
telegraf-os Documentation

Release 1

Guilhem Marchand

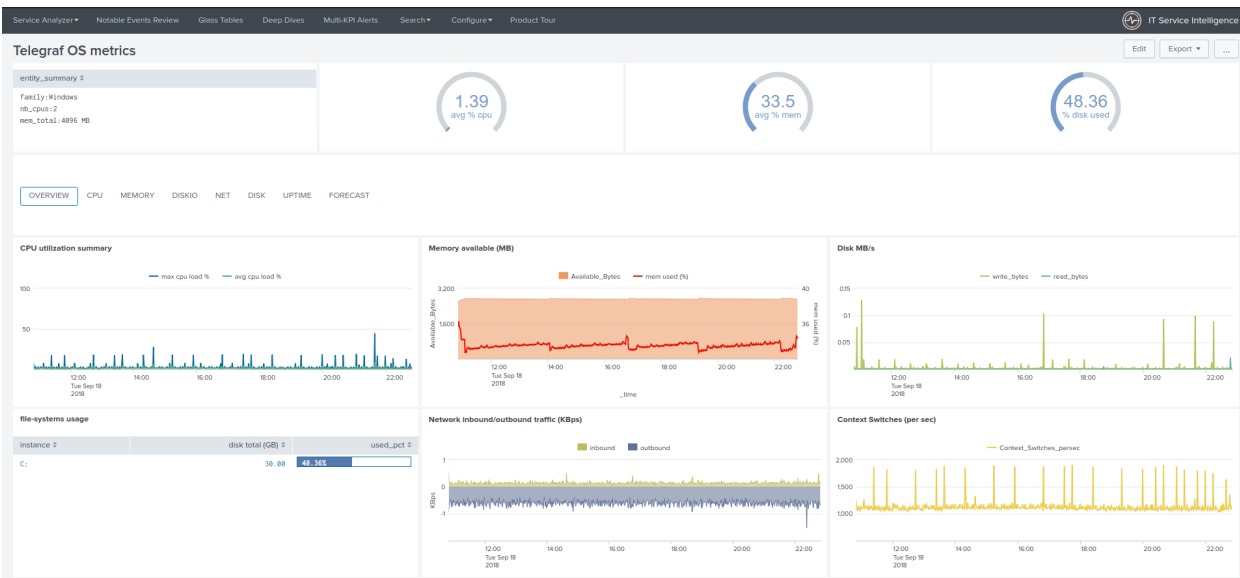
Aug 29, 2021

Contents

1	Overview:	3
1.1	About	3
1.2	Compatibility	3
1.3	Known Issues	4
1.4	Support & donate	4
1.5	Download	4
2	Deployment and configuration:	7
2.1	Deployment & Upgrades	7
2.2	Telegraf metrics ingestion	8
2.3	Splunk dashboards (health views)	17
3	Troubleshoot:	31
3.1	Troubleshoot & FAQ	31
4	Versioning and build history:	33
4.1	Release notes	33

The Splunk application for Operating System monitoring with Telegraf provides analytic and reporting for Linux and Windows metrics ingested in the Splunk metric store:





1.1 About

- Author: Guilhem Marchand
- First release published in October 2018
- Purposes:

The Splunk application for OS monitoring with Telegraf leverages the Influxdata Telegraf agent to provide key layer Operating System monitoring for Windows and Linux:

- Telegraf from Influxdata (<https://github.com/influxdata/telegraf>)

The Splunk application is backport of the ITSI module for Telegraf OS for non ITSI users.

<https://da-itsi-telegraf-os.readthedocs.io>

1.2 Compatibility

1.2.1 Splunk compatibility

All the metrics are ingested into the high performance Splunk metric store, Splunk 7.0.x or higher is required.

1.2.2 Telegraf compatibility

Telegraf supports various operating systems and process architectures including any version of Linux and Windows.

For more information:

- <https://portal.influxdata.com/downloads>

1.3 Known Issues

There are no known issues at the moment.

1.4 Support & donate

I am supporting my applications for free, for the good of everyone and on my own private time. As you can guess, this is a huge amount of time and efforts.

If you enjoy it, and want to support and encourage me, buy me a coffee (or a Pizza) and you will make me very happy!

The Splunk application for Operating System monitoring with Telegraf is community supported.

To get support, use of one the following options:

1.4.1 Splunk Answers

Open a question in Splunk answers for the application:

- <https://answers.splunk.com/app/questions/4271.html>

1.4.2 Splunk community slack

Contact me on Splunk community slack, or even better, ask the community !

- <https://splunk-usergroups.slack.com>

1.4.3 Open a issue in Git

To report an issue, request a feature change or improvement, please open an issue in Github:

- <https://github.com/guilhemmarchand/telegraf-os/issues>

1.4.4 Email support

- guilhem.marchand@gmail.com

However, previous options are far better, and will give you all the chances to get a quick support from the community of fellow Splunkers.

1.5 Download

1.5.1 Splunk Application for Operating System monitoring with Telegraf

The Splunk application can be downloaded from:

Splunk base

- <https://splunkbase.splunk.com/app/4271>

GitHub

- <https://github.com/guilhemmarchand/telegraf-os>

Deployment and configuration:

2.1 Deployment & Upgrades

2.1.1 Deployment matrix

Splunk roles	required
Search head	yes
Indexer tiers	no

If Splunk search heads are running in Search Head Cluster (SHC), the Splunk application must be deployed by the SHC deployer.

The deployment and configuration requires the creation of a dedicated metric index (by default called **telegraf**), see the implementation section.

2.1.2 Dependencies

The overview home page requires the deployment of the horizon chart and horseshoe third party visualisation addons on the search heads:

- <https://splunkbase.splunk.com/app/3117>
- <https://splunkbase.splunk.com/app/3166>

The metric workspace application is not required but highly recommended, if it is installed, the Metric link will be automatically available within the application bar:

- <https://splunkbase.splunk.com/app/4192/>

2.1.3 Initial deployment

The deployment of the Splunk application for OS monitoring with Telegraf is straight forward:

- Using the application manager in Splunk Web (Settings / Manages apps)
- Extracting the content of the tgz archive in the “apps” directory of Splunk
- For SHC configurations (Search Head Cluster), extract the tgz content in the SHC deployer and publish the SHC bundle

2.1.4 Upgrade

Upgrading the Splunk application is pretty much the same operation, use one of the techniques that matches your conditions / requirements.

2.2 Telegraf metrics ingestion

Implementing Telegraf and sending its metrics in Splunk is simple, and efficient.

It is not the purpose of this documentation to expose every piece of the installation and configuration of Telegraf or Splunk.

2.2.1 Telegraf installation and configuration

Telegraf standard installation (standalone and independant process)

The installation of Telegraf is really straightforward, consult:

- <https://github.com/influxdata/telegraf>

If you wish to deploy Telegraf as a Splunk TA application instead, consult the next step.

Telegraf deployment as Splunk application deployed by Splunk (TA)

You can publish Telegraf through a Splunk application that you push to your clients using a Splunk deployment server.

This means that you can create a custom Technology Addon (TA) that contains both the Telegraf binary and the telegraf.conf configuraton files.

This method has several advantages:

- If you are a Splunk customer already, you may have the Splunk Universal Forwarder deployed on your servers, you will NOT need to deploy an additional collector independently from Splunk
- You get benefit from the Splunk centralisation and deploy massively Telegraf from Splunk
- You can maintain and upgrade Telegraf just as you do usually in Splunk, all from your Splunk Deployment Server. (DS)

Linux 64 bits example:

- <https://github.com/guilhemmarchand/TA-telegraf-amd64>

Windows 64 bits example:

- <https://github.com/guilhemmarchand/TA-telegraf-windows64>

Telegraf minimal configuration

A minimal configuration to monitor Operating System metrics:

- <https://docs.influxdata.com/chronograf/latest/guides/using-precreated-dashboards/#system>

The output configuration depends on the deployment you choose to use to ingest metrics in Splunk, consult the next sections.

Example of a minimal telegraf.conf configuration that monitors Operating System metrics for Linux:

```
# Read metrics about cpu usage
[[inputs.cpu]]
  ## Whether to report per-cpu stats or not
  percpu = true
  ## Whether to report total system cpu stats or not
  totalcpu = true
  ## If true, collect raw CPU time metrics.
  collect_cpu_time = false
  ## If true, compute and report the sum of all non-idle CPU states.
  report_active = false

# Read metrics about disk usage by mount point
[[inputs.disk]]
  ## By default stats will be gathered for all mount points.
  ## Set mount_points will restrict the stats to only the specified mount points.
  # mount_points = ["/"]

  ## Ignore mount points by filesystem type.
  ignore_fs = ["tmpfs", "devtmpfs", "devfs"]

# Read metrics about disk IO by device
[[inputs.diskio]]
  ## By default, telegraf will gather stats for all devices including
  ## disk partitions.
  ## Setting devices will restrict the stats to the specified devices.
  # devices = ["sda", "sdb", "vd*"]
  ## Uncomment the following line if you need disk serial numbers.
  # skip_serial_number = false
  #
  ## On systems which support it, device metadata can be added in the form of
  ## tags.
  ## Currently only Linux is supported via udev properties. You can view
  ## available properties for a device by running:
  ## 'udevadm info -q property -n /dev/sda'
  # device_tags = ["ID_FS_TYPE", "ID_FS_USAGE"]
  #
  ## Using the same metadata source as device_tags, you can also customize the
  ## name of the device via templates.
  ## The 'name_templates' parameter is a list of templates to try and apply to
  ## the device. The template may contain variables in the form of '$PROPERTY' or
  ## '${PROPERTY}'. The first template which does not contain any variables not
  ## present for the device is used as the device name tag.
  ## The typical use case is for LVM volumes, to get the VG/LV name instead of
  ## the near-meaningless DM-0 name.
  # name_templates = ["$ID_FS_LABEL", "$DM_VG_NAME/$DM_LV_NAME"]

# Get kernel statistics from /proc/stat
[[inputs.kernel]]
```

(continues on next page)

(continued from previous page)

```

# no configuration

# Read metrics about memory usage
[[inputs.mem]]
# no configuration

# Get the number of processes and group them by status
[[inputs.processes]]
# no configuration

# Read metrics about swap memory usage
[[inputs.swap]]
# no configuration

# Read metrics about system load & uptime
[[inputs.system]]
# no configuration

# # Read metrics about network interface usage
[[inputs.net]]
# ## By default, telegraf gathers stats from any up interface (excluding loopback)
# ## Setting interfaces will tell it to gather these explicit interfaces,
# ## regardless of status.
# ##
# # interfaces = ["eth0"]
# ##
# ## On linux systems telegraf also collects protocol stats.
# ## Setting ignore_protocol_stats to true will skip reporting of protocol metrics.
# ##
# # ignore_protocol_stats = false
# ##

# # Read TCP metrics such as established, time wait and sockets counts.
[[inputs.netstat]]
# # no configuration

# # Monitor process cpu and memory usage
[[inputs.procstat]]
# ## PID file to monitor process
# # pid_file = "/var/run/nginx.pid"
# ## executable name (ie, pgrep <exe>)
# # exe = "nginx"
# ## pattern as argument for pgrep (ie, pgrep -f <pattern>)
# # pattern = "nginx"
# ## user as argument for pgrep (ie, pgrep -u <user>)
# # user = "root"
# ## Systemd unit name
# # systemd_unit = "nginx.service"
# ## CGroup name or path
# # cgroup = "systemd/system.slice/nginx.service"
#
# ## override for process_name
# ## This is optional; default is sourced from /proc/<pid>/status

```

(continues on next page)

(continued from previous page)

```
# # process_name = "bar"
#
# ## Field name prefix
# # prefix = ""
#
# ## Add PID as a tag instead of a field; useful to differentiate between
# ## processes whose tags are otherwise the same. Can create a large number
# ## of series, use judiciously.
# # pid_tag = false
#
# ## Method to use when finding process IDs. Can be one of 'pgrep', or
# ## 'native'. The pgrep finder calls the pgrep executable in the PATH while
# ## the native finder performs the search directly in a manor dependent on the
# ## platform. Default is 'pgrep'
# # pid_finder = "pgrep"
pattern = ".*"
```

Windows additional configuration (mem inputs)

For Windows memory management, the default win_mem inputs does not retrieve some of the metrics we need. You need to activate the memory inputs. (which on Windows uses WMI collection):

```
[[inputs.mem]]
# no configuration
```

Windows Active Directory Domain Controller

Follow instructions for “Active Directory Domain Controller”:

- https://github.com/influxdata/telegraf/tree/master/plugins/inputs/win_perf_counters#active-directory-domain-controller

Windows DNS server

Follow instructions for “DNS Server + Domain Controllers”:

- https://github.com/influxdata/telegraf/tree/master/plugins/inputs/win_perf_counters#dns-server-domain-controllers

Windows DFS server

For DFS Namespace, follow instructions for “DFS Namespace + Domain Controllers”:

- https://github.com/influxdata/telegraf/tree/master/plugins/inputs/win_perf_counters#dfs-namespace-domain-controllers

For DFS Replication, follow instructions for “DFS Replication + Domain Controllers”:

- https://github.com/influxdata/telegraf/tree/master/plugins/inputs/win_perf_counters#dfs-replication-domain-controllers

Microsoft IIS / ASP.NET

For IIS and ASP.NET, follow instructions for: IIS / ASP.NET

- https://github.com/influxdata/telegraf/tree/master/plugins/inputs/win_perf_counters#iis-aspnet

Linux processes monitoring (procstat)

In the linux views, the processes usage (both CPU and Memory) rely on the procstat inputs, which requires additional configuration depending on your context.

As for an example, the following configuration monitors all the processes owned by the “splunk” unix user:

```
[[inputs.procstat]]
#   ## PID file to monitor process
#   pid_file = "/var/run/nginx.pid"
#   ## executable name (ie, pgrep <exe>)
#   exe = "nginx"
#   ## pattern as argument for pgrep (ie, pgrep -f <pattern>)
#   pattern = "nginx"
#   ## user as argument for pgrep (ie, pgrep -u <user>)
#   user = "splunk"
```

2.2.2 HTTP Events Collector (HEC)

Splunk deployment with HEC (available with Telegraf starting version 1.8)

Telegraf agents -> HTTP over SSL -> Splunk HEC inputs

With Telegraf starting version 1.8, you can send metrics directly from Telegraf to HTTP Events Collector using the excellent serializer leveraging the http Telegraf output.

This is extremely simple, scalable and reliable.

Example of an HEC input definition:

Splunk inputs.conf:

```
[http://Telegraf]
disabled = 0
index = telegraf
indexes = telegraf
token = c386d4c8-8b50-4178-be76-508dca2f19e2
```

Telegraf configuration:

The Telegraf configuration is really simple and relies on defining your output:

Example:

```
[[outputs.http]]
# URL is the address to send metrics to
url = "https://mysplunk.domain.com:8088/services/collector"
# Timeout for HTTP message
# timeout = "5s"
# Optional TLS Config
# tls_ca = "/etc/telegraf/ca.pem"
# tls_cert = "/etc/telegraf/cert.pem"
```

(continues on next page)

(continued from previous page)

```
# tls_key = "/etc/telegraf/key.pem"
## Use TLS but skip chain & host verification
insecure_skip_verify = true
## Data format to output.
## Each data format has it's own unique set of configuration options, read
## more about them here:
## https://github.com/influxdata/telegraf/blob/master/docs/DATA_FORMATS_OUTPUT.md
data_format = "splunkmetric"
## Provides time, index, source overrides for the HEC
splunkmetrichec_routing = true
## Additional HTTP headers
[outputs.http.headers]
# Should be set manually to "application/json" for json data_format
Content-Type = "application/json"
Authorization = "Splunk c386d4c8-8b50-4178-be76-508dca2f19e2"
X-Splunk-Request-Channel = "c386d4c8-8b50-4178-be76-508dca2f19e2"
```

Push this configuration to your Telegraf agents, et voila.

Check data availability in Splunk:

```
| mcatalog values(_dims) as dimensions values(metric_name) as metric_name where_
↪index=telegraf metric_name=*
```

2.2.3 TCP / TCP-SSL Inputs

Splunk deployment with TCP inputs.

This deployment requires additional indexing time parsing configuration:

- <https://github.com/guilhemmarchand/TA-influxdata-telegraf>

The deployment is very simple and can be described as:

Telegraf agents → TCP or TCP over SSL → Splunk TCP inputs

In addition and to provide resiliency, it is fairly simple to add a load balancer in front of Splunk, such that you service continues to ingest metrics depending on Splunk components availability. (HAProxy, Nginx, F5, whatsoever...)

The data output format used by Telegraf agents is the “Graphite” format with tag support enable. This is simple, beautiful, accurate and allows the management of any number of dimensions.

Example of a tcp input definition:

Splunk inputs.conf:

```
[tcp://2003]
connection_host = dns
index = telegraf
sourcetype = tcp:telegraf:graphite
```

Telegraf configuration:

The Telegraf configuration is really simple and relies on defining your output:

Example:

```
[[outputs.graphite]]
  ## TCP endpoint for your graphite instance.
  ## If multiple endpoints are configured, the output will be load balanced.
  ## Only one of the endpoints will be written to with each iteration.
  servers = ["mysplunk.domain.com:2003"]
  ## Prefix metrics name
  prefix = ""
  ## Graphite output template
  ## see https://github.com/influxdata/telegraf/blob/master/docs/DATA_FORMATS_OUTPUT.
  ↪md
  # template = "host.tags.measurement.field"

  ## Enable Graphite tags support
  graphite_tag_support = true

  ## timeout in seconds for the write connection to graphite
  timeout = 2
```

Push this configuration to your Telegraf agents, et voila.

Check data availability in Splunk::

mcatalog values(_dims) as dimensions values(metric_name) as metric_name where index=telegraf
metric_name=*

2.2.4 SPLUNK file monitoring Ingestion

Splunk deployment with Splunk file monitoring.

This deployment requires additional indexing time parsing configuration:

- <https://github.com/guilhemmarchand/TA-influxdata-telegraf>

Telegraf has a “file” output plugin that allows writing metrics to a local file on the file-system, don’t say more that is much more than enough to Splunk it.

The Telegraf configuration is really simple and relies on defining your ouput:

Example:

```
[[outputs.file]]
  ## Files to write to, "stdout" is a specially handled file.
  files = ["/tmp/metrics.out"]

  ## Data format to output.
  ## Each data format has its own unique set of configuration options, read
  ## more about them here:
  ## https://github.com/influxdata/telegraf/blob/master/docs/DATA_FORMATS_OUTPUT.md
  data_format = "graphite"
  graphite_tag_support = true
```

Notes: this is a simplistic example, in real condition do not forget to manage the file rotation using a simple logrotate configuration for Linux, and relevant solution for other OS.

Splunk file input configuration:

I cannot say more, this is simple, very simple. Add the following configuration to any inputs.conf configuration file of your choice:

Example:

```
[monitor:///tmp/metrics.out]
disabled = false
index = telegraf
sourcetype = file:telegraf:graphite
```

Apply this simple input.conf, if you deploy thought the Splunk deployment server ensure splunkd is configured to restart in your serverclass configuration.

Et voila, Splunk ingests the metrics continuously and metrics are forwarded to the indexing layer using your Splunk infrastructure, be on-premise, private or Splunk Cloud.

2.2.5 KAFKA Ingestion

Splunk deployment with Kafka.

This deployment requires additional indexing time parsing configuration:

- <https://github.com/guilhemmarchand/TA-influxdata-telegraf>

If you are using Kafka, or consider using it, producing Telegraf metrics to Kafka makes a lot of sense.

First, Telegraf has a native output plugin that produces to a Kafka topic, Telegraf will send the metrics directly to one or more Kafka brokers providing scaling and resiliency.

Then, Splunk becomes one consumer of the metrics using the scalable and resilient Kafka connect infrastructure and the Splunk Kafka connect sink connector. By using Kafka as the mainstream for your metrics, you preserve the possibility of having multiple technologies consuming these data in addition with Splunk, while implementing a massively scalable and resilient environment.

On the final step that streams data to Splunk, the Kafka sink connector for Splunk sends data to Splunk HEC, which makes it resilient, scalable and eligible to all Splunk on-premise or Splunk Cloud platforms easily.

The deployment with Kafka can be described the following way:

Telegraf agents → Kafka brokers ← Kafka connect running Splunk sink connector → Splunk HTTP Event Collector (HEC)

Configuring Kafka connect:

- The Kafka connect properties needs to use the “String” converter, the following example start Kafka connect with the relevant configuration:

connect-distributed.properties:

```
# These are defaults. This file just demonstrates how to override some settings.
bootstrap.servers=kafka-1:9092,kafka-2:9092,kafka-3:9092
key.converter=org.apache.kafka.connect.storage.StringConverter
value.converter=org.apache.kafka.connect.storage.StringConverter
key.converter.schemas.enable=false
value.converter.schemas.enable=false
internal.key.converter=org.apache.kafka.connect.json.JsonConverter
internal.value.converter=org.apache.kafka.connect.json.JsonConverter
internal.key.converter.schemas.enable=false
internal.value.converter.schemas.enable=false
# Flush much faster (10s) than normal, which is useful for testing/debugging
offset.flush.interval.ms=10000
plugin.path=/etc/kafka-connect/jars
group.id=kafka-connect-splunk-hec-sink
config.storage.topic=__kafka-connect-splunk-task-configs
```

(continues on next page)

(continued from previous page)

```

config.storage.replication.factor=3
offset.storage.topic=__kafka-connect-splunk-offsets
offset.storage.replication.factor=3
offset.storage.partitions=25
status.storage.topic=__kafka-connect-splunk-statuses
status.storage.replication.factor=3
status.storage.partitions=5
# These are provided to inform the user about the presence of the REST host and port
↪ configs
# Hostname & Port for the REST API to listen on. If this is set, it will bind to the
↪ interface used to listen to requests.
# rest.host.name=
rest.port=8082
# The Hostname & Port that will be given out to other workers to connect to i.e. URLs
↪ that are routable from other servers.
# rest.advertised.host.name=kafka-connect-1

```

Apply the following command against Kafka connect running the Splunk Kafka sink connector: (<https://splunkbase.splunk.com/app/3862>)

- replace <ip address or host> by the IP address or the host name of the Kafka connect node, if you run the command locally, simply use localhost
- replace the port if required (default is 8082)
- replace the HEC token
- replace the HEC destination
- adapt any other configuration item up to your needs

Achieve the following command:

```

curl localhost:8082/connectors -X POST -H "Content-Type: application/json" -d '{
  "name": "kafka-connect-telegraf",
  "config": {
    "connector.class": "com.splunk.kafka.connect.SplunkSinkConnector",
    "tasks.max": "3",
    "topics": "telegraf",
    "splunk.indexes": "telegraf",
    "splunk.sourcetypes": "kafka:telegraf:graphite",
    "splunk.hec.uri": "https://myhecinput.splunk.com:8088",
    "splunk.hec.token": "fd96ffb6-fb3e-43aa-9e8b-de911356443f",
    "splunk.hec.raw": "true",
    "splunk.hec.ack.enabled": "true",
    "splunk.hec.ack.poll.interval": "10",
    "splunk.hec.ack.poll.threads": "2",
    "splunk.hec.ssl.validate.certs": "false",
    "splunk.hec.http.keepalive": "true",
    "splunk.hec.max.http.connection.per.channel": "4",
    "splunk.hec.total.channels": "8",
    "splunk.hec.max.batch.size": "1000000",
    "splunk.hec.threads": "2",
    "splunk.hec.event.timeout": "300",
    "splunk.hec.socket.timeout": "120",
    "splunk.hec.track.data": "true"
  }
}'

```

telegraf output configuration:

Configure your Telegraf agents to send data directly to the Kafka broker in graphite format with tag support:

```
[outputs.kafka]
  ## URLs of kafka brokers
  brokers = ["kafka-1:19092", "kafka-2:29092", "kafka-3:39092"]
  ## Kafka topic for producer messages
  topic = "telegraf"
  data_format = "graphite"
  graphite_tag_support = true
```

Et voila. Congratulations, you have built a massively scalable, distributable, open and resilient metric collection infrastructure.

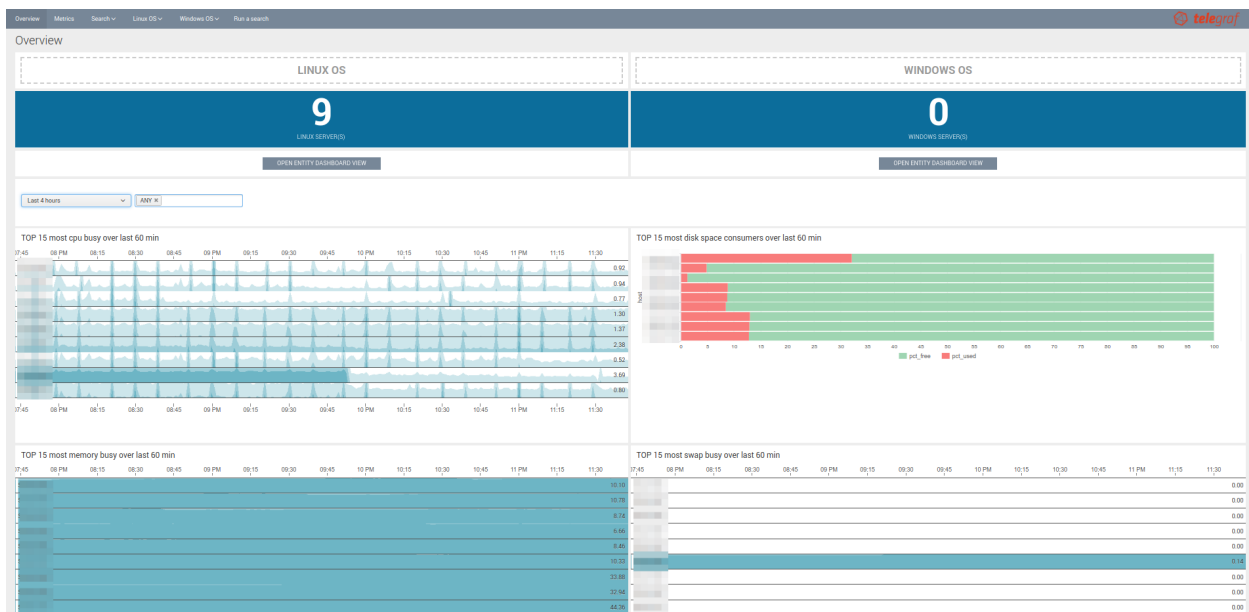
Check data availability in Splunk:

```
| mcatalog values(_dims) as dimensions values(metric_name) as metric_name where_
↪ index=telegraf metric_name=*
```

2.3 Splunk dashboards (health views)

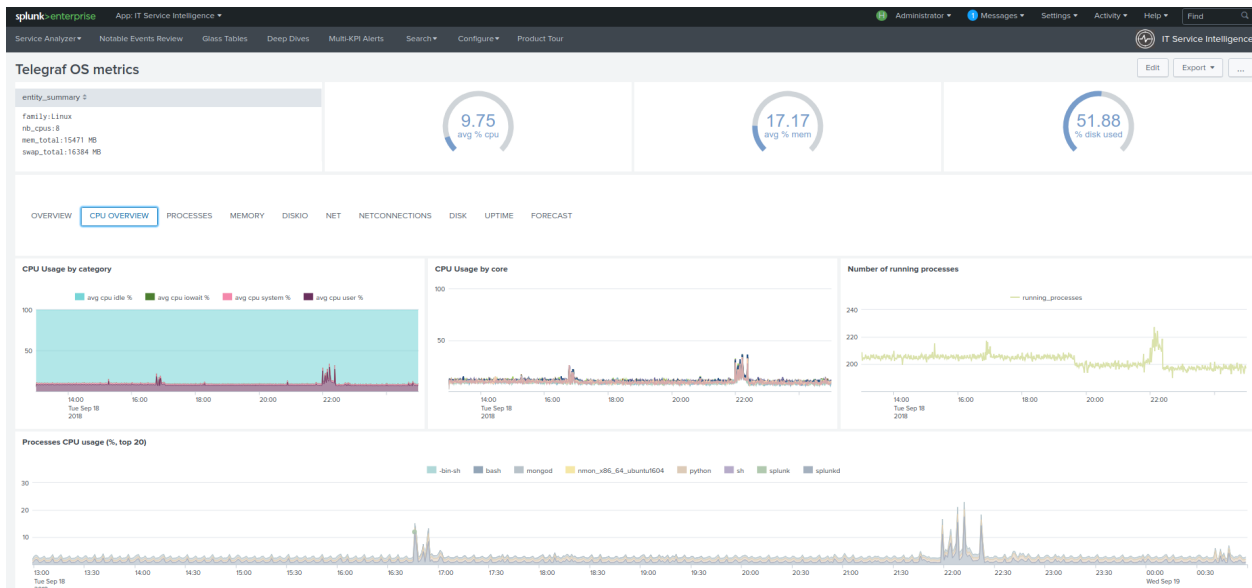
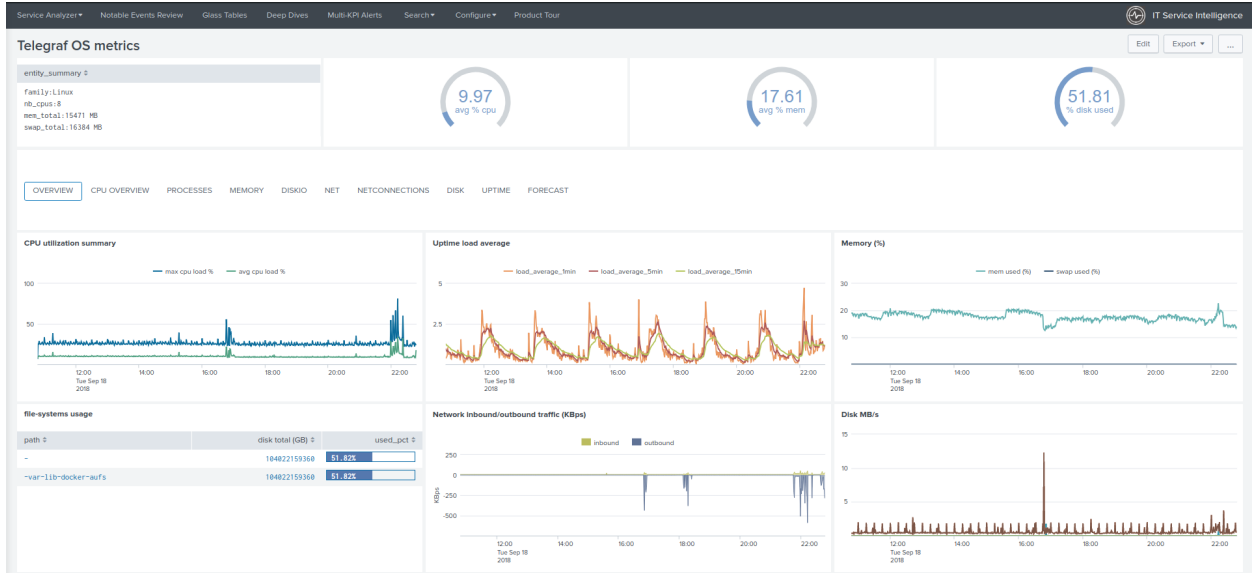
2.3.1 Overview (landing page)

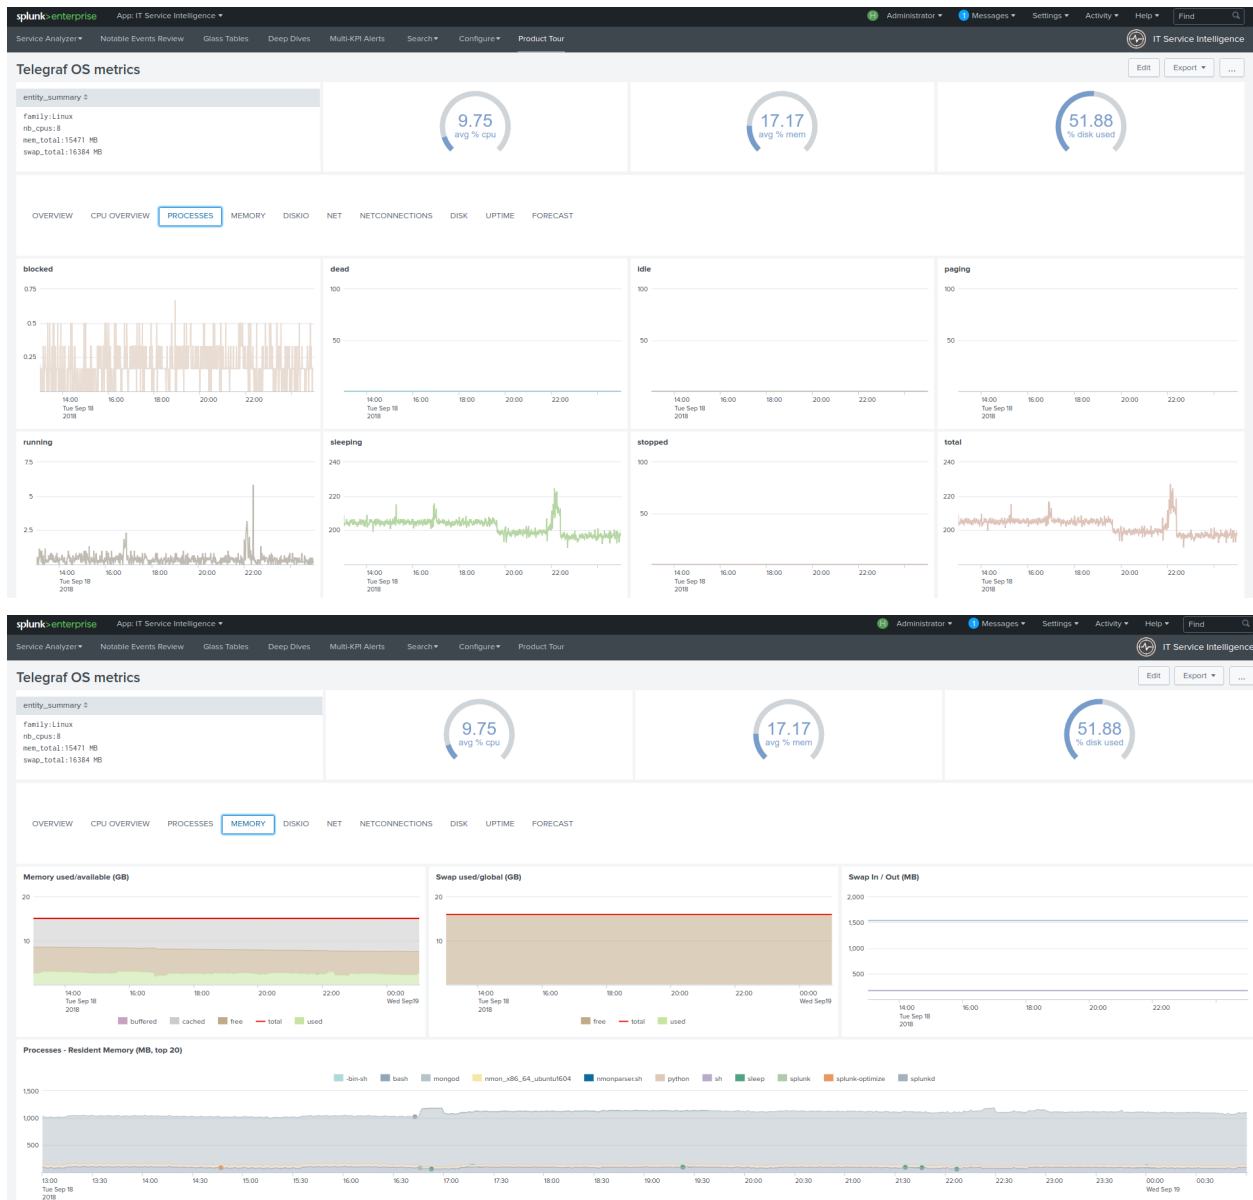
The Splunk application home page provides an overview of the Kafka infrastructure:

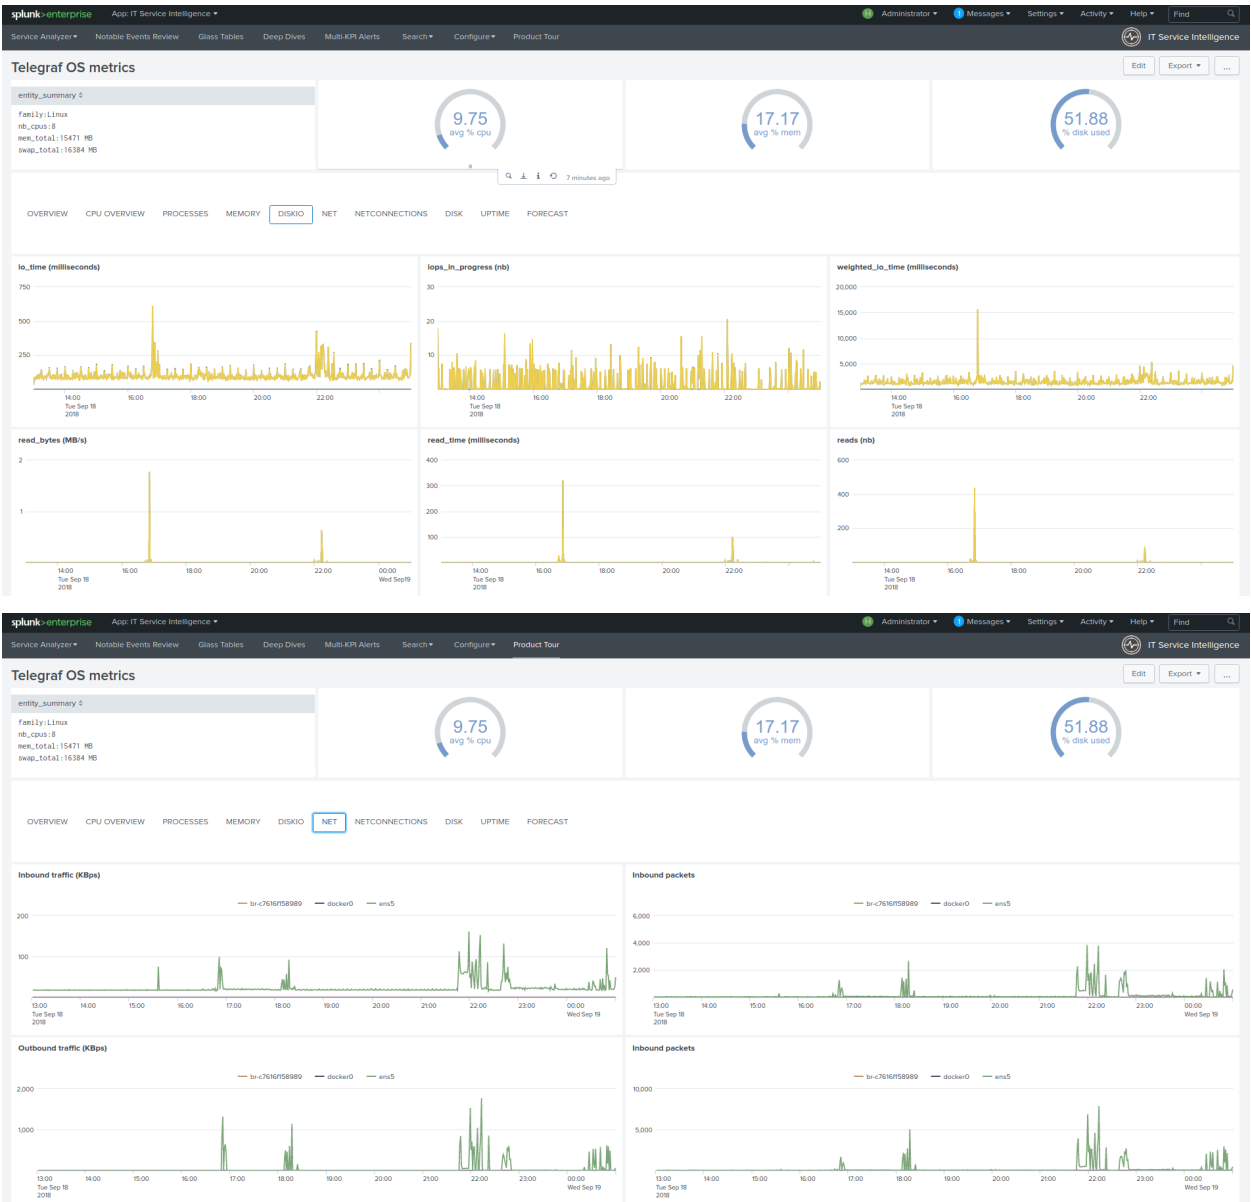


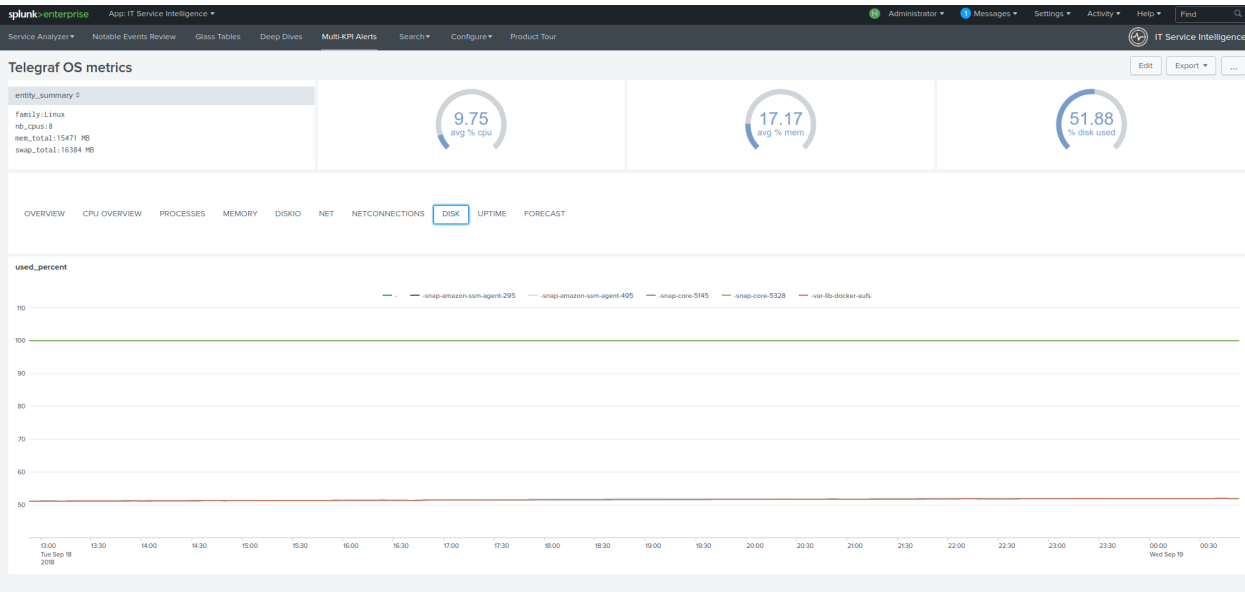
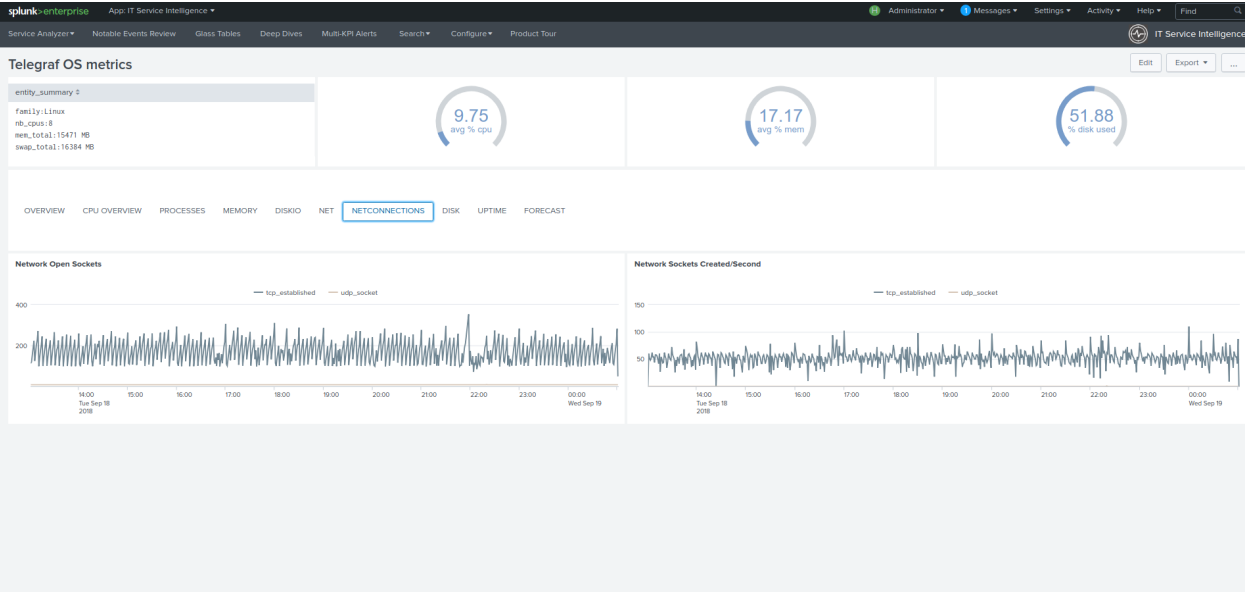
2.3.2 View for Linux OS

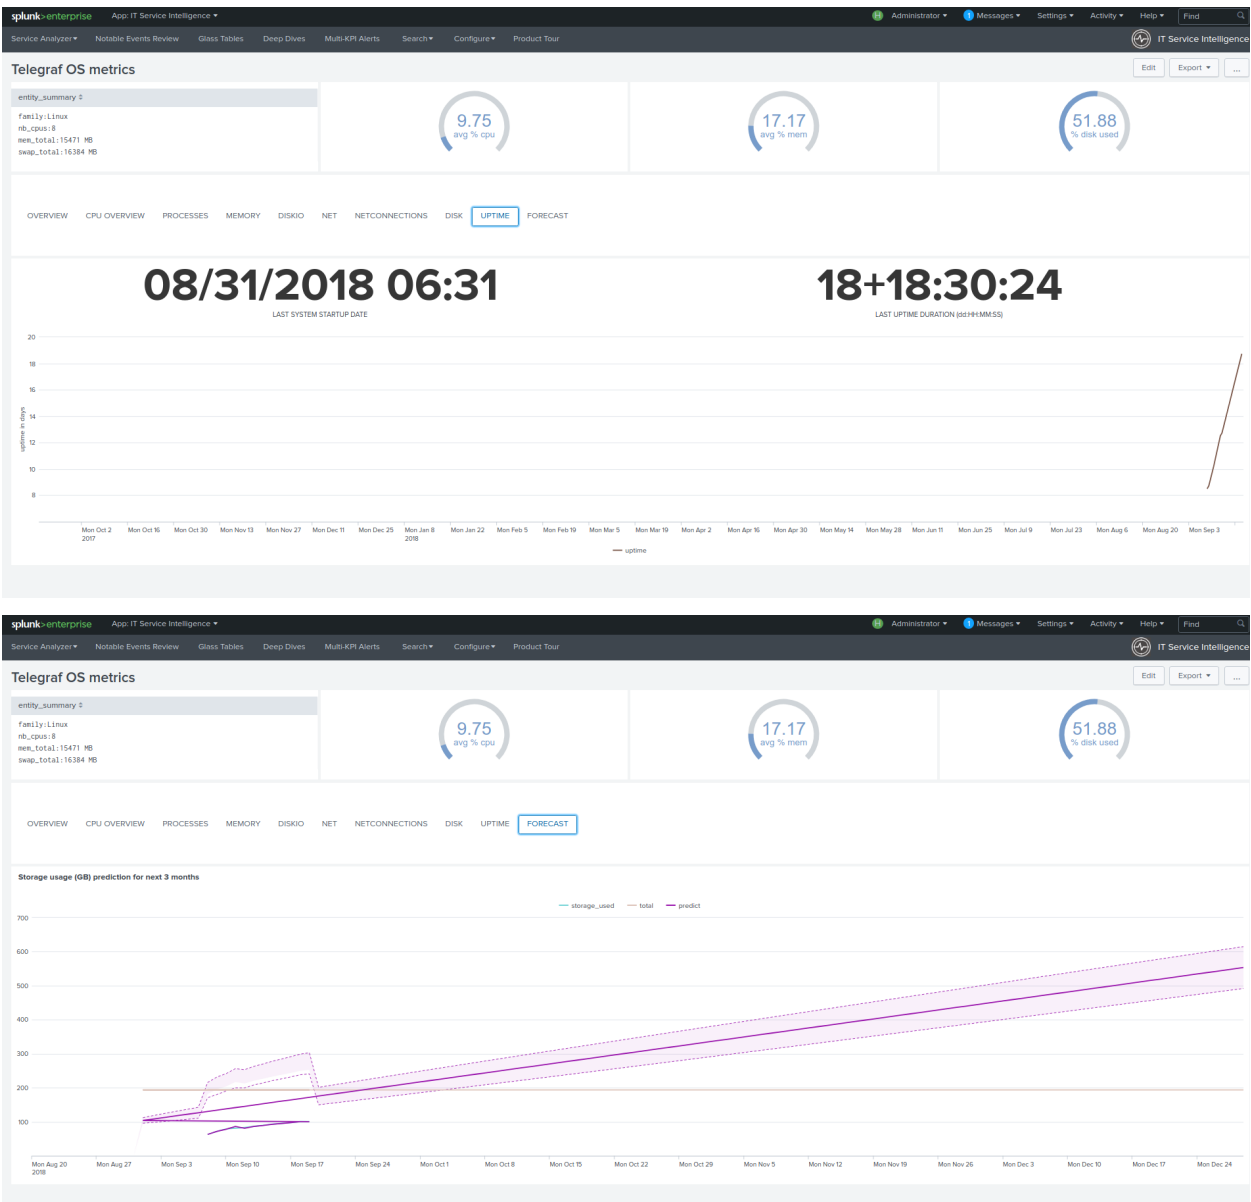
The Health view for Linux OS automatically appears as “Telegraf OS (Linux)” deepdive drilldown link when entities are discovered:





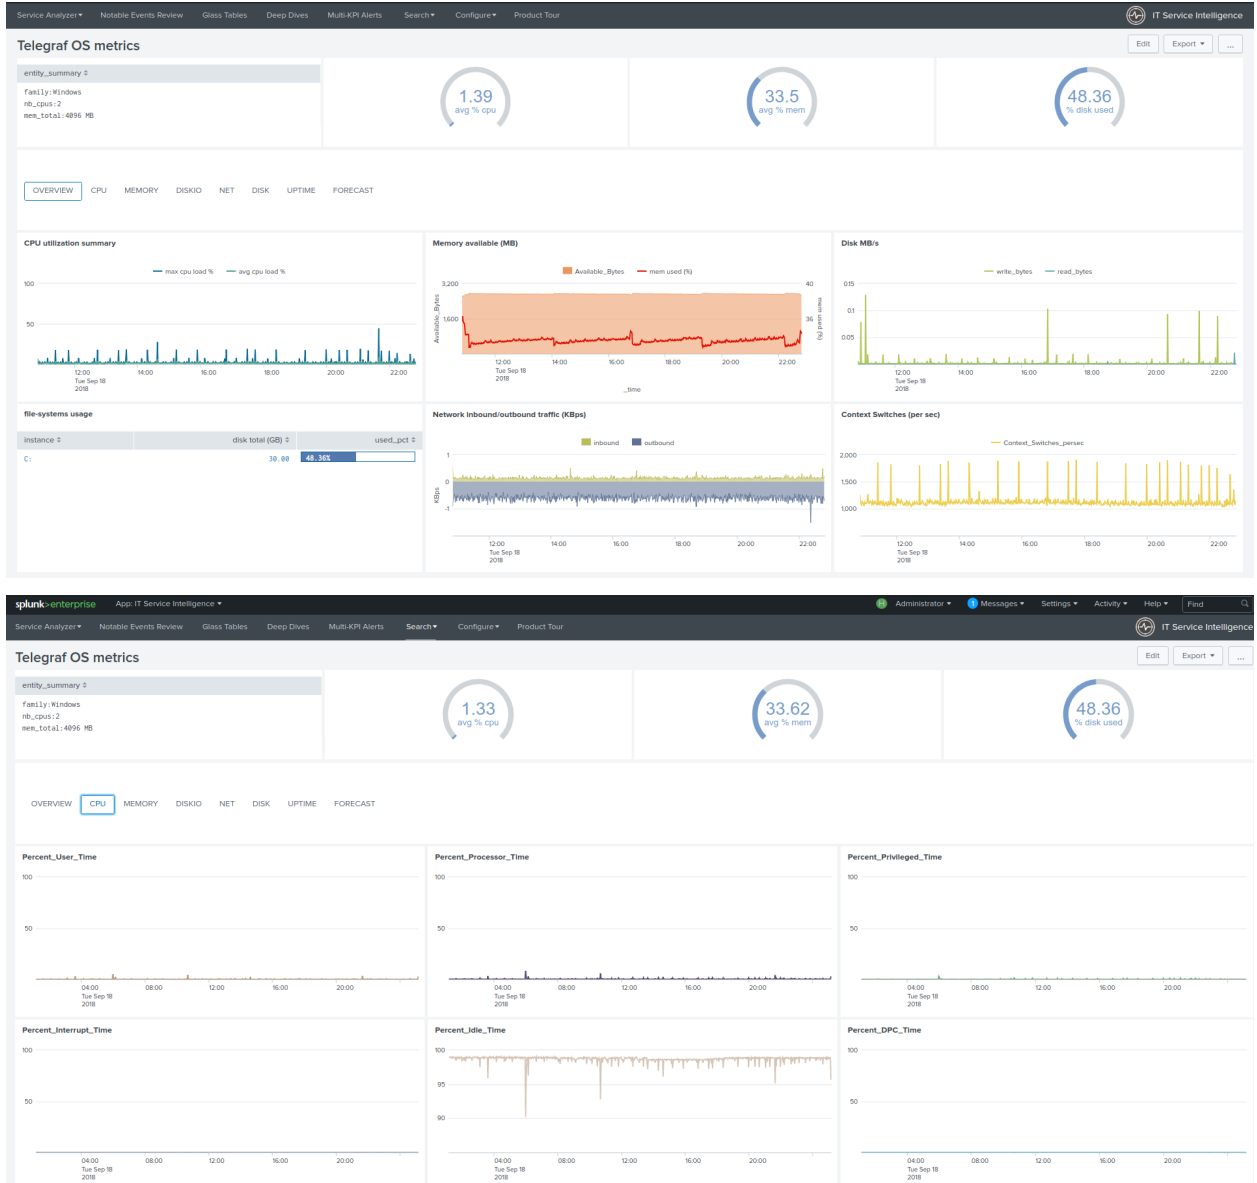


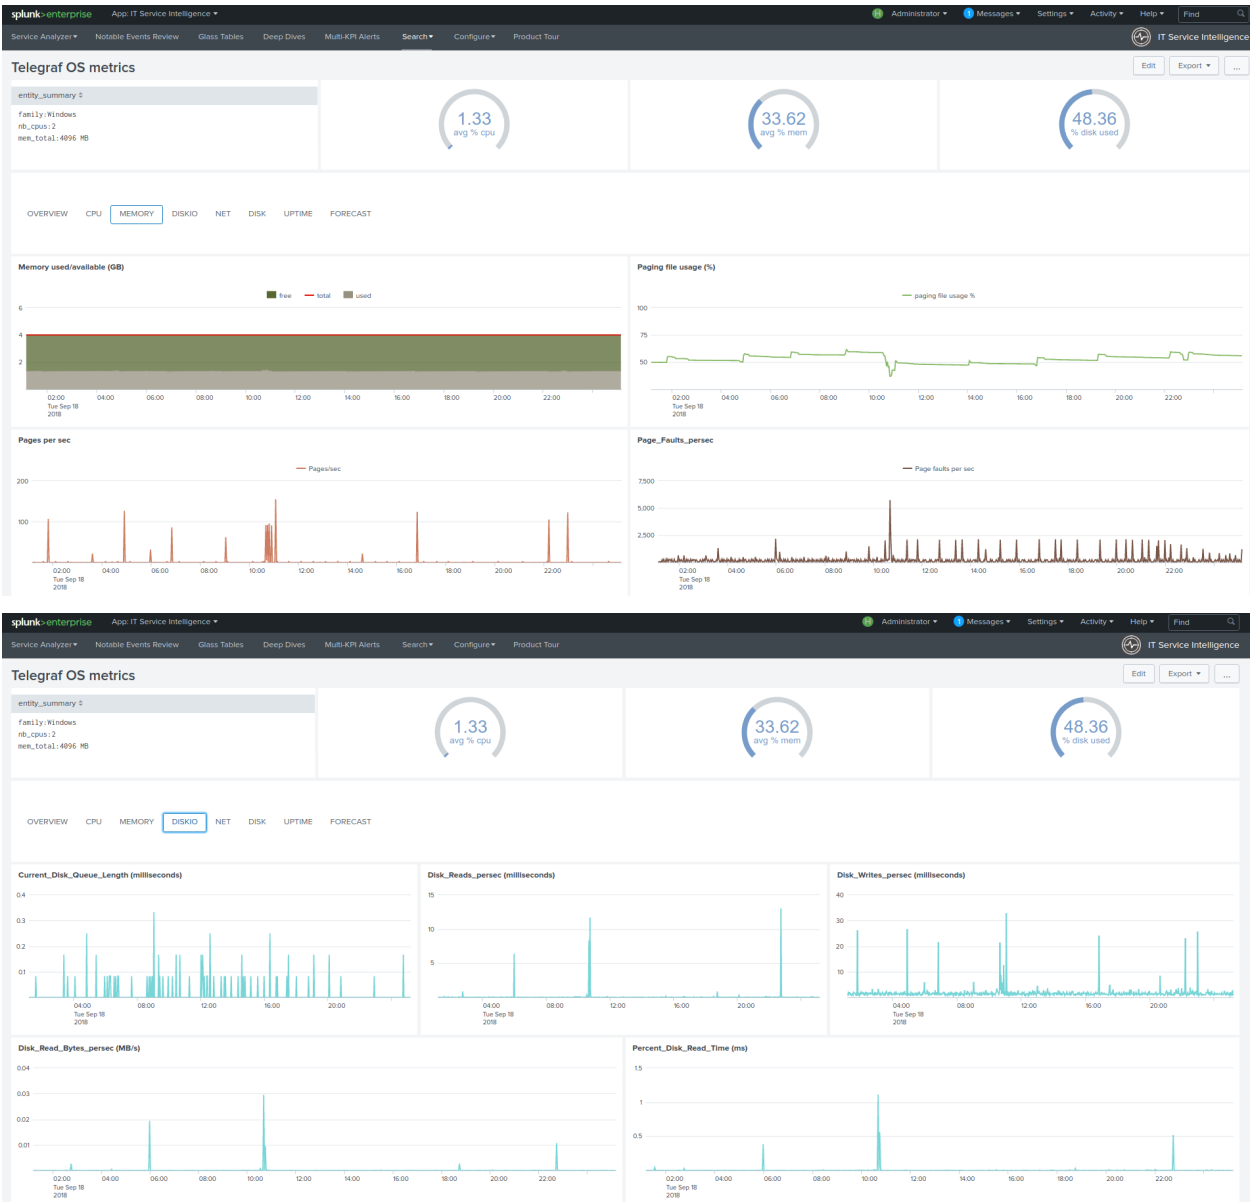


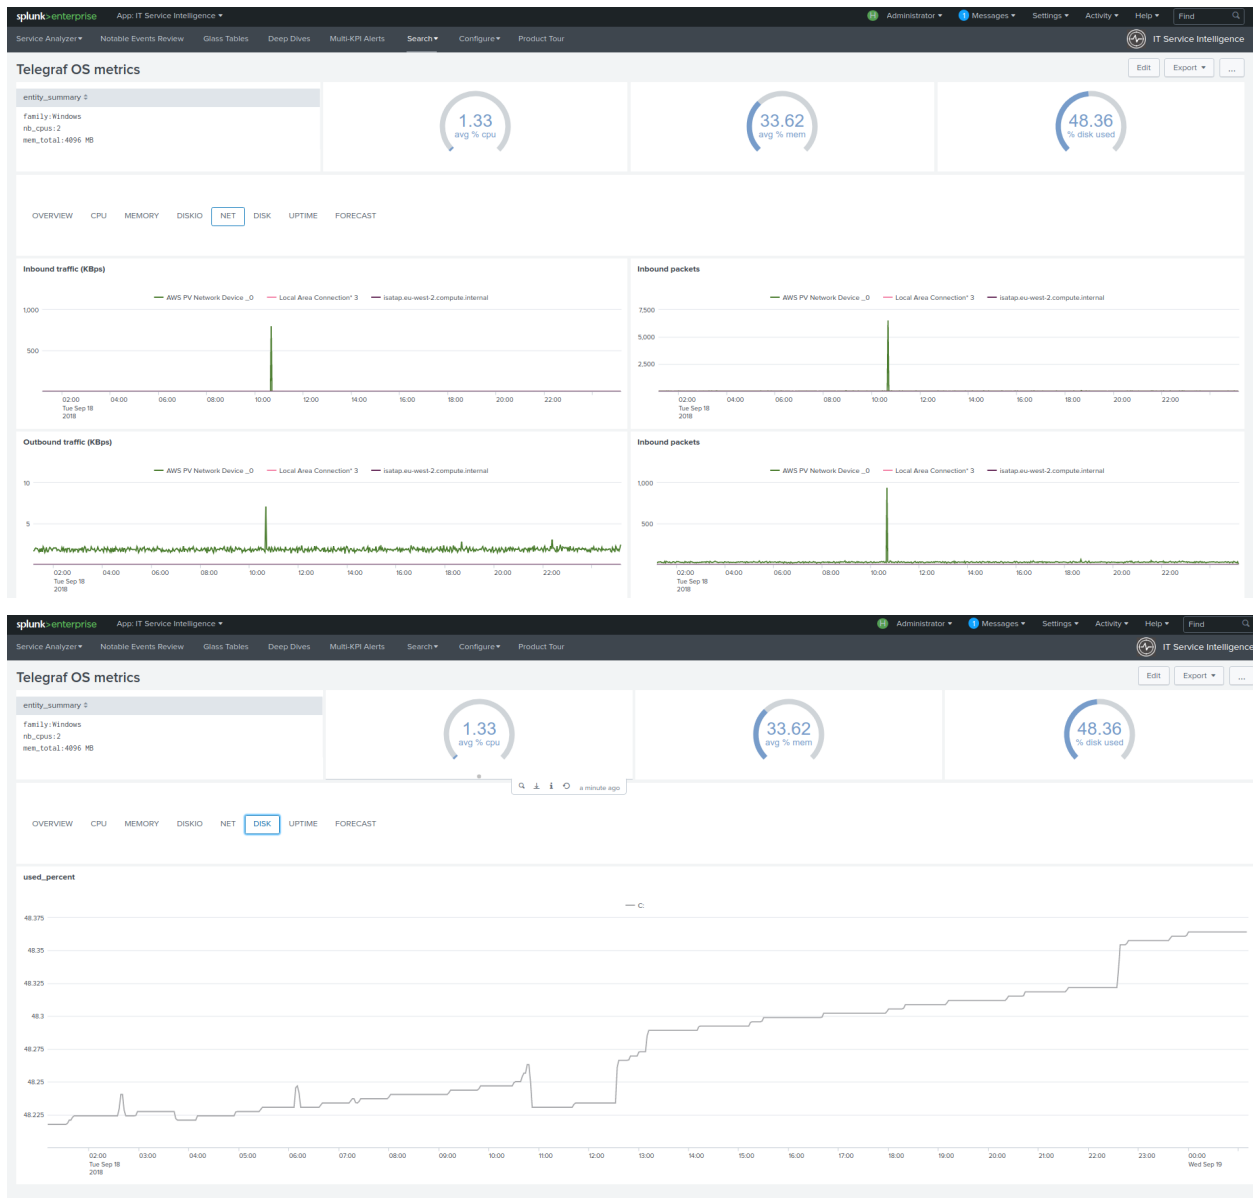


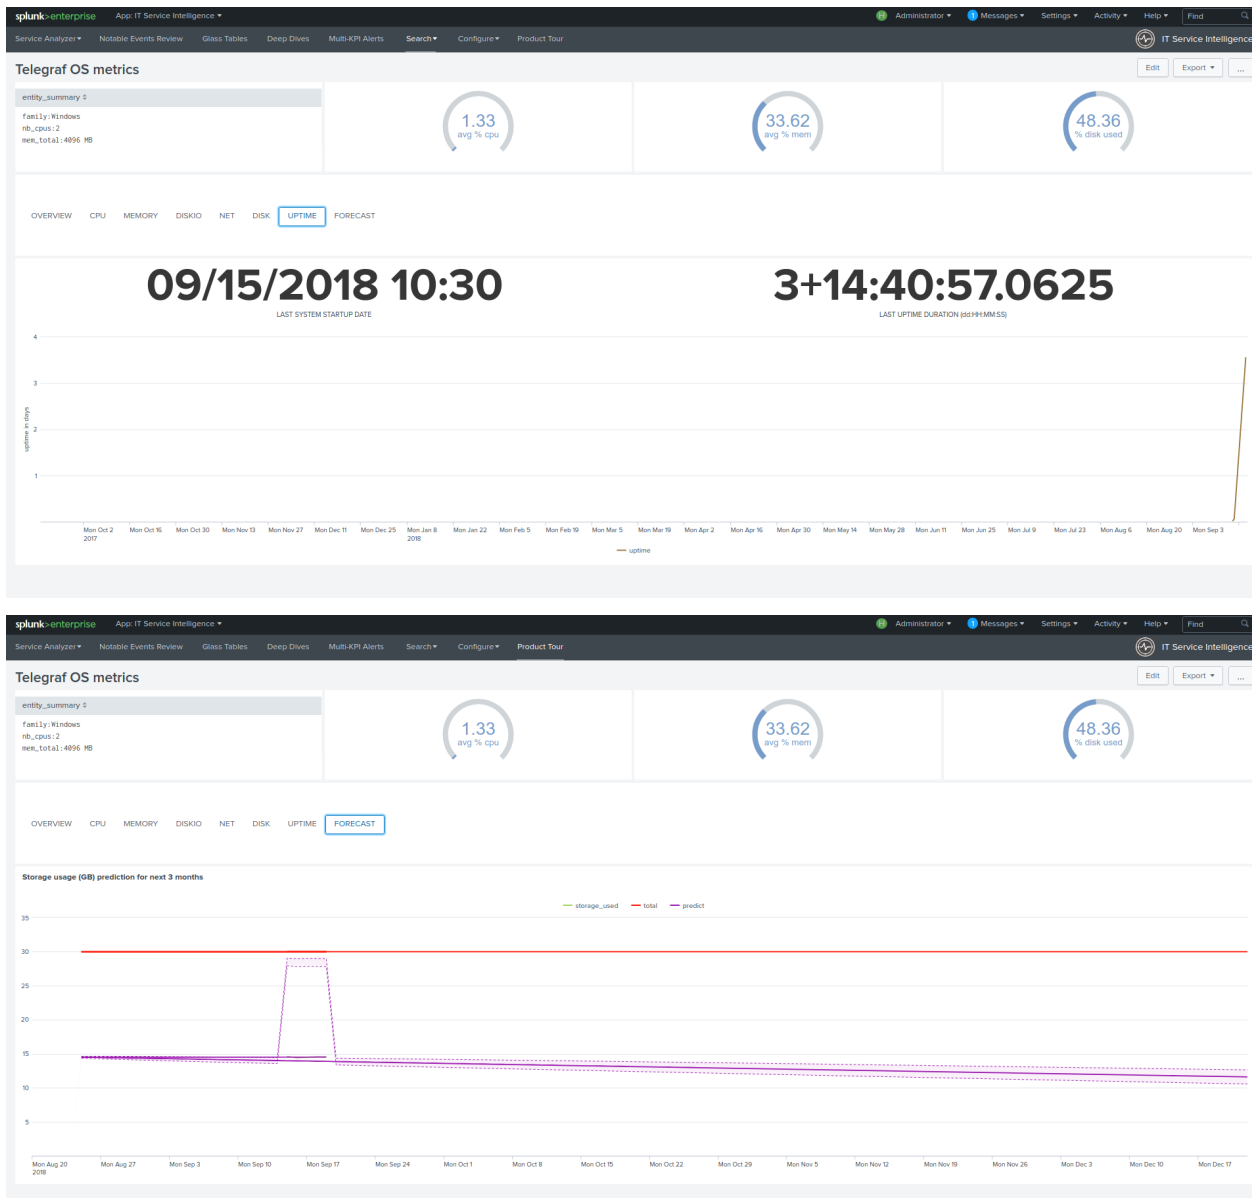
2.3.3 View for Windows OS

The Health view for Linux OS automatically appears as “Telegraf OS (Windows)” deepdive drilldown link when entities are discovered:



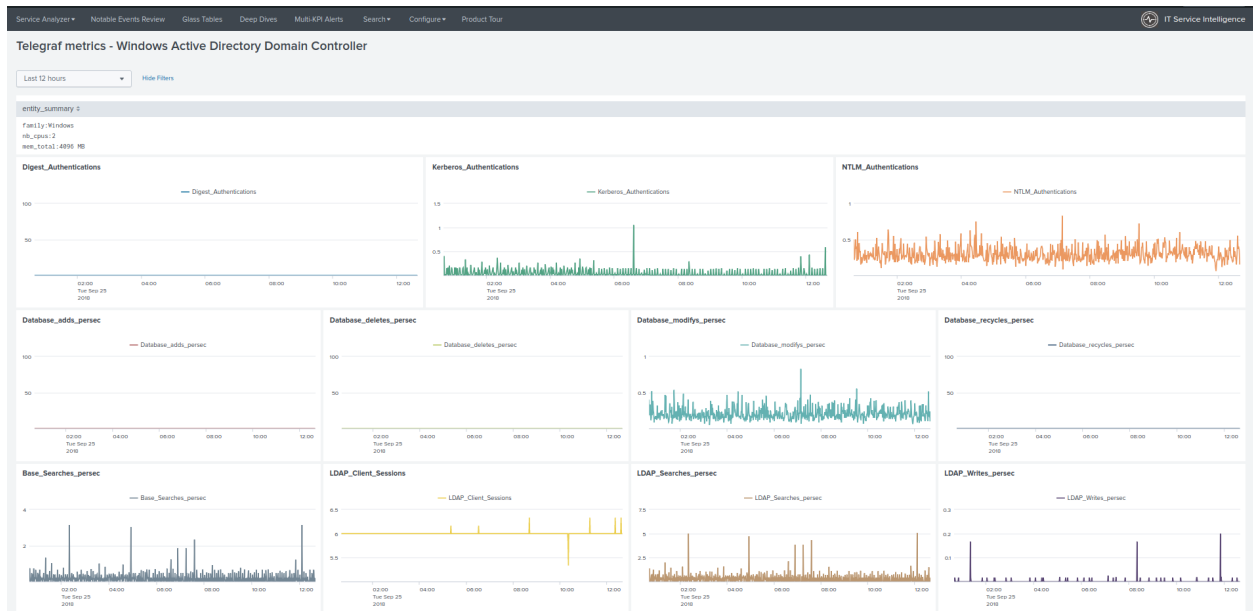






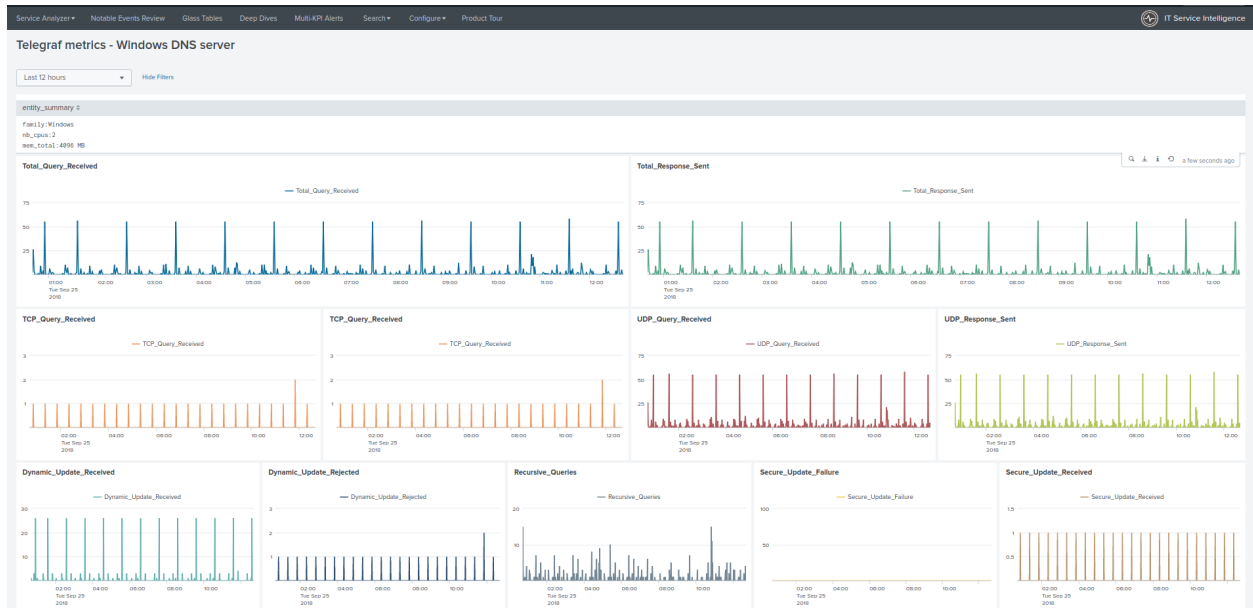
2.3.4 View for Windows Active Directory Domain Controller

The Health view automatically appears as “Telegraf Win AD-DC” deepdive drilldown link when entities are discovered:



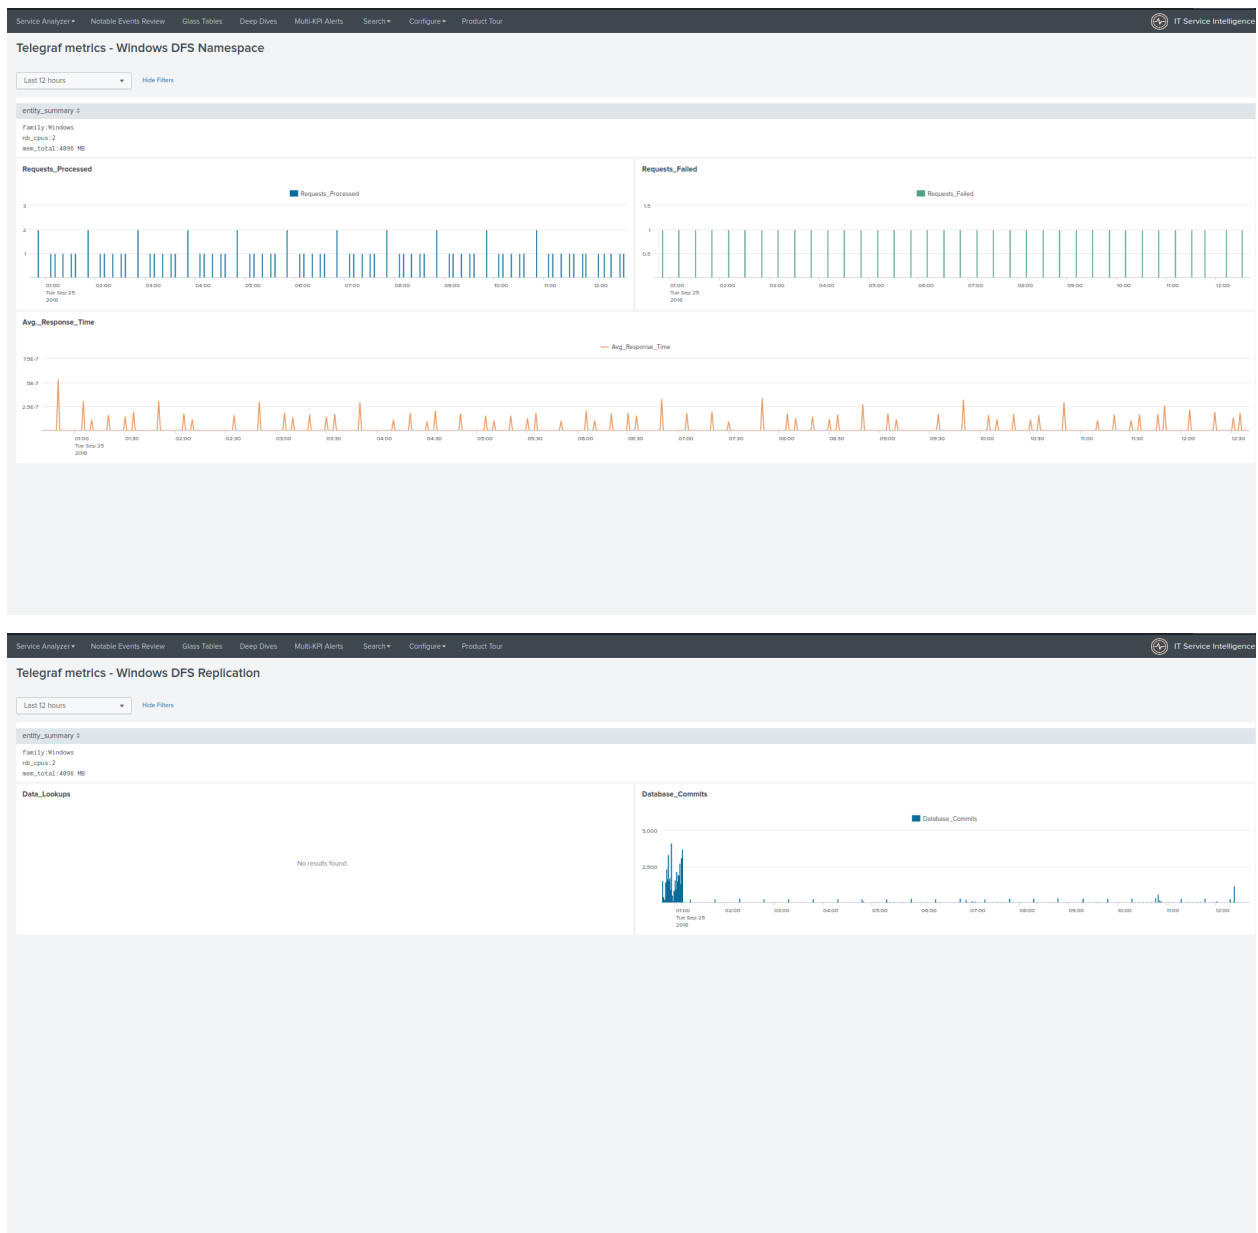
2.3.5 View for Windows DNS

The Health view automatically appears as “Telegraf Win AD-DC” deepdive drilldown link when entities are discovered:



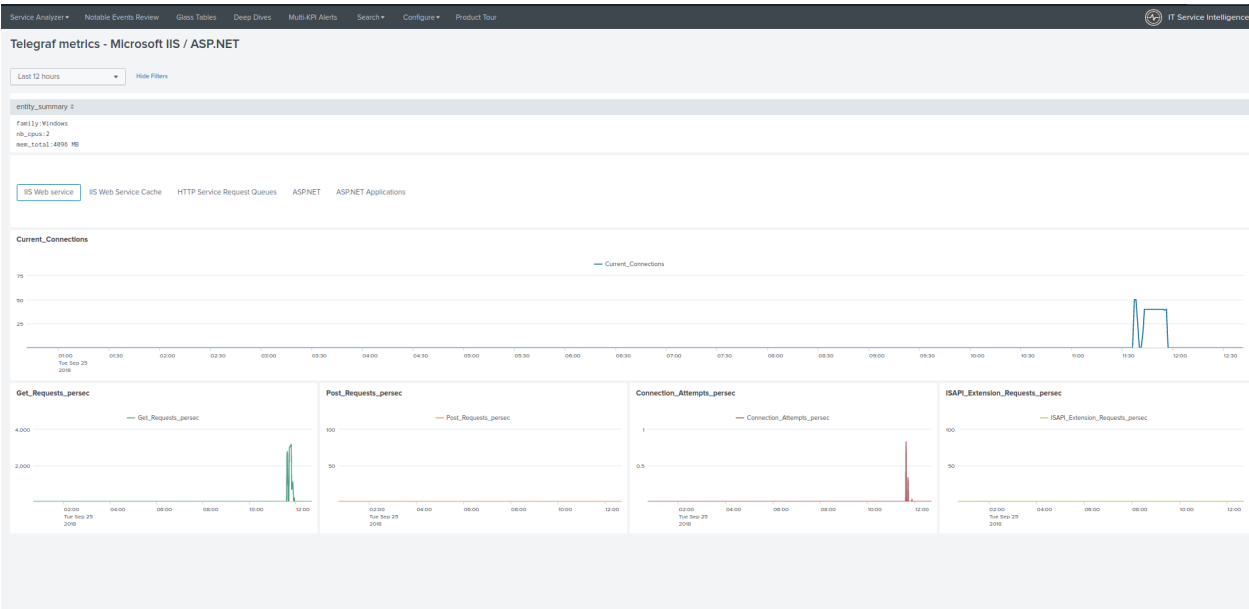
2.3.6 View for Windows DFS

The Health view automatically appear as “Telegraf Win DFS-NS” and “Telegraf Win DFS-REP” deepdive drilldown links when entities are discovered:



2.3.7 View for Microsoft IIS/ASP.NET

The Health view automatically appears as “Telegraf IIS / ASP.NET” deepdive drilldown links when entities are discovered:



3.1 Troubleshoot & FAQ

3.1.1 Missing metrics for Windows memory

For Windows memory management, the default win_mem inputs does not retrieve some of the metrics we need.

You need to activate the memory inputs. (which on Windows uses WMI collection):

```
[[inputs.mem]]  
# no configuration
```

3.1.2 Empty processes metrics (procstat)

In the linux views, the processes usage (both CPU and Memory) rely on the procstat inputs, which requires additional configuration depending on your context.

As for an example, the following configuration monitors all the processes owned by the “splunk” unix user:

```
[[inputs.procstat]]  
#   ## PID file to monitor process  
#   pid_file = "/var/run/nginx.pid"  
#   ## executable name (ie, pgrep <exe>)  
#   # exe = "nginx"  
#   ## pattern as argument for pgrep (ie, pgrep -f <pattern>)  
#   # pattern = "nginx"  
#   ## user as argument for pgrep (ie, pgrep -u <user>)  
#   user = "splunk"
```

Versioning and build history:

4.1 Release notes

4.1.1 Version 1.0.7

- Change: JQuery simple XML dashboard update

4.1.2 Version 1.0.6

- Feature: The App goes full dark mode
- Feature: Builtin out of the box alerts for basic generic system usage alerting (disk usage, cpu / memory saturation)
- Fix: Appinspect failure due to appLogo.png and appLogo_2X.png size

4.1.3 Version 1.0.5

- fix: improved nix family detection, avoid relying on kernel metrics which is Linux specific and might not be available if not configured

4.1.4 Version 1.0.4

- fix: Missing table_bar.cssl.js

4.1.5 Version 1.0.3

- fix: Missing objects and bad references

4.1.6 Version 1.0.2

- fix: hard coded references to savedsearches.conf

4.1.7 Version 1.0.1

- fix: hard coded references to index name in Overview

4.1.8 Version 1.0.0

- initial and first public release