Edge Intelligence Platform Overview

Preamble

This document provides an overview of the capabilities of the Edge Intelligence Platform.

Distributed Data Lake

An Edge Intelligence network provides a geographically distributed data lake and comprises of multiple edge servers in different locations all operating, managed and accessed as a single logical database. Users connect to a network and perform queries as though all of the data were in one central database.

Each edge server gathers and stores data created locally and allows the data to remain at source to avoid the restrictions and technical limitations imposed by shipping bulk data over public networks to a single central database. For example, network bandwidth may limit the volume of data that can be shipped; or the data may be too sensitive to be moved off-premise; or there may be geo-political reasons why the data cannot cross borders.

A network can be used to implement a hybrid cloud for data storage and analytics where some of the servers are on-premises and others are hosted in cloud data centres.

Network Topology

Each edge server is a leaf node in a tree of logical nodes where the depth and shape of the tree can be arranged as required. For example, edge servers may be distributed over various global locations and these can be grouped under nodes that represent different continents, countries or regions in any arbitrary hierarchy. Thus the topology of a network is used to reflect logical and/or spatial relationships between the data at the edges of the network and the shape and composition of that network is managed centrally and can be dynamically changed, such that any changes to the network shape take immediate effect.

At query time, queries can be submitted in the context of any logical node in the network - such as the root node of the whole network; or a node that denotes a particular region of the network; or a node that denotes an edge point. Hence users are able to perform analysis across the network as a whole or within a chosen area of the network.

Dynamic Schemas

Each network contains one or more schemas, where each schema is a collection of database objects such as tables, views and functions. All objects are managed centrally and can be dynamically changed, where any changes take immediate effect. Hence, schemas are agile and can be changed at will to meet the demands of evolving requirements as and when they occur. There is no need to "see around the corner" or make decisions with long-term consequences.

Even though a schema is relational, it is agile and dynamic enough to be considered almost schema-less.

Data Collection

Each edge server in a network can acquire messages from local data sources such as gateways, devices, probes or brokers and may be received in a variety of formats such as JSON with messages recorded as rows in tables and immediately available to SQL queries.

The parsing of messages and the mapping of message content to relational tables is defined and managed centrally and can be dynamically changed. For example, the values associated with specific keys in a JSON message can be mapped to specific columns in a table.

Data which is universally common across a network, such as dimension or reference data is defined and managed centrally too. For example, rows in a dimension table can be inserted, updated or deleted as required and those changes become immediately effective for queries performed in the network.

Managing Change

All definition and management operations in a network are performed centrally and any changes made are guaranteed to be both durable and to be consistent across the network.

Hence, each change only needs to be performed once and the network ensures that the change is universally applied on every server. Even though one or more edge servers may be unavailable at the time of a change, all changes immediately appear universally consistent to users connected to the network.

Data which is common across a network, such as dimension and reference data inevitably changes with time. The definition of a dimension today may be different from many months ago but when the dimension data is combined with event data we want the dimension data to be applied as it was at the time of the events. Hence the Edge Intelligence platform provides effective dating of universal data, so that changes can be applied from a given point in time and the history of those changes is retained indefinitely - so that queries can apply temporal conditions. Note that the use of temporal data is optional, such that changes can be applied across all time, where required; and corrections to invalid historic data can also be applied, if required.

Discrete Systems

A single network can contain multiple schemas where each schema contains a separate set of database objects and data. Discrete schemas allow the sharing of a common network topology among different user communities where users can be granted access to specific schemas and specific database objects within those schemas; and certain users can be granted sufficient privileges to perform queries across multiple schemas, if required.

A single installation of Edge Intelligence can also manage and operate multiple discrete networks where each network is wholly separate and has its own network topology and database schemas. In particular, separate networks cannot share data and are configured and managed entirely independently. Multiple independent networks are useful where disparate systems or different customers require separate infrastructure under a single installation.

Availability

The Edge Intelligence platform automatically manages replication, synchronisation and the selection of available resources. There is no need to configure master/slave operations or any need to manage fail-over or recovery operations. Meta data and universal data is automatically propagated around the network and message data is automatically replicated between co-located edge servers using multi-master synchronisation; while queries automatically choose relevant edge servers from those currently available.

The network topology and database object definitions provide sufficient information for the Edge Intelligence platform to coordinate activities for high-availability so that there is no need to manage it.

A query will only touch those edge servers that fall within the regions of the network being queried and will automatically choose the relevant edge servers from those currently reachable.

Given that not all edge servers can be completely up-to-date at all times, goals can be set for a query to prioritise server selection in order to optimise a goal. For example, it may be more important for a query to return the latest data or the most complete data.

The goals that can be set for a query are:

- Fastest response time (minimal latency)
- Highest availability (maximum success)
- Most recent data (choose servers with least stale data)
- Most complete data (choose servers with most data)
- Balanced workload (choose least busy servers)

Operations

Changes are all performed centrally and are automatically propagated to all servers in the network. Hence it is never necessary to apply the same change in multiple places. Changes are guaranteed to be durable, regardless of the availability of edge servers at the time of the change and immediately appear consistent across the network to all users of the network. Hence it is never necessary to apply the same change more than once.

All database objects can be dynamically changed, where any changes take immediate effect. Hence, schemas and objects are agile and can be changed at will to meet the demands of evolving requirements as and when they occur.

In conventional relational and non-relational systems, objects such as indexes, partitions and shards are the root of performance issues - because one or more structures are either missing or incorrect; and changes to these structures are notoriously difficult to implement because of the resources and time consumed by such changes. Edge Intelligence entirely removes any need to design, deploy or manage these structures and this is crucially important for dynamic change - especially when managing petabytes of data distributed across a population of widely geographically distributed servers.

Both structured and semi-structured data rich in textual content can be stored with fast text search available on all textual content, without the need to create or manage any text search indexing.

The depth and shape of a network can also be dynamically changed, with all changes to the network structure taking immediate effect. Hence, the network can always reflect the current perception of geographies and regions and how geo-distributed data is categorised or located.

Data which is universal across the network, such as dimension or reference data, is managed centrally via operations which permit rows to be inserted, updated, deleted or loaded from files in a variety of record formats. These operations are performed as ACID transactions allowing changes to multiple rows to happen atomically and durably and constraints can be created to enforce primary key, unique and foreign key integrity rules on universal data.

Message data is fact or event data and is treated as immutable, so that message data cannot be updated or deleted. This guarantees that message data maintains evidential quality and accurately reflects the data as it was at time of receipt.

Data acquired from messages arriving at network speed can become extremely voluminous when collected and retained for several weeks, months or years. Hence, Edge Intelligence allows message data which falls beyond a useful retention window to be automatically retired to reclaim storage space.

Each server operating at the edge of a network runs in a lights-out environment where there is no direct access to the server and where it may be deployed outside the convenience and security of a data centre. These edge servers run autonomously without supervision while collecting data in real-time and responding to any arbitrary queries submitted from the network centre without any prior knowledge of those queries. These queries can be very diverse. For example, a query reaching an edge point may require the edge server to aggregate a handful of columns across billions of rows or drill down and retrieve the

majority of columns from a handful of rows. Whatever the query, the server has to respond in reasonable time without any prior design to achieve that response time.

Moreover, messages being collected at edge point may be arriving at network speeds reaching many hundreds of thousands of messages per second from multiple gateways, devices or probes. These messages are materialised as rows in real-time and made available for queries.

More than one server may be operating at an edge point to provide resilience and high-availability for data collection and queries. Multiple servers at an edge point will collect data from the same data sources; and the servers will automatically synchronise their message data to provide a comprehensive set of messages on every server at the edge point.

Security

User authentication and access control information is held centrally and contained wholly within a secure data centre environment and never reaches the edge of the network.

Users can connect to a central server using these authentication methods:

- Password
- GSSAPI
- SSPI
- Kerberos
- LDAP
- RADIUS
- PAM
- SSL Certificate

Individual users can be allowed/denied login access as and when required.

Uses can be granted zero, one or more roles to permit them particular management abilities and/or query access to particular database objects and/or regions of the network and roles can be created to reflect particular query access profiles where a role can be granted query access to particular schemas, tables, columns and nodes in the network where query access can be granted or revoked from roles as and when required.

Management roles permit users granted those roles to perform a specific range of operations, such as:

- Manage users and roles
- Create and drop networks
- Grant/revoke roles to users
- Manage network topology
- Manage database objects
- Manage data

SSH access to an edge server is solely limited to the root user with a secure and secret password and no other access is permitted. This SSH access is only used during deployment and upgrade operations.

Communication with an edge server is restricted to only servers within the network which hold valid SSL certificate pairs, so that any direct connection to an edge server from an external system or client is prohibited. The creation and management of these SSL certificates is an automated part of the installation and server deployment process, so that there is no need to manage certificates manually or directly - deploying a new server will generate and configure the required certification. Certificates for an entire deployment can easily be refreshed periodically and as needed to maintain a high level of security.

Communication between servers is routinely encrypted to avoid any exposure of data passing over public networks; and disks on edge servers are typically encrypted to prevent physical theft of data, where a server is deployed within an insecure environment.

Particularly sensitive data, such as customer address information remains wholly contained within a secure data centre and never reaches the edge of the network.

Performance

The universally excellent performance of queries permits interactive exploration, analysis and root cause investigation of anomalies. The Edge Intelligence platform provides performance comparable to the best-in-class database technology for that class of query. For example, exploration can begin with a summary aggregation then follow with a forensic drill down into the detail behind a summary category and in both cases Edge Intelligence will provide excellent response times.

Crucially, this performance is delivered without any need to design for it or tune for it.

Fast text search is provided as standard on all text columns without any need to create search indexes in advance. Text searches can be performed in queries using logical search expressions which can include negation, conjunctive and disjunctive operators.

A single edge server can readily consume messages at sustained rates of hundreds of thousands of messages per second from probes, gateways, devices and brokers.

Across a network of tens or hundreds of servers, the overall message consumption rate can easily scale to many millions of messages per second.

Each edge server can handle tens to hundreds of terabytes of data - providing a retention window of granular message data that stretches to months or years; and with tens to hundreds of servers in a network, a single network can easily expand to petabyte scale.

Interfaces

Queries can be performed using standard ANSI SQL and the Edge Intelligence platform supports ODBC and JDBC protocols using the standard PostgreSQL driver so that most standard query clients and BI tools will operate with the Edge Intelligence platform.

Messages received from sources, such as gateways and devices, can be parsed and mapped to rows in a table and these messages can arrive in a variety of formats, such as JSON. The configuration for parsing and mapping messages to a table is managed centrally, such that a configuration for a particular device type can be defined once and used by many distributed agents at the edge of the network to process messages from their respective devices.

Messages are handled by a simple API that accepts a message as a text string and the API will parse and map the content of those messages according to the configuration in force. The API will generate one row for each message processed. Typically, this API is embedded in some code that runs at the edge to connect to message broker(s) to receive and handle messages from devices.

A standard agent utility is provided for injecting messages received through standard input and a standard file streamer utility is provided for discovering files and streaming their content.

Data can be read from files in a variety of formats that include delimited formats (such as CSV and TSV) and fixed width formats.

The file reader configuration provides complete control over aspects such as:

- Delimiter characters
- Quotation marking
- Escape characters
- Field widths
- Data formatting
- Record skipping, filtering and checking
- Field to column mapping