What's new in the Elastic Stack?

Alexander Reelsen
alex@elastic.co
@spinscale

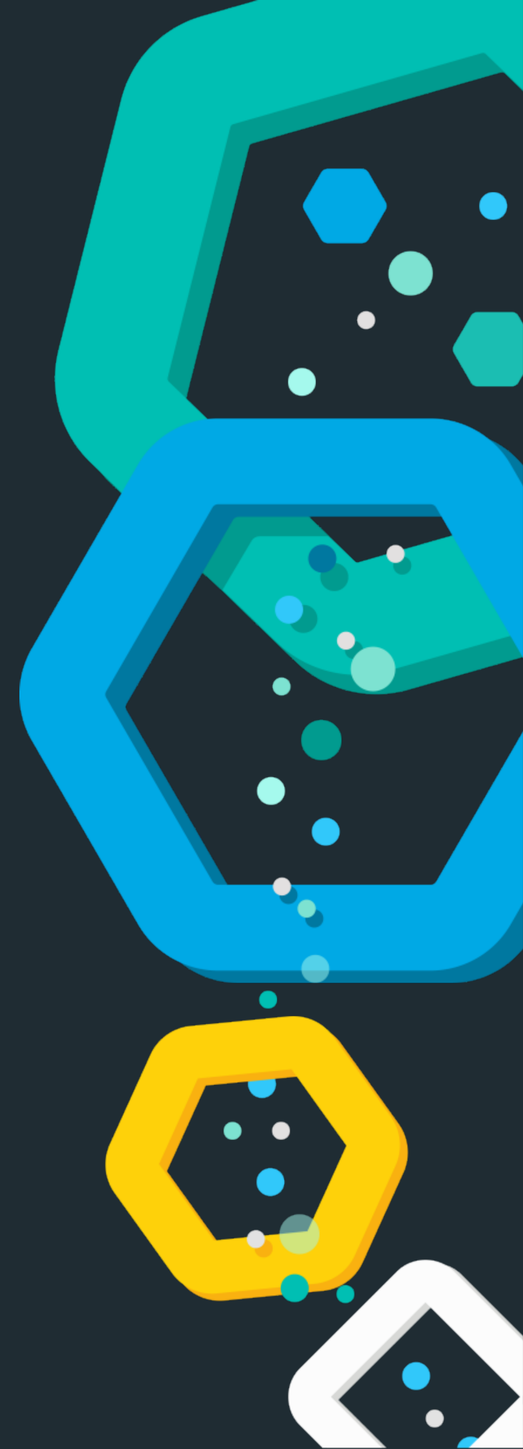

Agenda

- What's new in 6.x?
- What's new in 7.x?
- Q & A

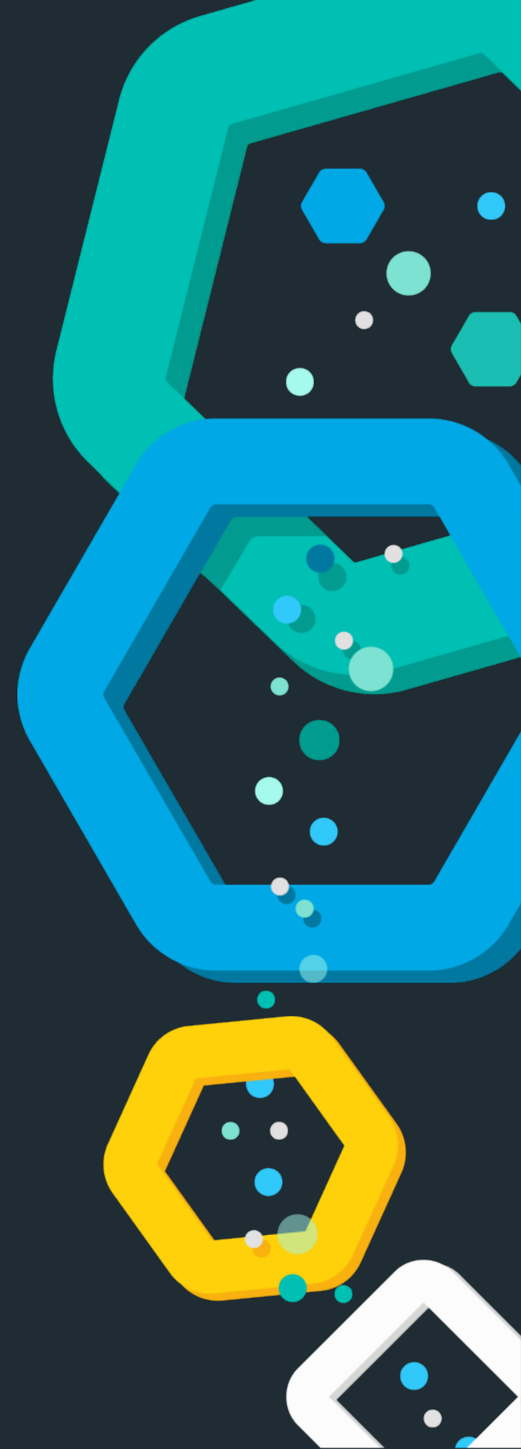

What's new in 6.x?

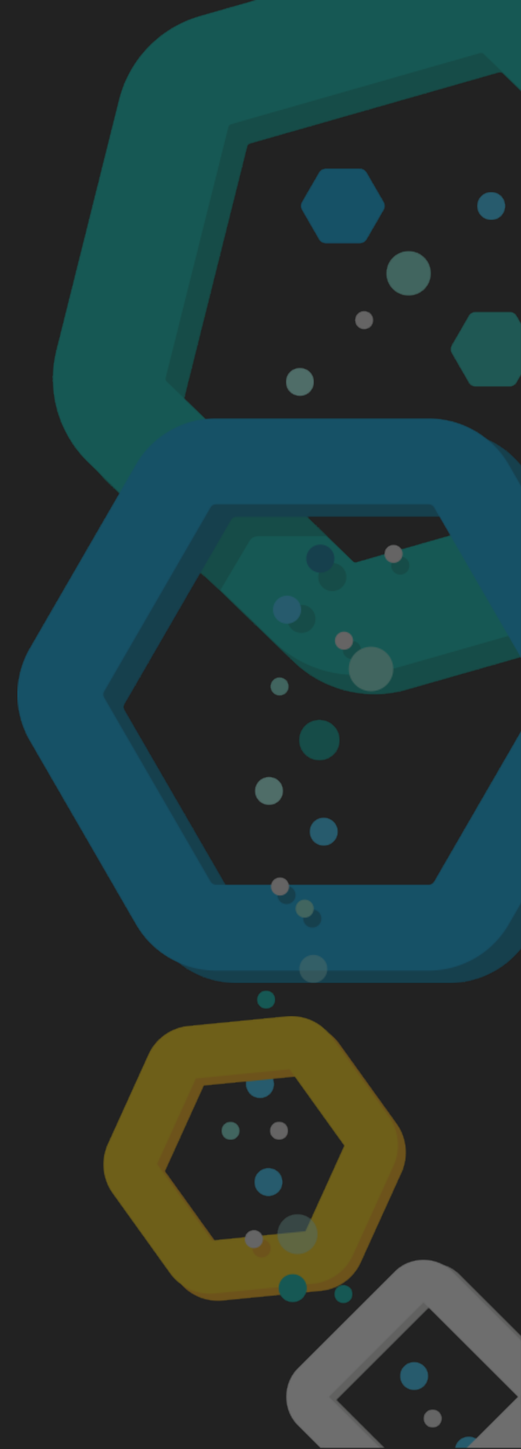

elastic

Elasticsearch 6.x

- 6.0
  - Zero downtime upgrades
  - Cross cluster search
  - Sequence id based recoveries
  - Index sorting
  - range based datatypes
- 6.1
  - Index splitting
- 6.2
  - Rank evaluation API
- 6.3
  - Rollup
  - Java 10 support
- 6.4
  - Reloadable secure settings
  - Field Aliases
  - Korean analyzer
- 6.5
  - G1GC support, Java 11
  - Minimal snapshots (50% less)
- 6.6
  - Frozen indices
  - BKD backed geoshapes
- 6.7
  - CCR
  - SQL
  - ILM
  - Upgrade Assistant
- 6.8
  - ECK (Elastic for Kubernetes)
  - Move security features into basic

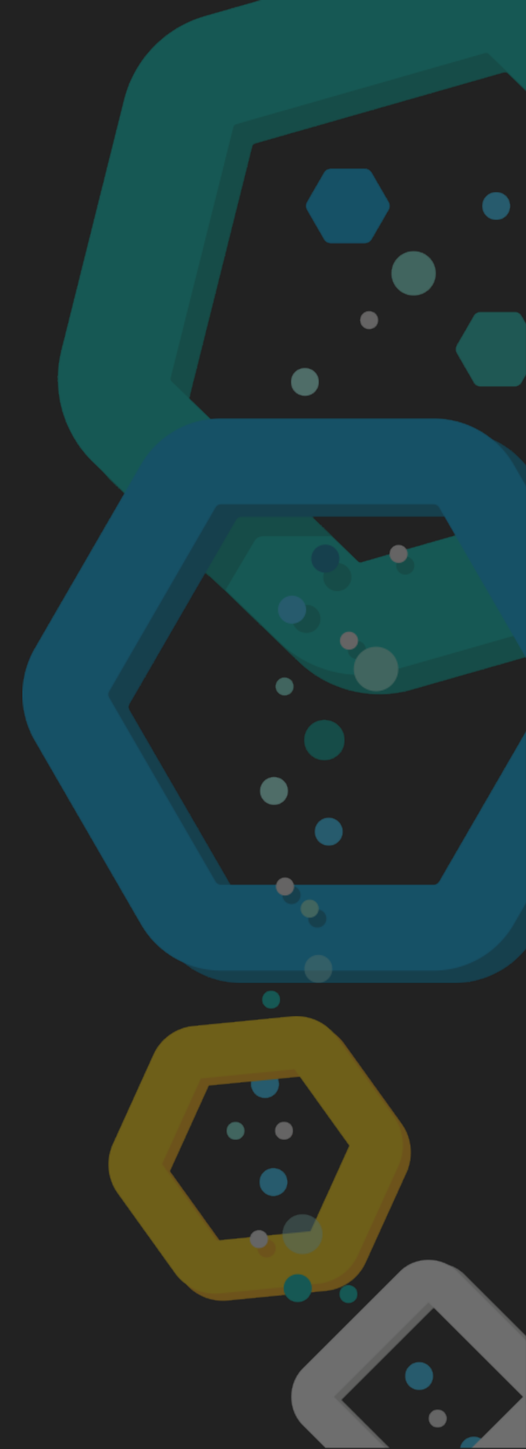

elastic

What's new in 7.x?

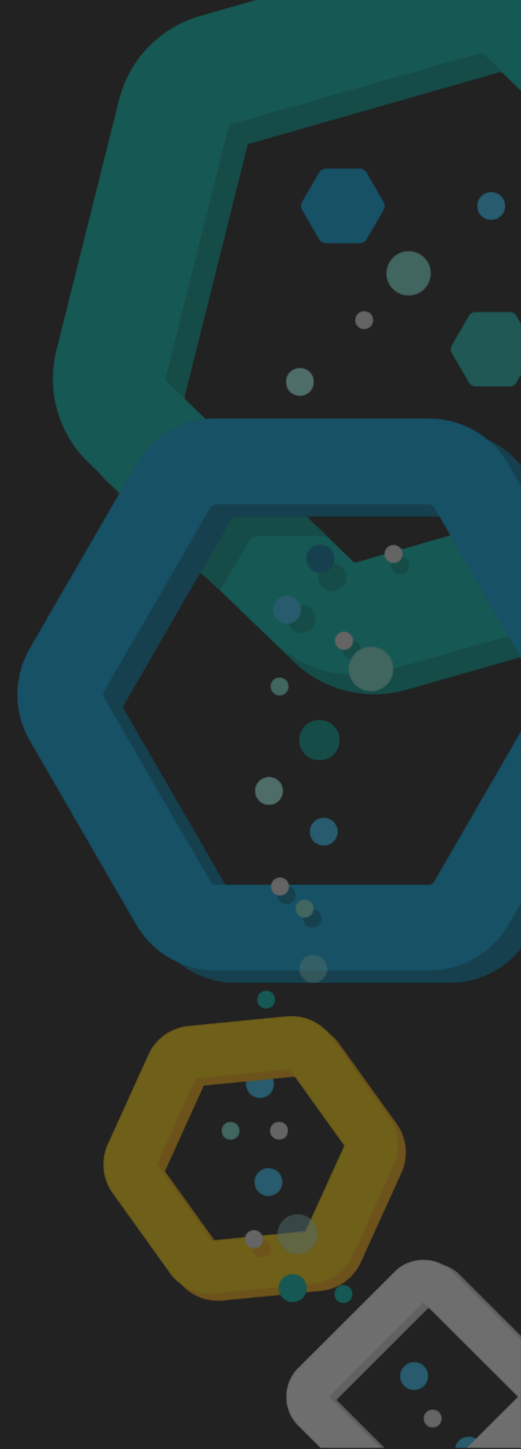

elastic

Kibana 7.x

- 7.0
  - Elastic UI Library
  - KQL by default (+ autocomplete)
  - Responsive dashboards
  - Dark mode
- 7.1
  - ECK (Elastic for Kubernetes)
  - Move security features into basic
- 7.2
  - Feature controls
  - Maps in dashboards
  - Export/Import saved objects
  - Metrics Explorer
- 7.3
  - Maps is now GA!
  - Kerberos support

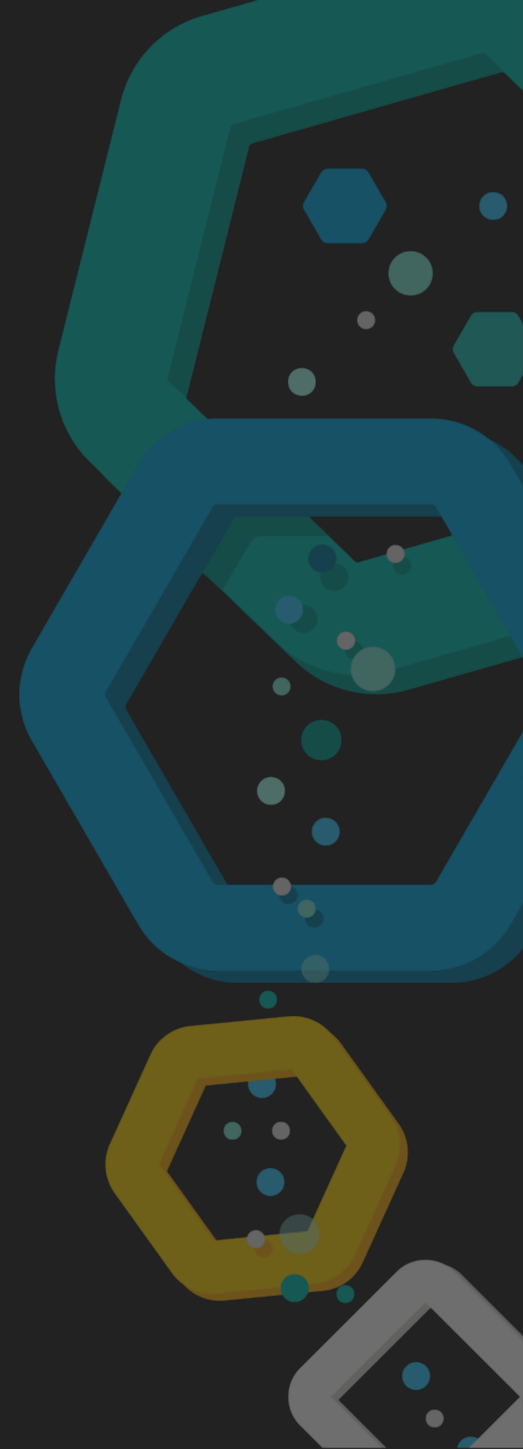

elastic

Beats 7.x

- 7.0
  - ECS
  - Filebeat: zeek, santa, netflow support, encodings
  - Auditbeat: system module
  - Metricbeat: Elasticsearch, Logstash & Kibana modules
  - Metricbeat: NATS, MSSQL, EC2, CouchDB
- 7.1
  - ECK (Elastic for Kubernetes)
  - Move security features into basic
- 7.2
  - script processor
  - security analytics: palo alto networks, cisco ASA, netflow
  - CoreDNS modules
  - windows: sysmon & security module
  - filebeat: container input
- 7.3
  - google cloud module, google pub/sub input
  - database support: oracle, RDS, cockroachdb
  - k8s monitoring
  - configuration only metricbeat modules

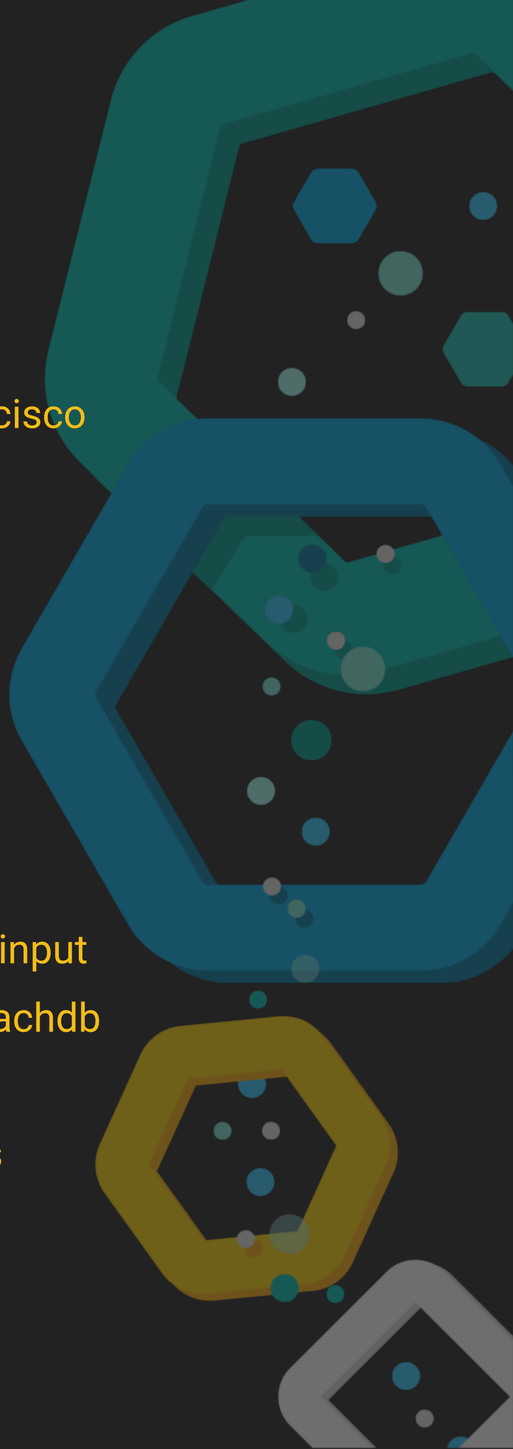

elastic

Logstash 7.x

- 7.0
  - Default java pipeline execution
  - ILM support
- 7.1
  - ECK (Elastic for Kubernetes)
  - Move security features into basic
- 7.2
  - Gradle based plugin workflow
  - JMS input
  - app search plugin
- 7.3
  - Pipeline to pipeline communication

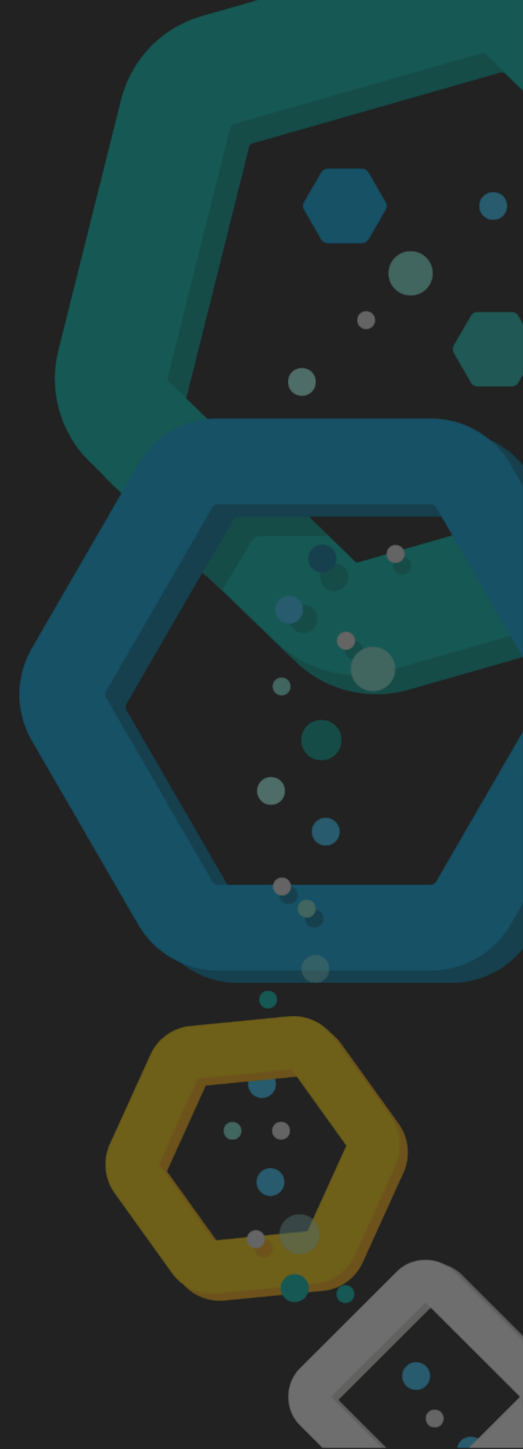

elastic

Stack 7.x

- 7.0
  - ECS
  - Hadoop: Kerberos, Java8, Cascading removed
  - Clients: rewritten JS client, new Go client, High Level REST client
- 7.1
  - ECK (Elastic for Kubernetes)
  - Move security features into basic
- 7.2
  - SIEM app
  - APM: improved Java agent metrics/ framework support
- 7.3
  - SIEM anomaly detection
  - APM: .NET agent is GA, SPA support, configure with APM UI

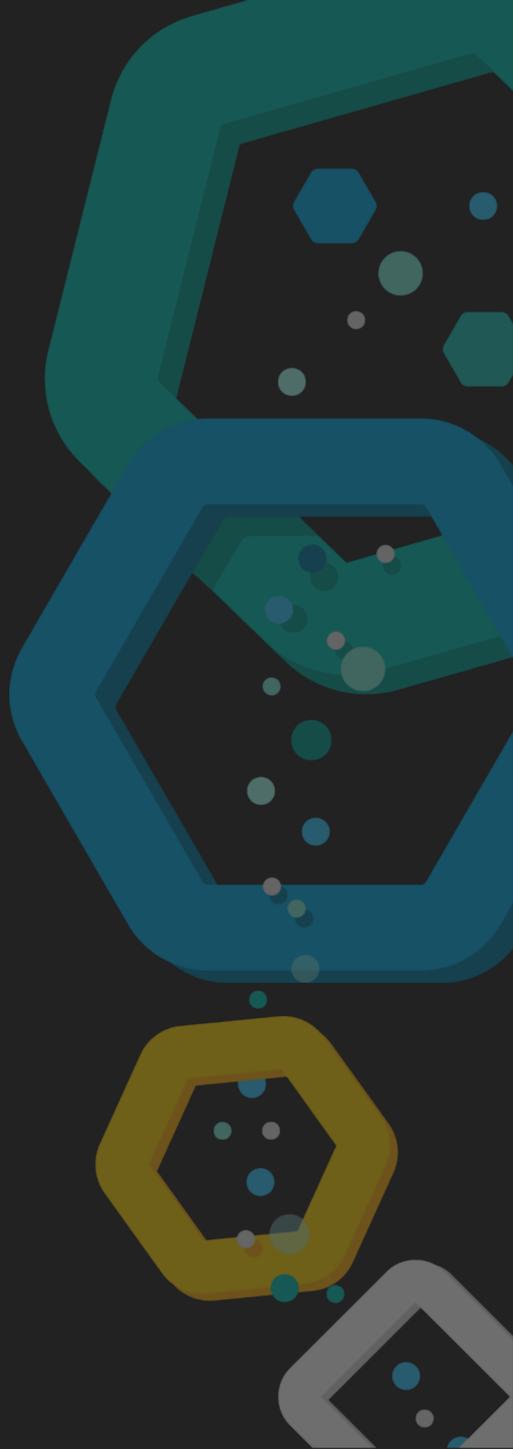

elastic

Elasticsearch 7.x

- 7.0
  - faster top-k retrieval
  - adaptive replica selection enabled by default
  - No refresh on idle shards (faster indexing)
  - `date_nanos` datatype
  - cluster coordination
  - `script_score` query
  - High Level REST client
  - Single shard index by default
  - ships with OpenJDK
- 7.1
  - ECK (Elastic for Kubernetes)
  - Move security features into basic
- 7.2
  - `search_as_you_type` datatype
  - replicated closed indices
  - `distance_feature` query
- 7.3
  - data frames
  - `rare_terms` aggregation
  - vector datatypes (`dense` & `sparse`)
  - voting only nodes
  - updateable synonyms
- 7.4
  - pinned queries
  - SLM

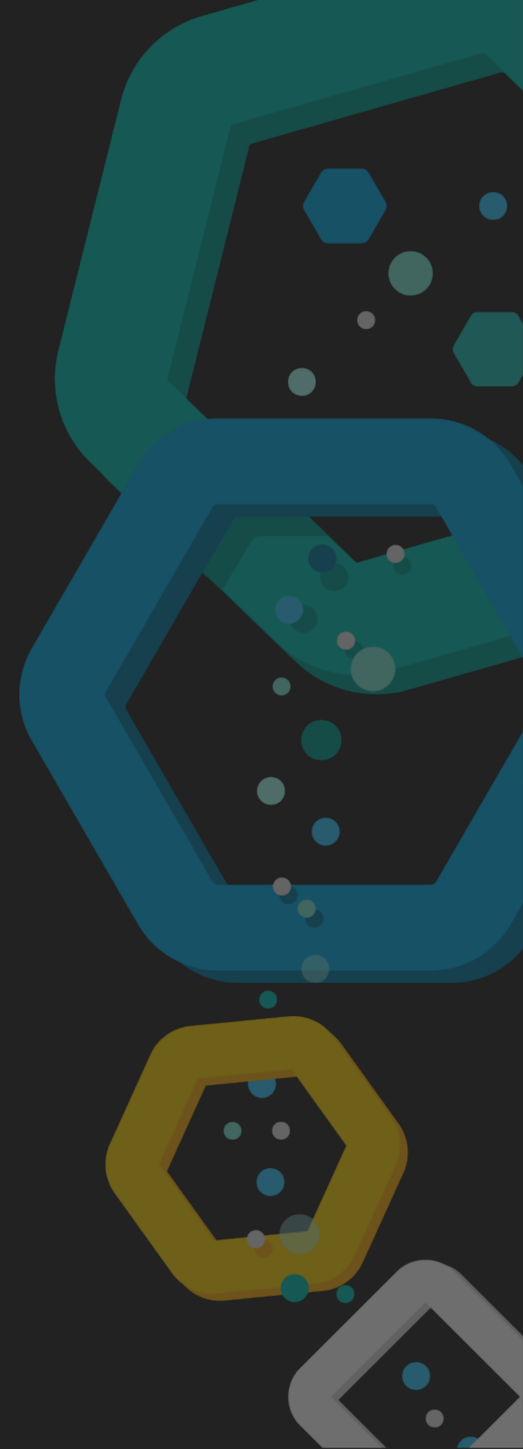

elastic

Elasticsearch 7.0 - Rewritten cluster coordination

- Gone: `discovery.zen.minimum_master_nodes`
- Sub-second master election
- Simplifying growing/shrinking of cluster
- Cluster bootstrapping/Voting configuration
- Rolling upgrades from 6 to 7 work
- Formal verification via TLA+

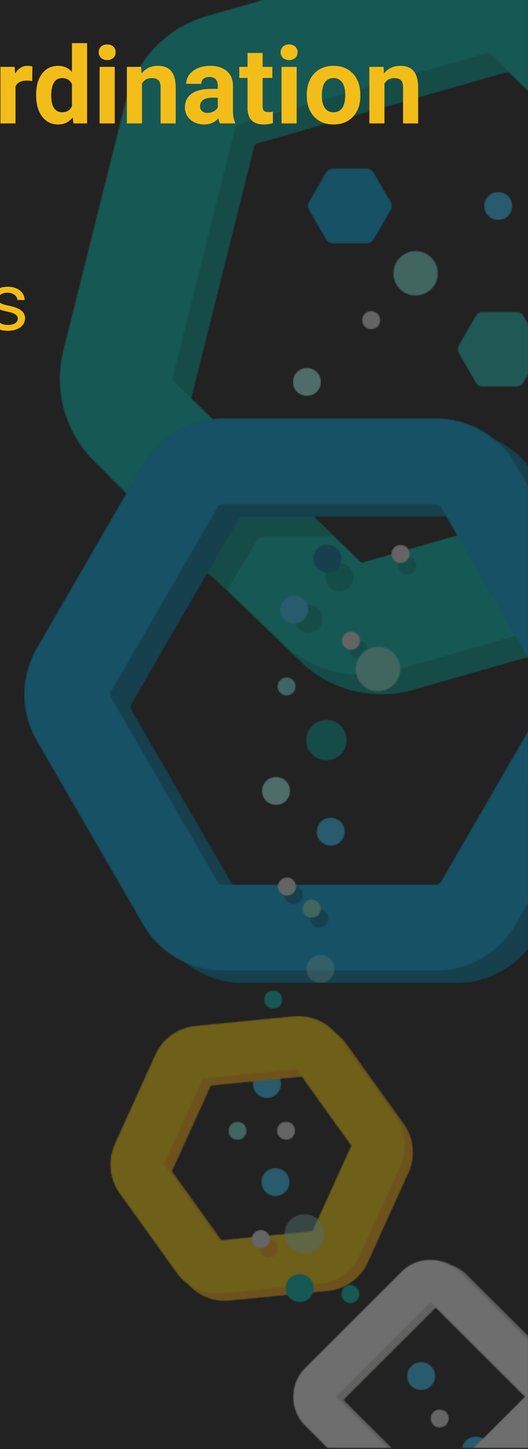

elastic

Elasticsearch 7.0 - Faster top-k retrieval

- While querying, exclude documents that cannot make it into the top hits
- Search: Elasticsearch OR Kibana
- Term 1: Elasticsearch (max score 5.0)
- Term 2: Kibana (max score 3.0)
- If first k results all have a score > 3.0, then documents only containing Kibana can be ignored
- Number of potential candidates is reduced while running

elastic

Elasticsearch 7.0 - Faster top-k retrieval

- Scores may no longer be negative
- Total hits are not counted by default

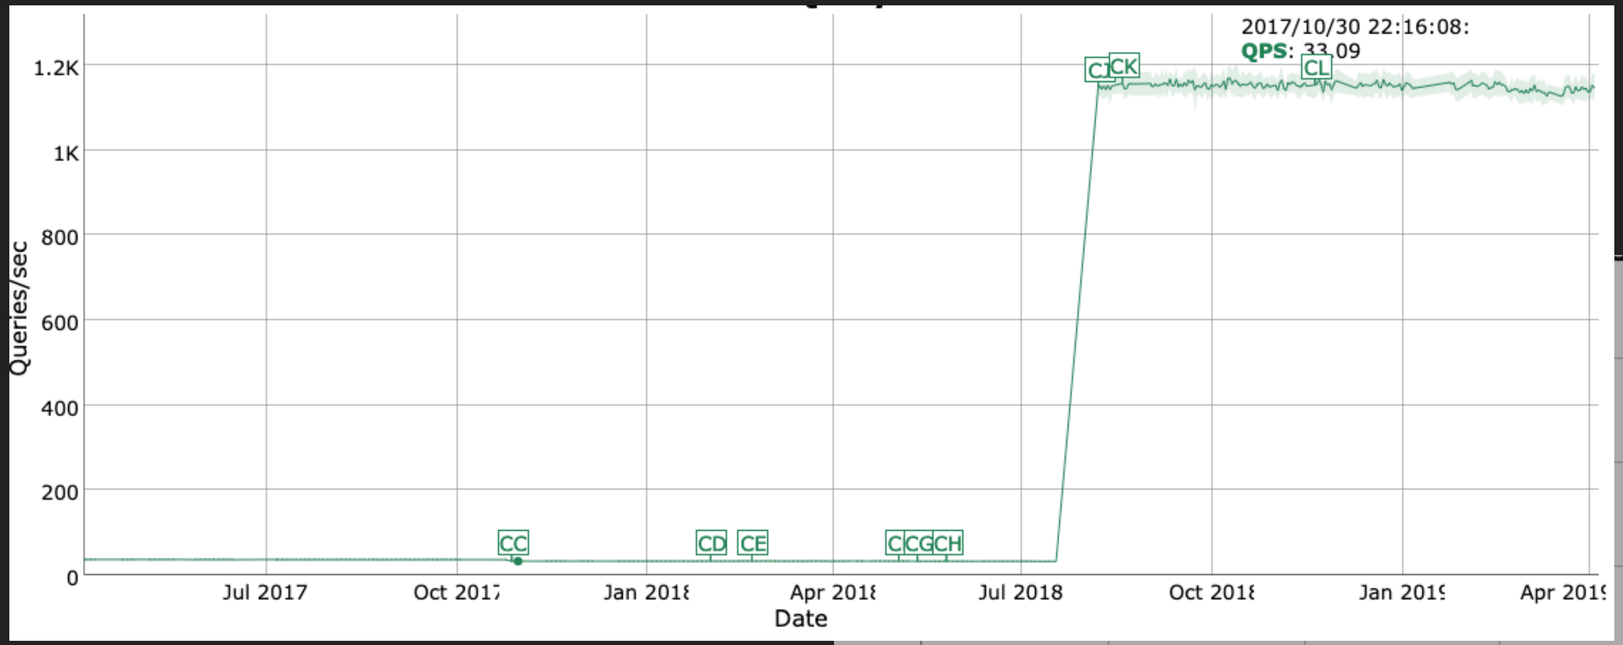

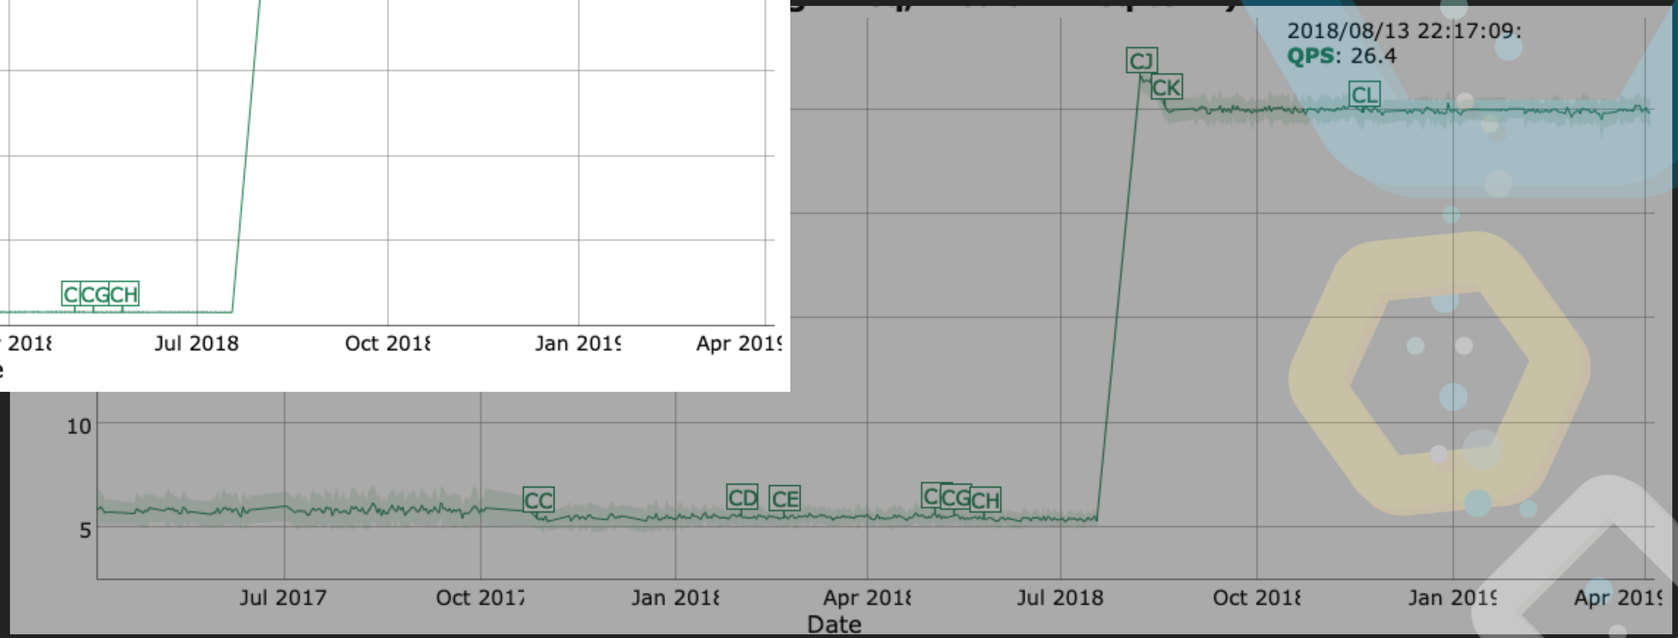

elastic

Elasticsearch - Adaptive Replica Selection

- Problem: Coordinating node round robins requests between data nodes
- Underperforming node harms the whole cluster
- Adaptive replica selection
  - Response time of previous requests
  - Search execution time of the data node
  - Queue size of the search threadpool on the data node

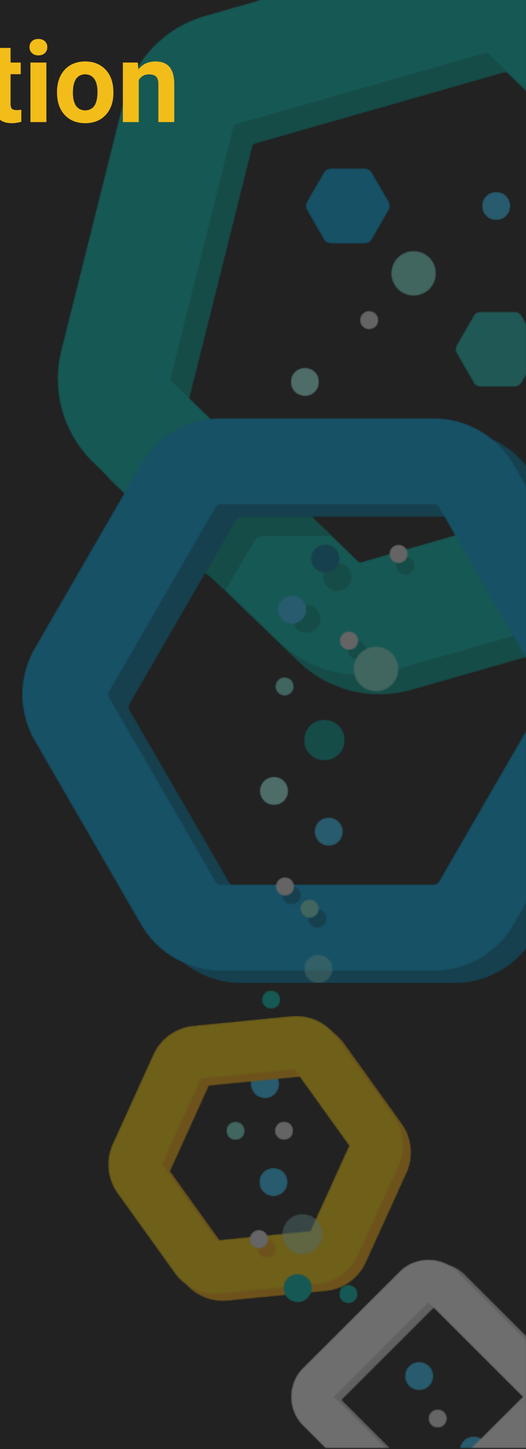

elastic

Elasticsearch - Nanosecond support

- new datatype: **`date_nanos`**
- stores nanoseconds since the epoch (reduced range!)
- internally: moved from Joda-Time to java time
- Aggregations: millisecond resolution!

- Beware: Upgrade path from 6.x!

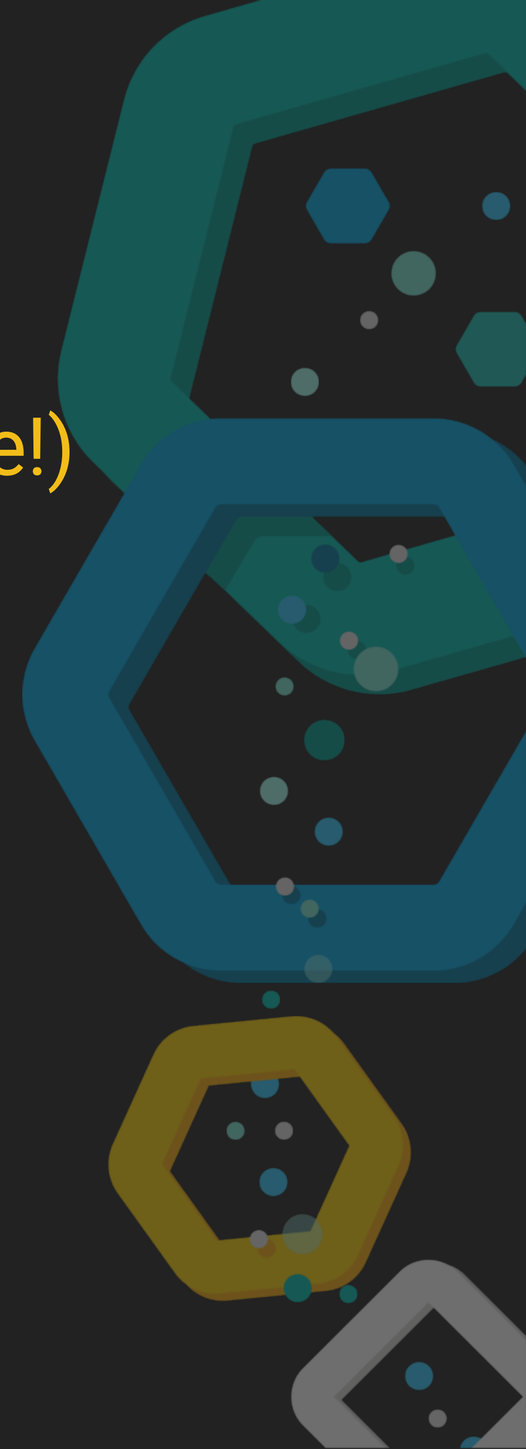

elastic

Discussion

… ask all the things!

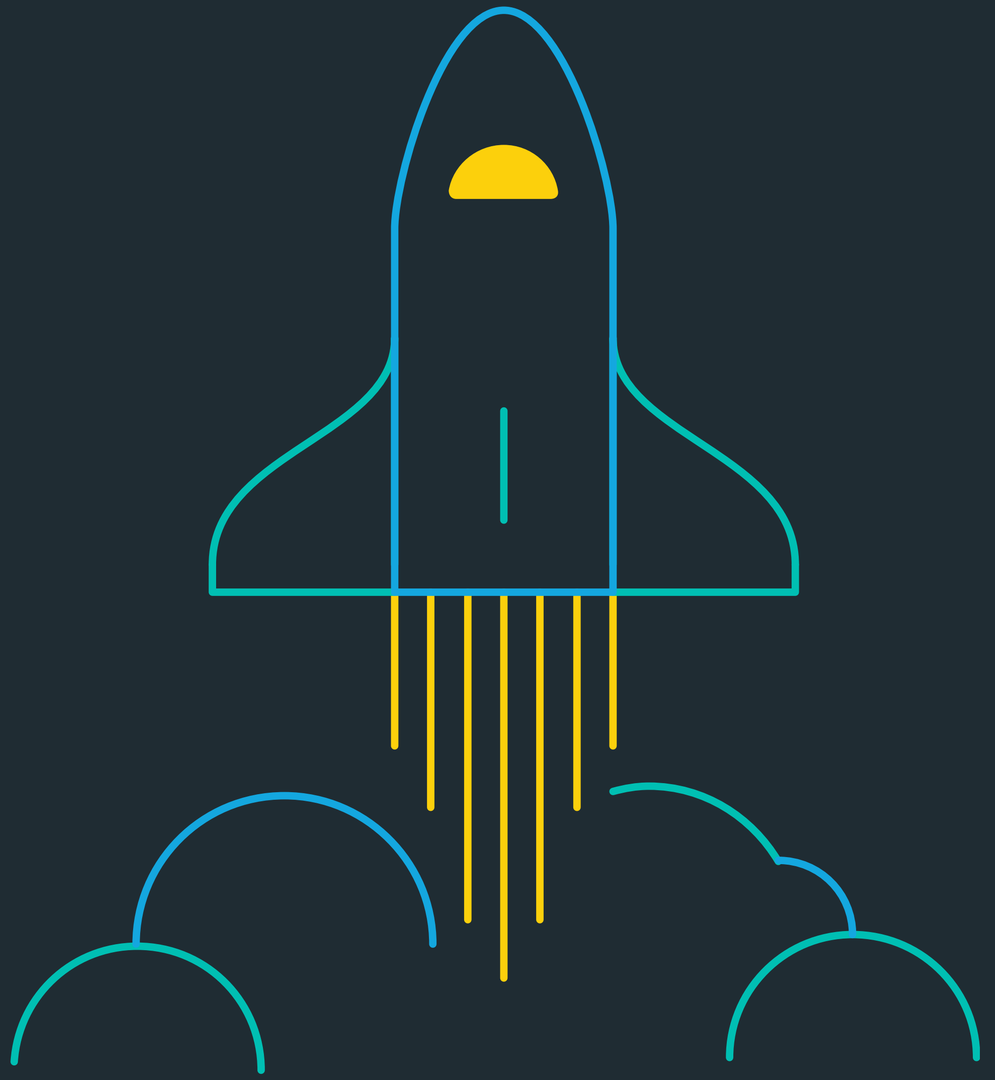

Links

- Elasticsearch
  - https://www.elastic.co/blog/easier-relevance-tuning-elasticsearch-7-0
  - https://www.elastic.co/blog/faster-retrieval-of-top-hits-in-elasticsearch-with-block-max-wand
  - https://www.elastic.co/blog/creating-frozen-indices-with-the-elasticsearch-freeze-index-api
  - https://www.elastic.co/blog/follow-the-leader-an-introduction-to-cross-cluster-replication-in-elasticsearch
  - https://www.elastic.co/blog/moving-from-types-to-typeless-apis-in-elasticsearch-7-0
  - https://www.elastic.co/blog/improving-node-resiliency-with-the-real-memory-circuit-breaker
  - https://www.elastic.co/blog/a-new-era-for-cluster-coordination-in-elasticsearch
  - https://www.elastic.co/elasticon/conf/2018/sf/reliable-by-design-applying-formal-methods-to-distributed-systems
  - https://github.com/elastic/elasticsearch-formal-models
  - C3: https://www.usenix.org/system/files/conference/nsdi15/nsdi15-paper-suresh.pdf
- Beats
  - https://www.elastic.co/blog/introducing-auditbeat-system-module

Links

- https://www.elastic.co/blog/security-for-elasticsearch-is-now-free
- https://www.elastic.co/blog/introducing-elastic-cloud-on-kubernetes-the-elasticsearch-operator-and-beyond

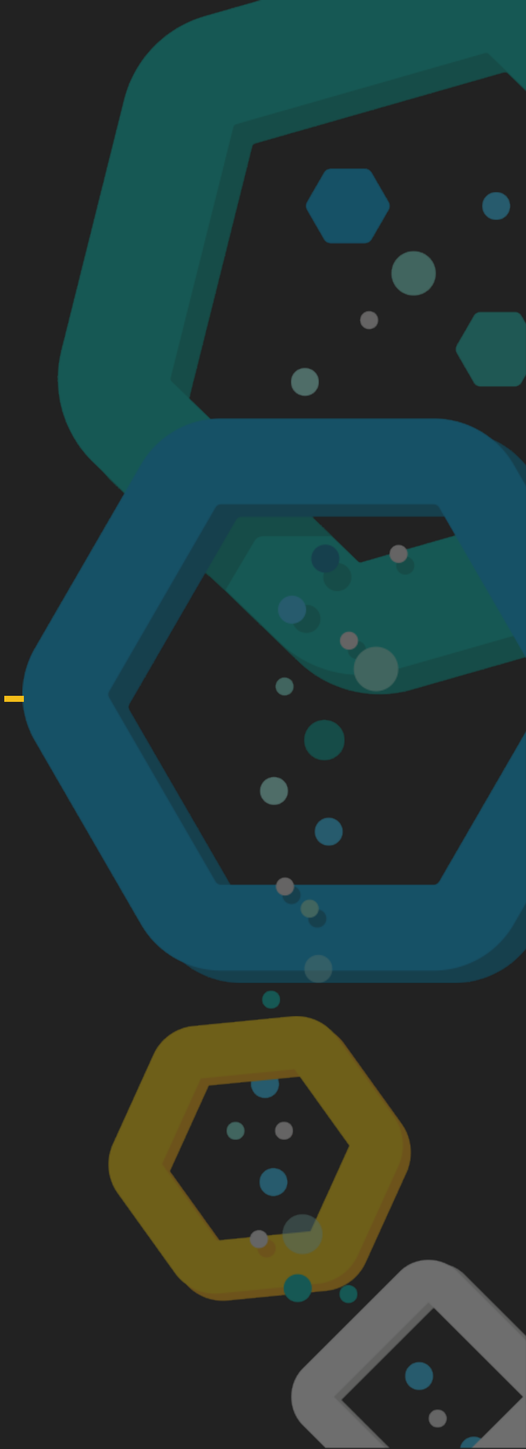

elastic