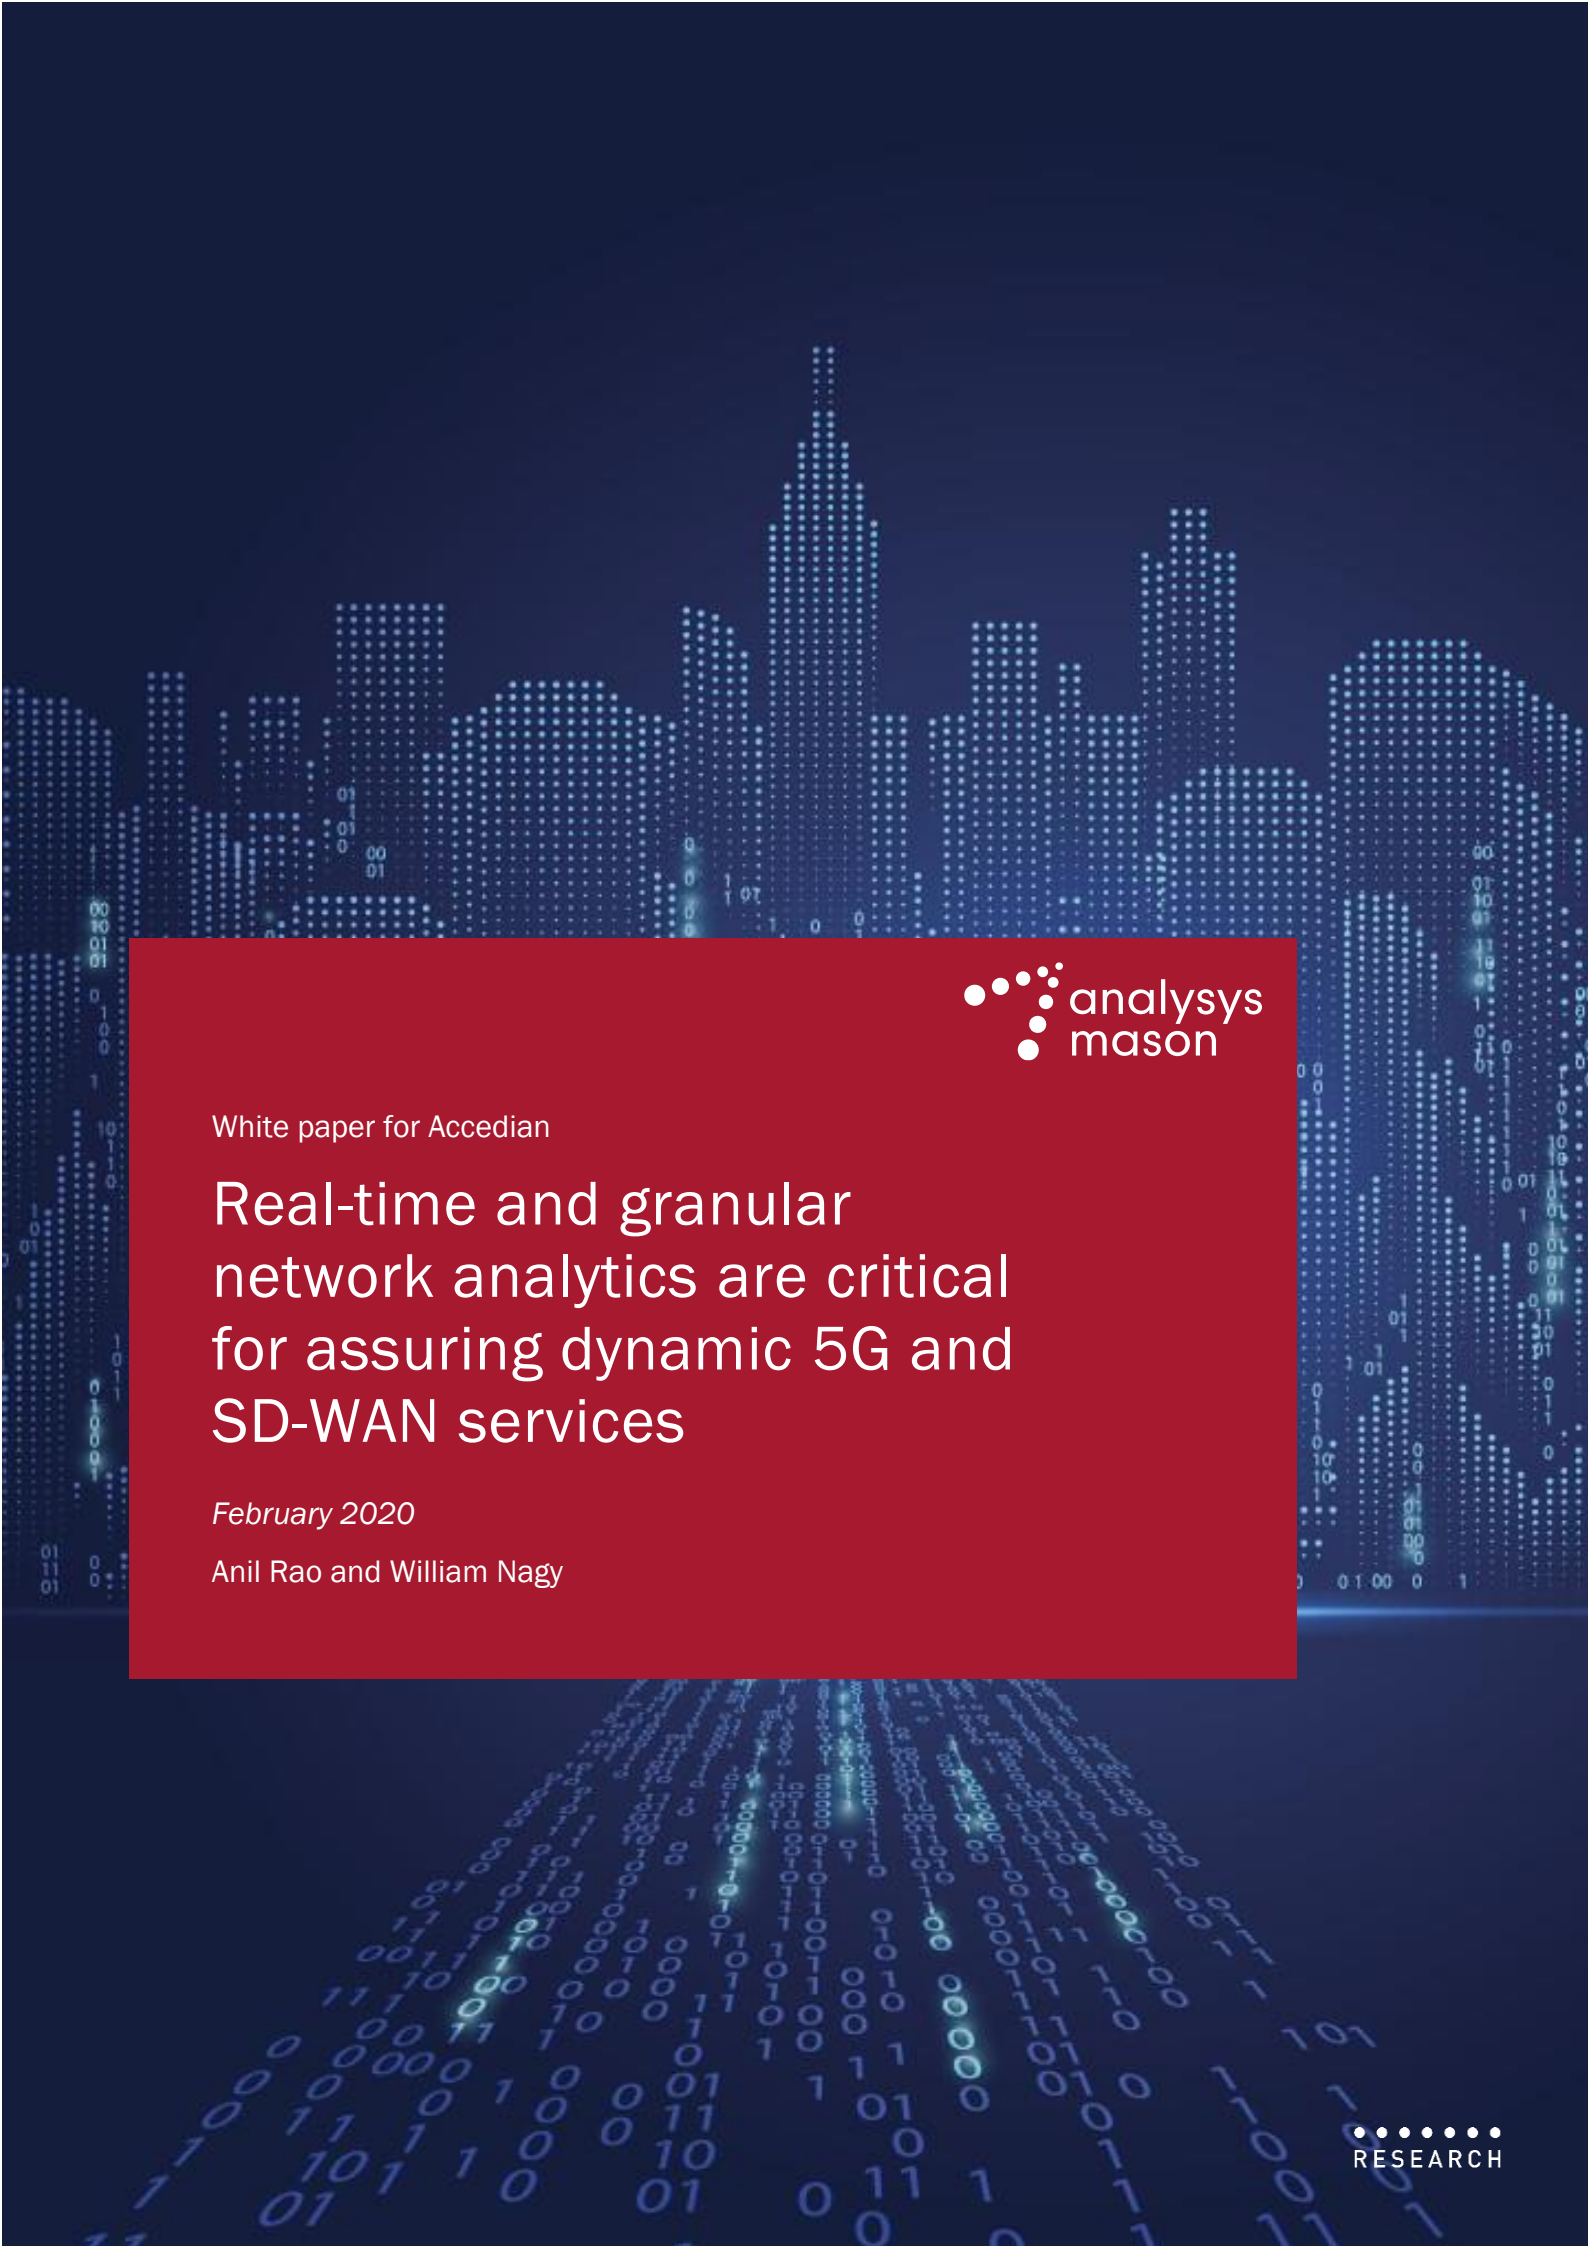

analysys mason

White paper for Accedian

Real-time and granular network analytics are critical for assuring dynamic 5G and SD-WAN services

February 2020

Anil Rao and William Nagy

RESEARCH

Contents

1. Executive summary

Enterprises across all industries are undergoing digital transformation. Their wide-ranging initiatives include delivering superior digital experiences, migrating to the cloud, embracing IoT and edge technologies, and automating their operations. Service providers are supporting enterprises' digital transformation efforts by launching enterprise services such as SD-WAN and by deploying edge cloud and 5G infrastructure to enable new service and business innovation.

In order to deliver new, dynamic services with faster launch and provisioning times, service providers are embracing virtualisation and networking technologies such as network functions virtualisation (NFV), software-defined networking (SDN) and cloud-native computing (CNC). Together, these technologies create a highly dynamic networking environment in which service providers can instantiate and provision services on demand, which enables enterprises to become agile businesses. However, these next-generation dynamic networks are highly disaggregated and complex to manage, which presents a new set of monitoring and operational challenges for service providers. Service providers must review their monitoring and service assurance strategies to deliver on-demand enterprise services with guaranteed quality and differentiated service level agreements (SLAs).

Traditional SNMP-based, poll-based network monitoring techniques rely on non-real-time network performance data and are insufficient to assure services delivered over highly dynamic networks. Service providers will struggle to detect common transient network issues such as microbursts because of the dynamic nature of these networks. Network monitoring must adapt with the changing network state and rapidly adjust to services that are provisioned on demand. To achieve this, performance monitoring solutions must use highly granular, real-time data in combination with advanced analytics techniques for correlation and issue isolation.

Monitoring techniques such as network telemetry (NT) are gaining popularity for precisely this reason – they can continuously stream real-time network data to northbound applications. NT provides network data at sub-second granularity that can be captured using standard modelling techniques such as YANG. Supplementing NT data with other data sources, such as wire data and active testing data, provides a powerful base for network analytics. A performance analytics platform must ingest data from multiple sources in a multi-vendor environment to generate comprehensive performance insights for the end-to-end network. These insights can then be used for network operations centre (NOC) and service operations centre (SOC) functions and customer reporting on self-service portals. Performance insights can also be used to drive closed-loop network automation in combination with policy rules and network orchestration systems.

Performance analytics can also be enhanced with machine learning (ML) and artificial intelligence (AI) techniques to establish predictive operations. Historical network data sets can be used to train ML models to identify the patterns that lead to degradations of network performance and service quality. When these models are applied to real-time network data, anomalies (such as microbursts, link saturations, misconfigurations or hardware issues) can be identified and future degradations can be predicted before they occur. Real-time network analytics – in combination with ML/AI – will play a critical role in delivering business-critical enterprise services with guaranteed service quality and form the basis of zero-touch automation in service provider networks.

2. Service providers are launching agile network services to support enterprise digital transformation

2.1 Enterprises worldwide are transforming operations, and service providers are evolving to enable their success

Enterprises across multiple industry verticals are transforming their business operations to meet internal and external market demands. A key business priority is to deliver superior digital experiences across online and offline channels and enable real-time customer interactions. From a technology standpoint, enterprises are migrating applications and software development to the cloud, adopting SaaS-based cloud services, and implementing IoT, edge and Industry 4.0 use cases to transform business models and operations. To achieve success with these initiatives, enterprises need agile on-demand services from service providers that supply the all-important WAN and cloud connectivity services for branches, as well as supply IoT and 5G connectivity outlined here.

- **SD-WAN**: enables on-demand policy-based traffic-steering technology at the edge, which enables enterprises to dynamically select a WAN path (for example, IP/MPLS or internet) based on factors such as the destination (including public cloud, private cloud, data centre or another branch) and the criticality of the applications. SD-WAN transforms the enterprise branches with higher levels of agility and dynamic control. When combined with multi-access techniques including LTE, 5G, Wi-Fi or fixed access, SD-WAN provides diverse connectivity choices for both static branches as well as a mobile workforce.

- **Edge clouds**: bring the power of compute and storage closer to the end user or application. Edge clouds reduce the latency of packet data because the user plane data does not need to be backhauled to a central cloud core. Edge clouds also allow service providers to support advanced enterprise IoT use cases, as well as big data and AI applications.

- **5G and network slicing:** uses the cloud-native 5G core to create end-to-end, isolated virtual logical networks that span the radio, transport, edge and the core. Each slice is capable of offering differentiated Quality of Service (QoS) and SLAs to meet service-based latency and reliability requirements. Network slicing paves the way for service providers to offer either differentiated slice-based services for whole industries or highly granular differentiated slices per use case, subscriber type, application – or to provide an enterprise with slice-level control, management and QoS.

SD-WAN offers service providers significant growth prospects

Analysys Mason forecasts that service provider spending on SD-WAN will reach about USD2.1 billion by 2023. Service providers are now offering managed and co-managed SD-WAN services to capitalise on the demand from enterprises. This allows service providers to continue working directly with enterprises to upgrade sites to SD-WAN and to guarantee performance. Furthermore, a managed SD-WAN service provides communications service providers (CSPs) with a strong foundation for strengthening the business relationship with an enterprise beyond connectivity and with the ability to add other value-added services such as security.

Figure 2.1: SD-WAN market forecast for service provider spend, worldwide, 2017–2023

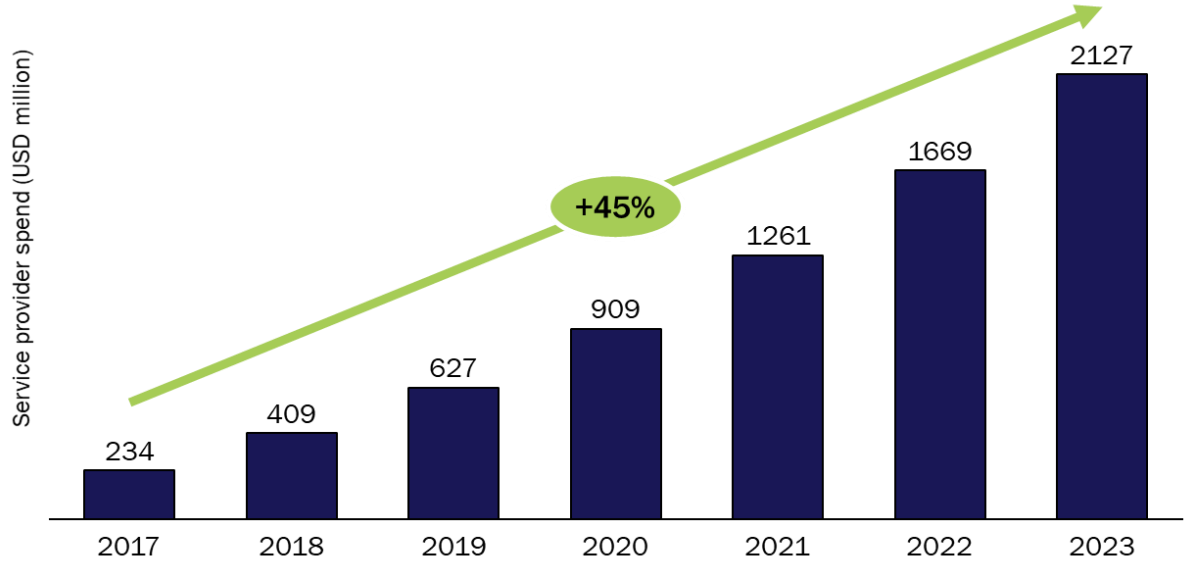

Source: Analysys Mason, 2019

Next-generation programmable cloud-native networks are driving service provider transformation

Service providers are disaggregating and virtualising their networks to achieve better economies of scale and reduced CAPEX, as well as to enable higher level of operations automation, service agility and reduced OPEX. SDN/NFV and containerisation all enable flexible and programmable networks that deliver on agility and rapid deployment requirements. Together, these technologies will make the network highly dynamic and complex with many moving parts that must be abstracted, managed and monitored so that service providers can deliver a superior customer experience and can guarantee QoS and SLAs.

Universal CPE (uCPE) platforms provide an example of a revolutionary approach to offering SD-WAN-based, on-demand business services. uCPE platforms reduce the complexity of the branch networks by replacing dedicated CPEs with equivalent software-based VNFs, which all run on cloud-based NFV infrastructure. By using an NFV orchestrator, CSPs can dynamically provision certified multi-vendor VNFs onto the uCPE platform and configure the end-to-end service. An SDN controller can make traffic-steering decisions to route traffic in the network based on network and security policies, and real-time traffic conditions. Service providers, such as AT&T, Bell Canada, BT, Telstra and Verizon, are preparing to support SD-WAN services on top of uCPE platforms.

5G will be a confluence of many cloud and networking paradigms

When implemented in full, a 5G standalone (SA) deployment will allow service providers to offer enterprises massive machine type communications (mMTC), and ultra-reliable low-latency communication (uRLLC) services. 5G SA will introduce a fully containerised next-generation core, which will help service providers to address a much broader 5G market opportunity than traditional telecoms services.

5G SA also allows automatic scaling-up of only those network function service modules that require more capacity. It also enables rapid service innovation through automated network and service orchestration by interlinking select network function services into end-to-end services. This will significantly reduce service creation timescales to mere minutes. 5G services that require network functions to be instantiated in microseconds will depend on cloud-native virtualisation and extreme automation technologies. Ultra-low latency requirements will

mandate the deployment of VNF instances in the edge clouds. Furthermore, to meet the scale and dynamicity of the service requirements, service providers may have to create and manage diverse network slices to cater for the many enterprises and use cases.

2.2 Automated assurance will be critical for delivering guaranteed service quality, superior customer experience and differentiated services

The increasing dynamicity and complexity of NFV/SDN and cloud-native 5G networks is placing unprecedented demand and pressure on service providers to deliver always-on, highly reliable and high-performance networks and services.

In a highly demanding business environment, enterprises cannot afford lapses in service because this can result in lost revenue, lost customers and a damaged reputation. Consequently, enterprises have significantly raised the bar of expectations around service quality, SLAs, customer experience and service differentiation. This, in turn, is compelling service providers to completely rethink how they monitor and assure their infrastructure, networks and services. Service providers need a fit-for-purpose automated assurance approach that meets the high expectations of digital enterprises and plays a key role in differentiating services. This approach must be taken into consideration from the outset when creating services and planning service delivery.

3. Service providers need a new operations approach to assure enterprise services in dynamic networks

3.1 Service providers need a way to see and react quickly to microbursts

Traditional monitoring techniques cannot detect common transient network issues such as microbursts, which (despite being short-lived) create significant problems within the network. Microbursts are surges of data packets that occur due to sudden spikes in network traffic. These anomalies have a domino effect on network performance, leading to a very short window of service quality degradation around latency, jitter and packet loss, and potentially causing a critical service failure. As such, microbursts are impossible to detect and trace in static physical networks that use monitoring techniques based on minute-level timescales.

It is even more difficult to detect microbursts in dynamic networks, which makes it much harder for the service provider to assess the impact on customer experience. Without a monitoring solution that is well-suited for a dynamic network, the customer is left exposed to the effects of microbursts and the service provider will have failed to assure the QoS that customers expect of the network.

Dynamic networks that deliver on-demand SD-WAN and 5G enterprise services will require sub-second level performance monitoring data (for example, in the order of milliseconds) to ensure rapid reaction times to service quality degradation caused by microbursts and also to deliver on customer expectations and SLAs. It will be important for service providers to prepare in advance so that once a problem is detected or predicted, action is taken quickly and efficiently to fix the problem before customers are impacted. This is not something that traditional monitoring approaches support.

3.2 Network performance monitoring must use real-time data

Traditional network monitoring techniques such as SNMP use polling mechanisms on network devices to gather network performance status, but this approach is not fit for purpose. Based on the 'pull' mechanisms, the polling intervals of these approaches can range between 5 and 15 minutes, which means that the network data used to identify performance issues and to perform root cause analysis or troubleshooting is 'old'. Furthermore, they provide a siloed view of the network and require a significant amount of systems integration, data aggregation, data cleaning, and post-processing to obtain an end-to-end view of the network performance.

With the dynamic nature of the NFV-/SDN-enabled networks where the network state can change at short notice, it will be essential to derive accurate inferences in order to make the right decisions about the next best action in order to resolve network performance issues. Advanced analytics and correlation techniques will be required to identify and isolate network performance issues. For example, in the dynamic uCPE-based SD-WAN or the 5G network, VNFs (or cloud-native NF) and service instances can be created and altered on-demand, including dynamic traffic flow changes based on SDN policies. Network monitoring must adapt and, if required, scale in line with the changing network, to monitor the portable VNFs and the modified service chains in real time. This can be achieved only by using near-real-time network data to discover network changes and correlate multiple network data sources (both historical and real-time streaming data) to isolate performance issues and to take decisive actions for remediation.

Use case: significant reduction in mean-time-to-resolution (MTTR)

With legacy performance monitoring approaches (discussed in section 3.2), if an issue arises just after the network has been polled, the service provider will not discover the issue until the next polling cycle or – in the worst case – after the enterprise customer has already experienced the performance degradation or a service outage. This means that the 'time to know' about an issue occurring can take as long as 30 minutes in some cases. Once an issue is detected, the service provider must then determine its root cause which will further prolong the 'time to discovery'.

The dynamic and transient nature of SD-WAN and 5G-based networks and services will only exacerbate the scale of the problem, which may lead to a significant surge in Mean Time to Respond and Repair (MTTR). Indeed, not all issues require real-time action, but it is important to know, discover and resolve issues when they occur. This is especially critical for the operations of dynamic networks: service providers cannot afford to wait for 30 minutes to simply identify the root cause because in that timeframe the network and/or the service state may have changed and resolution actions may no longer be relevant.

Analytics-driven service assurance will be key to significantly reducing the time and costs associated with assuring complex NFV- and SDN-enabled SD-WAN and 5G networks. Using ML and AI techniques, service providers can predict and prevent performance issues before they occur or impact service quality. The potential for cost avoidance and cost savings can be significant. Analysys Mason's models show that service providers that use a ML-/AI-based analytics approach can achieve a savings of 40–60% for service assurance in 5G networks compared with traditional assurance approaches.

3.3 Service assurance must evolve from a reactive to predictive mode

The largest CSPs typically have hundreds of systems for monitoring and reporting on network quality, as well as the impact that network quality has on subscribers of mobile data, broadband, and business services. These performance monitoring systems have been obtained from a combination of equipment vendors and independent

software suppliers and through internal ad-hoc development to support functions that are not available from a commercial software product. Integrating the myriad of performance monitoring systems to provide a unified view addresses this problem to some extent, but this involves high integration costs, resource and training burdens, bespoke development and scalability issues.

Consequently, service providers face two key problems: (a) they struggle to connect customer experience with the network and service quality issues, and (b) the siloed performance monitoring systems view results in highly manual and reactive operations.

Case study: Verizon, Predictive maintenance service

Verizon has created a multi-tiered approach to monetisation and operations efficiency using its AI capabilities. For example, the company offers AI services for specific use cases, including 'Verizon Condition-Based Maintenance (CBM)', as well as offering its AI platform for others to exploit. CBM uses data gathered from its own network to calculate optimum maintenance schedules for enterprise customers' networks. CBM works by gathering data from IoT devices deployed on customers' networks and by using the collected data in conjunction with its algorithms to predict and trigger maintenance orders. This enables low-cost scheduling of maintenance to prevent more-expensive reactive repairs if devices fail.

3.4 Assurance must become an integral part of service design from day one

Traditionally, service assurance has only been considered **after** building the networks and launching services. Service providers therefore do not prioritise investments in service assurance until after bringing the first customers onboard. Chronologically, CSPs first invest in the network build-out, then in network management systems and service fulfilment, charging and billing, and finally service assurance, just in time for handover to network operations.

Enterprise services are inherently dynamic: they can be created, altered and torn down at the flick of a switch. This means that assurance systems must align with these changing network and service behaviours from the outset. Making assurance an integral part of service design at the start will ensure that the service assurance modules are instantiated as and when the networks and services are created or modified.

4. Clean and granular network data must be at the heart of an assurance strategy for dynamic networks

Service providers require highly granular real-time data sources that provide them with the deep network visibility required to rapidly respond to microbursts and service quality degradation (for example, identify microbursts as they happen). Service providers also need a way to identify QoS issues from the perspective of the end user experience.

Figure 4.1: Diverse network data sources form a holistic view of the network

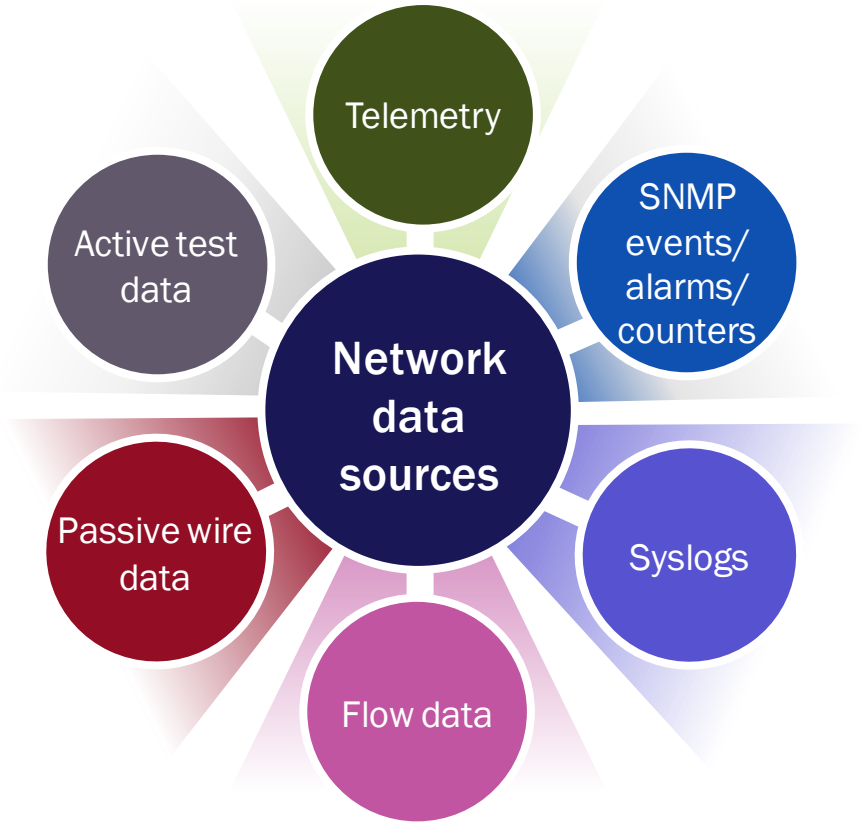

Source: Analysys Mason

4.1 Streaming network telemetry provides real time network performance data

Network telemetry (NT) is quickly emerging as a popular technology to obtain sub-second network performance data. NT-enabled devices such as routers, switches and firewalls can continuously stream real-time network data for network performance characteristics such as traffic and performance, I/O, error counters, queue statistics and so on. NT provides granular network data in real time, and when modelled using standard modelling language, the data structures of NT data can be conducive to applying analytics to generate insights.

Leading network equipment providers such as Cisco, Juniper and Arista Networks provide NT capability in their routers and switches and use a combination of vendor proprietary data models and/or industry standard data models such as YANG and OpenConfig. YANG is emerging as the de facto industry standard, but it has not yet cemented its position. It has gained support from platforms such as OpenConfig (a collaboration between service providers), which offers various off-the-shelf, vendor-neutral models for network operators to use.

4.2 Dynamic, active testing and monitoring provides a customer experience dimension to the network performance data

Active testing and monitoring provides service providers and enterprises with an additional dimension for monitoring networks. However, the active testing solutions must evolve with new capabilities to support the dynamic nature of SD-WAN and 5G networks. Service providers must consider dynamic active testing as part of a two-pronged strategy.

- Services such as the SD-WAN and other on-demand business services are most likely to be ordered via a self-service interface by the enterprise customer, so activation tests must be performed at the end of the provisioning process to validate the service configuration, performance and QoS, but before the service is made live and handed over to the customer. Similarly, active tests must also be performed to validate service configuration changes as well as during the troubleshooting process.

- Once operational, the dynamic services can be continuously monitored from an end-to-end perspective for any performance degradations. By simulating the user traffic on the SD-WAN paths, service providers can gain proactive insights into network performance for the chosen network paths. This can also help identify blind spots and congestion points in the network with respect to specific application types, which provides an extra layer of performance insights for remediation and troubleshooting

Independent virtual test agents or sensorscan be instantiated on-demand as part of the service provisioning process at key interfaces in the network and configured to perform active tests, either at predefined intervals (tests at more-frequent intervals provide real-time data) or on-demand. The solution is more scalable than traditional methods and provides real-time reporting on key performance metrics from the perspective of the end-to-end customer experience.

The highly granular data provided by tools such as network telemetry and virtual test agents enable service providers to identify transient network issues in real time. This ensures that network and service performance degradations are detected almost immediately and allows for rapid response times.

4.3 Combining and correlating data sources to create clean curated data is critical

Service providers usually already have many other sources of data gathered using well-established performance monitoring techniques – passive wire data, counters, syslogs, alarms/events and flow data are some of the examples. Each of these data sources, supplemented with telemetry and active test data, provide a different level of visibility and performance metrics. Service providers therefore need an extremely efficient way to correlate the diverse data sources to create a highly curated and clean data set for further processing. It is also very likely that the data sources come from multiple vendors, either network equipment vendors, independent assurance tool vendors or even service provider proprietary tools, which means that service providers also need a multi-vendor approach to data aggregation.

Insights generated from the network data will drive the actions and processes that are essential for supporting the operations of dynamic networks and services. Poor-quality network data will lead to inaccurate insights, and consequently, incorrect root cause analysis and troubleshooting actions on the network, which can lead to prolonged issue resolution times or service outage. For example, when providing a real-time aggregated view of the service performance of an enterprise at a regional level, it is critical to use an accurate view of the service performance KPIs at the site level, and in turn, the performance of the applications and network services within those sites. If the site level network performance data is inaccurate, the aggregated view at the regional level will be inaccurate. Therefore, the importance of clean accurate data cannot be overstated – the adage ‘garbage in, garbage out’ could not be more appropriate here.

5. A unified network performance analytics platform is the basis for assuring dynamic networks and services

Service providers must move away from a siloed approach for gathering network data to a horizontal unified network analytics platform that can ingest and aggregate network data from multiple sources in a multi-vendor environment. An illustrative reference architecture is presented in Figure 5.1 below.

Figure 5.1: Unified network analytics platform (illustrative reference architecture)

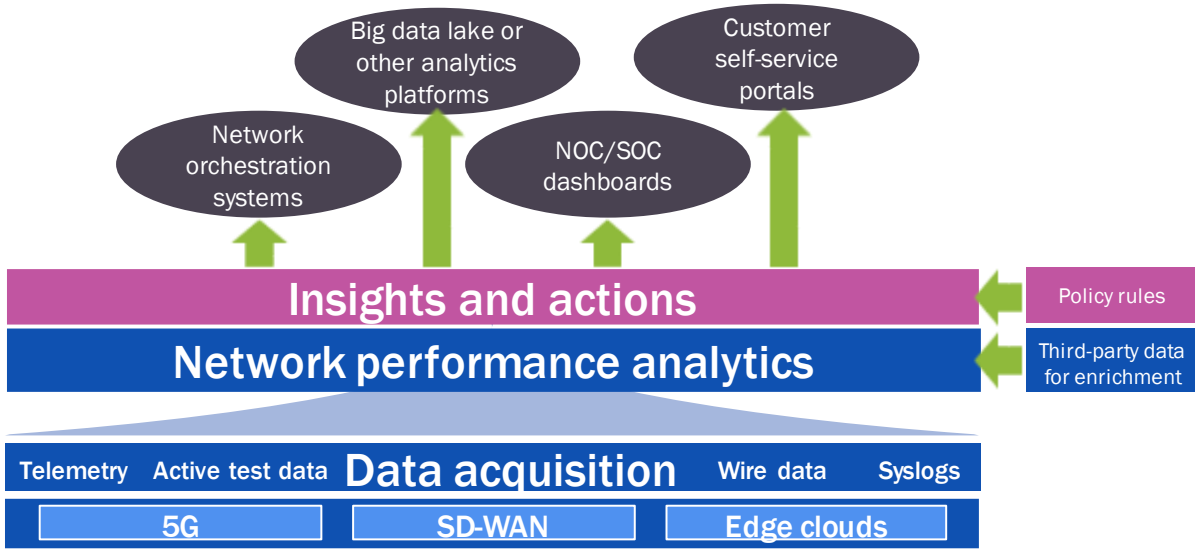

Source: Analysys Mason

5.1 Real-time insight based on network performance analytics provides the basis for assuring dynamic networks and services

From an architecture perspective, a unified platform will consist of the following key layers.

- **Data acquisition layer.** This layer provides the necessary abstractions to acquire data from the dynamic network infrastructure. Data from various network data sources such as the network telemetry, wire data and active test data is ingested into the platform. The network data is then curated and de-duplicated to create clean data for consumption by the performance analytics layer.

- **Performance analytics layer.** This layer correlates the clean network data to generate reusable data structures and metadata in standard formats, which drives algorithms and ML/AI models. The models generate various outputs such as historical KPI trends, anomalies, insights and predictions in near-real-time to be consumed by various higher layer applications.

- **Application layer.** This consumes the generated insights for further action, including creating the NOC/SOC dashboards for operations engineers; using insights as inputs to higher order big data analytics platforms; and generating performance reports for customer self-service portals. In addition, the insights can be combined with predefined policy rules to generate and trigger automated actions in network orchestration systems to drive zero-touch, closed-loop automation.

Use case: Enterprise self-service portals for reporting real time service performance

Enterprises are demanding near-real-time visibility into performance monitoring and reports of key performance metrics of their business-critical services. This will only increase with the proliferation of SD-WAN and, eventually, 5G-based enterprise services. Service providers that adopt the analytics-based service assurance approach outlined in this paper will already have recorded the real-time network performance for operations departments. Service providers can securely expose the same set of real-time data for reporting purposes through customer self-service portals for ease of consumption. Knowing that the service provider has the same view of the service performance as the enterprise customer provides complete transparency and improves customer trust.

5.2 Machine learning will enhance analytical insights and drive predictive assurance

The analytics function can be significantly enhanced through the application of ML techniques. Using reams of historical network data, supervised machine-learning algorithms can be trained to spot patterns (for example, degrading network performance) and trigger remediation routines (for example, supplement network capacity).

Case study: CenturyLink automatically detects issues with Ethernet and CDN services

CenturyLink has achieved a significant milestone for its SDN servers and services. Over 80% of issues are now automatically identified and resolutions are provided to restore services without human intervention. This has been possible using ML and AI for root cause analysis that can predict potential issues and suggest fixes.

ML also enables CenturyLink's SD-WAN Ethernet-based platform to identify and proactively fix potential issues with the service. One such issue was that customers were unable to change the amount of bandwidth that they purchase on the Ethernet service. This involved the need to manually change network requirements. A new managed bandwidth service was created that dynamically looked at typical usage and indicators of increases or decreases in data use. Once identified, bandwidth was automatically scaled up or down to better reflect customers' needs.

Continuous calibration of the algorithms can increase the accuracy of pattern matching and decisioning, to a point where there is enough confidence to establish predictive assurance. In a predictive assurance context, the models predict network or service issues, for hours, days or even weeks in advance, allowing enough time to take remediation action.

On the other hand, unsupervised learning algorithms have not received prior training on how to classify or label patterns but would employ a grouping or clustering model to organise data to understand potential structures and patterns before predicting outcomes. Reinforcement learning is when the machine-learning algorithm makes a single action and receives a notification on how good the decision was, and calibrates its next move based on the feedback. Of the three ML paradigms, supervised ML is the most-widely used technique and requires the skills of data scientists to set up and continuously calibrate the algorithms. All three machine-learning techniques are expected to play a critical role in achieving the vision of zero-touch automation on service provider networks.

Use case: Insights and actions for IoT

It is now widely accepted that the IoT revenue based on connectivity alone is not a good business model and additional value must be found. One such area is in the analysis of the data gathered from the huge number of IoT sensors and machines and providing intelligence and insight reporting services based on it.

Analytics technologies such as ML and AI lend themselves to some key characteristics that are inherently useful for business with IoT solutions. These include the processing of very large data sets to spot known patterns in data that trigger predefined actions. These can be 'learn-based' using ML technology, or preconfigured. AI helps with spotting variations without additional changes to rules. For monitoring applications, large quantities of data with no changes can be suppressed to prevent ongoing storage costs. Furthermore, trending data can be used to pre-empt threshold breaches, or to trigger actions or reporting. This may include reviewing maintenance schedules or performance metrics. In addition to known patterns, the ability to monitor what is 'normal' and report exceptions or unusual behaviour is of interest for security, faults or fraud.

6. Conclusion and recommendations

Service providers are enabling enterprise digital transformation by delivering on-demand business services such as SD-WAN and are preparing to launch innovative services based on edge clouds and 5G. To deliver these business-critical services with guaranteed service quality and differentiated SLAs – and at a fraction of the operational cost – service providers must embrace real-time network performance analytics. Traditional reactive performance monitoring approaches are not fit for purpose to monitor and assure dynamic services delivered on programmable networks. The complexity and dynamicity of these networks will demand highly responsive and automated operations that rely on real-time performance insights for rapid remediation and closed-loop automation.

Service providers should use highly curated clean network data with sub-second granularity (such as streaming network telemetry) to achieve this goal. The network telemetry data must be correlated with other real-time network data such as active test data and wire data in a real-time performance analytics platform to generate real-time performance insights. Applying ML/AI techniques will enable predictive operations and allow service providers to take pre-emptive action to prevent service quality degradations.

About the authors

Anil Rao (Principal Analyst) is the lead analyst for Analysys Mason's Automated Assurance and Service design and Orchestration research programs, covering a broad range of topics on the existing and new-age operational systems that will power telcos' digital transformation. His main areas of focus include: service creation, provisioning, and service operations in NFV/SDN-based networks, 5G, IoT, and edge clouds; the use of analytics, ML and AI to increase operations efficiency and agility; and the broader imperatives around operations automation and zero-touch networks. In addition to producing quantitative and qualitative research for both programs, Anil also works with clients on a range of consulting engagements such as strategy assessment and advisory, market sizing, competitive analysis and market positioning, and marketing support through thought-leadership collateral. Anil is also a frequent speaker and chair at industry events, and holds a BEng in Computer Science from the University of Mysore and an MBA from Lancaster University Management School, UK.

William Nagy (Analyst) is a member of the Telecoms Software and Networks research team in London, contributing primarily to the Automated Assurance programme. William holds a BSc in Physics from Queen Mary University of London and his dissertation was on investigating deconvolution and localisation imaging techniques in fluorescence microscopy.

This white paper was commissioned by Accedian. Analysys Mason does not endorse any of the vendor's products or services.

Published by Analysys Mason Limited • Bush House • North West Wing • Aldwych • London • WC2B 4PJ • UK
Tel: +44 (0)20 7395 9000 • Email: research@analysysmason.com • www.analysysmason.com/research

Registered in England and Wales No. 5177472