Questions to Highlight

- **I would like to understand how elastic search is different from competitors like Splunk etc**
    - A good starting resource of Elastic vs Splunk can be found at https://www.elastic.co/splunk-replacement

- **Is it possible to visualize data from multiple elasticsearch from a single kibana instance. Where can I find more resources regarding this?**
    - Yes, we support Cross Cluster Search (CCS) https://www.elastic.co/guide/en/elasticsearch/reference/current/modules-cross-cluster-search.html
    -
- **If we could send log data directly from Filebeat to Elasticsearch, what would be the role of Logstash?**
    - Logstash is more of a full featured ETL. It can offer queuing, data enrichment, it has over 200 inputs/outputs with very rich enrichment capabilities that go beyond what Beats can natively do with a direct Elasticsearch output.

- **Is it any specific agent to monitor kubernetes pods?**
    - https://www.elastic.co/integrations?solution=all-solutions&category=kubernetes

- **How can I mask the data in the logs?**
    - There are several ways, one way is to use our Ingest Pipelines to enrich, drop, mutate the data as desired. See https://www.elastic.co/guide/en/elasticsearch/reference/current/pipeline.html
    - Couple solutions to this, one is to use ingest pipelines to mask specific fields on ingest. Another is restricting access at a field level https://www.elastic.co/guide/en/elasticsearch/reference/current/field-level-security.html

- **How to correlate data from different sources and indices?**
    - Today we recommend using the Elastic Common Schema (ECS) to ensure data can easily be correlated from various services. We're also working on implementing Event Query Language (EQL) starting with 7.9 to do more advanced correlation activities. See https://www.elastic.co/guide/en/elasticsearch/reference/current/eql.html

- **How is security log data secured from deletion and an integrity perspective?**
    - Security can be defined down to the index (and even field level) so only users with specific access to that security index would have the access to delete/edit.

- **How can we push Azure diagnostic logs to elastic? Here we can see only Activity and audit logs.**
    - See if this module helps your use case https://www.elastic.co/guide/en/beats/filebeat/current/filebeat-module-azure.html

- **Can Kibana do something like each application team will have their own dashboard or logs view for their own applications?**
    - Yes, we support as many dashboards, visualizations, etc as desired. Kibana also supports a concept of spaces which are logically separated views into the data. See https://www.elastic.co/guide/en/kibana/current/xpack-spaces.html

- **How to collect log files from applications running in Docker containers?**
    - We have a Filebeat module for reading logs from Docker containers (separate one for metrics) https://www.elastic.co/guide/en/beats/filebeat/current/filebeat-input-docker.html

----

Elastic vs Competitor Questions

- **I would like to understand how elastic search is different from competitor like Splunk etc**
    - A good starting resource of Elastic vs Splunk can be found https://www.elastic.co/splunk-replacement
- **How is it different from the kafka streaming of logs into a big data log analytics?**
    - That is certainly one deployment option here. Kafka plays well with Elasticsearch and is a common implementation architecture for streaming data into the solution Amy is showing. Elasticsearch + Kibana provides the analytics capabilities.
    -

Requirement/Compatibility Questions

- **port requirements?**
    - Port Requirements - https://www.elastic.co/guide/en/elasticsearch/reference/current/modules-network.html walks through how to set up listening ports in the yml, this will however depend if you are using our Elasticsearch service, or if you are self hosting, etc.
- **Do you have ones for Cisco and Palo Alto?**
    - We do have cisco/palo alto modules and many other. You can check https://www.elastic.co/integrations as a quick search reference
- **Multi-Tenant ?**
    - We support multi-tenant

- **is it compatible for MAC system too ?**
  - yes, mac,linux, windows
- **Is it compatible for   RBAC ?**
  - RBAC is supported
-

Architecture/Recommendation Questions

- **In case of distributed system also should we use Filebeat directly to send data to Elastic search or use logstash as collector?**
  - Generally, you can send logs directly to elasticsearch from filebeat (*it handles backoff/retry for example*).Your specific use case may differ of course
- **There are several options for ingestion of AWS data, key/secret, role assumption, etc. What is the Elastic-recommended authentication method when using the AWS adapters?**
  - This would depend greatly on the overall deployment architecture. I'd recommend reaching out to our pre-sales folks or your account team to discuss your architecture specifically.

Kibana Questions

- **Can kibana do something like each application team will have their own dashboard or logs view for their own applications?**
  - Yes, we support as many dashboards, visualizations, etc as desired. Kibana also supports a concept of spaces which are logically separated views into the data. See https://www.elastic.co/guide/en/kibana/current/xpack-spaces.html
- **how can you discover logs across indices ?** using kibana
- **Is there a way , when we bring up Kibana , it starts with already existing Index patterns -> basically we want to predefine the Dashboards , filters and make it readily available for our users? so is there a config file where we can pre-define it ?**
  - In the advanced section of Kibana settings, you can define a "landing url" which
  - you can have user's open up a specific dashboard for example when they log in
  - You can also define a Space that only has certain dashboards, index patterns set up. used a lot for specific business groups. eg a marketing space for the marketing team
- **Is it possible to visualize data from multiple elasticsearch from a single kibana instance. Where can I find more resources regarding this.**
  - Yes, we support Cross Cluster Search (CCS) https://www.elastic.co/guide/en/elasticsearch/reference/current/modules-cross-cluster-search.html
- **You can use discover to search across indices with Kibana index patterns eg app_log-***
  - The "index" you select in discover is a pattern that matches one or more indices
- **are these dashboards default or customized??**

- **Hi, is it possible to be added an automated job, for example if some anomalies or any kind of issue is detected on some service on VM that we are monitor through Kibana, then alert get triggered, so that service will be auto restarted?**
  - Our alerting framework supports webhooks as an action, so you can trigger any action that can be invoked via a webhook. See https://www.elastic.co/guide/en/kibana/current/managing-alerts-and-actions.html
- **Does Kibana support SAML for Authentication**
- **I seem to remember that Spaces each require their own index pattern definitions. Is that accurate or do I misremember?**
  - Each space does have its own Kibana index pattern that points to the underlying index (data). You can copy index patterns, dashboards between spaces
- **Last time I tried, to move/copy visualizations and dashboards between spaces required export/import and the result is completely separate viz/dashboards.**
  - That was the earlier behavior. The new method is a option from management section to "copy to space" so you don't have to manually export import

Logs/Other Questions

- **How to collect log files from application running in docker container?**
  - We have a filtbeat module for reading logs from docker containers (separate one for metrics) https://www.elastic.co/guide/en/beats/filebeat/current/filebeat-input-docker.html
- **Is it possible to enhance data on query time? for example: lookup the country-name based on the country code in a phone number field (stored in a csv file)? Or has this to happen always on index time?**
  - Generally we enrich data at index time which is one of the ways we achieve great speed, scale, etc. We can enrich several ways, using logstash or using or ingest pipeline framework or two common ways.
- **kubernetes logs with EFK?**
- **i have an aws application and my elastic stack on premise. How can i collect logs from the aws application?**
- **I have an application in python that writes 5 different log files. About 10MB/min rate. How do I make filebeats to read these logs and show me the dashboard?**
- **one agent means, no need to install individual beats to get specific data ??**
- **elastic fluentd and kibara setup for kubernetes operator?**
- **Is there a module for Oracle DBs? Is there a module for Microsoft SQL?**
  - We have several modules available for logs/metrics. Beats is also an open and extensible. There is a strong community of contirbuters. https://www.elastic.co/guide/en/beats/metricbeat/current/metricbeat-modules.html and https://www.elastic.co/guide/en/beats/filebeat/current/filebeat-modules.html are good references.
- **how can we push Azure diagnostic logs to elastic. here we can see only Activity and audit logs .**
- **kubernetes clusters can be monitor ??**

- **I wanted to learn log analysis ...where to start...splunk or elastic ?**
- **For applications that are running in Azure App Service or Azure Function; how can we install beats?**
- **which parser, you people recommend for performance point of view, lagstash, fluentd and filebeat?**
  - If you need to parse/enrich data beyond the native capabilities of beats, we would recommend either our ingest pipeline feature within Elasticsearch or Logstash.
- **There are several options for ingestion of AWS data, key/secret, role assumption, etc. What is the Elastic-recommended authentication method when using the AWS adapters?**
- **can you please show an example for exporting and filtering Java exceptions**
- **who is doing logs parsing in your example?**
- **How file beat can be used to parse logs from individual batch processes?**
- **If we could send log data directly from Filebeat to Elasticsearch, what would be the role of Logstash?**
  - Logstash is more of a full featured ETL. It can offer queuing, data enrichment, it has over 200 inputs/outputs with very rich enrichment capabilities that go beyond what beats can natively do with a direct Elasticsearch output.
- **so u mean to search logs across indices -> we should keep name of all indices similar like Index_one, Index_two , so as to make patern as index* and search across indices ?**
  - If you have index like:
  - logs-app1
  - logs-app2
  - logs-app3
  - But in Kibana you define an index template for logs-*, you can search across multiple that way.
  - depends on your search use case but you can search with “name-*” pattern or “index1,index2” as an example
- **calculation of license (EPS) ?**
- **these pahses of rolling , thewould be on the same disk or they are s3 or some thng else**
- **So the anomaly detection would work only on timeseries part of the logs? Any deviation from the regular section of the flow would be flagged?**
- **Hi, is it possible to be added an automated job, for example if some anomalies or any kind of issue is detected on some service on VM that we are monitor through Kibana, then alert get triggered, so that service will be auto restarted?**
- **is one apm agent is enough to monitor e2e application??**
- **How can we push Azure diagnostic logs to elastic. Here we can see only Activity and audit logs .**
  - See if this module helps your use case https://www.elastic.co/guide/en/beats/filebeat/current/filebeat-module-azure.html

- **is there any functionality to automatically move indexes older than N days to an external storage such as Hadoop?**
  - Today the phases are really based on type of disk and disk to ram ratio, so that as the data ages it can age to slower/cheaper disk, rolled up for metric data, reduce replicas if desired, and ultimately delete if it has aged beyond it's desired age.
- **How is security log data secured from deletion and an integrity perspective?**
  - Security can be definined down to the index (and even field level) so only users with specific access to that security index would have the access to delete/edit …
- **How to correlate date from different sources and indices?**
  - Today we recommend using the Elastic Common Schema (ECS) to ensure data can easily be coorelated from various services. We're also working on implementing Event Query Language (EQL) starting with 7.9 to do more advanced correlation activities. See https://www.elastic.co/guide/en/elasticsearch/reference/current/eql.html
- **How can we read / parse logs for individual batch process deployed on linux using file beat?**
  - you can stream logs from the individual jobs, or if they are all in the same log you can group by the job name for example. lots of options depending on the specifics
- **How can I mask the data in the logs**
  - There are several ways, one way is to use our Ingest Pipelines to enrich, drop, mutate the data as desired. See https://www.elastic.co/guide/en/elasticsearch/reference/current/pipeline.html
  - Couple solutions to this, one is to use ingest pipelines to mask specific fields on ingest. another is restricting access at a field level
- **are these dashboards are inbuilt or customized?**
  - We have several out of the box dashboards and curated UI experiences. Dashboards can also be modified, custom built, and even branded to your own company using Canvas.
- **Is it any specific agent to monitor kubernetes pods?**
  - https://www.elastic.co/integrations?solution=all-solutions&category=kubernetes
- **is one apm agent is enough to monitor e2e application??**
  - https://www.elastic.co/ has more information about all the aspects of elastic stack, links to blogs and webinars, etc
- **When we send logs via Logstash or any forwarder to Elastic search, Dynamic creation of Fields happen , Is that a good way or shall we have a pre-define schema for our indices? which one is good approach in terms of memory , performance and ambiguity?**
- **Is there any new feature for logstash, which can automatically write grok patterns ?**
- **when the ilm rollover? my apm-server indices up to ilm policy level but ilm not work immediately**

- **I am using Kiana and Elastic, not for monitor or analizing logs, but for ingesting business data from our ERP and present information to business. Is there a place where I can find any about other use cases out of metrics, logs and security?**
    - This may help - https://www.elastic.co/what-is/elasticsearch-business-analytics
- **Is there any feature, to automatically write grok patterns based on the data input sending to ELK?**
    - In the ML section, there is a "Data Visualizer" where you can upload a sample csv or select an index pattern and part of that will attempt to create a grok pattern (you may need to tweak it because … grok) Once created you can attach it to an ingest pipeline that logstash can sent the data through (or just take the grok and put it in your config) blog post - https://www.elastic.co/blog/importing-csv-and-log-data-into-elasticsearch-with-file-data-visualizer
- when the ilm rollover? my apm-server indices up to ilm policy level but ilm not work immediately
    -