NoSQL approaches in GnpIS

Elasticsearch @ URGI Feedback

Phenomic Data Management & Sharing

- Data Discovery
  - Links to datasets, by metadata
  - Loose integration (keywords)
  - Data model
    - Simple
    - Easy to implement and feed
  - Ex:
    - GnpIS Google-like search
    - transPLANT search
    - WheatIS
    - IFB-Elixir
- Data set building, Data Mining
  - Combine data from different sources
  - Strong integration
  - Data model
    - Rich
    - Complex to model, implement and feed
  - Ex: GnpIS.Ephesis

Data discovery system

- GnpIS, transPLANT, WheatIS searches
- Google like, full text, filters
- http://www.transplantdb.eu

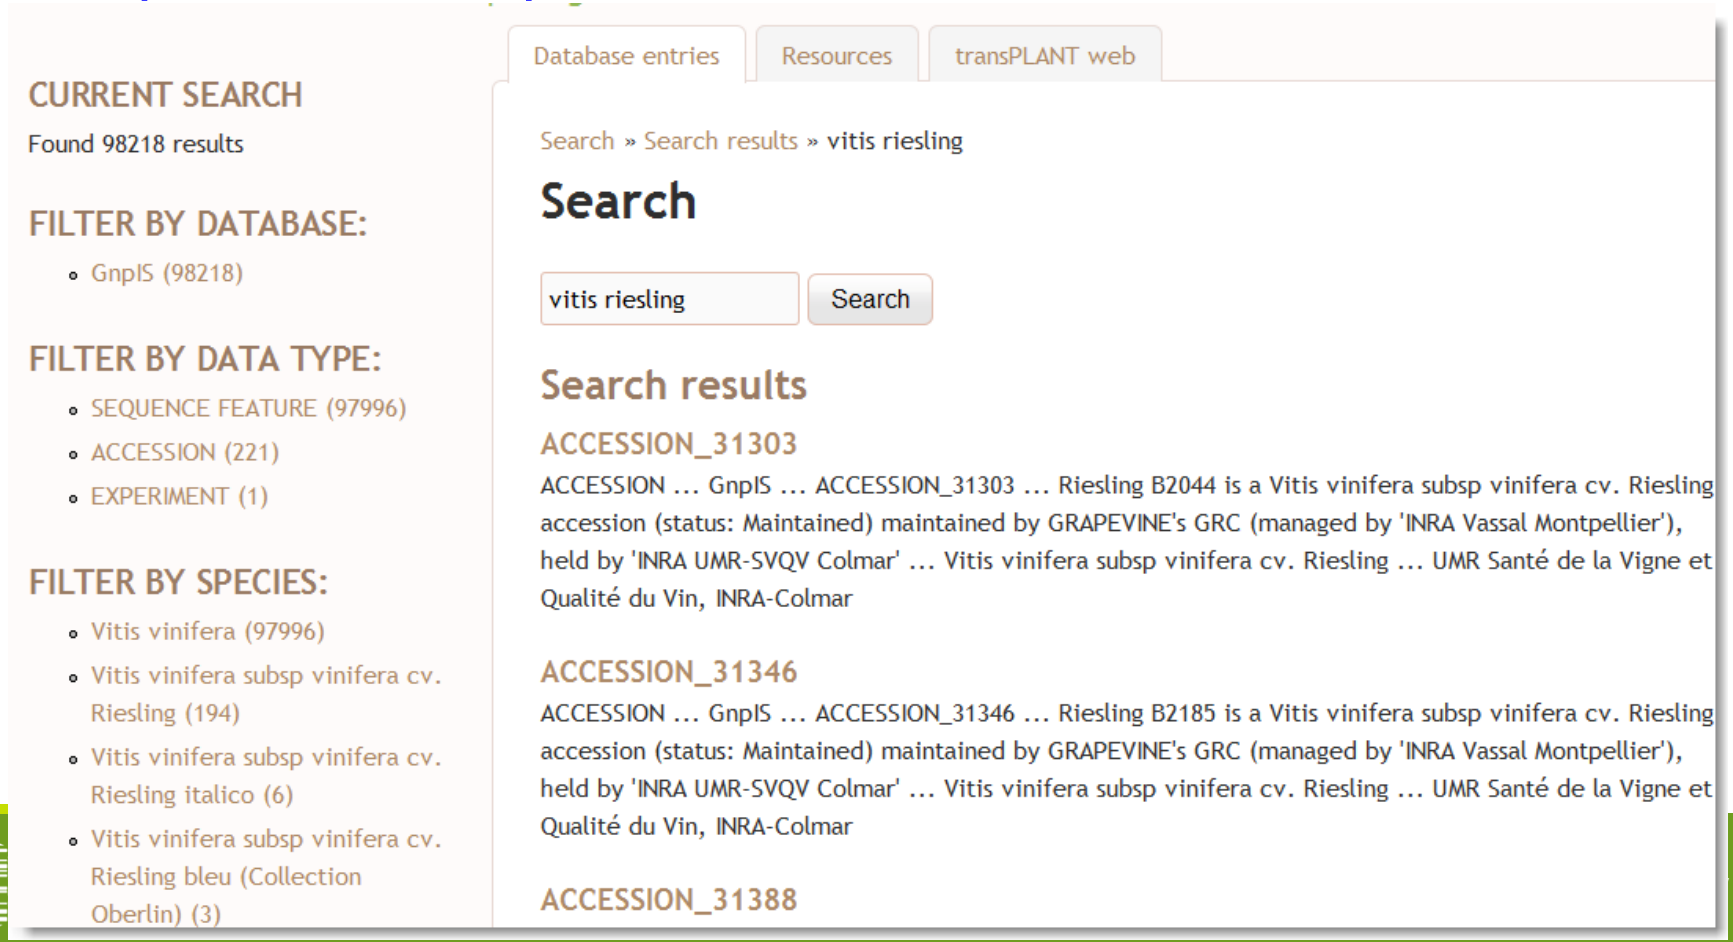

Data Set Building, Data Mining

- Phenotyping data
- Conceptual model

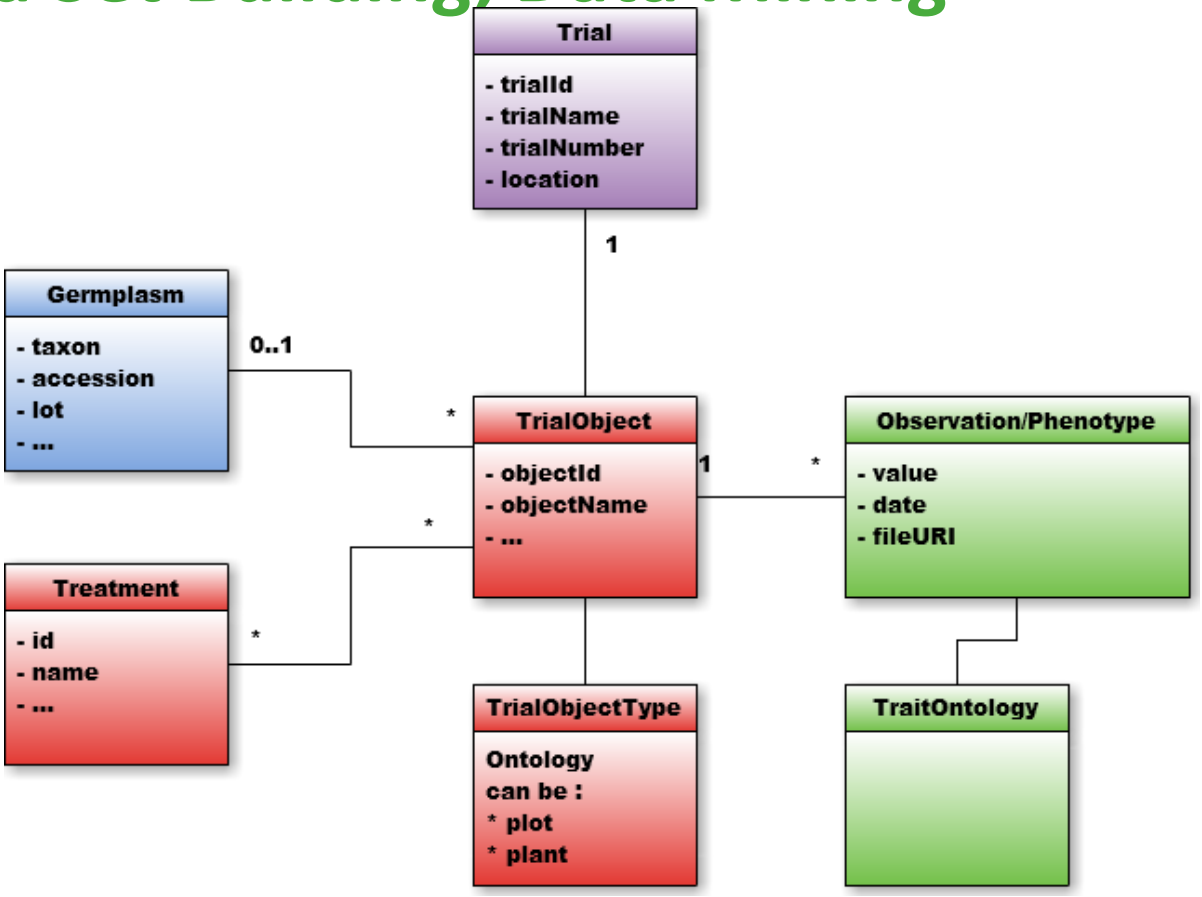

GnpIS.Ephesis Implementation

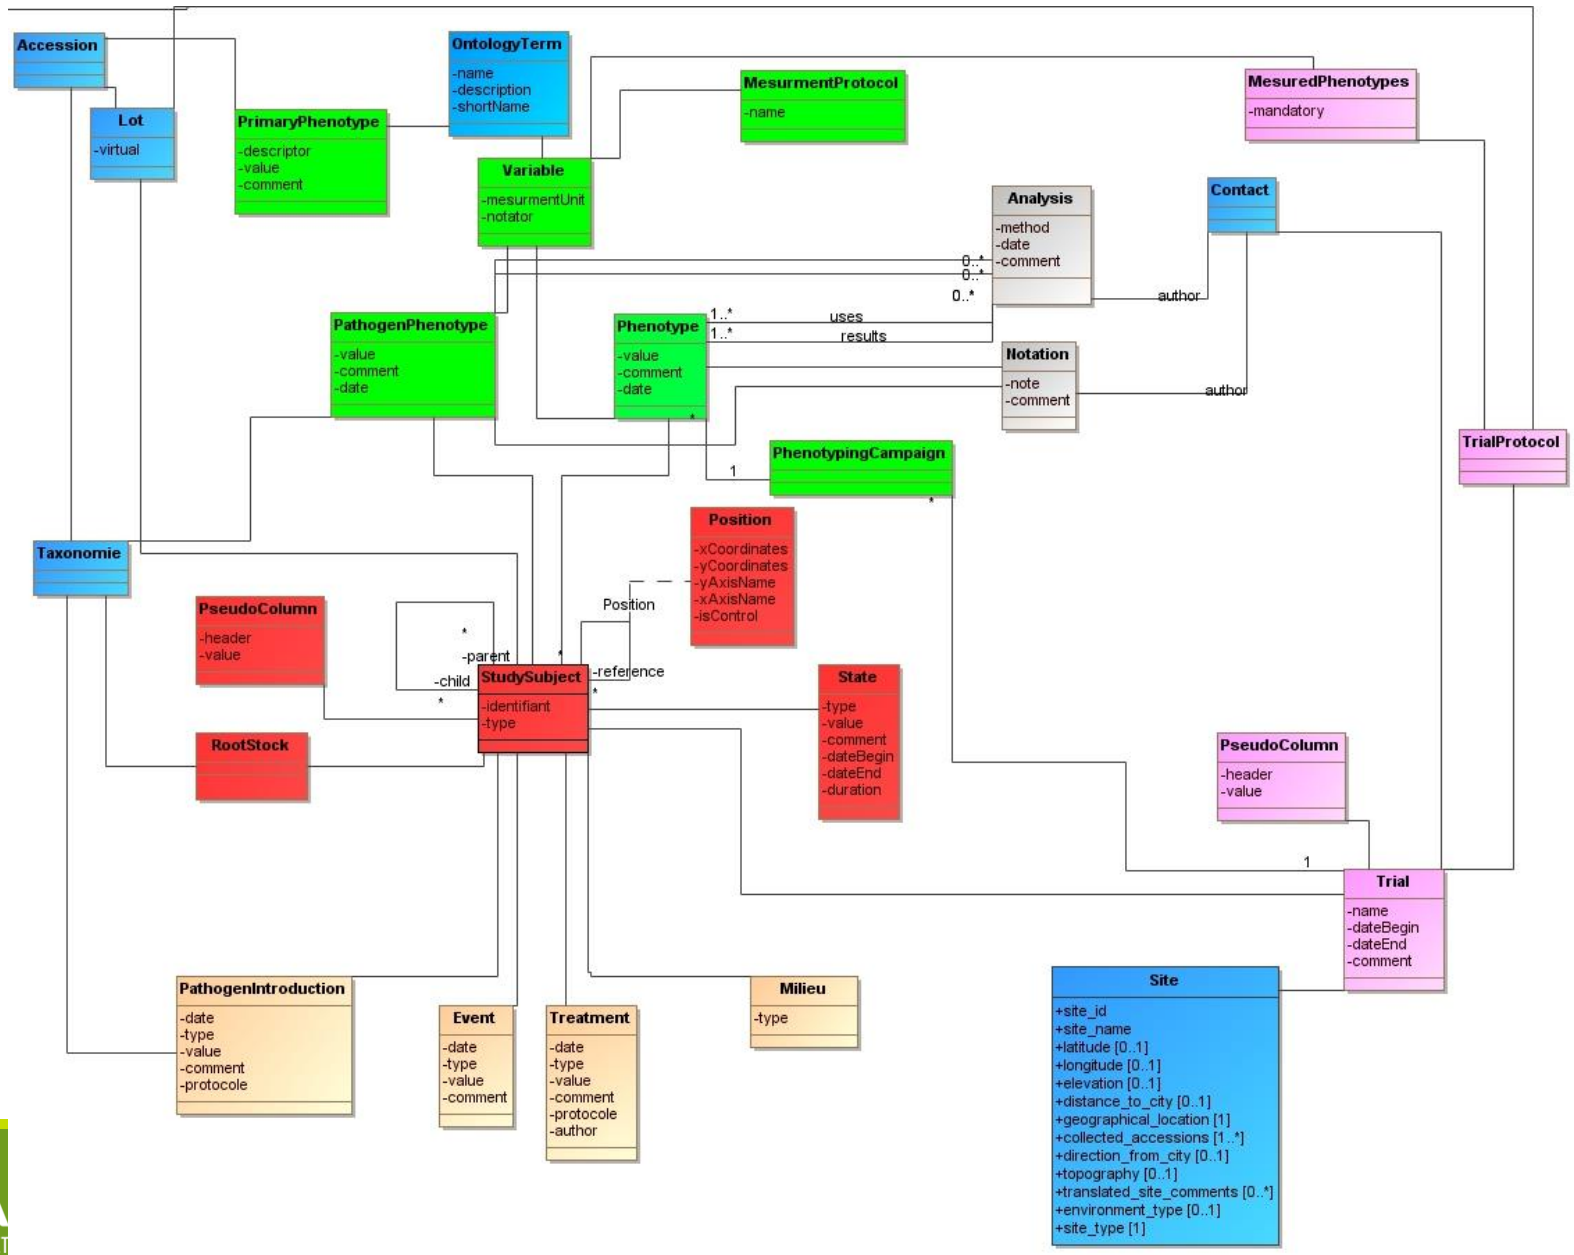

GnpIS.Ephesis V1 Performances

- StudySubject **join** Phenotype
  - **Too slow, too big**
- Solution
  - **Denormalisation, agregation**

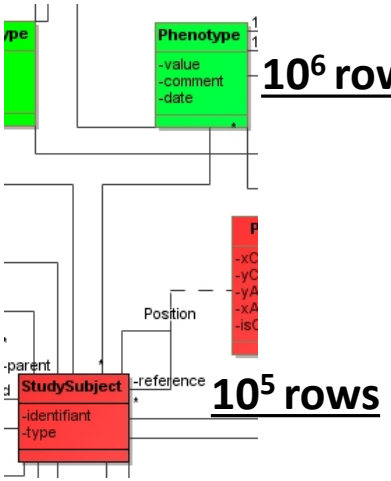

10^6 rows

10^5 rows

1-10 of 37,282 | Display 10 results per page

Lot Number	itk	Trial Name	Trial Site	Campaign	Rep	yield (rdt)
Alberic	t: treated	BTH_Clermont-Ferrand_2000_SetB1	Clermont-Ferrand	2000	1	NA
Alberic	t: treated	BTH_Clermont-Ferrand_2000_SetB1	Clermont-Ferrand	2000	2	93,4
DI00004	t: treated	BTH_Clermont-Ferrand_2000_SetB1	Clermont-Ferrand	2000	1	81,8
DI00004	t: treated	BTH_Clermont-Ferrand_2000_SetB1	Clermont-Ferrand	2000	2	95,1
DI00005	t: treated	BTH_Clermont-Ferrand_2000_SetB1	Clermont-Ferrand	2000	1	87,6
DI00005	t: treated	BTH_Clermont-Ferrand_2000_SetB1	Clermont-Ferrand	2000	2	89,7
EM00001	t: treated	BTH_Clermont-Ferrand_2000_SetB1	Clermont-Ferrand	2000	1	77,8
EM00001	t: treated	BTH_Clermont-Ferrand_2000_SetB1	Clermont-Ferrand	2000	2	91,8
EM00003	t: treated	BTH_Clermont-Ferrand_2000_SetB1	Clermont-Ferrand	2000	1	91,1
EM00003	t: treated	BTH_Clermont-Ferrand_2000_SetB1	Clermont-Ferrand	2000	2	85,4

GnpIS.Ephesis V2 Performances

- Study Subject
  - **No join**
- Phenotype: select N+1
- Ok now
- Not for the Future

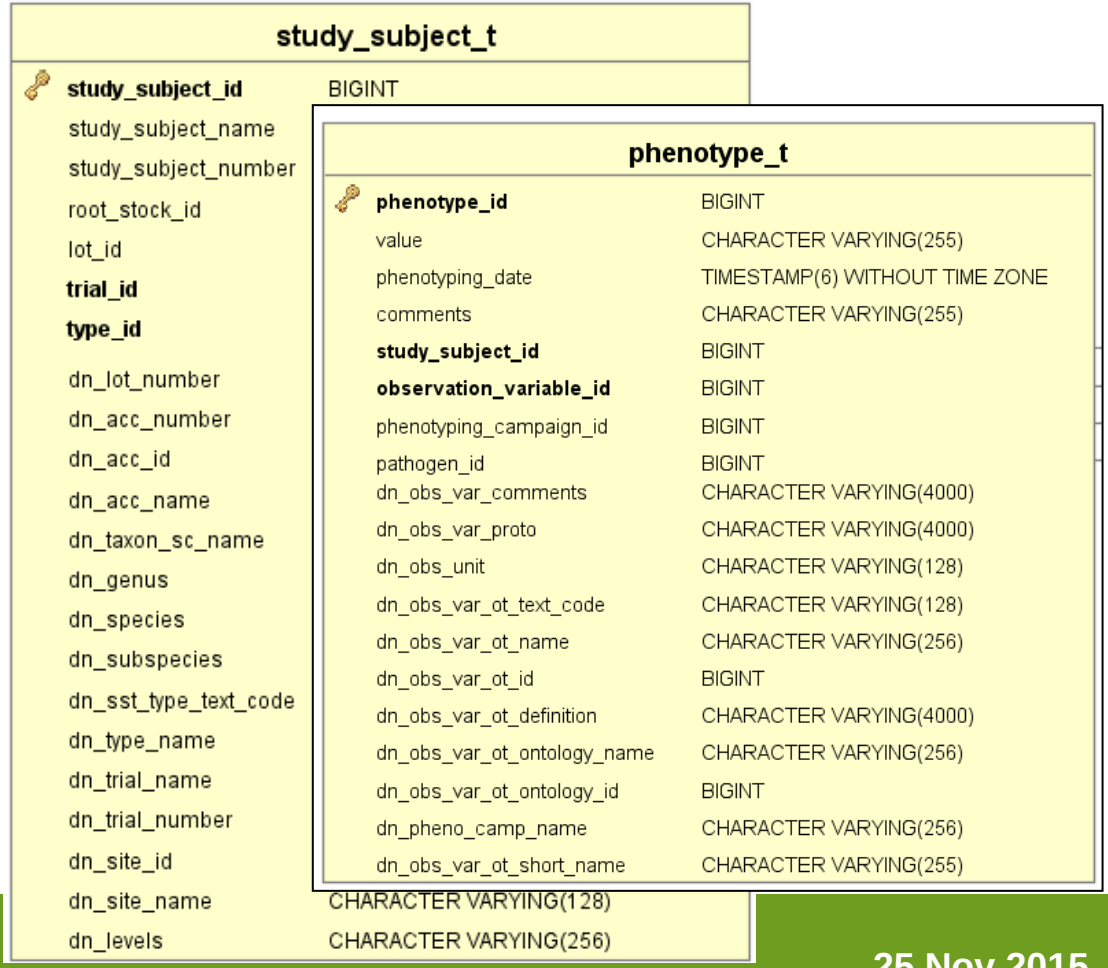

NoSQL Modelisation

- http://blog.palo-it.com/2013/06/10/modelisation-dun-schema-dune-base-de-donnees-nosql/
- NoSQL Document
- From Database to documents
  - **Denormalisation and agregation**
  - **3 methods**

Agregations

- Document interweaving

```
{
    'id': 10,
    'nom': 'dupont',
    'prenom': 'david',
    'email':'me@palo-it.com',
    'adresse':
        {
        'pseudo':'10 rue du test',
        'ville':'paris',
        'pays':'France',
        'code postal':75009
        }
}
```

- Field duplication
  - **Interweaved document as attributes of root document**
- Correlated documents (ie foreign Key)

```
Auteur
{
    "id": 10,
    "nom": "dupont",
    "prenom": "david",
    "livres": [
        101,503,339,342
    ]

}
Livre
{
    "id": 342,
    "titre": "NoSQL schema",
    "genre": "informatique",

    "tags":["informatique","bigdata","nosql"]
    "auteurs": [
        10,234
    ]

}
```

Application to phenotyping

- Breeding API
  - http://docs.brapi.apiary.io
- Study aka Trial


```
{
    "studyDbId": 123,
    "studyPUI": "http://phenome-fppn.fr/phenoarch/2014/1",
    "studyId" : "BRP-03",
    "studyName": "Blight Resistance in Phillipines",
    "studyObjective": "Test blight resistant cultivars",
    "studyType": "Trial",
    "studyLocation": "Phillipines",
    "studyProject": "Inovine",
    "dataSet": ["National Network", "Frost suceptibility network"],
    "studyPlatform": "Phenome",
    "startDate": "2015-06-01",
    "endDate": "2015-12-31",
    "programName": "RiceImprovementProgram",
    "designType": "RCBD",
    "keyContact": "Mr.PlantBreederA",
    "contacts":
    [{
        "type": "scientific coordinator","name": "John Doe","email":
    }]
    "meteoStationCode": "Anlez",
    "meteoStationNetwork": "OpenWheatherMap",
    "studyHistory": "Previous crop was pea, then maize",
    "studyComments",
    "seasons": ["2005", "2008"],
    "observationVariables": [
        {
            "observationVariableId": "CO_321:0000045",
            "observationVariableComment": "There might be a mistake
        },
        {
            "observationVariableId": "http://www.cropontology.org/r
            "observationVariableComment": ""
    }],
    "germplasms":[
        {
            "germplasmDbId": "01BEL084609",
            "germplasmPUI": "http://www.crop-diversity.org/mgis/acce
            "germplasmName": "Pahang"
    }]
}
```

Application to phenotyping

- Breeding API
- Phenotypes
- One document = one data matrix line
- Elasticsearch nested documents

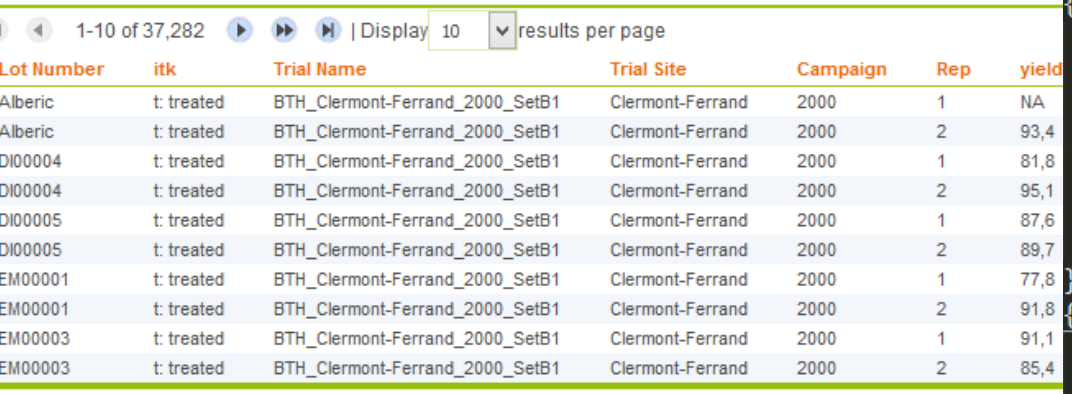

1-10 of 37,282 | Display 10 results per page

Lot Number	itk	Trial Name	Trial Site	Campaign	Rep	yield
Alberic	t: treated	BTH_Clermont-Ferrand_2000_SetB1	Clermont-Ferrand	2000	1	NA
Alberic	t: treated	BTH_Clermont-Ferrand_2000_SetB1	Clermont-Ferrand	2000	2	93,4
DI00004	t: treated	BTH_Clermont-Ferrand_2000_SetB1	Clermont-Ferrand	2000	1	81,8
DI00004	t: treated	BTH_Clermont-Ferrand_2000_SetB1	Clermont-Ferrand	2000	2	95,1
DI00005	t: treated	BTH_Clermont-Ferrand_2000_SetB1	Clermont-Ferrand	2000	1	87,6
DI00005	t: treated	BTH_Clermont-Ferrand_2000_SetB1	Clermont-Ferrand	2000	2	89,7
EM00001	t: treated	BTH_Clermont-Ferrand_2000_SetB1	Clermont-Ferrand	2000	1	77,8
EM00001	t: treated	BTH_Clermont-Ferrand_2000_SetB1	Clermont-Ferrand	2000	2	91,8
EM00003	t: treated	BTH_Clermont-Ferrand_2000_SetB1	Clermont-Ferrand	2000	1	91,1
EM00003	t: treated	BTH_Clermont-Ferrand_2000_SetB1	Clermont-Ferrand	2000	2	85,4


```
{
        "observationUnitDbId": 20,
        "observationUnitPUI": "http://phenome-fppn.fr/maugio/bloc/12/23
        "studyId": "RIGW1",
        "studyLocation": "Bergheim",
        "studyPUI": "http://phenome-fppn.fr/phenoarch/2014/1",
        "studyProject": "Inovine",
        "studyPlatform": "Phenome",
        "germplasmPUI": "http://inra.fr/vassal/41207Col0001E",
        "germplasmDbId": 3425,
        "germplasmName": "charger",
        "treatments":
        [
            {
                "factor" : "water regimen" ,
                "modality":"water deficit"
            }
        ],
        "X" : "",
        "Y" : "",
        "data": [
{
            "observationVariableId": "CO_321:0000045",
            "season": "2005",
            "observationValue" : "red",
            "observationTimeStamp": null,
            "quality": "reliability of the observation",
            "collectionFacilityLabel":  "phenodyne",
            "collector" : "John Doe and team"
},
{
            "observationVariableId": "http://www.cropontology.o
            "season": null,
            "observationValue" :  32,
            "observationTimeStamp": "2006-07-03::10:00",
            "quality": "8",
            "collectionFacilityLabel": null,
            "collector" : "userURIOrName"
```

GnpIS,Ephesis V3 performances

- 700 000 observation units
- 3 * 10^6 phenotypes / observations
- Elasticsearch implementation
- Response time
  - **https://urgi.versailles.inra.fr/ephesis/ephesis/viewer.do#dataResults/trialSetIds=6,5,7**
  - **770 Trials, 49 658 documents subset, 6 variables each**
  - **One page, Ajax call**
    - **300 – 600 ms**
  - **Full export**
    - **23 s**
- Genotyping :

Elasticsearch

- David Pilato
  - **http://david.pilato.fr/**
  - **Evangelist at elastic and creator of the Elastic French Speakers User Group.**
- http://fr.slideshare.net/*
- Elasticsearch
  - **https://www.elastic.co/**
  - **Créé en 2010**
  - **International**
  - **Applications et adoption croissante**

Elasticsearch

- Lucene based
- Search engine
  - **Not presented as a NoSQL database**
    - **NoSQL document search engine**
    - **BUT : it has backup systems**
  - **Not for long term storage**
    - **Aggregated documents, query oriented**
  - **Could be used as NoSQL DB ?**
    - **Storage**
- HTTP, REST, JSON
- Distributed, Scalable, Cluster and Cloud ready

ES vs SQL

Cherche moi un document
de **décembre 2011** portant sur la **france**
et contenant **produit** et **david**

En SQL :

```
SELECT
  doc.*, pays.*
FROM
  doc, pays
WHERE
  doc.pays_code = pays.code AND
  doc.date_doc > to_date('2011-12', 'yyyy-mm') AND
  doc.date_doc < to_date('2012-01', 'yyyy-mm') AND
  lower(pays.libelle) = 'france' AND
  lower(doc.commentaire) LIKE '%produit%' AND
  lower(doc.commentaire) LIKE '%david%';
```

http://fr.slideshare.net/dadoonet/elasticsearch-esme-sudria

Elasticsearch

```
curl -XGET 'http://localhost:9200/docindex/doc/_search' -d '{
    "query": {
        "bool" : {
            "filter" : {
                "term" : { "pays.libelle" : "france" },
                "match" : {
                    "doc.commentaire" : {
                        "query" : "produit david",
                        "operator" : "and"
                    },
                    "range" : {
                        "doc.date_doc" : {
                            "gte" : "2011-12",
                            "lt" :  "2012-01"
                        }
                    }
                }
            }
        }
    }
}'
```

Elasticsearch administration concepts

- URL HTTP : ES_instance/index/type/documents
- node :
  - Running instance of elasticsearch, belongs to a cluster.
- cluster :
  - one or more nodes with same cluster name
  - Single master node
    - chosen automatically
    - replaced if fails.
- shard
  - single Lucene instance
  - low-level “worker”
  - managed automatically by elasticsearch.
  - primary and replica shards.

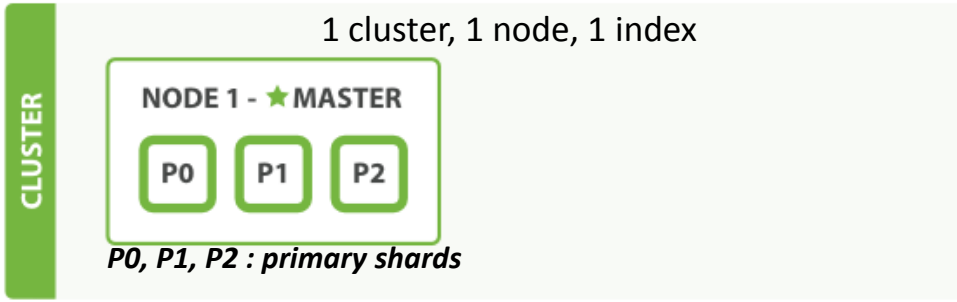

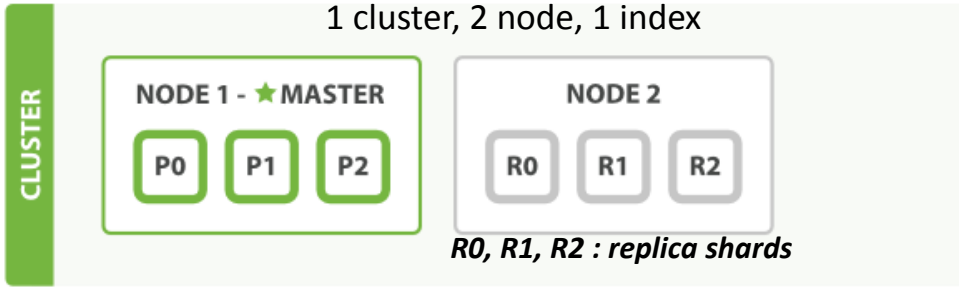

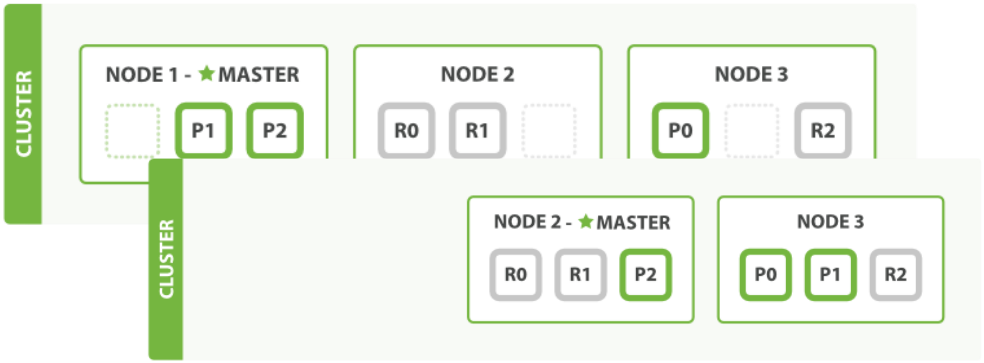

Elasticsearch developer concepts

- type
  - like a *table* in a relational database.
- index
  - like a *database* in a relational database.
  - has a **mappings** which defines multiple **types**.
  - types
    - Study
      - http://localhost:9200/phenoindex/studytype/_search?
    - Phenotypes
      - http://localhost:9200/phenoindex/phenotype/_search?
- Alias
  - like a SDGB view on multiple indices (possibly filtered via a query)

Development

- Very good query DSL/API documentation
  - https://www.elastic.co/guide/en/elasticsearch/reference/current/search.html
- Sense
  - Free, integrated into marvel (https://www.elastic.co/products/marvel)
  - https://www.elastic.co/guide/en/sense/current/sense-ui.html
  - query execution
  - autocompletion

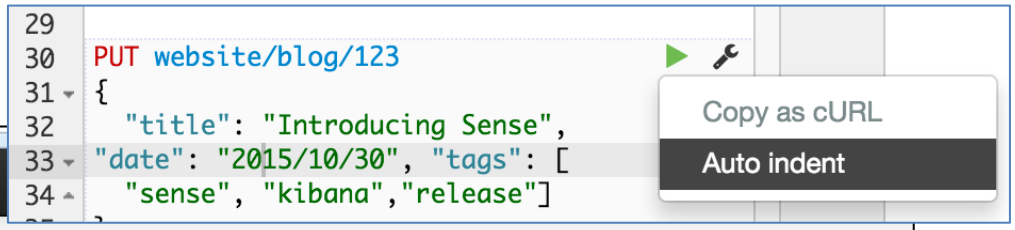

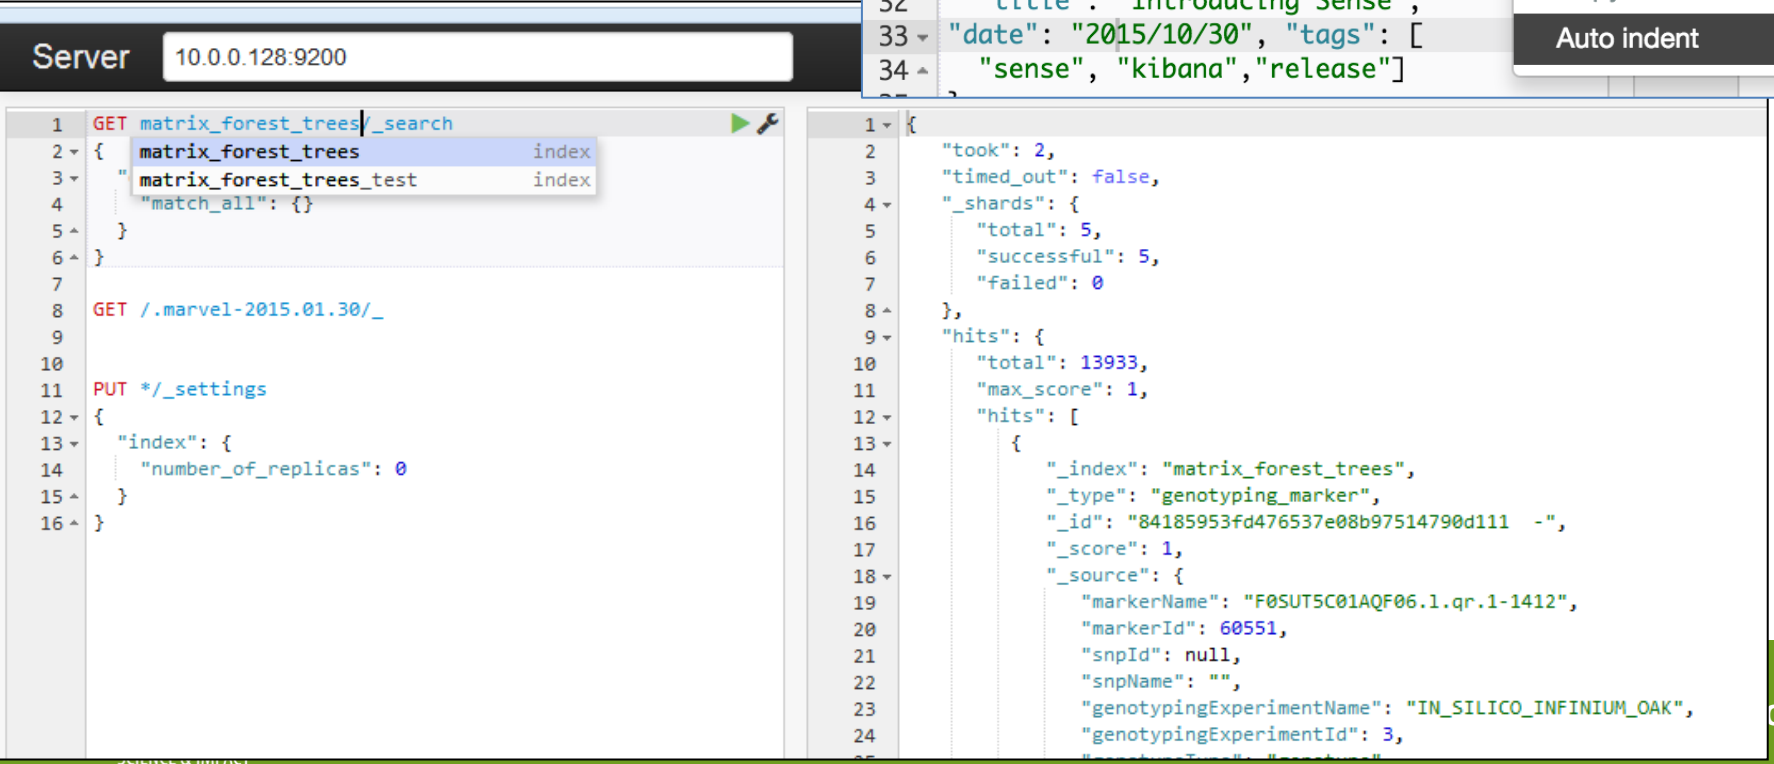

Agregations, ie Facets

- Analytic queries
- Two main concepts:
  - ***Buckets*** **Collections of documents that meet a criterion**
  - ***Metrics*** **Statistics calculated on the documents in a bucket**

```
SELECT COUNT(color) (1)
FROM table
GROUP BY color (2)
```

(1) `COUNT(color)` is equivalent to a metric.

(2) `GROUP BY color` is equivalent to a bucket.

Aggregation example

```
GET /cars/transactions/_search?search_type=count
{
    "aggs" : { (1)
        "colors" : { (2)
            "terms" : {
              "field" : "color" (3)
            }
        }
    }
}
```

```
{
...
   "hits": {
      "hits": [] (1)
   },
   "aggregations": {
      "colors": { (2)
         "buckets": [
            {
               "key": "red", (3)
               "doc_count": 4 (4)
            },
            {
               "key": "blue",
               "doc_count": 2
            },
            {
               "key": "green",
               "doc_count": 2
            }
         ]
      }
   }
}
```

1/ Display on result page as clickable facet
2/ on click on « red » filter color : red

Aggregation : Avg and imbrication

```
GET /cars/transactions/_search?search_type=count
{
   "aggs": {
      "colors": {
         "terms": {
            "field": "color"
         },
         "aggs": {
            "avg_price": { (1)
               "avg": {
                  "field": "price"
               }
            },
            "make": { (2)
                "terms": {
                    "field": "make" (3)
                }
            }
         }
      }
   }
}
```

```
{
...
   "aggregations": {
      "colors": {
         "buckets": [
            {
               "key": "red",
               "doc_count": 4,
               "make": { (1)
                  "buckets": [
                     {
                        "key": "honda", (2)
                        "doc_count": 3
                     },
                     {
                        "key": "bmw",
                        "doc_count": 1
                     }
                  ]
               },
               "avg_price": {
                  "value": 32500 (3)
               }
            },

...
}
```

Administration

- Monitoring : Marvel
  - **Gratuit dans ES 2.0**
- Installation
  - **Repository DEB and RPM**
  - **Unzip and run (development instance, tomcat like)**
- Configuration YAML

Elasticsearch architecture

- Scalability
  - Automatic index data distribution accross shards
  - Adding nodes to increase number of shards
- Shard Replication
  - Option
  - Index by index configuration
  - Data availability
  - Query performances

Elasticsearch@ URGI

- Version 1.7.3
- Data :
  - **65 Gb**
  - **600 millions documents, including nested**
- Cluster Prod (same for dev)
  - **2 nodes: VM, 16 Go RAM, 2To, 8 CPU**
- Each nodes
  - **Unlimited number of indices**
- Backup
  - **Weekly, retention 8 weeks**
  - **On the fly, without service interruption, incremental**
  - **Via a REST query (API HTTP)**
- Good data security
  - **Protection against index corruption**
  - **Near NoSQL DB state**

Insertion

- Generate JSON
  - Talend
  - Java
- Insert
  - Logstash
  - CURL

Thank you