By Shen Jie
Chief Architect of Huawei NFV Capability Center

Algorithms for intelligent and automated O&M

O&M may have come a long way with advances in scripting, automation, and legacy bug fixes, but having fewer problems hasn't diminished the burden of on-site maintenance on cloud core networks. In fact, many problems on today's live networks are entirely unexpected and leave O&M personnel at a complete loss with how to deal with them.

However, Huawei has developed algorithms to solve system change faults and routine silent faults in the cloud core network. Using intelligent O&M can solve live network problems and thus lower costs, boost efficiency, and meet specific KPI requirements.

The two main issues in O&M

The product scope of the cloud-based core network includes all pipes and voice switching hardware that come after wireless access. NEs are many and varied, with complex and diverse interfaces, complicated inter-NE signal interaction, and a huge number of monitored KPIs – in a typical VoLTE system, there can be as many as 30,000 or more KPIs.

Today's two most difficult issues are: (1) hidden dangers after upgrades that cannot be discovered in time to avoid subsequent incidents, and (2) silent faults that occur suddenly and cannot be quickly demarcated. When they aren't rapidly dealt with, the second type can cause major incidents.

The first type of issue primarily occurs during upgrades or changes in operations. Telecoms system upgrades are typically arranged in the early

hours of the morning when usage is low. But as service volumes after upgrades at that time are low, even if an incorrect operation occurs during the upgrade – be it a configuration error or a bug in the upgraded software itself – any system errors triggered won't always be discovered in the system's KPI monitoring.

In a recent case affecting one operator, an authentication problem was accidentally introduced during an HSS upgrade job. After the upgrade was completed, the authentication success rate KPI for the HSS was normal. However, other non-critical indicators, such as for the mobile switching center (MSC), showed that the authentication success rate was gradually decreasing. The anomaly was only identified by a comprehensive analysis of all indicators.

Discovering faults

How can these hidden dangers be discovered? Our analysis shows that two capabilities can help systems discover anomalies: (1) dynamically monitoring and recording the periodic dynamic threshold range of all indicators of the system, and (2) comparing all indicators in real time with the dynamic thresholds at the same point in time to discover indicators with large deviations.

The second type of issue principally occurs during routine monitoring scenarios. It mainly manifests as service anomalies with no alarms and normal KPIs.

In another case affecting an operator, an EOR soft failure led to the network-wide interruption of a VoLTE service for an extended period. Voice services were delivered through Circuit Switched Fallback (CSFB), user perception was low, and the main indicators were normal.

Problems of this type generally involve monitoring the indicators of many NEs and between NEs. The key to solving this type of issue is fast detection and fast data gathering to analyze and locate the problem, so it can be solved during the optimal timeframe.

Our analysis shows that two solutions can help systems quickly demarcate faults: (1) algorithms to compare all indicators with the dynamic thresholds at the same point in time to determine whether indicator fluctuations exceed learned ranges, which is then used for diagnosing whether there is a silent fault, and (2) a complete dataset that can lead to root cause analysis and solving detected faults.

Data processing and machine learning

Given the above, the following system capabilities are general requirements:

- Monitoring all indicators
- Learning and obtaining dynamic thresholds
- Diversifying and sorting indicator deviation calculations
- Correlating indicator deviation and system failures

To build these capabilities, operators need data processing platforms and machine learning (ML) algorithms.

Huawei has built a practical verification system that compares verification systems and on-site outcomes by reviewing past on-site incidents to verify the design and analysis results.

To build these capabilities, operators need data processing platforms and machine learning (ML) algorithms.

Huawei's verification system comprises two parts: a full-scale data processing framework and a single-indicator anomaly detection ML algorithm. The full-scale data processing framework uses real-time stream data processing to simultaneously collect and pre-process a large number of KPIs and metrics. The single-indicator anomaly detection algorithm targets gathered indicators and trains corresponding models based on their historical data. It then performs anomaly detection on the current indicator to quickly detect anomalies.

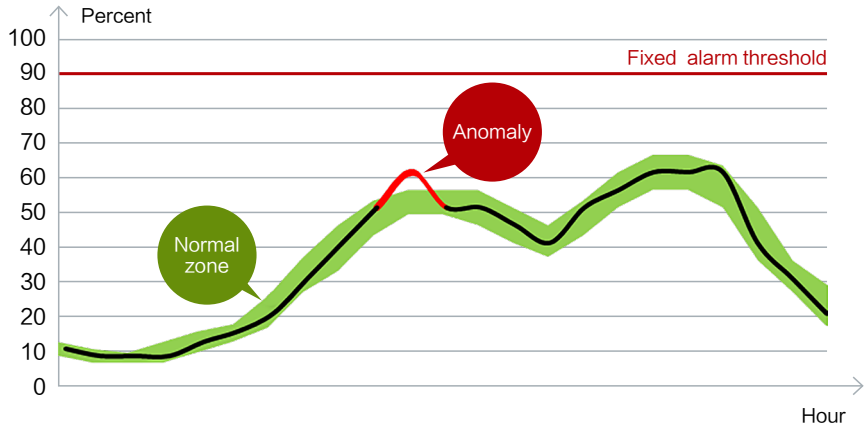

Considering system scalability, the data gathering framework supports multiple collection methods to collect data from different target systems, including logging, metrics, and tracing data. For example, traditional data gathering methods, such as SNMP and FTP, are used for legacy systems to support old NEs or EMS data collection.

For new microservices, mainstream open source data collection methods, such as Prometheus, are used as much as possible. The system is compatible with a range of collection methods to limit changes to the target system and support service system diversity.

For data processing, the system uses various pre-processing approaches due to the variety of data collected and different data extraction methods. To reduce the amount of post-development needed, service indicators are abstracted into generic data models. Collection and pre-processing functions are abstracted to form standalone, atomized functions that can be orchestrated. Doing so forms a service processing chain, which enables codeless service processing functions.

Single-indicator anomaly detection models

These steps facilitate data normalization on different data sources and types, making it possible to perform higher level analysis, which includes the fundamental full-scale anomaly detection approach: a single-indicator anomaly detection algorithm. This method focuses on full-scale data processing, rather than depending on domain knowledge to first perform metric screening. Thus, no potentially anomalous indicators are missed. After several iterations, single-indicator anomaly detection models for each on-site system can be formed.

Anomaly detection is an important part of O&M. In cloud services, which use white-box O&M, anomaly detection forms the basis of early detection, fault prediction, and other subsequent analysis approaches.

In the design of the single-indicator anomaly detection algorithm, the dynamic threshold model of the indicator is built from learning historical indicators. This model is applied in the actual environment to detect live network service data in real time and to determine if the data is anomalous.

The red horizontal line in the graph above represents the traditional static threshold detection method before

Basis of periodicity identification

- Periodicity/quasi-periodic sequence is reflected as a spike in the power spectrum
- If the period is not obvious or if it's a random sequence, the power spectrum is bell-shaped or flat

Fluctuation assessment

- Breaks KPI data down into low-frequency trends and high-frequency noise
- Global fluctuation: variance of a low-frequency tendency over a period of time, for example, a day
- Local fluctuation: amplitude of high-frequency noise

Tendency assessment

- KPI periodicity of usually one day
- The moving average based on one day
- The slope of the curve after the moving average is calculated

the ML algorithm is introduced. The green band is the learned dynamic threshold range. The "Anomaly" indicator cannot be detected by the traditional approach, but it can be accurately detected by the single-indicator detection algorithm.

In the flow process, the algorithm first identifies the periodicity of the indicator, and then selects the appropriate algorithm based on its volatility, tendency, and other characteristics.

One operator has deployed this algorithm model to perform live network PoC testing. Its network contains 50 IMS and 90 EPC NEs, with a total of 1,400 single indicators gathered. After two consecutive months of live network data verification and a relatively stable live network environment, four indicator anomalies were discovered. The operator was satisfied with the results and decided to commercialize the feature.

Single-indicator anomaly detection is the first step towards intelligent O&M. Later, the scope of detection can be expanded to cover more scenarios. Intelligent approaches such as multi-indicator correlation analysis, root cause localization algorithms, and alarm compression intelligent algorithms, covering processes such as fault detection, anomaly location, and root cause analysis, can also be explored. Necessary self-healing methods can also be added. The eventual direction of evolution is network automation and intelligence.

KPIs tend to be global, periodic, and real-time. The single-indicator anomaly detection algorithm is an effective way to quickly identify faults. For changes in operations, major event protection, and daily monitoring, KPI anomaly detection can also be used to quickly identify silent problems, reduce the scope of impact, shorten recovery time, and help operators shift from after-the-fact O&M to full-process monitoring, prediction, and automatic recovery.

KPI anomaly detection is a key means of improving core network reliability and MTTR.