then use self-join.
index=tz sourcetype=mip_deduct_charges source="D:\\MIP\\Logs\\Schedulers\\DeductionBatches\\AutoDeductions\\all_transactions.log" 0652001316 OR 0652002749 OR 0652003044 OR 0652006898 |table _time, msisdn, offer_id, pending_amount_attempted,is_balance_above_threshold|sort msisdn|selfjoin msisdn

index=tz sourcetype=mip_deduct_charges source="D:\\MIP\\Logs\\Schedulers\\DeductionBatches\\AutoDeductions\\all_transactions.log" 0652001316 OR 0652002749 OR 0652003044 OR 0652006898 |table _time, msisdn, offer_id, pending_amount_attempted,is_balance_above_threshold|sort msisdn|selfjoin overwrite=FALSE msisdn

_time	msisdn	offer_id	pending_amount_attempted	is_balance_above_threshold
2019-06-18 21:54:28.382	0652002749	3,4,5	5500	false
2019-06-18 18:53:05.673	0652003044	3,4,5	4000	false
2019-06-18 17:13:17.293	0652006898	3,4,5	4000	false

Overwrite with Joins – By default the field value is set to TRUE. If you do not want to overwrite the values of common field between both the data sets then use overwrite=FALSE.

So, wherever the amount field value is available in second search, it will overwrite the value from first search which is set to 0 for all the records. For, non-existing cx in second data set it will assign 0.
index=tz sourcetype=mip_deduct_charges offer_id|eval amount=0|table msisdn,offer_id,amount|join type=left msisdn[search index=tz sourcetype=mip_deduction_success|table msisdn,amount]

Module-15
Splunk Distributed Search

This is basically a separation of searching layer from the indexing layer. This can be done in a 3 different way –

- Splunk Web
- Splunk CLI
- The distsearch.conf configuration file

Please note you can still login to your indexer Splunk instance to get the data.

NOTE: Before adding a search peer, make sure distributed search is ON in Search head. Splunk search head is nothing but a full fledged Splunk installation.

Through Splunk Web:

Once installed, go to “settings” -> “Distributed search” –> “Search peers”. Add a new search peer by providing below info-

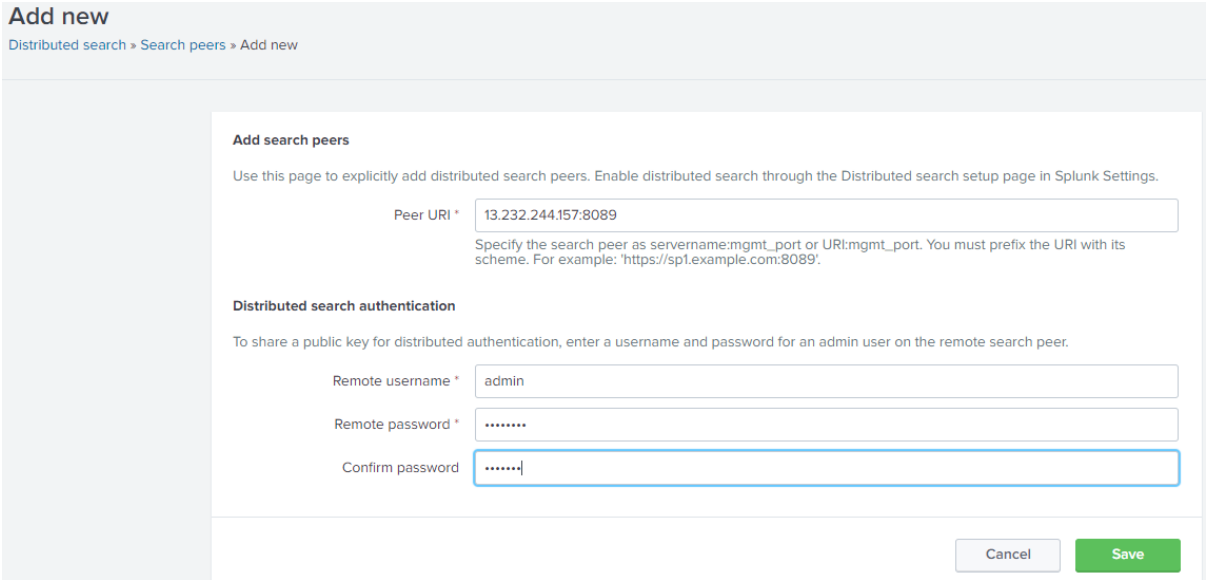

Through Splunk CLI:

You can add a search peer using CMI as well. Login to Search head server and go to the bin directory of Splunk installation. Execute below query –

./splunk add search-server X.X.X.X:8089 -auth admin:password -remoteUsername admin -remotePassword password

Once peer instance is added, it will be visible like below –

Peer URI	Splunk instance name	State	Replication status	Cluster label	Health status	Health check failures	Status	Actions
13.232.244.157:8089	ip-172-31-14-92.ap-south-1.compute.internal	Up	Successful	None	Healthy	None	Enabled | Disable	Quarantine | Delete

That's it. You have done it correctly. Go ahead and start searching the data in your search head.

Module-16:
Filter And Route Event Using HF

The events levels filtering, and routing can be easily gain using heavy forwarders.

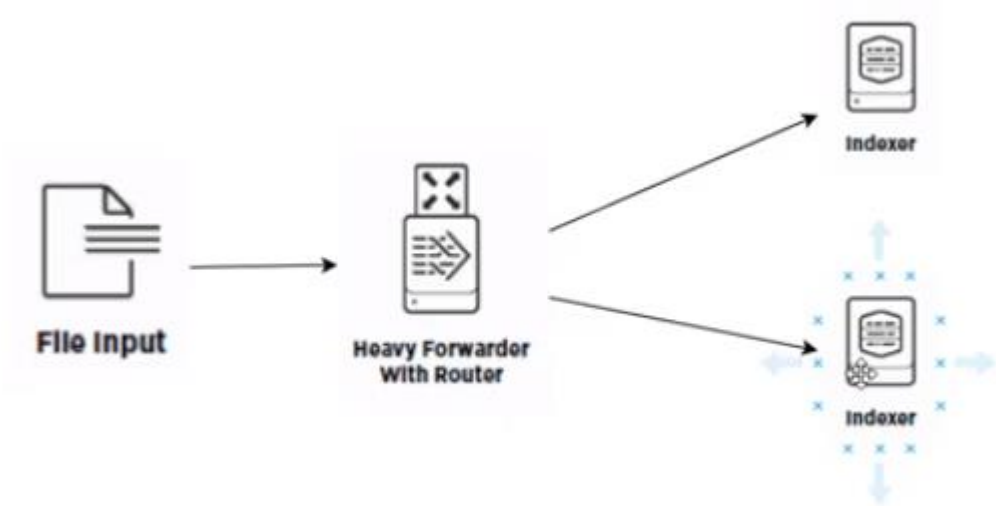

Module-17:
Deployment Server

To manage all your Splunk componenets you may use Splunk Deployment Server. It's very useful in deploying changes to the other Splunk components.You can create serverclass groups to deploy changes to set of identical Splunk components like only to Indexers, or search heads or forwarders, etc.

Go to your Splunk components that want to report to DS and execute below query from CLI. It can be done by adding deploymentsclient.conf file as well.

Set deploy-poll <IP of DS>:8089

Now, you can see all Splunk components reporting to your DS. You can manage all from DS only.

IMPORTANT:: To change hostname in Linux – **vi /etc/hostname** and put your convenient /human readable name for the server. Reboot post changes. See below for one of the server it's changed to IDX1

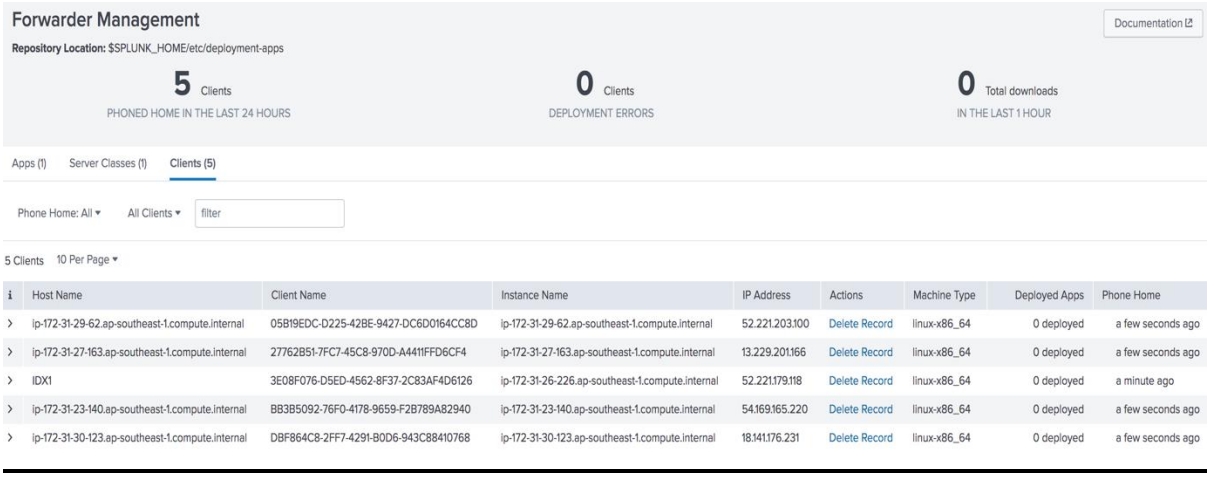

Module-18:
Indexer Clustering

This is to set-up single site indexer clustering. The two important terms in clustering are –
Replication Factor – Total copy of data
Search Factor – Searchable copy of data
All the Splunk instances should be running on separate machines in clustering. And, required n/w ports (22 SSH, 8000 Splunk web, 8080 or something else replication port, 9997 indexing port, 8089 mgmt port) should be open.

Cluster Master – it's to manage clustering.

Enable Master Node –

```
[clustering]
mode = master
replication_factor = 4
search_factor = 3
pass4SymmKey = whatever
cluster_label = cluster1
```

This example specifies that:

- the instance is a cluster master node.
- the cluster's replication factor is 4.
- the cluster's search factor is 3.
- the security key is "whatever". All nodes in the cluster use the same security key. It's an optional parameter.
- the cluster label is "cluster1." The optional cluster label is useful for identifying the cluster in the monitoring console. You set this attribute on the master node only.

Restart splunk

Enable peer nodes –

```
[replication_port://9887]

[clustering]
master_uri = https://<IP of CM>:8089
mode = slave
pass4SymmKey = whatever
```

This example specifies that:

- the peer will use port 9887 to listen for replicated data streamed from the other peers. You can specify any available, unused port as the replication port. Do not re-use the management or receiving ports.
- the peer's cluster master resides at `10.152.31.202:8089`.
- the instance is a cluster peer ("slave") node.
- the security key is "whatever". It requires only if you have defined it for master node.

Restart splunk on all peer nodes.

That's good job. You're indexer clustering is set-up. Login to cluster master and see. It should show something like this.

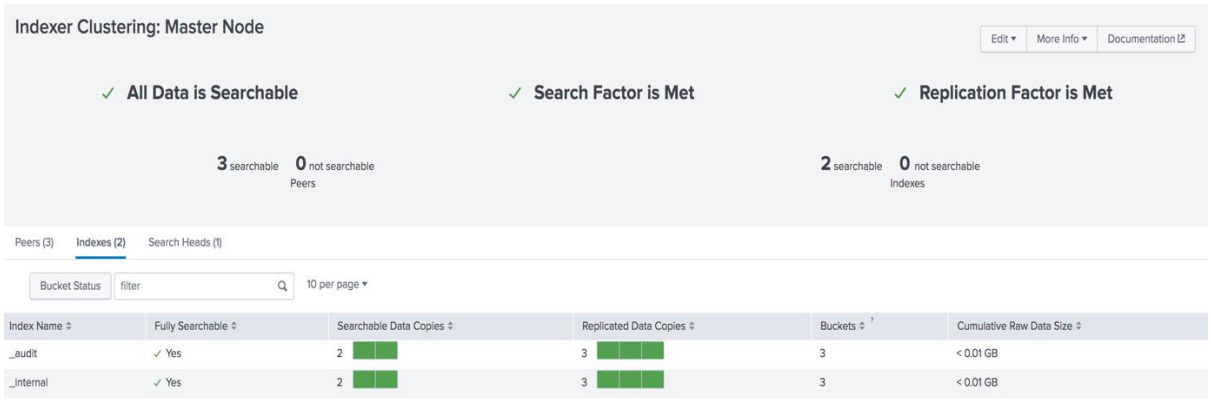

You can also check the indexer cluster status using CLI on CM server –

[root@ip-172-31-16-154 local]# /opt/splunk/bin/./splunk show cluster-bundle-status
Splunk username: admin
Password:

master
cluster_status=None
active_bundle
checksum=0DBF2F00FDA6661C3026CCA3D613D17C
timestamp=1589431875 (in localtime=Thu May 14 04:51:15 2020)
latest_bundle
checksum=0DBF2F00FDA6661C3026CCA3D613D17C
timestamp=1589431875 (in localtime=Thu May 14 04:51:15 2020)

last_validated_bundle
checksum=0DBF2F00FDA6661C3026CCA3D613D17C
last_validation_succeeded=1
timestamp=1589431875 (in localtime=Thu May 14 04:51:15 2020)
last_check_restart_bundle
last_check_restart_result=restart not required
checksum=
timestamp=0 (in localtime=Thu Jan 1 00:00:00 1970)

ip-172-31-29-62.ap-southeast-1.compute.internal 05B19EDC-D225-42BE-9427-DC6D0164CC8D default
active_bundle=0DBF2F00FDA6661C3026CCA3D613D17C
latest_bundle=0DBF2F00FDA6661C3026CCA3D613D17C
last_validated_bundle=0DBF2F00FDA6661C3026CCA3D613D17C
last_bundle_validation_status=success
restart_required_apply_bundle=0
status=Up

ip-172-31-26-226.ap-southeast-1.compute.internal 3E08F076-D5ED-4562-8F37-2C83AF4D6126 default
active_bundle=0DBF2F00FDA6661C3026CCA3D613D17C
latest_bundle=0DBF2F00FDA6661C3026CCA3D613D17C
last_validated_bundle=0DBF2F00FDA6661C3026CCA3D613D17C
last_bundle_validation_status=success
restart_required_apply_bundle=0
status=Up

ip-172-31-23-140.ap-southeast-1.compute.internal BB3B5092-76F0-4178-9659-F2B789A82940 default
active_bundle=0DBF2F00FDA6661C3026CCA3D613D17C
latest_bundle=0DBF2F00FDA6661C3026CCA3D613D17C
last_validated_bundle=0DBF2F00FDA6661C3026CCA3D613D17C
last_bundle_validation_status=success
restart_required_apply_bundle=0
status=Up

** Search for index="_internal" in CM and it will turn up the data from all peer indexers so it means replication is working properly.

Set-up Search Head from CM –

To search the cluster, you need to enable at least one search head in the indexer cluster. Add below stanza and params in server.conf on all Search Heads.

```
[clustering]
master_uri = https://10.152.31.202:8089
mode = searchhead
pass4SymmKey = whatever
```

This example specifies that:

- the search head's cluster master resides at `10.152.31.202:8089`.
- the instance is a cluster search head.
- the security key is "whatever". This is an optional and require only if you passed it any key while setting up Cluster master.

Super job. You search head is also configured. Go ahead and start searching, it will show the data from all peer nodes.

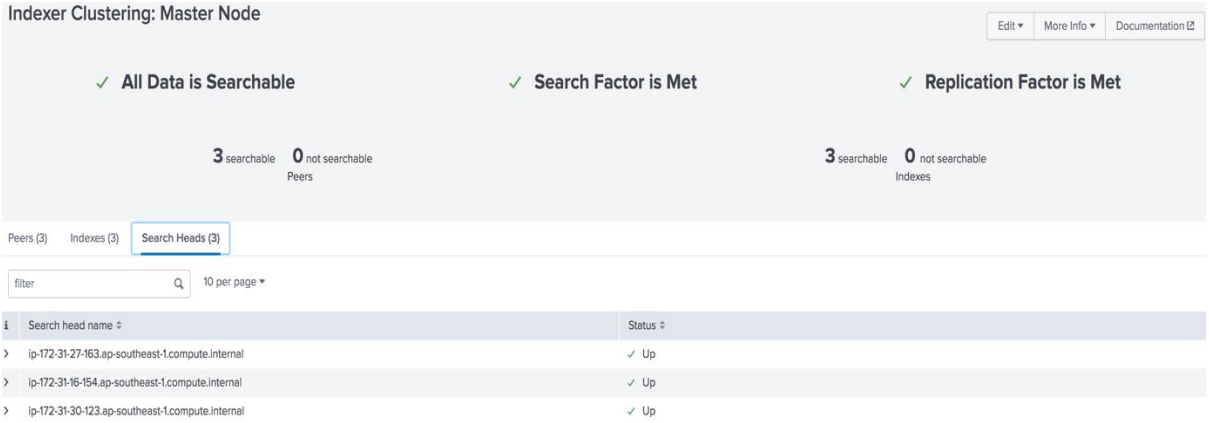

Note – If you're doing demo on free versios and come across disk space issue, then you can reduce the threshold on server.conf by addming below param. By default it's 5000 MB.

[diskUsage]
minFreeSpace = 500

Index Replication: By default indexer cluster replicates the data only for audit and internal indexes because repFactor is set to 0 in **indexes.conf** under default directory. Hence, in order to work it for all the indexes you want, you will have to set repFactor to auto on all the indexer peer nodes. Make the changes in local directory and not under default.

/opt/splunk/etc/system/local

repFactor=auto

Restart splunk

Also, it's very important to have **identical indexes.conf** on all peer nodes in order for replication work properly. If you're creating new index on one peer then deploy that in other peers as well. Can be under search app local or Splunk local but should be identical.

Module-19:
Search Head Clustering

Initiialize clustering on all search peers -

splunk init shcluster-config -auth admin:searchhead1 -mgmt_uri https://<IP of SH1>:8089 -replication_port 9080 -replication_factor 2 -shcluster_label shcluster1

splunk init shcluster-config -auth admin:searchhead1 -mgmt_uri https://<IP of Sh2>:8089 -replication_port 9080 -replication_factor 2 -shcluster_label shcluster1

Restart splunk

Now, define captain for SH clustering. You can set any search peer as captain and change it later. You need to mention all search peer IPs in the below query and execute it from the SH server you want to make captain.

splunk bootstrap shcluster-captain -servers_list
"https://18.141.176.231:8089,https://13.229.201.166:8089" -auth admin:searchhead1

That's it. Now, if you'll check settings option on your search heads, you'll notice change and get an new option of 'Search Head Clustering'

All the reports, alerts, dashboards, user access, etc, everything will be replicated on all the peers of search head clustering.

Important: - Now, you can set-up UF to send the data to one of the indexer and it will be. Automatically replicated over other indexers.

You may define a server class on DS and cfreate a app like base_config in order to push outputs.conf and inputs.conf to UFs. Please make sure indexing is enabled on indexer otherwise UFs will fail to connect to Indexer and apparently no data will be ingested.