



Azure Cognitive Search

Daniel Ward

Me



- ▶ Software consultant/coach at Lean TECHniques
- ▶ ~3 years at Lean TECHniques
- ▶ ~10 years in software
- ▶ Co-organizer of the San Antonio .NET User Group
- ▶ Blog: daninacan.com
- ▶  Twitter X: @danielwarddev
- ▶  LinkedIn: daniel-ward-dev



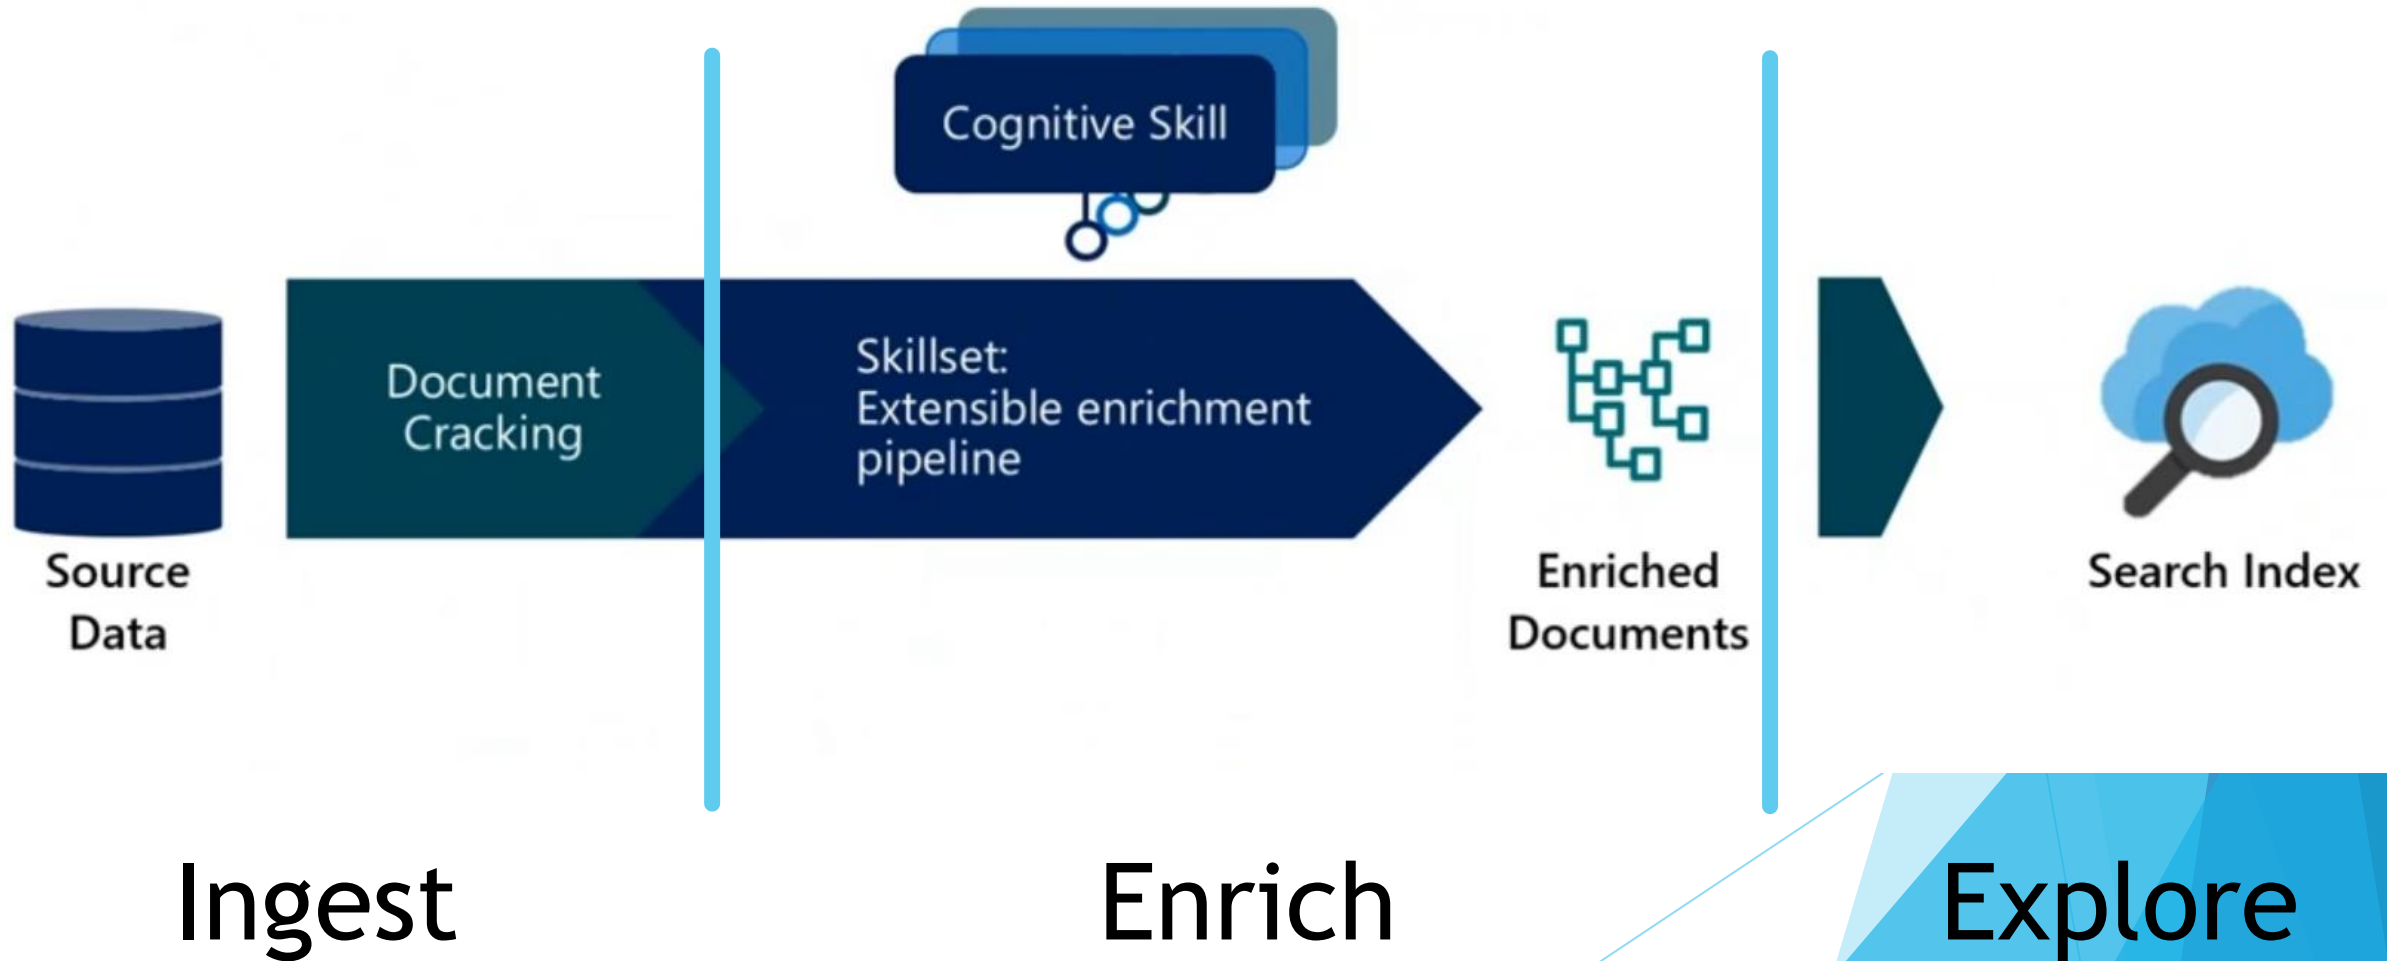
Overview

- ▶ What is Azure Cognitive Search?
- ▶ High-level overview
- ▶ Pricing
- ▶ Semantic search
- ▶ Vector search
- ▶ Demo
- ▶ Process recap
- ▶ Why use it? Why not?

What is Azure Cognitive Search?

- ▶ PaaS search service that utilizes cognitive services to build a rich search experience over private data
- ▶ You're not using AI, but the service is
- ▶ Ingests unstructured content, different types
- ▶ Semantic searching
- ▶ Using cognitive services to both ingest data and search it

High-level overview



Pricing

Example 1: Structured Data

You want to load 75GB of structured data into a search index. As your data is already structured, you will only require the main Azure Cognitive Search service (no additional features).

Pricing Tier: Standard S1 (search units of 25GB):

Standard S1 (Search units of 25GB)

Partitions: $75\text{GB} / 25\text{GB} = 3$

Replicas: 2 (requirement for 99.9% SLA for search queries)

Search Units: $3 \text{ partitions} \times 2 \text{ replicas} = 6$

Monthly Cost: $6 \text{ search units} \times (\$0.336 \times 730 \text{ monthly hours}) =$

\$1,471.68/month

[See this example in the pricing calculator >](#)

Example 2: Unstructured Data

You want to load 1,000 documents (assume 10GB) into a Cognitive Search service. Each document has 2 images. In addition to your main Azure Cognitive Search service, you'll use Document Cracking Image Extraction to extract the images, and Azure AI Services to tag images (to make them searchable).

Pricing Tier: Standard S1 (search units of 25GB):

Search Units: $1 \text{ partition} \times 2 \text{ replicas} = 2$

Monthly Search Unit Cost: $2 \text{ search units} \times (\$0.336 \times 730 \text{ monthly hours}) =$
\$490.56

One-Time AI Enrichment Costs: $\$490.56 + \$4 \text{ one-time enrichment cost}$

Document Cracking Image Extraction (2,000 images) = **\$2**

[Azure AI Services](#): Computer Vision (2,000 images) = **\$2**

\$490.56/month

[See this example in the pricing calculator >](#)

Pricing

	Free	Basic	Standard S1	Standard S2	Standard S3	Storage Optimized L1	Storage Optimized L2
Storage	50 MB	2 GB	25 GB (max 300 GB per service)	100 GB (max 1 TB per service)	200 GB (max 2 TB per service)	1 TB (max 12 TB per service)	2 TB (max 24 TB per service)
Max indexes per service	3	15	50	200	200 or 1,000/partition in high density ¹ mode	10	10
Scale out limits	N/A	Up to 3 units per service (max 1 partition; max 3 replicas)	Up to 36 units per service (max 12 partition; max 12 replicas)	Up to 36 units per service (max 12 partition; max 12 replicas)	Up to 36 units per service (max 12 partition; max 12 replicas) up to 12 replicas in high density ¹ mode	Up to 36 units per service (max 12 partition; max 12 replicas)	Up to 36 units per service (max 12 partition; max 12 replicas) up to 12 replicas in high density ¹ mode
Price per SU (Scale Unit)	\$0/month	\$73.73/month	\$245.28/month	\$981.12/month	\$1,962.24/month	\$2,802.47/month	\$5,604.21/month

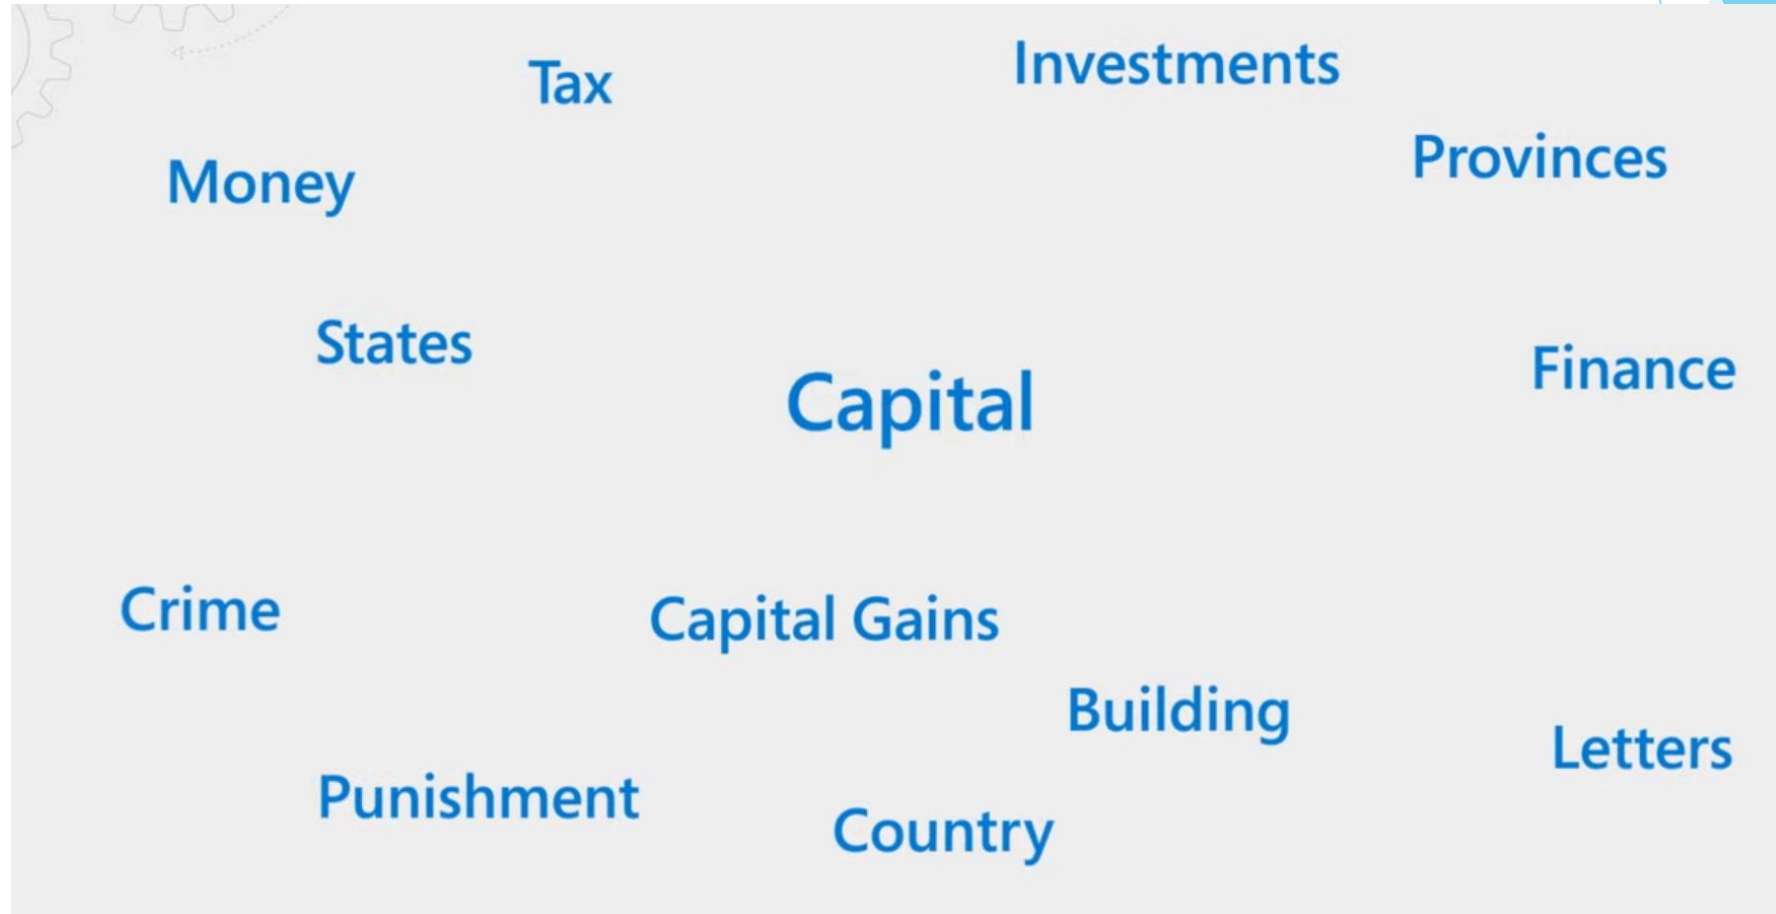
Usage Tiers	Free Try for free	Basic Dedicated environment for dev/test	Standard Run production workloads	Premium Enhanced performance and scale	Isolated High-Performance, Security and Isolation
Web, mobile, or API apps	10	Unlimited	Unlimited	Unlimited	Unlimited
Disk space	1 GB	10 GB	50 GB	250 GB	1 TB
Maximum instances	–	Up to 3	Up to 10	Up to 30*	Up to 100
Custom domain	–	Supported	Supported	Supported	Supported
Auto Scale	–	–	Supported	Supported	Supported
Hybrid Connectivity	–	Supported	Supported	Supported	Supported
Virtual Network Connectivity	–	Supported	Supported	Supported	Supported
Private Endpoints	–	Supported	Supported	Supported	Supported
Compute Type	Shared	Dedicated	Dedicated	Dedicated	Isolated
Pay as you go price	Free	\$13.140/month	\$69.35/month	\$81.03/month	\$277.40/month

Semantic search

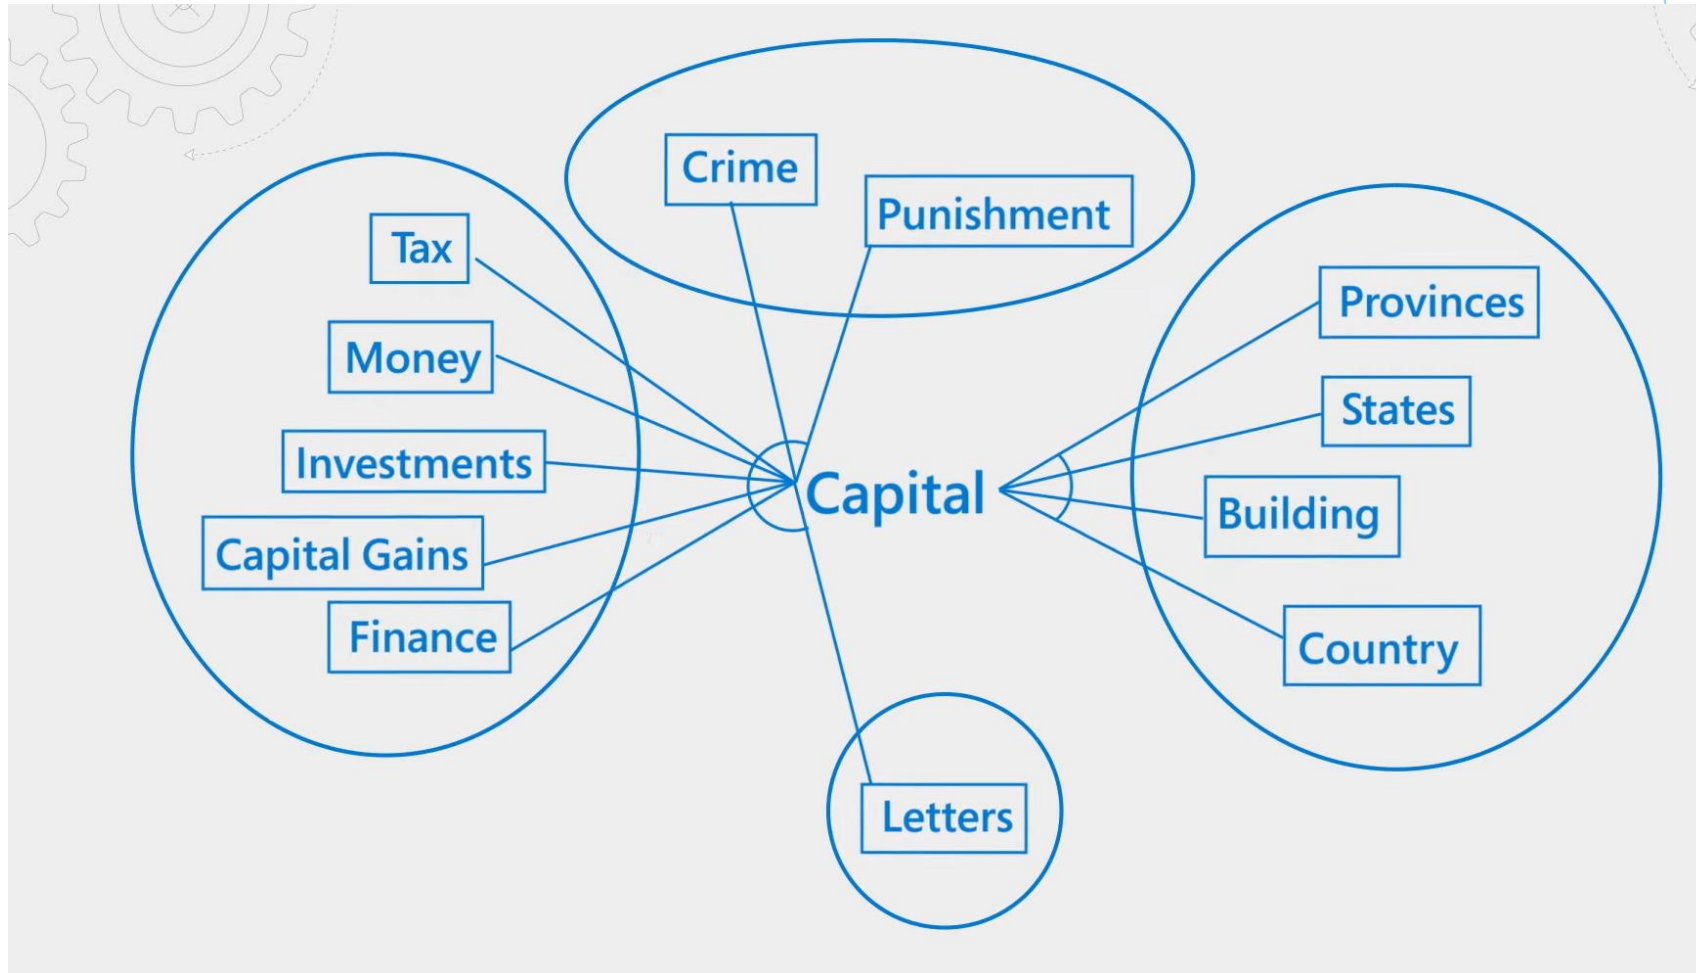


Capital

Semantic search



Semantic search





what's the capital of france?



Keyword

Capital punishment in France

https://en.wikipedia.org/wiki/Capital_punishment_in_France

"From Wikipedia, the free encyclopedia navigation search Capital punishment in France (French: peine de mort en France) is banned by Article 66-1 of the Constitution of the French Republic, voted as a constitutional amendment by the Congress of the French Parliament on 19 February 2007 and simply stating ""No one can be sentenced to death"" (French: Nul ne peut être condamné à mort). The death penalty was already declared illegal on 9 Octob...

Capital Gains Tax in France on Property

<http://www.frenchpropertylinks.com/essential/capital-gains-tax-in-france.html>

"Featured Properties House to renovate with small lake EUR €39,000More Info>Characterful stone house with attached barn and garden. EUR €21,600More Info>House on Large Flat Plot in Samoens EUR €297,000More Info>Central Dinan - Two Charming Stone Houses - One Habitable, One to Develop EUR €72,500More Info>Farmhouse with land and pool EUR €236,250More Info>Full list of property for sale in France Capital Gains...

What's the capital of Kazakhstan?

<http://www.kazakhstan.discovery.com/capital-of-kazakhstan.html>

What's the capital of Kazakhstan? Today the capital of Kazakhstan is Astana. During its history Kazakhstan had more than one capital. Let's start our list from the latest to the first. Astana literally means Capital in Kazakh language. It's the word primarily borrowed from Persian, meaning a place where the court is seated (capital city). It is the second largest city in Kazakhstan, after Almaty. It is located in north-central part of Kazakhstan, just under...

Removals to France - Ten things France is Most Famous For

<http://removals-to-france-paris.blogspot.com/2012/12/removals-to-france-ten-things-france-is.html>

Removals to France - Ten things France is Most Famous For If you require removals to France you are probably beginning to think about your new life in France, enjoying its renowned food and wine, visiting the many art galleries and monuments, and taking in the stunning scenery. Whatever your reasons for moving abroad, Hamiltons offer a full range of professional removal services, including full or part load transportation via road...

Semantic

Île-de-France

<https://www.britannica.com/place/Ile-de-France-region-France>

Île-de-France, région of France encompassing the north-central départements of Val-d'Oise, Seine-et-Marne, Seine-Saint-Denis, Ville-de-Paris, Hauts-de-Seine, Val-de-Marne, Essonne, and Yvelines. Île-de-France is bounded by the régions of Hauts-de-France to the north, Grand Est to the east, Bourgogne-Franche-Comté to the southeast, Centre to the south, and Normandy to the northwest. The capital is Paris. Area 4,637 square miles...

France

<https://www.britannica.com/place/France>

France, officially French Republic, French France or République Française, country of northwestern Europe. Historically and culturally among the most important nations in the Western world, France has also played a highly significant role in international affairs, with former colonies in every corner of the globe. Bounded by the Atlantic Ocean and the Mediterranean Sea, the Alps and the Pyrenees, France has long provided a geographic...

Where is France?

<https://www.worldatlas.com/eu/fr/where-is-france.html>

Where is France? Location of France on a map. What is the capital of France? Located in the continent of Europe, France covers 640,427 square kilometers of land and 3,374 square kilometers of water, making it the 43rd largest nation in the world with a total area of 643,801 square kilometers. France was founded as a distinct and original nation in 943. The population of France is 65,630,692 (2012) and the nation has a density of 102 people per...

Where is France?

<http://www.worldatlas.com/eu/fr/where-is-france.html>

Where is France? Location of France on a map. What is the capital of France? Located in the continent of Europe, France covers 640,427 square kilometers of land and 3,374 square kilometers of water, making it the 43rd largest nation in the world with a total area of 643,801 square kilometers. France was founded as a distinct and original nation in 943. The population of France is 65,630,692 (2012) and the nation has a density of 102 people per...

Semantic

Île-de-France

<https://www.britannica.com/place/Île-de-France-region-France>

Île-de-France, région of France encompassing the north-central départements of Val-d'Oise, Seine-et-Marne, Seine-Saint-Denis, Ville-de-Paris, Hauts-de-Seine, Val-de-Marne, Essonne, and Yvelines. Île-de-France is bounded by the régions of Hauts-de-France to the north, Grand Est to the east, Bourgogne-Franche-Comté to the southeast, Centre to the south, and Normandy to the northwest. The capital is Paris. Area 4,637 square miles...

France

<https://www.britannica.com/place/France>

France, officially French Republic, French France or République Française, country of northwestern Europe. Historically and culturally among the most important nations in the Western world, France has also played a highly significant role in international affairs, with former colonies in every corner of the globe. Bounded by the Atlantic Ocean and the Mediterranean Sea, the Alps and the Pyrenees, France has long provided a geographic...

Where is France?

<https://www.worldatlas.com/eu/fr/where-is-france.html>

Where is France? Location of France on a map. What is the capital of France? Located in the continent of Europe, France covers 640,427 square kilometers of land and 3,374 square kilometers of water, making it the 43rd largest nation in the world with a total area of 643,801 square kilometers. France was founded as a distinct and original nation in 943. The population of France is 65,630,692 (2012) and the nation has a density of 102 people per...

Where is France?

<http://www.worldatlas.com/eu/fr/where-is-france.html>

Where is France? Location of France on a map. What is the capital of France? Located in the continent of Europe, France covers 640,427 square kilometers of land and 3,374 square kilometers of water, making it the 43rd largest nation in the world with a total area of 643,801 square kilometers. France was founded as a distinct and original nation in 943. The population of France is 65,630,692 (2012) and the nation has a density of 102 people per...

Paris is the capital city of France, the capital of love and one of the principal tourist spots in modern Europe. The current population of the city approaches to 11 million people, and for many centuries **Paris** is considered to be a cultural and fashion capitals of the continent.

Paris, France

<http://www.latlong.net/place/paris-france-1666.html>

Île-de-France

<https://www.britannica.com/place/Île-de-France-region-France>

The capital is **Paris**. Area 4,637 square miles (12,011 square km). Pop. (1999) 10,952,011; (2014 est.) 12,027,565. The Seine River along the Île Saint-Louis, **Paris**.

France

<https://www.britannica.com/place/France>

The capital and by far the most important city of **France** is **Paris**, one of the world's preeminent cultural and commercial centres.

Where is France?

<https://www.worldatlas.com/eu/fr/where-is-france.html>

Paris is the **capital** city of **France**. It has a population of 2,138,551, and is located on a latitude of 48.85 and longitude of 2.35. **Paris** is also the political center of France, which is considered a Republic, and home to its Executive head of state.

Semantic w/
captions +
answer

Demo



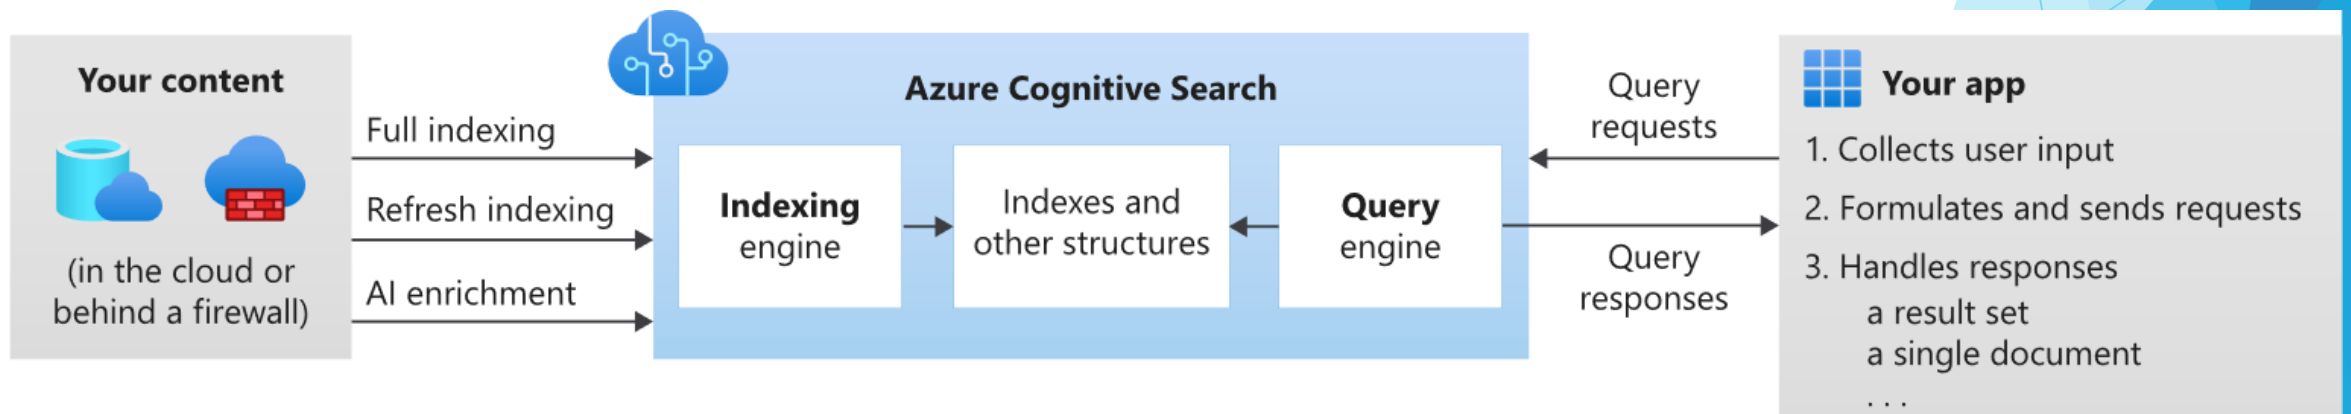
```
POST https://semantic.search.windows.net/indexes/msmarco/docs/search?api-version=2020-06-30-P
Content-Type: application/json
api-key: 9674A962D9C3C75A38C19B26CD49BB4D
```

```
{
  "search": "what's the capital of france?",
  "queryType": "semantic",
  "queryLanguage": "en-us",
  "searchFields": "Title,Body",
  "select": "Title",
  "top": 10
}
```

```
4 Content-Type: application/json; odata.metadata=minimal
5 Content-Encoding: gzip
6 Expires: -1
7 Vary: Accept-Encoding
8 request-id: 404166d9-7973-42cf-a738-bf66424eaca6
9 elapsed-time: 1266
10 OData-Version: 4.0
11 Preference-Applied: odata.include-annotations="*"
12 Strict-Transport-Security: max-age=15724800; includeSubDomains
13 Date: Thu, 18 Feb 2021 18:37:00 GMT
14 Connection: close
15 Content-Length: 574
16
17 {
18   "@odata.context": "https://semantic.search.windows.net/indexes\('msmarco'\)/\$metac
19   (*)",
20   "value": [
21     {
22       "@search.score": 23.705616,
23       "@search.semanticScore": 2.5378549098968506,
24       "Title": "FRANCE"
25     },
26     {
27       "@search.score": 22.167515,
28       "@search.semanticScore": 2.5243072248995304,
29       "Title": "France"
30     },
31     {
32       "@search.score": 23.295124,
33       "@search.semanticScore": 2.5225027650594711,
34       "Title": "Paris, France"
35     },
36     {
37       "@search.score": 21.426798,
38       "@search.semanticScore": 2.4771754592657089,
39       "Title": "Paris, France"
40     },
41     {
42       "@search.score": 23.809196,
43       "@search.semanticScore": 2.4743615612387657,
44       "Title": "France"
45     }
46   ]
47 }
```

Process recap

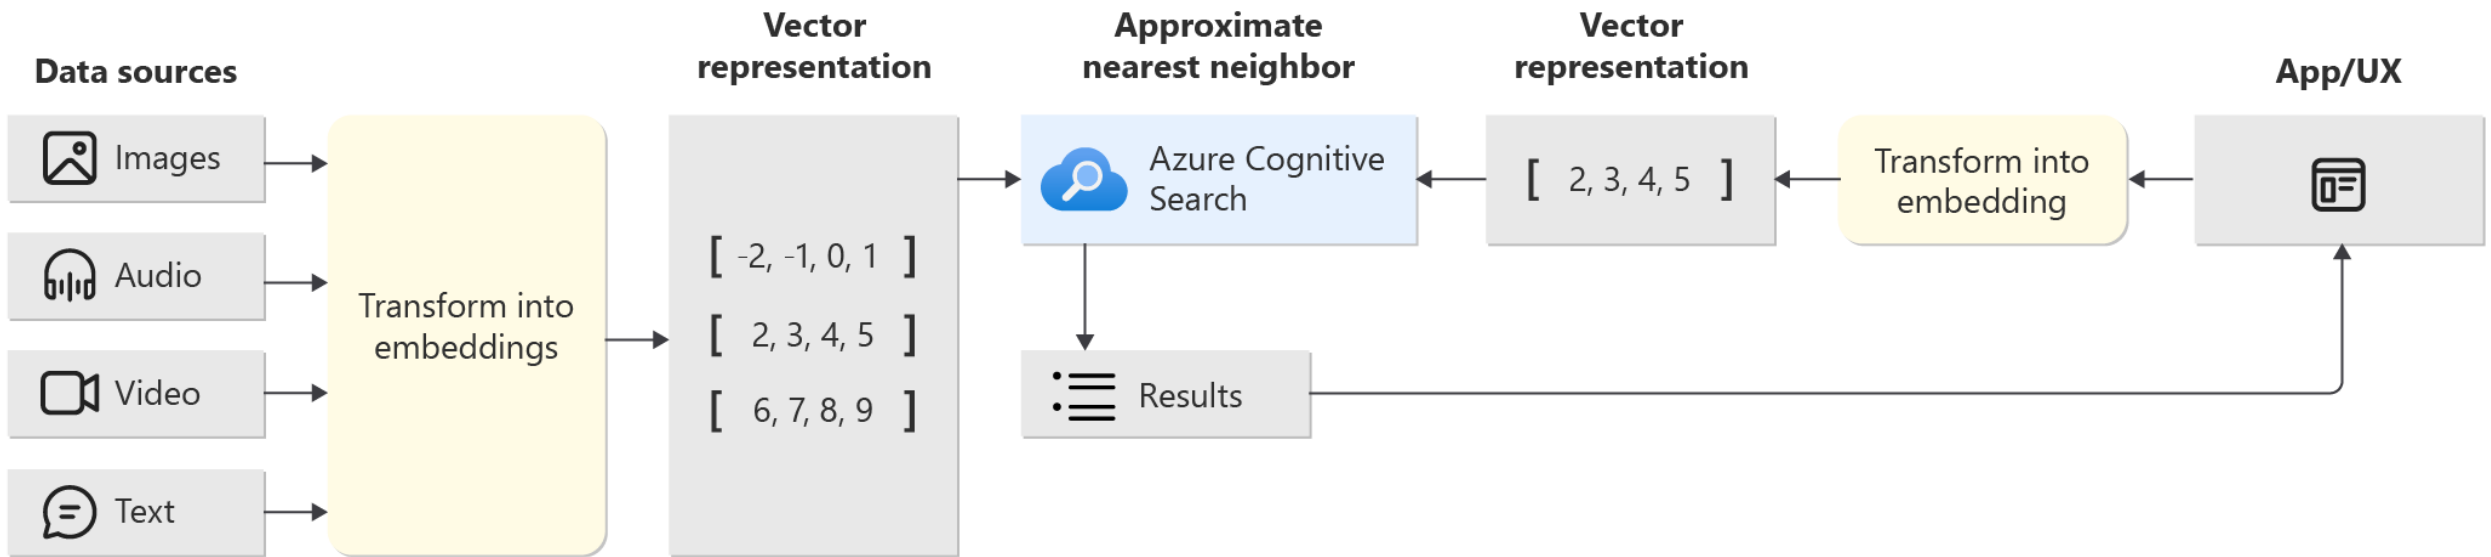
1. Bring a data source
2. Import data from it. Document cracking
3. Enrich with cognitive skills (OCR, translation, custom, etc)
4. Configure the index
5. Configure the indexer
6. Create the indexer and run it to create the index



Search strategy

- ▶ L1 - Retrieval. Quickly find all the documents from the index that satisfy the search criteria. These are scored and the top 50 are returned
 - ▶ Keyword (used in the demo) - traditional full-text search
 - ▶ Vector (next slide)
 - ▶ Hybrid - applies both keyword and vector retrieval and takes the best results from each
- ▶ L2 - Ranking. Semantic ranking takes the results from L1 and improves them. Only the top 50 are returned from L1 because this is relatively powerful

Vector search



Vector search

- ▶ Currently in public preview
- ▶ Encodes content into numeric representations based on semantic meaning
- ▶ Performs searches on those vectors
- ▶ Your content needs to have embeddings (their vector representation)
- ▶ Cognitive Search does not generate embeddings for you
- ▶ Azure OpenAI can generate them for you


Why use it?

- ▶ Don't need a data scientist at your org, don't need to create your own ranking models from scratch, etc.
- ▶ Your software needs advanced searching, but it may not be searching software
- ▶ You have different and/or unstructured data
- ▶ Use of cognitive services (ie. linguistic analysis)
- ▶ Other features
 - ▶ Keyword search
 - ▶ Geospatial support
 - ▶ Suggestions/auto-complete
 - ▶ Synonyms
 - ▶ Ranking
 - ▶ More

Why not use it?

- ▶ A simple traditional search suffices for your situation
- ▶ Sometimes, you do want traditional searching
- ▶ Example: Machine part numbers like ABC-123

Example






- ▶ Call center
- ▶ Estimate 700k calls per month, estimate average 20 minutes per call
- ▶ 14,000 minutes per month, 168,000 minutes per year
- ▶ Audio data sources

Recap

- ▶ Azure Cognitive Search service and overview
- ▶ Pricing
- ▶ Semantic search
- ▶ Vector search
- ▶ Demo
- ▶ Why/not use it?

Thank you!

- ▶ daninacan.com
- ▶   @danielwarddev
- ▶  daniel-ward-dev