EdgeIntelligence

Network Performance Use Case

December 2016

Executive Summary

The performance and features of Edge Intelligence enable the edge collection and the retention of massive volumes of probe data combined with the centralised collection and retention of performance counter data from telecommunications networks. Such a platform enables real-time network performance monitoring and dynamic root cause analysis in response to any observed performance anomalies.

This document presents the possibilities for such a solution and highlights the features of Edge Intelligence that make it feasible.

Use Case

Telecommunications network operators routinely gather information from their network infrastructure to measure the performance of the network in general and for specific services and specific regions of the network. However, it is common for network operators to gather only sufficient information to be made aware of failing KPIs yet be unable to accurately diagnose the root cause of a problem because the voluminous detailed operational data lying behind a failed KPI could not be speculatively retained in case of such a failure.

Edge Intelligence now permits the exhaustive collection and retention of detailed operational data from network probes distributed across the network so that the details that lie behind unforeseen issues can be immediately analysed and forensically investigated to identify their root cause.

Crucially, the modest hardware footprint required by Edge Intelligence, allows edge stores to be widely geographically distributed and co-located with network probes to avoid the need to ship terabytes of probe data to a central site for data analysis. Yet, the population of distributed edge stores can be managed and queried as a single logical database – to provide central performance monitoring, root-cause analysis and alerting in real-time.

Probe Data

Probe data is extremely voluminous and network operators are often forced to filter the data prior to collection to avoid overburdening both the network and the database collecting the probe data for analysis.

Edge Intelligence is able to acquire probe data at network speed and retain many terabytes of that data on a single commodity server – enabling the localised collection of probe data at source and avoiding the need to ship voluminous probe data over the network. An Edge Intelligence server can be located next to each probe in network and can operate autonomously while performing agile analytics and forensic investigative queries as part of a network of servers spread across the network.

A key feature of Edge Intelligence is its ability to achieve excellent performance from a small number of high-capacity high-latency disks (8TB HDDs for example), which allow a single server to readily manage many terabytes of probe data. Trials have demonstrated the ability of a single small server

to acquire probe data at rates of hundreds of thousands of messages per second while retaining that data for several days, weeks or months for fast exploratory root cause analysis of network performance problems and customer experience analysis.

Counter Data

Performance counter data is less voluminous than probe data and can be collected centrally from network OSS systems. However, counter data is extremely wide and may contain thousands of counters for a single measurement event. KPIs need to be evaluated in real-time against this data as it arrives; while subsequent unplanned investigatory and historical analytical queries need to be agile and responsive regardless of how they are phrased.

Trials have demonstrated the ability of an Edge Intelligence server to process millions of counter measurements per second while providing sub-second KPI evaluation and interactive response to ad-hoc analytical queries.

Moreover, the ability to retain the counter data for long periods of time allow new KPIs to be retrospectively applied and investigated.

Query Agility

A key characteristic of Edge Intelligence is its ability to rapidly respond to arbitrary queries without requiring any prior design for performance. Tables are defined as standard relational tables, but there is no need to design or implement any indexing or partitioning structures to achieve good performance – regardless of the nature of the queries submitted - from granular forensic queries to extensive analytical queries.

This characteristic removes any need to tune or optimise an edge server and allows the server to operate autonomously. This is important for edge stores operating at the edge of a network where a widely distributed deployment makes close supervision and regular intervention in the operation of each and every server impractical.

The agile nature of the solution permits all of the following:

- Drill into a KPI profile for particular components over a given period of time.
- Drill into the probe message detail related to a specific KPI anomaly.
- Verify the KPI anomaly and identify any network subscribers affected by it.
- Retrospectively apply new KPIs to historic data.

Journaling Data

KPIs provide pre-defined queries which need to be applied and updated as new data arrives and need to return results in sub-second time for real-time monitoring of the network.

The journaling feature of Edge Intelligence allows the results of queries to be materialised and dynamically updated as new data arrives. Journaling can be used to maintain rolling aggregates and/or alert to anomalous conditions while retaining those results for arbitrarily long period of time.

For example, the results of hourly aggregations can be retained for days, weeks, months or years and queried for analysis or additional aggregation over broader time intervals.

Purging Data

When data is arriving at network speed, it important to be able to remove obsolete or aged data as quickly as new data arrives. Edge Intelligence allows data to be retained in a rolling time window without having to explicitly define and manage time partitions.

Conclusion

Edge Intelligence provides a range of unique features which make it ideally suited to network performance management.

In particular, the edge servers can be geographically distributed and co-located with probes for the edge collection and analysis of massive volumes of probe data to avoid any unnecessary data filtering or back-hauling of voluminous data across the network. When combined with the centralised collection and retention of performance counter data, this can provide a real-time picture of network performance KPIs and allow immediate investigation of the root cause of any performance anomalies observed.