AUTOMATION IN THE CLOUD-NATIVE HYBRID NETWORK

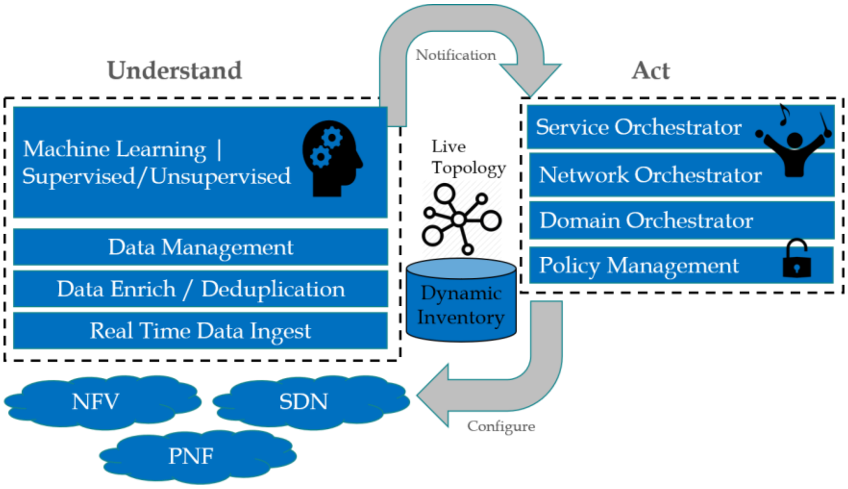

2/8/2018

Patrick Kelly, Appledore Research

www.appledoreresearch.com

Published by Appledore Research LLC • 44 Summer Street Dover, NH. 03820

Tel: +1 603 969 2125 • Email: Patrick.Kelly@appledorerg.com• www.appledoreresearch.com

Figures and projections contained in this report are based on publicly available information only and are produced by the Research Division of Appledore Research LLC independently of any client-specific work within Appledore Research LLC. The opinions expressed are those of the stated authors only.

Appledore Research LLC recognizes that many terms appearing in this report are proprietary; all such trademarks are acknowledged, and every effort has been made to indicate them by the normal USA publishing standards. However, the presence of a term, in whatever form, does not affect its legal status as a trademark.

Appledore Research LLC maintains that all reasonable care and skill have been used in the compilation of this publication. However, Appledore Research LLC shall not be under any liability for loss or damage (including consequential loss) whatsoever or howsoever arising because of the use of this publication by the customer, his servants, agents or any third party.

Contents

Automation in the Cloud-Native Hybrid Network

EXECUTIVE SUMMARY

Virtualization technologies, most commonly referred to as network function virtualization (NFV) and software defined networking (SDN), will require more automation in order to realize both the economic benefits and achieve the agility necessary to meet evolving cloud business models. The benefits of virtualization are that software and hardware can now be disaggregated, and CSPs can deploy commodity-based merchant silicon and network functions (NFs) independently. Ultimately, this means greater choice and freedom from vendor lock-in.

Existing workflow processes and systems were designed for fixed compute and network architectures. Physical networks are static and difficult to reconfigure. Physical network elements (NEs) contained custom high-performance, application-specific circuits and vendor-specific configurations. Despite the performance benefits, these NEs instigate high labor cost whenever there is a need to change, add, repair, and scale out to accommodate consumer demands. AT&T has estimated that the labor workforce used to support, repair, upgrade, and deploy its global network could be reduced by 30 percent over the next decade through virtualization and software control.

Legacy software systems mirror the static network and have been optimized around an inflexible set of assets. The rigidity means fewer choices and limits to operations as it becomes more difficult to improve with more modern equipment. The processes and the systems to support legacy systems require a high degree of human "touch points," limiting the ability to deliver on-demand and self-healing cloud-enabled services. This manually intensive process is further complicated by the existence of discrete software systems, which are deployed in technology silos. Further exacerbating the problem is the fact these same solutions have been deployed to satisfy single workflow processes in the service chain. These processes operate independently of the service lifecycle. All in all, these legacy products did what they were designed for well, but significant custom integration is necessary to link together order management, service fulfilment, network activation, inventory management, and service assurance.

Current methods, workflows, and software systems must evolve to support cloud-native architectures. Appledore Research asserts that control loops must be used in the cloud and virtualization domain to achieve automation and simplification. We acknowledge that many CSPs remain extremely skeptical about turning management and control of the network over to a machine.

This thinking is counterintuitive to the ways in which control loops have been applied to rocket navigation, and more recently self-driving cars. A properly designed control loop is self-correcting and eliminates redundant and potentially divergent logic and operations.

This white paper will provide the reader with a guide to the best practices most likely to achieve the benefits we cited above. We provide a framework that can be applied to the current operations environment of many CSPs. Our view is that legacy systems will not necessarily be replaced but instead be integrated into a cloud management framework. The cloud management system will be intent based. Intent-based networking allows operations and planning to define a desired state of the network and to automate network orchestration to implement and maintain a set of policies and rules that achieve that desired state. We will focus on the need for real-time operations data and assurance in the highly automated or complex service environments required by digital services and NFV.

CSP OPPORTUNITY IN CLOUD SERVICE DELIVERY MODELS

Enterprise businesses are migrating more of their workloads to the cloud. In the communication sector, more cloud service delivery models continue to proliferate as enterprises take advantage of IaaS, PaaS, and SaaS solutions to satisfy more complex application requirements.

In the telecommunication market most providers can provide variations of all three, but given the dominance of hyper-scale cloud providers such as Amazon and Google, most CSPs are finding more success in offering cloud services that focus on the network, security, and analytics. Examples include SD-WAN, firewall, threat detection, and universal CPE for branch and remote locations.

New revenue growth opportunities for CSPs combine cloud service delivery models with vertical solutions. ARG forecasts that the total addressable market for Internet of Things (IoT) will reach $1.2 trillion by 2020. CSPs can provide both the connectivity and management of IoT devices. Value-add applications can apply the vast amount of data in the CSP network and combine it with machine learning models to apply predictive analytics. Monitoring and processing such large amounts of data will require edge computing and analytics in the cloud data center. Global CSPs have both the infrastructure and footprint to win a large share of the IoT managed-services market.

Central Office Re-Architected as a Data Center (CORD) is gaining broader acceptance in the industry. It brings data center economies-of-scale and the agility of the cloud to residential, enterprise, and mobile customers. CORD will be built from commodity servers, white-box switches, disaggregated access technologies, and open-source software. CSPs can leverage a common hardware and software infrastructure to offer traditional connectivity as well as cloud services for residential, enterprise and mobile customers. CORD allows residential, mobile, and enterprise customers to configure and manage their service packages on-demand and in real time.

VIRTUALIZATION MARKET GROWTH AND ADOPTION

Appledore Research believes virtualization represents a sea change in the communication sector and it is a key technology enabler to fully realize the benefits of cloud services. Although the adoption of virtualization has not met the expectations of market pundits since the launch of ETSI NFV in November 2012, Appledore Research finds that deployments have been accelerating in the past six months. Appledore Research provides global deployment data for NFV and SDN, which is updated quarterly.

Appledore Research estimates that the virtualization market will be worth $120 Billion in infrastructure investments. We estimate that 40 percent of the network and outside plant can be virtualized. As virtualization technology (NFV and SDN) replaces legacy routers, switches, and gateways, automation will become vital in operational centers.

Appledore Research analysis of CSP financials last year finds that most operators’ annual operating expendutyre (OPEX) exceeds capital expenditure (CAPEX) by a ratio of 6:1. Savings in OPEX can fund a significant investment in new capital equipment only if in cases where it gives rise to substantial reductions in OPEX.

In the same way that applications are supported by dynamically configurable and fully automated cloud environments, virtualized network functions allow networks to be agile and responsive enough to automatically accommodate the needs of the traffic and services running over them. NFV and SDN complement each other but are also codependent. One without the other will not accrue the economic benefits of cloud-enabled virtualization.

CURRENT CHALLENGES IN THE OPERATIONAL ENVIRONMENT

Most network operation centers today require skilled technicians who can perform mostly manual tasks. At best, semi-automated workflows require swivel chair management in order to piece together insights and to achieve situational awareness. The business impact of this approach is slower response times, lost productivity, and unsatisfied customers. NFV and SDN technologies will overwhelm operational staff charged with assuring services, and situational awareness will degrade since workloads will be more dynamic — moving and changing according to demands and resource availability. The advent of containers and microservices will further complicate the situation and render existing management tools obsolete.

Hyper-scale cloud providers such as AWS, Google, Microsoft, and IBM recognize that new services will rely on many different technologies contained within complex multi-domain networks where digital services extend beyond the network boundaries of the CSP. These diverse ecosystems and the services they engender will require a much higher degree of automation. NFV, SDN, 5G and other emerging technologies will underpin the new infrastructure, which will reveal that current management systems have not been designed to support new, dynamic hyper-scale infrastructure.

The importance of the customer cannot be understated. To date, billions of dollars have been invested in customer care and service operations to improve Net Promoter Score (NPS), improve loyalty, and reduce churn. Table 1 highlights the effectiveness of a NOC/SOC and its ability to provide superior customer experience in the face of disruptive technology change. It reflects a change in priorities from today's operational performance metrics to future business outcomes driven by competition and customer demands.

Table 1: Evolution of Service Centric Business Models

Challenge	Target Outcome
Focus on Technology	Focus on Customer and Service
SLA Violation for VIP Customers	Predict and Remediate Proactively
T2R Cycle Times Hours / Days	Isolate, Test, and Identify Root Cause Minutes/Hours
Inaccurate Network Topology	Customer Context Live Topology
Network Capacity Constraints	"Infinite Scaling" Applied via Virtualization and Automation

Source: Appledore Research

A popular term used in the industry more recently is "service intent." This is the desired state of the network and infrastructure and its ability to maintain that condition as demand changes. To achieve this goal, we must look to control theory, which has been successfully applied in other industries.

Legacy service assurance systems must evolve to support closed-loop automation, necessitating coordination with network and service orchestrators. As workloads increase, move, and change, the management systems must perform in real time and close any gaps between the current state and desired state of service KPIs.

A critical piece in the system to achieve the benefits of closed-loop automation is a dynamic inventory that maintains a current state of the network. It should provide a consistent view of the customer and the underlay of network resources to maintain accuracy of a constantly changing network. Without dynamic inventory, the CSP is unable to correlate the customer experience to the actual performance and availability of the network.

The output of dynamic inventory is the ability to render a live topology graph. This graph can be utilized to both understand the dynamic network but also to provide context-sensitive data regarding customers and their relationship to the service-aware network.

CONTROL LOOPS APPLIED TO AUTOMATION IN THE HYBRID NETWORK

Appledore Research asserts that control loops must be used in the cloud and virtualization domain to achieve automation and simplification. We acknowledge that many CSPs are extremely skeptical of the idea of turning over management and control of the network to a machine for fear of catastrophic consequences, such as degraded performance or even complete failure of the network. This thinking is counterintuitive to how control loops have been applied to audio and video circuits, rocket navigation, and more recently self-driving cars. Control loops have several characteristics that make them more reliable and less risky than technologies that appear to be simpler. A properly designed control loop is self correcting and eliminates redundant and potentially divergent logic and operations. When applied correctly, service assurance and service fulfillment systems are affected by the same code so that changes to one workflow process is consistent and has the same desired effect to the other. Control loops provide a logical place to both identify a problem — a “state change” —and correct it.

Control loops can be integrated with policy rules, analytics, and active topology graphs. Figure 1 provides a mapping to classical closed-loop control systems and demonstrates how the network and management systems can be applied to this proven model.

Figure 1: Closed-loop control diagram applied to virtualized network and management systems

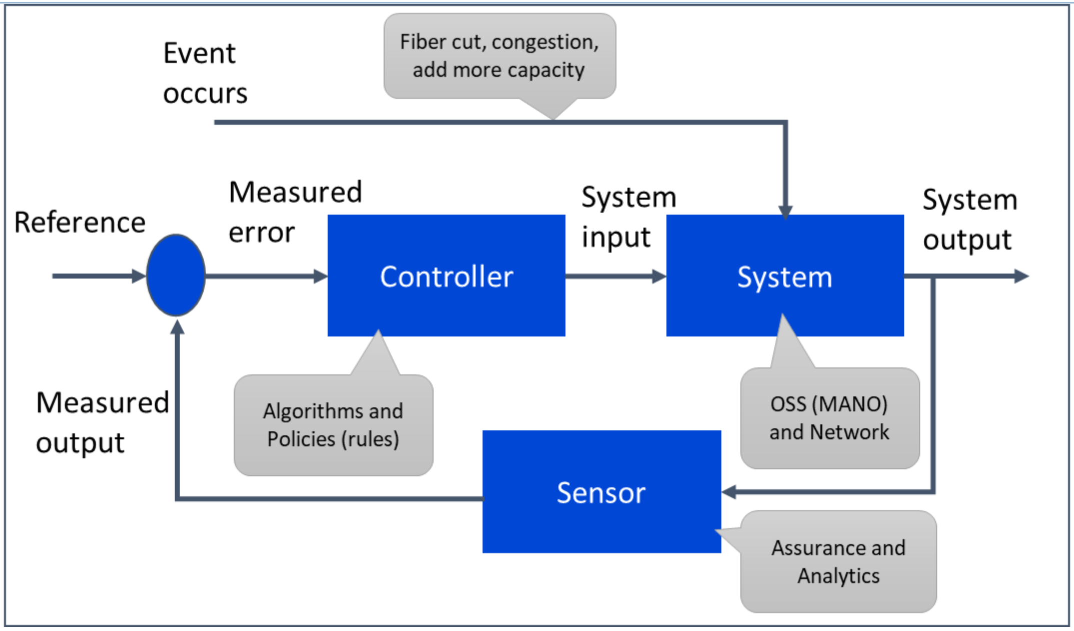

Source: Appledore Research

In closed-loop control, the control action from the controller is dependent on feedback from the process in the form of the value of the process variable (PV). In the case of a boiler analogy, a closed loop would include a thermostat to compare the building temperature (PV) with the temperature set on the thermostat (the set point - SP). This generates a controller output to maintain the building at the desired temperature by switching the boiler on and off. A closed-loop controller, therefore, has a feedback loop that ensures the controller exerts a control action that then manipulates the process variable so that it is the same as the "Reference input" or "set point."

In a hybrid virtualized network, any event-based changes such as a fiber cut, network congestion, or request for more capacity triggers a state change in the network. In the case of the fiber cut, a service assurance system would act as sensor to notify the SDN controller or service orchestrator of a state change. The controller, operating in tandem with the MANO, would implement a control action to reconfigure the network. At this point, multiple actions would be performed to determine network path reroute options, assess cost paths, and avoid links that are near capacity. All this activity would occur at near real time, autonomously, and perform configuration changes based on policy parameters defined by human experts.

Implementing software automation and virtualization technologies will yield vastly lower infrastructure and service lifecycle management unit costs as the network scales out. Appledore Research estimates this unit cost to be between 10-100x savings in capital and labor deployed. These [unit] cost reductions make it feasible for CSPs to consider new channels to market and more customized services that, until now, could not be justified. This approach also makes it possible to automatically "groom" thousands of NFs and flows to maximize the capacity utilization of infrastructure.

A closed control loop converges several functions that have traditionally been independent "stacks" in the OSS/BSS environment, specifically fulfillment and assurance. In this new scenario, assurance and analytics become the intelligence that guides orchestration to effect repair or scaling through a process that is in effect "re-fulfillment." What were two stacks now become a single, continuous, closed loop (Figure 2).

Figure 2: Orchestration, assurance, analytics role in the closed loop system

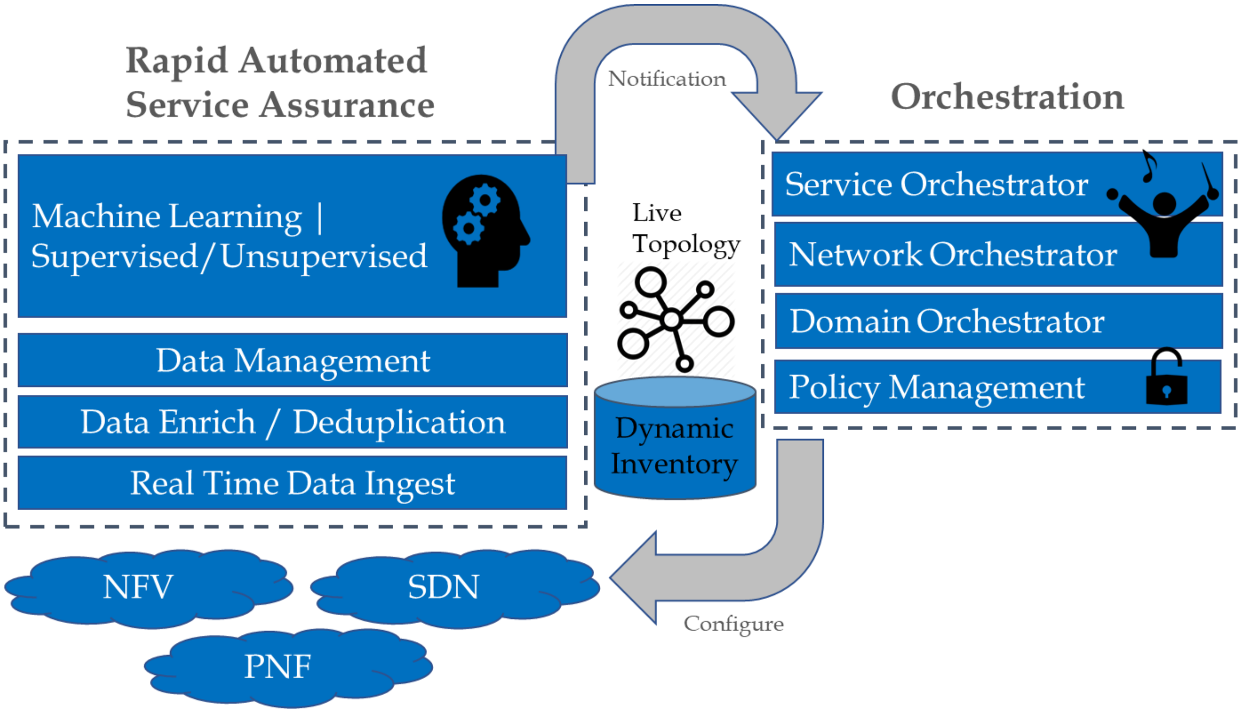

Source: Appledore Research

The importance of service assurance in this expanded role implies that architects and implementers must design assurance and analytics into the system on day 1. No longer is there a physical, deterministic relationship between facilities and services. That relationship is now highly dynamic and must be tracked, modeled and understood from the moment an order is processed. Without this mapping, customer and service impacts cannot be understood, SLAs cannot be enforced and healing cannot be performed — even manually. Assurance will be transformed from an independent process, intended to manage a network, into an integrated process — an integral part of orchestration. In fact, without assurance integrated into this new orchestration process, it will become difficult to fulfill the services that underpin the business case behind network and service virtualization.

BEST PRACTICES TO ACHIEVE AUTOMATION IN THE CLOUD NETWORK

Automation is already occurring in parts of the telecommunication network. Self-Organizing Network (SON) is an automation technology designed to make the planning, configuration, management, optimization and healing of mobile radio access networks simpler and faster. Radio configurations lend themselves to automation, since they have many parameters, and an optimum radio plan depends on complex interactions between cells, neighbors, frequency use, and a dynamic traffic load (mobile user equipment). It can't be done manually for obvious reasons, the main of which is adequate time to satisfy demand request. Centralized SON (C-SON) functions are typically concentrated closer to higher-order network nodes or the network OSS. C-SON, applied across regions of cells, have been proven in large-scale deployments. The benefits are periodic optimizations that occur more frequently; faster error notifications; and allowing for RANs to dynamically react to changes in offered load (traffic) — all things that would be impractical if performed manually.

No single control loop would be able to meet the requirements of even a simple set of services. Instead, Appledore Research believes that control loops will operate within each technology domain. We like to use the phrase “think global but act local.” The VIM/infrastructure, SDN, NFV, RAN (SON) and other technologies will perform autonomous actions at the local level first. Looking to the real world for confirmation, we see also that most VIMs and all “hyper scale clouds” operate successfully with this principle. Exceptions at the local level will expose an abstracted view to a higher level where service orchestration and end-to-end service model/assurance can take over. Note that this demands loose coupling, API exposure of lower-loop state, and an explicitly federated approach to E2E assurance. Table 2 provides a summary of best practices and the rationale for each of the six areas of the closed-loop methods and processes.

Table 2: Closed-loop Best Practices

Principal	Benefit / Rationale
Layered orchestration and nested loops	Simplification, latency reduction, abstraction based
Abstraction to higher levels of orchestration and control loops	Plug-n-play, high levels of re-use, minimize maintenance, simplify conflict management
Catalog driven micro service based	Essential to support layers and nesting, re-use, minimizes maintenance, allows loose coupling
Modular loop architecture both in domains (technology, subnets) and in software functionality	Multi-vendor support, easy to add components, innovate faster
Real-time operation	Seamless operation and improved SLA guarantees
Single orchestration method for fulfillment, healing, capacity expansion	Avoid duplication and divergence of methods

Source: Appledore Research

Each of these best practices reduces risk and aids in simplicity. Applying these principles results in a single, logical method to maintain that data is accurate, reduce logic to understandable components which may be re-used and de-bugged only once. It supports a set of rules that allow the logic to react to events and changes in the environment (infrastructure failure, congestion, SLA violation) avoiding the need to perform trend analysis and to answer to false alarms.

When we speak of “one continuous process,” there is an obvious tendency to expand this to “one single monolithic loop covering the service.” As noted earlier, experience shows that this is not a viable option. The best systems, from self-driving cars (e.g.: feedback systems) to robotics are all based on federated, hierarchical control loops. These implement the principle of beginning with the simplest, most deterministic action. Figure 3 represents an overlay of the best practice principles onto the Appledore Research reference architecture.

Figure 3: Closed Loop system overlay onto Appledore Research reference architecture

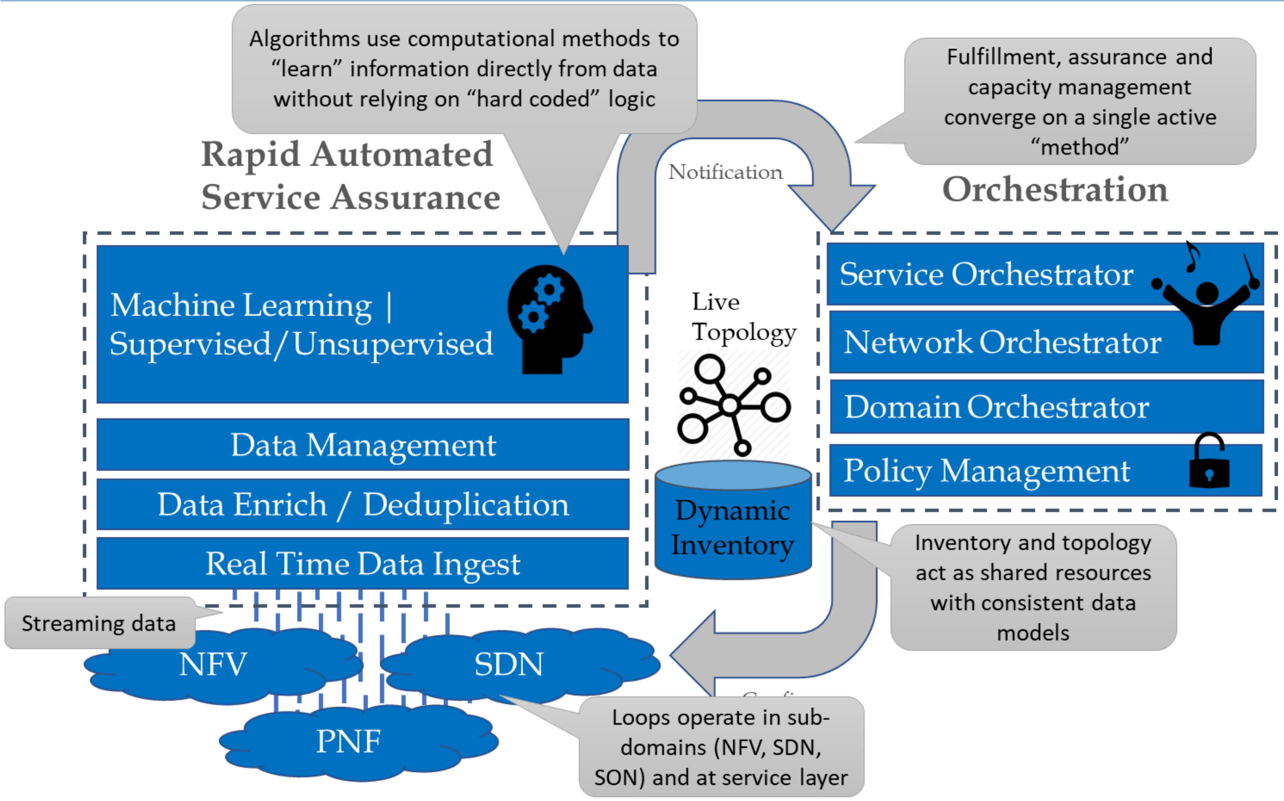

Algorithms use computational methods to "learn" information directly from data without relying on "hard coded" logic

Fulfillment, assurance and capacity management converge on a single active "method"

Inventory and topology act as shared resources with consistent data models

Streaming data

Loops operate in sub-domains (NFV, SDN, SON) and at service layer

Source: Appledore Research

THE ROLE OF SERVICE ASSURANCE IN A CLOSED-LOOP SYSTEM

Automation, the "merger" of fulfillment and assurance, and the emerging micro-services based "DevOps" methods all imply a new, and in many ways, expanded role for what we call Rapid Automated Service Assurance (RASA). Closed loops, and the automation they effect rely on timely, accurate, insightful intelligence that *identifies problems, correlates them with technical causes, and drives the original process to make automatic modifications*. Assurance and analytics combine to become complementary components of that intelligence platform.

RASA also takes on the role of federating and rolling up individual domains into an end-to-end view of the network and services. It applies service models and customer-service mappings to domain data and delivers complex analytics that add additional intelligence to otherwise technical decisions.

At Appledore Research, we also firmly believe that next-generation service assurance is a (functionally) converged platform that includes both assurance and analytics. Its function is to take all the relevant data — fault, performance data, streaming data, historical data, probed data — and bring them together in an informed view of what is occurring in the network and what the best actions should be. This intelligence can then be applied to the parametric (policy and rule-driven) orchestration functions capable of healing, scaling, optimizing or otherwise reconfiguring the network. Simultaneously, it delivers a more holistic infrastructure to understand customer experiences, service experiences and the best actions to take in the customer domain (and, therefore, drive *those* processes).

Assurance must now operate at the speed necessary for whatever process it is serving. For most virtualized network recoveries, this means it must be real-time and based on real-time reference data and correlation. Since virtualized networks are dynamic, it is critical that topology and inventory be updated in real time, and that the correlation or association of VNFs with infrastructure, flows, with underlay, and other dynamic mappings are resolved accurately and quickly. Inventory systems must also record a history of state changes that support state changes. These logs are useful to operations in the trouble to resolve process.

Similarly, many opportunities to interact with customers are fleeting and require real-time context. The system just observed X originating from a data stream and triggers a contextual event to a "listener," which may generate a corresponding offer that is not only relevant but also improves the customer experience.

At the same time, it is important to recognize that a) not all data can be real-time, and b) real-time and "data lake" analysis is complementary and must be part of a single, flexible process. An example is a real-time process that identifies user behaviors, combined with offline historical analytics to identify groups of subscribers to monitor. The combination yields qualified segmentation, and real-time triggers that allow for proactive assurance, upsell offers, or network reconfigurations to meet SLAs.

Appledore Research believes that as networks and services become more diversified and complex, and as services become more customized and personalized, the role of assurance and analytics will only grow. It will become increasingly important to recognizing trends and customer experiences, as well as to understanding technology in the context of customers. This intelligence can then be used to improve both the business (e.g.: offers, packages) and the underlying service quality (e.g.: network performance, service configuration).

We can also safely assume that CSPs will want the flexibility to use assurance and analytics to answer questions that arise in the day-to-day operation of the business. Tthe days of pre-ordained assurance "stove-pipes" are going to give way to an environment where data sources (e.g.: probes, collectors, etc.) are pooled and abstracted, and where CSPs quickly pull together analytical methods to address questions and opportunities as they arise. This demands that the assurance (and analytics) infrastructure is real-time capable, programmable, and exposes APIs to various operational processes. Figure 4 provides a more detailed diagram of the role of RASA in closed loop automation.

Figure 4: RASA subcomponents and data flows in a closed loop system

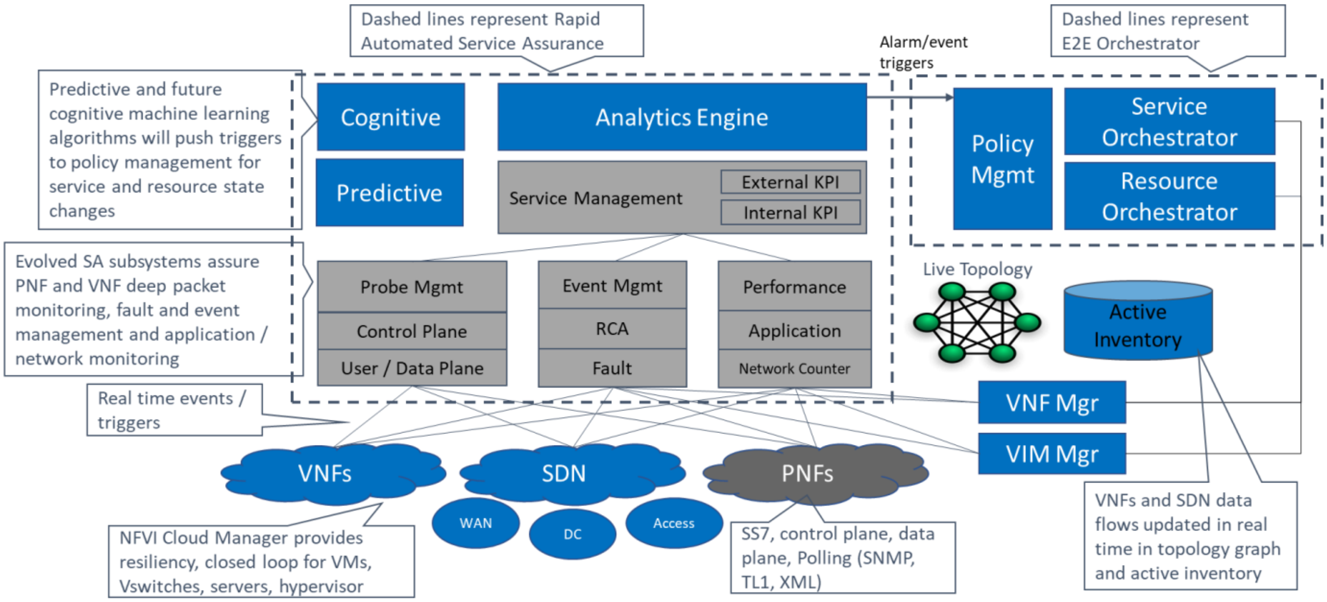
Dashed lines represent Rapid Automated Service Assurance
Dashed lines represent E2E Orchestrator
Alarm/event triggers
Predictive and future cognitive machine learning algorithms will push triggers to policy management for service and resource state changes
Cognitive
Analytics Engine
Predictive
Service Management
External KPI
Internal KPI
Policy Mgmt
Service Orchestrator
Resource Orchestrator
Evolved SA subsystems assure PNF and VNF deep packet monitoring, fault and event management and application / network monitoring
Probe Mgmt
Control Plane
User / Data Plane
Event Mgmt
RCA
Fault
Performance
Application
Network Counter
Live Topology
Active Inventory
Real time events / triggers
VNF Mgr
VIM Mgr
VNFs
SDN
PNFs
WAN
DC
Access
NFVI Cloud Manager provides resiliency, closed loop for VMs, Vswitches, servers, hypervisor
SS7, control plane, data plane, Polling (SNMP, TL1, XML)
VNFs and SDN data flows updated in real time in topology graph and active inventory

Source: Appledore Research

RECOMMENDATIONS

What do CSPs need to do to prepare for new services, network technologies in applying closed-loop automation? Experience shows that automation should not be “bolted on” but rather designed from the beginning. It must first be designed into local domain systems, such as those implementing NFV/MANO, SDN, SON, and others. It should also be designed into the end-to-end service views.

We believe there are six critical considerations when implementing successful, flexible automation loops and controlling risks:

1. Automation is critical to achieving the cost reductions, revenue growth, and operational agility that is the core value of virtualization.
2. Automation is not an afterthought; it must be designed into orchestration, assurance, policy and analytics — and throughout the OSS/BSS environment.
3. Simplicity is a virtue. Focus first on “local” loops and automation — e.g.: SDN, SON, VIM layer, etc. where complexity is more constrained.
4. Assurance and analytics provide the intelligent guidance and constraints on how to best fulfill, scale, heal, or optimize services. They must mesh seamlessly with orchestration to effect “re-fulfillment.”
5. Beware “use case automation” which is in effect hard coded to a specific use case or type. Such solutions will almost certainly not apply well to new use cases and may require significant maintenance over time. (Remember that automation is designed in…). Use cases should simply be parametric configurations of the basic process.

6. Efficient flexible automation is part of a modular "DevOps" construct through which micro-services (with implicit "micro-loops") roll up into more complex commercial services with little modification or development. This maintains loose coupling.

SUMMARY

Virtualization technology may be exciting on its own, but CSPs' real financial payoff is virtualization's ability to radically alter economics, lowering both OPEX and CAPEX, and to thereby enable new profitable services, business models and thus grow revenues and margins. This economic benefit, however, is entirely dependent on two pillars: a) automation and b) intelligence to guide that automation, as provided by RASA, which includes a convergence of traditional assurance and analytics capabilities — all capable of operating at high scale, in real-time.

The traditional silos of "fulfillment", "assurance" and "capacity expansion" will merge into a single, flexible process that can create new services, modify them, restore them and optimize the utilization of resources. This closed control loop-based process demands highly flexible orchestration at its core, its flexibility defined by policies, and the overall action directed by assurance and analytics. Successful deployments will have significant policy-driven flexibility, and clear separation between the orchestration role and the assurance/analytics roles.

We also believe that there should no longer be a single, monolithic control loop. Instead, we believe in modularity and federation of constituent domains — an approach that will greatly reduce technical and implementation risk. A "monolithic" approach to closed-loop automation, we believe, is fraught with risks and ignores the fact that deep understanding comes at the technology domain level.

Appledore Research strongly urge CSPs and suppliers alike to focus on modular solutions, based on nested loops, with deep expertise at the domain level, and that are guided by highly flexible "analytics and assurance" that can direct the most effective, and cost-effective utilization of a CSP's costly infrastructure.

The success of cloud services delivered in a hybrid virtualized network must apply the principals we outlined in this white paper. In summary:

1) Automation begins with strong local domain systems at the NFVi, SDN, MANO, and SON.

2) A single control process, based on highly flexible and parametric orchestration that is easily modified by intelligence from assurance and analytics.

3) A hierarchical model of "nested" control loops that operate first at the simplest, most local level possible, yet are federated into higher-level loops.

4) A federated end-to-end service model that "rolls up" local technologies and domains into an end-to-end view and exposes this model and its implicit customer-resource correlation to multiple business processes.

5) The ability to introduce additional intelligence on customers, overall network condition, and other trending factors such as those discovered over time through analytics.

6) Software that operates in real-time, able to accurately reflect the status of a constantly changing virtual network.

7) The network will remain a combination of virtual and physical for at least a decade, yet we cannot wait 10 years to improve operational efficiency. New processes must apply to hybrid environments.

8) Powerful and timely intelligence from Assurance and Analytics that can identify contextual trigger conditions, predict impending events, and ultimately help select "the optimal" solutions, considering multiple conditions.

In concluding designers and implementers of automation software systems should anticipate the uncertainty of humans and unforeseen circumstances that will occur. The system should be resilient to react and recover from potential catastrophic situations. This will move the industry forward and dispel some of the fear in turning over the proverbial operational keys to the machine.

ABOUT THE AUTHOR

Patrick Kelly is the Founder and Principal Analyst of Appledore Research. He has more than 25 years of experience in product management, business development, and technology consulting. He has advised executives and developed actionable business plans to help hundreds of technology companies profit in high growth software segments of the market. He is the leading authority and has published research in the areas of cloud economics, virtualization of the network, NFV, SDN, machine learning, orchestration, analytics, service management, and customer experience management.

Patrick founded Appledore Research Group in 2014 to focus on the business impact of cloud and virtualization in the telecommunication market. Prior to Appledore, he was Research Director at Analysys Mason, co-founder of OSS Observer (acquired by Analysys Mason in 2008), Director of Product Management for Aprisma (acquired by CA), and held many technical roles in the field supporting both enterprise and service provider customers.