

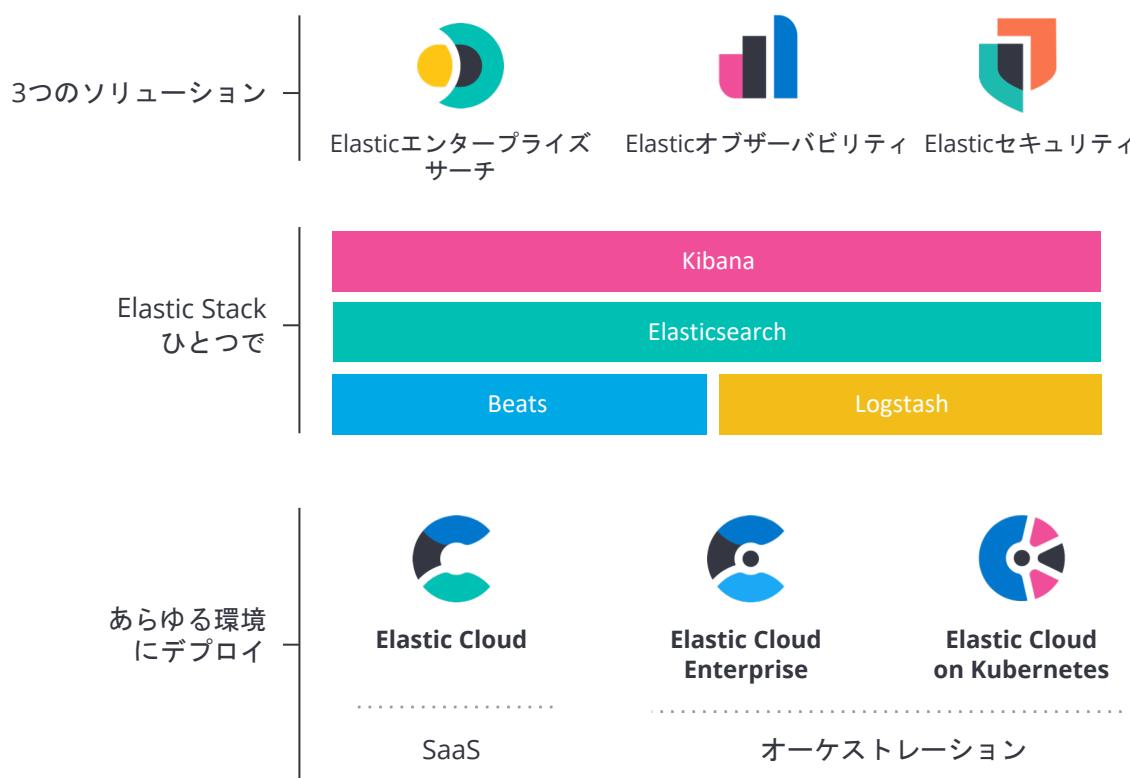


Elastic Stack ~7.6 Updates

Release Overview

Yukiya “Bruce” Shimizu
Solutions Architect

Elastic テクノロジー



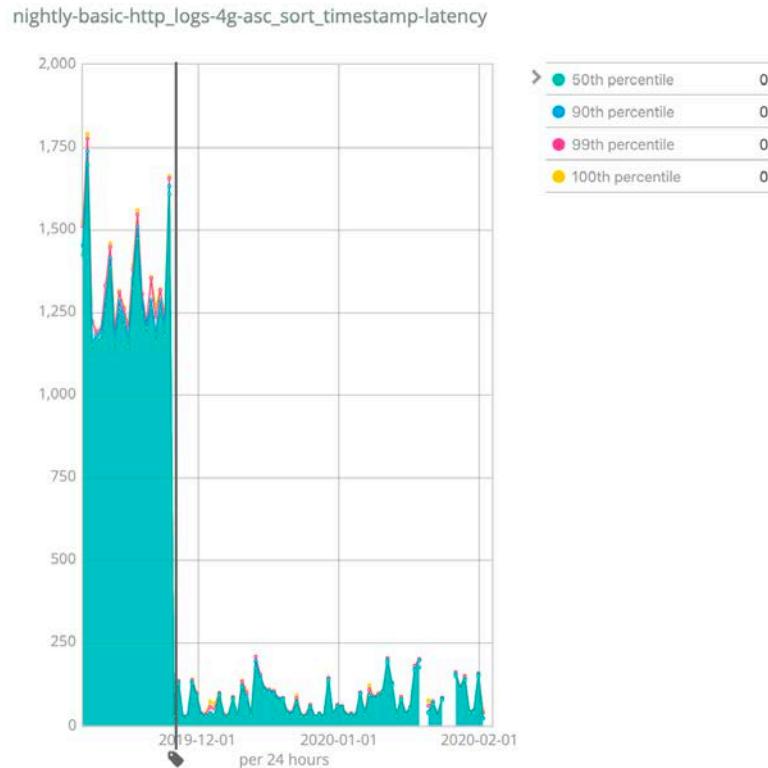


Elasticsearch

パフォーマンス最適化

DateとNumeric Data Typeのフィールドでソートが35倍高速化

- **Block-Max-WAND (7.0~)**
- 同様のアルゴリズムをDateや Numericフィールドのソートに適用
- 大量のTime Seriesデータの ソート処理などに有益



新機能

Index template mappings

- Index template mappings 作成をGUIで可能に
- 面倒なJSONでのmappings 作成から解放！

Mappings (optional)

Define how to store and index documents.

Configuration

Global settings for the index mappings

Dynamic field

true

Allow new fields discovery in document.

Date detection

Check if the string field is a date.

Numeric field

Check if the string field is a numeric value.

Dynamic dates format

Type and then hit "ENTER"

The dynamic_date_formats can be customised to support your own date formats.

Document fields

Define which fields the documents of your index will contain.

▼ user Object

└ name Text

└ address Object

 └ street Text

 └ city Text

 └ location Geo-point

date_created Date

 Add field

新機能

Histogram Datatype

```
PUT my_index
{
  "mappings": {
    "properties": {
      "my_histogram": {
        "type" : "histogram"
      },
      "my_text" : {
        "type" : "keyword"
      }
    }
  }
}
```



```
PUT my_index/_doc/1
{
  "my_text" : "histogram_1",
  "my_histogram" : {
    "values" : [0.1, 0.2, 0.3, 0.4, 0.5], ①
    "counts" : [3, 7, 23, 12, 6] ②
  }
}

PUT my_index/_doc/2
{
  "my_text" : "histogram_2",
  "my_histogram" : {
    "values" : [0.1, 0.25, 0.35, 0.4, 0.45, 0.5], ①
    "counts" : [8, 17, 8, 7, 6, 2] ②
  }
}
```

新機能

String Stats Aggregation

```
POST /twitter/_search?size=0
{
  "aggs" : {
    "message_stats" : {
      "string_stats" : {
        "field" : "message.keyword"
      }
    }
  }
}
```



```
{
  ...
  "aggregations": {
    "message_stats" : {
      "count" : 5,
      "min_length" : 24,
      "max_length" : 30,
      "avg_length" : 28.8,
      "entropy" : 3.94617750050791
    }
  }
}
```

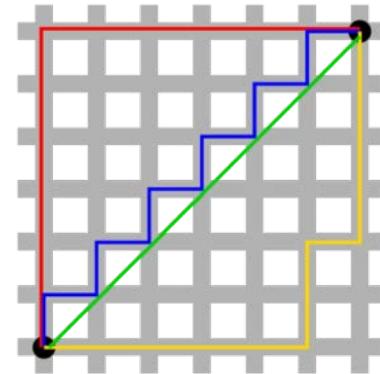
- Count
- Minimum length/Maximum length
- Average length
- Shannon entropy

<https://www.elastic.co/guide/en/elasticsearch/reference/7.6/search-aggregations-metrics-string-stats-aggregation.html>

New Vector Distance Functions

ベクトルを使った関連度ランキングの拡張

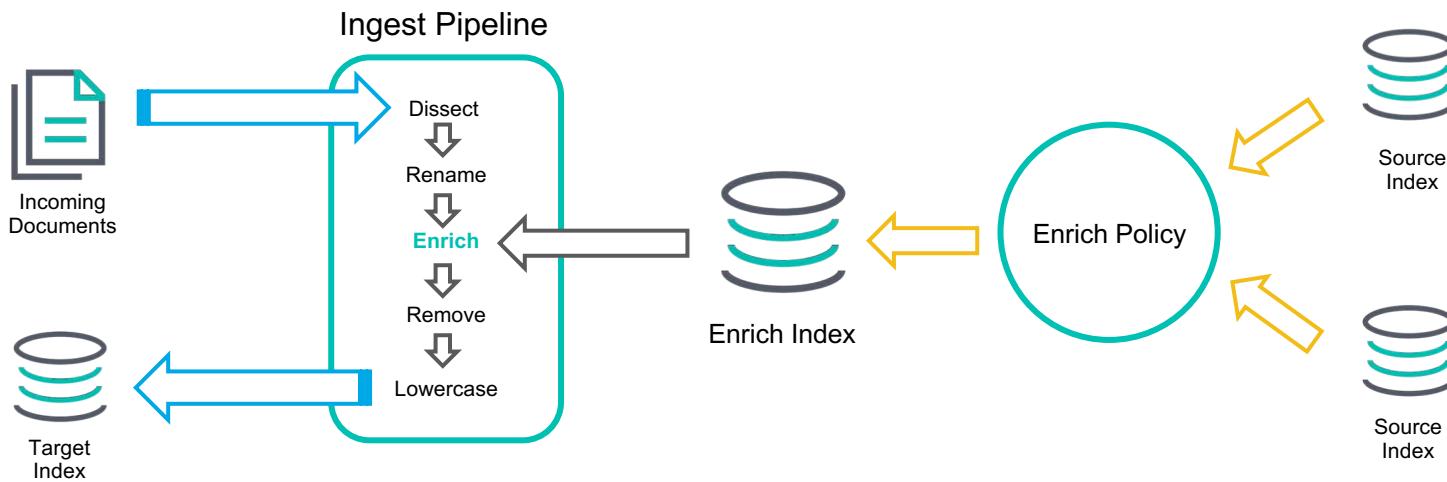
- 文書の類似性をベクトルを使って測定するシナリオは無数にある
- ベクトル化した文書から、文書（イメージやテキスト）を表現する様々なアルゴリズムを使ったベクトル化に至る幅広い応用
- 7.4から二つの類似性を測定するFunctionを導入
 - Manhattan Distance (L1 norm)
 - Euclidean Distance (L2 norm)
- 定義済みのPainless functionとして提供され、Script score queryの一部として他のクエリーに統合可能



<https://www.elastic.co/jp/blog/text-similarity-search-with-vectors-in-elasticsearch>

Enrich Processor

データ投入をよりシンプルに

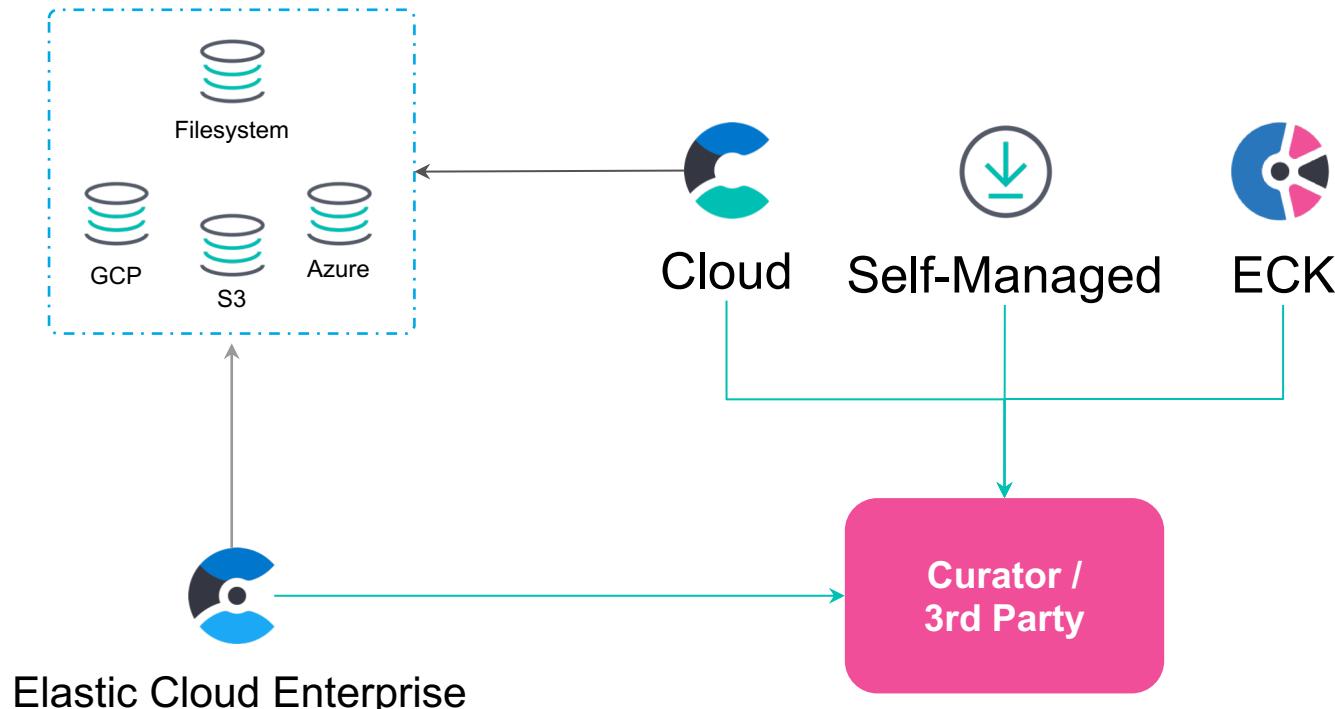


Snapshot Lifecycle Management

SLM

Snapshot Lifecycle Management (これまで)

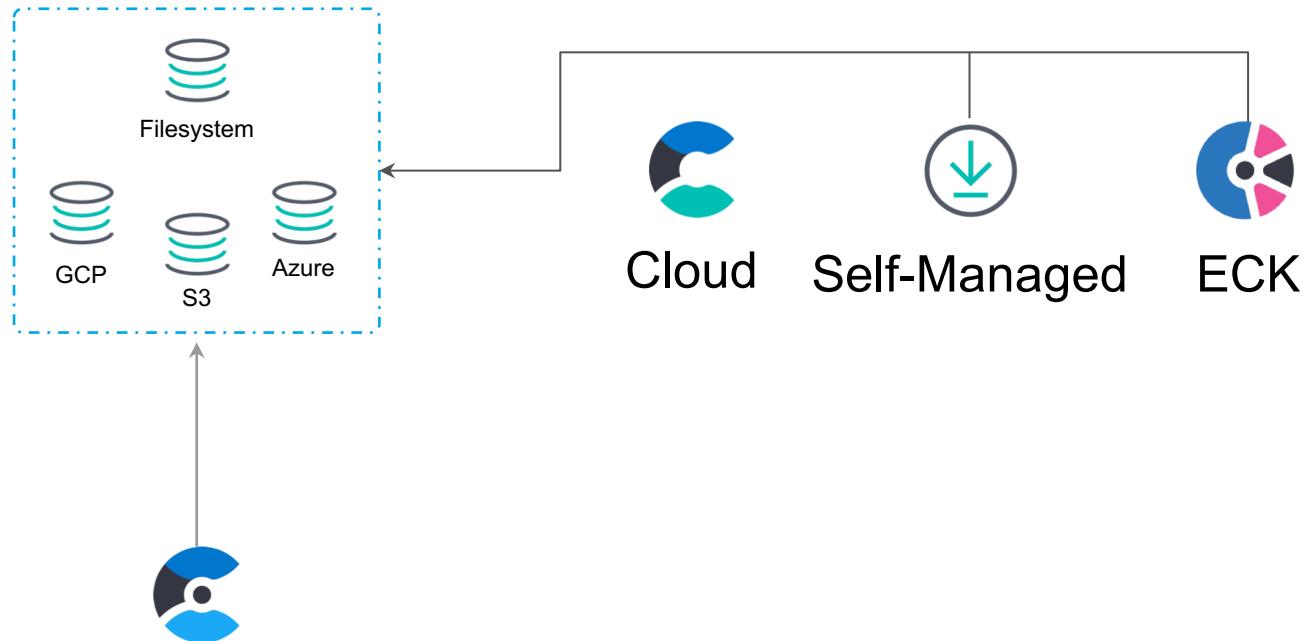
Snapshotの管理には、Curatorや3rd Partyツールが必要だった…



Elastic Cloud Enterprise

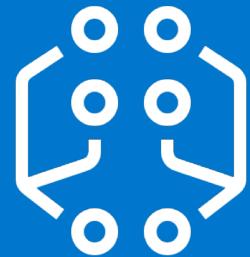
Snapshot Lifecycle Management (これから)

最早、Curatorや3rd Partyツールは必要なし



Elastic Cloud Enterprise

<https://www.elastic.co/guide/en/elasticsearch/reference/7.6/getting-started-snapshot-lifecycle-management.html>



Machine Learning

Elastic Transforms

目標

データを新しい形式に“transform”することを可能に

- 元々のデータを変更せずに違うデータ形式を提供し、機械学習に活用する
- Transformのステップ
 - Pivot
 - Aggregation

多次元の分析を可能に

- Outlier detection(はずれ値検知)
- Regression(回帰) と Classification(分類)
- 教師あり学習のモデル構築

Elastic Transforms

The screenshot shows the Elasticsearch Management interface with the 'Transforms' section selected. On the left, there's a sidebar with various management tools like Index Management, Rollup Jobs, and Watcher. The main area is titled 'Create transform' (BETA) and is divided into two main sections:

- Define pivot:** This section shows a saved search for '[eCommerce] Orders'. It has a 'Group by' section with 'customer_full_name.keyword' and 'customer_gender', and an 'Aggregations' section with 'products.base_price.max', 'products.base_price.avg', and 'products.base_unit_price.avg'. There are also dropdowns for 'category.keyword', 'currency', and 'customer_birth_date'.
- Transform pivot preview:** This section displays data from the source index 'kibana_sample_data_ecommerce'. It lists fields: category, currency, customer_first_name, customer_full_name, and customer_gender. Below this is a table showing transformed data with columns: customer_full_name.keyword, customer_gender, products.base_price.avg, products.base_price.max, and products.base_unit_price.avg. The table includes rows for Abigail Abbott, Abigail Adams, Abigail Austin, Abigail Bailey, and Abigail Baker, each with their gender and specific price metrics.

<https://www.elastic.co/guide/en/elasticsearch/reference/current/ecommerce-transforms.html>

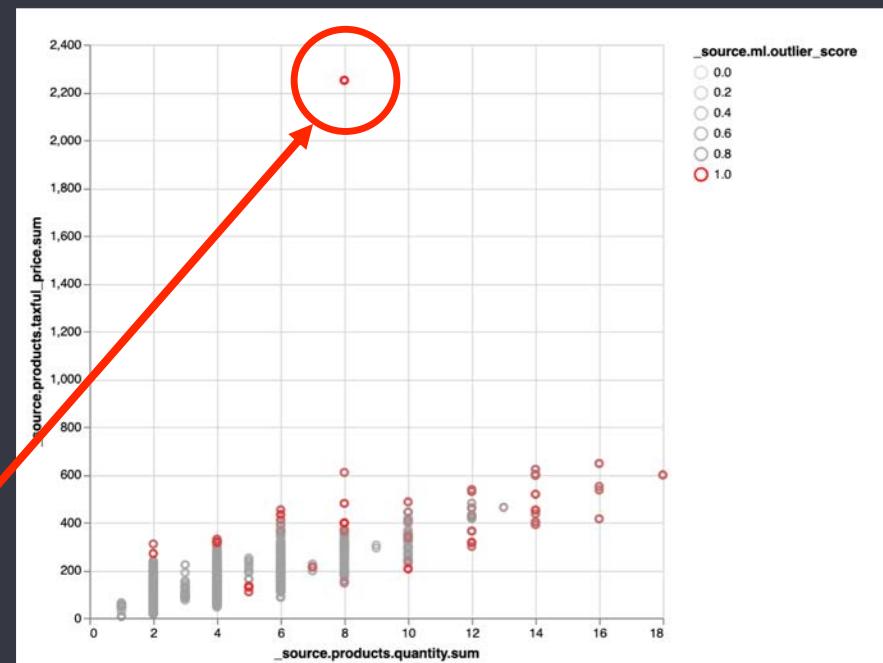
Outlier Detection



Elastic Data Frame Analytics

Outlier Detectionの活用

```
"customer_full_name" : {  
    "keyword" : "Wagdi Shaw"  
},  
"ml_id_copy" : "Vyu9e08pKNasT-9TLV9p3k0AAAAAAA",  
"products" : {  
    "taxful_price" : {  
        "sum" : 2250.0  
    },  
    "quantity" : {  
        "sum" : 8.0  
    }  
},  
"ml" : {  
    "outlier_score" : 0.9848338961601257,  
    "feature_influence.products.quantity.sum" : 0.007586637046188116,  
    "feature_influence.products.taxful_price.sum" : 0.992413341999054
```



Elastic Data Frame Analytics

Machine Learning | Data Frame Analytics

Overview Anomaly Detection Data Frame Analytics Data Visualizer

Analytics exploration EXPERIMENTAL

Job ID ecommerce_outliers_2

customer_gender	customer_last_name.keyword	ml.outlier_score ↓	products.base_unit_price.avg	products.category.keyword.value_count	products.price.avg
MALE	Shaw	0.9986892398613892	52.13790760869965	48	68.44225543478281
MALE	Frank	0.9984397888193594	54.68212890625	13	54.68212890625
MALE	King	0.9759089946746826	39.12682291666667	49	39.12682291666667
MALE	Palmer	0.973523996260498	52.43310546875	14	52.43310546875
MALE	Little	0.971738624572754	42.63755580357143	10	42.63755580357143
FEMALE	Byrd	0.8427810072898865	37.54624720982143	42	37.54624720982143
MALE	Bryan	0.8054728719856262	52.594140625	23	52.594140625
MALE	Elliott	0.796035230159759	21.28822916666668	14	21.28822916666668
FEMALE	Conner	0.7919033765792847	25.855234375	4	25.855234375
MALE	Simpson	0.70938408374748638	26.81274609375	5	26.81274609375
MALE	Wolfe	0.6986919661262512	49.14296875	12	49.14296875
FEMALE	Bryant	0.6833633780479431	32.653365384615384	53	32.653365384615384
FEMALE	Garcia	0.6529442667961121	41.35546875	29	41.35546875
MALE	Reyes	0.6303384332742532	42.01963404605263	30	42.01963404605263
MALE	Franklin	0.632892787456125	43.80517578125	11	43.80517578125
FEMALE	McDonald	0.6308591365814205	39.0875726744186	35	38.0517743255814
MALE	Tran	0.6244477033615112	35.817210477941174	52	34.317210477941174
FEMALE	Underwood	0.6104920506477356	33.45430715460526	53	33.45430715460526
FEMALE	Graham	0.6007459183433531	26.21580102040617	35	25.97092142677142
FEMALE	Hudson	0.5911657214164731	21.366373697916668	14	21.366373697916668
MALE	Hale	0.52949714696064453	50.0435546875	15	50.0435546875
MALE	Rivera	0.495784938335187	35.2421875	51	35.2421875
FEMALE	Gilbert	0.4156837463378906	26.59652549342105	28	26.59652549342105
MALE	Perkins	0.39029788109321594	33.76991102430556	51	33.76991102430556
MALE	Estrada	0.38506796956002317	49.8453125	15	49.8453125

Rows per page: 25 ▾

1 2 3 4 5 ... 15 >

<https://www.elastic.co/guide/en/machine-learning/current/ecommerce-outliers.html>

教師あり学習.....



Regression (回帰)

回帰とは、複数の変数、又は特徴量から連続値を予測する方法の一つ

例えば、ロンドン市内のアパートの広さと家賃の関係を知りたい場合

特徴量

連続値



Size (m ²)	Monthly rent (GBP)
44	1600
24	1055
63	2300

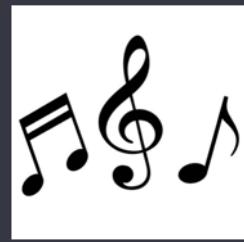
Classification (分類)

分類とは、所与のデータのクラス、又は種類を予測するプロセス



cancerous

benign



country

hip-hop

avant garde



default

safe

教師あり学習

Regression

ラベル付されたデータ

モデルを学習

dependent variable = how much?

analyzed fields: 分析に使うフィールド

テストデータを使ってモデルを評価

```
PUT _ml/_data_frame/analytics/model-flight-delays
{
  "source": {
    "index": [
      "kibana_sample_data_flights" ①
    ],
    "query": { ②
      "range": {
        "DistanceKilometers": {
          ...
          "DestRegion" : "UK",
          "OriginAirportID" : "LHR",
          "DestCityName" : "London",
          "FlightDelayMin" : 66,      ①
          "ml" : {
            "FlightDelayMin_prediction" : 62.527,   ②
            "is_training" : false  ③
          }
          ...
          ...
          ...
          "excludes": [ ⑥
            "Cancelled",
            "FlightDelay",
            "FlightDelayType"
          ]
        },
        "model_memory_limit": "100mb" ⑦
      }
    }
  }
}
```

教師あり学習

Classification

ラベル付されたデータ

モデルを学習

dependent variable = true/false

analyzed fields: 分析に使うフィールド

テストデータを使ってモデルを評価

```
PUT _ml/_data_frame/_analytics/_model-flight-delay-classification
{
  "source": {
    ...
    "FlightDelay" : false, ①
    ...
  },
  "d"
  "ml" : {
    "top_classes" : [ ②
      {
        "class_probability" : 0.939335365058496,
        "class_name" : "false"
      },
      {
        "class_probability" : 0.06066463494150393,
        "class_name" : "true"
      }
    ],
    "FlightDelay_prediction" : "false", ③
    "is_training" : false ④
  }
},
"model_memory_limit": "100mb" ⑤
}
```

モデルの評価

Regression example

```
{  
    "regression" : {  
        "mean_squared_error" : {  
            "error" : 3759.7242253334207  
        },  
        "r_squared" : {  
            "value" : 0.5853159777330623  
        }  
    }  
}
```

Mean Squared Error:
低い方が良い



R Squared (0-1):
高い方が良い



<https://www.elastic.co/guide/en/machine-learning/current/ml-dfanalytics-evaluate.html#ml-dfanalytics-classification>

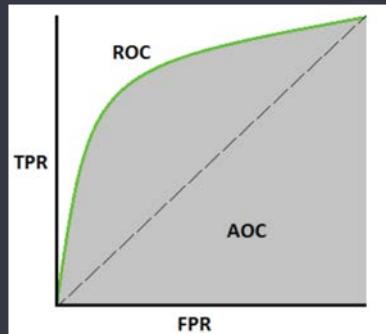
モデルの評価

Classification example

$$\text{accuracy} = \frac{TP + TN}{TP + TN + FP + FN}$$

Confusion Matrix:

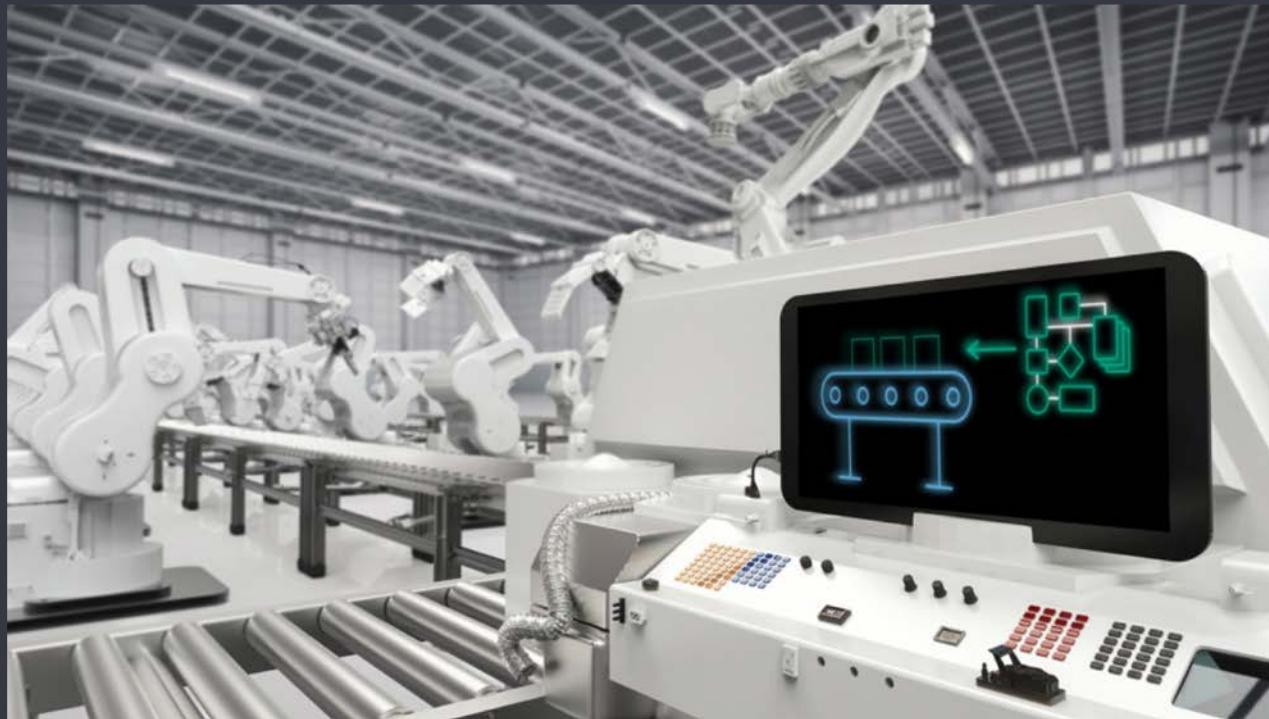
True/False
Positives/Negatives



```
{
  "classification": {
    "multiclass_confusion_matrix": {
      "confusion_matrix": [
        {
          "actual_class": "false", ①
          "actual_class_doc_count": 8778, ②
          "predicted_classes": [
            {
              "predicted_class": "false", ③
              "count": 7509 ④
            },
            {
              "predicted_class": "true",
              "count": 1269
            }
          ],
          "other_predicted_class_doc_count": 0
        },
        {
          "actual_class": "true",
          "actual_class_doc_count": 2939,
          "predicted_classes": [
            {
              "predicted_class": "false",
              "count": 1213
            },
            {
              "predicted_class": "true",
              "count": 1726
            }
          ],
          "other_predicted_class_doc_count": 0
        }
      ],
      "other_actual_class_count": 0
    }
  }
}
```

<https://www.elastic.co/guide/en/machine-learning/current/ml-dfanalytics-evaluate.html#ml-dfanalytics-classification>

さらに



運用可能に！
...推論

モデルの学習と推論

教師あり学習

学習/テスト/評価



```
PUT _ml/data_frame/analytics/churn
{
  "source": {
    "index": "customer_behaviour"
  },
  "dest": {
    "index": "customer_behaviour_churn"
  },
  "analysis": {
    "regression": {
      "dependent_variable": "churn_probability",
      "training_percent": 80
    }
  }
}
POST _ml/data_frame/analytics/churn/_start
```

推論



```
PUT _ingest/pipeline/predict_churn
{
  "description": "Predict customer churn",
  "processors": [
    {
      "inference": {
        "model": {
          "regression": {
            "model_id": "churn",
            "target_field": "churn_probability"
          }
        }
      }
    }
  ]
}
```



Kibana

Kibana Lens

Kibanaでデータを可視化する
簡単で直感的な新しい方法

データは目の前に

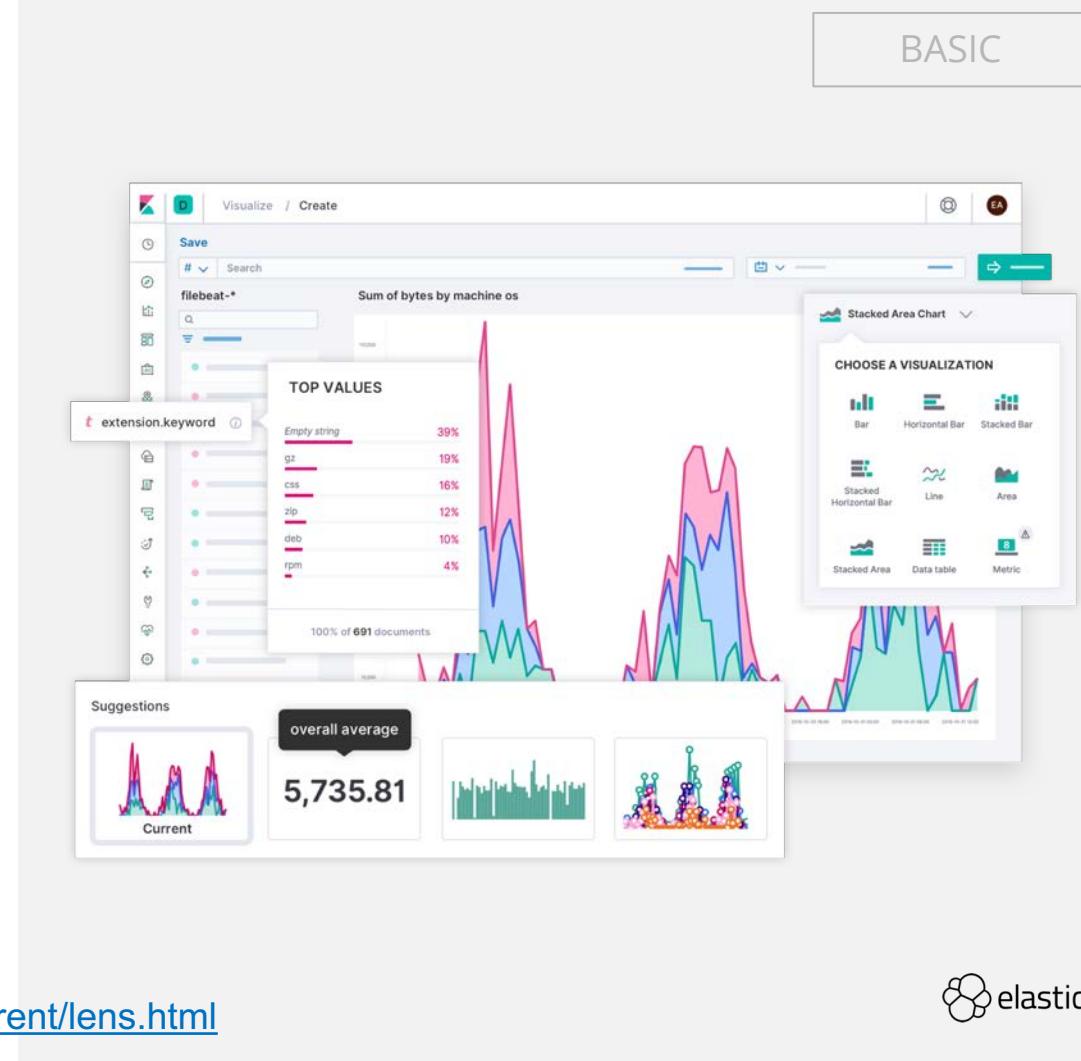
フィールドをクリックして探索を開始

ドラッグ&ドロップ

マウス一つの操作から目的地がわからない場合も、気軽に探索を

スマートなサジェストション

便利なチャートをサジェストして、Lensが分析を支援





Kibanaのイノベーションにおけるガイドライン

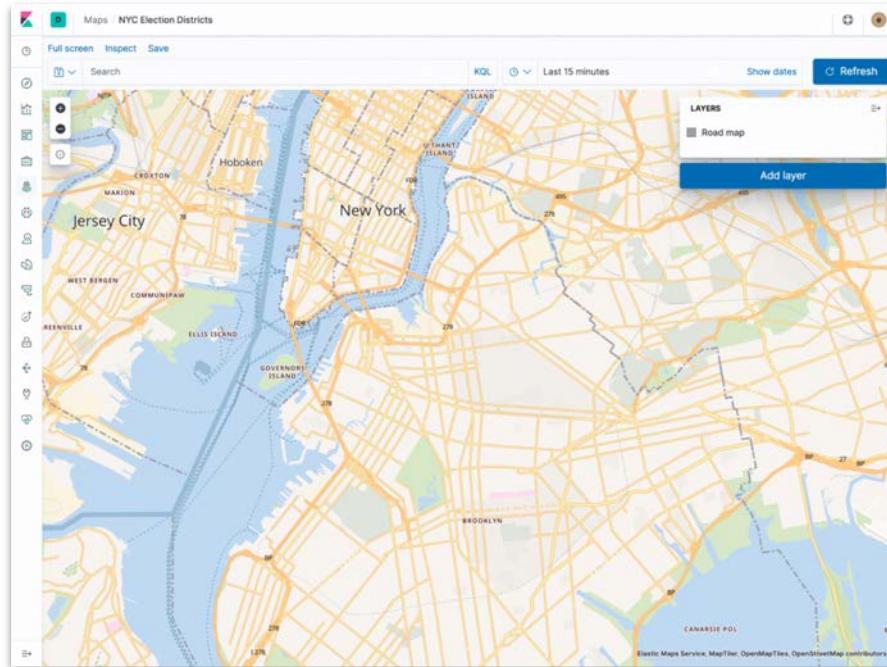
- 1 過去のKibanaの経験がないユーザーをエンパワー
- 2 簡単にデータフィールドの気付きや理解を促進
- 3 アドホックなシナリオに柔軟性をサポート
- 4 スピードやスケール、Elasticsearchのフィーチャーを活用
- 5 簡単に使って、しかもパワフルなインパクト



Elastic Maps

GeoJSONアップロードがGAに カスタムシェイプをMapsに追加する最速の方法

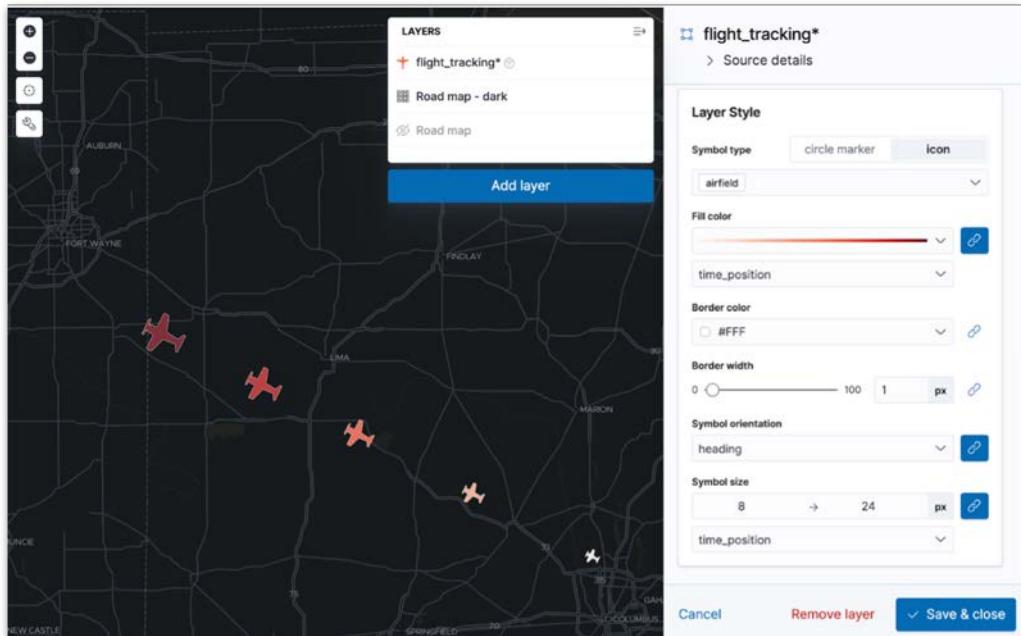
- 7.3からBeta提供され、現在GAとなり安定性も向上
- 営業テリトリーやオフィスの分散など、顧客固有のデータに最適
- 生のデータにフィルターやアグリゲーションといった分析を可能に



時系列データにスタイルを追加

日付/タイムフィールドと色やサイズを関連付ける

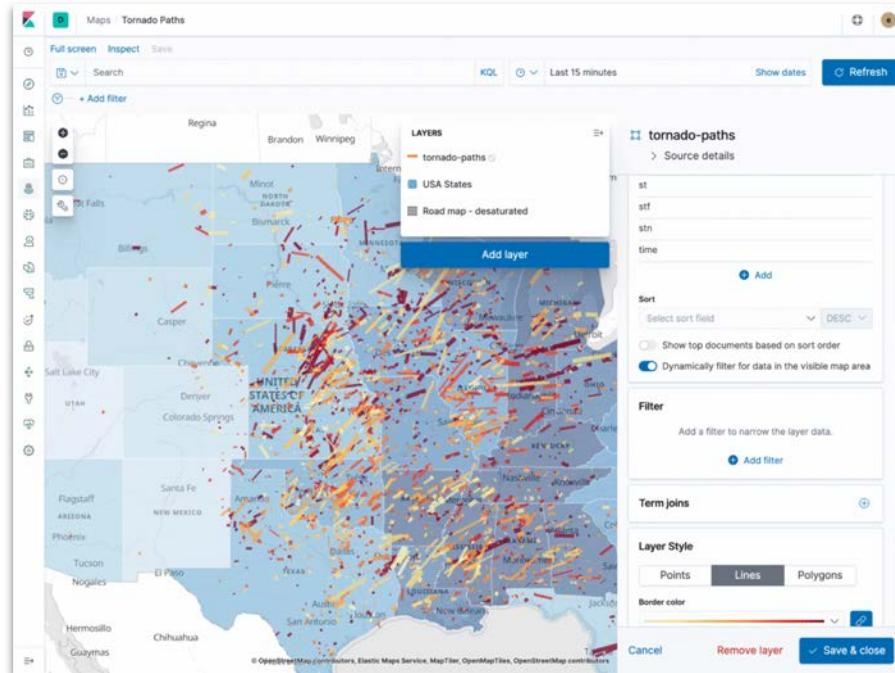
- アセットトラッキングやその他時系列データに最適
- 直近の車の位置、荷物の位置、人の位置などのトラッキング
- 従来は数値フィールドにのみ対応



レイヤーにおけるドキュメントのソート

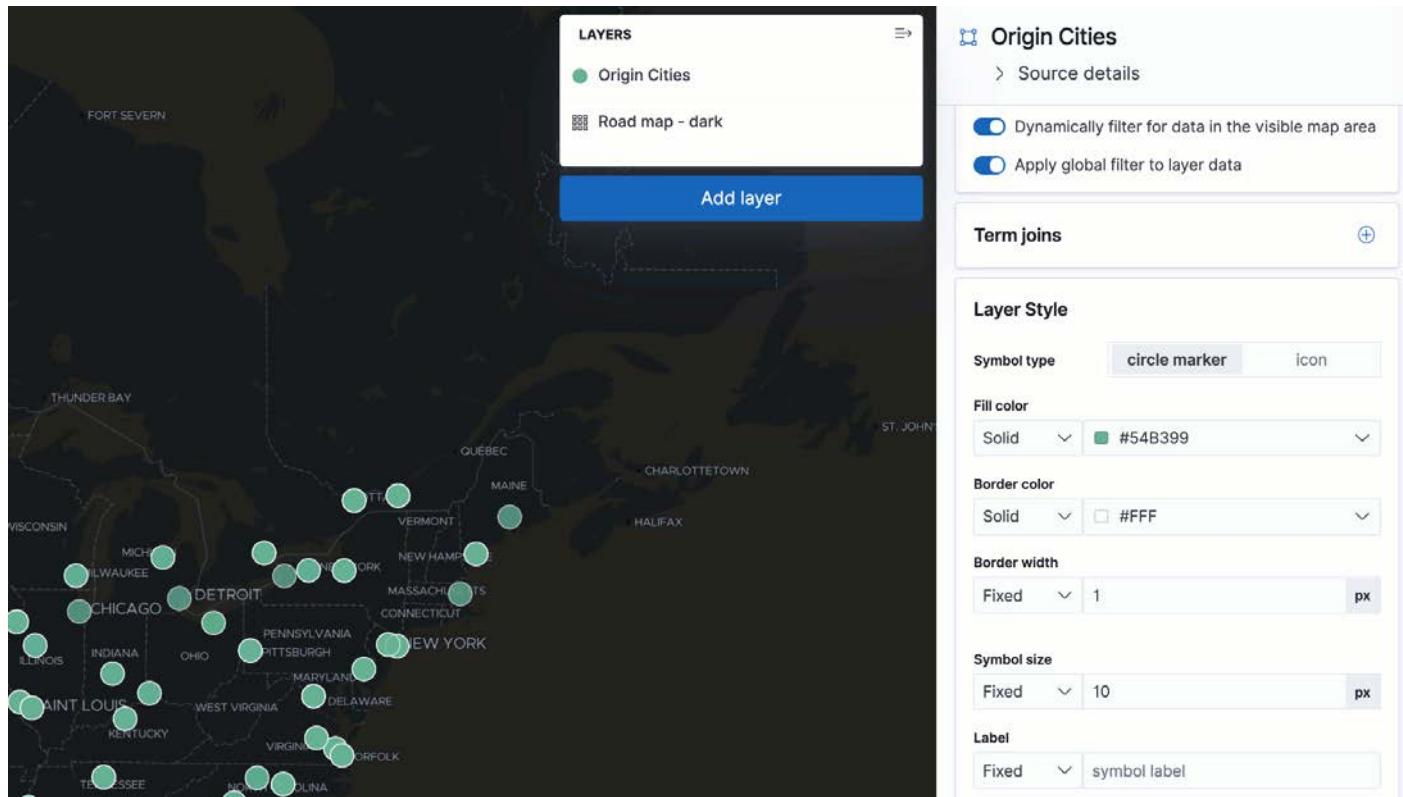
最も大切なデータをトップに

- 数値もしくは日付フィールドでソートした結果を、表示順に反映
- データの表現にさらなる一貫性
- 10k ドキュメントが上限



Categorical Styling

分類フィールドの値によって、データポイントを色付け





Beats

Beats アップデート

Added

Affecting all Beats

- Fail with error when autodiscover providers have no defined configs. [13078](#)
- Add autodetection mode for `add_docker_metadata` and enable it by default in included configuration files. [13374](#)
- Add autodetection mode for `add_kubernetes_metadata` and enable it by default in included configuration files. [13473](#)
- Use less restrictive API to check if template exists. [13847](#)
- Do not check for alias when `setupilm.check_exists` is false. [13848](#)
- Add support for numeric time zone offsets in timestamp process
- Add condition to the config file template for `add_kubernetes_metadata`
- Marking Central Management deprecated. [14018](#)
- Add `keep_null` setting to allow Beats to publish null values in events
- Added `shared_credential_file` option in aws related config for spec directory. [14157](#) [14178](#)
- Ensure that init containers are no longer tailed after they stop. [14201](#)
- Libbeat HTTP's Server can listen to a unix socket using the unix syntax. [13655](#)
- Libbeat HTTP's Server can listen to a Windows named pipe using syntax. [13655](#)
- Adding new Enterprise license type to the licenser. [14246](#)

Auditbeat

- Socket: Add DNS enrichment. [14004](#)

Filebeat

- Add support for virtual host in Apache access logs. [12778](#)
- Update CoreDNS module to populate ECS DNS fields. [13320](#) [13131](#)
- Parse query steps in PostgreSQL slowlogs. [13496](#) [13701](#)
- Add filebeat azure module with activitylogs, auditlogs, signinlogs
- Add support to set the document id in the json reader. [5844](#)
- Add input httpsjson. [13545](#) [13546](#)
- Filebeat Netflow input: Remove beta label. [13858](#)
- Remove event.timezone from events that don't need it in some log formats with and without timezones. [13918](#)
- Add ExpandEventListFromField config option in the kafka input.
- Add ELB fileset to AWS module. [14020](#)
- Add module for MISP (Malware Information Sharing Platform). [13410](#)
- Add filebeat azure module with activitylogs, auditlogs, signinlogs filesets. [13776](#) [14033](#)
- Add support for all the ObjectCreated events in S3 input. [14077](#)
- Add `source.bytes` and `source.packets` for uni-directional netflow events. [14111](#)
- Add Kibana Dashboard for MISP module. [14147](#)
- Add support for gzipped files in S3 input [13980](#)
- Add Filebeat Azure Dashboards [14127](#)

Heartbeat - Add non-privileged icmp on linux and darwin(mac). [13795](#) [11498](#) - Allow hosts to be used to configure http monitors [13703](#)

Added

Affecting all Beats

- Add a friendly log message when a request to docker has exceeded the deadline. [15336](#)
- GA the script processor. [14325](#)
- Add fingerprint processor. [11173](#) [14205](#)
- Add support for elasticsearch outputs. [14324](#)
- Add consumer lag metricset in Kafka consumer_group metricset [14822](#)
- Metrics in consumer_lag metricset in a dashboard [14863](#)
- Collector kubernetes autodiscovery: enable different resource based discovery [14738](#)
- Add `add_id` processor. [14524](#)
- Enable TLS 1.3 in all beats. [12973](#)
- Spooling to disk creates a lockfile on the platform. [15338](#)
- Enable DEP (Data Execution Protection) for windows packages. [15149](#)
- Users can now specify `monitoring.cloud_provider` to override `monitoring.elasticsearch.*` settings. [14399](#) [1554](#)
- Add support to kubernetes autodiscovery to add additional metadata from other source to events. [14875](#)

Logstash

- Expand data for the system/memory metricset [15492](#)
- Add azure storage metricset in order to retrieve metric values for storage accounts. [14548](#) [15342](#)

Logstash

- Add new fileset goop module for ingest pipeline. [15320](#)
- Add dashboard for the azureready2 module (ported from Logstash). [15321](#)
- Add experimental `list_from_field` support in `s3` fileset. [15322](#)
- Add azure-eventhub input which will use the a `list_from_field` support. [15323](#)
- Expose more metrics of harvesters (e.g. read, write, etc.) [15324](#)
- Include `log.source.address` for unparseable syslog entries. [15325](#)
- Release aws elb fileset as GA. [15426](#) [15380](#)
- Integrate the azure-eventhub with filebeat azure module. [15326](#)
- Release aws s3access fileset to GA. [15431](#) [154](#)
- Add cloudtrail fileset to AWS module. [14657](#) [15](#)
- New fileset googlecloud/firewall for ingesting firewall logs. [15328](#)
- google-pubsub input: ACK pub/sub message when acknowledged by publisher. [13346](#) [14715](#)

Functionbeat

- Remove Beta label from google-pubsub input. [13346](#) [14715](#)
- Add dashboard for AWS ELB fileset. [15804](#)
- Set `event.outcome` field based on googlecloud audit log output. [15731](#)
- Add dashboard for AWS vpcflow fileset. [16007](#)

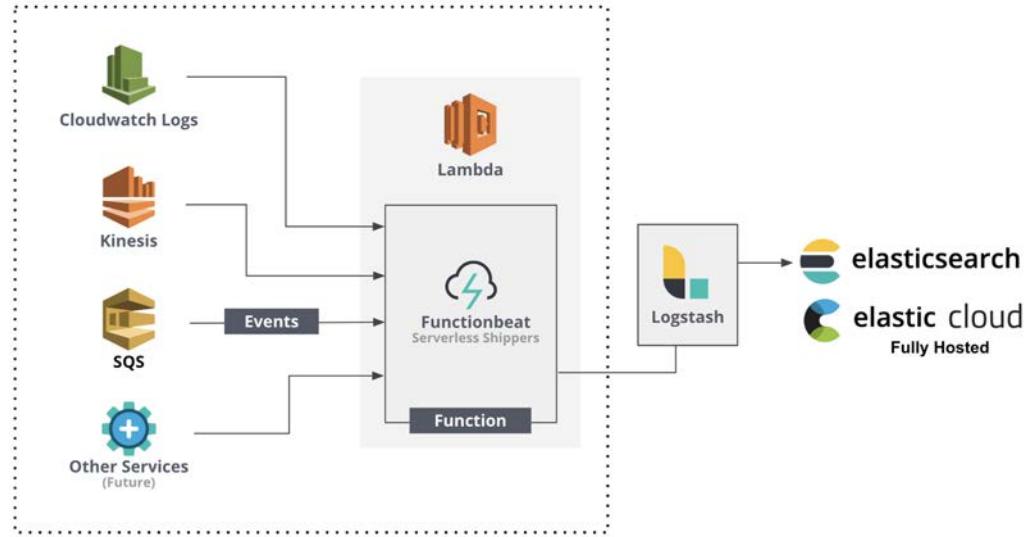
Functionbeat 機能向上

サーバレス環境のデータ収集

Logstash outputをダウンロード
のストリーム処理用に追加

構成可能なfunction tagsを追加

- グルーピングと
フィルタリング
- コスト賦課やチャージバック



```

# Tags are key-value pairs attached to the function.
#tags:
#  department: ops
  
```

functionbeat.yml

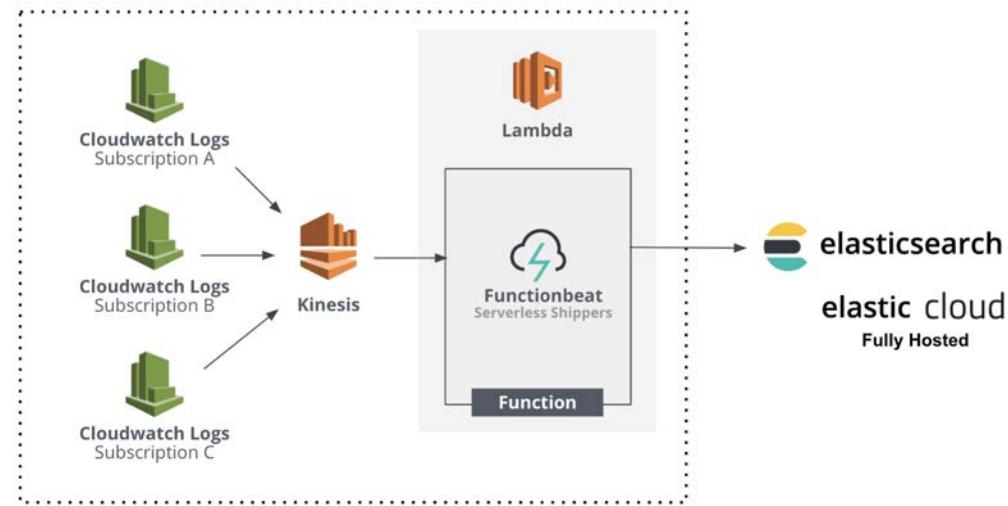
CloudWatch LogsのKinesis経由での投入

サーバレス環境のデータ収集

AWS CloudWatch LogsをKinesis経由で収集する、人気のあるクラウドモニタリングのアーキテクチャを可能に

新しい Cloudwatch Logs Kinesis function typeを追加

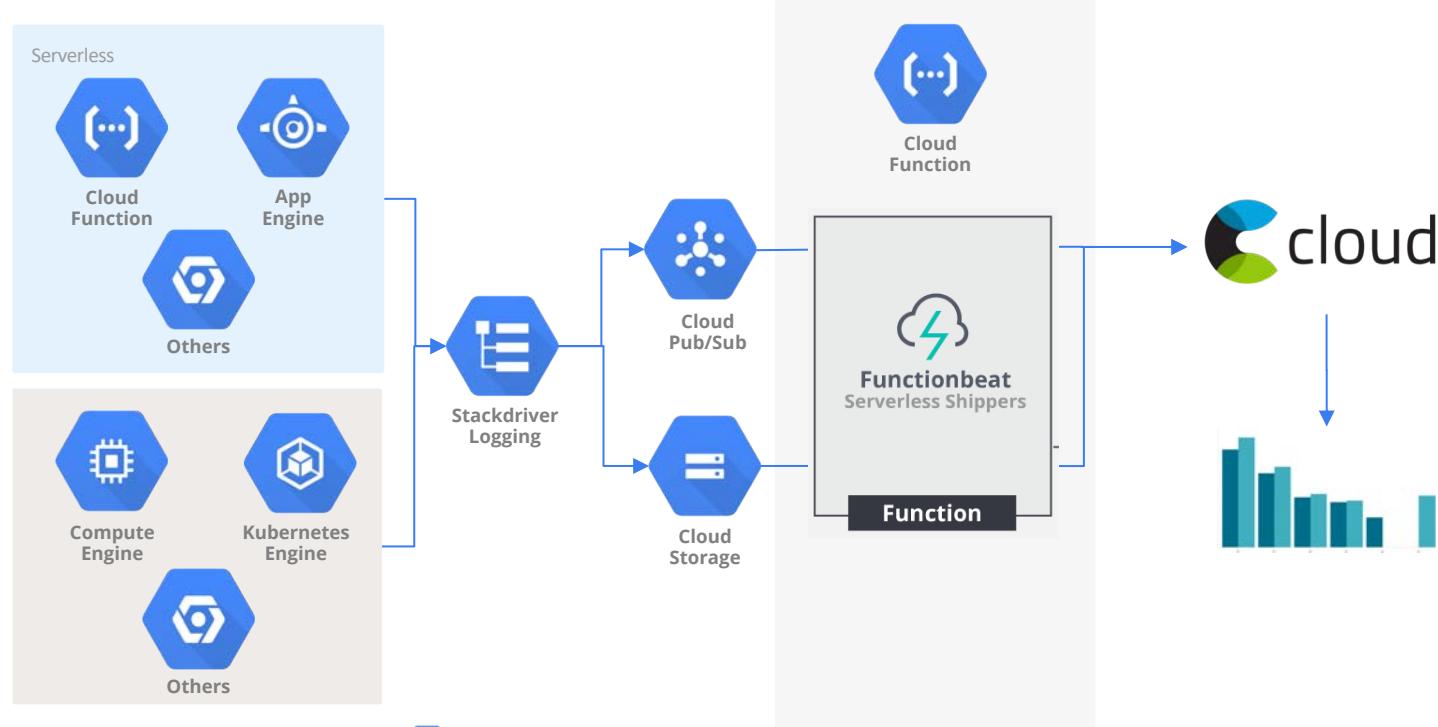
さらにマルチサブスクリプションの CloudWatch Logsモニタリングも可能に



Functionbeat for Google Function

サーバレス環境のデータ収集

<https://www.elastic.co/guide/en/beats/functionbeat/current/configuration-functionbeat-gcp-options.html>



プラットフォームサポートの拡張

Beats Feature ハイライト

新しい operating systems

- RHEL 8
- Amazon Linux 2
- Ubuntu 18.04
- Windows Server 2019



プラットフォームの拡張 = さらなるデータの拡張

Search. Observe. Protect.



Elastic エンタープライズサーチ



Elastic オブザーバビリティ



Elastic セキュリティ



Elastic Stack



Elastic オブザーバビリティ

Logs

Metrics

APM

Uptime

Elastic Approach to Observability

Dev & Ops Teams



Log Data

Web Logs
App Logs
Database Logs
Container Logs

Metrics Data

Container Metrics
Host Metrics
Database Metrics
Network Metrics
Storage Metrics

APM Data

Real User Monitoring
Txn Perf Monitoring
Distributed Tracing

Uptime Data

Uptime
Response Time

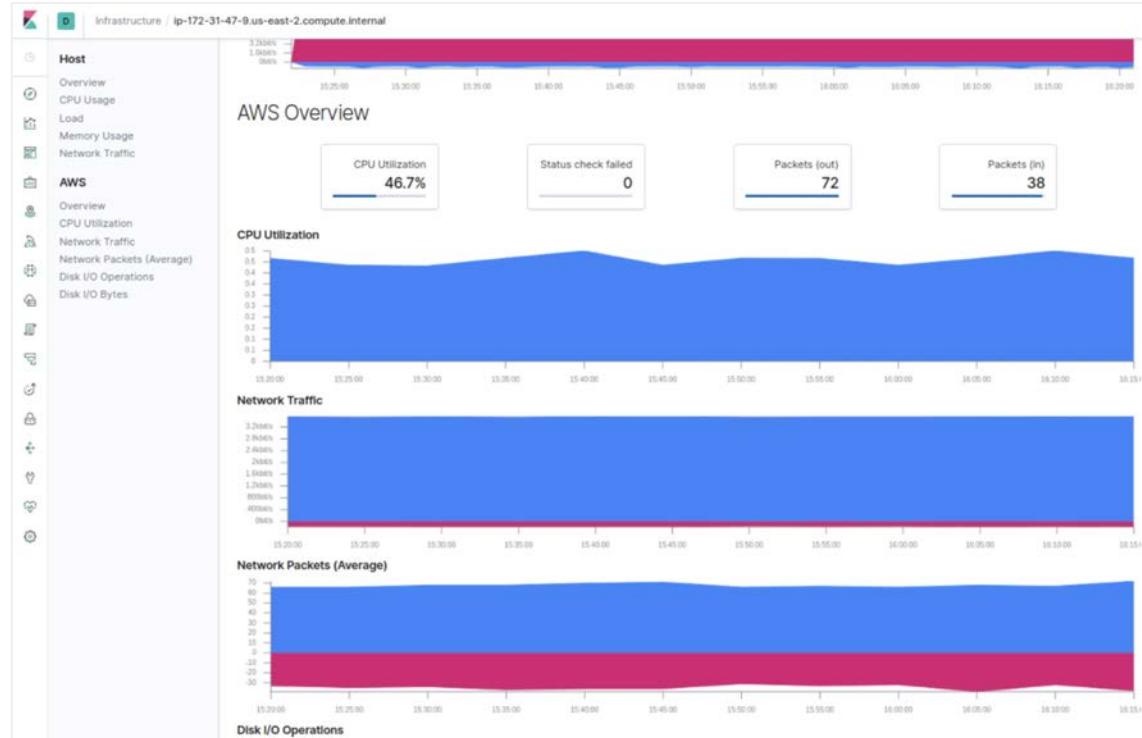
Elastic Common Schema



Cloud Monitoring

AWSオーバービューページ

AWS環境のヘルスを可視化



AWSメトリクスの統合

#welovecloud

Metricsets

- ec2
- sqs
- s3_requires/s3_daily_storage
- cloudwatch
- ebs
- elb
- rds
- sns
- sqs

CloudWatch Statistics(Metric Aggregation)



```
metricsets:
  - cloudwatch
metrics:
  - namespace: AWS/EC2
    name: ["DiskWriteOps"]
    statistic: ["Maximum", "Minimum"]
```

AWS Tags as Filters

ユーザーのコンテキストでモニタリング

Tagによってユーザーコンテキストをリソースに追加可能に

AWS moduleに新しい構成オプション

ユーザーのコンテキストでフィルター可能に

```
- module: aws
  period: 300s
  metricsets:
    - cloudwatch
  metrics:
    - namespace: AWS/EC2
      tags.resource_type_filter:ec2:instance
      statistic: ["Average"]
    tags:
      - key: "Organization"
        value: "Engineering"
```

AWS billing and usage

リソース利用状況と課金情報を素早く可視化

- billing:
AWSのEstimated Chargeを収集
 - usage:
AWS Cloudwatch API を使ってAWSリソースの利用状況を取得



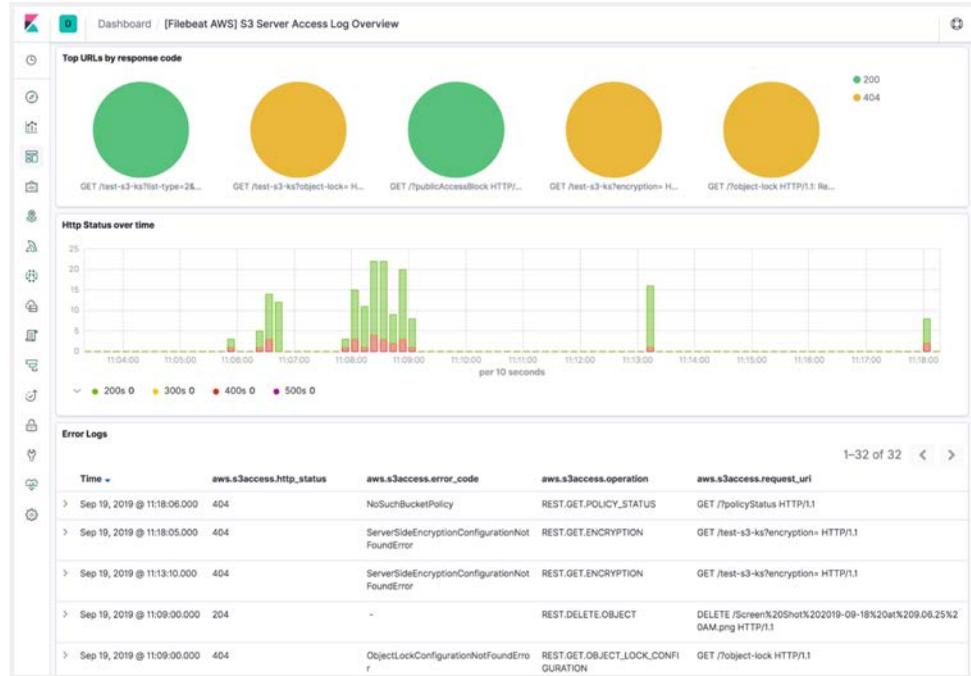
<https://www.elastic.co/guide/en/beats/metricbeat/7.6/metricbeat-metricset-aws-billing.html>

<https://www.elastic.co/guide/en/beats/metricbeat/7.6/metricbeat-metricset-aws-usage.html>

AWS S3 Server Log 向けモジュール

S3のアクセスログとS3に溜められたログを収集

- S3に溜められた各種サービスログ
 - VPC flow logs
 - ELB access logs
 - CloudTrail logs
- S3 Serverアクセスログ
 - Security Audits
 - Access Logs
 - S3の利用率を見るのに有効
- プリセットのDashboard



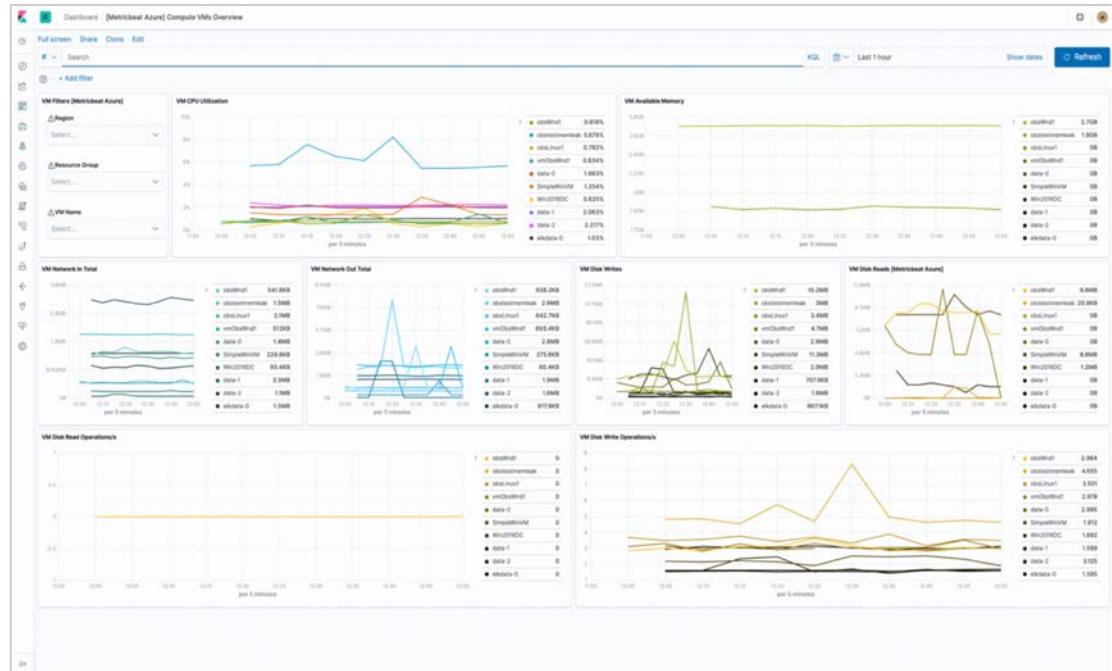
<https://www.elastic.co/guide/en/beats/filebeat/current/filebeat-input-s3.html>

<https://www.elastic.co/guide/en/beats/filebeat/current/filebeat-module-aws.html>

Azureモニタリング

Achievement unlocked!

- Azure環境をモニタリング
- Metricsets
 - monitor
 - compute_vm
 - compute_vm_scaleset
 - storage
- プリセットのDashboard
- Multi cloud monitoring
 - Azure | AWS | *



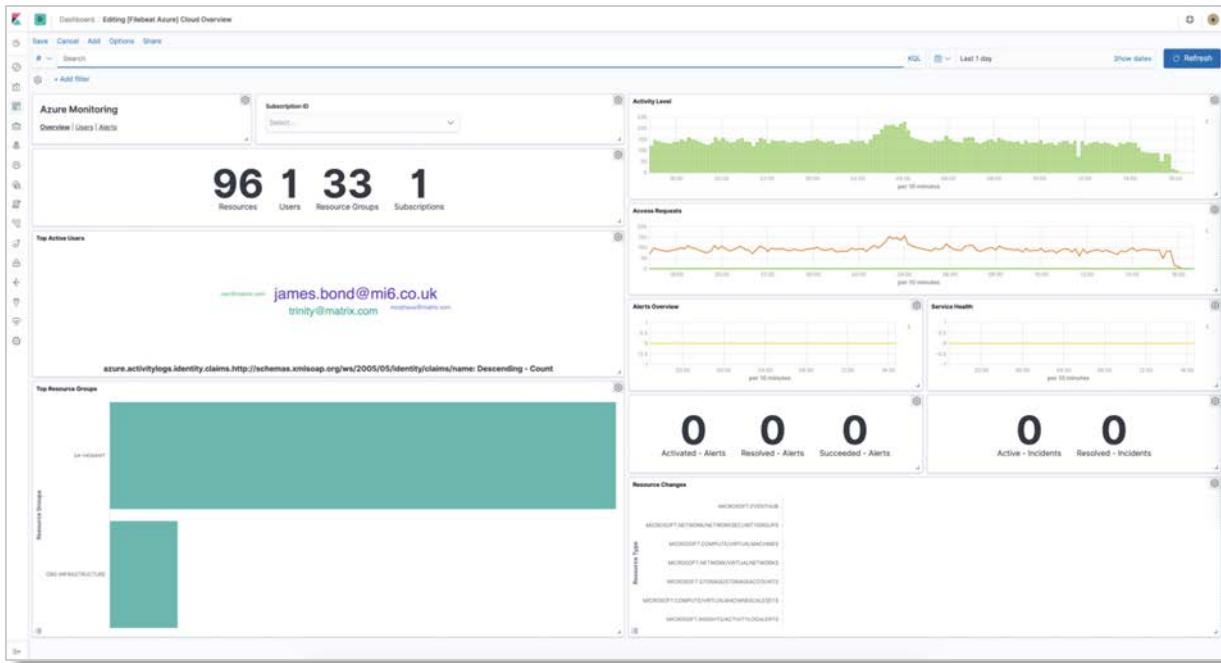
<https://www.elastic.co/blog/elastic-metrics-7-5-0-released>

<https://www.elastic.co/guide/en/beats/metricbeat/current/metricbeat-module-azure.html>

Azure Event Hub logs

サポートされるfilesets

- activitylogs
- signinlogs
- auditlogs



GCPモニタリング

Multi-Cloud環境を単一のインターフェースでモニタリング可能に

- GCP環境をモニタリング
- Metricsets
 - Compute
- Stackdriver APIsを活用
- FilebeatにもGoogle Cloud Moduleが追加
 - VPC flow logs
 - Firewall logs

```
metricbeat.modules:
  - module: googlecloud
metricsets:
  - compute
zone: "us-central1-a"
project_id: "your project id"
credentials_file_path: "your JSON credentials file path"
exclude_labels: false
period: 300s
```

<https://www.elastic.co/blog/elastic-observability-7-6-0-released>

<https://www.elastic.co/guide/en/beats/metricbeat/current/metricbeat-module-googlecloud.html>



Logs

Logs App 機能向上

さらなる機能追加

- カラム順序を自在に変更
 - ログカラムをドラッグ＆ドロップで
- キーワードのハイライト!
 - 固有のミニマップ
 - 出現数
 - 説明や推奨するアクションに活用
- UXの向上

The screenshot shows the Elastic Logs App interface. At the top, there are tabs for Stream, Analysis (Beta), and Settings. The Stream tab is selected. Below the tabs, there is a search bar and a timestamp range selector from 08/16/2019 to 08/16/2019 at 1:47:51 AM. A 'Customize' button is also present.

The main area displays log entries with columns for Timestamp and Message. The Timestamp column is currently sorted by Ascending Order (A-Z). The first few log entries are:

- Sep 16, 2019 @ 01:07:14.151 195.204.191.118 - [2018-07-30T08:07:14.151Z] "GET /styles/pretty-layout.css HTTP/1.1" 200 5448 "-" "Mozilla/5.0 (X11; Linux x86_64) AppleWebKit/534.24 (KHTML, like Gecko) Chrome/11.0.696.58 Safari/534.24"
- Sep 16, 2019 @ 01:09:27.145 230.49.131.128 - [2018-07-30T08:09:27.145Z] "GET /kibana/kibana-6.3.2-windows-x86_64.zip HTTP/1.1" 200 7185 "-" "Mozilla/5.0 (X11; Linux x86_64; rv:6.0a1) Gecko/20110842 Firefox/6.0a1"
- Sep 16, 2019 @ 01:15:21.692 239.98.179.175 - [2018-07-30T08:15:21.692Z] "GET /kibana/kibana-6.3.2-windows-x86_64.zip HTTP/1.1" 200 6249 "-" "Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.1; SV1; .NET CLR 1.1.4322)"
- Sep 16, 2019 @ 01:16:35.659 192.52.97.147 - [2018-07-30T08:16:35.659Z] "GET /beats/metricbeat HTTP/1.1" 200 9561 "-" "Mozilla/5.0 (X11; Linux x86_64; rv:6.0a1) Gecko/20110842 Firefox/6.0a1"

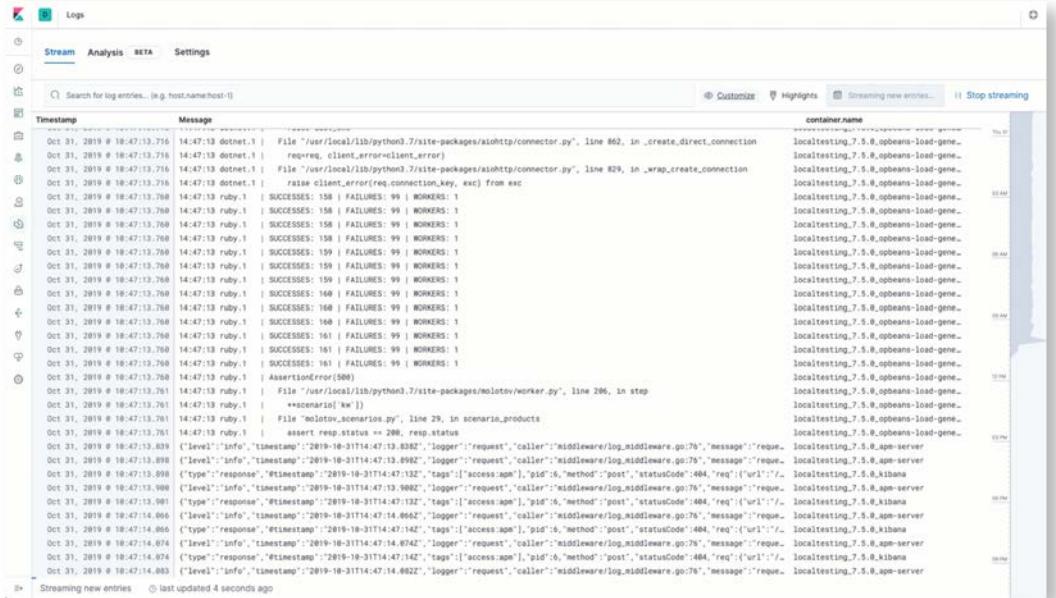
Below the logs, there is a 'Log Columns' section with a modal open. The modal shows three columns: 'Timestamp', 'Message', and 'Field'. The 'Timestamp' and 'Message' columns have descriptions and edit icons. The 'Field' column has a placeholder 'event.dataset' and a red 'X' icon for Discard and a blue 'Apply' button.

Log Rate 異常検知

どこにアドレスすべきなのか？

管理者やアナリストが、データセット毎にログレートから異常を検知することを支援

- 新しくリリースしたアプリのログが急に増えた(新規Appの影響?)
 - 既存アプリのログでスパイクが発生(アタック? プロモーション?)
 - 既存アプリのログが突如止まった(Appダウン? ログシッパーのダウン?)

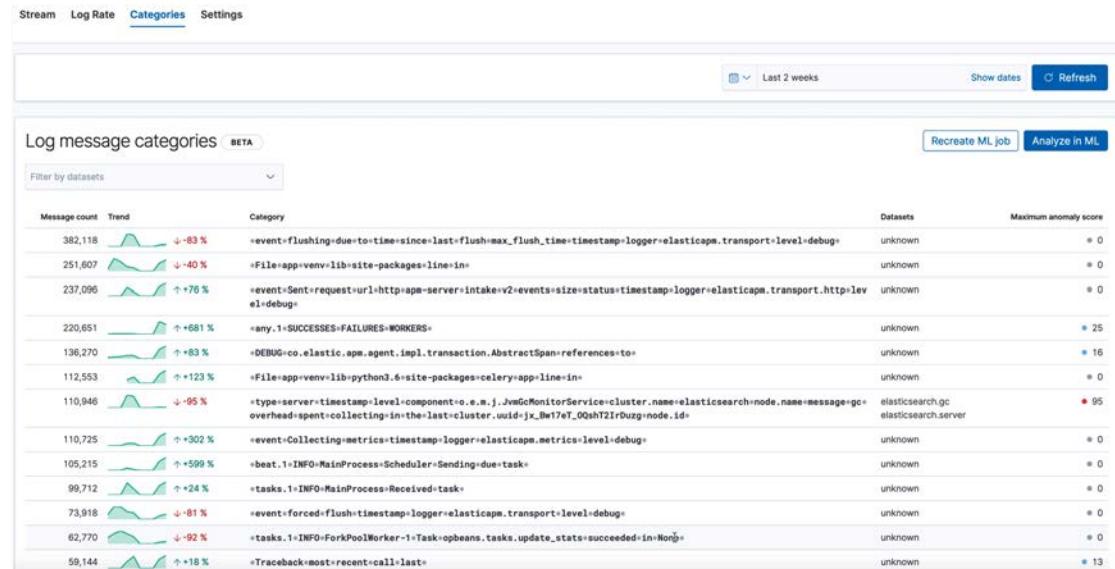


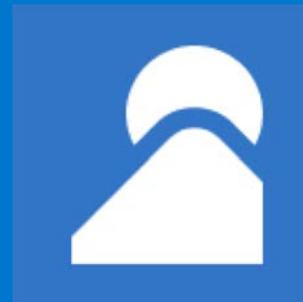
Logの自動分類

根本原因の発見をスピードアップ

ログパターンを識別し、機械学習により異常を検知

- 期待される動作を予測
機械学習を活用し異常を検知
- ログパターンを指示せずに異常を検知
ユーザーは正確なログのパターンを意識せどとも、機械学習がログを分類
- 機械学習の経験は不要
インデックスと学習用の期間を入力するだけでOK





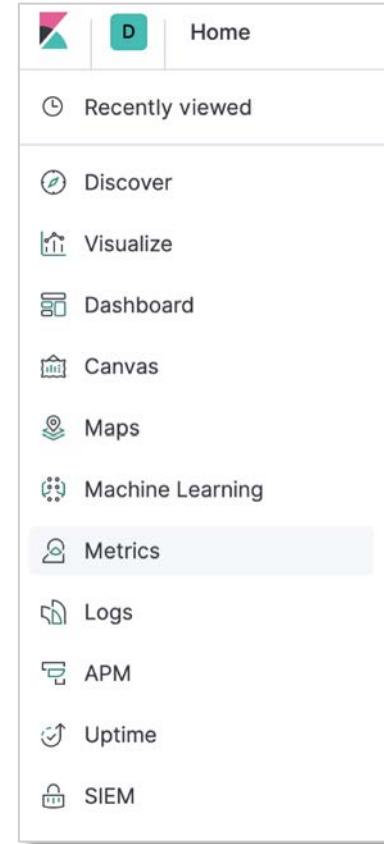
Metrics

Infrastructure app → Metrics app

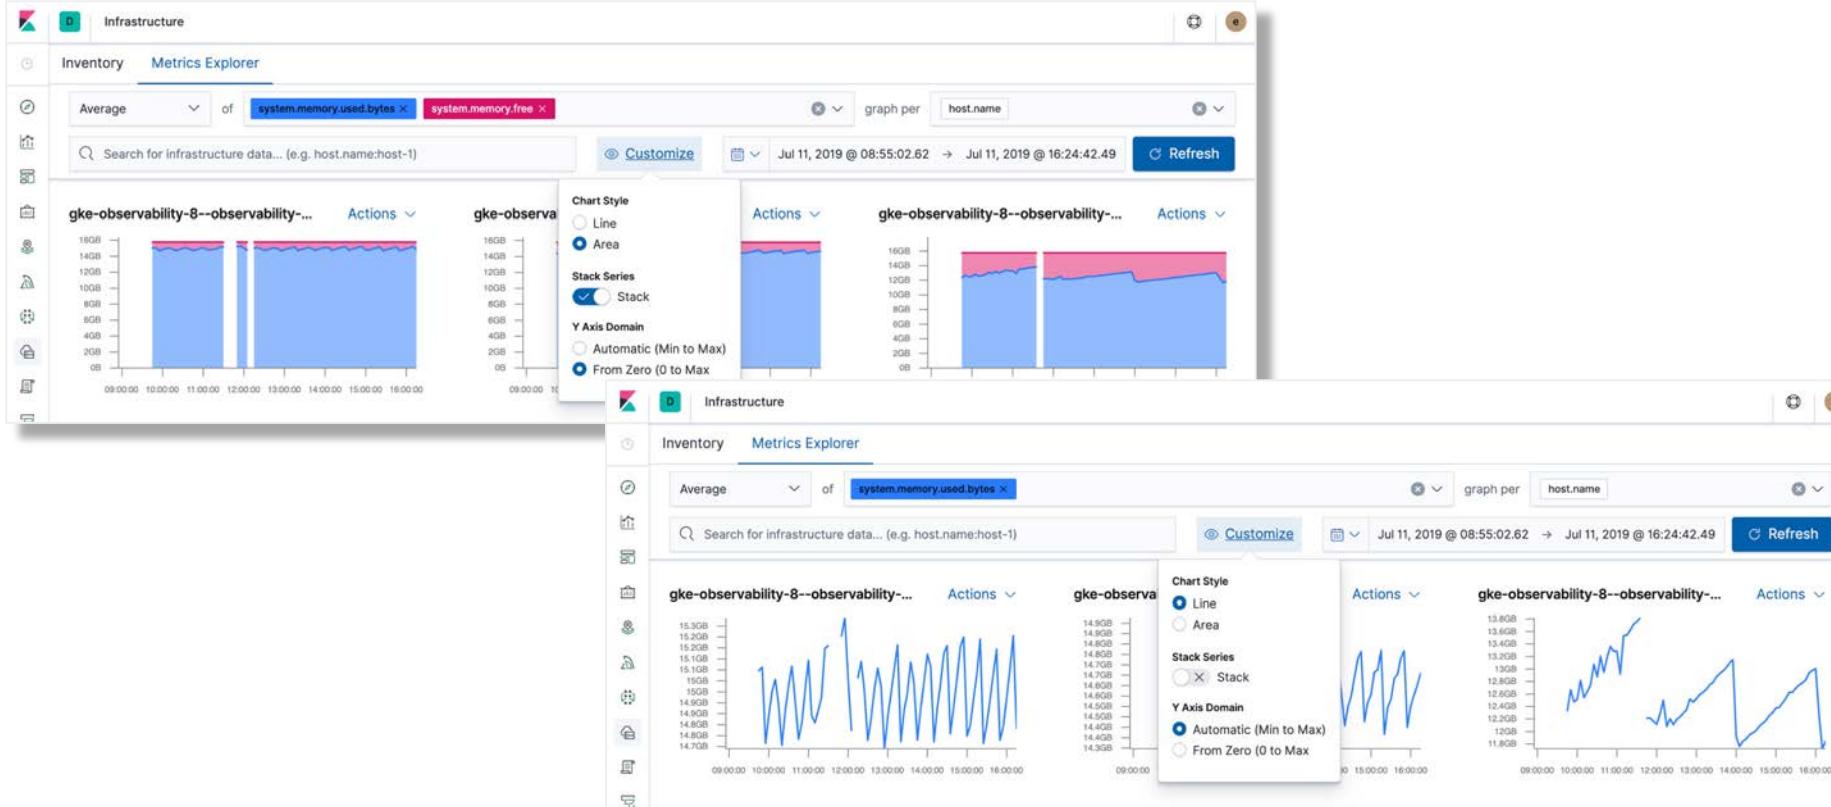
オブザーバビリティ

インフラのモニタリング以上の機能に！

The screenshot shows the Elastic website's product page. At the top, there are links for Products, Learn, Company, Pricing, Contact, and Try Free. Below this, there are three main sections: Elastic Enterprise Search, Elastic Observability, and Elastic Security. Under Elastic Enterprise Search, there are three items: App Search (Add search to your app), Site Search (Add search to your website), and Workplace Search (Search all of your content). Under Elastic Observability, there are four items: APM (Monitor your apps), Logs (Centralize, analyze logs), Metrics (Centralize, analyze metrics), and Uptime (Monitor availability). At the bottom, there is a link to 'View product story'.



Metrics Explorerがカスタマイズ可能に



Metrics Explorerのバーチャート

how do you data?

新しいチャートのオプション

カウンターや疎らな測定値を可視化するのに便利

積み上げチャートにして比較が容易に



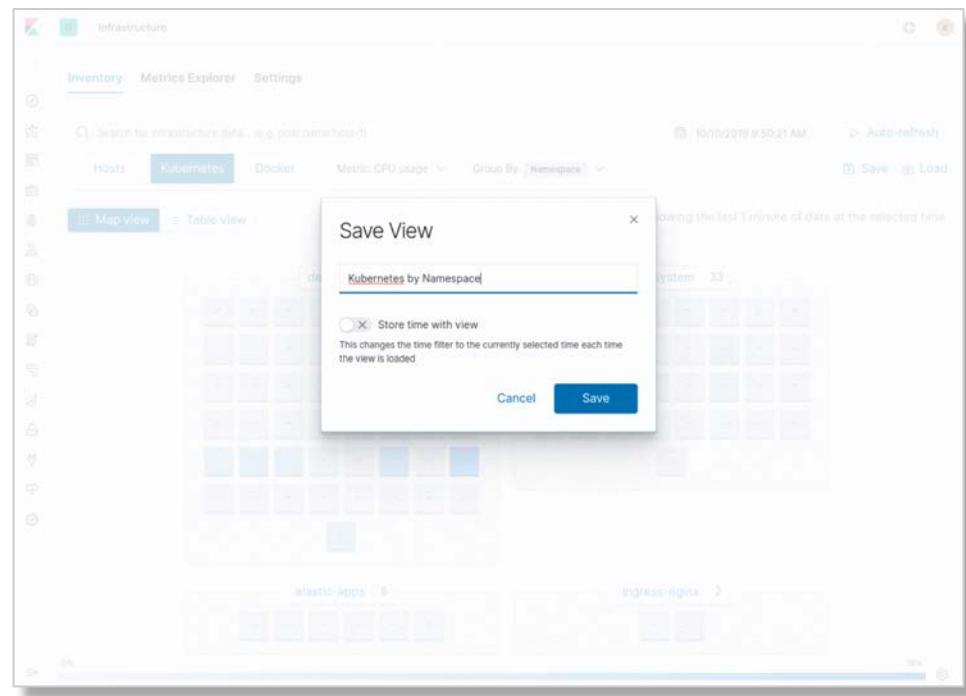
InventoryとMetrics Explorerが保存可能に

独自のプリセット

フィルター、検索、グルーピングを
保存可能に

過去の保存結果をいつでもロード可
能に

複数のリソースやメトリクスを見る
時に便利





APM

.NET Framework フルサポート

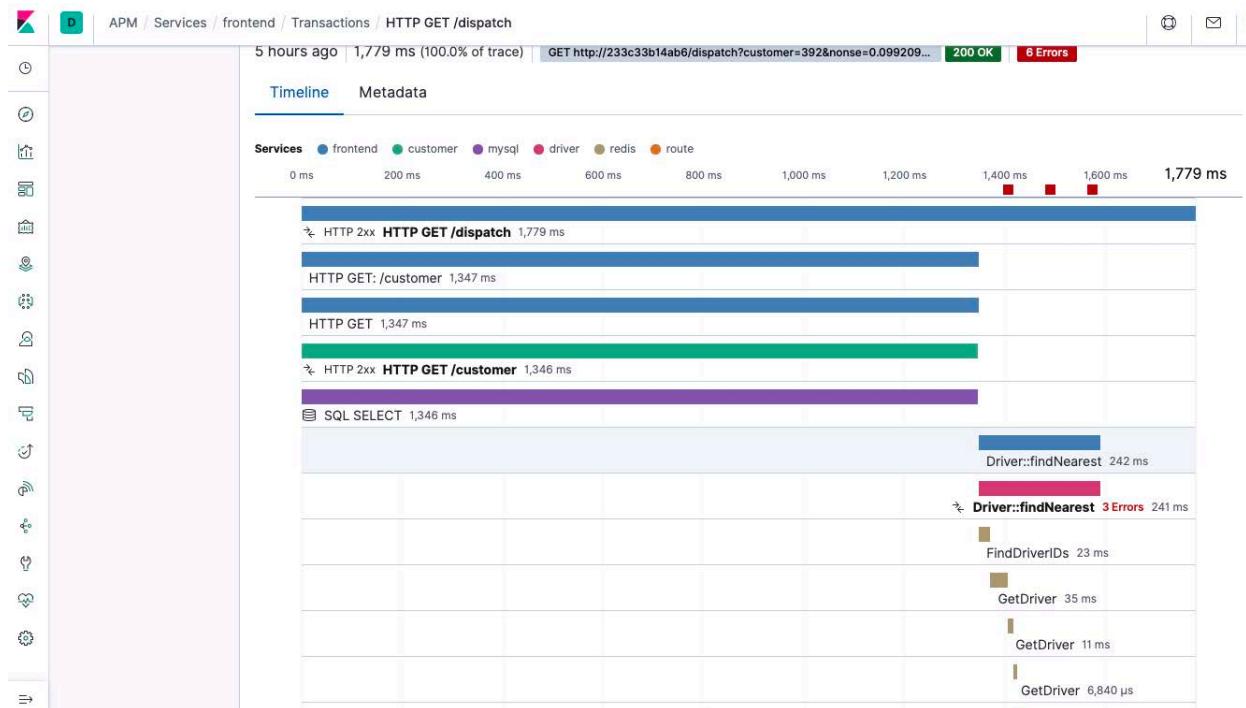
オートインストゥルメントをサポートし、プラグ&プレイの体験を

- アウトオブボックスのオートインストゥルメントを.NET Frameworkでサポート
 - トランザクションの作成し追跡するためのコードの変更は必要なし
- Supported Technologies
 - .NET Core 2.1~
 - .NET Framework 4.6.1~
 - Entity Framework Core 2.x
 - Entity Framework 6 6.2~
 - System.Net.Http.HttpClient on .NET Core
 - System.Net.Http.HttpClient on .NET Framework



Jaeger Intake

- Jaeger Agentで生成されたトレースをAPM Server経由でElasticに収集
- APM ServerのgRPCとThriftがサポートされる
- Jaegerから取り込まれたデータはElasticのデータモデルに自動的にマッピング
- トレースを受信するための新しいポートを開ける必要がある
- APM ServerはEnableに



<https://www.elastic.co/blog/elastic-apm-7-6-0-released>

<https://www.elastic.co/guide/en/apm/server/7.6/jaeger.html>

KibanaからAgent Configurationが可能に

- HTTP RequestのBodyを収集するか(CAPTURE_BODY)
- 分散トレースにおける Max. number of spans (TRANSACTION_MAX_SPANS)
- サンプルレート (TRANSACTION_SAMPLE_RATE)

Service name	Service environment	Sample rate	Capture body	Transaction max spa...	Last updated
opbeans-node	production	1 off		500	a few seconds ago
opbeans-python	production	0.5 errors		200	a few seconds ago
opbeans-ruby	All	0.02 all		100	a few seconds ago

✓ Configuration saved
The configuration for "opbeans-ruby" was saved. It will take some time to propagate to the agents.



Uptime

Kubernetes Servicesのステータス監視

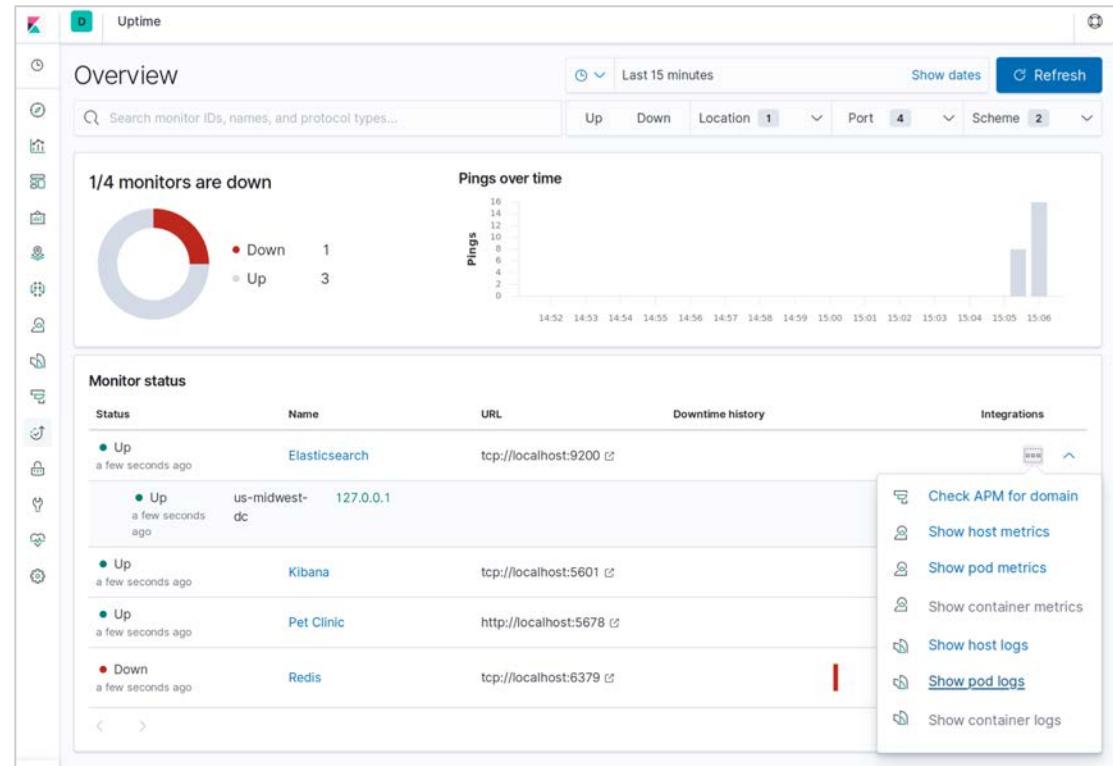
Hint-based auto-discovery

特別なラベルをPod/Containerに付与することで、HeartbeatがDockerおよびKubernetesから直接監視出来るように設定

PodもしくはContainerがスタートすると、HeartbeatがHintが設定されていないかをチェックし、適切な構成で開始

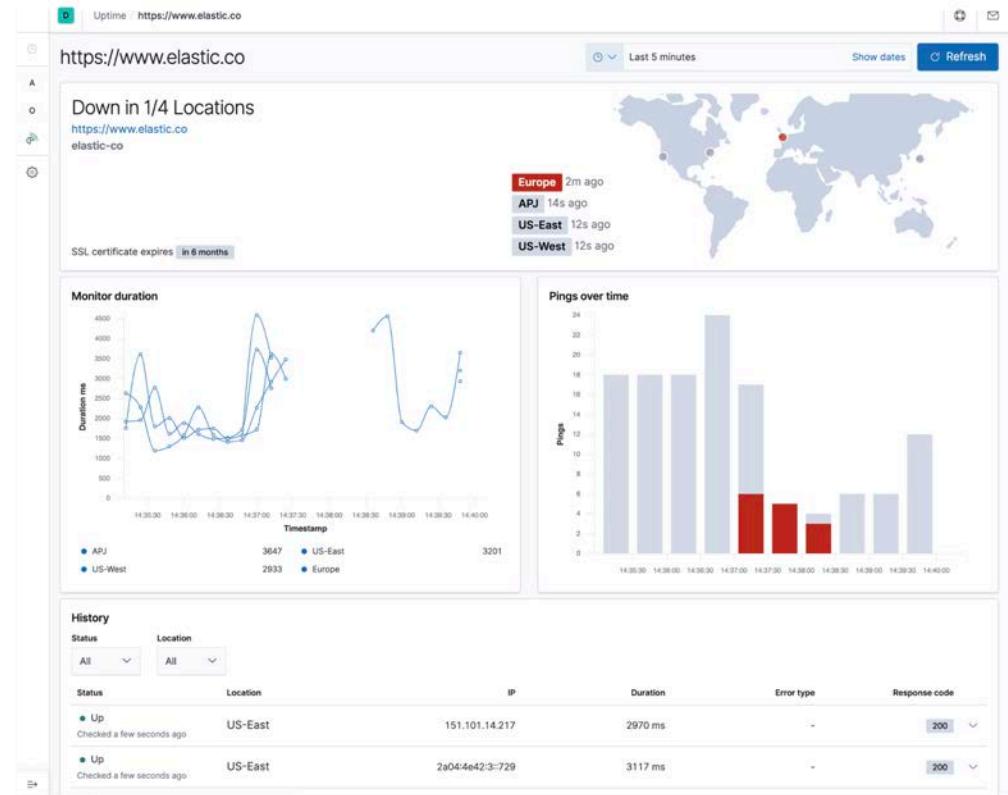


kubernetes



グローバルでの分散Uptimeチェックの可視化

異なるロケーションにおいて、
Uptimeを使って監視を行っている
場合に、一目で理解出来る
マップとステータスビューを
提供





Elastic セキュリティ

Endpoint SIEM



Elastic Endpoint Security

エンドポイントから始めるセキュリティ



ダメージを受け失う前に
マルウェアや
ランサムウェアを阻止



AIベースの
エンドポイントにおける
検知と対応



今日のハイブリッド・
クラウド環境を前提に設計

Elastic Endpoint Security

アンチウイルスソフトのように簡単で、もっとパワフル

Elastic Security

Sec Ops Teams



Block in real-time:

- Ransomware
 - Phishing
 - Exploits and malware
- Reflex custom preventions

Zero Trust data policy

Elastic Common Schema
Integrate any data source
ElasticSearch at the core

Simple alert triage

Incident visualization
ATT&CK alignment
Global ML detections
Customized detections

Instant automated response

Customized controls
One-click containment
Detect once, prevent many

Endpoint + SIEM





ENTERPRISE

Prevent

脅威をいち早くブロック

オンラインで自律的に防止

ランサムウェア、フィッシング、脆弱性への攻撃、マルウェアを厳格な第三者機関のテストで裏付けされた能力でブロック
クラウドでの分析は必要なし

MITRE ATT&CKマトリクスにマッピング

ただのペイロードではなく、何らかの破壊や損失の前に、敵対的な挙動を防止

完全にカスタマイズ可能なコントロール

固有の防止ポリシーを作成し、簡単にスケーラブルに適用可能

The screenshot shows the Elastic SIEM interface with the title "SIEM Hosts - Uncommon processes". The search bar contains "Search" and the time range is set to "Last 24 hours". Below the search bar are buttons for "All hosts", "Authentications", "Uncommon processes" (which is selected), "Anomalies", and "Events". The main area is titled "Uncommon processes" and shows a table with the following data:

Process name	Hosts	Instances	Host names	Last command	Last user
107fe9dc1222c1e	1	1	beats-ci-immutable-centos-7-1573252673329290817	/var/lib/jenkins/workspace/Beats_beats-beats-mbp_PR-14347/magefile/107fe9dc1222c1e3ee5f53192461177684059a5	jenkins
50-mold-news	1	1	internal-ci-immutable-ubuntu-1604-157331722001600487	/bin/sh	root
AM_Delta_Patch_1.305.1619.0.exe	1	1	siem-windows	C:\Windows\SoftwareDistribution\Download\insta	SYSTEM
CompatTelRunner.exe	1	1	siem-windows	C:\Windows\system32\compattelrunner.exe	SYSTEM
MpSigStub.exe	1	1	siem-windows	C:\Windows\system32\MpSigStub.exe	SYSTEM
TlWorker.exe	1	1	siem-windows	C:\Windows\winsxs\amd64_microsoft-windows-servicingstack_31bf3856ad364e35_10.0.17763.7	SYSTEM

Collect

全てのセキュリティデータ
をストアし、検索可能に

ゼロトラスト・ポリシー

耐タンパー性のためのカーネルレベルのデータ
収集とエンリッチメント

Elastic Common Schema (ECS)

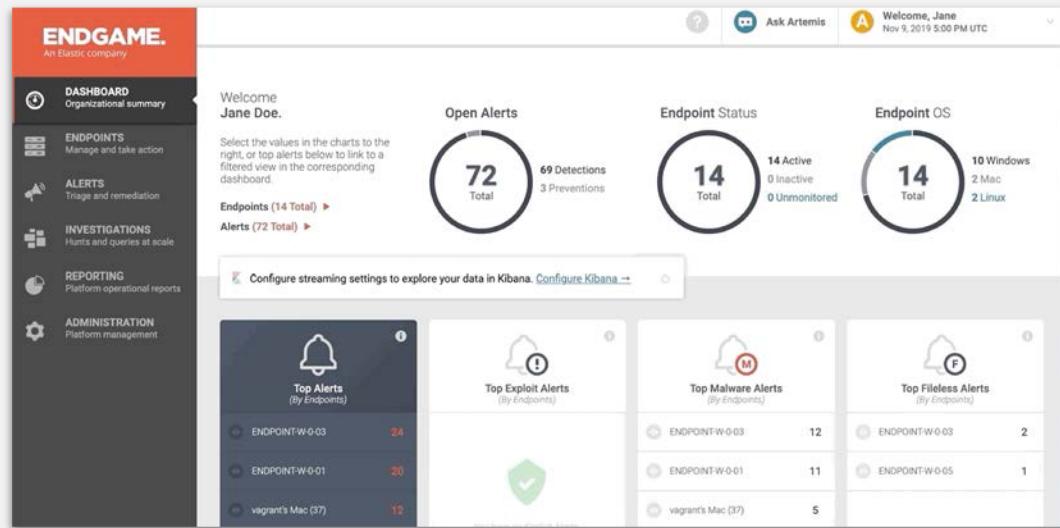
データモデルの統合のためのオープンソースベ
ースの仕様

全てのデータソースへの迅速なアクセス

制限なしに1つの製品で、セキュリティからオ
ペレーション、他のデータソースへアクセス

Elasticsearchを中心 に

全てのデータを迅速に検索



Detect

大規模に調査し、スコープを特定

容易なアラートのトリアージ

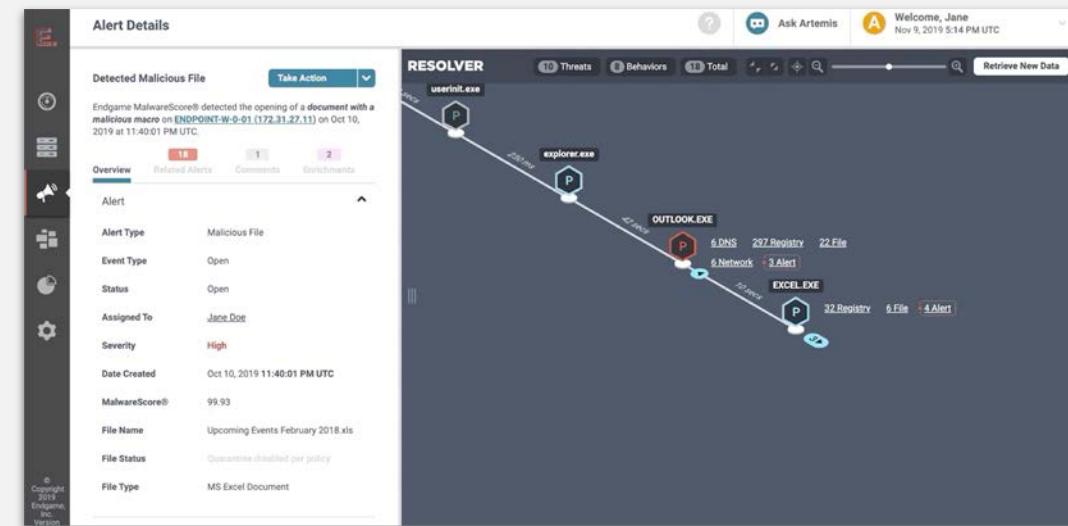
シンプルなワークフローでアラートをアサインし管理

自動的にアタックを可視化

Resolver™ Viewがアタックのスコープ特定や根本原因分析を加速し、ユーザー次の次元へ

カスタマイズした機械学習でグローバルに検知

全てのデータに対する、既にロード済の、ワンクリックの機械学習を使った分析





ENTERPRISE

Respond

修正、除外、実証

ワンクリックの封じ込め

さらなる敵対的な挙動を防ぐため、迅速にエンドポイントを隔離

リアルタイムかつ自動的なレスポンス

アタックのライフサイクルで、検知に対して自律的なミリ秒単位のレスポンス

一度検知したら何度も防止

検知した脅威を簡単に防止アクションに変換

既存のワークフローに組み込み

既存のビジネスプロセスに合わせて簡単に統合可能

Alerts							
72 All		47 Unread		Assign Alerts		Resolve Alerts	
Download JSON		Download CSV		Dismiss		Unactionable Alerts	
69	Detections	3	Preventions	72 Total	Assigned To Me	1 Quarantined	2
0 alerts currently selected	▼						
#	ALERT_TYPE	EVENT_TYPE	ASSIGNEE	OS	IP_ADDRESS	HOSTNAME	DATE
<input type="checkbox"/>	Malicious File Prevention	Open	Unassigned	Windows 7 (SP1)	172.31.27.12	ENDPOINT-W-0-02	Oct 23, 2019 9:32:56 AM UTC
<input type="checkbox"/>	Malicious File Detection	Open	Unassigned	Windows 7 (SP1)	172.31.27.13	ENDPOINT-W-0-03	Oct 19, 2019 1:51:31 PM UTC
<input type="checkbox"/>	Malicious File Detection	Execution	Anne.Anonymous	macOS High Sierra (10.13)	172.31.27.36	vagrant's Mac (37)	Oct 12, 2019 5:40:23 PM UTC
<input type="checkbox"/>	Malicious File Detection	Rename	Unassigned	macOS High Sierra (10.13)	172.31.27.36	vagrant's Mac (37)	Oct 12, 2019 5:40:23 PM UTC
<input type="checkbox"/>	Malicious File Detection	Execution	Unassigned	macOS High Sierra (10.13)	172.31.27.36	vagrant's Mac (37)	Oct 12, 2019 5:40:23 PM UTC
<input type="checkbox"/>	Malicious File Detection	Rename	Unassigned	macOS High Sierra (10.13)	172.31.27.36	vagrant's Mac (37)	Oct 12, 2019 5:40:19 PM UTC
<input type="checkbox"/>	Malicious File Detection	Execution	Unassigned	macOS High Sierra (10.13)	172.31.27.36	vagrant's Mac (37)	Oct 12, 2019 5:40:18 PM UTC
<input type="checkbox"/>	Malicious File Detection	Creation	Unassigned	Windows 7 (SP1)	172.31.27.15	ENDPOINT-W-0-05	Oct 12, 2019 5:27:31 PM UTC
<input type="checkbox"/>	Process Injection Detection	Shellcode Injection	Unassigned	Windows 7 (SP1)	172.31.27.15	ENDPOINT-W-0-05	Oct 12, 2019 5:27:31 PM UTC



Elastic SIEM

SIEM App Overview

SOCチーム向けの要約されたワークフロー

Elastic Endpointとの統合

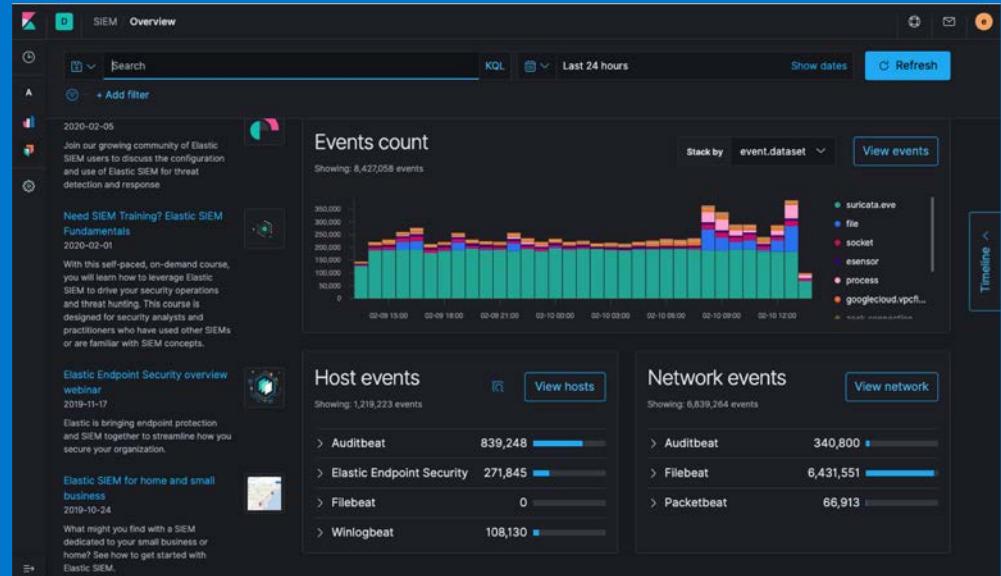
Endgameセンサーマネジメントプラットフォームからのセキュリティイベントがendgame-* indicesとして統合

新しいUI

トップソース/デスティネーション・カントリー、TLSウィジェット、イベントヒストグラムなど

新しい機械学習Job

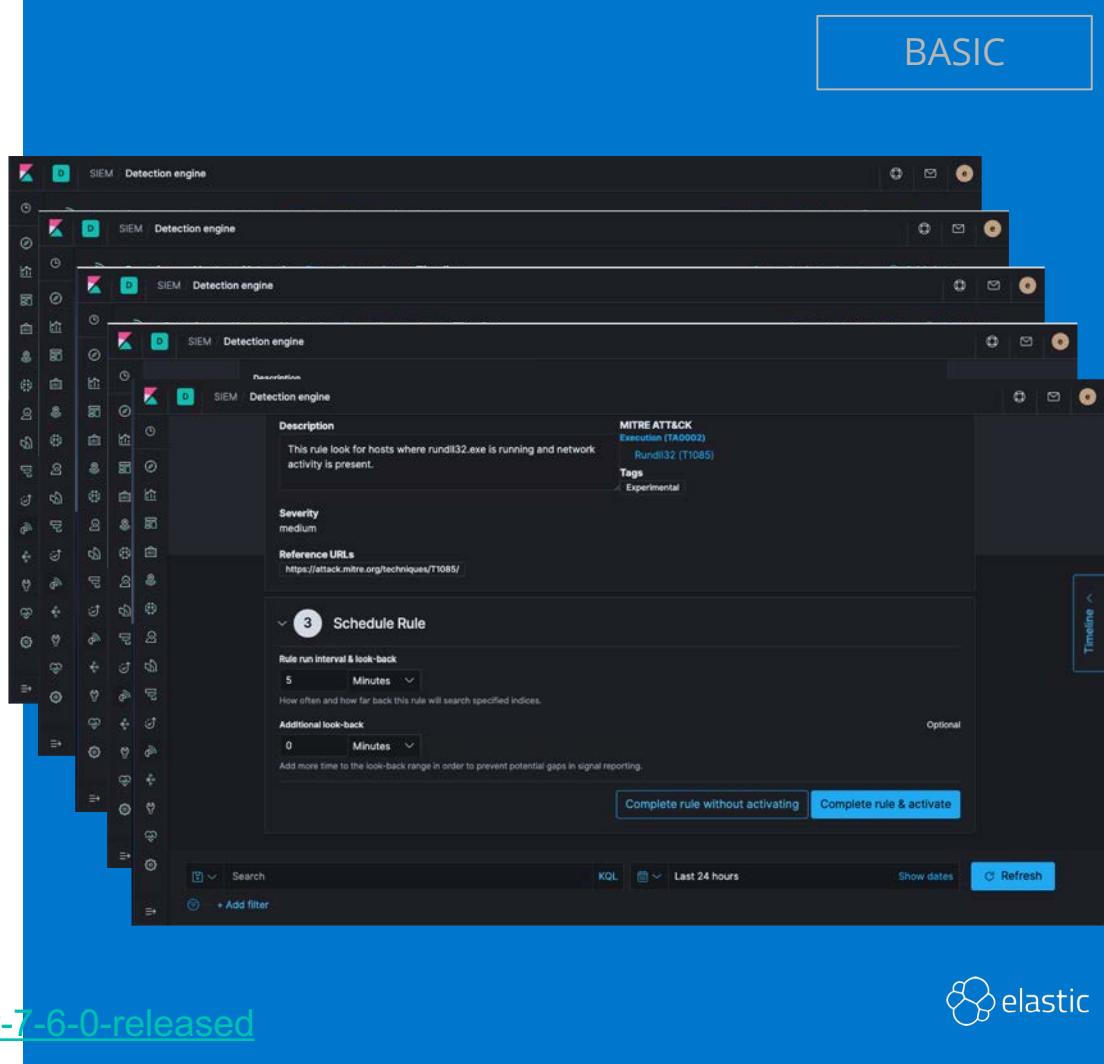
PacketbeatとWinlogbeatイベント向けの新しい機械学習Jobの追加



7.6 SIEM Detection Engine

ルールとシグナルの要約と管理

- MITRE ATT&CKに準拠したアウトオブボックスのカスタムルール
- 定期的にルール実行され、シグナルを生成
- SIEM app Timelineにてシグナルの調査が可能
- シグナルはさらに複雑な検知のためのビルディングブロック





Elastic エンタープライズサーチ

Workplace Search

App Search

Site Search



Elastic Site Search

Create and manage a tailored search experience for your website with world-class relevance, intuitive customization, and rich analytics.

Swiftype
My Account

Overview

Search Preview

Metrics

Analytics

Initiatives

Introducing Swiftype Conversion

Swiftype can track how search queries lead to any action you want — like reading an article or anything else you might imagine.

[Get Started](#) [Hide This](#)

Customize your search



Elastic Workplace Search

Instantly search across all the tools you use at work, including Dropbox, Salesforce, G Suite, Office 365, Zendesk, and more.

Recently Updated by You

KM World Notes

GOOGLE DOC Last updated 5 min

Current Active Accounts

GOOGLE DOC Last updated 10 min

Sprint Planning

in 30 minutes

Quin Hoxie, Brian Stevenson and 8 more

View on Google Calendar

Single letter search analytics

Submitted by Capsule Support

Hello! I've been going through our analytics and I was wondering about these single letter searches that result in no click-through.

Imgix - Proposals

4/7/17 at 3:00pm

HELPSCOUT ACTIVE Last updated 16 hours ago



Elastic App Search

A powerful set of APIs and developer tools designed for developers building rich, user-facing search applications.

app-prod-35

Overview

Analytics

Query Tester

Documents

Schemas

API Logs

SEARCH SETTINGS

Synonyms

Result Pinning

Weights & Boosts

ACCESS

API Keys

Analytics

Total Queries 2,021

Total Queries with No Results 251

200 400 600 800 1000 1200 1400 1600 1800 2000 2200 2400 2600 2800 3000



全てのESS regions
で利用可能に

デプロイメントの詳細

- App Search テンプレートを介して利用可能

The screenshot shows the 'Optimize your deployment' section of the Elastic Cloud interface. It lists several deployment types:

- I/O Optimized** (Recommended): Use for search and general all-purpose workloads. Includes a balance of compute, memory, and storage.
Default specs
- Compute Optimized**: Run CPU-intensive workloads or run smaller workloads cost-effectively when you need less memory and storage.
Default specs
- Memory Optimized**: Perform memory-intensive operations efficiently, including workloads with frequent aggregations.
Default specs
- Hot-Warm Architecture**: Use for time-series analytics and logging workloads that benefit from automatic index curation.
Default specs

A box highlights the **App Search** template, which allows creating search experiences with a refined set of APIs and tools. Below the templates, a note states: "Elastic Cloud supports many more options to cater to your specific use case such as hot-warm architecture optimized for logging, compute-focused setup optimized for analytics etc. [Learn more](#)".

Deployment pricing

Hourly rate	\$0.4253
-------------	----------

[Create deployment](#) [Customize deployment](#)

デプロイメントの構成

 App Search 1 configurations
Add refined search experiences to your applications.

gcp.appsearch.1 Application server Worker
A CPU optimized App Search instance.

Fault tolerance
 1 zones 2 zones 3 zones

RAM per Instance

2 GB 4 GB 8 GB

Instances ? RAM per Zone
1 × 2 GB = 2 GB

Summary
2 GB RAM × 1 instances × 2 zones = 4 GB RAM

> User setting overrides

[Select template](#) [Create deployment](#)



デプロイメントの詳細



Elasticsearch

Stores and queues everything for App Search

Based on high-CPU hardware

1–60GB x 32 nodes x 1–3 Availability Zones



App Search

Serves the application

Also based on high-CPU hardware

2–8GB x 16 nodes x 1–3 Availability Zones



Kibana

Needed for stack upgrades

Can also provide user management

1 GB x 1 AZ (free tier) sufficient

Elastic Workplace Search

is Elasticsearch and much more

A complete solution for rolling out a personalized, centralized, secure search experience

Search clients portfolio & workflow integrations

Showing results for email marketing pdfs.

Dropbox 12

Email Marketing Tools for the Modern Enterprise.pdf
in /Internal Operations/Vendor Material
Last updated by Peggie Labadie on 3/27/17 at 9:12am

Confluence 8

Email Marketing Education

Advanced Segmentation.pdf
in /Marketing/Email Marketing Education
Last updated by Swifttype Demo on 3/28/17 at 10:06am

People 6

Multivariate Testing.pdf
in /Marketing/Email Marketing Education
Last updated by Swifttype Demo on 3/28/17 at 10:06am

Developer Docs 5

Comparative Reports.pdf
in /Marketing/Email Marketing Education
Last updated by Jessica Smith on 3/28/17 at 10:06am

GitHub 3

Box 1

JIRA 1

Google Drive 1

EXTENSION

pdf 7

png 2

txt 1

svg 1

Show More

Showing results for iphone images in dropbox from Dropbox.

iPhone Front Gold.png
in /Design/Resources/Angle Mockups/PNG/**iPhone Front Gold.png**
Last updated by Swifttype Demo on 3/31/16 at 9:27pm

iPhone Front Silver.png
in /Design/Resources/Angle Mockups/PNG/**iPhone Front Silver.png**
Last updated by Swifttype Demo on 3/31/16 at 9:27pm

iPhone 04 Gold.png
in /Design/Resources/Angle Mockups/PNG/**iPhone 04 Gold.png**
Last updated by Swifttype Demo on 3/31/16 at 7:52pm

Management Interfaces

Manage Groups

Assign shared content sources and users to groups to create relevant search experiences for various internal teams.

Add a Group

Organization Groups

Showing 1-10 of 15 Groups

GROUP	CONTENT SOURCES	USERS	MANAGE
Marketing	No shared sources		Manage
Data Science		+ 5	Manage
Developer Experience Team			Manage
Design		+ 3	Manage
Mergers & Acquisitions	No shared sources	No users	Manage
Engineering		+ 22	Manage
SecOps	No shared sources	No users	Manage
Quality Assurance - Platform	No shared sources	No users	Manage
Product Management		No users	Manage
Investor Relations		+ 1	Manage

Out-of-the-box Connectors



and more

Deployment

Elastic Cloud on Kubernetes (ECK) 1.0 is now generally available

By Anurag Gupta

Share



Today, we're proud to announce that Elastic Cloud on Kubernetes (ECK) is moving out of beta and into general availability! With ECK, users now have a seamless way of deploying, managing, and operating the Elastic Stack on Kubernetes. Learn more on [our ECK product page](#), or [download ECK to get started](#).

As we announced with the alpha release of ECK back in May 2019, our vision for ECK is to provide an official way to orchestrate Elasticsearch on Kubernetes and provide a SaaS-like experience for Elastic products on Kubernetes. Kubernetes has continued to grow in popularity and has become the standard for orchestrating container workloads, and we've seen a growing number of users deploying the Elastic Stack on Kubernetes. That's why we've taken a number of steps to support container workloads, such as releasing [official Docker images](#) for Elasticsearch and Kibana, joining the CNCF, and [launching our Helm charts](#). Bringing ECK into general availability is the exciting next step on this journey.

The initial alpha release of ECK built on our years of operational knowledge gained from creating Elasticsearch and Elastic Cloud Enterprise and running our Elasticsearch Service. The community reception to the first alpha release (and the three early access releases that followed) has been extremely positive, and with the general availability of ECK we're excited to give our users a production-ready solution to deploy and



Elastic Cloud on Kubernetes

Orchestrate Elasticsearch on Kubernetes today.

[Download now](#)

Recommended Content

The screenshot shows a user interface for visualizing log data. At the top, there's a title "Visualize data from a log file" with a "More" link. Below that is a description: "The File Data Visualizer helps you understand the fields and metrics in a log file. Upload your file, analyze its data, and then choose whether to import the data into an Elasticsearch index." There are several input fields and checkboxes: "File to analyze", "Log file separator", "Log file encoding", "Log file offset", and "Log file size". A note says "You can upload files up to 100 MB. This feature is experimental, but awesome! Please leave us a review in GitHub." At the bottom, there's a "Visualize" button and a "Drop or drag and drop a file" placeholder.

ECKとは？

<https://www.elastic.co/products/elastic-cloud-kubernetes>

- Custom Resources Definitions (**CRD**)
 - Elasticsearch, Kibana, APM を追加
 - これらのリソースでKubernetes APIを拡張
- 一連の**controllers**
 - apiserver上のリソースを監視
 - 関連するリソースのCRUD操作
 - 実行中のESクラスタとの相互運用

→ Basicライセンスで使えます！

```
apiVersion:.elasticsearch.k8s.elastic.co/v1alpha1
kind: Elasticsearch
metadata:
  name: elasticsearch-sample
spec:
  version: 7.3.0
  nodes:
    - name: default
      nodeCount: 1
---

apiVersion: kibana.k8s.elastic.co/v1alpha1
kind: Kibana
metadata:
  name: kibana-sample
spec:
  version: 7.3.0
  nodeCount: 1
  elasticSearchRef:
    name: "elasticsearch-sample"
```

2020年2月19日

ニュース

EN JP

Microsoft Azure、東京でElasticsearch Serviceの提供を開始

著者 [Pieter Humphrey](#)

Share



Elastic CloudのElasticsearch Serviceの提供を、Azure東京（東日本）リージョンで開始いたしました。アジア太平洋で2番目、グローバルで5番目のAzureリージョンとなります。

現在サービスをご利用のお客様は、[ログイン](#)するだけで、すぐにAzure東京リージョンのElasticsearch Serviceをお使いいただくことができます。サービスをはじめてお使いになる方は、14日間の無料トライアルに[登録](#)してお試しいただけます。

パワフルなセキュリティやインデックスライフサイクル管理、機械学習など、Elasticならではのすぐれた機能が揃っています。Kibana Lensで直感的に可視化したり、Canvasでクリエイティブなプレゼンテーションを作成してみましょう。使いやすいスライダー、ボタンやデプロイテンプレートで、プロビジョニングから設定、デプロイのスケールまで手軽に実施できます。アプリ検索やロギング、メトリック、APM、SIEM、BI/分析など、多彩なソリューションをご用意しています。

Elasticで日本のカントリーマネージャーを務める川崎友和は、以下のように述べています。

「ITのモダナイゼーション、DXが進行する中、オンプレミスや従来型のデータセンターからパブリッククラウドに

Recommended Content

Best-in-Class Search Technology

Feature-rich search technology for your website.



Fully customizable user interface built on React.js.

Engineering team dedicated to making search accessible to everyone, from engineers to non-engineers.

Google Site Search Migration Video

We've made replacing your GSS installation easier than ever. Learn how to migrate from GSS to Elastic Site Search in four...

[Learn More](#)





Thank You !
