Understanding eoMind & Real-Time Streaming Analytics Insights

Introduction - The Automation of Insight

One of the largest challenges in telecoms today is making sense of the huge amounts of data available to management, operations, marketing, engineering and customer care in order to understand the subscribers' Quality of Experience (QoE). There is simply too much data making it almost impossible to actually identify the issues in a short period of time.

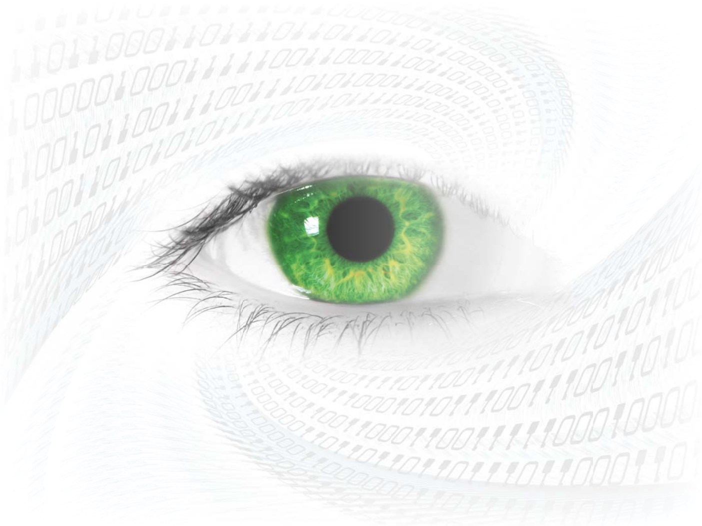

Big data in telecoms is characterized by the large volumes of data (often many billions of event records per day are generated) and the many varieties of that data (the different interfaces, network and application technologies that are monitored). Data can be extracted from network elements, probes, sensors, log files and even from social networks.

For a variety of reasons - mostly a lack of time and resources - this priceless data is often left lying dormant; opportunities for improved service, cost, and customer retention are lost. Allowing this data to sit idle is clearly not an effective use of a CSP's information asset. Identifying the insights inside the data streams is therefore key to identifying issues that affect your customers, and also to respond to threats and opportunities for service delivery.

A continuous process of extraction of information is required to maximize the investment in data sources that CSPs have already made. Let's look at how this data has been exploited up until now.

The Extraction of Insight - Traditional Methods

For a long time, CSPs have been aware that their raw data assets can be analyzed in order to derive:

Corrective Action Insights

- To take corrective actions, quickly
- To react to threats (to the network or the customer)

Decision-based Insights

- Take improve qualitative decision making - such as planning, forecasting and provisioning
- To improve quality of the understanding of services, customers, devices, networks and applications
- To react to opportunities - such as offers and upsells

Because the volume of data is so high - and due to the complexity of the content of the data - analysis has been difficult and costly, leading to a frustrating execution gap between what you want to achieve with the data, and the tools and technologies sitting between you and that goal.

First generation BI systems involved a complicated and cumbersome process of extracting the source data, enriching it, loading it into a database and then presenting this information in some form of reporting tool at the top layer. A specialist would use this set of tools to derive insights.

This often led to data warehouses being used for a very limited number of functions, and often serving a limited number of users in the organization (a limited set of experts who understood and had access to the reporting tools on top). The original aim for data warehouses to provide democratic access to the insights in the data for all by serving multiple departments often never materialized. There is, in fact, too much data and too many types of data to realistically achieve this by manual techniques.

First generation BI systems employed an architecture of extract-transform-load: here data was pulled off source systems, transformed or mediated in some way and then loaded into a database. Once in a database another layer was added on top in order to provide end users with an ability to execute reports on the data.

This approach has the following drawbacks:

- One or more specialists are needed to continually be on hand to analyze the data
- A large number of reports/data views are needed in order to cover all scenarios where insight might be derived. This often leads to a reporting treadmill – one report after another is created, used for a single purpose, then discarded, and so the cycle continues with limited insights coming to light
- Data moves between the layers in the processing are required – this is therefore costly in terms of CPU, LAN, WAN and storage. It is above all else too slow, leading to delays in identifying the issues as they happen (often called a “too-late” architecture)

The technology shift that allowed us to get around these problems came from some of the developments in the world of big data and the general need to analyze large volumes of data quickly and more importantly automatically.

Take the example of real-time advertising on social network sites: recommendations need to be formed, bid-on and communicated to users in milliseconds. Moving data around in the old fashioned way is not going to work. Automatic trading algorithms in the banking sector cannot wait for data to be analyzed by a specialist before a decision can be made: it has to happen without any delay. It is this world of “split-second decision making” that largely drove the development of automatic streaming analytics.

Automated Streaming Analytics - A New Approach to Mining your Data Assets in Real-Time

What CSPs need now is the combination of extremely low latency insight discoveries (within milliseconds of an event taking place) and a cost effective way of analyzing this data automatically without human intervention.

This gives rise to the need for what we call "self-operating analytics"; analytics that run in an unattended fashion, analyzing the data for you continuously and alerting you to these key insights.

The advantages of self-operating analytics are:

- No specialists need to be on hand to analyze the data, as streaming algorithms do this for you
- There is no need for reports to be interpreted - just the results of the analytics
- The need for data to be moved multiple times or stored on a costly file system or database is removed: it is in contrast with batch data warehousing an architecture built around speed

It is this philosophy of moving the analysis towards the data, and not the other way around, that embodies much of the thinking behind big data and streaming self-operating analytics in particular, and enables low latency insights at lower cost and higher utility than previous technologies. Anritsu has adopted this approach for the development of the eoMind Streaming Analytics platform.

eoMind Real-Time Streaming Analytics

What is eoMind?

eoMind combines a series of algorithms that automatically inform you of the impact of problems and the quality of experience of your subscribers: when things go wrong, or when unusual events take place, in true millisecond real-time. It communicates these findings to you in natural language on a social network like wall – and allows you to interact with the algorithms to fine-tune and personalize the output to your needs. It is backed up by dynamic visualizations that are created as supporting case notes on each finding.

eoMind takes data from your own data sources; your own probes (not only Anritsu probes), network elements, OSS/BSS systems and from social networks. The eoMind platform runs a series of pre-defined learning algorithms that analyze the streams of data in real-time.

These algorithms are designed by Anritsu to give you the kind of insight you would get if you employed a team of experts to continuously monitor every possible data source. The difference is that these algorithms analyze all the data, all the time, in real-time.

eoMind not only flags up issues for you, but it also makes recommendations as to what steps should be taken next. This combination of detect and recommend is at the heart of eoMind's algorithms.

What does eoMind show?

eoMind publishes its findings to a personalized wall that allows you to subscribe to the areas of interest that are important to you: voice, data, 2G, 3G, 4G, VoIP, for example.

You interact with these posts in the same way you would with postings from a social network. You can:

- Unfollow the conversation (show no more insights of this type)
- Only show when scale exceeds X (I'm still interested but don't show me anything unless, for instance a certain number of subscribers are impacted by this insight)
- Analyze further the data using eoSight multidimensional reporting tool
- Share (with colleagues – adding comments)
- View time-series history of the event and the size of the impact to subscribers

In this way you can interact and fine-tune the types of insight eoMind posts to your wall. This kind of feedback loop is key to giving eoMind's algorithms the information they need to ensure that what you see on your wall is always of interest and, importantly, correct.

eoMind's Automatic Streaming Architecture

Anritsu's eoMind solution is a genuine big data processing platform for telecoms data. It performs as much processing of data as possible in memory, without the need to store it on a (slow) file-system. It is a streaming analytics platform, which is why eoMind can process millions or rows per second on a very small computing platform compared to the traditional method of passing data via mediation into database storage.

eoMind data processing is highly scalable and runs on a virtualized single-node or distributed cluster.

As more and more algorithms are added, eoMind can deploy these modules on to the cluster. We have developed an expandable in-memory computation platform that can grow as algorithms are added over time.

Nowhere in the processing steps is the data stored in a database or moved onto a files system. It is this scalable in-memory processing architecture that makes eoMind unique.

Anritsu's eoMind SIRCA Algorithm

A highlight algorithm we are delivering on the eoMind platform is SIRCA (Subscriber Impact and Root Cause Analysis).

This algorithm attempts to automatically answer the following questions:

- How do I detect issues that really affect customers quickly (within milliseconds)?
- How do I localize the cause of issues that affect customers quickly?
- How do I do this across every network technology and all the data all the time?
- How do I do this without setting up complex thresholds and other baseline analyses?

In short, SIRCA looks for negative event clouds – issues that are affecting a certain amount of users and then tries to look for common 'co-occurrence' features in the source data. We will show you some examples of these later on.

When a cluster of users where a problem is taking place is detected, we then try to automatically identify the common factors in the event. This is something that can take a human 15 minutes of slicing and dicing in a suitably powerful business intelligence tool. In eoMind's case, SIRCA is performing this type of analysis for you continuously on all your data sources.

This type of subscriber impact analysis is becoming very popular within CSPs. The idea is to identify problems based on the number of subscribers impacted by the problem, as opposed to identifying them based on the number of incidents of the problem.

This is a huge step forward, as this type of analysis was previously too computationally intensive to be carried out inside databases. Since eoMind uses streaming analytics, we can perform this type of analysis all the time and in real-time.

This new detect and recommend workflow is driven by a need to move beyond first generation customer experience management solutions, where every single subscriber had a set of low-level metrics calculated for them and stored all the time, and a set of alarms raised based on thresholds. This type of heavy-iron approach to instrumenting the customers' experience was not scalable and required extremely large (and costly) databases.

SIRCA Output Examples

Some examples of the SIRCA algorithm output are shown below. Our aim is to make the output of the algorithms completely human-readable. When you see this posted on your wall, you should immediately be able to understand what SIRCA is trying to say and then decide whether you like the post and want to take a further action, or whether you want to suppress it – to unfollow the thread.

In Figure 1 in page 6, we can see that there are two findings SIRCA has posted to a user's wall. The first one shows an issue with MAP (roaming) interconnect records, where a spike of 1.148 individual subscribers are being affected by a roaming error. In the second one, you can see SIRCA has identified the common co-occurrence features (the root causes it thinks might be). You can see the trends for each possible cause, and also a list of them. By looking at these you can quickly tell what is likely to be related to the cause of the problem.

MAP / Outbound

At 16:28 there are **1148** users affected by MAP Error Code '**34 System Failure**'.
Issue has been happening for 31 minutes, affecting a total of **9372** users, and peak was at 16:23 with **1198** affected users.
The issue is primarily related to the following causes:

- *Linkset Name:* **3_BRSTP1-ODSTP1** (59% ▲59%) **3_OGSTP2-ODSTP1** (57% ▲57%) **3_OGSTP2-OOSTP1** (29% ▲29%) **3_BRSTP1-OOSTP1** (25% ▲25%)
- *Originating Point Code:* **NI3 1283** (85% ▲85%) **NI3 1284** (43% ▲43%)
- *MAP Operation Code:* **2 updateLocation** (74% ▲74%) **23 updateGprsLocation** (57% ▲57%)

MAP / Outbound

At 16:28 there are **645** users affected by MAP Error Code '**1 Unknown Subscriber**'.
Issue has been happening for 31 minutes, affecting a total of **2634** users, and peak was at 16:06 with **741** affected users.
The issue is primarily related to the following causes:

- *Linkset Name:* **3_OGSTP2-ODSTP1** (41% ▲41%) **3_BRSTP1-ODSTP1** (40% ▲40%) **3_OGSTP2-OOSTP1** (27% ▲27%) **3_BRSTP1-OOSTP1** (27% ▲27%)
- *Originating Point Code:* **NI3 1283** (69% ▲69%) **NI3 1284** (48% ▲48%)
- *MAP Operation Code:* **56 sendAuthenticationInfo** (100% ▲100%)

Figure 1: SIRCA Output Example I

In Figure 2 below, we can see users affected by LTE congestion issues.

IUPS / Retainability

From 13:54 to 13:57 there have been **2376** users affected by Iu Release Cause '**14 Failure in the Radio Interface Procedure**'.
Peak was at 13:56 with **1711** (▲2.3x) affected users.
The issue was primarily related to the following causes:

- *Last RAI:* **GERAN/UTRAN RAC: 801-1** (10% ▲6%) **GERAN/UTRAN RAC: 501-1** (10% ▲6%) **GERAN/UTRAN RAC: 201-1** (9% ▲6%) **GERAN/UTRAN RAC: 801-1** (8% ▲4%) **GERAN/UTRAN RAC: 601-1** (7% ▲4%)
- *Last Serving RNC:* **2** (9% ▲6%) **0** (9% ▲6%) **2** (9% ▲6%) **1** (8% ▲4%) **6** (7% ▲4%)

IUPS / Retainability

From 13:54 to 13:57 there have been **6351** users affected by Iu Release Cause '**31 No remaining RAB**'.
Peak was at 13:56 with **4487** (▲2.4x) affected users.
The issue was primarily related to the following causes:

- *Access Point Name:* **tele** (20% ▲12%) **mart** (7% ▲4%) **apple** (5% ▲3%)
- *Last RAI:* **GERAN/UTRAN RAC: 201-1** (8% ▲5%) **GERAN/UTRAN RAC: 801-1** (8% ▲5%) **GERAN/UTRAN RAC: 001-1** (7% ▲5%) **GERAN/UTRAN RAC: 501-1** (7% ▲5%) **GERAN/UTRAN RAC: 801-1** (7% ▲4%)
- *Last Serving RNC:* **18** (7% ▲4%) **20** (6% ▲4%) **02** (6% ▲4%) **05** (6% ▲4%) **8** (5% ▲3%)

Figure 2: SIRCA Output Example II

You have control over the posts on your wall – this is so that your wall only contains the important information for you. When you see a post on your wall, the actions you can take are:

- **Unfollow this thread:** Shows no more instances of this problem as it gets worse – this teaches SIRCA that the problem isn't in fact a problem
- **Unfollow All:** Unsubscribe from this instance of SIRCA, e.g. "show me no more MAP errors, I'm not interested in roaming"
- **Share:** Allows you to annotate and share to another colleague e.g. "Ignore this, it was a network intervention/ upgrade we carried out last night at midnight"
- **Analyze:** Allows you to analyze in eoSight, our data visualization and exploration platform

Conclusion - Anritsu's Path to Insight

It is obvious that if CSPs are to exploit the value of data from probes, networks and indeed existing OSS/BSS systems, then manual data discovery techniques alone will not be enough. The need for split-second decision making and real-time detection of anomalies requires an automatic data processing architecture.

The too-late architecture of moving data around multiple times needs to be replaced by a new streaming analytics approach. Anritsu has responded to this with a platform for the execution of streaming algorithms and for a series of pre-packaged algorithms that are designed to bring this information to light for you, automatically.

CSPs expect the expertise that has been built up over the years to be encapsulated inside these algorithms. In the years to come, as we move towards automation and network auto-correction, these algorithms will play an increasing part in reducing the effort required by us all to deploy, monitor and react to events on these complex systems.

The combination of eoMind and eoSight allows you to master to data inside your network.

About Anritsu Service Assurance

The Anritsu Multi-Dimensional Service Assurance suite for customer experience analytics provides complete solutions for existing and next-generation wireline and wireless communication systems and service providers. Anritsu's Analytics platforms consist of eoSight, a big data platform and next generation visualization solution and eoMind, a machine learning solution consisting of a series of packaged up algorithms that self-operate on the data for you, with each algorithm being aimed at solving a particular business problem. Our low cost storage combined with class leading visualization and analytics means that you can make the most of your telecommunications data asset.

Anritsu Corporation (www.anritsu.com) has been a provider of innovative communications solutions for more than 120 years. Anritsu sells in over 90 countries worldwide with approximately 4,000 employees. For more information, please contact us at info@anritsu.com.

Anritsu A/S
Service Assurance Division
Kay Fiskers Plads 9
2300 Copenhagen S
Denmark

Tel: +45 72 11 22 00
Fax: +45 72 11 22 10
E-mail: info@anritsu.com
Japan +81-46-223-1111
USA +1-972-644-1777

Anritsu
envision : ensure