




Elasticsearch Query Language

ES|QL

David Pilato - @dadoonet
Developer | Evangelist

slides & demo





Elastic and Kibana support a number of query languages

A brief history of Elasticsearch's analytical capabilities





ES|QL

- Language
- Engine
- Visualization



ES|QL

the language

ES|QL Features

- Unstructured and structured data
- Piped query language
- SQL-like filtering and data manipulation
- Lookups

ES|QL commands

Source (From, Row)

Filter (Where)

Processing (Eval)

Aggregation (Stats)

TopN (Sort + Limit)

Expansion (Enrich , MV_Expand)

Extraction (Dissect, Grok)

75+ functions:

- 10 aggregate
- 20+ math
- 10+ string
- 7 date-time
- 15 conversion
- 4 conditionals
- 12 multi-value / mv_



ES|QL

the engine




66

The new ES|QL execution engine was designed **with performance in mind** — it **operates on blocks** at a time instead of per row, **targets vectorization** and cache locality, and embraces specialization and **multi-threading**. It is a separate component from the existing Elasticsearch aggregation framework with different performance characteristics.

Query planner

- ✓ Flexible distributed execution
- ✓ Allow multiple roundtrips



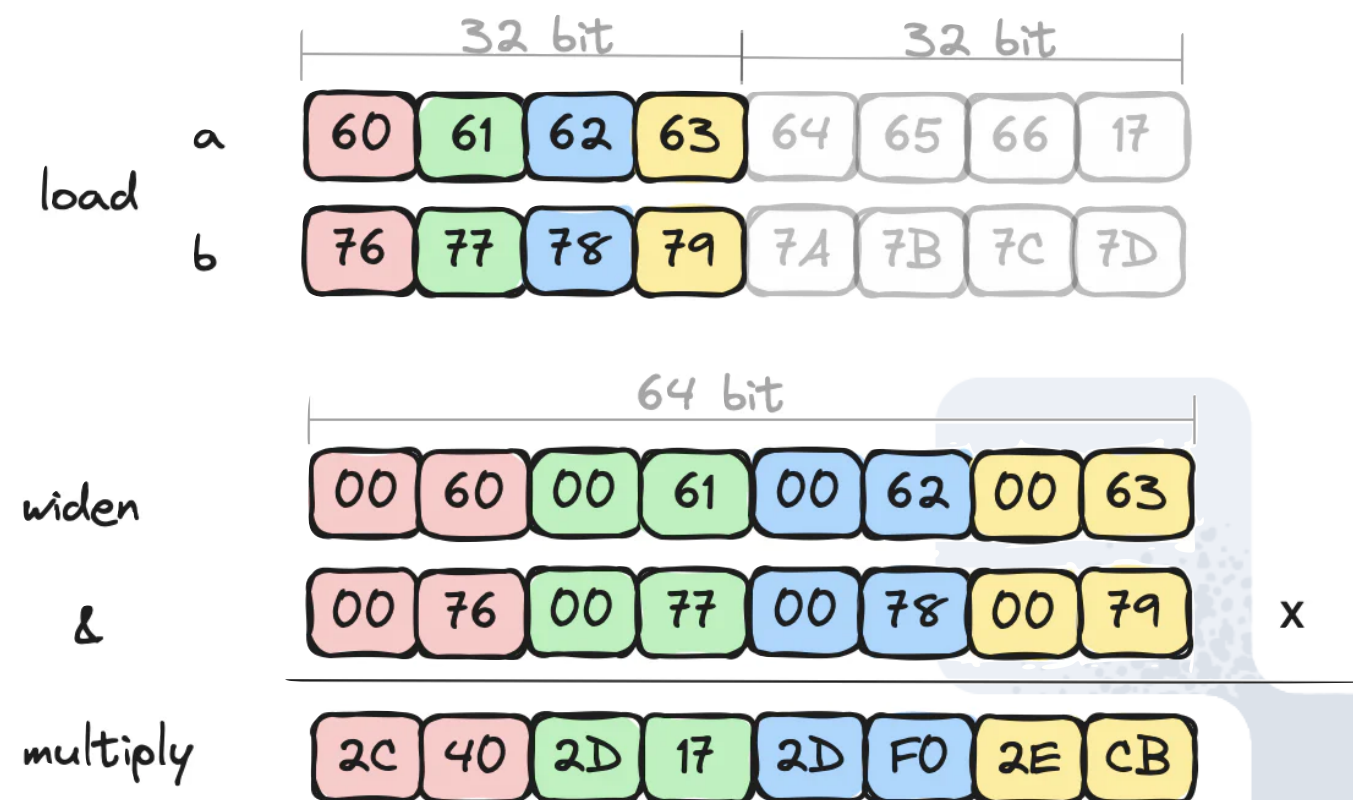
Compute engine

- ✓ Tabular data representation
- ✓ From 1 thread per shard to many
- ✓ Spilling to disk if needed
- ✓ Streaming of data across nodes

Vectorization


“convert from a scalar implementation, which processes a single pair of operands at a time, to a vector implementation, which processes one operation on multiple pairs of operands at once.”

```
for (i = 0; i < n; i++)  
    c[i] = a[i] + b[i];
```



Benchmarks

<https://elasticsearch-benchmarks.elastic.co/#tracks/esql/nightly/default/30d>



ES|QL

in action

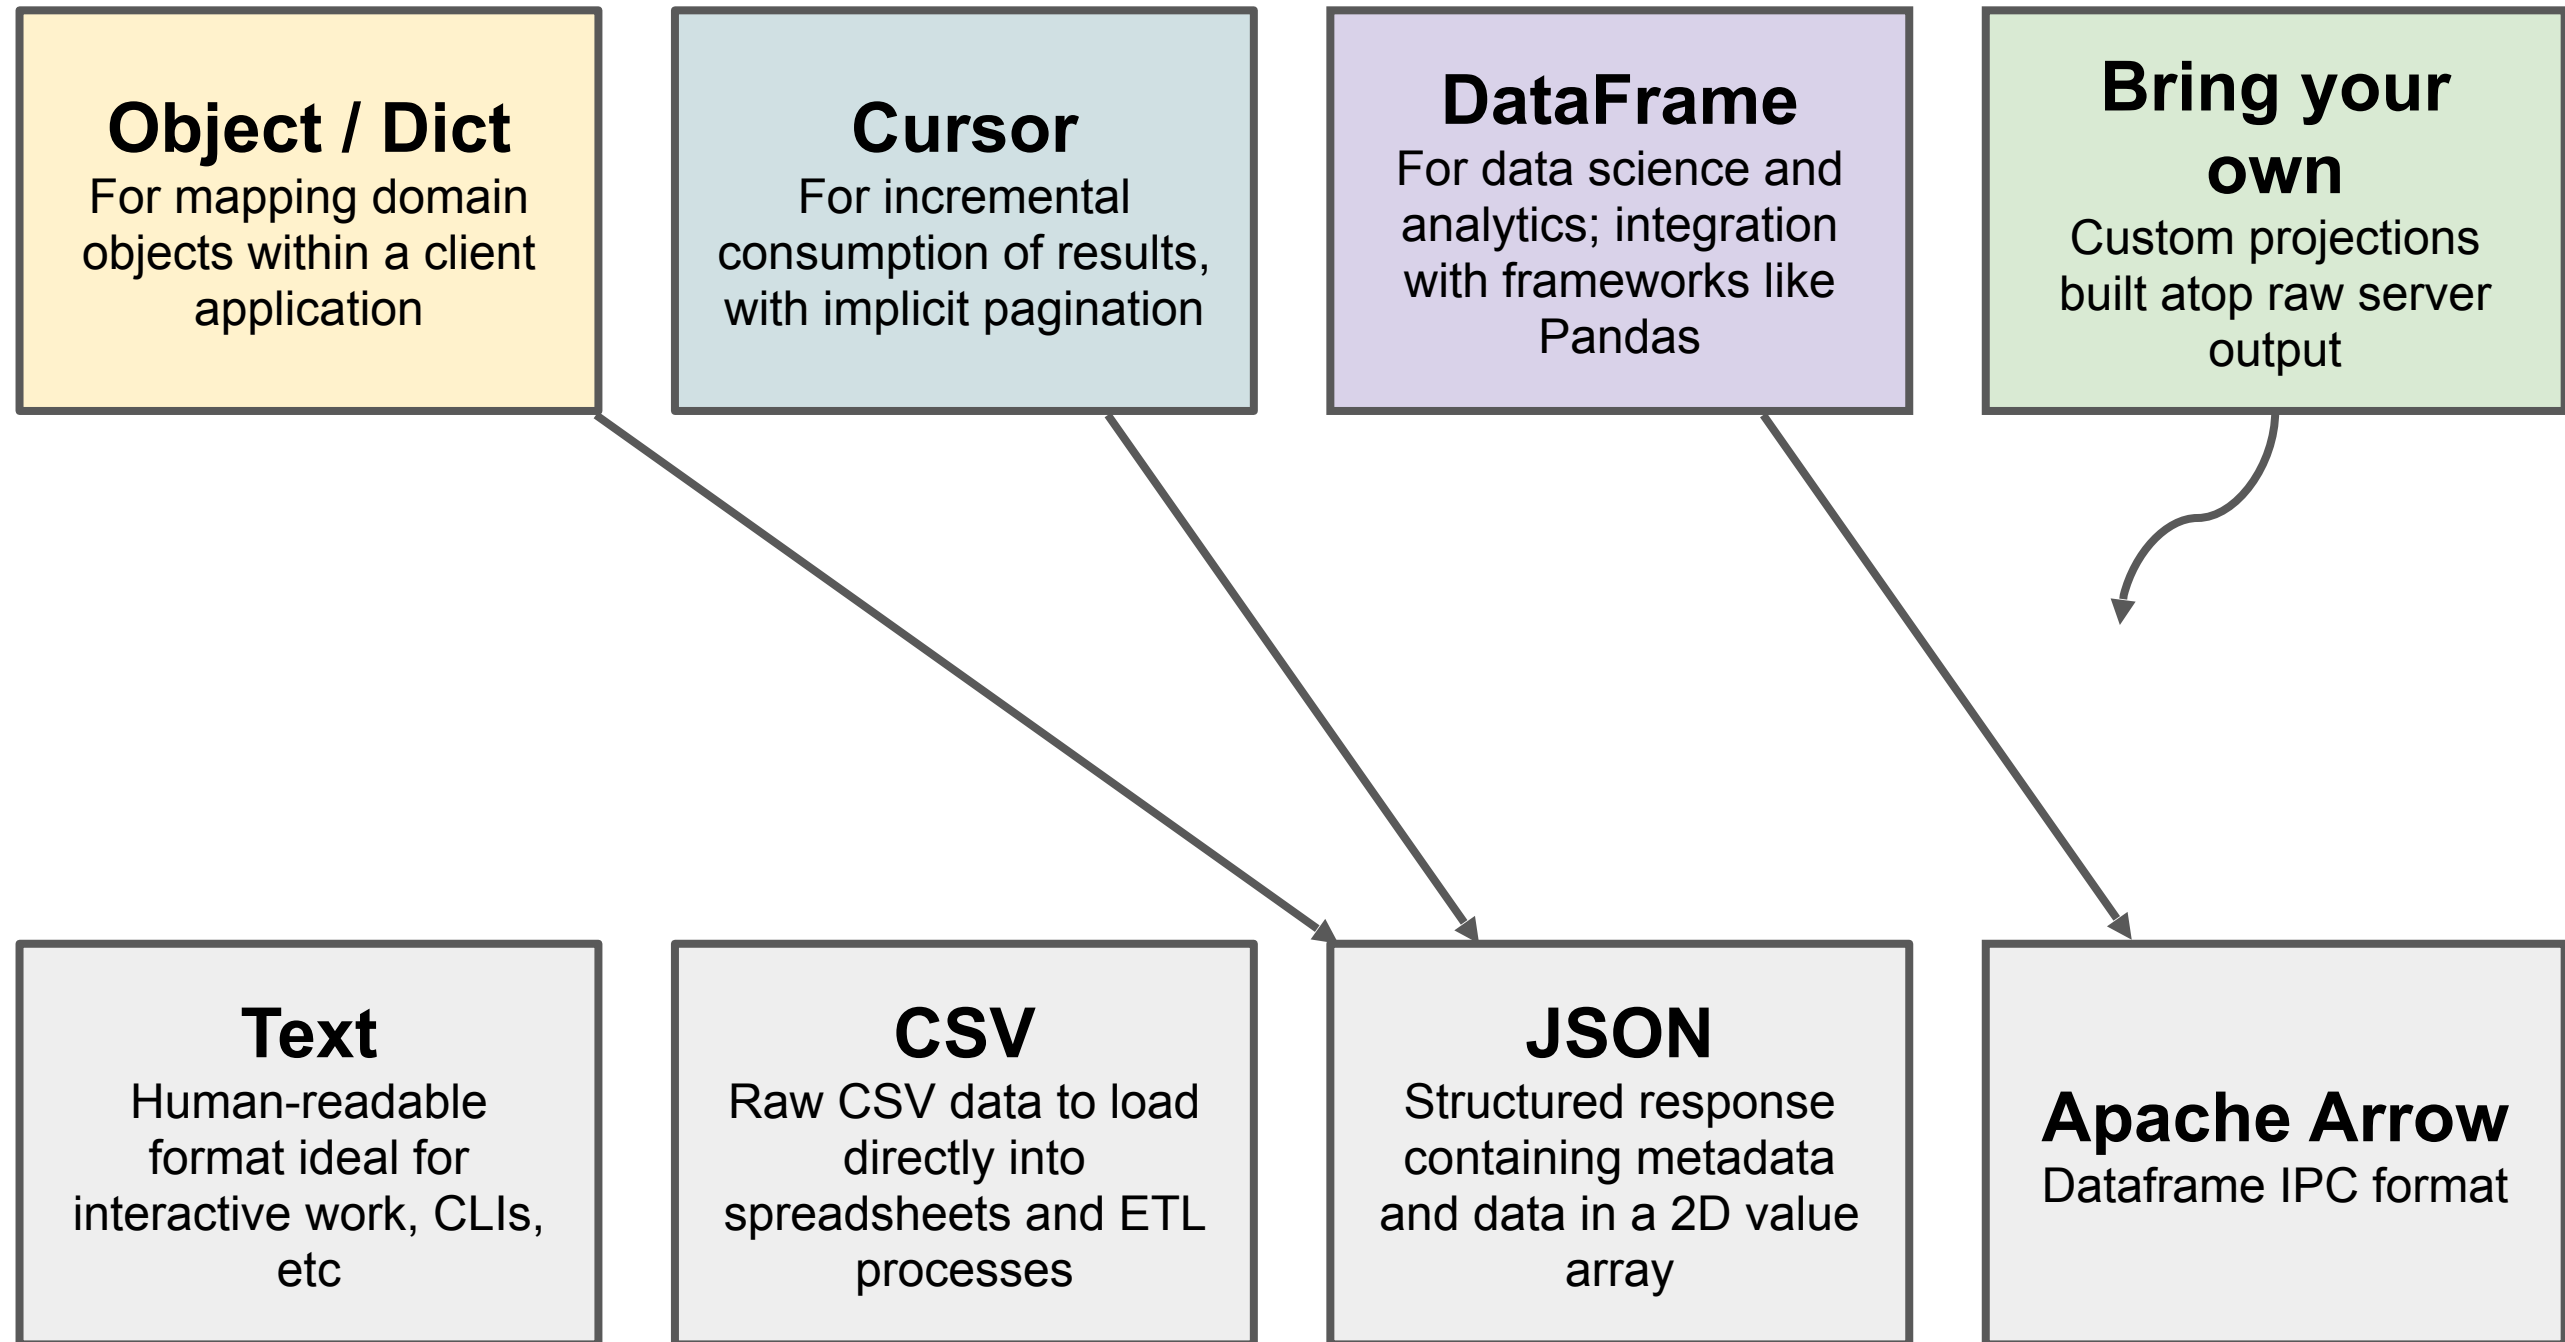
<https://github.com/dadoonet/esql-demo>

slides & demo



Ways to consume ES | QL results

Each language client will offer a selection of projections relevant to that language ecosystem.



Users can consume raw data directly from the server output in one of several formats.

Object API

<https://github.com/dadoonet/elasticsearch-java-client-demo>

```
String query = """
    FROM persons
    | WHERE name == "David"
    | KEEP name
    | LIMIT 1
    """;

Iterable<Person> persons = client.esql()
    .query(ObjectsEsqlAdapter.of(Person.class), query);
for (Person person : persons) {
    assertNull(person.getId());
    assertNotNull(person.getName());
}
```

ResultSet JDBC API

<https://github.com/dadoonet/elasticsearch-java-client-demo>

```
String query = """
    FROM persons
    | WHERE name == "David"
    | KEEP name
    | LIMIT 1
    """;

try (ResultSet resultSet = client.esql()
    .query(ResultSetEsqAdapter.INSTANCE, query)) {
    assertTrue(resultSet.next());
    assertEquals("David", resultSet.getString(1));
}
```


Named parameters

```
POST /_query
{
  "query": ""
    from logs-*
    | stats x = ?function(?field) by ?breakdownField
    | where x >= ?value
  "",
  "params": [
    {"function": {"identifier": "avg"}},
    {"field": {"identifier": "network.bytes"}},
    {"breakdownField": {"identifier": "agent.name"}},
    {"value": 1000}
  ]
}
```

Java API with named parameters

<https://github.com/dadoonet/elasticsearch-java-client-demo>

```
String query = """
    FROM persons
    | WHERE name == ?name
    | KEEP name
    | LIMIT 1
    """;

Iterable<Person> persons = client.esql()
    .query(ObjectsEsqlAdapter.of(Person.class), query,
        Map.of("name", "David"));
for (Person person : persons) {
    assertNull(person.getId());
    assertNotNull(person.getName());
}
```

Coming next

```
WHERE KQL("bytes>=1024")
```

JOINS!

```
joinType JOIN indexName (AS qualifier)? condition?
```

```
joinType: LOOKUP | LEFT | RIGHT | INNER
```

```
condition:
```

```
  ON identifier == identifier
```

```
  | USING identifier
```

```
FROM person
```

```
| INLINESSTATS avg_children = AVG(children)
```

```
| WHERE children > avg_children
```

```
| LIMIT 1000
```

TBD

```
FROM search-movies METADATA _score, _id
| WHERE imdbrating > 7
| FORK [WHERE title:"Shakespeare" | SORT _score DESC | LIMIT 10] // fork1
      [WHERE semantic_title:"Shakespeare" | SORT _score DESC | LIMIT 10] // fork2
      [WHERE plot:"Shakespeare" | SORT _score DESC | LIMIT 10] // fork3
      [WHERE writers:"Shakespeare" | SORT _score DESC | LIMIT 10] // fork4
| KEEP title, semantic_title, _fork, _score, imdbrating, plot, writers
| DROP writers, semantic_title, plot
| RRF
```



_score	_fork	title	imdbrating
0.04918032786885246	[fork0, fork1, fork2]	Shakespeare in Love	7.199999889265137
0.031544957774465976	[fork1, fork3]	Othello	7.800000190734863
0.0315136476426799	[fork1, fork3]	Macbeth	7.5
0.030330882352941176	[fork1, fork3]	Much Ado About Nothing	7.400088095367432
0.016129032258064516	fork3	Julius Caesar	7.400088095367432
0.015873015873015872	fork1	Juliet of the Spirits	7.599999984632568
0.015873015873015872	fork3	Looking for Richard	7.40008895367432
0.015625	fork1	The King's Speech	8.100000381469727
0.015384615384615385	fork1	O Brother, Where Art Thou?	7.800088190734863
0.015151515151515152	fork3	Falstaff - Chimes at Midnight	7.900008895367432
0.014925373134328358	fork3	Maqbool	8.300000198734863
8.014925373134328358	fork1	Only Fools and Horses....	8.899999618530273
8.014785882352941176	fork3	Ran	8.300000198734863
0.014492753623188406	fork1	The Lovers on the Bridge	7.5
8.014285714285714285	fork1	Horatio Hornblower: The Duel	8.199999809265137



Elasticsearch Query Language

ES|QL

David Pilato - @dadoonet

Developer | Evangelist

slides & demo

