Intelligence Analytics as a Service

Service Definition Document, version 1.0

Service Overview

This service provides AI-driven insight and intelligence for customers through the development and operation of tailored entity resolution, data engineering, data analysis, analytics and AI. The service combines customer data and context with external data to provide actionable intelligence on things the customer cares about

Intelligence Analytics as a Service combines a series of powerful capabilities together to enable customers to fully understand and leverage their data. It then enables customers to focus on identifying actionable insight and intelligence from their data for a range of problem sets and challenges. The core of the system is the development, deployment and management of analytics tailored to the client's specific use cases, operational context, and data holdings. This could be supply chain vulnerability and threat, fraud detection and investigation, physical security, and threats as well as more broader domain awareness, threat identification and management.

The service combines detailed customer data and accompanying context with Data Edge Analytics' wealth of experience and understanding of relevant data types, data feature identification and their operational use. Customer data combined with broader external data is resolved to provide a richer feed for analytics to identify a range of indicators which highlight and monitor things that the customer cares about.

The Component Mix

The service brings together a range of third-party commercial products (which will require their own license agreements) with existing Crown Copyright capabilities developed specifically for government departments in the defence, national security, law enforcement, and central government sectors.

Note that Data Edge Analytics are not a reseller and do not provide licenses on behalf of third-party commercial vendors. Data Edge Analytics believe that these products are currently the most appropriate for the service. The service is inherently flexible, and Data Edge Analytics recognise that customer choice is essential. What today is the optimal technology choice may not be in a few months' time.

Senzing Entity Resolution is the principle entity resolution capability offered. This sits at the core of the service, providing a store of resolved entities (such as people, organisations, vessels) as a basis for downstream analysis. Changes to entities can then be monitored in real time (if required) as all new incoming records are then matched against all the current entities, the graph is constantly being re-assessed.

Records used in the entity resolution service must contain certain identifiers (e.g. names, addresses, IDs, phone numbers, etc.) for the service to work. The service is particularly adept at resolving human entities as it is optimised to work on the messiest identifiers – names and addresses. Furthermore, the system has been extended to work with various types of events,

including travel and commercial/financial data (e.g. purchase orders). The service can work at billion row scale (e.g. bulk data). It is a real time system.

The service incorporates a core knowledge graph from which analytics are developed, entities of interest identified and alerting and watch listing activity is driven. The service includes a variety of Crown Copyright visualisation and workflow tools. The knowledge graph can be explored directly as well as being interrogated via a Graph RAG LLM chat interface. This enables a variety of different user skill types to be accommodated and enables easy interrogation of the base entities with the knowledge graph.

Deployment model: Cloud Software as a Service (SaaS)
We provide a mixture of proprietary, open source and Crown-IP applications running on a cloud infrastructure run by either the client or by their cloud service provider. We will manage (install; set up; configure; run; maintain) all intelligence analytics related services and tools.

Information assurance ¨C Impact Level (IL) at which the G-Cloud Service is accredited to hold and process information
As our services are run in the Client's cloud, the accreditation level is dependent on Client's cloud and the requirements of the datasets to be exploited by the service.

Details of the level of backup/restore and disaster recovery that will be provided on-boarding and off-boarding processes/scope etc.
As our services are run in the Client's cloud, accreditation, back up, DR etc. is dependent on the Client's cloud and associated services available therein.

Service constraints
Maintenance windows are dependent on software upgrades. These will be ad hoc and are not scheduled. Customisation of the stack (e.g. feature extensions to the tools) can be provided but this will be at additional cost and via mutual agreement. The service requires specific hardware; the specification of the hardware depends on the data volumes to be processed.

The service requires certain datasets to be made available. It is the Buyer's responsibility to ensure that such data is available (we can advise on the selection and suitability of datasets once the requirements are fully understood.)

Service levels
Operational service availability is dependent on client needs given that data engineering, data science and devsecops support will be required to run the service. We are able to provide support services within normal working hours, i.e. 0900-1700, Monday – Friday. Due to the highly technical nature of the analytical services, we only offer a blended technical support package that includes all first, second and third-line support. Once the buyer's specific support requirements are understood, we will be able to agree recompense arrangements should support levels not be met.

Associated data engineering and analysis services
Dependent on the buyer's aims and objectives, the service will include a variety of advanced data engineering and data science services. These cover dataset evaluation, preparation, transformation, ingest and load related tasks, cataloguing, advice on dataset selection, feature

identification, and analytics development. For more complex engagements this will also include assistance with data verification, validation and quality assessments.

Training

The service includes a variety of bespoke in-house training courses. This training will be tailored to the level of skill and expertise possessed by the customer and their specific needs. Data Edge Analytics want to focus on knowledge transfer and upskilling to the customer. This enables Data Edge to assist the customer on evermore complex aspects of their challenge and at the same time develop the customer capability to effectively run, maintain and exploit what they have developed.

Business change, consultancy services

Dependent on the delivery model, we can offer business change services. These aim to ensure that new ways of working consistent with the overall service offer are embedded to maximise the chances of successful delivery where the client elects not to outsource all services to us. These services embrace process development, policy formulation, change planning and communications.

Ordering and invoicing process

Due to the bespoke nature of the service, the Buyer will need to advise us of their specific needs via a formal service request. These will include client's aims; the datasets they have available; hosting deployment options; business uses cases and the service outcomes required.

Set up, configuration

We can set up, configure and install all required services in the buyer's cloud environment. (A costed technical evaluation will need to be undertaken prior to this).

Technical requirements for Senzing Entity Resolution

Below are example technical specifications for the Senzing deployment for guidance purposes only. A full technical evaluation will be conducted prior to service delivery.

Single Node Deployments

- Up to 10 million records
- Requirements: 8 cores, 48GB of RAM, 200GB of direct-attached SSD NVMe storage
- Example AWS instance: i3.2xlarge ~$0.63/hr
- Approximate ingestion rate of typical data at 100 records per second, 10 million would load in about 1 day.

Up to 50 million records

- Requirements: 16 cores, 96GB of RAM, 1TB of direct-attached SSD or NVMe storage
- Example AWS instance: i3.4xlarge ~$1.25/hr
- Approximate ingestion rate of typical data at 200 records per second,
- 50 million would load in under 3 days.

Up to 100 million records

- Requirements: 32 cores, 192GB of RAM, 2TB of direct-attached SSD or NVMe storage
- Example AWS instance: i3.8xlarge ~$2.50/hr

- Approximate ingestion rate of typical data at 400 records per second, 100 million would load in under a week.

Multi-Node Deployments
multi-node deployments easily support billions of input records. Technical specifications for this volume of records are available on re□quest.

Storage Guidelines
The general guideline on storage planning is to allocate 20kB of flash-based storage per input record. This equates to approximately 1TB of storage per 50M records.

20kB is the baseline estimate, but it can increase if your data is very feature-rich and varied. We suggest that you base your storage requirements on an analysis after loading a large sample data source.

You will also need to account for additional system software, logs, source record files, etc. This can be placed on general-purpose storage.

Scaling Systems
A deployment of Senzing includes both a database server(s) and Senzing, the single node deployments above are indicative of having both the database and Senzing installed on a single machine. In addition to hardware configurations that can achieve significantly faster performance and handle even larger data sets - into the billions - Senzing can scale both the database and Senzing horizontally.

Cloud Environment Considerations
- Senzing performance is sensitive to latency and IOPS.
- Co-located: If your systems are sitting far away from each other in the data centre, or in different data centres, the network latencies are going to be much higher than if they are co-located in the same switch.
- Local Flash on the database: A single locally attached NVMe will achieve more than 100k IOPS where a remote SAN may only achieve 2k IOPS. In cloud environments, be particularly aware of the random read/write IOPS capabilities on your database modes.
- Burstable and tier limits: note, especially with IO systems, any burstable or tier limits that may apply to resources you provision.

Real time operations considerations:

To perform real-time entity resolution, you must read from the database more than you write to it.
-Flash storage is much faster than traditional spinning disks and has become so affordable they are becoming standard equipment. You can still use spinning disks, but it may take 10x longer to load your data.

Performance Considerations
The performance expectations above are based upon typical person or company data sets such as master customer lists, prospect lists, employee lists, watch lists, national registries, etc.

You can run into data sets that have extra-large records or highly related data – meaning everybody is related to everybody else. You can increase performance by adding more cores and RAM at 6GB per core. Regarding sizing per thread, you will need 1.5GB RAM per thread with a minimum of 6GB per node.

Offboarding: data removal

All data is stored in non-proprietary formats in an Elastic Search index. This data will be stored in the Client's Cloud and will always be available. At point of contract termination or for any other reason (e.g. specific policy driven schedules), data will be removed from all stores and from the Senzing engine itself.

Security

All data is secured as per Buyer's cloud in conformity with any associated security policies. We can provide resources with SC and DV clearances.