VIAVI Solutions

White Paper

Open Assurance and Analytics

Enabling Data Science and Business Automation

Introduction

Technology is constantly changing and never so much is this evident as today than with telecommunications. Now more than ever, society is relying on the speed, reliability, and flexibility of telecommunications services to enable them to work remotely and communicate reliably in seemingly endless new ways. From enabling telemedicine to automated manufacturing to the sudden shift to remote working and online learning, telecommunications technology is the critical enabler for service providers to addres s the rapidly changing needs of society. Enabling technologies such as virtualization and cloud, network automation, 5G networks with the implementation of CUPS and a service based architecture, new RF frequencies with distinct physics, the movement from a cell-centric to beam-centric environment adding the element of the 3rd dimension for RF coverage, ultra-dense, massively scalable networks for ultra-low latency and high bandwidth services such as AR/VR, Mobile Edge Computing, and Network Slicing are all needed now and operators are responding by deploying these technologies as quickly as possible.

Technology evolutions create significant challenges; however, they also can enable huge opportunities for service providers' businesses. Operators must constantly innovate to take full advantage of these opportunities that new technology creates but do so in a fiscally responsible manner. Partners can assist, but also must keep up with the pace of technology.

Vendors that assist operators to leverage the most out of a given technology turn must provide valuable innovation to allow service providers to maximize the technology opportunity. One of the more challenging areas in which to accomplish this is in the area of telecommunications big data – that is, achieving adequate visibility to assure performance of these new networks while gathering relevant and timely "actionable insights" to leverage the monetization opportunity presented by the exponentially growing data that modern telecommunications networks create.

Innovation Challenges

With respect to 5G, realizing the potential of the network is as much about architecture of the 5G network as it is about advances in telecommunications technology. The introduction and wide adoption of virtualization and cloud has demonstrated the importance, variability, and complexity of leveraging the 5G opportunity in its stand-alone (SA) architecture end-state.

Solutions that provide operators actionable insights and access to data for assurance, monetization, and analytics purposes also need to be architected so the cost of the solution does not outpace the rate of current and future data growth in the networks.

For both service provider networks and the assurance, monetization, and analytics solutions that provide access to these actionable insights and data sources, the lack of standards, or the abundance of standards, in some cases, has led to more variability and complexity for operators and vendors alike. As the assurance and analytics solution migrates into a pure virtual environment in parallel with the evolution of operators to 5G SA networks with its potential hundreds of thousands of network function instances, it is critical to maintain cost effectiveness. Additionally, the importance of selected and/or re-selecting relevant and important data is much higher than ever before.

In many cases, this does not mean the most performance, but the ability to deterministically demonstrate value and performance to the various stakeholders within operators that need insights from **the** mobile data. As the opportunity that V-RAN and O-RAN presents and operators begin to leverage, it is clear that a unified view of RF, RAN, Fronthaul, Midhaul, Backhaul (aka ‘xHaul’), and Core, along with services provided to subscribers is a must for manageability of big data visibility. It also suggests a move away from ‘swivel-chair’ visibility to unified views of different functions in the network, such as RAN, traditional Core, and IMS.

This paper will lay out the aspects of the assurance and analytics solution requirements (heretofore referred to as “assurance solutions”) and describe what a target architecture should strive for in order to achieve the necessary “breathable” characteristics required by the new network paradigms introduced as part of the wholesale migration to NFV/SDN and 5G.

Important Considerations for Assurance

When building or expanding networks by leveraging virtualized infrastructure, the assurance solution can no longer be addressed during later integration phases. It must be designed or at least planned into the solution from day one to ensure that adequate capacity of the infrastructure is allocated to accommodate assurance in the appropriate forms.

This brings the challenge of using the same infrastructure for assurance as the network itself. The tradeoff of complete isolation vs. flexibility and cost savings is a fait accompli. The assurance solution must perform beyond reproach such that network administrators maintain confidence and can rely on provided observations as fact during complex diagnosis.

Also, the need to transition from bespoke assurance solutions and the diversity of network designs made possible by the new technology innovations necessitate early design considerations for the assurance solution to be deployed. It is no longer reasonable to expect there will be an "out of box" bespoke solution available that will meet the needs of these new network architectures that can be selected "after the fact".

During the design phase there are 3 key ODE (Operations, Diagnostics, Exploration) use case categories that will need to be addressed by new assurance systems in this emerging network environment:

- **Operations** – Engineering, Operations, and Optimization
- **Diagnostics** – Diagnosis and Troubleshooting
- **Exploration** – Monetization and Data Science

We will look at each of these areas separately to review how they were addressed by legacy systems and the best approach to provide a solution that meets the requirements of the new network paradigms introduced by NFV/SDN and 5G.

It's important to understand that a traditional assurance solution designed to address the use cases defined here were as complex as a traditional network itself. When observing network communications, all messages and context had to be held longer than either endpoint would normally in order to ensure that all possible failure conditions were handled and reported effectively. Furthermore, complex correlations between disparate interfaces that are required to bisect problems and troubleshoot nuanced interoperability deficiencies can be significantly more complex than two elements exchanging transactions.

While the traditional approach worked for voice only circuit switched networks in the past, the introduction of the "data network" very quickly proved that this model is not sustainable nor cost effective for the virtualized and orchestrated all-data networks that are now emerging.

Operations

Developing network plans from investment decisions, through design and planning, and finally into operation, is the fundamental business of a service provider. Validating designs through observed performance is necessary to qualify the network investment itself. A properly designed assurance systems must be able to measure performance and confirm that service provider assets are operating at maximum efficiency. The solution can also identify necessary optimizations and help to prioritize those activities.

Visibility to observed subscriber and network performance and associated metrics are used to support business processes and often tie directly to important milestones with financial incentives in many operators. Whether these metrics are used to support ongoing network engineering, internal operational efficiencies, or drive direct business value to end customers, a properly architected and configured assurance solution provides the best source of truth.

In the dynamic virtualized world of 5G SA, "swivel-chair" or non-unified full census solutions are no longer a cost-effective solution to support these activities. A representative subset of the subscriber population can cost effectively deliver the same business results.

A unified subscriber experience view, from RF, RAN, xHaul, to Core is now critical for full context – that is, to understand and evaluate various conditions and compare their relative importance and priority.

This is especially important for an operator's key customers, such as enterprise accounts. SLAs are common and using key performance metrics that are already monitored to drive optimization efforts for best customer experience and most efficient use of resources is a natural approach to differentiate the offered network service to end customers.

Diagnostics

Contingent captures are now and have been a hallmark of traditional assurance. For a given interface or individual subscriber that has been determined important enough to observe, every packet is recorded by the probe.

When a network investigation occurs, the user or analysis application initiates a mining of the recorded data. Whether this is brute force, or via indexed structures, the read mechanism is still significantly less often than the write — sometimes as much as 95% resources consumed by the write with 5% or less for the reads.

The technology evolution of storage has not maintained the same level of improvements as other compute resources such as CPU and memory. Early assurance systems storage cost was about 20% of the solution with a deep historical component, but today, to maintain this same historical aspect can turn the cost of storage to become roughly 50% of the assurance solution.

As diagnostics transition into a dynamic assurance-based system, operators will need to consider this historical component and the move to more dynamic captures to address this use case. Of course, assurance vendors will need to evolve their solutions to enable the flexibility and dynamic nature of the 5G SA end state to maximize operators' success in highly cost-effective means.

To meet today's new requirements when designing an assurance solution to support real-time diagnosis and troubleshooting, it is critical to enable the solution to capture only the most relevant and important observations.

Each observation point must be ultimately reliable in terms of coverage and time synchronization — i.e. it's important that responses do not get recorded before the request itself. This would result in loss of faith in the solution that must be relied on to diagnose fine grain challenges within protocol state machines, not just message sequencing.

With a rapidly evolving network, historical views are important for diagnosing problems. In the new network paradigm this will need to be accomplished via data model/policy driven on-demand data capture.

Determining when an anomaly was either first introduced into the network, or where it is and isn't observed currently are important facts for consideration. This also involves trend analysis. That is, if a specific correlation or metric is not currently enabled when the network engineer needs to evaluate it, the assurance solution needs to be able to enable it and begin generating a history for trend analysis.

The benefit of this on-demand analysis is the avoidance of large-scale or full census collection, recording, computation, and archival of contingent measures.

Exploration

In addition to the core function of the assurance platform previously defined with the introduction and ubiquity of today's data networks, access to valuable data for the purpose of monetization and data science are a key function that a properly architected assurance solution uniquely addresses. Flexible exploration of big data for these use cases is key for identifying patterns and trends that operator's businesses can leverage for incremental revenue and value differentiation versus peers and the so-called 'FAANG' technology companies.

Monetization

Knowing the location of subscribers and what they are doing with their devices is of great importance to mobile operators. This information is critical for managing and optimizing networks, but has even more value in the additional revenue it can provide as operators bridge the gap between the digital and physical lives of subscribers for the benefit of third parties in a privacy/GDPR compliant manner. Third parties could include any business that has a need for the movement patterns of subscribers, such as banks trying to assess the best location for a new ATM machine, retailers trying to attract potential customers in an area, or governments trying to enhance public transportation options, among others.

Mobile networks hold a tremendous amount of data on where people go throughout their day, which applications they use, and the topics that matter most to those people at any given time. Until now, this data was under-used as the bespoke assurance and analytics systems used by operators could not enable flexible exploration of the big data gathered nor could they handle the amount of granularity (subscriber apps categorized as "web browsing" vs. "Facetime Audio", "WhatsApp Video", etc.) needed to gain valuable insights for monetization.

The flexible analysis needed by operator marketing and other research teams require deep granularity, without restrictions to the queries, trends, and dimensions explored. As people and machines move from one location to another and use various apps throughout a day, the underlying data reveals large-scale movement patterns and trends that can benefit all sectors in need of reliable — and precise — information about the population and what matters most to those people. This information is a valuable commodity that creates new revenue opportunities for operators.

Data Science

With the rise of big data showing unique value across many industries and situations, it is natural for every data centric organization to seek opportunity to leverage existing data. In many cases the concept of "Data Exhaust" is the first element that is sought to capture for possible value.

In the case of assurance, which has been extremely data rich for generations, the challenge isn't necessarily about capturing the "Data Exhaust", but developing techniques for intelligently experimenting with the data to demonstrate value, i.e. data science, before triggering significant expense to expand observation, collection, and reporting infrastructure.

This selective experimentation with the data is like early gold rush adventurers panning for gold before they invested in a mine. As operator's data science teams mature, the pattern of their usage and scale of commitments change. This usage is extremely difficult to predict or accommodate in a deterministic fashion. Additionally, historic assurance solutions are optimized to protect the ROI of the solution and not sized to accommodate this new use case.

Going forward, it's important that any assurance solution supports this new use case in a progressive fashion to enable access to the observations without detrimentally impacting the existing users. As the value of the new insights are understood, the service provider can plan appropriately to add enough capacity to the assurance solution to support the ongoing mining of that value. In some cases, that capacity may be short lived or even bound by a specific project timeline. Therefore, the assurance solution must be capable of expanding capacity to support data insight efforts in a protected fashion, as well as releasing those expanded resources when necessary.

As the data science teams begin to leverage the assurance solution effectively, the velocity in which insights are uncovered will increase. Whether they are for opportunities to leverage new data streams or for assisting with the operation and planning of the network itself, these scenarios will be impractical to predict. The assurance system must be mature enough to accommodate this dynamic new usage pattern.

Key Aspects of Assurance Solution Design

Having addressed the three major categories (ODE) of assurance and their distinctions, we will now review 4 key aspects of the solution design that will be fundamental to consider when planning the required network assurance solution deployment in the new NFV/SDN 5G network paradigm:

- Virtualization
- Telemetry
- Continuous Integration and Development
- Usability and Interactions with the Data

Virtualization

Virtualization in this context is a fully orchestrated horizontally 'breathable' network functions virtualization (NFV) environment. From a service assurance perspective, the primary impact of virtualization is the same as the network elements themselves: a mandated software-based solution delivery for a broader variety of hardware and for orchestration. The nuanced impact to assurance is the transition from being a secondary consideration, essentially an after-thought, to a co-resident initial consideration with all aspects of the network. Solutions move to become network functions on the same general-purpose hardware as the network.

By using the same infrastructure as the network elements, the relative cost aspect of the network is no longer considered in only monetary terms which can be adapted for business considerations. Now relative evaluations are primarily based on what infrastructure resources are consumed; CPU cores, memory, storage, and I/O. Most importantly there becomes an opportunity cost of those infrastructure resources that are spent for operations (cost) and not revenue generating applications. This has led to many discussions about the relative complexity of assurance and the evaluation of leveraging more self-reported metrics over directly observed packets.

Finally, an operational impact to assurance is that the complex solution involves more variables as the software is supported in different infrastructure environments, adding more parties to the discussion for planning, operation, and diagnosing the assurance system itself. Because the assurance system must maintain the confidence of the user, the significant risk of additional infrastructure variables must be mitigated through procedural training and defensive telemetry. Most operators are hesitant to give full access to the necessary elements of the infrastructure to properly diagnose or even optimize the assurance solution, so it is incumbent on the assurance system itself to provide adequate telemetry such as self-test results for functionality and performance.

Telemetry

In addition to maintenance and optimization of the assurance solution, telemetry offers insight into the specific usage of the assurance system. This information is of extreme interest to determine the effectiveness of the system across the organization based on its usage. This usage can easily inform to the investment decision by demonstrating usage of specific workflows by specific users and groups.

Additionally, the workflow logs can also be mined directly, in conjunction with the broader data lake concept to extract insight into what type of problems are diagnosed and where they are detected vs resolved. With hundreds of thousands of network functions that ramp up and down on demand, protecting the tools that give eyes on these functions, their performance, and the subscribers they service is table stakes for management of the business.

It's important that the assurance solution is significantly instrumented for utmost data integrity. If or when data is lost for any reason, the telemetry system must report it immediately, and a provide a confidence score transparently to the end user or administrator. Confidence is lost in an assurance system if the specific data is not available for the time period for whatever reason, or if the user must wait endlessly only to be told that the data doesn't exist when they know it should.

It's crucial for the assurance solution to be significantly instrumented to constantly maintain the utmost data integrity. It will be incumbent on the operator to prioritize the resources and activities to ensure that this is accomplished.

Continuous Integration and Development

The software industry has embraced and lauded iterative development with focus on velocity. Continuous development concepts allow complex solutions to adapt quickly based on immediate feedback. Network operators have embraced the concept of high velocity iterations focused on deployment and integration to enable new features and applications for their end users.

For a network operator to embrace continuous development, it must start with continuous integration. Bringing multiple modules from multiple vendors increases the complexity and risk by orders of magnitude. The first element of any continuous integration is baseline regression testing for known use cases. The assurance solution offers unique insights for the integration and regression testing to compare results using the exact same observation criteria as the live network.

Perhaps the most critical feature of the assurance system, when participating in continuous deployment and integration, is the ability to adapt quickly to new complexities introduced by other vendors attempting to bring solutions into the network. Using a model driven development method for assurance solutions allows the assurance vendor to deliver a functional solution for these discovered complexities as quickly as possible.

Another benefit of the model driven approach is that as new functional features mature in usage, they can be migrated from functional to performant through a repatriation effort that allow assurance engineers to optimize appropriately based on the nature of the new function. Essentially, the model driven approach allows functional features to be introduced quickly and grow in performance as necessary.

Usability & Interactions with the Data

As the networks have evolved, and as complexity continues to increase, the availability of expert engineers has become a challenge.

Intuitive solutions that provide clear, transparent, and unified, RF, RAN, xHaul, Core subscriber experience views — in short, unified views of an entire given situation, provide the best means of augmenting the existing experts while guiding less experienced users through the methods of maintaining the network. These guided workflows allow the experts within the assurance vendor to distribute common lessons learned across other users based on demonstrated values — essentially allowing the operator to purchase real world experience that is embedded in the tooling through guided workflows.

With respect to expert users, the assurance solution itself is in a unique position to assist with building intuition with the actual network being operated. Using the assurance system to capture and archive specific events allows follow-on engineers to replay those events with the actual tooling to obtain firsthand experience on how to diagnose and respond to specific challenges. Vendors adding UX elements like annotations will further enhance this capability by allowing senior expert users to augment the captured scenario with specific notes and guidance for training purposes.

Users of an assurance solution have different needs — i.e. there are different use cases for each type of user. As most network administrators rely on quantitative data to drive decisions, when hypothesizing on root cause of anomalies, they require validation with reproducible data. The challenge comes when the cost to acquire and process the desired data rises beyond an easy decision. Assurance vendors developing an approach to determine what observations are important for which activities is therefore critical. The following diagram presents a framework to cover aspects of the ODE use cases:

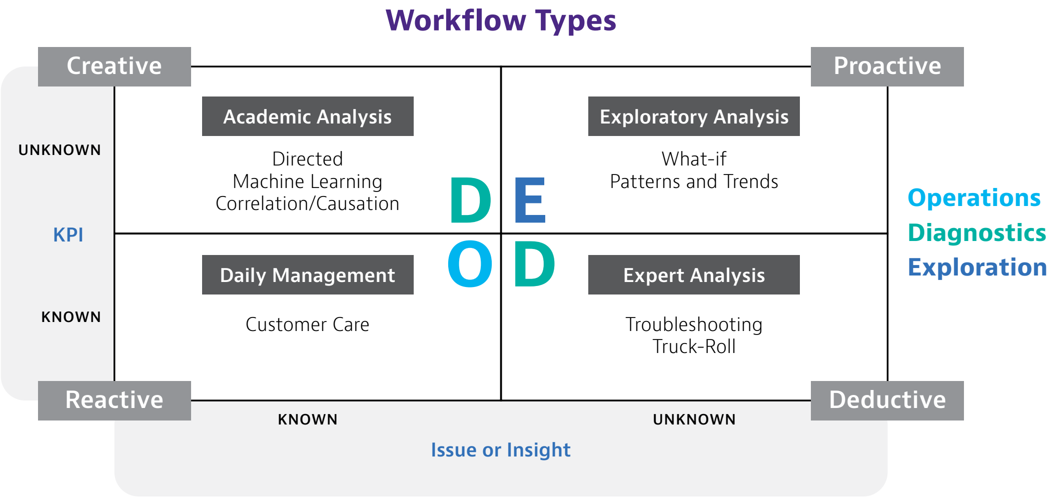

Assurance User ODE Use Cases

Daily Management — Operations

In this use case, there is a well-known relationship between known issues or insights and the appropriate metrics or KPIs used to react. This situation is most appropriate for automation with a clear understanding of the cause-effect relationship and track record of correlated results. While the primary role of assurance in this quadrant is "Enable Automation", the opportunity to provide intuitive workflows with detailed information enables an additional role of "Build Intuition".

Expert Analysis — Diagnostics

Here, a known KPI has triggered activity that must be resolved for a less-likely scenario unable to be covered by automation and expert analysis is required. This expert will leverage all the data available to quickly resolve the issue and may enable new observations like dedicated traces, or elevated logging levels to accelerate resolution. The primary role of the assurance solution in this quadrant is "Expert Augmentation". It is important, however, to recognize that any result from this quadrant must be evaluated for migration to daily management to avoid re-diagnosing recurrent problems.

Academic Analysis — Diagnostics

An issue is known, but a correlated KPI or Metric is not known in the Academic Analysis quadrant. Perhaps a significant problem was found during Expert Analysis but was not correlated to the initiating KPI. Here the analysis is exploring KPIs and metrics to find a proper correlation to confidently move the resolution into daily management and automation. The primary role of assurance in this quadrant is like Expert Analysis: "Expert Augmentation"

Exploratory Analysis — Data Opportunity Exploration

The final quadrant involves exploring data for new correlations and following the thread to verify causation. This approach works with both opportunities to improve (i.e. optimization) as well as new opportunities to monetize insights derived from the data. It is important to note, that the flexibility of a given vendor's assurance system is critical to lower the barrier to entry for exploration. Experimenting with a subset of high-quality data before expanding across the entire network is significantly more cost effective than over-reaching and incurring significant cost. The primary role of assurance in this quadrant is "Observation Broker" – a user who controls access to observations in a manner that does not impact the OD portion of the system, making the right data available to the right people, at the right time and for the right cost. This user would be deterministic with respect to cost so that the user/admin can decide if it is the right cost before executing the study. Decisions here will be based on concurrent requested exploratory workflows

Next Generation Assurance Methodology

We have identified and discussed the three key ODE use cases for network assurance and have also reviewed the design considerations for new assurance solutions in NFV/SDN 5G network environments. Now we will review in detail the methodologies needed to enable and support both.

Self-Reporting and Direct Observation

Traditionally the key challenges for self-reported observations were the trust and validation of the data, impact to the availability of resources for the core function of the element in question, and potential unavailability of the self-reported data in the event of an element outage or critical network event. These considerations have been the primary motivation for direct observation.

Direct observation via passive network monitoring involves a non-intrusive third-party observer monitoring the network communications between two or more endpoints. Traditionally, to properly observe the communication between endpoints, the observer was required to maintain significant hold periods to account for all possible failure scenarios.

With the introduction of LTE and all the all-data core network, it has now been demonstrated that the sheer volume of data as well as the complexity of physically tapping the interfaces pushes the cost and risk of direct observation beyond a reasonable ROI. As a result, the emerging proposal for a solution is to leverage the elements themselves to collect and report relevant observations including raw packets. However, these self-reporting feeds come at a cost of resources and/or performance of the network element itself along with potentially significant financial license costs as well.

There is now an opportunity in the new NFV/SDN 5G network environments to strike a "balance" between direct observation and self-reporting to achieve the required visibility to support the previously described use cases, provided the correct methodology is employed.

Virtualized Passive Monitoring and Validation with the 'Truth in Packets'

With the introduction of virtual infrastructure and the 5G architectural changes, the cost and benefit evaluation of self-reporting and passive monitoring should be reconsidered for many interfaces. Additionally, operation of network elements within virtualized infrastructure makes self-reporting a potential as a standalone function independent of the network function itself — a potential complement to direct observation (passive monitoring).

To support the premise of an assurance solution that integrates self-reporting and direct observation there is a concept that has arisen recently in the assurance industry that requires consideration. Namely that there is a "truth in packets". While logs and self-reported measures are generally a good starting point for analysis, when all else fails, there is always "truth in packets", leading engineers to analyze full on-the-wire or per subscriber Wireshark packet captures to finally debug a problem occurring with a solution that involves wireless network communications.

Within the virtualized architecture, there are several distinct methods of providing virtual tapping for packet observation. A given virtualized assurance/passive monitoring solution must be developed to support the relevant methods for a specific network. However, the primary consideration for assurance supporting virtual taps is to consider direct observation vs. self-reported metrics. This usually means that the virtual tap must reside collocated to the network element in question and have complete visibility to the relevant packets directly and quickly, while maintaining dedicated resources to avoid any degradation such as with self-reported observations.

Self-reporting in a Comprehensive Assurance Strategy

As operators become more cost conscious, and with the ability/opportunity for NEMs to incorporate self-reporting as a standalone virtual function, it is logical to consider the idea of utilizing more self-reported observations.

As soon as self-reporting is considered, it is important to focus on maintaining the trust in the observation. It is important to validate any self-reported data feed through itinerant and campaign based direct observations or more intensive and complex testing scenarios.

The simplest approach is to replay known scenarios into a system and validate the self-reported results. This may however limit the ability to quickly add scenarios for test because the expected results end up being crafted by hand.

The alternative to the crafting test & results scenario is to enable the on-demand passive monitoring of the element being validated. Upon completion of the scenario, the self-reported results can be compared to the passive observations for validation. This is generally more effective but does carry the complexity of interpreting the self-reported results accurately, which is why standardization of those self-reported observations is encouraged.

Comparing results between self-reported and direct observations is complex enough — without the reality of the variance existing between vendors as well. Most operators strive to maintain at least two supply chain sources for all elements of their network. This provides negotiating power as well as good network hygiene to avoid homogenous elements with a unified risk profile.

The challenge of comparing self-reported results from network element vendor A to vendor B is complicated by the lack of transparency on the generation of these metrics by the vendors. The use of passive monitoring to effectively reverse engineer critical measures based on observed scenarios enables the operator's assurance solution to understand and rationalize the variant measure reported by disparate vendors.

Topology

Networks have grown and will continue to grow significantly in both scale and complexity. The challenge of accurately maintaining the topology of network functions increases as a result. Several options for dynamic discovery have been utilized to varying affect and the opportunity to use the assurance solution as a consolidation and validation of the topology records is a recurring discussion.

The virtualized infrastructure reality of networks today offers the key benefit of dynamic resource allocation which allows elements to be "spun up", and "spun down" with horizontal scaling. Many of these nodes/network functions end up being itinerant. The topology representation of a virtualized assurance system must account for the itinerant nature of some elements while maintaining the proper historical view to support trend analysis appropriately. Therefore, it is critical that the observation and real-time portion of the assurance solution is isolated from the historic portion when targeting a virtualized network.

Capture Everything or Just Perform Data Sampling?

Additionally, when facing the unique challenge of both virtualization and the evolution of beam based 5G, the complexity of the network increases dramatically. The already pressured ROI of the general-purpose assurance solution must adapt to this complexity while maintaining use case functionality and minimizing the cost. The best way to approach this is by disaggregating assurance across multiple dimensions.

The first method is to consider the three primary user workflow categories and provide a solution that can adaptively accommodate those needs. Traditional assurance solutions have focused on a "full census" coverage which indicates all interfaces that need to be monitored are always fully monitored as well as historically archived. However, each of the ODE user workflow profiles all have distinct requirements resulting in a mix of methods for coverage within assurance that can be considered for future assurance deployments.

Operations Coverage

Traditionally the operation use cases benefitted from the full census conditions that existed to support diagnosis. It's rarely been considered whether the operation workflows could function with less. Based on anecdotal evidence collected over a number of years, it is now presumed that *most of the operational workflow can remain functional with the representative subscriber sampling as far down as 5%*. Keep in mind that the KPI must be extrapolated based on the sampling model, but the resultant KPIs remain directionally correct when used for network optimization and even priority assessments. The key factor to support the operations workflow is to ensure that the sampling model is based on subscribers directly and that the results are extrapolated to 100% to maintain consistency when different selection percentages are used.

Diagnostics Coverage

Troubleshooting needs to have visibility to all aspects of the network. Traditionally full census aspired to provide rapid response and reasonable forensic investigations but failed due to the challenges associated with indexing and storing the vast amounts of data in a format and schema that allowed for rapid search and retrieval of device, subscriber, network, or service events in context of the issue under investigation. The promised potential of these bespoke solutions was never fully realized, and they essentially became very expensive "call-trace" and "wire-shark" RCA tools never truly providing full "assurance" capabilities.

As the relative cost of storage has gone up through technology trends, the forensic history for assurance solutions has reduced as well. Early deployments offered 30 days of packet captures which had to be reduced to 7 days and has often now been challenged to stay at 4 days to support the "long weekend" scenario. It's evident that while historical coverage is beneficial, both the depth of buffering and alternative resolution options through real-time and future-capture become viable for cost.

Today's primary consideration for diagnosis should be to maintain interface coverage of both RAN and core to be able to see anything in a unified, in real-time manner. In addition, it is important to place appropriate capture criteria like a subscriber identifier to ensure that relevant traffic is captured to facilitate troubleshooting moving forward. Techniques like IMSI white list are critical to support progressive diagnosis vs. capture everything for <x> amount of time.

Exploration Coverage

Exploration is the most undefined of the assurance workflow profiles, but it has the most potential benefit. The ability to select specific observations, observation elements, metrics, and selection criteria will enable the data scientist to adequately explore the vast possibilities from network data. The ability to start with a small representative sample is important to facilitate fast iterations as the data scientists apply agile principles and iterate quickly. As confidence in both the relevance and quality of the data increases, the flexible assurance system can grow in capacity and allocation to provide expanded coverage. The result is a solution that is deterministic in expected resources for network wide rollout of any analysis and facilitates logical evaluation of a situation dependent business case to roll out specific analytics.

Another method of disaggregating assurance across multiple dimensions is to consider how sampling is performed. Should sampling be by packet? This is the simplest form, however interfering with state-machines and complex DPI pattern matching and heuristics have follow-on effects in metrics that cannot be extrapolated. Random packet sampling can also present as significant transport errors.

Should sampling be by flow? This is more reasonable; however, it is still very difficult to extrapolate the sample effectively into usable data. While transport/xHaul metrics are not nearly as negatively affected as with packet sampling, the DPI impact for heuristic service classification detection continues to be large.

By session? Session sampling is much better than packet or flow sampling, however, it is difficult to manage the large mix of varied session characteristics such as long running sessions. Selecting too many of these impacts more representative sessions' availability. Additionally, it is an incomplete view of subscriber activity. Finally, is the difficulty of maintaining a proper selection criteria archive to inform users if a specific observation would have been captured or not.

The best option for sampling for selective scaling is subscriber sampling. There are several selection criteria such as by IMSI/SUPI hashing to maintain a representative random sample that can be extrapolated in a straightforward manner. Here, the selection criteria archive can be queried by a workflow to determine if a specific subscriber was in the monitored population for the respective time window being queried. When a subscriber is selected, all traffic for that subscriber is captured and processed.

Assurance Slicing Per Workflow Using Subscriber Selection

In short, there are multiple approaches to disaggregating assurance across multiple dimensions to maintain use case functionality while minimizing the assurance solution cost. The assurance solution needs the flexibility for the operator to determine what approach is best and it should be in line with the use case required. To maximize this, operators need to consider that for example, if a subscriber is selected for diagnosis, that does not guarantee that that data is processed for the network operation use case. Only when a subscriber is selected to be included in the operation data population will it be processed in the operation analysis modules and archives. This means that it is very possible that single subscriber could be selected for both operation and diagnosis modules and the population participation must be maintained separately.

Selecting a distinct population is a valuable capability to enable the assurance solution to scale down and still provide valuable insights to the operation use case. While the concept of populations is relatively new to assurance, the benefits continue to grow when applied to the exploratory workflow that could very effectively use many populations. The working assumption is to limit the concurrent population sets to 5 to minimize per observation overhead as well as performance.

By introducing the multiple population and selection criteria, the required capacity of the assurance solution can be reduced well below full census and achieve valuable results. With the system able to function with a wide variety of resource mixes and priorities, it facilitates the ability for an individual use case to justify and expand the assurance footprint and capacity. This expansion would be on the merit of the single case and avoid having to shop other budgets to share the large cost to achieve full census

In essence, as in 5G with network slicing, assurance is sliced per ODE workflow as in these new environments, it now makes sense to dynamically assign resources for assurance activities according to priority, need, and use case as defined via the methodologies detailed in this whitepaper.

Resilience

As the usage of assurance evolves into more departments and use cases, the importance of the solution subsequently and continuously increases. The primary concern regarding resilience is to avoid losing data and functionality. Partial loss of functionality is far preferable to total degradation.

Graceful Degradation

One key approach to maintain functionality is to ensure deterministic degradation. With the introduction of population selection, the primary factor of selected percentage can be adjusted to control the demand on the system. This approach degrades the load on the assurance system on a per module basis while maintaining relevant functionality.

High Availability (HA) vs. N + M

Traditionally, the cost and complexity added to the assurance system to maintain the telco standard of five 9s of reliability became limiting factors. Proper HA involves fully redundant active-standby solutions with checkpointing state-machines to verify zero data loss on resource restart. When assurance solutions started adding this, the level of cost quickly escalated, and designs shifted to N+M. The intent of N+M is to provide a pool of resources in which lost functionality in the (N) resources is recovered by allocating from the (M) pool. Theoretically this approach allows for up to M failures with various levels of degradation during switch over and restart. Additionally, load balancing, in a resilient manner, with all resources mostly loaded, allows for the migration of load to the standby capacity in a more direct manner without requiring a resource configuration and startup.

Model Driven Development for Assurance Solutions & Data Abstraction

Considering the additional use cases around data exploration, there is an opportunity to pivot the architecture of the assurance platform to enable faster feature velocity. The use of model driven development (MDD) allows the observations to be processed into metrics and stored into a flexible data model. The requirement to support dynamic data models and metric processing for data exploration makes the adoption of MDD principles important.

The key complexity when applying MDD to assurance, is to ensure that the architecture model accommodates the location of processing and resources In some cases, such as low latency user-plane analysis at the edge, the computation resources must be allocated at the edge to avoid the large scale backhaul of observed user-plane traffic. However, when building a data exploration model, the direct observations are preferred in the central archive because it allows more complex explorations during query. The backhaul cost must be modeled and allocated to the use case to ensure transparent demonstration of the cost of an exploratory model.

MDD for assurance should include an observation catalog of types of observations: direct, proxy, or synthetic (reconstructed). The processing elements then need appropriate models to detail the operations for generated measures across selected dimensions. Additionally, the selection model needs to accommodate the various selection criteria and processing models to ensure the population survives the statistical calculations to support the distinct workflows and user workflow categories. Finally, when a model is operating at less than full census, the historical selection criteria is maintained within the model as well as KPI projections to support historical workflows across all three ODE user classes: operations, diagnostics, and exploration.

Like the model driven development and its applicability to the data exploration use case, a highly flexible data model to allocate data archive resources is highly beneficial. When we consider enabling the data exploration use case with a sliced assurance model, the size of the data landscape is no longer a contributor to the minimum viable assurance solution. The data model challenge becomes the dynamic nature of the changes to the data model to enable experimentation without requiring network wide data schema changes and associated resource allocation. The ability to experiment with a solution and then allocate resources to expand coverage as deemed necessary is the fundamental requirement from the data model to support exploration.

In order to support this level of experimentation, it is important that the data model not only be dynamic but support prioritized interactions. High priority interactions with the data coming from diagnostics and operations must be maintained while experiment usage is served opportunistically. It should be minimally possible for an experimental query to impact the stability or performance of the data archives even if it is at the cost of performance for the experimental query. Additionally, even within the operational use cases, interactive workflows should be served with a higher priority than periodic reports and exports.

In addition to managing priorities amongst data access, the data archive system should provide transparent usage and cost metrics. This should help plan the resource allocation of the assurance solution directly to support normal expansions with network growth as well as the deterministic rollout of experimental data sets across the entire network.

Suggested Architectural Requirements

When building or expanding networks by leveraging virtualized infrastructure, the assurance solution can no longer be addressed during later integration phases. It must be designed or at least planned into the solution from day one to ensure that adequate capacity of the infrastructure is allocated to accommodate assurance in the appropriate forms. This brings the challenge of using the same infrastructure for assurance as the network itself. The tradeoff of complete isolation vs. flexibility and cost savings is a fait accompli. The assurance solution must perform beyond reproach such that the network administrators maintain confidence and can rely on provided observations as fact during complex diagnosis.

Containerized Linux

With the migration to virtualized infrastructures, the best solution to support the widest range of carrier environments is utilizing containerized Linux instances orchestrated by Kubernetes. Using this as a fundamental building block provides top performance with dramatic flexibility. While each permutation of infrastructure may require slightly different procedures and techniques, containers can easily be adapted to carrier IT policies effectively. These techniques align with the broader IT industry that is rapidly leveraging these technologies to their advantages.

Keep in mind that there are still tradeoffs and is not a magic bullet. Anytime a solution is leveraging third-party hardware, performance nuances will arise. Running within an infrastructure that is administered without transparency proves to be difficult to diagnose when something does eventually happen to the running solution. It is more important to ensure that all critical elements that are optimized for performance implement appropriate telemetry including self-tests to assist in diagnosing issue in a third-party environment.

Microservices and APIs

Microservices imply strong APIs. Additionally, the microservices approach requires modular component driven implementations. This modularity facilitates not only higher development velocity, but higher flexibility in the operational system as well. A primary example is to consider the scenario where complex KPIs need to be calculated at the edge, while another set need to be calculated globally after observation are collected centrally. The KPI module configured to run at the edge for one set of observations, while running centrally for another set avoids the necessity to spend engineering resources to code the solution and can often be achieved with basic configuration.

Performance Telemetry and Regression

Use of the modules and strong APIs is required of microservices. It's also important to build unit tests and regression tests that highlight and maintain performance. When specific components fundamentally require performance levels, those telemetry metrics can be gathered during development to ensure that no degradation is allowed. Additionally, when an early phase feature is developed, the expected performance can be managed and then evolved based on the operator's priority and updated with end-users. Essentially start with functionality that is often critical to a customer and then follow with highly optimized performant modules.

Configuration

Configuration of anything within the system can easily grow complexity beyond the original intention. It's important for a dynamic solution to support custom configuration options, but it is equally important to have a clear view and history of all configuration changes. Any configurable element in the system must maintain a template to define the configuration, purpose, options, ranges, recommendations with context, as well as a full history as configured for a deployed system. To support changing defaults between releases, all historical records should be labeled with release information by the given assurance vendor. While this may seem simple, enforcement is critical to maintain deterministic service and support.

Conclusions

As the industry embraces the technology disruptions and innovations associated with the latest trends driven by NFV/SDN and 5G, it is imperative that service providers recognize it will no longer be acceptable to treat assurance and analytics as an afterthought to be addressed once the target network design has been finalized and deployed.

In order to properly equip the teams that will be responsible for operating and managing the new network once deployed, they will need to be enabled out of the gate with the Operations, Diagnostics, and Exploration capabilities that a properly designed and integrated assurance solution must provide. It is no longer practical to assume there will be a bespoke solution available for implementation post-deployment to provide these capabilities.

In addition, it is incumbent on every operator to embrace the opportunity that access to the vast insights that mobile data provides via a properly designed assurance system to capitalize on monetization opportunity and to support data science activities that can contribute to the development of new services, identification of critical trends and patterns, and longer-term statistical correlations.

Selection of a capable assurance and analytics partner as you embark on this journey is as important as is the selection of your network vendors. It is imperative that the selected assurance solution partner have the requisite domain expertise to understand the implications and nuances of these new data collection approaches with respect to:

- Identification, collection, mediation, and correlation of the data required
- proactive fault identification
- automated workflow creation and enablement
- comprehensive reporting and analytics
- ...given all the paradigm shifts and their implications outlined in this whitepaper

Equally important as the assurance vendor's domain expertise will be their willingness to be flexible and truly work as a trusted partner with you as you embark on this journey. We must all acknowledge that there are too many variables and still many potential unknowns to claim we have all the answers before comprehensive review of any new network design that fully incorporates the assurance solution as a core component and not as an afterthought.

VIAVI Solutions

Contact Us **+1 844 GO VIAVI**
(+1 844 468 4284)

To reach the VIAVI office nearest you,
visit viavisolutions.com/contact

openassurance-analytics-wp-xpf-nse-ae
30191086 900 0420