

LAB 4

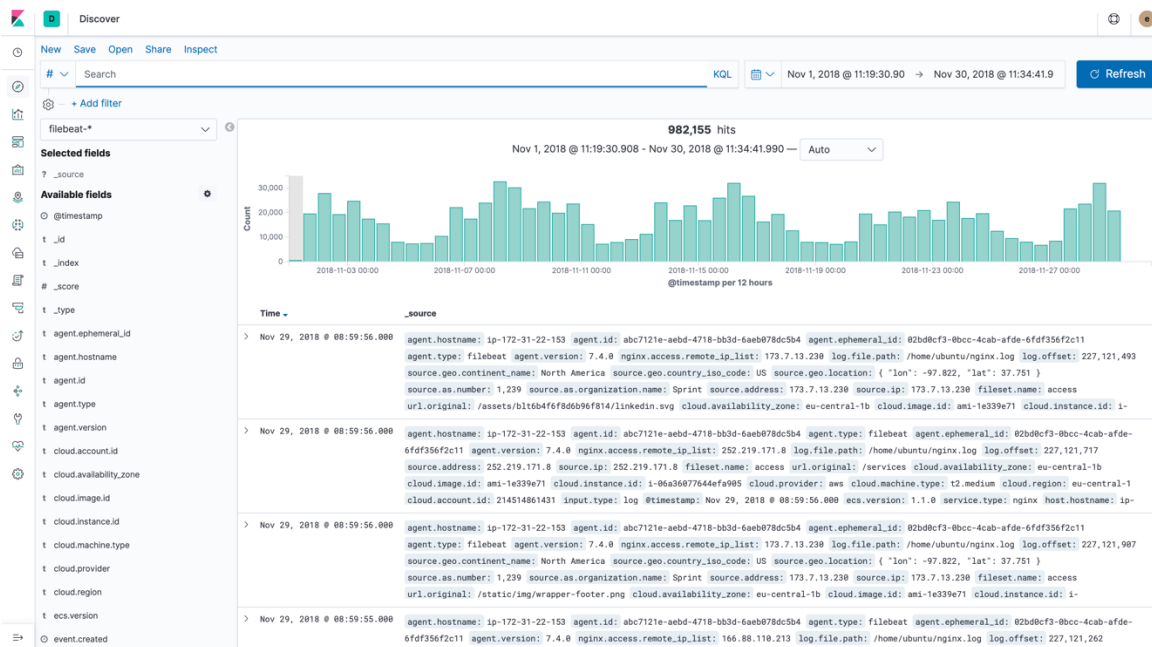
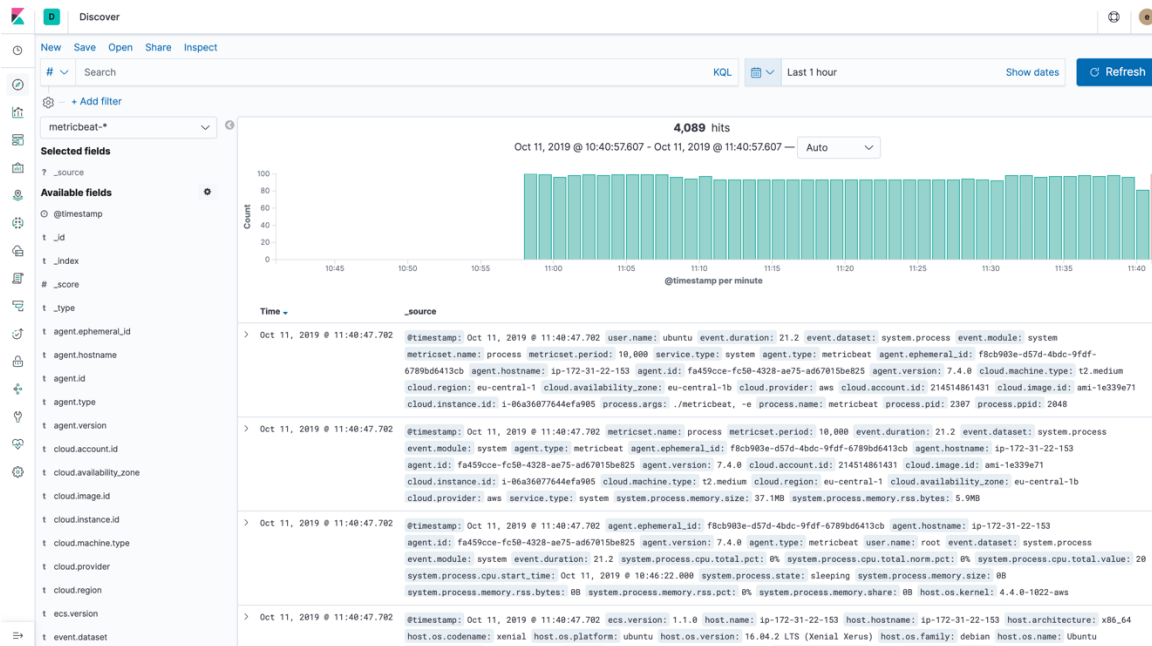
はじめに

本 Lab では、Elastic Stack のそのほかの高度な機能について体験していただきます。実際に体験していただくのはセキュリティ、アラート、機械学習になります。

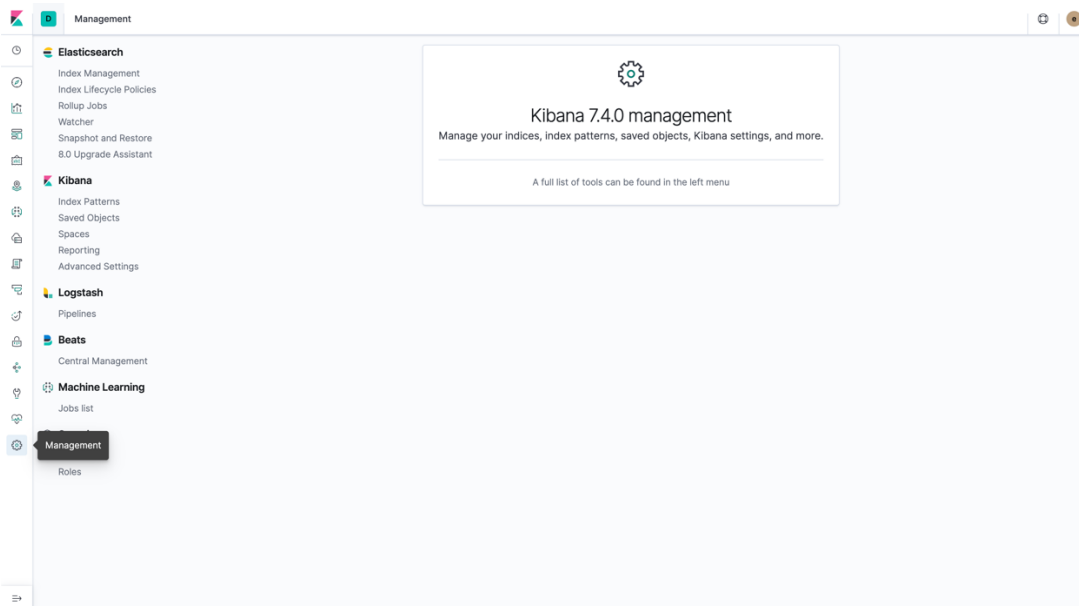
機能	概要
セキュリティ	適切なユーザーに適切なアクセス権限を付与する機能、それが Elastic Stack の Security です。たとえば IT 部門や事業部、アプリ開発チームごとに善意のユーザーを管理することで、不正な侵入を防ぐことができます。また、お客様や経営陣にも個別の権限を付与して、Elastic Stack のデータを安全に、安心して共有できます。
アラート	Elasticsearch のクエリ機能をフルパワーで活用する Elastic Stack の Alerting なら、データの重要な変化を見逃しません。つまり、Elasticsearch でクエリできるものは何でも通知可能です。
機械学習	トレンドや周期性などからデータの振る舞いを自動的に、リアルタイムにモデル化し、すばやく問題を特定して原因分析を手助けします。さらに、誤検出を防ぎます。

Security

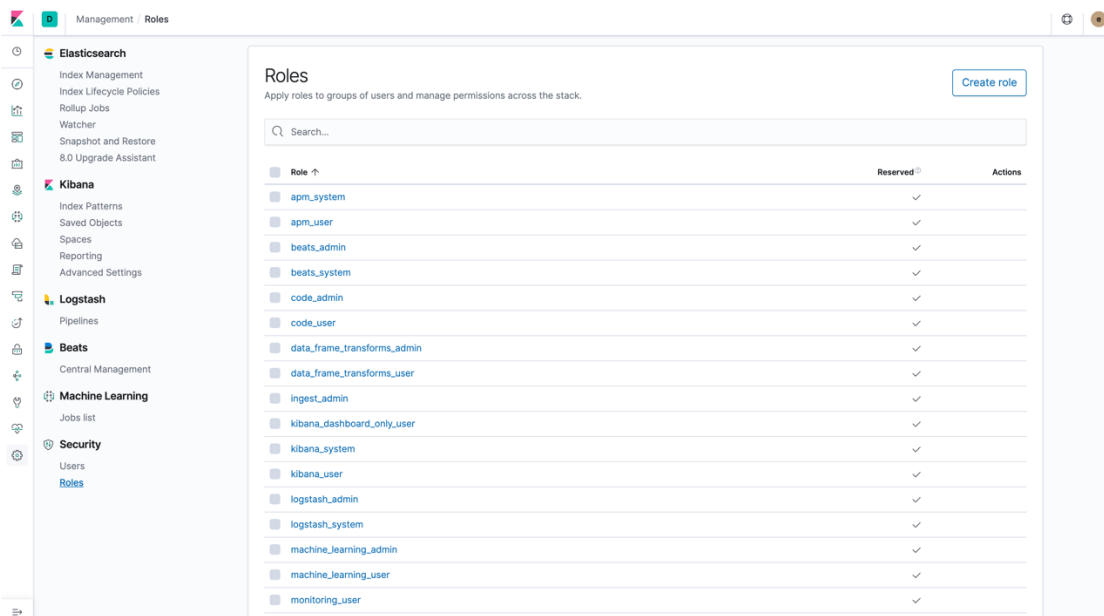
1. Kibana の menu から“Discover” を選択します。“filebeat-*” か “metricbeat-*” index pattern を選択します。“filebeat-*” を選択した場合、日付のレンジを “Nov 1st, 2018 – Nov 28th, 2018” として下さい。“metricbeat-*” を選択した場合、適切な日付のレンジに調整して下さい。



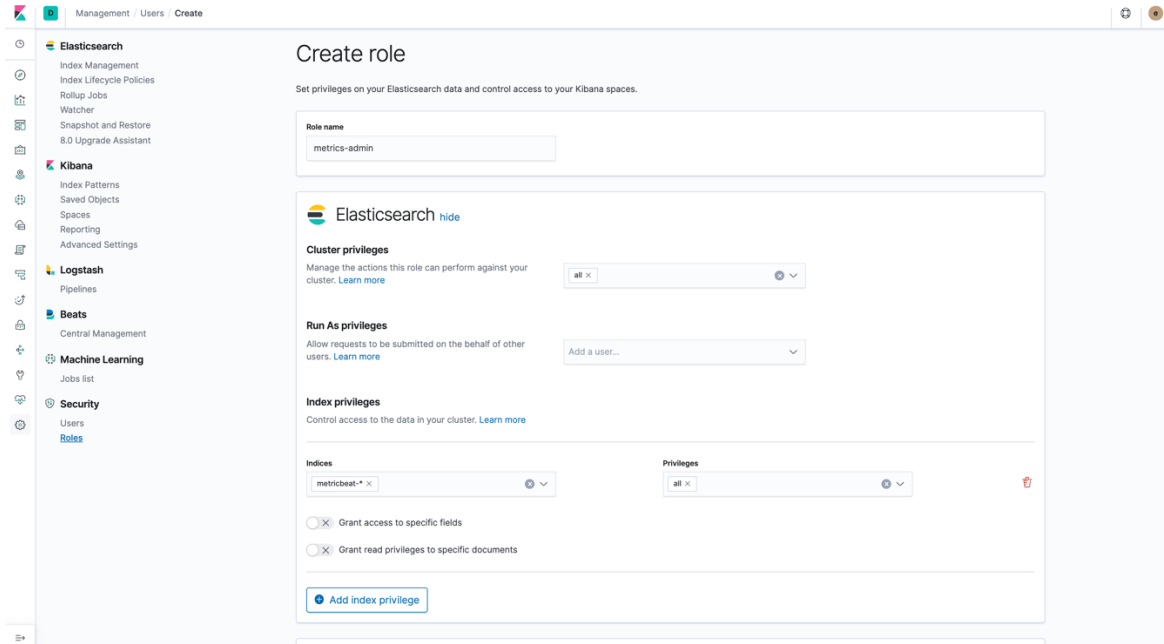
- もしこれらのユースケースが異なるグループ（メトリクスグループとロググループ）に別れていて、例えばコンプライアンス上、メトリクスグループにはログを見せてはいけないとしたら、どうしますか？そこで Elastic Security の出番です。“Management” を Kibana menu から選択します。



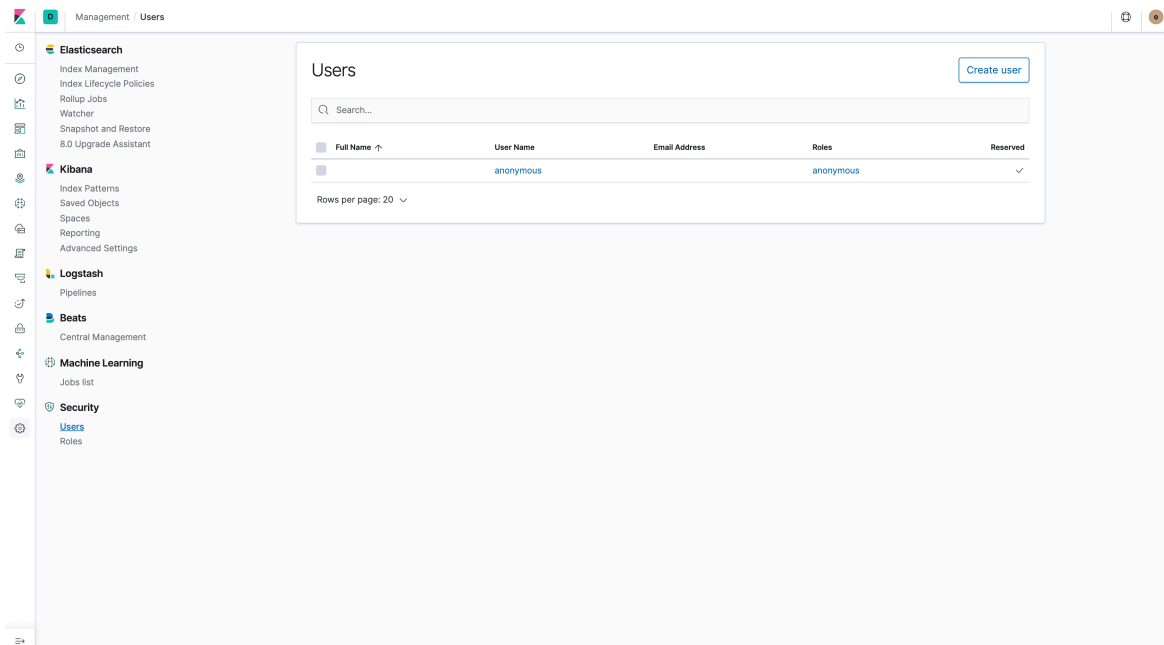
- Security > Roles から“Create Role”をクリックします。



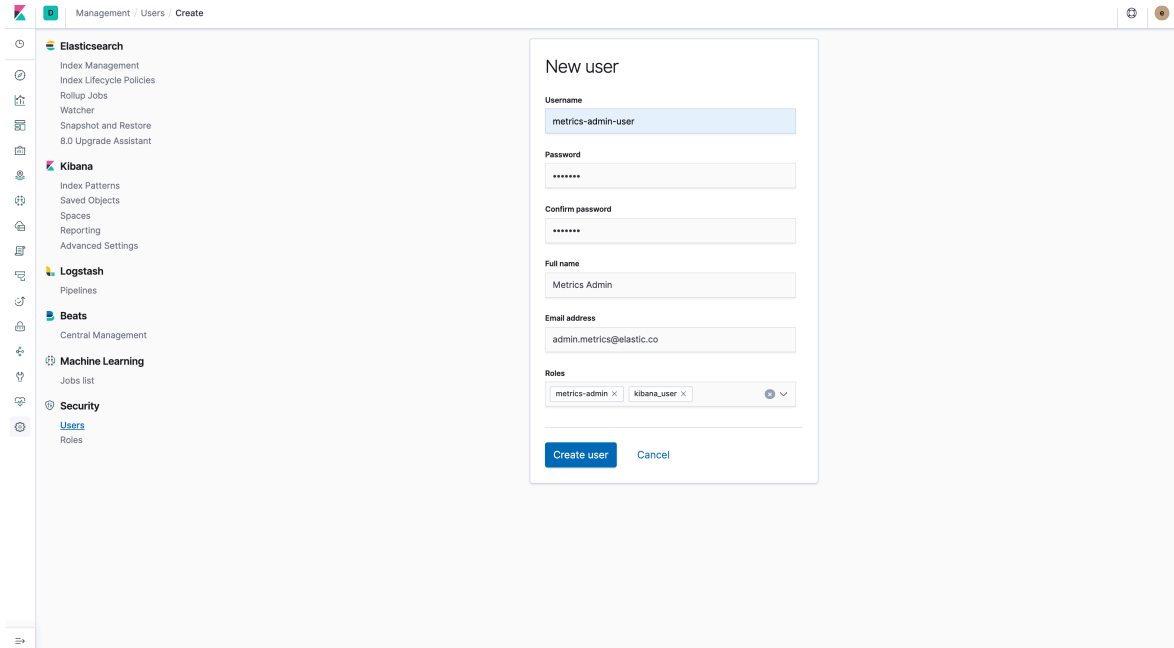
- 新しい role に名前をつけます (“metrics-admin” for example)。Cluster privileges で “all” を選択し、Indices は “metricbeat-*” を選択します。Index privileges は “all” を選択します。最後に “Create role” をクリックします。



- Role が作成されたら、“Management” タブで Security > Users から “Create new user” をクリックします。



6. Username と password, full name, email を入力します。Roles には、“kibana-user” と先ほど作成した“metrics-admin”を追加します。最後に “Create User” をクリックします。

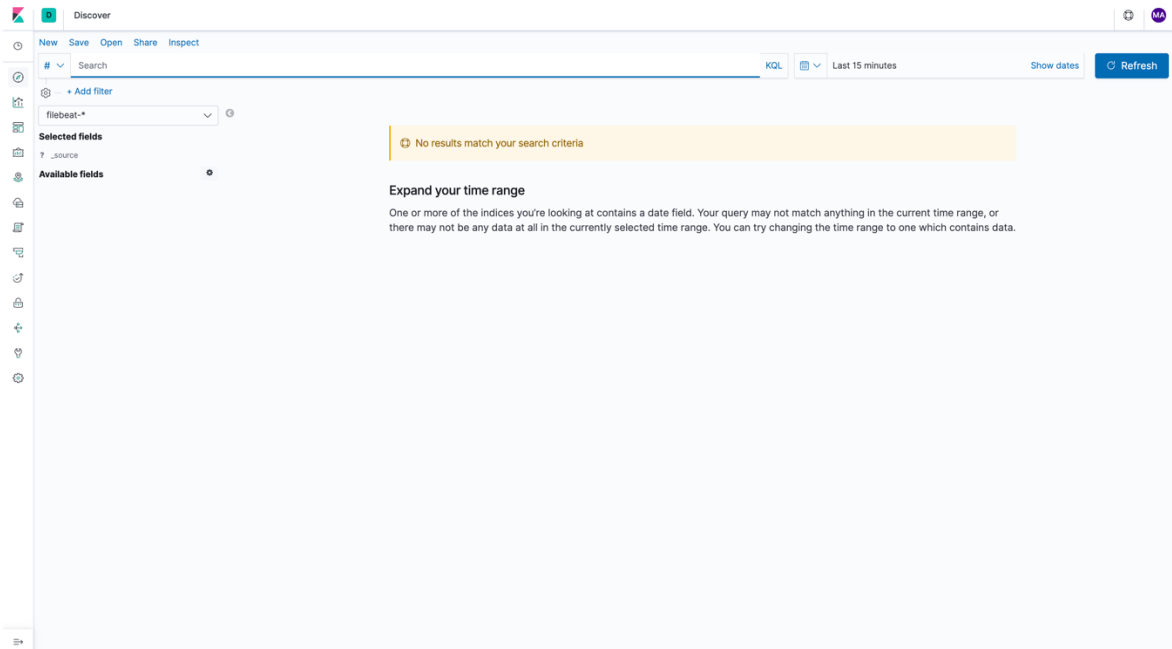


The screenshot shows the 'New user' form in the Elastic Management console. The form is titled 'New user' and is located in the 'Management / Users / Create' section. The form fields are as follows:

- Username:** metrics-admin-user
- Password:** [masked with dots]
- Confirm password:** [masked with dots]
- Full name:** Metrics Admin
- Email address:** admin.metrics@elastic.co
- Roles:** metrics-admin x kibana_user x

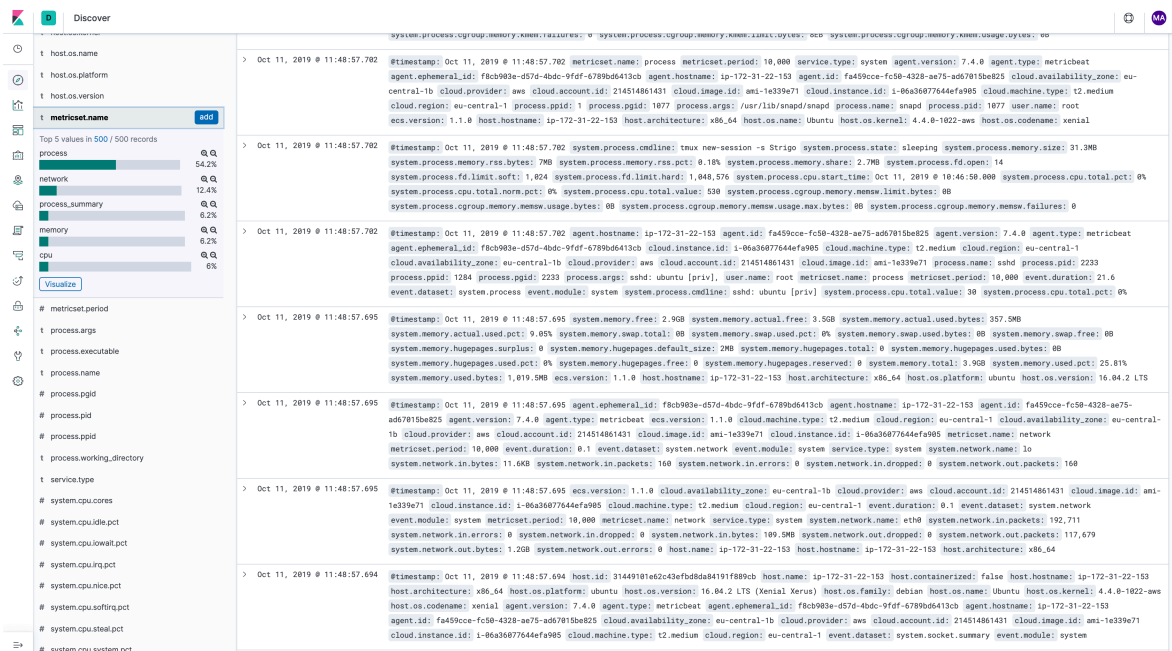
At the bottom of the form, there are two buttons: 'Create user' and 'Cancel'.

7. 別のブラウザ画面を開き、新しく作成したユーザーで Kibana にログインします。“Discover” タブを開いてみましょう。Index pattern で“metricbeat-*” しか全てのデータを見ることのできないことに気づきましたか？ “filebeat-*” index pattern では、データを見ることはできません。



index pattern “filebeat-*” 自体は見る事ができます。

- “metricbeat-*” index pattern を再度選択します。“metricset.name” field をクリックし、収集したメトリクスの種類を確認します。

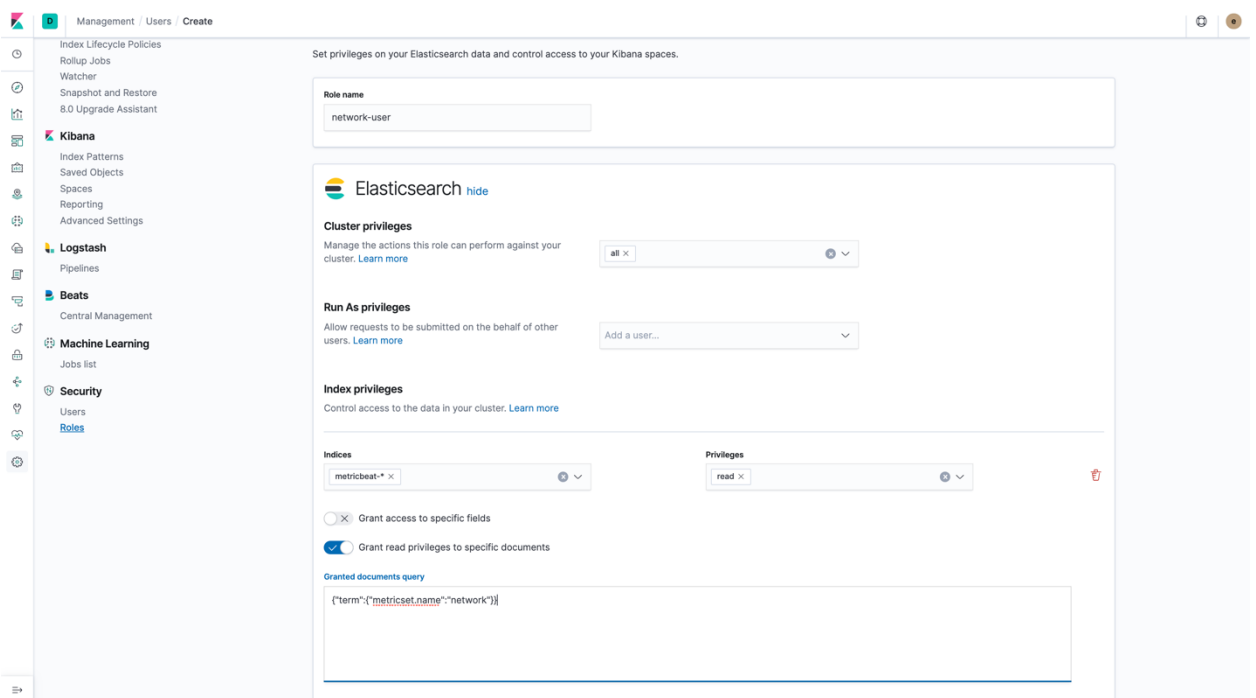


9. 一例として“network”というものがあります。例えば、ネットワークオペレータはネットワークデータに関するデータだけ見ることができて、他を見ることはできない、というシナリオを想定してみましょう。ドキュメントレベルの属性ベースのアクセスコントロール(document level security)を提供すべきでしょうか？再度“elastic”ユーザーでログインし（または、既にかいているブラウザ画面に戻り）、“network-user”という role を作成しましょう。

以下の設定を行います：

- Cluster privileges : all
- Indices: metricbeat-*
- Index Privileges: read
- Click on: **Grant read privileges to specific documents**
- 次の文字列を入力: `{"term":{"metricset.name":"network"}}`

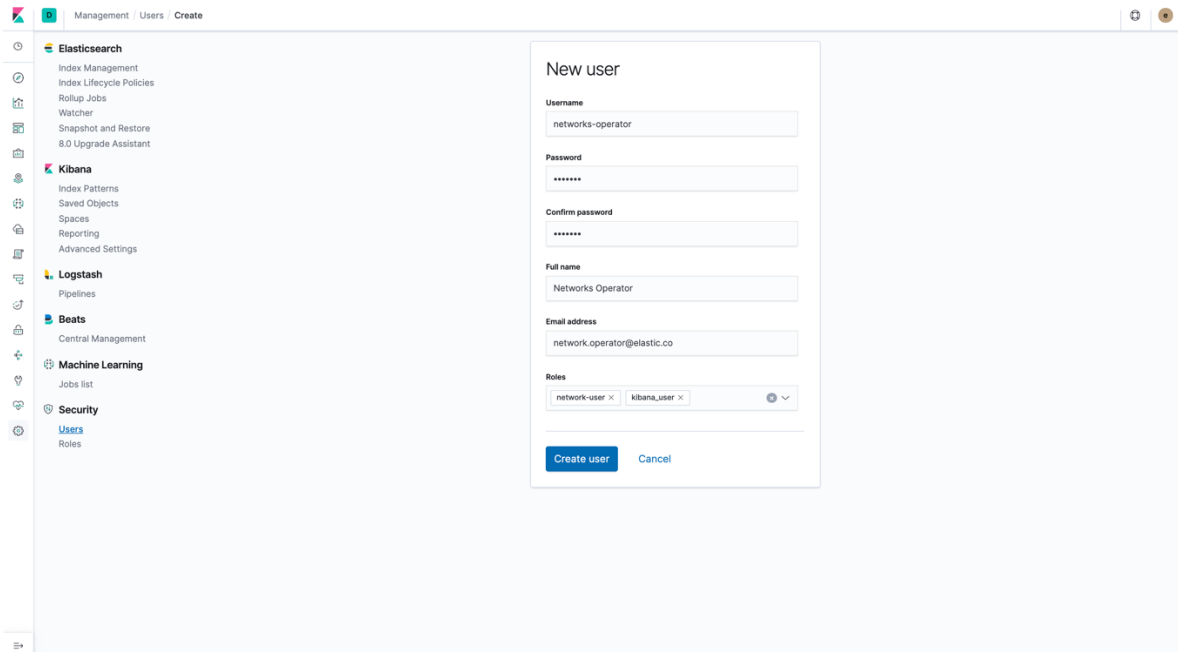
最後に“Create Role”をクリックします。



The screenshot shows the Kibana 'Create Role' configuration page. The role name is 'network-user'. The configuration includes:

- Cluster privileges:** all
- Run As privileges:** Add a user...
- Index privileges:** metricbeat-* with read privileges
- Index privileges options:**
 - Grant access to specific fields
 - Grant read privileges to specific documents
- Granted documents query:** `{\"term\":{\"metricset.name\":\"network\"}}`

10. 新しく作成した role に紐づく user を作成します。“kibana-user” を必ず追加して下さい。これを忘れると、新しく作成した user は Elasticsearch APIs にアクセスすることができません。



11. 新しく作成した user で再度ログインします。Discover から “metricbeat-*” index pattern を選択します。“metricset.name” field をクリックして展開してみましょう。“network”.だけが見えることがわかりますか？他の値、metrics-admin では見ることができたものは、この user では見ることができません。



12. (Optional) Elastic Stack の Security 機能、Field Level Security をもう少し見てみましょう。elastic user でログインし、“filebeat-*” index pattern を選択します。日付の範囲を between Nov 1st, 2018 - Nov 28th, 2018 とします。“nginx.access.remote_ip_list”, “source.address”, “source.ip” をクリックしてみてください。これらの fields はアクセスした user の IP です。この情報は admin user のみがアクセスできるようにしたい場合を考えてみましょう。

The screenshot shows the Kibana Discover interface. On the left, the 'nginx.access.remote_ip_list' field is selected, and a bar chart displays the top 5 values in 500 records. The main panel shows a list of log entries with their corresponding JSON documents. The selected field 'source.address' is highlighted in blue.

どのように解決できるでしょうか？今回は、Elasticsearch の API の API を使って新しい role を作成してみましょう。Kibana から“Dev Tools”をクリックし、以下のコードをペーストして下さい。

```
PUT _security/role/filebeat-restricted
{
  "cluster": [
    "all"
  ],
  "indices": [
    {
      "names": [
        "filebeat-*"
      ],
      "privileges": [
        "read"
      ],
      "field_security": {
        "grant": [
          "*"
        ],
        "except": [
          "source.address",
          "source.ip",
          "nginx.access.remote_ip_list"
        ]
      }
    }
  ],
  "run_as": [],
  "metadata": {},
  "transient_metadata": {
    "enabled": true
  }
}
```

API call を実行します。

```
1 PUT _security/role/filebeat-restricted
2 {
3   "cluster": [
4     "all"
5   ],
6   "indices": [
7     {
8       "names": [
9         "filebeat-*"
10      ],
11      "privileges": [
12        "read"
13      ],
14      "field_security": {
15        "grant": [
16          "*"
17        ],
18        "except": [
19          "source.address",
20          "source.ip",
21          "nginx.access.remote_ip_list"
22        ]
23      }
24    },
25  ],
26  "run_as": [],
27  "metadata": {},
28  "transient_metadata": {
29    "enabled": true
30  }
31 }
32
```

```
1- {
2-   "role": {
3-     "created": true
4-   }
5- }
6-
```

Role を作成したら、“site-operator”という user を作成し、先ほど作成した role “filebeat-restricted”を割り当てます。

Management / Users / Create

Elasticsearch
Index Management
Index Lifecycle Policies
Rollup Jobs
Watcher
Snapshot and Restore
8.0 Upgrade Assistant

Kibana
Index Patterns
Saved Objects
Spaces
Reporting
Advanced Settings

Logstash
Pipelines

Beats
Central Management

Machine Learning
Jobs list

Security
Users
Roles

New user

Username:

Password:

Confirm password:

Full name:

Email address:

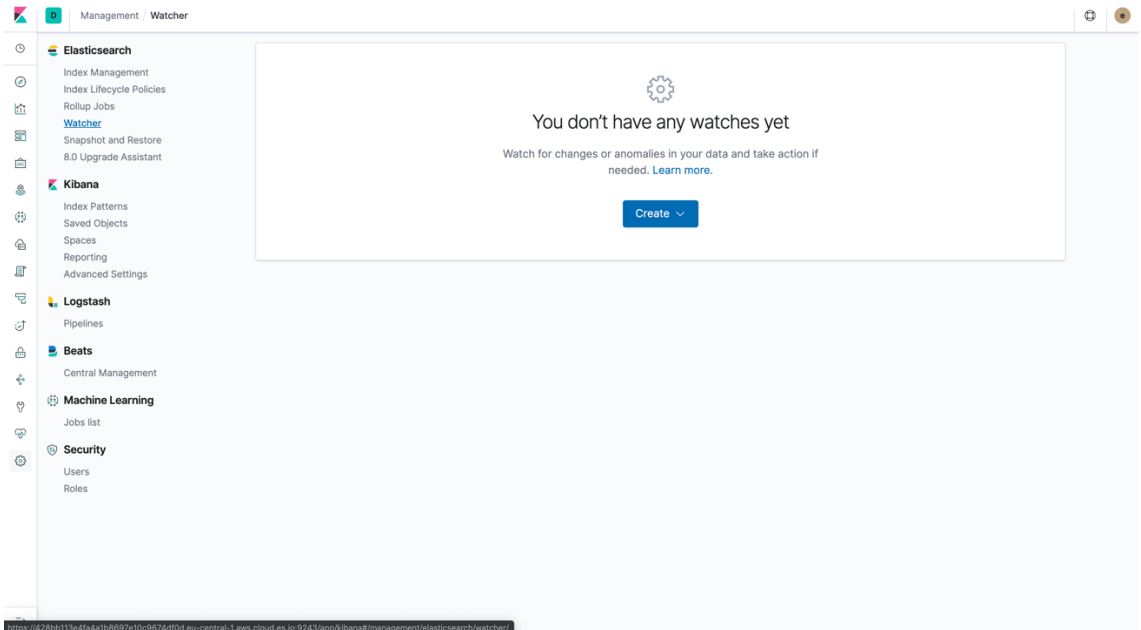
Roles:

新しく作成した user でログインし、“Discover” から “filebeat-*” index pattern を選択します。IP fields が見えなくなっています。

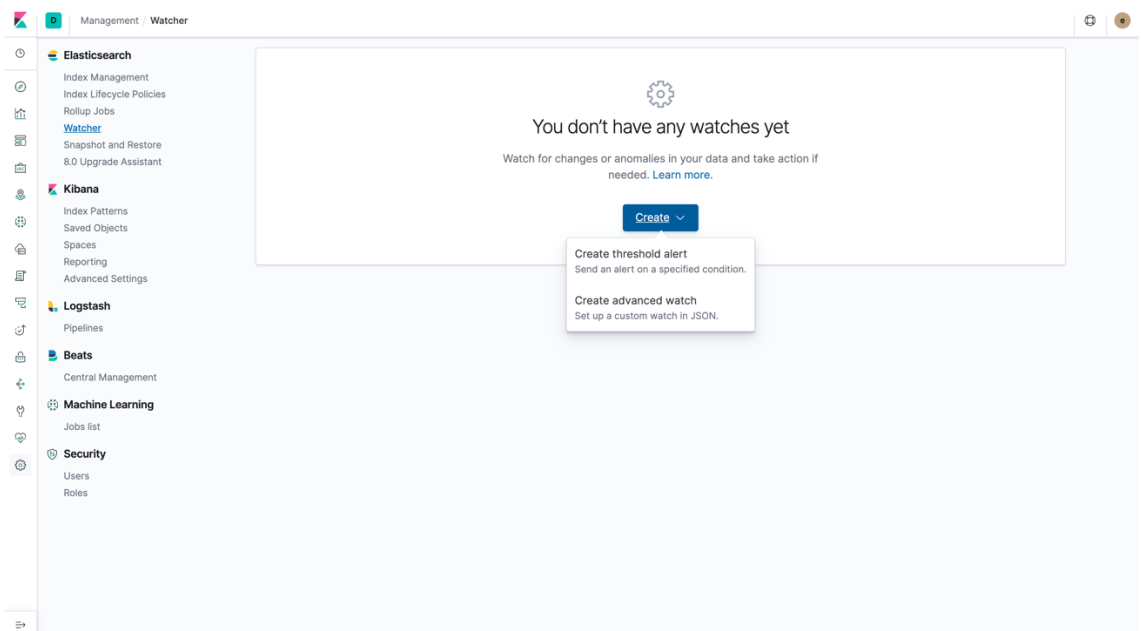
Field	Value
t host.os.name	url.original: /assets/blta76b2ae2b2f1e998/xing.svg cloud.availability_zone: eu-central-1b cloud.image.id: ami-1e339e71 cloud.instance.id: i-06a36077644ef995 cloud.provider: aws cloud.machine.type: t2.medium cloud.region: eu-central-1 cloud.account.id: 214514861431 input.type: log
t host.os.platform	@timestamp: Nov 29, 2018 @ 08:59:45.000 ecs.version: 1.1.0 service.type: nginx host.hostname: ip-172-31-22-153 host.os.kernel: 4.4.0-1022-aws
t host.os.version	> Nov 29, 2018 @ 08:59:45.000 agent.hostname: ip-172-31-22-153 agent.id: abc7121e-aebd-4718-bb3d-6aeb078dc5b4 agent.type: filebeat agent.ephemeral_id: 02bd0cf3-0bcc-4cab-afde-6fdf356f2c11 agent.version: 7.4.0 log.file.path: /home/ubuntu/nginx.log log.offset: 227,120,575 fileset.name: access
t http.request.method	url.original: /assets/bltfd3b1511512f4632/twitter.svg cloud.availability_zone: eu-central-1b cloud.image.id: ami-1e339e71 cloud.instance.id: i-06a36077644ef995 cloud.provider: aws cloud.machine.type: t2.medium cloud.region: eu-central-1 cloud.account.id: 214514861431 input.type: log
t http.request.referrer	@timestamp: Nov 29, 2018 @ 08:59:45.000 ecs.version: 1.1.0 service.type: nginx host.hostname: ip-172-31-22-153 host.os.kernel: 4.4.0-1022-aws
# http.response.body.bytes	> Nov 29, 2018 @ 08:59:41.000 agent.hostname: ip-172-31-22-153 agent.id: abc7121e-aebd-4718-bb3d-6aeb078dc5b4 agent.ephemeral_id: 02bd0cf3-0bcc-4cab-afde-6fdf356f2c11 agent.type: filebeat agent.version: 7.4.0 log.file.path: /home/ubuntu/nginx.log log.offset: 227,120,117 source.geo.continent_name: North America
# http.response.status_code	agent.type: filebeat agent.version: 7.4.0 log.file.path: /home/ubuntu/nginx.log log.offset: 227,119,683 source.geo.continent_name: Asia source.geo.region_iso_code: US-MO source.geo.city_name: Cameron source.geo.country_iso_code: US source.geo.region_name: Missouri
t http.version	source.geo.location: { "lon": -94.2228, "lat": 39.734 } source.as.number: 393,442 source.as.organization_name: United Electric Cooperative
t input.type	fileset.name: access url.original: /static/img/wrapper-footer.png cloud.availability_zone: eu-central-1b cloud.image.id: ami-1e339e71
t log.file.path	> Nov 29, 2018 @ 08:59:40.000 agent.hostname: ip-172-31-22-153 agent.id: abc7121e-aebd-4718-bb3d-6aeb078dc5b4 agent.ephemeral_id: 02bd0cf3-0bcc-4cab-afde-6fdf356f2c11 agent.type: filebeat agent.version: 7.4.0 log.file.path: /home/ubuntu/nginx.log log.offset: 227,119,683 source.geo.continent_name: Asia
# log.offset	source.geo.country_iso_code: IN source.geo.location: { "lon": 77, "lat": 20 } source.as.number: 7,633 source.as.organization_name: Software Technology Parks of India - Bangalore fileset.name: access url.original: /assets/blt6e0d3e570896cfb/logo-elastic-1000x343.png cloud.availability_zone: eu-central-1b cloud.image.id: ami-1e339e71 cloud.instance.id: i-06a36077644ef995 cloud.provider: aws
t service.type	cloud.availability_zone: eu-central-1b cloud.image.id: ami-1e339e71 cloud.instance.id: i-06a36077644ef995 cloud.provider: aws
# source.as.number	> Nov 29, 2018 @ 08:59:40.000 agent.hostname: ip-172-31-22-153 agent.id: abc7121e-aebd-4718-bb3d-6aeb078dc5b4 agent.type: filebeat agent.ephemeral_id: 02bd0cf3-0bcc-4cab-afde-6fdf356f2c11 agent.version: 7.4.0 log.file.path: /home/ubuntu/nginx.log log.offset: 227,119,894 source.geo.continent_name: Asia
t source.as.organization.name	source.geo.country_iso_code: IN source.geo.location: { "lon": 77, "lat": 20 } source.as.number: 7,633 source.as.organization_name: Software Technology Parks of India - Bangalore fileset.name: access url.original: /assets/bltfd3b1511512f4632/twitter.svg cloud.availability_zone: eu-central-1b cloud.image.id: ami-1e339e71 cloud.instance.id: i-06a36077644ef995 cloud.provider: aws cloud.machine.type: t2.medium cloud.region: eu-central-1
t source.geo.city_name	> Nov 29, 2018 @ 08:59:36.000 agent.hostname: ip-172-31-22-153 agent.id: abc7121e-aebd-4718-bb3d-6aeb078dc5b4 agent.ephemeral_id: 02bd0cf3-0bcc-4cab-afde-6fdf356f2c11 agent.type: filebeat agent.version: 7.4.0 log.file.path: /home/ubuntu/nginx.log log.offset: 227,119,295 source.geo.continent_name: Asia
t source.geo.continent_name	source.geo.region_iso_code: CN-BJ source.geo.country_iso_code: CN source.geo.region_name: Beijing source.geo.location: { "lon": 116.3883, "lat": 39.9289 } source.as.number: 17,964 source.as.organization_name: Beijing Dian-Xin-Tong Network Technologies Co., Ltd. fileset.name: access
t source.geo.country_iso_code	url.original: /assets/blt6e0d3e570896cfb/logo-elastic-1000x343.png cloud.availability_zone: eu-central-1b cloud.image.id: ami-1e339e71
t source.geo.location	> Nov 29, 2018 @ 08:59:36.000 agent.hostname: ip-172-31-22-153 agent.id: abc7121e-aebd-4718-bb3d-6aeb078dc5b4 agent.type: filebeat agent.ephemeral_id: 02bd0cf3-0bcc-4cab-afde-6fdf356f2c11 agent.version: 7.4.0 log.file.path: /home/ubuntu/nginx.log log.offset: 227,119,439 source.geo.continent_name: Asia
t source.geo.region_iso_code	source.geo.region_iso_code: CN-BJ source.geo.country_iso_code: CN source.geo.region_name: Beijing source.geo.location: { "lon": 116.3883, "lat": 39.9289 } source.as.number: 17,964 source.as.organization_name: Beijing Dian-Xin-Tong Network Technologies Co., Ltd. fileset.name: access
t source.region_name	url.original: /static/img/wrapper-footer.png cloud.availability_zone: eu-central-1b cloud.image.id: ami-1e339e71 cloud.instance.id: i-
o suricata.event.timestamp	
t url.original	
t user.name	
t user_agent.device.name	
t user_agent.name	
t user_agent.original	

Alerting

1. Kibana の “Management” から、“Watcher” をクリックします。



2. まず、シンプルな threshold alert を作成しましょう。“Create threshold alert” ボタンをクリックします。



3. 名前を付けて、以下を入力します。

Indices to query: metricbeat-*

Time filed: @timestamp

Run watch every: 5 min

Matching condition に count() を選択して、直近 5 分間の全てのドキュメントに対して、threshold ラインを上回るように閾値(IS ABOVE)を調整してみましょう。

Management / Watcher / Create

Elasticsearch

- Index Management
- Index Lifecycle Policies
- Rollup Jobs
- Watcher
- Snapshot and Restore
- 8.0 Upgrade Assistant

Kibana

- Index Patterns
- Saved Objects
- Spaces
- Reporting
- Advanced Settings

Logstash

- Pipelines

Beats

- Central Management

Machine Learning

- Jobs list

Security

- Users
- Roles

Create threshold alert

Send an alert when your specified condition is met. Your watch will run every 5 minutes.

Name: metrics-threshold

Indices to query: metricbeat-*

Time field: @timestamp

Run watch every: 5 minutes

Match the following condition

WHEN count() OVER all documents IS ABOVE 50 FOR THE LAST 5 minutes

Perform 0 actions when condition is met

Create alert Cancel

次に average of system.network.in.bytes を直近 5 分間の全てのドキュメントに対して選択し、threshold ラインを上回るように閾値(IS ABOVE)を調整してみましょう。

Management / Watcher / Create

Elasticsearch

- Index Management
- Index Lifecycle Policies
- Rollup Jobs
- Watcher
- Snapshot and Restore
- 8.0 Upgrade Assistant

Kibana

- Index Patterns
- Saved Objects
- Spaces
- Reporting
- Advanced Settings

Logstash

- Pipelines

Beats

- Central Management

Machine Learning

- Jobs list

Security

- Users
- Roles

Create threshold alert

Send an alert when your specified condition is met. Your watch will run every 5 minutes.

Name: metrics-threshold

Indices to query: metricbeat-*

Time field: @timestamp

Run watch every: 5 minutes

Match the following condition

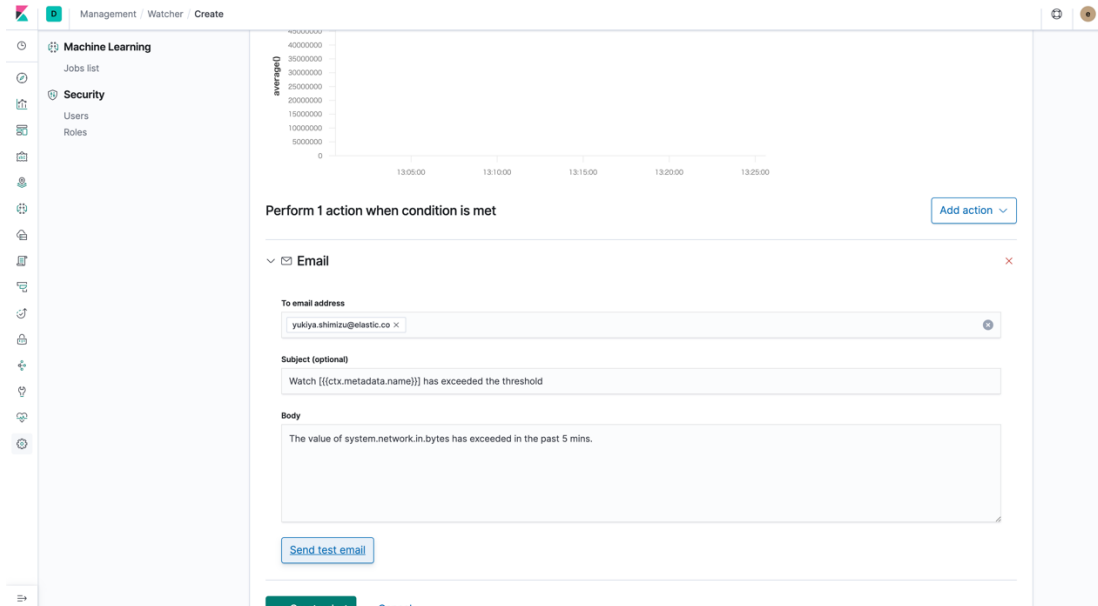
WHEN average() OF system.network.in.bytes OVER all documents IS ABOVE 55000000 FOR THE LAST 5 minutes

Perform 0 actions when condition is met

Create alert Cancel

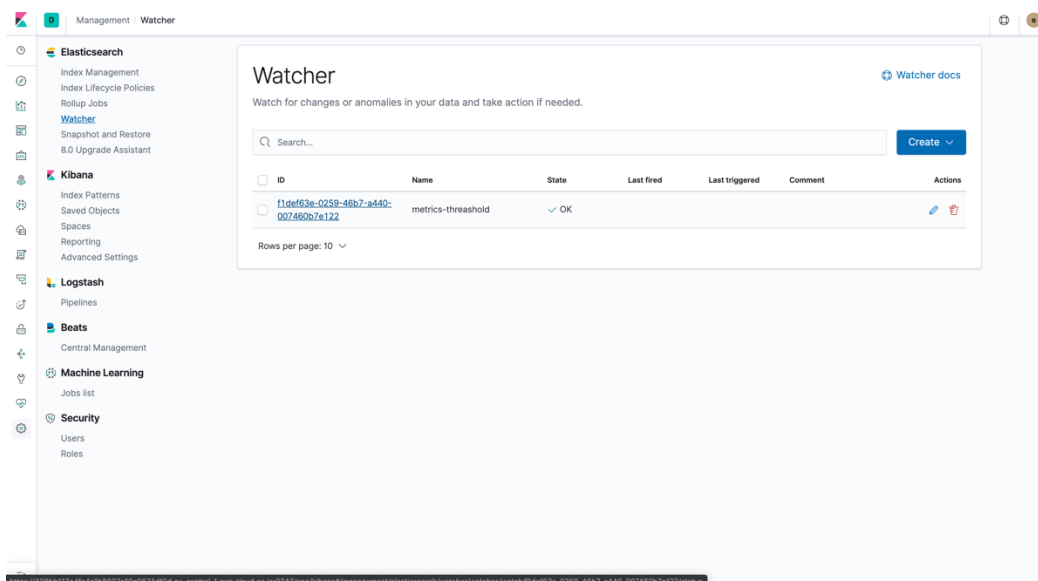
閾値を超えるシナリオに対して action を定義します。Email action を選択し、Lab0 で入力した whitelist の email アドレスを入力し、Body に email のテキストを入力します。

“Send test email”をクリックすると、email が配信されます。



“Create alert”ボタンを押し、Watcher の設定を保存します。

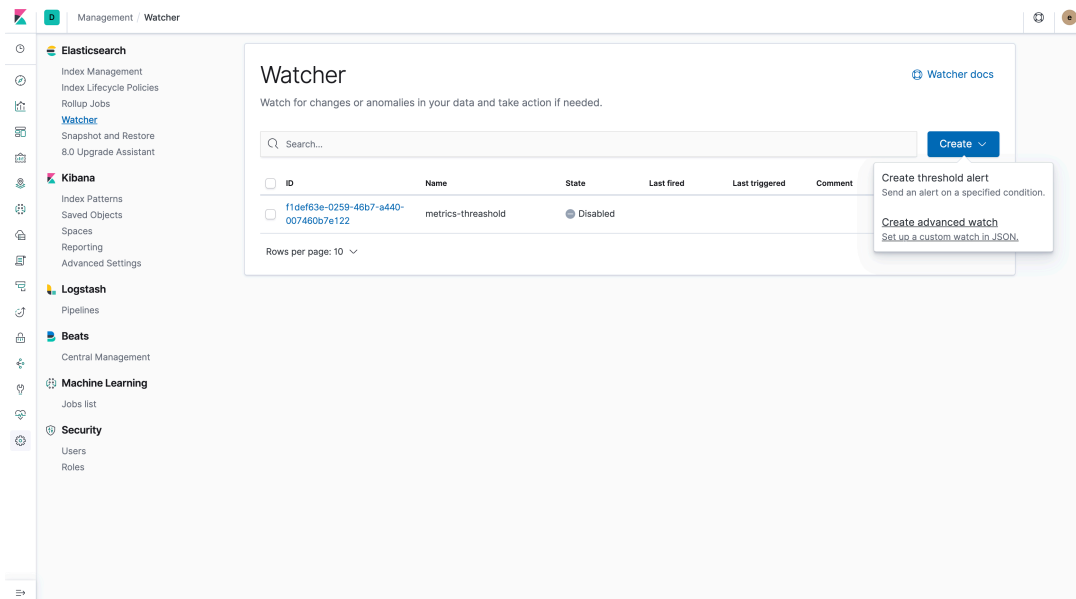
4. この状態のままだと、5 分毎に email を受信することになります。これは望まれないことかもしれません。当該の watch をクリックします。



Deactivate をクリックし、Watch を無効化します。

The screenshot shows the Elasticsearch Watcher interface. The breadcrumb navigation at the top reads "Management / Watcher / Status". The left sidebar contains a navigation menu with categories: Elasticsearch (Index Management, Index Lifecycle Policies, Rollup Jobs, Watcher, Snapshot and Restore, 8.0 Upgrade Assistant), Kibana (Index Patterns, Saved Objects, Spaces, Reporting, Advanced Settings), Logstash (Pipelines), Beats (Central Management), Machine Learning (Jobs list), and Security (Users, Roles). The main content area is titled "Current status for 'metrics-threshold'" and includes "Deactivate" and "Delete" buttons. Below the title are tabs for "Execution history" (selected) and "Action statuses". A dropdown menu is set to "Last one hour". A table with columns "Trigger time", "State", and "Comment" is shown, but it is empty with the message "No execution history to show". At the bottom of the table area, it says "Rows per page: 10".

5. Watcher スクリーンに戻ります。次に “Create advanced watch” をクリックします。



ここでは、Elasticsearch に query し、アラートを上げる方法を試してみます。データに基づいてどのような alert も作成できる柔軟性を提供します。「検索できるものは、アラートできる」と言えるでしょう。

ある IP アドレスから 1000 以上の悪意あるアクセスがあって、それを検知したいとしましょう。

最初に頻度を決めます：

```
{
  "trigger": {
    "schedule": {
      "interval": "1d"
    }
  }
}
```

次に search :

```
"input": {
  "search": {
    "request": {
      "search_type": "query_then_fetch",
      "indices": [
        "filebeat*"
      ],
      "types": [],
      "body": {
        "size": 0,
        "query": {
          "bool": {
            "must": {
              "match": {
                "nginx.access.remote_ip_list": "1.190.172.233"
              }
            },
            "filter": {
              "range": {
                "@timestamp": {
                  "gte": "now-290d/d",
                  "lt": "now/d"
                }
              }
            }
          }
        }
      }
    }
  }
}
```

次に condition :

```
"condition": {
  "compare": {
    "ctx.payload.hits.total": {
```

```
    "gte": 1000
  }
}
```

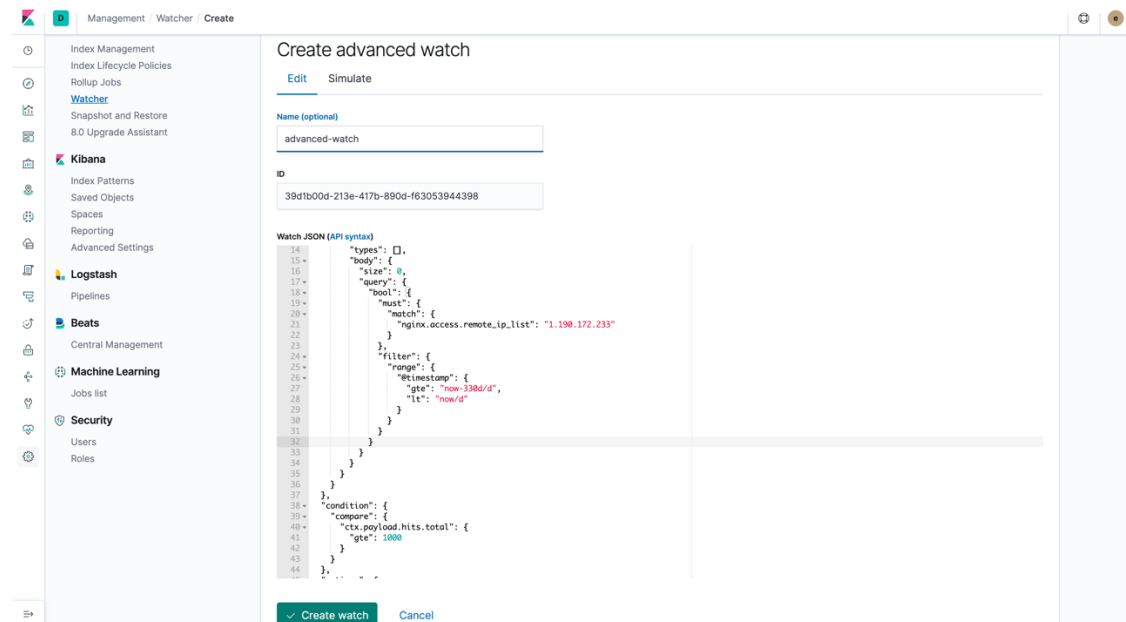
最後に action :

```
"actions": {
  "send_email": {
    "email": {
      "profile": "standard",
      "to": [
        "yukiya.shimizu@elastic.co"
      ],
      "subject": "Watcher Notification",
      "body": {
        "text": "There are {{ctx.payload.hits.total}} documents in
your index. Threshold is 1000."
      }
    }
  }
}
```

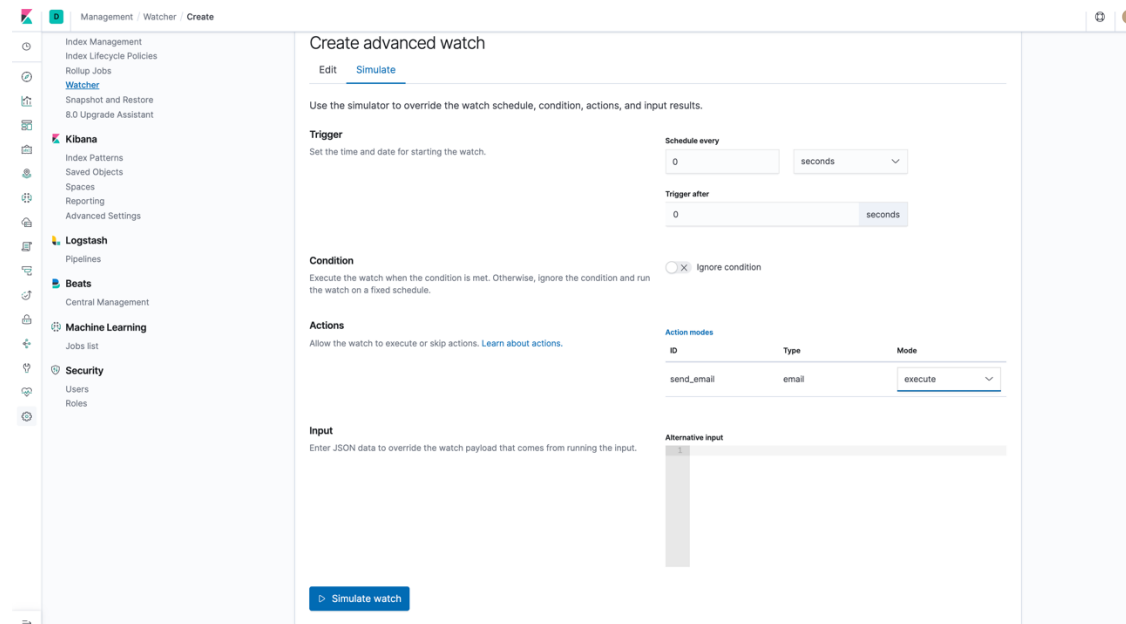
上記を1つにまとめます：

```
{
  "trigger": {
    "schedule": {
      "interval": "1d"
    }
  },
  "input": {
    "search": {
      "request": {
        "search_type": "query_then_fetch",
        "indices": [
          "filebeat*"
        ],
        "types": [],
        "body": {
          "size": 0,
          "query": {
            "bool": {
              "must": {
                "match": {
                  "nginx.access.remote_ip_list": "1.190.172.233"
                }
              },
              "filter": {
                "range": {
                  "@timestamp": {
                    "gte": "now-330d/d",
                    "lt": "now/d"
                  }
                }
              }
            }
          }
        }
      }
    }
  },
  "condition": {
    "compare": {
      "ctx.payload.hits.total": {
        "gte": 1000
      }
    }
  },
  "actions": {
    "send_email": {
      "email": {
        "profile": "standard",
        "to": [
          "yukiya.shimizu@elastic.co"
        ],
        "subject": "Watcher Notification",
        "body": {
          "text": "There are {{ctx.payload.hits.total}} documents in your index. Threshold is 1000."
        }
      }
    }
  }
}
```

上記 JSON を Watch JSON にカット&ペーストし、名前を付けます。



セーブの前に、Simulate タブをクリックし、Simulate Watch をクリックします。その際、Simulation Mode は、"Execute" とします。



Watch は “Firing” と表示されます。これは query がヒットしたことを意味します。

The screenshot shows the 'Create advanced watch' configuration page in Kibana. The 'Simulate' tab is active, and the 'Simulation results' panel on the right shows a 'Firing' event. The simulation output is as follows:

```
{
  "watch_id": "...",
  "mode": "elastic",
  "state": "executed",
  "user": "elastic",
  "status": {
    "state": {
      "active": true,
      "timestamp": "2019-10-11T04:34:19.654Z"
    },
    "last_checked": "2019-10-11T04:34:19.678Z",
    "last_met_condition": "2019-10-11T04:34:19.678Z",
    "actions": {
      "send_email": {
        "ack": {
          "timestamp": "2019-10-11T04:34:19.678Z",
          "state": "ackable"
        },
        "last_execution": {
          "timestamp": "2019-10-11T04:34:19.678Z",
          "successful": true
        },
        "last_successful_execution": {
          "timestamp": "2019-10-11T04:34:19.678Z",
          "successful": true
        }
      }
    },
    "execution_state": "executed",
    "version": -1
  },
  "trigger_event": {
    "type": "manual",
    "triggered_time": "2019-10-11T04:34:19.652Z",
    "manual": {
      "schedule": {
        "scheduled_time": "2019-10-11T04:34:19.652Z"
      }
    }
  }
}
```

Edit Tab から “Create alert” ボタンをクリックし、保存します。

The screenshot shows the 'Watcher' configuration page in Kibana. The 'Create alert' button is visible, and the table below shows the list of watches:

ID	Name	State	Last fired	Last triggered	Comment	Actions
f1df63e-0259-46b7-a440-007460b7e122	metrics-threshold	Disabled				
39d1b00d-213e-417b-890d-f63053944398	advanced-watch	OK				

Machine Learning

1. Kibana 上で “Machine Learning” を選択します。Job リストから “filebeat-nginx_ecs-access-visitor_rate_ecs” を選択し、“Start datafeed” をクリックします。

The screenshot shows the Kibana Machine Learning interface. At the top, there are navigation tabs: Job Management, Anomaly Explorer, Single Metric Viewer, Transforms, Analytics, Data Visualizer, and Settings. Below the tabs, it displays 'Active ML Nodes: 0 Total jobs: 10 Open jobs: 0 Closed jobs: 10 Active datafeeds: 0'. A search bar and a 'Create new job' button are visible. The main area is a table of jobs. The job 'filebeat-nginx_ecs-access-visitor_rate_ecs' is selected, and a context menu is open over it, showing options: Start datafeed, Clone job, Edit job, and Delete job.

2. 時間軸の指定で、“Start at beginning of data” を選択し、“Search end time”は、30th November 2018 を “Specify end time” で指定し、“Start” をクリックします。

The screenshot shows the Kibana Machine Learning interface with a modal dialog open. The dialog is titled 'Start filebeat-nginx_ecs-access-visitor_rate_ecs'. It has two sections: 'Search start time' with options 'Start at beginning of data', 'Start from now', and 'Specify start time'; and 'Search end time' with options 'No end time (Real-time search)' and 'Specify end time'. A calendar is shown for November 2018, with the 30th selected. The 'Start' button is highlighted.

- Datafeed state が “stopped” に変わるのを待ち、結果を確認するため “Single Metrics viewer” をクリックします。

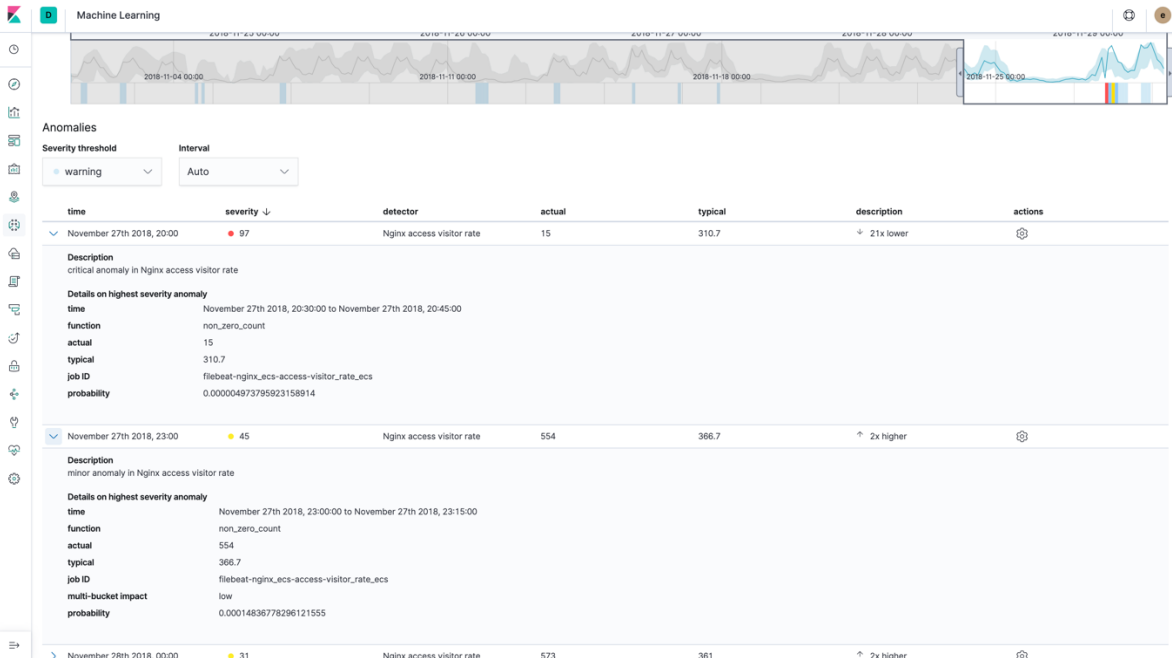
The screenshot shows the AWS SageMaker Job Management console. At the top, there are navigation tabs: Job Management, Anomaly Explorer, Single Metric Viewer, Transforms, Analytics, Data Visualizer, and Settings. Below the navigation, there are statistics: Active ML Nodes: 0, Total jobs: 10, Open jobs: 0, Closed jobs: 10, Active datafeeds: 0. A search bar and a 'Refresh' button are visible. A table lists various datafeeds with columns for ID, Description, Processed records, Memory status, Job state, Datafeed state, Latest timestamp, and Actions. The datafeed 'filebeat-nginx_ecs-access-visitor_rate_ecs' is highlighted, and a tooltip indicates it can be opened in the Single Metric Viewer.

ID	Description	Processed records	Memory status	Job state	Datafeed state	Latest timestamp	Actions
filebeat-apache_ecs-access-low_request_rate_ecs	HTTP Access Logs: Detect low request rates (ECS)	0	ok	closed	stopped		
filebeat-apache_ecs-access-source_ip_request_rate_ecs	HTTP Access Logs: Detect unusual source IPs - high request rates (ECS)	0	ok	closed	stopped		
filebeat-apache_ecs-access-source_ip_url_count_ecs	HTTP Access Logs: Detect unusual source IPs - high distinct count of URLs (ECS)	0	ok	closed	stopped		
filebeat-apache_ecs-access-status_code_rate_ecs	HTTP Access Logs: Detect unusual status code rates (ECS)	0	ok	closed	stopped		
filebeat-apache_ecs-access-visitor_rate_ecs	HTTP Access Logs: Detect unusual visitor rates (ECS)	0	ok	closed	stopped		
filebeat-nginx_ecs-access-low_request_rate_ecs	HTTP Access Logs: Detect low request rates (ECS)	0	ok	closed	stopped		
filebeat-nginx_ecs-access-source_ip_request_rate_ecs	HTTP Access Logs: Detect unusual source IPs - high request rates (ECS)	0	ok	closed	stopped		
filebeat-nginx_ecs-access-source_ip_url_count_ecs	HTTP Access Logs: Detect unusual source IPs - high distinct count of URLs (ECS)	0	ok	closed	stopped		
filebeat-nginx_ecs-access-status_code_rate_ecs	HTTP Access Logs: Detect unusual status code rates (ECS)	0	ok	closed	stopped		
filebeat-nginx_ecs-access-visitor_rate_ecs	HTTP Access Logs: Detect unusual visitor rates (ECS)	2,688	ok	closed	stopped	2018-11-29 08:59:56	

- モデルがどのように構築されたかを確認しましょう。終わりの方に、visitor rate の急激な落ち込みと上昇という異常がわかるはずです。



- 上位2つの異常を展開してみましょう。1つ目は21倍低い落ち込みで、2つ目は2倍高い上昇です。何を異常とするかは、実際の値(actual)が期待値より高い/低いではなく、発生確率(probability)です。Probabilityは、異常を展開すると表示されます。



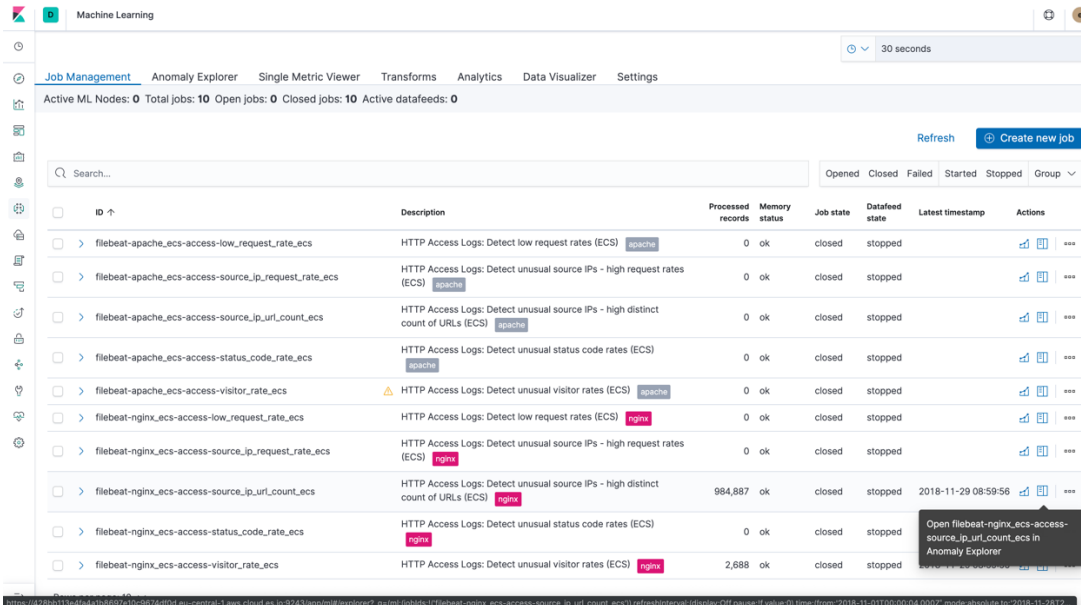
- 再度、“Machine Learning”メニューから“filebeat-nginx_ecs-access-source_ip_url_count_ecs” job を選択し、“Start datafeed”をクリックします。Start at the beginning of data を選択し、“Search end time”は、30th November 2018 を“Specify end time”で指定し、開始します。

The screenshot shows the 'Machine Learning' Job Management page. A table lists various jobs. The job 'filebeat-nginx_ecs-access-source_ip_url_count_ecs' is highlighted. A context menu is open over this job, with 'Start datafeed' selected. The table columns include ID, Description, Processed records, Memory status, Job state, Datafeed state, Latest timestamp, and Actions.

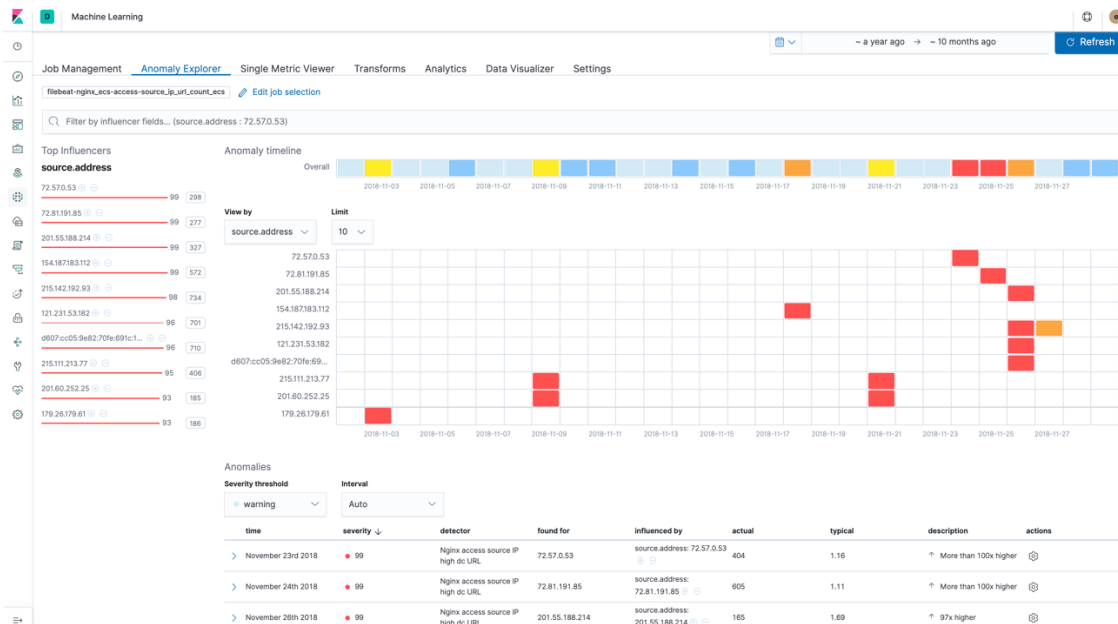
ID	Description	Processed records	Memory status	Job state	Datafeed state	Latest timestamp	Actions
filebeat-apache_ecs-access-low_request_rate_ecs	HTTP Access Logs: Detect low request rates (ECS) apache	0	ok	closed	stopped		[Start datafeed] [Clone job] [Edit job] [Delete job]
filebeat-apache_ecs-access-source_ip_request_rate_ecs	HTTP Access Logs: Detect unusual source IPs - high request rates (ECS) apache	0	ok	closed	stopped		[Start datafeed] [Clone job] [Edit job] [Delete job]
filebeat-apache_ecs-access-source_ip_url_count_ecs	HTTP Access Logs: Detect unusual source IPs - high distinct count of URLs (ECS) apache	0	ok	closed	stopped		[Start datafeed] [Clone job] [Edit job] [Delete job]
filebeat-apache_ecs-access-status_code_rate_ecs	HTTP Access Logs: Detect unusual status code rates (ECS) apache	0	ok	closed	stopped		[Start datafeed] [Clone job] [Edit job] [Delete job]
filebeat-apache_ecs-access-visitor_rate_ecs	HTTP Access Logs: Detect unusual visitor rates (ECS) apache	0	ok	closed	stopped		[Start datafeed] [Clone job] [Edit job] [Delete job]
filebeat-nginx_ecs-access-low_request_rate_ecs	HTTP Access Logs: Detect low request rates (ECS) nginx	0	ok	closed	stopped		[Start datafeed] [Clone job] [Edit job] [Delete job]
filebeat-nginx_ecs-access-source_ip_request_rate_ecs	HTTP Access Logs: Detect unusual source IPs - high request rates (ECS) nginx	0	ok	closed	stopped		[Start datafeed] [Clone job] [Edit job] [Delete job]
filebeat-nginx_ecs-access-source_ip_url_count_ecs	HTTP Access Logs: Detect unusual source IPs - high distinct count of URLs (ECS) nginx	0	ok	closed	stopped		[Start datafeed] [Clone job] [Edit job] [Delete job]
filebeat-nginx_ecs-access-status_code_rate_ecs	HTTP Access Logs: Detect unusual status code rates (ECS) nginx	0	ok	closed	stopped		[Start datafeed] [Clone job] [Edit job] [Delete job]
filebeat-nginx_ecs-access-visitor_rate_ecs	HTTP Access Logs: Detect unusual visitor rates (ECS) nginx	2,688	ok	closed	stopped	2018-11-29 08:59:56	[Start datafeed] [Clone job] [Edit job] [Delete job]

The screenshot shows the same Job Management page with a modal dialog open. The dialog title is 'Start filebeat-nginx_ecs-access-source_ip_url_count_ecs'. It has two sections: 'Search start time' with options 'Start at beginning of data' (selected), 'Start from now', and 'Specify start time'; and 'Search end time' with options 'No end time (Real-time search)' and 'Specify end time' (selected). A calendar is displayed for November 2018, with the 30th selected. 'Cancel' and 'Start' buttons are at the bottom.

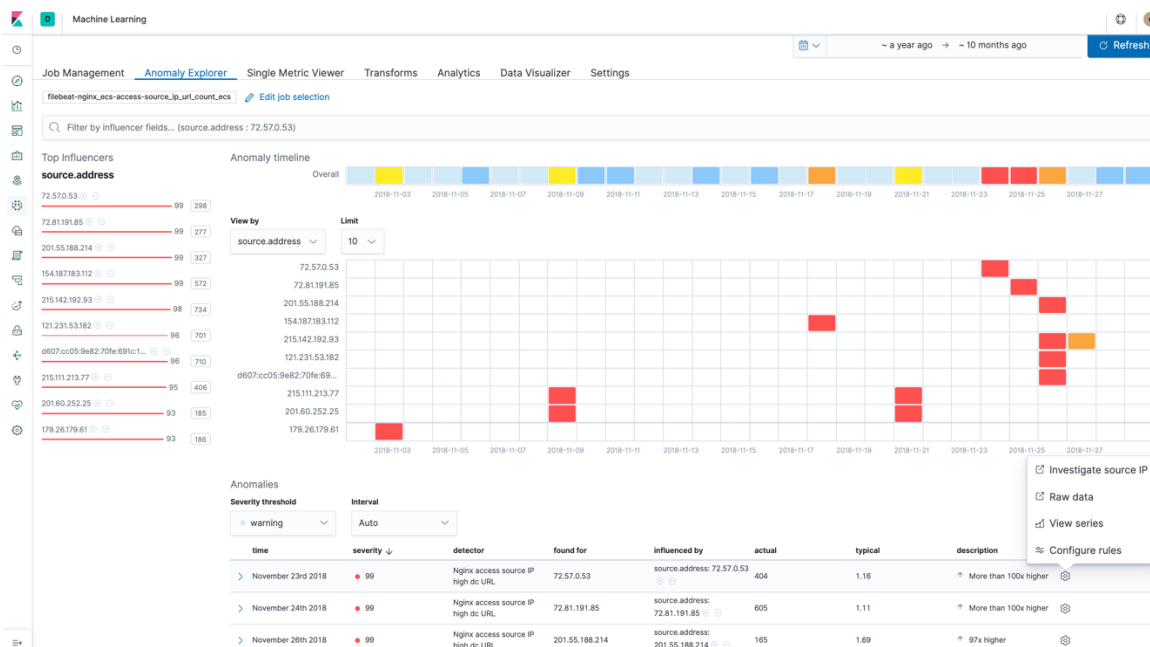
7. Datafeed state が “stopped” に変わるのを待ち、今度は、結果を確認するため “Anomaly Explorer” をクリックします。



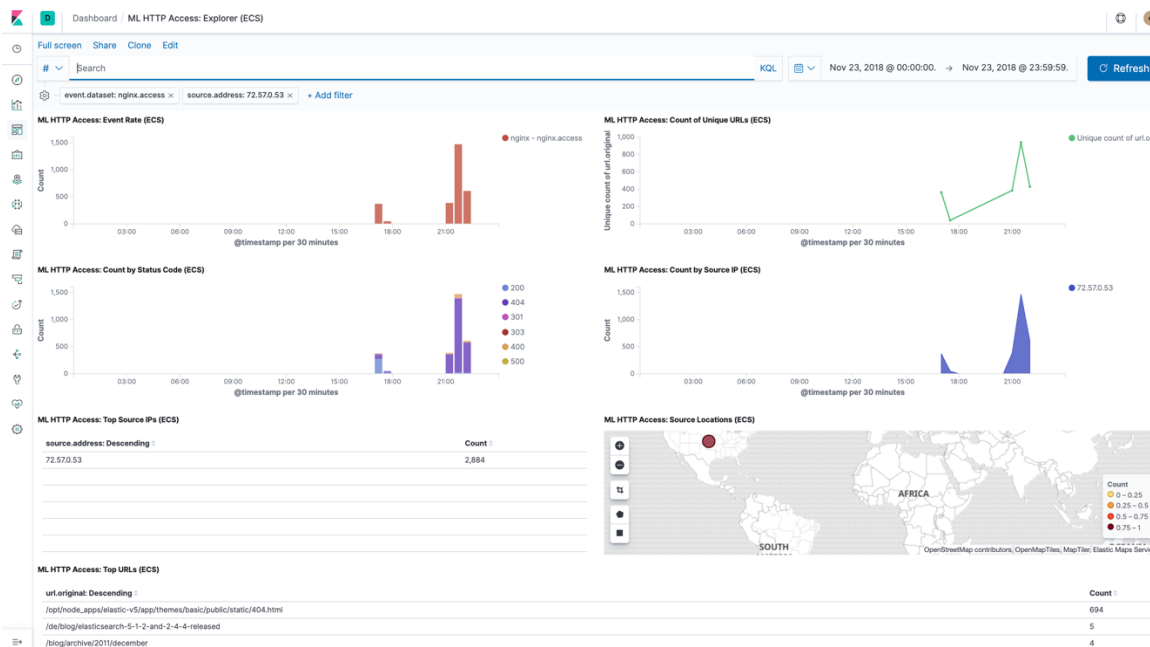
8. Anomaly explorer から特定の IP でブレークダウンされた結果を確認することができます。Influencers はデータセットの中で異常として見なされるための、統計的に顕著な要素です。機械学習 job を定義する際に、オプションとして特定のフィールドを contributor として定義することができます。Anomaly Timeline で job の結果の概要を見ることができ、マトリクスではそれぞれの IP の結果を見ることができます。Overall の結果と個別の IP の結果が同じでないことに気付きましたか？ IP “154.187.183.112” は 99 に対して、overall anomaly score は 52 です。



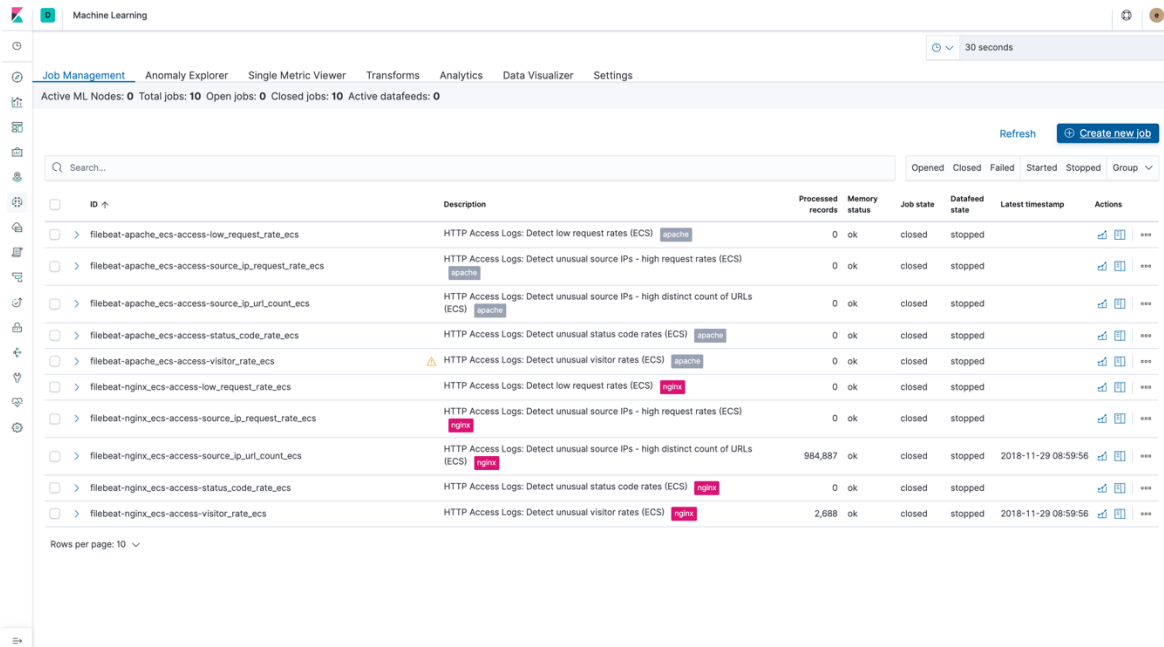
- 機械学習 job を設定する際に、結果と特定の view を紐づけることができます。View を開く時に、当該 job のメタデータが渡されます。“actions”から “Investigate source IP”を開いてみましょう。



ML HTTP Access: Explorer (ECS) dashboard が表示されます。



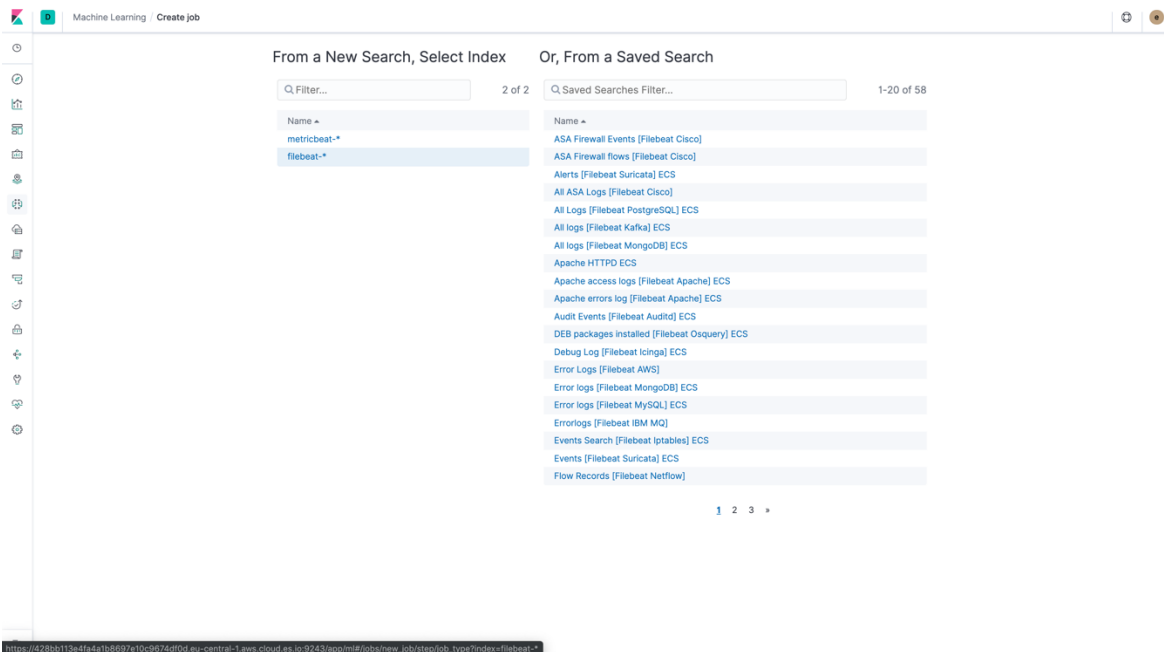
10. (Option) 自身の Single Metrics Machine Learning job を作ってみましょう。 Machine Learning をクリックし、 Create Job をクリックします。



The screenshot shows the 'Machine Learning' dashboard. At the top, there are navigation tabs: 'Job Management', 'Anomaly Explorer', 'Single Metric Viewer', 'Transforms', 'Analytics', 'Data Visualizer', and 'Settings'. Below the tabs, a summary bar indicates 'Active ML Nodes: 0', 'Total Jobs: 10', 'Open jobs: 0', 'Closed jobs: 10', and 'Active datafeeds: 0'. A search bar is present, along with 'Refresh' and 'Create new job' buttons. The main area is a table of jobs with columns: ID, Description, Processed records, Memory status, Job state, Datafeed state, Latest timestamp, and Actions. The table lists various jobs related to HTTP Access Logs, such as 'filebeat-apache_ecs-access-low_request_rate_ecs' and 'filebeat-nginx_ecs-access-status_code_rate_ecs'. The 'filebeat-nginx_ecs-access-source_ip_url_count_ecs' job shows 984,887 processed records.

ID	Description	Processed records	Memory status	Job state	Datafeed state	Latest timestamp	Actions
filebeat-apache_ecs-access-low_request_rate_ecs	HTTP Access Logs: Detect low request rates (ECS) apache	0	ok	closed	stopped		
filebeat-apache_ecs-access-source_ip_request_rate_ecs	HTTP Access Logs: Detect unusual source IPs - high request rates (ECS) apache	0	ok	closed	stopped		
filebeat-apache_ecs-access-source_ip_url_count_ecs	HTTP Access Logs: Detect unusual source IPs - high distinct count of URLs (ECS) apache	0	ok	closed	stopped		
filebeat-apache_ecs-access-status_code_rate_ecs	HTTP Access Logs: Detect unusual status code rates (ECS) apache	0	ok	closed	stopped		
filebeat-apache_ecs-access-visitor_rate_ecs	HTTP Access Logs: Detect unusual visitor rates (ECS) apache	0	ok	closed	stopped		
filebeat-nginx_ecs-access-low_request_rate_ecs	HTTP Access Logs: Detect low request rates (ECS) nginx	0	ok	closed	stopped		
filebeat-nginx_ecs-access-source_ip_request_rate_ecs	HTTP Access Logs: Detect unusual source IPs - high request rates (ECS) nginx	0	ok	closed	stopped		
filebeat-nginx_ecs-access-source_ip_url_count_ecs	HTTP Access Logs: Detect unusual source IPs - high distinct count of URLs (ECS) nginx	984,887	ok	closed	stopped	2018-11-29 08:59:56	
filebeat-nginx_ecs-access-status_code_rate_ecs	HTTP Access Logs: Detect unusual status code rates (ECS) nginx	0	ok	closed	stopped		
filebeat-nginx_ecs-access-visitor_rate_ecs	HTTP Access Logs: Detect unusual visitor rates (ECS) nginx	2,688	ok	closed	stopped	2018-11-29 08:59:56	

“filebeat*” index pattern を選択し、次に Single metric を選択します。



The screenshot shows the 'Create job' interface. It has two main sections: 'From a New Search, Select Index' and 'Or, From a Saved Search'. The 'From a New Search, Select Index' section has a search filter 'metricbeat*' and a list of index patterns including 'filebeat*'. The 'Or, From a Saved Search' section has a search filter and a list of saved searches such as 'ASA Firewall Events (Filebeat Cisco)', 'Apache HTTPD ECS', and 'Error Logs (Filebeat AWS)'. A URL bar at the bottom shows a link to a specific job configuration.

Machine Learning / Create job

Create a job from the index pattern filebeat-*

Use a supplied configuration
The fields in your data have been recognized as matching known configurations. Select to create a set of machine learning jobs and associated dashboards.

- Nginx access logs**
Find unusual activity in HTTP access logs from filebeat (ECS)

Use a wizard
Use one of the wizards to create a machine learning job to find anomalies in your data.

- Single metric**
Detect anomalies in a single time series.
- Multi metric**
Detect anomalies in multiple metrics by splitting a time series by a categorical field.
- Population**
Detect activity that is unusual compared to the behavior of the population.
- Advanced**
Use the full range of options to create a job for more advanced use cases.

Learn more about your data
If you're not sure what type of job to create, first explore the fields and metrics in your data.

- Data Visualizer**
Learn more about the characteristics of your data and identify the fields for analysis with machine learning.

https://4285b113e4f441b58697e10c06724f0d-aw-central-1.amazonaws.com/elastic/9243/app/ml/jobs/new_job/recognize?d=nginx_ecs&index=filebeat-*

Time range はデフォルトで識別されたままにして、Next をクリックします。

Machine Learning / Create job / Single metric

- Time range
- Pick fields
- Job details
- Validation
- Summary

Time range

Nov 1, 2018 @ 09:00:04.000 → Nov 29, 2018 @ 08:59:56.000 [Use full filebeat-* data](#)

> Next

Field に Count(Event data) を選択して、Next をクリックします。

The screenshot shows the 'Pick fields' step of a five-step wizard. The progress bar at the top indicates that 'Time range' (step 1) is complete, and 'Pick fields' (step 2) is the current step. Below the progress bar, a dropdown menu is set to 'Count(Event rate)'. A line chart displays a time series of data points from 11-01 09:00 to 11-27 09:00, with values ranging from approximately 1000 to 7000. Below the chart, there are settings for 'Bucket span' (set to 15m) and 'Sparse data' (unchecked). At the bottom, there are '< Previous' and '> Next' navigation buttons.

Job ID を入力し、Next をクリックします。

The screenshot shows the 'Job details' step of the wizard. The progress bar indicates that 'Time range' (step 1) and 'Pick fields' (step 2) are complete, and 'Job details' (step 3) is the current step. The 'Job ID' field contains the text 'custom-job1'. The 'Job description' field is empty. Below these fields, there is a 'Groups' dropdown menu. At the bottom, there are '< Previous' and '> Next' navigation buttons.

Validation で Valid と出力されたら、Next をクリックします。

The screenshot shows the 'Validation' step of a machine learning job creation process. At the top, a progress bar indicates five steps: 'Time range' (checked), 'Pick fields' (checked), 'Job details' (checked), 'Validation' (active, with a '4' in a blue circle), and 'Summary' (with a '5' in a grey circle). Below the progress bar, the title 'Validation' is followed by a green success message: 'Time range' with a checkmark and the text 'Valid and long enough to model patterns in the data.' Below this message are two buttons: '< Previous' and '> Next'. The left sidebar contains various icons for navigation and settings.

最後に、Summary で Create job をクリックします。

The screenshot shows the 'Summary' step of a machine learning job creation process. The progress bar at the top shows all five steps completed: 'Time range', 'Pick fields', 'Job details', 'Validation', and 'Summary' (active, with a '5' in a blue circle). The main content area displays the title 'New job from index pattern filebeat-*' and a line graph showing data over time from 11-01 09:00 to 11-27 09:00. Below the graph, job configuration details are listed in three columns: 'Job ID' (custom-job), 'Job description' (No description provided), and 'Groups' (No groups selected); 'Bucket span' (15m), 'Influencers' (No influencers selected); and 'Enable model plot' (True), 'Use dedicated index' (False), and 'Model memory limit' (10MB). At the bottom, there are four buttons: '< Previous', 'Create job' (highlighted in blue), 'Preview job JSON', and 'Convert to advanced job'. The left sidebar contains various icons for navigation and settings.

Job が終了するのを待ち、View Result をクリックします。

Machine Learning / Create job / Single metric

1 Time range 2 Pick fields 3 Job details 4 Validation 5 Summary

New job from index pattern filebeat-*

Job ID: custom-job
Job description: No description provided
Groups: No groups selected

Bucket span: 15m
Influencers: No influencers selected

Enable model plot: True
Use dedicated index: False
Model memory limit: 10MB

View results Reset job Start job running in real time Create watch

Machine Learning

Nov 1, 2018 @ 09:00:04.0 → Nov 29, 2018 @ 08:59:56.0 Refresh

Job Management Anomaly Explorer **Single Metric Viewer** Transforms Analytics Data Visualizer Settings

custom-job Edit job selection

Detector: count

Single time series analysis of count

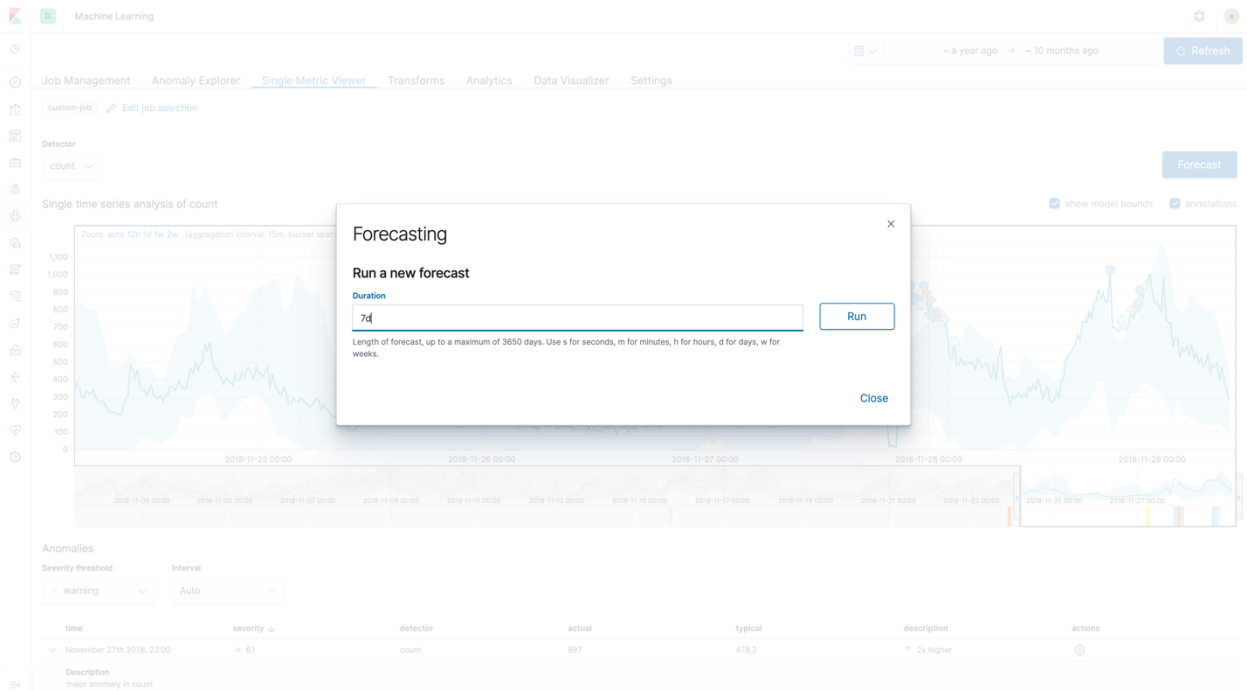
show model bounds annotations

Zoom: auto 12h 1d 1w 2w (aggregation interval: 15m, bucket span: 15m)

time	severity ↓	detector	actual	typical	description	actions
> November 27th 2018, 23:00	61	count	897	478.2	2x higher	
> November 27th 2018, 05:00	44	count	470	121.1	4x higher	
> November 28th 2018, 00:00	24	count	855	471.1	2x higher	

Single Metrics viewer で、異常をクリックしてみましょう。異常はどのようなものですか？

次に Forecast ボタンを押して、期間を 7d に設定して Run をクリックします。



Forecasting 結果を見てみましょう。Forecast パートは茶色で表示されています。

