elastic

Workshop Elasticsearch @Ecole 42

David Pilato
Elastic, @dadoonet

Elastic **Stack**

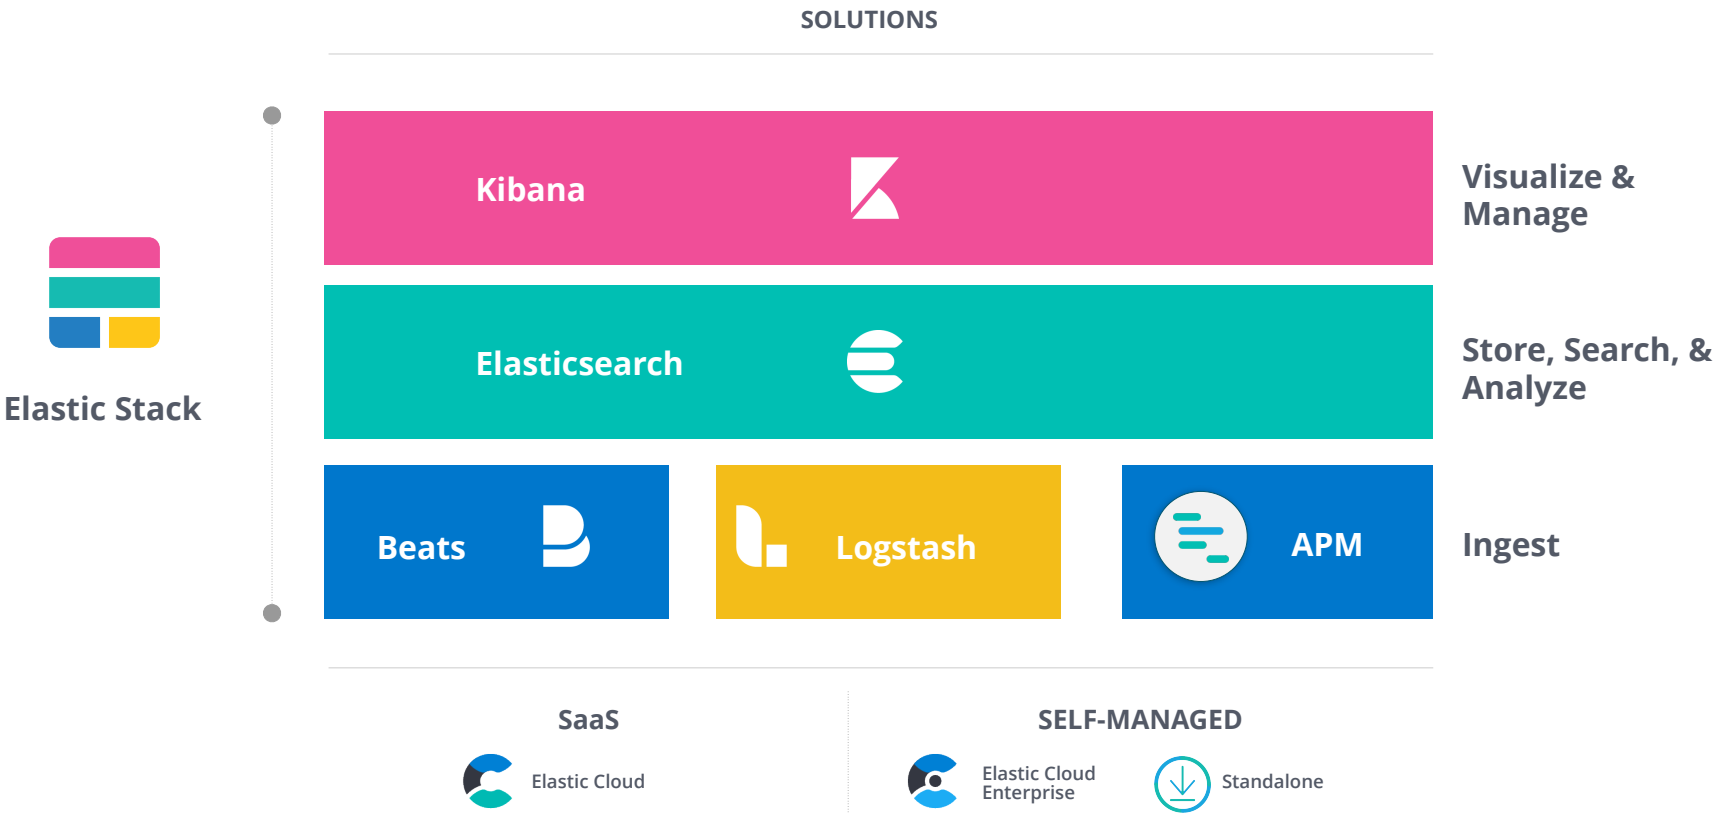

Deployment in the Enterprise

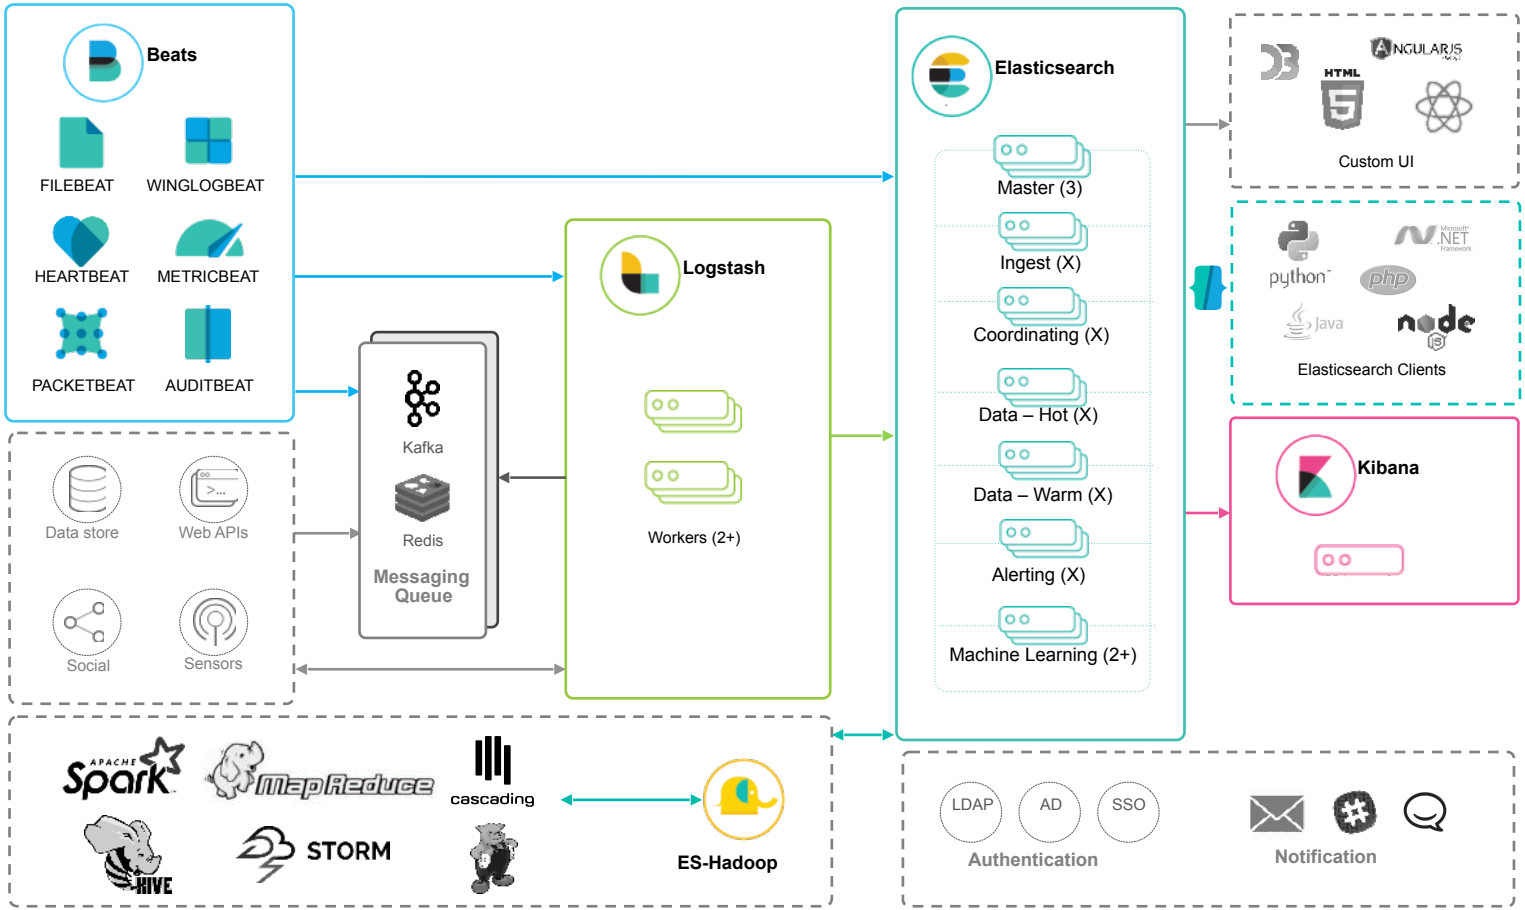

elastic

Open source

Apache 2.0 : aujourd'hui comme demain.

Téléchargement gratuit

Entre autres fonctionnalités :

- Clustering et haute disponibilité
- Recherche et analyse ultra-performantes
- Visualisation des données et tableaux de bord
- Et plus encore

Basic

L'offre gratuite qui le restera toujours.

Téléchargement gratuit

Tous les avantages de l'open source, plus :

- Principales fonctionnalités de sécurité
- Solutions telles qu'APM, SIEM, Maps et d'autres encore
- Canvas
- Et plus encore

Gold

Plus de fonctionnalités. Support technique dédié.

Nous contacter

Tous les avantages de l'offre Basic, plus :

- Alerting
- Reporting
- Gestion de l'ingestion
- Support technique aux heures ouvrées
- Et plus encore

Platinum

L'expérience ultra-complète.

Nous contacter

Tous les avantages de l'offre Gold, plus :

- Fonctionnalités de sécurité avancées
- Machine Learning
- Réplication inter-clusters
- Support technique 24 h/24, 7 j/7, 365 j par an
- Et plus encore

https://www.elastic.co/fr/subscriptions

elastic

Agenda

- Elasticsearch overview
- Workshop 0: getting started
- Workshop 1: let’s index some documents
- Workshop 2: let’s search them
- Workshop 3: let’s pull some analytics
- Workshop 4: let’s add a powerful live UI on top


elasticsearch

think document!

- Change your mindset:
  - Forget SQL!
  - Index what you want to find
- A document
  - A JSON object
  - Core field types (string, numbers, booleans, dates)
  - Complex field types (arrays, objects)
  - Additional field types (dates, geo points, geo shapes)

A document

```
{
  "name" : "elastic",
  "website" : "http://www.elastic.co",
  "category" : "software",
  "founded_year" : 2012,
  "overview" : "The company behind elasticsearch",
  "tags" : ["search", "datastore", "analytics"],
  "location" : {
    "city" : "Amsterdam",
    "country_code" : "NL",
    "geo" : {
      "lat" : 52.370176,
      "lon" : 4.895008
    }
  }
}
```

workshop 0

setup

docker compose (elasticsearch)

```
---
version: '3'
services:

  elasticsearch:
    image: docker.elastic.co/elasticsearch/elasticsearch:$ELASTIC_VERSION
    environment:
      - bootstrap.memory_lock=true
      - discovery.type=single-node
      - cluster.routing.allocation.disk.threshold_enabled=false
    ulimits:
      memlock:
        soft: -1
        hard: -1
    ports:
      - 9200:9200
    networks: ['stack']
```

https://gist.github.com/dadoonet/f3c67bebf2cf604df02f78ac2cb2fbde

elastic

docker compose (kibana)

```
  kibana:
    image: docker.elastic.co/kibana/kibana:$ELASTIC_VERSION
    ports: ['5601:5601']
    networks: ['stack']
    links: ['elasticsearch']
    depends_on: ['elasticsearch']

networks:
  stack: {}
```

https://gist.github.com/dadoonet/f3c67bebf2cf604df02f78ac2cb2fbde

elastic

start

- run

```
docker-compose up
```

- open Kibana

```
open http://0.0.0.0:5601/
```

workshop 1

we index persons

workshop 1: index some documents

- Load demo-console.txt file in Kibana dev console

https://gist.github.com/dadoonet/f3c67bebf2cf604df02f78ac2cb2fbde

elastic

workshop 1: 500 000 persons

- use injector script

```
java -jar injector-7.0.jar --debug --nb 500000
```

- see effect with _cat API

```
GET _cat/indices/person?v
```

https://ela.st/injector

workshop 2

we search for persons

workshop 3

we compute persons

workshop 4

clic and play!