

SUBSCRIBER ANALYTICS

TOP 10 CONSIDERATIONS WHEN EVALUATING A SOLUTION

Analytics solutions able to provide an in-depth understanding of subscriber activity can be one of the most strategic competitive advantages for communication service providers. The ability to analyze and retain granular details of all network activity, by subscriber, allows for improvement to all aspects of the customer lifecycle – such as new customer acquisition, enforcing tiered usage plans, self-service support, monetizing OTT services, proactive upgrades and anticipating churn. In addition, subscriber analytics drives down network costs by enabling wise decisions about how to best deploy new capital and to optimize utilization of existing resources throughout the network.

Not all analytics solution providers are the same. Analytics required by network operators continue to push the envelope in-terms of the architecture and scale needed to provide instant insight for millions of subscribers and rapidly increasing content consumption. To help you evaluate analytics solutions providers and maximize customer revenue and operating profit, we suggest you consider asking the following:

1. FULLY RELATIONAL SQL

Does the provider offer the ease of SQL with industry standard tools/interfaces?

Edge Intelligence – the architecture is based on a fully relational model, using standard SQL. This makes it easy to combine with reference data and access all of the underlying data on your own accord. Industry standard JDBC/ODBC interfaces compatible with standard developer and BI tools allow you to use your tools of choice, while leveraging existing knowledge and skills. There are no limitations to accessing data stored within our system – issue any query at any time. The system also integrates seamlessly into existing developer workflows, real-time dashboards and internal/customer portals.

2. DATABASE ARCHITECTURE

Does the provider provide flexibility for exploratory and highly iterative analytics?

Edge Intelligence - analytics has to be exploratory and highly iterative – because it is not possible to anticipate how data needs to be analyzed in the future. Our patented, database architecture reinvents the way data is stored and accessed by pre-optimizing all requests for performance. In doing so, it is future proof for the analytics you require today, but also might need later – including a combination of broad aggregate analytics, needle in the haystack forensics, standard reporting, ad hoc discovery, etc. Fast performance across all query types is achieved in a way that can't be achieved using Hadoop or databases optimized for either column-store or row-store format.

3. HARDWARE REQUIREMENTS

What kind of hardware is required, and how cost effective is it when deployed at-scale?

Edge Intelligence – no special hardware is required as the architecture leverages modern hardware characteristics. Standard servers using general purpose CPU, DDR memory, and HDD disk is what's used to achieve network scale, in real-time, for millions of subscribers and billions of daily new records. Software can be deployed within physical, virtual or container-based environments. Commodity resources and scalability associated with our solution typically help reduce overall hardware cost by at least 50% when compared to alternative approaches.

4. DATA RETENTION TIME

How long can data be retained for granular analysis and compliance/forensics?

Edge Intelligence – there is no limitation on how long data can be retained. Some of our active deployments retain data for several years including one with more than a trillion records. The architecture decouples storage and compute resources, making it possible to add additional storage, cost effectively, without impacting the operations of the system. Sampling frequency can also be increased without having to compromise data retention time. Data can be automatically deleted from the system upon a programmable aging criterion.

5. PRIVACY & SECURITY

Does the provider require access to the network operators personally identifiable data?

Edge Intelligence – the system doesn't require access to the network provider's data. All operating aspects including software updates, monitoring, etc. never requires having access or visibility to the underlying subscriber data. This is in sharp contrast to other architectures which require full-time access to both the software and underlying data. Encryption is always applied to all in-flight and at-rest data and SSL/TLS is applied for all server-to-server communications across the system to ensure security.

6. SCALABLE FOR GROWTH

How scalable is the architecture to accommodate subscriber and content growth?

Edge Intelligence – the architecture has been developed in a way that gracefully scales for millions of subscribers, with the ability to input, correlate and analyze data in real-time. A highly resilient, messaging & file-based queuing scheme immediately persists any collected data and correlates across all sources of data in real-time. The system scales upward to support millions of NetFlow or IPFIX events/second, thousands of DHCP/Radius messages/second and millions of DNS responses/second. Petabytes of data can be stored in aggregate. All the while being able to respond to any query, including ad-hoc and forensic, in a matter of seconds.

7. DEPLOYMENT FLEXIBILITY

Is the system designed to work equally well when deployed either on-premise or in the cloud?

Edge Intelligence – the system’s overall performance and scale is agnostic to deployment model and applies the same security whether it operates in on-premise, cloud or a hybrid environment model. Due to the massive data volumes associated with subscriber analytics for network operators, an on-premise deployment is the most likely model for the foreseeable future. With an on-premise model, servers can be de-centralized and placed wherever rack space is available and to preserve bandwidth. De-centralization of storage will not impact query performance. Portal for command and control can also be either hosted on-premise or cloud-based depending on your own security preferences.

8. EASE OF OPERATION

How much effort is required for software installation, software upgrades, performance optimization, data synchronization & replication?

Edge Intelligence – none, as the system has been designed for zero-touch operation. Software is installed automatically from a central location to all machines using an Ansible environment. Software upgrades are propagated thru a single centralized command without requiring manual intervention. Data synchronization, replication and certificate management are all done in an autonomous manner. Performance optimization is also fully automated, requiring no manual indexing, partitioning, sharding or other forms of tuning.

9. CENTRALIZED ACCESS

How easy is it to gain access to data sources stored across a variety of locations?

Edge Intelligence – the system allows for easy, centralized access to distributed data sources. This makes it simple to query and analyze data on a per-subscriber, regional and global scale to help unlock new business insights. This is different than legacy systems which are unable to “federate” across different data sources. Without federation, an independent analysis must be conducted within each source location of data followed-by manual work in order to correlate all of the data sources together.

10. FORWARD LOOKING

Can the same platform be used not only for subscriber data, but also for machines/sensors?

Edge Intelligence – the volumes of data generated by machines & sensors will be much larger than those generated by subscribers and require a de-centralized approach to analytics. The platform can be used for data analytics generated by machines/sensors including real-time SQL streaming and machine learning. This is accomplished through distributed data collection and retention, at the edge close to where the data is generated, making the platform ideal for new revenue generating services in emerging areas such as IoT & edge computing services.