What's new in the Elastic Stack?

8.x edition & AMA

Philipp Krenn | @xeraa

Alexander Reelsen | @spinscale

elastic

Late 7.x

- Observability: Google Cloud Logs Integration, User Experience & Synthetic Monitoring
- Enterprise Search: Kibana Integration, Crawler
- Security: Osquery integration
- Stack: Runtime fields & Searchable Snapshots

elastic

8.x

- Observability: CI/CD, AWS Lambda visibility
- Enterprise Search: SharePoint Connector
- Stack: **Security on by default**, ANN & NLP, Field usage & Disk usage APIs, Lucene 9
- Geo: Hex tiles, Vector tiles
- ARM support

elastic

Observability

CI/CD

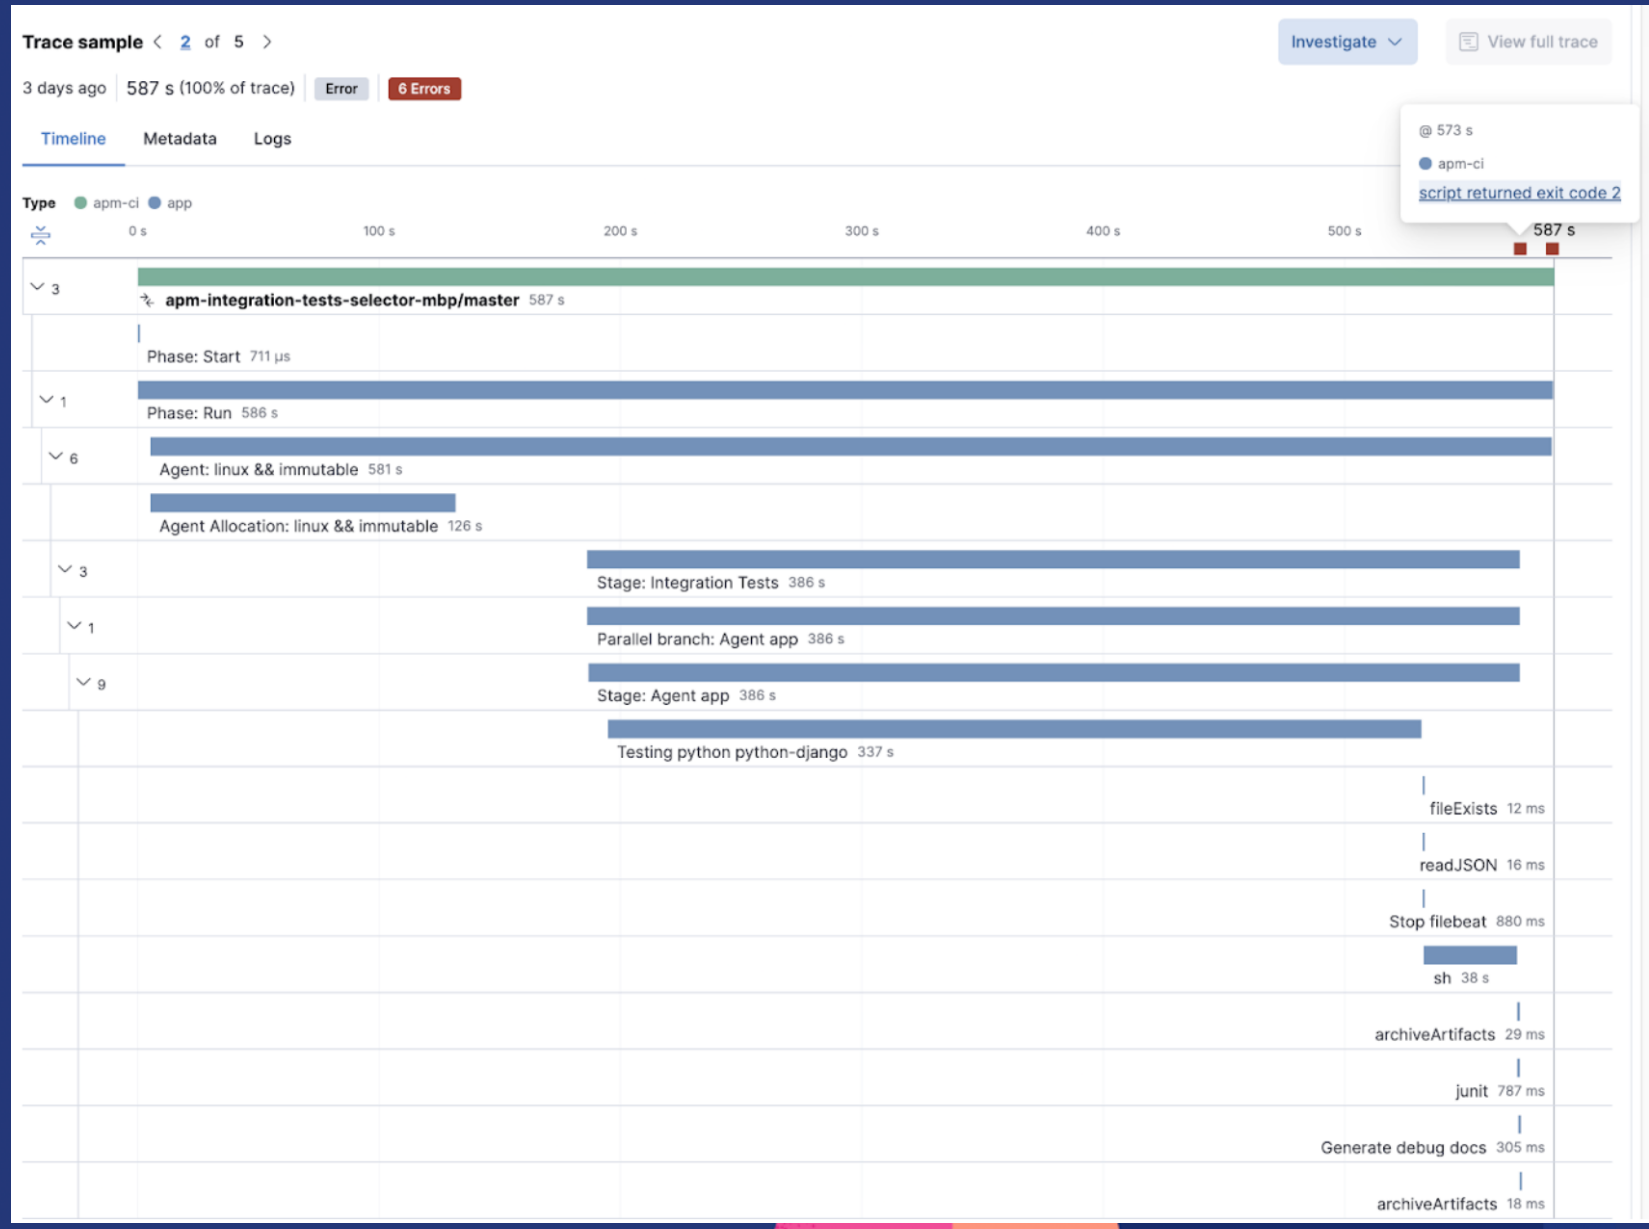

elastic

Security-on-by-default

- Whole stack!
- TLS
- Authorization + Authentication
- Running with Testcontainers

elastic

knn search

- Using `dense_vector` fields for efficient vector search
- HNSW (Hierarchical Navigable Small World)
- `8.2`: Filtering support
- Attention: **Dedicated endpoint**!

elastic

NLP - Language detection

```
POST _ingest/pipeline/_simulate
{
  "pipeline": {
    "processors": [
      {
        "inference": { "model_id": "lang_ident_model_1", "inference_config": { "classification": {} } }
      }
    ]
  },
  "docs": [
    {
      "_source": {
        "text": "This is an english text, albeit rather short."
      }
    },
    {
      "_source": {
        "text": "Bitte melden Sie sich schnellstmöglich bei uns. Wir sind jederzeit telefonisch zu erreichen."
      }
    }
  ]
}
```

elastic

NLP - Language detection


```
{
  "docs" : [
    {
      "doc" : {
        "_source" : {
          "text" : "This is an english text, albeit rather short.",
          "ml" : {
            "inference" : {
              "prediction_score" : 0.997576998993641, "model_id" : "lang_ident_model_1", "prediction_probability" : 0.997576998993641,
              "predicted_value" : "en"
            }
          }
        }
      }
    },
    {
      "doc" : {
        "_source" : {
          "text" : "Bitte melden Sie sich schnellstmöglich bei uns. Wir sind jederzeit telefonisch zu erreichen.",
          "ml" : {
            "inference" : {
              "prediction_score" : 0.9999990788535671, "model_id" : "lang_ident_model_1", "prediction_probability" : 0.9999990788535671,
              "predicted_value" : "de"
            }
          }
        }
      }
    }
  ]
}
```


elastic

NLP - Named Entity Recognition

```
# sentence transformers require cmake
pip3 install eland tqdm torch transformers sentence_transformers
```

```
~/Library/Python/3.8/bin/eland_import_hub_model \
  --url "https://elastic:s3cr3t@my-elasticsearch-endpoint.europe-west3.gcp.cloud.es.io:9243" \
  --hub-model-id elastic/distilbert-base-cased-finetuned-conll03-english \
  --task-type ner \
  --start
```

elastic

NLP - Named Entity Recognition

```
POST _ingest/pipeline/_simulate
{
  "pipeline": {
    "processors": [
      {
        "inference": {
          "model_id": "elastic__distilbert-base-cased-finetuned-conll03-english",
          "inference_config": { "ner": { } }
        }
      }
    ]
  },
  "docs": [
    { "_source": { "text_field": "I've been living in Munich for 15 years." } },
    { "_source": { "text_field": "One of the nicest places in Munich is the Englischer Garten." } },
    { "_source": { "text_field": "Been working at Elastic for more than 9 years." } }
  ]
}
```

elastic

NLP - Named Entity Recognition


```
{
  "doc" : {
    "_index" : "_index", "_id" : "_id",
    "_source" : {
      "text_field" : "I've been living in Munich for 15 years.",
      "ml" : {
        "inference" : {
          "model_id" : "elastic__distilbert-base-cased-finetuned-conll03-english",
          "entities" : [
            {
              "start_pos" : 20, "end_pos" : 26,
              "class_name" : "LOC", "class_probability" : 0.9992382591441749,
              "entity" : "Munich"
            }
          ],
          "predicted_value" : "I've been living in [Munich](LOC&Munich) for 15 years."
        }
      }
    }
  }
}
```


elastic

NLP - Named Entity Recognition


```
{
  "doc" : {
    "_index" : "_index", "_id" : "_id",
    "_source" : {
      "text_field" : "One of the nicest places in Munich is the Englischer Garten.",
      "ml" : {
        "inference" : {
          "model_id" : "elastic__distilbert-base-cased-finetuned-conll03-english",
          "entities" : [
            {
              "start_pos" : 28, "end_pos" : 34,
              "class_name" : "LOC", "class_probability" : 0.9993211907123687,
              "entity" : "Munich"
            },
            {
              "start_pos" : 42, "end_pos" : 59,
              "class_name" : "ORG", "class_probability" : 0.9521962770340935,
              "entity" : "Englischer Garten"
            }
          ],
          "predicted_value" : "One of the nicest places in [Munich](LOC&Munich) is the [Englischer Garten](ORG&Englischer+Garten)."
        }
      }
    }
  }
}
```


elastic

NLP - Named Entity Recognition


```
{
  "doc" : {
    "_index" : "_index", "_id" : "_id",
    "_source" : {
      "text_field" : "Been working at Elastic for more than 9 years.",
      "ml" : {
        "inference" : {
          "model_id" : "elastic__distilbert-base-cased-finetuned-conll03-english",
          "entities" : [
            {
              "start_pos" : 16, "end_pos" : 23,
              "class_name" : "ORG", "class_probability" : 0.9996217159395848,
              "entity" : "Elastic"
            }
          ],
          "predicted_value" : "Been working at [Elastic](ORG&Elastic) for more than 9 years."
        }
      }
    }
  }
}
```


elastic

More NLP...

- Check out the annotated text field type for the `predicted_value` field
- Check out the third party model documentation

elastic

Upgrade Assistant

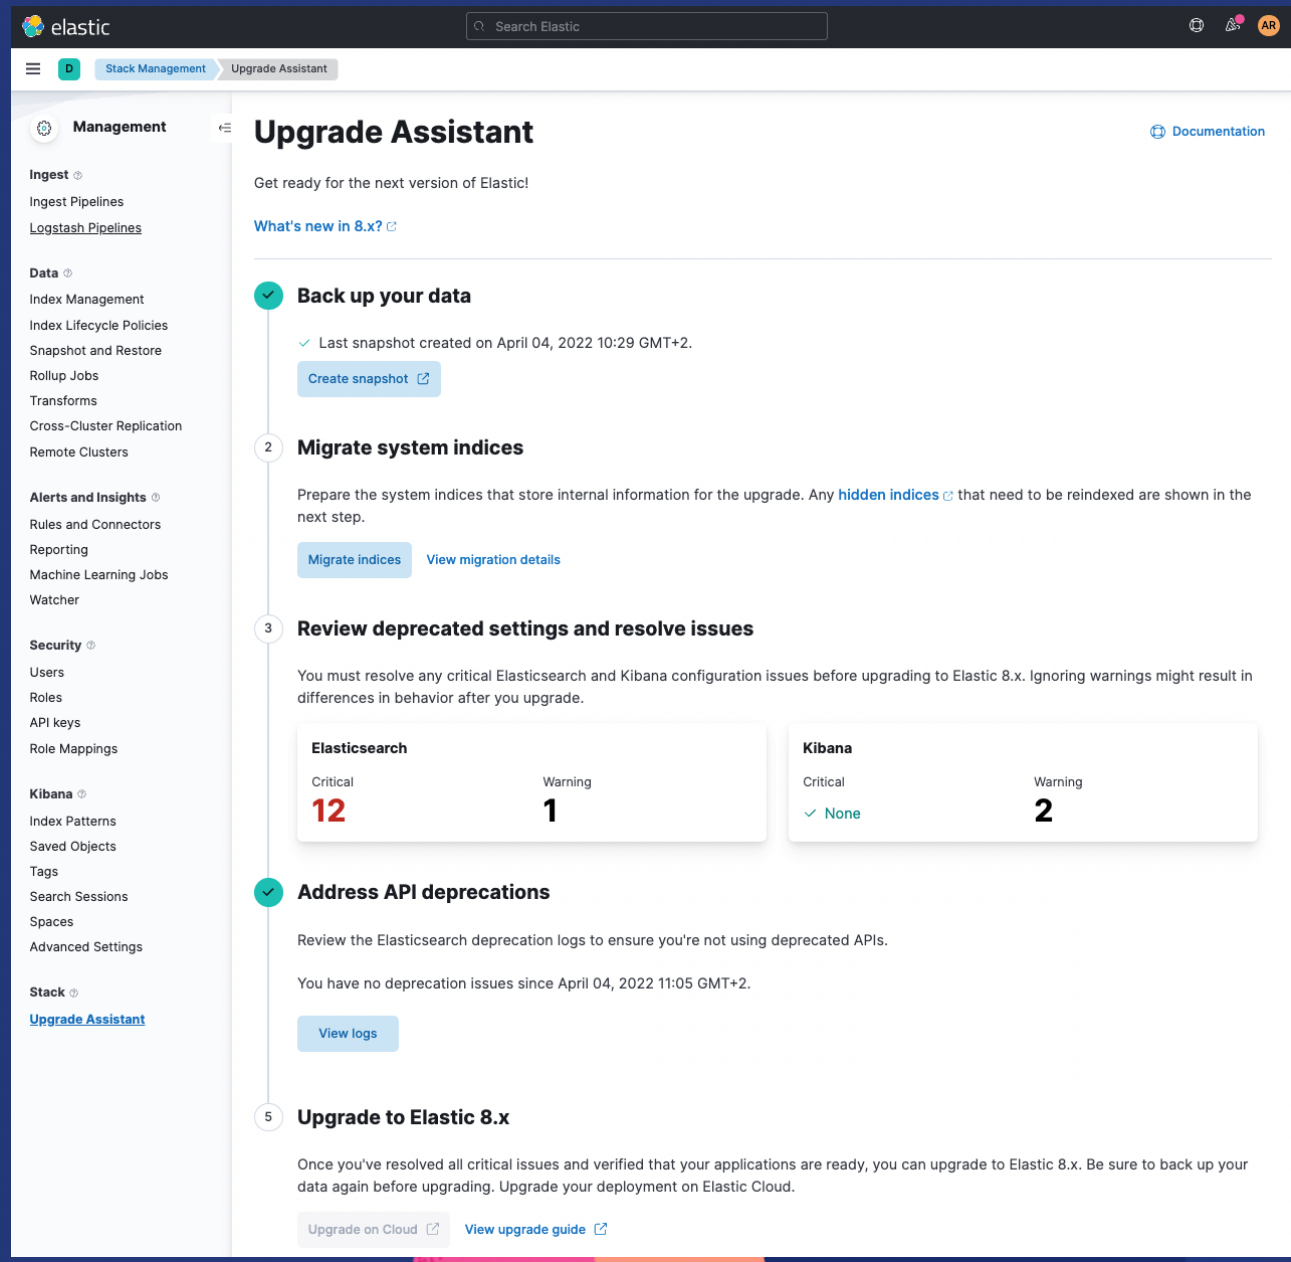

elastic

Upgrade Assistant

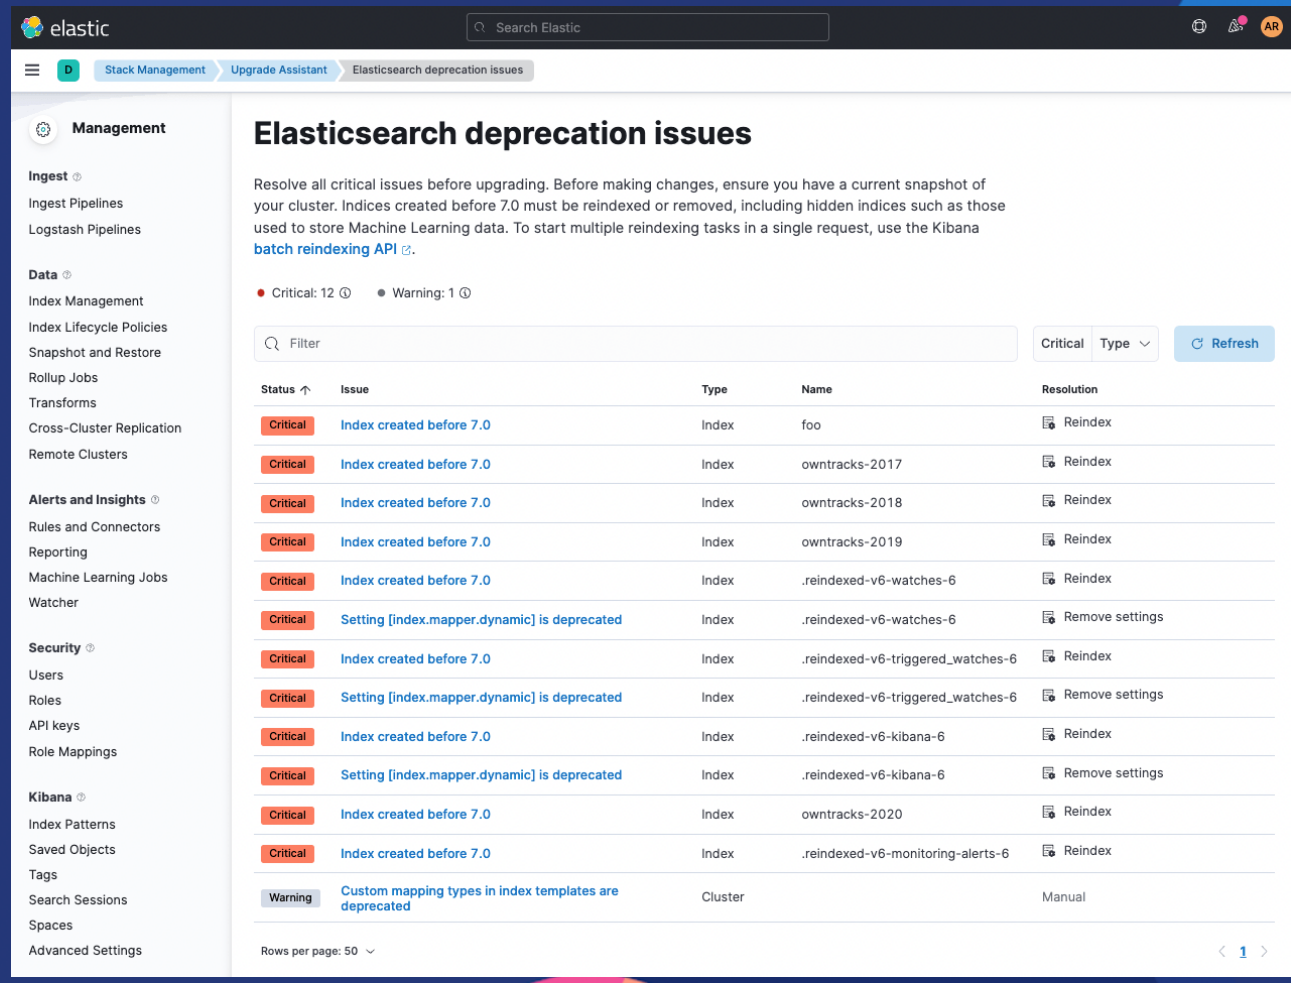

elastic

Elasticsearch Java Client

- Created from spec like all other clients
- Also exists for 7.x allowing for smooth migration.

elastic

Terraform Provider

- Elastic Cloud Provider
- Elasticstack Provider

elastic

Terraform Provider

```
resource "ec_deployment" "spring-boot-app-search" {
  name                  = "spring-boot-app-search"

  region                = "azure-westeurope"
  version               = "8.1.0"
  deployment_template_id = "azure-memory-optimized"

  elasticsearch {}

  kibana {}

  integrations_server {}

  enterprise_search {}
}
```

elastic

What's next

- ANN search filtering
- Lookup runtime fields
- doc-values only fields (for more efficient searchable snapshots)
- random sampler aggregation

elastic

Thanks for listening!

Q & A

Philipp Krenn | @xeraa

Alexander Reelsen | @spinscale

elastic